A role for hippocampal CA1 in structural learning in mice

Svenja Nierwetberg

A dissertation submitted in partial fulfilment

of the requirements for the degree of

Doctor of Philosophy

of

University College London

Sainsbury Wellcome Centre for Neural Circuits and Behaviour

University College London

15th July 2024

I, Svenja Nierwetberg, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

The meaning of individual events or cues in the environment is often dependent on their position relative to other cues surrounding them. The ability to learn about relationships between such ambiguous cues – often called structural learning – enables us to recognise common underlying structures of events and is thought to form the basis of episodic memory.

One area implicated in structural learning is the hippocampus. Specifically, neurons in the CA1 area of the hippocampus have been shown to represent variables such as cue configurations and their order in space and time. To investigate the neural basis of structural learning, we designed an odour-based task that requires mice to learn not only about sets of odour cues, but about their relative order in time.

Importantly, the task design allows for manipulation of the temporal structure and identity of cues separately, enabling dissociation of their neural mechanisms. Using this task, we found that mice can flexibly use previously learnt relational structures and adapt to both changes in the temporal pattern as well as in cue identity.

In line with a role for hippocampal circuitry, optogenetic inactivation of ventral CA1 (vCA1) markedly impaired task performance. Using in vivo calcium imaging, we found that vCA1 neurons encode a wide variety of task-relevant information, including maintaining odour identity across the delay and exhibiting context-specific responses to odours. Furthermore, population-level analysis revealed that neurons in vCA1 encode cues and cue combinations more robustly than outcomes.

Impact statement

Flexible behaviour is essential for our ability to adapt to novel situations as well as to inhibit inappropriate maladaptive behaviours. This behavioural flexibility relies on the brain's capacity to derive meaning from ambiguous cues by taking into account the temporal and spatial context surrounding them. Notably, this ability, often called structural, contextual or relational learning has been shown to be impaired in some of the most debilitating neural disorders – including anxiety disorders, major depression, bipolar disorder, and schizophrenia. Therefore, understanding the neural substrates and cellular mechanisms of functional structural learning provides a first step towards identifying some of the changes underlying neuropsychiatric disease.

In this thesis, I investigated hippocampal contributions to structural learning, focusing on how the hippocampus, specifically the ventral CA1 area, processes and encodes relationships between cues in order to make optimal decisions. To do this, I used an innovative odour-based task in mice as well as optogenetic tools and genetically encoded calcium indicators.

To the academic community, this research provides insight into hippocampal function and contributes to the understanding of episodic memory and behavioural flexibility, potentially informing future neuroscience research. Furthermore, the development of a non-spatial, odour-based structural learning task might be adopted and refined in future studies looking at structural learning in other brain regions, or in genetic mouse models of different disorders.

Outside of academia, this thesis and work derived from it might have implications for our conceptual understanding of neurodegenerative

conditions such as Alzheimer's disease, which are characterised by impairments in episodic memory, or into neuropsychiatric disorders such as schizophrenia that are associated to impairments in structural learning and behavioural flexibility.

More broadly, insights from this thesis might inform science policy and might be of use for funding bodies when deciding about the potential of future research programmes related to memory and mental health.

However, the impact of this research is likely to unfold incrementally, contributing to the field of neuroscience over many years. It provides a partial answer to fundamental questions about how the hippocampus contributes to the understanding of cues within context that can be built upon by further studies, potentially finally leading to changes in the diagnosis and treatment of neuropsychiatric and neurodegenerative diseases, and thereby indirectly supporting improvements in cognitive health and quality of life.

Acknowledgements

The aim of this document is to convince readers that over the past 5 years, I have gained a deep understanding of circuit neuroscience in general, and my project in specific, and to demonstrate my ability to use the scientific method to further the field. However, it would be amiss not to mention that this achievement would have been impossible without the support and advice I have been lucky to receive.

First and foremost, I would like to express my immense gratitude to Andrew for his supervision, encouragement and relentless enthusiasm. Thank you for your support and vision, and for changing my view on numerous results that I was sure would signal the failure of my project.

Secondly, to Marcus for his mentorship and support. Thank you for adapting to the meandering course of this project (away from the basal ganglia) and for always making me feel that I was part of the SJ lab nonetheless.

Then, of course, I want to thank the members of the MacAskill and SJ lab, both past and present. An undertaking like this is impossible without the companionship of people who understand the highs and lows of experimental science, and I am grateful to have had such a lovely group of people to complain to and laugh with. Special thanks are due to Fred for his counsel that has doubtlessly made me a better scientist, to Rawan for letting me live with her during the pandemic and saving my sanity, and finally to Karyna for coming with me on various adventures – from Cosyne, to plunges into ice cold ponds – and for always having an open ear.

The Sainsbury Wellcome Centre has been a wonderful place to spend the last 5 years (sometimes full-time, sometimes part-time). This supportive and collaborative environment has not only made it easy for me to learn from experts on so many techniques and topics, but also has given me the opportunity to take part in organising scientific symposia and science outreach opportunities that have enriched my PhD immeasurably.

In particular, I want to thank Sarah, Peter and Quentin who have made SWC feel like a home, even as I moved my experiments to mainland UCL. From the interviews all the way to editing this thesis, and all the times in between, you've been invaluable.

I want to also express my gratitude to Saielle, Timo, and Apri, who, despite having no professional stake in this thesis or even the field, have supported me in more ways than I could have imagined. Your unwavering belief in me has made all the difference.

Finally, to Dammy, whose contributions are too numerous to list. Thank you for everything.

Contents

Table of Figures

List of Tables

1 Introduction

In this chapter, I will give a definition of the type of learning that is the topic of this thesis: structural learning. I'll describe its importance as a building block of essential cognitive function and give an overview of the literature placing the hippocampus as a possible circuit supporting the computation underlying structural learning. Finally, I will outline the experimental conditions required to investigate the role of hippocampal circuits in structural learning.

1.1 Structural learning as a fundamental building block of cognition

Learning is an essential function of the brain, fundamental to the adaptability and independent survival of people and animals. Defined as a change in behaviour as a result of gaining information about the world (Kandel, 2021), learning is the association process connecting this information about the external world to an action. How exactly this process is implemented on the neuronal level has been a focus of neuroscientific research from the molecular level up to whole-network connectivity studies for decades and yet it is still not fully understood today.

Much of the research into learning has used classical conditioning as a paradigm, a type of learning in which a neutral stimulus is presented immediately before or simultaneously with an outcome, and the success of learning can be measured as the development of actions indicating prediction of such outcomes (e.g. Pavlovian salivation, auditory fear conditioning) (Blanchard & Blanchard, 1972; Davis, 1992; Maren, 1996; Pavlov, 1927; Pearce & Hall, 1980; Rescorla & Wagner, 1972). However, in our everyday lives as well as in the natural habitats of most animals, cues aren't always uniquely associated to one outcome – often the meaning of individual events or cues is dependent on other neutral cues

surrounding them. For example, understanding the meaning of an event may depend on its spatial or temporal relationships to other neutral cues.

This type of learning has been given multiple names, among those structural, relational, hierarchical or contextual learning. For the purposes of this thesis, I will discuss this type of learning under the term structural learning and define it as any learning that requires a directional association of two or more neutral cues in space or time (e.g. A before B, C to the right of D; Aggleton et al., 2007) required to correctly predict an outcome. Unlike classical conditioning, structural learning captures the directional relationships between neutral cues as opposed to cues and valued events, such as outcomes or actions.

Given this definition, contextual learning might be seen as a special case of structural learning in which many cues, often from different modalities, have to be integrated in order to correctly predict appropriate actions and likely outcomes (Aggleton et. al, 2007). While it is not currently clear whether structural learning and contextual learning are indeed supported by the same neural substrates, they both require the ability to associate multiple neutral stimuli and combine them in a directed way. Similarly, episodic memory can be conceptualised as relying on the foundational process structural learning: here, instead of ordered cues in space, specific events are ordered into temporal structures (Eichenbaum, 2013).

In addition to its putative function in supporting episodic memory and contextual spatial behaviour, structural learning has been suggested to have more general advantages such as enabling detection of commonalities between distinct experiences. By integrating cues into a relational structure, the meaning of new cues with a similar structure can be inferred, thereby enabling generalisation and facilitating not only flexible behaviour in novel contexts but possibly also the uniquely human ability to imagine impossible scenarios (Hassabis et al., 2007). Taken together, structural learning is put forward as a building block for the

essential task of organising knowledge into structures, an essential part required for flexible behaviour.

However, despite its essential role in many everyday behaviours, there is only limited insight into how this form of learning is achieved within the brain. Crucially, structural learning cannot be explained by the same type of learning rules that have been put forward as the basis of classical associative learning, e.g. reinforcement learning in the basal ganglia (Schultz et al., 1997; Sutton & Barto, 1998). Thus, investigating the neural substrates underlying structural learning is essential to understand its distinct mechanisms.

1.2 Brain regions associated to structural learning

Multiple brain regions have been proposed to perform processes tied to structural learning as well as derived cognitive functions that are thought to rely on such learning, e.g. categorisation, inference and generalisation. Inference relies on expanding learnt relational structures to include new items or situations, while categorisation involves grouping similar items or experiences together according to a rule which is often defined by structural relations. Finally, generalisation allows for the application of learned structure to entirely new, yet similar situations. Thus, brain regions associated with either of these abilities likely contain or process representations of structure. While I cannot describe the full breadth of the research into all regions within the constraints of this thesis, I will give a brief overview of the three regions most often associated to structural learning.

One region connected to structural learning in a large body of studies in humans, primates and rodents is the prefrontal cortex (PFC). The PFC can be anatomically subdivided into three main areas: the dorsolateral prefrontal cortex (dlPFC), the medial prefrontal cortex (mPFC) and the orbitofrontal cortex (OFC). Structurally,

the two latter areas are highly interconnected, and it is these areas that are associated to functions related to structural learning.

1.2.1 mPFC, working memory and rule-based learning

Prominent accounts of mPFC function state that single neurons in mPFC maintain representations of task-relevant stimuli over multiple seconds, thereby allowing their integration with subsequent stimuli in working memory tasks (Funahashi et al., 1993; Levy & Goldman-Rakic, 2000; Miller et al., 1996; Vogel et al., 2022; F. A. W. Wilson et al., 1993). This sustained activity has two key properties. Firstly, it is specific to the stimulus being remembered, i.e. different neuronal ensembles encode different stimuli (Funahashi et al., 1993; Miller et al., 1996). Secondly, this sustained stimulus response is task-dependent, i.e. distractor stimuli are not maintained in the same way (Rainer et al., 1998}. Thus, sustained cue responses in mPFC are not purely a buffer for sensory perception but are specific to the current task and goals. This is further supported by the finding that some neurons in mPFC respond selectively to the temporal context of a task (e.g. fire specifically in the inter-trial interval)(Jung et al.,1998), suggesting an abstract representation of task structure that is either present in mPFC, or can be accessed by it. This idea is further supported by the finding that the population activity in mPFC changes when task rules change (Durstewitz et al., 2010; Rich & Shapiro, 2009) and that an intact mPFC is required for rule-based categorisation (Monchi et al., 2001; Wallis et al., 2001).

Taken together, mPFC activity thus relates both to fundamental components to structural learning as well as higher cognitive processes that require a structural understanding of the environment.

1.2.2 OFC and representations of task state

The same is also true for its neighbour, OFC, but while mPFC seems to have stronger representations of cues and contexts, OFC is suggested to be especially

involved in representing and updating the value of choices and outcomes (Balleine et al., 2011).

This is especially important in partially observable scenarios with high uncertainty, such as bandit tasks, in which contingencies reverse suddenly. In line with this, lesion studies show that OFC lesions impact behavioural flexibility, reducing the ability to switch between behavioural strategies even though learning of the initial stimulus-outcome pairing is unaffected (Butter, 1969; Dias et al., 1997; McAlonan & Brown, 2003; Schoenbaum et al., 2003; Sul et al., 2010). Similarly, tasks in which the sensory features of different outcomes inform behaviour are dependent on functional OFC circuits (McDannald et al., 2005). This has led to the hypothesis that OFC might contain representations of task states (Wikenheiser & Schoenbaum, 2016; R. C. Wilson et al., 2014). Specifically, it is suggested that OFC integrates sensory experience with choice and reward history to form representations of states in which a specific rule (e.g. "right lever is rewarded") are true.

