Okinawa Institute of Science and Technology Graduate University

Thesis submitted for the degree
Doctor of Philosophy

Investigation of the Bayesian Sensory-Motor Integration in the Cerebral Cortex

---

by

Sergey Zobnin

Supervisor: Kenji Doya

March 2024

Declaration of Original and Sole Authorship

I, Sergey Zobnin, declare that this thesis entitled "**Investigation of the Bayesian Sensory-Motor Integration in the Cerebral Cortex**" and the date presented in it are original and my own work.

I confirm that:

- No part of this work has previously been submitted for a degree at this or any other university.
- Reference to the work of others have been clearly acknowledged. Quotations from the wok of others have been clearly indicated, and attributed to them.
- In cases where others have contributed to part of this work, such contribution has been clearly acknowledged and distinguished from my own work.
- None of this work has been previously published elsewhere.

Signature:

Date: March 2024

Abstract

Animals need to sense their environment and differentiate sensory states to optimize behavior. For example, a force too weak may fail to pull a desired object closer, while too much effort will be unnecessarily exhausting. The prior experience interacting with similar objects can help generate the required action and produce an expectation of the sensory feedback. Combining prior expectation with the actual sensory observation is formalized with Bayes' rule. The neural mechanisms for such probabilistic computations remain unclear but, in mammals, can be attributed to the specific anatomical patterns of the cerebral cortex. In particular, the pathways carrying bottom-up sensory signals mostly terminate in layer 4 and activate pyramidal neurons of the layer 2/3 of the sensory cortex, while the neurons of deeper layers 5 and 6 mostly receive projections from other cortical areas including the motor cortex. Based on this, I hypothesized that the pyramidal neurons of the superficial layers encode sensory evidence while the deeper layer neurons encode prior and posterior estimations. I conducted a series of lever-pulling experiments on mice, where the lever provided a variable tactile feedback. Calcium imaging of the S1 cortical area using a prism lens enabled me to analyze how matches or mismatches between action-dependent prediction and actual sensory inputs are encoded by neurons in superficial and deep layers.

I have identified functional asymmetry between superficial and deep cortical layers in the context of the probabilistic sensory-motor task. The deep layers were more associated with the prior information about the expected tactile stimulus in terms of the number of task coding neurons and the amount of information per given population size. In comparison with the deep layers, superficial layers were stronger associated with the sensory information about the actual tactile stimulus, but also encoded prior information. This research contributes to the study of the neural mechanisms underlying probabilistic estimation in the cerebral cortex.

Acknowledgements

Abbreviations

BMI - brain-machine interface
GRIN - gradient refractive index
HE - heavy expected trial
HT - heavy actual trial
LE - light expected trial
LED - light-emitting diode
LT - light actual trial
PBS - phosphate-buffered saline
ROP - roving oddball paradigm
SWI - size-weight illusion

Glossary

AAV - adeno-associated viruses
AIP - anterior intraparietal area
ANOVA - analysis of variance
ChR2 - channelrhodopsin-2
DAPI - 4′,6-diamidino-2-phenylindole DNA fluorescent stain
df - degrees of freedom
dF/F - change in calcium fluorescent activity
GCaMP6f - genetically encoded calcium (Ca^{2+}) sensor
ICA - independent component analysis
M1 - primary motor area
p - *p*-value
P(S|X) - sensory likelihood of *X* given *S*
P(X) - prior probability of *X*
P(X|S) - posterior probability of *X* given *S*
PCA - principal component analysis
PMv - ventral premotor area
PPC - posterior parietal cortex
S1 - primary somatosensory area
S2 - secondary somatosensory area
S1FL - somatosensory forelimb area
X - chi
z - *z*-value

Contents

List of Figures

List of Tables

Chapter I. Introduction

1.1 Bayesian inference in the brain

Animals need to sense the environment and differentiate sensory states to optimize their behavior. For example, to bring an object closer, such as food, the weight of it has to be estimated to apply sufficient muscle strength. In real life, instantaneous sensory observations are frequently partial and, therefore, are associated with uncertainty, which makes estimations of the environmental variables probabilistic. Such an estimation can be influenced by the previous experience, which can be formalized with Bayes' rule. The rule prescribes to update a belief in a hypothesis X upon receiving a sensory evidence S, such that the posterior probability $P(X|S)$ for a hypothesis is proportional to the prior probability $P(X)$ for that hypothesis multiplied by the likelihood $P(S|X)$ to observe that sensory evidence, given that the hypothesis is true.

$$P(X|S) = \frac{P(X)*P(S|X)}{P(S)} \quad (1)$$

The idea that the brain is capable of probabilistic estimation by using and updating the internal model of the world, was first introduced by Helmholtz in the 19th century and was later applied to neural networks (Dayan et al. 1995; Lochmann and Deneve 2011). Behavioral studies demonstrated that humans can combine prior knowledge with new sensory evidence in a nearly Bayesian manner (Ernst and Banks 2002; Körding and Wolpert 2004). Animals were also shown to use information about their own action to update their posterior belief about the environmental state (Funamizu, Kuhn, and Doya 2016). This process can be described as dynamic Bayesian inference and involves integration of sensory inputs with information about the animal's own actions.

To understand how the brain can be capable of combining probabilities in a Bayesian fashion, I will review the biologically inspired models of neural networks.

1.2 Anatomical and functional features of the cerebral cortex

Anatomical and functional features of the cerebral cortex provide important clues for investigating biological mechanisms of Bayesian inference. Many of these features are remarkably uniform throughout the cerebral cortex (Douglas and Martin 2004), the most prominent one being laminar (layered) structure. Different layers have characteristic distributions of cell types and patterns of connectivity.

Within the connections among cortical and subcortical areas, feedforward (ascending) and feedback (descending) pathways are distinguished. A typical trajectory of the neural activity propagating through feedforward axonal bundles is from receptor neurons through thalamus into the primary and, then, secondary sensory cortical areas. Feedback pathways are represented by axonal projections in the opposite direction.

By sheer numbers, the long range feedforward, feedback and lateral connections between cortical regions are very few. Retrograde tracing in multiple cortical areas of the macaque brain consistently showed that 95% of neurons projecting to a given cortical area lie within 2 mm from the injection site, and the number of the labeled neurons rapidly drops with distance (N. T. Markov et al. 2011). The neurons providing extrinsic (long range) inputs were 1.3% for subcortical and 3% for corticocortical connections.
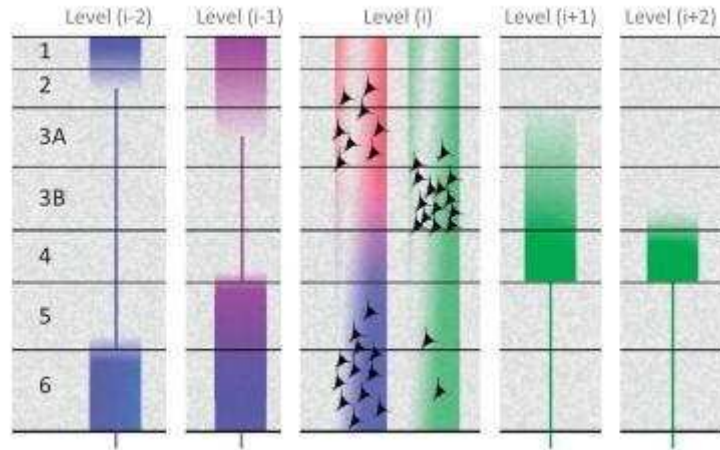
Figure 1.1. Asymmetry between extrinsic feedforward and feedback connections. Color bars of Level(i) represent the laminar position of feedback projecting neurons (red and blue) and feedforward projecting neurons (green). The color bars of other levels show the termination pattern of their axons. The distinct laminar preferences increase gradually with distance. Outer layers receive feedback mostly from deep layers. Layer 4 receives feedforward mostly from superficial layers (Adams, Shipp, and Friston 2013).

There exists an anatomical and functional asymmetry between long range feedforward and feedback connections in the cortical hierarchy (Felleman and Van Essen 1991; Bastos et al. 2012; Nikola T. Markov et al. 2014; Kok et al. 2016). The cell bodies of the neurons that give rise to feedforward corticocortical projections, tend to reside in superficial layers of the cortex while their postsynaptic target is typically located in the middle layers, in particular the granular layer 4 (L4). For feedback connections the source and target neurons tend to reside in deep and outer layers respectively (**Figure 1.1**).

Most intrinsic (short range) connections are intralaminar, providing inputs to neurons within the same layers (Douglas and Martin 2004). Local intralaminar inhibition is a common circuit motif for all layers throughout V1, S1 and M1 (Kätzel et al. 2011). A persistent pattern of interlaminar intrinsic connectivity found in sensory and motor areas includes strong excitatory connections from L4 to L2 (L2/3) and from L3 (L2/3) to L5.

To summarize, cortical features include: much denser local connectivity compared to extrinsic connections, similar innervation patterns shared among feedforward projections and among feedback projections, common neural motif of intrinsic connections. These properties have led to an idea of a canonical microcircuit concept (Douglas and Martin 2004). A canonical microcircuit attempts to capture computational complexity within identical elementary units that collectively can model the function of the cerebral cortex. The cortical column was proposed as a likely candidate for this role (George and Hawkins 2009; Bastos et al. 2012). The model quantitatively summarizing the canonical circuitry of a cortical column is shown in **Figure 1.2** (Haeusler and Maass 2007).

Canonical microcircuitry suggests that local neuronal ensembles perform identical computations on different levels of the hierarchy in sensory cortical areas. Throughout these ensembles each cortical layer might carry out an individual computational role. Since the theory indicates that local ensembles can implement Bayesian inference, L4 and L2/3 might participate in feeding forward the sensory evidence from hierarchically lower areas to update probabilistic estimation of an environmental variable. Feedback projections from higher associative cortical areas tend to target apical dendrites of L2/3 and L5 pyramidal cells in L1, and can communicate the prior expectations about the variable.

In dynamic Bayesian inference, the prior expectation can come from the knowledge of which sensory state the current action has led to in the past. This information could be signaled by cortico-cortical projections from motor to sensory areas. Uncertainties associated

with actions and sensory input can be combined in a Bayesian fashion (Orbán and Wolpert 2011), which was demonstrated in visuo-motor tasks built around the prior knowledge of the task as well as current sensory uncertainty (Körding and Wolpert 2004; Tassinari, Hudson, and Landy 2006). Because the somatosensory system is strongly interconnected with the motor system, an appropriate alternative to investigate the mechanisms of sensory-motor integration might be a behavioral task that is based on somatosensory signals.
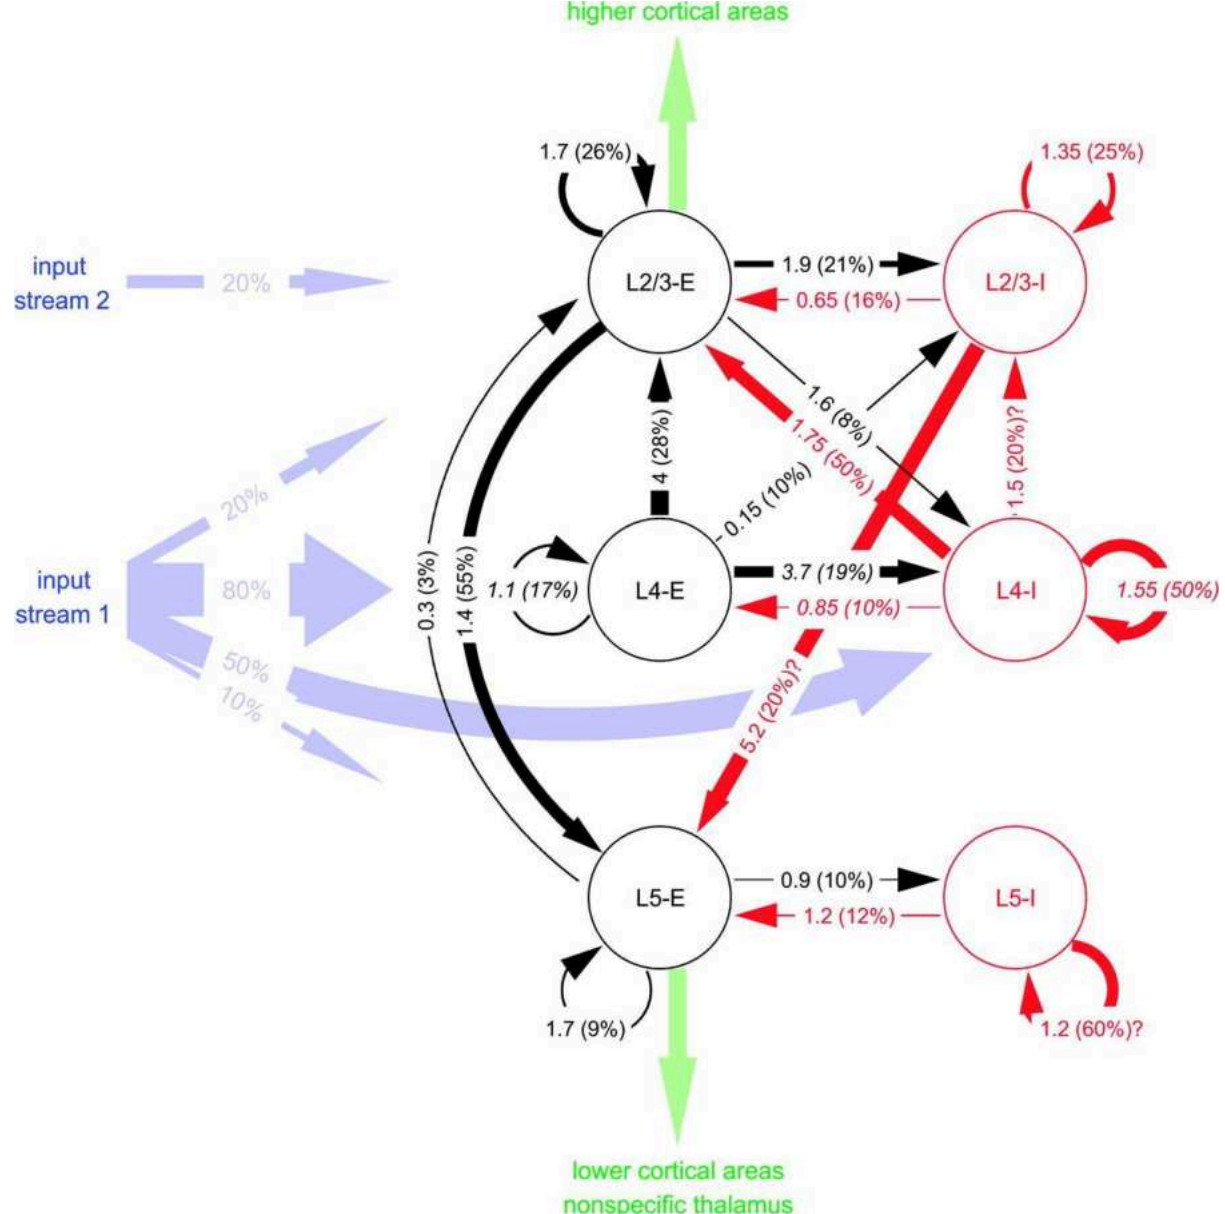
Figure 1.2. Cortical microcircuit template. Numbers at arrows represent connection strength (mean amplitude of postsynaptic potential at soma in mV) and connection probabilities (in parentheses). Connection probability denotes a ratio of pairs of neurons in corresponding layers that share a monosynaptic connection. (I) Inhibitory and (E) excitatory populations of each layer within 100 μm horizontal distance are shown. The blue arrows represent synaptic inputs from distal neurons-feedforward input stream 1 and lateral or top-down stream 2 (Haeusler and Maass 2007).

1.3 Sensory-motor integration

To understand what computations can be expected during integration of sensory and motor signals, I will review here how somatosensory sensation and motor output are formed, and then see how these signals interact.

In generating optimal behavior, knowing one's own body position and forces that act on it, is crucial. Mechanoreceptors serve as the primary source of information for these sensations with muscle primary endings, secondary endings and tendon organs preferentially signaling limb movement, limb position and applied force, respectively (Matthews 1982). Cutaneous receptors, aside from the sense of touch, mostly participate in kinesthesia of fingers in primates, while joint receptors tend to signal only about threshold joint positions.

Somatosensory signals ascend from the first order sensory neurons located in dorsal root ganglion through the second order cell that is located in either dorsal horn or medulla oblongata. The muscle afferents branch in the brainstem, with one branch going to the cerebellum and the other continuing via medulla and the thalamus to the somatosensory and motor cortical areas (Proske and Gandevia 2009). Cerebellum is mostly involved in learning motor skills and coordinating voluntary movements by means of Purkinje cells providing inhibition to vestibular and cerebellar nuclei (Latash 2008). The cerebral cortex, on the other hand, is associated with planning, initiation and execution of voluntary movements (J.-Z. Guo et al. 2015).

In the vibrissal part of the somatosensory system in rodents, the sensory signals about the force applied to the whiskers are transmitted to the barrel part of the primary somatosensory cortex. The signals arrive to L4 of a barrel column that mainly excites L2/3 neurons, which in turn excite L5A and L5B (Mao et al. 2011). The barrel cortex contains cortical columns of densely packed neurons that are somatotopically connected with whiskers.

In both primates and rodents, the topographic property of the primary somatosensory area is characteristic for the tactile afferents that extend from cutaneous mechanoreceptors. Stimulations of limb muscle afferents demonstrated that they also target, sometimes somatotopically, primary somatosensory area in cats (Oscarsson and Rosen 1963; Landgren and Silfvenius 1969), and primates (Landgren and Silfvenius 1969; Phillips, Powell, and Wiesendanger 1971), along with motor and secondary somatosensory cortex.

A number of cortical regions interacting through cortico-cortical connections were shown to be involved to optimize motor control in primates performing grasping tasks (reviewed in Davare et al. 2011). For example, a part of the posterior parietal cortex (PPC), anterior intraparietal area (AIP) contains visuomotor neurons that are activated during a particular type of grasp. AIP makes reciprocal connections with the ventral premotor area (PMv) that is reciprocally connected to M1. Transcranial magnetic stimulation of AIP disrupted the adjustments of the grip force.

To investigate the interaction between sensory and motor cortices in rodents, Mao et al. photostimulated axons of ChR2 expressing neurons of vibrissal S1 (vS1) , and recorded whole-cell postsynaptic activity in vibrissal M1 neurons, in the method called subcellular Channelrhodopsin-2-assisted circuit mapping (Mao et al. 2011). A subset of L2/3 and L5 neurons of the barrel cortex projected to vibrissal M1 where they strongly targeted upper layer neurons in L2/3 and L5A and only weakly deep layers L5B and L6 (**Figure 1.3**). The targeted neurons in L2/3 and L5A of M1 projected strongly back to S1 mainly to the deeper L5B layer. Using the anatomical features of long range connections in sensory cortical areas (**Figure 1.1**) as the criteria for establishing hierarchical relationship among other areas, the motor cortex is considered to occupy a higher level in the cortical hierarchy relatively to somatosensory areas (Felleman and Van Essen 1991).

The motor cortex differs from the sensory areas in that it lacks clearly defined granular layer 4, thus it's called agranular (Weiler et al. 2008). M1 layers also occupy different depths compared to S1 (**Figure 1.3**). During uncaging of glutamate in the mouse motor cortex, the strongest excitatory connection was found to be from L5A to L2/3 and from L2/3 to L5A/B (Weiler et al. 2008).
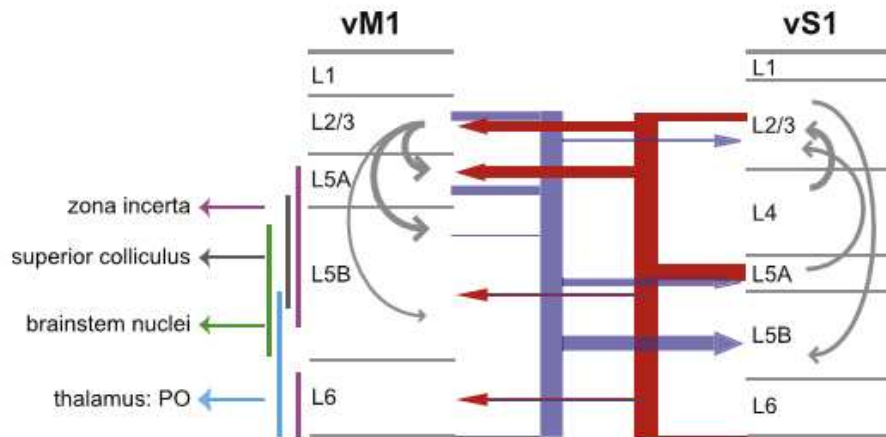
Figure 1.3. Long-range circuits connecting vibrissal S1 and vibrissal M1 (Mao et al. 2011). Line thickness is proportional to the strength of connection.

The motor signals generated in the cortex can be passed via corticospinal fibers through the medullary pyramid at which they decussate. In primates, the corticospinal neurons located in motor cortex can project directly to the lower motor neurons in the anterior lamina of the dorsal horn, and are involved in control of forelimb movement (Evarts 1968; Boudrias et al. 2010; Quallo, Kraskov, and Lemon 2012). In rodents these projections terminate in the base of the dorsal horn and do not directly innervate motor neurons (Watson, Paxinos, and Puelles 2012). Despite the absence of direct projections, involvement of the sensorimotor cortex of mice in initiating and executing a learned movement was shown to be crucial. Cortical suppression of primary sensorimotor (S1 and M1) area blocked prehension movements, but didn't influence naturalistic grooming behavior (J.-Z. Guo et al. 2015). In addition to motor parameters such as exerted force (Evarts 1968), the neurons of primary motor cortex in monkeys were also shown to code for tactile load during control of a handle (Kurtzer, Herter, and Scott 2005).

Sensory input can influence behavior at different levels. The simplest and the fastest mechanism, called a reflex arc, involves an excitatory monosynaptic connection that a muscle efferent can make with an alpha-motoneuron in the spinal cord. Oligosynaptic and polysynaptic reflexes involve more central synapses and can have a more complex effect on the motor output. On top of that, parameters of these reflexes can be modulated with central commands signaling voluntary movements in a similar fashion with the inverse model of motor control (Latash 2008; Forgaard et al. 2015). The timescales of muscle responses to sensory perturbations were revealed with EMG (electromyography) recordings. For example, an abrupt change in the handle position required the monkey to return the handle to the target area to get reward (Evarts 1968). Multiple phases of EMG activity were recorded at around 12, 35 and 80 ms, that involved monosynaptic reflexes as well as cortical sensorimotor responses.

Information about one's actions is crucial for estimating the body position and interaction with the environment. For example, muscle spindle afferents can increase their firing rate in result of both muscle stretch caused from externally applied load, or upon an increase in fusimotor activity of gamma motor neurons. A corollary discharge that sends an efference copy about the exerted effort from motor to sensory centers at the same time as the dispatch of the motor command, can disambiguate muscle spindle signals (Matthews 1982). This is analogous to the forward model of motor control, where the actions are used to predict future sensory state. Different brain areas, including cerebellum and parietal cortex, were shown to be involved in a forward model for multisensory comparison between the feedback during movement and its expected profile, based on past experience (Proske and Gandevia

2009). To understand the role of motor projections from M2 to S1 in mice, Manita et al. used optogenetics in a behavioral somatosensory differentiation task to manipulate connection from M2 to S1 (Manita et al. 2015). Inactivation of the top-down connection from M2 to S1 revealed its functional importance for accurate perception of tactile stimuli. Another clue to the nature of motor involvement in sensation was provided in an experiment involving ischaemic paralysis of an arm, in which a perceived wrist position was shifted from the real position in the direction of the effort exerted by the human subject (Gandevia et al. 2006). This also provided the evidence that corollary discharge alone can cause perceptions even without sensory input from muscle afferents (Gandevia et al. 2006; Proske and Gandevia 2009).

1.4 Neural implementation of probabilities

Traditionally, a neuron's firing rate has been used as a marker of the internal representation of a variable's value. Linear models based on firing rate allowed for a reliable extraction of motor commands parameters in brain-machine interface (BMI) setups (Carmena et al. 2003). However, for the brain to generate behaviors in line with Bayesian inference (Summerfield and Koechlin 2008), it should be capable of representing probabilities and have the faculties to operate them (Pouget et al. 2013).

For explaining this capacity, multiple computational frameworks based on biologically inspired neurons were developed. For example, Deneve presented a model of a single Bayesian neuron modeling a hidden variable that has some temporal statistics. The neuron generates a spike whenever it receives an input that increases certainty in a value of the modeled variable (Deneve 2008). Such a neuron exhibits properties similar to integrate-and-fire neurons, which indicates that individual biological cells can be equipped with the ability to optimally integrate evidence about events. For example, pyramidal neurons in layer 5 (L5) of the mammalian cortex are equipped with the apical dendrites that reach up to layer 1 (L1). Such a dendritic tree receives long range connections and can integrate signals from multiple brain areas(M. E. Larkum et al. 2009; M. Larkum 2013).

Contrary to the model of a single neuron modeling a single variable, it has been shown experimentally that activity of ensembles of central neurons is more representative of a behavioral variable than an activity of individual cells (Nicolelis and Lebedev 2009). The accuracy of motor parameters predicted by brain-machine interface (BMI) was shown to quickly improve with increased number of the analyzed cortical neurons in primates, but grows ever slower when the size of the ensemble reaches 20 or 30 cells (Carmena et al. 2003). Neurons from different cortical areas contributed to encoding parameters differently, which was in line with the notion of functional differentiation of cortical regions. For example, gripping force was better represented in posterior parietal cortex than in supplementary motor area, while for velocity of arm movement the relationship was the opposite. Both motor parameters were best encoded by neurons of M1.

Probabilistic inference based on neuronal networks was proposed in predictive coding models such as the one by Rao and Ballard, in which neurons can compute the most likely properties of a sensory stimulus (Rao and Ballard 1999). Predictive coding framework postulates the hierarchy of neural networks as the main principle allowing for probabilistic inference through the interaction of top-down predictions and bottom-up prediction errors. Prediction error signals have been reported across multiple brain regions, specifically in the mammalian cortex and its sensory areas. Studies observed an increased neuronal activity in response to a prediction violating stimulus. Multiple cortical areas were shown to contain neurons that respond with increased firing rate in response to prediction violation (Meyer and Olson 2011; Todorovic et al. 2011).

Later theoretical developments enabled model neurons to represent not only the most likely (mean) values, but also the degree of certainty, or variance in a Gaussian distribution (Friston 2005). Finally, it has been argued that neuronal populations encode the whole probability distribution and that this distribution can have an arbitrary non-Gaussian shape (Ma et al. 2006). The idea of neuronal populations encoding distributions was called probabilistic population coding. One model used spike-based population coding to implement dynamic Bayesian inference (Boerlin and Denève 2011).

Predictive coding models were initially inspired by the features of the sensory cortical areas. Recently, attempts were made to introduce a common computational basis under both sensory and motor functions, based on anatomical similarity between the pathways of sensory and motor cortical areas (Adams, Shipp, and Friston 2013). This way, the descending motor commands receive a new interpretation, called active inference, which states that motor signals are, in fact, top-down predictions of the causes of proprioceptive sensations. A copy of the prediction about the expected sensory input is sent from motor to somatosensory cortex, where it's integrated with the actual bottom-up sensory input, caused by the movement.

Similarly, based on the concept of the canonical cortical circuit, the motor and sensory cortices were hypothesized to realize analogous computations (Doya 2021), in which each layer of the sensory cortex is associated with a distinct probabilistic value in the Bayesian inference process (**Figure 1.4**).

These theoretical developments suggest mechanisms for how biological neural networks can represent probabilities and implement computations that approximate Bayesian inference. They indicate that ensembles of cortical neurons, in particular pyramidal cells, could be involved in computing probabilistic inference. Experiments showed involvement of multiple cortical areas in this process (Jones et al. 2012; Funamizu, Kuhn, and Doya 2016; Schröder, Schmidt, and Blankenburg 2019) but the exact biological mechanisms remain unclear. Specifically, the theorized computational roles of the different types of neurons and cortical layers have not been directly tested and quantified. To address this question and engage with the numerous existing models, an experiment should characterize functional asymmetries between different populations of neurons in the context of a probabilistic task.

