Technische Universität München
TUM School of Medicine and Health

Behavioural and Neuronal Signatures of Executive Control in Monkey and Mouse Prefrontal Cortex

Daniel Benjamin Hähnke

Vollständiger Abdruck der von der TUM School of Medicine and Health der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Thomas Korn

Prüfer*innen der Dissertation:

1. Prof. Dr. Simon Nikolas Jacob
2. Prof. Julijana Gjorgjieva, Ph.D.

Die Dissertation wurde am 10.07.2023 bei der Technischen Universität München eingereicht und durch die TUM School of Medicine and Health am 08.11.2023 angenommen.

Difficult to see; always in motion is the future.

Yoda
The Empire Strikes Back
George Lucas, 1980

Abstract

Prefrontal cortex (PFC) plays a vital role in goal-directed behaviours, by enabling executive functions like working memory, planning, decision-making, and inhibitory control. This thesis comprises two projects that explored different aspects of prefrontal executive function in conjunction with other brain regions.

The first project investigated how PFC categorises sensory information and maintains it in working memory (WM) for future action, and how coordinated cross-regional neuronal activity enables the storage of multiple WM items while preserving individual representation and protecting against distractions. Using paired recordings in the prefrontal and parietal cortex of monkeys performing a numerosity discrimination task, WM content was found to be organised by frequency-specific oscillatory activity. Beta-band signalling from the parietal to frontal cortex conveyed information about the most recent numerical input, while theta-band coupling from the frontal to parietal cortex differentiated between multiple memorised numerosities. Task-relevant and distracting stimuli were embedded within the spiking activity of single prefrontal neurons, but could be segregated by decoding spikes at specific phases of parietal theta oscillations. These findings suggest that the fronto-parietal network utilises frequency-specific communication channels to facilitate both sequential bottom-up and parallel top-down information transmission, providing a crucial mechanism for safeguarding WM against interference.

The second project examined how sensory information is used by the prefrontal cortex to guide and prepare future action in decision-making in mice. Auditory movement instructions indicated the direction in which animals should rotate a response ball in a two-alternative forced choice task. Prior context cues were either predictive or not for the upcoming movement instruction. Predictive context cues improved task performance and response times. Submovements detected in context epoch ball rotation indicated motor preparation, and their quantity and characteristics suggested a left-default behavior. Electrophysiological recordings from the prelimbic cortex (PL) and mediodorsal thalamic nucleus (MD) were obtained during task performance. Neuronal responses to the non-predictive context cue were immediate, while responses to predictive cues were delayed, indicating differential processing of the context types. Individual neuronal units varied in their context responsiveness, with most responding to the non-predictive cue and fewer to the cue allowing left-response preparation, potentially linked to the observed default behavior. Context information was high in both brain regions, but MD encoded upcoming instructions earlier and for a longer duration compared to PL. Prior contexts minimally influenced neuronal responses to and encoding of movement instruction, but partially shifted the subpopulations of responsive neurons and altered the timing of activation for context-invariant units. These findings suggest that mice employ efficient asymetric behavioural strategies, supported by prefrontal top-down inhibitory control and MD-driven movement preparation. Additionally, despite varying neuronal implementations induced by contexts, population-based mechanisms could provide robust encoding for task-relevant information.

Zusammenfassung

Der präfrontale Cortex (PFC) spielt eine entscheidende Rolle bei zielgerichteten Verhaltensweisen, indem er exekutive Funktionen wie Arbeitsgedächtnis, Planung, Entscheidungsfindung und inhibitorische Kontrolle ermöglicht. Diese Dissertation umfasst zwei Projekte, die verschiedene Aspekte der präfrontalen exekutiven Funktionen in Verbindung mit anderen Hirnregionen untersucht haben.

Das erste Projekt untersuchte, wie der Präfrontalcortex sensorische Informationen kategorisiert und im Arbeitsgedächtnis für zukünftige Handlungen aufrechterhält, und wie koordinierte neuronale Aktivität zwischen verschiedenen Hirnregionen die Speicherung mehrerer Elemente im Arbeitsgedächtnis ermöglicht, während individuelle Repräsentationen erhalten bleiben und vor Ablenkungen geschützt werden. Durch simultane Datenaufnahme im präfrontalen und parietalen Cortex von Affen, die eine Aufgabe zur Unterscheidung von Anzahlen (Numerositäten) durchführten, wurde festgestellt, dass der Inhalt des Arbeitsgedächtnis durch frequenzspezifische oszillatorische Aktivität organisiert ist. Die Signalübertragung im Beta-Band vom parietalen zum frontalen Cortex lieferte Informationen über den aktuellsten numerischen Eingabewert, während die Kopplung im Theta-Band vom frontalen zum parietalen Cortex zwischen den verschiedenen gemerkten Numerositäten unterschied. Für die Aufgabe relevante als auch ablenkende Reize wurden zur gleichen Zeit durch die Raten von Aktionspotentialen einzelner präfrontaler Neuronen repräsentiert, konnten jedoch durch das Dekodieren von Aktionspotentialen zu bestimmten Phasen der parietalen Theta-Oszillationen getrennt werden. Diese Ergebnisse legen nahe, dass das frontoparietale Netzwerk frequenzspezifische Kommunikationskanäle nutzt, um sowohl eine sequenzielle Bottom-up- als auch eine parallele Top-down-Informationsübertragung zu ermöglichen, was somit einen entscheidenden Mechanismus zur Absicherung des Arbeitsgedächtnisses gegen Störungen bereitstellt.

Das zweite Projekt untersuchte, wie sensorische Informationen vom präfrontalen Cortex verwendet werden, um zukünftige Handlungen bei der Entscheidungsfindung von Mäusen zu lenken und vorzubereiten. Auditorische Bewegungsanweisungen gaben die Richtung vor, in der die Tiere einen Antwortball in einer Aufgabe mit zwei Antwortmöglichkeiten drehen sollten. Vorherige kontextuelle Hinweise sagten die bevorstehende Bewegungsanweisung entweder voraus (prädiktiv) oder nicht (nicht-prädiktiv). Prädiktive kontextuelle Hinweise verbesserten die Ausführung und Reaktionszeiten der Aufgabe. In der Kontextphase der Ballrotation wurden Teilbewegungen festgestellt, die auf eine motorische Vorbereitung hinwiesen, und ihre Menge und Merkmale deuteten auf ein standardmäßiges Verhalten nach links hin. Während des Verhaltens wurden elektrophysiologische Ableitungen aus dem prälimbischen Cortex (PL) und dem mediodorsalen Thalamus (MD) durchgeführt. Die neuronale Reaktion auf den nicht-prädiktiven kontextuellen Hinweis erfolgte sofort, während die Reaktionen auf prädiktive Hinweise verzögert waren, was auf eine unterschiedliche Verarbeitung der Kontextarten hinweist. Individuelle Neuronen zeigten unterschiedliche Reaktionen auf den Kontext, wobei die meisten auf den nicht-prädiktiven Hinweis reagierten und weniger auf den Hinweis, der eine Vor-

bereitung der Linksbewegung ermöglichte, was möglicherweise mit dem beobachteten Standardverhalten zusammenhängt. Die Kontextinformation war in beiden Hirnregionen hoch, aber der MD kodierte bevorstehende Anweisungen früher und für eine längere Dauer im Vergleich zum PL. Vorherige Kontexte beeinflussten neuronale Reaktionen auf und die Kodierung von Bewegungsanweisungen minimal, änderten jedoch teilweise die Subpopulationen reagierender Neuronen und änderten den Zeitpunkt der Aktivierung für kontextinvariante Einheiten. Diese Ergebnisse legen nahe, dass Mäuse effiziente asymmetrische Verhaltensstrategien einsetzen, die durch präfrontale top-down inhibitorische Kontrolle und MD-gesteuerte Bewegungsvorbereitung unterstützt werden. Trotz unterschiedlicher neuronaler Implementierungen, die durch den Kontext induziert werden, könnten populationsbasierte Mechanismen eine robuste Repräsentation der für die Aufgabe relevanten Informationen bereitstellen.

Contents

List of Figures

List of Tables

Glossary

2AFC	two-alternative forced choice. 44
AIC	Akaike Information Criterion. 159
alpha	EEG/LFP frequency band from 10 Hz to 12 Hz. 8, 43, 122
ANOVA	analysis of variance. 47–49, 54–63, 75, 84, 139, 159
beta	EEG/LFP frequency band from 12 Hz to 32 Hz. iii, 8, 14, 25–32, 34–37, 39, 41–43, 119–122, 132, 161, 162
BIC	Bayesian Information Criterion. 156
CDF	cumulative distribution function. 76, 91, 93, 94, 96, 97
CI	confidence interval. 71, 72, 74, 113, 115, 165, 166
CNV	contingent negative variation. 6
delta	EEG/LFP frequency band from 2 Hz to 4 Hz. 25, 27–32, 34, 36, 39, 41–43, 120–122, 162
dlPFC	dorsolateral prefrontal cortex. 11
DMS	delayed match-to-sample. 6, 7, 11, 19
DR	delayed response. 6, 11
DREADD	designer receptors exclusively activated by a designer drug. 14
ECoG	electrocorticography. 8
EEG	electroencephalography. 6, 8, 10
ERP	event-related potentials. 139, 140
FEF	frontal eye field. 18
gamma	EEG/LFP frequency band from 32 Hz. 8, 25, 27, 29, 34, 37, 119, 122
GC	Granger Causality. 32, 112–116, 165, 166
GP	Gaussian Process. 155

hodology	Study of pathways. Also called connectomics.. 12, 18
IPS	intraparietal sulcus. 12, 13, 137, 138
IQR	inter-quartile range. 160
IT	inferior temporal area. 11
KDE	kernel density estimate. 84, 85, 160
LFP	local field potential. 8–10, 19, 25–29, 32, 34–42, 119, 121–123, 138–140, 143, 144, 161, 162
lPFC	lateral prefrontal cortex. 11
LSM	left submovement. 56–63
M2	secondary motor cortex. 18
MD	mediodorsal thalamic nucleus. iii, 14–16, 18, 21, 65, 66, 68–75, 77–86, 88–104, 106–110, 112–115, 117, 124, 127, 128, 131, 132, 135, 163, 165, 166
MEG	magnetoencephalography. 8, 10
mPFC	medial prefrontal cortex. 14, 16, 17, 21, 129, 132
MST	medial superior temporal area. 12
MT	middle temporal area. 12
PFC	prefrontal cortex. iii, 5–8, 10–15, 18–21, 24–36, 38–42, 68, 120–122, 125, 128–130, 132, 137, 138, 144, 162
PL	prelimbic cortex. iii, 65–68, 70–75, 77–86, 88–99, 101–110, 112–115, 117, 124, 127, 131, 132, 135, 163, 165, 166
PLSI	phase-locking selectivity index. 36, 141
PLV	phase-locking value. 29, 30, 32, 34–37, 39, 42, 140, 141, 161
PPC	posterior parietal cortex. 12, 19
PSI	phase-slope index. 30–32, 142, 161
RBF	Radial Basis Function. 155
RSM	right submovement. 56–63
SEM	standard error of the mean. 28, 38–42, 71, 73, 143, 144, 162
SF	spike-field. 32, 34, 36, 37
SM	submovement. 52–54, 56–58, 60, 62, 64
SNr	substantia nigra pars reticulata. 14
SPL	sound pressure level. 133

SVM — Support-Vector Machine. 98, 99, 101, 102, 104, 106

theta — EEG/LFP frequency band from 4 Hz to 10 Hz. iii, 8, 25–27, 29, 32, 34–37, 39, 41–43, 120–122, 161, 162

ToL — Tower of London. 14

V4 — visual area V4. 10, 121

VIP — ventral intraparietal cortex. 12, 13, 19, 20, 24–34, 36–39, 41, 42, 120–122, 137, 138, 144, 162

VTA — ventral tegmental area. 14, 17

WCST — Wisconsin Card Sorting Test. 5, 14

WM — working memory. iii

ZETA — Zenith of Event-based Time-locked Anomalies. 76, 78, 79, 81, 82, 91

Part I

Introduction

1 Introduction

1.1 Perception-Action Cycle and Goal-directed Behaviour

Organisms interact with their environment in an endless cybernetic loop of sensation and action. In its simplest form the loop comprises an organism that senses external stimuli from the environment, leading to an action by the organism, thus modifying the environment. The modified environment leads to a modified stimulation of the organism, and so on (Fig. 1.1a). When involving the cerebral cortex, sensations lead to perceptions and the loop becomes the perception-action cycle.

A key feature of animals with more complex brains is their ability to show goal-directed behaviours which transcend simple stimulus-response associations and innate behaviours. Goal-directed behaviours intuitively integrate into the cycle of reciprocal interaction between organisms and their environment (Fig. 1.1b). Here, a perception doesn't trigger an immediate action anymore but is processed in the light of internally represented future goals. In order to attain these goals, individual actions are then orchestrated in new ways.

As an example, consider something as simple as perceiving the presence of potential food. A simple response might be to immediately eat it. However, dependent on internal goals, with potentially different scopes, this might not be the most optimal action. For example, one might defer immediate eating and temporarily store the food in order to satiate expected future hungriness. In a different context abstract rules could require holding off until certain other factors are fulfilled, e.g. so as to display characteristics desirable by other members of a social group one wants integrate in (in order to gain future benefits of being member of that group). Regardless of the if and when, the act of eating might require a more or less complex sequence of actions, e.g. to extract edible parts. For this, an internal representation of the food's anatomy, acquired by similar prior experiences and updated by the currently perceived instance, is beneficial. This representation can be manipulated and used to decide which actions in which temporal order are most favourable for the attainment of a goal.

1.2 Executive Functions

Goal-directed behaviours are made possible by a set of cognitive processes called *executive functions*. Here, the term *executive* means that these functions are suited to support the organisation of action sequences used to carry out (i.e. execute) deliberate plans (Norman & Shallice, 1986; E. K. Miller & Cohen, 2001; Tanji & Hoshi, 2008; E. K. Miller & Wallis, 2009; Diamond, 2013). They are the components that integrate into

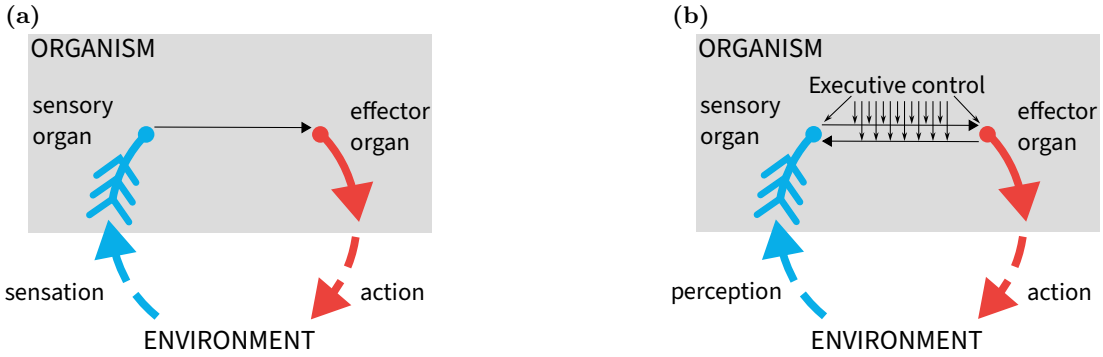

Figure 1.1: The cybernetic loop between organisms and their environment. (a) Simple stimulus-response mapping within the sensation-action cycle. After von Uexküll (1926). **(b)** Executive control at the top of the perception-action cycle allows organisms to organise behaviour according to internal goals. Executive control does not only influence action but also perception. After Fuster (2022).

the perception-action cycle, mainly by influencing actions after perception but also by influencing perception itself (Fig. 1.1b). Typically, goal-directed behaviour is dependent on the orchestrated interplay of several executive functions. Thus, no executive function works in isolation. It is nevertheless useful to conceptualise them into distinct components. Dependent on the goal and the current progression towards it, the deployment of different components is more or less pronounced. In the following the executive functions that are especially important for this thesis will be briefly introduced.

Working memory The ability to temporarily persist information in the absence of external stimuli is central to goal-directed behaviour (A. D. Baddeley, 1983). Without it, only simple stimulus-response associations with stereotyped behaviour would be possible. In contrast, working memory fills the temporal gap between stimuli and delayed actions, providing a flexible substrate to guide behaviour according to internal plans (A. Baddeley, 2012; Fuster, 2015). Importantly, working memory as it is thought of today is not simply short-term memory on its way to long-term storage. Rather, it is an active process that likely involves the activation of preexisting long-term memories of percepts along with those of rules and possible actions (Fuster, 2022). This co-activation, malleability and future-directedness by goals are important characteristics that distinguish working memory from simple short-term memory.

As a practical example, consider being told the address of a location. This might activate memories of the location's neighbourhood, along with memories of going there previously (possible actions) and rules (traffic regulations, urgency of arriving).

Planning The previous example of working memory will naturally transition into a phase of planning ("Which streets do I have to go along?", "When do I need to take a turn and to what side?", ...).

Planning is thus the process of laying out an ordered sequence of temporally separated actions for attaining a goal (Diamond, 2013; Fuster, 2015). That is, the process of planning creates new cross-temporal contingencies between the goal and expected sensory inputs as well as individual acts along the way. The individual components of a plan can be rooted in memory, and thus, planning can be conceptualised as a rearrangement of these memories with new objectives, orders and time schedules. In addition, plans can be recalled and retained and, as such, planning has been aptly referred to as the "memory of the future" (Ingvar, 1985).

Decision-making Similar to the proverbial truth that many roads lead to Rome, there are many possible sequences of action to attain a goal. Intimately linked to planning, decision-making is the selection of a plan or action amongst multiple possible alternatives, with the intention of execution (Fuster, 2015). During this process, different factors like basic biological drives, expected but uncertain rewards and costs need to be reconciled in order to arrive at a suitable decision (Tanji & Hoshi, 2008).

Continuing with the previous example, possible factors that influence the choice of route would be expected benefits ("Do I want to arrive fast or enjoy a walk?"), costs and consequences ("What if I arrive late?").

Attentional set Cognitive resources are finite and need to be allocated to relevant processes to most efficiently serve an animal for its survival. Attentional set primes sensory and perceptual structures in expectation of sensory stimulation (perceptual set), as well as motor and higher executive structures in the preparation of an action or a plan (executive set) (Tanji & Hoshi, 2008). Intuitively, a concentration of resources always brings about a selection of a smaller number of effectors at the cost of others. Thus, attentional set has both facilitatory and suppressive aspects.

A common life example of this is the failure of immediately recognising people in uncommon contexts. Concretely, if one is in a leisure time context meeting colleagues from work might be surprising because perceptual set would have increased expectation of one group of persons while suppressing expectations of others.

Inhibitory control All functions mentioned above entail or require some degree of suppression of inappropriate processes (Fuster, 2015; Diamond, 2013). Working memory, for example, is a function that focuses on relevant internal representations and thus is inherently susceptible to interference by distracting external stimuli. Here, interference can originate from very salient but irrelevant stimuli or from stimuli that elicit percepts that are similar to current memoranda. Since working memory has been suggested to involve activation of previous long-term memories (Fuster, 2022), it must also be resilient to activation of previously relevant information that in the current context is inappropriate.

Lack of this cross-temporal distractor resilience leads to perseverant behaviour, as it is famously tested in the Wisconsin Card Sorting Test (WCST) (Milner, 1963). In the test, the subjects need to sort cards according to shape, quantity or colour of symbols printed

on the cards. The rule that governs the sorting is changed periodically. Thus, rule changes require inhibition of prior rules and a shift to attend to the currently relevant feature.

Another standard test to assert inhibitory control is the Stroop task (Stroop, 1935). Here, subjects are asked to either read colour words or name the colour in which they are presented. If the word and the colour of the letters conflict it is especially difficult to name the colour. This is because when seeing words we learned to attend the meaning of the letters and ignore their colour. Reversing the rule thus requires the suppression of strong habitual but inappropriate responses and an attentional shift towards a less salient feature.

1.3 Prefrontal Cortex: Neuronal Substrate of Executive Functions

Executive functions integrate multisensory information and internal states to guide behaviour. Thus, a structure that supports these functions needs to be connected to

1. sensory areas to receive sensory input and influence perception,
2. subcortical regions to receive input about internal states like motivation, affect and other biological drives, and
3. motor areas to implement and execute plans of action.

The prefrontal cortex (PFC) receives visual, auditory and somatosensory input from association cortices in parietal, temporal and occipital lobes (Barbas et al., 2011) and is also connected to motor structures like supplementary motor area, cerebellum and superior colliculus (E. K. Miller & Cohen, 2001). Most of these connections are reciprocal and not to primary sensory or motor areas but association areas. In addition, the prefrontal cortex (PFC) is connected to many other subcortical structures, either directly or via thalamus, including hypothalamus, subthalamus, mesencephalon, amygdala and hippocampus (E. K. Miller & Cohen, 2001; E. K. Miller & Wallis, 2009; Tanji & Hoshi, 2008; Fuster, 2015). Thus, it receives information about movement execution and coordination, internal state and motivation. Together, the numerous reciprocal connections to these brain structures make the PFC a prime candidate to be involved in executive functions (Figs. 1.2a and 1.2b).

Indeed, lesion studies suggested the PFC to be prominently involved in executive functions. Probably the most famous case is that of a man called Phineas Gage, who in the mid-19$^{\text{th}}$ century suffered from impaired planning and impulse control after an iron bar had penetrated his skull during a work accident, lesioning most of his left orbitomedial prefrontal cortex (Harlow, 1868). More intentional and scientifically rigorous studies with more focused lesions supported the notion that PFC is critical for behaviours involving working memory (Malmo, 1942; Spaet & Harlow, 1943; Mishkin, 1957; Glick et al., 1969; Goldman & Rosvold, 1970; Levy & Goldman-Rakic, 1999; Milner et al., 1985; Marshuetz, 2005), distractor resilience (Malmo, 1942; Chao & Knight, 1995; Stuss

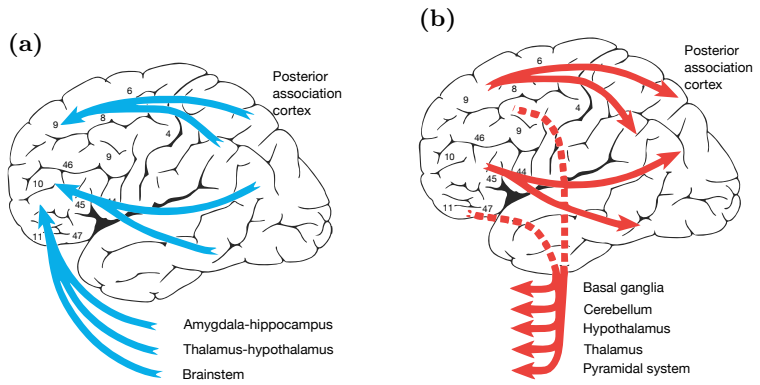

Figure 1.2: Prefrontal connectivity. Lateral aspect of the human brain. Afferents to (a) and efferents from (b) PFC. Vectorised from Fuster (2015).

et al., 1982), attentional set (Milner, 1963; Milner, 1964; Dias et al., 1996), inhibitory control (Stuss et al., 1982; Chao & Knight, 1995; Mishkin, 1957; Butter, 1969; Goldman & Rosvold, 1970; Dias et al., 1996; Dias et al., 1997; Vendrell et al., 1995), planning (Goel & Grafman, 1995; Volle et al., 2011; Shallice, 1982; Karnath & Wallesch, 1992), and decision-making (Levy & Goldman-Rakic, 1999; Bechara et al., 1994; Goldberg & Podell, 2000; Baudry & Monperrus, 2022).

Neurophysiology While lesion studies were useful to demonstrate the involvement of PFC in executive function, they are undoubtedly limited in terms of temporal precision, validity in non-pathological conditions, and the discovery of neural mechanisms. Electrophysiological recordings of prefrontal activity at different spatial levels have been of paramount importance to characterise the neurophysiology associated with executive functions.

The earliest studies which used electroencephalography (EEG) in a cognitive paradigm found that human frontal cortex generates a characteristic negative evoked potential during the delay between a "warning" stimulus and an "imperative" stimulus that required an action (Walter et al., 1964; Walter, 1964). Importantly, this so-called contingent negative variation (CNV) occurred irrespective of the sensory modality (visual, auditory, tactile) and only if the paradigm involved an upcoming action, i.e. the stimuli were associated with goal-directed behaviour. The CNV is generally considered to be a manifestation of selective attention and expectancy (Tecce, 1972; Chennu et al., 2013; Holmes et al., 2021).

Soon after the discovery of the CNV single neuronal units in the prefrontal cortex of monkeys were found to have elevated activity during the presentation of behaviourally relevant stimuli or the ensuing delay epoch in delayed response (DR) (Fuster & Alexander, 1971) and delayed match-to-sample (DMS) tasks (Fuster et al., 1982). Again, this cue-invariant sustained activity can be interpreted as a reflection of general behavioural relevance of stimuli and attention.

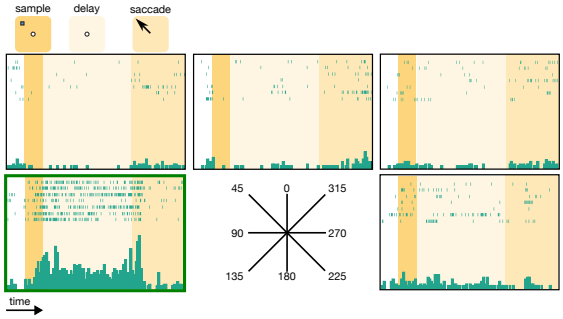

Figure 1.3: Delay epoch activity of a prefrontal neuron in an oculomotor delayed response task. Top: At the beginning of the task one of nine visuospatial targets is shown on a screen. After this a delay epoch follows, during which the target location must be remembered. At the end of the delay epoch a go cue instructs the animal to make a saccade to the cued location on the screen. Bottom: Spike rasters and peri-stimulus time histograms for five target locations are shown (135° to 225° omitted for brevity). The example neuron specifically elevates its firing rate in trials when the target is at 90° (plot with green border). This can be interpreted as the a neuronal representation of target memory and its contingent action plan. Modified from Arnsten (2009) after a classical experiment by Funahashi et al. (1989).

Behavioural tasks that require the retention of some properties of distinct external stimuli for later use, e.g. delayed comparison with test stimuli (DMS), typically find an additional kind of prefrontal neuronal activity: cells whose delay activity varies with different stimuli, i.e. they *encode* a memorandum (Fig. 1.3). In this respect PFC can encode manifold properties of stimuli of diverse modalities, as long as the properties are important for solving the task, e.g. colour (Fuster et al., 1982; Kubota et al., 1980; M. Watanabe, 1986), location (Rainer et al., 1998), motion (Zaksas & Pasternak, 2006), shape (E. K. Miller et al., 1996), vibration frequency (Romo et al., 2002), tone pitch (Fuster et al., 2000), category (Freedman et al., 2001), and numerosity (Nieder et al., 2002). Thus, this prefrontal delay activity is a neuronal correlate of working memory which is required to achieve a cross-temporal transfer of task-relevant information.

Whereas tasks involving working memory typically require PFC, their neural correlates are also found in many areas of posterior cortex (Christophel et al., 2017; Bisley et al., 2004; Chelazzi et al., 2001; Fuster & Jervey, 1982). Importantly, though, working memory information in those areas is generally specific to modality or certain properties that the respective area is specialised in. Conversely, PFC activity is, by and large, invariant and thus represents the highest abstractions. This suggests that PFC plays a central role in working memory by its interaction with other cortices, by coordinating their activation.

These interactions have been shown, for instance, by reversibly decreasing neuronal activity through cooling either PFC or another cortex. In addition to generally decreasing delay activity in the other, non-cooled cortex, this also decreases working memory information, for example in a DMS task requiring the retention of colour information (Fuster et al., 1985). This has often been interpreted as evidence for recurrent excitation between PFC and the respective association areas, leading to reverberating delay activ-

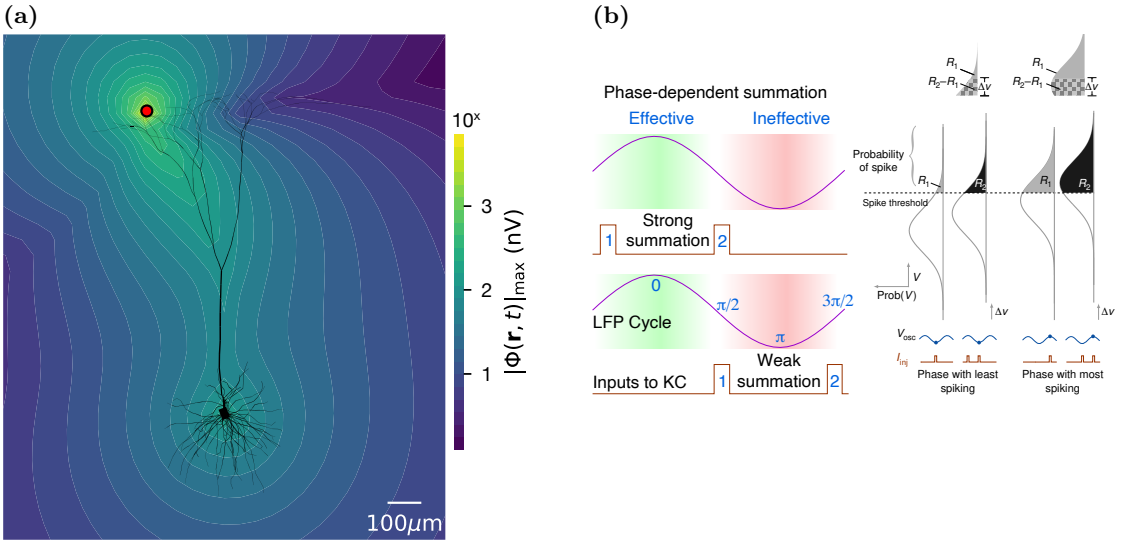

Figure 1.4: Extracellular potentials, single cell (a) Extracellular potential in the vicinity of a pyramidal neuron after a single excitatory postsynaptic current at the location of the red dot. Modified after Lindén et al. (2014, Fig. 3) via Binder. **(b)** Oscillations of the extracellular potential lead to windows of effective and ineffective integration of synaptic input (*left*) and, thus, to increased or decreased spiking probabilities (*right*). Copied from Gupta et al. (2016).

ity (Amit, 1995; Goldman-Rakic, 1995; X. J. Wang, 2001). Indeed, many computational studies support this idea using recurrently connected spiking units that can sustain information during delays without external stimulation (Zipser et al., 1993; Durstewitz et al., 2000; Carter & Wang, 2007; Pulvermüller & Garagnani, 2014; Bouchacourt & Buschman, 2019; J. X. Wang et al., 2018).

Oscillations It is plausible that re-entry between PFC and associated structures during working memory synchronously activates large groups of neurons, thus leading to oscillatory activity (for an excellent review of possible mechanisms see X.-J. Wang, 2010).

In brief, every synaptic input leads to currents across the neuronal membrane (Kandel et al., 2021). This not only changes the voltage potential of the neuronal membrane but the entire neighbourhood (Fig. 1.4a). Parallel orientation of large numbers of neurons and synchronous synaptic input to them subsequently leads to a summation of the extracellular potential (Fig. 1.5) which can be recorded at different spatial resolutions in the EEG, electrocorticography (ECoG) and local field potential (LFP) (Fig. 1.6).

Indeed, studies investigating executive functions commonly find task-related modulation of prefrontal oscillatory activity as measured by the LFP, EEG and magnetoencephalography (MEG) (Roux & Uhlhaas, 2014). In humans, frontal EEG/MEG power is modulated at theta, alpha, beta or gamma frequency bands (Fig. 1.7) during retention of verbal (Jensen & Tesche, 2002; Onton et al., 2005; Meltzer et al., 2008), somatosensory (Spitzer et al., 2010), and visual information (Tallon-Baudry et al., 1999; Roberts et al.,

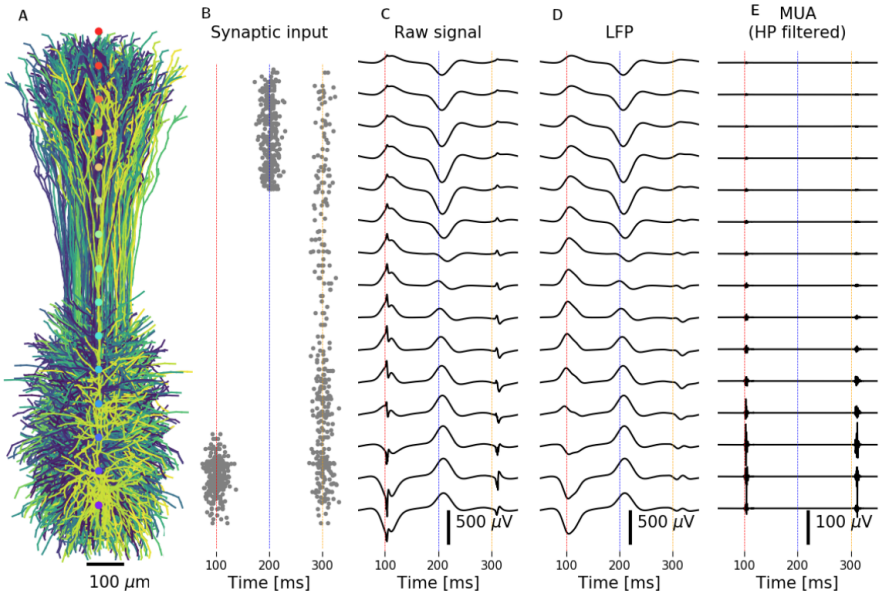

Figure 1.5: Extracellular potentials, neuronal population. Simulation of how synaptic input leads to extracellular potentials. A: Column of 10 000 cortical layer V pyramidal cells with recording locations represented by dots along the vertical axis. B: Synaptic input at different locations and times. C: Resulting raw signal after synaptic inputs. D: Corresponding lowpass filtered LFP signal. E: Corresponding highpass filtered MUA spiking activity. Copied from Ness et al. (2021, Fig. 4).

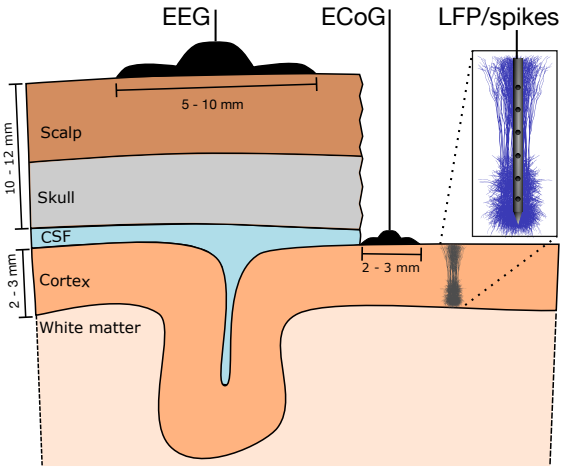

Figure 1.6: Sources of neural signals. Schematic cross-section through human scalp, skull and superficial brain to show the recording locations of brain signals that capture summed oscillatory activity. EEG: electroencephalography; ECoG: electrocorticography or intracranial EEG; LFP: local field potentials. Modified after Ness et al. (2021, Fig. 1).

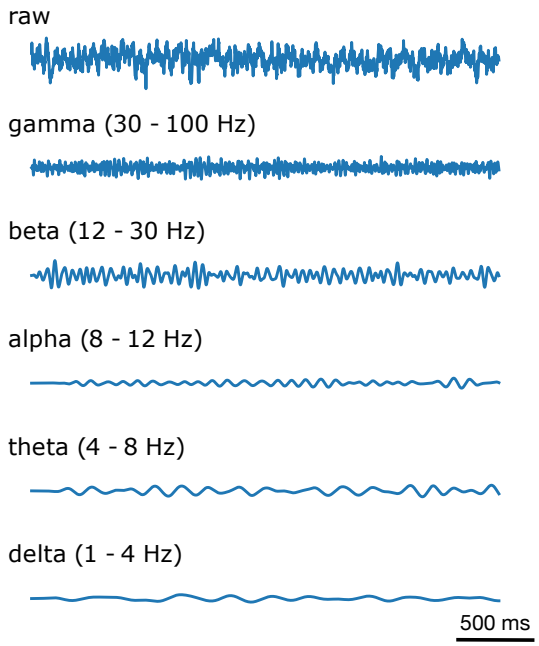

Figure 1.7: Frequency bands. Raw LFP and signals derived by bandpass filtering within commonly used frequency bands in LFP, EEG and MEG research.

2013). Similarly, in monkeys working memory tasks modulate prefrontal LFP power at various frequencies (Liebe et al., 2012; Lundqvist et al., 2016; Buschman & Miller, 2007).

In addition to intra-regional effects the relationship between oscillations in different regions has been of particular interest. Oscillations of the LFP lead to modulation of local neurons' excitability and thus to windows of opportunity for integration of information (Gupta et al., 2016, Fig. 1.4b). In the recent years evidence accumulated for the *communication through coherence* hypothesis (Fries, 2005) which posits that distant brain regions can communicate more efficiently during windows of enhanced excitability by alignment (i.e. synchronisation) of their oscillatory activity. For instance, working-memory related synchronisation has been found between PFC and visual area V4 (V4) (Liebe et al., 2012), parietal cortex (Sauseng et al., 2005; Salazar et al., 2012; Antzoulatos et al., 2016; Buschman & Miller, 2007), temporal cortex (Fiebach et al., 2006), hippocampus (Brincat & Miller, 2015), and other brain regions (Helfrich & Knight, 2016). Similar to modulation of local power, cross-regional phase-synchronisation occurs at various frequencies, possibly reflecting distinct resonant frequencies of different neuronal networks activated by different attributes of a memorandum, as well as possible upcoming actions (X.-J. Wang, 2010; Fuster, 2015).

Effects of interference/distraction Cognitive functions that are temporally detached from external stimuli require internal representations of memoranda and/or plans. By some viewed as essentially being forms of attention (Fuster, 2015), these self-sustained processes are prone to distraction by external and internal interference.

Effects of interference on planning and working memory can be shown in DR and DMS tasks. For instance, an early study by Fuster (1973) using a DR task showed unstructured external visual and auditory distractors during the delay epoch decreased firing of prefrontal units in monkeys. Multi-tasking can also be regarded as a source of interference: Embedding a DR task within a change detection task diminished information about the upcoming response location in monkey lateral prefrontal cortex (lPFC) (K. Watanabe & Funahashi, 2014).

In contrast, a study using a visuospatial DR task with intervening distractor locations found suppression of distractor information in dorsolateral prefrontal cortex (dlPFC) but not parietal cortex (Suzuki & Gottlieb, 2013). Parthasarathy et al. (2017) added to this the finding that the presentation of the distractor severely changed the neuronal population code in dlPFC across the two delay epochs. Further analyses of this dataset found that information stability and consequently resistance to interference was achieved within a neuronal subspace that was consistent throughout the trial (Parthasarathy et al., 2019).

Resilience of PFC to distractors was also found in working memory and selective attention tasks. For example, in a DMS task involving the matching of visual shapes the delay epoch presentation of non-matches did not affect activity of PFC neurons but inferior temporal area (IT) activity was disturbed (E. K. Miller et al., 1996). Additionally, PFC seems to suppress responses to distracting stimuli in sensory cortices, as for example focal dlPFC lesions lead to larger responses evoked by auditory distractors in auditory cortex (Chao & Knight, 1998). Similarly, in a visuospatial DMS task the presentation of non-match locations left dlPFC neurons largely unaffected (Qi et al., 2010). Consistent with this notion of distractor resilience in PFC, in visual search and change detection tasks PFC neurons did not represent information about the distracting item (Lennert & Martinez-Trujillo, 2011; Cosman et al., 2018).

However, some studies using different stimuli found notable exceptions. In a study that required human participants to match faces in a DMS task the presentation of interfering faces during the retention period did not affect PFC encoding of the sample face. Conversely, presentation of categorically different pictures (natural scenes) selectively degraded sample information (Yoon et al., 2006). PFC has also been shown to encode both relevant and irrelevant, i.e. distracting, information at the same time. For example, in a DMS task that required monkeys to remember the number of items (numerosity) in a visual display PFC encoded both the sample and interfering numerosities (Jacob & Nieder, 2014, Fig. 1.10a). Another task that required the discrimination of stimuli based on one of two features whose relevance was signalled by a context cue found the simultaneous encoding of both relevant and irrelevant features in dlPFC via a multidimensional code (Aoi et al., 2020).

In sum, keeping in mind those exceptions it appears that PFC is mostly robust to interfering information. If distractors are encoded, different mechanisms ensure that relevant information is preserved.

1.4 Related Brain Structures Investigated in this Thesis

Executive functions are inherently integrative. Therefore, as a key structure that enables those functions PFC cannot work in isolation but requires input from and exerts influences onto other brain regions. In the following, two areas that are strongly associated with PFC and that are important for this thesis will be briefly introduced.

1.4.1 Ventral intraparietal cortex

In the macaque brain the ventral intraparietal cortex (VIP) is located in the fundus of the intraparietal sulcus (IPS) and thus part of posterior parietal cortex (PPC) (Fig. 1.8a). While the location of a homologue in humans has been less clear, in a recent review Foster et al. (2022) suggested the presence of a human "VIP complex" comprising three parietal areas based on functional homologies.

Initially, macaque VIP was found to receive inputs from middle temporal area (MT) and medial superior temporal area (MST) (Maunsell & van Essen, 1983; Boussaoud et al., 1990). Later retrograde tracing found that in addition to those visual association areas VIP receives input from numerous areas of posterior cortex associated with somatosensory, vestibular, auditory and multisensory processing, as well as from frontal cortex including motor association and prefrontal cortices (J. W. Lewis & Van Essen, 2000). Consistent with what this hodological data suggests, VIP has been shown to be involved in unimodal higher sensory processing, multisensory integration, and notably, cognitive functions (Foster et al., 2022).

Within the latter, ventral intraparietal cortex (VIP) is especially relevant for processing of discrete quantities, namely numerosities and enumeration, i.e. the quantity of concurrent and sequential items ("counting"), respectively. Monkeys can memorise the number of dots in a visual display and match it to a test quantity (Fig. 1.8b; Nieder et al., 2002; Nieder & Miller, 2004a). Besides prefrontal cortex (PFC), VIP contains an especially high number of neurons whose firing rates during stimulus presentation and memory delays systematically vary with numerosity (Fig. 1.8c; Nieder & Miller, 2004b). Furthermore, VIP neurons are spontaneously selective for numerosity in naïve animals that were not trained to discriminate numerosities (Viswanathan & Nieder, 2013) and even encode empty sets (Ramirez-Cardenas et al., 2016). Hierarchically, VIP appears to be at an earlier stage of number processing, as neurons in VIP have a shorter latency than PFC to encode numerosity, and are more sensitive to lower level visual properties (Nieder & Miller, 2004b). In addition, VIP neurons also encode the accumulated number of visual or auditory items shown in a sequence, but compared to PFC a smaller number is invariant to the sensory modality, suggesting a lower level of abstraction (Nieder et al., 2006; Nieder, 2012).

All in all, there is a large body of evidence that suggests VIP is part of a fronto-parietal network that is critically important for temporal integration of quantity, two aspects of which are temporal summation and cross-temporal bridging via working memory.

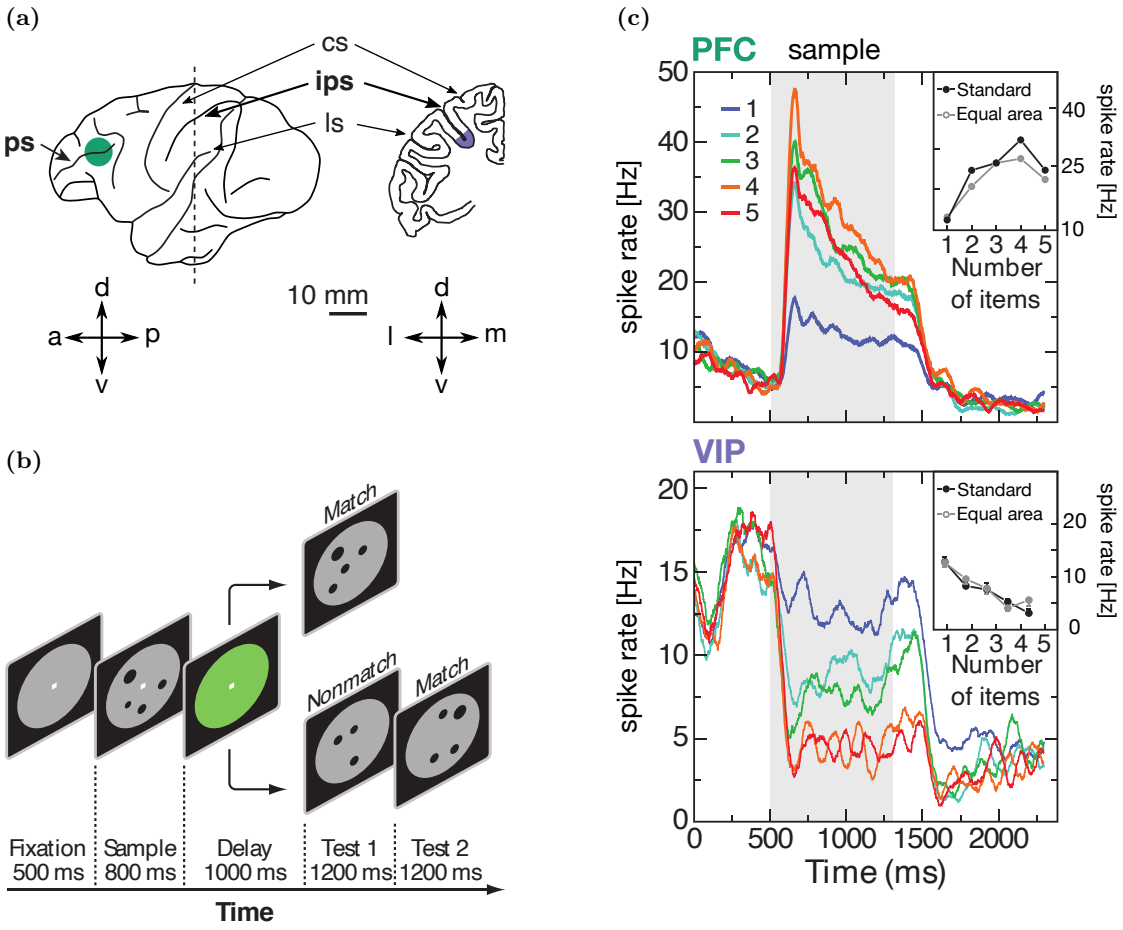

Figure 1.8: Numerosity coding in PFC and VIP. (a) Locations of PFC and VIP in the macaque brain. (Left): Lateral view. PFC is labelled in green. Dashed line: Position of coronal slice shown on the right. (Right): Coronal slice. Fundus of intraparietal sulcus (IPS) (i.e. VIP) is labelled in purple. cs: central sulcus; ips: intraparietal sulcus; ls: lateral sulcus; ps: principal sulcus; a: anterior; p: posterior; d: dorsal; v: ventral; l: lateral; m: medial. Modified from Jacob et al. (2018). **(b)** Delayed match-to-numerosity task. Modified from Nieder and Dehaene (2009). **(c)** Firing rates of example neurons varied with the shown numerosity. (Top): Example neuron from PFC. (Bottom): Example neuron from VIP. Modified from Nieder and Dehaene (2009).

1.4.2 Mediodorsal thalamic nucleus

PFC has commonly been defined as the cortex projected to by the mediodorsal thalamic nucleus (MD) (Rose & Woolsey, 1948; Uylings & van Eden, 1991; Fuster, 2015). Reciprocal connections between the two areas are strong and are in fact PFC's most prominent ones to subcortical areas (Fuster, 2015).

In contrast to the classical view of thalamus as a simple relay of sensory information on its way to neocortex (Fig. 1.9a; Hubel & Wiesel, 1979), MD with its main driver inputs from cortical layer 5 neurons has more recently been designated a higher-order relay nucleus that is used for trans-thalamic cortico-cortical communication (Fig. 1.9b; Sherman & Guillery, 1998; Sherman, 2016). Notably, in the latter view, the presence of parallel trans- and extra-thalamic cortico-cortical pathways has been interpreted as efference copies. Similar to PFC, MD receives input from numerous subcortical areas, namely from the reticular thalamic nucleus, and structures in the striatum (dorsal and ventral striatum), midbrain (ventral tegmental area (VTA), substantia nigra pars reticulata (SNr)), hypothalamus, and brainstem (Mitchell, 2015; Ouhaz et al., 2018) and thus has access to behavioural parameters like motivation, state and drive.

Lesions of MD often lead to behavioural deficits similar to PFC lesions. Due to its location deep in the brain, studies of human MD are limited to lesions inflicted by infarctions. Standard testing batteries of executive functions (including tests of attention, working memory, interference control and planning like WCST, Tower of London (ToL) and Stroop) commonly find impairments in patients whose MD is damaged (Van der Werf et al., 2003; Carlesimo et al., 2011; Danet et al., 2015). Consistent with these findings, animal studies, which allow for targeted lesions or inactivation of MD found deficits in tasks requiring response alternation (Isseroff et al., 1982; Chauveau et al., 2005), reversal learning (Dolleman-van der Weel et al., 2009; Chudasama et al., 2001; Parnaudeau et al., 2013) or some sort of working memory (Isseroff et al., 1982; Stokes & Best, 1990; Parnaudeau et al., 2013; Bolkan et al., 2017).

Connection-specific analyses and experiments demonstrate the close relationship between MD and PFC. For example, Parnaudeau et al. (2013) found in mice that MD neurons had increased phase-locking to medial prefrontal cortex (mPFC) beta oscillations during the delay epoch of a T-maze task. Further, a reversible selective inhibition of MD using the designer receptors exclusively activated by a designer drug (DREADD) system led to a decrease in beta band phase-locking selectively during the delay epoch. In a later study by the same research group using the same behavioural paradigm direction-specific optogenetic inhibition was used to further dissect reciprocal communication of MD and PFC (Bolkan et al., 2017). Inhibition of MD terminals in mPFC or vice-versa during an entire trial decreased behavioural performance in trials with long, but not short delays, suggesting a role of both regions for cross-temporal retention of information. In addition, trial-epoch specific inhibition of the two pathways showed that reciprocal MD-mPFC communication during the delay epoch, and unidirectional influences from mPFC during the choice epoch, are important for successful behaviour. Electrophysiological recordings showed that functional directionality shifts in an epoch-dependent manner from MD-to-mPFC during the delay to mPFC-to-MD during the choice. Fi-

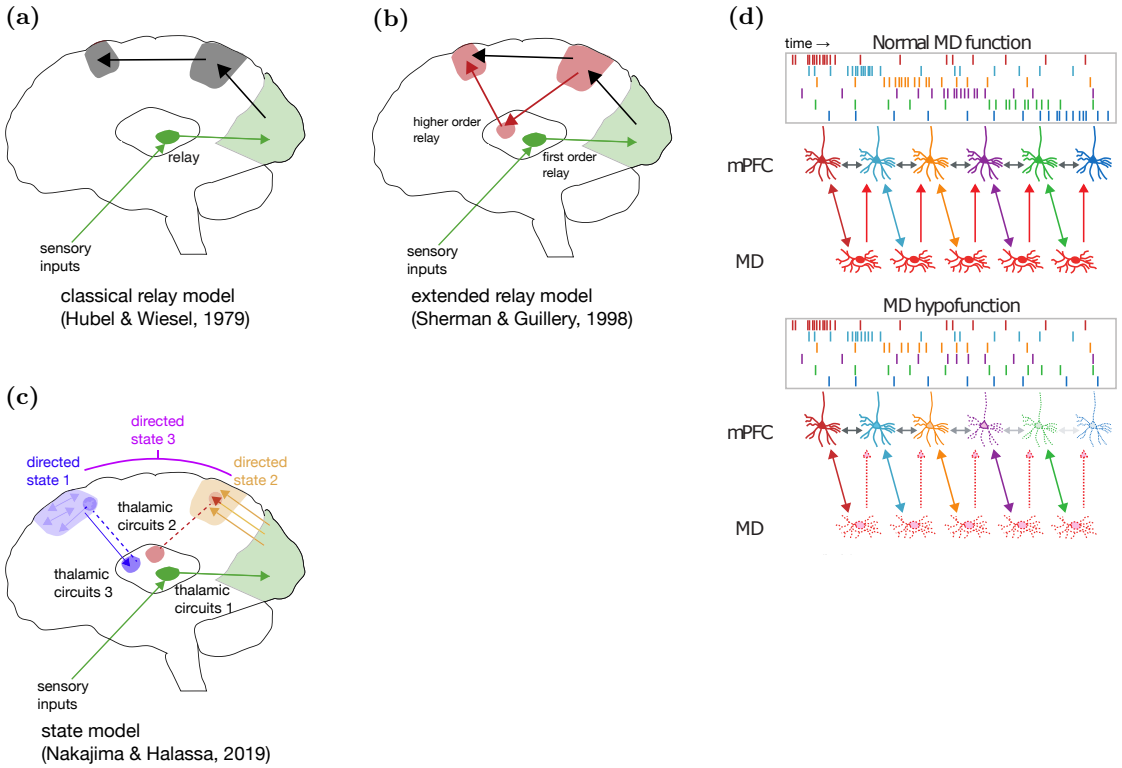

Figure 1.9: Models of thalamic functions (modified from Nakajima and Halassa (2017) after Sherman (2016)). **(a)** In the classical relay model sensory input is simply relayed via thalamic nuclei, e.g. visual information from the retina to LGN to V1. **(b)** Based on the driver input to thalamic nuclei, Sherman and Guillery (1998) proposed that first-order relays are driven by sensory input while higher-order relays are driven by cortical input. Cortico-thalamo-cortical pathways parallel to cortico-cortical pathways were proposed to contain efference copies. **(c)** Based on findings that MD can sustain prefrontal delay activity ((d)) it was proposed that some higher order thalamic nuclei maintain neuronal states in different cortical areas by for example influencing local cortical functional connectivity. **(d)** PFC neurons' activity can tile the memory epoch in behavioural tasks. (top): Recent findings in mice suggest that under normal conditions MD sustains prefrontal functional connectivity that is required for intra-prefrontal hand-over of information during delays. (bottom): If MD function is impaired functional connectivity in PFC degrades, thus abolishing delay epoch tiling. Copied from Parnaudeau et al. (2018).

nally, the authors found that inhibition of MD terminals in mPFC disrupted the delay epoch tiling of mPFC units' firing peaks, suggesting that MD is important for sustained mPFC delay activity (Fig. 1.9d).

A similar tiling of mPFC unit activity throughout a memory epoch was found in studies by the Halassa lab, when mice were instructed by rule cues to selectively attend to one of two sensory modalities to guide their subsequent choice. Enhancing MD excitability increased delay epoch functional connectivity between putative prefrontal pyramidal neurons and mPFC rule encoding and improved behavioural performance (Schmitt et al., 2017). Furthermore, a follow-up study found that MD activity is important for context-switching by increasing neuronal representations in mPFC of currently relevant contexts and conversely suppressing representations of the irrelevant one (Rikhye et al., 2018).

Together, these findings led to the recent proposal of a refined model of thalamic function in addition to sensory and cross-cortical relay functions: Some thalamic nuclei, amongst them MD, might support cognitive function by enabling/sustaining certain local cortical states, characterised e.g. by modulation of local functional connectivity, that are conducive to cognitive functions (Nakajima & Halassa, 2017, Fig. 1.9c).

1.5 Model Organisms in Systems Neuroscience

1.5.1 Monkey

Neuroscientists have utilised various model organisms to unravel the complexities of brain function, ranging from invertebrates such as worms and fruit flies to vertebrates like fishes and rodents (Bovenkerk & Kaldewaij, 2015; Friedrich et al., 2010; Kazama, 2015; Sengupta & Samuel, 2009). However, non-human primates, particularly macaque monkeys, have played a pivotal role as model organisms for several decades because they offer a unique opportunity to study the neural mechanisms that underlie complex cognitive and behavioral functions, such as decision-making, attention, memory, and language.

One of the main advantages of using macaque monkeys as a model organism is their close evolutionary relationship to humans (Gibbs et al., 2007). They possess highly developed brain structures that closely resemble those of humans, including intricate cortical organization and connectivity (Caminiti et al., 2015; Mantini et al., 2013).

Another notable aspect is their advanced sensory and motor systems. Monkeys exhibit fine motor skills, dexterity, and hand-eye coordination comparable to humans. This similarity allows researchers to investigate the neural underpinnings of motor control and sensorimotor integration, crucial for understanding movement disorders and developing neuroprosthetic devices (Castiello & Dadda, 2019; Morissette & Di Paolo, 2018). Additionally, monkeys' visual system closely resembles that of humans, making them invaluable in unraveling the neural mechanisms of visual perception and processing (Felleman & Van Essen, 1991; Markov, Vezoli, et al., 2014; Orban et al., 2004).

Beyond their neurobiological similarities, monkeys also exhibit complex social behaviors and cognitive capabilities that parallel certain aspects of human behavior. They

live in hierarchical social structures, engage in cooperative interactions, and display intricate social dynamics. These traits make them particularly well-suited for studying the neural basis of social cognition, empathy, decision-making, and social reward processing (Testard et al., 2021). Moreover, monkeys can be trained to perform complex behavioral tasks, facilitating the investigation of higher cognitive functions such as attention, memory, and learning (Fuster, 2015).

However, it is important to acknowledge the drawbacks associated with using monkeys as model organisms. They require large housing facilities and specialised care from highly trained staff. Moreover, practical and ethical considerations often limit the number of animals used in studies involving monkeys, leading to low statistical power and scientific generalisability. Additionally, methods for manipulating neuronal circuits in monkeys are often lagging behind those found in other model organisms (Roelfsema & Treue, 2014).

In conclusion, monkeys have played a crucial role in the investigation of highly complex brain functions. While their use as model organisms has some limitations, it is important to consider other species alongside monkeys to overcome these drawbacks.

1.5.2 Mouse

Whereas neuroscientific research in monkeys is arguably an important pillar in our understanding of the brain (Gray & Barnes, 2019), the field of system neuroscience has seen a steady rise in the number of rodent studies concerned with cognition, especially using mice. Mice are relatively easy to care for and readily available, thus facilitating standardisation of experimental protocols and consequently increasing statistical robustness (The International Brain Laboratory et al., 2021; The International Brain Laboratory et al., 2022). Equally importantly, short development cycles of tools compared to primates promoted the availability of an extensive molecular and genetic toolkit in mice (van der Weyden et al., 2011; Kaczmarczyk & Jackson, 2015). Most notably for systems neuroscience, the combination of mouse lines with robust conditional (cell-type specific or driven by recombinases) expression of genetic constructs, such as markers, reporters and promotors, and viruses with antero- or retrograde axonal transport can provide the opportunity for detailed manipulations and measurements of defined neuronal circuits (Nassi, Cepko, et al., 2015; Kim et al., 2017).

As a concrete example, the influence of dopaminergic cells from the VTA onto mPFC could be investigated using a double virus injection strategy. First, a DAT-Cre virus injected into VTA would lead to local expression of Cre-recombinase. Second, a retrograde virus injected into mPFC would be transported to VTA cells projecting to mPFC. If that retrograde virus contained a Cre-dependent DIO opsin the opsin would only be expressed in VTA cells projecting to mPFC. Thus, illumination of the injection site in mPFC would locally modulate dopaminergic neuron activity.

In addition to optogenetic modulators, another kind of widely used genetic constructs in the mouse are reporter proteins whose fluorescence is dynamically modulated by biophysical processes. In this way, researchers can optically measure neuronal activity with high spatial resolution (Emiliani et al., 2015), using sensors for Ca^{2+} (Grienberger

& Konnerth, 2012) or, more recently, for voltage (Knöpfel & Song, 2019; Bando et al., 2021), neurotransmitters and neuromodulators (Sabatini & Tian, 2020).

In contrast to these advantages, the use of mice arguably comes with downsides that concern the generalisability of findings to humans. First, behavioural tasks that mice can be trained to need to be catered to the ecological niche into which mice have evolved. On the one hand this means that tasks involving spatial, olfactory and auditory modalities are expected to be easier to learn. On the other hand, mice's cognitive capacity limits complexity of tasks, so that concreteness, both temporally (i.e. short or no memory delays) and in response modalities (e.g. licking, moving to locations), are favoured over abstractness. Even so, recent years have seen a rise in the development and use of more complex behavioural tasks in mice (Carandini & Churchland, 2013; Liu et al., 2014; Wimmer et al., 2015; Rikhye et al., 2018), albeit at the price of increasing the duration of animal training up to multiple months (Halassa, 2017, Q&A). Second, because of their evolutionary distance there is no definite consensus on functional homologues of primate PFC in mice. In contrast to primates, what has been called PFC in rodents is lacking a granular cortical layer IV, precluding the definition based on cytoarchitecture (Laubach et al., 2018). Furthermore, using the MD-projection criterion can also be problematic in so far as MD subregions in rodents project to secondary motor cortex (M2). Based on hodological and functional arguments, the latter has been proposed to be a rodent analogue of primate frontal eye field (FEF) (Reep et al., 1987), an area that is regarded not to be part of PFC by some researchers (E. K. Miller & Cohen, 2001). Possibly due to these under-specifications, the use of the term PFC and its subregions is inconsistent across studies in mice. To alleviate these issues, recent efforts have been made to more robustly define mouse PFC (Laubach et al., 2018; J. A. Harris et al., 2019; Le Merre et al., 2021).

Despite certain shortcomings, the tools that are available in mice have proven useful to gain more mechanistic insights into neuronal processes. Along with methodological advances in non-human primates (Stauffer et al., 2016; Nassi, Avery, et al., 2015; El-Shamayleh & Horwitz, 2019; Tremblay et al., 2020) systems neuroscience research in mice is therefore an important piece in the puzzle of understanding brain function.

1.6 Aims of the Thesis

This thesis comprises two projects with the overarching goal to further elucidate the role of the prefrontal cortex (PFC) in executive functions, particularly in the translation of sensory information into future action and the safeguarding of that information in working memory against interference.

1.6.1 Structuring of Abstract Working Memory Content by Fronto-Parietal Synchrony in Primate Cortex

Working memory is a distributed process that involves the interaction of PFC with association areas in posterior cortex. Extensive research suggests that a distributed network comprising PFC and PPC as central hubs plays a crucial role in working memory

(Chafee & Goldman-Rakic, 2000; Dotson et al., 2014; Rottschy et al., 2012; Salazar et al., 2012; Wendelken et al., 2008). The complexity of behaviour requires the flexible and structured composition of transient widespread neuronal networks that enable different aspects of working memory (Bressler & Menon, 2010; Mesulam, 1990; Tononi et al., 1998). How typically weak long-distance cortico-cortical connections allow controlled cross-areal communication is an open question (Markov, Ercsey-Ravasz, et al., 2014). A growing number of studies suggests that synchronous large-scale neural activity at distinct oscillation frequencies might functionally couple distributed brain areas and thus support integrative processes (Buschman & Miller, 2009; Buzsáki & Draguhn, 2004; Fell & Axmacher, 2011). Individual neurons in PFC show mixed selectivity, i.e. their activity can simultaneously contain information about different, possibly conflicting, task variables (Parthasarathy et al., 2017; Rigotti et al., 2013). Thus, while the formation of cross-area functional networks is one aspect of long-range communication, it is equally important to segregate relevant from irrelevant information (Tognoli & Kelso, 2014). However, the neuronal mechanisms that enable this kind of filtering are not clear. Local field potentials (LFPs) modify the excitability of individual neurons in an oscillating manner and could thus provide temporal windows to disambiguate multiplexed information (Gupta et al., 2016). Indeed, one study that explored this hypothesis found that multiple behaviourally relevant memory items that were simultaneously represented by neuronal activity in PFC could be separated by considering action potentials at different phases of the LFP (Siegel et al., 2009), suggesting that downstream receiver populations could "tune in" to phase-defined groups of more informative spikes in order to extract respective information.

To further explore the mechanisms of distributed working memory and the selection of behaviourally relevant information, the first project in this thesis explored the role of LFPs in frontoparietal communication in a task that required the resilience of abstract working memory to interference by distractors.

Non-human primates had been trained in a delayed match-to-sample (DMS) task that required the retention of the number of visual items (numerosity) across a delay epoch of several seconds (Figs. 1.10 and 2.1, Jacob & Nieder, 2014). To test the influence of distractors, the delay epoch was interrupted by the display of a behaviourally irrelevant numerosity. While the animals were performing the task microelectrodes were used to record electrophysiological signals from lateral PFC and VIP. Contrary to the expectation that PFC would mostly suppress responses to irrelevant information, as suggested by studies involving less abstract memoranda (Everling et al., 2002; di Pellegrino & Wise, 1993; Lennert & Martinez-Trujillo, 2011; Qi et al., 2010; Suzuki & Gottlieb, 2013), a previous analysis of single-unit activity (i.e. action potentials) showed that in PFC the interfering numerosity strongly decreased the information about the sample numerosity and led to representation of both relevant and irrelevant items in working memory. Despite initial breakdown of relevant information, PFC recovered it towards the end of the delay epoch. At the same time, VIP neurons also encoded both items but representations of the relevant item were stronger (Jacob & Nieder, 2014, Figs. 1.10a and 1.10b). However, from the examination of the two regions in isolation it was unclear how this recovery and filtering of relevant information could have been achieved mechanistically.

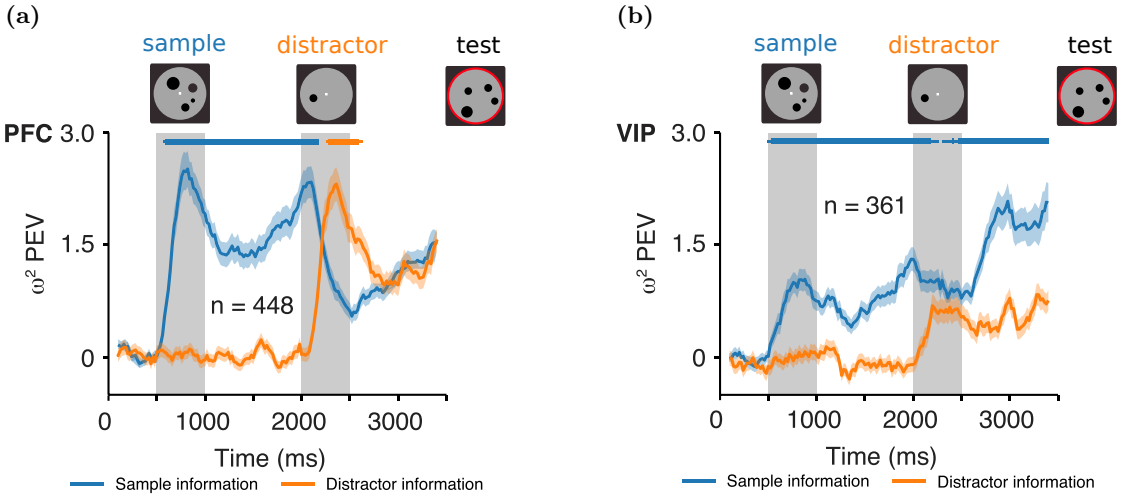

Figure 1.10: Neuronal coding for relevant and irrelevant numerosities in PFC **(a)** and VIP. **(b).** Animals were required to remember a sample numerosity and match it to a later test numerosity. An irrelevant distractor numerosity was shown during the memory epoch. Modified from Jacob and Nieder (2014).

1.6.2 Neuronal Signatures of Contextual Decision-Making in Mouse Prefrontal Cortex and Mediodorsal Thalamus

The nervous system's ability to predict future events based on prior information is fundamental to optimising upcoming behavior by preparing actions or allocating attention to sensory structures (Teufel & Fletcher, 2020). Predictive cues, in comparison to neutral cues, generally improve performance and decrease response times in behavioural tasks (Dalmaso et al., 2019; Mukherjee et al., 2021).

Representations of action planning in the form of preparatory behaviours prior to responses have mainly been studied in oculomotor tasks in humans. For example, only predictive cues lead to pre-saccades to the cued location (Jones et al., 2008) and rates and latencies of micro-saccades are modulated by predictiveness (Denison et al., 2019; Weaver et al., 2014). However, research into equivalent preparatory behaviours for non-primate species is lacking.

Prefrontal cortex has been shown to be involved in the processing of predictiveness. For example, ramping preparatory neuronal activity of single units towards an upcoming motor act, which is often found in timed behavioural tasks (Guo et al., 2017), exhibits smaller gradients in some monkey PFC neurons after ambiguous sensory cues that predict the required prospective behavioural response with some degree of uncertainty (Quintana & Fuster, 1999). In an auditory discrimination task in mice, PFC single units were shown to selectively respond to predictive but not unpredictive sensory cues (Nakajima et al., 2019). While predictive cues improved sensory discrimination on a behavioural level,

little is known about how varying states of uncertainty impact prefrontal processing of response instructions.

Efficient switching of cognitive states involving selective attention is linked to an MD-mediated increase of intra-prefrontal functional connectivity (Nakajima et al., 2019; Schmitt et al., 2017). Similarly, the shift of behavioural strategies in predictive and ambiguous contexts is associated with increased inter-regional functional connectivity between PFC and MD, the strength of which scales with the degree of current ambiguity (Hummos et al., 2022; B. A. Wang & Pleger, 2020). Consistent with this, mouse MD has recently been shown to track task uncertainty (Mukherjee et al., 2021). Thus, the structures and processes that enable rapid switching of ambiguity-dependent states are becoming increasingly well studied. However, it is unclear how various states themselves are characterised, for example in terms of neuronal recruitment within the brain areas of interest.

In order to investigate preparatory behaviour in a non-primate species, the effect of varying degrees of ambiguity on neuronal coding and neuronal states, the second project in this thesis combined a response device that allowed moment-to-moment readout of movements at millisecond precision (Sanders & Kepecs, 2012) with an auditory discrimination task that embedded the animals into predictive or ambigous cognitive contexts. One kind of context enabled the preparation of an action plan for a delayed response, while the other required deferring of planning until ambiguity was resolved by an instruction cue.

Thus, in *unambiguous* trials it would be beneficial to prepare motor structures for an efficient response (i.e. executive set). The enforced delay until the response requires two aspects of inhibitory control in these kinds of trials. Since a selected plan doesn't lead to immediate action but is future directed ("memory of the future"), it must be resilient against interference from other alternative plans, including but not limited to routine action sequences. In addition, this sustained kind of inhibitory control should be accompanied by a more immediate inhibition of action at inappropriate times — even if those actions are part of an appropriate plan.

Similarly, *ambiguous* trials in which animals had to defer their planning required resilience against the selection and immediate execution of routine plans. Furthermore, in contrast to the other kind of trials in which motor structures could be prepared for specific actions, ambiguous trials would benefit from preparing perceptual structures to more efficiently discriminate instruction cues (i.e. perceptual set).