In summary, OFC activity is associated with tracking distinct states within a task and updating the associated values of specific actions, especially in situations where these task states are unobservable (i.e. uncued by the sensory environment (Wikenheiser & Schoenbaum, 2016)). OFC therefore is thought to represent task states at multiple levels of abstraction (within trial as well as between trials and even blocks), showing that OFC constructs or/and processes structural information.

1.2.3 Hippocampus and the cognitive map

Another brain region involved in structural learning is the hippocampus (HC). Ever since the discovery of place cells (O'Keefe & Nadel, 1979), HC has been firmly associated with the idea of a cognitive map. The underlying idea, first proposed by Tolman (Tolman, 1948), refers not to exclusively to a map of the

spatial environment, but rather to an mental representation that binds external sensory features with internal factors to form an abstract relational structure that facilitates flexible behaviour.

Indeed, in addition to an important role in spatial memory and navigation, hippocampal circuits have also been shown to be required for autobiographical memory (Burgess, 2002; Vargha-Khadem et al., 1997), contextual learning (Blanchard et al., 1977; Fanselow, 1990), scene construction (Graham et al., 2010), and temporal sequences (MacDonald et al., 2013; Omer et al., 2022). While these functions might seem unrelated at first glance, it has been suggested that HC performs a common computation underlying all of the above processes: encoding relationships between external and internal cues implicitly, even in absence of a task; then integrating cues and events that are likely to co-occur into states (or contexts) and finally predicting likely future states based on the current experience (Behrens et al., 2018; Dudchenko & Wood, 2014; Eichenbaum & Cohen, 2014; Niv, 2019). If we return to our definition of structural learning as the learning of directional relations between multiple stimuli, this theory of hippocampal function posits HC as an essential part of the putative structural learning network.

1.2.4 Summary

Taken together, these data support a role in structural learning for mPFc and OFC as well as HC. In line with a synergistic integration of each of their contributions, electrophysiological and histological studies have identified pathways that allow for interaction between all regions. HC projects directly to mPFC from CA1 and the subiculum (Cenquizca & Swanson, 2007; Jay & Witter, 1991; Sánchez-Bellot et al., 2022), and mPFC connects to HC via a disynaptic route through entorhinal cortex (EC) or through nucleus reuniens (NR) (Burwell & Amaral, 1998; Vertes et al., 2007). Similarly, OFC and HC are connected via EC and NR (McKenna & Vertes, 2004; Witter et al., 2000). Finally, mPFC and OFC are tightly

interconnected (Haber et al., 2022), thus making inter-regional coordination of different processes required for structural learning possible.

How exactly the contributions to structural learning might be subdivided across PFC, OFC and HC is not yet fully understood. Differences in task design and data acquisition make it difficult to directly compare the role of each of these areas to structural learning. From the literature, it seems that that OFC and PFC represent structural information with a higher level of abstraction and task-specificity (Miller & Cohen, 2001; Sul et al., 2010; Wikenheiser & Schoenbaum, 2016), while HC encodes both task-relevant as well as task-irrelevant stimuli and has been shown to be especially sensitive to the statistics of the environment even without the requirement of a task (Eichenbaum, 2017; Niv, 2019). These subtle differences might hint at a role for hippocampus as a first step in the chain of processes leading from the construction of structural cues all the way to abstract categorisation.

In this thesis, my goal was to investigate the potential role of the hippocampus in structural leaning. Therefore, in the following sections, I will describe theories of hippocampal function put forward in the literature, and how they, taken together with considerations of the circuitry and connectivity of hippocampus, can support the unifying hypothesis of hippocampus as a possible first step in the hierarchy of structural learning and related processes.

However, this is not meant to imply that the hippocampus is the only area in the brain implicated in or necessary for this kind of learning. An interesting future direction would be to contrast and compare the contributions of OFC and PFC to the findings described in this thesis.

1.3 Theories of hippocampal function

1.3.1 Episodic and episodic-like memory

The hippocampus has been implicated in the encoding and retrieval of episodic memory since the mid-20th century, largely due to evidence from human case studies in neurology and neuropsychiatry. Episodic memory can be defined as a type of long-term memory that encodes unique events in sequential order together with their temporal, spatial and emotional context (Eichenbaum, 2017; Panoz-Brown et al., 2016). Famously, patient H.M. suffered severe memory impairments after a bilateral medial temporal lobe resection, apparently without any perceptual disorder and general intellectual loss (Scoville & Milner, 1957). This was the first in a series of studies that suggested a role for the hippocampal formation in the encoding and retention of memories – but only in some, not all cases. Specifically, short-term memory as well as procedural and semantic memory were shown to be unaffected by damage to the hippocampus (Burgess, 2002; Vargha-Khadem et al., 1997).

Further evidence for the hippocampal involvement in episodic memory comes from the progression of neurodegenerative diseases such as Alzheimer's disease and Lewi body dementia: these diseases which lead to episodic memory loss tend to damage the hippocampus as well as the entorhinal cortex well before other brain areas (Rao et al., 2022).

But also in healthy human subjects, hippocampal activity is linked to memory: fMRI studies report that the hippocampus is activated when subjects are asked to remember the order of objects from a virtual reality driving game, or when reconstructing the correct order of scenes previously seen in a clip (Ekstrom & Bookheimer, 2007; Lehn et al., 2009). Interestingly, the hippocampus was shown to be preferentially active when the subjects were asked to order events in time, as opposed to making simple recency judgements (Eichenbaum, 2013).

The technical limitations of researching episodic memory in humans leave many questions on the cellular and circuit level unanswered, but while animal models offer a richer toolkit to study neural circuits and computations, it is unclear whether the concept of episodic memory can be applied to rodents. Thus, researchers have had to explore creative ways to investigate episodic-like memory in rodent models.

One such paradigm that has been widely used to interrogate memory for individual experiences in rodents is Pavlovian fear conditioning. In this paradigm, neutral stimuli such as tones, lights, or entire environments are arranged to predict aversive outcomes such as electric shocks, and then the display of learned fear behaviours is quantified in order to measure mnemonic retention.

Foundational work found that damage to the amygdala as well as the hippocampus could damage the expression of learned fear (Blanchard et al., 1977; Blanchard & Blanchard, 1972; Phillips & LeDoux, 1992), corroborating the role of hippocampus in memory. However, while fear conditioning to a discrete cue was found to be mediated by the amygdala (Maren, 2005), the role of the hippocampus in this paradigm turned out to be more complicated: hippocampus seems to be essential to the behaviour only in cases where an entire environment (context) predicts the shock (as opposed to a singular cue) and only if the animal had sufficient time exploring this context prior to the shock (Fanselow, 1990). But even conditioning to a discrete cue can be hippocampus-dependent – if a short temporal delay, often called a “trace interval”, is introduced between the cue and the aversive outcome, hippocampal lesions impair retention and expression of the fear memory (Quinn et al., 2002; Sellami et al., 2017). Interestingly, some studies show that hippocampal involvement increases as delays become longer in duration (Sellami et al., 2017).

These results suggest that the role of hippocampus in memory might lie in binding together multiple stimuli across space (in a context) or time (across a

trace interval), thereby performing the computation underlying structural learning.

1.3.2 Scene imagination and mental time travel

In a line of research related to episodic memory, it was found in the early 2000s that damage to the hippocampus in humans does not solely impact their ability to recall past experiences, but that patients were also impaired at describing scenes in the present (Graham et al., 2010) and even constructing fictitious and future scenes in their imagination (Hassabis, Kumaran, Vann, et al., 2007).

The construction of scenes is of course intricately linked with autobiographical memory: for most people, recalling the past involves imagining the spatial and temporal context of the memory ("mental time travel"). It is therefore not surprising that in patients with hippocampal damage, impairments in autobiographical memory predict difficulties with scene construction ability, while in patients without memory deficits, scene construction is generally intact (Squire et al., 2010). It is exactly this variability in impairments however that makes it difficult to conclusively describe the function of hippocampus based on case studies – some patients suffer no lapse in autobiographical memory, even though their hippocampus has been substantially damaged.

Functional magnetic resonance imaging (fMRI) studies in healthy subject have tried to clear up some of these inconsistencies and have shown that the hippocampus is indeed engaged when imagine future scenes (Addis et al., 2007) as well as imagined scenarios (Hassabis, Kumaran, & Maguire, 2007). It has been proposed that the common computation underlying all of these functions is the construction of rich spatial contexts (Maguire & Mullally, 2013).

However, these results could also be satisfactorily explained by the theory that hippocampus is required for the binding of arbitrary relations among individual

elements within a context, irrespective of whether these elements are framed as an autobiographical episode, therefore once again tying back to the computation underlying our definition of structural learning.

1.3.3 Spatial navigation

Indeed, in research largely parallel to the studies linking the hippocampus to episodic memory, another strand of hippocampal literature began with the recording of principal neurons in hippocampus that fire when a rat is in a particular location in its environment – the discovery of "place cells" (O'Keefe & Dostrovsky, 1971).

Since then, spatial navigation has been the predominant paradigm to study memory in rodents and has yielded an astonishing amount of insight into the neural properties that allow the hippocampus to map out environments and guide navigation. Cell types such as grid cells in medial entorhinal cortex (cells that fire at regular intervals as an animal traverses an environment), border cells in subiculum (cells that respond to an environmental boundary in a specific position relative to the animal) as well as splitter cells in CA1 (cells whose firing rates are modulated depending on the animals past and future trajectory) have been described and, together, comprise a rich model of the world represented in hippocampal circuits (Ainge, Tamosiunaite, et al., 2007; Byrne et al., 2007; Deshmukh & Knierim, 2013; Hafting et al., 2005; Lever et al., 2009; McNaughton et al., 2006; O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1979; Taube et al., 1990; Tennant et al., 2022).

These findings exemplify how the neurons in hippocampus and surrounding areas are uniquely suited to encode the different components of a multisensory environment as well as the events that take place within it (Albasser et al., 2013; Hartley et al., 2014; Jezek et al., 2011). More than simply representing spatial variables, lesion and inactivation studies have also suggested that functional

hippocampal circuits are necessary for the learning and retention of spatial maze tasks (Jarrard, 1978; Morris et al., 1982; M. B. Moser et al., 1995), and the firing of spatially selective neurons has found to be predictive of spatial memory performance (Dupret et al., 2010). It was thus suggested that the unifying factor between autobiographical memory, imagination and navigation might be the underlying spatial component (Maguire & Mullally, 2013).

1.3.4 Temporal representations in hippocampus

This view, however, was complicated by the discovery of a further type of neuron in the hippocampus CA1 area in the early 2000s: “time cells”.

These cells were first reported in animals running on a treadmill, their activity tiling the time spent running on a wheel in a specific sequence, just like place cells firing one after the other on a linear path (Pastalkova et al., 2008). Further research showed that this type of cellular activity could even be found in the absence of running in stationary rats (MacDonald et al., 2013), suggesting that hippocampus encodes time alongside space, and not solely as a secondary effect of movement (e.g. for velocity tracking or path integration).

Furthermore, several studies have shown that the activity of principal neurons in hippocampus during a delay period in a behavioural task can reflect past events. In two studies using different versions of an odour matching task, the neural activity not only tiled the delay period, but involved a different population of neurons depending on which odour had preceded the delay (MacDonald et al., 2013; Taxidis et al., 2020).

Taken together, this data shows that hippocampal circuits can encode cues and events within their temporal context as well as within their spatial context, thus encoding information about the item itself as well as its relative position along those dimensions.

1.3.5 A common computation unifying different theories of hippocampus

In summary, evidence from studies in humans and rodents suggests that hippocampus is fundamentally involved in the encoding, but not storage, of long-term autobiographical memory, is required for spatial navigation as well as the tracking of elapsed time, and is linked to the perception and construction of scenes.

Numerous theories have attempted to describe how the hippocampus may support such a diverse range of functions. While each theory draws on evidence from a different set of studies and contributes nuanced and specific interpretations, many can be summarised as the idea that hippocampus is required for the binding of arbitrary elements within an experience or across items within a context, regardless of whether these relationships are embedded within an autobiographical memory or a spatial framework (Aggleton & Pearce, 2001; Eichenbaum & Cohen, 2014; Ranganath, 2010; Rudy & Sutherland, 1995; Whittington et al., 2020).

Specifically, it has been suggested that even time and space may be just two examples of dimensions along with the hippocampus can structure knowledge, i.e. build neural representations that contain both the individual items, events or cues as well as the relation between them, thereby performing the fundamental operation of structural learning.