The following research questions can shed new light on the cortical implementation of probabilistic estimations. Do cortical layers respond to the probabilistic task and its variables, and, if yes, is there any asymmetry between deep and superficial layers in terms of which variables they encode? Is there a neural signal evolution during the task in line with the hypothesized probabilistic computations? For example, the signals corresponding to the prior expectations about a sensory variable can be obtained before the sensory stimulus, while the sensory evidence and error detection signals should be available during and after the stimulus. My study aims to quantify these effects and provide new details for developing the theories of how the brain implements probabilistic computations.
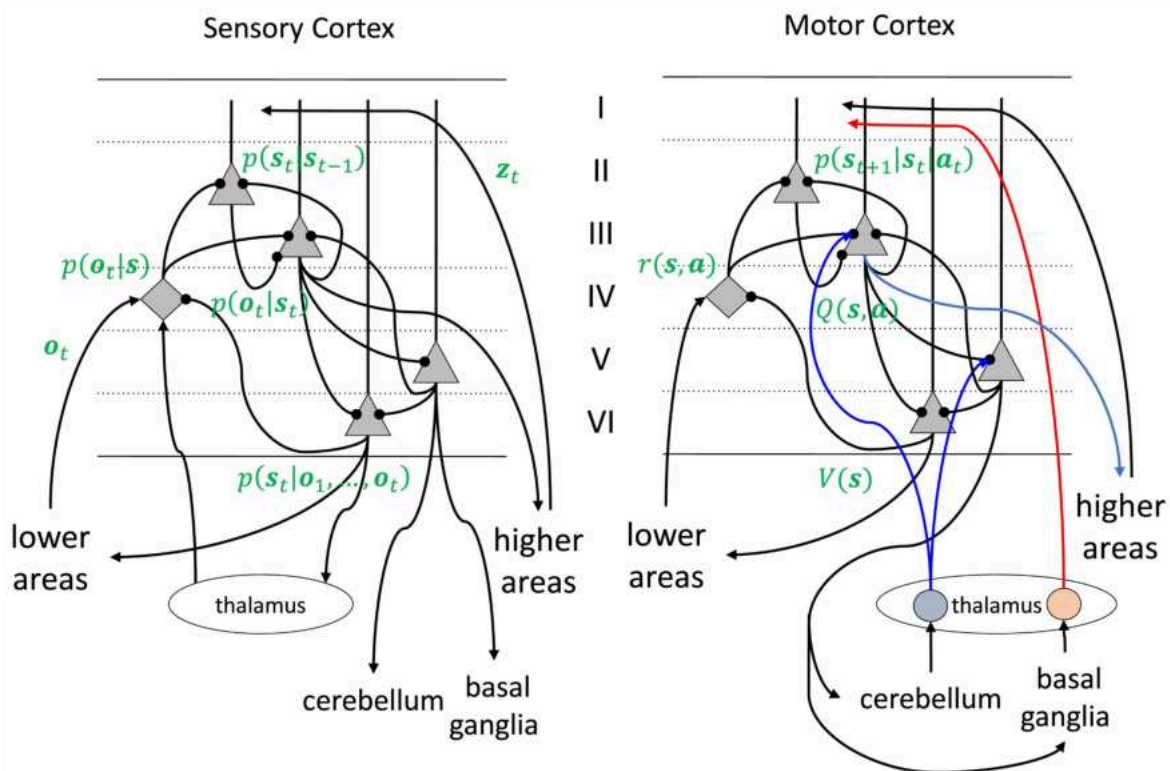
Figure 1.4. **Hypothetical realizations of the dynamic Bayesian inference by cortex**. Left- the model of dynamic Bayesian inference realization by the sensory cortex. Right- the optimal control model realization by the motor cortex.

1.5 Approaches to investigate inference in the cortical sensorimotor loop

Investigation of the cerebral cortex requires experimenting on mammals, with mice being one of the most common model organisms. The capability of associative somatosensory learning was already demonstrated in mice using electrical and mechanical stimulation of a hindpaw (Hirasawa, Yamada, and Murayama 2016), thus suggesting it can be an appropriate model. The most prominent features of the cerebral cortex, discussed in the previous sections, are shared among mammals, including humans. One distinguishing feature of the rodent cortex is the barrel area, a large part of the primary somatosensory cortex. It has distinct structural cortical columns that have receptive fields topographically aligned with whiskers. This makes the research of barrel cortex convenient, but the findings might be more difficult to translate to other mammals. Simple behavioral tasks based on lever control can be a good alternative that enables us to work directly with the sensorimotor loop to investigate whether sensory and motor signals are integrated in a Bayesian manner.

The neurophysiological recordings have to be conducted during the performance of the task to identify the neural mechanisms of Bayesian inference. It has recently been advocated that populations of neurons can encode probabilistic distributions of sensory variables (Beck, Latham, and Pouget 2011; Boerlin and Denève 2011). Thus, simultaneous recording from individual neurons in multiple cortical layers is the appropriate research tool for the project.

In vivo electrophysiology has been widely used for studying the function of the cerebral cortex. It is frequently conducted by using multiple microelectrodes implanted into the region of interest with their tips exposed. This puts the tips in the extracellular space enabling detection of small electric current caused by ionic flow through membranes of nearby neurons. This way, action potentials from multiple neurons are recorded on each

channel, therefore spike sorting is later applied to the recorded data to distinguish the cells based on each cell's spike waveform profile. Tetrodes are attempting to ease this data processing by guaranteeing that four electrodes are close enough that they record electrical activity of the same cells. For example, tetrode recordings from a small population of neurons allowed Erlich et al. to identify the cells whose firing rate predicted later movement, demonstrating preparation and planning skills in mice (Erlich, Bialek, and Brody 2011). While electrophysiology provides superior temporal resolution, identifying cells from one day of experiments to another has been difficult.

One solution to the limitations of electrophysiology is neuroimaging. It provides the spatial information necessary to identify cells throughout arbitrarily long chronic recordings. Voltage-sensitive dyes represent one such exciting perspective and work by changing their fluorescent properties based on the surrounding electrical field. Manita et al. used wide-field cortical voltage-sensitive dye (cVSD) imaging to establish a reciprocal functional connection between S1 and M2 (Manita et al. 2015). By administering sodium channel blocker, tetrodotoxin, they influenced neural response to a tactile stimulation. This allowed them to establish that the top-down extrinsic connection from M2 to S1 is responsible for late activation in S1, which some authors suggest to be responsible for conscious perception of a sensory stimulus. However, it is currently challenging to achieve imaging with voltage sensitive dyes that has both cellular spatial resolution and high temporal resolution in vivo.

Intracellular calcium concentration is a proxy for the spiking activity of neurons. Among other mechanisms, action potentials trigger rapid calcium influx from the extracellular space through voltage-gated calcium channels. The resulting increase in the number of free calcium ions can be further amplified by the release of calcium storages from the endoplasmic reticulum (Grienberger and Konnerth 2012; Ross 2012). Spike-evoked calcium transients last from 20 to 100 ms. Genetically encoded calcium indicators can measure changes in cytosolic free calcium concentration. GCaMP indicators are easier to target to cell populations than synthetic calcium dyes and can potentially detect individual action potentials, including spikes within burst spike trains (Chen et al. 2013). Among them, GCaMP6f has a fast kinetics with the rise time to peak around 50 ms and half decay time less than 200 ms. Adeno-associated viruses (AAV) allow for an effective widespread transduction in a number of cell types of the central nervous system, including cortical neurons and glia (Samaranch et al. 2013; Watakabe et al. 2015). AAV9 capsid serotype and the neuron specific synapsin promoter driving the GCaMP expression allows targeting of pyramidal neurons and interneurons in cortical layer 2/3 and 5 in mammalian cortex. Calcium imaging using genetically encoded indicators such as GCaMP does not provide high temporal resolution to detect individual action potentials, but delivers cellular and subcellular spatial resolution in vivo.

The temporal resolution of calcium indicators puts a limitation on the behavioral protocols, compared to electrophysiological recordings. For example, to find neural correlates of each of the two sensory stimuli, these stimuli should be presented for hundreds of milliseconds at a time with hundreds of milliseconds separation in between. In a lever task for a mouse, this would mean the animal should be exposed to one stimulus throughout the whole pull/push movement. One way to implement this is to have a sequence of trials, each with one sensory stimulus.

In vivo calcium imaging has been used with two-photon microscopy that can image deep layers of the cerebral cortex (up to 1 mm from the surface). Funamizu et al. built an acoustic virtual reality setup and conducted goal-reaching experiment on mice head-fixed under a two-photon microscope (Funamizu, Kuhn, and Doya 2016). They obtained calcium signal from different cortical layers by changing the depth of the imaging plane from one experimental day to another. The study demonstrated that populations of cortical neurons

implement prediction of hidden environmental states based on performed actions and update this prediction by sensory evidence. Limited depths and slow readjustment of the imaging plane limit the usage of two-photon microscopy for investigation of intrinsic cortical interactions between layers.

To overcome the limitations of two-photon microscopy, miniature endoscopic GRIN (gradient refractive index) lenses can be implanted into the neural tissue to arbitrary depth. Straight rod-like GRIN lenses attached to the charge-coupled device (CCD) cameras, allow imaging the deep brain structures with minimal displacement of the brain tissue volume, unachievable with traditional optical microscopy (Yoshizawa, Ito, and Doya 2018). The same lens is typically used to both deliver the laser or LED light from outside the brain to excite the neural fluorescent indicator, and to guide the fluorescent signal to the microscope. For monitoring intrinsic cortical activity, prism microendoscope is a novel and effective research tool allowing recording from multiple layers simultaneously (Resendez et al. 2016; Gulati, Cao, and Otte 2017). For example, Murayama and Larkum conducted calcium imaging by implanting a prism lens attached to a GRIN lens (Murayama and Larkum 2009). They obtained a significantly higher calcium activity in the dendrites of L5 pyramidal cells in the awake state, compared to anesthetized state.

The implantable GRIN lenses have to target a specific area of the brain in order to maximize the number of recorded task-relevant neurons and, therefore, statistical significance of the experimental results. Based on the anatomical and functional data reviewed above, the forelimb area of S1 cortex is suitable for finding neurons that encode tactile variables of a lever task.

To summarize, a probabilistic lever control behavioral task on mice with microendoscopic GCaMP6f calcium imaging of the somatosensory cortex can be an effective tool to investigate Bayesian sensory-motor integration.

Chapter II. Methods

2.1 Behavioral setup

The head-fixed setup was developed to include a custom-made 3D-printed mouse body holder with an armrest, the mechanism for fixating the head frame, a water spout, and a lever. The mouse body was placed in the body holder with its left arm on the armrest and the right arm controlling the lever. For neural imaging, the microscope was plugged into the baseplate at the top of the mouse head (**Figure 2.1**). The setup was equipped with an infrared licking sensor and an infrared camera for overall observation of the animals' behavior.
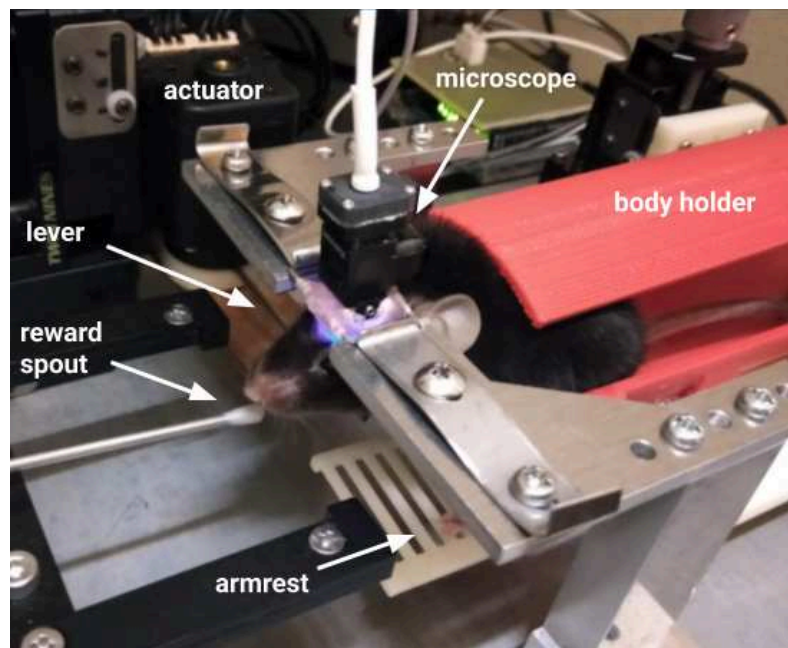
Figure 2.1. Head-fixed experimental setup for lever-pulling task

The lever contained a stainless steel handle that was placed close to the animal's right forelimb, and a 3D-printed flexible part that allowed elastic deformations of the lever (**Figure 2.2A**). Before experiments, the lever was calibrated to map the output voltage of the magnetic sensor to the applied force. The hall-effect magnetic sensor provides a linear relationship between the magnetic field intensity and the output voltage, within a certain range of magnetic field intensity. I adjusted the lever thickness, length, and the magnet distance from the sensor such that the forces approximately ±200 mN remained within that range.

To control all lever movements, I used a rotational actuator (Dynamixel AX-12A) that was mechanically attached to both the flexible lever and a custom-made circuit board. The lever hosted a permanent magnet placed close to the hall-effect analog magnetic sensor located on the circuit board to measure its deflection. The output voltage signal of the magnetic sensor was communicated through the Arduino chip to the PC that controlled the motor. This way, the setup could amplify weak force applied to the lever by generating motor torque with a variable amplification.

Figure 2.2. Force-sensing lever. **(A)** The diagram representing main components of the force-sensing lever. The magnet is attached to the lever and is suspended over the magnetic sensor, all of which mounted on a rotatable lever platform. Force applied horizontally to the handle (Arrow 1) causes elastic bending of the flexible part of the lever and a small misalignment between the magnet and the magnetic sensor. The measured magnetic field change is used to generate motor torque that rotates the lever platform in the direction of the applied force (Arrow 2). **(B)** Signal flow diagram.

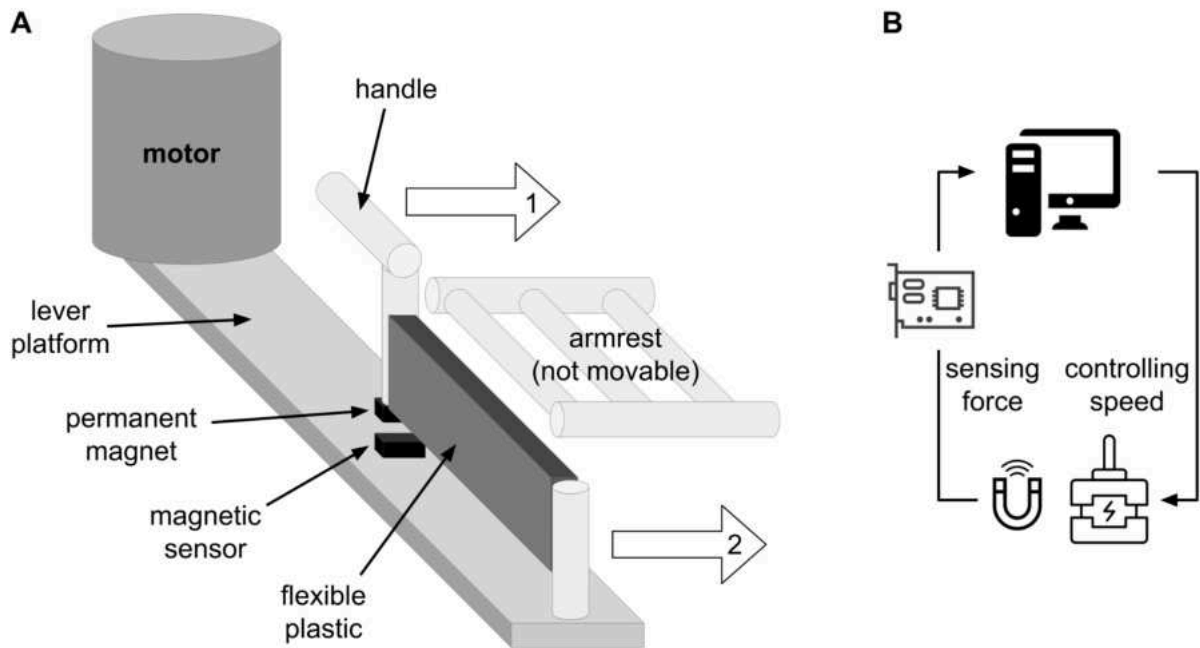I programmed the lever to maintain a linear relationship between applied force and speed (**Formula 2**). The linear coefficient for each trial type depended on the force ("resistance") required to overcome to move the lever to the reward position within around 1 s (**Formula 3**). The typical values for the resistance force for light/heavy trials were 8/16 mN.

$$S = F \star Coefficient_{TRIAL} \quad (2)$$

$$Coefficient_{TRIAL} = \frac{Reward\ Position - Initial\ Position}{Resistance \times 1\,s} \quad (3)$$

where: F - applied force

S - lever speed

$Coefficient_{TRIAL}$ - amplification coefficient

$Resistance_{TRIAL}$ - resistance

The typical reward lever position was set to correspond to the handle displacement from the initial position by around 8 mm. Upon reaching this threshold distance, the motor would stop rotating and, thus, the lever would get fixed in place, such that no force applied to it in either direction would result in the lever movement. The hardware and software of the system had a small latency, with each cycle of measurements and the motor control being 6.5 msec on average. This enabled a seemingly seamless operation of the lever.

2.2 Neural imaging preparations

All animal procedures were conducted in accordance with guidelines of the Okinawa Institute of Science and Technology. I used eight C57BL/6J male mice, from 2 to 3 months of age. The neural imaging preparation procedures included a virus injection surgery, a prism probe implantation surgery, and a baseplate attachment, and were done as described before (Gulati, Cao, and Otte 2017). For surgical procedures and the baseplate attachment

procedure, the animals were anesthetized with isoflurane (3% for induction and 1-2% for maintenance) and placed on the stereotactic device to immobilize the head. Eye ointment was applied to prevent them from drying and intense lights during procedures. After surgeries, carprofen (Rodent MD's Rimadyl, 5 g per day) was given to the animals as food for the duration of a week. The lens implantation surgery was done 1-2 weeks after the virus injection surgery. The baseplate attachment procedure was done 1-2 weeks following the lens implantation.

During surgeries, the scalp was trimmed between the ears and eyes and disinfected with iodine. After making the skin incision and opening the skin, the skull surface was cleaned and polished with cotton swabs. The coordinates for the virus injection were determined to be the center of the S1 forelimb area of the left hemisphere, 0.3 mm anterior and 2 mm lateral from the bregma (Paxinos and Franklin 2019). The craniotomy was performed by drilling a 2-mm hole in the skull over the injection site. For labeling the cortical neurons with calcium indicator, I injected the viral solution of AAV9.Syn.GCaMP6f.WPRE.SV40 with a glass capillary tube. 500-1000 nl was delivered across two to three sites, for 10-30 minutes at each site, starting from the ventral-most position of about 850 μm from the brain surface and gradually retracting the capillary dorsally.

The 1-mm wide prism lens, 4.3 mm length and 0.36 NA (PPL-1043, Inscopix Inc.) was implanted 0.2 mm medially from the injection site with its flat edge aligned with the sagittal plane and facing the injection site (**Figure 2.3**). Before implantation, an insertion tract of 1 mm wide and 1 mm depth for the probe's sharp edge was made sagittally in the brain using the incision knife mounted to the stereotaxic manipulator. After the lens insertion, I covered the exposed brain tissue around the lens with elastomer adhesive. After the adhesive was cured, I fixed the probe to the skull surface surrounding the cranial window with ultraviolet cyanoacrylate adhesive.

At least 1 week following lens implantation, I installed the baseplate in which the microscope gets plugged in for imaging. The prism lens produces a vertical flip and the ½ pitch introduces 180 degrees rotation, while the microscope camera records a cropped image 1280 ✕ 800 pixels. Taking this into account, the baseplate was oriented so that the longest side of the image corresponded to the dorsal-ventral axis. During baseplate installation, the distance from the implanted lens was adjusted such that the larger number of neural cells remained in focus. For head fixation during the experiment, a squared chamber frame was placed over the implanted probe and fixed with dental acrylic.

During experiments, the blue excitation light of 475±7 nm wavelength and power 0.4-1.8 mW/mm^2 was used to excite the neural tissue for fluorescent imaging. Images were recorded 20 frames per second at the absorption peak at 535±25 nm corresponding to green light. For each mouse, the focal plane was electronically adjusted using Inscopix Acquisition Software within the range 100-300 μm to include the largest number of visible cells.
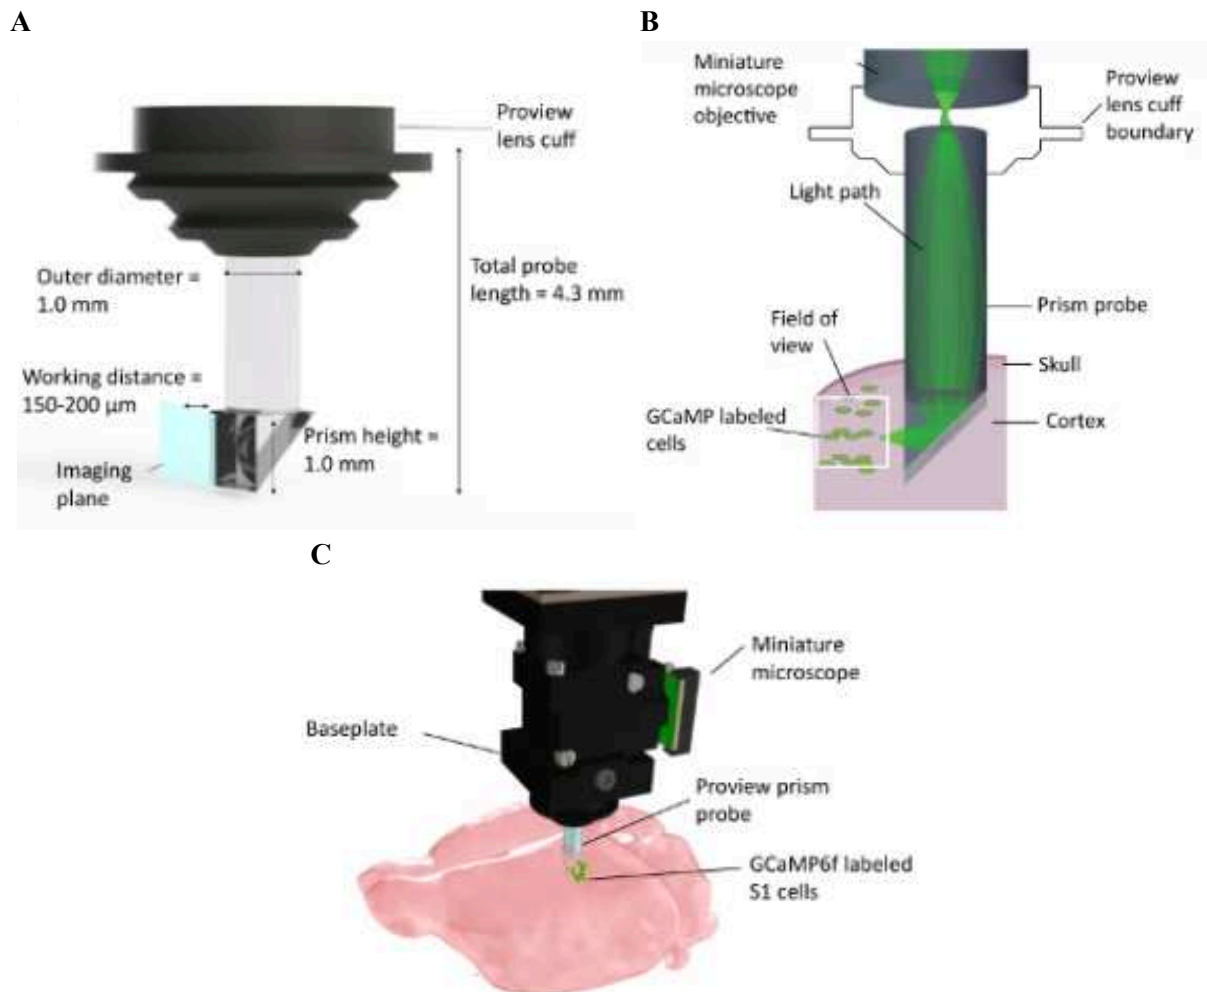
A

B

C

Figure 2.3. *In vivo* calcium imaging across multiple cortical layers with prism lens. (A) Prism probe specifications and depiction of the imaging plane. The reflective coating on the inside of the hypotenuse allows for imaging 90° from the insertion plane of the prism probe. **(B)** Illustration of in-vivo calcium imaging setup showing the light path for a small area within the full field of view through prism probe implanted in mouse cortex. **(C)** Integration of microscope with prism probe for multi-layer cortical imaging of GCaMP6f labeled S1 cells (Gulati et al. 2017).

2.3 Behavioral training

After recovery from the surgery (typically 1 week), habituation and behavioral training was started. In all mice, there were no visible alterations in gate or other functional deficiencies that could potentially be caused by the nervous tissue damage from the implant and the surgery. Before habituation, training, and recording sessions, the animals were placed on water restriction protocol to maintain their weight within 75-80% range of the initial weight (Z. V. Guo et al. 2014). In my experience, a weight higher than 80% decreased animals' motivation.

Each habituation session took 15-20 minutes once a day during around 14 days. During the entire habituation session, the mouse was given water using a syringe or the water spout. During the first few days the mouse was held on the hand during the entire habituation session. Once the mouse seemed relaxed, as indicated by grooming, I placed the mouse on the experimental setup without restraining. Next, I would head-fix the mouse for a progressively longer time, while giving water from the spout (**Figure 2.4**). Habituation would be considered completed once the animal could remain head-fixed on the setup for at least 10 minutes without resisting. The water spout was placed in front of the mouse in the head-fixed

position to deliver 2.5 µL droplets of tap water per pump. When triggered, the reward pump produced a distinct sound that could be heard by the mouse. By the end of habituation, all mice learned to lick after hearing the pump sound.
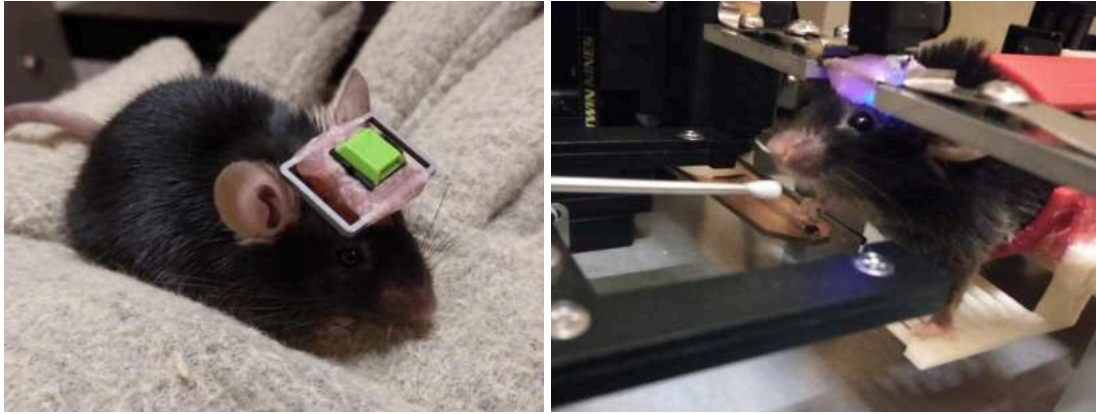

Figure 2.4. Mice habituation and training. Left- habituation to hands. Right- task training.

To train lever-pulling behavior, I provided reward manually whenever the mouse placed the arm on the lever. After the mouse kept tactile contact with the lever, I ran the experimental session with the lever resistance and the threshold distance for reward delivery set to their smallest values. During the training, I gradually increased the lever resistance, the threshold distance, while reducing the number of manual rewards.

Before each trial, the lever is placed in the initial position and is not responsive to the force acting on it. Each trial started with the sound cue that consisted of consecutive 1 s tone of 3 kHz, followed by a 0.25 s tone of 3.5 kHz (**Figure 2.6A**). Immediately after the sound cue was finished, the force measurements from the force sensor translated into the motor's torque, so the mouse could pull the lever. If the lever reached the threshold position, the pump delivered 3 x 2.5 µL of water over 1 s time period. If within 8 seconds the mouse did not sufficiently pull the lever, the omission sound pulse of 1.5 kHz was produced for 0.4 seconds. After the end of the trial, the lever remained fixed in its position for 2.5 s after which it returned to the initial position. The intertrial interval was randomly generated with a uniform distribution from 2.5 to 6 seconds.