Taken together, the different contexts not only deploy particular executive functions on the level of individual trials, but also require the flexible switching of cognitive states across trials.

Detailed analysis of the behavioural data revealed preparatory behaviour that varied with ambiguity. In addition, extracellular recordings from mPFC and MD during behaviour was used to characterise different neuronal states and investigate effects of ambiguity on processing of movement instructions.

Part II

Results

2 Structuring of Abstract Working Memory Content by Fronto-Parietal Synchrony in Primate Cortex[1]

2.1 Behavioural Task

In order to investigate how distractor resilience could be implemented for abstract working memory information, two monkeys were trained to remember the quantity of black circles (numerosity) shown in a visual display (Fig. 2.1). After a memory delay of 2500 ms this sample numerosity had to be matched to a test numerosity. An outcome-irrelevant numerosity (*distractor*) was shown during the memory delay in 80 % of trials. The remaining 20 % of trials showed a blank grey circle instead of an interfering stimulus, thus serving as a control. Thus, each trial comprised five distinct epochs: sample presentation (*S*) from 0 ms to 500 ms, first delay/memory (*M1*) from 500 ms to 1500 ms, distractor presentation (*D*) from 1500 ms to 2000 ms, second delay/memory (*M2*) from 2000 ms to 2500 ms, and the tests in which the animals had to signal a match or non-match.

2.1.1 Open questions from previous analyses

Wideband neuronal signals were recorded from PFC and VIP (Figs. 2.2a and 2.2b), brain regions in which neurons are tuned to abstract numerosities (Foster et al., 2022). In a previous analysis of this dataset both areas were shown to contain information about both the behaviourally relevant sample numerosity as well as the irrelevant distractor numerosity (Fig. 1.10, Jacob & Nieder, 2014). While most numerosity-selective single neurons in both areas were sample-selective, a larger number showed mixed selectivity for sample and distractor in PFC than VIP. Furthermore, representation of the distractor was in general lower in VIP. Because the animals had been previously trained on a simpler task without an interfering distractor in which all numerosities were behaviourally relevant (Jacob & Nieder, 2014; Nieder et al., 2006), it is possible that a low-level computational module extracted abstract numerosities and made that information available to upstream regions. The resolution of the relatively new introduction of a conflict of behavioural relevance was likely delegated to downstream areas.

This dual representation of relevant and irrelevant information did not prevent the animals from performing the task successfully (monkey R: 72 %, monkey W: 71 %). Through some mechanism, sample information and its behavioural relevance was protected from interference and could be read out at the appropriate time in the trial. In

[1] The data shown in this chapter was published in Jacob et al. (2018).

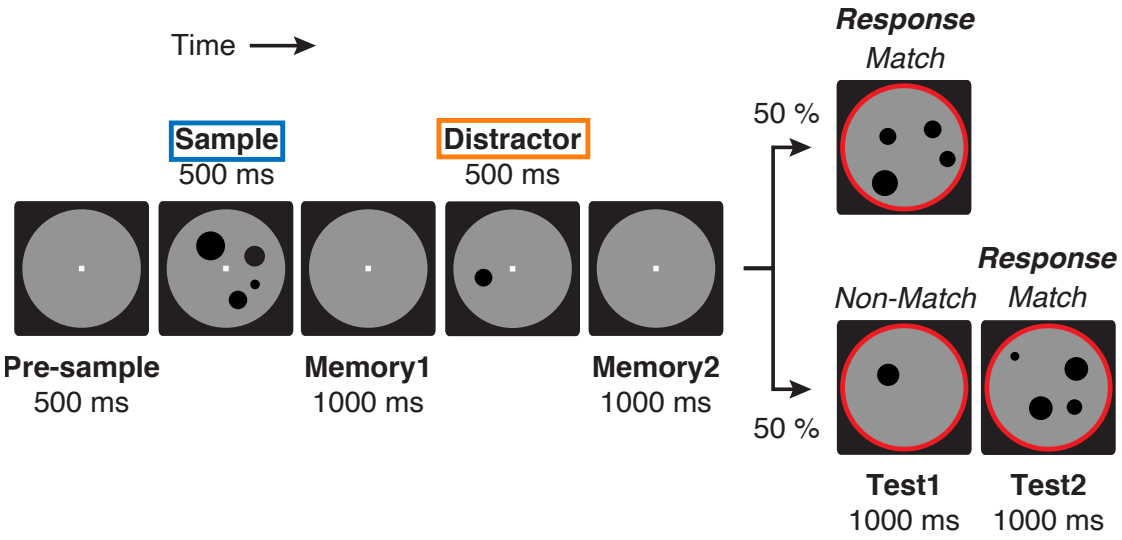

Figure 2.1: Delayed match-to-numerosity task. A **Sample** numerosity had to be memorised for 2500 ms and matched with test numerosities **Test1** or **Test2**. An irrelevant interfering **Distractor** numerosity during the memory delay had to be resisted. Modified from Jacob et al. (2018).

the following, I explored how cross-regional communication could have been established via LFPs and how mixed information could have been demixed.

2.2 Intra-regional Strength of LFP Oscillations: Power

2.2.1 Spectro-temporal modulation by task events

The spectral power of a signal is associated with the amplitude of oscillations at specific frequencies and thus quantifies how strongly different frequencies are represented in the combined signal. To quantify if and how different LFP frequencies varied with the task events, a sliding-window based approach was used to calculate LFP power relative to a baseline epoch (−500 ms to 0 ms before sample presentation, Figs. 2.2c and 2.2d).

Power was modulated in frequency-specific dynamics throughout the trial in both PFC and VIP. In PFC the onsets of the sample and distractor stimuli were accompanied by strong transient wideband increases in the theta, beta, and high gamma bands (Fig. 2.2c). Notably, increases in the theta and beta bands were markedly stronger for the distractor onset. Power in the high gamma band stayed relatively high during stimulus presentation, while it returned to baseline for theta and decreased below baseline for beta. The memory delays (**M1** and **M2**) were both characterised by elevated delta (2 Hz to 4 Hz) to theta power and below-baseline power in beta. All in all, in PFC the spectro-temporal dynamics of power modulation were virtually identical in the two trial halves (Sample and Memory1 from 0 ms to 1500 ms, Distractor and Memory2 from 1500 ms to 3000 ms).

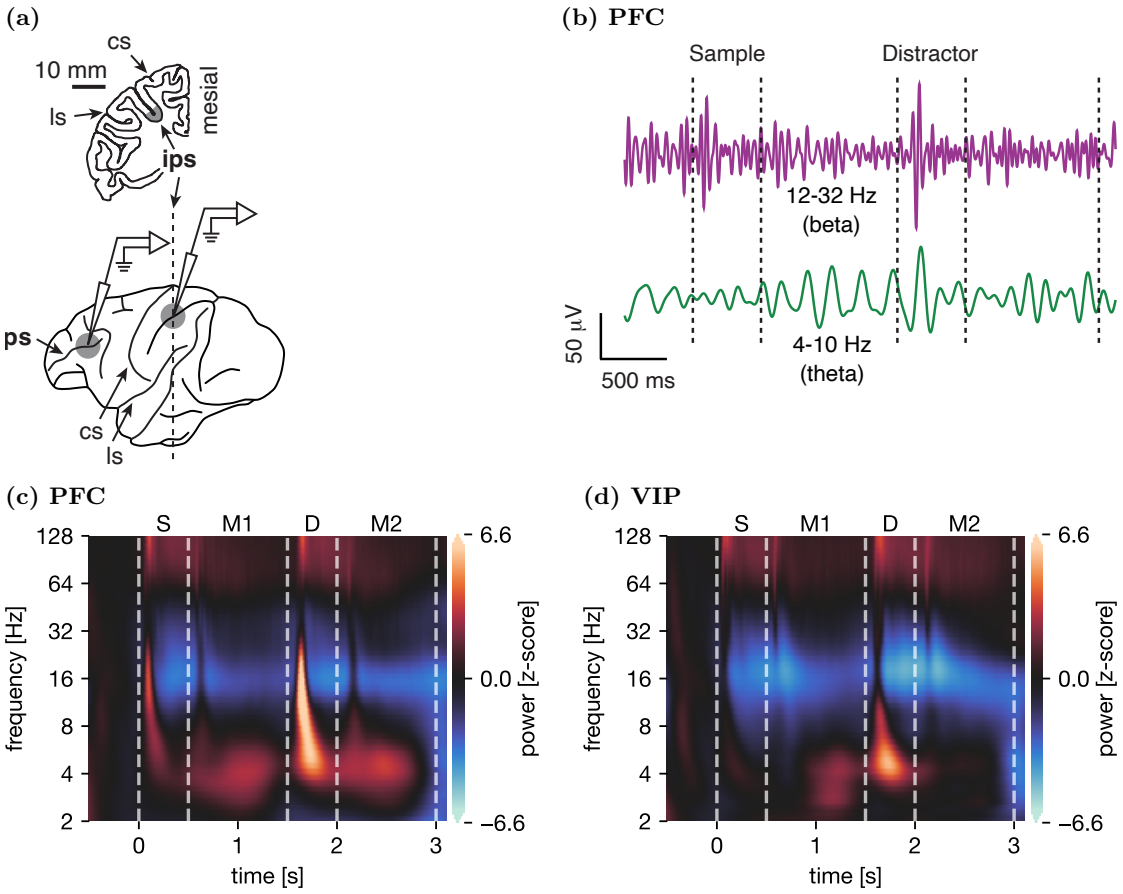

Figure 2.2: Local field potentials (LFPs) in monkey prefrontal and ventral intraparietal cortex. **(a)** Recording location of extracellular neuronal signals. **(b)** Time-domain LFP in PFC, filtered at beta 12 Hz to 32 Hz and theta 4 Hz to 10 Hz frequency ranges. **(c, d)** Sliding-window analysis of spectro-temporal LFP power in PFC and LFP, z-scored to pre-sample task epoch (−0.5 s to 0 s). Dashed vertical lines: trial epoch boundaries, S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2. Modified from Jacob et al. (2018).

At first glance, LFP power in VIP showed similar spectro-temporal dynamics like PFC. However, a few noteworthy differences were observed (Fig. 2.2d). During the sample presentation only high gamma power increased. In contrast to PFC, the presentation of the distractor numerosity was associated with a very strong increase of theta power in addition to high gamma. While in PFC delta band power sustained increased levels throughout both memory delays, in VIP it only increased in the second half of the *first* memory delay and stayed at baseline levels after the presentation of the distractor numerosity.

Taken together, power was tightly regulated by the task. Distinct recruitment at functionally different task epochs suggests functional specialisation of frequency bands. Furthermore, intra-band differences between the first and second trial halves suggest that information about the numerosities' behavioural relevance was present already at the power level.

2.2.2 Numerosity-dependent oscillation intensity

The previous section demonstrated that LFP power modulation was dependent on the time in the trial and behavioural relevance of the most recent numerosity. To test if power also contained information about sample and distractor numerosity, I used Naïve Bayes classifiers. Classifiers were trained to find numerosity-dependent differences in power at individual frequencies at specific times (i.e. a single spectro-temporal "pixel"). Classifier training was performed on a random class-matched subset of 75 % of trials. The remaining 25 % were used to test the classifiers' prediction accuracy for sample and distractor numerosities.

Sample and distractor numerosities could be decoded with above-chance accuracy in multiple frequency bands (Figs. 2.3a to 2.3d). Prediction accuracy was generally higher in PFC than in VIP. During sample presentation, first memory delay and distractor presentation, sample numerosity could be decoded from power in the lower (delta and theta) as well as higher frequencies (beta, gamma) (Figs. 2.3a and 2.3b). In PFC decoding was especially strong in the delta and beta bands. Interestingly, during the presentation of the distractor numerosity, sample information in the beta band was especially high. In contrast, during the second memory delay sample information was almost lost in the beta band and shifted towards lower frequencies. The distractor numerosity could be decoded from power in the delta, theta and beta bands. Distractor information was especially strong during the presentation of the distractor stimulus and was almost lost during the second memory delay.

Classifier analyses can discover numerosity-dependent covariation of LFP power but don't provide information about the concrete organisation of different classes. During the second memory delay power in the low-frequency delta band was in fact ordered as a monotonically increasing function of sample numerosity (Figs. 2.3e and 2.3f, left). Distances between subsequent sample numerosities decreased in accordance with the Weber-Fechner law (Nieder, 2016). In contrast, there was no such order when sorted by the distractor numerosity. In the beta band power was less ordered and separation

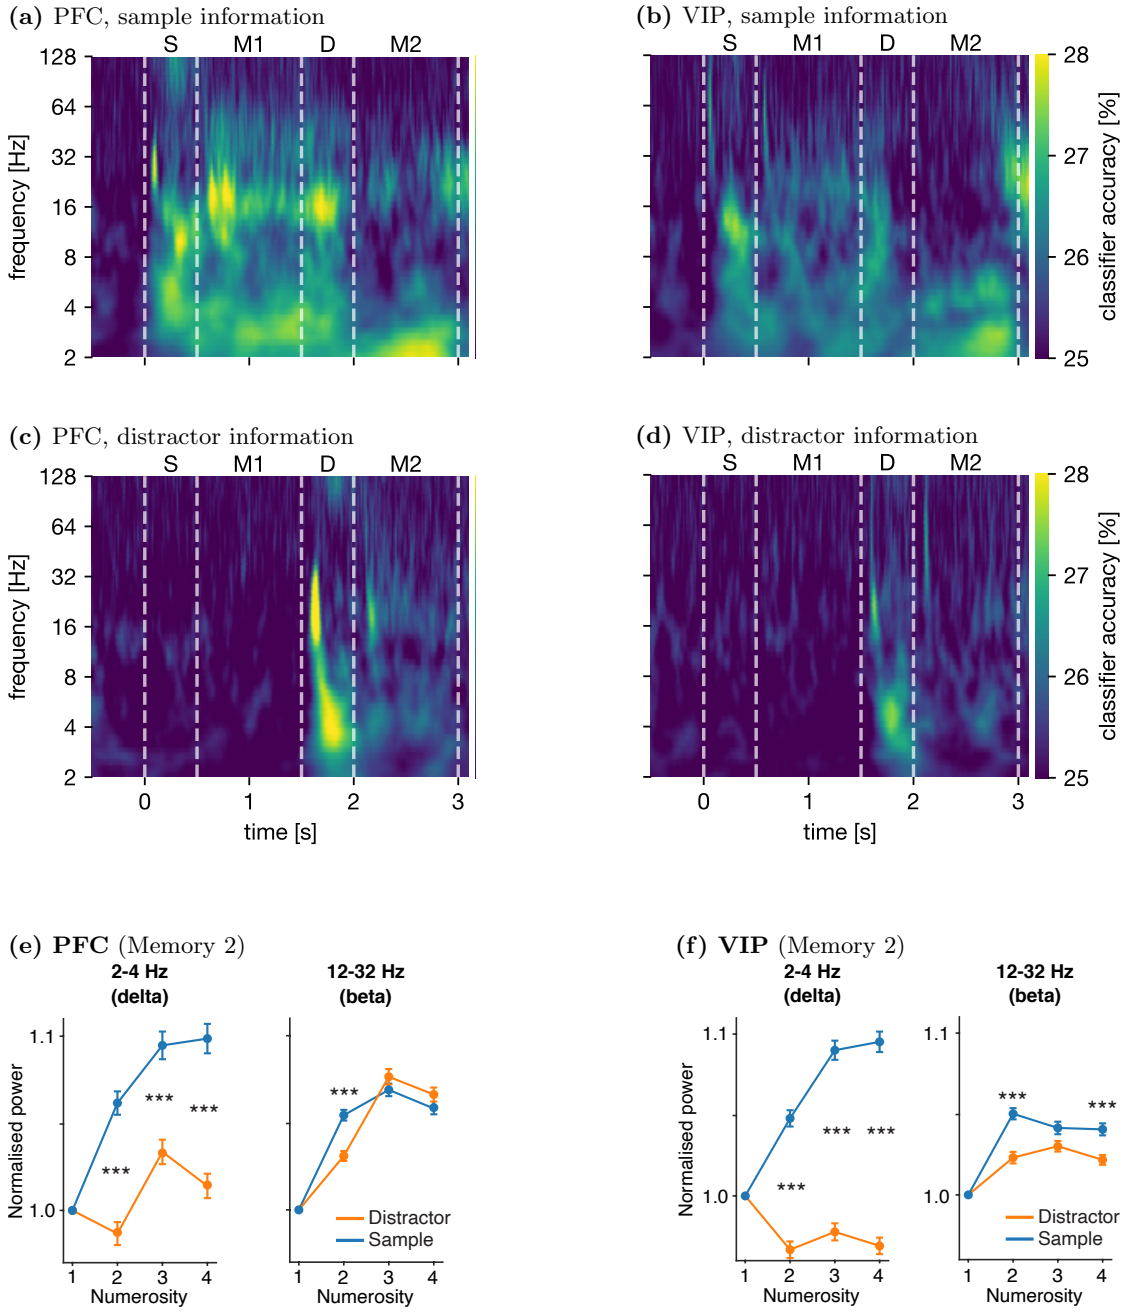

Figure 2.3: Numerosity dependent separation of power (a) Classification accuracy for sample (chance level: 25 %) of Naïve Bayes classifiers using LFP power in PFC as features. **(b)** Same analysis like in (a) for classifiers trained and tested on VIP LFP power. **(c, d)** same as (a, b) for the distractor numerosities. Dashed vertical lines: trial epoch boundaries. S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2. **(e)** PFC power in the second memory delay (M2) as a function of sample/distractor numerosity in the 2 Hz to 4 Hz (delta, left) and 12 Hz to 32 Hz (beta, right) bands. **(f)** same analysis like (e) for VIP. Power values were normalised to values in trials with numerosity 1. Error bars: SEM across sites. $^{***}p < 0.001$ (Wilcoxon signed-rank test, sample vs. distractor). Modified from Jacob et al. (2018).

between numerosities less pronounced, possibly explaining the lower decoding accuracy in this frequency band (Figs. 2.3e and 2.3f, right).

In sum, LFP power did not only vary with task epochs but distinct frequencies carried information about working memory content. Spectro-temporal dynamics of information suggest a functional specialisation of distinct frequency bands. Importantly, during the second memory delay LFP power was preferentially organised by the behaviourally relevant sample numerosity, as demonstrated by numerosity-dependent ordering as well as increased sample information.

2.3 Cross-regional Communication: Functional Connectivity

The findings in the previous section suggest that within PFC and VIP LFP oscillations were organised into functionally separate frequency channels. Because of inter-regional similarities in the task-related activation of those channels it is likely that LFP signals were coordinated. Synchronised activity is thought to be conducive to long-range communication (Fries, 2005) and might therefore support the fronto-parietal exchange of working memory information in order to solve the task.

2.3.1 Task-related fronto-parietal synchrony in distinct frequency bands

As a first step in order to find out if oscillations in PFC and VIP synchronised during the task, I quantified cross-regional synchrony using the phase-locking value (PLV). Briefly, the PLV quantifies how consistent the phase difference of two oscillatory signals was across observations (e.g. at a certain time in the task across trials, Lachaux et al. (1999)), irrespective of the direction of the interaction.

High cross-regional synchronisation was found for the delta, theta and beta frequency bands (Fig. 2.4). In contrast to the increases in gamma band power that were associated with stimulus presentations (Figs. 2.2c and 2.2d), synchronisation at higher frequency bands was absent. PFC and VIP synchronised strongest in the beta frequency band at the onsets of the numerosity displays. Synchrony during the presentation of the distractor was stronger than during sample presentation (Fig. 2.4a), a trend that was consistent with higher LFP power during distractor presentation (Figs. 2.2c and 2.2d). These synchronisation peaks were additional to an elevated beta synchrony that persisted throughout the trial, also in trials without an interfering numerosity (Fig. 2.4b). In comparison, fronto-parietal synchrony in the lower frequencies (delta, 2 Hz to 4 Hz) was highest during the memory delays, i.e. when the sample information in the power in these frequencies was highest (compare with Figs. 2.3a and 2.3b). Notably, omission of the distractor still resulted in a precisely timed reduction of delta band synchrony. In these trials delta band synchrony during the second memory delay was especially high and shifted to earlier times (Fig. 2.4b).

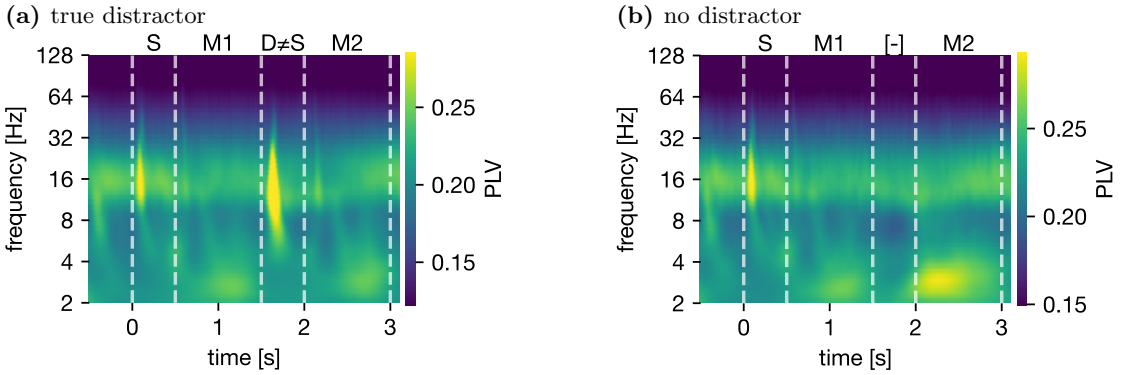

Figure 2.4: Task- and frequency specific cross-regional synchrony. Phase-locking value (PLV) PFC-VIP for different distractor conditions. **(a)** Trials in which the distractor numerosity was different from the sample numerosity. **(b)** Trials in which a blank circle was shown instead of an interfering numerosity. S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2. Dashed vertical lines represent trial epoch boundaries. Pseudocolour scales across plots were normalised to pre-sample ranges for visual comparability. This was necessary because of varying numbers of trials for individual conditions and the PLV's bias to 1 for small sample sizes. Modified from Jacob et al. (2018).

In summary, PFC and VIP showed task-dependent cross-regional synchronisation in distinct frequency channels. Synchrony was precisely timed and also varied with behavioural relevance of the numerosities.

2.3.2 Distinct frequency bands for direction-specific synchrony

To assert if cross-regional synchronisation was dominated by fronto-parietal (PFC to VIP) or parieto-frontal (VIP to PFC) cross-regional communication, I extended the non-directed PLV analysis to directed metrics.

The phase-slope index (PSI) quantifies the dominant direction of information flow across regions (Nolte et al., 2008). Interestingly, information flow was not symmetrical but strongly directed (Fig. 2.5a). Directionality was segregated between lower and higher frequencies. Positive values indicated predominant information flow from PFC to VIP in the delta band during the memory delays. The opposite net direction (VIP to PFC) was characteristic of the beta band, both for the presentation of the stimuli as well as the sustained synchrony throughout the trial. Thus, compared to baseline, fronto-parietal information flow increased during the memory delays and parieto-frontal flow increased during sample and distractor presentation (Fig. 2.5b).

In trials without a distracting numerosity the fronto-parietally dominated flow in the delta band was interrupted during the omission of the distractor numerosity (Fig. 2.5c), similar to what was found for the non-directed synchrony (compare Fig. 2.4b). Similarly, fronto-parietal net information flow increased earlier than in trials with a distractor, but not as early as the non-directed synchrony.

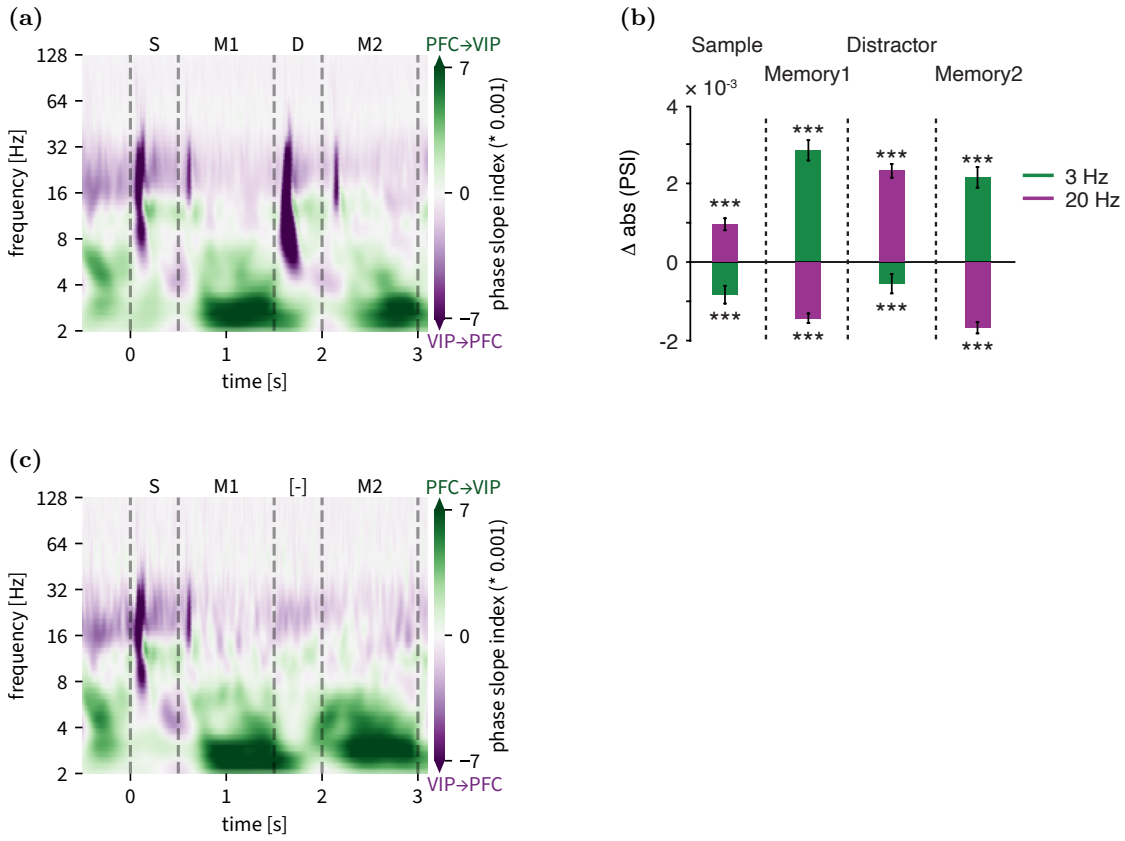

Figure 2.5: Net cross-regional information flow (phase-slope index (PSI)). (a) Time-frequency representation of net information flow in trials with a distracting numerosity. Positive values (green colours) indicate net flow from PFC to VIP, while negative values (purple) indicate the opposite direction. **(b)** Modulation of net direction in comparison to baseline (−0.5 s to 0 s) at 3 Hz (delta band) and 20 Hz (beta band) in (a). **(c)** same analysis like in (a). for trials without a distracting numerosity. S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2. Dashed vertical lines represent trial epoch boundaries. Modified from Jacob et al. (2018).

Since the PSI quantifies the *net* or *dominant* direction of information flow, it is not clear if cross-regional information flow seen in Figure 2.5 was in fact bi-directional and dominated by one direction or exclusively uni-directional. To answer this question, I extended the analysis of directional functional connectivity to Granger Causality (GC), which allows the quantification of uni-directional information flow. This analysis shows that most information flow was indeed uni-directional. In the delta band high GC values in Figure 2.6a indicate high fronto-parietal information flow, while low values in Figure 2.6b suggest almost none from VIP to PFC. On the other hand, the beta band was characterised by almost exclusively uni-directional parieto-frontal flow. In addition to the PSI analysis during the presentation of the distractor there was some feedback information flow from PFC in the beta band.

In the trials without an interfering numerosity the non-directed PLV had found increased delta band synchrony immediately after the omission of the distractor (Fig. 2.4b), whereas the PSI analysis had found fronto-parietal net information flow slightly later (Fig. 2.5c). GC analysis resolves this apparent discrepancy. In fact, fronto-parietal information flow in the delta band was high immediately after distractor omission (Fig. 2.6c). Of note, however, in these trials an equally strong *parieto-frontal* flow was discovered at the same time (Fig. 2.6c), that was totally absent in trials *with* a distractor. Because this information flow from VIP decreased more rapidly during the second memory delay, the net dominance of fronto-parietal flow was lagged.

LFPs predominantly measure spatially and temporally coordinated synaptic activity close to the recording electrode (Kajikawa & Schroeder, 2011). To investigate directed communication in a more explicit way, in the following I quantified the synchronisation of spikes from a sending population (since action potentials are the unequivocal output of neurons) to the LFP signal of a receiving region (Pesaran et al., 2008; Salazar et al., 2012; Liebe et al., 2012). A total of 4956 $\mathrm{PFC_{unit}}$-$\mathrm{VIP_{LFP}}$ and 3525 $\mathrm{VIP_{unit}}$-$\mathrm{PFC_{LFP}}$ pairs were tested for significant spike-field (SF) locking, i.e. whether spikes from a neuronal unit did not occur at random LFP phases[2]. During the second memory delay, the delta to theta band and the beta band displayed especially high numbers of locked SF pairs, thus dividing the frequency scale into separate channels similar to what was found with the LFPs alone (Fig. 2.7a). Furthermore, more neuronal PFC units were locked to VIP LFP phases in the delta band than VIP units to PFC LFPs, while the opposite was true for the beta band[3]. Similarly, the locking *strength* across all neuronal units, as quantified by SF-PLV, was higher for the fronto-parietal direction in the delta band and parieto-frontal signalling dominated in the beta band[4] (Fig. 2.7b).

All in all, distinct frequency bands were separated by largely uni-directional cross-regional communication. Whereas higher frequencies in the beta band were dominated by information flow from VIP to PFC, fronto-parietal signalling was present at the lower

[2] $p < 0.05$, Rayleigh test for circular uniformity

[3] locked pairs at 3 Hz: 1745 vs. 973
locked pairs at 20 Hz: 990 vs. 1149
odds-ratio 2.081, $p < 10 \cdot 10^{-35}$, Fisher's exact test

[4] Wilcoxon rank-sum test: $p = 0.007$ (3 Hz), $p = 5.107 \cdot 10^{-7}$ (20 Hz)
total number of neurons: 444 (PFC), 359 (VIP)

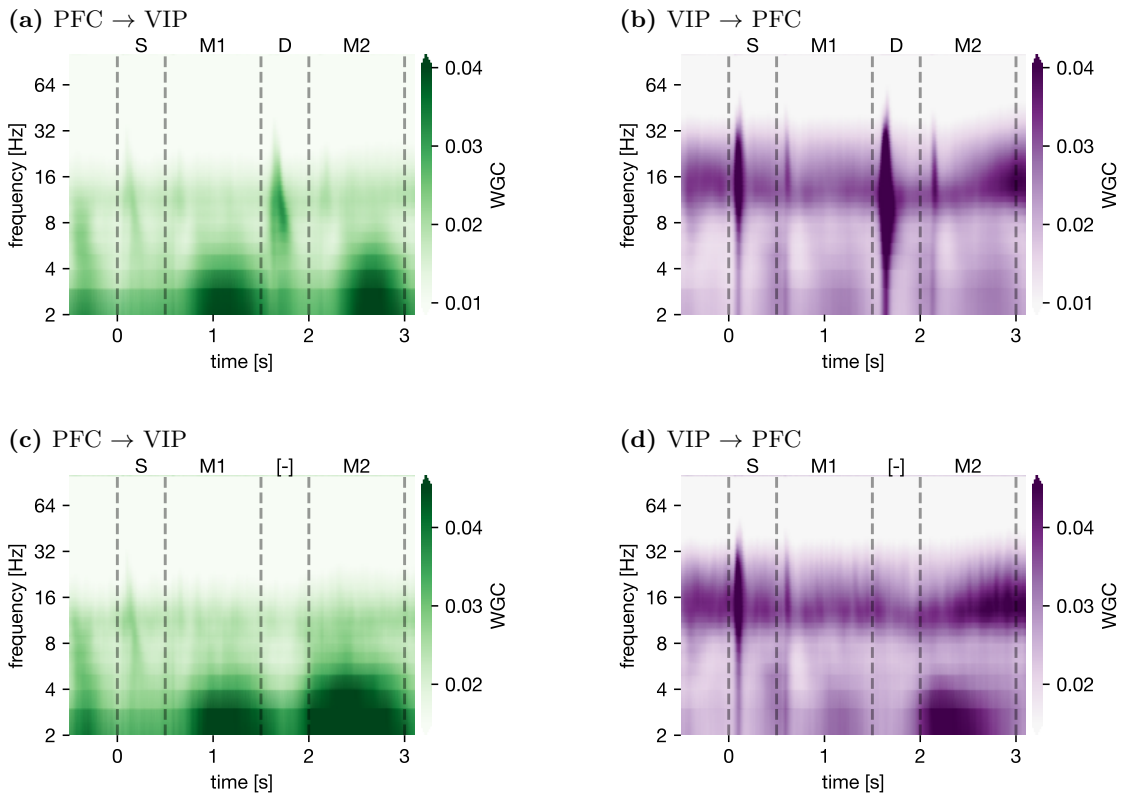

Figure 2.6: Granger Causality with and without distractor numerosity. Granger Causality for directed top-down (PFC to VIP, (a, c)) or bottom-up (VIP to PFC, (b, d)) synchrony in trials with (a, b) and without (c, d) interfering numerosity. S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2, WGC: Wiener-Granger Causality. Dashed vertical lines represent trial epoch boundaries. Note that the visual appearance differs because of the transformation from a linear to a logarithmic frequency scale (due to the computational algorithm). Modified from Jacob et al. (2018).

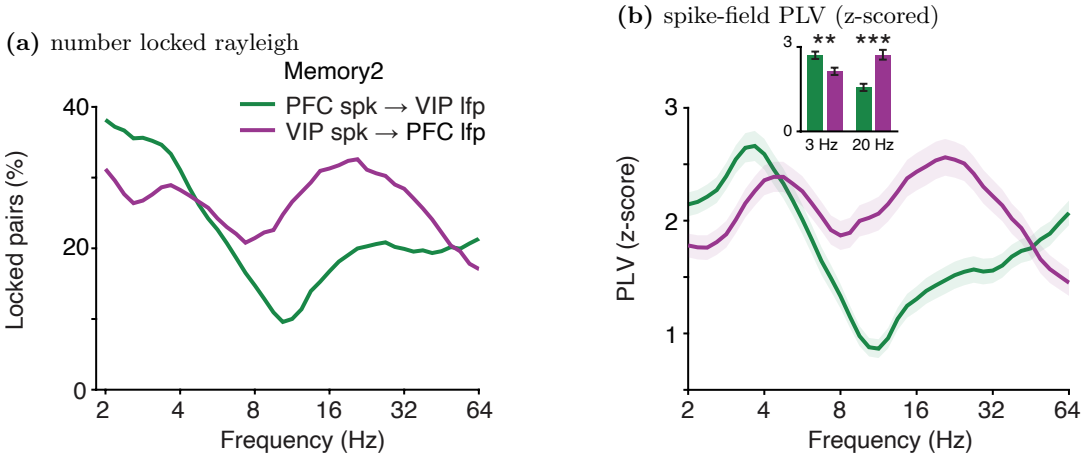

Figure 2.7: Interregional spike-field locking. **(a)** Percentage of locked spike-LFP pairs (from total: PFC_{unit} - VIP_{LFP} $n = 4956$, VIP_{unit} - PFC_{LFP} $n = 3525$) as a function of LFP frequency quantifying if spikes of single neuronal units occurred at specific phases of individual sites' LFPs (Rayleigh test for circular uniformity $p < 0.05$ per combination of single unit, LFP site and frequency). **(b)** Spike-field locking strength quantified by the PLV shows how accurately spikes occurred at specific LFP phases (z-score from shuffled null distribution and sem across neurons). Modified from Jacob et al. (2018).

frequencies that comprise the delta and theta bands. Analyses with locally referenced LFP signals that controlled for sensitivity of synchrony measures to the referencing scheme produced virtually indistinguishable results (Fig. S1).

2.3.3 Task- and memory-content dependent fronto-parietal synchrony

To test if the frequency-specific directed connectivity was a fixed network property (e.g. "fitted" by training) or if it varied with task demands, SF-PLV for numerosity-coding PFC units to parietal LFP sites was computed for the memory delays. During the first memory delay, i.e. when the sample numerosity was the only memory item, phase-locking of sample-coding PFC units was highest in the beta band (Fig. 2.8a). After the presentation of the distracting numerosity, however, high locking of sample-selective units shifted to the lower theta and higher gamma frequencies. Thus, locking in the theta band was higher in the second memory delay, while for the beta band it was higher in the first memory delay[5]. PFC units whose spiking coded the second memory item, i.e. the distractor, showed SF-PLV that was especially strong in the beta frequency range, where it was stronger than for sample-selective units[6] (Fig. 2.8b). Thus, PFC units that encoded the most recent memory item were phase-locked most strongly to parietal beta frequencies. Since phase-estimation is less robust with lower LFP power,

[5] Wilcoxon rank-sum $p < 0.01$
[6] Wilcoxon rank-sum $p < 0.01$

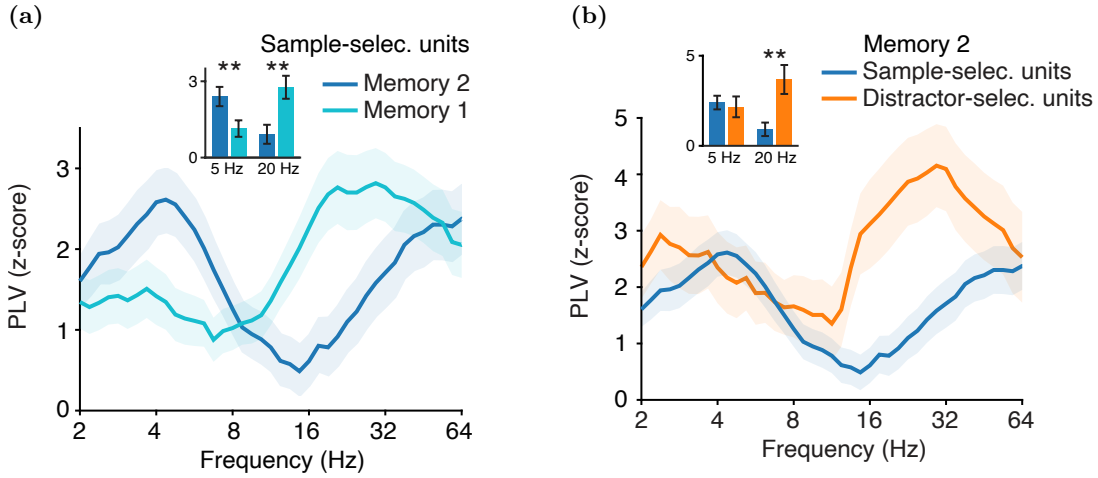

Figure 2.8: Spike-field locking strength (PLV) for fronto-parietal pairs with significant sample or distractor coding. (a) Locking strength during the first and second memory delays for pairs of frontal sample-numerosity coding single units and parietal LFPs at sites where sample-coding units were recorded (memory 1: $n = 47$; memory 2: $n = 46$ PFC units; z-score from shuffled null-distribution, shading across neurons). Inset: Cross-epoch comparisons of locking strength at the 5 Hz (theta) and 20 Hz (beta) frequencies (Wilcoxon rank-sum test, $^{**}p < 0.01$). **(b)** Same analysis like in (a) for pairs of frontal distractor-coding single units and parietal LFPs at site where distractor-coding units were recorded during the second memory delay ($n = 19$ PFC units; sample-coding units from (a) as comparison). Inset: Comparisons of locking strength at the 5 Hz (theta) and 20 Hz (beta) frequencies across sample and distractor-coding pairs (Wilcoxon rank-sum test, $^{**}p < 0.01$). Modified from Jacob et al. (2018).

this analysis was repeated for power-stratified observations, which produced the same trends (Figs. S2a and S2b).

Fronto-parietal synchrony changed if the distractor was omitted (Figs. 2.4b, 2.5c, 2.6c and 2.6d). Did it also change for different sample and distractor numerosities, i.e. did synchrony also contain abstract numerosity information? The possibility of consistent variation of $\mathrm{PFC_{LFP}}$-$\mathrm{VIP_{LFP}}$ PLV with numerosities was explored using the phase-locking selectivity index (PLSI).

Information about the sample numerosity was very high in the delta/theta and beta frequency bands during stimulus presentation (Fig. 2.9a). During the first memory delay and distractor presentation sample information was still above baseline levels. In the second memory delay, after the presentation of the distractor, sample information in the delta/theta bands increased, while it virtually vanished in the beta band. In contrast, information about the distractor was only found during distractor presentation in the delta/theta bands and, especially strong, in the beta band (Fig. 2.9b).

To summarise, task requirements changed fronto-parietal synchrony, with working memory item specific spike-field locking and strength of LFP-LFP synchrony.

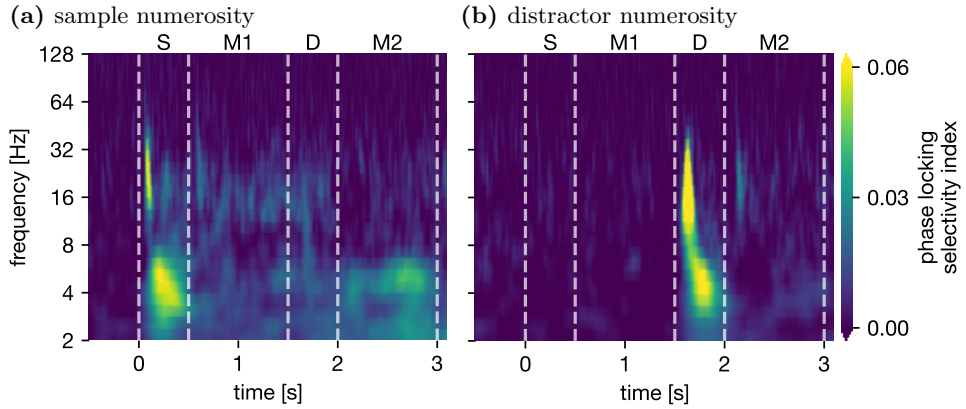

Figure 2.9: Phase-locking selectivity index (PLSI) PFC-VIP. **(a)** PLSI quantifying how much and consistently fronto-parietal synchrony varied with sample numerosity. **(b)** Same as (a) but quantifying variability across distractor numerosities. S: sample presentation, M1: memory 1, D: distractor presentation, M2: memory 2. Dashed vertical lines represent trial epoch boundaries. Modified from Jacob et al. (2018).

2.3.4 Phase- and frequency dependent information

The previous analyses suggested that during the second memory delay the delta/theta bands were especially important for communication from PFC to VIP and maintaining sample information. Prefrontal neuronal units' spiking activity during the second memory delay was informative of both the sample and distractor numerosities (Jacob & Nieder, 2014, Fig. 1.10). LFP oscillations change a neuron's excitability and can modulate presynaptic spikes' impact on the postsynaptic membrane potential, thus acting as filtering reading frame (Gupta et al., 2016). To test if a receiving parietal population could demix multiplexed sample and distractor information from prefrontal neurons, I quantified how much information prefrontal spikes binned into ranges of phases of VIP LFPs contained. The rationale behind this approach is that some (groups of) spikes might be more informative than others (Fig. 2.10).

During the first memory delay sample information was not distributed uniformly across spikes from prefrontal sample-selective units ($n = 98$, Figs. 2.11a and 2.11c). Rather, spikes that occurred around optimal phases of theta and beta/gamma oscillations in VIP carried significantly more information than at other phases[7]. Notably, prefrontal sample information was distributed even more non-uniformly across VIP phases during the second memory delay ($n = 73$, Figs. 2.11b and 2.11c). Whereas phase-dependent coding during the first memory delay was higher in the beta than the theta bands, in the second delay it was strongly shifted to lower frequencies[8] (Fig. 2.11c). Additionally, the optimal phase for sample readout in the theta band shifted from $-\frac{1}{4}\pi$ in the first to $-\frac{3}{4}\pi$ in the second memory delay. The strong phase-dependent concentration of aug-

[7] $p < 0.01$, permutation test
[8] $p < 0.05$, permutation test

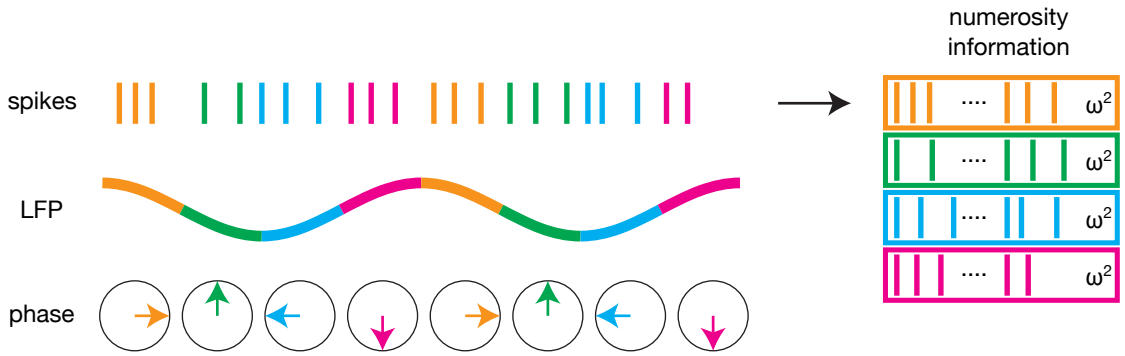

Figure 2.10: Schematic for phase-resolved spike information. Left: spikes (horizontal lines) were grouped by the LFP phase at which they occurred (represented by the corresponding colouring). Right: individual "phase groups" of spikes were then subjected to an information analysis.

mented sample information suggests that other phases in the theta band could enhance readout of the distractor numerosity. Indeed, in the same memory epoch distractor information from distractor-selective prefrontal neurons was non-uniformly distributed across parietal theta phases ($n = 73$, Figs. 2.11d and 2.11e). Of note, the optimal theta phases for sample and distractor readout were well separated[9] (Fig. 2.11f). Strikingly, the distractor numerosity occupied the same optimal readout theta phase as the sample numerosity did in the first memory delay. This was reminiscent of the "taking-over effect" that was observed in the spike-field phase-locking value (Fig. 2.8). To control for possible effects of preprocessing on phase estimation, the same analysis was performed without the subtraction of evoked potentials (Fig. S3). Consistent with the previous results, phase-dependency of information was similarly distributed across frequencies and optimal readout theta readout phases were significantly different[10].

Within PFC numerosity information was also enhanced at specific phases of PFC LFP oscillations (Fig. 2.12). Similar to the cross-regional perspective, both sample and distractor information were phase-dependent in the theta band and their optimal readout phases were significantly separated[11] (Figs. 2.12c and 2.12d). Additionally, distractor but not sample information was significantly phase-dependent in the beta band (Fig. 2.12c). Accordingly, optimal readout phases in that frequency band did not differ[12] (Fig. 2.12d).

The phase differences of optimal theta readout phases between cross- and intra-regional were around half a period. Since this was also the average phase difference between cross-regional LFPs (Fig. 2.13a), which in turn is a consequence of common-average referencing (Shirhatti et al., 2016), a control analysis is warranted. Complex mediation analysis (Pascual-Marqui et al., 2017) was used to test if the LFP signal in VIP was mediated by frontal LFP. Indeed, the highest contribution was found in the

[9] sample: $134 \pm 24°$, distractor: $26 \pm 25°$, $p = 0.03$, permutation test
[10] sample: $126 \pm 27°$, distractor: 32 ± 22 , $p = 0.03$, permutation test
[11] sample: $34 \pm 20°$, distractor: $134 \pm 31°$, $p = 0.03$, permutation test
[12] sample: $151 \pm 66°$, distractor: $130 \pm 52°$, $p = 0.64$, permutation test

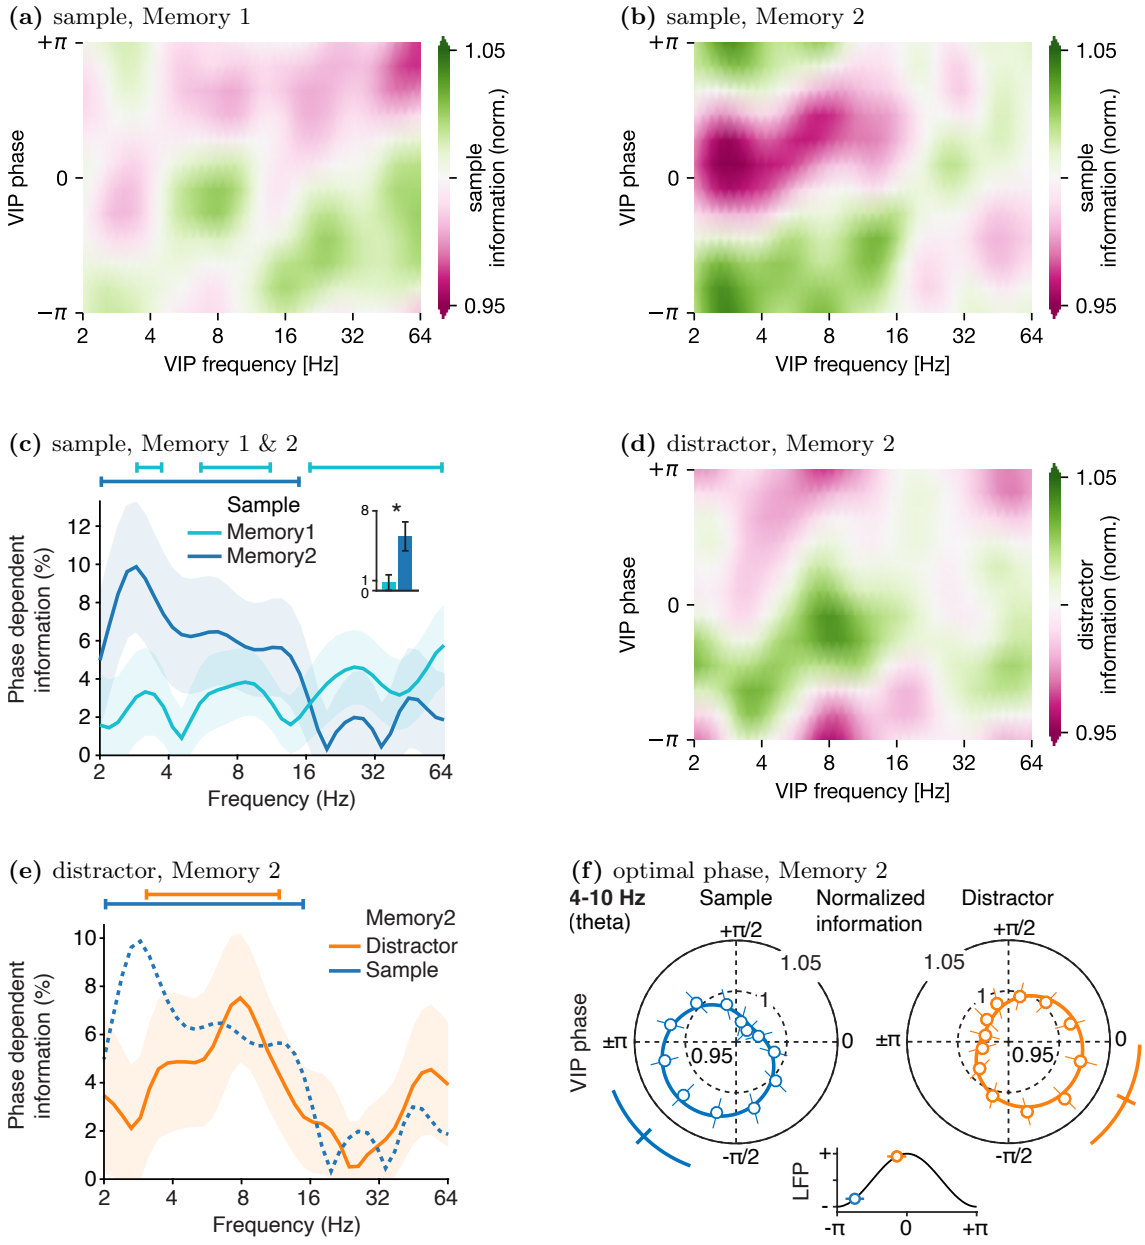

Figure 2.11: Prefrontal numerosity information dependent on parietal oscillation phase. (a) Normalised sample information (ω^2) in spike counts from sample-selective prefrontal neurons at specific phases of VIP LFP during the first memory delay. **(b)** same analysis like (a) during the second memory delay. **(c)** Percentage of phase-dependent sample information in PFC spikes as a function of parietal LFP frequency. Shading: bootstrap SEM. Horizontal bars: significant phase-dependency ($p < 0.01$, permutation test). *(caption continued on next page)*

delta/theta bands during the memory delays (Fig. 2.13b). Notably, however, mediation by PFC oscillations never exceeded 25 %, suggesting that most of the fronto-parietal phase dependent coding was not an artefact of intra-frontal phase-dependency.

All in all, both sample and distractor information was non-uniformly distributed across VIP LFP phases, especially in the theta band. Phase-dependence was especially strong during the second memory delay, when both numerosities were held in working memory and possibly competed for resources. This competition could theoretically be resolved by a downstream receiver population because optimal readout phases were well separated.

2.3.5 Behavioural relevance

Did fronto-parietal connectivity impact behaviour? If so, the presented metrics should co-vary with behavioural performance. Figure 2.14a shows the difference in fronto-parietal LFP synchrony PLV for the fastest and slowest 25 % of correct trials. During the second memory delay, theta synchrony was higher for faster trials, while delta or beta synchrony were not predictive of the response time ($p < 10^{-4}$, Wilcoxon rank sum test, Fig. 2.14b).

Next, the phase-dependence of sample and distractor spike information was assessed for error trials. To ensure statistical comparability trials numbers were matched with the analysis for correct trials (compare Figs. 2.11b and 2.11d). Due to the relatively small number of error trials, a random subset of correct trials was replaced with all available error trials in respective behavioural sessions, thus possibly underestimating effects (on average 37 % of trials were replaced, Siegel et al. (2009)). For sample information the replacement with error trials did not have an effect on the phase dependence in delta and alpha bands (Fig. 2.14c). Optimally encoding phases for these bands as well as the percentage of phase-dependent information were similar to the same analyses that used only correct trials (Fig. 2.14e). Notably, however, sample information in the theta band was less and in the beta band was more phase-dependent than in purely correct trials[13]. The same analysis for the distractor numerosity, again, shows similar phase dependency and optimal phase in the delta band (Figs. 2.14d and 2.14f). In contrast, increased

[13] $p < 0.05$, permutation test

Figure 2.11 ***(continued from previous page)*:** Inset: mean ratio of phase-dependency in the delta and beta bands (whiskers: bootstrap SEM, $p < 0.05$ permutation test). **(d)** same analysis like (b) for distractor information contained in spike counts of distractor-selective neurons. **(e)** same analysis like (c) for the distractor information during the second memory delay (sample-information from (c) for reference). **(f)** Top, polar plots: Normalised prefrontal sample and distractor information as a function of VIP theta phase. Circle markers and associated error bars indicate mean normalised information and bootstrap SEM at the individual phase bin. Solid closed traces are cosine fits. Short solid lines along the polar dimension and orthogonal partial circles indicate optimal readout phases and bootstrap SEM. Bottom: Location of optimal phases with bootstrap SEM on the LFP oscillation in the time domain (cosine). Modified from Jacob et al. (2018).

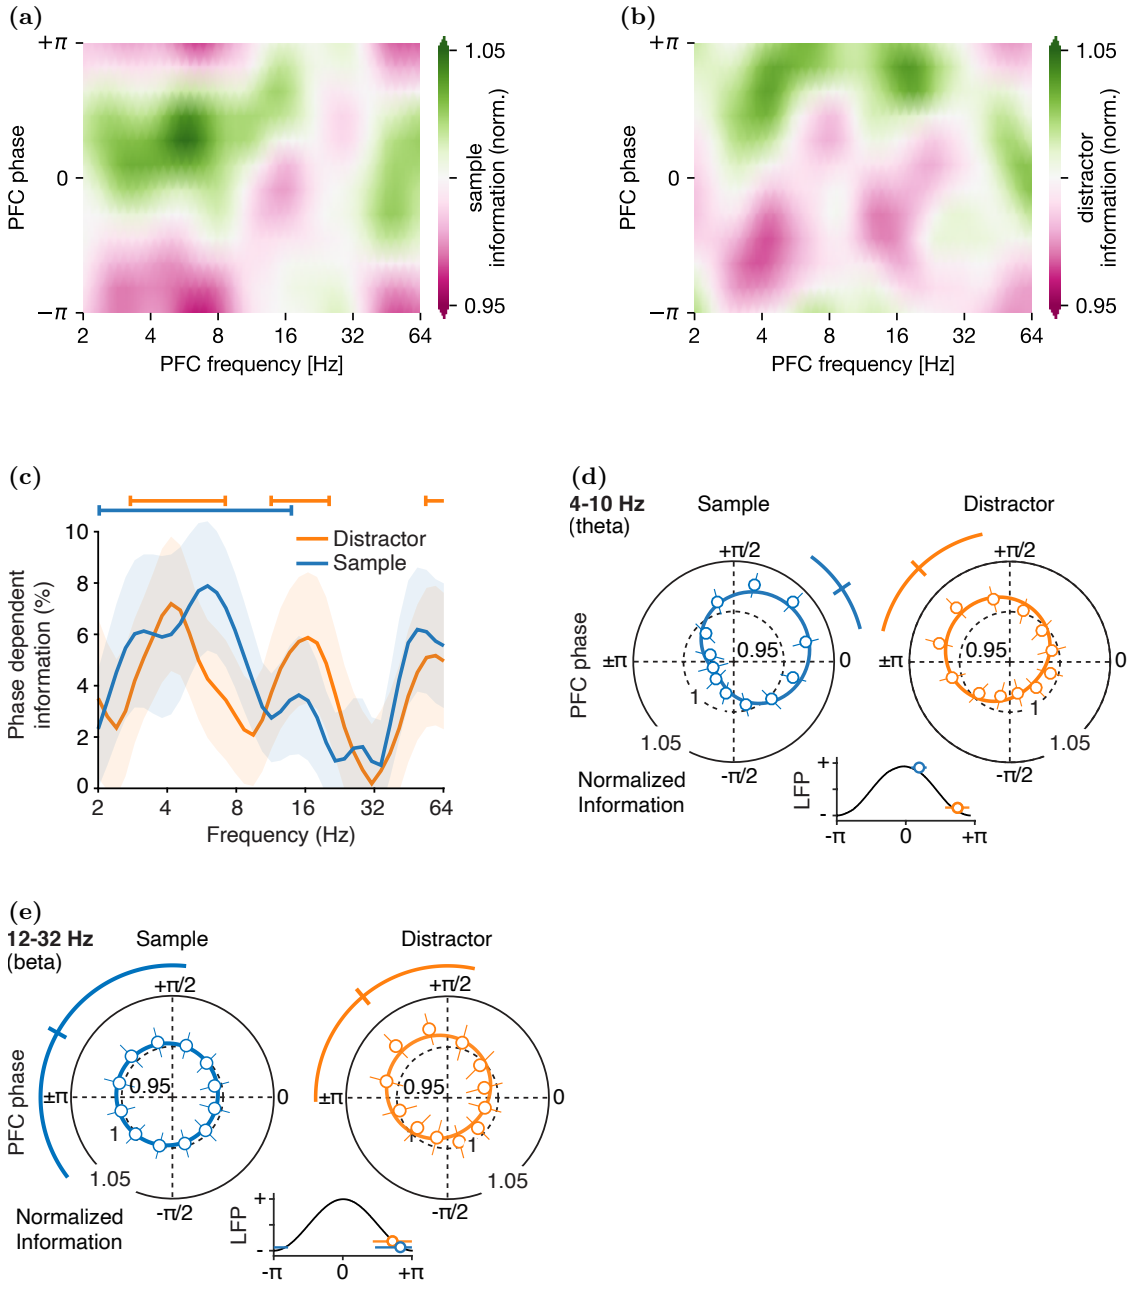

Figure 2.12: Prefrontal numerosity information dependent on prefrontal oscillation phase. (a) Normalised sample information (ω^2) in spike counts from sample-selective prefrontal neurons at specific phases of PFC LFP during the second memory delay. **(b)** same analysis like (a) for distractor information contained in spike counts of distractor-selective neurons. **(c)** Percentage of phase-dependent sample or distractor information in PFC spikes as a function of prefrontal LFP frequency. Shading: bootstrap SEM. *(caption continued on next page)*

Figure 2.12 ***(continued from previous page)*****:** Horizontal bars: significant phase-dependency ($p < 0.01$, permutation test). Inset: mean ratio of phase-dependency in the delta and beta bands (whiskers: bootstrap SEM, $p < 0.05$ permutation test). **(d)** Top, polar plots: Normalised prefrontal sample and distractor information as a function of PFC theta phase. Circle markers and associated error bars indicate mean normalised information and bootstrap SEM at the individual phase bin. Solid closed traces are cosine fits. Short solid lines along the polar dimension and orthogonal partial circles indicate optimal readout phases and bootstrap SEM. Bottom: Location of optimal phases with bootstrap SEM on the LFP oscillation in the time domain (cosine). **(e)** same analysis like (d) for beta frequency band. Modified from Jacob et al. (2018).

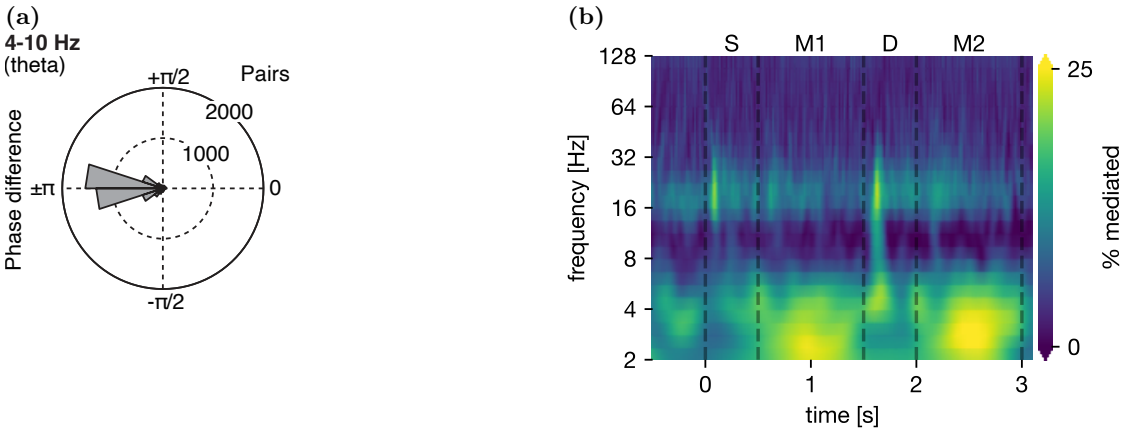

Figure 2.13: Fronto-parietal interactions. (a) Mean phase differences between LFPs from PFC and VIP in the theta frequency band. **(b)** Complex mediation analysis. Percentage of VIP-LFP signal mediated by PFC-LFP and not mediated by PFC-spikes. Modified from Jacob et al. (2018).

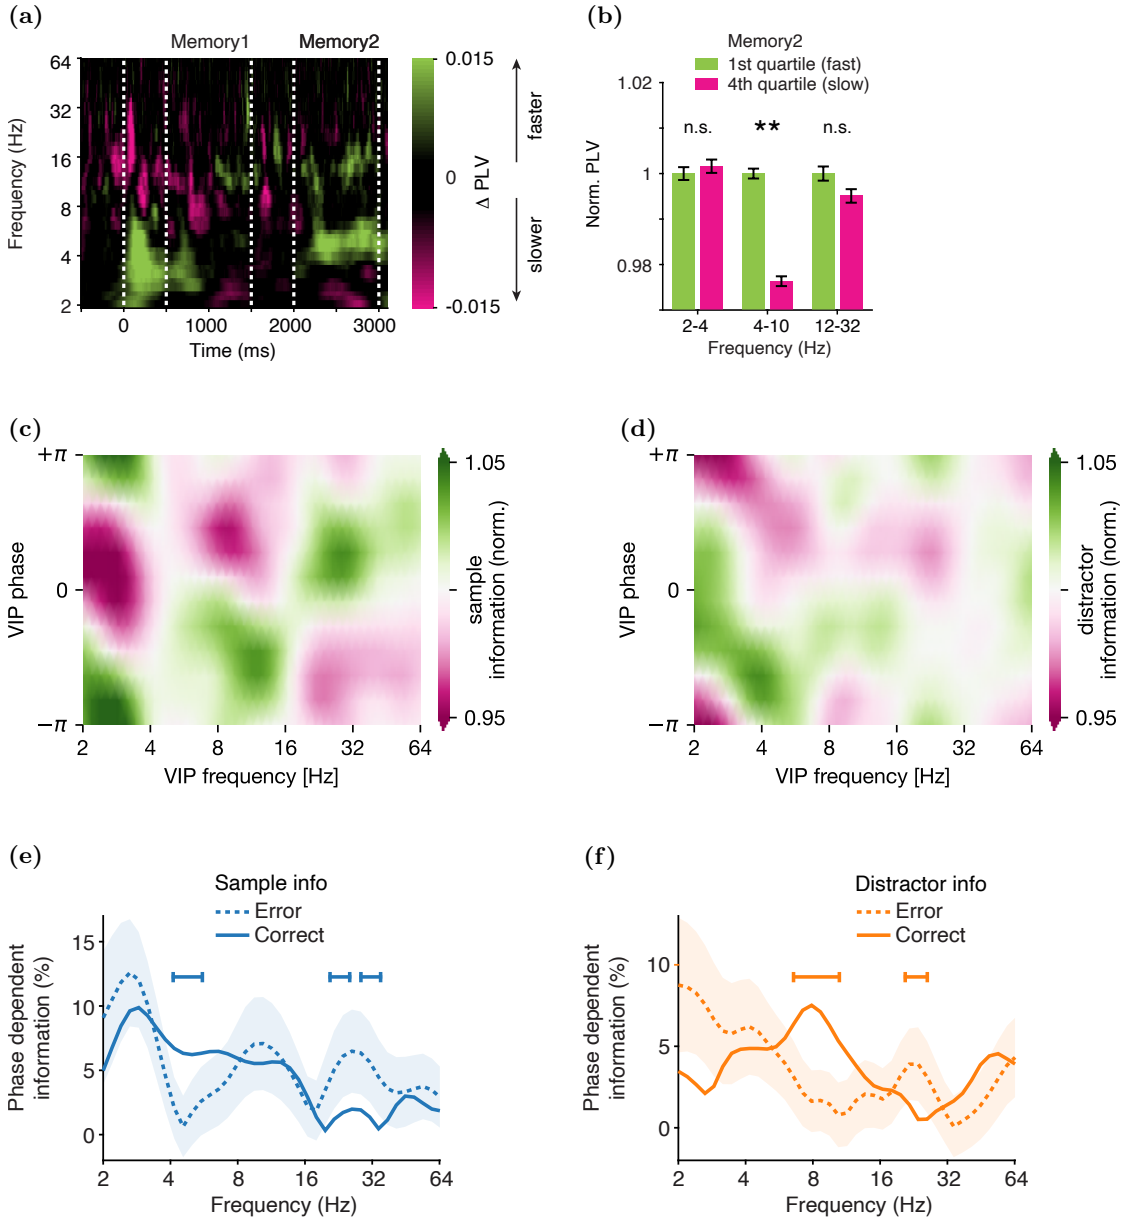

Figure 2.14: Behavioural relevance of fronto-parietal connectivity. **(a)** Difference in fronto-parietal synchrony (PLV) for the fastest and slowest 25 % of trials. **(b)** Normalised fronto-parietal synchrony in the delta, theta and beta bands during the second memory delay for the fastest and slowest 25 % of trials. $**: p < 10^{-4}$, Wilcoxon rank sum test. **(c)** Normalised sample information (ω^2) in spike counts from sample-selective prefrontal neurons at specific phases of VIP LFP during the second memory delay in error trials (compare Fig. 2.11b). **(e)** Percentage of phase-dependent sample information in PFC spikes in error trials as a function of parietal LFP frequency. Correct trials for reference (compare Fig. 2.11e). Shading: bootstrap SEM. Horizontal bars: significant phase-dependency ($p < 0.01$, permutation test). **(d, f)** same analysis like (c, e) for the distractor numerosity. Modified from Jacob et al. (2018).

phase-dependence in the alpha and decreased in the beta band were predictive of error trials.

In summary, fronto-parietal connectivity, especially in the theta band, was predictive of behaviour.

3 Neuronal Signatures of Contextual Decision-Making in Mouse Prefrontal Cortex and Mediodorsal Thalamus

3.1 Behaviour

3.1.1 Response ball task to test executive function in mice

To investigate how sensory information is translated into future action, including its representation in working memory and its utilisation in decision-making, as well as how these processes are influenced by different contexts, a behavioral setup was constructed to train and test head-fixed mice. This setup extended a response device that was originally developed by Sanders and Kepecs (2012) (Fig. 3.1). While in the setup the animals rested their forepaws on a ping pong ball that was the manipulandum of the response device (henceforth called *response ball*). The ball was fixed on an axis parallel to the animal's anterior-posterior axis, thus allowing only left or right rotations via paw movements. Binary choices were indicated by rotating the response ball by a pre-defined angle in one or the other direction (i.e. left or right) until it exceeded a threshold which was equal for correct and wrong choices (see example traces in Fig. 3.5). Distance to the threshold was not explicitly signalled to the animal but rather learned implicitly during training and varied across animals. The distance to the threshold was dependent on the cumulative rotated angle to both sides, i.e. if there was a rotation to the wrong side the required rotation angle to the correct side was effectively increased.

Mice were trained in an auditory two-alternative forced choice (2AFC) task that can be conceptualised into two major trial epochs: a context epoch followed by an instruction epoch (Fig. 3.2).

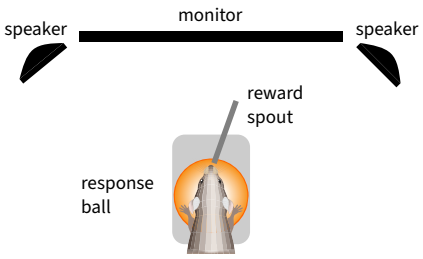

Figure 3.1: Core elements of the behavioural testing setup An animal was head-fixed in front of a computer monitor. Its forepaws rested on the response ball. Auditory stimuli were played from two speakers next to the monitor. The response ball could be rolled left or right. Correct rotation resulted in water dispensal from the reward spout.