In support of this theory, an fMRI study found “grid-cell-like” activity in entorhinal cortex, a major input to hippocampus, when subjects were presented with images of birds with varying leg and neck lengths, thus suggesting that the hippocampal formation may represent the birds in a 2D space spanned by these

two variables (Constantinescu et al., 2016). In mice, it has been found that hippocampal neurons can represent a “place” in sound space if sound frequency is an important part of the task, analogous to place cells representing a location in the environment (Aronov et al., 2017). Similarly, it has been shown that depending on task requirements, single cells in hippocampus can respond selectively to specific spatial contexts, landmarks and even goals, indicating that representations in hippocampus are driven by behavioural requirements and can construct relations between items of different modalities as well as internally generated variables (Ekstrom et al., 2007; Wood et al., 1999; Gauthier.2018).

In summary, structural learning is a unifying theory of hippocampus that is consistent with experimental data from a wide range of studies using different techniques, paradigms and model organisms. In the next section, I will outline how the anatomy and circuitry of the hippocampus supports its putative function of performing this common computation over a set of diverse inputs.

1.4 Anatomy and connectivity of the hippocampus

1.4.1 Pattern separation, completion and comparison in the trisynaptic loop

The hippocampus consists of a largely unidirectional transverse loop (termed the “trisynaptic loop”) in which cortical information enters the hippocampal formation through the entorhinal cortex (EC) and is then passed via excitatory connections through the dentate gyrus (DG), the CA3 and CA1 region, and then to subiculum from where information is distributed back to cortex (Amaral & Witter, 1989; Andersen et al., 1971; Cenquizca & Swanson, 2007; Valero & Prida, 2018).

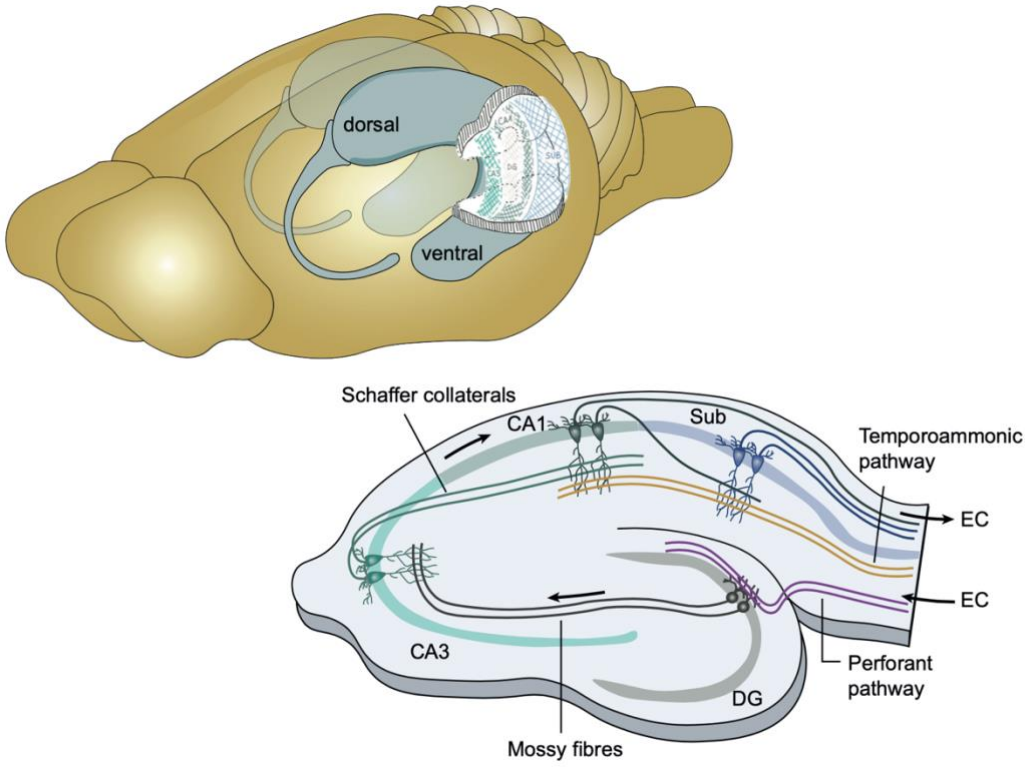

Figure 1.1: Structure of Hippocampus. Top, schematic of the position of hippocampus (blue) within the rodent brain with a transparent inlay indicating the different regions of hippocampus (CA1, CA3, DG, Sub). Bottom, Schematic of the hippocampal trisynaptic loop. Figure adapted from the 2020 PhD thesis by Candela Sánchez Bellot (UCL)

The specific neural anatomy of the trisynaptic loop, especially of the CA3 and DG, has already in the 1970s led to the hypothesis that this circuit might be responsible for the weighing of novelty and similarity in the storage and retrieval of memory (Marr, 1971). Specifically, the DG contains principal neurons ("granule cells") that outnumber the input cells from EC by an order of magnitude (Yassa & Stark, 2011). These granule cells show next to no recurrent connectivity within DG, and therefore are proposed to be well suited to separate patterns of similar inputs, converting them into a unique pattern of granule cell activity.

In the next step of the model, this proposed differentiated signal gets sent to CA3, where it causes powerful depolarisation due to large synapses (McNaughton & Morris, 1987) and is thought to drive the distinct encoding of new memories. However, CA3 also receives input directly from EC and further has strong recurrent connectivity, which might be used to auto-associate previously

stored representation from partial cues. Since its proposal, this theory of pattern separation and completion in DG and CA3 has found support from experimental data (Lee & Kesner, 2004; Leutgeb et al., 2007; Neunuebel & Knierim, 2014).

While the theory was first suggested with the function of encoding and retrieval of autobiographic memory in mind, learning about ensembles of overlapping cues while maintaining both separability and similarity lends itself equally well to the representation of spatial environments and abstract relational structures.

The putative purpose of recalling past experience similar to a present situation is of course the ability to predict likely future events and outcomes. In the trisynaptic pathway, CA1 is the recipient of the "pattern completed" CA3 activity as well as direct projections from the EC as well as other regions, and CA1 was thus put forward as a candidate to compare these recalled experiences to the present sensory input (Kumaran & Maguire, 2006). Indeed, data from fMRI studies and intracranial EEG studies in humans has shown increased CA1 activity in relation to prediction violation (Axmacher et al., 2010; Kumaran & Maguire, 2006). Interestingly, this mismatch seems to be stronger in tasks where implicit sequences were violated rather than tasks in which subjects were asked to make explicit match or mismatch judgements (Chen et al., 2015).

Taken together, the circuit pattern along the trisynaptic pathway with its additional direct projections provides a cellular basis for a common hippocampal computation, and hints at the possible parts of this computation: pattern separation, pattern completion to generate predictions and finally compare these predictions to the present sensory data. These processes place the hippocampus as a possible site for structural learning: in order to recognise the shared structure across experiences, pattern separation and completion are necessary.

But this alone is not sufficient to make structural learning possible. In order to integrate events and cues that are in different spatial locations and/or occur in different moments in time, any brain region underpinning structural learning must have a mechanism to connect distant events or stimuli. The redundancy within hippocampal networks as well as connectivity between sub-regions of hippocampus as well as to prefrontal areas makes it a good candidate to fulfil this role of maintaining information (discussed in more detail in **Chapter 6.4.1**).

1.4.2 Functional differentiation along the dorso-ventral axis

While the circuit motif described in the previous section repeats itself along the longitudinal axis of the hippocampus, afferent and efferent connectivity change along the dorso-ventral axis of the hippocampus (Swanson & Cowan, 1977), thereby allowing the same computation to be executed on variable information.

Specifically, the dorsal HC receives inputs from the visual, auditory and somatosensory cortices via EC, while the parts of EC projecting to ventral HC receive more inputs from the amygdala, olfactory cortices as well as the hypothalamus. This pattern of inputs is largely maintained on the efferent side: axons projecting from CA1 and the subiculum maintain their dorso-ventral position in the lateral-to-medial axis of the EC (T. V. Groen & Wyss, 1990; T. van Groen et al., 1986; Köhler, 1985), and direct projections from dCA1 and vCA1 show a different ensemble of targets, e.g. dCA1 but not vCA2 project to anterior thalamus and retrosplenial cortex, vCA1 but not dCA1 project to VTA and amygdala, (Fanselow & Dong, 2010).

Given these differences anatomical projection patterns as well as distinct signatures of gene expression, dorsal and ventral hippocampus have been proposed to be functionally distinct (Fanselow & Dong, 2010; M. Moser & Moser, 1998). However, what exactly the respective role of dorsal and ventral hippocampus might be is not yet fully understood. This is partly due to the

general difficulty of comparing results acquired through different behavioural paradigms, but also due to a bias in anatomical targets between the fields: the navigation field has historically focused on dorsal regions (Bittner et al., 2017; M. B. Moser et al., 1995; O'Keefe & Nadel, 1979) whereas fear conditioning studies have more frequently targeted ventral hippocampus (Kjelstrup et al., 2002; Maren & Holt, 2004; Twining et al., 2020). Anecdotally, studies using tasks paradigms that require structural learning have shown a role for ventral hippocampus (Pennartz et al., 2011; Riaz et al., 2017).

In summary, the ventral and dorsal hippocampus have been suggested to be functionally distinct units given their different afferent and efferent connectivity. However, the same circuit motif is present throughout the entire hippocampus. A parsimonious hypothesis combining these two views is that the hippocampus might carry out a common computation across the entire dorso-ventral axis, but that the functional use of this computation varies based on the inputs it is performed on, and the projections it is sent to.

1.4.3 Odour information in the hippocampus

The olfactory system provides information about odours in the environment to the hippocampus, with the olfactory bulb (OB) serving as the first relay for incoming olfactory signals. The OB processes sensory input from olfactory receptor neurons and transmits this information via the mitral and tufted cells to the piriform cortex (PC), anterior olfactory nucleus, and other downstream targets (Shepherd, 2003).

From the piriform cortex, projections travel to the lateral entorhinal cortex (LEC), which provides a major input to the hippocampus, specifically targeting the dentate gyrus (DG) and CA3 regions through the perforant pathway (Diodato et al., 2016). This part of entorhinal cortex projects more robustly to ventral than to the dorsal hippocampus (Kerr et al., 2007). Furthermore, there is evidence for

sparse direct connections from the OB (via mitral and tufted cells) to the ventral portion of CA1 and DG (Luskin & Price, 1983). Additionally, the piriform cortex, as a secondary processing centre for olfactory information, has been shown to provide direct input to the ventral hippocampus (Kesner et al., 2011). These direct projections may play a critical role in associating specific odours with contexts and outcomes, since, as discussed in **Section 1.4.2**, ventral hippocampus also projects more strongly to regions associated with value such as the amygdala and the NAc.

In summary, while the indirect pathway via the entorhinal cortex is the primary route for olfactory information to reach the hippocampus, direct projections from the olfactory bulb and piriform cortex provide a complementary route for processing olfactory-contextual associations. Both of these routes target the ventral hippocampus more robustly than dorsal regions which makes the ventral hippocampus a promising target to investigate odour related behaviours.

1.5 Investigating structural learning in a rodent model

As outlined above, data from functional as well as anatomical studies supports the hypothesis of the hippocampus as an essential hub in the network for structural learning. However, to understand how neural activity in hippocampus can support such a function, it is necessary to leverage the tools available in the rodent model to gain insight into the way hippocampus represents the components required for structural learning both within individual neurons as well as on a population level.

Both fear conditioning and spatial navigation paradigms have been successfully used to ask questions that contain structural learning aspects, as for example how hippocampus supports the construction of a spatial map comprising features and goals in navigation, or the how the multisensory context triggering a temporal sequence that leads to aversive outcomes in fear conditioning is

encoded in hippocampal circuits. While these paradigms have pushed our understanding of hippocampal contribution to ethological behaviour forward, they have drawbacks with regards to narrowing down the precise neural computation taking place within hippocampal circuits.

Specifically, spatial navigation tasks make it difficult to disambiguate spatial information from movement signals or representations of specific cues or goals, and, since in these tasks animals often visit locations more than once, it is not always possible to distinguish neural representations of past (memory) from present (Eichenbaum & Cohen, 2014).

Similarly, fear conditioning as a paradigm only allows matching single-cell neural activity to external or internal variables to a very limited extent (Maren et al., 2013).

Therefore, a number of attempts have been made to create tasks that allow to ask questions about structural learning in hippocampus more directly.