Approximately one week after starting the training, all mice could perform the task with around a hundred hits within one session without manual facilitation of the reward (**Figure 2.5**), after which I started the experiments.
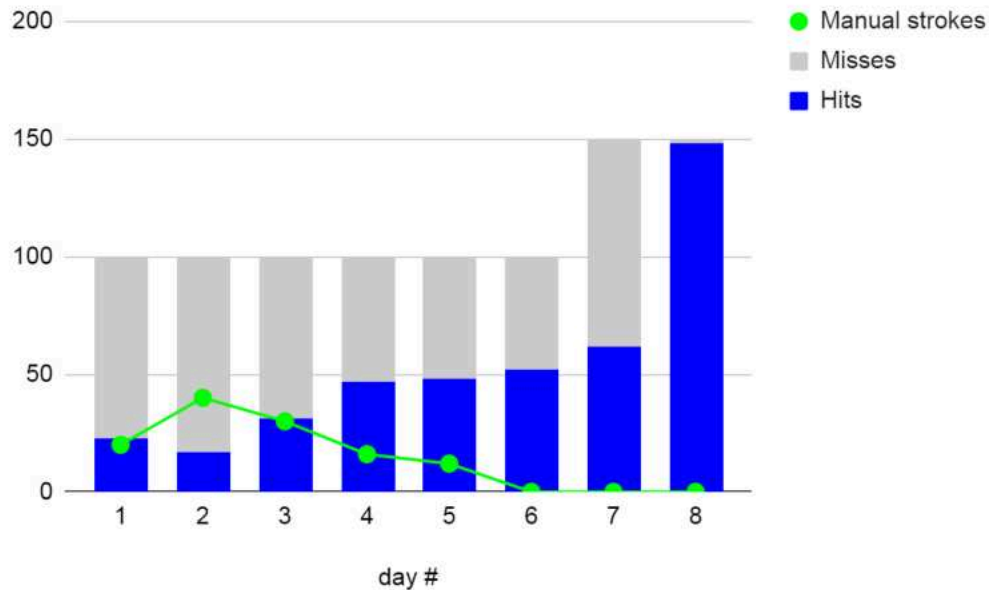
Figure 2.5. Sample task performance of a single mouse. Within 8 days, all mice were capable to reach sufficient performance of around 100 hits.

There was one session conducted on each mouse in one day. During the session, mice typically licked around 1 ml of water from the reward spout and satisfy its thirst during the task. The mouse would be motivated to do the task again the next day. If the mouse's performance was low or its weight decreased below 75%, I manually provided an extra amount of water 0.5-1 ml.

The sessions consisted of sequences of trials with small and high resistances, i.e. light and heavy trials. In one type of session, a probabilistic session, each resistance level occurred with a certain probability P determined before the session, with the other resistance occurring with probability 1 - P. There were three types of probabilistic sessions: light, heavy, and uniform. The value P for light (heavy) trials was set as 0.9, 0.85, 0.8 for light (heavy) sessions and 0.5 for uniform sessions. I also ran several roving oddball sessions, where the trains of one trial type alternated with the trains of another type. Each train had a random length of 5-15 trials (**Figure 2.6B**).

Each session consisted of 150 trials which took around 20 minutes to complete. I manually prematurely aborted some sessions if the mouse looked no longer motivated to perform the task to get water reward. For behavioral analysis, I collected the recordings of the lever position, applied force, licking, and reward pump (**Figure 2.7**).
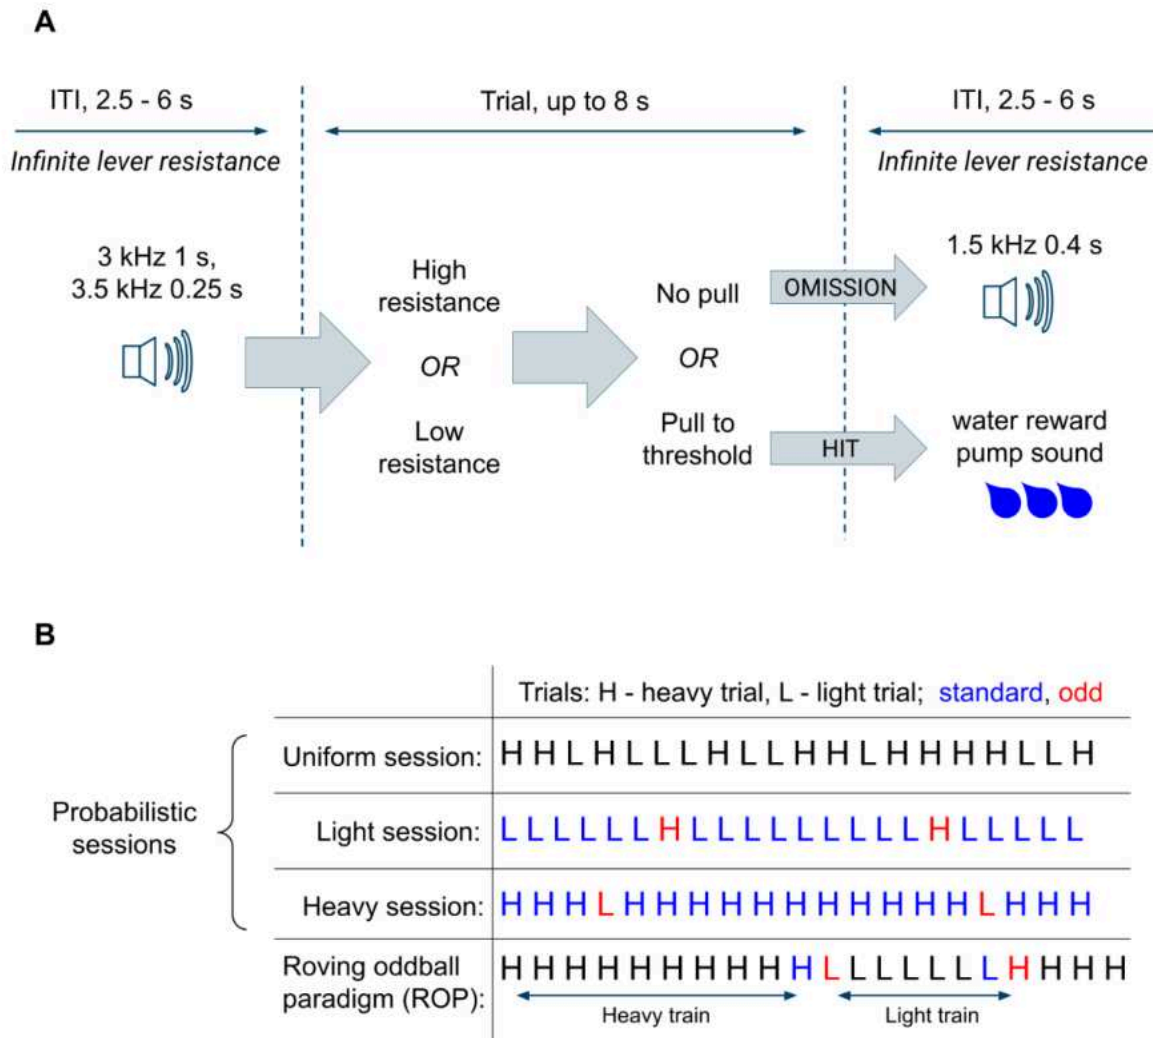
Figure 2.6. Lever-pulling task. (A) Single trial diagram. During Intertrial intervals (ITI) the lever is unresponsive to the applied force. At each trial the lever becomes responsive to force and its resistance is set to either high or low, i.e. heavy (H) or light (L) trials, respectively. If the mouse pulls the lever to the threshold distance, reward is provided, otherwise the trial is considered an omission. **(B)** Session types are determined as the sequence of trials of different types (one example sequence per session type is shown).
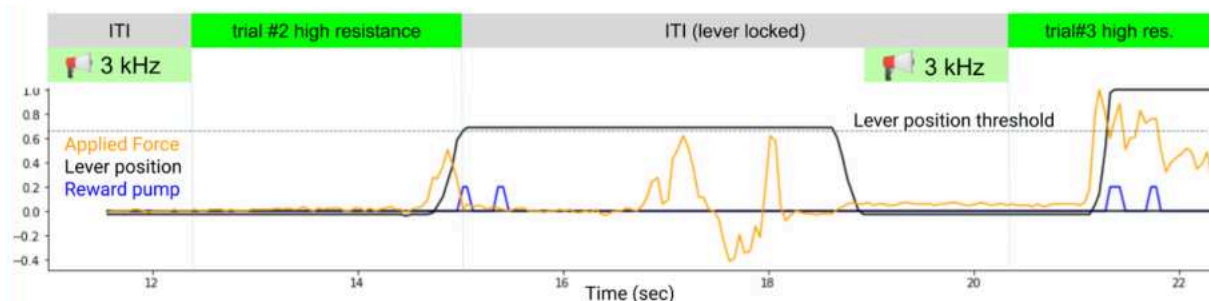
Figure 2.7. Behavioral recording. The normalized traces of lever are shown: position (black), applied force (orange), and reward pump (blue). During the trial #2 and trial #3, the applied force quickly caused the lever displacement and the following water reward, while during ITI the lever remained fixed in place.

2.4 Histology

After experiments, I anesthetized the mice with lidocaine and transcardially perfused them with 4% paraformaldehyde in PBS. During the brain separation from the cranial bone, the lens implant remained attached to the surface of the bone with its lower end protruding 1 mm into the brain. I carefully extracted the brain avoiding any damage the lower end of the lens could cause to the tissue. The coronal sections were cut at 120 μm or 60 μm thickness with a vibratome, immunolabeled with DAPI, mounted on glass plates, and imaged with a confocal fluorescent microscope (LSM780, Zeiss).

2.5 Behavioral analysis

The data within each session were analyzed around the time of pull onset. The time of pull onset was determined as the time when the applied force value exceeded 5 mN during a trial and if it led to the successful trial completion within < 2.5 s. In some trials, mice kept pulling the lever before the trial onset and such premature pull trials were excluded from analysis. The 5 mN force threshold also allowed to filter out task-irrelevant forces caused by mouse leaning on the lever as well as minimize false readings influenced by the noise of the force sensor and other nonlinearities in the lever. The recorded behavioral and neural traces were analyzed within the time window of ±2.5 s around the pull onset.

2.6 Neural imaging

I processed the recorded neural imaging videos using Inscopix Data Processing Software. The videos were motion-corrected and normalized to obtain dF/F signal before applying PCA-ICA algorithm to identify components corresponding to individual cells. For each mouse, the video data from different days were then aligned and the cells present across days were matched.

2.7 Encoding analysis

For every neuron, all recorded traces dF/F were normalized and aligned by the pull initiation time. I quantified the neural responses to task-relevant variables at different time windows around the pull onset. These quantified response features included the averaged and peak calcium signal values of the neuron before pull start, during the pulling behavior and after pull start. I defined a neuron as task variable coding if at least one feature of the neural response significantly changed with the variable. For binary explanatory variables, such as trial type, I conducted unequal variances t-test for independent samples ($p<0.05$). For numeric explanatory variables, such as resistance and probability of a stimulus, I estimated the Pearson linear correlation coefficient and conducted a t-test on it ($p<0.05$).

I used the number of task variable coding cells and overall number of recorded cells in each layer to test whether each task variable is better represented by deep or superficial layer cells (two independent sample z-test for proportions).

To identify any functional asymmetry among layers, I compared the layered distribution of cells responding to one task variable with the layered distribution of cells responding to another task variable (chi-square test for independence). I centered the analysis around the comparison between representations of couples of variables: (1) expected trial type and the actual trial type, (2) expected resistance level and the actual resistance level. In addition, I analyzed neural representations of the following probabilistic variables of the task: (1) the ratio of the main trial type within the session, which can be interpreted as certainty

with the lowest value being 0.5 within uniform sessions, (2) the likelihood of the current trial type as defined by its ratio within the session.

To further test for the presence of a functional asymmetry among layers, I carried out a two-way ANOVA test with 2x2 groups to select the ensembles of cells that responded to the following factors: actual trial type, expected trial type, the interaction between actual and expected trial types. For these cells, I also tested whether there is any relation of task variables with the layers using test for proportions and test for independence.

2.8 Decoding analysis

To further investigate how task-relevant variables are represented by each population of neurons, I conducted a regression analysis with regularization to characterize the ability to decode each variable using the neural ensembles of cortical layers. I conducted logistic regression for modeling binary variables and linear regression for numeric variables. I performed leave-one-out cross-validation by consecutively excluding the data of each trial, training the model on the rest of the trials, and predicting the variable at the excluded trial. During regression analysis, I gradually modified the strength of regularization to arrive at an ensemble of neurons that produced the best prediction, as was measured by the R^2 coefficient of determination for numeric predicted variables and the accuracy score for binary variables.

$$Accuracy = \frac{Correct\ Predictions}{Total\ Predictions} \quad (4)$$

For the mice in which both deep and superficial layers were observed, I conducted comparative analysis of the prediction quality for the important task variables.

To further contrast the computational roles of the deep and superficial cortical layers, I analyzed their prediction performance during the odd trials, when the lever resistance was unexpectedly too high or too low. I trained the regularized logistic regression models on the standard light and heavy trials and used them to predict the trial type. I compared how often each layer-specific neural decoder produced a prediction biased towards the actual trial type or the expected trial type. I conducted this analysis for different ensemble sizes, different timesteps of the pulling task, and separately for probabilistic and ROP sessions.

Chapter III. Results

3.1 Behavioral analysis

Overall, I recorded over 17,721 trials from 122 sessions conducted on 8 mice. 13,287 trials were successful hits.

I aimed to investigate whether the mouse behavior changed with the task conditions. The task design did not require the force applied to the lever to be differential across session types and across trial types. If the force was large enough, the lever could be moved to the reward position in both heavy and light trials. However, strong force on light trials can come at a cost of unnecessary effort and fatigue.

Mice exhibited different strategies in the lever-pulling task which can be roughly classified into two categories: premature pulling and timely pulling. I identified a pull as premature when the mouse started pulling before the trial start, when the lever was still locked (**Figure 3.1**). Such trials were frequently associated with behaviors indiscriminate of the task, such as static positioning of the body or resisting head fixation. After filtering out premature pulls, there were 4,254 trials left, on which the mice pulled properly after the trial start, and they were used for further behavioral and neural analysis. About half of these trials were completed within the first 0.2 s after pull onset (**Figure 3.2**).

The comparison of behavioral variables across conditions revealed that the mice applied stronger force in heavy trials than in light trials. The expected trial type also influenced the mouse behavior. First, the mice started licking earlier, even before pulling, if the expected trial was light. In addition, there was a statistically significant increase in force when mice expected heavy trials (**Figure 3.3**). Notably, this motor effect was found during later stages of pulling (after +0.375 s of pull onset) and mostly after successfully completing the trial and receiving reward.
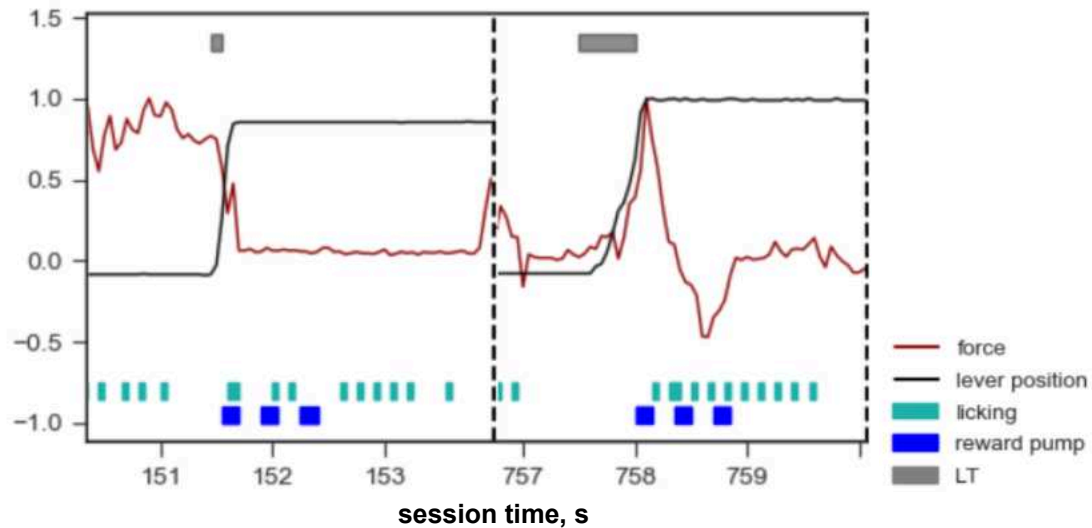
Figure 3.1. Examples of two distinct strategies for completing the lever-pulling task. Left-premature pull, the mouse applied force before the trial start and the force decreased when the lever moved away. Right- proper pull, the pulling behavior is initiated after trial start. Both pulling fragments are from light trials (LT, onset and offset illustrated with gray bars at the top) from the uniform session (mouse #6). Vertical axis represents normalized values of force and lever position.
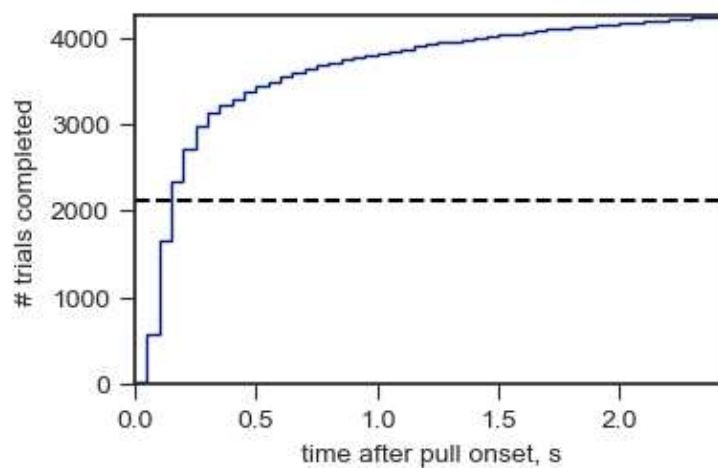
Figure 3.2. Trial completion after pull onset. A half of all trials (horizontal dashed line) were finished within the first 0.2 s after pull initiation.

A

B

C   **D**
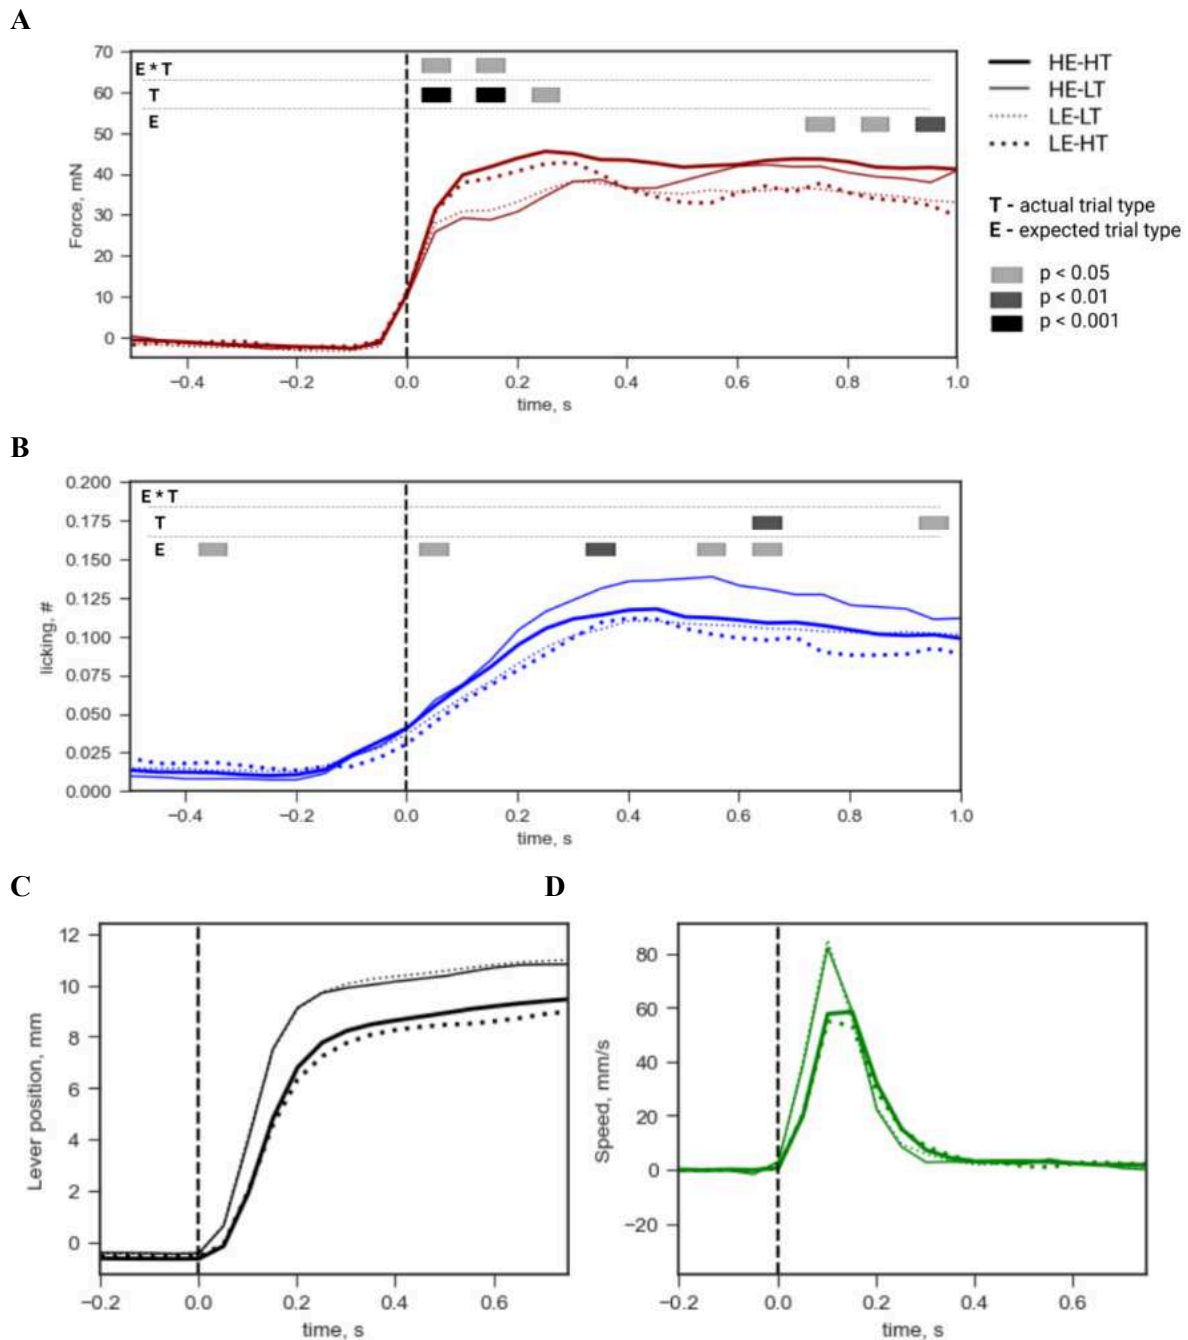
Figure 3.3. Lever-pulling behavior in probabilistic sessions. (A) Averaged force profiles across experiment conditions (0 s timestep is the moment of pull start). The mice applied stronger force in heavy trials (HT) early after pull onset (the trace was binned by 0.1 s time windows and tested with two-way ANOVA). When expecting heavy trials the mice also applied stronger force but during later stages of the pull. **(B)** Licking. The mice licked prematurely more frequently when expecting light trial (LE) than when expecting heavy trial (HE) (at 0.35 s before pull onset, two-way ANOVA). **(C, D)** During the light trials the lever was moved faster with a slight overshooting compared to the heavy trials.

3.2 Neural data analysis

3.2.1 Calcium imaging of cortical layers

I imaged neuronal activity across cortical layers of somatosensory forelimb area (S1FL) using genetically encoded calcium indicator GCaMP6f. The resulted field of view of the imaged tissue was around 1188 $\times$ 750 µm with the long side oriented parallel to dorsal-ventral axis. Collectively, using PCA-ICA analysis on the calcium signal recordings, I extracted 3107 and 1085 components corresponding to deep and superficial individual neurons, respectively (**Figure 3.4)**. The imaged cells were sorted into deep and superficial in coronal slices immunostained with DAPI by visually estimating the relevant position of the unlabeled layer 4 and the track of the lens implant (**Figure 3.5**). Immunostaining and layer identification for one mouse (#8) could not be performed due to premature death and I excluded this mouse from the analysis that required comparison between layers.
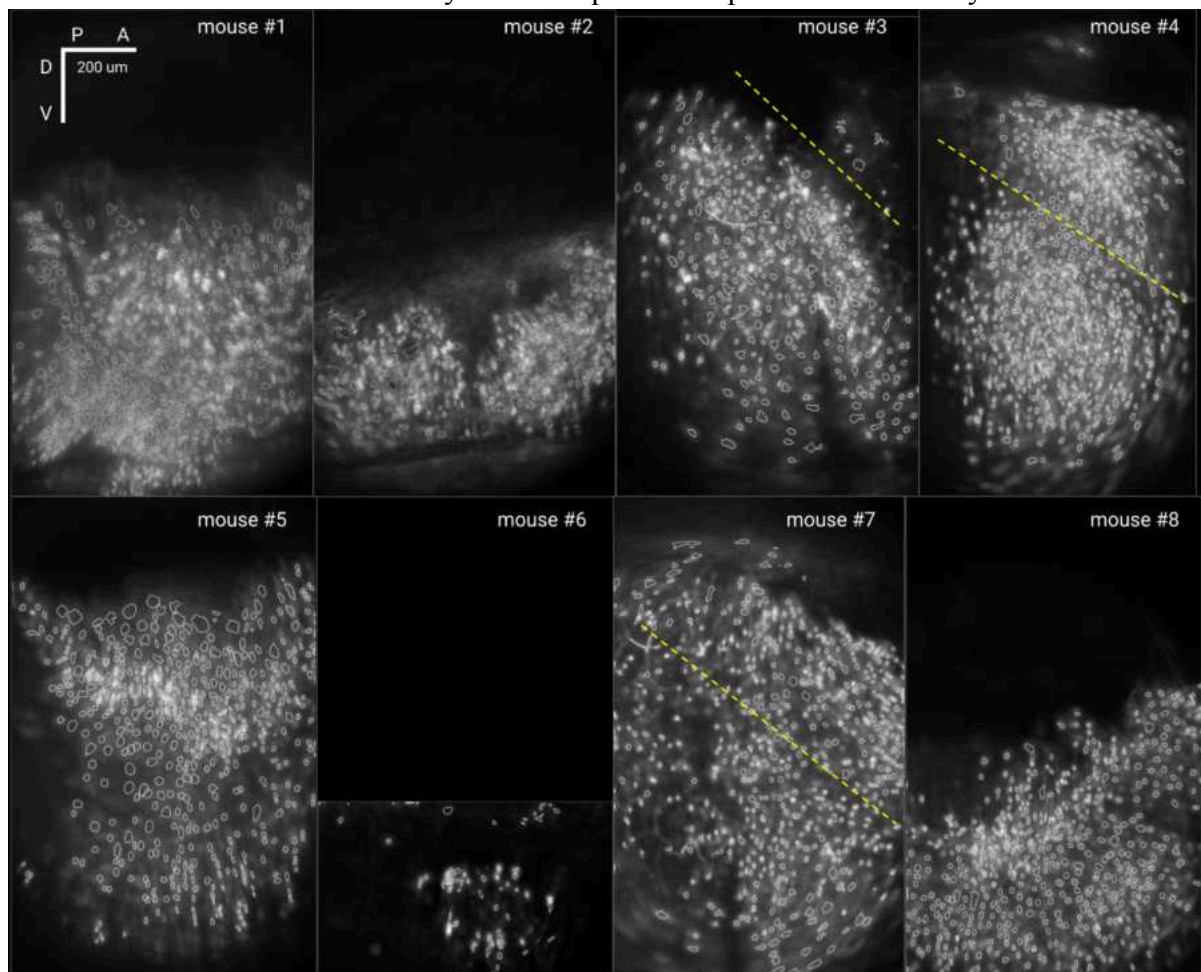

Figure 3.4. Results of PCA-ICA calcium imaging analysis. Data from 8 mice are shown (scale bar-200 µm). From each mouse, except one, from 150 to 550 neurons were extracted (highlighted with white contour). For mice where both superficial and deep layers could be observed, the yellow dashed line indicates the boundary between layers.
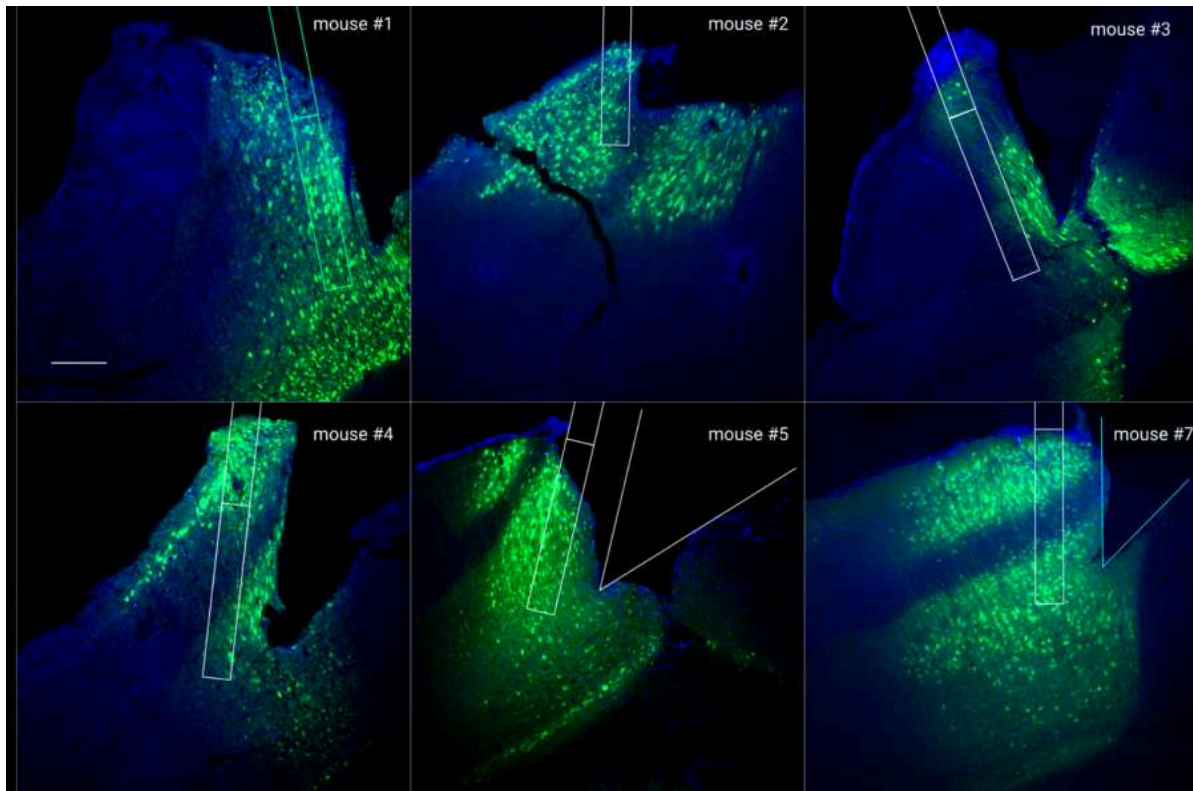
Figure 3.5. Estimation of the imaged layers. Location of the layer 4 unlabeled with GCaMP6f (green color) can be visually identified. The brain slices were made coronally, perpendicular to the plane of the recorded calcium signal videos in **Figure 3.4**. Blue color- DAPI. The data from 6 different mice are shown. The white rectangle depicts the estimated location of the focal plane of the prism lens and its ventral-dorsal midpoint. Scale bar - 200 microns. Mice #3, #4, and #7 have both superficial and deep layers within the focal view. Mice #1, and #5- only deep layers. Mouse #2- only superficial layer.