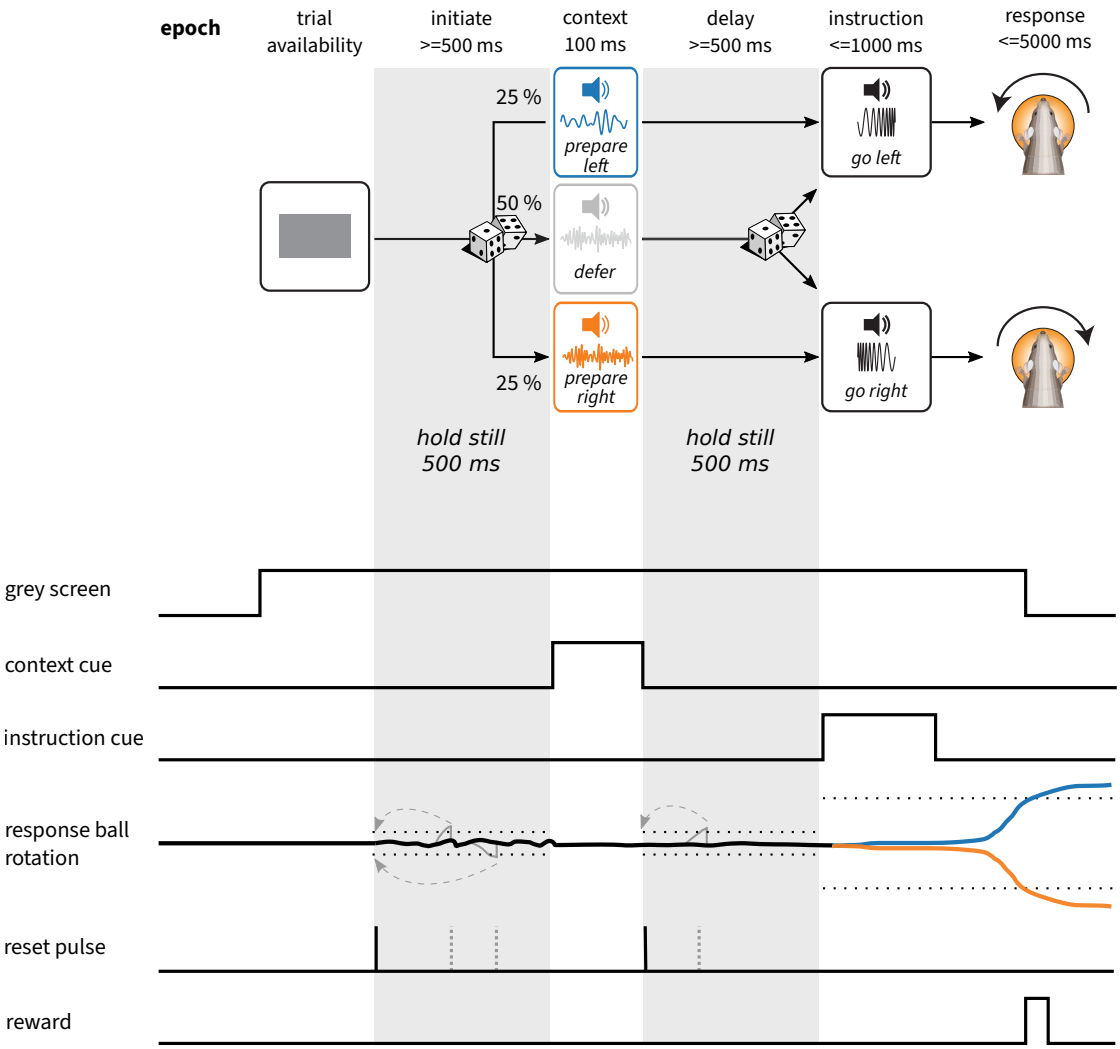

Figure 3.2: Two-alternative forced choice task with continuous read-out of response ball rotations. Head-fixed animals rested their forepaws on a ping-pong ball while being located in front of a computer screen. A grey screen signalled the availability of and being in a trial. To initiate a trial animals were required to hold the ball still for 500 ms. If during hold-still epochs the admissible range (thresholds: dotted lines) was exceeded (grey rotation traces), the recorded angle as well as the hold still timer were reset (grey arrows and dashed reset pulses). In these cases the epochs were extended until a continuous stretch of quiescence was registered. After initiation one of three randomly selected auditory cues was played for 100 ms to inform animals if the upcoming response could be prepared (coloured rectangles) or the choice must be deferred to later in the trial. This was followed by a delay epoch in which animals were again required to hold the ball still for 500 ms. After this delay one of two instruction cues was played for up to 100 ms. If the context was non-predictive, the instruction cue was selected randomly. On hearing the instruction cue the animal could turn the ball left or right. If a threshold of required rotation angle was hit, the potentially still playing instruction cue stopped and the grey screen turned off. Hitting the correct side's threshold lead to dispensal of a water reward. After an incorrect choice as well as not rotating the ball sufficiently far within 5000 ms ("missing") reward was omitted.

Table 3.1: Possible trial conditions and respective percentages of total trials.

	instruction	
context	*go-left* (50 %)	*go-right* (50 %)
prepare-left (25 %)	*left-prepare* (25 %)	- (0 %)
prepare-right (25 %)	- (0 %)	*right-prepare* (25 %)
defer (50 %)	*left-defer* (25 %)	*right-defer* (25 %)

The primary task was an auditory discrimination during the so-called instruction epoch. One of two auditory movement instructions (*go-left* or *go-right*, sinusoidal up- (11 kHz to 14 kHz) or down-sweeps (15 kHz to 12 kHz), respectively, maximum duration 1000 ms, Fig. 7.2) signalled the required response side. Animals could respond as soon as the instruction was started. A correct response by ways of exceeding a rotation threshold was followed by a water reward. An incorrect response immediately aborted the trial.

The instruction epoch was preceded by the so-called context epoch. It comprised a stimulus presentation of fixed duration of 100 ms followed by a delay of at least 500 ms (dependent on the animal's behaviour, see below). Short auditory context cues played at the beginning of the context epoch informed the animal about the upcoming instruction. Two contexts, referred to as *prepare-left* (white-noise lowpass-filtered at 8 kHz) and *prepare-right* (white-noise highpass-filtered at 14 kHz), were predictive of the instruction, giving the animal the possibility to prepare its upcoming action after the instruction. In contrast, the third context cue (unfiltered white-noise) was uninformative, thus I refer to it as the *defer* context because the animal had to defer its choice until the instruction epoch. Predictive and non-predictive contexts were shown pseudo-randomly with equal probability (i.e. 50 % non-predictive, 50 % predictive, of which half were predictive of either the *go-left* or the *go-right* instruction). Together, this constitutes four different trial conditions dependent on combination of context and instruction (Table 3.1).

Importantly, the response ball was required to be held still during the context epoch; animals could not simply continuously rotate it throughout the delay. This was controlled by smaller thresholds set on the cumulative rotation angle. These thresholds were set as small as possible but were not at 0 for practical reasons (physiological tremor, noise in response device). If either threshold was exceeded, the timer for the delay epoch was reset to 0, thus effectively extending the context epoch. Since this delayed the possibility for reward, animals were thus incentivised to keep the delay as short as possible.

3.1.2 Mice learned abstract goal-directed behaviour and used context advantageously

In order to assert if animals had successfully learned the behavioural task and if context cues were used, task performance (percent correct) and response times (time from instruction to successful crossing of rotation threshold) were investigated per condition (instruction × context).

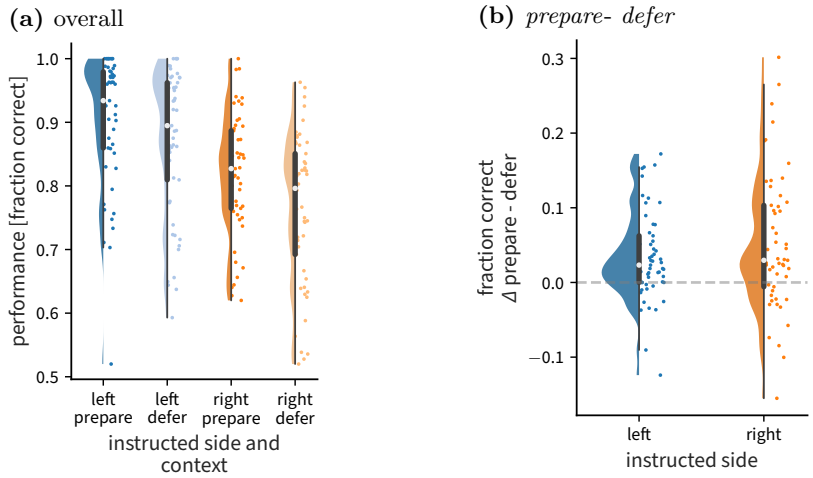

Figure 3.3: Task performance as fraction correct (excluding missing trials) in valid recording sessions. **(a)** Performance per condition (instruction × context). Statistically higher performance in *go-left* trials: repeated-measures two-way ANOVA, main factor instruction $F(1, 52) = 21.21, p = 2.7 \cdot 10^{-5}$; statistically higher performance in *prepare* trials: repeated-measures two-way ANOVA, main factor context $F(1, 52) = 29.50, p = 1 \cdot 10^{-6}$. **(b)** Performance difference of *prepare* and *defer* per instructed side. Context effect was similarly strong in *go-left* and *go-right* trials: Wilcoxon signed-rank test, $T = 611, p = 0.35$. Variance was higher in *go-right* trials: Levene test for equal variances, $W = 5.48, p = 0.02$. Every dot is the mean performance in one session; see methods for description of raincloud plots.

Overall, animals showed a high task performance of 83.98 % (Fig. 3.3a). Of note, animals were significantly better in performing *go-left* trials in comparison to *go-right* trials[1] (Fig. 3.3a). Furthermore, they consistently showed higher performance in trials in which they received a *prepare* context as compared to when they had to defer their decision[2] (Figs. 3.3a and 3.3b).

While the average effect of the *prepare* context on task performance was similar for both instructed sides[3], the range of values was higher for *go-right* trials[4] (Fig. 3.3a).

The response time is the duration from the onset of the instruction to the first point in time when the cumulative response ball angle exceeded the threshold value for the instructed side. Overall, animals responded very fast (Fig. 3.4a, 310.82 ± 104.25 ms, correct trials, session medians). Since this value is well below the maximum duration of the instruction cue (1000 ms), this means that in most correct trials the instruction cue was not fully played.

[1] 88.83 ± 10.39 % (*go-left*) vs. 79.20 ± 9.65 % (*go-right*)
repeated-measures two-way ANOVA, main factor instruction $F(1, 52) = 21.21, p = 2.7 \cdot 10^{-5}$

[2] 86.15 ± 9.77 % (*prepare*) vs. 81.81 ± 11.65 % (*defer*)
repeated-measures two-way ANOVA, main factor context $F(1, 52) = 29.50, p = 1 \cdot 10^{-6}$

[3] Wilcoxon signed-rank test, $T = 611, p = 0.35$

[4] means within instructed side: 3.41 ± 6.20 % (*go-left*) vs. 5.27 ± 9.23 %
Levene test for equal variances, $W = 5.48, p = 0.02$

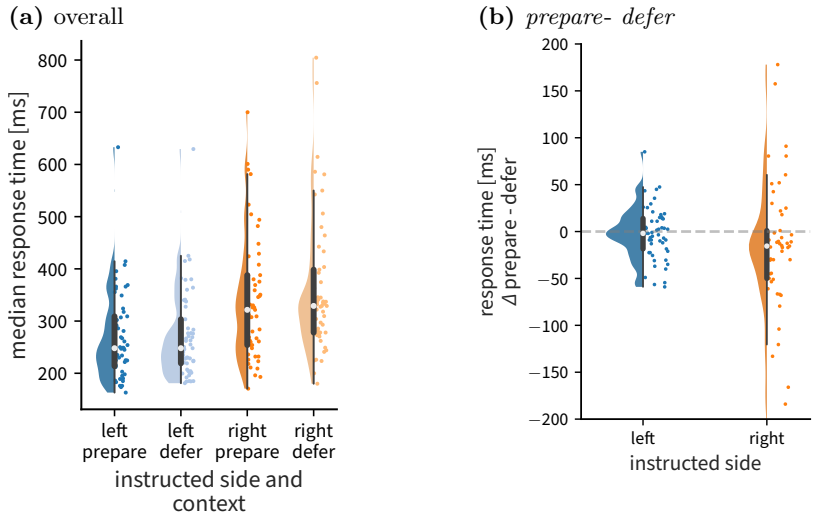

Figure 3.4: Response time (time from instruction onset to threshold crossing) for correct trials in valid recording sessions. **(a)** Response time per condition (instruction × context). Response times were statistically shorter in *go-left* trials: repeated-measures two-way ANOVA, main factor instruction $F(1, 52) = 24.92, p = 7 \cdot 10^{-6}$ **(b)** Response time difference of prepare and defer per instructed side. Context only had an effect on response times in *go-right* trials: Wilcoxon signed-rank tests against 0, $T = 454.5, p = 0.02$ (*go-right*), $T = 650.0, p = 0.56$ (*go-left*). Variance was statistically higher in *go-right* trials: Levene test for equal variances, $W = 10.58, p = 1.54 \cdot 10^{-3}$. Every dot is the median response time in one session; see methods for description of raincloud plots.

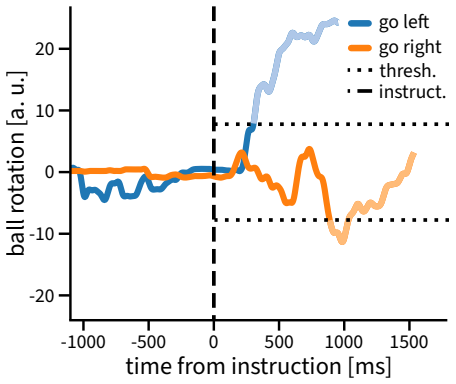

Figure 3.5: Response ball example traces. Traces coloured blue and orange represent cumulative rotation (y-axis) of the response ball across time (x-axis) in one example trial each with *go-left* and *go-right* movement instruction, respectively. Vertical dashed line represents the time of the onset of the movement instruction (*go-left* or *go-right*), before which ($t < 0$) was the context epoch and after which ($t > 0$) was the instruction epoch. Horizontal dotted lines represent the (implicit, hidden to the animal) thresholds to which the animal was required to rotate the response ball. Parts of the response ball traces after the first crossing of the thresholds were desaturated to signify that this was not relevant for the outcome of the task.

Animals were faster in *go-left* trials, consistent with the higher task performance in those trials[5] (Fig. 3.4a). A prior *prepare* context decreased response times in trials with *go-right*[6] but not *go-left*[7] instructions[8] (Fig. 3.4b). Similar to task performance the effect of context was more variable across sessions for *go-right* trials[9] (Fig. 3.4b).

Together, this shows that animals successfully learned the mapping of an abstract behavioural response and external stimuli. Context cues were used advantageously and increased trial performance and, to a lesser extent, response times. Notably, behaviour in trials with *go-left* and *go-right* responses significantly differed.

3.1.3 Covert movements during context epoch correlated with task performance and response times

Using the continuous read-out of the response ball's rotation throughout the entire trial, the decision-making process can be dissected in more detail. Figure 3.5 shows to cumulative response ball rotation across time for one trial each with a *go-left* and *go-right* movement instruction. In the figure the traces were aligned to movement instruction onset. Thus, parts of the traces that are at times $t < 0$ are part of the context epoch, while parts at times $t >= 0$ are part of the instruction epoch.

In the instruction epoch ($t >= 0$) the two traces not only differ in the side at which the threshold was crossed, but also in the movements that were made beforehand. While the blue trace for the *go-left* trial shows a relatively straight rotation towards the threshold after a short delay, the orange trace for the *go-right* trial is more complex. Here, the ball was slowly rotated towards the right threshold, then moved back quickly and, finally,

[5] 269.82 ± 83.58 ms (*go-left*) vs. 344.72 ± 144.68 ms (*go-right*)
repeated-measures two-way ANOVA, main factor instruction $F(1, 52) = 24.92, p = 7 \cdot 10^{-6}$

[6] Wilcoxon signed-rank test against 0 (*go-right*), $T = 454.5, p = 0.02$

[7] Wilcoxon signed-rank test against 0 (*go-left*), $T = 650.0, p = 0.56$

[8] per side differences prepare - defer of session medians: -1.89 ± 27.85 ms (*go-left*) vs. -21.78 ± 83.01 ms (*go-right*)

[9] Levene test for equal variances, $W = 10.58, p = 1.54 \cdot 10^{-3}$

quickly moved to cross the right threshold. The desaturated part of the trace shows that after the initial threshold crossing the ball was slowly rotated back left. In contrast, the post-threshold trace of the *go-left* trial (desaturated blue) shows that in that trial the ball was further rotated left even after threshold crossing. Since the trial outcome was determined by the side of the first threshold crossing the rotation after threshold crossing was not subject to any reinforcement (hence the desaturated colouring).

The parts of the traces before the onset of the movement instruction ($t < 0$) are parts of the context epoch. Animals were required to hold the response ball still during this epoch, which translated into smaller measured rotations. Due to technical reasons movements within a very small "grace threshold" were not penalised, which is demonstrated by the fact that at $t = 0$ the ball rotation values of two example traces are not at exactly 0. Furthermore, the two examples demonstrate a relatively active context epoch for the *go-left* trial, while there was almost no movement in the *go-right* trial.

The small grace thresholds of the response ball during the context epoch open up the opportunity to investigate micro-movements and their influence on trial outcome and response time. Thus, in the following these psychometric parameters were investigated in subsets of trials. First, the relative movement activity 0 ms to 600 ms from context cue onset was used to classify trials by active or calm context epoch. Second, the dominant *side* of movement 0 ms to 600 ms from context cue onset was captured by the movement tendency that defines subsets of trials with dominant left-ward movement (*leftish*) and dominant right-ward movement (*rightish*).

Movement activity Figure 3.6a shows that, on average, the task performance was not affected by how much an animal moved during the context epoch[10]. In contrast, response times differed dependent on context epoch activity (Fig. 3.6b): Responses in the instruction epoch were consistently faster if animals had been more active before[11].

Movement side tendency Context movement side tendency could be either consistent or inconsistent with the upcoming movement instruction (Table 3.2). Concretely, if during the context epoch there was more covert movement to the left (*leftish* context epoch) and the later movement instruction was *go-left*, context movement tendency was consistent with instruction. Conversely, had context movement tendency been *rightish* in *go-left* trials, they would be inconsistent.

Figure 3.7a shows a clear effect of context epoch tendency on task performance. Except for *go-left-prepare* trials[12] all conditions had a higher task performance if the tendency was *consistent* with the upcoming instruction[13].

[10] Wilcoxon signed-rank test: $p_{go\text{-}left-prepare} = 0.118$, $p_{go\text{-}left-defer} = 0.193$, $p_{go\text{-}right-prepare} = 0.539$, $p_{go\text{-}right-defer} = 0.901$

[11] Wilcoxon signed-rank test: $p_{go\text{-}leftprepare} = 0.002$, $p_{go\text{-}leftdefer} = 9.37 \cdot 10^{-7}$, $p_{go\text{-}rightprepare} = 1.84 \cdot 10^{-8}$, $p_{go\text{-}rightdefer} = 5.24 \cdot 10^{-6}$

[12] Wilcoxon signed-rank test, $p_{go\text{-}left-prepare} = 0.7$

[13] Wilcoxon signed-rank test, $p_{go\text{-}left-defer} = 0.022$, $p_{go\text{-}right-prepare} = 0.018$, $p_{go\text{-}right-defer} = 0.016$

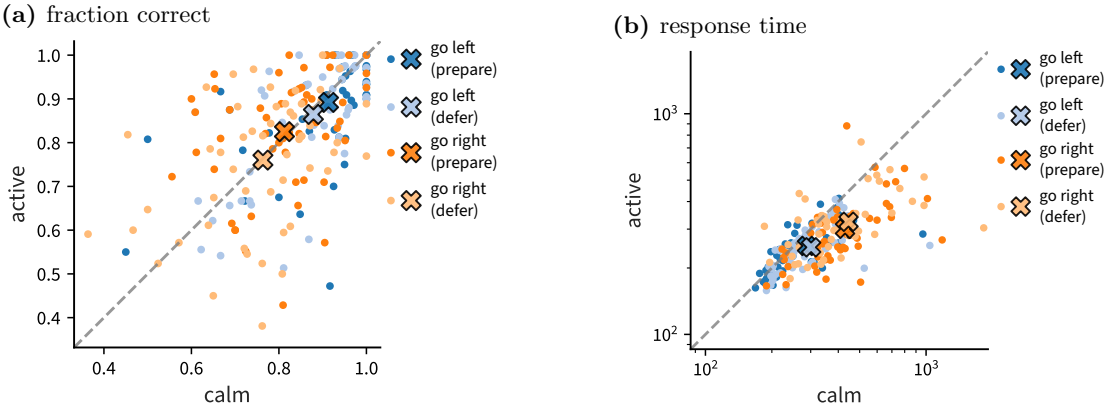

Figure 3.6: Performance and response times dependent on movement *activity* (active and calm) during the context epoch (0 ms to 600 ms). (a) Mean task performances per trial condition in calm and active trials. Performance was not significantly influenced by context movement activity: Wilcoxon signed-rank test: $p_{go\text{-}left-prepare} = 0.118$, $p_{go\text{-}left-defer} = 0.193$, $p_{go\text{-}right-prepare} = 0.539$, $p_{go\text{-}right-defer} = 0.901$. **(b)** Median response times per trial condition in calm and active trials. Response time was significantly shorter in context-active trials: Wilcoxon signed-rank test: $p_{go\text{-}leftprepare} = 0.002$, $p_{go\text{-}leftdefer} = 9.37 \cdot 10^{-7}$, $p_{go\text{-}rightprepare} = 1.84 \cdot 10^{-8}$, $p_{go\text{-}rightdefer} = 5.24 \cdot 10^{-6}$. Every dot represents the average per session and trial condition. Colours represent trial conditions. Big × markers represent the mean across sessions.

In contrast to performance, the response times were not consistently affected by movement tendency during the context epoch[14] (Fig. 3.7b).

Together, these findings show that very simple trial classifications based on the continuous read-out of response ball roll provide a more detailed view onto the decision-making process.

[14] Wilcoxon signed-rank test, $p_{go\text{-}left-prepare} = 0.750$, $p_{go\text{-}left-defer} = 0.646$, $p_{go\text{-}right-prepare} = 0.004$, $p_{go\text{-}right-defer} = 0.311$

Table 3.2: Consistency of context epoch movement side tendency with upcoming movement instruction.

	instruction	
side tendency (context epoch)	*go-left*	*go-right*
leftish	consistent	inconsistent
rightish	inconsistent	consistent

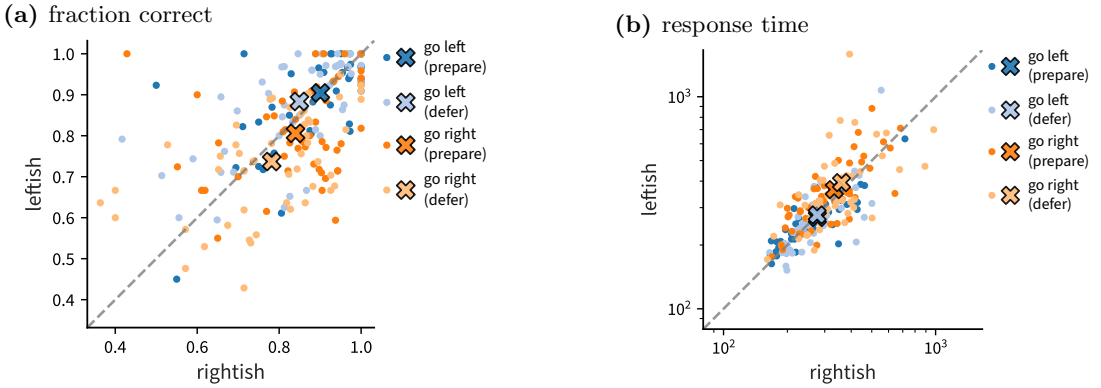

Figure 3.7: Performance and response times dependent on movement *tendency* (leftish and rightish) during the context epoch (0 ms to 600 ms). (a) Mean task performances per trial condition in leftish and rightish trials. Performance was significantly better in trials in which tendency was consistent with instruction, except for *go-left prepare* trials: Wilcoxon signed-rank test, $p_{go\text{-}left-prepare} = 0.7$, $p_{go\text{-}left-defer} = 0.022$, $p_{go\text{-}right-prepare} = 0.018$, $p_{go\text{-}right-defer} = 0.016$. **(b)** Median response times per trial condition in leftish and rightish trials. Response times were not significantly different depending on tendency, except for *go-right prepare* trials: Wilcoxon signed-rank test, $p_{go\text{-}left-prepare} = 0.750$, $p_{go\text{-}left-defer} = 0.646$, $p_{go\text{-}right-prepare} = 0.004$, $p_{go\text{-}right-defer} = 0.311$. Every dot represents the average per session and trial condition. Colours represent trial conditions. Big × markers represent the mean across sessions.

3.1.4 Submovement decomposition

3.1.4.1 Response ball rotation traces can be decomposed into individual submovements

While the simple approach described in the previous Section already demonstrated the usefulness of a continuous readout of movement, the ball rotation data allowed for even finer dissection of the ongoing processes. Previous work in humans, non-human primates and, more recently, mice (Flash & Hogan, 1985; Flash & Hochner, 2005; Bollu et al., 2018) has established that complex limb movements are composed of multiple simple kinematic primitives. These discrete primitives or submovements (SMs) can be conceptualised as realisations of a more continuous, ongoing decision making process that is comprised of micro-decisions (Fishbach et al., 2007).

Leveraging this conceptual framework, I developed an efficient algorithm that fits bell-shaped so-called minimum-jerk SMs to the continuous ball rotation traces (see *Submovement decomposition* for detailed description). While minimum-jerk SMs were chosen for their simplicity, other primitives are viable alternatives: e.g. the support-bounded log-normal (Plamondon et al., 1993; Rohrer & Hogan, 2006) or be based on extrema of temporal derivatives of trajectories speed minima and curvature (Viviani & Terzuolo, 1982; speed minima and curvature: Bollu et al., 2018; jerk and snap extrema ("soft-symmetry"): Fishbach et al., 2005).

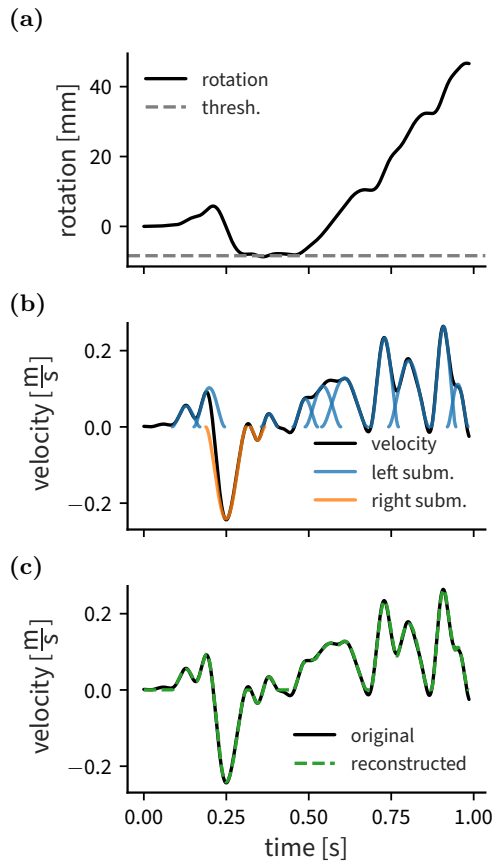

Figure 3.8: Decomposition of the ball rotation time-series into kinematic primitives. Shown is a short segment from an example trial, starting from the instruction cue onset. **(a)** Rotation or path length of moved ball circumference in mm (positive and negative values for left and right rotations). Horizontal dashed line represents the (correct) threshold that had to be crossed. Trace after the first crossing is when the animal received and consumed its reward. **(b)** Ball velocity with minimum-jerk submovements fitted (blue and orange). **(c)** Original velocity and the reconstruction from summed submovements in (b)

In the approach chosen here, the response ball rotation trace (Fig. 3.8a) is first converted into a (rotational) velocity trace (Fig. 3.8b, black trace). The algorithm then iteratively fits minimum-jerk SMs to the velocity trace (coloured bell-shaped curves in Fig. 3.8b). Since the response ball could be rotated left or right around the anterior-posterior rotational axis, individual SMs can be either leftward or rightward. In Fig. 3.8b this is represented by colouring of SMs in blue and orange, respectively, as well as positive and negative velocities. More complex velocity profiles are realised by the partial overlapping of individual SMs of different strength (peak velocity) and duration. Finally, super-positioning all extracted SMs reconstructs (green dashed curve in Fig. 3.8c) the original velocity trace with high fidelity.

3.1.4.2 Prior context decreased delay to response initiation but did not affect duration of execution

Response times are often used as a proxy for the duration of decision-making. However, they measure a mixture of at least two components: First, the time after a stimulus that is required to start a motor response, and, second, the time it takes to execute

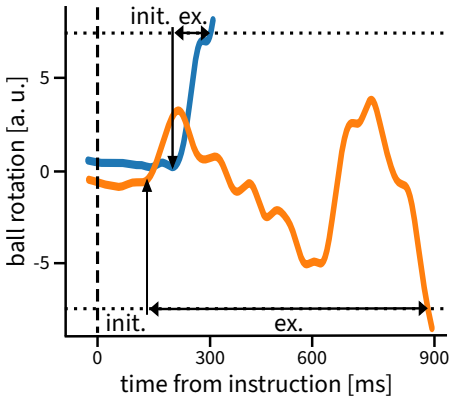

Figure 3.9: Explanation of response initiation and execution. Two example response ball rotation traces across time (compare Fig. 3.5). Response initiation (*init.*) is the point in time after movement instruction onset at which the first submovement occurred. Response execution (*ex.*) is the duration from response initiation until the first crossing of the response threshold. In the examples response initiation is earlier in the orange than in the blue trace. Conversely, response execution is much shorter in the blue than in the orange trace. Vertical dashed line: movement instruction onset. Horizontal dotted lines: rotation thresholds.

the movement that is part of the response. Arguably, the first component, the response initiation, is more concerned with internal processes of decision-making, while the second component, the response execution, is more a reflection of mechanical processes (Dotan et al., 2019).

Therefore, to better understand the response after a movement instruction was shown, and to investigate if prior contexts differentially affected different components of the response, I used the extracted SMs to define the response initiation and execution. Response initiation is the time of the onset of the first SM, while the response execution is the duration from that onset to the first crossing of the rotation thresholds.

In the explanatory Figure 3.9 both components differ quantitatively for the two examples. The blue rotation trace of a *go-left* example trial has slightly later response initiation but a very short, i.e. fast, response execution. In contrast, the orange trace of a *go-right* example trial has an earlier response initiation time but a lengthy, convoluted response, and thus, a long response execution time.

In most sessions the median response *initiation* was within the first 98.40 ± 27.29 ms after the instruction onset. Responses in *go-left* trials were initiated earlier than in *go-right* trials[15] (Fig. 3.10a).

Similarly, the response *execution* time was faster in *go-left* trials[16] (Fig. 3.10b). A prior *prepare* context made response *initiation* faster in *go-left*[17] but not *go-right* trials[18]. Conversely, response *execution* was not affected by prior context[19].

In sum, the continuous response ball data extends simple response times by enabling the dissection of the response into initiation and execution. Consistent with the compound response times, it was found that responses in *go-left* trials were initiated and executed earlier and faster. Additionally, the dissection found a distinctive effect of prior context only on the initiation of *go-left* trials.

[15] 89.88 ± 38.46 ms (*go-left*) vs 106.88 ± 35.35 ms (*go-right*)
repeated-measures two-way ANOVA, main factor instruction $F(1, 49) = 19.41, p = 5.7 \cdot 10^{-5}$

[16] 173.48 ± 53.72 ms (*go-left*) vs 222.950 ± 85.315 (*go-right*))
repeated-measures two-way ANOVA, $F(1, 49) = 13.89, p = 5.03 \cdot 10^{-4}$

[17] Wilcoxon signed-rank test against 0: $T = 334.5$, $p = 0.003$

[18] Wilcoxon signed-rank test against 0: $T = 589.0$, $p = 0.815$

[19] Wilcoxon signed-rank test against 0: *go-left*: $T = 612.0$, $p = 0.806$; *go-right*: $T = 498.0$, $p = 0.178$

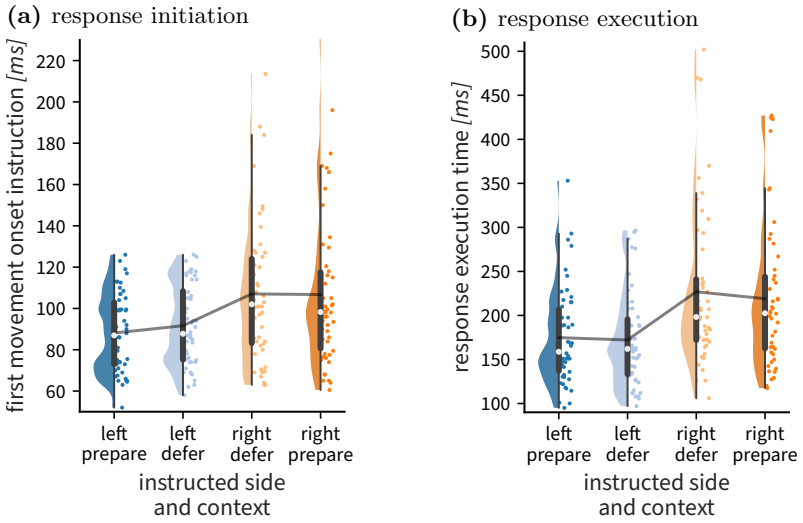

Figure 3.10: Response initiation and execution for correct trials (during the instruction epoch). **(a)** Response initiation, the time from instruction cue onset to the first extracted sub-movement. Initiation was significantly earlier in *go-left* than *go-right* trials: repeated-measures two-way ANOVA, main factor instruction $F(1, 49) = 19.41, p = 5.7 \cdot 10^{-5}$. Prior *prepare* context was associated with significantly earlier initiation of *go-left* but not *go-right* responses: Wilcoxon signed-rank test against 0: *go-left*: $T = 334.5$, $p = 0.003$ *go-right*: $T = 589.0$, $p = 0.815$. **(b)** Response execution, response initiation to exceeding the rotation threshold. Execution was significantly earlier in *go-left* than *go-right* trials: repeated-measures two-way ANOVA, $F(1, 49) = 13.89, p = 5.03 \cdot 10^{-4}$). Prior context did not significantly affect the execution: Wilcoxon signed-rank test against 0: *go-left*: $T = 612.0$, $p = 0.806$; *go-right*: $T = 498.0$, $p = 0.178$. Every dot is the median time in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

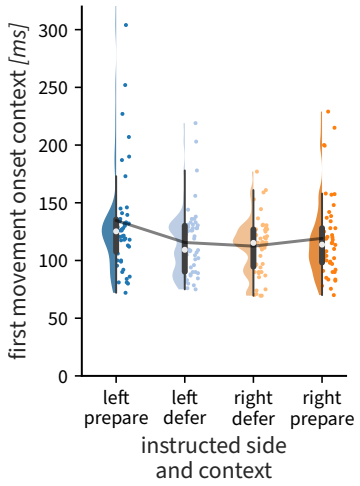

Figure 3.11: Movement initiation (context epoch), the time of onset of the first extracted submovement, for correct trials. Every dot is the median time in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

3.1.4.3 Informative contexts delayed initiation of covert submovements in the context epoch

As stated earlier, the task required animals to keep the response ball still during the context epoch, in order to proceed to the instruction epoch. Small movements that were nevertheless acceptable by the task logic could be decomposed into SMs. To see if timing of covert SMs was affected by the context, the onsets of the first SMs were quantified.

Figure 3.11 shows that in most sessions the median initiation of a detectable submovement was at 120.46 ± 37.41 ms, that is, shortly after the offset of the context cue. Variability of initiation times across sessions was comparable to that in the instruction epoch, demonstrating good detectability of sub-threshold SMs. The presentation of a *prepare* context consistently delayed the onset of SMs[20].

3.1.4.4 Default surplus of left submovements and context-modulation of only right submovements

Movements to both sides could occur in any condition, e.g. the *go-left* trials were not restricted to movements that only went to the left side. What was the balance of SMs to either side at different trial epochs and how did context and instruction influence it? Using the decomposition of velocity traces, one can independently investigate the occurrence of these events.

In the following, the aforementioned balance was quantified per trial by the difference of left submovements (LSMs) and right submovements (RSMs) normalised by their sum. This side index is defined from -1 to 1, with 0 indicating perfect balance and positive values indicating a surplus of SMs to the left.

During the delay after presentation of the context cue animals were required to hold the response ball still for 500 ms in order to proceed to the instruction epoch. Rotation

[20] 126.92 ± 43.53 ms (prepare) vs. 114.01 ± 27.28 ms (defer)
repeated-measures two-way ANOVA, main factor context $F(1, 49) = 7.72, p = 7.7 \cdot 10^{-3}$

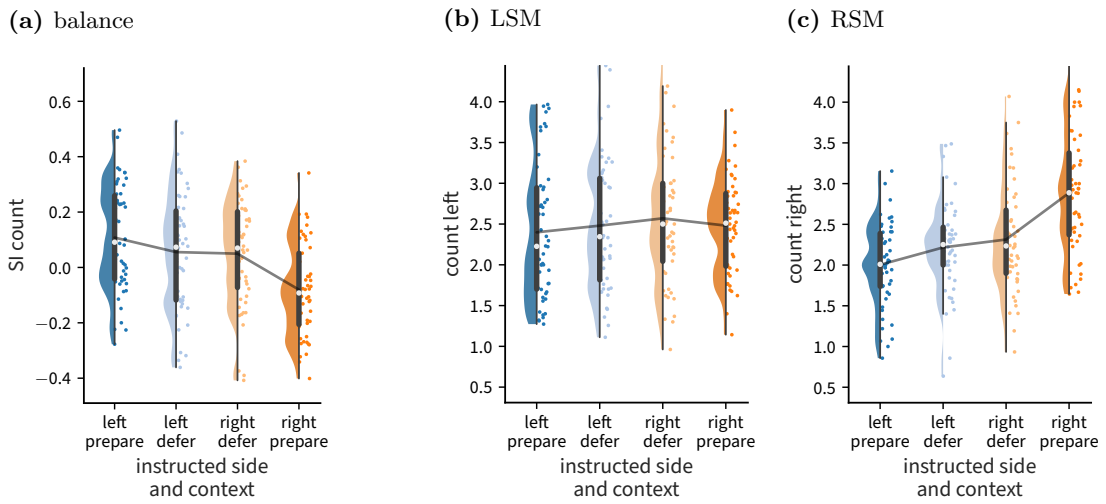

Figure 3.12: Submovement counts (context epoch, 0 ms to 600 ms) for correct trials. **(a)** Movement balance (side index, SI), trial-wise difference of LSMs and RSMs normalised by their sum. Positive values indicate LSM surplus, negative values RSM surplus. Side index was significantly modulated by upcoming instruction and context: repeated-measures two-way ANOVA: main factor upcoming instruction, $F(1, 49) = 24.94, p = 8 \cdot 10^{-6}$, interaction context $\times$ upcoming instruction, $F(1, 49) = 21.91, p = 2.3 \cdot 10^{-5}$. **(b)** Number of LSMs. Number was statistically invariant to conditions: repeated-measures two-way ANOVA, main factor upcoming instruction, $F(1, 49) = 1.478, p = 0.230$, interaction context $\times$ upcoming instruction, $F(1, 49) = 0.004, p = 0.948$ **(c)** Number of RSMs. Number was statistically varied with prepare-side and context: repeated-measures two-way ANOVA, main factor upcoming instruction, $F(1, 49) = 43.35, p = 2.93 \cdot 10^{-8}$, interaction context $\times$ upcoming instruction, $F(1, 49) = 34.61, p = 3.55 \cdot 10^{-7}$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

of the response ball across small thresholds would have extended the delay after the presentation of the context cue (Fig. 3.2).

Notably, in this epoch (0 ms to 600 ms from context cue onset) the side index varied with the conditions: both context and prepare-side significantly modulated the balance[21] (Fig. 3.12). Both *prepare* contexts pushed that balance to their respective sides, that is, the side index for *prepare-left* was the most positive (0.11 ± 0.19) while the negative values in *prepare-right* trials (-0.08 ± 0.17) indicate a surplus of RSMs. Moreover, positive average values in the *defer* conditions (0.05 ± 0.20) indicate a surplus of LSMs in the non-informative context. Interestingly, the extra modulation from this supposed baseline in *prepare-left* trials was weaker than for *prepare-right* trials (differences prepare-defer per side: 0.05 ± 0.18 (left) vs -0.13 ± 0.20 (right)).

How were these different balances implemented in terms of SMs to individual sides? Was the occurrence of SMs to one or the other side differentially modulated within var-

[21] repeated-measures two-way ANOVA: main factor upcoming instruction, $F(1, 49) = 24.94, p = 8 \cdot 10^{-6}$, interaction context $\times$ upcoming instruction, $F(1, 49) = 21.91, p = 2.3 \cdot 10^{-5}$

ious contexts? Remarkably, the number of LSMs was invariant to the type of context[22] (Fig. 3.12b). Instead, all the aforementioned modulations of balance were implemented by changes in the number of SMs to the right, which varied with context as well as with prepare-side[23] (Fig. 3.12c). This detailed view also shows that the increase of the number of RSMs in the *prepare-right* context was stronger than the decrease in numbers in the *prepare-left* context, consistent with what was found for the balances.

Together, quantification of the number of SMs in the context epoch discovered a default imbalance towards more left SMs that is only inverted in *prepare-right* trials. Furthermore, the balance between SMs was only modulated by changes in the numbers of right SMs, while the number of LSMs remained unaffected by the different conditions.

3.1.4.5 Submovement balance consistent with instructed side but distinct modulation of left and right submovements

While the previous subsection showed a possibly unexpected preparatory effect on behaviour during the context epoch, correct responses in the instruction epoch are expected to be correlated with large movement balances to the instructed side. Figure 3.13a shows that this was the case for the time from instruction onset to the detection of a threshold crossing.

In trials with the *go-left* instruction the balance was shifted to very positive (i.e. leftish) values (0.580 ± 0.208), while, conversely, it was shifted to very negative (rightish) values in *go-right* trials[24] (Fig. 3.13a, -0.60 ± 0.15). As opposed to what was found in the context epoch, however, there was no effect of context in the instruction epoch[25]. Notably, the absolute values clustered around 0.59, indicating that responses were not purely composed of SMs to the instructed side.

Absolute values of movement balances were similar for *go-left* and *go-right* trials, which could be a sign of similar implementations by mixing left and right SMs. Zooming in to the occurrence of SMs to individual sides, to the contrary, indicated a slightly different picture. As expected, both the occurrence of LSMs (Fig. 3.13b) and RSMs (Fig. 3.13c) varied with the instruction side: more LSMs were made in *go-left* than in *go-right* trials[26], and vice-versa for SMs to the right[27]. However, the modulation was stronger for RSMs (Fig. 3.13c) than for LSMs. This was achieved by a higher number of RSMs in *go-right* trials than the number of LSMs in *go-left* trials. In addition, the *prepare-right* context was associated with a specific further down-regulation of SMs to the left[28].

---

22 repeated-measures two-way ANOVA, main factor upcoming instruction, $F(1, 49) = 1.478, p = 0.230$, interaction context × upcoming instruction, $F(1, 49) = 0.004, p = 0.948$

23 repeated-measures two-way ANOVA, main factor upcoming instruction, $F(1, 49) = 43.35, p = 2.93 \cdot 10^{-8}$, interaction context × upcoming instruction, $F(1, 49) = 34.61, p = 3.55 \cdot 10^{-7}$

24 repeated-measures two-way ANOVA, main factor instruction: $F(1, 49) = 995.31, p = 3.26 \cdot 10^{-34}$

25 repeated-measures two-way ANOVA, interaction instruction × context: $F(1, 49) = 0.08, p = 0.37$

26 3.11 ± 0.58 (*go-left*) vs. 1.39 ± 0.77 (*go-right*), repeated-measures two-way ANOVAs, main factor instruction: LSMs: $F(1, 49) = 236.24, p = 2.25 \cdot 10^{-20}$

27 1.24 ± 0.72 (*go-left*) vs. 3.70 ± 0.82 (*go-right*) repeated-measures two-way ANOVAs, main factor instruction: RSMs: $F(1, 49) = 467.70, p = 1.03 \cdot 10^{-26}$

28 repeated-measures two-way ANOVA, interaction instruction × context: $F(1, 49) = 5.17, p = 0.027$

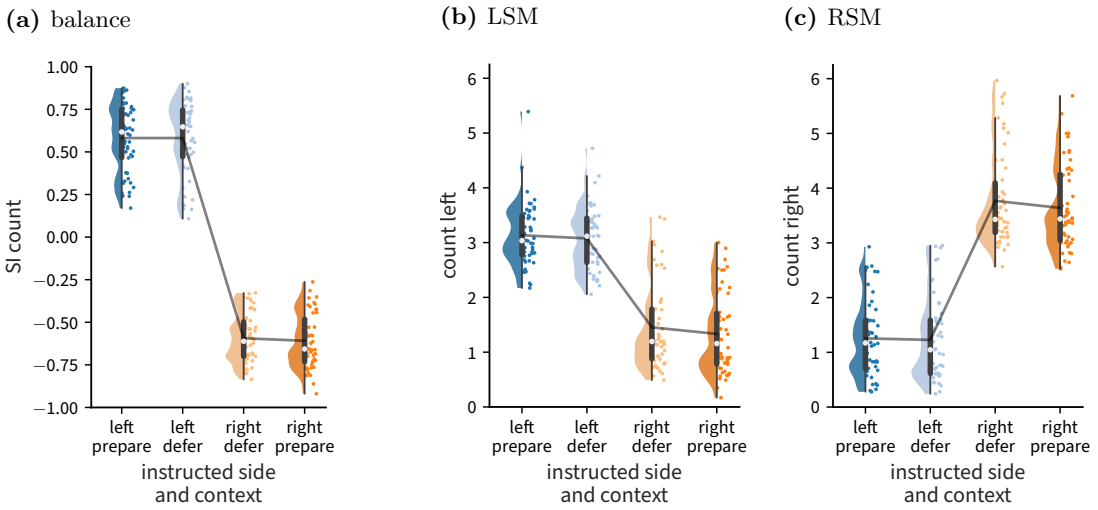

Figure 3.13: Submovement counts (instruction epoch, from instruction onset to threshold crossing) for correct trials. **(a)** Movement balance, trial-wise difference of LSMs and RSMs normalised by their sum. Positive values indicate LSM surplus, negative values RSM surplus. Balances significantly varied with instruction but not with prior context: repeated-measures two-way ANOVA, main factor instruction, $F(1, 49) = 995.31, p = 3.26 \cdot 10^{-34}$), interaction instruction $\times$ context, $F(1, 49) = 0.08, p = 0.37$. **(b)** Number of LSMs. **(c)** Number of RSMs. Both number of LSMs and RSM significantly varied with instruction: repeated-measures two-way ANOVAs, main factor instruction: LSMs: $F(1, 49) = 236.24, p = 2.25 \cdot 10^{-20}$; RSMs: $F(1, 49) = 467.70, p = 1.03 \cdot 10^{-26}$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

Together, this shows that the balance of SM numbers during the instruction epoch was consistent with the instructed side and not influenced by prior context. While modulation happened for both the number of left and right SMs it was stronger for the latter.

3.1.4.6 Submovement side-specific effects of context on strength and duration

Adjusting the number of discrete submovements to sway the balance to one or the other side is one way of motor preparation or achieving the goal of rotating the response ball to the desired side. Another way is to employ SMs with different scales. In order to investigate if the animals also modulated the SMs properties, in the following they were quantified on the level of strength and duration.

Similar to their count (Fig. 3.12b) the strength of SMs to the left was largely invariant to trial conditions (Fig. 3.14a), with the exception of a weak increase for both *prepare* contexts[29].

On the other hand, their duration did covary significantly with different contexts (Fig. 3.15a): LSMs were longer-lasting if the animals were informed that the upcoming instruction would be left and, conversely, shorter-lasting if they were informed of the opposite[30].

In contrast, the strength of RSMs varied consistently with the effects on their counts[31] (Fig. 3.14b). They were strongest in the *prepare-right* context and weakest in the *prepare-left* context[32]. Of note, the strength in the *prepare-left* context was more similar to the *defer* trials, reminiscent of the effect for the counts of SMs.

RSM *durations* also varied with conditions, but in a different pattern (Fig. 3.15b). As expected, duration was longest for the *prepare-right* trials (71.98 ± 6.05 ms). Notably, however, trials in the *prepare-left* context did not show the shortest-lasting RSMs (69.11 ± 7.69 ms). Instead, there was a pattern which only involved the main factors, i.e. *prepare* context trials and the trials with upcoming *go-right* instruction had longer durations[33].

Together, this data shows that preparation in the context epoch was not only implemented by executing varying numbers of SMs but also had an effect on how these were realised. In addition, there was a dissociation of the affected parameters. While LSMs were mainly tuned by varying durations RSMs were modulated both in strength and duration.

[29] repeated-measures two-way ANOVA, main effect context, $F(1, 49) = 4.78, p = 0.03$

[30] 71.93 ± 6.07 mm/s (prepare-left) vs. 69.29 ± 6.41 mm/s (prepare-right); repeated-measures two-way ANOVA, interaction context × upcoming instruction $F(1, 49) = 8.60, p = 5.11 \cdot 10^{-3}$

[31] repeated-measures two-way ANOVA, interaction context × upcoming instruction $F(1, 49) = 39.85, p = 7.69 \cdot 10^{-8}$

[32] 50.13 ± 12.37 mm/s (prepare-right) vs. 38.23 ± 7.67 mm/s (prepare-left)

[33] repeated-measures two-way ANOVA, main factor upcoming instruction $F(1, 49) = 9.24, p = 3.79 \cdot 10^{-3}$, main factor context $F(1, 49) = 5.12, p = 0.028$

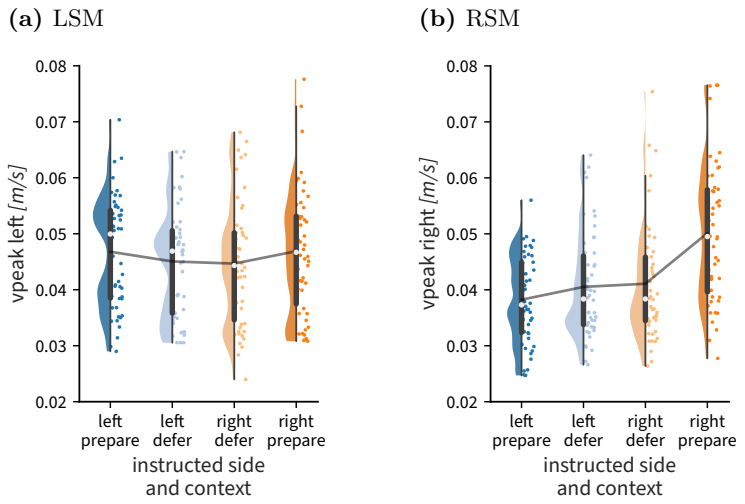

Figure 3.14: Submovement strength in the context epoch (0 ms to 600 ms from context cue onset) for correct trials. **(a)** Submovement strength, quantified as submovements' peak velocity, for left submovements (LSMs). LSMs were significantly modulated by context: repeated-measures two-way ANOVA, main factor context, $F(1, 49) = 4.78, p = 0.03$. **(b)** same like (a) for right submovements (RSMs). RSMs were significantly modulated by combination of context and upcoming instruction: repeated-measures two-way ANOVA, interaction context × upcoming instruction $F(1, 49) = 39.85, p = 7.69 \cdot 10^{-8}$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

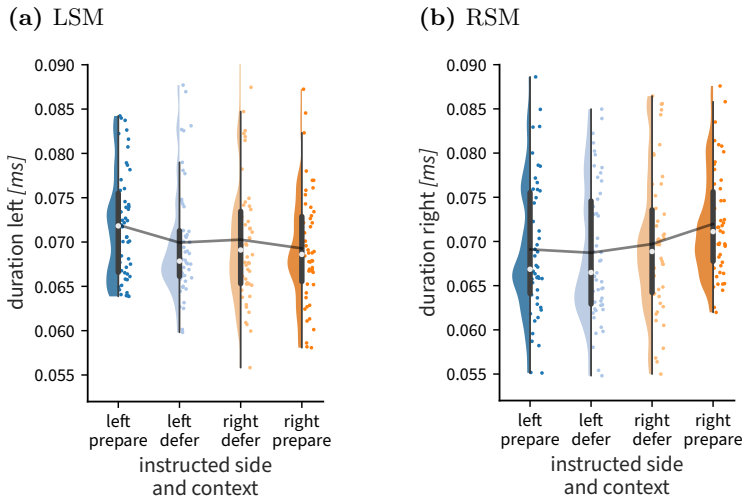

Figure 3.15: Submovement duration in the context epoch (0 ms to 600 ms from context cue onset) for correct trials. **(a)** Submovement duration for left submovements (LSMs). LSMs were significantly modulated by combination of context and upcoming instruction: repeated-measures two-way ANOVA, interaction context × upcoming instruction $F(1, 49) = 8.60, p = 5.11 \cdot 10^{-3}$. **(b)** same like (a) for right submovements (RSMs). RSMs were significantly modulated by only by context and upcoming instruction: repeated-measures two-way ANOVA, main factor upcoming instruction $F(1, 49) = 9.24, p = 3.79 \cdot 10^{-3}$, main factor context $F(1, 49) = 5.12, p = 0.028$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

3.1.4.7 Instruction epoch submovement parameters modulated by instruction but not prior context

Because in the instruction epoch animals were required to rotate the response ball to the instructed side, they greatly increased the frequency of the corresponding submovements (Fig. 3.13). This alone could have been sufficient to reach the goal. However, a more efficient solution would be to also increase strength and duration of SMs.

Indeed, SMs comprising the response were not only modulated in number but also in strength and duration (Figs. 3.16 and 3.17). Accordingly, LSMs were stronger[34] and longer-lasting[35] in *go-left* trials than in *go-right* (Figs. 3.16a and 3.17a). In *go-right* trials, both parameters were downregulated to levels as low as during the context epoch.

The same was true for RSMs: they were stronger[36] and longer-lasting[37] in *go-right* than in *go-left* trials (Figs. 3.16b and 3.17b). Strength and duration in *go-left* trials were also close to context epoch levels.

Consistent with SM frequency prior context had no effect on either strength nor duration of LSMs or RSMs.

In sum, SMs that were consistent with the movement instruction were strongly upregulated not only in terms of frequency but also in their properties.

3.1.5 Summary behaviour

Mice learned to use an abstract goal-directed behaviour, namely, directed rotations of a response ball, to obtain rewards. In addition, they successfully acquired the association of arbitrary sounds as movement instructions or preparatory contexts. Effects of preparation were present at different behavioural levels. Informative contexts improved task performance and response times (Figs. 3.3 and 3.4).

The use of continuously measured response ball rotation greatly extended the psychometric repertoire. Using simple trial subsetting I showed that consistent movement tendencies during the context epoch improved task performance (Fig. 3.7a) and that activity in that epoch propagated to the instruction epoch (Fig. 3.6b). An even more detailed view of the response ball data revealed that the measured rotations were composed of discrete submovements (SMs) (Fig. 3.8). During both trial epochs these SMs were modulated in numbers (Figs. 3.12 and 3.13), strength (Figs. 3.14 and 3.16) and duration (Figs. 3.15 and 3.17), thus providing multiple behavioural substrates affected by contexts and instructions.

Of note, there was a consistent behavioural left-preference, indicated by both classic psychometrics as well as detailed dissection of the response ball data. The findings

[34] $122.68 \pm 30.21\,\mathrm{mm/s}$ (*go-left*) vs. $47.28 \pm 10.12\,\mathrm{mm/s}$ (*go-right*); repeated-measures two-way ANOVAs, main factor instruction $F(1,49) = 352.74, p = 4.95 \cdot 10^{-24}$

[35] $83.36 \pm 8.64\,\mathrm{ms}$ (*go-left*) vs. $61.86 \pm 6.62\,\mathrm{ms}$ (*go-right*); repeated-measures two-way ANOVAs, main factor instruction $F(1,49) = 256.64, p = 4.10 \cdot 10^{-21}$

[36] $43.34 \pm 7.90\,\mathrm{mm/s}$ (*go-left*) vs. $96.12 \pm 12.71\,\mathrm{mm/s}$ (*go-right*), repeated-measures two-way ANOVAs, main factor instruction $F(1,49) = 863.76, p = 8.85 \cdot 10^{-33}$

[37] $59.81 \pm 8.99\,\mathrm{ms}$ (*go-left*) vs. $79.88 \pm 6.98\,\mathrm{ms}$ (*go-right*), repeated-measures two-way ANOVAs, main factor instruction $F(1,49) = 157.99, p = 6.02 \cdot 10^{-17}$

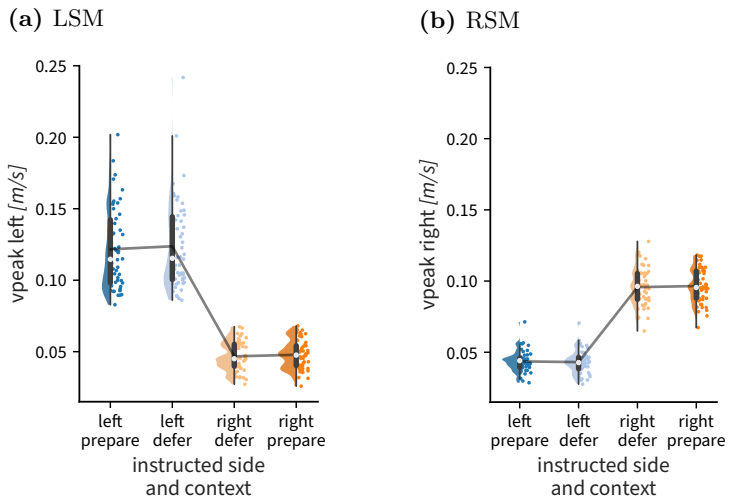

Figure 3.16: Submovement strength in the instruction epoch (instruction onset to threshold crossing) for correct trials. **(a)** Submovement strength, quantified as submovements' peak velocity, for left submovements (LSMs). LSMs were significantly modulated by instruction: repeated-measures two-way ANOVAs, main factor instruction $F(1, 49) = 352.74, p = 4.95 \cdot 10^{-24}$. **(b)** same like (a) for right submovements (RSMs). RSMs were significantly modulated by instruction: repeated-measures two-way ANOVAs, main factor instruction $F(1, 49) = 863.76, p = 8.85 \cdot 10^{-33}$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

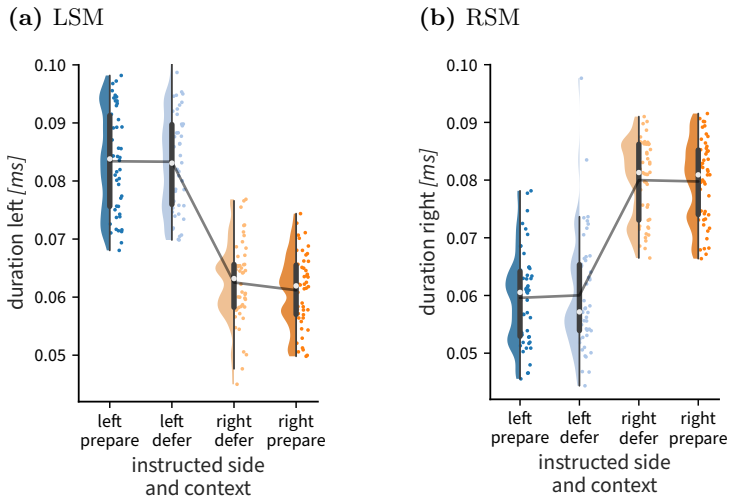

Figure 3.17: Submovement duration in the instruction epoch (instruction onset to threshold crossing) for correct trials. **(a)** Submovement duration for left submovements (LSMs). LSMs were significantly modulated by instruction: repeated-measures two-way ANOVAs, main factor instruction $F(1, 49) = 256.64, p = 4.10 \cdot 10^{-21}$ **(b)** same like (a) for right submovements (RSMs). RSMs were significantly modulated by instruction: repeated-measures two-way ANOVAs, main factor instruction $F(1, 49) = 157.99, p = 6.02 \cdot 10^{-17}$. Every dot is the mean in one session; see methods for description of raincloud plots. Black lines across the conditions connect arithmetic means.

from SMs decomposition indicate that in many cases the behaviour in non-informative context epochs was closer to that in informative *prepare-left* ones. Conversely, behaviour in *prepare-right* context epochs was modulated more. Consistent with this, left SMs were more invariant to contexts, while all parameters of right SMs displayed context-dependent variations.

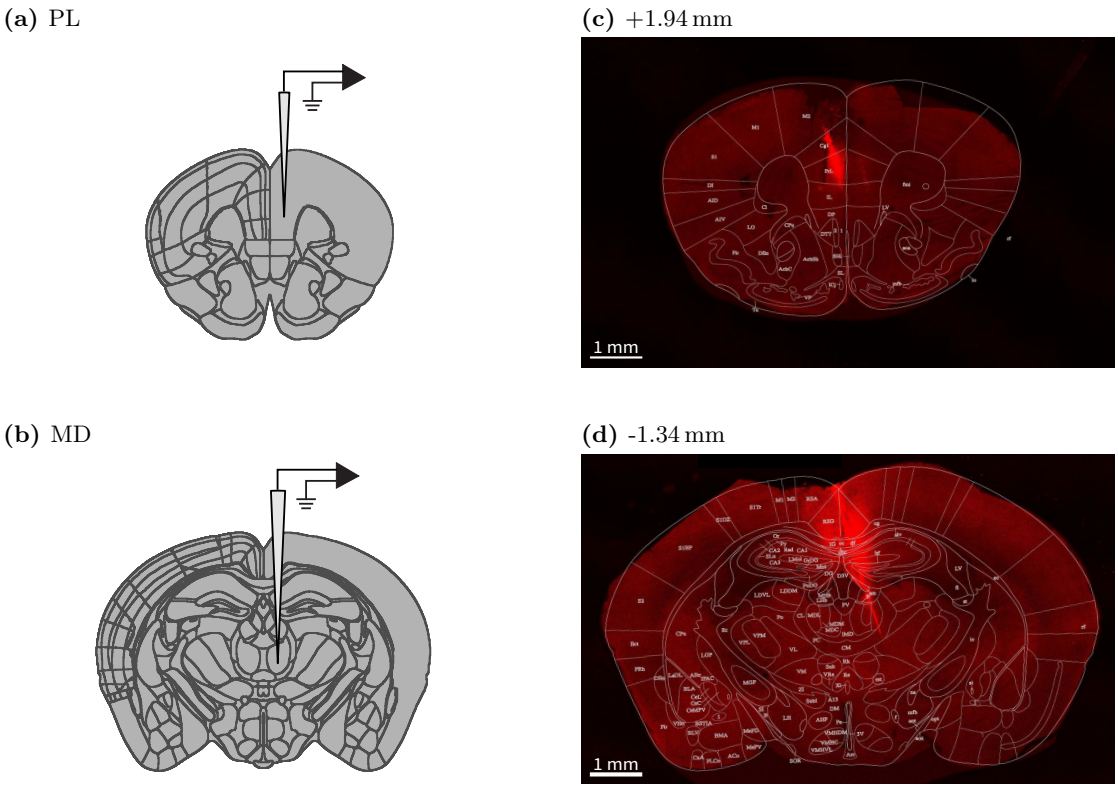

Figure 3.18: Recording locations
(a, b) Mouse brain coronal schematics of recording sites for (a) PL and (b) MD. Modified from Allen Brain Atlas (Q. Wang et al., 2020). **(c, d)** Histological example brain slices with DyeI labelling of silicon probe tracts for (c) PL (at bregma +1.94 mm) and (d) MD (at bregma -1.34 mm). Slicing and microscopy images courtesy Ajit Ranganath.

3.2 Neurophysiology

3.2.1 Firing rates in single neuronal units and regional populations

A total of 1020 prelimbic cortex (PL) and 711 MD single units[38] were recorded acutely in 38 sessions. Of these, 629 (PL) and 420 (MD) units matched the inclusion criteria in 37 sessions from 4 animals (Table 3.3; correct probe locations could not be verified for one of the five trained animals). Simultaneous recordings from PL and MD were made in 13 sessions (156 PL units, 312 MD units, 2 animals), with the remainder of the sessions being single-region recordings (22 sessions / 473 units and 2 sessions / 108 units in PL and MD, respectively).

[38] In this section the term *single unit* is used interchangeably with the term *unit* and *neuron*. In all cases this refers to the same putative biological entity that is thought to be the source of well isolated action potentials.

Table 3.3: Number of single neuronal units, sessions and animals. Note that the grand total of sessions does not sum across rows because of the inclusion of sessions with simultaneous PL and MD recordings. Similarly, numbers of animals do not add across rows because single-region and bi-regional recordings were performed in the same animals.

region	recording	units	sessions	animals
PL	single region	473	22	3
	simultaneous	156	13	2
	total	629	35	3
MD	single region	108	2	1
	simultaneous	312	13	2
	total	420	15	2
	grand total	1049	37	4

To get a first impression of the neuronal responses to contexts and instructions, in the following I will start by showing example single neuronal units from PL and MD. Afterwards, to identify trends in the populations and possible inter-regional differences between PL and MD, I will show regional-averaged activity profiles, as well as firing rate peak latencies.

3.2.1.1 Single unit activity in PL and MD varied with context cue and instruction and showed heterogeneous response profiles

Figures 3.19 and 3.20 show the firing rates in the different contexts and instructions for two example single units each in PL and MD. Neurons both in PL and MD responded to context and/or instruction. Typically, neurons quickly increased their firing rates after the onset of a stimulus. After an early peak, many neurons showed decreasing sustained activity throughout the context delay or during the instruction epoch. Peak values as well as latencies within a neuron often varied with context and/or instruction.

Example single units

The first sample PL neuron showed maximum activity in trials with the *prepare-right* context, peaking after the offset of the context stimulus presentation and showing lower but sustained firing during the context delay (orange colour, Figs. 3.19a and 3.19b). In *defer* context trials (grey colour) this neuron's activity peaked during stimulus presentation and slowly decayed its activity to baseline levels afterwards. The *prepare-left* context (blue colour) elicited the lowest and most delayed peak response in this neuron. As a consequence of this the context cue information contained in the spiking activity is high throughout the context epoch (Fig. 3.19c).

On the other hand, the *go-left* instruction, which is contingent on the latter context, led to a very fast response within the first 100 ms which, notably, varied with previous

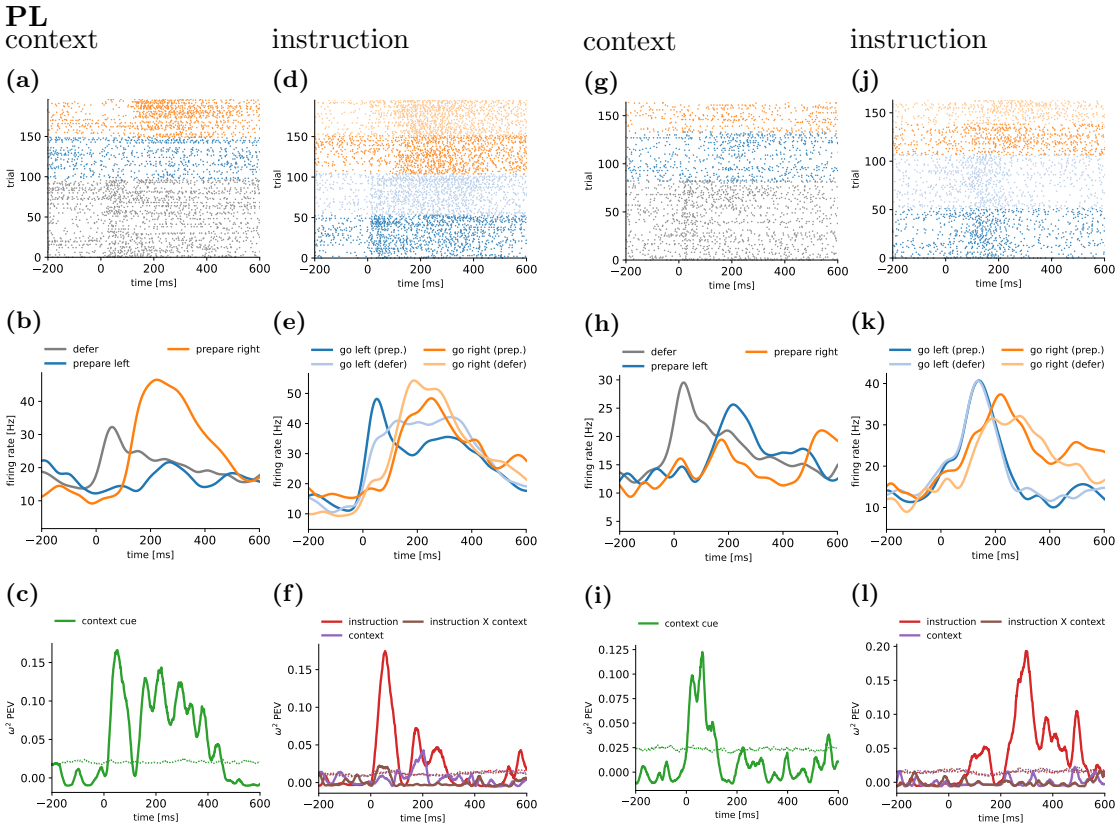

Figure 3.19: Firing activity and contained information of examples single units from PL. **(a)** Raster plot of action potentials aligned to context cue onset at 0 ms. Trials sorted and coloured by condition. **(b)** Smoothed firing rates of the single unit from (a). **(c)** Information about the context cue contained in firing variations of this single unit. **(d–f)** same single unit as in (a–c), but aligned to instruction onset and sorted by instruction side and prior context. (f) shows information about the instruction side (*ins*), context (*ctx*), and their interaction (*i*c*). **(g–l)** same as in (a–f), but for a different example single unit in PL.

context (dark and light blue colours, Figs. 3.19d and 3.19e): While the *prepare* context led to a transient early activity increase, followed by decreased but sustained activity, in trials following a *defer* context the *go-left* response only showed a plateau with higher sustained activity. Conversely, the onsets of the response to the *go-right* instruction were delayed and more similar in *prepare* and *defer* trials.

Another neuron recorded in PL showed a response profile that shared some properties with the previously described example neuron but differed in others. Response profiles to the different contexts were qualitatively similar to the previous neuron (Figs. 3.19g and 3.19h): The peak latency for the *defer* trials was earlier than for *prepare* contexts and the activity slowly decayed after the respective peaks. Quantitatively, however, this neuron did not maximally respond to the *prepare-right* but the *defer* context. In fact, the response to the *prepare-right* context was the lowest in this neuron.