1.5.1 Configural learning theory

From the above, it is clear that in order to understand the neural computation performed in hippocampus, task designs with more control over the specific stimuli are required.

One attempt at developing such paradigms was made with the configural learning theory proposed by Sutherland and Rudy in the late 1980s (Sutherland & Rudy, 1989). In their framework, tasks that can be solved by a simple association of one cue (e.g., a blue light, or a specific tone) to an outcome should not require hippocampal activity, whereas tasks that are solved by association between a configural cue (a cue with two or more elements, e.g., both a light and a tone, or several visual cues) to an outcome should be impaired by hippocampal

damage. To test their theory, they proposed several tasks that they predicted would be hippocampus-dependent: Negative Patterning tasks in which both cues (A and B) are rewarded when occurring on their own, but a combination of both cues (AB) is not rewarded (formalised as A+, B+, AB-), or Biconditional Discrimination tasks, in which four cues (A, B, X, and Y) are presented in pairs and correlated with reward and a neutral or negative outcome in such a way that the configurations but not the individual cues reliably signal trial outcomes (AX+, AY-, BX-, BY+).

While the data initially seemed to support their theory (Sutherland & Rudy, 1989), with time, conflicting evidence emerged: multiple studies failed to replicate the negative patterning results of the 1989 study, or could only partly reproduce them (Bussey et al., 2000; Davidson et al., 1993; McDonald et al., 1997). Several of these studies showed not only that the lesioned rats learned the task normally, but the results of transfer trials (behavioural controls) even suggested that the animals were indeed still using a configural strategy to solve the tasks (Rudy & Sutherland, 1995). Even in more difficult task paradigms testing the same hypothesis, results of hippocampal lesions remained variable, depending on the spatial distance between the individual elements to be included in the configural cue (Albasser et al., 2013) or the inclusion of an entire context as one “cue” (Riaz et al., 2017). Taken together, this suggests that hippocampal activity is only required for these tasks under specific, but as yet unknown conditions.

1.5.2 Transitive Inference tasks

Another attempt to make the investigation of the common hippocampal computation tractable uses the previously mentioned advantages of remembering relations between items: the possibility of inference.

Specifically, if the hippocampus encodes not only the items that occur together, but also their relation or order in space, time, or possibly value, this might enable

inference about items that are only indirectly related. For example, if a rodent learns that in a choice between A and B, A should be selected over B (A > B) and further learns that in a choice between B and C, B is the correct option (B > C), the subject can infer that A > C, even if they have never previously encountered that choice (Dusek & Eichenbaum, 1997).

Experiments using this paradigm showed however that this approach has pitfalls: in the scenario described above, A is statistically rewarded 100% of the trials, choosing B equals 50% chance of reward while C is never rewarded. Therefore, rodents can circumvent the need for inference by relying purely on associative learning (Dusek & Eichenbaum, 1997). To avoid this, the chain of overlapping premise pairs thus needs to be at least 5 items long (A > B > C > D > E). While A and E still have different total reward probabilities, B and D are now a test pair that requires transitive inference. While inactivation data indeed shows that BD associations are sensitive to hippocampus lesions while AE trials can still be performed (Dusek & Eichenbaum, 1997; Elzakker et al., 2003; Johnston et al., 2021), only a limited number of trials can be performed on the “test pairs” before they become a learnt association in their own right rather than relying on inference.

Therefore, inference in rodents quickly reaches a level of complexity that goes beyond the scope of feasibility.

Between this example and the aforementioned configural learning, two considerations emerge for future structural learning paradigms. Firstly, these tasks need to be carefully designed to avoid any possible solution of the task through alternative strategies such as tracking of total reward probability. Secondly, since multiple brain regions can be recruited to solve structural learning, even after careful task design the requirement of hippocampus for the task can be sensitive to details such as the specific temporal structure of the task.

1.6 Overview

In summary, in this introduction I have defined structural learning and given an overview of neural substrates that may contribute to it. I've further summarised how hippocampus is placed at a crucial position within this network both from a functional as well as from an anatomical perspective. Finally, I've described past attempts to characterise the specific role of hippocampus in structural learning and detailed some of their limitations.

To understand if the hippocampus is indeed crucial to organising information and enabling inference and generalisation, a carefully controlled structural learning task in a rodent model is needed. With that, it will be possible to interrogate not only what precise external and internal variables are represented by the hippocampus at any point in the task, but to perturb these representations and determine their behavioural relevance.

Therefore, my **first aim** was to design and test a task that can be used to directly test structural learning in mice. In **Chapter 3**, I introduce an odour sequence task for mice that requires subjects to learn about both the identity of an odour and its temporal position within a sequence and show data from control experiments and logistic regression that indicate that this new task indeed fulfils the criteria to test structural learning.

As described above, it has been proposed that structural learning is the underlying mechanism of abstraction and generalisation and conveys greater behavioural flexibility than other types of learning. The **second aim** of my thesis thus was to test whether these skills could be observed in animals performing our task. Since our task allows for manipulation of the temporal structure and identity of cues separately, in **Chapter 4**, I use this ability to probe generalisation to novel cues in the same temporal structure, as well as adaptation to changes

in the temporal patterns associated with a learnt relational structure. Consistent with the proposed advantages of structural learning, I find that after initial learning, mice could rapidly adapt to manipulations of cue identity and temporal structure, suggesting flexible use of previously learnt relational structures.

The ventral CA1 area of the hippocampus is strongly implicated in structural learning. Thus, the **third aim** of my thesis was to test the necessity of vCA1 activity for performance of this non-spatial structural task. To this end, I present data from optogenetic inactivations of vCA1 neurons during task performance in expert mice in **Chapter 4**, showing that optogenetic inactivation impaired task performance only in task paradigms where the delay between the odour cues exceeds 10s, suggesting that vCA1 activity in this task is necessary specifically to maintain and bind information about individual stimuli across long delays.

Driven by this, the **fourth aim** of my thesis was to investigate how neurons in vCA1 integrate distinct sensory and internal variables to form a coherent representation of relational structure that can be used to support behaviour. In **Chapter 5**, I recorded vCA1 activity during the task using calcium imaging through implanted gradient-refractive index (GRIN) lenses. Consistent with the optogenetic data, I found robust encoding of task variables in vCA1 neurons, both at the single cell level as well as when quantifying encoding across the entire population.

Finally, in **Chapter 6** I summarise my findings and discuss how they fit into the current theories of hippocampal function and structural learning in mice.

2 Material and Methods

2.1 Animals

Adult C57BL/6 male and female mice (9 - 11 weeks old) provided by Charles River were used for all experiments. Three groups of animals were used for experiments described in this thesis:

1. Mice used for behavioural proof of concept (implanted with only a lightweight metal head holder (headbar))

2. Mice used for optogenetic manipulations (implanted with headbar and optical fibre, injected with an AAV)

3. Mice used for calcium imaging (implanted with headbar and miniscope base with GRIN lens (ø1mm, 3.8mm long) and injected with an AAV)

All animals underwent stereotaxic surgery and returned to their home cage for at least 1 week to allow full recovery. Animals were housed in cages of 1 to 4 and kept in a controlled environment under a 12h light/dark cycle with *ad libitum* access to food and water (unless stated otherwise). All experiments followed Home Office and University College London guidelines and were in accordance with the Project License (PP2254048) and the Establishment License of University College London (X7069FDD2).

Listed below are all animals whose data is presented in this thesis, split by which figures this data is used in. Additionally, an asterisk marks the four animals whose data was collected by other lab members. The animal marked as “excluded” never reached the performance criterion (“70% correct trials”) in any task with a delay longer than 5s.

Table 2.1 Animals included in Chapter 3: All animals whose data has contributed to the figures shown in Chapter 3. Animals marked with an asterisk indicate that these experiments were carried out by someone other than me.

Index	Mouse ID	Fig 3.2	Fig 3.3a	Fig 3.3b	Fig 3.4	Fig 3.5
1	AM_O*	☑	☐	☐	☐	☑
2	AM_L*	☑	☐	☐	☐	☑
3	DO1*	☑	☐	☐	☐	☑
4	DO2*	☑	☐	☐	☐	☑
5	SN69	☑	☐	☐	☐	☑
6	SN70	☑	☐	☐	☐	☑
7	SN71	☑	☑	☑	☐	☑
8	SN72	☑	☑	excluded	excluded	☑
9	SN73	☑	☑	☑	☑	☑
10	SN74	☑	☑	☑	☑	☑
11	SN75	☑	☑	☑	☑	☑
12	SN76	☑	☐	☐	☐	☐
13	SN78	☑	☐	☐	☐	☐
14	SN79	☑	☐	☐	☐	☐
15	SN80	☑	☐	☐	☐	☑

Table 2.2 Animals included in Chapter 4: All animals whose experimental data is presented in Chapter 4.

Index	Mouse ID	Fig 4.1	Fig 4.2	Fig 4.3	Fig 4.4	Fig 4.5	Fig 4.6	Fig 4.7	Fig 4.8
1	SN74	☑	☑	☑	☑	☑	☐	☐	☐
2	SN75	☑	☑	☑	☑	☑	☐	☐	☐
3	SN76	☑	☑	☑	☑	☑	☐	☐	☐
4	SN78	☑	☑	☑	☑	☑	☐	☐	☐
5	SN79	☑	☑	☑	☑	☑	☐	☐	☐
6	SN80	☑	☑	☑	☑	☑	☐	☐	☐
7	SN87	☑	☑	☑	☑	☑	☐	☐	☐
8	SN89	☑	☑	☑	☑	☑	☐	☐	☐
9	SN91	☑	☑	☑	☑	☑	☐	☐	☐
10	SN92	☑	☑	☑	☑	☑	☐	☐	☐
11	SN96	☐	☑	☑	☐	☐	☐	☐	☐
12	SN97	☐	☑	☑	☐	☐	☐	☐	☐
13	SN98	☐	☑	☑	☐	☐	☐	☐	☐
14	DO3*	☐	☐	☐	☐	☑	☐	☐	☐
15	DO4*	☐	☐	☐	☐	☑	☐	☐	☐
16	DO5*	☐	☐	☐	☐	☑	☐	☐	☐
17	SN71	☐	☐	☐	☐	☐	☑	☐	☐
18	SN73	☐	☐	☐	☐	☐	☑	☐	☐
19	SN74	☐	☐	☐	☐	☐	☑	☐	☐
20	SN75	☐	☐	☐	☐	☐	☑	☐	☐
21	SN78	☐	☐	☐	☐	☐	☑	☐	☐
22	SN79	☐	☐	☐	☐	☐	☐	☑	☑
23	SN80	☐	☐	☐	☐	☐	☐	☑	☑
24	SN69	☐	☐	☐	☐	☐	☐	☑	☑
25	DO6*	☐	☐	☐	☐	☐	☐	☑	☑

Table 2.3 Data used in Chapter 5: List of animals whose experimental data is presented in Chapter 5.

Index	Mouse ID	Fig 5.2 - 5.10
1	SN87	☑
2	SN96	☑
3	SN98	☑

2.2 Stereotaxic surgery

Stereotaxic surgeries were carried out on 7 - 12 week old mice anaesthetised with isofluorane (Isofluorane 100%; Piramal Critical Care) according to previously described protocols. For induction mice were placed in a red Perspex chamber (AN010ASR; VetTech) with 1L/min flow of 4% vaporised isofluorane (in medical oxygen, 99.5% minimum purity). Following induction, fur on scalp was shaved (ChroMini Pro; MOSER) and the mice were placed into a stereotactic frame (Model 902 Small Animal Stereotaxic Instrument; KOPF) onto a feedback-controlled thermal control unit (50-7001; Harvard Apparatus) which was maintained between 35 and 37°C throughout the surgery. During induction and throughout the surgery, the induction chamber and the stereotaxic frame were connected to an activated carbon scavenging filter (Cardiff Aldasorber; Shirley Aldred & Co) and an active scavenging unit (Model AN005; VetTech). The animal's eyes were protected from desiccation using artificial tear ointment (Viscotears Liquid Gel).

The scalp was sterilised with HiBiSCRUB®. An incision was made to expose the skull from bregma to lambda. After application of a few drops of the local anaesthetic Marcain (0.025% in sterile saline), the connective tissue was removed by applying hydrogen peroxide with sterile cotton buds. After ensuring horizontal alignment using bregma and lambda skull landmarks, small holes were drilled in the skull at the coordinates of interest (see **Table 2.1**) using a stainless steel bur (19008-07; Meisinger) attached to a miniature drill (Ideal

Micro-Drill™; CellPoint Scientific). For the duration of the surgery, anaesthesia was maintained at the same flow rate as listed above, with the isofluorane concentration was brought down to 1-2%. Stereotaxic coordinates are listed in Table 2.4.