3.2.2 Neural response to the task

To examine if the neurons respond to the lever-pulling task, I aligned the calcium activity traces by the pull initiation moment (applied force > 5 mN). I found robust neuronal responses to the lever-pulling task in all mice. Activities of multiple cells were correlated with the applied force (**Figure 3.6, Table 1**).

Figure 3.7, 3.8A, 3.8B and **Appendix 1** visualize the activities of the neurons in four different conditions (Heavy/Light Expected Heavy/Light Trials). In both deep and superficial layers, there was an increased number of neurons that were active within proximity to the time of pull. In addition, the cells tended to maintain their peak time across different days and the experiment conditions. Cells that were more active before pull onset tended to have a wider and less stable time tuning compared to the cells with the peak activity during and after the pull onset. This was especially evident for the deep layers. During probabilistic sessions, the neuronal tuning was more perturbed during odd conditions (**Figure 3.7, Figure 3.8A, Appendix 1A**). This is mainly due to the smaller number of odd trials than standard trials. During ROP sessions, the number of odd and standard trials is nearly the same and the perturbations in neuron tuning across task conditions were similar (**Figure 3.8B**).

I quantified *dF/F* calcium signals of the neurons and estimated the change in neural response across behavioral conditions, such as expected and actual trial types. I averaged calcium activity across neurons of each layer in different experiment conditions to check for different responses of the layers (**Figure 3.9**). Superficial cells were more active than deep cells after pull onset, especially in light expected light trials and heavy expected heavy trials.

In addition, deep cells were more active than superficial cells before pull onset when expecting heavy trials.
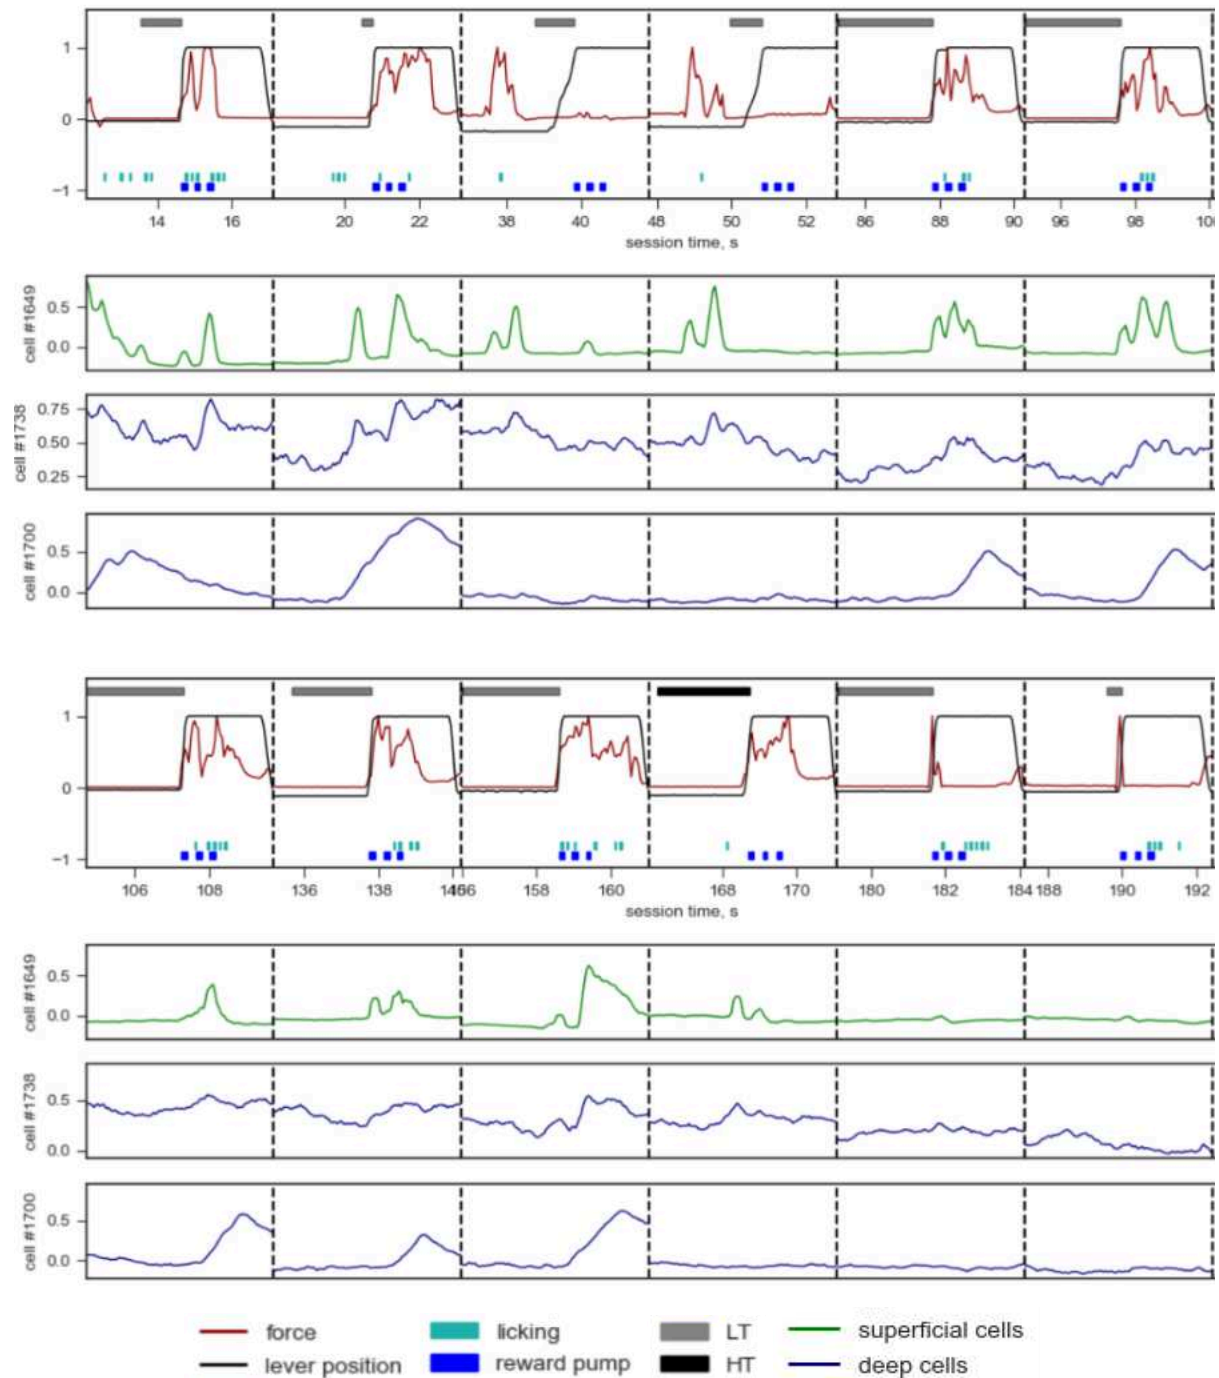
Figure 3.6. Behavioral traces and calcium traces correlated with the applied force. The neural activity of three example cells are shown. Behavioral and calcium activity traces were extracted from a single light session (mouse #7), trimmed at ±2.5 s around the pull onsets and normalized. Force-correlated neurons were identified by conducting a t-test on Pearson linear correlation coefficient (p<0.001). Light gray and dark gray bars represent the duration of light and heavy trials, respectively. Vertical axis represent normalized values of the applied force and the lever position.
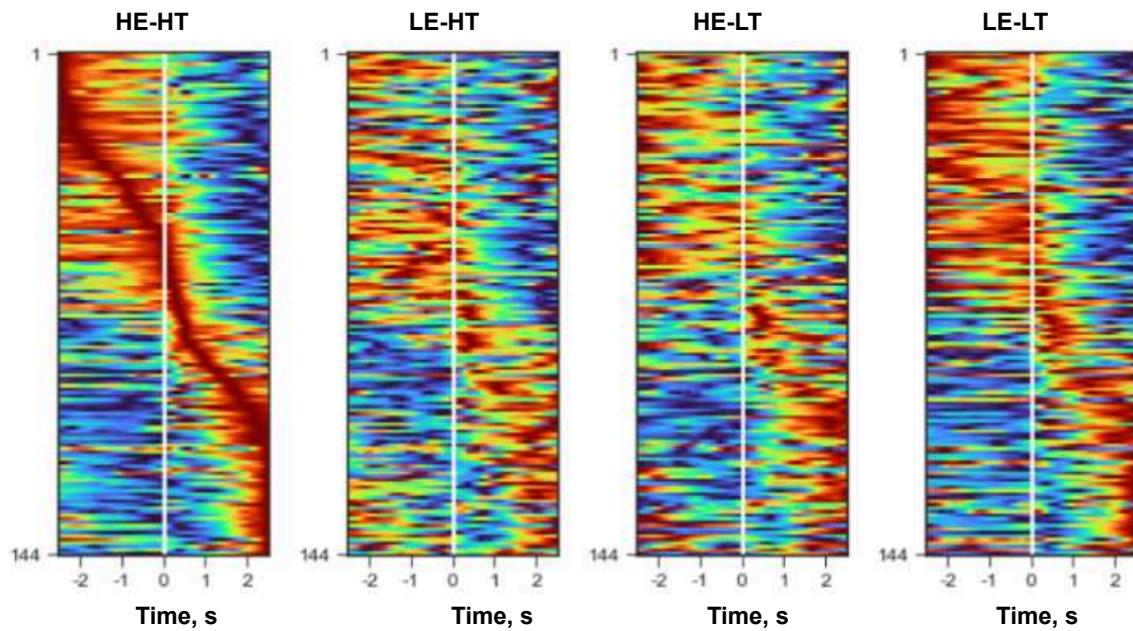
Figure 3.7. Pull-relevant neural activity in deep neurons. Activity traces of neurons are shown, aligned by the time of pull onset (at t=0 s). Each panel represents a single experiment condition, across which the data was time aligned, trial averaged, time-averaged with a moving window of 200 ms, and normalized. Each horizontal line corresponds to a single cell. Within all panels, cells are sorted by the time of their peak activity during heavy trials in heavy sessions (left panel). The cells roughly maintained their time tuning to the lever-pulling task across light and heavy trials (LT, HT) and light and heavy sessions (LE, HE). Deep neurons of mouse #3 during probabilistic sessions are shown.

Figure 3.8A. Pull-relevant neural activity during probabilistic session. Same as **Figure 3.7** for mouse #7 with both deep cells (bottom) and superficial cells (top).

Figure 3.8B. Pull-relevant neural activity during ROP sessions. Same as **Figure 3.7** for mouse #5 with only deep layers.

Figure 3.9. Layer-specific neuronal response to the lever-pulling task. Averaged activity traces across all neurons of all mice are shown.

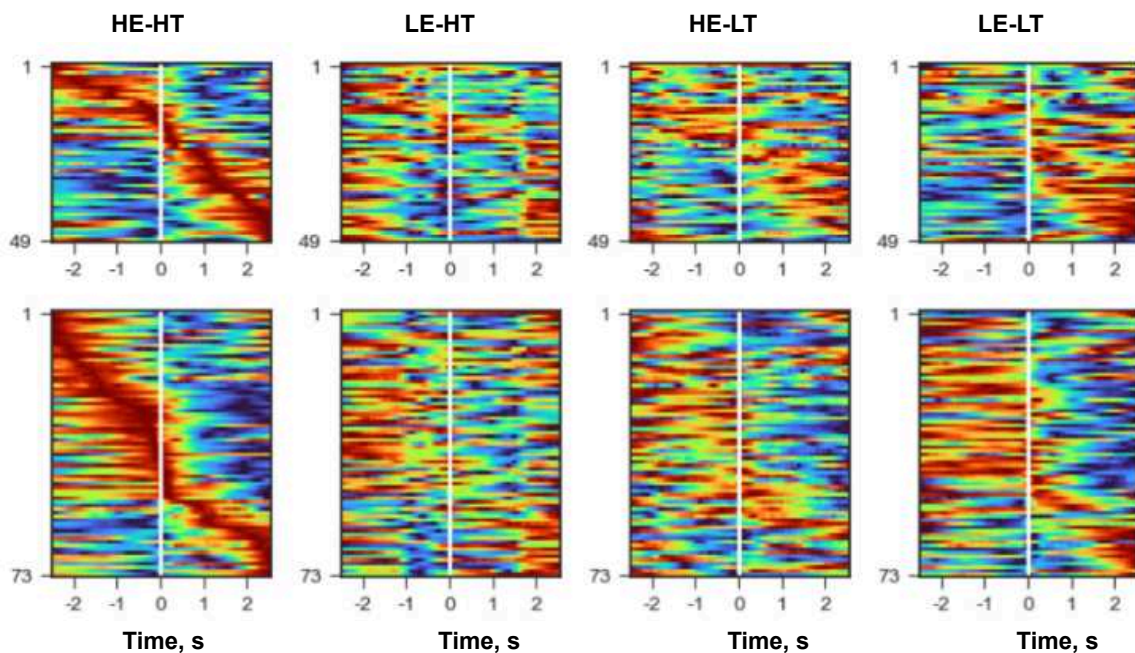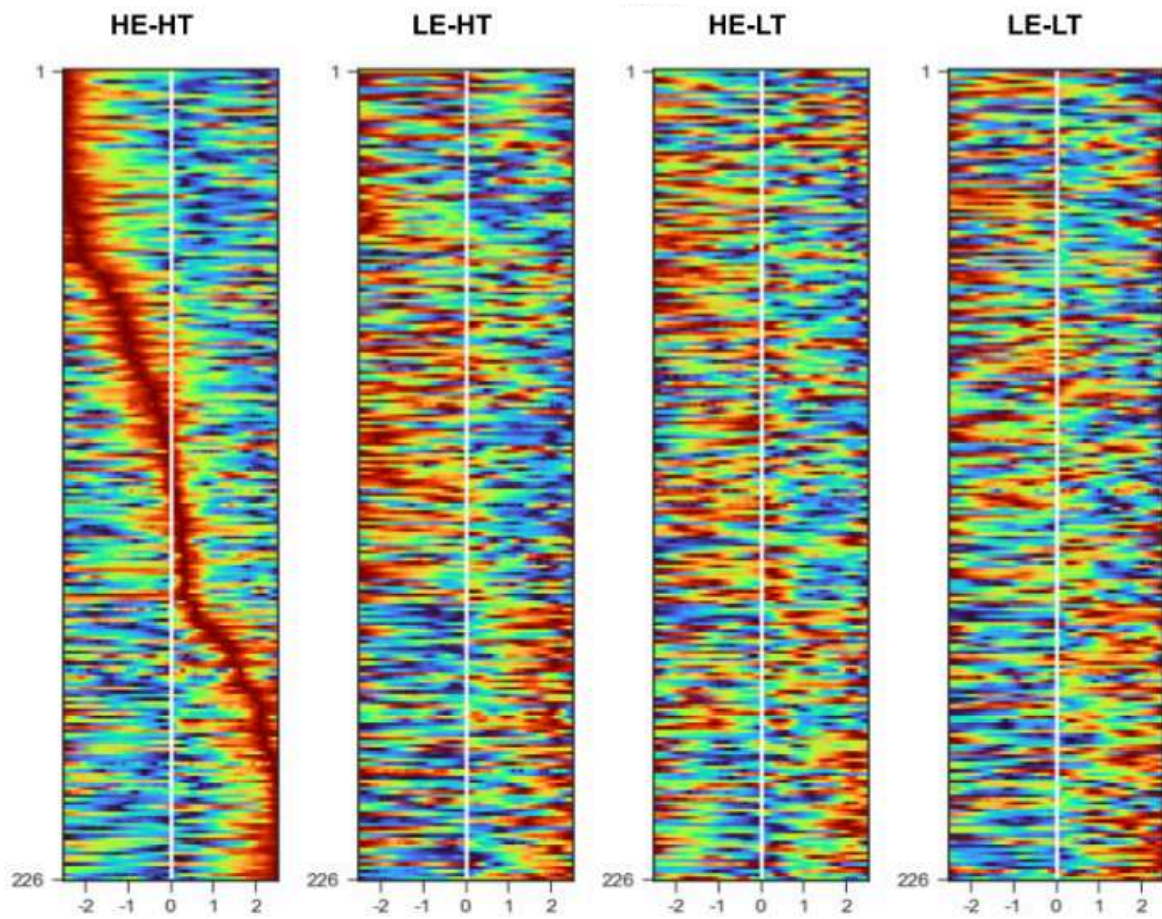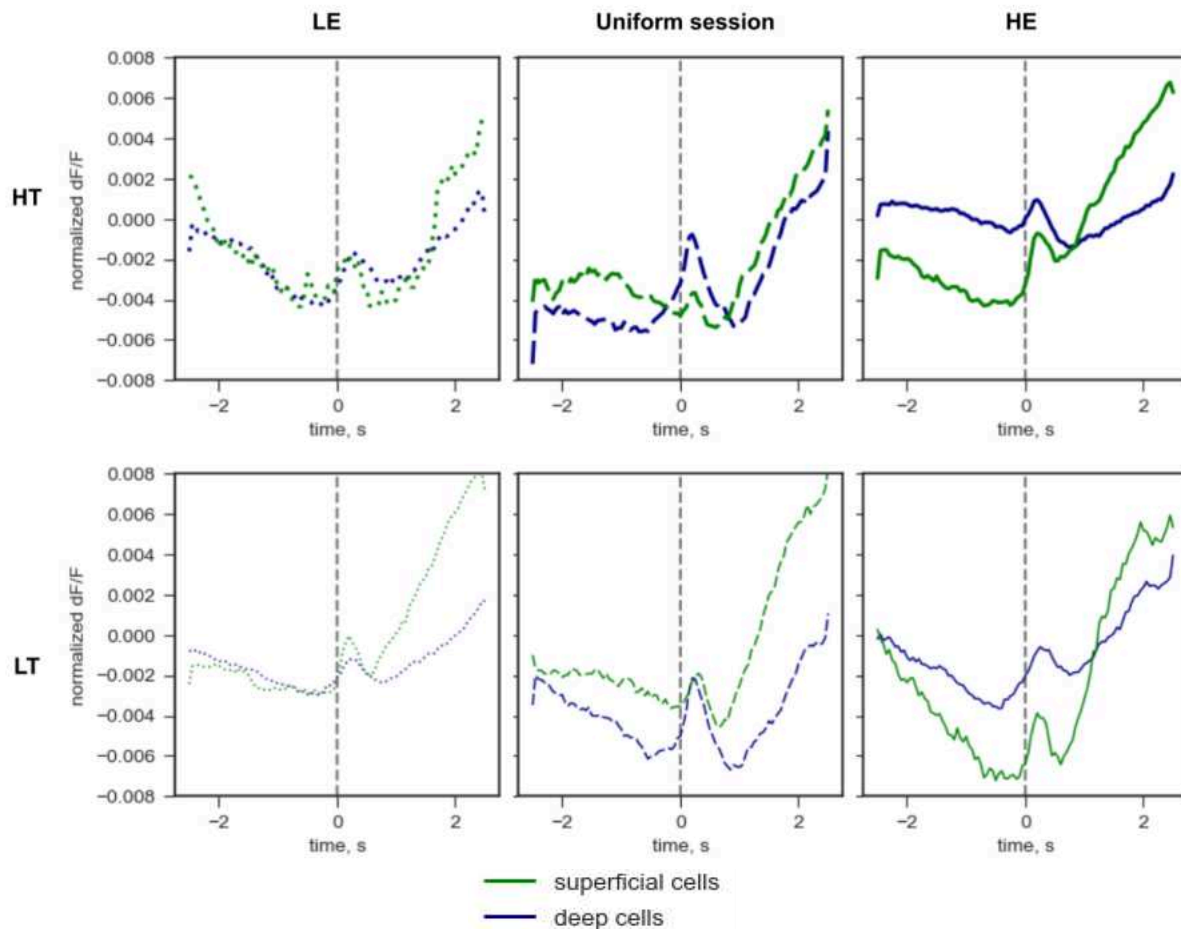The timing characteristics of the calcium activity varied largely from cell to cell and many neurons had a response much slower than the lever-pulling behavior in a single trial (**Figure 3.6**). In order to find the neurons tuned to the task variables, I extracted several features from the neuronal response and quantified their change across task conditions. The features included the average response of a cell before and after pull onset, as well as its peak amplitude ("after-pull peak").

I found a large number of cells whose peak amplitude after pull onset changed in response to the expected trial type (**Figure 3.10**), as well as the actual trial type (**Figure 3.11**). The numbers of cells with the major task variable correlations are shown in **Table 1**. In addition, the number of cells tuned to the actual and expected trial types changed at different stages of pulling. Both superficial and deep populations tuned to the expected trial type increased in size during later stages of the pull (**Figure 3.12A**). Among deep cells, there were more neurons tuned to the expected trial type than actual trial type and their number kept increasing after pull onset with the peak at around 2.25 s after pull onset. Only superficial cells increased in number in response to the actual trial type, while the number of deep cells remained roughly constant. Superficial cells also responded to the expected trial type.

To identify if the superficial and deep layers differ in the type of information they encode, I compared the distributions of the cells tuned to the expected trial type and the cells tuned to the actual trial type across layers (**Figure 3.13**). The after-pull peak amplitude had a significant association between the layers and the type of information about the trial. Superficial layers were associated with the actual trial type, while deep layers were more associated with the expected trial type. This heterogeneity was present during later stages of the pull and not during early stages (**Figure 3.12B, Appendix 2A, Appendix 2B**).
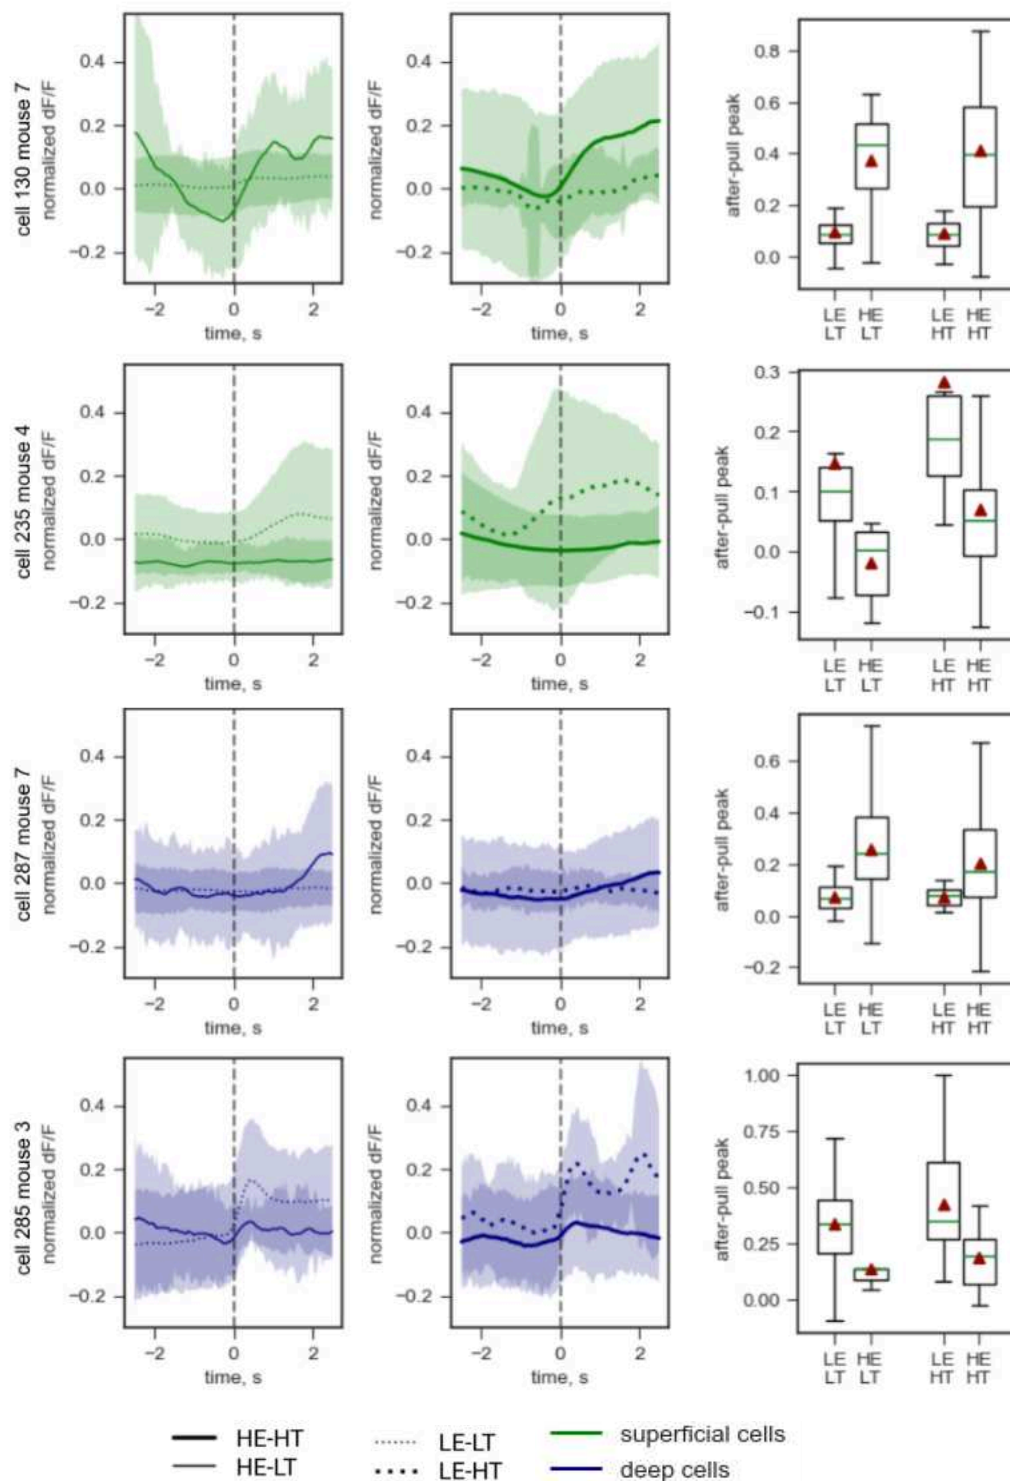
Figure 3.10. **Expected trial type coding neurons**. Example neurons whose after-pull peak amplitude changed in response to the expected trial type (unequal variances t-test, p<0.05). Left and middle panels- averaged activity traces across task conditions (error bar- standard deviation). Right panel- averaged peak values across types of expected trial type. Both superficial and deep layers contained both large numbers of neurons that stronger responded to light expected trials and large numbers of neurons that stronger responded to heavy trials.