Similar to the first example neuron, responses to the *go-left* instructions peaked earlier than for the *go-right* instructions (Fig. 3.19k). In contrast, however, the instruction responses did not vary with prior context. Furthermore, the response to the *go-left* instruction was more transient and did not show a sustained plateau.

Neurons in the MD often show delay activity that is similar to PFC (Fuster & Alexander, 1971; Schmitt et al., 2017). It is therefore instructive to also investigate the firing rate profiles of MD single units (Fig. 3.20).

The first example MD single unit quickly increased activity after onset of the stimuli for *defer* and *prepare-right* contexts, peaking during their presentation (Figs. 3.20a and 3.20b). While activity for both contexts peaked around the same time, the magnitude was three times larger in *defer* trials. After peak activity in those trials firing decreased rapidly to levels in the *prepare-right* context, and then, similarly slowly decayed until about 300 ms to 400 ms. Trials in the *prepare-left* context, on the other hand, slowly increased this neuron's activity until peaking between 100 ms to 200 ms, i.e. after the offset of the cue. Thereafter it decayed with the same speed as the initial increase. Consistent with this response profile during the context epoch, this unit had very high context information during the first 100 ms (Fig. 3.20c).

In the instruction epoch this neuron showed context-modulated responses to the instructions (Figs. 3.20d and 3.20e). Trials with the *go-left* instruction showed low latency peak activity within the first 50 ms after instruction onset. Importantly, trials with a prior *prepare* context had higher *go-left* responses than trials with a prior *defer* context. *go-right* trials, on the other hand, had a slower activity increase and decay, but still showed differential activity dependent on the prior context. Thus, the the neuron had high information about the instruction, as well as some context information in the first 100 ms (Fig. 3.20f).

The last example single unit in MD again showed varying peak latencies for different contexts and instructions. During the context epoch the strongest response was in *defer*, followed by *prepare-right* trials (Figs. 3.20g and 3.20h). Responses for the *prepare-left*

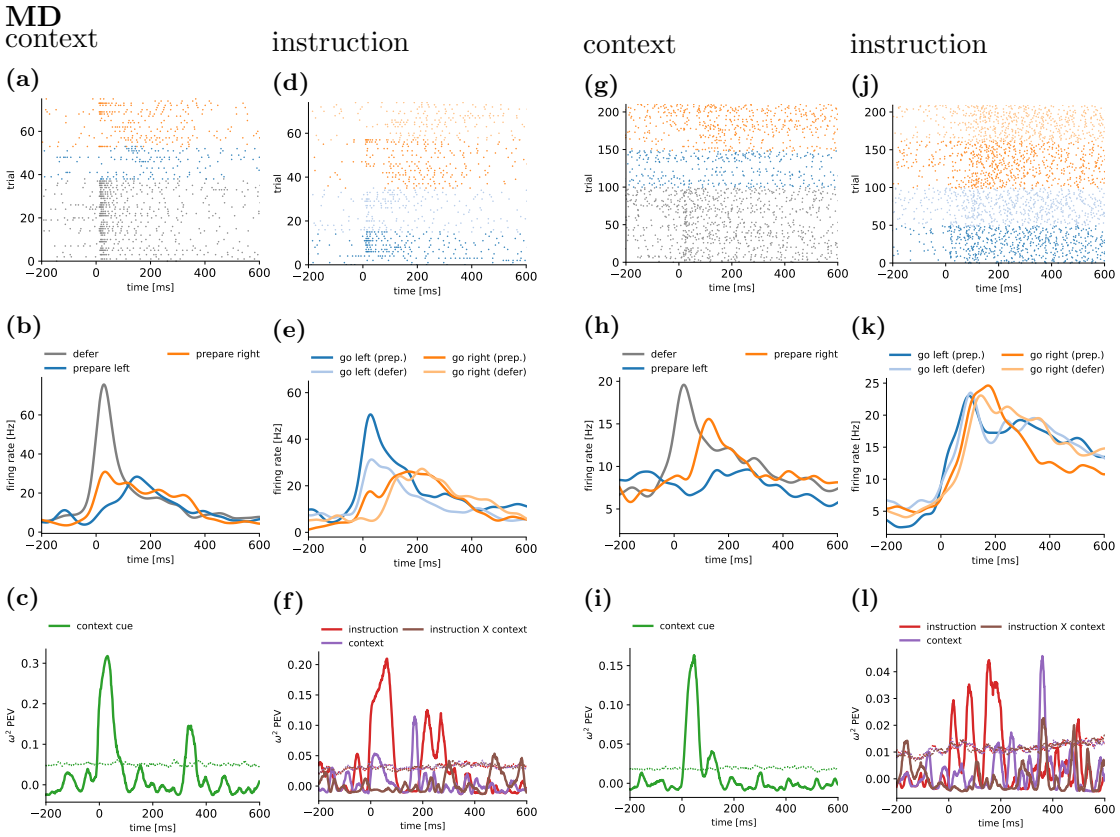

Figure 3.20: Firing activity and contained information of examples single units from MD. **(a)** Raster plot of action potentials aligned to context cue onset at 0 ms. Trials sorted and coloured by condition. **(b)** Smoothed firing rates of the single unit from (a). **(c)** Information about the context cue contained in firing variations of this single unit. **(d–f)** same single unit as in (a–c), but aligned to instruction onset and sorted by instruction side and prior context. (f) shows information about the instruction side (*ins*), context (*ctx*), and their interaction (*i*c*). **(g–l)** same as in (a–f), but for a different example single unit in MD.

trials were virtually at baseline. The peak latency for the *defer* context again was very short, during the presentation of the stimulus, while the peak for *prepare-right* context was shortly after stimulus offset. Accordingly, the context information was high throughout the first 100 ms (Fig. 3.20i).

Responses to the different instruction cues were similar to each other in this unit (Figs. 3.20j and 3.20k). Activity rapidly increased after instruction onset, peaking between 100 ms to 200 ms. However, a small latency difference was present: Response onset as well as peak activity for *go-right* trials was slightly offset, thus resulting in reasonably good instruction information (Fig. 3.20l).

Together, these example neurons show that single-unit activity realised heterogeneous response profiles for different contexts and instructions. Typically, neurons showed low-latency activity peaks that were followed by a slower decay. Peak times and magnitudes for different contexts and instructions varied within and across the example cells. Spiking activity in single-units showed information about behaviourally relevant parameters.

3.2.1.2 Firing rate averages and peak latencies across single units show distinct response profiles for context and instruction in PL and MD

In order to detect general trends in the two regions' activity profiles, it is informative to investigate population activity. To this end, two complementary views incorporating raw firing rates are shown. Firstly, the standardised firing rates across all included single units of a region were plotted (Figs. 3.21 and 3.23). Since those profiles are suggestive of latency differences, secondly, single unit peak latencies were considered (Figs. 3.22 and 3.24).

Context epoch In the context epoch the PL population had its peak activity for the *defer* context during the presentation of the context cue around 50 ms (Fig. 3.21a, grey). In trials with either of the two *prepare* contexts, on the other hand, average activity increased later and had its peak activity in the delay after context cue presentation, shortly before 200 ms (Figs. 3.21a and 3.22a, blue and grey). Notably, peak magnitudes for *prepare-left* context (blue line) were much lower than those in *defer* and *prepare-right* trials. Peak magnitudes orderly descended from *defer* to *prepare-right* to *prepare-left*.

After the initial peak, average activity in all conditions was still elevated for another 400 ms to 500 ms, indicating sustained delay epoch activity of numerous PL units. Interestingly, in this late context epoch average firing rates for *prepare-left* and *defer* trials were similar to each other (Fig. 3.21a, blue and grey, starting at around 300 ms). Activity for *prepare-right* trials (orange), on the other hand, was more distinct.

In the MD population the averaged activity during the context epoch showed similar trends as observed in PL (Fig. 3.21b). Consistent with PL, peak magnitudes in response to the *defer* and *prepare-right* contexts were clearly higher than in response to the *prepare-left* context. Likewise, the activity peaks for the *prepare* contexts occurred later

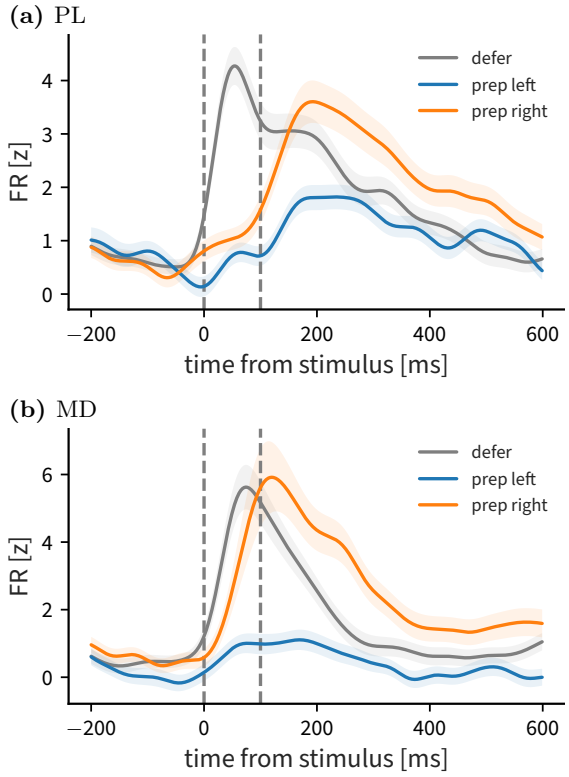

Figure 3.21: Average firing rates across single neuronal units (context epoch). Plotted are the standardised (z-scored) firing rates averaged across single units in PL ($n = 629$, **(a)**) and in MD ($n = 420$, **(b)**) from −200 ms before to 600 ms after context cue onset. Different colours represent the different context cues for *defer*, *prepare-left* and *prepare-right*. Shading: SEM across single units. Horizontal dashed lines: Onset and offset times of context cue presentation.

than for the *defer* context (Fig. 3.22b). Furthermore, the magnitude of the sustained activity in the later context epoch was highest for the *prepare-right* context.

Conversely, notable differences to PL activity were found in MD. The range of peak magnitudes was greater in MD than in PL, ranging from approximately 1 z to 6 z vs. 2 z to 4 z. For the *defer* and *prepare-right* contexts they were higher in MD than in PL, while the magnitude for the *prepare-left* was only about half of the one in PL (Fig. 3.21b). In contrast to the clear ordering of magnitudes observed in PL, the peak magnitudes for *defer* and *prepare-right* were virtually identical. Consistent with similar magnitudes, and in contrast to more heterogeneous profiles in PL, the temporal activity profiles in *defer* and *prepare-right* trials were very similar in MD, almost like time-shifted copies.

The comparison of peak latencies showed clear cross-regional differences. In the population-averaged activity, peaks for the *prepare* contexts in the MD preceded those in PL by about 100 ms to 150 ms (Figs. 3.21a and 3.21b). By comparison, peak latencies for the *defer* context were almost the same in PL and MD. Notably, however, the small latency difference that was present indicated that MD activity was *lagging* PL. This was further corroborated by analysis of single units' peak latencies[39] (Figs. 3.22a to 3.22c).

[39] population peak difference PL - MD: *defer*: −10 ms to −5 ms, $p = 0.04$; *prepare-left*: 50 ms to 170 ms, $p = 0.004$; *prepare-right*: 50 ms to 240 ms, $p = 0.002$; bootstrap 95 % confidence interval (CI), permutation test

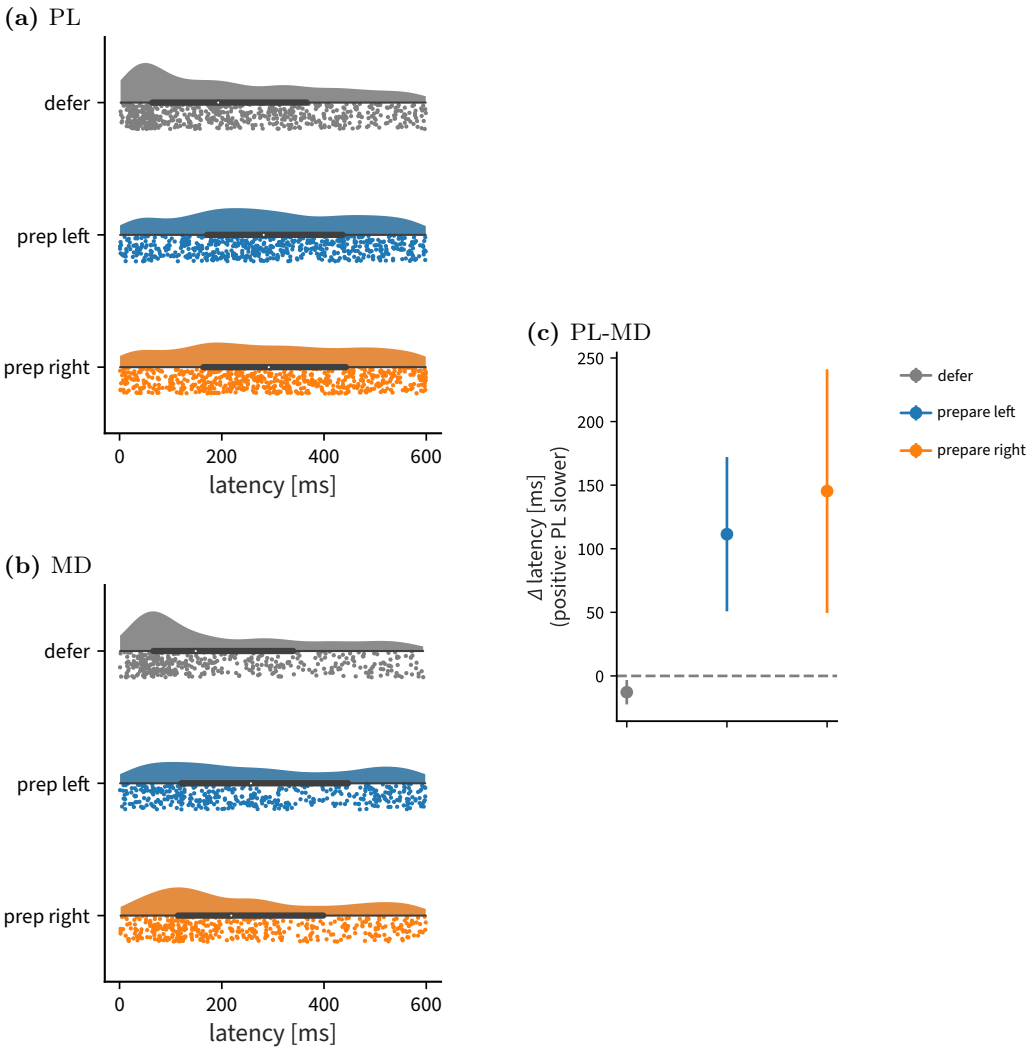

Figure 3.22: Single units' peak latencies during the *context* epoch split by conditions **(a)** ζ-test derived latencies for firing rates in PL. Each coloured dot represents the peak latency of one single unit. See methods for explanation of the remainder of raincloud plots. Distributions were statistically not different from uni-modal: Hartigan's dip test for unimodality: *defer*: 0.011, $p = 0.91$ *prepare-right*: 0.011, $p = 0.90$ *prepare-left*: 0.013, $p = 0.78$. **(b)** same as (a) but for single units from md. Distributions of *defer* and *prepare-right* condition were statistically not different from uni-modal, but *prepare-left* was: Hartigan's dip test for unimodality: *defer*: 0.012, $p = 0.97$, *prepare-right*: 0.020, $p = 0.28$, *prepare-left*: 0.028, $p = 0.021$. **(c)** Differences of the distributions' kdes' modes from (a, b). Dots represent differences (PL - MD): whiskers represent bootstrapped 95 % CI. Values of CI and p-value permutation test): *defer*: -10 ms to -5 ms, $p = 0.04$; *prepare-left*: 50 ms to 170 ms, $p = 0.004$; *prepare-right*: 50 ms to 240 ms, $p = 0.002$.

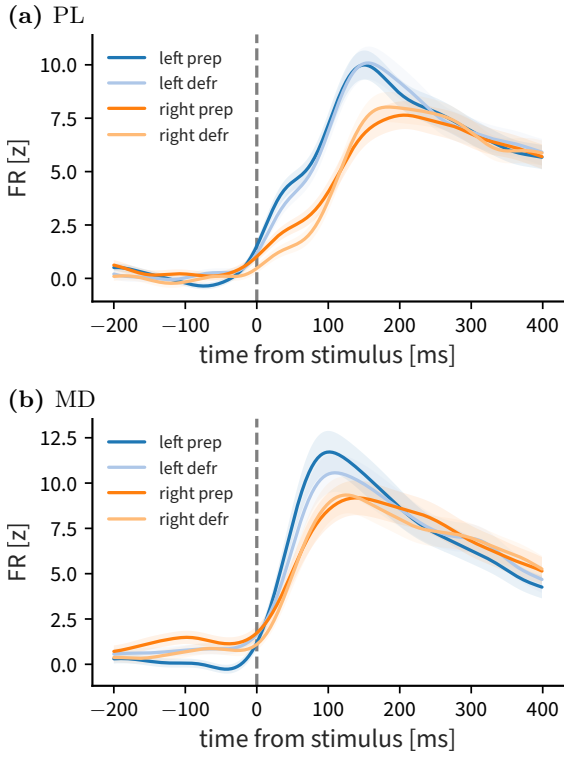

Figure 3.23: Average firing rates across single neuronal units (instruction epoch). Plotted are the standardised (z-scored) firing rates averaged across single units in PL ($n = 629$, **(a)**) and in MD ($n = 420$, **(b)**) from −200 ms before to 400 ms after movement instruction onset. Colours represent the movement instructions (blue *go-left*, orange *go-right*) and the prior context (saturated *prepare*, desaturated *defer*). Shading: SEM across single units. Horizontal dashed lines Onset time of movement instruction presentation.

Instruction epoch In terms of the behavioural task, the instruction stimulus resolved ambiguity in the trials with a prior *defer* context and served as a go-cue after *prepare* contexts (Fig. 3.2). In general, neuronal activity during the instruction epoch was considerably higher than during the context epoch. Furthermore, both regions' activities varied with the instruction side, with higher and earlier responses to the *go-left* instruction (Figs. 3.23a and 3.23b).

The temporal dynamics of activity increase in PL showed a two-component process (Fig. 3.23a). In the first 50 ms activity rapidly increased until a short saddle of about 20 ms characterised by smaller increases was reached. After this saddle the activity increased again more rapidly, until the activity arrived at its maximum around 150 ms to 200 ms. Following this, activity slowly decreased. Interestingly, the majority of the differences between *prepare* and *defer* contexts occurred during this first component. Here, activity in trials within a *defer* context was lower[40]. Conversely, at subsequent times context-dependent modulation could not be reliably detected[41].

This two-component activity profile could be obtained if individual neurons displayed two-component profiles or if the latencies of individual neurons' peak activities concentrated at two distinct times. Analysis of individual neurons' peak latencies provided evidence for the latter hypothesis: In all but one condition (*right-prepare*) the distributions

[40] *prepare* vs. *defer*, 0 ms to 100 ms: $p_{go\text{-}left} = 4.636 \cdot 10^{-5}$, $p_{go\text{-}right} = 0.010$, Wilcoxon signed-rank test
[41] *prepare* vs. *defer*, 100 ms to 200 ms: $p_{go\text{-}left} = 0.090$, $p_{go\text{-}right} = 0.146$, Wilcoxon signed-rank test

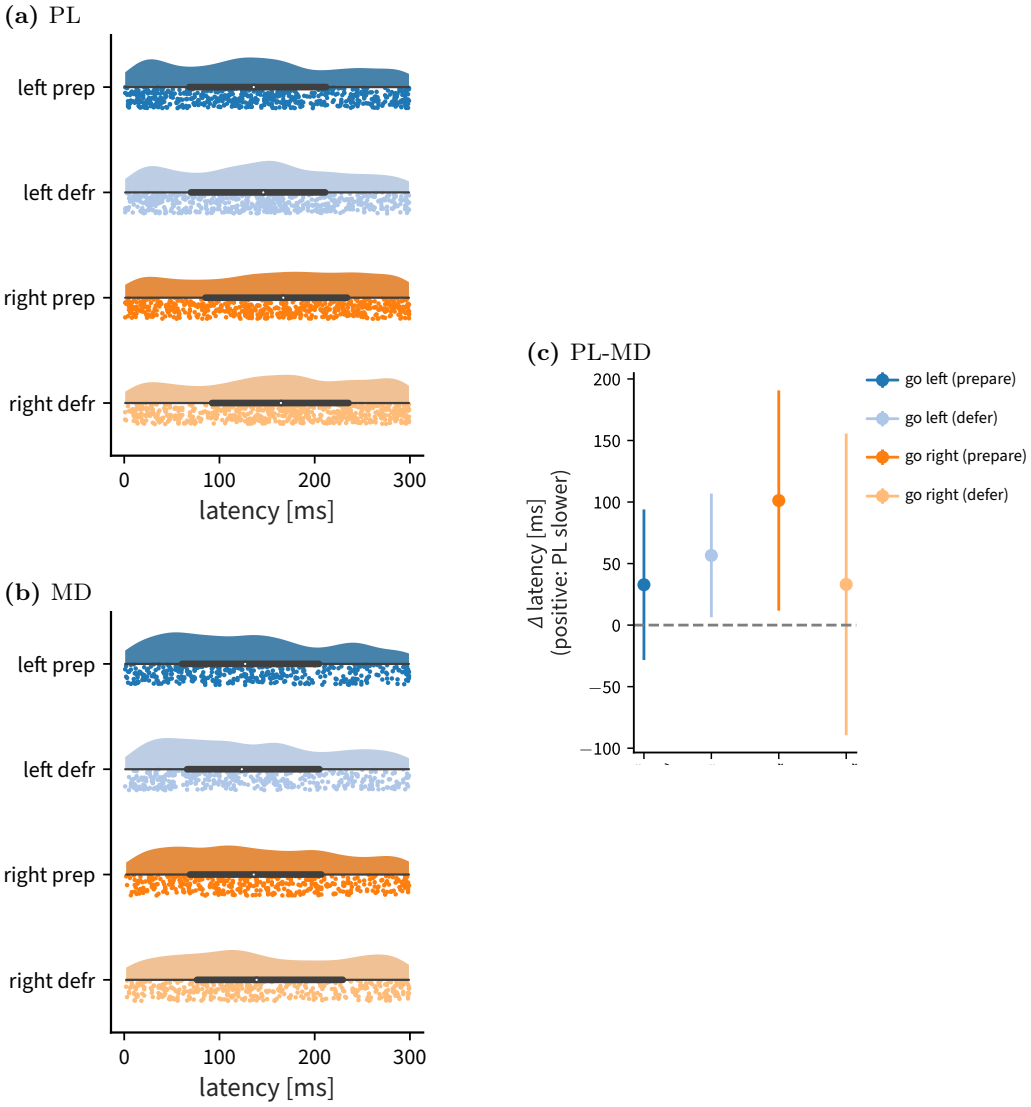

Figure 3.24: Single units' peak latencies during the *instruction* epoch split by conditions

(a) ζ-test derived latencies for firing rates in PL. Each coloured dot represents the peak latency of one single unit. See methods for explanation of the remainder of raincloud plots. Distributions were significantly different from unimodal for *go-left* (*prepare*), *go-left* (*defer*) and *go-right* (*defer*): Hartigan's dip test for unimodality: *go-left* (*prepare*): $p = 4.28 \cdot 10^{-3}$, *go-left* (*defer*): $p = 2.67 \cdot 10^{-2}$, *go-right* (*prepare*): $p = 0.305$, *go-right* (*defer*): $p = 4.60 \cdot 10^{-2}$. **(b)** same as (a) but for single units from MD. Distributions were not significantly different from unimodal except for *go-right* (*prepare*): Hartigan's dip test for unimodality: *go-left* (*prepare*): $p = 0.561$, *go-left* (*defer*): $p = 0.474$, *go-right* (*prepare*): $p = 4.47 \cdot 10^{-2}$, *go-right* (*defer*): $p = 0.861$. **(c)** Differences of the distributions' kdes' modes from (a, b). Dots represent differences (PL - MD): whiskers represent bootstrapped 95 % CI. CI (p-value permutation test): *left-prepare*: $p = 0.30$; *left-defer*: $p = 0.28$; *right-prepare*: $p = 0.12$; *right-defer*: $p = 0.17$.

of peak latencies from PL single units were significantly non-unimodal[42] (Fig. 3.24a). Furthermore, no single units with a saddle-like activity profile were found by visual inspection.

In contrast to the two components present in the PL population, the temporal dynamics across MD single units were simpler (Fig. 3.23b). Following instruction onset, firing rates increased rapidly but steadily and peaked between 100 ms to 150 ms. Consistent with the simpler one-component response profiles the distributions of individual single units' peak latencies were statistically not different from unimodal except for one condition[43] (Fig. 3.24b). The average peaks appeared to be temporally embedded between the first and second component of PL activity increase. Due to the differences in distribution shapes (unimodal vs. multimodal) no statistical evidence for which region was leading could be found (Fig. 3.24c). Finally, and also in contrast to PL, no robust context-dependent modulation of the instruction response was found in MD on a population level[44].

In summary, the investigation of firing rates and peak latencies found distinct response profiles in PL and MD that shared some condition-specific relationships. Peak responses for the uninformative *defer* context always preceded the peak responses to the predictive *prepare* contexts. MD's peak responses to *prepare* contexts clearly preceded PL's, whereas, in contrast, PL was minimally leading MD for the *defer* context. Magnitudes for *defer* and *prepare-right* contexts were considerably higher than for the *prepare-left* context. This effect was especially pronounced in MD.

Responses to the instruction were greater than to the contexts and responses to *go-left* were higher than to *go-right*. In PL the response to instruction was a two-component process that might have been achieved by different neuronal subgroups. Prior contexts had effects on the first component. MD's responses were simpler and not modulated by prior context.

3.2.2 Neuronal responsiveness and preference

The population-averaged magnitudes of firing rates in the previous section can be suggestive of differential activity modulation dependent on the external stimuli. However, different magnitudes could have arisen differently. For example, one could envision subpopulations of single neuronal units that were exclusively modulated by one context cue and/or movement instruction. Different quantitative population averages in the firing rates would then be achieved by differently sized subpopulations. Another possibility would be that single neuronal units had overall similar firing rate profiles with quantitatively distinct activations consistent with the population averages.

---

[42] Hartigan's dip test for unimodality: *go-left* (*prepare*): $p = 4.28 \cdot 10^{-3}$, *go-left* (*defer*): $p = 2.67 \cdot 10^{-2}$, *go-right* (*prepare*): $p = 0.305$, *go-right* (*defer*): $p = 4.60 \cdot 10^{-2}$

[43] Hartigan's dip test for unimodality: *go-left* (*prepare*): $p = 0.561$, *go-left* (*defer*): $p = 0.474$, *go-right* (*prepare*): $p = 4.47 \cdot 10^{-2}$, *go-right* (*defer*): $p = 0.861$

[44] *prepare* vs. *defer*, 0 ms to 100 ms: $p_{go\text{-}left} = 0.108$, $p_{go\text{-}right} = 0.398$, Wilcoxon signed-rank test; *prepare* vs. *defer*, 100 ms to 200 ms: $p_{go\text{-}left} = 0.135$, $p_{go\text{-}right} = 0.446$, Wilcoxon signed-rank test

Furthermore, neuronal information metrics like analysis of variance (ANOVA)-derived ω^2 can be indicative of consistent firing rate differences across conditions. However, they lack information about the source of the variation. For example, in the case with two trial conditions a neuron's firing could carry high information because it is only activated in one or it is activated to different extents in the two conditions. The "magnitude" of activation or *responsiveness* to individual conditions, however, cannot be extracted from neuronal information measures.

To further dissect single unit activity in different task conditions, in the following I employed measures derived from the parameter-free Zenith of Event-based Time-locked Anomalies (ZETA) test (Montijn et al., 2021). Briefly, ZETA is calculated as follows. Individual neuronal units' spike times are aligned to events of interest, e.g. stimulus presentations. Event-relative spike times pooled across multiple instances of the event of interest, e.g. across trials, are used to generate an empirical cumulative distribution as a function of time relative to the event. Positive or negative deviations of this cumulative distribution function (CDF) from a linear baseline indicate the modulation of a neuron's firing rate by the event. Statistical significance is assessed against a null distribution that is generated by randomly jittering event onset. Thus, the ZETA test does not require the specification of arbitrary analysis bins like the standard approach of comparing firing rates before and after an event using a two-sample t-test (or its nonparametric alternatives) would.

First, the numbers of single units that were statistically responsive to individual conditions were quantified. This was done in order to ascertain various questions about specialised subpopulations:

1. *Epoch specialisation*: Were single units responsive both to context cues and movement instructions, or were distinct subpopulations activated?

2. *Condition specialisation*:
    a) *single vs. mixed responsiveness*: were there specialised subpopulations of neurons that were exclusively responsive to a single context/instruction or did neurons show mixed responsiveness?
    b) *distributions*: if condition-specialised subpopulations existed, were they similarly sized?

3. *Context-dependent instruction-responsiveness*
    a) *prepare vs. defer*: were the same single units instruction-responsive after both context types (context-stable)?
    b) *quantitative modulation*: was instruction-responsiveness modulated by prior context?

Secondly, the neuronal preferences were analysed in order to see if there were neuronal subpopulations defined by behavioural contingency. As a concrete example, this asked whether a single unit that preferred the *prepare-right* context also preferred the behaviourally contingent *go-right* instruction; and similarly for the other contingencies.

3.2.2.1 Partially overlapping subpopulations with epoch-specific and -invariant responsiveness

A large proportion of PL and MD neurons was responsive to one or more contexts (0 ms to 600 ms from context cue onset) or instructions (0 ms to 300 ms from movement instruction onset) (PL: 74.1 %, MD: 81.2 %, bright green and pink, as well as the dark green overlap areas in Fig. 3.25 and Table 3.4). The proportion of single units satisfying this most inclusive criterion was significantly higher in MD than in PL[45]. Not all of those cells were responsive in both task epochs: Only half of responsive single units were so both to context *and* instruction (PL: 35.6 % of total, MD: 38.1 % of total, dark green area in Fig. 3.25, Table 3.4). The single units that were responsive in only one task epoch were not distributed equally. In both regions significantly more units were responsive to the movement instructions than to the context cues[46] (Fig. 3.25, bright green and pink areas, respectively).

Together this shows that responsiveness to context and instruction was realised by partially overlapping, but not distinct, groups of single neuronal units.

3.2.2.2 Most single units were exclusively responsive to one context; responsiveness was inhomogeneously distributed

In both regions about half of all cells were responsive to one or more contexts (Figs. 3.26a and 3.26b and Table 3.5, PL: 48.8 %, MD: 52.1 %). Interestingly, most of these cells were responsive to only a single context. As a consequence, only a quarter of context-responsive cells showed mixed responsiveness (PL: 35.8 % vs. 13.0 %; MD: 36.0 % vs. 16.2 %; single-responsive vs. multi-responsive).

Responsiveness to different contexts was inhomogeneously distributed. Similar to the graded magnitudes in the firing rates (Fig. 3.21) the number of responsive cells showed a clear descending order with most units being responsive to *defer*, less to *prepare-*

[45] $\chi^2(1, N = 1049) = 6.78, p = 0.009$

[46] PL: 48.8 % context vs. 60.9 % instruction: $\chi^2(1) = 18.11, p = 2.087 \cdot 10^{-5}$
MD: 52.1 % context vs. 67.1 % instruction: $\chi^2(1) = 19.01, p = 1.304 \cdot 10^{-5}$

Table 3.4: Responsive cells. Percentage of cells that were responsive to context (0 ms to 600 ms from context cue onset) and/or instruction (0 ms to 300 ms from movement instruction onset). A cell was responsive if $p_{zeta} < 0.05$ for Zenith of Event-based Time-locked Anomalies (ZETA) test. Total number of cells: $n_{PL} = 629$, $n_{MD} = 420$

	% total	
responsiveness to	PL	MD
context *or* instruction	74.1	81.2
context *and* instruction	35.6	38.1
context (one or more)	48.8	52.1
instruction (*go-left* or *go-right*)	60.9	67.1

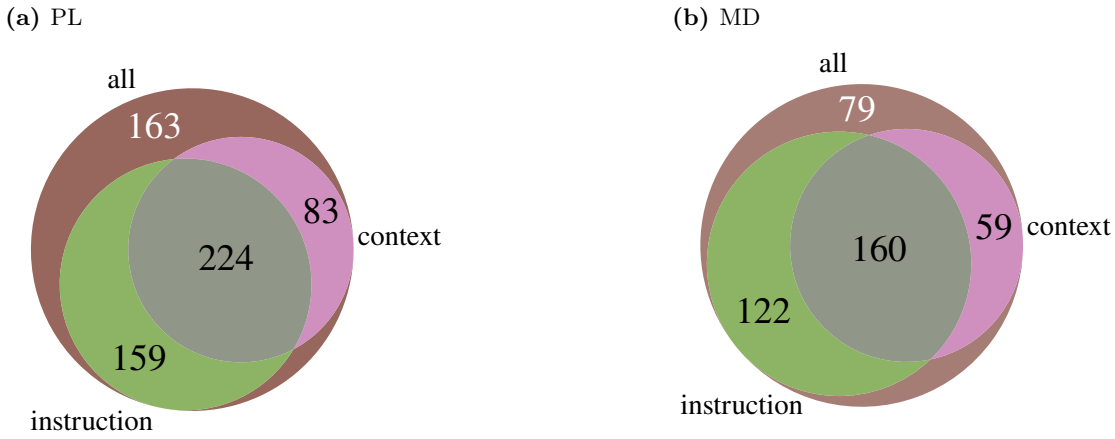

Figure 3.25: Single units responsive to context (0 ms to 600 ms from context cue onset) or instruction. (a) Area-proportional Euler diagrams showing the numbers of responsive cells within the total set of cells. brown: non-responsive; pink: exclusively context-responsive; green: exclusively instruction-responsive; grey overlap: context- and instruction-responsive. **(b)** Same as (a) but for MD. Responsiveness is defined as a $p_\zeta < 0.05$. Proportion of context- or instruction-responsive cells within set of total number of cells significantly higher in MD than in PL: $\chi^2(1, N = 1049) = 6.78, p = 0.009$. Proportion of context- *and* instruction-responsive cells (intersect) within set of all responsive cells not significantly different across regions: $\chi^2(1, N = 807) = 0.60, p = 0.807$.

Table 3.5: Context-responsive cells. Percentage of cells that were responsive to individual contexts (0 ms to 600 ms from context cue onset). A cell was responsive if $p_{zeta} < 0.05$ for ZETA test. Total number of cells: $n_{PL} = 629$, $n_{MD} = 420$

	% total	
responsiveness to	PL	MD
context (one or more)	48.8	52.1
one context only	35.8	36.0
multiple contexts	13.0	16.2
defer (one or more)	36.7	38.3
prepare-left (one or more)	9.5	6.9
prepare-right (one or more)	16.9	24.8
defer and *prepare-left*	4.3	3.8
defer and *prepare-right*	8.9	13.1
prepare-left and *prepare-right*	2.4	2.6

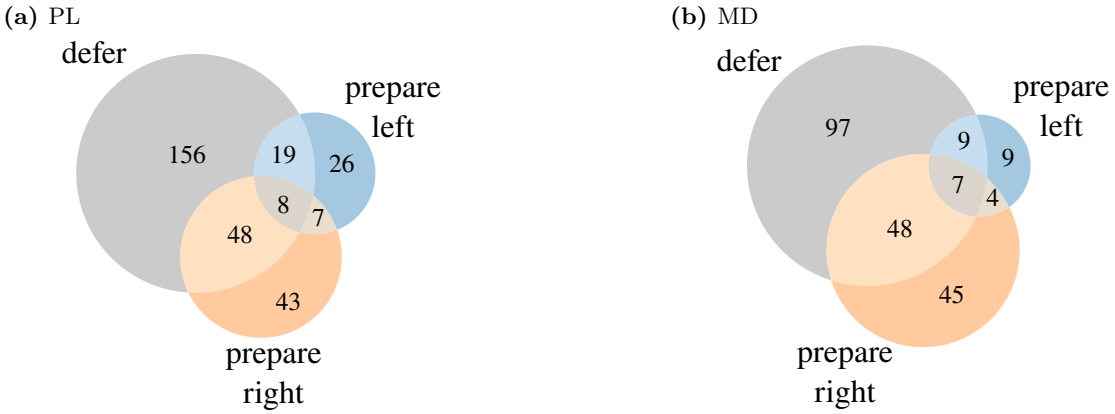

Figure 3.26: Units responsive to individual contexts (0 ms to 600 ms from context cue onset).
(a, b) Area proportional Euler diagrams showing the numbers of cells that were responsive to individual contexts. grey: *defer*; blue: *prepare-left*; orange: *prepare-right*; Overlaps represent cells that were responsive to multiple contexts. Responsiveness is defined as a $p_\zeta < 0.05$. Relative proportions of responsiveness to individual contexts (*defer*, *prepare-left*, *prepare-right*) significantly different across regions: $\chi^2(2, N = 526) = 8.14, p = 0.017$. Proportionally significantly more *prepare-right* than *prepare-left* responsive cells in MD vs. PL: $\chi^2(1, N = 298) = 7.18, p = 0.007$. Proportions of *defer* and *prepare* responsive cells not significantly different across regions: $\chi^2(1, N = 658) = 1.26, p = 0.262$.

right and the fewest to *prepare-left*[47] (Figs. 3.26a and 3.26b and Table 3.5). While the descending order was qualitatively similar in PL and MD the relative proportions of context-responsive cells were significantly different across regions[48]. Of note, the different proportions were realised primarily by a shift in the balance of *prepare* responsive single units: In MD there were proportionally more *prepare-right* than *prepare-left* responsive cells than in PL[49].

As a consequence, in MD the difference between the number of cells responsive to *defer* or *prepare-right* contexts was smaller than in PL. This is reminiscent of the less clear difference between these contexts in MD firing rate magnitudes (Fig. 3.21b).

Finally, relative multi-responsiveness (i.e. overlap/intersect) was not proportionally distributed across contexts. About half of the *prepare*-responsive cells each were also responsive to the *defer* context. In contrast, cells that were responsive to both *prepare* contexts were far less common[50].

Together, this shows that there was a large fraction of cells with single context responsiveness, i.e. a strong specialisation. General responsiveness was inhomogeneously distributed across the three contexts and sizes of subpopulations followed a clear trend, matching the firing rate magnitudes (Fig. 3.21).

3.2.2.3 More single units responsive to *go-left* instruction; less specialisation than for contexts

Approximately two thirds of the cells were responsive to either of the two instructions (Table 3.6 and Figs. 3.27a and 3.27b; PL: 60.9 %, MD: 67.1 %). The *go-left* instruction modulated significantly more cells' activities in both regions (PL: 50.2 % vs. 40.2 %[51]; MD: 52.6 % vs. 38.8 %[52]; *go-left* vs. *go-right*). Proportionally, PL and MD displayed a similar asymmetry[53]. In contrast to relatively strong specialisation of context-responsive cells (Figs. 3.26a and 3.26b), many cells showed significant responsiveness to *both* instructions.

Even so, the fraction of single units that were specialised, i.e. exclusively responsive to a single instruction, was about half amongst all instruction-responsive cells[54]. Notably, specialisation was stronger for the *go-left* instruction: The quantities of exclusively *go-left*-responsive single-units were about twice as many as exclusively *go-right*-responsive units[55].

---

47 PL: 36.7 %, 16.9 %, 9.5 %; MD: 38.3 %, 24.8 %, 6.9 %; *defer*, *prepare-right*, *prepare-left*, respectively

48 $\chi^2(2, N = 691) = 8.14, p = 0.017$

49 Proportions *prepare-right* to *prepare-left* responsive cells in MD vs. PL: $\chi^2(1, N = 298) = 7.18, p = 0.007$. Proportions *defer* to *prepare* responsive cells in MD vs. PL: $\chi^2(1, N = 658) = 1.26, p = 0.262$.

50 PL: 4.3 %, 8.9 %, 2.4 %; MD: 3.8 %, 13.1 %, 2.6 %; *defer AND prepare-left*, *defer AND prepare-right*, *prepare-left AND prepare-right*, respectively

51 $\chi^2(1 - way, N = 529) = 6.98, p = 0.008$

52 $\chi^2(1 - way, N = 384) = 8.76, p = 0.003$

53 $\chi^2(1, N = 953) = 0.30, p = 0.583$.

54 PL: 29.3 % vs. 32.0 % MD: 29.3 %) vs. 32.9 %; multi-responsive vs. single-responsive, respectively

55 single-responsive *go-left* as fraction of *go-left* vs. single-responsive *go-right* as fraction of *go-right*: PL: 41.8 % vs. 27.3 %; MD: 44.3 % vs. 24.5 %

Table 3.6: Instruction-responsive cells. Percentage of cells that were responsive to individual instructions (0 ms to 300 ms from movement instruction onset). A cell was responsive if $p_{zeta} < 0.05$ for ZETA test. Total number of cells: $n_{PL} = 629$, $n_{MD} = 420$

	% total	
responsiveness to	PL	MD
instruction (*go-left* or *go-right*)	60.9	67.1
one instruction only	32.0	32.9
both instructions (intersect)	29.3	29.3
go-left (with intersect *go-right*)	50.2	52.6
go-left (both contexts)	26.2	23.6
go-left (*prepare* only)	13.0	12.9
go-left (*defer* only)	11.0	16.2
go-right (with intersect *go-left*)	40.2	38.8
go-right (both contexts)	16.9	16.0
go-right (*prepare* only)	10.2	11.2
go-right (*defer* only)	13.2	11.7

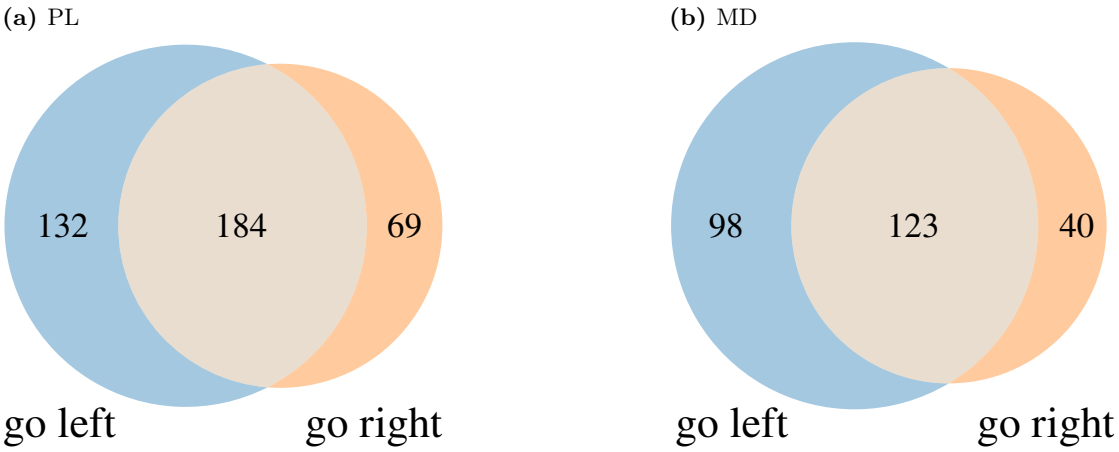

Figure 3.27: Units responsive to individual instructions (0 ms to 300 ms from movement instruction onset).
(a, b) Area proportional Euler diagrams showing the numbers of cells that were responsive to individual instructions. blue: *go-left*; orange: *go-right*; Overlaps represent cells that were responsive to both instructions. Responsiveness is defined as a $p_\zeta < 0.05$. More cells were responsive to *go-left* than *go-right* instruction: $\chi^2(1-way, N = 529) = 6.98, p = 0.008$ (PL), $\chi^2(1-way, N = 384) = 8.76, p = 0.003$ (MD). Similar proportions across PL and MD: $\chi^2(1, N = 953) = 0.30, p = 0.583$.

Table 3.7: Instruction-responsive single units dependent on prior context. Percentage of all cells that were responsive to instruction (0 ms to 300 ms from movement instruction onset), dependent on prior context. A cell was responsive if $p_{zeta} < 0.05$ for ZETA test. Total number of cells: $n_{PL} = 629$, $n_{MD} = 420$

	% total	
responsiveness to	PL	MD
go-left (both contexts)	26.2	23.6
go-left (*prepare* only)	13.0	12.9
go-left (*defer* only)	11.0	16.2
go-right (both contexts)	16.9	16.0
go-right (*prepare* only)	10.2	11.2
go-right (*defer* only)	13.2	11.7

Together this shows that responsiveness to movement instructions was asymmetrically distributed towards more *go-left*-responsive cells, matching the firing rate magnitudes (Fig. 3.23). Subpopulations overlapped more than for contexts, i.e. specialisation was less pronounced. While responsiveness to *prepare* contexts was differentially distributed for contexts, proportions for instruction-responsive cells were remarkably similar across PL and MD.

3.2.2.4 Context-dependent shift of instruction-responsive single-unit subpopulations

The stimuli for movement instruction were always presented after a context cue had been shown (Fig. 3.2). Therefore, to ascertain the effect of prior context on neuronal processing of the instruction, the responsiveness of instruction-modulated single units (see Table 3.6) was investigated in trials with different contexts.

Notably, in about half of instruction-responsive cells both in PL and MD the prior context determined whether a single unit's firing rate would be modulated by the instruction (Table 3.7, symmetric differences in Figs. 3.28a to 3.28d). These context-exclusive cells, which were only instruction-responsive if either a *prepare* or *defer* context had been shown, existed in contrast to context-stable cells, which were instruction-responsive regardless of prior context (intersections in Figs. 3.28a to 3.28d). The proportions of context-exclusive to -stable cells were significantly different in PL. While only 47.8 % of *go-left*-responsive cells were context-exclusive, a higher proportion of *go-right*-responsive PL cells was specialised to responses after only one of the contexts[56] (58.1 %, Figs. 3.28a and 3.28b). This difference was not observed in MD: Here, the fractions of context-exclusive single units were consistently higher than that of the context-stable groups[57] (*go-left*: 55.2 %; *go-right*: 58.9 %, Figs. 3.28c and 3.28d).

[56] $\chi^2(1, N = 569) = 5.59, p = 0.018$
[57] $\chi^2(1, N = 384) = 0.38, p = 0.537$

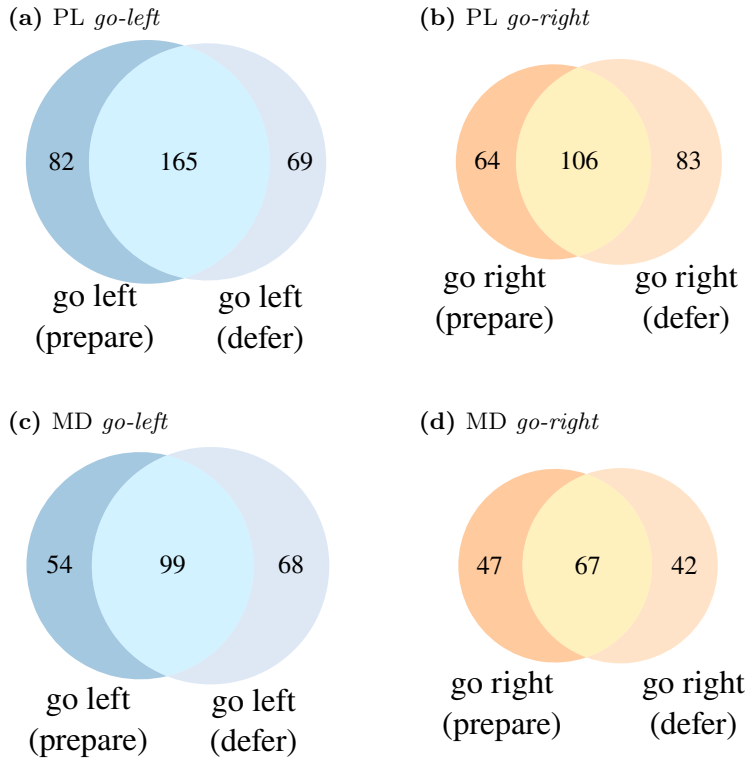

Figure 3.28: Instruction-responsive single units dependent on prior context. **(a)** Area-proportional Euler diagram depicting the numbers of single units in PL that were responsive to *go-left* instruction after the contexts *prepare-left* (darker blue circle to the left), *defer* (lighter blue circle to the right), or after both (intersection in the middle). No significant difference between the proportions of *prepare*-exclusive and *defer*-exclusive units: $\chi^2(1-way, N = 151) = 1.12, p = 0.290$. **(b)** same as (a) for single units in PL that were responsive to *go-right* instruction after the contexts *prepare-right* (darker orange circle to the left), *defer* (lighter orange circle to the right), or after both (intersection in the middle). No significant difference between the proportions of *prepare*-exclusive and *defer*-exclusive units: $\chi^2(1 - way, N = 147) = 2.46, p = 0.117$.
In the main text the single units within the symmetric set difference are called *context-exclusive*, while those within the intersection are called *context-stable*. Proportion of context-exclusive vs. -stable was higher for *go-right*- than for *go-left*-responsive units: $\chi^2(1, N = 569) = 5.59, p = 0.018$. **(c)** same as (a) for single units in MD. No significant difference between the proportions of *prepare*-exclusive and *defer*-exclusive units: $\chi^2(1-way, N = 124) = 1.61, p = 0.205$. **(d)** same as (b) for single units in MD. No significant difference between the proportions of *prepare*-exclusive and *defer*-exclusive units: $\chi^2(1 - way, N = 96) = 0.04, p = 0.838$. Proportion of context-exclusive vs. -stable was *not* significantly different for *go-right*- than for *go-left*-responsive units: $\chi^2(1, N = 384) = 0.38, p = 0.537$.

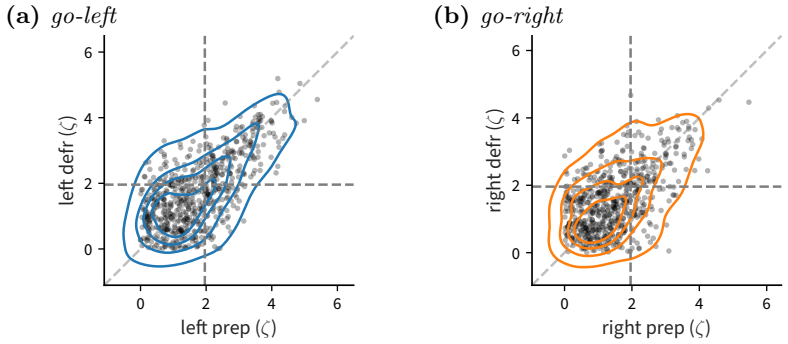

Figure 3.29: Modulation of instruction-responsiveness by context (PL). (a, b) Responsiveness magnitudes (ζ, 0 ms to 300 ms after instruction onset) for single cells in *prepare* and *defer* trials. Every dot is one unit. Overlay is the KDE across units. Horizontal and vertical lines at 1.96 specify the threshold of deeming a unit responsive. Quadrants created by these lines divide the population into non-responsive (lower-left), *defer*-exclusive (upper-left), *prepare*-exclusive (lower-right), context-stable (upper-right). Every dot one unit.

Cells that were exclusively responsive after one context were equally distributed. In other words, neither after *prepare* nor *defer* context were statistically more cells responsive to any of the instructions or in any of the regions[58] (Figs. 3.28a to 3.28d). Rather, the exact cells that responded changed, as evidenced by the overlapping sets described above.

Functional gradients across subregions are a common feature in the brain (Fuster, 2015; Hardung et al., 2017). In order to test if the observed partial shift of responsive subpopulations was accomplished by employing different subregions, the channel depths along the recording probe of the different single unit subsets were investigated. Figures S4a to S4d shows the channel depths for single units that were part of different subsets in responsiveness to instructions. One-way ANOVAs with pairwise Tukey's HSD post-hoc tests found no statistically significant differences across the subsets.

Together this shows that prior contexts partially shifted the neuronal populations that were responsive to the instructions. However, no evidence was found for spatial segregation. In addtion, neither of the contexts led to a higher number of responsive cells.

3.2.2.5 Prior context smoothly modulated and did not bias magnitude of instruction-responsiveness

Although there was no evidence for specialised localisation of context-stable and -exclusive single units it is possible that there was a further distinction on the functional

[58] exclusively responsive after *prepare* or *defer*; PL: 13.0 % vs. 11.0 % (*go-left*), $\chi^2(1-way, N = 151) = 1.12, p = 0.290$ and 10.2 % vs. 13.2 % (*go-right*) $\chi^2(1-way, N = 147) = 2.46, p = 0.117$ MD: 12.9 % vs. 16.2 % (*go-left*) $\chi^2(1-way, N = 124) = 1.61, p = 0.205$ and 11.2 % vs. 11.7 % (*go-right*) $\chi^2(1-way, N = 96) = 0.04, p = 0.838$

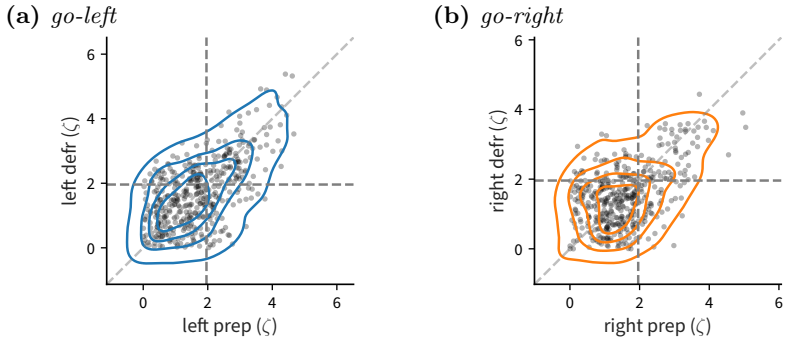

Figure 3.30: Modulation of instruction-responsiveness by context (MD). (a, b) Responsiveness magnitudes (ζ, 0 ms to 300 ms after instruction onset) for single cells in *prepare* and *defer* trials. Every dot is one unit. Overlay is the KDE across units. Horizontal and vertical lines at 1.96 specify the threshold of deeming a unit responsive. Quadrants created by these lines divide the population into non-responsive (lower-left), *defer*-exclusive (upper-left), *prepare*-exclusive (lower-right), context-stable (upper-right). Every dot one unit.

level. In order to follow this idea, in the following I investigated the *magnitude* of single units' responsiveness to the movement instruction after different contexts.

Previously, a single unit's responsivity was threshold dependent. Thus, context-exclusive units' responsiveness magnitude exceeded the threshold for trials with only one context, while the same metric was statistically significant after both contexts in context-stable units. To find out if these categories were distinct and formed clusters, single units' magnitudes of instruction-responsiveness in trials after *prepare* and *defer* contexts was plotted in Figures 3.29a, 3.29b, 3.30a and 3.30b.

Responsiveness magnitudes were smoothly, but not uniformly, distributed and no apparent clusters were found. Especially the groups of context-stable cells were noticeably outside of the bulk of the distribution. However, neither context-stable nor context-exclusive cells formed distinct clusters.

Did one of the contexts bias the populations' instruction-responsiveness magnitudes? For example, even though the same numbers of single units were instruction-responsive exclusively after the *prepare* context, was their instruction-responsiveness higher than that of units that were exclusively responsive after *defer* contexts? Figures S5a, S5b, S6a and S6b show that neither context consistently activated context-exclusive cells that were more or less instruction-responsive[59].

Similarly, instruction-responsiveness was not consistently swayed in context-stable single units (Figs. S5c and S6c). Only in context-stable *go-right*-responsive MD cells was the magnitude significantly higher in trials with a prior *prepare* context[60].

[59] PL: *go-left*: Mann-Whitney-U: $p = 0.584$; *go-right*: Mann-Whitney-U: $p = 0.266$
MD: *go-left*: Mann-Whitney-U: $p = 0.565$; *go-right*: Mann-Whitney-U: $p = 0.433$

[60] PL: *go-left*: 0.001 ± 0.102, $p = 0.941$ (one-sample Wilcoxon test against 0 mean); *go-right*: -0.002 ± 0.110, $p = 0.903$ (one-sample Wilcoxon test against 0 mean)

In sum, the data presented in Sections 3.2.2.1 to 3.2.2.5 shows very similar trends in neuronal responsiveness in PL and MD. Many single units were responsive to at least one of the context cues or movement instructions. In general, populations were not distinct but overlapping. A higher distinction of subpopulations was found for responsivity to contexts, while for movement instructions the overlaps were higher. Prior contexts partially changed which cells were instruction-responsive, but this was not realised via different brain locations nor by different responsiveness magnitudes.

3.2.2.6 Preference

The analyses in Sections 3.2.2.1 to 3.2.2.5 were primarily focussed on whether single units were responsive to one or multiple contexts and/or instructions. In case of multi-responsiveness (i.e. statistically responsive to multiple contexts or instructions) it is unclear which condition a single unit "preferred". However, the information about which condition most strongly modulated a neuron's activity and, thus, which one it mostly represented, might be useful to further evaluate differential processing of those cues. Therefore, in the following I defined a significantly responsive single unit's *preference* as the context or instruction at which maximum responsiveness (argmax of ζ value) was achieved. Analysis windows were the same as in the previous section, i.e. 0 ms to 600 ms after context cue onset for the context-preference, and 0 ms to 300 ms after movement instruction onset for the instruction-preference.

The analysis is comprised of two parts. First, I investigated if preferences were distributed disproportionately to responsiveness observed in the previous section. Second, in order to see if single units' preferences represented behavioural contingencies, the preferences across trial epochs were tested.

3.2.2.7 Higher preference for *go-left* instruction than expected from responsiveness

As shown in Sections 3.2.2.2 and 3.2.2.3 responsiveness to contexts and instructions were inhomogeneously distributed. Briefly, in the group of context-responsive cells most were responsive to the *defer* context and a decreasing number of cells were responsive to *prepare-right* and *prepare-left* contexts (Table 3.5 and Fig. 3.26). Similarly, most instruction-responsive cells were responsive to the *go-left* instruction (Table 3.6 and Fig. 3.27). It doesn't follow, however, that the same ranking was realised via preference. For example, because of the high number of multi-responsive single units it would be possible that even though more single units were responsive to the *go-left* than to the *go-right* instruction, more preferred the latter.

This theoretically possible scenario was not observed. In both regions the rankings of preferred context and instruction were consistent with the proportions of responsive

MD: *go-left*: 0.010 ± 0.107, $p = 0.457$ (one-sample Wilcoxon test against 0 mean); *go-right*: 0.045 ± 0.107, $p = 0.002$ (one-sample Wilcoxon test against 0 mean)

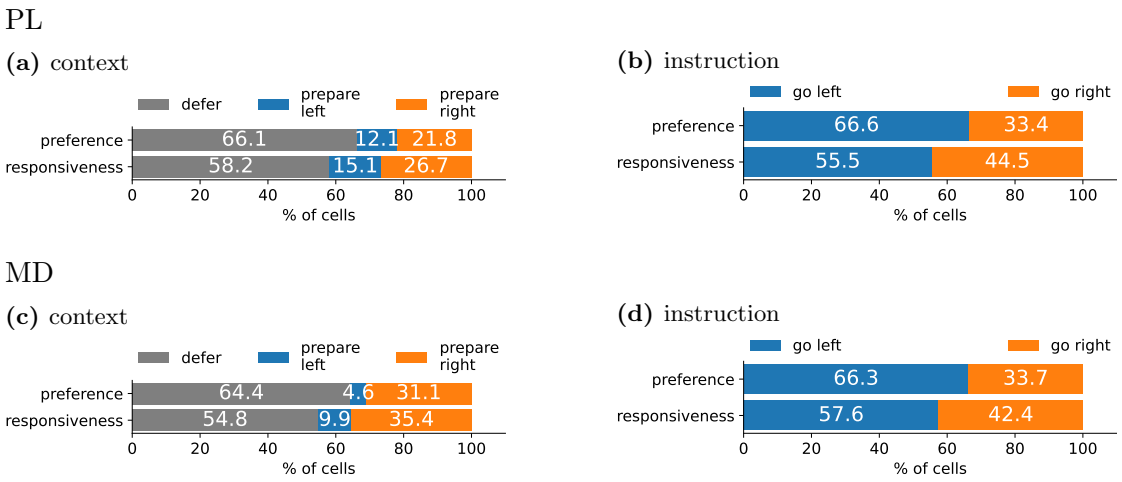

Figure 3.31: Proportions of single units' responsiveness and preference. **(a)** Left: Percentages of context-responsive PL single units that were responsive to the individual contexts. Right: Percentage of context-responsive PL single units that preferred an individual context. Context-preferences were not equally distributed: $\chi^2(1-way, N=307) = 152.94, p = 6.16 \cdot 10^{-34}$. Proportions of context-preference were not significantly different from those of context-responsiveness: ($\chi^2(2, N=704) = 4.62, p = 0.099$. **(b)** Left: Percentages of instruction-responsive PL single units that were responsive to the individual instructions. Right: Percentage of instruction-responsive PL single units that preferred an individual instruction. Instruction-preferences were not equally distributed: $\chi^2(1-way, N=383) = 42.11, p = 8.62 \cdot 10^{-11}$. More single units preferred *go-left* than expected from responsiveness: $\chi^2(1, N=952) = 11.18, p = 8.29 \cdot 10^{-4}$; **(c)** same as (a) but for MD single units. Context-preferences were not equally distributed: $\chi^2(1-way, N=219) = 118.05, p = 2.32 \cdot 10^{-26}$. Proportions of context-preference were significantly different from those of context-responsiveness: all contexts: $\chi^2(2, N=513) = 7.31, p = 0.026$; Higher fraction of *defer*-preferring than *prepare-left*-preferring than expected from responsiveness: *defer* and *prepare-left*: $\chi^2(2, N=341) = 5.38, p = 0.020$. **(d)** same as (b) but for MD single units. Instruction-preferences were not equally distributed: $\chi^2(1-way, N=282) = 30.01, p = 4.29 \cdot 10^{-8}$); More single units preferred *go-left* than expected from responsiveness: $\chi^2(1, N=666) = 4.89, p = 0.027$).

cells. That is, most context-responsive cells preferred *defer*, followed by *prepare-right* and *prepare-left*[61] (Figs. 3.31a and 3.31c). Similarly, the *go-left* instruction was preferred by more units than the *go-right* instruction[62] (Figs. 3.31b and 3.31d).

Although the extreme scenario of preference rankings distinct from those in responsiveness was not observed, it is still possible that preferences differently distributed than what would be expected from the proportions of responsive cells. That is, the inhomogeneous distribution of preferences could be more or less extreme than what was observed for the responsiveness. To investigate that possibility, proportions from preference and responsiveness were compared.

In PL context preference was indeed similarly distributed as expected from the proportions of responsiveness[63] (Fig. 3.31a). In contrast, MD showed a higher number of *defer*-preferring neurons over *prepare-left*-preferring ones than expected from responsiveness[64] (Fig. 3.31c).

In comparison, neurons in both PL and MD disproportionately preferred the *go-left* instruction, thus deviating from what was expected from the proportions of responsiveness[65] (Figs. 3.31b and 3.31d). In other words, in the majority of cells that were responsive to both instructions the activity was more strongly modulated by the *go-left* instruction.

3.2.2.8 Cross-epoch preferences: Most neurons likely to prefer *go-left*; *prepare-right*-preferring most "loyal"

The two *prepare* contexts were predictive and, thus, much of their information was consistent with that conferred by the instructions (Fig. 3.2). This begs the question if single units that preferred one *prepare* context were more likely to prefer the behaviourally contingent movement instruction (e.g. did cells that preferred the *prepare-right* context also prefer to *go-right* instruction?). Complementary to this, which instruction was preferred by the majority of cells that preferred the neutral *defer* context?

Figure 3.32 shows the flows of preferences across epochs. At first glance it is already evident that single units that preferred one context did not uniformly prefer one of the instructions. Rather, the *go-left*- and *go-right*-preferring groups were heterogenously constituted.

To answer the question about consistency in cross-epoch preference, the instruction preferences of single units preferring individual contexts were tested. In PL significantly more *prepare-left*-preferring single units also preferred the behaviourally contingent *go-left* instruction[66]. For the *prepare-right*-preferring cells about equal numbers of cells were

[61] PL: 32.3 %, 10.7 %, 5.9 %, $\chi^2(1-way, N=307) = 152.94, p = 6.16 \cdot 10^{-34}$ MD: 33.6 %, 16.2 %, 2.4 %, $\chi^2(1-way, N=219) = 118.05, p = 2.32 \cdot 10^{-26}$), *defer*, *prepare-right*, *prepare-left*

[62] PL: 40.5 %, 20.3 %, $\chi^2(1-way, N=383) = 42.11, p = 8.62 \cdot 10^{-11}$), MD: 44.5 %, 22.6 %, $\chi^2(1-way, N=282) = 30.01, p = 4.29 \cdot 10^{-8}$), *go-left*, *go-right*

[63] $\chi^2(2, N=704) = 4.62, p = 0.099$).

[64] all contexts: $\chi^2(2, N=513) = 7.31, p = 0.026$ *defer* and *prepare-left*: $\chi^2(2, N=341) = 5.38, p = 0.020$

[65] PL: $\chi^2(1, N=952) = 11.18, p = 8.29 \cdot 10^{-4}$; MD: $\chi^2(1, N=666) = 4.89, p = 0.027$

[66] $\chi^2(1-way, N=23) = 5.26, p = 0.022$

go-left- or *go-right*-preferring. Thus, for these cells the proportion of later preference for *go-left*/*go-right* instruction was less extreme than for the previous group[67]. A similar trend was found for *prepare*-context preferring MD cells: the proportion of *go-left* to *go-right* preferring units was less imbalanced in *prepare-right*-preferring than in *prepare-left*-preferring units. However, due to small sample size statistical power was lower and both proportions were significantly not different[68].

Single units that preferred one of the *prepare* contexts were the minority (Figs. 3.31a and 3.31c). The majority of context-responsive neurons was most strongly modulated by the *defer* context. Furthermore, an even higher fraction was not significantly context-responsive (Fig. 3.25). From a behavioural point of view these cells can be categorised as neutral. Did these neutral cells prefer one or the other movement instruction? Indeed, *defer*-preferring and cells that had been non-responsive to context were about twice as likely to prefer the *go-left* than the *go-right* movement instruction. This was consistent between PL and MD[69].

Since these cells can be conceived of as being context-neutral, it is possible that the likelihood to become *go-left*-preferring was smaller than for cells that preferred the *prepare-left* context. In other words, was the likelihood for *go-left*-preference even higher in *prepare-left* cells, similar to the more extreme imbalance for movement tendency in *prepare-left* trials (Fig. 3.12)? No evidence was found for this hypothesis: the proportions of *go-left* vs *go-right* were statistically not different in *prepare-left*-, *defer*-preferring or non-context-responsive cells[70].

As already demonstrated by the only partial overlap of single units that were both context- *and* instruction-responsive (Fig. 3.25), many cells that had been context-responsive became non-responsive to instructions and vice versa. Was one of the context-preferring groups more involved in cross-epoch representation of task variables? To test this, the proportion of cells that became instruction-responsive was compared with the one that become non-responsive to instructions. In both PL and MD the single units that preferred the *prepare-left* context were equally likely to be responsive or non-responsive to the instructions[71]. In contrast, cells that showed preference for *prepare-right* or *defer* contexts were at least 2.5 times more likely to also be responsive to an instruction than

[67] $\chi^2(1-way, N=52)=0.69, p=0.405$

[68] *prepare-left*-preferring units: $\chi^2(1-way, N=5)=1.8, p=0.180$; *prepare-right*-preferring units: $\chi^2(1-way, N=54)=3.63, p=0.057$;

[69] PL: to *go-left* vs. to *go-right* (from *defer*-preferring) $\chi^2(1-way, N=149)=30.13, p=4.046\cdot10^{-8}$; to *go-left* vs. to *go-right* (from non context-responsive) $\chi^2(1-way, N=159)=11.63, p=6.49\cdot10^{-4}$; MD: to *go-left* vs. to *go-right* (from *defer*-preferring) $\chi^2(1-way, N=101)=20.05, p=7.54\cdot10^{-6}$; to *go-left* vs. to *go-right* (from non context-responsive)

[70] proportions *go-left*- vs *go-right*-preferences; PL from *prepare-left* vs from *defer* $\chi^2(1, N=172)=0.01, p=0.91$; $\chi^2(1, N=182)=0.55, p=0.48$; from *prepare-left* vs from non-context-responsive MD from *prepare-left* vs from *defer* $\chi^2(1, N=106)=0.02, p=0.89$; from *prepare-left* vs from non-context-responsive $\chi^2(1, N=127)=0.11, p=0.74$;

[71] PL $\chi^2(1-way, N=37)=2.19, p=0.139$; MD $\chi^2(1-way, N=10)=0, p=1$;

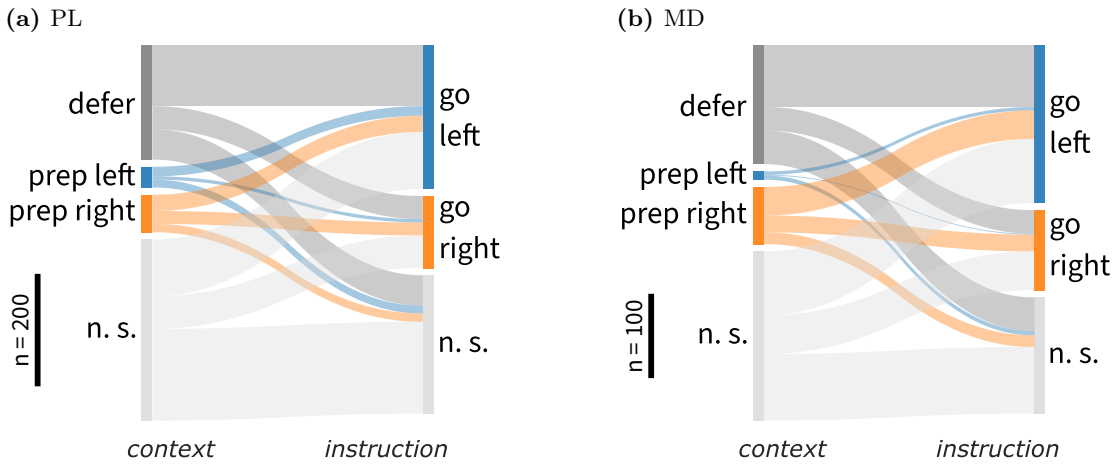

Figure 3.32: Cross-epoch flow of single-unit preferences. **(a)** Bar heights proportional to number of PL units that preferred the corresponding context or instruction. Left: Preferences for context, or lack thereof. Right: Preferences for instruction, or lack thereof. Black bar on the left: scale bar (number of units). **(b)** same as (a) for MD.

not being responsive[72]. Between *prepare-right* and *defer* preferring cells the probabilities were statistically comparable[73].