Region	RC	ML	DV
Dorsal hippocampus	−3.2	±1.7	−1.3
Ventral hippocampus	−3.7	±3.2	−4.5

Table 2.4: Stereotaxic injection coordinates
Injection coordinates taken from Mouse Brain Atlas (Paxinos & Franklin, 2019). All coordinates given in mm calculated relative to bregma. RC: rostrocaudal, ML: mediolateral, DV: dorsoventral

Injections were carried out with a Nanoject II (Drummond Scientific) using longshaft borosilicate glass pipettes with a tip diameter of ~ 10 - 50µm, back-filled with mineral oil and front-filled with ~ 1µL of the substance to be injected. A total volume of 250 - 500nL of each virus was injected in increments of 14nL or 28nL in 15s intervals. Following infusion of the virus, the pipette was left in place for an additional 5 minutes before being slowly retracted. Viruses used in the experiments described throughout the thesis are listed in Table 2.5.

Virus	Source
AAV1-CaMKII-Cre	Addgene, 105558
pAAV- FLEX-ArchT-tdTomato	Addgene, 28305
pGP-AAV-syn-FLEX-jGCaMP7f-WPRE	Addgene, 104492

Table 2.5: Viral constructs
Constructs used for the experiments described in this thesis

For animals undergoing optogenetic manipulations, a total of volume of 300nL of a mixture of pAAV-FLEX-ArchT-tdTomato, and AAV2/1-CamKII-Cre in a ratio of 1:1 was injected into the target region. Then, a fiber optic cannula (ø200µm core diameter, 0.39 NA, uncleaved, cut to approximately 4.5 mm; CFML12U-20, Thorlabs) was implanted bilaterally directly following virus injection in the same surgery.

For calcium imaging, a mixture of AAV2/1-CamKII-Cre, AAV-syn-FLEXjGCaMP7f and sterile saline was injected (ratio 1:1:2). Then, a GRIN (Gradient Index) lens (ø1mm, length 3.8mm; G1P10, Thorlabs) was implanted directly over the injection site. To do this, the craniotomy for injection was made larger: approximately 1mm in diameter, centring around the injection coordinates. After carefully removing bone fragments, the dura was punctured with a needle tip and dura shreds as well as any bone fragments were removed with forceps. During this time, the brain surface was kept moist by applying sterile saline. With a blunted 30-G needle connected to a vacuum pump (Compton Compressors), the cortical tissue above hippocampus was aspirated while simultaneously the craniotomy was kept irrigated with sterile saline from a 10ml syringe with a 30-G needle tip. Aspiration was performed slowly, layer by layer, until the white colour and typical striation of corpus callosum changed to the darker colour of CA1. Once the desired depth was reached, sterile silicone sponges soaked in saline were placed on the tissue for 5-10 minutes to allow any bleeding to subside while keeping the tissue from drying out.

Then, a miniscope base with the GRIN lens fixed in place (see **Chapter 2.4.5.1**) was attached to a stereotaxic arm, and lowered slowly until it was gently seated on the surface of the exposed hippocampus. Over several minutes, the lens was lowered to the final DV position.

To aid cement attachment, the skull was roughened, and two metal screws were inserted into the skull. Both fibre implants for optogenetic manipulation as well as the miniscope base for calcium imaging were secured to the skull by applying two layers of adhesive dental cement (Superbond, C&B). On all animals, an additional custom-made lightweight metal head holder (headbar, designed by the Svoboda lab (Guo et al., 2014) and manufactured by the UCL mechanical workshop) was attached to the skull surface. For animals only undergoing behavioural training, a ø1.25 mm stainless ferrule was attached to the middle of

the headbar during the adhesion process to facilitate handling of the animal later.

During the surgical process, mice were subcutaneously injected with carprofen (0.05mg/mL in saline) to reduce inflammation risk and provide pain relief. After the procedure, they were allowed to recover from anaesthesia in a heated chamber for a minimum of 30 minutes, or until they were fully ambulatory, before they were returned to their home cage. Animals received carprofen in their drinking water (0.05 mg/mL) for 48 hrs post-surgery.
For optogenetic manipulations as well as calcium imaging, viral expression was allowed to occur for at least 3 weeks before the experiment. The locations of injection sites were verified for each experiment histologically.

2.3 Anatomy

2.3.1 Histology

Mice were anaesthetised with 100mg/kg ketamine (KetaVet) and 10mg/kg xylazine (Zoetis) in 0.5mL sterile saline (BAYER). Following confirmation of deep anaesthesia, animals were transcardially perfused with ice-cold 4% paraformaldehyde, the brains were dissected and fixed in 4% paraformaldehyde overnight at 4ºC. Brain samples were transferred to phosphate buffered saline (PBS, pH 7.2) after overnight fixation.

Viral expression in the animals injected with AAVs was confirmed by imaging the brain using serial section 2-Photon microscopy (BrainSaw) in which whole brains are embedded in agar. In this technique, the face of the sample is automatically sliced and each coronal section is imaged, repeating these steps until the whole sample has been imaged. The microscope is controlled by ScanImage Basic (Vidrio Technologies, USA) using BakingTray (available on the BakingTray github), a custom-written software wrapper for setting up the imaging parameters.

Images were assembled using StitchIt (available on the StitchIt github). Both the microscope as well as the software wrappers were designed and assembled by Rob Campbell (Head of Advanced Microscopy Facility at the Sainsbury Wellcome Centre). Use of the BrainSaw system and the attached interfaces was done with assistance of trained staff from the Microscopy Facility and/or senior members of the lab of Athena Akrami.

2.4 Behavioural Studies

2.4.1 Experimental Setup

All animals were trained on an olfactory paired-associates task. For this task, animals were placed on a cylindrical treadmill consisting of a 3D printed wheel (ø15cm, width: 7cm) suspended through a metal axis, allowing for 1D rotation (**Figure 2.1**) This setup including parts numbers and 3D printing files has been described in detail before (Marbach & Zador, 2017). Mice were headfixed on the treadmill by attaching the implanted headbar to a custom-made metal holder with adjustable height and angle.

The odour delivery system was based on a modular design adapted from the Schaefer lab (Ackels et al., 2021). A constant stream of clean air (~0.7L/min) was carried to the behavioural rig via tubing through a custom-made odour manifold and was then supplied to the mouse through the right side of a custom-made 3D printed tube holder. From the left side of the same holder, air was constantly removed from the setup via tubing connected to a vacuum pump (Compton Compressors), thus creating a constant stream of flowing air in front of the animal's head.

The odour stimuli themselves were supplied from an odour manifold which consisted of a 12.2 x 3.2 x 1.5cm stainless steel block with four milled indentations (ø2cm). Within each of these indentations was a threaded through hole for the installation of an input flow controller (AS1211F-M5-04, SMC) and an

output connector (IKTX0322170A & IKTX0322190A, The Lee Company). For each inset, the cap of a 15ml glass vial (#27003, Sigma Aldrich) with the centre removed was pushed into the indentation and sealed with epoxy resin (Araldite Rapid, Huntsman Advanced Materials), allowing for 15ml glass vials to be screwed in and out of the insets for rapid replacement.

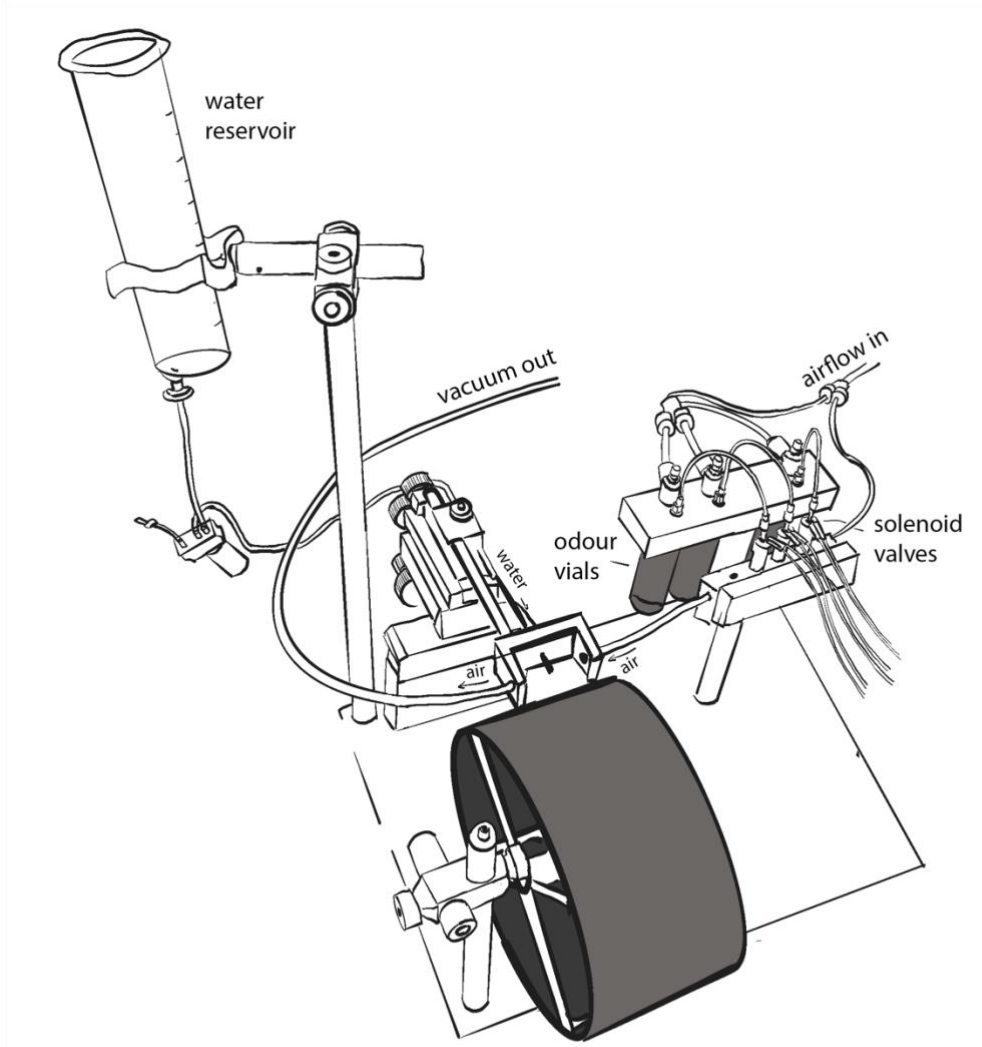

Figure 2.1: Schematic of Experimental Setup. Cylindrical treadmill on top of which a custom-made tube holder supplies clean or odourised air from the right which is then sucked out to a vacuum pump from the left. From the front, water rewards can be delivered.

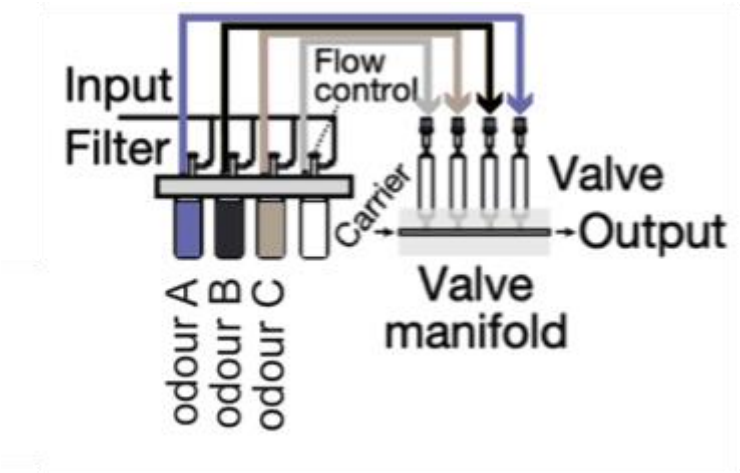

Figure 2.2: Schematic of odour manifold. The system for odour delivery consists of a carrier airflow in which odourised air from three different odour vials can be injected.