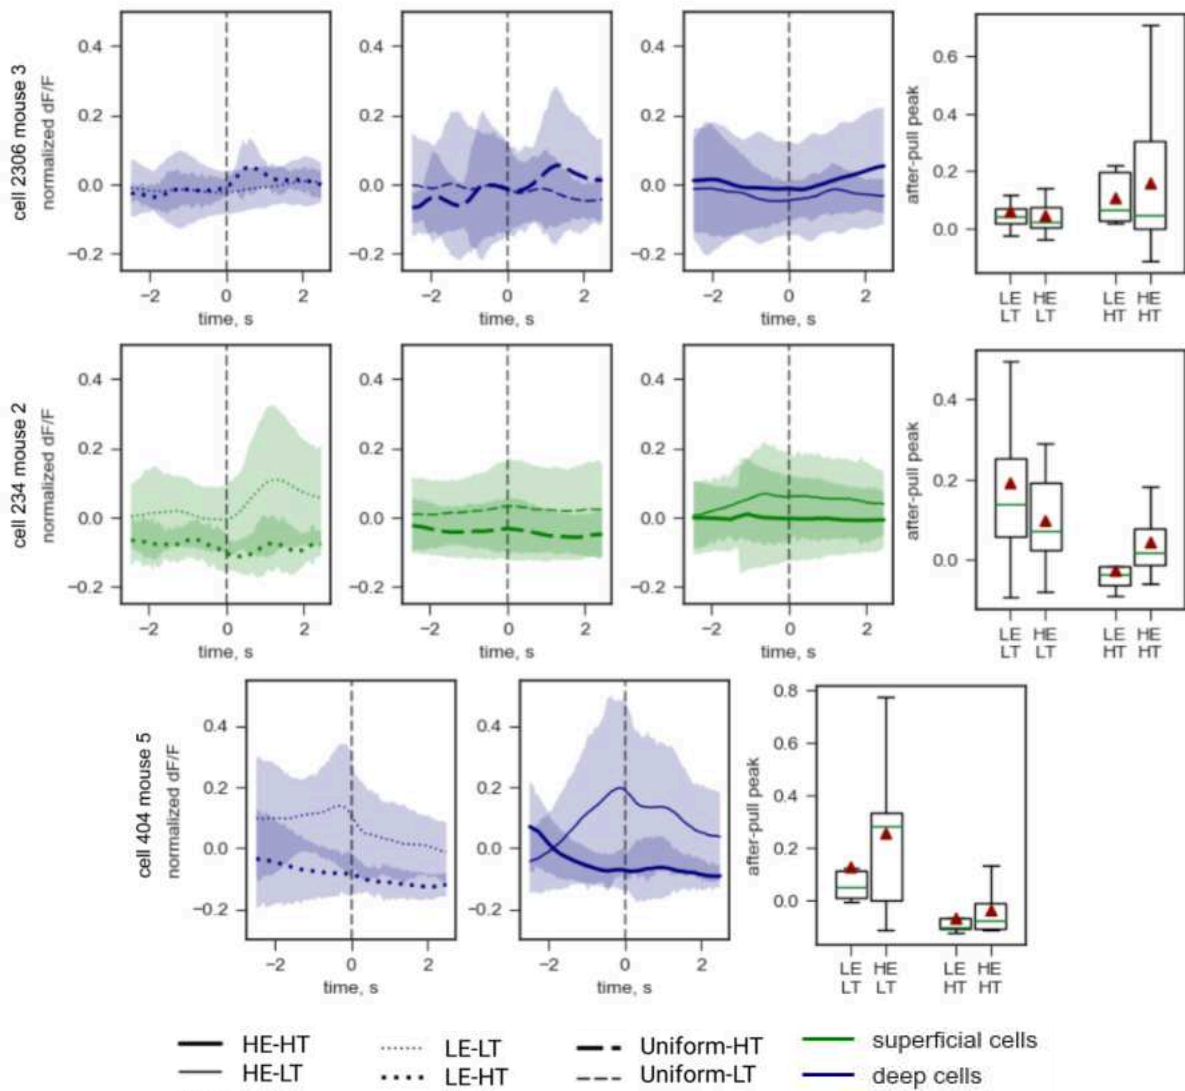**Figure 3.11. Actual trial type coding neurons**. Example cells with the after-pull peak tuned to the actual trial type (unequal variances t-test, p<0.05). First two rows, left to right- light session, uniform session, heavy session. Third row corresponds to the data from ROP sessions. Right panel only uses peak values across expected trial types, i.e. from light, heavy, and ROP sessions.

	Mouse ID	cell #	HT	LT	HE	LE	Standard	Odd	-Force	+Force	licking	reward pump	lever position	speed
			after-pull peak dF/F						dF/F(t)					
Series 1 Probabilistic sessions	1	720	43	25	24	26	69	8	254	106	29	31	180	40
	2	370	9	22	17	21	35	2	124	91	23	24	111	19
Series 2 Probabilistic sessions	3	483	15	29	35	24	33	1	132	151	52	79	184	36
		8	1	0	0	0	2	0	3	2	1	1	3	0
	4	462	18	6	32	4	38	0	146	130	20	21	132	5
		284	10	3	6	11	16	1	58	119	13	7	133	5
	5	620	24	20	22	23	29	5	114	175	26	15	144	5
	6	68	4	1	2	2	8	1	12	20	9	9	25	7
	7	289	25	4	27	7	15	3	109	53	32	14	64	14
		336	31	17	31	23	19	2	73	106	49	41	148	19
series 1-2 superficial		628	42	20	37	34	37	3	134	227	63	49	284	24
series 1-2 deep		1922	86	60	118	60	123	10	513	529	139	138	549	67
series 1-2 superficial, %			6.7	3.2	5.9	5.4	5.9	0.5	21.3	36.1	10	7.8	45.2	3.8
series 1-2 deep, %			4.5	3.1	6.1	3.1	6.4	0.5	26.7	27.5	7.2	7.2	28.6	3.5
Series 3 ROP sessions	3	454	32	14	45	28	36	4	90	160	151	94	186	34
		8	1	0	0	0	3	0	0	2	2	3	5	2
	5	945	48	51	108	53	65	9	180	247	133	91	289	58
	6	59	5	5	8	3	3	1	34	14	14	17	22	10
	7	1101	68	61	43	50	73	11	306	228	67	14	276	32
		683	39	49	26	42	38	9	169	178	50	18	232	26
series 3 superficial		691	40	49	26	42	41	9	169	180	52	21	237	28
series 3 deep		2559	153	131	204	134	177	25	610	649	365	216	773	134
series 3 superficial, %			5.8	7.1	3.8	6.1	5.9	1.3	24.5	26	7.5	3	34.3	4.1
series 3 deep, %			6	5.1	8	5.2	6.9	1	23.8	25.4	14.3	8.4	30.2	5.2
Series 4 Probabilistic sessions	3	365	8	4	6	4	3	2	129	115	50	67	118	69
		6	1	0	0	0	0	0	3	1	0	1	1	4
	5	569	19	7	10	22	2	1	229	136	65	68	169	97
	6	57	2	3	1	2	1	1	19	15	20	18	21	14
	7	279	12	9	13	3	1	1	102	53	12	7	68	17
		325	17	8	10	3	4	2	110	84	12	9	89	14
overall superficial		1650	100	77	73	79	82	14	416	492	127	80	611	70
overall deep		5751	280	214	352	225	307	40	1602	1497	651	514	1698	398
overall superficial, %			6.1	4.7	4.4	4.8	5	0.8	25.2	29.8	7.7	4.8	37	4.2
overall deep, %			4.9	3.7	6.1	3.9	5.3	0.7	27.9	26	11.3	8.9	29.5	6.9

deep cells   superficial cells

Table 1. Task variable coding neurons. The series 1-2 of the experiment contain unique cells from all mice. I repeatedly used mice for series 3 and 4. For each mouse the cells were extracted and matched within each series but not across different series. Correlation with after-pull peak calcium activity value was computed using unequal variances t-test for two independent samples each corresponding to the binary value of the task variable. The cells that had a statistically significant increase ($p<0.05$) in response to the task variable value are shown. Correlations with calcium activity trace *dF/F* were computed using t-test on linear correlation coefficient ($p<0.05$).
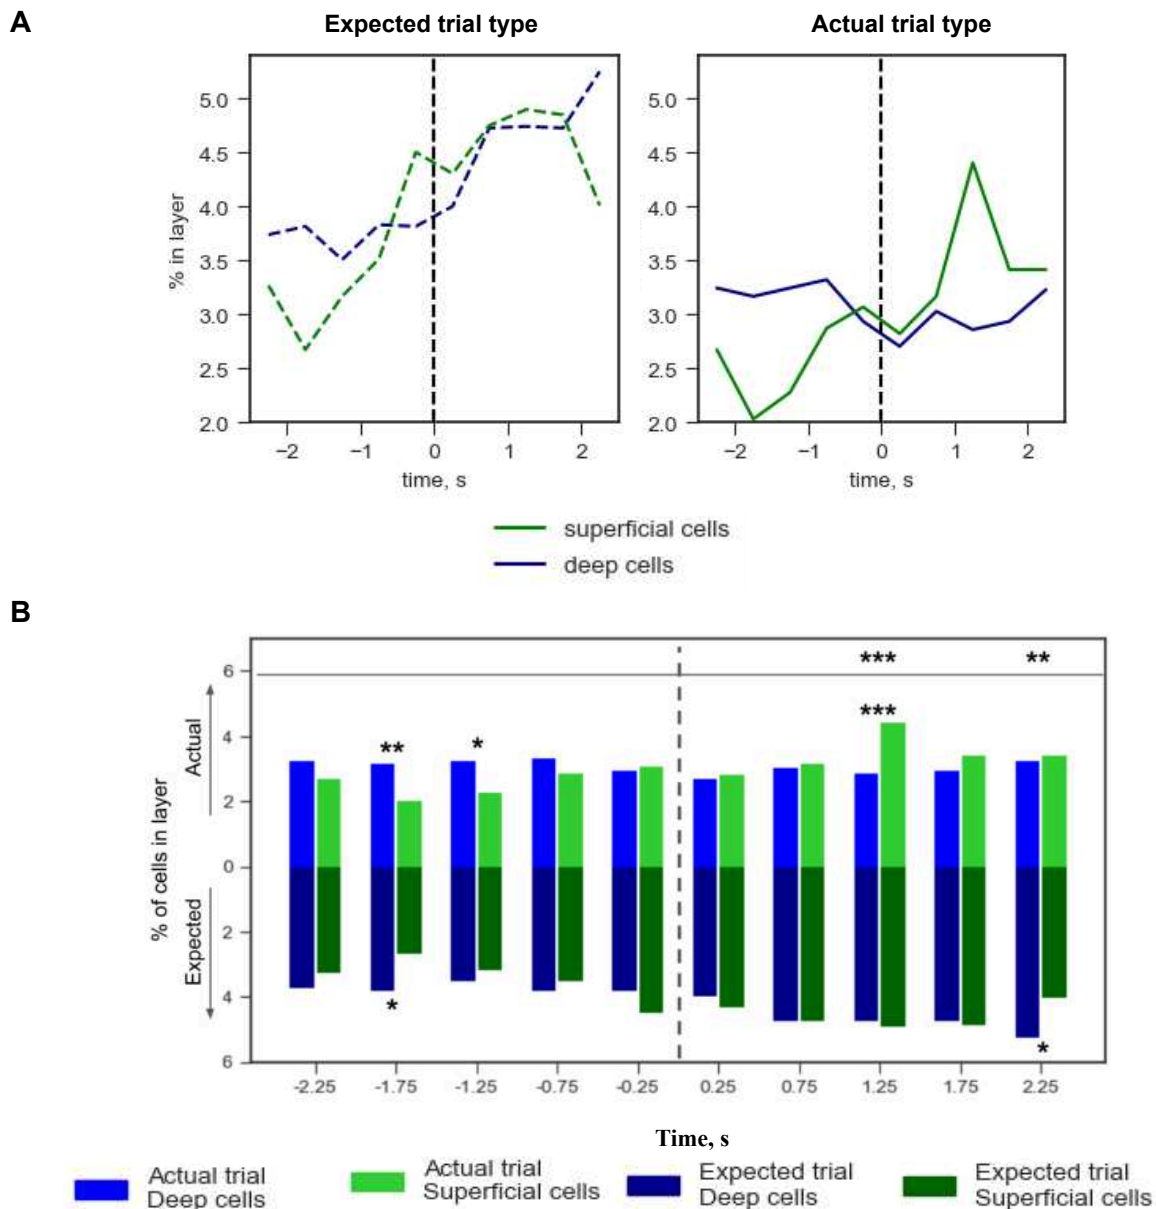
Figure 3.12. The sizes of task-relevant ensembles during lever-pulling task. **(A)** The proportion of cells per layer responding to the expected (left) and the actual trial type (right) are shown. Each cell trace was split into 0.5 s periods that were checked for the influence from expected and actual trial type (two-way ANOVA, $p < 0.05$). **(B)** Association between cortical layers and the type of information about the trial based on comparative analysis of the population sizes of the task variable coding cells. After pull onset, the proportion of the neurons tuned to the actual trial type was higher in superficial layers than in deep layers within 1..1.5 timeframe (asterisks adjacent to the bars, two-sample z-test for proportions, $p < 0.001$). There were more deep cells tuned to the expected trial type than superficial cells during later stages (two-sample z-test for proportions, $p < 0.05$). Based on the chi-square test of independence, there was a significant association between the cells' layers and the type of information to which the cells responded during later stages of the pull (asterisks at the top, $df = 1$, $p < 0.01$). Before the time defined as the pulling onset, deep cells responded stronger to both expected and actual trial types. The data from all mice and all sessions is shown.
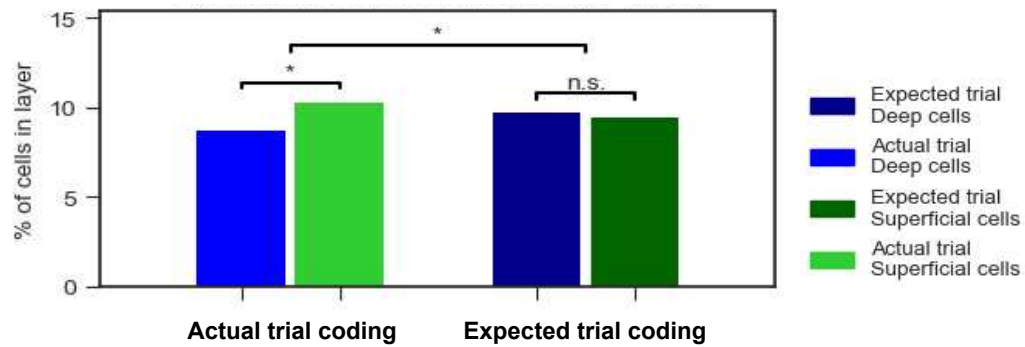
Figure 3.13. Association of after-pull peak amplitude between layers and the type of information about the trial. The cells responding with their after-pull peak to expected and actual trial type are shown (unequal variances t-test, *p*<0.05). The proportion of the actual trial type cells was higher in superficial layers (two-sample z-test for proportions, *z*=0.027 *p*=0.016). Based on chi-square tests of independence, there was a significant association between the cells' layers and the type of information about the trial to which the cells responded. The deep layers were more associated with expected trial type and superficial layers were more associated with actual trial type (X^2=5.8, *df*=1, *p*=0.016).

I tested whether individual neurons responded in line with the Bayesian computation. I identified a cell as a Bayesian posterior neuron if its response in unexpected trial types fell in-between the responses to the actual trial type and the expected trial type (**Figure 3.14A**). I used neuronal averaged activity after pull onset to identify this relationship. Despite the number of these cells being low, there was a statistically higher proportion of them in the deep layers than in superficial layers (**Figure 3.14B**).

Additionally, I identified the cells whose response was in line with the prediction-error coding in which responses to unexpected sensory input are enhanced. In this case, the difference between responses to trial types should be amplified when the trial is unexpected. There was no significant difference in the proportion of such cells among cortical layers (**Figure 3.15**).
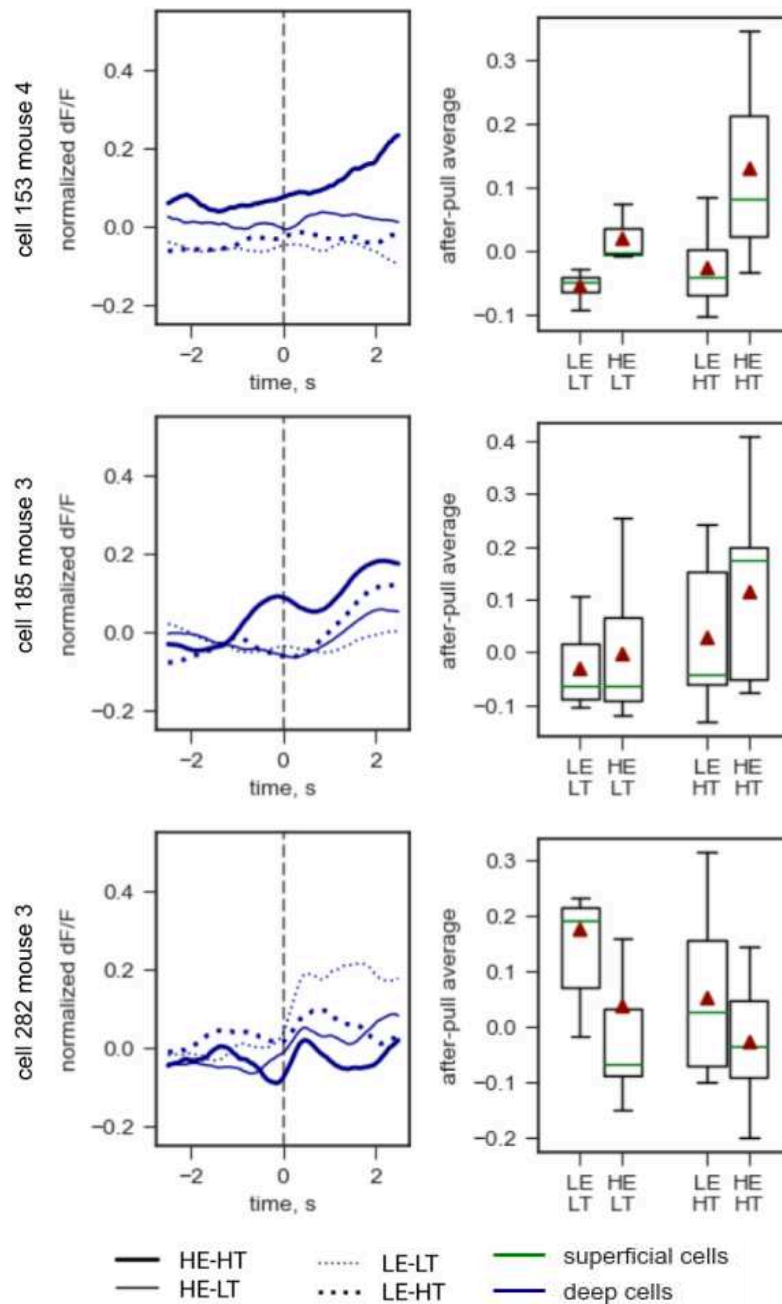
Figure 3.14A. Neurons with Bayesian posterior response. Averaged activity traces of three sample cells. Expected heavy and expected light trials correspond to the extreme values of the averaged activity after pull onset (> 0 s timestep). Unexpected trial responses lie in-between, in line with the Bayesian posterior estimation. Top- there is a strong bias towards the sensory input; middle panel- strong bias towards the prior; bottom panel- reverse relationship, i.e. the cell was more active in light trials.
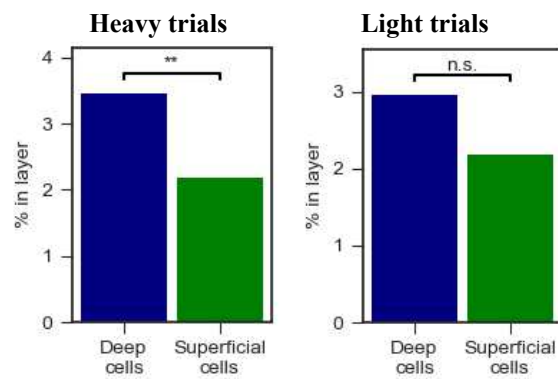
Figure 3.14B. Layer distribution of neurons with Bayesian posterior response. Distribution of Bayesian posterior cells across layers. Compared to superficial layers, deep layers contained a higher proportion of neurons whose averaged after-pull response was in line with Bayesian posterior estimation (two-sample z-test for proportions, z=3.44, p<0.001). Left panel- cells responding stronger to heavy trials (z=2.89, p=0.004); Right- cells responding stronger to light trials (z=1.88, p=0.060).

A
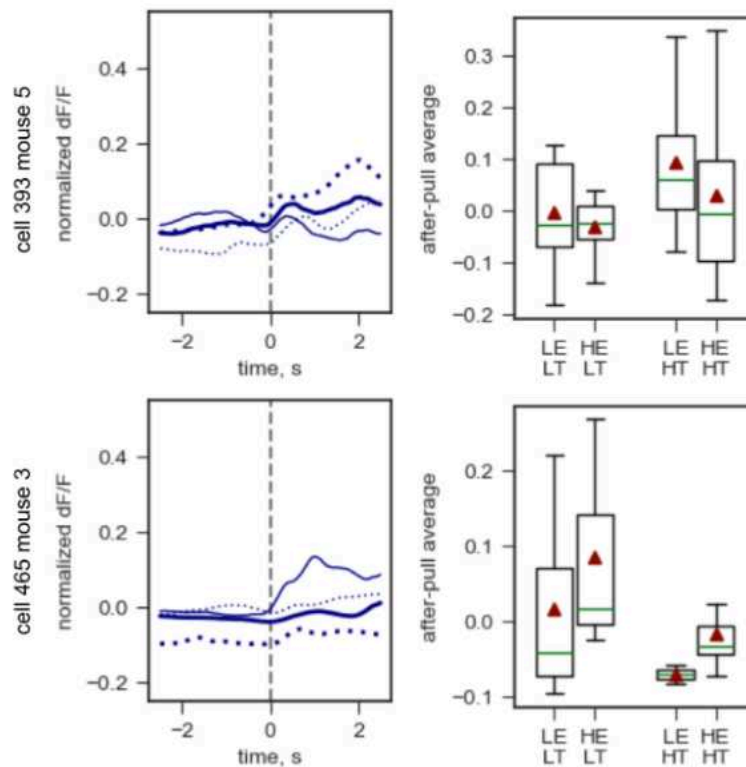
B

Figure 3.15. Neurons with prediction-error response. **(A)** Averaged activity of two example cells. Averaged after-pull response in an unexpected trial is exaggerated compared t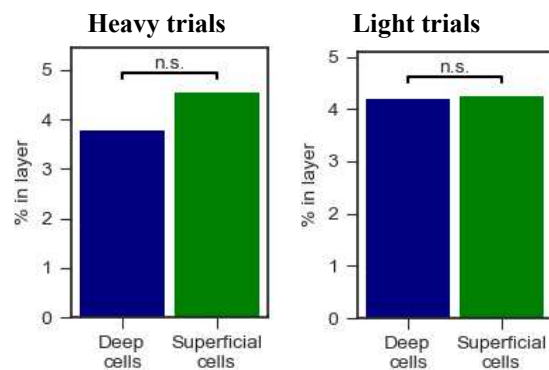o the activity in expected trial of the same type. Top- cell responds stronger to heavy trials; Bottom- cell responds stronger to light trials. **(B)** Prediction-error cell distributions across layers. Both layers contained similar proportions (two-sample z-test for proportions).

3.3 Neural population decoding analysis

To estimate the computational roles of the cortical layers in a probabilistic sensorimotor task, we need to analyze not only the task variables coding of single cells, but also the amount of information contained within the neural signal of each population of cells. To do so, I estimated the capability of different numbers of neurons in deep and superficial layers to predict the task variables.

I identified the neurons whose activity traces were correlated with the force applied to the lever by conducting a t-test on Pearson linear correlation coefficient ($p<0.05$). I used these cells to build a layer-specific predictive model to predict the value of force at each timestep. Using the leave-one-out cross-validation, I iteratively excluded each trial, trained the

decoding model on the rest of the trials, and predicted the force profile during the excluded trials. Lasso linear regression with regularization was used to modify the number of predicting cells to identify the ensemble of force-predictive cells. The model prediction was characterized by the coefficient of determination:

$$R^2 = 1 - \frac{residual\ sum\ of\ squares}{total\ sum\ of\ squares} = 1 - \frac{\Sigma(Y_t - M_t)^2}{\Sigma(Y_t - \overline{Y})^2} \quad (6)$$

where Y_t was the measured value of the variable at each timestep, and M_t was the value predicted by the model.

Prediction performance of the decoders built on calcium activity traces could reach up to 30% or higher for most mice when the model was trained and used on the data from the same day (example data is shown on **Figure 3.16A, 3.17**). The decoders' prediction performances across multiple days of probabilistic sessions was limited, however, with the coefficient of determination being negative. This is probably due to neural plasticity and other factors that affected neuronal signals and changed the cells' tunings from one day to another. In general, the poor force prediction was most likely due to slow calcium dynamics of many neurons. While being correlated with the applied force, their response was often delayed and smoothened compared to force (for example, the deep cell #1700 on **Figure 3.6**).

During ROP sessions, the task conditions change within each session, which can make it easier to examine force predictions in different conditions. I used regularized regression model to identify the optimal population size in each layer that produced the best force predictions during standard trials in ROP sessions of the mouse #7 that had both superficial and deep cells (**Figure 3.16B**). I then used this model to produce force predictions during odd heavy and odd light trials (**Figure 3.18**).

To further investigate computational roles of the cortical layers, I used the task coding neurons selected as described in the previous section for building the layer-specific regularized logistic regression model and predicting the types of expected and actual trials. The model's performance was estimated as the proportion of expected trial types correctly predicted using the neuronal activity. I considered the model to be predictive of a binary variable if it performed better than predicting always light trial or always heavy trial (no-prediction model) (**Figure 3.19, Appendix 3**).

I analyzed the quality of predictions produced by the populations of different sizes and at different timesteps of the task during a leave-one-out cross-validation (**Figure 3.20**). Based on this analysis, the deep cortical layers produced better predictions for the actual trial type and, especially, the expected trial type after pull onset. The deep cell populations of the same sizes predicted the expected trial type better than the actual trial type. The superficial cells better predicted the expected trial type.

Additionally, I used the after-pull peak amplitude of each neuron to train the model and predict the task variables. In comparison with the time-dependent neural activity, the after-pull peak amplitude of the calcium response allowed for overall better prediction performances, with the peak at 0.95. These results also contain a more pronounced difference between the decoders (**Figure 3.21**).