In sum, if a cell was responsive to an instruction it was much more likely to prefer the *go-left* instruction, irrespective of the preference for a context. Notably, this was less pronounced in cells that had preferred the *prepare-right* context. In contrast, *prepare-left*-preferring cells were not more likely to become *go-left*-preferring than *defer*-preferring or context-non-responsive cells. Thus, weak behavioural contingency was only observed for *prepare-right*-preferring single units, while others had a shared probability for movement instruction preference.

3.2.3 Neuronal recruitment

So far, the time resolution of analyses of single units' activity was limited to trial epochs. We now know per trial epoch how many neurons were significantly responsive to individual contexts and instructions, which conditions they preferred and what individual neurons' preferences were *across* epochs. In addition to these trial-epoch resolved measures, it is of interest *when* cells were recruited, in order to investigate if external cues might have been processed differently.

[72] PL: (from *prepare-right*-preferring) to instruction-responsive vs. to non-responsive $\chi^2(1-way, N = 67) = 20.43, p = 6.18 \cdot 10^{-6}$; (from *defer*-preferring) to instruction-responsive vs. to non-responsive $\chi^2(1-way, N = 203) = 2.19, p = 2.60 \cdot 10^{-11}$; MD: (from *prepare-right*-preferring) to instruction-responsive vs. to non-responsive $\chi^2(1-way, N = 68) = 2.19, p = 1.23 \cdot 10^{-6}$; (from *defer*-preferring) to instruction-responsive vs. to non-responsive $\chi^2(1-way, N = 141) = 2.19, p = 2.79 \cdot 10^{-7}$;

[73] instruction-responsive vs. non-responsive: PL: from *prepare-right* vs. from *defer* $\chi^2(1, N = 270) = 0.27, p = 0.60$; MD: from *prepare-right* vs. from *defer* $\chi^2(1, N = 209) = 1.07, p = 0.30$;

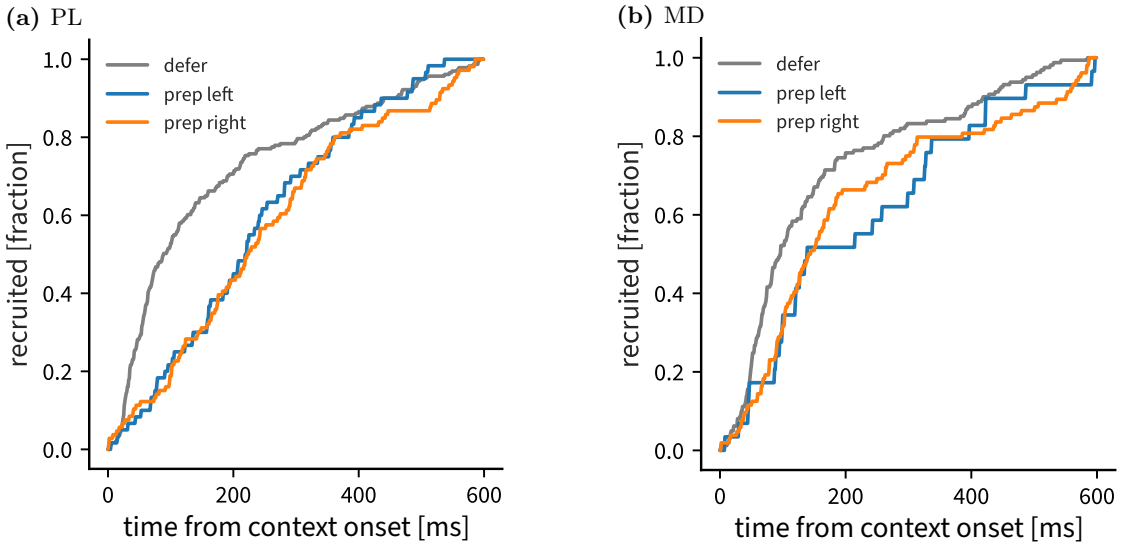

Figure 3.33: Neuronal recruitment profiles (context epoch). **(a)** Empirical CDFs for the recruitment ($p_\zeta < 0.05$) of new single units in PL throughout the epoch in different contexts. Distributions were different between *defer* and *prepare* contexts, but not between *prepare-left* and *prepare-right*: *defer* vs. *prepare-left* $W_2 = 0.35, p = 1.09 \cdot 10^{-5}$, *defer* vs. *prepare-right* $W_2 = 0.36, p = 1.17 \cdot 10^{-8}$, *prepare-right* vs. *prepare-left* $W_2 = 0.12, p = 0.63$, (Epps-Singleton two-sample test). **(b)** same as (a) for MD cells. Distributions were different between *defer* and *prepare* contexts, but not between *prepare-left* and *prepare-right*: *defer* vs. *prepare-left* $W_2 = 0.29, p = 0.02$, *defer* vs. *prepare-right* $W_2 = 0.24, p = 1.23 \cdot 10^{-3}$, *prepare-right* vs. *prepare-left* $W_2 = 0.15, p = 0.66$; (Epps-Singleton two-sample test).

In the following, the latency of a single unit's recruitment was derived from the ZETA-test. First, the recruitment profiles for individual contexts and instructions were compared. Afterwards, the effect of prior context on instruction epoch recruitment was assessed. Here, the prior identification of context-stable and context-exclusive single units was leveraged.

3.2.3.1 Distinct recruitment profiles in *defer* and *prepare* contexts

Figures 3.33a and 3.33b show CDFs of the recruitment times of significantly responsive PL and MD cells for various contexts during the context epoch. For better visual comparison the CDFs were normalised to the total number of cells recruited for the corresponding context (compare Figs. 3.26a and 3.26b). In both regions cells were recruited throughout the context epoch, as indicated by the steady increase of the curves.

The temporal dynamics of recruitment for the *defer* context were clearly visually and statistically different from the dynamics in the *prepare* contexts[74]. Across the two *prepare* contexts, in contrast, the recruitment profiles were similar and statistically not different[75]. In both regions the majority of *defer* responsive cells were recruited within the first 100 ms. Afterwards the recruitment of cells slowed down but didn't entirely stop. Conversely, recruitment for the *prepare* contexts was slower. PL units were almost uniformly recruited throughout the first 400 ms (50 % after 200 ms, Fig. 3.33a). The recruitment speed was halved in the latter third of the context epoch. While the *prepare* recruitment speed in MD was slower than for *defer*, it was not as slow as in PL (Fig. 3.33b). Here, most cells were recruited within the first 200 ms, followed by a slower recruitment speed from 200 ms to 600 ms.

In sum, recruitment profiles during the context epoch were clearly distinct between *defer* and *prepare* conditions, while with the *prepare* conditions no difference was found. Across PL and MD recruitment profiles were very similar apart from only minor deviations.

3.2.3.2 Instruction-epoch recruitment distinct in PL but similar in MD

Similar to the context epoch recruitment profiles during the instruction epoch showed recruitment throughout the entire epoch. The rates of recruitment differed between PL MD. In PL, recruitment speeds showed several phases (Fig. 3.34a). During the first 50 ms there was a fast increase of newly recruited cells. This was followed by a brief phase of about 50 ms with slower speeds, which evolved into a phase with higher recruitment speeds. This pattern is reminiscent of the firing rate dynamics in PL (Fig. 3.23a). In contrast, responsive cells in MD were initially recruited fast (50 % in the first 100 ms) and then recruitment rates slowly tapered off (Fig. 3.34b). Interestingly, recruitment dynamics for *go-left* and *go-right* were statistically different in PL but not in MD[76].

Together, recruitment profiles in the context epoch were similar across regions but different within, i.e. across contexts. Conversely, recruitment in the instruction epoch differed more between the regions.

3.2.3.3 Prior contexts led to similar recruitment profiles but distinct recruitment sequences of context-stable cells

At least in PL the recruitment profiles differed already at early times of the instruction epoch. This could have been caused by prior *prepare* contexts. Thus, in order to

[74] PL: *defer* vs. *prepare-left* $W_2 = 0.35, p = 1.09 \cdot 10^{-5}$, *defer* vs. *prepare-right* $W_2 = 0.36, p = 1.17 \cdot 10^{-8}$; MD: *defer* vs. *prepare-left* $W_2 = 0.29, p = 0.02$, *defer* vs. *prepare-right* $W_2 = 0.24, p = 1.23 \cdot 10^{-3}$, Epps-Singleton two-sample test for equality of distribution functions

[75] PL: *prepare-right* vs. *prepare-left* $W_2 = 0.12, p = 0.63$; MD: *prepare-right* vs. *prepare-left* $W_2 = 0.15, p = 0.66$; Epps-Singleton two-sample test for equality of distribution functions

[76] Epps-Singleton two-sample test for equality of distribution functions; PL: $W_2 = 0.229, p = 5.43d - 7$, MD: $W_2 = 0.106, p = 0.22$

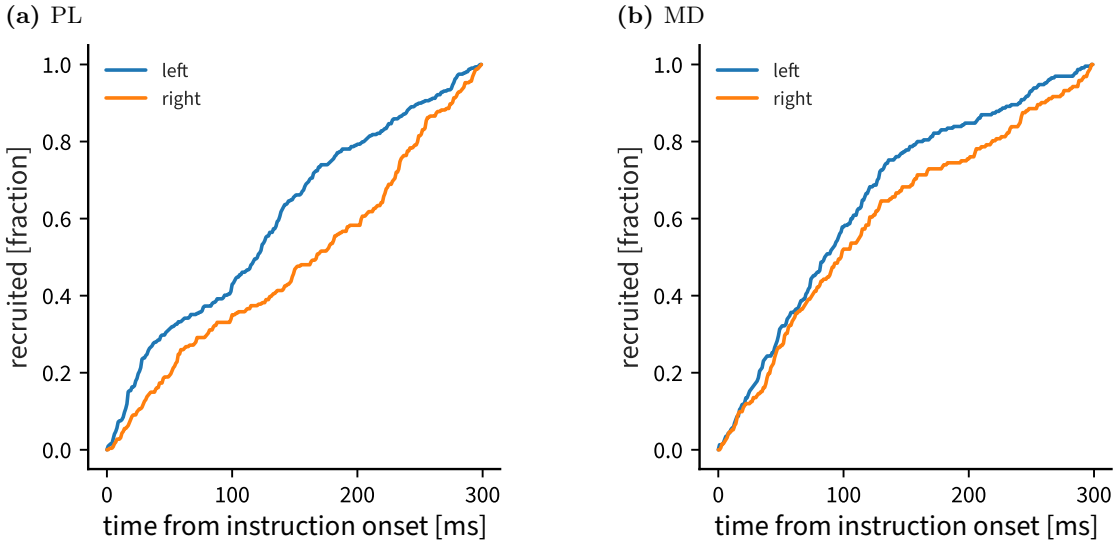

Figure 3.34: Neuronal recruitment profiles (instruction epoch). **(a)** Empirical CDFs for the recruitment ($p_\zeta < 0.05$) of new single units in PL throughout the epoch in different instructions. Distributions were different between *go-left* and *go-right* instructions: $W_2 = 0.229, p = 5.43 \cdot 10^{-7}$ (Epps-Singleton two-sample test). **(b)** same as (a) for MD cells. Distributions were *not* different: $W_2 = 0.106, p = 0.22$ (Epps-Singleton two-sample test).

see if there was an effect of prior context on instruction epoch recruitment that was possibly diluted by collective examination, recruitment was investigated separately for trials after *prepare* and *defer* contexts. Prior contexts could have affected recruitment in different neuronal subpopulations: either large population of cells that were significantly instruction-responsive after both contexts (context-*stable*, Figs. 3.28a to 3.28d) or the non-overlapping populations of cells that were instruction-responsive after only one type of context (context-*exclusive*, Table 3.7).

In context-stable single units the prior contexts did, by definition, not affect if a cell was significantly responsive to an instruction (Table 3.7). Furthermore, the magnitude of responsiveness was not or only marginally influenced by prior contexts (Figs. 3.29 and 3.30). Were these cells also "stable" in terms of recruitment during the instruction epoch?

Figures 3.35a to 3.35d show that there were striking cross-context similarities in the recruitment dynamics for the different instructions in PL and MD. Concretely, the cumulative numbers of recruited cells in *prepare* trials were indistinguishable from *defer* trials in any condition during the epoch[77].

[77] PL: *go-left*: $W_2 = 0.067, p = 0.858$, *go-right*: $W_2 = 0.066, p = 0.872$; MD: *go-left*: $W_2 = 0.141, p = 0.276$, *go-right*: $W_2 = 0.090, p = 0.953$; Epps-Singleton two-sample test for equality of distribution functions (*prepare* vs. *defer*)

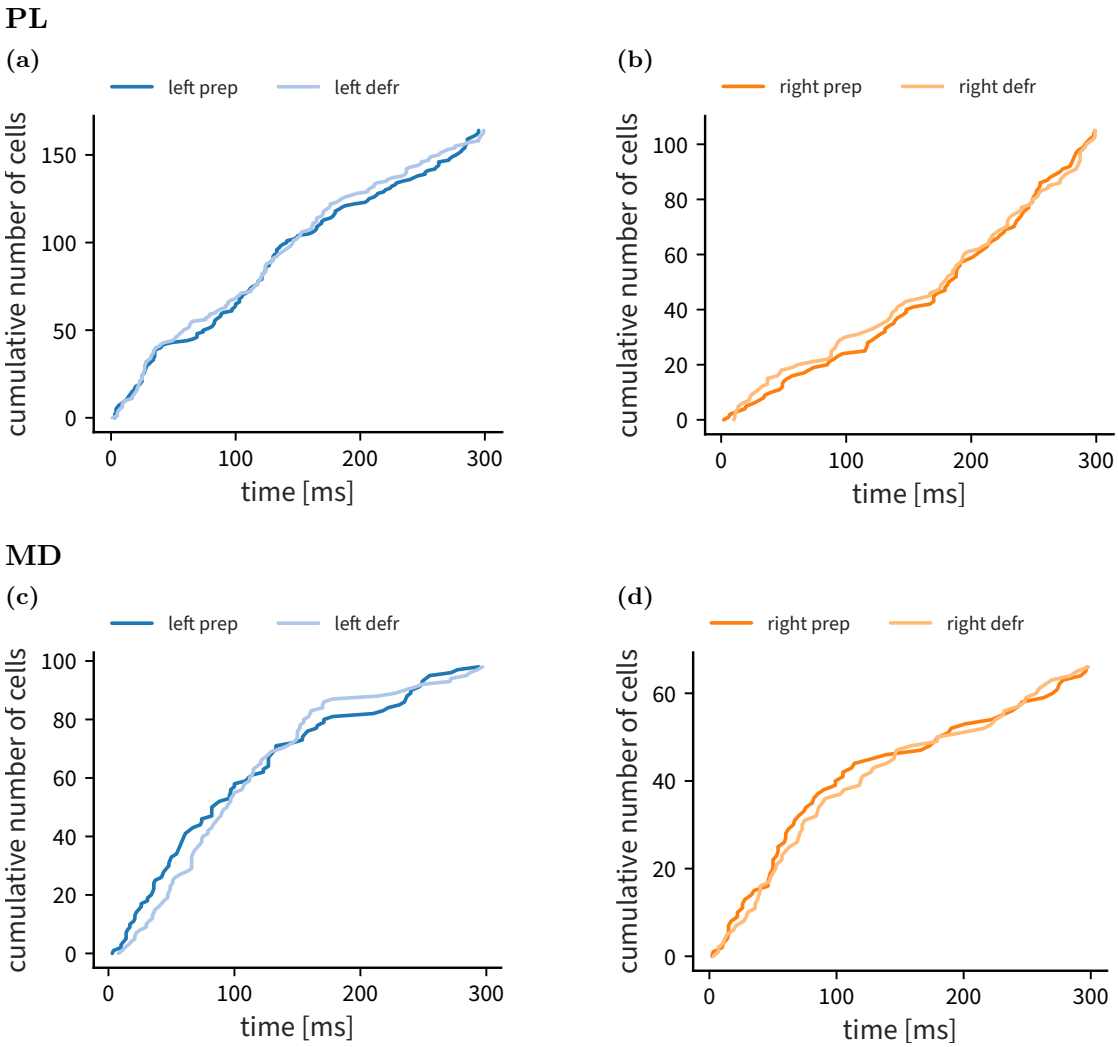

Figure 3.35: Recruitment profiles of context-*stable* single units during instruction, split by prior context. **(a)** CDF of recruitment times of context-stable single units from PL that were significantly responsive to *go-left* instruction. Sorting was done within the different conditions, y-axis doesn't reflect cell identity. **(b)** same as (a) for *go-right*-responsive single units. Distributions were not different between *prepare* and *defer*: *go-left*: $W_2 = 0.067, p = 0.858$, *go-right*: $W_2 = 0.066, p = 0.872$ (Epps-Singleton two-sample test). **(c, d)** same as (a, b) for context-stable single units in MD. Distributions were not different between *prepare* and *defer*: *go-left*: $W_2 = 0.141, p = 0.276$, *go-right*: $W_2 = 0.090, p = 0.953$). (Epps-Singleton two-sample test).

PL

(a)

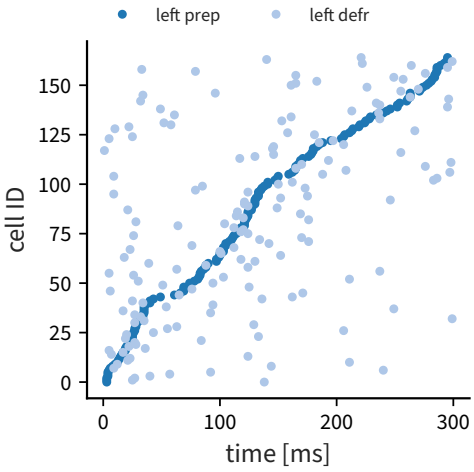

(b)

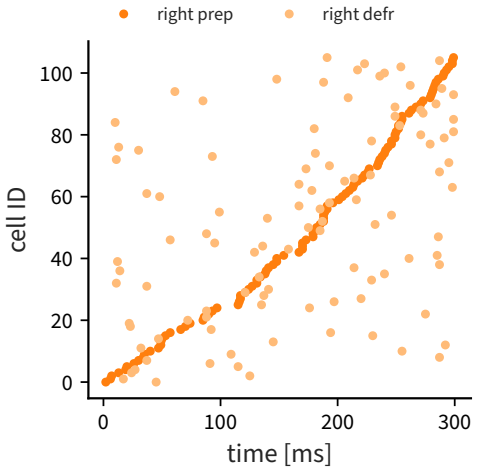

MD

(c)

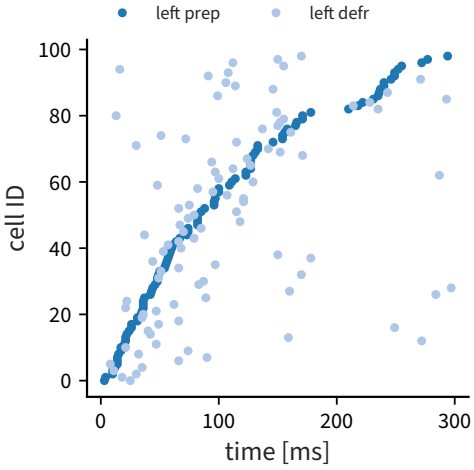

(d)

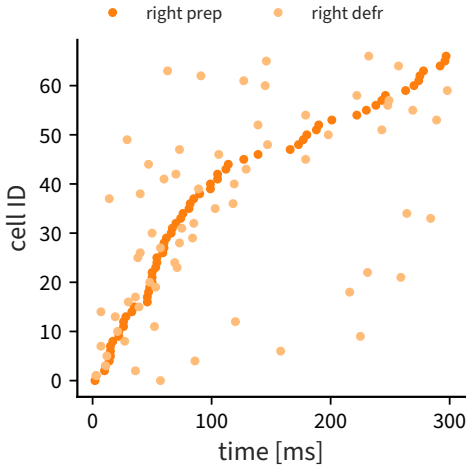

Figure 3.36: Recruitment times, sorted by *prepare* trials, in context-stable instruction-responsive single units. **(a)** Recruitment times of context-stable *go-left*-responsive single units from PL, sorted by times in *prepare* trials. Per definition, the recruitment sequence in *prepare* trials is reproduced. Recruitment sequence was statistically not the same in *defer* trials: $\tau = -0.010, p = 0.846$. **(b)** same as (a) for context-stable *go-right*-responsive PL cells. Recruitment sequences were statistically not the same: *go-right*: $\tau = -0.009, p = 0.894$. **(c, d)** same as (a, b) for MD single units. Recruitment sequences were statistically not the same: *go-left*: $\tau = 0.033, p = 0.631$, *go-right*: $\tau = 0.046, p = 0.585$. All statistical tests: Kendall's τ for correspondence between rankings.

Sequential activity (*"tiling"*) is a phenomenon that has been observed in mouse PL (Parnaudeau et al., 2018). The implicit assumption in studies that show this is that information is sequentially passed on from one neuron to the next. Since the recruitment profiles of context-stable cells were virtually indistinguishable between trials with prior *defer* and *prepare* context, it seemed likely that recruitment sequences did not substantially change across contexts. To test this idea, neurons were ordered according to their recruitment time sorted in the respective *prepare* trials (Figs. 3.36a to 3.36d). Notably, applying the same sorting on the recruitment times in *defer* trials showed that recruitment sequences considerably varied across contexts[78].

In contrast to the previous neuronal subpopulations, context-exclusive cells were non-overlapping groups. Prior context *did*, by definition, modulate their instruction-responsiveness. Did the context also influence when single units were recruited?

Visually, in PL recruitment in *prepare* trials appeared to be faster than in *defer* trials (Figs. 3.37a and 3.37b). However, this trend was statistically significant only for *go-right* trials[79]. Conversely, recruitment in MD had less clear dynamics. Some *defer*-exclusive neurons in *go-left* trials were recruited at a faster pace after 100 ms, while for *go-right* recruitment for *prepare*-exclusive units tended to be faster (Figs. 3.37c and 3.37d). Overall, however, the CDFs were statistically not different[80].

Together, the data shown here demonstrated that while recruitment profiles were stable across contexts in single units with context-stable instruction-responsiveness, their times of recruitment were not context-stable, thus constituting different recruitment sequences. In addition, recruitment profiles for context-exclusive cells were more distinct, especially in PL.

All in all, the data described in Section 3.2.3 shows similar neuronal recruitment properties in PL and MD. Recruitment profiles during the context epoch were mainly distinguished by the abstract split of *defer* against *prepare* contexts. The biggest differences between the regions was found for the recruitment profiles during the instruction epoch. Here, profiles differed within PL but not within MD. Further investigation of cross-context within-instruction recruitment again showed similar characteristics across regions. While cross-context recruitment profiles for instructions were virtually identical, recruitment sequences differed.

[78] PL: *go-left*: $\tau = -0.010, p = 0.846$, *go-right*: $\tau = -0.009, p = 0.894$; MD: *go-left*: $\tau = 0.033, p = 0.631$, *go-right*: $\tau = 0.046, p = 0.585$; ranking *prepare* vs. ranking *defer*, Kendall's τ for correspondence between rankings

[79] *go-left*: $W_2 = 0.179, p = 0.157$, *go-right*: $W_2 = 0.244, p = 0.021$ (Epps-Singleton two-sample test for equality of distribution functions)

[80] *go-left*: $W_2 = 0.190, p = 0.197$, *go-right*: $W_2 = 0.178, p = 0.375$ (Epps-Singleton two-sample test for equality of distribution functions)

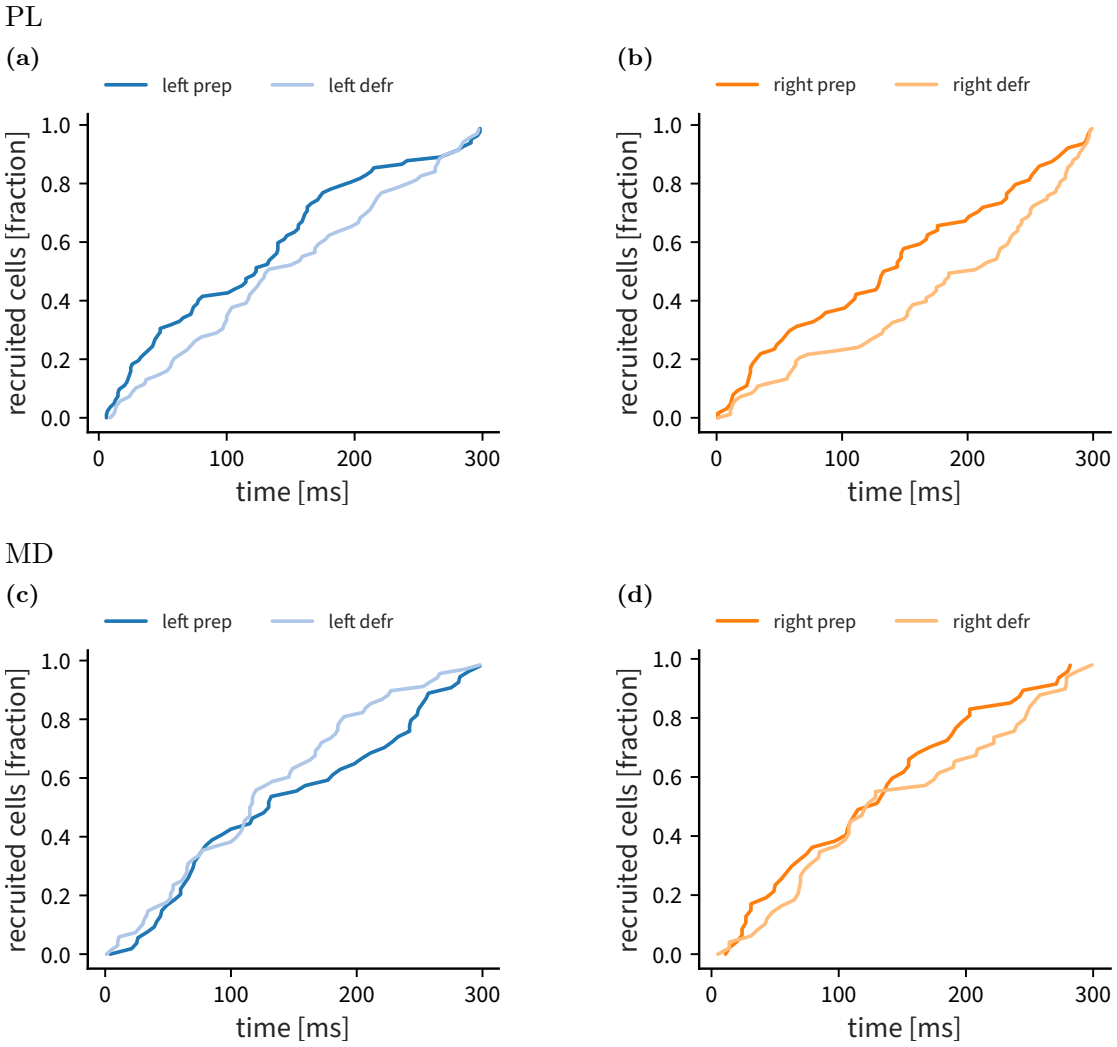

Figure 3.37: Recruitment profiles of context-exclusive single units during instruction, split by prior context. **(a)** CDF of recruitment times of context-exclusive single units from PL that were significantly responsive to *go-left* instruction. Distributions were not different between *prepare* and *defer*: $W_2 = 0.179, p = 0.157$. **(b)** same as (a) for *go-right*-responsive cells. Distributions were statistically different between *prepare* and *defer*: $W_2 = 0.244, p = 0.021$. **(c, d)** same as (a, b) for MD cells. Distributions were not different between *prepare* and *defer*: *go-left*: $W_2 = 0.190, p = 0.197$, *go-right*: $W_2 = 0.178, p = 0.375$. All tests: Epps-Singleton two-sample test for equality of distribution functions.

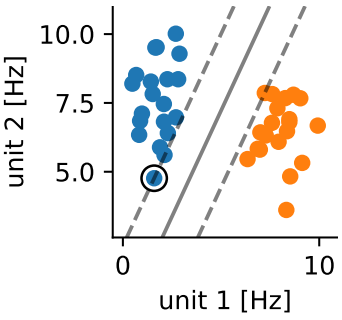

Figure 3.38: Schematic of a classification using a support vector machine in 2D space. Dots represent firing rates of two neuronal units in different trials. Different colours represent firing rates in different conditions. The support vector machine finds the hyperplane (solid line) that separates the two conditions with the biggest margins (dashed lines) to the next data point (circled blue dot). Different neuronal coding would be represented by the groups of dots occupying different regions and a different resultant hyperplane. Thus, the hyperplane is a proxy of neuronal coding. Adding more neurons' firing rates results in a high-dimensional hyperplane.

3.2.4 Population coding

So far, the analysis focused on activity of single neurons in response to the contexts and instructions. This disregards the fact that neurons do not work in isolation but in wider networks. By ways of interplay between neurons, neuronal networks can thus lead to emergent effects. One of these effects can be improved / more robust encoding of task variables.

In this section I employed Support-Vector Machine (SVM) classifiers in order to investigate coding of task variables in neuronal populations. Classifiers were trained on the combined spiking activity of hundreds of single units. Thus, a classifier trained at one point in time represents a high-dimensional fingerprint of the neuronal population code at that time (Fig. 3.38). This fingerprint can be used in two complementary ways. First, testing a classifier at the same point in time that it was trained in shows how robust the population code was *across trials*. Secondly, using the same classifier but testing at points in time in which it was not trained shows the cross-temporal stability of the neuronal code (E. M. Meyers, 2018). In other words, this shows if and how fast the population code changed.

In the following, the classifier approach was applied to check coding strength and stability for context cues and movement instruction. Furthermore, to find out if prior contexts changed robustness and cross-temporal generalisability, population codes for movement instructions were dissected by context. Finally, to leverage behavioural data recorded via the response ball, I investigated if population coding was affected by micro-movements during the context epoch.

3.2.4.1 PL and MD populations strongly encoded context cues

During the context epoch three different contexts were shown (*defer*, *prepare-left*, *prepare-right*). Thus, a random "guess" of a trial's context would on average lead to an accuracy of 33.3 %. This would be the case if the classifier would not have been able to extract any information from the population's activity. Figures 3.39a and 3.39c show the temporal evolution of information about the context contained in the firing rates of PL and MD populations within the context epoch. In both populations the information rapidly increased after the onset of the context cue and peaked at around 80 %. The peak occurred

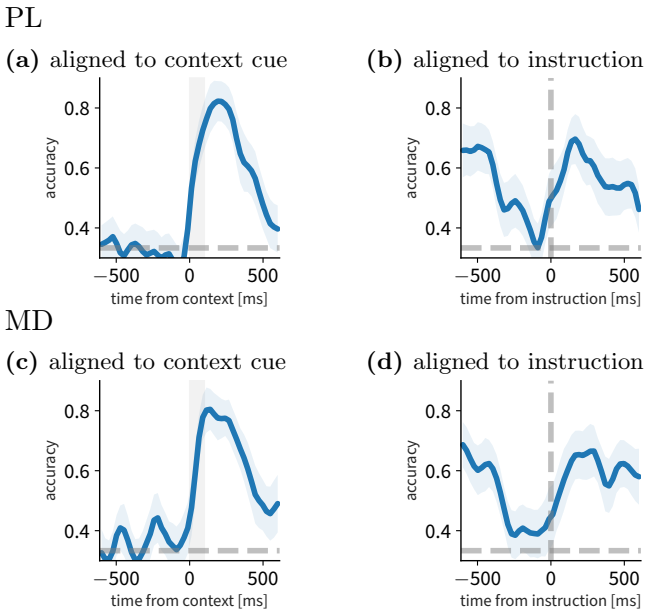

Figure 3.39: Context cue information in neuronal populations. (a) Classification accuracies of SVM classifiers trained on the neuronal activity of the PL pseudopopulation and tested at the same points in time, aligned to context cue onset. Saturated line: Mean across pseudotrial shuffles ($n = 100$). Blue shading: Standard deviation across pseudotrial shuffles. Grey shading: Time of context stimulus presentation. Horizontal dashed line: Chance level. **(b)** same as (a) but with data aligned to instruction cue. **(c, d)** same as (a, b) but for the MD pseudopopulation.

about 100 ms after cue offset in PL and at offset in MD. After the peak the information in the populations decayed close to chance level.

The length of the context epoch varied across trials. At the minimum it was 600 ms if an animal did not move the response ball more than a certain small "grace rotation". If, however, an animal moved the response ball across the acceptable threshold the context epoch was extended by 100 ms increments (Fig. 3.2 and Section 7.4.3). Thus, progression into the instruction epoch was dependent on terminating response ball movements. In order to see if context information kept decaying or if it was reactivated shortly before the instruction epoch, classifiers were also trained on population activity that was aligned to the onset of the instruction epoch. Figures 3.39b and 3.39d show that both regions' context information was almost absent shortly before the instruction epoch (negative times, left parts of the plots). Only small traces of context information were observed for MD. Context information increased after the onset of the instruction epoch, as expected by the behavioural contingency of *prepare* contexts with the respective movement instructions.

To further explore how preparative information was processed the accuracies to distinguish between the two *prepare* contexts was assessed (Fig. 3.40). Peak coding of the to-be-prepared side was significantly earlier in MD than in PL[81]. Furthermore, in MD

[81] difference: 78 ms, $p < 0.01$ (permutation test, $n = 1000$)

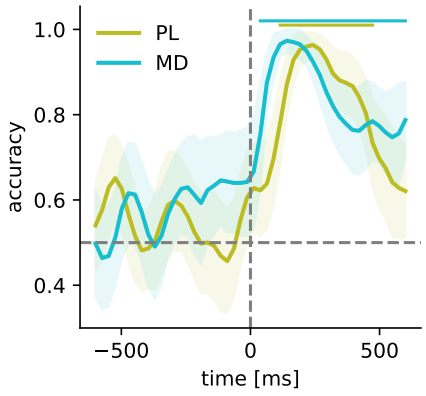

Figure 3.40: Comparison of population coding of upcoming action in PL and MD. Classification accuracies of classifiers trained on the neuronal activity of the pseudopopulations in *prepare-left* and *prepare-right* trials and tested at the same points in time, aligned to context cue onset. Saturated line: Mean across pseudotrial shuffles ($n = 100$). Shading: Standard deviation across pseudotrial shuffles. Vertical dashed line: Time of context stimulus presentation. Horizontal grey dashed line: Chance level. Coloured lines at the top: significant coding ($p < 0.01$, permutation test).

significant coding had a lower latency and was sustained throughout the context epoch[82].

3.2.4.2 Population code for context was partially generalisable within the context epoch but distinct from code in the instruction epoch

As alluded to earlier in the introductory paragraph of this section, classifier analyses can be used to assess the stability of the neuronal code. If neuronal code was stable (i.e. did not change) across a certain time interval, high decoder accuracy is expected for classifiers trained and tested on different time points within the interval. On the other hand, if the code rapidly changed, only classifiers that were trained and tested on the same point in time would achieve high classification accuracy.

Cross-temporal coding stability of context cue information was in investigated via cross-temporal testing. Figure 3.41 shows the classifier accuracies represented by pseudocolours. Accuracies for classifiers tested at the same time points that they were trained in are plotted at the main diagonal that spans from the top left to the bottom right. Off-diagonal pixels represent cross-temporal accuracies. Onset of the context and instruction epochs are represented by vertical and horizontal dashed lines. If the code was entirely stable within the context epoch, the lower right quadrant of Figures 3.41a and 3.41c would be expected to have uniform colouring. In contrast, the experimentally observed code stability varied within the context epoch. Cross-temporal generalisability was relatively low in the first half of the context epoch and subsequently increased, as represented by the funnel-like structure.

Due to the behavioural contingency of *prepare* contexts and movement instructions it is possible that the same population code was used across epochs. In Figures 3.41b and 3.41d the top right quadrants show cross-temporal coding from context epoch to instruction epoch, and the bottom left quadrants show the inverse. "Cool" pseudocolours in those quadrants show that in general population codes were distinct across the epochs. One notable exception were the first 100 ms of the instruction epoch in PL. At that time the same code that was found during the context epoch was weakly reactivated.

[82] PL: 115 ms to 472 ms; MD: 38 ms to 600 ms

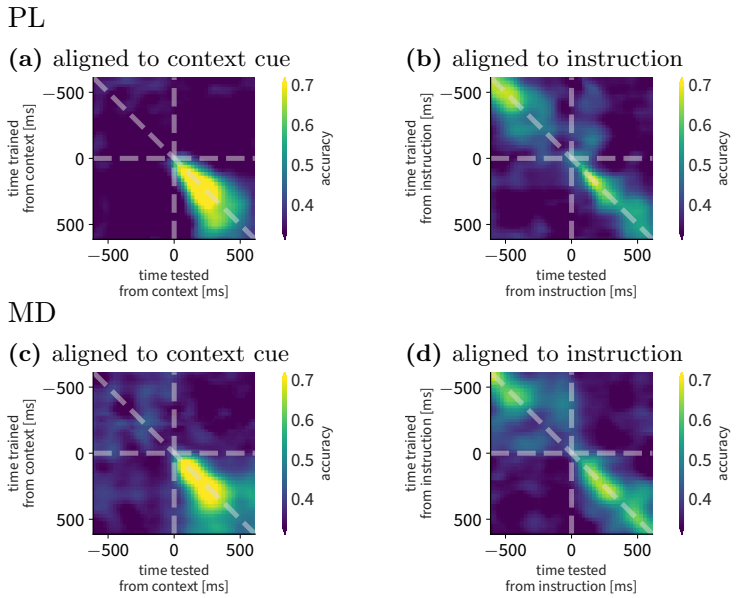

Figure 3.41: Cross-temporal code stability of context cue information in neuronal populations. **(a)** Cross-temporal classification accuracy (mean across pseudotrial shuffles), represented as pseudocolours. Decoders were trained on data from time points shown on the y-axis and tested on data from time points shown on the x-axis. Vertical and horizontal dashed lines represent time of stimulus alignment. Diagonal dashed line represents equal train and test times (see Figure 3.39). Classification accuracies of SVM classifiers trained on the neuronal activity of the PL pseudopopulation and tested at the same points in time, aligned to context cue onset. **(b)** same as (a) but with data aligned to instruction cue. **(c, d)** same as (a, b) but for the MD pseudopopulation.

In sum, Sections 3.2.4.1 and 3.2.4.2 show that the context was strongly encoded in both regions but decayed towards the end of the context epoch. Upcoming instructions were encoded earlier and longer in MD than in PL. Analysis of cross-temporal coding revealed an evolution from highly dynamic to more generalisable codes. Generalisability was, however, mostly restricted to the context epoch and did not propagate substantially into the instruction epoch.

3.2.4.3 Movement instruction was very robustly encoded and persisted after trial outcome

As shown in previous sections, both PL and MD displayed high firing rates (Fig. 3.23) and contained a higher number of single units that were responsive (Fig. 3.25) to the two movement instructions *go-left* and *go-right*. In order to investigate how this translated into the population code, the classifier approach was extended to these task variables.

Using the population activity the instructions of single trials could be decoded almost perfectly, as represented by accuracies close to 100 % (Figs. 3.42a and 3.42c). Due to the exact behavioural contingencies of *prepare* contexts and respective movement instructions the classifier accuracies were above chance level ($> 50\,\%$) before the instruction epoch started. Shortly after the onset of the instruction accuracies increased rapidly and reached peak values at around 150 ms (PL, Fig. 3.42a) and 200 ms (MD, Fig. 3.42c). Subsequently, information decreased only slowly and retained high levels.

This persistence of information might be a side effect of the varying durations across trials that were needed to rotate the response ball across the target threshold. Therefore, to investigate the evolution of movement instruction information up to and after this important trial event, data was also aligned to threshold crossing. Figures 3.42b and 3.42d show that, indeed, information built up to maximum levels shortly before threshold crossing. Although instruction information decreased subsequently it remained high even during the times when the animal received and consumed reward. Notably, while in PL information decayed immediately after threshold crossing, information in MD persisted at peak levels for about 200 ms before decaying.

Together, this shows that movement instruction was very rapidly and robustly encoded by both PL and MD. Information built up until threshold crossing and subsequently persisted for long durations.

3.2.4.4 Code stability for movement instruction varied strongly in PL but not MD

To assess the code stability of the population code for movement instruction, cross-temporal decoding was used, similar to what was done for decoding the context.

In PL coding stability varied considerably within the instruction epoch (Fig. 3.43a). Coinciding with the initial information increase and peak after instruction onset the code was very dynamic in the first 150 ms, as represented in the plot by the very narrow band of high accuracy off the diagonal. Subsequently, however, the code became much more generalisable. Interestingly, aligning to the crossing of the target threshold revealed

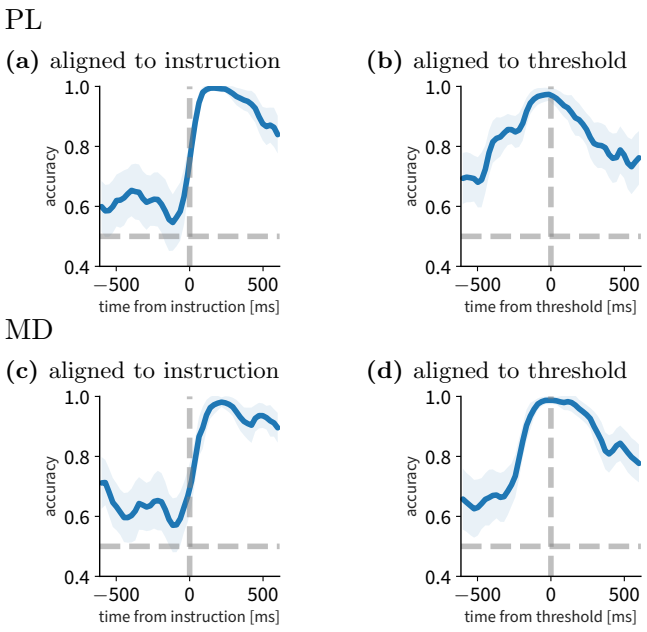

Figure 3.42: Movement instruction information in neuronal populations. (a) Classification accuracies of SVM classifiers trained on the neuronal activity of the PL pseudopopulation and tested at the same points in time, aligned to onset of movement instruction. Saturated line: Mean across pseudotrial shuffles ($n = 100$). Blue shading: Standard deviation across pseudotrial shuffles. Vertical dashed line: Time of instruction stimulus presentation. Horizontal dashed line: Chance level. **(b)** same as (a) but with data aligned to threshold crossing. **(c, d)** same as (a, b) but for the MD pseudopopulation.

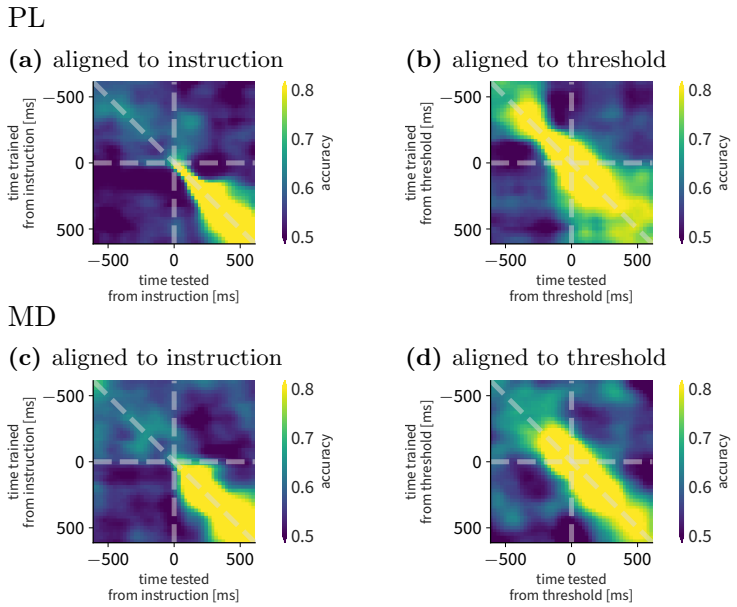

Figure 3.43: Cross-temporal code stability of movement instruction information in neuronal populations. **(a)** Cross-temporal classification accuracy (mean across pseudotrial shuffles), represented as pseudocolours. Decoders were trained on data from time points shown on the y-axis and tested on data from time points shown on the x-axis. Vertical and horizontal dashed lines represent time of trial event alignment. Diagonal dashed line represents equal train and test times (see Figure 3.42). Classification accuracies of SVM classifiers trained on the neuronal activity of the PL pseudopopulation and tested at the same points in time, aligned to onset of movement instruction. **(b)** same as (a) but with data aligned to threshold crossing **(c, d)** same as (a, b) but for the MD pseudopopulation.

that the code in PL was very stable within the epoch after threshold crossing, since the same classifier achieved good accuracies at virtually every time point (Fig. 3.43b). Furthermore, this alignment uncovered that there were two other phases of code stability in PL: up to 250 ms before threshold crossing was a phase of relatively high stability, which was followed by a highly dynamic phase in the remaining time before the target was hit.

In contrast to PL code stability in MD did not vary substantially. Here, the code generalised to around 250 ms almost from the beginning of the instruction epoch and remained at that bandwidth at later times (Fig. 3.43c). Alignment to the crossing of the threshold was consistent with this: generalisability was at the same bandwidth and did not change (Fig. 3.43d). Consequently, this means that overall code stability was not boosted after threshold crossing like it was observed for PL.

In sum, code stability for movement instruction differed between the regions. While there were two to three metastable periods in PL, the stability bandwidth of the code in MD did not change.

3.2.4.5 Prior context only weakly influenced strength of instruction coding but shifted its onset

Figure 3.42 showed that the neuronal populations encoded movement instruction almost perfectly, even if no distinction was made between trials of different prior context. Does that mean that movement instruction coding was so robust that prior context did not have any influence? Intuitively, *prepare* contexts might have been able to boost coding levels because these contexts were perfectly predictive of the upcoming movement instructions. To further assess this, classifiers were trained and tested exclusively on either *prepare* or *defer* trials.

As expected, in *prepare* trials there was high instruction information before the onset of the instruction (Figs. 3.44a and 3.44d). Similarly consistent with expectation, no significant instruction coding was found in *defer* trials (Figs. 3.44b and 3.44e). Directly comparing both kinds of trials shows that indeed information content was consistently significantly higher in *prepare* trials[83] (Figs. 3.44c and 3.44f). However, this was mainly an effect of earlier coding onset in *prepare* trials than in *defer* trials, while temporal dynamics and peak information were similar.

3.2.4.6 Population code for movement instruction conserved across contexts

In the previous section it was shown that maximum decoder accuracy, as a proxy for information content in the neuronal population, was only weakly influenced by prior context. Conversely, similar information profiles appeared to be simply offset in time.

Did the contexts only shift instruction coding *onsets* or was the population code *itself* changed? The similar but shifted accuracy profiles in Figures 3.44c and 3.44f open up the possibility that the same code was economically reused but time-shifted in trials with different prior contexts. In order to test this possibility, classifiers trained in one context were compared to classifiers trained in the other context. As stated earlier, classifiers extract a kind of "fingerprint" of the neuronal activity across the entire population. The parametrisation of this profile specifies a high-dimensional hyperplane which can be thought of as a proxy of the neuronal code. Thus, comparing the hyperplane angles across different classifiers indicates how similar the neuronal code was in the trials that were used for classifier training. If the codes were merely shifted in time, this would be visible as an off-diagonal line of low hyperplane angles (i.e. high similarity) in a plot comparing cross-context codes at different points in time. Concretely, if the code in *defer* trials was simply the code of *prepare* trials but delayed by 100 ms, the angles of *defer* classifier hyperplanes at time t would be smallest in comparison to *prepare* classifier hyperplanes at time $t - 100\,\text{ms}$.

In contrast to this time-shift hypothesis, Figures 3.45a and 3.45d show that for the majority of the instruction epoch, the codes across contexts were most similar at the same points in time (i.e. on the main diagonals of the respective plots). Notably, the codes' similarity varied during the epoch. In the first 100 ms codes started out dissimilar

[83] Cluster-corrected permutation tests $p < 0.01$

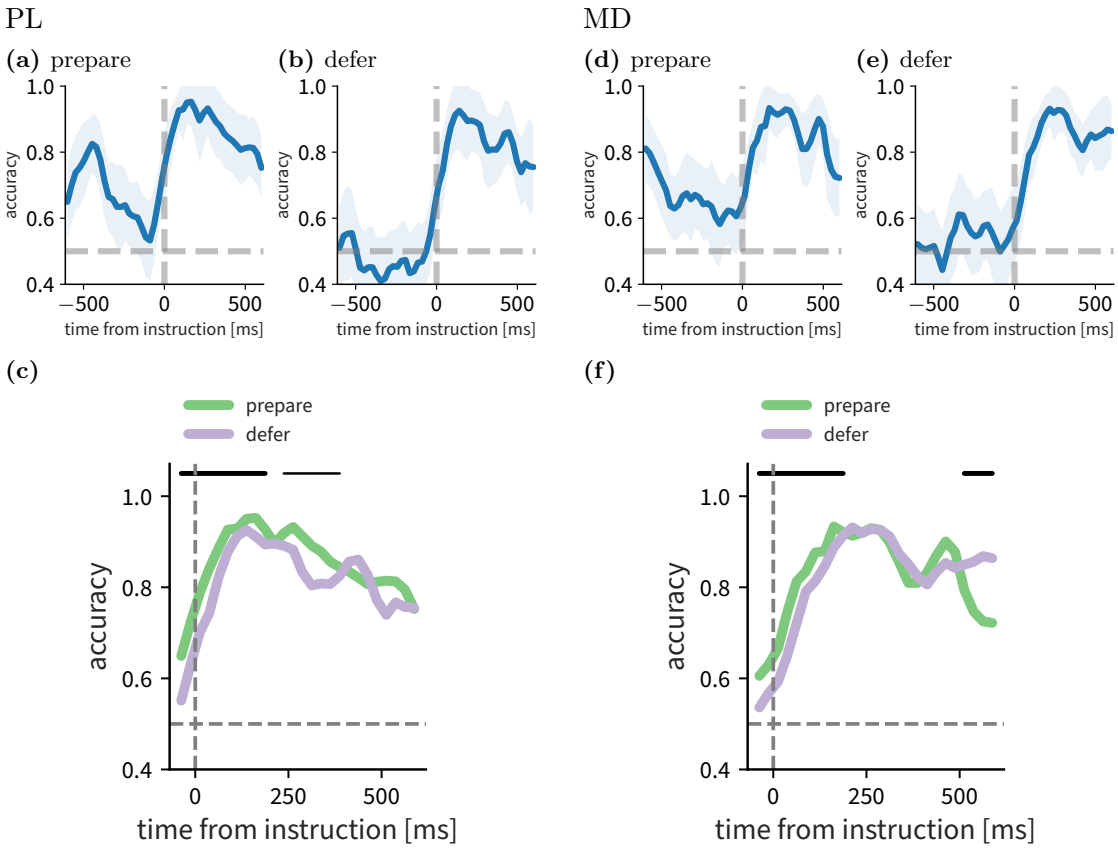

Figure 3.44: Effect of prior context on population coding of movement instruction. **(a)** Classification accuracies of SVM classifiers in *prepare* trials trained on the neuronal activity of the PL pseudopopulation and tested at the same points in time, aligned to onset of movement instruction. Saturated line: Mean across pseudotrial shuffles ($n = 100$). Blue shading: Standard deviation across pseudotrial shuffles. Vertical dashed line: Time of instruction stimulus presentation. Horizontal dashed line: Chance level. **(b)** same as (a) for *defer* trials. **(c)** Comparison of classifier accuracies from *prepare* (a) and *defer* (b). Cluster-corrected permutation test, horizontal bars: thick $p < 0.01$, thin $p < 0.05$. Included in the comparison were times −50 ms to 600 ms from instruction to avoid spurious clusters from the context epoch. **(d–f)** same as (a–c) for the MD pseudopopulation.

and became more similar. After a period of high code similarity from 100 ms to 300 ms the hyperplane angles increased again.

Since the above suggests that neuronal codes for instruction were conserved across contexts it should be possible to use a classifier trained in one context to decode movement instruction in trials with the other prior context. Indeed Figures 3.45b, 3.45c, 3.45e and 3.45f show that cross-contextual decoding (solid lines) was almost as good as using intra-contextual classifiers (dashed lines, data from Figs. 3.44c and 3.44f). Notably, in PL the classifier trained on *defer* trials was worse than the intra-context classifier in the first 100 ms of *prepare* trials[84] (Fig. 3.45b). The opposite way, decoding instruction in *defer* trials using a classifier trained on *prepare* trials, achieved accuracies that were indistinguishable from the intra-context classifier during that time interval (Fig. 3.45c). Interestingly, this "uni-directional" effect was absent in MD. Here, even though a delayed coding onset had been detected in *defer* trials (Fig. 3.44f), during the first 200 ms the *defer* classifier achieved similar decoding accuracies in *prepare* trials like the intra-classifier and vice-versa (Figs. 3.45e and 3.45f). Strong differences between the accuracies were only present in *defer* trials *after* 200 ms, when the intra-classifier (trained on *defer* trials) decoded instruction significantly better than the cross-classifier.

In sum, this shows that the population code for movement instruction was strongly conserved across trials with different prior contexts. This conservation also extended to the temporal evolution of code.

3.2.4.7 Neuronal coding was degraded in wrong trials

So far, Sections 3.2.4.1 to 3.2.4.6 showed data from correct trials. If that neuronal information was relevant for the animals' behaviour, it should correlate with it. Because of the animals' very good task performance (Fig. 3.3) only a small number of error trials in which the response ball was turned to the wrong side was available. In order to have comparable statistical power, I therefore replaced correct trials with the available error trials before training the decoders (see e.g. Jacob et al., 2018; Siegel et al., 2009). Due to the high number of correct trials that remain part of the trial set this serves as a lower bound on the effect of error trials.

Figures 3.46a and 3.46b show that in both PL and MD context information was consistently significantly higher in correct trials than in datasets substituted with error trials.[85]

Similarly, instruction information after instruction onset was consistently higher in datasets containing only correct trials (Fig. 3.47). Notably, the degradation effect appeared to be stronger in trials with prior *defer* context (Figs. 3.47b and 3.47d) than in those with prior *prepare* context (Figs. 3.47a and 3.47c).

Together this shows that response errors correlated with degraded coding of the relevant variables. Information seemed to be more robust in trials with prior *prepare* context.

[84] Cluster-corrected permutation tests $p < 0.01$
[85] Cluster-corrected permutation tests $p < 0.01$

PL

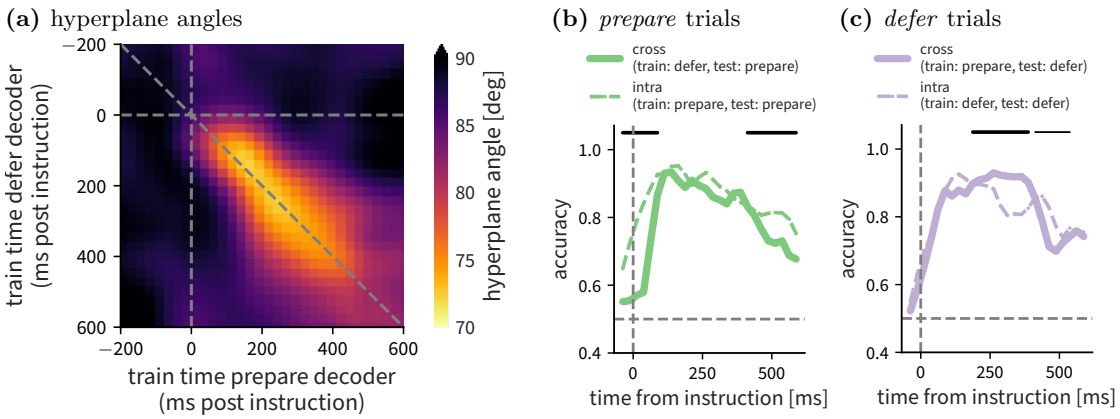

MD

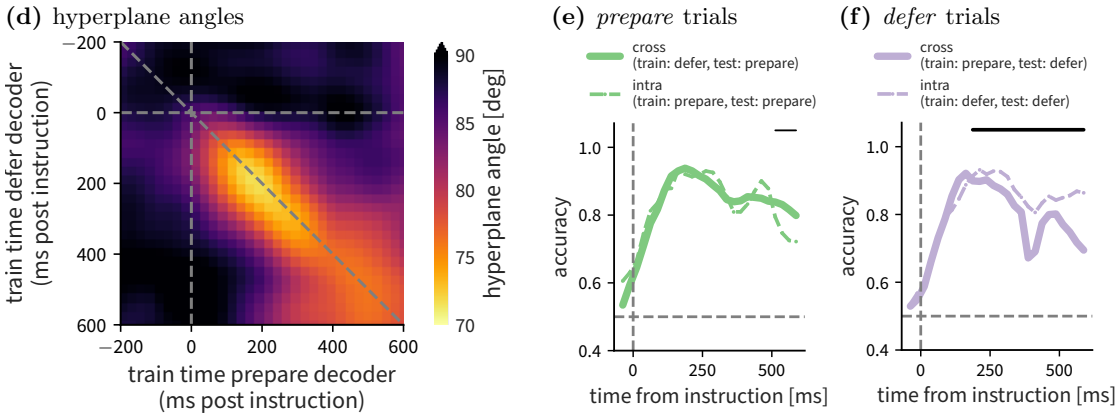

Figure 3.45: Stability of movement instruction coding across prior contexts. (a) Angles between hyperplanes of classifiers trained on PL population activity in *defer* trials vs. *prepare* trials as a proxy of code similarity. Warmer colours indicate higher similarity. Hyperplanes were not only compared for classifiers that were trained at the same time (main diagonal dashed line) but also at different times (off-diagonal pixels). Horizontal and vertical dashed lines: instruction onset. **(d)** same as (a) for MD population. **(b)** Decoder accuracies in PL for instruction (*go-left* vs. *go-right*) for held-out *prepare* test trials. Dashed time-course: Decoder was trained and tested on *prepare* trials (intra-context). Solid time-course: Decoder was trained on *defer* and tested on *prepare* trials (cross-context). **(c)** Same as (b) but tested on held-out *defer* trials. Dashed and solid time courses are intra- and cross-context decoders, thus identities are flipped. **(e, f)** Same as (b, c) but for MD. Vertical dashed line: instruction onset. Horizontal dashed line: chance level. Horizontal bars above line plots: significant clusters ($p < 0.01$, cluster-corrected permutation test).

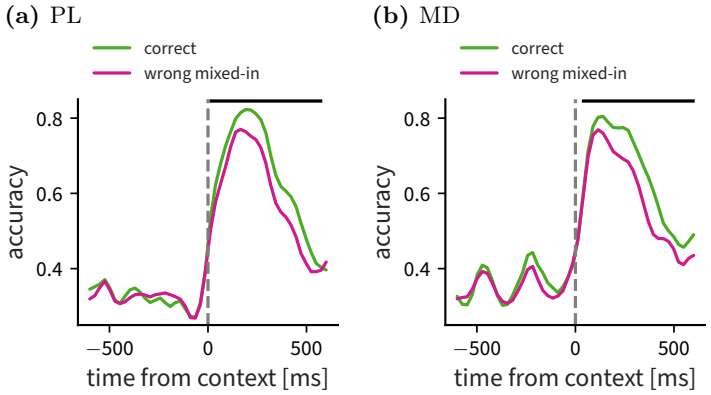

Figure 3.46: Degraded context information in error trials. (a) Decoding accuracies in PL for correct trials (green) and after replacing correct trials with error trials (magenta, "wrong mixed-in"), aligned to context cue onset. **(b)** same as (a) for MD. Horizontal bars: Significant differences between correct and error-trial-substituted data ($p < 0.01$, cluster-corrected permutation test). Only data from after the context cue onset was used for statistical tests.

3.2.4.8 Higher and more generalisable context information in trials with covert rightwards movements

Although the context epoch was enforced to be a time during which the animals should not make large response ball movements, the dissection of the behaviour in that epoch (0 ms to 600 ms after context cue onset) showed that micro-movements within the acceptable range of rotation did occur (Section 3.1.3).

In order to test if covert movements during the context epoch correlated with neuronal coding, in the following I investigated classifier accuracies for trials with relatively more left ("leftish") and right ("rightish") submovements (compare Section 3.1.3). Importantly, the subsetting of relative movement tendency was done per context, so that in each trial subset the same number of trials was present for each context.

Figures 3.48a, 3.48b, 3.49a and 3.49b show that indeed context information differed dependent on response ball movement. If animals covertly moved the ball relatively more to the left side, as was their more common behaviour (Figs. 3.3 and 3.12), context information was high in the beginning of the context epoch but quickly degraded and did not generalise well across time (Figs. 3.48a and 3.49a). In contrast, context information was higher, longer lasting and more temporally generalisable, especially after 200 ms into the epoch, if animals moved the response ball relatively more to the right (Figs. 3.48b and 3.49b).

Together, this shows that covert right-wards movements occurred especially if context information was high.

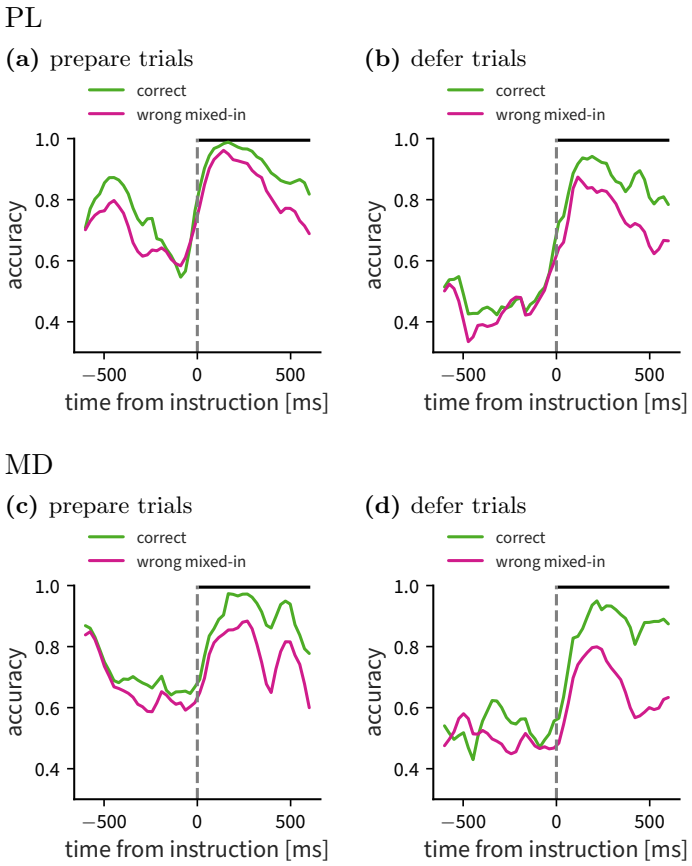

Figure 3.47: Degraded instruction information in error trials. (a) Decoding accuracies in *prepare* trials in PL for correct trials (green) and after replacing correct trials with error trials (magenta, "wrong mixed-in") aligned to movement instruction onset. **(b)** same as (a) for *defer* trials. **(c, d)** same as (a, b) for MD. Horizontal bars: Significant differences between correct and error-trial-substituted data ($p < 0.01$, cluster-corrected permutation test). Only data from after the instruction was used for statistical tests.

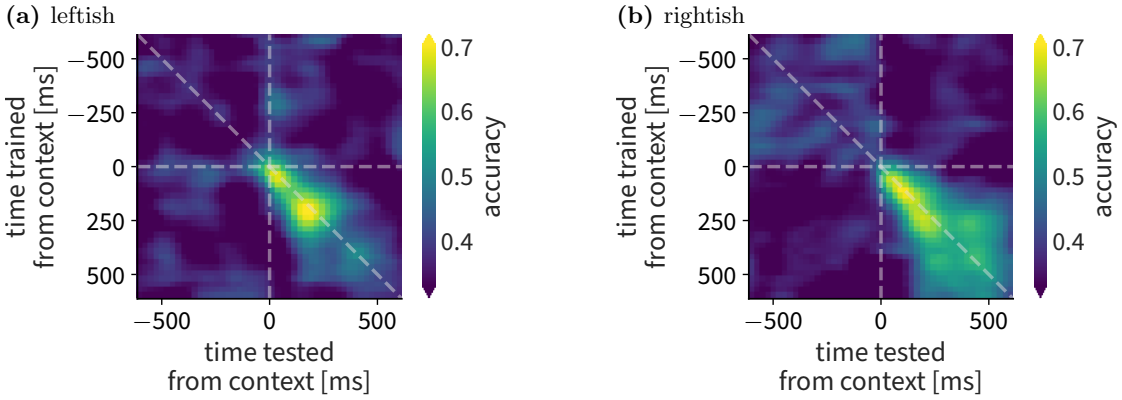

Figure 3.48: Neuronal population coding of context information dependent on covert movement tendency (PL). Cross-temporal classifier accuracies for context from activity in PL pseudopopulation in trials with relatively more covert left-wards (a) and right-wards (b) movements (0 ms to 600 ms after context cue onset).

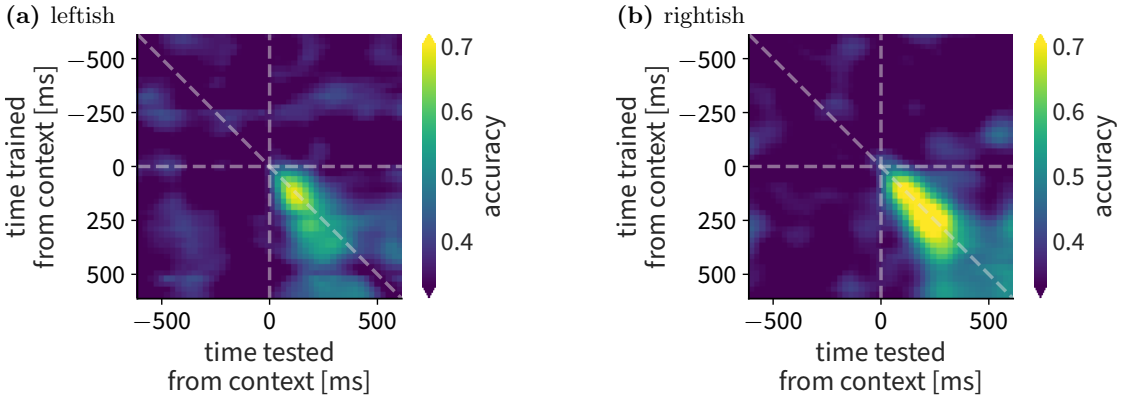

Figure 3.49: Neuronal population coding of context information dependent on covert movement tendency (MD). Cross-temporal classifier accuracies context from activity in for MD pseudopopulation in trials with relatively more covert left-wards (a) and right-wards (b) movements (0 ms to 600 ms after context cue onset).

In sum, Sections 3.2.4.7 and 3.2.4.8 showed that both context and movement instruction could be robustly decoded from the population activity in PL and MD. This coding was relevant for successful behaviour. The temporal evolution of population code for movement instruction was conserved across trials with different prior contexts.

Table 3.8: Significant GC links during the context epoch (0 ms to 600 ms).

link direction		$n_{significant}/n_{total}$ (% total)		
source	target	*defer*	*prepare-left*	*prepare-right*
PL	PL	2701/14 728 (18.3)	1642/11 442 (14.4)	1469/9192 (16.0)
MD	MD	1218/12 858 (9.5)	593/8684 (6.8)	897/11 084 (8.1)
PL	MD	283/3457 (8.2)	156/2654 (5.9)	180/2640 (6.8)
MD	PL	271/3457 (7.8)	110/2654 (4.1)	193/2640 (7.3)

3.2.5 Functional connectivity

Many analyses shown in the previous sections assumed more or less independent neurons that, as a whole, give rise to different phenomena. Obviously, neurons within areas are actually tightly linked into neuronal systems and their communication, which can vary with task requirements, is of major importance for proper brain functions. In addition to within-area communication, PL and MD are well known to be anatomically densely interconnected and functionally dependent (Fuster, 2015). Thus, it is very likely that functional connectivity within and across these areas varied with task requirements.

To gain insight into this topic, therefore, I employed the analytical framework of GC (Barnett & Seth, 2015). Briefly, GC assumes that a cause A precedes the effect B and contains unique information about it. Spike trains of pairs of simultaneously recorded single units within and across regions were thus tested for significant functional links.

3.2.5.1 Fractions of significant functional connections were consistently ordered by context

In the context epoch from 0 ms to 600 ms after context stimulus onset the fraction of significant GC links was asserted. Fractions of significant functional links varied significantly with context, i.e. fractions were not equal (Table 3.9). The highest fraction of links was found in the context epoch of trials with *defer* context, followed by *prepare-right* and, lastly and leastly, *prepare-left* (Fig. 3.50 and Table 3.8). Although this order was found in all directions, it was only statistically significant amongst the links *within* PL or MD (Table 3.9, source PL, target PL; source MD, target MD), not fronto-thalamic or thalamo-frontal (Table 3.9, source PL, target MD; source MD, target PL).

Intra-regional links were more common than cross-regional ones, with PL and MD ranging from 14.35 % to 18.34 % and 6.8 % to 9.47 %, respectively, while across they ranged from 5.88 % to 8.19 % and 4.14 % to 7.84 %, for fronto-thalamic and thalamo-frontal, respectively (Table 3.8).