To generate airflow through the olfactometer, a pressurised air source was split into two separate lines: the input line and the carrier airflow (Figure 2.2). The carrier airflow was connected to the input of the odour manifold. The input line was split further into three lines connected to each odour position of the odour manifold. Odourised air was created by opening the VHS valves (INKX0514750A, The Lee Company) connected to one of the odour positions, thus supplying air to either of three glass vials containing three of the odorants listed in Table 2.3, undiluted unless stated otherwise.

During odour stimulation, opening a VHS valve would inject odourised air (air that travelled through either of the odour vials) into the airflow. After 1s, the valve would close, returning the air flow to only clean air, thus quickly clearing out any residual odours. Both the clean air as well as odourised air was supplied at a flow rate of ~0.7L/min, measured with a flowmeter (2510 Flowmeter, Brooks Instrument).

All odours were chosen from a list of odorants that have been shown to be neither appetitive nor aversive to mice (Root et al., 2014), and the combinations were selected in a way that ensured sensory separability (Pashkovski et al., 2020).

Odour	Source	Product Code
Acetophenone	Sigma-Aldrich	A10701-1L
4-Allyanisol	Sigma-Aldrich	A29208-100G
Amyl Acetate	Sigma-Aldrich	W504009
Eucalyptol	Sigma-Aldrich	C80601
2-Phenolacetate	Sigma-Aldrich	77861-1L

Table 2.6: Odour components used as behavioural stimuli in paired-associates task**.**

Odorant concentration in the open air was measured through a mini photoionization detector (PID; 200B mini PID; Aurora Scientific) located in the position of the mouse snout. The PID device ionises the volatile components in the air with high-energy UV light. The now charged particles register as a current in an electric field, giving a readout of their concentration in the air. Different odour components have different ionisation potentials, therefore the amplitude of the current measured by the PID cannot be compared between different odours. It can however be used to compare concentrations of the same odour across time, as well as measure how long it takes for the concentration of odour to return to negligible levels after odour delivery.

To compensate for this limitation, we additionally used an airflow sensor (AWM-5000 series; Honeywell) to make sure that the amount of odourised air injected into the airflow was the same for each odour. This sensor works on the principle that the airflow across the sensor causes heat transfer, which is then output as an analogue voltage signal. An example measurement taken with the PID as well as example traces measured by the airflow sensor are shown in **Figure** 2.3.

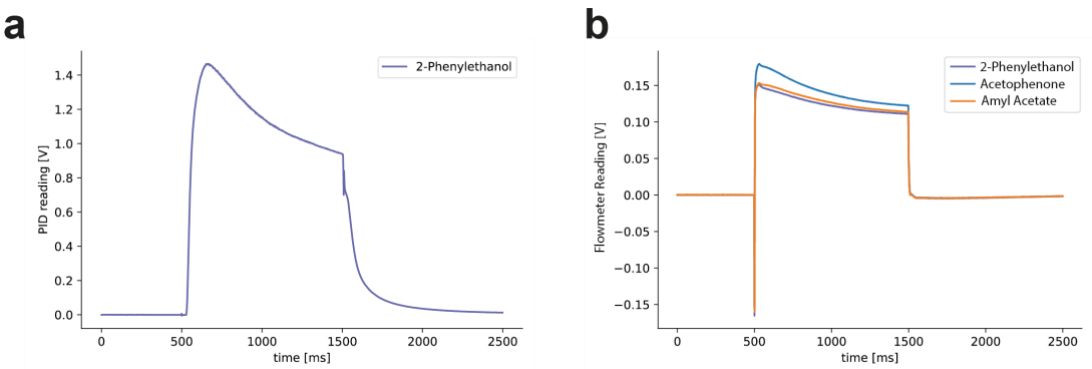

Figure 2.3: Example odour traces from Photoionisation device and airflow sensor. a. Average mini-photoionisation detector signal in response to 1s injection of air odourised with 2-Phenylethanol into the carrier airflow. **b**. Average change in air flow in response to the injection of all three odour cues into the air stream. Note that the sharp downwards spike is an electrical artifact caused by the electrical noise of the valve opening.

In the middle of the custom-made tube holder, a metal lick spout was positioned at a suitable distance from the headbar holder. Licking was detected using a printed circuit board operating as a lickometer (Capacitive Breakout Board, SparkFun). One end of the circuit was attached to the metal lick spout and the other to ground. The board was supplied with 5V. Whenever the mouse would lick the tube, an electrical circuit would close, creating a voltage drop that was recorded as a continuous analogue signal (RSE: Referenced Single-Ended, i.e. voltage is measured against ground provided by the device). Water droplets (~10µl) were released by a 3-way solenoid valve (LFRA1220370D, The Lee Company) and delivered to the mouse through the lick spout from a reservoir installed at 15cm height to ensure gravity flow (i.e. no pump was required).

The behavioural rig was controlled with custom-written software (Arduino), through a microcontroller (Arduino Uno, Arduino) and a data acquisition board (NI USB-6001, National Instruments). Odour valves were opened by sending a 5V TTL pulse from the Arduino to a custom-built spike-and-hold driver. Each driver could provide a 0.5ms 24V pulse to cause the valve to open and then maintained its opened position with a 3.3V holding voltage of 1s duration. Similarly, for the opening of the water valve, another custom-made driver converted the 5V Arduino TTL pulse into a 12V square pulse of a variable length (duration

calibrated to equal ~10µL water reward). The licking as well as the odour pulses were recorded through the NI board. Both stimuli and licks were visualised in LabView (Version 20.0f1, National Instruments).

2.4.2 Behavioural training and recording protocol

After recovering from the surgical procedures for a minimum of 7 days, mice were water-restricted to approximately 85% of their ad libitum body weight. After at least a week of water-restriction and habituation to manual handling by experimenter, training for the olfactory paired-associates task began. Throughout the training period, mice were provided with 1-1.5mL of water daily, either throughout or after each session.

Mice were initially habituated to head fixation on a static surface for two consecutive days, followed by two days of habituation to head fixation on the treadmill, starting at 1-2 minutes of head fixation and gradually increasing the time to 10 minutes. On the fifth day, the lick spout was placed in its normal position and mice were given water drops through the spout at randomly chosen time intervals of 1s, 5s or 7s (**Figure 2.4** left).

In the first stage of training (shaping stage, **Figure 2.4** middle), mice were presented only with the rewarded pairs of odours (AB, BC, CA), separated by a 5s delay. A water reward was delivered 1.5s after the second odour cue in a Pavlovian fashion, i.e. water was delivered irrespective of the mouse's licking behaviour. Between trials, there was a 10s pause (inter-trial interval, ITI). Mice were kept on this training protocol until they started licking in anticipation of the reward in the response window between odour cue and water delivery (1.5s) in at least 70% of trials.

Mice were then trained on the full task: All six types of trials (3 rewarded odour pairs, 3 unrewarded odour pairs) were presented pseudorandomly, ensuring that

all trial types were presented in each block of 6 trials (**Figure 2.4,** right). Mice were not punished when they licked on trials with non-rewarded odour pairs, but no reward was presented. Within ~8-10 days, mice learned to refrain from licking in non-rewarded trials. During training, mice gradually refrained from licking during the delay period between odours as well, and well-trained mice would typically initiate licking right after the second odour.

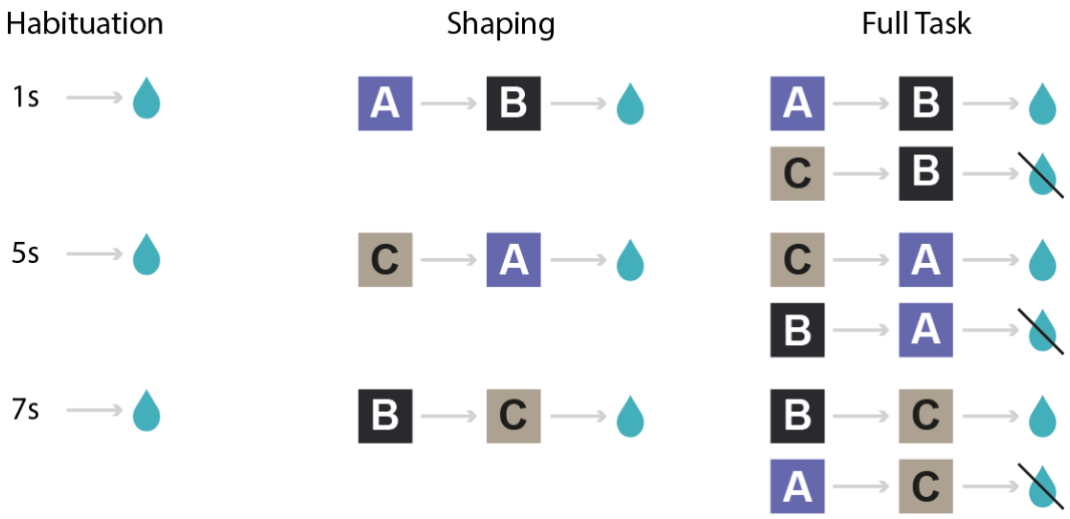

Figure 2.4: Training Protocol. After habituation to the process of head fixation and the treadmill, mice were given water rewards after pseudo-random time intervals of 1s, 5s, or 7s. When they reliably consumed the water reward, they were exposed to a shaping stage in which only rewarded odour pairs were presented, followed by a reward 1.5s after the second odour. Once the mice would predictively lick before the reward delivery in at least 70% of trials, they were moved on to the full task with counterbalanced presentation of rewarded and unrewarded odour pairs.

Responses were assessed based on licking during the response window only. As previously described, licks were detected with a capacitive breakout board.
The board records a binary analogue signal. Any signal above the threshold voltage (4V) is therefore counted as a lick, and any signal below is recorded as an absence of licking. If any licking occurred during the response window (1.5s after presentation of the second odour of a rewarded pair), the trial was counted as a hit. If licking occurred in the response window after presentation of an unrewarded odour pair, the trial was counted as a false alarm instead. In the case where no licking was detected in the response window, the trials were labelled as miss or correct rejection, respectively.

After mastering the first version of the task with a 5s delay between odour stimuli, the mice were then taken through variations of the task. One variation consisted of changes in the delay time separating the two odour cues. For this manipulation, after the mice had acquired the full task, we extended the delay to 10s. After they successfully performed at this longer delay, we then trained them on a version with a 20s and finally a 30s delay. For all versions of the task, the Inter-Trial Interval (ITI) was kept at twice the length of the delay time, such that in trials with a 10s delay between odours, the ITI would be 20s for example.

A second behavioural manipulation consisted of changes in odour identity. In this paradigm, we exposed mice that had acquired the full task with a delay of 10s to a new set of odours, replacing odour C with an unfamiliar odour C2.

Furthermore, multiple control experiments were conducted (see **Chapter 2.4.3**). Performance on each day was quantified as the percentage of correct trials (sum of hits and correct rejections divided by total trials) over the entire session. A session would typically run for approximately 80 trials, corresponding to a duration of 20min in the protocol with 5s delay between odours. In protocols with longer delays, the sessions would be longer, ensuring a similar number of trials in each session.

2.4.3 Behavioural control experiments

For control sessions with auditory cues only, the airflow was turned off completely and the tube delivering odours to the mouse was disconnected. Therefore, the only cues available to the mouse were the clicks of the valves opening. For control sessions with sensory and auditory cues, the carrier airflow was active, but the odour vials were removed and replaced with empty glass vials. This way, the mice experienced the changes in air pressure associated with opening and closing the odour valves as well as the resulting clicks, but no odour

cues were present. The system was flushed of any residual odour for at least 30min before each behavioural session.

To control for potential mixing of successive odour cues, a 1:10000 dilution of odour C was filled into the 4th port of the odour manifold. On some trials ("dilution trials"), instead of a sequence of "odour C – delay – odour A or B", the diluted odour was instead added into the airflow throughout the entire time of the trial up to the presentation of the second odour ("diluted odour C – odour A/B"). Thus, in these trials the mouse was presented with a mixture of a low concentration of odour C, directly followed by a larger concentration of either odour A or B (indicating a rewarded and non-rewarded trial, respectively, according to the same rules as in the standard task). If the animals had indeed learned to pay attention to a configural cue made up of a mixture of residual first odour combined with the second odour cue, they should be able to perform well on those dilution trials.

2.4.4 Optogenetic manipulations

In behaviour sessions with optogenetic inhibition, the subject was head-fixed on the treadmill as usual. Then, a bifurcated optic fiber (core = 200 µm, NA = 0.22, Doric) was attached to the ferrules at the end of the implanted fibre optic cannulae using interconnects (ADAL3, Thorlabs). The fiber was connected to a fiber-optic rotary joint (FRJ 1x1 FC, Doric) that was in turn connected to a 532nm green laser (Shanghai Laser & Optics Century Co.).