While informative, the difference in predictions generated with leave-one-out cross-validation may be limited when (1) the two variables are correlated with each other, and, (2) there is a similar number of cells tuned to each of the variables. So, to further contrast the functional roles of the deep and superficial layers, I trained logistic regression decoders on standard heavy and light trials, and used them to predict the trial type during odd

trials, when the actual and expected trial types were mismatched. Since we use only two types of trials, this prediction is a binary classification process where:

$$R_{expected} + R_{actual} = 1 \quad (5)$$

$R_{expected}$ - the proportion of predictions (both heavy and light) in line with the expected trial type,

R_{actual} - the proportion of predictions (both heavy and light) in line with the actual trial type.

The performance at 0.5 corresponded to no affinity of a layer to either actual or expected trial type.

The prediction of mismatched trials during ROP sessions revealed a consistent effect in line with the Bayesian hypothesis, i.e. the expected trial type was better predicted by the deep cell population, especially before and during pulling, while the actual trial type was better predicted by the superficial cell population, especially during and after pull (**Figure 3-22A, 3-22B, Appendix 3**). Despite consistency, the difference in performance was not drastic, with the peak value being only at around 0.65. This is probably due to the limited number of ROP trials used for training the model, since only the last trial of each trial train was used for training (for each mouse around 30 overall). In addition to the identified asymmetry between the layers, superficial layers represented information about the Actual trial type before pull onset, with the peak at approximately 0.5 s before pulling (**Figure 3-22A, 3-22B**). This was due to the fact that sometimes mice could sense the resistance by applying small subthreshold force after trial start and before producing a stronger pull. When predicting the trials with the short period between trial start and pulling, there was no premature prediction bias towards Actual trial type (**Appendix 4**).

The same analysis for the probabilistic sessions revealed a high variance across mice and the series of sessions. The overall effect during probabilistic sessions is shown in **Figure 3-23**. The deep layer predicted the trial type before pull onset and expected trial type after pull onset. The superficial cells predicted the expected trial type better, especially after the pull onset and actual trial type before pull onset. For both layers, the smaller populations of cells predicted the expected trial type while the larger ensembles mostly predicted the actual trial type.
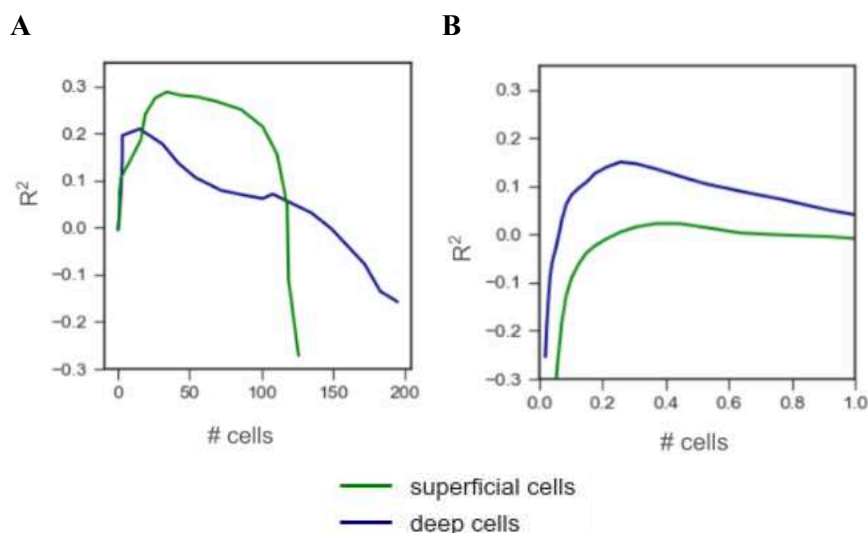

Figure 3.16. Force prediction using variable number of force-tuned cells. (A) Decoders' predictions within one (heavy) probabilistic session (mouse #3). **(B)** Decoders' prediction of force within multiple ROP sessions of mouse #7.
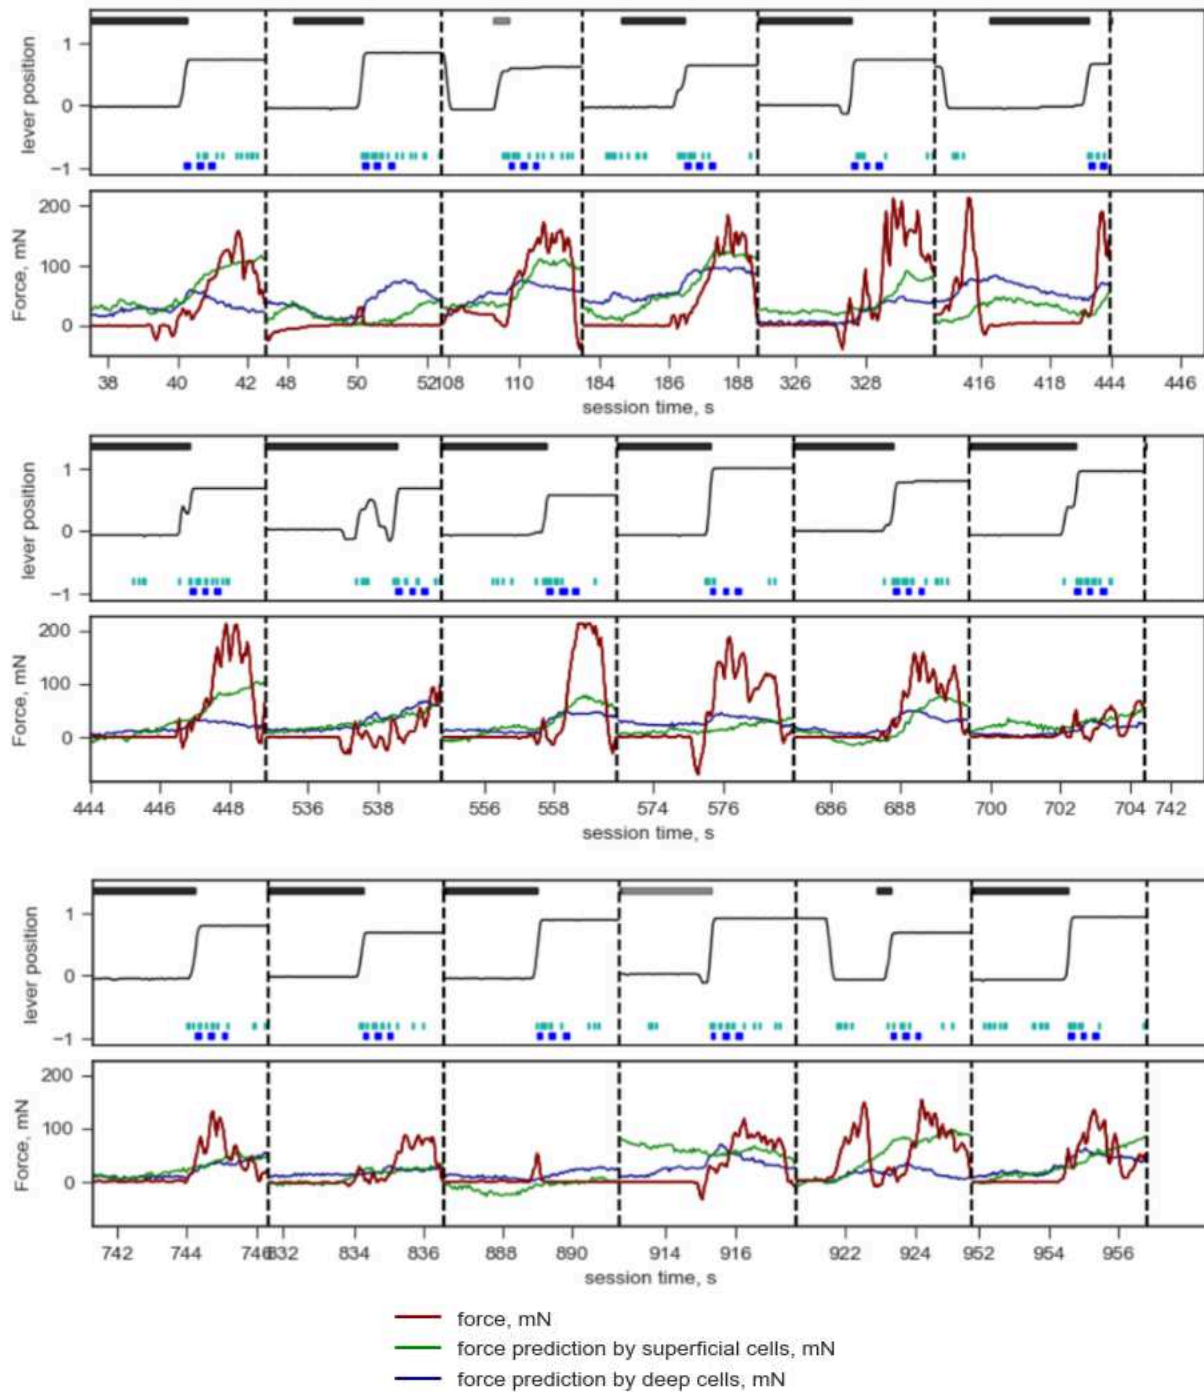
Figure 3.17. Prediction of force with layer-specific decoders. Consecutive pulls from a single heavy session (mouse #3) are shown.
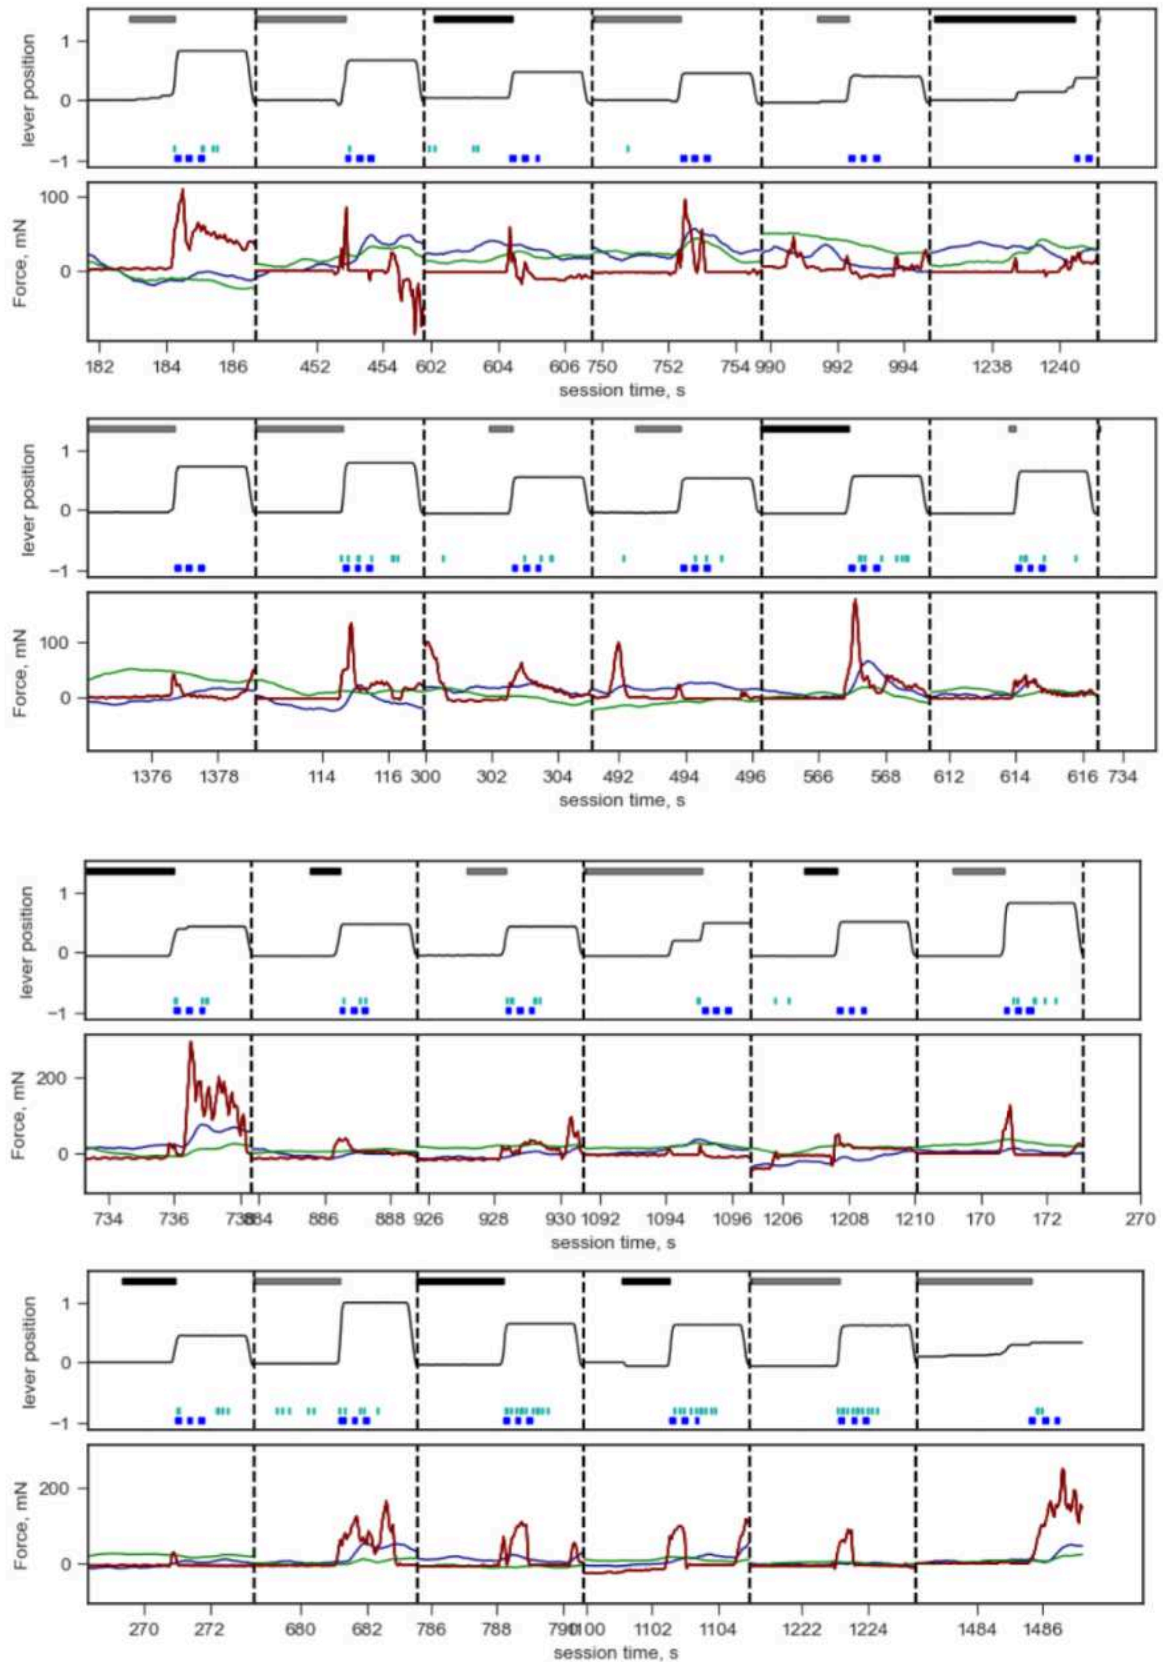
Figure 3.18. Prediction of force during odd trials with layer-specific decoders. Prediction of force during odd trials of three different ROP sessions (mouse #7).

A   **B**

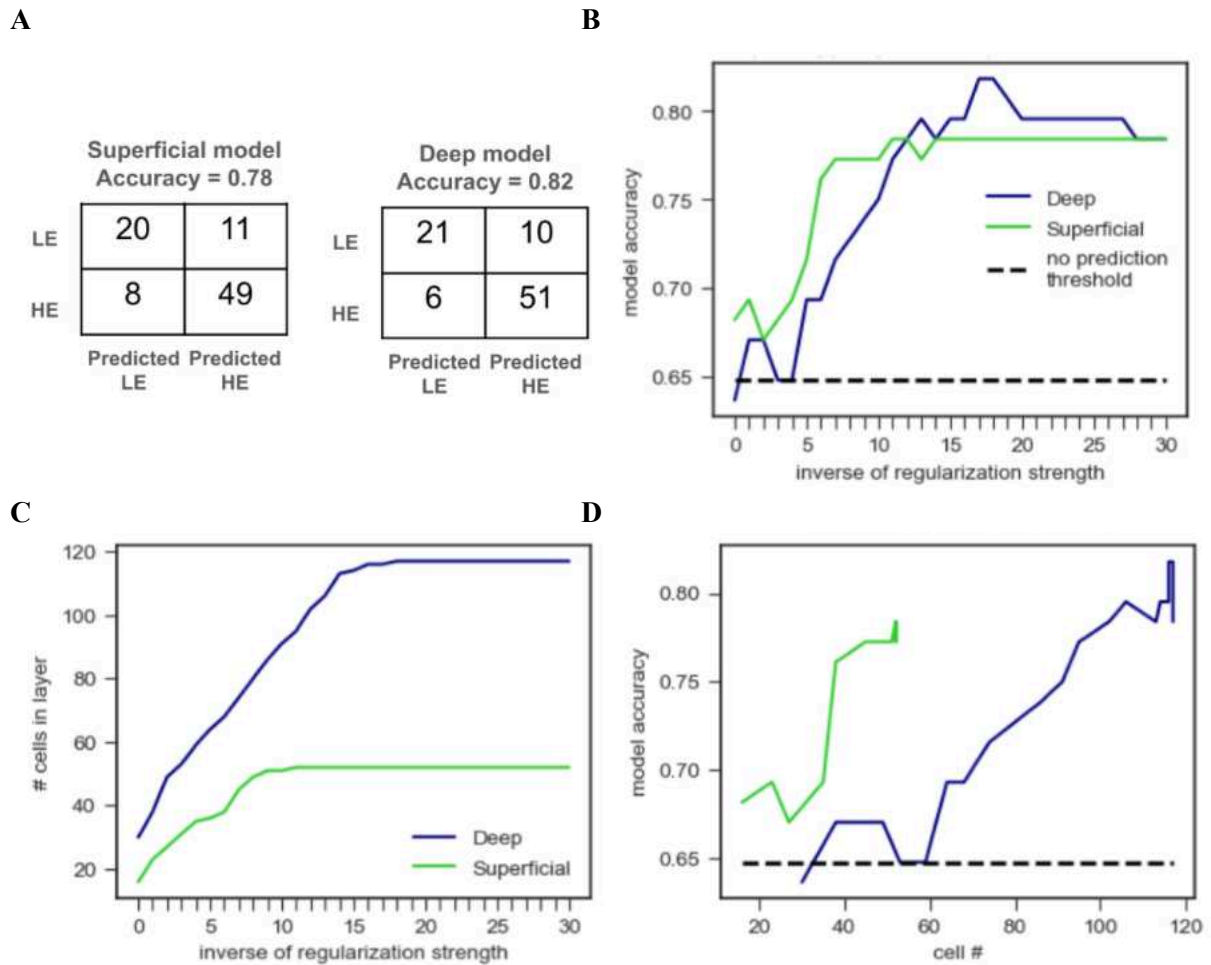**C**   **D**

Figure 3.19. Predicting expected trial type with layer-specific models. The amplitude of the after-pull peak was used to train and predict the expected trial type during excluded trials. The data from one mouse (mouse #4) is shown. **(A)** Accuracy assessment (the best case for each layer of the presented mouse is shown). **(B)** Regularization influence on the model accuracy. The horizontal axis represents the inverse of the regularization strength, i.e. the number of cells used as predictors increased from left to right. **(C)** Regularization influence on the number of cells. **(D)** Model accuracy as the function of the number of predictive cells. For the presented mouse, both superficial (green) and deep (blue) decoders performed better than no-prediction, but superficial cells predicted the expected trial type better than deep cells with the same numbers of cells.

A
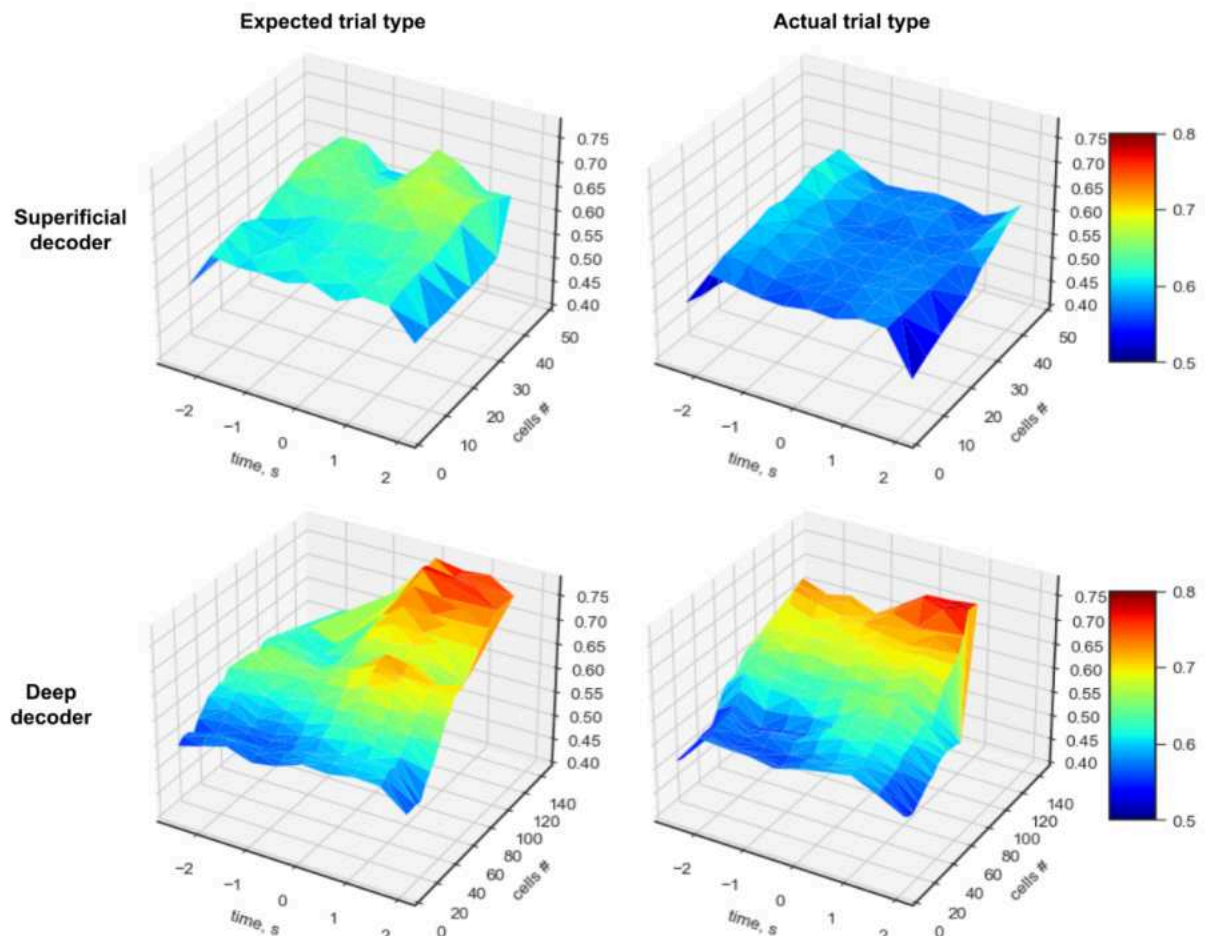
B
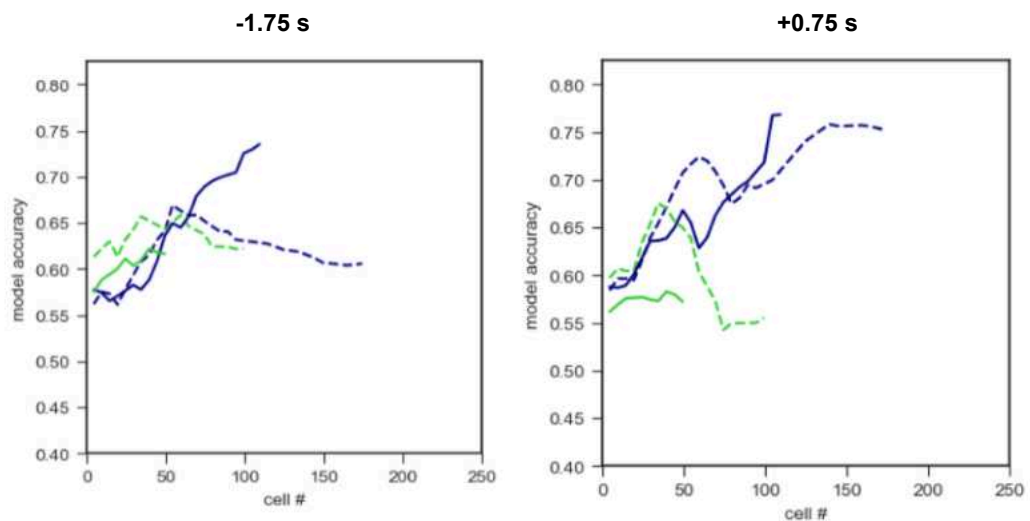
Figure 3.20. Decoders' performances at predicting expected and actual trial type. **(A)** Deep decoders performed best after pull onset while predicting the expected trial type better than the actual trial type (see red colors extending to smaller ensembles). Superficial cells predicted expected type better than actual trial type. The averaged data for all mice is shown. **(B)** The slices of the data on (A) before and after pull onset. Per ensemble size, the deep neurons predicted expected trial better than the actual trial and predicted it better than did superficial layers.
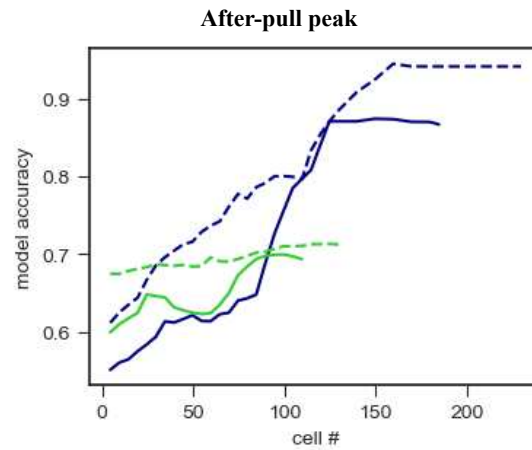
After-pull peak

Figure 3.21. Predicting expected and actual trial types with layer-specific models. Per ensemble size, the after-pull peak amplitude of the deep neurons predicted expected trial better than actual trial and predicted it better than did superficial layers.

A

Figure 3.22A. Prediction of the odd trials' type during ROP sessions by the same neuronal populations. Because of the binary outcome of the trial type classification, the shapes of the left and the right surfaces are mirrored around the horizontal plane at 0.5 prediction level. The data points below 0.5 are marked with dark blue.