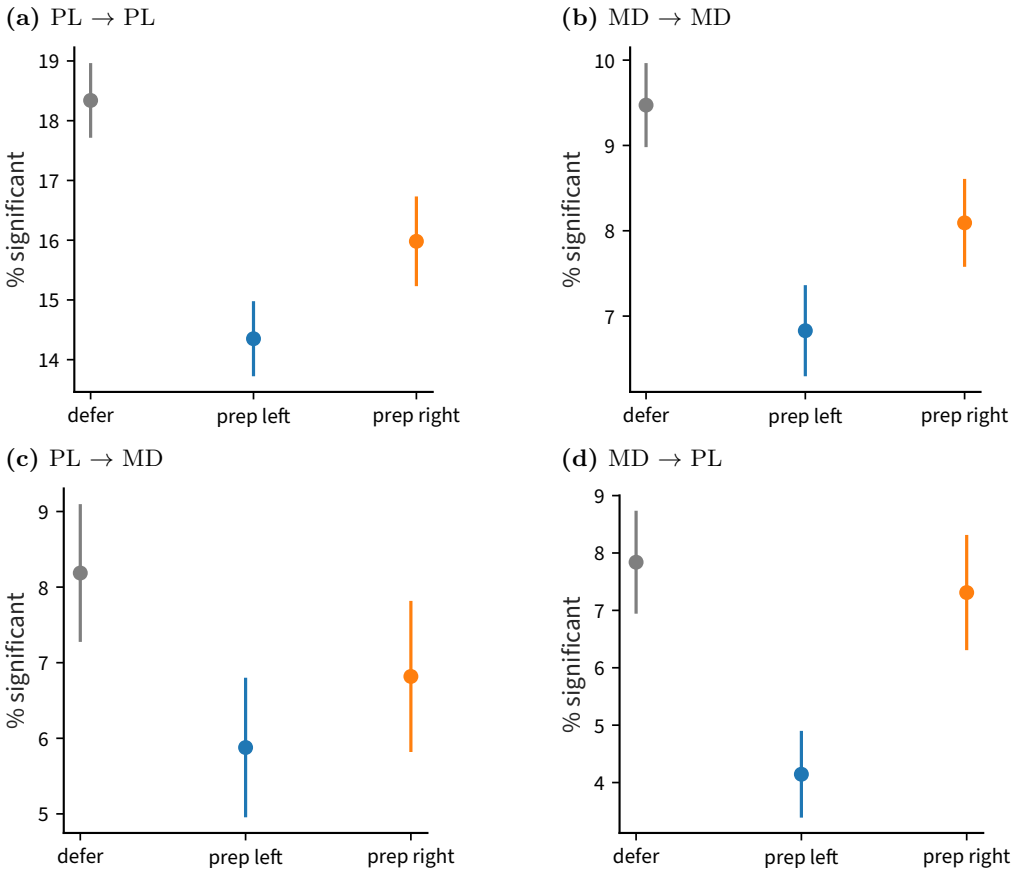

Figure 3.50: Significant GC links during context epoch. **(a)** Percentage of significant functional links between source cells in PL and target cells in PL. **(b–d)** same as (a) but for MD to MD, PL to MD and MD to PL, respectively. Vertical error bars show 95 % bootstrap CI.

Table 3.9: χ^2 contingency tests for significant GC links during context epoch.

source	target	comparison	χ^2	p	df
PL	PL	all	54.900	$1.198 \cdot 10^{-12}$	2
		defer/prepare-left	52.946	$3.429 \cdot 10^{-13}$	1
		defer/prepare-right	15.308	$9.132 \cdot 10^{-5}$	1
		prepare-left/prepare-right	7.692	$5.545 \cdot 10^{-3}$	1
MD	MD	all	41.194	$1.135 \cdot 10^{-9}$	2
		defer/prepare-left	39.640	$3.053 \cdot 10^{-10}$	1
		defer/prepare-right	11.646	$6.434 \cdot 10^{-4}$	1
		prepare-left/prepare-right	9.446	$2.116 \cdot 10^{-3}$	1
PL	MD	all	10.887	0.004	2
		defer/prepare-left	10.101	0.001	1
		defer/prepare-right	3.257	0.071	1
		prepare-left/prepare-right	1.589	0.208	1
MD	PL	all	32.775	$7.640 \cdot 10^{-8}$	2
		defer/prepare-left	30.498	$3.341 \cdot 10^{-8}$	1
		defer/prepare-right	0.444	$5.053 \cdot 10^{-1}$	1
		prepare-left/prepare-right	21.376	$3.775 \cdot 10^{-6}$	1

3.2.5.2 Functional links during instruction epoch varied with prior context within PL but were context-invariant within MD

Similar to the context epoch, during the instruction epoch the highest fractions of links were found again in PL (Table 3.10 and Fig. 3.51). In contrast to the context epoch, the proportions of GC-significant links during the instruction epoch (0 ms to 300 ms from instruction) varied with link direction (Fig. 3.51a).

Within PL the highest fraction of significant links were in the *go-left* instruction after the *prepare* context. The *go-right* instruction after the *prepare* condition, in contrast, had the fewest significant links. Finally, if an instruction had been preceded by a *defer* context the fractions of links were between those extrema. Furthermore, in those trials there was no difference between the instructions (Table 3.11). Thus, the fraction of functional links within PL was dependent on which of the three contexts had been shown.

Conversely, within MD the fraction of significant links only varied with instruction, while the prior context did not have an effect (Fig. 3.51b and Table 3.11). More GC links were found in *go-left* trials. Looking across regions the fraction of significant links ranged from 6.32 % to 8.10 % and 6.16 % to 7.12 %, for PL to MD and MD to PL, respectively (Figs. 3.51c and 3.51d). The proportions of cross-regional links did not vary significantly with context or instruction (Table 3.11).

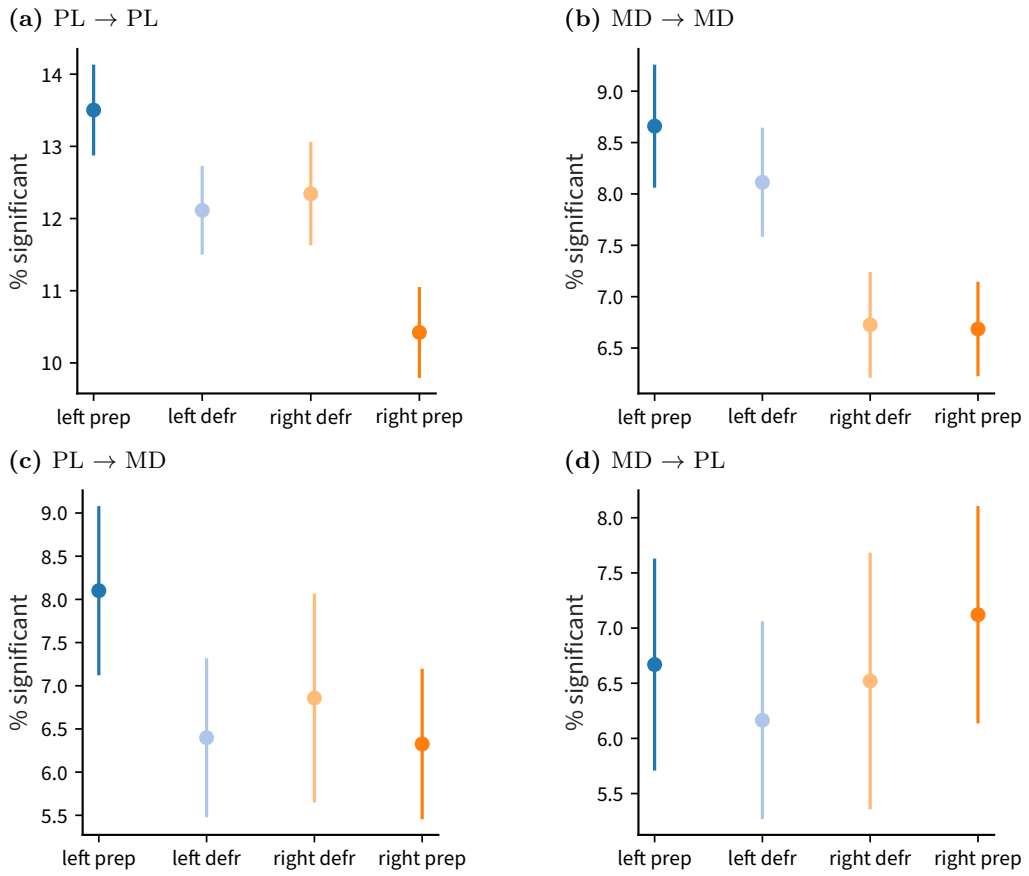

Figure 3.51: Significant GC links during instruction epoch. (a) Percentage of significant functional links between source cells in PL and target cells in PL. **(b–d)** Same as (a) but for MD to MD, PL to MD and MD to PL, respectively. Vertical error bars show 95 % bootstrap CI.

Table 3.10: Significant GC links during the instruction epoch (0 ms to 300 ms).

link direction		$n_{significant}/n_{total}$ (% total)			
source	target	*go-left defer*	*go-left prepare*	*go-right defer*	*go-right prepare*
PL	PL	1293/10 674 (12.1)	1545/11 442 (13.5)	1065/8628 (12.3)	958/9192 (10.4)
MD	MD	764/9416 (8.1)	752/8684 (8.7)	643/9560 (6.7)	741/11 084 (6.7)
PL	MD	164/2563 (6.4)	215/2654 (8.1)	122/1779 (6.9)	167/2640 (6.3)
MD	PL	158/2563 (6.2)	177/2654 (6.7)	116/1779 (6.5)	188/2640 (7.1)

Table 3.11: χ^2 contingency tests for significant GC links during instruction epoch.

source	target	comparison	χ^2	p	df
PL	PL	all	35.824	$8.159 \cdot 10^{-8}$	3
		go-left: *prepare*/*defer*	7.258	$7.060 \cdot 10^{-3}$	1
		go-right: *prepare*/*defer*	12.823	$3.424 \cdot 10^{-4}$	1
		prepare: *go-left*/*go-right*	35.426	$2.649 \cdot 10^{-9}$	1
		defer: *go-left*/*go-right*	0.166	$6.840 \cdot 10^{-1}$	1
		go-left/*go-right* (pooled)	15.732	$7.299 \cdot 10^{-5}$	1
		defer/*prepare* (pooled)	0.047	$8.283 \cdot 10^{-1}$	1
MD	MD	all	35.104	$1.158 \cdot 10^{-7}$	3
		go-left: *prepare*/*defer*	1.418	$2.338 \cdot 10^{-1}$	1
		go-right: *prepare*/*defer*	0.007	$9.355 \cdot 10^{-1}$	1
		prepare: *go-left*/*go-right*	23.054	$1.575 \cdot 10^{-6}$	1
		defer: *go-left*/*go-right*	11.286	$7.808 \cdot 10^{-4}$	1
		go-left/*go-right* (pooled)	33.240	$8.148 \cdot 10^{-9}$	1
		defer/*prepare* (pooled)	0.211	$6.459 \cdot 10^{-1}$	1
PL	MD	all	7.157	0.067	3
		go-left: *prepare*/*defer*	4.618	0.032	1
		go-right: *prepare*/*defer*	0.354	0.552	1
		prepare: *go-left*/*go-right*	5.149	0.023	1
		defer: *go-left*/*go-right*	0.249	0.618	1
		go-left/*go-right* (pooled)	1.593	0.207	1
		defer/*prepare* (pooled)	1.184	0.276	1
MD	PL	all	1.721	0.632	3
		go-left: *prepare*/*defer*	0.410	0.522	1
		go-right: *prepare*/*defer*	0.438	0.508	1
		prepare: *go-left*/*go-right*	0.304	0.581	1
		defer: *go-left*/*go-right*	0.146	0.703	1
		go-left/*go-right* (pooled)	0.642	0.423	1
		defer/*prepare* (pooled)	1.065	0.302	1

In sum, this analysis shows two trends. First, functional connectivity during the context epoch consistently varied with the context shown, a pattern that is reminiscent of responsiveness (Table 3.5). Second, connectivity within PL was context-dependent, while in MD it was context-invariant and instruction-dependent. Control analyses that equalised trial counts across conditions (stratification) yielded qualitatively similar results (Figs. S7 and S8).

Part III

Discussion

4 Structuring of Abstract Working Memory Content by Fronto-Parietal Synchrony in Primate Cortex

The presented data demonstrate that local field potentials (LFPs) in the prefrontal and parietal cortices can be involved in the maintenance, long-range transmission and selection of abstract working memory information. Distinct frequency bands showed functional specialisation and served as direction-specific communication channels. A phase-dependent code could be used to disentangle simultaneously encoded sample and distractor numerosities from prefrontal spiking activity.

4.1 Working memory dependent balance of beta and gamma power

Recently, beta and gamma oscillations in prefrontal cortex have been suggested to reflect working memory maintenance and retrieval, and stimulus encoding, respectively (Lundqvist et al., 2016). Similarly, in the presented behavioural task power in these bands was anti-correlated (Figs. 2.2c and 2.2d). During task epochs of putative information encoding, gamma power was up-regulated and beta power down-regulated. Conversely, during the memory delays, the balance of these bands was reversed and trial-to-trial fluctuations in the beta band carried information about the sample and distractor numerosities (Fig. 2.3), consistent with a role of the beta frequencies in maintenance and retrieval of working memory content.

4.2 Dominant parieto-frontal (feedforward) communication in the beta band

The data presented here extend those findings to those frequency bands' roles in long-range communication within the wider fronto-parietal network. Gamma band oscillations were found to be only involved in local signalling. This is consistent with their short cycle periods which might preclude the simultaneous recruitment of neuronal populations that are located further apart (Harmony, 2013; but see Liebe et al., 2012, for long-range synchrony in PFC and V4).

Beta oscillations, on the other hand, established a sustained parieto-frontal feedforward communication channel (Figs. 2.4, 2.5, 2.6b, 2.6d and 2.7). Strong beta band synchrony between prefrontal and parietal cortices was found in earlier studies using

behavioural tasks that required the maintenance of a single item in working memory (Antzoulatos et al., 2016; Salazar et al., 2012). Similar to the latter study, I found increasing phase-synchrony and parieto-frontally directed communication towards the end of the last memory delay, temporally close to the upcoming test stimulus (Figs. 2.6b and 2.6d). Together this suggests that not only local increases in beta oscillations, but synchronisation of inter-regional beta phases is conducive to working memory retrieval, extending Lundqvist et al. (2016)'s hypothesis.

4.3 Delay epoch prefrontal-to-parietal (feedback) communication via low frequency oscillations

The lower frequency delta and theta bands were prominently present during the memory delays. Power was strongly elevated in those epochs and carried ordered information about the relevant sample numerosity (Figs. 2.2 and 2.3). Furthermore, these bands served as the primary transmission channel for feedback information from PFC to VIP (Figs. 2.4 to 2.7). Together, this suggests a role of these bands in maintenance and/or retrieval of working memory information. In a prior analysis of this dataset it was shown that information about the sample numerosity contained within spiking activity increased towards the end of the memory delays (Jacob & Nieder, 2014). The primarily fronto-parietal communication found in the present analyses now suggests that prefrontal sample information could have been transmitted via the delta/theta band and thus be responsible for a large amount of information in VIP. It's possible that mechanistically this represents a re-entry loop that continuously reactivates numerosity representations in VIP, sustained by driving input from PFC (Bodner et al., 2005; X.-J. Wang, 2006). If distractors were omitted, frontal drive still stopped after memory delay 1 at the low frequencies, which is an example of perceptual set in VIP in expectation of a numerosity display that could have been caused by PFC at the end of the first memory delay. Interestingly, when distractors where omitted this channel also enabled parieto-frontal (feedforward) along with strong fronto-parietal communication (Figs. 2.6c and 2.6d), thus demonstrating functionality that can flexibly adapt to different task demands.

A recent study found that VIP neurons represent empty sets categorically different from non-zero numerosities (Ramirez-Cardenas et al., 2016). Although the authors of that study didn't explicitly test for it, it is possible that, in addition, different neuronal sub-populations in VIP are activated. In contrast to that study, an empty set was not explicitly signalled to the animals in the present task. However, the fixed time frame of the task and, thus, the expectation of seeing a set of dots make the omission of a distractor numerosity comparable to that study's empty set. It is thus likely that VIP neurons transmitted information about the presence of an empty set to prefrontal cortex. The presence of this categorically different signal might have elicited a strong feedback drive from prefrontal cortex, "pushing" relevant sample numerosity information to VIP.

4.4 Delta/theta phase separates multiplexed parallel information

Previously, it was shown that during the second memory delay of this task information about both sample and distractor numerosities increased to equal levels in PFC. In VIP, on the other hand, information about the sample during that epoch was remarkably higher than about the distractor, suggesting a process that filtered out irrelevant information (Jacob & Nieder, 2014).

The analyses presented here provide a possible mechanism for this by leveraging local field potentials. During this task epoch, sample and distractor information from prefrontal neurons showed strong dependence on the phase of parietal delta/theta oscillations (Fig. 2.11). Interestingly, phase-dependency was low in the first memory delay, suggesting that the fronto-parietal network flexibly adapts the coding mechanism according to task demands. Different levels of neuronal excitability in VIP caused by LFPs (Gupta et al., 2016) could serve as a reading frame enabling the selective readout of relevant information at an "optimal" phase.

This mechanism was previously proposed in a task involving two behaviourally relevant ordered working memory items (Siegel et al., 2009). There, phase-dependent information within PFC was prominent in the delta and beta bands and working memory items could be distinguished by phase only in the latter. In the present task phase-dependency of sample and distractor information was also found intra-PFC for the delta/theta bands but only for the distractor in the beta band. Different from the task in Siegel et al. (2009) our task is critically dependent on VIP's role in numerosity processing and, arguably, its connectivity with PFC. Long-range feedback communication was prominent in the lower frequencies, which might have required a shift of phase-dependent coding to that band.

Possibly due to their prominent association with non-REM slow-wave sleep (Léger et al., 2018), and intractability in fast-paced behavioural tasks, little is known about the role of delta oscillations in cognition. Studies have shown functional synchronisation of prefrontal and parietal delta rhythms in decision-making tasks involving the comparison of a sample memory item with a test stimulus (Antzoulatos et al., 2016; Nácher et al., 2013). More generally, prefrontal delta oscillations have been hypothesised to be involved in inhibition of interfering motor behaviours and inappropriate neural activity via "functional deafferentiation" (Harmony, 2013). While this mechanism might appear inconsistent with prominent synchronisation in the task presented here, delta oscillations might be helpful in the isolation of the entire fronto-parietal network involved in numerosity processing. However, while clear phase-dependency was found in both delta and theta bands, only phase-dependency of the sample in the theta band was predictive of task performance. Together with the response-acceleration for higher theta band synchrony during this epoch, this suggests a critical role for theta signalling, similar to a previous study investigating connectivity of PFC and V4 (Liebe et al., 2012). Consistent with this, theta phase has been suggested to represent a frame to store sequentially presented items (Voytek et al., 2015; Heusser et al., 2016). More generally, theta-band

connectivity of the prefrontal and parietal cortices is associated with cognitive control, goal-directed attention and working memory maintenance and manipulation (Cooper et al., 2015; Fellrath et al., 2016; Sarnthein et al., 1998; Sauseng et al., 2004; Sauseng et al., 2005; Schack et al., 2005).

Interestingly, error trials were associated with lower phase-dependent coding of the distractor in the alpha band (8 Hz to 12 Hz). Alpha oscillations have been implied in the selective suppression of to-be-ignored items (Noonan et al., 2018; Schroeder et al., 2018; Sghirripa et al., 2021; van Zoest et al., 2021; Wöstmann et al., 2019). The presented data suggest a possible mechanism for this: Broadcasting the distractor numerosity in this band might help to guard working memory by for example preparing downstream circuits.

4.5 Direction-specific communication via separate frequency bands

Oscillations in the low frequency ranges (delta, theta) were largely dominated by feedback (fronto-parietal) communication, while the higher frequency alpha/beta bands subserved feedforward communication from VIP to PFC. According to studies made in primate visual cortical areas feedforward communication predominantly occurs in the gamma band while feedback communication is associated with the beta band (Bastos et al., 2015; Michalareas et al., 2016; van Kerkoerle et al., 2014). While this general concept of higher frequency bottom-up and lower frequency top-down communication is consistent with the presented data, a general shift to lower frequencies was observed for communication between PFC and VIP. Lower frequency oscillations with their larger temporal windows by cycle might be required to synchronise brain regions that are farther apart than various visual areas (Harmony, 2013). In addition, higher-cognitive areas might in general use lower frequencies for inter-regional communication, as demonstrated by other studies (Johnson et al., 2017; Phillips et al., 2014). Future work should investigate if this shift to lower frequencies is a general feature of long-range communication in the fronto-parietal network.

4.6 Outlook

Many studies have demonstrated flexible frequency-specific modulation and inter-regional synchronisation of local field potentials in a multitude of cognitive tasks (Liebe et al., 2012; Siegel et al., 2009; Buschman et al., 2012; Brincat & Miller, 2015; Buschman & Miller, 2007). Extending those observational studies, in recent years an increasing number of studies was published that investigated mechanistic roles of LFPs in cognition, for example by computational modelling (Sherfey et al., 2020) or exogenously establishing or disrupting inter-regional phase synchronisation (Polanía et al., 2012). In addition to the notion that perception and memory are "more than spikes" (Watrous et al., 2015), local field potentials are more than spatial and temporal averages, but have specific features in distinct cortical layers (Bastos et al., 2018), travel spatially across areas (Zhang

et al., 2018), differentially engage distinct cell-types (Tseng & Han, 2019), and are actually short-lived, distinct, and heterogeneous events (Lundqvist et al., 2016; E. K. Miller et al., 2018; Cole & Voytek, 2017). Future work should therefore investigate how these features are related to the findings in the present study.

5 Neuronal Signatures of Contextual Decision-Making in Mouse Prefrontal Cortex and Mediodorsal Thalamus

This project investigated executive functions in mice at a finer temporal scale than traditional studies that employ simple binary responses. To this end, I examined the influence of external contexts on the concrete expression of goal-directed behaviour and neuronal representations in PL and MD, two brain areas that are strongly interconnected and have been associated with executive control.

In an auditory decision-making task mice successfully learned to use an abstract goal-directed behaviour to receive reward. Prior informative and non-informative context cues were incorporated into the animals' behaviour. Leveraging continuous read-out of a behavioural variable, the behavioural effects of contexts could be analysed at sub-trial precision, showing distinct motor implementations.

Extracellular electrophysiological recordings in PL and MD during the behavioural task showed neuronal signatures of preparation and default-overriding signals in activity profiles, responsiveness and preference, neuronal recruitment and population coding.

5.1 Response abstractness via the lack of feedback coupling

In this study a ping-pong ball that was fixed in one axis was used as a response device. Mice successfully learned to rotate the ball left or right in response to auditory movement instructions in order to receive water reward, thus demonstrating goal-directed behaviour (Fig. 3.2). The use of a directional response as behavioural paradigm is advantageous to a go/no-go task because it avoids the inherent asymmetries of action suppression and execution (Carandini & Churchland, 2013). This particular response device was chosen because it enabled the continuous read-out of a behavioural variable, thus extending the possibility to dissect responses. Similar to the original study by Sanders and Kepecs (2012) that introduced this response device, animals in this study did not receive any feedback regarding ball position while they rotated the response ball. This makes the behaviour they exhibited quite abstract because mice have to learn that the cumulative signed rotation, i.e. all movements, early and late, influence trial outcome. While the abstractness of the response likely limited the number of animals that could be included in the task (Section 5.10), it is possible that it ensured involvement of higher-order cognitive brain areas even after extensive training (Fuster, 2015). Nevertheless, a study that coupled the response ball's rotation to visual feedback reported partial

dependence on mouse PFC to inhibit incorrect choices (Huda et al., 2018). Due to the decreased training time when using direct feedback (Section 5.10), it is likely that a possible outsourcing of cognitive processes to lower-order brain areas was not yet completed in that study.

5.2 Multiple levels of behaviour affected by context cues

In the behavioural task that was used in this project animals were instructed about the required movement side by one of two auditory movement instructions. Prior to the instructions one of three context cues was presented. Two of these cues were predictive of the upcoming movement instruction, whereas one was non-predictive (Fig. 3.2). Because no comparison with the context cues was required the task can also be solved by solely attending to the movement instructions. However, the animals used the context cues to their advantage: task performance (Fig. 3.3) as well as response times (Fig. 3.4) were improved in trials with a *prepare* context. This is in line with previous research that usually shows better performance and response times if facilitatory, i.e. predictive, cues are presented (Blaukopf & DiGirolamo, 2005; Hoffmann & Sebald, 2005; Lim et al., 2019; Moon et al., 2016). However, mice have also been shown to sporadically ignore facilitatory cues (Toptaş et al., 2022).

The use of the continuous data of the response ball rotation provided a richer insight into behaviour. Small movements during the context epoch that were consistent with the upcoming movement instruction improved task performance (Fig. 3.7). This is reminiscent of effects in primate research using smooth visual pursuit tasks: incompatible (i.e. inconsistent with the target) anticipatory eye movements usually lead to lower task performance (de Hemptinne et al., 2006).

Response ball rotation data could be further unpacked by decomposing movement into kinematic primitives, called submovements (Fig. 3.8). Quantity of submovements during the context epoch was shown to be modulated in a way that was consistent with predictive cues (Fig. 3.12). Similarly, strength and duration of consistent movements were upregulated in those trials (Figs. 3.14 and 3.15). This is further evidence for the successful use of the context cues and suggests a rehearsal-like behaviour during the context epoch. Again, this is similar to oculovisual smooth pursuit tasks: cuing the direction of upcoming movement leads to anticipatory eye movements in the cued direction, scaled by predictiveness of the cue (Kowler et al., 2019; Santos & Kowler, 2017).

Interestingly, context effects were not observed on submovements during the instruction epoch (Figs. 3.13, 3.16 and 3.17). Similarly, response execution times were invariant to prior context (Fig. 3.10b). This suggests that the exact motor implementations of the responses were very stable. It is possible that this is an effect of the relatively late introduction of context cues during training. By that time the learned motor responses might be consolidated and resistant to change. Another possibility is that the reported metrics are not suitable to uncover context effects during the instruction epoch. Further

analyses could investigate motor implementations at an even finer scale, for example by investigating sequences (motifs) of different types of submovements.

5.3 Distinct motor implementations and *go-left* as a default behaviour

Despite the expected equivalence of responses to the two sides a recurring finding were disparities in the two conditions. Task performance was higher and responses were faster in *go-left* than in *go-right* trials (Figs. 3.3 and 3.4). Consistent with this, the beneficial effect of a *prepare* context on task performance was more variable in *go-right* trials (Fig. 3.3b) and absent for response times in *go-left* trials (Fig. 3.4b). This suggests that left responses were more optimised and less pliable than right responses.

The trend was also apparent on a finer scale when investigating the response ball rotation. Not only the compound response times were faster in *go-left* trials, but correct left responses were initiated earlier and executed faster than correct right responses (Fig. 3.10). Furthermore, consistent "anticipatory" movements during the context epoch increased the probability of a correct trial in all but the "*go-left-prepare-left* " trials (Fig. 3.7a), again supporting the hypothesis that *go-left* responses were optimised.

Moreover, motor implementations in the context epoch were distinct. Informative *prepare* context cues shifted the balance of left and right submovements towards the cued side (Fig. 3.12a), similar to studies in primates that show a reduction of incompatible anticipatory movements in direction-cued trials (de Hemptinne et al., 2006). Notably, in trials with the non-informative context cue the balance of submovements was left-shifted (Fig. 3.12a). Additionally, the effect of *prepare* context in pushing the balance to the cued side was weaker for the *prepare-left* than for the *prepare-right* trials. Together, this suggests that in *defer* trials there was a left-default that was overridden by the *prepare-right* context cue.

Further unpacking of the submovement balances revealed that context effects only applied to submovements to the right (Fig. 3.12c), whereas the number of left submovements was invariant to trial condition (Fig. 3.12b). Combined with a similar effect for the submovement strength (Figs. 3.14a and 3.14b) this suggests that the motor implementation for left movements was more robust and less pliable. This might also explain why the responses in *go-right* trials were slower: If left submovements were the default, additional right submovements would increase the path length towards a successful threshold crossing.

Why was there such a left-default in the behaviour in the first place? Direction biases in mouse behaviour have been reported in earlier studies and often behavioural training involves strategies that are aimed at reducing them (Huda et al., 2018; The International Brain Laboratory et al., 2021; Timothy et al., 2021). The small number of response sides and external stimuli in the current task might have favoured a strategy in which a default behaviour was established. If required by the current trial this default might be overridden in a graded way. This might be more economical and easier to implement for animals with limited cognitive capacities like mice. For example, lower neuronal activ-

ity is observed in mouse prefrontal cortex for habitual than for goal-directed behaviour (Barker et al., 2018). Furthermore, a strategy with an implicit default might increase detectability of the context cue for the preparation of the "non-default" behaviour, because mice respond well to "unexpected" stimuli (Rogalla et al., 2020). Thus, if the *prepare-left* context would be considered the "expected" stimulus the decreased detectability would need to be offset by a higher baseline probability of executing left behaviour.

However, those economical and ecological considerations don't explain the consistent effect of a *left* default behaviour. One explanation could be intrinsic characteristics of the batch of animals that were used. Mice are known to have strong paw lateralisation, similar to handedness in humans (Manns et al., 2021). In contrast to humans, however, who show an overall prevalence of around 90 % right-handedness (Papadatou-Pastou et al., 2020), no population asymmetries have been observed in mice. This even holds for strongly inbred and laterality-selected mice, precluding a population asymmetry in the small in-house breeding facility from which our mice originated (Collins, 1985; Manns et al., 2021; Signore, Chaoui, et al., 1991; Signore, Nosten-Bertrand, et al., 1991). Despite those arguments on the population-level, it is still possible that the small sample in this study had a strong lateralisation asymmetry (binomial probability $p_{x=5} = 0.03125$). If that was the case, responses to one side might have been easier to execute for those animals, similar to how humans show better motor performance if using their preferred hand.

Other possible explanations of behavioural left-preference are based in methodology. First, the auditory stimuli that were used as context cues and movement instructions might have been perceived with different salience (Section 5.10). One reason for different perceptability could be the age-related hearing loss of high frequency auditory stimuli in C57BL/6 mice (Ison et al., 2007; Walton et al., 1995).

In addition to stimulus properties, asymmetries in the behavioural hardware setup and training might have played a role in establishing a preference for left movements. For practical reasons, the setup was accessible to the experimenter from the left side. This might have introduced a behavioural asymmetry caused by favouring the less "disruptive" interior of the setup that was on the animals' right.

Finally, because the training regime started with large blocks of repeated movements to one side, movements to the left might have been more extensively practised in the beginning of training. Easy-to-implement controls for those issues would be rotation of the setup and randomising movement sides from the very beginning of the training.

5.4 Processing of context and instruction in PL and MD

Extracellular recordings of neuronal activity were made in PL and MD while the animals performed the behavioural task. The majority of included single neuronal units in both brain regions exhibited activity that was modulated by one of the contexts or movement instructions (Fig. 3.25). On a population level both regions strongly encoded both context and instruction (Figs. 3.39 and 3.42). Furthermore, population coding deteriorated in incorrect trials (Figs. 3.46 and 3.47). Together this suggests that both PL and MD

were important for the processing of those signals and that the neuronal representation contributed to successful behaviour. This is consistent with a growing body of literature that shows that variables related to goal-directed behaviour are encoded in rodent PFC (Le Merre et al., 2021; Pinto & Dan, 2015; Rikhye et al., 2018) and MD (Courtiol & Wilson, 2016; Kawagoe et al., 2007; R. L. A. Miller et al., 2017; Mukherjee et al., 2021; Rikhye et al., 2018).

5.5 Behavioural importance of instruction reflected by neuronal representation

Although both the context and movement instruction were represented neuronally, the dominant factor seemed to be the movement instruction. More single units' activity was modulated by the instruction than by the context (Fig. 3.25), firing rate was about twice as high for instruction than for context (Figs. 3.21 and 3.23), and the population encoded it virtually perfectly (Fig. 3.42).

Why did the instruction consistently have a higher neuronal impact? A possibility is the difference in the two stimulus types' behavioural importance. The *prepare* contexts could be used to solve the task without listening to the instruction. However, they only appeared in half of the trials, which made them a less reliable source of information for goal-directed behaviour. Furthermore, the contexts were more temporally detached from the actual required behaviour, whereas the instruction was a signal to take a concrete and immediate action.

Nevertheless, two alternative explanations cannot be excluded. Since the the context cues were introduced relatively late in the behavioural training the different amounts of neural representation could be a function of time (see Section 5.10). Other than that, the use of a single set of instruction stimuli did not allow to distinguish between cognitive rules and movement-related activity. Indeed, it has recently been shown that mouse cortical activity, including PFC, is dominated by movements (Musall et al., 2019). The authors of that study found that especially uninstructed movements, i.e. those that had not been reinforced by training, resulted in strong widespread cortical signals. In contrast, instructed movements had very weak contributions. However, in their task the instructed movements were directional licks, while whole-limb movements were an example of uninstructed ones. Arguably, tongue movements are smaller and smoother and thus have smaller contributions to cortical activity than movements of entire body parts (Bollu et al., 2021, compare low number of tongue-movement selective units). In contrast, the instructed movements in the behavioural task presented in this thesis required the use of at least one forelimb in order to sufficiently rotate the response ball. Often, mice would not only move their forelimbs but use their entire body as a counterweight for stabilisation (data not shown). Thus, our instructed movement might have indeed dominated the neuronal signal, explaining the stronger representation of the instruction.

For these reasons, since the neuronal signal was likely a mixture of go-cue, the stimulus properties themselves, the animal's choice and its movement, it would be interesting

for follow-up projects to investigate the individual contributions. Possible methods for decoupling would be generalised linear models (Goltstein et al., 2021; Park et al., 2014), demixed principal component analysis (Kobak et al., 2016), or vector orthogonalisation (Inagaki et al., 2022).

5.6 Neuronal signatures of uncertainty and action preparation

Contexts were distinctly neuronally signalled. Signalling of the *defer* context was strikingly different from the *prepare* contexts. For instance, peak latencies were shortest (Figs. 3.22a and 3.22b) and neuronal recruitment fastest (Figs. 3.33a and 3.33b) in *defer* trials, while for the two *prepare* conditions these metrics were similar, suggesting a categorical difference in the processing of those cues. Furthermore, for *defer* condition peak firing rates were highest (Fig. 3.21), most context-responsive neurons showed responsivity (Figs. 3.26a and 3.26b) and an even higher fraction preferred this condition (Figs. 3.31a and 3.31c).

Why was this non-informative cue represented so strongly? It is possible that this is a way that uncertainty was explicitly signalled, which might have had top-down preparatory effects on sensory systems to improve sensory discrimination of the upcoming action-determining instruction (i.e. *perceptual set*, Fuster, 2015). A common finding in humans is that prefrontal neural activity increases with increasing predictive uncertainty (Catena et al., 2012; Huettel et al., 2005; Kéri et al., 2004). Similarly, mice mPFC has been shown to shape downstream neural responses specifically in tasks involving uncertainty (Hamm et al., 2021; Starkweather et al., 2018).

The two *prepare* contexts are in opposition to the *defer* context: since they allow for the preparation of an upcoming action they don't signal uncertainty. This might explain why activity was lower, fewer neuron were responsive and preferred those contexts. However, since they signal the correct action it is possible that PFC readied motor structures for the upcoming response.

Thus, for future studies using this task it would be interesting to first, modulate behavioural uncertainty to probe graded activation of PFC, and second, record from or silence prefrontal input to possible effector regions in order to investigate preparatory top-down effects, e.g. in auditory cortex in uncertain trials and premotor cortex in *prepare* trials.

5.7 Neuronal representation of default behaviour

The two *prepare* contexts had similar neuronal recruitment dynamics and peak latencies (Figs. 3.22 and 3.33). However, although *prepare-left* and *prepare-right* contexts could have been expected to have equivalent behavioural relevance, namely preparation to move to one or the other direction, striking asymmetries were found in other metrics. Many more single units were responsive to and preferred the *prepare-right* over the

prepare-left context (Figs. 3.31a and 3.31c). Similarly, average firing activity was very low for *prepare-left* (Fig. 3.21).

This suggests a differential processing of the *prepare* contexts. Apart from distinct sensory properties of the stimuli (Fig. 7.2), this asymmetry might be best viewed in the light of the observed behaviour. As described above, the animals showed behaviour that suggested left-moving as a default behaviour. It is known that mice react especially well to unexpected stimuli (Rogalla et al., 2020) and that mouse prefrontal cortex shows reduced activity during default behaviour (Barker et al., 2018) as well as down-regulates neuronal responses to expected stimuli in sensory cortex (Hamm et al., 2021). Thus, if left-behaviour was considered the default for the animals, it is possible that the context cue that signalled this default as the upcoming required action was only weakly encoded neuronally. In contrast, the context cue that prepared for the right-movement action would be "surprising" and would also require a behavioural override signal. This might be represented by the high neuronal activity and number of responsive neurons. Strong override signals to abandon current default behaviour have been found previously in neuronal subtypes of mouse PFC in a foraging task (Kvitsiani et al., 2013; Pearson et al., 2014). Reducing activity to cues for default behaviour and only signalling overrides might be an economical way of encoding different behavioural options in task with low degrees of freedom.

At first glance this interpretation might be in conflict with the findings for instructions. If the default left-behaviour does not need to be signalled strongly, why was there dominant neuronal responsivity and preference for the *go-left* over the *go-right* instruction (Figs. 3.31b and 3.31d)? One possibility is that the low representation of *prepare-left* is a compensation of the strong *go-left* representation. Another interpretation of this discrepancy is the superposition of different events during the instruction epoch. This epoch simultaneously involved the resolution of previous inhibition of large movements, the behavioural rule of the required motor action and its immediate execution. Future experiments should therefore try to more clearly disambiguate these events.

In any case, the default behaviour and signalling interpretation raises the question of what came first: was the behavioural strategy first implemented and the neuronal signals followed or was the observed left default a consequence of the distinct signalling? Following the evolution of neuronal signals during learning of the task might shed more light on this question.

5.8 The effect of contexts on instruction signalling

The neuronal signalling of movement instruction was partially dependent on the prior context. Not all neuronal units that were responsive to an instruction were so after both *prepare* and *defer* contexts. Rather, populations of context-stable and context-exclusive cells were observed (Figs. 3.28a to 3.28d). This suggests that, dependent on prior context, the same instruction was processed by different sets of cells: One context-invariant "core" population and distinct additional, smaller subpopulations that were only responsive after one or the other context. Notably, the context-exclusive sub-populations were of

comparable sizes. This might be surprising considering the predictive nature of the *prepare* contexts, after which one might expect more cells that were responsive to the cued instruction. However, it is consistent with the observation that on a population level none of the contexts biased responsiveness to higher magnitudes (Figs. 3.29 and 3.30).

Consistent with the lack of a contextual responsiveness bias, recruitment profiles for context-stable cells were not affected by prior context (Fig. 3.35), indicating that the speed of activation of individual cells was similar. Previously, it was hypothesised that task-relevant information is handed over by sequential sparse activation (also referred to as "tiling") of mouse prefrontal neurons (Bolkan et al., 2017; Parnaudeau et al., 2018; Schmitt et al., 2017). Thus, the recruitment speed of context-stable cells was suggestive of a similar effect. Notably, however, the points in time when individual cells were activated were not invariant to prior context (Fig. 3.36), dismissing the idea of stable information hand-over. It is possible that this randomness in activation delay was introduced by the additional recruitment of the distinct subpopulations of context-exclusive cells, the inclusion of which could have led to distinct sequences. In addition, the instruction epoch did not involve any behavioural delay, whereas tiling had previously been shown during task epochs without sensory stimulation.

Although the prior contexts activated partially distinct subsets of cells the populational coding of instruction was exceptionally good in both *prepare* and *defer* contexts (Fig. 3.44). Indeed the only difference was that early in the instruction epoch the coding was slightly worse in *defer* than in *prepare* contexts. This could be explained by residual coding that was carried over from the late context epoch. While the population code of contexts and instructions was different, as evidenced by the lack of cross-epoch code stability (Figs. 3.43a and 3.43c), it is possible that a small number of cells exhibited similar firing activity during both epochs.

Furthermore, across contexts the population code for instruction was remarkably stable (Fig. 3.45). This might be surprising considering the random activation delays of individual context-stable cells, as described above. However, in recent years a more integrated view has highlighted the importance of neuronal networks in coding over individual neurons. Thus, it is possible that different configurations of individual neurons encoded instruction in stable, low-dimensional subspaces (Gallego et al., 2017; Gallego et al., 2018; P. Gao et al., 2017; Murray et al., 2017; Recanatesi et al., 2022; Rule et al., 2020). In future analyses it would be interesting to see if the contextual subpopulations have distinct contributions to coding.

5.9 Cross-regional timing differences suggestive of ordered interplay

Neuronal responses in PL and MD exhibited overall similarities. Both regions displayed comparable variations in responsiveness and preference towards individual contexts and instructions. Furthermore, the influence of prior contexts on instruction responsiveness and coding showed similar patterns. These cross-regional consistencies are in line with a recent study that showed that more than 90 % of MD cells exhibited activity that

was consistent with those of mPFC neurons, i.e. MD cells "mirrored" PFC cells, in a spatially-guided delayed response task in rats (R. L. A. Miller et al., 2017).

Other than the mentioned consistencies I found cross-regional timing differences. Peaks for the *defer* context were earlier in PL than MD and the converse was found for the *prepare* contexts (Fig. 3.22c). This further supports the interpretation that the *prepare* contexts were processed as one category distinct from the *defer* context. Similarly, the lead of PL in *defer* context suggests that there might have been a top-down signal for the inhibition of motor preparation. The clear lead of MD for the *prepare* contexts, on the other hand, suggests a primary role of MD in the preparation of the action plans. This is further corroborated in population coding: The response side to-be-prepared was encoded earlier in MD than PL. MD also sustained this information longer than PL, up until the onset of the instruction (Fig. 3.40). This putative higher involvement of MD in action planning might also carry over into the instruction epoch: Intra-regional Granger-causality links in MD were only modulated by instruction, while in PL they were dependent on context.

Together, these observations are reminiscent of previous findings. For example, two studies from the Kellendonk lab showed that in a delayed response task in mice beta oscillations in mouse MD lead those in PFC specifically in the delay epoch (Bolkan et al., 2017; Parnaudeau et al., 2013), suggesting a role in the coding of upcoming action plans. Similarly, a study in monkeys that performed a conceptually similar task found that MD neurons contributed more and earlier to processing of prospective information (Y. Watanabe & Funahashi, 2012; Y. Watanabe et al., 2009). Finally, the afore-mentioned study by R. L. A. Miller et al. (2017) that found MD neurons mirroring PFC neurons also found a higher proportion of MD cells modulated by upcoming motor responses.

5.10 Limitations and recommendations for future studies

Like every study this project has several limitations (Ross & Bibler Zaidi, 2019), some of which have briefly been alluded to before.

The data and interpretation is purely descriptive because no manipulations were made. Therefore, future studies should use drug interventions and/or targeted optogenetic manipulation of neuronal activity to gain mechanistic insights.

Behavioural training, especially initial instrumental learning, required a large amount of time. This is most probably a consequence of using the ball rotations as responses. The lack of continuous feedback with regards to goal-progression introduced a high level of abstraction. This abstraction delayed or even prevented animals from learning the responses. Consequentially, the animals' ages might not have been optimal for cognitive tasks. Furthermore, the implant for head-fixation might have led to abnormal skull growth during this extended period of time, thus reducing target accuracy of neuronal recordings. A possible improvement in training time could be achieved by the introduction of direct coupling of ball rotations with an external stimulus. Indeed, more recent studies that use a variant of the response ball use visual feedback coupled to the instantaneous response device rotation and stress its importance for robust and re-

producible training of high numbers of mice (Aoki et al., 2017; Huda et al., 2018; The International Brain Laboratory et al., 2021). Similar advantages of coupling of response device rotation with external visual, auditory and multimodal stimuli in contrast to its omission have been observed in a comparative study from our laboratory (Ranganath et al., 2022). A different approach that also reduces response abstractness is the use of a more concrete response device. For example, Bollu et al. (2018) used a paw-operated response *joystick* that has "built-in" feedback, via both applied force and absolute, i.e. directly observable, device position during the response.

During the majority of the progression of behavioural training only movement instruction stimuli were presented. In contrast, context cues were introduced only after animals had very high response accuracy in discrimination of movement instructions. This means that instructions had a longer time to impact behaviour and neuronal implementations than contexts. In addition, because contexts were not required to solve the behavioural task, their possible impact was reduced further. Other than changing the behavioural paradigm it could be advantageous to introduce context cues earlier in training.

The choice of external stimuli might have had unanticipated effects on the animals' behaviour and might have limited the interpretability of the data. Age-related hearing loss of the C57BL6 mouse strain primarily affects the perception of high-frequency sounds (Ison et al., 2007; Walton et al., 1995). Thus, highpass-filtered white noise sound for *prepare-right* and high-starting-frequency sinusoidal downsweeps for *go-right* might have been perceived less clearly and/or later. However, neuronal representation of contexts is converse to perception bias expected from high-frequency hearing loss: *prepare-right* context (orange lines in Fig. 3.21), which was signalled by supposedly badly perceptible highpass filtered white noise sounds (Fig. 7.2c), was represented strongly, whereas *prepare-left* context, (blue lines in Fig. 3.21), which was signalled by supposedly well-perceived lowpass filtered white noise sounds (Fig. 7.2a), was weakly represented. Furthermore, similar peak latencies (Fig. 3.21) suggest that the two *prepare* cues were perceived at the same times, further excluding sensory effects. With regards to movement instruction the beginning frequencies of the signalling sounds were only spaced apart by a few kHz (Figs. 7.2d and 7.2e), thus probably reducing the influence of selective hearing loss. Furthermore, the cited studies investigated hearing *thresholds*, meaning they determined at what low sound pressure level (SPL) animals could not *detect* the sound. The auditory stimuli used in this project, on the other hand, were presented at much higher SPLs to ensure robust detection.

Nevertheless, with the current choice of a single set of stimulus-response contingencies it is not possible to determine if observed effects were dependent on stimulus features or the rules encoded by them. Follow-up studies should therefore randomise stimulus-response contingencies and/or introduce more than one stimulus per rule, possibly in a different sensory modality, to separate cognitive from sensory variables.

As described above, the "blind" ball rotations that were required as responses increased complexity for the animals. On the other hand, only binary choices, i.e. whether the left or the right threshold was exceeded, were tested by the behavioural paradigm. These binary choices could have also been tested using more concrete behaviours, like directional licking (Z. Gao et al., 2018; Inagaki et al., 2022) or lever pressing (Caran-

dini & Churchland, 2013). One argument for using the response ball is the continuous readout of a behavioural variable, thus possibly giving insight into an evolving decision. However, modern analyses of videos of behaving animals can provide similar or richer insights, identifying distinct behavioural states (Graving & Couzin, 2020) and movement progressions (Bollu et al., 2021) in behaviours that might be more natural for mice. Even so, the response ball could be used for other types of tasks that require non-binary responses, like parameterised accuracy tests. Because of the abstractness of the responses using the response ball, it might still be valuable to instead use more concrete response devices, like paw-operated choice-sticks (Bollu et al., 2018) that also increase response data to two dimensions. In any case, combination with recording of behavioural videos should be considered in those tasks, because of the low additional effort and high potential return.

In addition to the general behavioural paradigm more concrete methodological challenges of this project have to be considered.

The processing of the continuous behavioural data was impacted by transformations of ball rotation data. Concretely, raw digital data was transformed into an analog signal which was then used as an input to the measurement devices that digitised the signal for use in the computer (Fig. 7.1). The device that did the digital-to-analog transformation required a "resetting" of the analog signal to baseline levels at task epoch boundaries and during times of enforced calm periods. During the resetting no behavioural data was recorded, leading to gaps in the data that needed to be inferred by interpolation using Gaussian processes (Section 7.6.3). In addition, the software that was used for orchestrating the presentation of task-related stimuli limited the recording of behavioural data to the trial itself, thus ignoring behaviour in inter-trial-intervals. These discontinuities precluded analysis methods that require long continuous data, like certain GLMs or VAE-SNE. Future projects should avoid discontinuities and conversion noise by directly connecting raw digital sensor signals to computer peripherals and should use more recent versions of the behaviour orchestration software that allow continuous recording across trials.

With regards to the recording of neurophysiological signals the choice of acute extracellular recordings should be challenged. First, movements might have deteriorated signal. Even though the animal's skull was fixed in the setup, body movements led to brain tissue movements. Because the neural probes were more stationary the movements might have caused additional noise in the recordings that would make analysis of local field potentials difficult. Furthermore, a well-known problem in acute extracellular recordings is the neural drift (Hill et al., 2011). Small movements of the surrounding brain tissue lead to different relative positions of neurons to the recording electrodes, thus changing the signal. This is typically observed in the first minutes after the initial insertion of a neural probe into the brain tissue. To ease this, one needs to wait for the brain tissue to "settle" in order to receive stable neural signals. Unfortunately, this waiting period had a high impact on our animal's motivation, thus decreasing data yield. Furthermore, possibly due to large body movements we still observed neuronal unit drift, so that most units were not observed for the entirety of a session. This also resulted in a low number of simultaneously recorded units, which precluded more elaborate analyses

of functional connectivity, especially across regions. Second, acute recordings required to first train the animals on the behavioural task and then perform surgery for craniotomy. All surgical procedures entail the risk of losing animals due to complications or at the least damaging brain tissue. Furthermore, several days of recovery are required after a surgery with general anaesthesia. During this time, animals received water *ad libitum* and were not trained in the task. Consequently, on the first days of recording, animals were only weakly motivated to perform the task and behavioural yield was low.

Due to these disadvantages of acute recordings it should be considered to instead perform chronic implantation of neural probes in the initial surgery before training. This would make a second surgery obsolete, make initial targeting of the probes more accurate, probably lead to more stable recordings, and would allow for interruption-free behaviour and even the investigation of learning effects. Potential downsides would be the increased effort in implantation and *post-mortem* extraction of the neural probe, and build-up of neural scar tissue that decreases signal. However, the latter could be alleviated by using movable microdrives (Wimmer et al., 2015).

To sum up, I recommend follow-up projects to implement more behavioural controls, more fully embrace the strengths of the response device in accuracy-testing paradigms, and perform chronic neural recordings to increase signal stability.

5.11 Concluding summary

All in all, this project showed both behavioural and neuronal signatures of action preparation and its inhibition. To the best of my knowledge, this is the first study that dissected with sub-trial resolution the continuous behavioural data that is recorded with the kind of response device that was used. The combination of this finer-grained behavioural data with neuronal data suggested a cognitive strategy that involves an implicit default response that can be overridden. Many similarities between neuronal data from PL and MD suggest a similar involvement of the two brain areas in solving the behavioural task. Furthermore, cross-regional timing differences suggest cross-regional collaboration and higher contribution of MD in action preparation. Future experiments are needed to investigate if similar cognitive strategies are typical of mice performing behavioural tasks with low degrees of freedom and if the observed neuronal signatures are a mechanism of implementing those strategies.

Part IV

Methods

6 Monkey Project[1]

6.1 Experimental model and subject details

Two adult male rhesus monkeys (Macaca mulatta, monkey R and monkey W, 12 and 13 years old) were used in this study. All experimental procedures were in accordance with the guidelines for animal experimentation approved by the national authority, the Regierungspräsidium Tübingen.

6.2 Method details

6.2.1 Surgical procedures

Monkeys were implanted with two right-hemispheric recording chambers centered over the principal sulcus of the lateral prefrontal cortex (PFC) and the ventral intraparietal cortex (VIP) in the fundus of the intraparietal sulcus (IPS).

6.2.2 Task and stimuli

A detailed description of the monkeys' task and behavioral performance is provided elsewhere (Jacob & Nieder, 2014). The animals grabbed a bar to initiate a trial and maintained eye fixation (ISCAN, Woburn, MA) within 1.75° of visual angle of a central white dot. Stimuli were presented on a centrally placed gray circular background subtending 5.4° of visual angle. Following a 500 ms pre-sample (pure fixation) period, a 500 ms sample stimulus containing 1 to 4 dots was shown. The monkeys had to memorize the sample numerosity for 2500 ms and compare it to the number of dots (1 to 4) presented in a 1000 ms test stimulus. Test stimuli were marked by a red ring surrounding the background circle. If the numerosities matched (50 % of trials), the animals released the bar (correct Match trial). If the numerosities were different (50 % of trials), the animals continued to hold the bar until the matching number was presented in the subsequent image (correct Non-match trial). Match and non-match trials were pseudo-randomly intermixed. Correct trials were rewarded with a drop of water. In 80 % of trials, a 500 ms interfering numerosity of equal numerical range was presented between the sample and test stimulus. The interfering numerosity was not systematically related to either the sample or test numerosity and therefore not useful for solving the task. In 20 % of trials, a 500 ms gray background circle without dots was presented instead of an interfering stimulus, i.e. trial length remained constant (control condition, blank). Trials with and without interfering numerosities were pseudo-randomly intermixed. Stimulus

[1] This methods chapter was modified from Jacob et al. (2018).

presentation was balanced: a given sample was followed by all interfering numerosities with equal frequency, and vice versa. Throughout the monkeys' training on the distractor task, there was never a condition where a stimulus appearing at the time of the distractor was task-relevant.

Low-level, non-numerical visual features could not systematically influence task performance (Jacob & Nieder, 2014; Nieder et al., 2002): in half of the trials, dot diameters were selected at random. In the other half, dot density and total occupied area were equated across stimuli. CORTEX software (NIMH, Bethesda, MD) was used for experimental control and behavioral data acquisition. New stimuli were generated before each recording session to ensure that the animals did not memorize stimulus sequences.

6.2.3 Electrophysiology

Up to eight 1 MΩ glass-insulated tungsten electrodes (Alpha Omega, Israel) per chamber and session were acutely inserted through an intact dura with 1 mm spacing. To access VIP, electrodes were passed along the course of the intraparietal sulcus to a depth of 9 mm to 13 mm below the cortical surface (Jacob & Nieder, 2014; Nieder et al., 2006; Ramirez-Cardenas et al., 2016). Correct positioning of the electrodes in VIP was verified by physiological criteria (responses to moving visual stimuli and tactile stimulation). A total of 616 PFC sites (monkey R: 368, monkey W: 248) and 614 VIP sites (monkey R: 376, monkey W: 238) was recorded from. Extracellular neuronal signals were acquired with a unity-gain headstage and hardware bandpass-filtered into spiking activity (0.1 kHz to 8 kHz, sampling rate 40 kHz) and local field potentials (LFPs) (0.7 Hz to 170 Hz, sampling rate 1 kHz). Single units were recorded at random; no attempt was made to preselect for particular response properties. Signal amplification, filtering and digitalization were accomplished with the MAP system (Plexon, Dallas, TX). Waveform separation was performed offline (Plexon Offline Sorter).

6.3 Quantification and statistical analysis

Data analysis was performed with MATLAB (Mathworks, Natick, MA) using custom scripts, the FieldTrip toolbox (Oostenveld et al., 2011) and the CircStat toolbox (Berens, 2009). Figures were replotted for this thesis using matplotlib 3.4.2 (Caswell et al., 2021) in python 3.8.10 (Van Rossum & Drake, 2009). To account for neuronal response latencies in prefrontal and parietal cortex (Jacob & Nieder, 2014), the starting and end points of all trial epochs were offset by 100 ms for analysis.

6.3.1 Preprocessing

Single units were included in the analysis if the following criteria were met: (*a*) their average firing rate across trials was at least 1 Hz; (*b*) they were recorded for at least one correct trial in all 20 conditions (4 sample numerosities x 5 interfering numerosities including the control [0] condition); (*c*) they modulated their firing rate in the course of the trial (task-related neurons, one-way ANOVA with average firing rates in the

presample [fixation], sample, first memory, interfering stimulus, and second memory periods; evaluated at $p < 0.05$).

LFP traces were mean-centered, filtered for line noise removal (4^{th} order Butterworth notch at 50 Hz and first and second harmonics), and re-referenced to the average of all prefrontal and parietal electrodes within a session, unless stated otherwise. Signals phase-locked to stimulus presentation (i.e. event-related potentialss (ERPs)) were removed by subtracting the average across trials from individual trials, unless stated otherwise. ERP subtraction was performed separately for all analyzed trial subsets.

6.3.2 Spectral transformation

Complex time-frequency representations X of single LFP trials were estimated by convolution of signal x with complex kernels k:

$$X(t, \omega) = x(t) * k(t, \omega) \tag{6.1}$$

where t and ω are time and frequency, respectively, $*$ is the convolution operator and k represents frequency-dependent Hanning-tapered complex sinusoids:

$$k(t, \omega) = A \left(1 - \cos(\frac{2\pi t\omega}{q})\right) e^{2i\pi t\omega} \tag{6.2}$$

where A is a constant normalizing k to unit power, q is the kernel width in number of cycles. Except for the analysis of Wiener-Granger Causality (see below), which required a linear frequency axis for algorithmic reasons, frequencies ranging from 2 Hz to 128 Hz, logarithmically spaced in steps of $2^{1/8}$, and kernel widths of 3 were used. For spike-LFP measures complex Fourier coefficients were obtained from X at the individual spike times.

6.3.3 Power

The time-varying power pow_x of signal x at frequency ω was computed as the squared norm of its time-frequency transformation:

$$pow_x(t, \omega) = |(X(t, \omega))|^2 \tag{6.3}$$

Power was averaged across trials, sessions and electrodes and z-scored to baseline (500 ms pre-sample epoch):

$$z_{pow_x}(t, \omega) = \frac{pow_x(t, \omega) - \mu_{BL}(\omega)}{s_{BL}(\omega)} \tag{6.4}$$

where $\mu_{BL}(\omega)$ and $s_{BL}(\omega)$ are the mean and standard deviation, respectively, of $pow_x(t, \omega)$ during baseline at frequency ω.

6.3.4 Phase

The time-varying phase φ_x of signal x at frequency ω is the argument of the complex Fourier coefficients X:

$$\varphi_x(t, \omega) = \arg(X(t, \omega)) \tag{6.5}$$

6.3.5 Cross-spectrum

The complex time- and frequency dependent cross-spectrum crs_{xy} of signals x and y is the product of their conjugated spectral decompositions X and Y:

$$crs_{xy}(t,\omega) = X(t,\omega)Y(t,\omega)^* \tag{6.6}$$

where * is the complex conjugate.

6.3.6 Phase Locking Value

The frequency-dependent PLV was computed as the norm of the average across observed phases:

$$PLV(\omega) = \left| \frac{1}{R} \sum_{r=1}^{R} e^{i\varphi(r,\omega)} \right| \tag{6.7}$$

where $\varphi(r,\omega)$ is the cross-channel phase $\arg(crs_{xy}(\omega))$ at trial r and frequency ω (LFP-LFP PLV) or the spike phase of spike r at frequency ω (spike-LFP PLV), and R is the number of observations (trials or spikes, respectively). Because the PLV is biased towards 1 for low sample sizes the minimum number of observations was set to 50.

Spike-field PLVs for individual units were computed for all available unit-LFP pairs, averaged across pairs and then z-scored using the mean and standard deviation of a null distribution obtained by randomly shuffling the association of single-trial spike train and corresponding LFP trace within a session before spike phase estimation ($n = 1000$) (Buschman et al., 2012). In an additional analysis to control for effects of varying LFP power on the robustness of spike phase estimation, spikes were stratified according to their associated LFP magnitude. At both 5 Hz and 20 Hz, histograms of the spike-associated LFP magnitudes were computed in the two conditions that were compared (e.g. first and second memory delay). This yielded a ratio of spike counts at every magnitude bin, which were used as probabilities for random subsampling of spikes. A spike was included in the PLV analysis if a sample drawn from a uniform distribution on $[0, 1]$ was less than or equal to the magnitude-associated probability.

For the PLV analysis in fast and slow trials, correct match trials were separated into quartiles based on reaction times (non-match trials were not included because the second test image following the non-match was always a match and therefore predictable). Trial types (no distractor, repeat sample, true distractor) were matched across subsets, i.e. a given condition appeared equally often in the first and fourth quartile. ERPs were computed and subtracted separately for each subset.

6.3.7 Phase Locking Selectivity Index

To determine whether neuronal synchrony was stimulus-specific, a PLSI was computed, based on an approach reported in (Salazar et al., 2012), exchanging coherence for PLV. PLSI is a mutual information measure and was computed as a function of time and

frequency:

$$PLSI(t,f) = \sum_r \sum_s P(s)P(r|s)\log_2 \frac{P(r|s)}{P(r)} \tag{6.8}$$

with

$$P(r) = \sum_s P(r|s)P(s) \tag{6.9}$$

where $P(r)$, $P(s)$ and $P(r|s)$ are the probabilities of the response (i.e. PLVs), the stimulus (i.e. numerosities) and the conditional probability of the response given the stimulus (i.e. PLVs at individual numerosities). Since the PLV is an average measure across trials, $P(r|s)$ was estimated in increments of 0.01 on $[0,1]$ using a normal distribution with mean $\mu = PLV_s$ and standard deviation σ estimated using a jackknife procedure:

$$\sigma = \sqrt{\frac{n-1}{n}\sum_{t=1}^{n}(x_t - \bar{x})^2} \tag{6.10}$$

with n being the number of trials for a given stimulus, x_t the jackknife PLV with trial t left out and $\bar{x}$ the mean across jackknife samples. For statistical assessment a semi-generalized surrogate distribution of the $PLSI_{surr}$ was computed. At each of 1000 iterations condition labels were randomly shuffled across trials and PLSI was computed for one randomly chosen combination of channels. To mitigate underestimation of the bias of PLSI due to the use of a semi-generalized procedure, PLSI values were corrected by a factor α which sets the average PLSI across frequency and time bins during the baseline to 0:

$$PLSI_{corr} = \alpha E_{PLSI} \tag{6.11}$$

with

$$\alpha = \frac{\sum_{t,f} PLSI}{\sum_{t,f} E_{PLSI}} \tag{6.12}$$

where t are the time bins during the baseline -0.5 s to 0 s, f are all frequency bins $[2, 128]$ and E_{PLSI} is the expected value from the surrogate distribution:

$$E_{PLSI} = \sum_n PLSI_{surr}(n)P(PLSI_{surr}(n)) \tag{6.13}$$

evaluated at $n = 500$ bins.

6.3.8 Phase Slope Index

The PSI is used to infer dominant unidirectional interactions, i.e. the net flow of information. It is based on the idea that the time lag required for a signal to travel from one location to another constitutes a frequency-dependent inter-regional phase difference which results in a phase slope across frequencies. The frequency-dependent PSI_{xy} of

signals x and y was computed from the conjugated complex coherencies of neighbouring frequencies:

$$PSI_{xy}(\omega_f) = \Im\left(\sum_{f}^{f+1} coh^*_{xy}(\omega_f) coh_{xy}(\omega_{f+1})\right) \tag{6.14}$$

where $\Im(\cdot)$ is the imaginary part, and

$$coh_{xy}(\omega_f) = crs_{xy}(\omega_f)/\sqrt{pow_x(\omega_f) pow_y(\omega_f)} \tag{6.15}$$

the complex coherency and * the complex conjugate.

6.3.9 Wiener-Granger Causality

Wiener-Granger Causality (WGC) quantifies how the inclusion of past values of signal y improves the prediction of signal x in comparison to an autoregressive model of x (Granger, 1969) frequency-dependent $WGC_{x\to y}(\omega)$ from signal x to y at frequency ω was calculated as follows:

$$WGC_{x\to y}(\omega) = \ln\left(\frac{pow_y(\omega)}{pow_y(\omega) - \left(\Sigma_{xx} - \frac{\Sigma^2_{yx}}{\Sigma_{yy}}\right)|H_{yx}(\omega)|^2}\right) \tag{6.16}$$

where $pow_y(\omega)$ is the power of signal y at frequency ω, Σ_{xx} and Σ_{yy} noise variances of signal x and y, respectively, and Σ_{yx} the noise covariance of signals x and y in their auto-regressive models and $H_{yx}(\omega)$ is the spectral transfer matrix. The noise covariance matrix Σ and $H_{yx}(\omega)$ was derived by spectral factorisation of the trial-averaged cross-spectral density crs_{xy} using the Wilson-Burg algorithm as implemented in FieldTrip.

6.3.10 Naïve Bayes decoding

The amount of information about the sample or distractor numerosity contained in local oscillatory activity was quantified using a Naïve Bayes classifier (MATLAB Machine Learning Toolbox) (C. M. Lewis et al., 2016). For each time-frequency bin a classifier was trained on the power of 75 % of randomly chosen trials. The feature vector comprised data from all prefrontal or parietal recording sites of an individual session. The remaining 25 % of trials were used to test the classifier. The numerosity with the highest posterior probability was interpreted as the classifier's prediction for a given trial. Accuracy was calculated as the ratio of correctly predicted trials to all test trials. For decoding of the distractor numerosity, the control condition trials (i.e. no distractor presented) were excluded. Accuracy was z-scored to baseline and averaged across sessions.

6.3.11 Phase Dependent Coding

To quantify the phase-dependent information about the sample and distractor numerosity that was carried by a neuron's spiking activity, spikes were grouped by LFP phase

into 12 equally spaced bins (Siegel et al., 2009). The information $I(\omega, b, s)$ of spikes at frequency ω and in phase bin b about stimulus s was computed using explained variance measure:

$$I(\omega, b, s) = \frac{SS_{eff} - df_{eff} MS_{error}}{SS_{total} + MS_{error}} \tag{6.17}$$

where the individual terms are derived from a one-way ANOVA across trials: SS_{eff} is the sum of squares of the effect, df_{eff} are the degrees of freedom for the effect, MS_{error} is the mean square of the error and SS_{total} is the total sum of squares. $I(\omega, b, s)$ of individual units with LFP channels of interest were combined into "sites" by averaging across pairs. Subsequently, $I(\omega, b, s)$ was smoothed with a two-dimensional Hanning kernel (0.5 octave by 90°, FWHM). $I(\omega, b, s)$ was averaged across sites and then normalized to the average information across phase bins per frequency:

$$I_{norm}(\omega, b, s) = \frac{\frac{1}{U}\sum_u^U I_u(\omega, b, s)}{\frac{1}{B}\sum_b^B \frac{1}{U}\sum_u^U I_u(\omega, b, s)} \tag{6.18}$$

where U is the number of sites and B is the number of phase bins.

Phase-dependent information (PDI) was quantified by the peak-to-mean modulation of a cosine fit to the phase-binned information $I_u(\omega, b, s)$ which was normalized to the average information across phase bins. The modulation is a measure of the non-uniformity of the magnitudes of the phase-binned information around an origin-centred circle in the complex plane:

$$PDI(\omega, s) = 4\frac{|\sum_u^U \sum_b^B I_u(\omega, b, s) e^{i\lambda_b}|}{\sum_u^U \sum_b^B I_u(\omega, b, s)} \tag{6.19}$$

The standard error of the mean (SEM) was estimated from bootstrapping across sites ($n = 1000$) Significant PDI against zero was calculated by shuffling phase labels within frequencies of $I(\omega, b, s)$ before smoothing ($n = 10000$). Observed values that exceeded the 99$^{\text{th}}$ percentile and that were present for more than one frequency bin were labelled significant.

To investigate the relevance of phase-dependent coding on behaviour, a subset of correct trials was replaced with the available error trials before computing $I(\omega, b, s)$. Differences in PDI were statistically tested using a null distribution of differences which was generated by shuffling outcome labels of smoothed $I_u(\omega, b, s)$ ($n = 1000$, $\alpha = 0.05$, two-tailed).

The optimally encoding phase $\varphi_{opt}(\omega, s)$ is the peak phase of a cosine fit to $I(\omega, b, s)$ or the phase of the average vector across phases and sites:

$$\varphi_{opt}(\omega, s) = \arg\left(\sum_u^U \sum_b^B I_u(f, b, s) e^{i\lambda_b}\right) \tag{6.20}$$

SEM of phase and $I_{norm}(\omega, b, s)$ was computed after smoothing by bootstrapping across sites ($n = 10000$).

6.3.12 Complex mediation

The influence of prefrontal LFPs on the synchrony between prefrontal spikes and parietal LFPs was investigated using complex mediation analysis. This allows us to quantify the influence of an independent variable x on a dependent variable y taking into account the influence of a third, mediating variable z. Classically, in the univariate real case this has been calculated by using regression coefficients c and c' of y regressed on x and x and z: $y = cx$ and $y = c'x + bz$ (MacKinnon et al., 2007). Here, c represents the overall effect of x on y, while c' represents the partial effect of x on y taking into account the partial effect of z on y. Thus, the proportion mediated is given by $1-(c'/c)$. An alternative approach to mediation analysis uses (partial) correlation coefficients, which is advantageous because the mediation effect is not dependent on the variance of the underlying data (Boca et al., 2014). Similarly, the proportion mediated would be $1-(r_{xy|z}/r_{xy})$, with r_{xy} and $r_{xy|z}$ the correlation coefficient of x and y and the partial correlation coefficient of x and y conditioned on z, respectively.

A more generalized approach to mediation analysis of complex valued multivariate data is the use of the RV coefficient of Escoufier (Escoufier, 1973; Pascual-Marqui et al., 2017). In the univariate real case the RV coefficient is equivalent to the squared correlation coefficent, which is why the proportion mediated θ_m was computed as

$$\theta_m = 1 - \sqrt{RV_{xy|z}/RV_{xy}} \tag{6.21}$$

with the RV coefficient

$$RV_{x,y} = \frac{tr(S_{x,y}S_{x,y}^*)}{\sqrt{tr(S_{x,x}^2)}\sqrt{tr(S_{y,y}^2)}} \tag{6.22}$$

and the partial RV coefficient

$$RV_{xy|z} = \frac{tr[(S_{zx} - S_{zy}S_{yy}^{-1}S_{yx})(S_{zx} - S_{zy}S_{yy}^{-1}S_{yx})^*]}{\sqrt{tr[(S_{zz} - S_{zy}S_{yy}^{-1}S_{yz})^2]}\sqrt{tr[(S_{xx} - S_{xy}S_{yy}^{-1}S_{yx})^2]}} \tag{6.23}$$

where S_{xy} is the variance-covariance matrix of matrices x and y, $tr(\cdot)$ is the trace of a matrix and * denotes the complex conjugate. Here, the matrices x, y and z are complex valued Fourier coefficients grouped as blocks of subcomponents, i.e. all channels PFC-LFP, all channels VIP-LFP and all channels PFC spike per session.

7 Mouse Project

7.1 Animal procedures

All animal procedures were performed in accordance with approvals by the local government (Regierung von Oberbayern). Animal health was monitored and scored daily. Male mice (wild type C57BL/6, in-house breeding, Institute for Neuroscience, Technical University Munich) were used for experiments. At surgery, animals were around 8-10 weeks old. After surgery, animals were singly housed under a regular 12 h/12 h light-dark cycle. Access to food and water was *ad libitum*. When in behavioural training, a controlled water schedule was applied that aimed for the animals to receive their daily water intake during the training session. If an animal did not succeed to receive at least 1500 µL of water during training, animals were supplemented to that amount.

7.2 Surgery

7.2.1 Pre-surgery procedures

Animals were anaesthetised with isoflurane (2 % (vol/vol) in O_2) and transferred onto a heating pad. Body temperature was continuously monitored via a rectal probe and kept to 37.5 °C via adjustments to the heating pad's temperature. The animal's upper front teeth were then fitted into the tooth bar of a stereotaxic frame (Model 900LS, *Kopf Instruments, USA*, extended with computer-controlled motors for each axis: Drill and Nanoinjection Robot, *Neurostar, Germany*). Inhalation anaesthetic was then delivered via the snout component of the stereotaxic frame and isoflurane concentration was lowered to 0.8 % to 1.5 %, and adjusted depending on the animal's breathing rate and foot pinch reflex. Analgesia (metamizole, 200 mg/kg body weight) was administered subcutaneously. Eyes were covered with eye cream (Augen- und Nasensalbe, *Bepanthen, Germany*) to avoid drying.