Using TTL pulses from the Arduino microcontroller, the laser would be controlled at a pre-determined time relative to the behavioural cues, turning on in a square pulse pattern, followed by a 500ms long linear downward ramp to prevent a post-illumination rebound burst of action potentials (Chuong et al., 2014). The square pulse would start 500ms before the first odour cue, continuing until the end of the response window and reward delivery, at which time the light would ramp

down over the course of 500ms. The laser power was calibrated to lie at approximately 5mW/mm^2 at the end of the optical fiber.

2.4.5 Recording of GCaMP signals

2.4.5.1 Preparation of implants

To measure neuronal activity in CA1 we recorded calcium activity in GCaMP7f-expressing neurons with a head-mounted miniature microscope with epifluorescent light source (miniscope; Open Ephys/UCLA), as described in a study by Zhang et al. (Zhang et al., 2019). To image fluorescence in a deeper structure such as ventral hippocampus, the implantation of a Gradient Refractive Index (GRIN) lens is required. This type of lens is most often rod-shaped and uses gradual variation of the refractive index within the lens material to focus light from one side of the lens at a specific point on the other side of the lens. For this approach to work, it is essential that the imaging plane, i.e. the lens of the miniscope (OpenEphys/UCLA), is placed at a pre-defined and reproducible distance from the GRIN lens. This is ensured by the miniscope baseplate, a threaded cylinder that is fixed on the animal's skull directly over the implanted GRIN lens which allows to reversibly mount the miniscope at a specific distance from the GRIN lens with a set screw. Furthermore, whenever the miniscope is not mounted, the baseplate can be capped with an acrylic lid which protects the GRIN lens from dust and scratching.

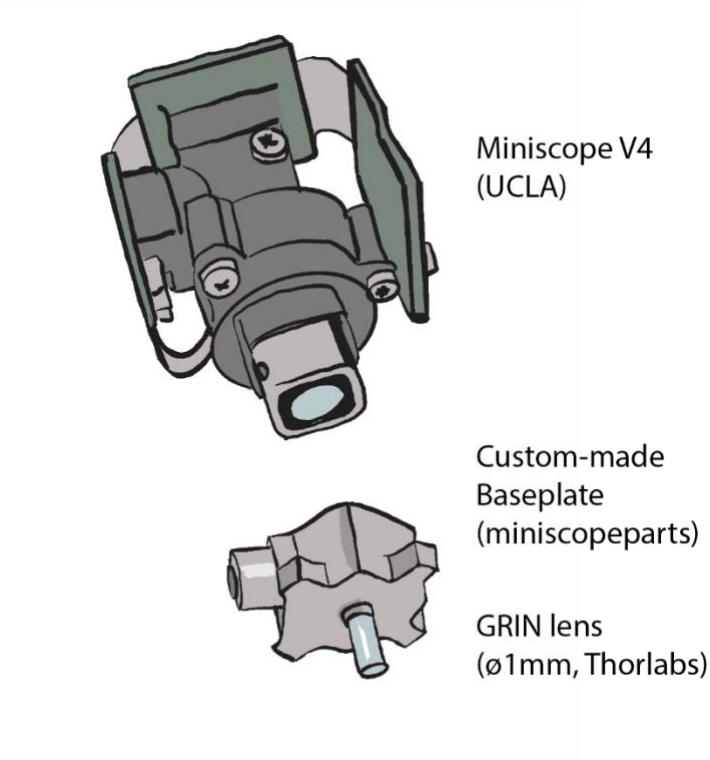

Figure 2.4: Schematic of miniscope, baseplate and GRIN lens. The GRIN lens was fixed into the miniscope base and the construct was then implanted onto the mouse's skull. The miniscope could be reversibly mounted and unmounted onto this base, allowing for a consistent imaging window.

In all experiments described below, we used a novel type of baseplate (V4-V2GrinBasePlate, made to order by miniscopeparts) where the GRIN lens (ø1mm, length 3.8mm; G1P10, Thorlabs) is fixed into the miniscope base prior to the surgery such that both GRIN lens as well as the baseplate can be implanted in the same surgical step. To ensure that the optical path of the screw-mounted miniscope is aligned with the GRIN lens at the correct distance, a slide with fluorescent material is used when glueing the GRIN lens into the baseplate. Only when a sharp image of the material on the slide is formed at the right focal distance, the lens is fixed in place with superglue (Loctite). The baseplate with the GRIN lens now fixed in place was stored in a padded box for optical components until it was implanted.

2.4.5.2 Recording setup

As mentioned above, we used a combination of a miniature microscope (miniscope, Open Ephys/UCLA) and a GRIN lens to record calcium activity in ventral hippocampus. To briefly summarise the technique, the miniscope uses blue LED light to excite green-fluorescing fluorophores such as GCaMP7f and detects the emitted fluorescence light with a C-MOS sensor. In our setup, the miniscope was connected to a data acquisition board (Miniscope DAQ v3.3; OpenEphys) that in turn was connected to a computer running the miniscope software (Miniscope-DAQ-QT-Software; Aharoni Lab). The DAQ board was also connected to an Arduino microcontroller running the behaviour protocol, allowing for miniscope recordings to be triggered in synchrony with task events. Specifically, the microcontroller sent a TTL pulse to the miniscope DAQ both 500ms before the first stimulus presentation and upon ending the behaviour protocol, thus establishing two reference points to match the timestamps of behavioural events to the frames of the miniscope recording.

In behaviour sessions with calcium imaging, the mouse was head-fixed on the treadmill as usual. After loosening the set screw on the previously implanted miniscope base, the protective cap could be removed, and the GRIN lens surface was cleaned with a sterile Q-tip dipped in methanol. Then, the miniscope was screwed onto the base. Using the miniscope software, the LED power was adjusted for each animal such that fluorescent structures were visible, but not saturating. Furthermore, the electronic focus (+/- 200 µm) and imaging gain was selected to optimise the sharpness of fluorescent structures in the field of view. These settings were conserved for each animal across recording sessions. The behavioural protocol was started as usual and automatically triggered the start of the miniscope recording via the previously described TTL pulse.

2.5 Data Analysis

2.5.1 Behavioural analysis

In all experiments, behavioural performance was quantified as the percentage of correct trials (p correct), i.e. trials in which the mouse either correctly displayed anticipatory licking after the presentation of a rewarded odour pair, or correctly withheld licking in response to an unrewarded pair.

As a measure of sensitivity, d prime (d') was used additionally in some experiments. Defined as

$$d' = z(H) - z(FA) \qquad (2.1)$$

with H = P (lick | rewarded odour pair), FA = P (lick | unrewarded odour pair) and z as the gaussian z-transform.

In contrast to p correct, d' takes into account whether there are large discrepancies between the percent of correct licking and correct withholding responses and controls for the overall level of licking behaviour.

2.5.1.1 Error analysis

Given the structure of our task (go/no-go), there are four possible outcomes for any trial: when a rewarded odour pair is presented, the mouse can either lick in anticipation ('Hit') or fail to do so ('Miss'). Conversely, when a non-rewarded pair occurs, the mouse can withhold licking ('Correct Rejection') or lick anyway ('False Alarm').

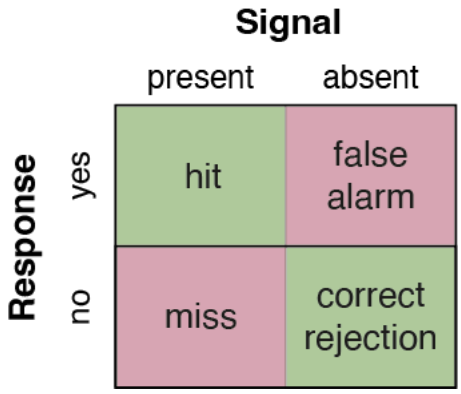

Figure 2.5: Possible outcomes in go/no-go task. There are two components that determine the outcome of any trial: the presented odour pair (signal) and the response (lick/no lick). These can be classed into correct trials (hit and correct rejection, green) and errors (miss and false alarm, red).

By distinguishing between the two types of errors laid out in this model, we can further characterise the effects of any manipulation that reduces the animal's performance of the task, such as changes in task structure (**Chapter 3**) as well as optogenetic manipulations (**Chapter 4**).

2.5.1.2 Logistic Regression

To quantify the influence of different task components (single odours, odour configurations, past choice, etc.) we used logistic regression (Akam et al., 2021; Parker et al., 2016). These models try to classify trials using a logit transformation to combine differently weighted predictor variables into a single probability (target variable), in our case describing whether the mouse is likely to display anticipatory licking. This transformation uses the following formula:

$$log\frac{L(i)}{1-L(i)} = \beta_{X1}X1 + \beta_{X2}X2 + \cdots + \beta_0 \qquad (2.2)$$

with *L(i)* as the probability of licking on a given trial i, *X1, X2*, … as the predictor variables (e.g. "odour A present" or "previous trial rewarded"), β_{X1}, β_{X2}… as the regression coefficients of each predictor and β_0 as a constant bias-term.

The predictor variables used in our task were either related to the presence of a given odour in the trial in any position (A, B, C), to the odour configuration disregarding order (A and B, B and C, or C and A), the odour pairs taking order into account (AB, AC, BA, BC, CA, CB), trial history (previous reward, previous lick) as well as a general bias term.

In addition to a complete model using all the above predictor variables, we also split these predictor variables into categories to build models that only had access to some of the variables. Our "elemental" model (E) predicts licking behaviour only based on the odour identity without any information on the position. A "configural" model (C) uses the configuration predictors, whereas the "history" model (h) tries to predict licking solely based on previous outcome and previous choice. Lastly, the "structural" model (s) uses the six ordered odour pairs as predictor variables.

For all models, we used the *scikit-learn* Python library. We first split the data into training data (75% of data) and test data (25% of data). We then used the *GridSearchCV* function to find the best hyperparameters for the regression. With those parameters, we instantiated the Logistic Regression model with these hyperparameters and trained it on the labelled training data (i.e. both the matrix containing the trial events used as predictors as well as the actual behaviour of the mouse). We then tested the performance of the model by comparing the model predictions on the unlabelled test data to the actual mouse behaviour.

To compare the performance between classifiers, we used the area under the curve (AUC) as a quantifiable measure. This measure comes from the receiver operating characteristic curve (ROC curve) which plots the sensitivity (true positive rate) against the probability of false alarm (false positive rate). The AUC, that is the integral of the curve over the entire ROC space, ranges from 0 to 1. A perfect classification would amount to an AUC of 1, and a random binary

classifier would yield an AUC of 0.5. This measure is therefore useful to compare the accuracy of models that use different sets of predictor variables.

The second metric we used to compare predictive value across models using different sets of predictor variables is delta AUC (ΔAUC), a measure derived from the single-model AUC described above. For ΔAUC, we take a complete model using all aforementioned predictor variables and calculate its AUC. We then remove the predictor variables of interest (e.g. all those associated with reward history) and compare the difference between the AUC of the resulting model and the AUC of the complete model. This difference is called ΔAUC and measures the change in model performance that results from removing specific features.

Lastly, to understand how specific variables affect the overall prediction, we look at the beta coefficients (seen in formula 2.2). The absolute value of a beta coefficient is a measure of how much weight the classifier assigns to this variable, and the direction of correlation is evident from the coefficient's sign. Positive beta coefficients denote a positive correlation of the variable in question with the target behaviour (licking) whereas negative beta coefficients indicate a negative correlation. In other words, a positive beta value for the predictor variable "AB" denotes that the presence of this pair increases the likelihood of mice licking, whereas a negative value would denote that the presence of AB decreases it.

2.5.2 Analysis of GCaMP signals

All miniscope recordings were conducted using the miniscope software (Miniscope-DAQ-QT-Software; Aharoni Lab). The data was then processed, and the calcium signals extracted using the semi-automated Minian pipeline, (Dong et al., 2022) the steps of which are outlined below.

2.5.2.1 Data processing

As a preparatory step for the analysis, raw data (**Figure 2.7a**) first underwent noise removal using a script provided by the Aharoni lab which utilises several low-pass filters and 2-D fast Fourier transforms (2D-fft) to filter out high-frequency electrical noise introduced by hardware components of the V4 Miniscope (Aharoni Lab).