B

Figure 3.22B. Binary classification of the odd trials during ROP sessions at different timepoints. The shown timepoints roughly correspond to decoders' performance peaks.
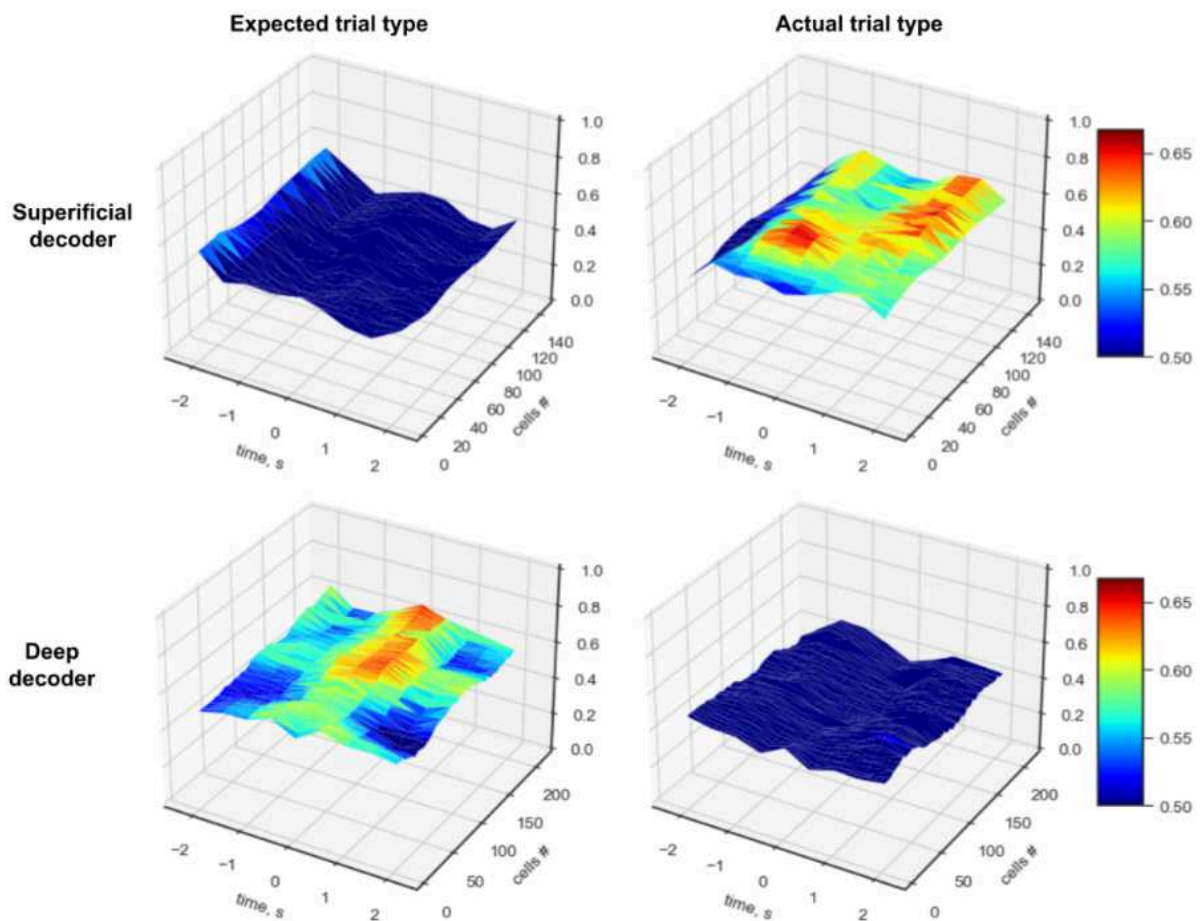
A

B

Figure 3.23. Prediction of the odd trials' type during probabilistic sessions by the same neuronal populations. Because of the binary outcome of the trial type classification, the shapes of the left and the right surfaces are mirrored around the horizontal plane at 0.5 prediction level. The data points below 0.5 are marked with dark blue.
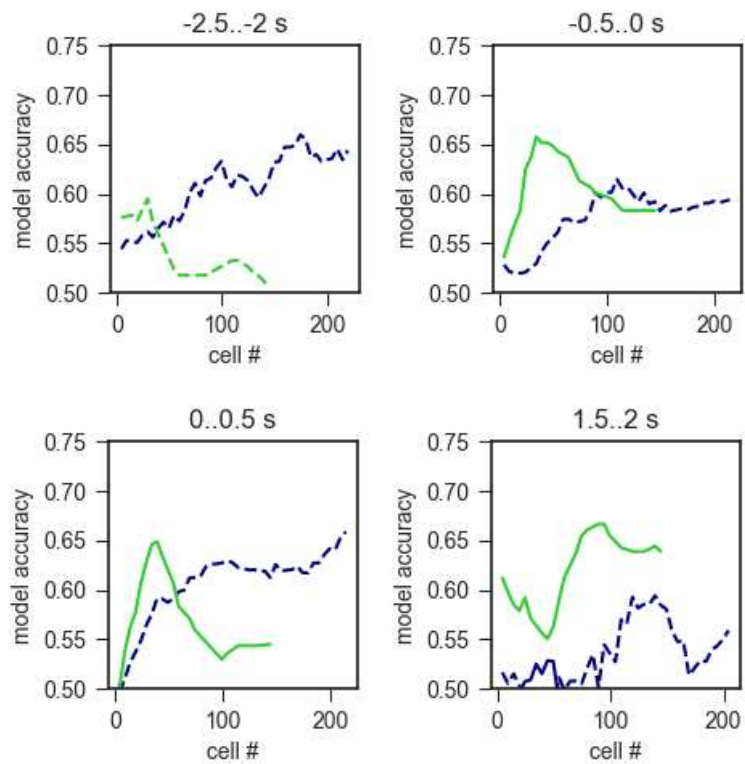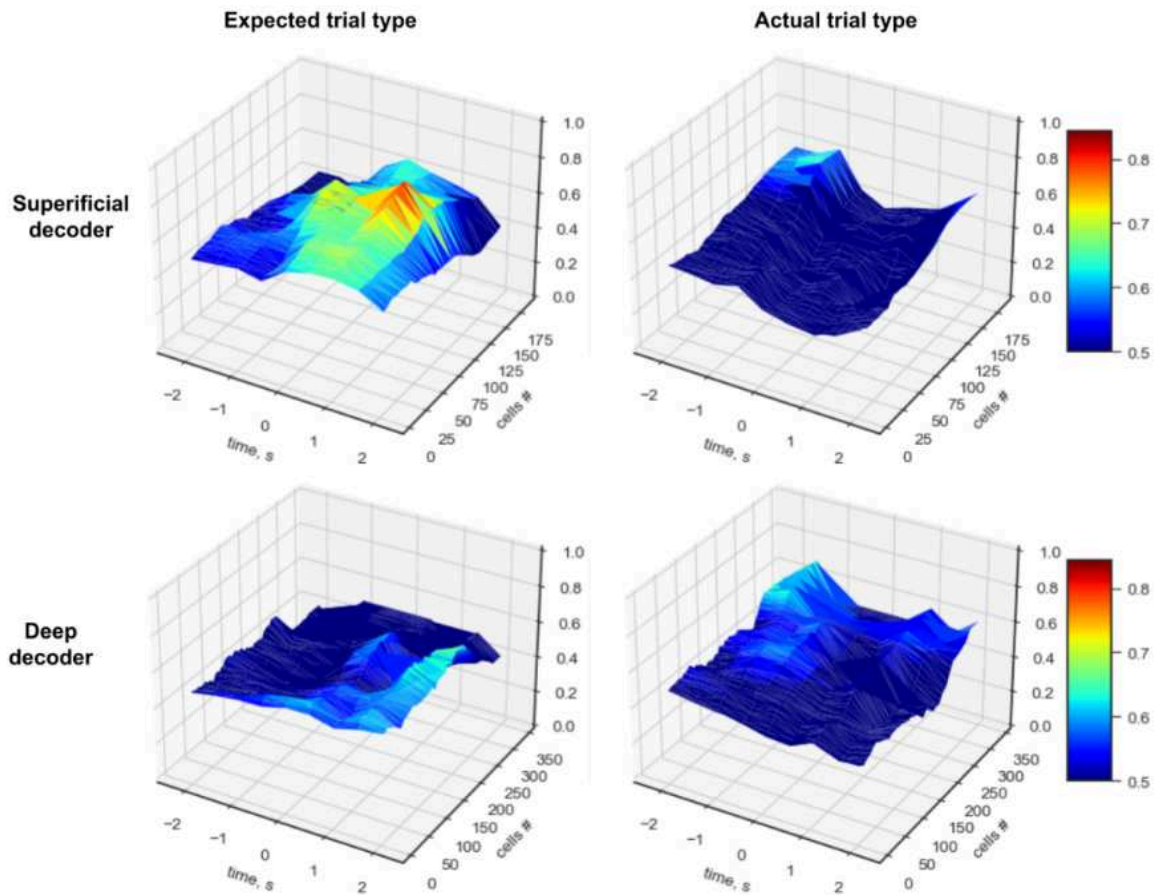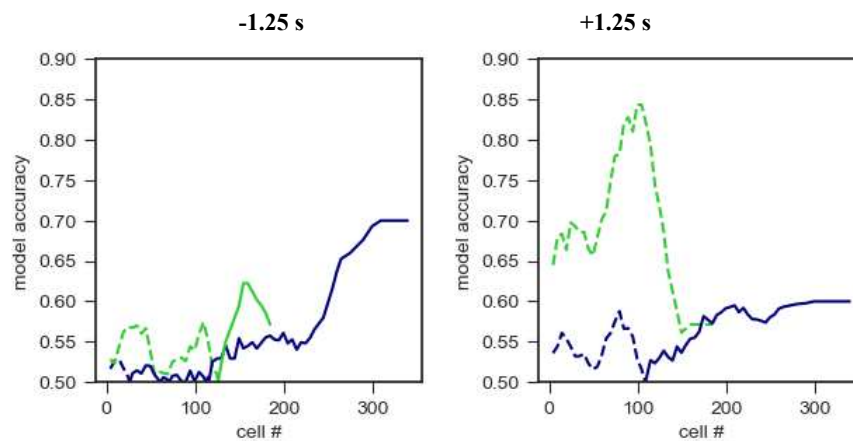
Chapter IV. Discussion

I have identified behavioral changes and neural signal changes in response to different conditions of the probabilistic lever-pulling task. The mice applied more force during actual heavy trials, but also when only expecting heavy trials. Deep layers were associated with the information about the expected trial type in terms of both the numbers of task variable coding neurons within the layer and across layers, and the neural decoders' prediction quality per population size. Superficial layers were more associated with the information about the actual trial type in comparison with the deep layers, while also encoding the expected trial type. This functional asymmetry was present after and not before pull onset.

Previous studies have suggested that the brain uses Bayesian computation during sensory integration and motor command generation. The neural signals corresponding to this computation were found in different areas of the cerebral cortex in mammals. Previous studies, however, primarily focused on the distinct functions of the cortical areas and were limited in assessing the roles of the different cortical layers during probabilistic computations (Schröder, Schmidt, and Blankenburg 2019; Funamizu, Kuhn, and Doya 2016). To my best knowledge, the current study represents the first neuroimaging experiment that directly tests computational roles of the superficial and deep neural populations during a probabilistic sensory-motor task. For this experiment, I used a GRIN lens implanted in S1FL cortical area of mice, capable of imaging multiple cortical areas at the same time. I custom-built a force-sensing lever with modifiable resistance to adapt a well-established sensory mismatch paradigm to an active lever-pulling task.

The current study has certain limitations. First, the timing characteristics of the calcium signal may not always be sufficiently fast enough to capture the neural activity relevant to lever-pulling. To mitigate it, I extracted and analyzed the features of the neural response within the 5 s time window around the pull onset, which included the peak amplitude after pulling and averaged activity at different timesteps. Second, the lens implant can disrupt the axonal projections and damage the cell bodies that are crucial for the studied neural mechanisms. To minimize the damage of the neural pathways involved in a sensory-motor task, I positioned the implant medially from the imaged cortical area such that projections between S1 and M1 area were less damaged. Third, the immunohistological study revealed, in some mice, damage of the superficial layers caused by the implant or the surgery, that limited the number of animals where both superficial and deep cells could be observed.

During both actual and expected heavy trials the mice applied stronger force. This is in line with the previous studies that demonstrated Bayesian-like adjustments in probabilistic motor tasks (Brayanov and Smith 2010). Contrary to what can be expected from the Bayesian computation, however, the force changed in response to the actual trial type before responding to the expected trial type (**Figure 3.22**). There can be a number of explanations for this behavior. One is that while anticipating a certain resistance level, the mice may have associated reward with the amount of the motor effort required by that resistance. Meeting an unexpectedly low resistance could cause mice to pull stronger to complete the planned motor act, while meeting an unexpectedly high resistance could cause them to release the lever faster.

Many neurons of the imaged S1FL cortical area had a robust tuning to the lever-pulling task (**Figure 3.6, 3.7**). The layer distribution analysis of the task variable coding neurons revealed heterogeneous response by the superficial and deep layers. It was preferentially the superficial cells that responded to the actual trial type and deep cells that responded to the expected trial type (**Figure 3.13, Appendix 2A**). Within each layer-specific population of cells there were more cells tuned to expected trial type than cells tuned to actual

trial type (**Figure 3.12A**), while my hypothesis predicts this only for deep neurons. In addition, the neural response was larger after pull onset for both types of information about the trial, while in a Bayesian computation we should expect to see the evidence of the Bayesian prior information before pulling and Bayesian posterior after pull onset. The obtained results are in line with the previous studies, however, that reported the influence of the motor cortex projections on the late calcium activity of L5 pyramidal neurons in S1 (Manita et al. 2015).

Based on the averaged prediction performances of the layer-specific decoders, the deep populations of cells contained more information per population size about the expected trial type than the actual trial type, including before pulling (**Figure 3.20, 3.21, 3-22A, 3-22B**). The early representation of the expected value was consistent with the increased overall calcium response of the deep layers before pulling onset when only expecting heavy trials and regardless of the actual trial type (**Figure 3.9**). In comparison with the superficial populations of cells, also, they better encoded the expected trial type and worse encoded the actual trial type. The preferential computational role of the superficial cells was more complicated. During ROP sessions, the superficial cells predicted the actual trial type, especially after pull onset. Within probabilistic sessions, their role was inconsistent from one animal to another. The assumption behind the design of the current study is that the ROP and probabilistic sessions are analogous in how the introduced prior and its violation affect the neuronal signaling. This assumption can be violated if neurons change their tuning from one day to another, thus making the decoding models based on probabilistic sessions less practical.

My study did not show evidence for the prediction-error coding and, specifically, the involvement of the superficial layers in it. In all mice, many more neurons had a statistically significant increase in calcium activity in response to standard trials and not in response to odd trials (**Table 1**). The overall after-pull response of the superficial layers was larger than the response of the deep layers during standard trials, and not during odd trials (**Figure 3.9**). In addition, I have identified the cells whose individual calcium activity traces could encode prediction error. The difference in the distribution of such neurons across deep and superficial layers was not statistically significant (**Figure 3.15**).

Conclusions

In recent years, multiple frameworks were proposed to address the neural mechanisms underlying Bayesian estimation in animals. While the anatomy or the neural circuits and neuroimaging studies suggest the involvement of the cerebral cortex, the precise neural mechanisms remain unknown. I found an asymmetry in computational roles of the L2/3 and L5/6 neurons during probabilistic sensory-motor task. Both in terms of the number of task variable coding neurons and the performance of the neural signal decoding models, the information about the sensory input was associated with L2/3 and the information about the prior expectation was associated with L5/6. The calcium neural signal encoded both of these types of information during and after completion of the sensory-motor task. Future research could take advantage of the current study to take a deeper look at the neural mechanisms of the brain underlying probabilistic computations.

Bibliography

Adams, Rick A., Stewart Shipp, and Karl J. Friston. 2013. "Predictions Not Commands: Active Inference in the Motor System." *Brain Structure & Function* 218 (3): 611–43.

Bastos, Andre M., W. Martin Usrey, Rick A. Adams, George R. Mangun, Pascal Fries, and Karl J. Friston. 2012. "Canonical Microcircuits for Predictive Coding." *Neuron* 76 (4): 695–711.

Beck, Jeffrey M., Peter E. Latham, and Alexandre Pouget. 2011. "Marginalization in Neural Circuits with Divisive Normalization." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31 (43): 15310–19.

Boerlin, Martin, and Sophie Denève. 2011. "Spike-Based Population Coding and Working Memory." *PLoS Computational Biology* 7 (2): e1001080.

Boudrias, Marie-Hélène, Rebecca L. McPherson, Shawn B. Frost, and Paul D. Cheney. 2010. "Output Properties and Organization of the Forelimb Representation of Motor Areas on the Lateral Aspect of the Hemisphere in Rhesus Macaques." *Cerebral Cortex* 20 (1): 169–86.

Brayanov, Jordan B., and Maurice A. Smith. 2010. "Bayesian and 'Anti-Bayesian' Biases in Sensory Integration for Action and Perception in the Size-Weight Illusion." *Journal of Neurophysiology* 103 (3): 1518–31.

Carmena, Jose M., Mikhail A. Lebedev, Roy E. Crist, Joseph E. O'Doherty, David M. Santucci, Dragan F. Dimitrov, Parag G. Patil, Craig S. Henriquez, and Miguel A. L. Nicolelis. 2003. "Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates." *PLoS Biology* 1 (2): E42.

Chen, Tsai-Wen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, et al. 2013. "Ultrasensitive Fluorescent Proteins for Imaging Neuronal Activity." *Nature* 499 (7458): 295–300.

Davare, Marco, Alexander Kraskov, John C. Rothwell, and Roger N. Lemon. 2011. "Interactions between Areas of the Cortical Grasping Network." *Current Opinion in Neurobiology* 21 (4): 565–70.

Dayan, P., G. E. Hinton, R. M. Neal, and R. S. Zemel. 1995. "The Helmholtz Machine." *Neural Computation* 7 (5): 889–904.

Deneve, Sophie. 2008. "Bayesian Spiking Neurons II: Learning." *Neural Computation* 20 (1): 118–45.

Douglas, Rodney J., and Kevan A. C. Martin. 2004. "Neuronal Circuits of the Neocortex." *Annual Review of Neuroscience* 27: 419–51.

Doya, Kenji. 2021. "Canonical Cortical Circuits and the Duality of Bayesian Inference and Optimal Control." *Current Opinion in Behavioral Sciences* 41 (October): 160–67.

Erlich, Jeffrey C., Max Bialek, and Carlos D. Brody. 2011. "A Cortical Substrate for Memory-Guided Orienting in the Rat." *Neuron* 72 (2): 330–43.

Ernst, Marc O., and Martin S. Banks. 2002. "Humans Integrate Visual and Haptic Information in a Statistically Optimal Fashion." *Nature* 415 (6870): 429–33.

Evarts, E. V. 1968. "Relation of Pyramidal Tract Activity to Force Exerted during Voluntary Movement." *Journal of Neurophysiology* 31 (1): 14–27.

Felleman, D. J., and D. C. Van Essen. 1991. "Distributed Hierarchical Processing in the Primate Cerebral Cortex." *Cerebral Cortex* 1 (1): 1–47.

Forgaard, Christopher J., Ian M. Franks, Dana Maslovat, Laurence Chin, and Romeo Chua. 2015. "Voluntary Reaction Time and Long-Latency Reflex Modulation." *Journal of Neurophysiology* 114 (6): 3386–99.

Friston, Karl. 2005. "A Theory of Cortical Responses." *Philosophical Transactions of the*

Royal Society of London. Series B, Biological Sciences 360 (1456): 815–36.

Funamizu, Akihiro, Bernd Kuhn, and Kenji Doya. 2016. "Neural Substrate of Dynamic Bayesian Inference in the Cerebral Cortex." *Nature Neuroscience* 19 (12): 1682–89.

Gandevia, Simon C., Janette L. Smith, Matthew Crawford, Uwe Proske, and Janet L. Taylor. 2006. "Motor Commands Contribute to Human Position Sense." *The Journal of Physiology* 571 (Pt 3): 703–10.

George, Dileep, and Jeff Hawkins. 2009. "Towards a Mathematical Theory of Cortical Micro-Circuits." *PLoS Computational Biology* 5 (10): e1000532.

Grienberger, Christine, and Arthur Konnerth. 2012. "Imaging Calcium in Neurons." *Neuron* 73 (5): 862–85.

Gulati, Srishti, Vania Y. Cao, and Stephani Otte. 2017. "Multi-Layer Cortical Ca2+ Imaging in Freely Moving Mice with Prism Probes and Miniaturized Fluorescence Microscopy." *Journal of Visualized Experiments: JoVE*, no. 124 (June). https://doi.org/10.3791/55579.

Guo, Jian-Zhong, Austin R. Graves, Wendy W. Guo, Jihong Zheng, Allen Lee, Juan Rodríguez-González, Nuo Li, et al. 2015. "Cortex Commands the Performance of Skilled Movement." https://doi.org/10.7554/eLife.10774.

Guo, Zengcai V., S. Andrew Hires, Nuo Li, Daniel H. O'Connor, Takaki Komiyama, Eran Ophir, Daniel Huber, et al. 2014. "Procedures for Behavioral Experiments in Head-Fixed Mice." *PloS One* 9 (2): e88678.

Haeusler, Stefan, and Wolfgang Maass. 2007. "A Statistical Analysis of Information-Processing Properties of Lamina-Specific Cortical Microcircuit Models." *Cerebral Cortex* 17 (1): 149–62.

Hirasawa, Naoto, Kazuyuki Yamada, and Masanori Murayama. 2016. "Brief Hind Paw Stimulation Is Sufficient to Induce Delayed Somatosensory Discrimination Learning in C57BL/6 Mice." *Behavioural Brain Research* 301 (March): 102–9.

Jones, Joshua L., Guillem R. Esber, Michael A. McDannald, Aaron J. Gruber, Alex Hernandez, Aaron Mirenzi, and Geoffrey Schoenbaum. 2012. "Orbitofrontal Cortex Supports Behavior and Learning Using Inferred but Not Cached Values." *Science* 338 (6109): 953–56.

Kätzel, Dennis, Boris V. Zemelman, Christina Buetfering, Markus Wölfel, and Gero Miesenböck. 2011. "The Columnar and Laminar Organization of Inhibitory Connections to Neocortical Excitatory Cells." *Nature Neuroscience* 14 (1): 100–107.

Kok, Peter, Lauren J. Bains, Tim van Mourik, David G. Norris, and Floris P. de Lange. 2016. "Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback." *Current Biology: CB* 26 (3): 371–76.

Körding, Konrad P., and Daniel M. Wolpert. 2004. "Bayesian Integration in Sensorimotor Learning." *Nature* 427 (6971): 244–47.

Kurtzer, Isaac, Troy M. Herter, and Stephen H. Scott. 2005. "Random Change in Cortical Load Representation Suggests Distinct Control of Posture and Movement." *Nature Neuroscience* 8 (4): 498–504.

Landgren, S., and H. Silfvenius. 1969. "Projection to Cerebral Cortex of Group I Muscle Afferents from the Cat's Hind Limb." *The Journal of Physiology* 200 (2): 353–72.

Larkum, Matthew. 2013. "A Cellular Mechanism for Cortical Associations: An Organizing Principle for the Cerebral Cortex." *Trends in Neurosciences* 36 (3): 141–51.

Larkum, Matthew E., Thomas Nevian, Maya Sandler, Alon Polsky, and Jackie Schiller. 2009. "Synaptic Integration in Tuft Dendrites of Layer 5 Pyramidal Neurons: A New Unifying Principle." *Science* 325 (5941): 756–60.

Latash, Mark L. 2008. *Neurophysiological Basis of Movement*. Human Kinetics.

Lochmann, Timm, and Sophie Deneve. 2011. "Neural Processing as Causal Inference." *Current Opinion in Neurobiology* 21 (5): 774–81.

Manita, Satoshi, Takayuki Suzuki, Chihiro Homma, Takashi Matsumoto, Maya Odagawa, Kazuyuki Yamada, Keisuke Ota, et al. 2015. "A Top-Down Cortical Circuit for Accurate Sensory Perception." *Neuron* 86 (5): 1304–16.

Mao, Tianyi, Deniz Kusefoglu, Bryan M. Hooks, Daniel Huber, Leopoldo Petreanu, and Karel Svoboda. 2011. "Long-Range Neuronal Circuits Underlying the Interaction between Sensory and Motor Cortex." *Neuron* 72 (1): 111–23.

Markov, Nikola T., Julien Vezoli, Pascal Chameau, Arnaud Falchier, René Quilodran, Cyril Huissoud, Camille Lamy, et al. 2014. "Anatomy of Hierarchy: Feedforward and Feedback Pathways in Macaque Visual Cortex." *The Journal of Comparative Neurology* 522 (1): 225–59.

Markov, N. T., P. Misery, A. Falchier, C. Lamy, J. Vezoli, R. Quilodran, M. A. Gariel, et al. 2011. "Weight Consistency Specifies Regularities of Macaque Cortical Networks." *Cerebral Cortex* 21 (6): 1254–72.

Matthews, P. B. C. 1982. "Where Does Sherrington's' Muscular Sense' Originate? Muscles, Joints, Corollary Discharges?" *Annual Review of Neuroscience* 5 (1): 189–218.

Ma, Wei Ji, Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget. 2006. "Bayesian Inference with Probabilistic Population Codes." *Nature Neuroscience* 9 (11): 1432–38.

Meyer, Travis, and Carl R. Olson. 2011. "Statistical Learning of Visual Transitions in Monkey Inferotemporal Cortex." *Proceedings of the National Academy of Sciences of the United States of America* 108 (48): 19401–6.

Murayama, Masanori, and Matthew E. Larkum. 2009. "Enhanced Dendritic Activity in Awake Rats." *Proceedings of the National Academy of Sciences of the United States of America* 106 (48): 20482–86.

Nicolelis, Miguel A. L., and Mikhail A. Lebedev. 2009. "Principles of Neural Ensemble Physiology Underlying the Operation of Brain-Machine Interfaces." *Nature Reviews. Neuroscience* 10 (7): 530–40.

Orbán, Gergo, and Daniel M. Wolpert. 2011. "Representations of Uncertainty in Sensorimotor Control." *Current Opinion in Neurobiology* 21 (4): 629–35.

Oscarsson, O., and I. Rosen. 1963. "PROJECTION TO CEREBRAL CORTEX OF LARGE MUSCLE-SPINDLE AFFERENTS IN FORELIMB NERVES OF THE CAT." *The Journal of Physiology* 169 (December): 924–45.

Paxinos, George, and Keith B. J. Franklin. 2019. *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates*. Academic Press.

Phillips, C. G., T. P. Powell, and M. Wiesendanger. 1971. "Projection from Low-Threshold Muscle Afferents of Hand and Forearm to Area 3a of Baboon's Cortex." *The Journal of Physiology* 217 (2): 419–46.

Pouget, Alexandre, Jeffrey M. Beck, Wei Ji Ma, and Peter E. Latham. 2013. "Probabilistic Brains: Knowns and Unknowns." *Nature Neuroscience* 16 (9): 1170–78.

Proske, Uwe, and Simon C. Gandevia. 2009. "The Kinaesthetic Senses." *The Journal of Physiology* 587 (Pt 17): 4139–46.

Quallo, Marsha M., Alexander Kraskov, and Roger N. Lemon. 2012. "The Activity of Primary Motor Cortex Corticospinal Neurons during Tool Use by Macaque Monkeys." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 32 (48): 17351–64.

Rao, R. P., and D. H. Ballard. 1999. "Predictive Coding in the Visual Cortex: A Functional Interpretation of Some Extra-Classical Receptive-Field Effects." *Nature Neuroscience* 2 (1): 79–87.

Resendez, Shanna L., Josh H. Jennings, Randall L. Ung, Vijay Mohan K. Namboodiri, Zhe Charles Zhou, James M. Otis, Hiroshi Nomura, Jenna A. McHenry, Oksana Kosyk, and Garret D. Stuber. 2016. "Visualization of Cortical, Subcortical and Deep Brain Neural

Circuit Dynamics during Naturalistic Mammalian Behavior with Head-Mounted Microscopes and Chronically Implanted Lenses." *Nature Protocols* 11 (3): 566–97.

Ross, William N. 2012. "Understanding Calcium Waves and Sparks in Central Neurons." *Nature Reviews. Neuroscience* 13 (3): 157–68.

Samaranch, Lluis, Ernesto A. Salegio, Waldy San Sebastian, Adrian P. Kells, John R. Bringas, John Forsayeth, and Krystof S. Bankiewicz. 2013. "Strong Cortical and Spinal Cord Transduction after AAV7 and AAV9 Delivery into the Cerebrospinal Fluid of Nonhuman Primates." *Human Gene Therapy* 24 (5): 526–32.

Schröder, Pia, Timo Torsten Schmidt, and Felix Blankenburg. 2019. "Neural Basis of Somatosensory Target Detection Independent of Uncertainty, Relevance, and Reports." *eLife* 8 (March). https://doi.org/10.7554/eLife.43410.

Summerfield, Christopher, and Etienne Koechlin. 2008. "A Neural Representation of Prior Information during Perceptual Inference." *Neuron* 59 (2): 336–47.

Tassinari, Hadley, Todd E. Hudson, and Michael S. Landy. 2006. "Combining Priors and Noisy Visual Cues in a Rapid Pointing Task." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 26 (40): 10154–63.

Todorovic, Ana, Freek van Ede, Eric Maris, and Floris P. de Lange. 2011. "Prior Expectation Mediates Neural Adaptation to Repeated Sounds in the Auditory Cortex: An MEG Study." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 31 (25): 9118–23.

Watakabe, Akiya, Masanari Ohtsuka, Masaharu Kinoshita, Masafumi Takaji, Kaoru Isa, Hiroaki Mizukami, Keiya Ozawa, Tadashi Isa, and Tetsuo Yamamori. 2015. "Comparative Analyses of Adeno-Associated Viral Vector Serotypes 1, 2, 5, 8 and 9 in Marmoset, Mouse and Macaque Cerebral Cortex." *Neuroscience Research* 93 (April): 144–57.