7.2.2 Head plate implantation

The animal's scalp was exposed by removing head hair using electric clippers. Local anaesthetic (lidocaine, 2 % in 0.9 % NaCl solution) was administered subcutaneously at the scalp. The animal's head was then fixed by fitting stereotaxic ear bars. The scalp was disinfected using 70 % ethanol, then excised. The periosteum was removed using forceps and the skull's surface was roughened using a scalpel blade. Afterwards, the skull surface was thoroughly cleaned with 0.9 % NaCl solution. A small craniotomy posterior and lateral to the skull lambda landmark was made using a dental drill and a stainless

steel screw attached to a metal pin, that was later used as an electrical reference, was implanted.

Next, skull alignment was optimised using the calibration functionality of Neurostar's stereotaxic navigation software. First, the three-dimensional location of skull landmarks bregma and lambda were measured to determine skull yaw, pitch and scale. Two points 500 µm left and right to the line described by the two skull landmarks were used to determine skull roll. If rotation was not optimal the skull was realigned and calibration repeated.

After calibration the skull was covered with light-curing adhesive (Optibond All-in-One, *Kerr, Germany*) and hardened using blue light. In order to have accurate entry sites for later acute recordings, stereotaxically-guided markings were made. The target regions were selected in the navigation software, as if planning an injection. A drill was automatically moved vertically above the target site and lowered until the skull was reached. The drill was used to make clearly visible markings into the solid adhesive and the superficial layers of the skull.

A small bar made of aluminium was then implanted as the component for head-fixation in the behavioural setup. Dental cement (Tetric EvoFlow, *Ivoclar Vivadent, Germany*) was applied onto the skull at the site of implantation, anterior of lambda. The stereotaxic frame was used to position the metal bar horizontally into the liquid dental cement. After solidifying the cement using blue light, the metal bar was detached from the stereotaxic frame. To ensure a stable implantation the entire exposed skull was then covered with multiple layers of dental cement, excluding the marked entry points for later electrophysiological recordings. The recording entry points were instead covered with rapid-curing silicone elastomeres (Kwik-Cast, *World Precision Instruments, United Kingdom*).

7.2.3 Pre-recording craniotomy

After animals reached behavioural criterion they were prepared for craniotomy by giving *ad libitum* access to water the day prior to surgery. Pre-surgery procedures were done as described above. Silicone elastomeres covering the recording entry points were removed and the exposed skull was thoroughly cleaned using 70 % ethanol, followed by 0.9 % NaCl solution. The markings made in the first surgery were used as a visual guide. Craniotomies were made using a dental drill and covered with rapid-curing silicone elastomeres (Kwik-Cast, *World Precision Instruments, United Kingdom*).

7.2.4 Post-surgery procedures

After both head plate implantation and pre-recording craniotomy, long-acting analgesic (meloxicam, 1.5 mg/kg body weight) was administered subcutaneously. The animal was then removed from the surgery setup and transferred into its home cage, which was heated. The following three days long-acting analgesic (meloxicam, 1.5 mg/kg body weight) was administered daily.

7.3 Behavioural setup

7.3.1 High-level overview

Behavioural training and experiments were performed in a custom-made sound-attenuating chamber. Components for behavioural procedures were mounted on an optical breadboard. Animals were head-fixed at 10 cm elevation using the head bar that was implanted during surgery. The animal's body was located in a cut plastic tube, which allowed free movement. Forepaws were resting on the response ball (see below). A blunt cannula was placed in front of the animal's mouth for water delivery. The cannula was connected to a water-filled 5 mL syringe that was dispensed by an electronically operated syringe pump (NE-500, *New Era Pump Systems, USA*). A 10 inch LCD monitor (FT10TMB, *Faytech, Germany*). was located 15 cm in front of the animal. One electrostatic speaker (ES1, *Tucker Davis Technology, USA*) each was located to the left and right of the monitor for delivery of auditory stimuli. Speakers were driven by the accompanying electrostatic speaker driver (ED1, *Tucker Davis Technology, USA*). Interfacing of inbound and outbound signals to a general purpose personal computer was done via a breakout panel (BNC-2090A, *National Instruments, USA*) connected to a data acquisition device (PCIe-6321, *National Instruments, USA*).

7.3.2 Response device

The response device was a ping-pong ball, the rotation of which was fixed by one polar axis using a small metal rod. This rotation axis was positioned parallel to the ground and to the animal's anterior-posterior axis. Thus, only left and right rotations (i.e. roll) were possible. A diametrically magnetised disc magnet was attached to the anterior pole of the metal rod. Rotations of the ball rotated the magnet. The rotation was measured using a Hall-effect rotary encoder (MIB22H, *Megatron, Germany*), the digital signals of which were transformed to an analog signal via a custom-made converter (design Christian Obermayer) before being transmitted into the breakout panel. Briefly, the transformation was done as follows (Fig. 7.1). Two-channel digital signals from the rotary encoder were used to determine rotation direction of the current step. The value of one of two digital counters was incremented dependent on the step rotation direction. Using digital-to-analog converters, each digital counter's sum was then mapped to an analog voltage. Finally, the two analog voltage signals were subtracted, to receive an analog signal that scaled with cumulative rotation. TTL pulses were required to reset the digital counters, in order to reset the analog signal back to baseline and to avoid counter overflows. During these reset TTL pulses no rotation data could be acquired.

7.3.3 Orchestration of behavioural task

Monkeylogic 1 (version 2014-10, Asaad et al. (2013)) with NIMH toolbox (version 2017-01-05, Hwang et al. (2019)) running in MATLAB 2017a 64 bit (*Mathworks, USA*), Windows 8 64-bit (*Microsoft, USA*) was used for behavioural control. The software used the incoming analog voltage signal from the response device to position a cursor between

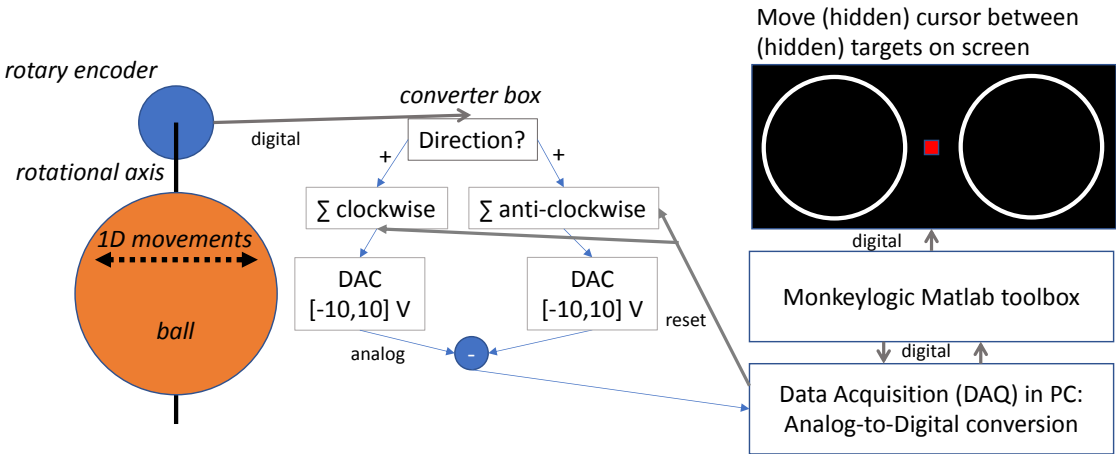

Figure 7.1: Transformations of behavioural signals.

two targets. The targets' distance from the centre scaled with the cumulative rotation that was needed to be registered as a left- or right-response. Outgoing digital (i.e. TTL) signals via the breakout-panel controlled water dispensal via the syringe pump and reset the converter circuit for the incoming response device rotation data. Outgoing analog signals were voltage waveforms for the electrostatic speaker driver. Task dependent control of timing and waveform selection was done via custom-written MATLAB scripts, which can be found at https://gitlab.lrz.de/jacob_lab/behavioral_tasks/-/tree/master/ML1/Ball.

7.4 Behavioural training

7.4.1 Habituation

Prior to instrumental training mice were habituated, similar to published protocols (Guo et al., 2014). They were manually handled by the experimenter and received up until 1500 µL water from a pipette. Afterwards, familiarisation with the body tube was supported by baiting it with water droplets, so that animals would be motivated to enter the tube themselves. Next, mice were habituated to head-fixation. Whilst inside the body tube they were moved into the setup, so that the body tube was just posterior to the response ball. Self-initiated exits from the front of the body tube were motivated by baiting the response ball with water droplets. Mice were then head-fixed using the implanted head bar. The water spout was moved in front of their snout and water was dispensed from it, in order to associate head-fixation with reward. These steps were typically done in different sessions on subsequent days.

7.4.2 Pre-training

In order to teach animals to use the response ball to receive reward, small rotations of the ball triggered water dispensal from the reward spout. We typically started with small

rotation thresholds of around 3.5 mm of the ball's circumference and large reward sizes of 8 µL. Movement training was done in blocks, i.e. in one block animals had to move the ball left, in others right. Multiple side blocks were presented within one session. Trial availability was signalled by showing a grey screen (RGB $0.1, 0.1, 0.1$), that was on throughout a trial and turned off after a threshold crossing. Monitor was black at background luminance between trials. Incorrect rotations (i.e. to the wrong side) were not rewarded and corresponding trials were immediately repeated.

7.4.3 Auditory decision making task

First stage: movement instructions Behavioural training with auditory stimuli was done in stages. The task from pre-training was extended as follows. After the onset of the grey screen animals had to hold the ball still for 50 ms. Trial progression was stalled until that duration. Afterwards, one of two auditory stimuli, sinusoidal up- and down-sweeps (11 kHz to 14 kHz, 15 kHz to 12 kHz, at around 78 dB sound pressure level Fig. 7.2) of duration of 1000 ms, was played. These sinusoidal sweeps were the movement instructions, with upsweeps signalling *go-left* and downsweeps signalling *go-right*. The ball response could be made immediately after onset of the sweep. Animals initially had 20 s to respond. If the ball was rotated across one of the two thresholds, playback of the auditory stimulus was aborted, and water was dispensed if the correct side's threshold was crossed. Afterwards the grey screen was turned off and an inter-trial-interval of 5000 ms followed. If no response within the allowed time frame was made the trial was ended without punishment. Initially, this task was presented in directional blocks of about 20 to 30 trials. With increasing task proficiency block sizes were reduced until trial types were presented randomly. Similarly, initial hold-time and rotation thresholds were gradually increased to 500 ms and 8 mm, and reward size was gradually decreased to 4 µL.

Second stage: context cues After animals had reached more than 90 % session performance on subsequent days in the randomised version of the previous task stage the context cues were introduced. Auditory context cues were 100 ms long and either pure or frequency-filtered white noise at around 61 dB sound pressure level (Fig. 7.2). White noise was not informative of the upcoming movement instruction, white-noise highpass-filtered at 14 kHz predicted the *go-right* instruction and white-noise lowpass-filtered at 8 kHz predicted the *go-left* instruction. Thus, the cues are referred to as *defer*, *prepare-right* and *prepare-left* contexts. Within the trial, context cues were presented after the initial calm period of 500 ms and before the movement instructions. Similar to the initial calm period, animals were required to hold the ball still between the offset of the context cue and the onset of the movement instruction. This time interval was initially at 50 ms and gradually increased to 500 ms if the animal's performance allowed it.

Full task Thus, the complete task was as follows (Fig. 3.2). First, a grey screen (see above) was shown on the monitor in front of the animal, in order to signal trial

availability. Then, the response ball's virtual position was reset to baseline and the animal was required to hold the ball still within ±2.5 mm rotation for 500 ms (initial calm period). If the small thresholds were crossed, i.e. the ball was not kept still enough, the virtual ball position and timer were reset, until the required calm time was registered. If the ball was kept still, one of three context cues was played for 100 ms from the speakers left and right to the monitor. The ball's virtual position was reset to baseline and another period (context delay epoch) ensued that required the animal to hold the ball still for 500 ms. Once again, if the ball was not kept still, the delay was extended similar to the initial calm period. If the ball had been kept still, one of two movement instruction was played for up to 1000 ms. Animals could immediately respond by rotating the ball to the left or right. Responses were allowed for up to 5000 ms after movement instruction onset. If one of the two thresholds of ±8 mm was crossed, a potentially still playing sound was stopped. If the response side was correct, water reward was dispensed from the water spout for 500 ms. If the the response side was incorrect the trial was aborted. The grey screen was turned off at the beginning of the inter-trial interval of 5000 ms.

7.5 Neurophysiology

7.5.1 Manual procedures

Animals were aged around 10 months to 12 months at the time of recording. Prior to recording animals were head-fixed in the setup as usual. A custom-made reference cable was connected to the implanted skull screw pin and the reference electrode of the adapter that held the silicon probe. Silicone elastomere that covered the craniotomies was removed and craniotomies and brain surface were rinsed using 0.9 % saline. One silicon probe (A1x32-Poly2-10mm-50s-177-A32, *NeuroNexus Technologies, USA*) per site was moved to the brain surface using a micromanipulator (*Luigs-Neumann, Germany*). Using the brain surface as the baseline, the neural probe was slowly lowered to the desired depth. The well surrounding the craniotomy was filled with 2 % agarose in 0.9 % saline to increase stability and avoid drying of the brain's surface. After insertion, the probe was allowed to settle for 10 min to 30 min. Recording of neural signals was started shortly before starting the behavioural session. Behavioural session was controlled as it was during training. After the behavioural session the neural probe was retracted, the craniotomy rinsed with 0.9 % saline and covered with silicone elastomere.

7.5.2 Recording

Recordings were made using the OmniPlex™ Neural Recording Data Acquisition System by *Plexon, USA*. A silicon probe (A1x32-Poly2-10mm-50s-177-A32, *NeuroNexus Technologies, USA*) was mounted onto a a probe adaptor (Connect HST/32V, *Plexon, USA*), onto which an analog headstage (HST/32o25-GEN3-36P-G1, *Plexon, USA*) was mounted. The headstage was connected to a digitising amplifier (DigiAmp™, *Plexon,*

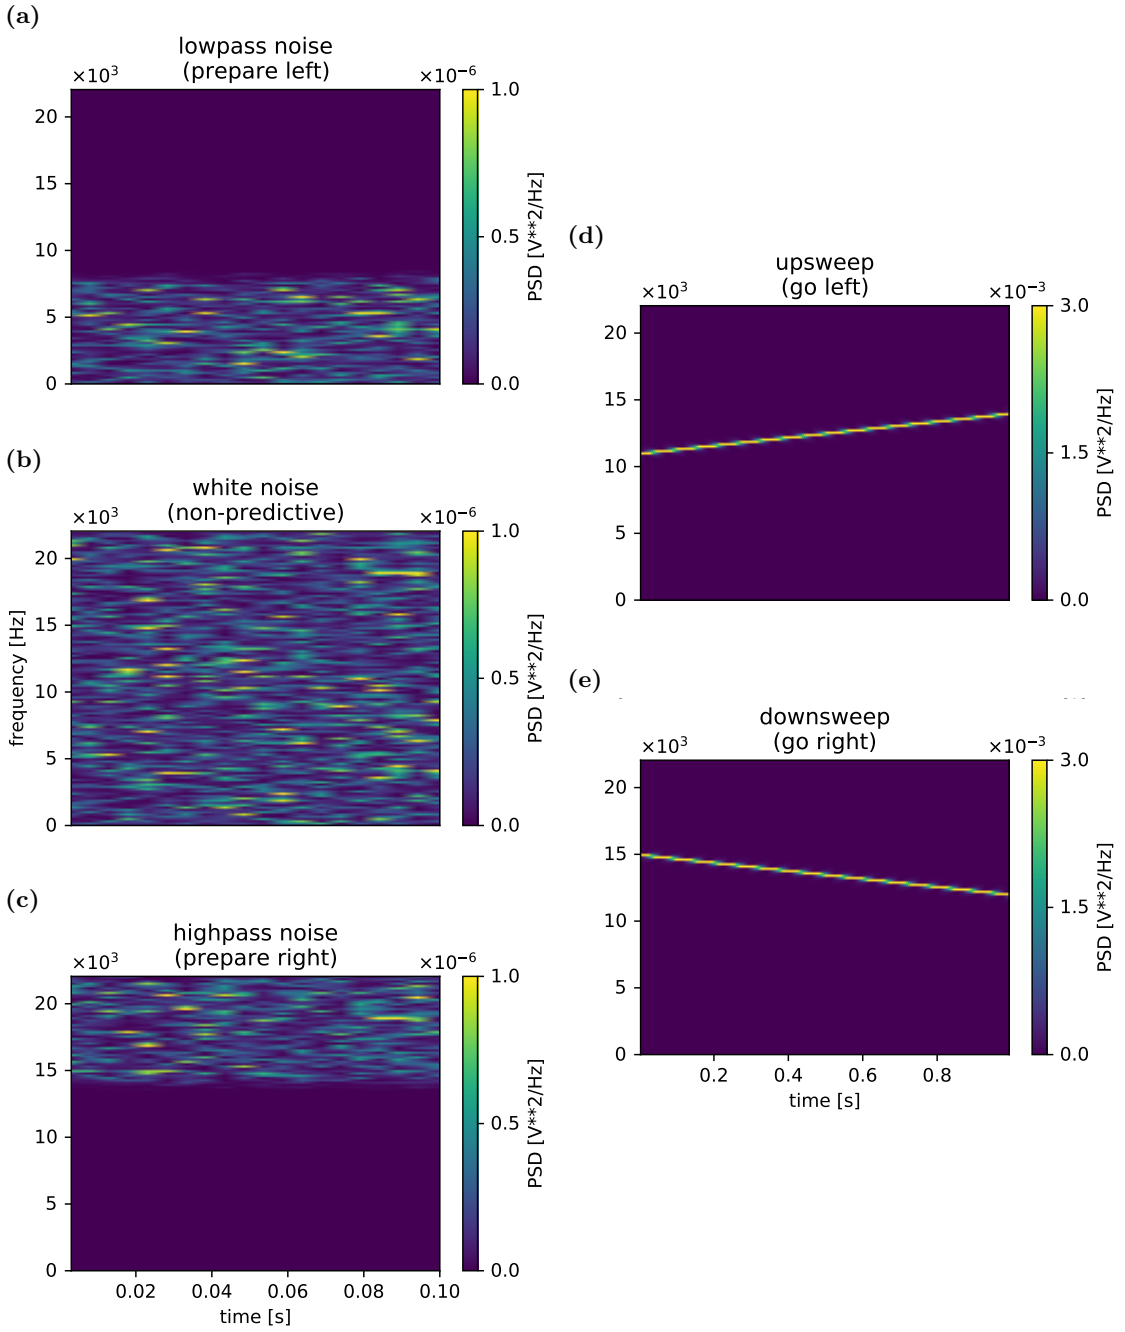

Figure 7.2: Spectrograms for auditory stimuli used in the experiment. **(a)** White noise lowpass-filtered at 8 kHz was used as a context cue to signal an upcoming *go-left* trial. **(b)** White noise was used as a non-informative context cue. **(c)** White noise highpass-filtered at 14 kHz was used as a context cue to signal an upcoming *go-right* trial. **(d)** A sinusoidal upsweep (11 kHz to 14 kHz) was used as the movement instruction for *go-left*. **(e)** A sinusoidal downsweep (15 kHz to 12 kHz) was used as the movement instruction for *go-right*.

USA), which was connected to an OmniPlex™ chassis. The chassis was connected to a personal computer (*Dell, USA*), running Windows 7 (*Microsoft, USA*).

Data acquisition was controlled using the OmniPlex™ software. Signal was sampled at a rate of 40 kHz. Wideband signal was split into field potentials (lowpass filtered at 500 Hz and downsampled to 1000 Hz) and continuous spike signal (highpass filtered at 300 Hz) and stored to disk in Plexon's proprietary `.pl2` format.

7.5.3 Spike sorting

Spikes were assigned to clusters, i.e. putative single neuronal units and multi neuronal units, using a semi-automated approach. First, the highpass filtered continuous spike signal stored in `.pl2` format was converted into a flat binary file with length $n_{channels} \cdot n_{timepoints}$. Data in that file was automatically clustered based on spatiotemporal spike waveforms via Kilosort 1 (Pachitariu et al., 2016) using the parameters shown in Table 7.1.

This over-splits spikes into 512 clusters and requires manual merging. Manual curation of the pre-clustered data was done in Phy (https://github.com/cortex-lab/phy). This included merging, splitting, labelling and cutting. For a given cluster Phy shows the most similar clusters in descending order. Clusters were merged if their spatio-temporal waveforms or autocorrelation histograms were very similar. Sometimes clusters drifted in terms of amplitude, but also spatially towards neighbouring channels, resulting in separated clusters. This was clearly visible as sudden partitions in plots of first principal components against time. These clusters were merged. Splitting was very rarely necessary. It was only done when upon visual inspection of principal components scatter plots there were clearly distinct clusters within the same time window. Labelling was based on waveform and the autocorrelation histograms. If there was a distinct, often large-amplitude, waveform and / or a physiologically plausible refractory period, a cluster was assigned the "good" (i.e. "single unit") label. If there were clear violations of refractory period but a physiologically plausible waveform, this indicates a recording from multiple units, and the cluster was labelled "MUA". Note that, contrary to threshold-based MUA detection, we can have multiple MUA clusters per channel, based on distinct waveforms. Finally, clearly non-physiological waveforms were labelled as "noise". For subsequent analysis we cut off periods of sparse firing at the beginning and end of a unit's lifetime to ensure stationary firing rates. This was necessary because units drifted in and out of our measurements because of the acute recording paradigm.

7.6 Data analysis

Data was analysed using custom scripts in python and MATLAB, with the support of various libraries (Table 7.2).

parameter	value
Nfilt	512
Nrank	3
th	4 10 10
crit	0.65
criterion_noise_channels	0.2
fs	40000
fshigh	300
initialize	no
lam	5 20 20
loc_range	3 1
long_range	30 6
mask_max_channels	5
maxfr	20000
merget	0.1
momentum	0.05 0.0025
nfiltmax	10000
nneigh	16
nneighpc	12
nskipcov	1
nannealpasses	4
nfullpasses	6
nt0	61
nt0min	20
ntbuff	64
scaleproc	200
shuffleclusters	1
spkth	-6
splitt	0.1
whitening	full

Table 7.1: Parameters used for Kilosort spike clustering

Name	Version	Reference
python	3.9.4	Van Rossum and Drake (2009)
numpy	1.20.3	C. R. Harris et al. (2020)
scipy	1.8.0	Gommers et al. (2022)
scikit-learn	1.0.1	Grisel et al. (2021)
lmfit	1.0.2	Newville et al. (2021)
numba	0.52.0	Lam et al. (2020)
pandas	1.2.4	Reback et al. (2021)
xarray	0.19.0	Hoyer et al. (2021)
matplotlib	3.4.2	Caswell et al. (2021)
ptitprince	3.4.2	https://github.com/pog87/PtitPrince
MATLAB	R2020a	Mathworks, Inc.
MVGC	1.3	Barnett and Seth (2015)

Table 7.2: Software used for data analysis

7.6.1 Valid trials and sessions

To avoid spurious effects of motivation only the bulk of a behavioural session was used for further analyses, thus cutting off the warm-up and cool-down periods, using the following algorithm. Contiguous groups of trials were designated "bad runs" if the running average correct performance across 15 trials was below 50 %. The first non-"bad" run after 20 trials into the session (warm-up) was designated the beginning of the valid trials. The start of the first "bad run" after the beginning of the valid trials was used as the cutoff for the end of the session. The end of a session was signalled by the animal via missing or repetitive incorrect responses. Finally, all *missing* trials (i.e. no response) were invalidated.

A valid session was defined as having at least 100 valid trials overall, 70 % overall behavioural performance, 50 % performance per trial condition, and 5 correct trials per trial condition.

Furthermore, for the analysis of spiking data only those sessions were included for which the location of the neural probe was histologically verified to be in the target area. Thus, spiking data of one animal was excluded.

7.6.2 Psychometrics

Session performance was defined as the fraction of correct trials with respect to overall valid trials. Response time was defined as the time interval from the onset of the movement instruction to threshold crossing. Differences between *prepare* and *defer* trials were calculated per session.

Setup number	Animals	Noise variance σ^2
0	A30, A32	$5 \cdot 10^{-4}$
1	D13, D15, T12	$1 \cdot 10^{-3}$

Table 7.3: Setup numbers for animals and empirically determined noise variance.

7.6.3 Preprocessing response ball traces

The electrical elements that were used to count rotations from the rotary encoder had an upper limit and thus needed to be reset / cleared for new trials, using an external reset pulse signal sent by *Monkeylogic*. Furthermore, some epochs in our behavioural task required the animals to hold the ball within a certain range close to 0 V. Therefore, reset pulses were triggered within the task if the animals exceeded this range. During the reset pulse, which lasted 1 frame (about 33 ms at 30 Hz refresh rate), the net voltage signal was brought back to a baseline, close to, but not exactly at 0 V. In this time windows the actual ball rotations (or lack thereof) were masked and could not be observed.

To fill in these gaps and filter other noise I reconstructed ball rotation velocity traces using Gaussian Process (GP) regression with a 2-component kernel, using `scikit-learn`'s `gaussian_process.GaussianProcessRegressor`. In the first component, gap interpolation was achieved by using an Radial Basis Function (RBF) kernel with the length scale parameter ℓ fixed for all sessions. The RBF length scale ℓ determines the smoothing of a function, as well as the maximum steps away from data points that the function can be extrapolated to. In general, it's not possible to extrapolate more than ℓ units away from the data (Duvenaud, 2014). The RBF can be thought of as a weighting function of neighbouring data points, putting more weight to closer data points and less to more distant ones. Thus, to retain maximum fidelity (i.e. small movements) of the original signal ℓ should be chosen as small as possible, while still ensuring that the RBF kernel covers data points on both sides of a gap to be filled. With a minimum of 2 data points on each side of the gap and a maximum reset pulse length of 33 ms, ℓ was chosen as 37 ms. Time points of reset pulses were determined algorithmically on the raw signal. The second component of the GP regression involves a white noise kernel that is employed to filter noise in the voltage dimension. This kind of noise arose from repeated digital-to-analog and analog-to-digital conversions, as well as outlier peaks presumably produced by insufficiently filtered signals acquisition pacemaker signals. White noise kernels are parameterised by their variance σ^2. σ^2 was determined per behavioural setup. Pooled values of manually chosen segments of ball rotation traces that showed no movement were used to fit normal distributions per session. The standard deviations σ derived from these distributions were pooled per setup and the peak of a kernel density estimate on these pooled σs was used to determine σ^2 per setup. Table 7.3 shows the mapping of setup and animals.

7.6.4 Trial subsetting dependent on movement

Trial subsetting was done for movement activity and movement tendency during the first 600 ms of the context epoch. Movement activity was defined as the absolute path length, and, similarly, movement side tendency was defined as the cumulative (signed) path length normalised by the absolute path length. To allow for fair comparisons, equal trial subsets were defined separately within conditions. Each subset contained the upper or lower two quintiles of the respective metric within the condition (i.e. 0 % to 40 % and 60 % to 100 %).

7.6.5 Submovement decomposition

Minimum-jerk submovements were fitted onto single trials' response ball velocity traces using a custom iterative algorithm, which can be found at https://gitlab.lrz.de/jacob_lab/submo. On a high level the algorithm consists of the steps *proposal*, *fitting* and *evaluation*.

proposal During proposal the residual of data and current best model is used to find a new candidate submovement. In the first iteration the residual is just the raw velocity because the model is initialised to all zeroes. First, the local velocity maxima and their neighbouring local minima were found using `scipy-1.5.3`'s `signal.argrelextrema` (Virtanen et al., 2020). Difference between left and right local minima were used as the initial guess for submovement duration. The proposed duration d and velocity peak v_{peak} were used to calculate the corresponding path length a

$$a = v_{peak}/1.875 \cdot d \tag{7.1}$$

The proposed submovement with the longest path length was used for the fitting stage if it wouldn't be entirely nested within another submovement and it wasn't tried to fit in a previous iteration.

fitting Non-linear least-squares fitting was done via `lmfit` using custom functions (just-in-time compiled using `numba`) for minimum jerk velocity

$$v(t) = \frac{a}{d} \cdot (-60 \cdot nt(t)^3 + 30 \cdot nt(t)^4 + 30 \cdot nt(t)^2) \tag{7.2}$$

and its corresponding Jacobian

$$J(t) = [\frac{\partial f}{\partial t_0}(t), \frac{\partial f}{\partial d}(t), \frac{\partial f}{\partial a}(t)] \tag{7.3}$$

with the partial derivatives

$$\frac{\partial f}{\partial t_0}(t) = \frac{-a}{d^2} \cdot (120 \cdot nt(t)^3 - 180 \cdot nt(t)^2 + 60 \cdot nt(t)) \tag{7.4}$$

$$\frac{\partial f}{\partial d}(t) = \frac{a}{d^2} \cdot (-150 \cdot nt(t)^4 + 240 \cdot nt(t)^3 - 90 \cdot nt(t)^2) \tag{7.5}$$

$$\frac{\partial f}{\partial a}(t) = \frac{1}{d} \cdot (-60 \cdot nt(t)^3 + 30 \cdot nt(t)^4 + 30 \cdot nt(t)^2) \tag{7.6}$$

with v as velocity, a as amplitude (or path length), d as duration, t_0 as movement onset time, $nt(t)$ as the normalised time $\frac{t-t_0}{d}$.

evaluation The model was evaluated after every fitting step. If Bayesian Information Criterion (BIC) was improved in comparison to the previous iteration the fit was accepted. Otherwise the proposed submovement was put on a deny-list and the fitting procedure was repeated for the next longest proposed movement. If no more movements could be fitted the iteration was exited.

7.6.6 Response/movement initiation and execution time

Response/movement initiation time was defined as the time of onset of the first fitted submovement after the onset of the movement instruction or context cue, respectively. Response execution time was defined as the time interval from response initiation to threshold crossing.

7.6.7 Submovement aggregated metrics

Aggregated submovement metrics were calculated from fitted submovements in the context epoch (first 0 ms to 600 ms) and in the instruction epoch (from movement instruction onset to threshold crossing). Submovement count was the number of fitted submovements in the time of interest. Movement balance SI was the normalised signed count per trial as in

$$SI = \frac{n_{left} - n_{right}}{n_{left} + n_{right}}. \tag{7.7}$$

Submovement strength was v_{peak}, and submovement duration was d, as defined in Equation (7.1).

7.6.8 Inclusion criteria neuronal units

Only spiking activity which was deemed to originate from single units (Section 7.5.3) were considered for analysis. Single units were further limited to have at least an average firing rate of 1 Hz and be present for at least 5 correct trials per trial condition. Unit lifetime was determined as the interval of trials in which a unit showed not more than 5 consecutive trials of firing rate below 1 Hz.

7.6.9 Standardised firing rate

Binary spike trains were first downsampled to 1000 Hz and then smoothed to continuous firing rate time series by convolving with a Gaussian window of $\sigma = 25$ ms and length of $n = 9 \cdot \sigma$. Individual single units' time series $fr(t)$ were standardised to baseline as

$$fr_z(t) = \frac{fr(t) - \overline{fr_{bl}}}{\widehat{fr_{bl}}} \tag{7.8}$$

with $\overline{fr_{bl}}$ as the mean baseline firing rate, $\widehat{fr_{bl}}$ as the standard deviation of the baseline firing rate and the baseline defined as $-200\,\text{ms}$ to $0\,\text{ms}$ to the presentation of the trial availability grey screen.

7.6.10 Zeta-test derived metrics

Neuronal responsiveness to stimuli was computed from the *raw* spike timestamps using the parameter-free ZETA-test (Montijn et al., 2021). The test computes the magnitude of responsiveness (ζ) and a ζ-derived instantaneous firing rate. A single unit was deemed *responsive* to a stimulus if its ζ value (responsiveness magnitude) exceeded 1.96. Latency was defined as the point in time of maximum ζ-derived instantaneous firing rate. Region-wide recruitment profiles were based on these latencies. Preference was defined as the stimulus to which a neuron's activity was maximally modulated ($argmax([\zeta_{defer}, \zeta_{prepare\text{-}left}, \zeta_{prepare\text{-}right}])$ or $argmax([\zeta_{go\text{-}left}, \zeta_{go\text{-}right}])$).

7.6.11 Population coding: support vector machine (SVM)

In order to assess encoding of the task variables in the neuronal populations L_2-regularised linear support vector machines were used (`scikit-learn`'s `svm.LinearSVC`) with hinge-loss, stopping tolerance $tol = 1 \cdot 10^{-5}$, regularisation $C = 100$, and maximum number of iterations $n_{iter} = 10\,000$. In the case of decoding context, i.e. a 3-class classification, 1-vs-1 classifiers were used.

Binary spike trains were smoothed by convolution with a Gaussian window ($\sigma = 25\,\text{ms}$) and downsampled to 40 Hz.

Activity from simultaneously recorded neurons was not robustly available for the dataset due to the recording scheme (daily acute recordings with neuronal drift). Thus, region-specific pseudopopulations containing all valid neurons across all sessions were created for classifier training and testing (E. Meyers & Kreiman, n.d.). Train and test trials were created by randomly sampling 50 trials (with replacement) from individual neurons' trials and concatenating across neurons. This procedure was done for 100 shuffles of random train/test splits (train fraction 0.75), to increase the probability that all trials were used for the analysis. One caveat here is that trials from units with a small number of trials are duplicated, so that the variances of those units' features is low. These so-called pseudotrials abolish trial-wise neuronal covariation that might contain information. Therefore, the results shown represent a lower bound on the achievable information. Information was quantified as the mean accuracy of class predictions of trained classifiers on held-out test trials. Means and standard deviation across pseudo-trial shuffles.

Cross-temporal decoding was tested by training a classifier on data of one time bin and testing at other time bins, similar to as reviewed in E. M. Meyers (2018).

Two approaches were used to assess context-dependent differences of the population code for movement instruction. Each fitted linear support vector classifier defines a multi-dimensional hyperplane that optimally separates data points of two classes. Thus, intuitively the hyperplane can be viewed as the "fingerprint" of population activity at

specific times. The angle ϕ between two hyperplanes $\vec{\omega_a}$ and $\vec{\omega_b}$ was calculated via

$$\phi = \arccos \vec{\omega_a} \bullet \vec{\omega_b} \tag{7.9}$$

with $\bullet$ representing the dot-product operation. All angles of a hyperplane fitted for one pseudo-trial shuffle at one time point with hyperplanes from the other context's trials were computed.

Cross-context decoding was done similar to cross-temporal decoding. Classifiers were trained on a training subset of trials in one context. Testing was done with held out trials of the same and the other context.

7.6.12 Functional connectivity: Granger Causality

Binary spike trains were convolved with half-Gaussian windows ($\sigma = 50\,\mathrm{ms}$). For each session, concurrent trials for pairs of single units were calculated. Multivariate Granger Causality in state space was then computed for pairs of preprocessed time-series, using containerised `MATLAB` and `MVGC` library connected to python. Optimal model order was selected using the Akaike Information Criterion (AIC). Significance was assessed using the F-test integrated in the library, evaluated at $\alpha = 0.05$.

7.6.13 Comparative statistics

The following statistical tests were used for comparison of groups.

Repeated-measures two-way ANOVA Comparison of the four trial conditions *prepare-go-left*, *defer-go-left*, *prepare-go-right*, *defer-go-right*. Main factors: movement instruction (*go-left* vs. *go-right*), abstract context (*prepare* vs. *defer*); Interaction: (movement instruction) $\times$(abstract context).

Wilcoxon signed-rank test

- **two-sample** Comparison of context-differences across movement instructions: diff(*go-left-prepare*, *go-left-defer*) vs. diff(*go-right-prepare*, *go-right-defer*).
- **one-sample** Comparison of context-differences against 0: diff(*go-left-prepare*, *go-left-defer*) vs. 0; diff(*go-right-prepare*, *go-right-defer*) vs 0.

Levene's test for equal variances Comparison of distribution variances in context-differences across movement instructions: diff(*go-left-prepare*, *go-left-defer*) vs. diff(*go-right-prepare*, *go-right-defer*).

Hartigan's dip test for unimodality Within conditions, check distribution for unimodality.

χ^2 tests

- **two-way/contingency** Comparison of equal fractions across regions, or within regions across sub-groups.

one-way Comparison of equal fractions within regions.

permutation tests

neuronal latencies across regions Random shuffling of region assignments for single unit's peak latencies.

significant neuronal coding Decoder scores from pseudotrial runs were shuffled in time and mean across pseudotrials was compared against observed means.

differences in neuronal coding Clusters of siginifcant differences were found as follows. Input data were scores within pseudo-trials. Candidate clusters were contiguous runs of observed differences between conditions (e.g. *prepare* vs. *defer*; correct vs. wrong). Cluster statistic was the absolute sum within clusters. Assignment of condition and decoder score was then shuffled and cluster statistic was calculated for each shuffle. A cluster was deemed significant if across those shuffles the observed cluster statistic was greater than in 99 % of shuffles.

Epps-Singleton two-sample test for equal distributions Used to compared recruitment dynamics across conditions.

Kendall's τ test for equal sequences Use to compared recruitment times of context-stable units across contexts.

7.6.14 Raincloud plots

Raincloud plots are a combination of scatter plots, box plots and violin plots (Allen et al., 2021) and thus represent single data points as well as their distributions. In this thesis colours represent trial conditions, every dot is the average of a metric in one session, half-violins are kernel density estimate (KDE) smoothed distributions, box-plots show quartiles (white dots are medians), whiskers include 1.5 times inter-quartile range (IQR).

References

Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., van Langen, J., & Kievit, R. A. (2021). Raincloud plots: A multi-platform tool for robust data visualization. *Wellcome Open Research*, *4*, 63. https://doi.org/10.12688/wellcomeopenres.15191.2

Amit, D. J. (1995). The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, *18*(4), 617–626. https://doi.org/10.1017/S0140525X00040164

Antzoulatos, E. G., Miller, E. K., & Pasternak, T. (2016). Synchronous beta rhythms of frontoparietal networks support only behaviorally relevant representations. *eLife*, *5*, e17822. https://doi.org/10.7554/eLife.17822

Aoi, M. C., Mante, V., & Pillow, J. W. (2020). Prefrontal cortex exhibits multi-dimensional dynamic encoding during decision-making. *Nature neuroscience*, *23*(11), 1410–1420. https://doi.org/10.1038/s41593-020-0696-5

Aoki, R., Tsubota, T., Goya, Y., & Benucci, A. (2017). An automated platform for high-throughput mouse behavior and physiology with voluntary head-fixation. *Nature Communications*, *8*(1), 1196. https://doi.org/10.1038/s41467-017-01371-0

Arnsten, A. F. T. (2009). Stress signalling pathways that impair prefrontal cortex structure and function. *Nature Reviews. Neuroscience*, *10*(6), 410–422. https://doi.org/10.1038/nrn2648

Asaad, W. F., Santhanam, N., McClellan, S., & Freedman, D. J. (2013). High-performance execution of psychophysical tasks with complex visual stimuli in MATLAB. *Journal of Neurophysiology*, *109*(1), 249–260. https://doi.org/10.1152/jn.00527.2012

Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual Review of Psychology*, *63*, 1–29. https://doi.org/10.1146/annurev-psych-120710-100422

Baddeley, A. D. (1983). Working memory. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *302*(1110), 311–324. https://doi.org/10.1098/rstb.1983.0057

Bando, Y., Wenzel, M., & Yuste, R. (2021). Simultaneous two-photon imaging of action potentials and subthreshold inputs in vivo. *Nature Communications*, *12*(1), 7229. https://doi.org/10.1038/s41467-021-27444-9

Barbas, H., Zikopoulos, B., & Timbie, C. (2011). Sensory pathways and emotional context for action in primate prefrontal cortex. *Biological Psychiatry*, *69*(12), 1133–1139. https://doi.org/10.1016/j.biopsych.2010.08.008

Barker, J. M., Glen, W. B., Linsenbardt, D. N., Lapish, C. C., & Chandler, L. J. (2018). Habitual Behavior Is Mediated by a Shift in Response-Outcome Encoding by

Infralimbic Cortex. *eNeuro*, *4*(6), ENEURO.0337–17.2017. https://doi.org/10.1523/ENEURO.0337-17.2017

Barnett, L., & Seth, A. K. (2015). Granger causality for state space models. *Physical Review E*, *91*(4), 040101. https://doi.org/10.1103/PhysRevE.91.040101

Bastos, A. M., Loonis, R., Kornblith, S., Lundqvist, M., & Miller, E. K. (2018). Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. *Proceedings of the National Academy of Sciences*, *63*, 201710323–6. https://doi.org/10.1073/pnas.1710323115

Bastos, A. M., Vezoli, J., Bosman, C. A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J. R., De Weerd, P., Kennedy, H., & Fries, P. (2015). Visual areas exert feedforward and feedback influences through distinct frequency channels. *Neuron*, *85*(2), 390–401. https://doi.org/10.1016/j.neuron.2014.12.018

Baudry, B., & Monperrus, M. (2022). *Exhaustive Survey of Rickrolling in Academic Literature*. https://doi.org/10.48550/arXiv.2204.06826

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*(1), 7–15. https://doi.org/10.1016/0010-0277(94)90018-3

Berens, P. (2009). **CircStat** : A *MATLAB* Toolbox for Circular Statistics. *Journal of Statistical Software*, *31*(10). https://doi.org/10.18637/jss.v031.i10

Bisley, J. W., Zaksas, D., Droll, J. A., & Pasternak, T. (2004). Activity of Neurons in Cortical Area MT During a Memory for Motion Task. *Journal of Neurophysiology*, *91*(1), 286–300. https://doi.org/10.1152/jn.00870.2003

Blaukopf, C. L., & DiGirolamo, G. J. (2005). The automatic extraction and use of information from cues and go signals in an antisaccade task. *Experimental Brain Research*, *167*(4), 654–659. https://doi.org/10.1007/s00221-005-0125-8

Boca, S. M., Sinha, R., Cross, A. J., Moore, S. C., & Sampson, J. N. (2014). Testing multiple biological mediators simultaneously. *Bioinformatics*, *30*(2), 214–220. https://doi.org/10.1093/bioinformatics/btt633

Bodner, M., Shafi, M., Zhou, Y.-D., & Fuster, J. M. (2005). Patterned firing of parietal cells in a haptic working memory task. *European Journal of Neuroscience*, *21*(9), 2538–2546. https://doi.org/10.1111/j.1460-9568.2005.04085.x

Bolkan, S. S., Stujenske, J. M., Parnaudeau, S., Spellman, T. J., Rauffenbart, C., Abbas, A. I., Harris, A. Z., Gordon, J. A., & Kellendonk, C. (2017). Thalamic projections sustain prefrontal activity during working memory maintenance. *Nature Neuroscience*, *20*(7), 987–996. https://doi.org/10.1038/nn.4568

Bollu, T., Ito, B. S., Whitehead, S. C., Kardon, B., Redd, J., Liu, M. H., & Goldberg, J. H. (2021). Cortex-dependent corrections as the tongue reaches for and misses targets. *Nature*, *594*(7861), 82–87. https://doi.org/10.1038/s41586-021-03561-9

Bollu, T., Whitehead, S. C., Prasad, N., Walker, J., Shyamkumar, N., Subramaniam, R., Kardon, B., Cohen, I., & Goldberg, J. H. (2018). Automated home cage training of mice in a hold-still center-out reach task. *J Neurophysiol*, 13. https://doi.org/10.1152/jn.00667.2018

Bouchacourt, F., & Buschman, T. J. (2019). A Flexible Model of Working Memory. *Neuron*, *0*(0). https://doi.org/10.1016/j.neuron.2019.04.020

Boussaoud, D., Ungerleider, L. G., & Desimone, R. (1990). Pathways for motion analysis: Cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque. *The Journal of Comparative Neurology*, *296*(3), 462–495. https://doi.org/10.1002/cne.902960311

Bovenkerk, B., & Kaldewaij, F. (2015). The use of animal models in behavioural neuroscience research. *Current Topics in Behavioral Neurosciences*, *19*, 17–46. https://doi.org/10.1007/7854_2014_329

Bressler, S. L., & Menon, V. (2010). Large-scale brain networks in cognition: Emerging methods and principles. *Trends in cognitive sciences*, *14*(6), 277–290. https://doi.org/10.1016/j.tics.2010.04.004

Brincat, S. L., & Miller, E. K. (2015). Frequency-specific hippocampal-prefrontal interactions during associative learning. *Nature Neuroscience*, *18*(4), 576–581. https://doi.org/10.1038/nn.3954

Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D., & Miller, E. K. (2012). Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex. *Neuron*, *76*(4), 838–846. https://doi.org/10.1016/j.neuron.2012.09.029

Buschman, T. J., & Miller, E. K. (2007). Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science (New York, N.Y.)*, *315*(5820), 1860–1862. https://doi.org/10.1126/science.1138071

Buschman, T. J., & Miller, E. K. (2009). Serial, covert shifts of attention during visual search are reflected by the frontal eye fields and correlated with population oscillations. *Neuron*, *63*(3), 386–396. https://doi.org/10.1016/j.neuron.2009.06.020

Butter, C. M. (1969). Perseveration in extinction and in discrimination reversal tasks following selective frontal ablations in Macaca mulatta. *Physiology & Behavior*, *4*(2), 163–171. https://doi.org/10.1016/0031-9384(69)90075-4

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929. https://doi.org/10.1126/science.1099745

Caminiti, R., Innocenti, G. M., & Battaglia-Mayer, A. (2015). Organization and evolution of parieto-frontal processing streams in macaque monkeys and humans. *Neuroscience and Biobehavioral Reviews*, *56*, 73–96. https://doi.org/10.1016/j.neubiorev.2015.06.014

Carandini, M., & Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature Neuroscience*, *16*(7), 824–831. https://doi.org/10.1038/nn.3410

Carlesimo, G. A., Lombardi, M. G., & Caltagirone, C. (2011). Vascular thalamic amnesia: A reappraisal. *Neuropsychologia*, *49*(5), 777–789. https://doi.org/10.1016/j.neuropsychologia.2011.01.026

Carter, E., & Wang, X.-J. (2007). Cannabinoid-mediated disinhibition and working memory: Dynamical interplay of multiple feedback mechanisms in a continuous attractor model of prefrontal cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, *17 Suppl 1*, i16–26. https://doi.org/10.1093/cercor/bhm103

Castiello, U., & Dadda, M. (2019). A review and consideration on the kinematics of reach-to-grasp movements in macaque monkeys. *Journal of Neurophysiology*, *121*(1), 188–204. https://doi.org/10.1152/jn.00598.2018

Caswell, T. A., Droettboom, M., Lee, A., De Andrade, E. S., Hunter, J., Hoffmann, T., Firing, E., Klymak, J., Stansby, D., Varoquaux, N., Nielsen, J. H., Root, B., May, R., Elson, P., Seppänen, J. K., Dale, D., Lee, J.-J., McDougall, D., Straw, A., … Ivanov, P. (2021, May 8). *Matplotlib/matplotlib: REL: V3.4.2* (Version v3.4.2). https://doi.org/10.5281/ZENODO.4743323

Catena, A., Perales, J. C., Megías, A., Cándido, A., Jara, E., & Maldonado, A. (2012). The brain network of expectancy and uncertainty processing. *PloS One*, *7*(7), e40252. https://doi.org/10.1371/journal.pone.0040252

Chafee, M. V., & Goldman-Rakic, P. S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *Journal of Neurophysiology*, *83*(3), 1550–1566. https://doi.org/10.1152/jn.2000.83.3.1550

Chao, L. L., & Knight, R. T. (1998). Contribution of human prefrontal cortex to delay performance. *Journal of Cognitive Neuroscience*, *10*(2), 167–177. https://doi.org/10.1162/089892998562636

Chao, L. L., & Knight, R. T. (1995). Human prefrontal lesions increase distractibility to irrelevant sensory inputs. *Neuroreport: An International Journal for the Rapid Communication of Research in Neuroscience*, *6*(12), 1605–1610. https://doi.org/10.1097/00001756-199508000-00005

Chauveau, F., Célérier, A., Ognard, R., Pierard, C., & Béracochéa, D. (2005). Effects of ibotenic acid lesions of the mediodorsal thalamus on memory: Relationship with emotional processes in mice. *Behavioural Brain Research*, *156*(2), 215–223. https://doi.org/10.1016/j.bbr.2004.05.026

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (2001). Responses of neurons in macaque area V4 during memory-guided visual search. *Cerebral Cortex (New York, N.Y.: 1991)*, *11*(8), 761–772. https://doi.org/10.1093/cercor/11.8.761

Chennu, S., Noreika, V., Gueorguiev, D., Blenkmann, A., Kochen, S., Ibáñez, A., Owen, A. M., & Bekinschtein, T. A. (2013). Expectation and Attention in Hierarchical Auditory Prediction. *Journal of Neuroscience*, *33*(27), 11194–11205. https://doi.org/10.1523/JNEUROSCI.0114-13.2013

Christophel, T. B., Klink, P. C., Spitzer, B., Roelfsema, P. R., & Haynes, J.-D. (2017). The Distributed Nature of Working Memory. *Trends in cognitive sciences*, *0*(0). https://doi.org/10.1016/j.tics.2016.12.007

Chudasama, Y., Bussey, T. J., & Muir, J. L. (2001). Effects of selective thalamic and prelimbic cortex lesions on two types of visual discrimination and reversal learning. *The European Journal of Neuroscience*, *14*(6), 1009–1020. https://doi.org/10.1046/j.0953-816x.2001.01607.x

Cole, S. R., & Voytek, B. (2017). Brain Oscillations and the Importance of Waveform Shape. *Trends in cognitive sciences*, *0*(0). https://doi.org/10.1016/j.tics.2016.12.008

Collins, R. L. (1985). On the Inheritance of Direction and Degree of Asymmetry. In *Cerebral Lateralization in Nonhuman Species* (pp. 41–71). Elsevier. https://doi.org/10.1016/B978-0-12-286480-3.50009-4

Cooper, P. S., Wong, A. S. W., Fulham, W. R., Thienel, R., Mansfield, E., Michie, P. T., & Karayanidis, F. (2015). Theta frontoparietal connectivity associated with proactive and reactive cognitive control processes. *NeuroImage*, *108*, 354–363. https://doi.org/10.1016/j.neuroimage.2014.12.028

Cosman, J. D., Lowe, K. A., Zinke, W., Woodman, G. F., & Schall, J. D. (2018). Prefrontal Control of Visual Distraction. *Current biology: CB*, *28*(3), 414–420.e3. https://doi.org/10.1016/j.cub.2017.12.023

Courtiol, E., & Wilson, D. A. (2016). Neural Representation of Odor-Guided Behavior in the Rat Olfactory Thalamus. *Journal of Neuroscience*, *36*(22), 5946–5960. https://doi.org/10.1523/JNEUROSCI.0533-16.2016

Dalmaso, M., Castelli, L., & Galfano, G. (2019). Anticipation of cognitive conflict is reflected in microsaccades: Evidence from a cued-flanker task. *Journal of Eye Movement Research*, *12*(6), 10.16910/jemr.12.6.3. https://doi.org/10.16910/jemr.12.6.3

Danet, L., Barbeau, E. J., Eustache, P., Planton, M., Raposo, N., Sibon, I., Albucher, J.-F., Bonneville, F., Peran, P., & Pariente, J. (2015). Thalamic amnesia after infarct: The role of the mammillothalamic tract and mediodorsal nucleus. *Neurology*, *85*(24), 2107–2115. https://doi.org/10.1212/WNL.0000000000002226

de Hemptinne, C., Lefèvre, P., & Missal, M. (2006). Influence of Cognitive Expectation on the Initiation of Anticipatory and Visual Pursuit Eye Movements in the Rhesus Monkey. *Journal of Neurophysiology*, *95*(6), 3770–3782. https://doi.org/10.1152/jn.00007.2006

Denison, R. N., Yuval-Greenberg, S., & Carrasco, M. (2019). Directing Voluntary Temporal Attention Increases Fixational Stability. *Journal of Neuroscience*, *39*(2), 353–363. https://doi.org/10.1523/JNEUROSCI.1926-18.2018

Diamond, A. (2013). Executive Functions. *Annual review of psychology*, *64*, 135–168. https://doi.org/10.1146/annurev-psych-113011-143750

Dias, R., Robbins, T. W., & Roberts, A. C. (1996). Primate analogue of the Wisconsin Card Sorting Test: Effects of excitotoxic lesions of the prefrontal cortex in the marmoset. *Behavioral Neuroscience*, *110*(5), 872–886. https://doi.org/10.1037//0735-7044.110.5.872

Dias, R., Robbins, T. W., & Roberts, A. C. (1997). Dissociable forms of inhibitory control within prefrontal cortex with an analog of the Wisconsin Card Sort Test: Restriction to novel situations and independence from "on-line" processing. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *17*(23), 9285–9297.

di Pellegrino, G., & Wise, S. P. (1993). Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *13*(3), 1227–1243.

Dolleman-van der Weel, M. J., Morris, R. G. M., & Witter, M. P. (2009). Neurotoxic lesions of the thalamic reuniens or mediodorsal nucleus in rats affect non-mnemonic aspects of watermaze learning. *Brain Structure and Function*, *213*(3), 329–342. https://doi.org/10.1007/s00429-008-0200-6

Dotan, D., Pinheiro-Chagas, P., Al Roumi, F., & Dehaene, S. (2019). Track It to Crack It: Dissecting Processing Stages with Finger Tracking. *Trends in cognitive sciences*, *23*(12), 1058–1070. https://doi.org/10.1016/j.tics.2019.10.002

Dotson, N. M., Salazar, R. F., & Gray, C. M. (2014). Frontoparietal correlation dynamics reveal interplay between integration and segregation during visual working memory. *The Journal of Neuroscience*, *34*(41), 13600–13613. https://doi.org/10.1523/JNEUROSCI.1961-14.2014

Durstewitz, D., Seamans, J. K., & Sejnowski, T. J. (2000). Neurocomputational models of working memory. *Nature Neuroscience*, *3 Suppl*, 1184–1191. https://doi.org/10.1038/81460

Duvenaud, D. (2014). Automatic model construction with Gaussian processes. https://doi.org/10.17863/CAM.14087

El-Shamayleh, Y., & Horwitz, G. D. (2019). Primate optogenetics: Progress and prognosis. *Proceedings of the National Academy of Sciences*, *116*(52), 26195–26203. https://doi.org/10.1073/pnas.1902284116

Emiliani, V., Cohen, A. E., Deisseroth, K., & Häusser, M. (2015). All-Optical Interrogation of Neural Circuits. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(41), 13917–13926. https://doi.org/10.1523/JNEUROSCI.2916-15.2015

Escoufier, Y. (1973). Le Traitement des Variables Vectorielles. *Biometrics*, *29*(4), 751. https://doi.org/10.2307/2529140

Everling, S., Tinsley, C. J., Gaffan, D., & Duncan, J. (2002). Filtering of neural signals by focused attention in the monkey prefrontal cortex. *Nature Neuroscience*, *5*(7), 671–676. https://doi.org/10.1038/nn874

Fell, J., & Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature Reviews Neuroscience*, *12*(2), 105–118. https://doi.org/10.1038/nrn2979

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, *1*(1), 1–47. https://doi.org/10.1093/cercor/1.1.1-a

Fellrath, J., Mottaz, A., Schnider, A., Guggisberg, A. G., & Ptak, R. (2016). Theta-band functional connectivity in the dorsal fronto-parietal network predicts goal-directed attention. *Neuropsychologia*, *92*, 20–30. https://doi.org/10.1016/j.neuropsychologia.2016.07.012

Fiebach, C. J., Rissman, J., & D'Esposito, M. (2006). Modulation of inferotemporal cortex activation during verbal working memory maintenance. *Neuron*, *51*(2), 251–261. https://doi.org/10.1016/j.neuron.2006.06.007

Fishbach, A., Roy, S. A., Bastianen, C., Miller, L. E., & Houk, J. C. (2005). Kinematic properties of on-line error corrections in the monkey. *Experimental Brain Research*, *164*(4), 442–457. https://doi.org/10.1007/s00221-005-2264-3

Fishbach, A., Roy, S. A., Bastianen, C., Miller, L. E., & Houk, J. C. (2007). Deciding when and how to correct a movement: Discrete submovements as a decision making process. *Experimental Brain Research*, *177*(1), 45–63. https://doi.org/10.1007/s00221-006-0652-y

Flash, T., & Hogan, N. (1985). The coordination of arm movements: An experimentally confirmed mathematical model. *The Journal of Neuroscience*, *5*(7), 1688–1703. https://doi.org/10.1523/JNEUROSCI.05-07-01688.1985

Flash, T., & Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Current Opinion in Neurobiology*, *15*(6), 660–666. https://doi.org/10.1016/j.conb.2005.10.011

Foster, C., Sheng, W.-A., Heed, T., & Ben Hamed, S. (2022). The macaque ventral intraparietal area has expanded into three homologue human parietal areas. *Progress in Neurobiology*, *209*, 102185. https://doi.org/10.1016/j.pneurobio.2021.102185

Freedman, D. J., Riesenhuber, M., Poggio, T., & Miller, E. K. (2001). Categorical representation of visual stimuli in the primate prefrontal cortex. *Science (New York, N.Y.)*, *291*(5502), 312–316. https://doi.org/10.1126/science.291.5502.312

Friedrich, R. W., Jacobson, G. A., & Zhu, P. (2010). Circuit neuroscience in zebrafish. *Current biology: CB*, *20*(8), R371–381. https://doi.org/10.1016/j.cub.2010.02.039

Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, *9*(10), 474–480. https://doi.org/10.1016/j.tics.2005.08.011

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*(2), 331–349. https://doi.org/10.1152/jn.1989.61.2.331

Fuster, J. M. (1973). Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *Journal of Neurophysiology*, *36*(1), 61–78. https://doi.org/10.1152/jn.1973.36.1.61

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science (New York, N.Y.)*, *173*(3997), 652–654. https://doi.org/10.1126/science.173.3997.652

Fuster, J. M., Bauer, R. H., & Jervey, J. P. (1985). Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain Research*, *330*(2), 299–307. https://doi.org/10.1016/0006-8993(85)90689-4

Fuster, J. M., Bodner, M., & Kroger, J. K. (2000). Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature*, *405*(6784), 347–351. https://doi.org/10.1038/35012613

Fuster, J. M., & Jervey, J. P. (1982). Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *Journal of Neuroscience*, *2*(3), 361–375. https://doi.org/10.1523/JNEUROSCI.02-03-00361.1982

Fuster, J. M. (2022). Cognitive Networks (Cognits) Process and Maintain Working Memory. *Frontiers in Neural Circuits*, *15*. Retrieved September 27, 2022, from https://www.frontiersin.org/articles/10.3389/fncir.2021.790691

Fuster, J. M. (2015). *The Prefrontal Cortex*. elsevier.

Fuster, J. M., Bauer, R. H., & Jervey, J. P. (1982). Cellular discharge in the dorsolateral prefrontal cortex of the monkey in cognitive tasks. *Experimental Neurology*, *77*(3), 679–694. https://doi.org/10.1016/0014-4886(82)90238-2

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, *94*(5), 978–984. https://doi.org/10.1016/j.neuron.2017.05.025

Gallego, J. A., Perich, M. G., Chowdhury, R. H., Solla, S. A., & Miller, L. E. (2018). A stable, long-term cortical signature underlying consistent behavior. *bioRxiv*, 447441. https://doi.org/10.1101/447441

Gao, P., Trautmann, E., Yu, B. M., Santhanam, G., Ryu, S., Shenoy, K., & Ganguli, S. (2017). A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 1–50. https://doi.org/10.1101/214262

Gao, Z., Davis, C., Thomas, A. M., Economo, M. N., Abrego, A. M., Svoboda, K., De Zeeuw, C. I., & Li, N. (2018). A cortico-cerebellar loop for motor planning. *Nature*, *563*(7729), 113–116. https://doi.org/10.1038/s41586-018-0633-x

Gibbs, R. A., Rogers, J., Katze, M. G., Bumgarner, R., Weinstock, G. M., Mardis, E. R., Remington, K. A., Strausberg, R. L., Venter, J. C., Wilson, R. K., Batzer, M. A., Bustamante, C. D., Eichler, E. E., Hahn, M. W., Hardison, R. C., Makova, K. D., Miller, W., Milosavljevic, A., Palermo, R. E., … Zwieg, A. S. (2007). Evolutionary and Biomedical Insights from the Rhesus Macaque Genome. *Science*, *316*(5822), 222–234. https://doi.org/10.1126/science.1139247

Glick, S., Goldfarb, T., & Jarvik, M. (1969). Recovery of delayed matching performance following lateral frontal lesions in monkeys. *Communications in Behavioral Biology*, *3*, 299–303.

Goel, V., & Grafman, J. (1995). Are the frontal lobes implicated in "planning" functions? Interpreting data from the Tower of Hanoi. *Neuropsychologia*, *33*(5), 623–642. https://doi.org/10.1016/0028-3932(95)90866-P

Goldberg, E., & Podell, K. (2000). Adaptive decision making, ecological validity, and the frontal lobes. *Journal of Clinical and Experimental Neuropsychology*, *22*(1), 56–68. https://doi.org/10.1076/1380-3395(200002)22:1;1-8;FT056

Goldman, P. S., & Rosvold, H. E. (1970). Localization of function within the dorsolateral prefrontal cortex of the rhesus monkey. *Experimental Neurology*, *27*(2), 291–304. https://doi.org/10.1016/0014-4886(70)90222-0

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron*, *14*(3), 477–485. https://doi.org/10.1016/0896-6273(95)90304-6

Goltstein, P. M., Reinert, S., Bonhoeffer, T., & Hübener, M. (2021). Mouse visual cortex areas represent perceptual and semantic features of learned visual categories. *Nature Neuroscience*, 1–11. https://doi.org/10.1038/s41593-021-00914-5

Gommers, R., Virtanen, P., Burovski, E., Weckesser, W., Oliphant, T. E., Cournapeau, D., Haberland, M., Reddy, T., Alexbrc, Peterson, P., Nelson, A., Wilson, J., Endolith, Mayorov, N., Polat, I., Van Der Walt, S., Laxalde, D., Brett, M., Larson, E., … Kern, R. (2022, February 5). *Scipy/scipy: SciPy 1.8.0* (Version v1.8.0). https://doi.org/10.5281/ZENODO.5979747

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, *37*(3), 424. https://doi.org/10.2307/1912791

Graving, J. M., & Couzin, I. D. (2020). VAE-SNE: A deep generative model for simultaneous dimensionality reduction and clustering. *bioRxiv*, 2020.07.17.207993. https://doi.org/10.1101/2020.07.17.207993

Gray, D. T., & Barnes, C. A. (2019). Experiments in macaque monkeys provide critical insights into age-associated changes in cognitive and sensory function. *Proceedings of the National Academy of Sciences*, *116*(52), 26247–26254. https://doi.org/10.1073/pnas.1902279116

Grienberger, C., & Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, *73*(5), 862–885.

Grisel, O., Mueller, A., Lars, Gramfort, A., Louppe, G., Prettenhofer, P., Blondel, M., Niculae, V., Nothman, J., Joly, A., Fan, T. J., Vanderplas, J., Kumar, M., Lemaitre, G., Qin, H., Hug, N., Estève, L., Varoquaux, N., Layton, R., … Eustache. (2021, October 25). *Scikit-learn/scikit-learn: Scikit-learn 1.0.1* (Version 1.0.0). https://doi.org/10.5281/ZENODO.5596244

Guo, Z. V., Hires, S. A., Li, N., O'Connor, D. H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D., Peron, S., Xu, N.-l., Cox, J., & Svoboda, K. (2014). Procedures for Behavioral Experiments in Head-Fixed Mice (S. A. Simon, Ed.). *PLOS ONE*, *9*(2), e88678. https://doi.org/10.1371/journal.pone.0088678

Guo, Z. V., Inagaki, H. K., Daie, K., Druckmann, S., Gerfen, C. R., & Svoboda, K. (2017). Maintenance of persistent activity in a frontal thalamocortical loop. *Nature*, *39*, 1062. https://doi.org/10.1038/nature22324

Gupta, N., Singh, S. S., & Stopfer, M. (2016). Oscillatory integration windows in neurons. *Nature Communications*, *7*, 13808. https://doi.org/10.1038/ncomms13808

Halassa, M. (2017, October 18). *Thalamic amplification of cortical connectivity sustains attentional control*. LMU Biocenter, B01.019, Small Lecture Hall Großhaderner-str. 2 82152 Martinsried.

Hamm, J. P., Shymkiv, Y., Han, S., Yang, W., & Yuste, R. (2021). Cortical ensembles selective for context. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(14), e2026179118. https://doi.org/10.1073/pnas.2026179118

Hardung, S., Epple, R., Jäckel, Z., Eriksson, D., Uran, C., Senn, V., Gibor, L., Yizhar, O., & Diester, I. (2017). A Functional Gradient in the Rodent Prefrontal Cortex Supports Behavioral Inhibition. *Current Biology*, *27*(4), 549–555. https://doi.org/10.1016/j.cub.2016.12.052

Harlow, J. (1868). Recovery from passage of an iron bar through the head. B. Mass. *Med. Soc*, *2*, 3–20.

Harmony, T. (2013). The functional significance of delta oscillations in cognitive processing, 1–10. https://doi.org/10.3389/fnint.2013.00083/abstract

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., … Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

Harris, J. A., Mihalas, S., Hirokawa, K. E., Whitesell, J. D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., Feiner, A., Feng, D., Gaudreault, N., Gerfen, C. R., Graddis, N., Groblewski, P. A., Henry, A. M., Ho, A., Howard, R., ... Zeng, H. (2019). Hierarchical organization of cortical and thalamic connectivity. *Nature*, *575*(7781), 195–202. https://doi.org/10.1038/s41586-019-1716-z

Helfrich, R. F., & Knight, R. T. (2016). Oscillatory Dynamics of Prefrontal Cognitive Control. *Trends in cognitive sciences*, *0*(0). https://doi.org/10.1016/j.tics.2016.09.007

Heusser, A. C., Poeppel, D., Ezzyat, Y., & Davachi, L. (2016). Episodic sequence memory is supported by a theta-gamma phase code. *Nature Neuroscience*. https://doi.org/10.1038/nn.4374

Hill, D. N., Mehta, S. B., & Kleinfeld, D. (2011). Quality Metrics to Accompany Spike Sorting of Extracellular Signals. *Journal of Neuroscience*, *31*(24), 8699–8705. https://doi.org/10.1523/JNEUROSCI.0971-11.2011

Hoffmann, J., & Sebald, A. (2005). Local contextual cuing in visual search. *Experimental Psychology*, *52*(1), 31–38. https://doi.org/10.1027/1618-3169.52.1.31

Holmes, E., Parr, T., Griffiths, T. D., & Friston, K. J. (2021). Active inference, selective attention, and the cocktail party problem. *Neuroscience & Biobehavioral Reviews*, *131*, 1288–1304. https://doi.org/10.1016/j.neubiorev.2021.09.038

Hoyer, S., Roos, M., Hamman, J., Keewis, Cherian, D., Fitzgerald, C., Fujii, K., Hauser, M., Maussion, F., Crusaderky, Clark, S., Kleeman, A., Kluyver, T., Munroe, J., Nicholas, T., Amici, A., Illviljan, Barghini, A., Banihirwe, A., ... Mühlbauer, K. (2021, July 23). *Pydata/xarray: V0.19.0* (Version v0.19.0). https://doi.org/10.5281/ZENODO.5130718

Hubel, D. H., & Wiesel, T. N. (1979). Brain mechanisms of vision. *Scientific American*, *241*, 150–162. https://doi.org/10.1038/scientificamerican0979-150

Huda, R., Sipe, G. O., Adam, E., Breton-Provencher, V., Pho, G., Gunter, L., Wickersham, I. R., & Sur, M. (2018). Bidirectional control of goal-oriented action selection by distinct prefrontal cortex circuits. *bioRxiv*, 307009. https://doi.org/10.1101/307009

Huettel, S. A., Song, A. W., & McCarthy, G. (2005). Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(13), 3304–3311. https://doi.org/10.1523/JNEUROSCI.5070-04.2005

Hummos, A., Wang, B. A., Drammis, S., Halassa, M. M., & Pleger, B. (2022). Thalamic regulation of frontal interactions in human cognitive flexibility. *PLoS Computational Biology*, *18*(9), e1010500. https://doi.org/10.1371/journal.pcbi.1010500

Hwang, J., Mitz, A. R., & Murray, E. A. (2019). NIMH MonkeyLogic: Behavioral control and data acquisition in MATLAB. *Journal of Neuroscience Methods*, *323*, 13–21. https://doi.org/10.1016/j.jneumeth.2019.05.002

Inagaki, H. K., Chen, S., Ridder, M. C., Sah, P., Li, N., Yang, Z., Hasanbegovic, H., Gao, Z., Gerfen, C. R., & Svoboda, K. (2022). A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell*, *185*(6), 1065–1081.e23. https://doi.org/10.1016/j.cell.2022.02.006

Ingvar, D. H. (1985). "Memory of the future": An essay on the temporal organization of conscious awareness. *Human neurobiology*, *4*(3), 127–136.

Ison, J. R., Allen, P. D., & O'Neill, W. E. (2007). Age-Related Hearing Loss in C57BL/6J Mice has both Frequency-Specific and Non-Frequency-Specific Components that Produce a Hyperacusis-Like Exaggeration of the Acoustic Startle Reflex. *Journal of the Association for Research in Otolaryngology*, *8*(4), 539–550. https://doi.org/10.1007/s10162-007-0098-3

Isseroff, A., Rosvold, H. E., Galkin, T. W., & Goldman-Rakic, P. S. (1982). Spatial memory impairments following damage to the mediodorsal nucleus of the thalamus in rhesus monkeys. *Brain Research*, *232*(1), 97–113. https://doi.org/10.1016/0006-8993(82)90613-8

Jacob, S. N., Hähnke, D., & Nieder, A. (2018). Structuring of Abstract Working Memory Content by Fronto-parietal Synchrony in Primate Cortex. *Neuron*, *99*(3), 588–597.e5. https://doi.org/10.1016/j.neuron.2018.07.025

Jacob, S. N., & Nieder, A. (2014). Complementary roles for primate frontal and parietal cortex in guarding working memory from distractor stimuli. *Neuron*, *83*(1), 226–237. https://doi.org/10.1016/j.neuron.2014.05.009

Jensen, O., & Tesche, C. D. (2002). Frontal theta activity in humans increases with memory load in a working memory task. *The European Journal of Neuroscience*, *15*(8), 1395–1399. https://doi.org/10.1046/j.1460-9568.2002.01975.x

Johnson, E. L., Dewar, C. D., Solbakk, A.-K., Endestad, T., Meling, T. R., & Knight, R. T. (2017). Bidirectional Frontoparietal Oscillatory Systems Support Working Memory. *Current Biology*, *27*(12), 1829–1835.e4. https://doi.org/10.1016/j.cub.2017.05.046

Jones, C. M., Gray, R., Spence, C., & Tan, H. Z. (2008). Directing visual attention with spatially informative and spatially noninformative tactile cues. *Experimental Brain Research*, *186*(4), 659–669. https://doi.org/10.1007/s00221-008-1277-0

Kaczmarczyk, L., & Jackson, W. S. (2015). Astonishing advances in mouse genetic tools for biomedical research. *Swiss Medical Weekly*, *145*, w14186. https://doi.org/10.4414/smw.2015.14186

Kajikawa, Y., & Schroeder, C. E. (2011). How local is the local field potential? *Neuron*, *72*(5), 847–858. https://doi.org/10.1016/j.neuron.2011.09.029

Kandel, E. R., Koester, J. D., Mack, S. H., & Siegelbaum, S. A. (Eds.). (2021). *Principles of neural science* (6th ed.). McGraw Hill.

Karnath, H. O., & Wallesch, C. W. (1992). Inflexibility of mental planning: A characteristic disorder with prefrontal lobe lesions? *Neuropsychologia*, *30*(11), 1011–1016. https://doi.org/10.1016/0028-3932(92)90052-N

Kawagoe, T., Tamura, R., Uwano, T., Asahi, T., Nishijo, H., Eifuku, S., & Ono, T. (2007). Neural correlates of stimulus-reward association in the rat mediodorsal thalamus. *Neuroreport*, *18*(7), 683–688. https://doi.org/10.1097/WNR.0b013e3280bef9a6

Kazama, H. (2015). Systems neuroscience in Drosophila: Conceptual and technical advantages. *Neuroscience*, *296*, 3–14. https://doi.org/10.1016/j.neuroscience.2014.06.035

Kéri, S., Decety, J., Roland, P. E., & Gulyás, B. (2004). Feature uncertainty activates anterior cingulate cortex. *Human Brain Mapping*, *21*(1), 26–33. https://doi.org/10.1002/hbm.10150

Kim, C. K., Adhikari, A., & Deisseroth, K. (2017). Integration of optogenetics with complementary methodologies in systems neuroscience. *Nature Reviews Neuroscience*, *18*(4), 222–235. https://doi.org/10.1038/nrn.2017.15

Knöpfel, T., & Song, C. (2019). Optical voltage imaging in neurons: Moving from technology development to practical tool. *Nature Reviews Neuroscience*, *20*(12), 719–727. https://doi.org/10.1038/s41583-019-0231-4

Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F., Qi, X.-L., Romo, R., Uchida, N., Machens, C. K., & van Rossum, M. C. (2016). Demixed principal component analysis of neural population data. *eLife*, *5*, e10989. https://doi.org/10.7554/eLife.10989

Kowler, E., Rubinstein, J. F., Santos, E. M., & Wang, J. (2019). Predictive Smooth Pursuit Eye Movements. *Annual Review of Vision Science*, *5*, 223–246. https://doi.org/10.1146/annurev-vision-091718-014901

Kubota, K., Tonoike, M., & Mikami, A. (1980). Neuronal activity in the monkey dorsolateral prefrontal cortex during a discrimination task with delay. *Brain Research*, *183*(1), 29–42. https://doi.org/10.1016/0006-8993(80)90117-1

Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J. Z., & Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature*, *498*(7454), 363–366. https://doi.org/10.1038/nature12176

Lachaux, J.-P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, *8*(4), 194–208. https://doi.org/10.1002/(SICI)1097-0193(1999)8:4⟨194::AID-HBM4⟩3.0.CO;2-C

Lam, S. K., Pitrou, A., Stuartarchibald, Florisson, M., Seibert, S., Markall, G., Anderson, T. A., Haenel, V., Rjenc29, Bourque, J., Luk-F-A, Meurer, A., Leobas, G., Oliphant, T. E., Densmirn, Njwhite, Pronovost, E., Seefeld, S., Grecco, H., … Sanjay, R. (2020, December 17). *Numba/numba: Version 0.52.0* (Version 0.52.0). https://doi.org/10.5281/ZENODO.4343231

Laubach, M., Amarante, L. M., Swanson, K., & White, S. R. (2018). What, If Anything, Is Rodent Prefrontal Cortex? *eneuro*, *5*(5), ENEURO.0315–18.2018. https://doi.org/10.1523/ENEURO.0315-18.2018

Le Merre, P., Ährlund-Richter, S., & Carlén, M. (2021). The mouse prefrontal cortex: Unity in diversity. *Neuron*, *109*(12), 1925–1944. https://doi.org/10.1016/j.neuron.2021.03.035

Léger, D., Debellemaniere, E., Rabat, A., Bayon, V., Benchenane, K., & Chennaoui, M. (2018). Slow-wave sleep: From the cell to the clinic. *Sleep Medicine Reviews*, *41*, 113–132. https://doi.org/10.1016/j.smrv.2018.01.008

Lennert, T., & Martinez-Trujillo, J. (2011). Strength of response suppression to distracter stimuli determines attentional-filtering performance in primate prefrontal neurons. *Neuron*, *70*(1), 141–152. https://doi.org/10.1016/j.neuron.2011.02.041

Levy, R., & Goldman-Rakic, P. S. (1999). Association of Storage and Processing Functions in the Dorsolateral Prefrontal Cortex of the Nonhuman Primate. *Journal*

of Neuroscience, *19*(12), 5149–5158. https://doi.org/10.1523/JNEUROSCI.19-12-05149.1999

Lewis, C. M., Bosman, C. A., Brunet, N. M., Lima, B., Roberts, M. J., Womelsdorf, T., De Weerd, P., Neuenschwander, S., Singer, W., & Fries, P. (2016). Two frequency bands contain the most stimulus-related information in visual cortex. *bioRxiv*, 049718. https://doi.org/10.1101/049718

Lewis, J. W., & Van Essen, D. C. (2000). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *The Journal of Comparative Neurology*, *428*(1), 112–137. https://doi.org/10.1002/1096-9861(20001204)428:1⟨112::aid-cne8⟩3.0.co;2-9

Liebe, S., Hoerzer, G. M., Logothetis, N. K., & Rainer, G. (2012). Theta coupling between V4 and prefrontal cortex predicts visual short-term memory performance. *Nature Neuroscience*, *15*(3), 456-62- S1–2. https://doi.org/10.1038/nn.3038

Lim, A., Eng, V., Osborne, C., Janssen, S. M. J., & Satel, J. (2019). Inhibitory and Facilitatory Cueing Effects: Competition between Exogenous and Endogenous Mechanisms. *Vision*, *3*(3), 40. https://doi.org/10.3390/vision3030040

Lindén, H., Hagen, E., Leski, S., Norheim, E., Pettersen, K., & Einevoll, G. (2014). LFPy: A tool for biophysical simulation of extracellular potentials generated by detailed model neurons. *Frontiers in Neuroinformatics*, *7*. Retrieved September 28, 2022, from https://www.frontiersin.org/articles/10.3389/fninf.2013.00041

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R., Chen, Z., Chen, Y., & Li, C. T. (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science*, *346*(6208), 458–463. https://doi.org/10.1126/science.1256573

Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., & Miller, E. K. (2016). Gamma and Beta Bursts Underlie Working Memory. *Neuron*, 1–14. https://doi.org/10.1016/j.neuron.2016.02.028

MacKinnon, D. P., Fairchild, A. J., & Fritz, M. S. (2007). Mediation Analysis. *Annual Review of Psychology*, *58*(1), 593–614. https://doi.org/10.1146/annurev.psych.58.110405.085542

Malmo, R. B. (1942). Interference factors in delayed response in monkeys after removal of frontal lobes. *Journal of Neurophysiology*, *5*(4), 295–308. https://doi.org/10.1152/jn.1942.5.4.295

Manns, M., Basbasse, Y. E., Freund, N., & Ocklenburg, S. (2021). Paw preferences in mice and rats: Meta-analysis. *Neuroscience and Biobehavioral Reviews*, *127*, 593–606. https://doi.org/10.1016/j.neubiorev.2021.05.011

Mantini, D., Corbetta, M., Romani, G. L., Orban, G. A., & Vanduffel, W. (2013). Evolutionarily Novel Functional Networks in the Human Brain? *Journal of Neuroscience*, *33*(8), 3259–3275. https://doi.org/10.1523/JNEUROSCI.4392-12.2013

Markov, N. T., Ercsey-Ravasz, M. M., Ribeiro Gomes, A. R., Lamy, C., Magrou, L., Vezoli, J., Misery, P., Falchier, A., Quilodran, R., Gariel, M. A., Sallet, J., Gamanut, R., Huissoud, C., Clavagnier, S., Giroud, P., Sappey-Marinier, D., Barone, P., Dehay, C., Toroczkai, Z., … Kennedy, H. (2014). A weighted and directed in-

terareal connectivity matrix for macaque cerebral cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, *24*(1), 17–36. https://doi.org/10.1093/cercor/bhs270

Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., Barone, P., Dehay, C., Knoblauch, K., & Kennedy, H. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *The Journal of Comparative Neurology*, *522*(1), 225–259. https://doi.org/10.1002/cne.23458

Marshuetz, C. (2005). Order Information in Working Memory: An Integrative Review of Evidence From Brain and Behavior. *Psychological Bulletin*, *131*(3), 323–339. https://doi.org/10.1037/0033-2909.131.3.323

Maunsell, J. H., & van Essen, D. C. (1983). The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *3*(12), 2563–2586.

Meltzer, J. A., Zaveri, H. P., Goncharova, I. I., Distasio, M. M., Papademetris, X., Spencer, S. S., Spencer, D. D., & Constable, R. T. (2008). Effects of working memory load on oscillatory power in human intracranial EEG. *Cerebral Cortex (New York, N.Y.: 1991)*, *18*(8), 1843–1855. https://doi.org/10.1093/cercor/bhm213

Mesulam, M.-M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Annals of Neurology*, *28*(5), 597–613. https://doi.org/10.1002/ana.410280502

Meyers, E., & Kreiman, G. (n.d.). Tutorial on Pattern Classification in Cell Recording. In *Visual Population Codes*.

Meyers, E. M. (2018). Dynamic population coding and its relationship to working memory. *Journal of Neurophysiology*, *120*(5), 2260–2268. https://doi.org/10.1152/jn.00225.2018

Michalareas, G., Vezoli, J., van Pelt, S., Schoffelen, J.-M., Kennedy, H., & Fries, P. (2016). Alpha-Beta and Gamma Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron*, *89*(2), 384–397. https://doi.org/10.1016/j.neuron.2015.12.018

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202. https://doi.org/10.1146/annurev.neuro.24.1.167

Miller, E. K., & Wallis, J. D. (2009). Executive function and higher-order cognition: Definition and neural substrates. *Encyclopedia of neuroscience*, *4*(99-104). https://doi.org/10.1016/B978-008045046-9.00418-6

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *Journal of Neuroscience*, *16*(16), 5154–5167. https://doi.org/10.1523/JNEUROSCI.16-16-05154.1996

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*(2), 463–475. https://doi.org/10.1016/j.neuron.2018.09.023

Miller, R. L. A., Francoeur, M. J., Gibson, B. M., & Mair, R. G. (2017). Mediodorsal Thalamic Neurons Mirror the Activity of Medial Prefrontal Neurons Responding

to Movement and Reinforcement during a Dynamic DNMTP Task. *eNeuro*, *4*(5), ENEURO.0196–17.2017. https://doi.org/10.1523/ENEURO.0196-17.2017

Milner, B., Petrides, M., & Smith, M. L. (1985). Frontal lobes and the temporal organization of memory. *Human Neurobiology*, *4*(3), 137–142.

Milner, B. (1963). Effects of Different Brain Lesions on Card Sorting: The Role of the Frontal Lobes. *Archives of Neurology*, *9*(1), 90–100. https://doi.org/10.1001/archneur.1963.00460070100010

Milner, B. (1964). Some effects of frontal lobectomy in man. *The frontal granular cortex and behavior.*, 313–334.

Mishkin, M. (1957). Effects of small frontal lesions on delayed alternation in monkeys. *Journal of Neurophysiology*, *20*(6), 615–622. https://doi.org/10.1152/jn.1957.20.6.615

Mitchell, A. S. (2015). The mediodorsal thalamus as a higher order thalamic relay nucleus important for learning and decision-making. *Neuroscience and biobehavioral reviews*, *54*, 76–88. https://doi.org/10.1016/j.neubiorev.2015.03.001

Montijn, J. S., Seignette, K., Howlett, M. H., Cazemier, J. L., Kamermans, M., Levelt, C. N., & Heimel, J. A. (2021). A parameter-free statistical test for neuronal responsiveness. *eLife*, *10*, e71969. https://doi.org/10.7554/eLife.71969

Moon, H.-S., Baek, J., & Seo, J. (2016). Effect of Redundant Haptic Information on Task Performance during Visuo-Tactile Task Interruption and Recovery. *Frontiers in Psychology*, *7*, 1924. https://doi.org/10.3389/fpsyg.2016.01924

Morissette, M., & Di Paolo, T. (2018). Non-human primate models of PD to test novel therapies. *Journal of Neural Transmission (Vienna, Austria: 1996)*, *125*(3), 291–324. https://doi.org/10.1007/s00702-017-1722-y

Mukherjee, A., Lam, N. H., Wimmer, R. D., & Halassa, M. M. (2021). Thalamic circuits for independent control of prefrontal signal and noise. *Nature*, *600*(7887), 100–104. https://doi.org/10.1038/s41586-021-04056-3

Murray, J. D., Bernacchia, A., Roy, N. A., Constantinidis, C., Romo, R., & Wang, X.-J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(2), 394–399. https://doi.org/10.1073/pnas.1619449114

Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S., & Churchland, A. K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience*, *22*(10), 1677–1686. https://doi.org/10.1038/s41593-019-0502-4

Nácher, V., Ledberg, A., Deco, G., & Romo, R. (2013). Coherent delta-band oscillations between cortical areas correlate with decision making. *Proceedings of the National Academy of Sciences*, *110*(37), 15085–15090. https://doi.org/10.1073/pnas.1314681110

Nakajima, M., & Halassa, M. M. (2017). Thalamic control of functional cortical connectivity. *Current Opinion in Neurobiology*, *44*, 127–131. https://doi.org/10.1016/j.conb.2017.04.001

Nakajima, M., Schmitt, L. I., & Halassa, M. M. (2019). Prefrontal Cortex Regulates Sensory Filtering through a Basal Ganglia-to-Thalamus Pathway. *Neuron*. https://doi.org/10.1016/j.neuron.2019.05.026

Nassi, J. J., Avery, M. C., Cetin, A. H., Roe, A. W., & Reynolds, J. H. (2015). Optogenetic Activation of Normalization in Alert Macaque Visual Cortex. *Neuron*, *86*(6), 1504–1517. https://doi.org/10.1016/j.neuron.2015.05.040

Nassi, J. J., Cepko, C. L., Born, R. T., & Beier, K. T. (2015). Neuroanatomy goes viral! *Frontiers in Neuroanatomy*, *9*. Retrieved September 29, 2022, from https://www.frontiersin.org/articles/10.3389/fnana.2015.00080

Ness, T. V., Halnes, G., Næss, S., Pettersen, K. H., & Einevoll, G. T. (2021, May 2). *Computing extracellular electric potentials from neuronal simulations*. arXiv: 2006.16630 [q-bio]. Retrieved September 28, 2022, from http://arxiv.org/abs/2006.16630

Newville, M., Otten, R., Nelson, A., Ingargiola, A., Stensitzki, T., Allan, D., Fox, A., Carter, F., Michał, Pustakhod, D., Lneuhaus, Weigand, S., Osborn, R., Glenn, Deil, C., Mark, Hansen, A. L. R., Pasquevich, G., Foks, L., … Maier, B. F. (2021, February 7). *Lmfit/lmfit-py 1.0.2* (Version 1.0.2). https://doi.org/10.5281/ZENODO.4516651

Nieder, A. (2012). Supramodal numerosity selectivity of neurons in primate prefrontal and posterior parietal cortices. *Proceedings of the National Academy of Sciences*, *109*(29), 11860–11865. https://doi.org/10.1073/pnas.1204580109

Nieder, A. (2016). The neuronal code for number. *Nature Reviews. Neuroscience*, *17*(6), 366–382. https://doi.org/10.1038/nrn.2016.40

Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual review of neuroscience*, *32*(1), 185–208.

Nieder, A., Diester, I., & Tudusciuc, O. (2006). Temporal and spatial enumeration processes in the primate parietal cortex. *Science (New York, N.Y.)*, *313*(5792), 1431–1435. https://doi.org/10.1126/science.1130308

Nieder, A., Freedman, D. J., & Miller, E. K. (2002). Representation of the Quantity of Visual Items in the Primate Prefrontal Cortex. *Science*, *297*(5587), 1708–1711. https://doi.org/10.1126/science.1072493

Nieder, A., & Miller, E. K. (2004a). Analog numerical representations in rhesus monkeys: Evidence for parallel processing. *Journal of cognitive neuroscience*, *16*(5), 889–901.

Nieder, A., & Miller, E. K. (2004b). A parieto-frontal network for visual numerical information in the monkey. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(19), 7457–7462. https://doi.org/10.1073/pnas.0402239101

Nolte, G., Ziehe, A., Nikulin, V. V., Schlögl, A., Krämer, N., Brismar, T., & Müller, K.-R. (2008). Robustly Estimating the Flow Direction of Information in Complex Physical Systems. *Physical Review Letters*, *100*(23), 234101. https://doi.org/10.1103/PhysRevLett.100.234101

Noonan, M. P., Crittenden, B. M., Jensen, O., & Stokes, M. G. (2018). Selective inhibition of distracting input. *Behavioural Brain Research*, *355*, 36–47. https://doi.org/10.1016/j.bbr.2017.10.010

Norman, D. A., & Shallice, T. (1986). Attention to Action. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and Self-Regulation: Advances in Research and Theory Volume 4* (pp. 1–18). Springer US. https://doi.org/10.1007/978-1-4757-0629-1_1

Onton, J., Delorme, A., & Makeig, S. (2005). Frontal midline EEG dynamics during working memory. *NeuroImage*, *27*(2), 341–356. https://doi.org/10.1016/j.neuroimage.2005.04.014

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience*, *2011*, 1–9. https://doi.org/10.1155/2011/156869

Orban, G. A., Van Essen, D., & Vanduffel, W. (2004). Comparative mapping of higher visual areas in monkeys and humans. *Trends in Cognitive Sciences*, *8*(7), 315–324. https://doi.org/10.1016/j.tics.2004.05.009

Ouhaz, Z., Fleming, H., & Mitchell, A. S. (2018). Cognitive Functions and Neurodevelopmental Disorders Involving the Prefrontal Cortex and Mediodorsal Thalamus. *Frontiers in Neuroscience*, *12*, 33. https://doi.org/10.3389/fnins.2018.00033

Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., & Harris, K. D. (2016). Fast and accurate spike sorting of high-channel count probes with KiloSort, 4448–4456.

Papadatou-Pastou, M., Ntolka, E., Schmitz, J., Martin, M., Munafò, M. R., Ocklenburg, S., & Paracchini, S. (2020). Human handedness: A meta-analysis. *Psychological Bulletin*, *146*(6), 481–524. https://doi.org/10.1037/bul0000229

Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience*, *17*(10), 1395–1403. https://doi.org/10.1038/nn.3800

Parnaudeau, S., O'Neill, P.-K., Bolkan, S. S., Ward, R. D., Abbas, A. I., Roth, B. L., Balsam, P. D., Gordon, J. A., & Kellendonk, C. (2013). Inhibition of Mediodorsal Thalamus Disrupts Thalamofrontal Connectivity and Cognition. *Neuron*, *77*(6), 1151–1162. https://doi.org/10.1016/j.neuron.2013.01.038

Parnaudeau, S., Bolkan, S. S., & Kellendonk, C. (2018). The Mediodorsal Thalamus: An Essential Partner of the Prefrontal Cortex for Cognition. *Biological Psychiatry*, *83*(8), 648–656. https://doi.org/10.1016/j.biopsych.2017.11.008

Parthasarathy, A., Herikstad, R., Bong, J. H., Medina, F. S., Libedinsky, C., & Yen, S.-C. (2017). Mixed selectivity morphs population codes in prefrontal cortex. *Nature Neuroscience*, *20*(12), 1770–1779. https://doi.org/10.1038/s41593-017-0003-2

Parthasarathy, A., Tang, C., Herikstad, R., Cheong, L. F., Yen, S.-C., & Libedinsky, C. (2019). Time-Invariant Working Memory Representations in the Presence of Code-Morphing in the Lateral Prefrontal Cortex. *bioRxiv*, 563668. https://doi.org/10.1101/563668

Pascual-Marqui, R. D., Faber, P., Milz, P., Kochi, K., Kinoshita, T., Nishida, K., Yoshimura, M., Kitaura, Y., Ikeda, S., & Ishii, R. (2017). The Cross-Frequency Mediation Mechanism Of Intracortical Information Transactions. *bioRxiv*, 119362. https://doi.org/10.1101/119362

Pearson, J. M., Watson, K. K., & Platt, M. L. (2014). Decision making: The neuroethological turn. *Neuron*, *82*(5), 950–965. https://doi.org/10.1016/j.neuron.2014.04.037

Pesaran, B., Nelson, M. J., & Andersen, R. A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. *Nature*, *453*(7193), 406–409. https://doi.org/10.1038/nature06849

Phillips, J. M., Vinck, M., Everling, S., & Womelsdorf, T. (2014). A Long-Range Fronto-Parietal 5- to 10-Hz Network Predicts "Top-Down" Controlled Guidance in a Task-Switch Paradigm. *Cerebral Cortex*, *24*(8), 1996–2008. https://doi.org/10.1093/cercor/bht050

Pinto, L., & Dan, Y. (2015). Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed Behavior. *Neuron*, *87*(2), 437–450. https://doi.org/10.1016/j.neuron.2015.06.021

Plamondon, R., Alimi, A. M., Yergeau, P., & Leclerc, F. (1993). Modelling velocity profiles of rapid movements: A comparative study. *Biological Cybernetics*, *69*(2), 119–128. https://doi.org/10.1007/BF00226195

Polanía, R., Nitsche, M. A., Korman, C., Batsikadze, G., & Paulus, W. (2012). The Importance of Timing in Segregated Theta Phase-Coupling for Cognitive Performance. *Current Biology*, *22*(14), 1314–1318. https://doi.org/10.1016/j.cub.2012.05.021

Pulvermüller, F., & Garagnani, M. (2014). From sensorimotor learning to memory cells in prefrontal and temporal association cortex: A neurocomputational study of disembodiment. *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, *57*, 1–21. https://doi.org/10.1016/j.cortex.2014.02.015

Qi, X., Katsuki, F., Meyer, T., Rawley, J., Zhou, X., Douglas, K., & Constantinidis, C. (2010). Comparison of neural activity related to working memory in primate dorsolateral prefrontal and posterior parietal cortex. *Frontiers in Systems Neuroscience*, *4*. Retrieved August 27, 2022, from https://www.frontiersin.org/articles/10.3389/fnsys.2010.00012

Quintana, J., & Fuster, J. M. (1999). From Perception to Action: Temporal Integrative Functions of Prefrontal and Parietal Neurons. *Cerebral Cortex*, *9*(3), 213–221. https://doi.org/10.1093/cercor/9.3.213

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature*, *393*(6685), 577–579. https://doi.org/10.1038/31235

Ramirez-Cardenas, A., Moskaleva, M., & Nieder, A. (2016). Neuronal Representation of Numerosity Zero in the Primate Parieto-Frontal Number Network. *Current Biology*, *26*(10), 1285–1294. https://doi.org/10.1016/j.cub.2016.03.052

Ranganath, A., Hähnke, D., & Jacob, S. N. (2022, November 13). *Within-trial sensory feedback during instrumental learning enables rapid and efficient task acquisition.* San Diego.

Reback, J., McKinney, W., Jbrockmendel, Van Den Bossche, J., Augspurger, T., Cloud, P., Hawkins, S., Gfyoung, Sinhrks, Roeschke, M., Klein, A., Petersen, T., Tratner, J., She, C., Ayd, W., Naveh, S., Patrick, Garcia, M., Schendel, J., ... H-Vetinari. (2021, April 12). *Pandas-dev/pandas: Pandas 1.2.4* (Version v1.2.4). https://doi.org/10.5281/ZENODO.4681666

Recanatesi, S., Pereira-Obilinovic, U., Murakami, M., Mainen, Z., & Mazzucato, L. (2022). Metastable attractors explain the variable timing of stable behavioral action sequences. *Neuron*, *110*(1), 139–153.e9. https://doi.org/10.1016/j.neuron.2021.10.011

Reep, R. L., Corwin, J. V., Hashimoto, A., & Watson, R. T. (1987). Efferent connections of the rostral portion of medial agranular cortex in rats. *Brain Research Bulletin*, *19*(2), 203–221. https://doi.org/10.1016/0361-9230(87)90086-4

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), 585–590. https://doi.org/10.1038/nature12160

Rikhye, R. V., Gilra, A., & Halassa, M. M. (2018). Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nature Neuroscience*, 1–18. https://doi.org/10.1038/s41593-018-0269-z

Roberts, B. M., Hsieh, L.-T., & Ranganath, C. (2013). Oscillatory activity during maintenance of spatial and temporal information in working memory. *Neuropsychologia*, *51*(2), 349–357. https://doi.org/10.1016/j.neuropsychologia.2012.10.009

Roelfsema, P. R., & Treue, S. (2014). Basic neuroscience research with nonhuman primates: A small but indispensable component of biomedical research. *Neuron*, *82*(6), 1200–1204. https://doi.org/10.1016/j.neuron.2014.06.003

Rogalla, M. M., Rauser, I., Schulze, K., Osterhagen, L., & Hildebrandt, K. J. (2020). Mice tune out not in: Violation of prediction drives auditory saliency. *Proceedings of the Royal Society B: Biological Sciences*, *287*(1919), 20192001–10. https://doi.org/10.1098/rspb.2019.2001

Rohrer, B., & Hogan, N. (2006). Avoiding Spurious Submovement Decompositions II: A Scattershot Algorithm. *Biological Cybernetics*, *94*(5), 409–414. https://doi.org/10.1007/s00422-006-0055-y

Romo, R., Hernández, A., Zainos, A., Brody, C., & Salinas, E. (2002). Exploring the cortical evidence of a sensory-discrimination process. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *357*(1424), 1039–1051. https://doi.org/10.1098/rstb.2002.1100

Rose, J. E., & Woolsey, C. N. (1948). The orbitofrontal cortex and its connections with the mediodorsal nucleus in rabbit, sheep and cat. *Research Publications - Association for Research in Nervous and Mental Disease*, *27 (1 vol.)* 210–232.

Ross, P. T., & Bibler Zaidi, N. L. (2019). Limited by our limitations. *Perspectives on Medical Education*, *8*(4), 261–264. https://doi.org/10.1007/s40037-019-00530-x

Rottschy, C., Langner, R., Dogan, I., Reetz, K., Laird, A. R., Schulz, J. B., Fox, P. T., & Eickhoff, S. B. (2012). Modelling neural correlates of working memory: A coordinate-based meta-analysis. *NeuroImage*, *60*(1), 830–846. https://doi.org/10.1016/j.neuroimage.2011.11.050

Roux, F., & Uhlhaas, P. J. (2014). Working memory and neural oscillations: A-γ versus ϑ-γ codes for distinct WM information? *Trends in cognitive sciences*, *18*(1), 16–25. https://doi.org/10.1016/j.tics.2013.10.010

Rule, M. E., Loback, A. R., Raman, D., Driscoll, L. N., Harvey, C. D., & O'Leary, T. (2020). Stable task information from an unstable neural population. *eLife*, *9*, e51121. https://doi.org/10.7554/eLife.51121

Sabatini, B. L., & Tian, L. (2020). Imaging Neurotransmitter and Neuromodulator Dynamics In Vivo with Genetically Encoded Indicators. *Neuron*, *108*(1), 17–32. https://doi.org/10.1016/j.neuron.2020.09.036

Salazar, R. F., Dotson, N. M., Bressler, S. L., & Gray, C. M. (2012). Content-specific fronto-parietal synchronization during visual working memory. *Science*, *338*(6110), 1097–1100. https://doi.org/10.1126/science.1224000

Sanders, J. I., & Kepecs, A. (2012). Choice ball: A response interface for two-choice psychometric discrimination in head-fixed mice. *Journal of Neurophysiology*, *108*(12), 3416–3423. https://doi.org/10.1152/jn.00669.2012

Santos, E. M., & Kowler, E. (2017). Anticipatory smooth pursuit eye movements evoked by probabilistic cues. *Journal of Vision*, *17*(13), 13. https://doi.org/10.1167/17.13.13

Sarnthein, J., Petsche, H., Rappelsberger, P., Shaw, G. L., & von Stein, A. (1998). Synchronization between prefrontal and posterior association cortex during human working memory. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(12), 7092–7096. https://doi.org/10.1073/pnas.95.12.7092

Sauseng, P., Klimesch, W., Doppelmayr, M., Hanslmayr, S., Schabus, M., & Gruber, W. R. (2004). Theta coupling in the human electroencephalogram during a working memory task. *Neuroscience Letters*, *354*(2), 123–126. https://doi.org/10.1016/j.neulet.2003.10.002

Sauseng, P., Klimesch, W., Schabus, M., & Doppelmayr, M. (2005). Fronto-parietal EEG coherence in theta and upper alpha reflect central executive functions of working memory. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, *57*(2), 97–103. https://doi.org/10.1016/j.ijpsycho.2005.03.018

Schack, B., Klimesch, W., & Sauseng, P. (2005). Phase synchronization between theta and upper alpha oscillations in a working memory task. *International Journal of Psychophysiology*, *57*(2), 105–114. https://doi.org/10.1016/j.ijpsycho.2005.03.016

Schmitt, L. I., Wimmer, R. D., Nakajima, M., Happ, M., Mofakham, S., & Halassa, M. M. (2017). Thalamic amplification of cortical connectivity sustains attentional control. *Nature*, *522*, 50. https://doi.org/10.1038/nature22073

Schroeder, S. C. Y., Ball, F., & Busch, N. A. (2018). The role of alpha oscillations in distractor inhibition during memory retention. *The European Journal of Neuroscience*, *48*(7), 2516–2526. https://doi.org/10.1111/ejn.13852

Sengupta, P., & Samuel, A. D. T. (2009). Caenorhabditis elegans: A model system for systems neuroscience. *Current Opinion in Neurobiology*, *19*(6), 637–643. https://doi.org/10.1016/j.conb.2009.09.009

Sghirripa, S., Graetz, L., Merkin, A., Rogasch, N. C., Ridding, M. C., Semmler, J. G., & Goldsworthy, M. R. (2021). The Role of Alpha Power in the Suppression of Anticipated Distractors During Verbal Working Memory. *Brain Topography*, *34*(1), 102–109. https://doi.org/10.1007/s10548-020-00810-4

Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *298*(1089), 199–209. https://doi.org/10.1098/rstb.1982.0082

Sherfey, J. S., Ardid, S., Miller, E. K., Hasselmo, M. E., & Kopell, N. J. (2020). Prefrontal oscillations modulate the propagation of neuronal activity required for working memory. *bioRxiv*, 531574. https://doi.org/10.1101/531574

Sherman, S. M. (2016). Thalamus plays a central role in ongoing cortical functioning. *Nature Neuroscience*, *16*(4), 533–541. https://doi.org/10.1038/nn.4269

Sherman, S. M., & Guillery, R. W. (1998). On the actions that one nerve cell can have on another: Distinguishing "drivers" from "modulators". *Proceedings of the National Academy of Sciences*, *95*(12), 7121–7126. https://doi.org/10.1073/pnas.95.12.7121

Shirhatti, V., Borthakur, A., & Ray, S. (2016). Effect of Reference Scheme on Power and Phase of the Local Field Potential. *Neural computation*, *28*(5), 882–913. https://doi.org/10.1162/NECO_a_00827

Siegel, M., Warden, M. R., & Miller, E. K. (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(50), 21341–21346. https://doi.org/10.1073/pnas.0908193106

Signore, P., Chaoui, M., Nosten-Bertrand, M., Perez-Diaz, F., & Marchaland, C. (1991). Handedness in mice: Comparison across eleven inbred strains. *Behavior Genetics*, *21*(4), 421–429. https://doi.org/10.1007/BF01065977

Signore, P., Nosten-Bertrand, M., Chaoui, M., Roubertoux, P. L., Marchaland, C., & Perez-Diaz, F. (1991). An assessment of handedness in mice. *Physiology & Behavior*, *49*(4), 701–704. https://doi.org/10.1016/0031-9384(91)90305-8

Spaet, T., & Harlow, H. F. (1943). Problem solution by monkeys following bilateral removal of the prefrontal areas. II. Delayed reaction problems involving use of the matching-from-sample method. *Journal of Experimental Psychology*, *32*(5), 424–434. https://doi.org/10.1037/h0058008

Spitzer, B., Wacker, E., & Blankenburg, F. (2010). Oscillatory Correlates of Vibrotactile Frequency Processing in Human Working Memory. *Journal of Neuroscience*, *30*(12), 4496–4502. https://doi.org/10.1523/JNEUROSCI.6041-09.2010

Starkweather, C. K., Gershman, S. J., & Uchida, N. (2018). The Medial Prefrontal Cortex Shapes Dopamine Reward Prediction Errors under State Uncertainty. *Neuron*, *98*(3), 616–629.e6. https://doi.org/10.1016/j.neuron.2018.03.036

Stauffer, W. R., Lak, A., Yang, A., Borel, M., Paulsen, O., Boyden, E. S., & Schultz, W. (2016). Dopamine Neuron-Specific Optogenetic Stimulation in Rhesus Macaques. *Cell*, *166*(6), 1564–1571.e6. https://doi.org/10.1016/j.cell.2016.08.024

Stokes, K. A., & Best, P. J. (1990). Mediodorsal thalamic lesions impair "Reference" and "working" memory in rats. *Physiology & Behavior*, *47*(3), 471–476. https://doi.org/10.1016/0031-9384(90)90111-G

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662. https://doi.org/10.1037/h0054651

Stuss, D. T., Kaplan, E. F., Benson, D. F., Weir, W. S., Steven, C., & Sarazin, F. F. (1982). Evidence for involvement of orbitofrontal cortex in memory functions: An interference effect. *Journal of Comparative and Physiological Psychology*, *96*(6), 913–925. https://doi.org/10.1037/0735-7036.96.6.913

Suzuki, M., & Gottlieb, J. (2013). Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. *Nature Neuroscience*, *16*(1), 98–104. https://doi.org/10.1038/nn.3282

Tallon-Baudry, C., Kreiter, A., & Bertrand, O. (1999). Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. *Visual Neuroscience*, *16*(3), 449–459. https://doi.org/10.1017/s0952523899163065

Tanji, J., & Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. *Physiological Reviews*, *88*(1), 37–57. https://doi.org/10.1152/physrev.00014.2007

Tecce, J. J. (1972). Contingent negative variation (CNV) and psychological processes in man. *Psychological Bulletin*, *77*(2), 73–108. https://doi.org/10.1037/h0032177

Testard, C., Tremblay, S., & Platt, M. (2021). From the field to the lab and back: Neuroethology of primate social behavior. *Current Opinion in Neurobiology*, *68*, 76–83. https://doi.org/10.1016/j.conb.2021.01.005

Teufel, C., & Fletcher, P. C. (2020). Forms of prediction in the nervous system. *Nature Reviews Neuroscience*, 1–12. https://doi.org/10.1038/s41583-020-0275-5

The International Brain Laboratory, Aguillon-Rodriguez, V., Angelaki, D., Bayer, H., Bonacchi, N., Carandini, M., Cazettes, F., Chapuis, G., Churchland, A. K., Dan, Y., Dewitt, E., Faulkner, M., Forrest, H., Haetzel, L., Häusser, M., Hofer, S. B., Hu, F., Khanal, A., Krasniak, C., … Zador, A. M. (2021). Standardized and reproducible measurement of decision-making in mice (N. Uchida & M. J. Frank, Eds.). *eLife*, *10*, e63711. https://doi.org/10.7554/eLife.63711

The International Brain Laboratory, Banga, K., Benson, J., Bonacchi, N., Bruijns, S. A., Campbell, R., Chapuis, G. A., Churchland, A. K., Davatolhagh, M. F., Lee, H. D., Faulkner, M., Hu, F., Hunterberg, J., Khanal, A., Krasniak, C., Meijer, G. T., Miska, N. J., Mohammadi, Z., Noel, J.-P., … Witten, I. B. (2022, May 12). *Reproducibility of in-vivo electrophysiological measurements in mice*. https://doi.org/10.1101/2022.05.09.491042

Timothy, S., Malka, S., Jesse, K., Gabriela, M.-N., & Conor, L. (2021). Prefrontal deep projection neurons enable cognitive flexibility via persistent feedback monitoring. *Cell*, *184*(10), 2750–2766.e17. https://doi.org/10.1016/j.cell.2021.03.047

Tognoli, E., & Kelso, J. A. S. (2014). The metastable brain. *Neuron*, *81*(1), 35–48. https://doi.org/10.1016/j.neuron.2013.12.022

Tononi, G., Edelman, G. M., & Sporns, O. (1998). Complexity and coherency: Integrating information in the brain. *Trends in Cognitive Sciences*, *2*(12), 474–484. https://doi.org/10.1016/S1364-6613(98)01259-5

Toptaş, P., Gür, E., & Balcı, F. (2022). Count-based decision-making in mice: Numerosity vs. stimulus control. *Animal Cognition*, *25*(6), 1621–1630. https://doi.org/10.1007/s10071-022-01652-2

Tremblay, S., Acker, L., Afraz, A., Albaugh, D. L., Amita, H., Andrei, A. R., Angelucci, A., Aschner, A., Balan, P. F., Basso, M. A., Benvenuti, G., Bohlen, M. O., Caiola, M. J., Calcedo, R., Cavanaugh, J., Chen, Y., Chen, S., Chernov, M. M., Clark, A. M., … Platt, M. L. (2020). An Open Resource for Non-human Primate Optogenetics. *Neuron*, *108*(6), 1075–1090.e6. https://doi.org/10.1016/j.neuron.2020.09.027

Tseng, H.-a., & Han, X. (2019). Distinct Oscillation Dynamics Selectively Coordinate Excitatory and Inhibitory Neurons in Prefrontal Cortex during Sensory Discrimination. *bioRxiv*, *86*, 629659. https://doi.org/10.1101/629659

Uylings, H. B., & van Eden, C. G. (1991). Chapter 3 Qualitative and quantitative comparison of the prefrontal cortex in rat and in primates, including humans. In *Progress in Brain Research* (pp. 31–62, Vol. 85). Elsevier. https://doi.org/10.1016/S0079-6123(08)62675-8

Van der Werf, Y. D., Scheltens, P., Lindeboom, J., Witter, M. P., Uylings, H. B. M., & Jolles, J. (2003). Deficits of memory, executive functioning and attention following infarction in the thalamus; a study of 22 cases with localised lesions. *Neuropsychologia*, *41*(10), 1330–1344. https://doi.org/10.1016/s0028-3932(03)00059-9

van der Weyden, L., White, J. K., Adams, D. J., & Logan, D. W. (2011). The mouse genetics toolkit: Revealing function and mechanism. *Genome Biology*, *12*(6), 224. https://doi.org/10.1186/gb-2011-12-6-224

Van Rossum, G., & Drake, F. L. (2009). *Python 3 reference manual*. CreateSpace.

van Kerkoerle, T., Self, M. W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., & Roelfsema, P. R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences*, *111*(40), 14332–14341. https://doi.org/10.1073/pnas.1402773111

van Zoest, W., Huber-Huber, C., Weaver, M. D., & Hickey, C. (2021). Strategic Distractor Suppression Improves Selective Control in Human Vision. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *41*(33), 7120–7135. https://doi.org/10.1523/JNEUROSCI.0553-21.2021

Vendrell, P., Junqué, C., Pujol, J., Jurado, M. A., Molet, J., & Grafman, J. (1995). The role of prefrontal regions in the Stroop task. *Neuropsychologia*, *33*(3), 341–352. https://doi.org/10.1016/0028-3932(94)00116-7

Virtanen, P., Gommers, R., Burovski, E., Oliphant, T. E., Weckesser, W., Cournapeau, D., Alexbrc, Peterson, P., Reddy, T., Wilson, J., Haberland, M., Mayorov, N., Endolith, Nelson, A., Walt, S. V. D., Laxalde, D., Brett, M., Polat, I., Larson, E., … Feng, Y. (2020, October 17). *Scipy/scipy: SciPy 1.5.3* (Version v1.5.3). https://doi.org/10.5281/ZENODO.4100507

Viswanathan, P., & Nieder, A. (2013). Neuronal correlates of a visual "sense of number" in primate parietal and prefrontal cortices. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(27), 11187–11192. https://doi.org/10.1073/pnas.1308141110

Viviani, P., & Terzuolo, C. (1982). Trajectory determines movement dynamics. *Neuroscience*, *7*(2), 431–437. https://doi.org/10.1016/0306-4522(82)90277-9

Volle, E., Gonen-Yaacovi, G., de Lacy Costello, A., Gilbert, S. J., & Burgess, P. W. (2011). The role of rostral prefrontal cortex in prospective memory: A voxel-based lesion study. *Neuropsychologia*, *49*(8), 2185–2198. https://doi.org/10.1016/j.neuropsychologia.2011.02.045

von Uexküll, J. (1926). *Theoretical biology*. K. Paul, Trench, Trubner & co. ltd.

Voytek, B., Kayser, A. S., Badre, D., Fegen, D., Chang, E. F., Crone, N. E., Parvizi, J., Knight, R. T., & D'Esposito, M. (2015). Oscillatory dynamics coordinating human frontal networks in support of goal maintenance. *Nature Neuroscience*, *18*(9), 1318–1324. https://doi.org/10.1038/nn.4071

Walter, W. G. (1964). The Convergence and Interaction of Visual, Auditory, and Tactile Responses in Human Nonspecific Cortex. *Annals of the New York Academy of Sciences*, *112*(1), 320–361. https://doi.org/10.1111/j.1749-6632.1964.tb26760.x

Walter, W. G., Cooper, R., Aldridge, V. J., McCALLUM, W. C., & Winter, A. L. (1964). Contingent Negative Variation : An Electric Sign of Sensori-Motor Association and Expectancy in the Human Brain. *Nature*, *203*(4943), 380–384. https://doi.org/10.1038/203380a0

Walton, J. P., Frisina, R. D., & Meierhans, L. R. (1995). Sensorineural hearing loss alters recovery from short-term adaptation in the C57BL/6 mouse. *Hearing Research*, *88*(1-2), 19–26. https://doi.org/10.1016/0378-5955(95)00093-J

Wang, B. A., & Pleger, B. (2020). Confidence in Decision-Making during Probabilistic Tactile Learning Related to Distinct Thalamo-Prefrontal Pathways. *Cerebral Cortex (New York, N.Y.: 1991)*, *30*(8), 4677–4688. https://doi.org/10.1093/cercor/bhaa073

Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., & Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, *21*(6), 860–868. https://doi.org/10.1038/s41593-018-0147-8

Wang, Q., Ding, S.-L., Li, Y., Royall, J., Feng, D., Lesnar, P., Graddis, N., Naeemi, M., Facer, B., Ho, A., Dolbeare, T., Blanchard, B., Dee, N., Wakeman, W., Hirokawa, K. E., Szafer, A., Sunkin, S. M., Oh, S. W., Bernard, A., … Ng, L. (2020). The

Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell*, *181*(4), 936–953.e20. https://doi.org/10.1016/j.cell.2020.04.007

Wang, X. J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences*, *24*(8), 455–463. https://doi.org/10.1016/s0166-2236(00)01868-3

Wang, X.-J. (2006, November 13). A microcircuit model of prefrontal functions: Ying and yang of reverberatory neurodynamics in cognition. In J. Risberg & J. Grafman (Eds.), *The Frontal Lobes* (1st ed., pp. 92–127). Cambridge University Press. https://doi.org/10.1017/CBO9780511545917.006

Wang, X.-J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Physiological Reviews*, *90*(3), 1195–1268. https://doi.org/10.1152/physrev.00035.2008

Watanabe, K., & Funahashi, S. (2014). Neural mechanisms of dual-task interference and cognitive capacity limitation in the prefrontal cortex. *Nature Neuroscience*, *17*(4), 601–611. https://doi.org/10.1038/nn.3667

Watanabe, M. (1986). Prefrontal unit activity during delayed conditional Go/No-Go discrimination in the monkey. I. Relation to the stimulus. *Brain Research*, *382*(1), 1–14. https://doi.org/10.1016/0006-8993(86)90104-6

Watanabe, Y., & Funahashi, S. (2012). Thalamic mediodorsal nucleus and working memory. *Neuroscience and biobehavioral reviews*, *36*(1), 134–142. https://doi.org/10.1016/j.neubiorev.2011.05.003

Watanabe, Y., Takeda, K., & Funahashi, S. (2009). Population vector analysis of primate mediodorsal thalamic activity during oculomotor delayed-response performance. *Cerebral Cortex (New York, N.Y.: 1991)*, *19*(6), 1313–1321. https://doi.org/10.1093/cercor/bhn170

Watrous, A. J., Fell, J., Ekstrom, A. D., & Axmacher, N. (2015). More than spikes: Common oscillatory mechanisms for content specific neural representations during perception and memory. *Current opinion in neurobiology*, *31*, 33–39. https://doi.org/10.1016/j.conb.2014.07.024

Weaver, M. D., Paoletti, D., & van Zoest, W. (2014). The impact of predictive cues and visual working memory on dynamic oculomotor selection. *Journal of Vision*, *14*(3), 27. https://doi.org/10.1167/14.3.27

Wendelken, C., Bunge, S. A., & Carter, C. S. (2008). Maintaining structured information: An investigation into functions of parietal and lateral prefrontal cortices. *Neuropsychologia*, *46*(2), 665–678. https://doi.org/10.1016/j.neuropsychologia.2007.09.015

Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nakajima, M., Deisseroth, K., & Halassa, M. M. (2015). Thalamic control of sensory selection in divided attention. *Nature*, *526*(7575), 705–709. https://doi.org/10.1038/nature15398

Wöstmann, M., Alavash, M., & Obleser, J. (2019). Alpha Oscillations in the Human Brain Implement Distractor Suppression Independent of Target Selection. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *39*(49), 9797–9805. https://doi.org/10.1523/JNEUROSCI.1954-19.2019

Yoon, J. H., Curtis, C. E., & D'Esposito, M. (2006). Differential effects of distraction during working memory on delay-period activity in the prefrontal cortex and the visual association cortex. *NeuroImage*, *29*(4), 1117–1126. https://doi.org/10.1016/j.neuroimage.2005.08.024

Zaksas, D., & Pasternak, T. (2006). Directional Signals in the Prefrontal Cortex and in Area MT during a Working Memory for Visual Motion Task. *Journal of Neuroscience*, *26*(45), 11726–11742. https://doi.org/10.1523/JNEUROSCI.3420-06.2006

Zhang, H., Watrous, A. J., Patel, A., & Jacobs, J. (2018). Theta and Alpha Oscillations Are Traveling Waves in the Human Neocortex. *Neuron*, *98*(6), 1269–1281.e4. https://doi.org/10.1016/j.neuron.2018.05.019

Zipser, D., Kehoe, B., Littlewort, G., & Fuster, J. (1993). A spiking network model of short-term active memory. *Journal of Neuroscience*, *13*(8), 3406–3420. https://doi.org/10.1523/JNEUROSCI.13-08-03406.1993

Appendix

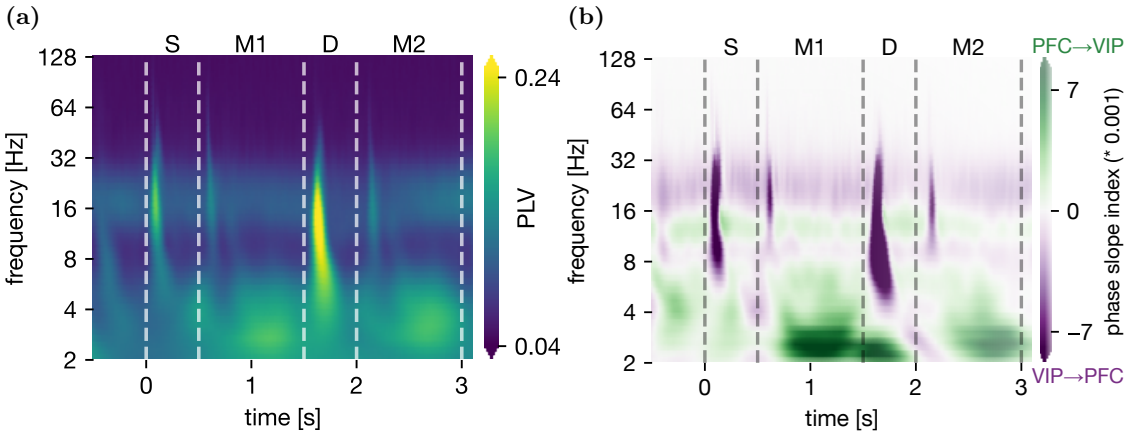

Figure S1: Control analyses of cross-regional LFP synchrony measures with local referencing of signals for trials with a distractor numerosity. (a) Phase-locking value (PLV) (compare to Fig. 2.4a). (b) Phase-slope index (PSI) (compare to Fig. 2.5a). Modified from Jacob et al. (2018).

Figure S2: Spike-field locking strength (PLV) for fronto-parietal pairs with significant sample or distractor coding, power stratified control. (a) (related to Fig. 2.8a) Cross-epoch comparisons of locking strength at the 5 Hz (theta) and 20 Hz (beta) frequencies (Wilcoxon rank-sum test, $^{**}p < 0.01$).
(b) (related to Fig. 2.8b) Comparisons of locking strength at the 5 Hz (theta) and 20 Hz (beta) frequencies across sample and distractor-coding pairs (Wilcoxon rank-sum test, $^{**}p < 0.01$). Modified from Jacob et al. (2018).

No subtraction of evoked potentials
PFC spikes, VIP LFPs

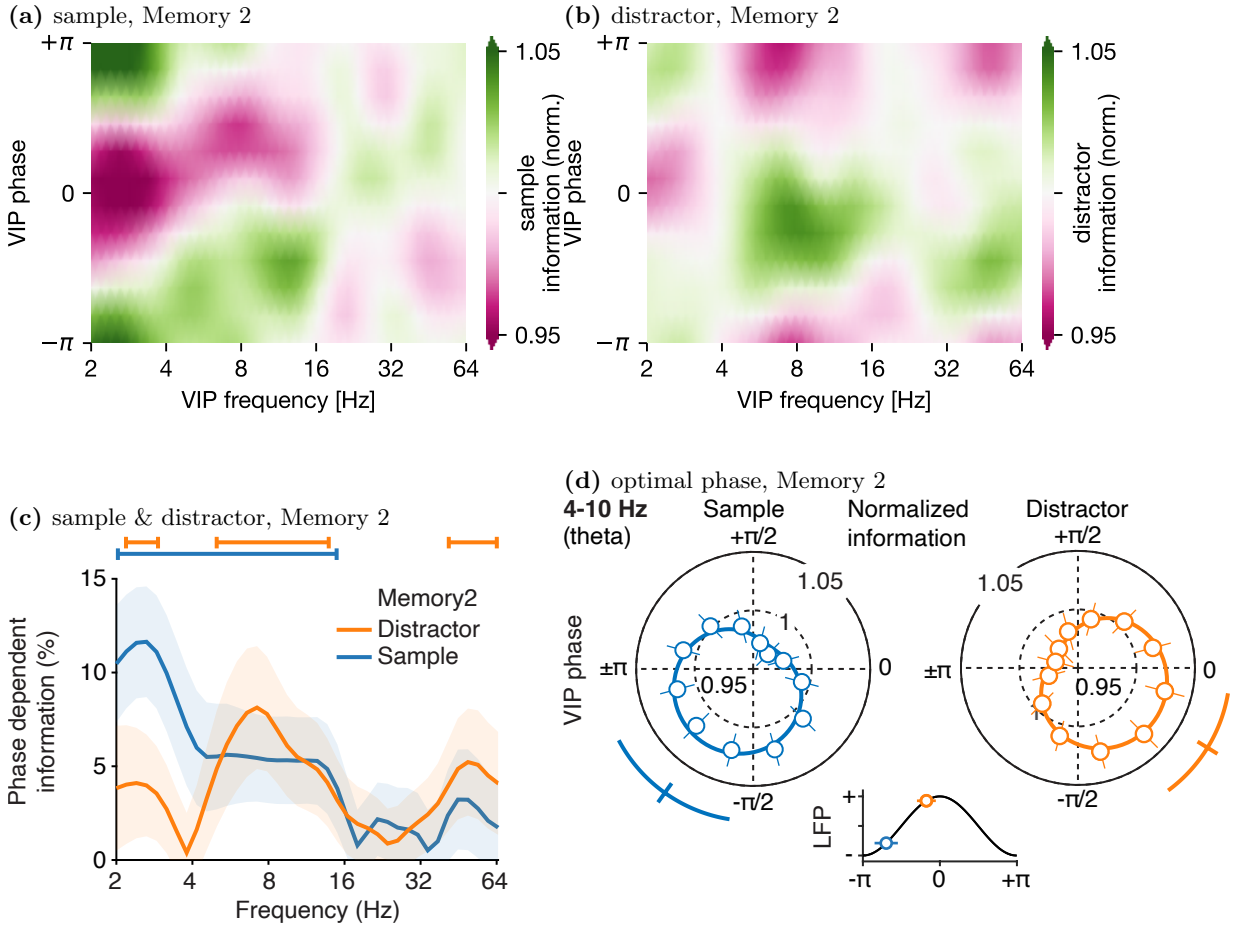

Figure S3: Prefrontal numerosity information dependent on parietal oscillation phase (no subtraction of evoked potentials). (a) Normalised sample information (ω^2) in spike counts from sample-selective prefrontal neurons at specific phases of VIP LFP during the second memory delay. (b) same analysis like (a) for distractor information contained in spike counts of distractor-selective neurons. (c) Percentage of phase-dependent sample or distractor information in PFC spikes as a function of parietal LFP frequency. Shading: bootstrap SEM. Horizontal bars: significant phase-dependency ($p < 0.01$, permutation test). Inset: mean ratio of phase-dependency in the delta and beta bands (whiskers: bootstrap SEM, $p < 0.05$ permutation test). (d) Top, polar plots: Normalised prefrontal sample and distractor information as a function of VIP theta phase. Circle markers and associated error bars indicate mean normalised information and bootstrap SEM at the individual phase bin. Solid closed traces are cosine fits. Short solid lines along the polar dimension and orthogonal partial circles indicate optimal readout phases and bootstrap SEM. Bottom: Location of optimal phases with bootstrap SEM on the LFP oscillation in the time domain (cosine). Modified from Jacob et al. (2018).

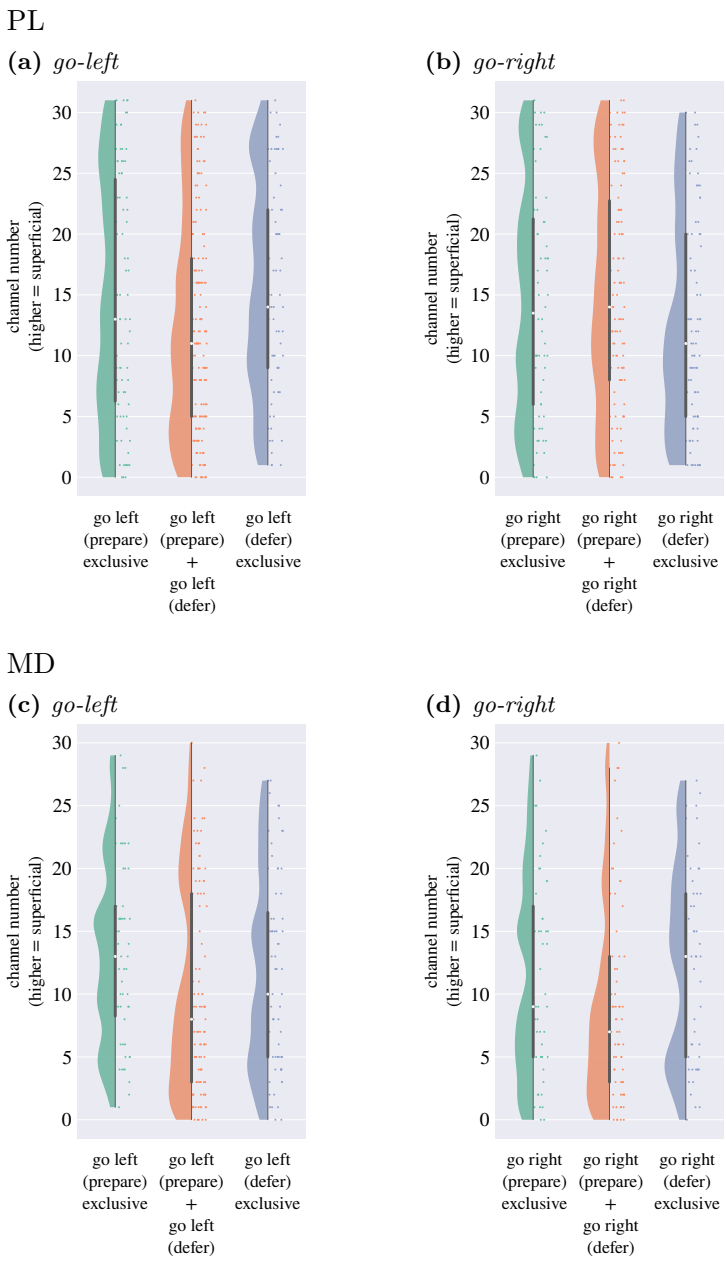

Figure S4: Channel depths of instruction-responsive subsets of single units. (a) Channel depths of *go-left*-responsive units in PL. Green (*left prep without left defr*): *go-left*-responsive exclusively after *prepare* context. Orange (*left prep intersect left defr*): *go-left*-responsive in both *prepare* and *defer* contexts (context-stable). Purple (*left defr without lefr prep*): *go-left*-responsive exclusively after *defer* context. **(b)** same as (a) but for *go-right*-responsive single units. **(c, d)** same as (a, b) but for single units recorded in MD. All differences statistically non-significant: 1-way ANOVAs with pairwise Tukey's HSD post-hoc tests.

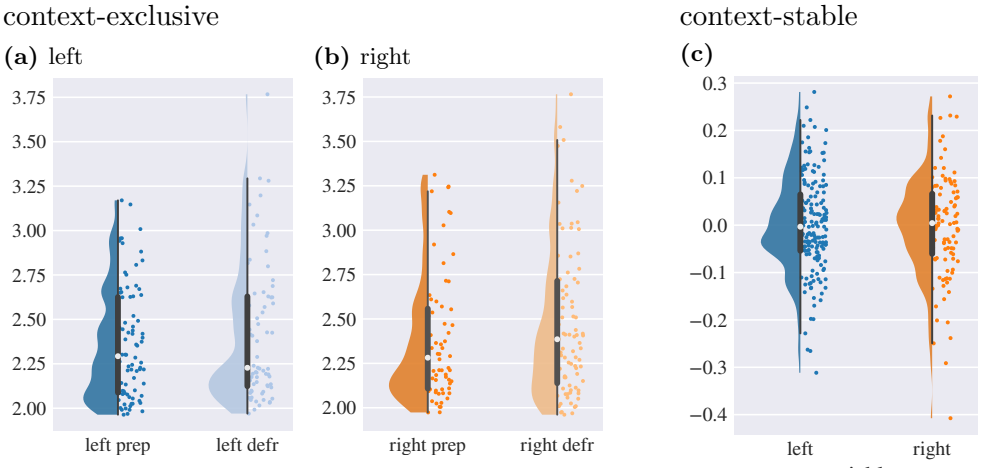

Figure S5: Influence of prior context on magnitude of instruction-responsiveness (PL). **(a)** Raw magnitude of *go-left*-responsiveness in single units that were exclusively instruction-responsive after *prepare* or *defer* context. No statistically significant difference was found (Mann-Whitney-U: $p = 0.584$). **(b)** same as (a) for *go-right*-responsive units. No statistically significant difference was found (Mann-Whitney-U: $p = 0.266$). **(c)** Context-index ($\frac{prepare-defer}{prepare+defer}$) of instruction-responsiveness magnitudes for context-stable units. No statistically significant modulations by context were found (0.001 ± 0.102, $p = 0.941$, (*go-left*); -0.002 ± 0.110, $p = 0.903$, (*go-right*); one-sample Wilcoxon test against 0 mean).

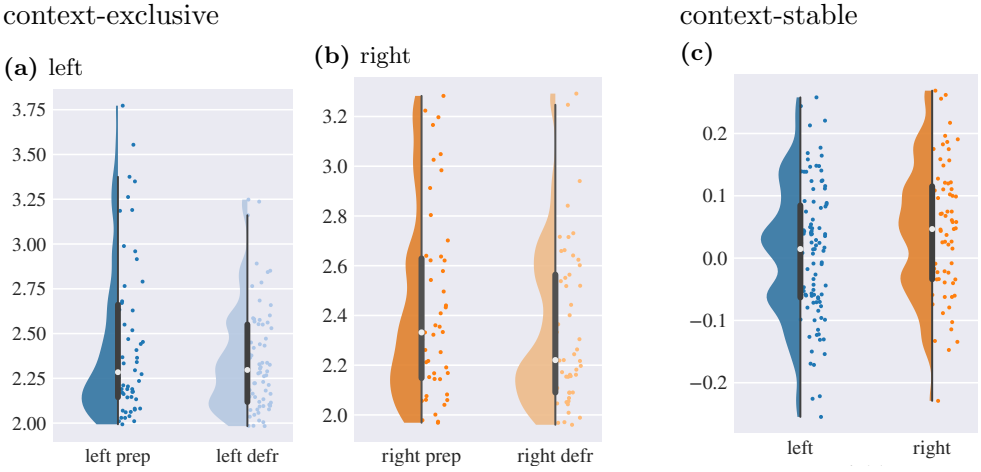

Figure S6: Influence of prior context on magnitude of instruction-responsiveness (MD). **(a)** Raw magnitude of *go-left*-responsiveness in single units that were exclusively instruction-responsive after *prepare* or *defer* context. No statistically significant difference was found (Mann-Whitney-U: $p = 0.565$). **(b)** same as (a) for *go-right*-responsive units. No statistically significant difference was found (Mann-Whitney-U: $p = 0.433$). **(c)** Context-index ($\frac{prepare-defer}{prepare+defer}$) of instruction-responsiveness magnitudes for context-stable units. Statistically significant modulation by context was not found for *go-left* but for *go-right* (0.010 ± 0.107, $p = 0.457$, (*go-left*); 0.045 ± 0.107, $p = 0.002$, (*go-right*); one-sample Wilcoxon test against 0 mean).

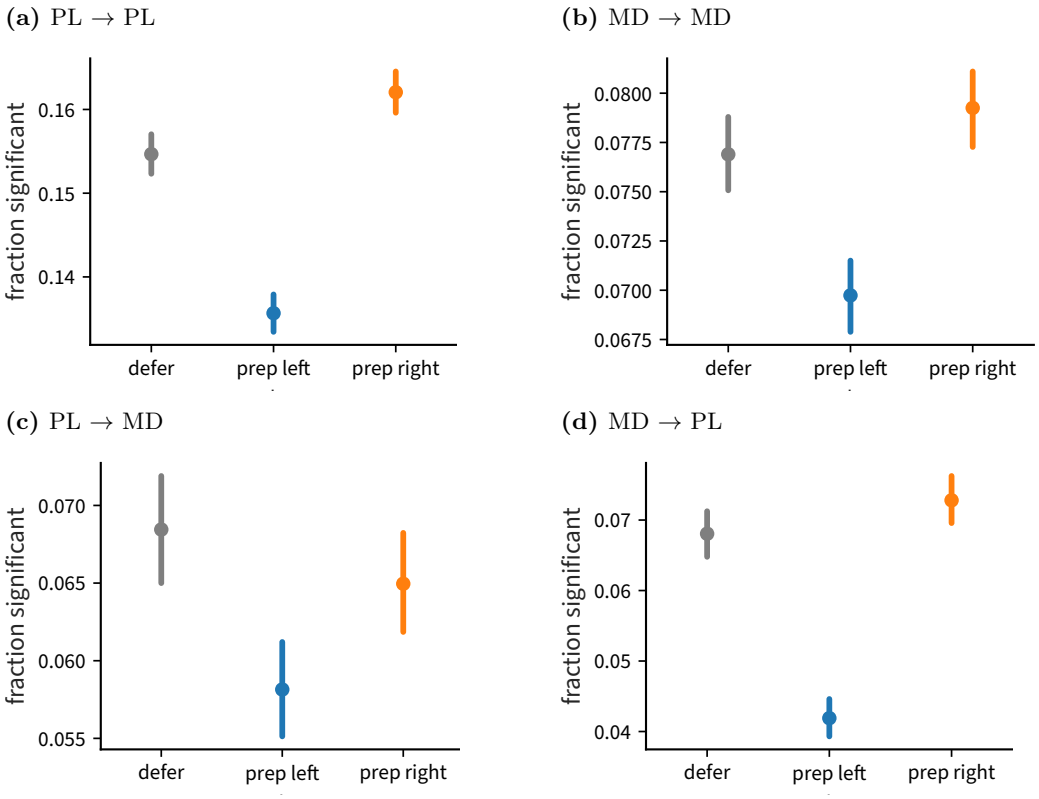

Figure S7: Significant GC links during context epoch (control for Fig. 3.50, number of trials equalised across conditions). **(a)** Fraction of significant functional links between source cells in PL and target cells in PL. **(b–d)** same as (a) but for MD to MD, PL to MD and MD to PL, respectively. Vertical error bars show 95 % bootstrap CI. Before calculating the GC metric, trials were drawn randomly so that all conditions had the number of trials that the condition with the lowest number of trials had (i.e. stratification). This was done for $n = 10$ times and the number of significant links was averaged across these iterations.

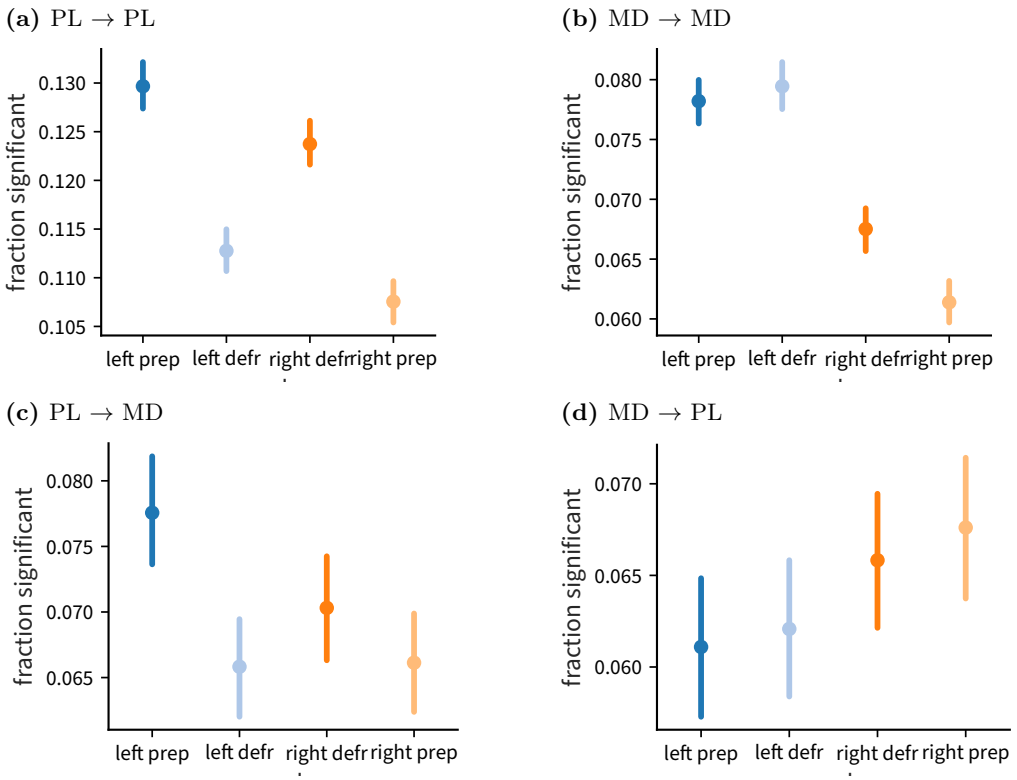

Figure S8: Significant GC links during instruction epoch (control for Fig. 3.51, number of trials equalised across conditions). **(a)** Fraction of significant functional links between source cells in PL and target cells in PL. **(b–d)** Same as (a) but for MD to MD, PL to MD and MD to PL, respectively. Vertical error bars show 95 % bootstrap CI. Before calculating the GC metric, trials were drawn randomly so that all conditions had the number of trials that the condition with the lowest number of trials had (i.e. stratification). This was done for $n = 10$ times and the number of significant links was averaged across these iterations.

Acknowledgements

I would like to thank my supervisor Simon Jacob for the trust that he placed in me to lay foundations for infrastructure and processes in the Jacob Lab, providing the idea and dataset for the monkey project, envisioning the mouse project, guidance for data analysis in both projects, and offering feedback on earlier versions of this thesis. Thanks to my Thesis Advisory Committee members Anton Sirota, who was part of the first iteration of my TAC, Helmuth Adelsberger, who participated in both iterations, and Julijana Gjorgjieva, who was involved in the second iteration, for feedback on the projects.

I would like to express my gratitude to all members of the Jacob Lab who made my time in the lab more enjoyable. A special mention goes to the members of the "Biederstein epoch", which constituted the majority of my experimental lab time - Ajit Ranganath, Tobias Bernklau, Viktor Eisenkolb and Leonie Mehrke. I am particularly grateful to Ajit and Tobias for their contributions to the mouse project dataset through animal training and recording.

Speaking of Biederstein: Thanks to Christian Obermayer for the design of the electronics of the mouse response device and general electronics debugging sessions, Felix Beyer for constructing the initial iterations of the mouse response device, and for his support in all matters related to making things. Additionally, I am grateful to Hsing-Jung Chen-Engerer and Arjan Dijke for their cheerfulness.

This entire endeavour would not have been possible without the moral support of friends and family. I would like to especially thank my friends Aljoscha Kindermann and Rafael Torrejon Torres for their compassion in response to my many rants. Additionally, I am deeply grateful to my mother for her unwavering support throughout my academic journey. Her constant love and encouragement have been instrumental in my success.

Finally, I am indebted beyond words to my partner Katrin Gerrer. Your endurance throughout the countless sacrificed weekends, holidays and times off work, which should have been for our mutual recreation, has been a true testament to your unwavering support. Especially in the past year, *defer*ing immediate action for future rewards has been my ultimate challenge, and you have lifted me out of the valley of despair and sorrow, time and again. Your words of encouragement and valuable advice provided me with the motivation to persevere. Thank you.

I may not have gone where I intended to go,
but I think I have ended up where I needed to be.

Dirk Gently
The Long Dark Tea-Time of the Soul
Douglas Adams, 1988