Watson, Charles, George Paxinos, and Luis Puelles. 2012. *The Mouse Nervous System*. Academic Press.

Weiler, Nicholas, Lydia Wood, Jianing Yu, Sara A. Solla, and Gordon M. G. Shepherd. 2008. "Top-down Laminar Organization of the Excitatory Network in Motor Cortex." *Nature Neuroscience* 11 (3): 360–66.

Yoshizawa, Tomohiko, Makoto Ito, and Kenji Doya. 2018. "Reward-Predictive Neural Activities in Striatal Striosome Compartments." *eNeuro* 5 (1). https://doi.org/10.1523/ENEURO.0367-17.2018.

Appendices

Appendix 1. Task-relevant neural activity

Activity traces of neurons are shown, aligned by the time of pull onset (at t=0 s). Each panel represents a single experiment condition, across which the data was time aligned, averaged over trials, time averaged with a moving window of 200 ms, and normalized. Each horizontal bar corresponds to a single cell. **(A)** Within all panels, cells are sorted by the time of their peak activity during heavy trials in heavy probabilistic sessions (left panel). Cells roughly maintained their tuning to lever-pulling task. Superficial and deep cells of the same mouse (mouse #4) are shown. **(B)** Activity traces averaged across all task conditions in probabilistic sessions (7 different mice with corresponding layers). **(C)** Activity traces averaged across all task conditions in ROP sessions (4 different mice with corresponding layers).

B

C

mouse #3 deep

mouse #5 deep

mouse #6 deep

mouse #7 deep

mouse #7 superficial
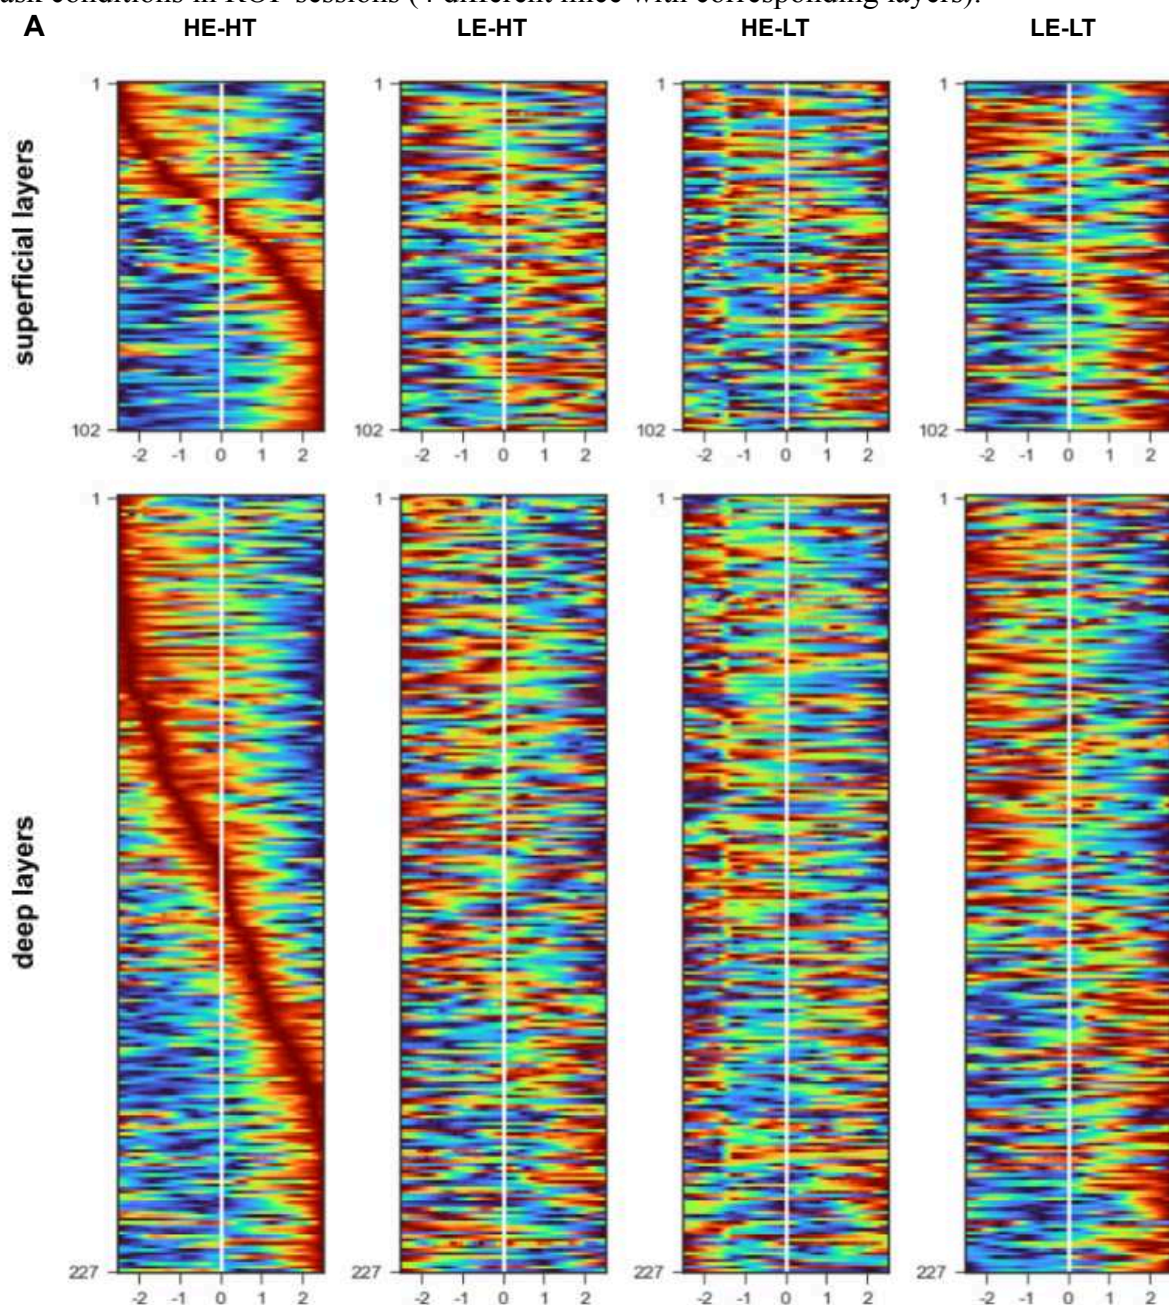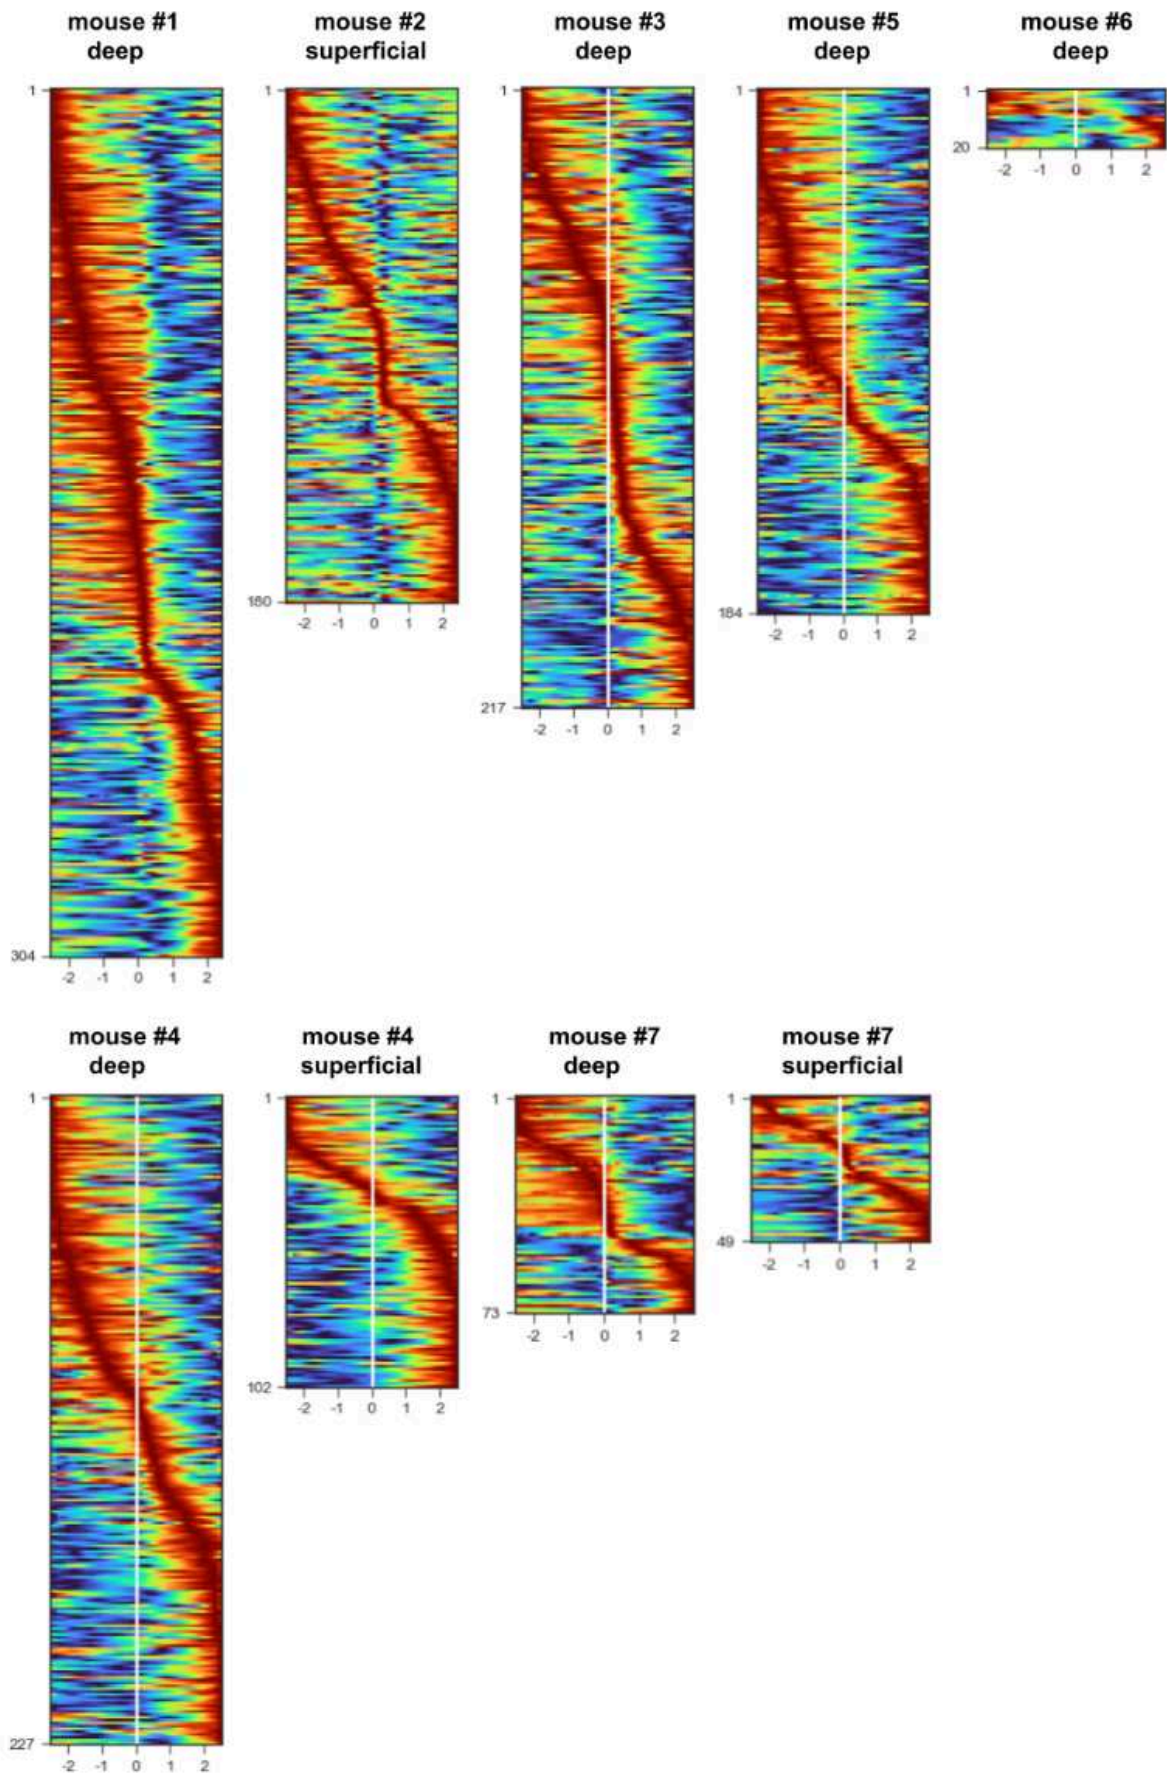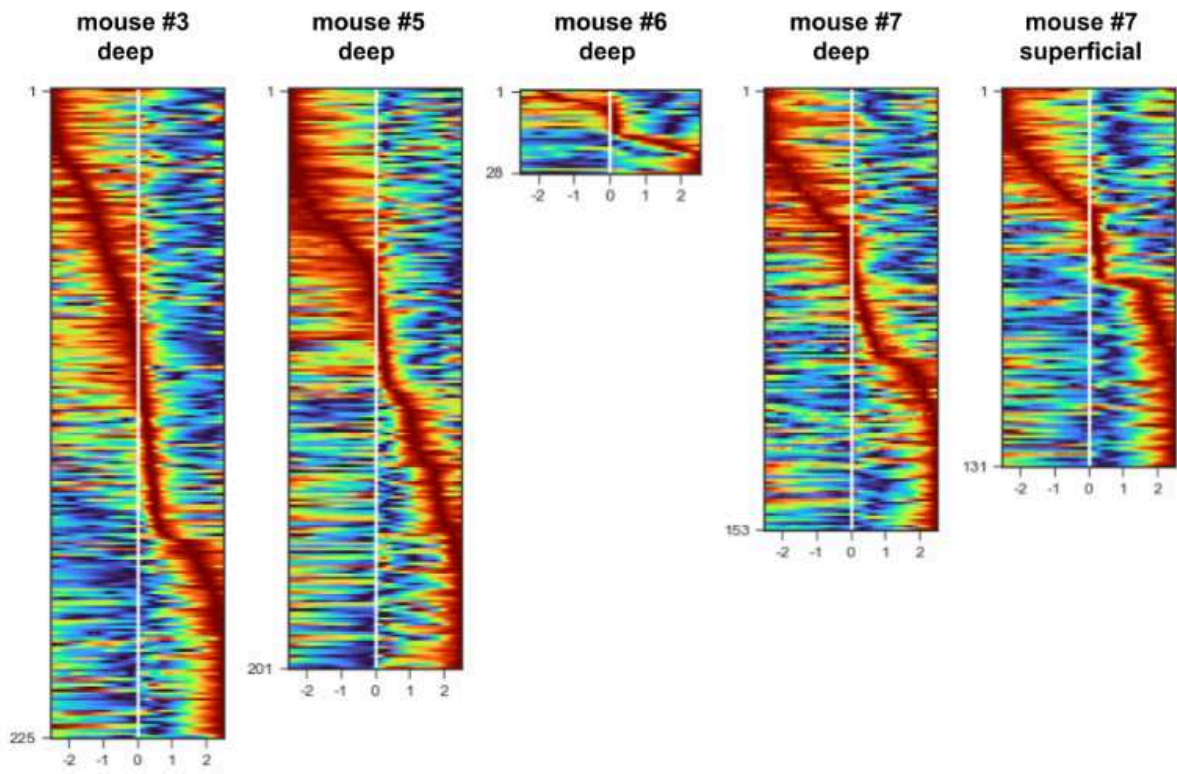
Appendix 2A. Association between cortical layers and the type of information about the trial type.

The distribution of cells responding to the expected and actual trial types are shown. The traces were binned into 0.5 time intervals, averaged and tested for a change across conditions with unequal variances t-test ($p<0.05$). After pull onset, the proportion of the neurons tuned to the expected trial type was higher among deep layers than among superficial layers (asterisks adjacent to the bars, two-sample z-test for proportions, $p<0.05$). Based on the chi-square test of independence, there was a significant association between the cells' layers and the type of information to which the cells responded during later stages of the pull (asterisks at the top, $df=1$, $p<0.01$).
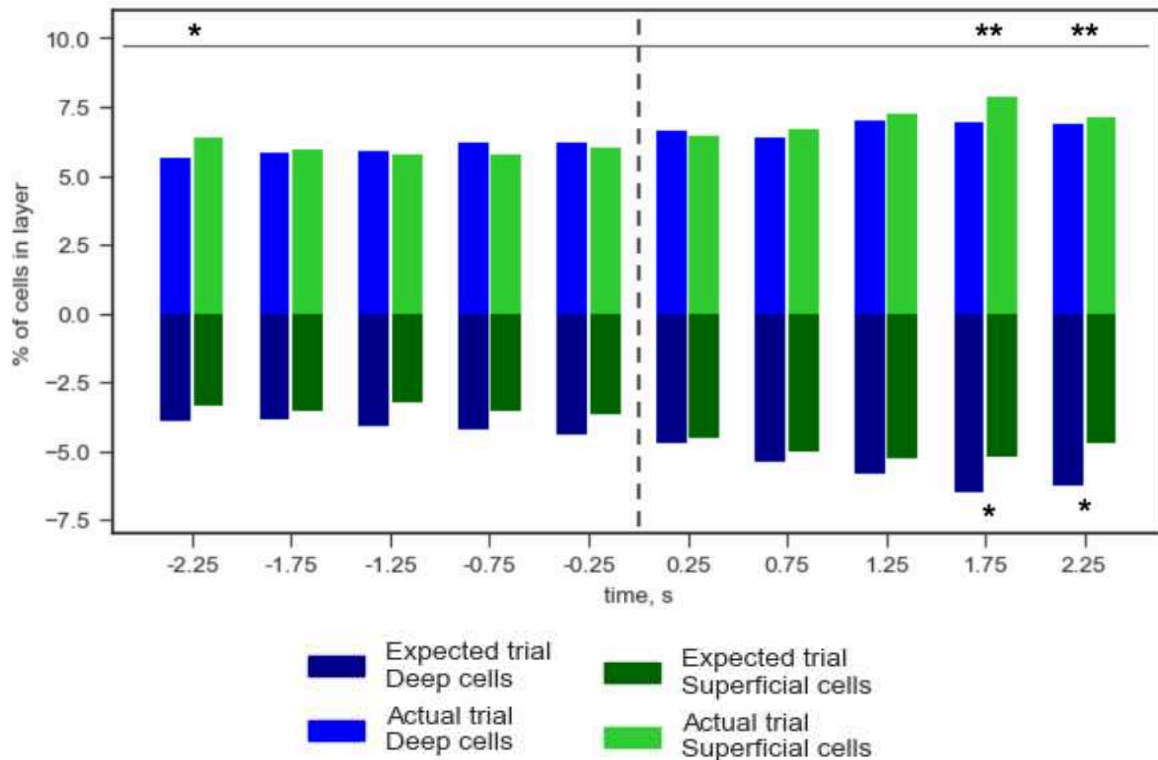
Appendix 2B. Neuronal tuning to light and heavy trials.

The distribution of cells tuned to light and heavy trials, both expected and actual, are shown (unequal variances t-test, $p<0.05$). Among deep neurons there were more cells that increased their activity in response to heavy trials than neurons that responded positively to light trials, both expected and actual. The numbers of superficial and deep cells tuned to both light and heavy expected trials increased after pull onset. There were slightly more superficial cells that were tuned to light actual trials than heavy actual trials.
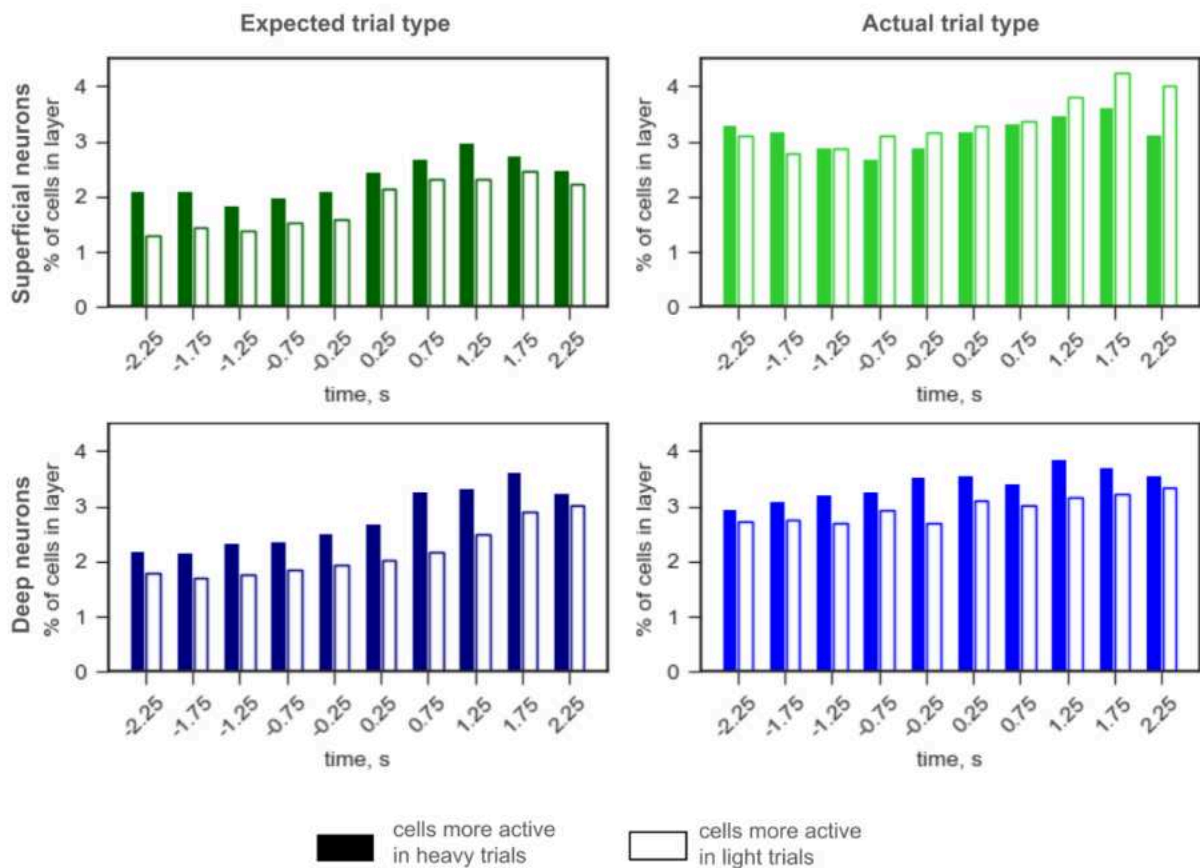
Appendix 3. Prediction performance of expected and actual trial type by layer-specific ensembles within mice.

The amplitude of the after-pull peak was used to train and predict the expected trial type during excluded trials. Rows correspond to three different series of sessions collected from two mice that contained large numbers of observable superficial (green) and deep (blue) cells. Dashed line shows the no-prediction level of predicting the trial type. Per ensemble size, the deep cells from middle and bottom rows better encode the expected trial type, while superficial layers better encode the actual trial type. The data in the first row, however, has the opposite relationship.
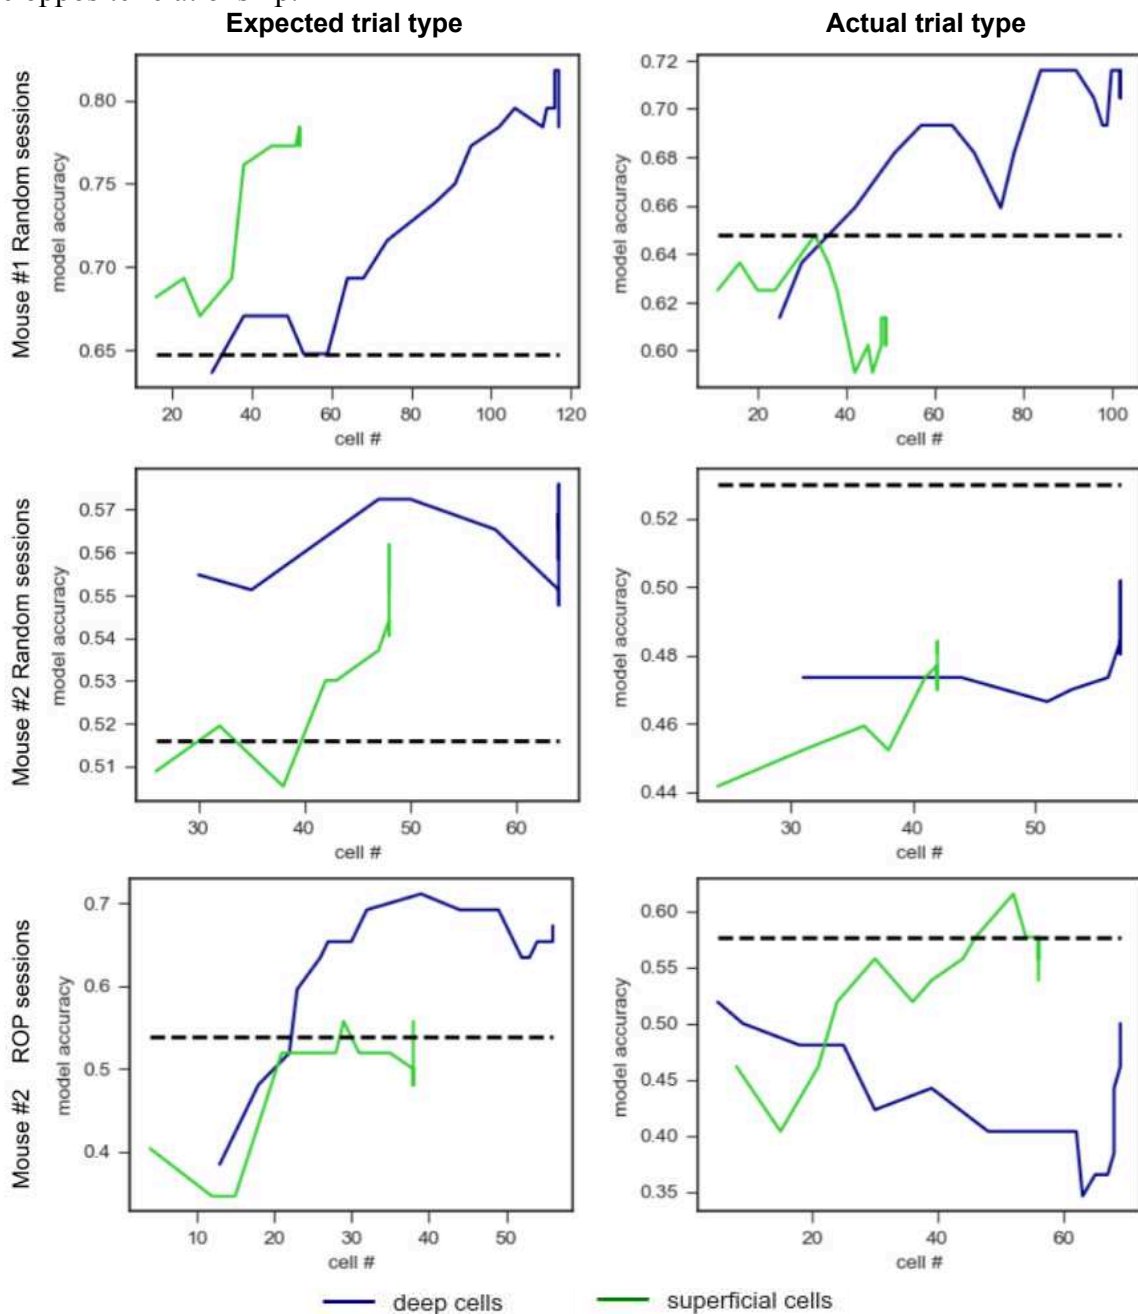
Appendix 4. Binary classification of the odd trials with small time lag during ROP sessions at different timepoints.

The model was trained using the standard trials in ROP sessions and used to predict the trials only when a pull was initiated within 0.8 seconds after trial start.