Then, the denoised videos were passed through a pre-processing stage (**Figure 2.7b**) that included downsampling the video by sub-setting and averaging the frames. The video was then cropped to only include parts with bright signal inside the GRIN lens, and the frame rate was dropped from 15 fps to 7.5fps. In the next part of the script, the minimum fluorescence value for each pixel was subtracted from every frame to remove background glow and vignetting. Salt-and-pepper noise on each frame was then removed by passing the data through a median filter. The last step of pre-processing was a background removal step introduced in MIN1PIPE (Lu et al., 2018).

The pre-processed video was then passed through a standard template-matching algorithm, thereby correcting for lateral motion by evaluating cross-correlation between each frame and a reference frame (**Figure 2.7c**).

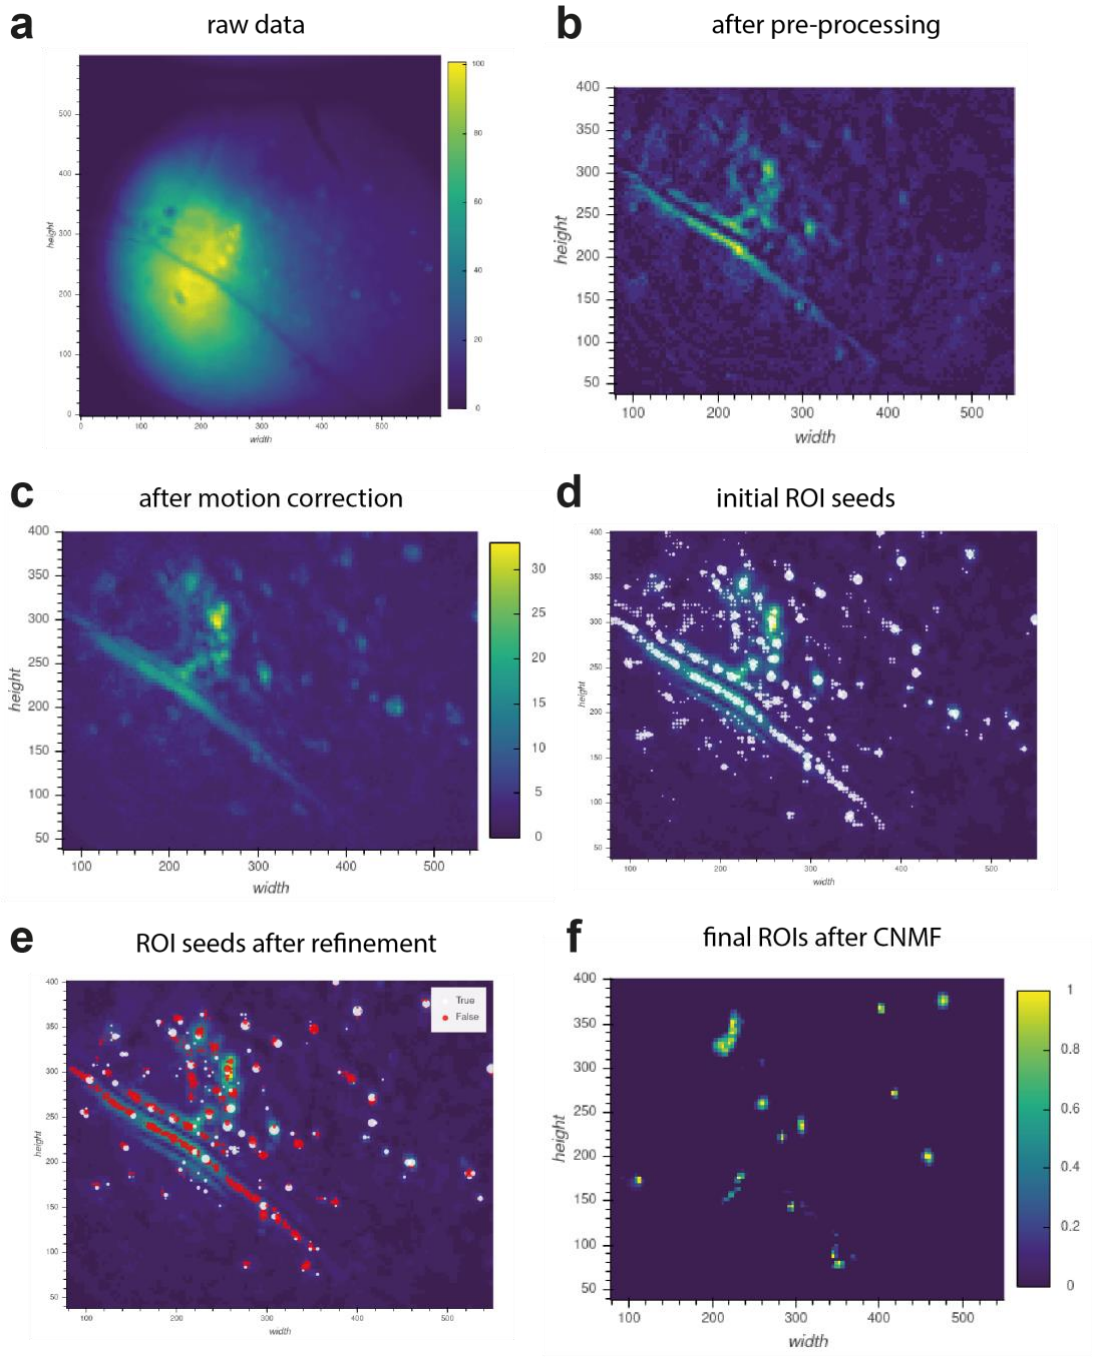

Figure 2.6: Minian analysis pipeline. a. Raw images from miniscope recording. **b.** Image after denoising and glow-removal. **c.** Image averaged over whole session after motion correction. **d.** Initial ROI seeds (white dots) **e**. ROI seeds after correlation- and signal-to-noise-based refinement (red: dropped seeds, white: true seeds) f. Final ROIs overlaid with average calcium signal after constrained non-negative matrix factorization (CNMF) algorithm.

To generate possible locations for neurons, the Minian script then looked for local maxima in the fluorescence across frames, assigning these as "seeds". This set of seeds was over-complete by design, containing many seeds that weren't biologically relevant regions of interest (**Figure 2.7d**). The seeds were then refined based on various metrics, such as a minimum change of fluorescence over time as well as the signal to noise ratio and correlation with neighbouring pixels (**Figure 2.7e**).

After pruning and merging the seeds, the putative ROIs were passed through a constrained non-negative matrix factorization (CNMF) algorithm. The algorithm refines the spatial outline of the cells (spatial update) and then denoises and deconvolves the individual traces (temporal trace). Those two steps are then performed again, until a satisfactory result is produced upon visual inspection (**Figure 2.7e**).

The calcium traces exported after this step were the basis of all further analysis.

2.5.2.2 Analysing single-cell selectivity

To evaluate whether individual neurons respond selectively to specific task events, we calculated a selectivity index using a formula adapted from (Ahmed et al., 2020):

$$SI = f_+ - f_- \tag{2.3}$$

Here, f_+ represents the average activity of the neuron during the examined trial time period across all trials containing the task event of interest; f_- represents the average activity during all trials without the event. For example, to evaluate the selectivity of a neuron for odour A as the first odour, we would take the average activity on each trial during the presentation of the first odour, and then subtract the mean response in trials with odour A as a first odour from the mean response in all other trials.

To mitigate the influence of spurious differences in calcium activity due to limited trial numbers, we compared this selectivity index (SI) to indices derived from 1000 shuffled datasets, where the labels of trial types were randomly reassigned. To compare selectivity indices across cells and task events, we calculated a measure sigma using the below formula:

$$\sigma = \frac{m_s - SI_t}{std_s} \qquad (2.4)$$

Where m_s is the mean of the shuffled SI distribution, std_s is its standard deviation and SI_t is the true selectivity index. We considered all cells showing $\sigma > 2.5$ as selective for a given event.

2.5.2.3 Population encoding of behavioural events

In order to test whether the population activity encodes task events (such as “Odour A First” or “Rewarded trial”), we used Support Vector Machines (SVMs). SVMs are a set of supervised learning methods used for classification. In our case, we trained classifiers to predict the presence/absence of the task event (e.g. “Odour A First”) from the neural data of either a single session or multiple sessions.

To that end, we first calculated the population activity at a given time point by averaging the activity of each neuron within the period of interest (e.g. during the presentation of the first odour). Thus, for each trial, every neuron only had one value, and the entire population could be represented as an n-dimensional vector with n as the number of recorded neurons.

In order to combine data from several mice into one model, we built pseudo-populations by matching trials from different sessions according to trial type. In order to avoid biases caused by spurious correlations in the neural activity, we

built 100 different pseudo-populations and combined the trials differently each time.

To then implement the SVM, we split the data (either from a single session or from the pseudopopulation) into training data (75%, labelled) and test data (25%, unlabelled). As before, we used the *GridSearchCV* function to find the best set of hyperparameters. We then used these parameters and instantiated a SVM with a linear Kernel and trained it on the part of the data designated training data. The model was then fit to best separate the data according to labels (in our case either trial type, or presence/absence of a specific odour) by finding the optimal hyperplane in the n-dimensional space to maximally separate the neuronal data according to the experimenter-defined categories.

The SVM was then tested by having it predict the labels of the test data, and its performance was measured as the accuracy at which the model correctly predicted the task event of interest.

2.5.3 Statistical analysis

All statistics were calculated using the Python packages *scipy* and *pingouin*. Summary data are reported as mean ±s.e.m. (standard error of the mean). Normality of data distributions was determined by visual inspection of the data points. Test statistics are detailed in the main text. Threshold for statistical significance was defined as 0.05.

No power analysis was run to determine sample size a priori. The sample sizes chosen are similar to those used in previous publications. Throughout the figures the * symbol represents $p < 0.05$.

3 An olfactory paired-associates task to probe structural learning

3.1 Introduction

Learning about the relationships between cues enables us to identify common underlying structures of events and is fundamental to adaptive behaviour (Tolman, 1932). This type of learning requires recognising the structured relationship between distinct cues in our environment, as the meaning of cues can differ dramatically dependent on their order in time, or their position in space. Importantly, learning about structured relationships between cues is not only foundational to our everyday lives, but generates the complex associations and relationships that form the basis of episodic and semantic memory (Eichenbaum, 2001). Impairments in the use of such memory are a consistent hallmark of the most debilitating neural disorders from Alzheimer's disease to schizophrenia, depression, and generalised anxiety, further emphasising the importance of structural learning (Chamberlain & Sahakian, 2006; Öngür et al., 2006; Rao et al., 2022).

Despite its importance, there is limited insight into how structural learning is achieved within the brain on a cellular and circuit level. To address this question, different types of paradigms have been put forward to study structural learning in rodent models (Aggleton et al., 2007; Albasser et al., 2013; Eichenbaum et al., 1987; Sutherland & Rudy, 1989). What all of them have in common is that they require the subject to learn not about individual cues associated with outcomes, but about sequences of cues, often containing overlapping or similar elements.

One extensively studied paradigm that fulfils these criteria is spatial navigation. When animals learn to navigate to a specific goal within an environment, they

need to process a variety of visual, tactile, and olfactory cues and commit them to memory. In order to recall the same trajectory in the future, the individual locations must be encoded in the right sequence and become associated to the goal that lies at the end. Rodents easily achieve this, and more than that, we know that they can adapt to changes in an environment, infer new routes to a known goal, therefore firmly establishing navigation as an example of structural learning (Tolman, 1948).

However, spatial navigation tasks have some important limitations with regards to understanding the neural substrates of structural learning. As they rely on animals moving through arenas (virtual or real) it is difficult to disambiguate abstract spatial information from movement signals or representations of specific cues or goals, which is important when trying to understand how the binding of elements within a context is achieved in the brain. Furthermore, in these tasks animals often visit locations more than once, making it hard to distinguish neural representations of past (memory), present, and even future (planning) (Eichenbaum & Cohen, 2014).

In research largely parallel to the study of spatial navigation, special cases of associative learning have been proposed as suitable paradigms to test structural learning. For example, “configural cues” (built from two or more elements, e.g., both a light and a tone) as well as “occasion setters” (cues that precede an ambiguous conditioned stimulus) might fulfil the criteria for structural learning (Schmajuk & Buhusi, 1997; Sutherland & Rudy, 1989). In these tasks, the animal can predict the outcome of a trial if it correctly recalls the cues separated in space or modality (for configural cues) or in time (for occasion setters).

Subsequent studies however often found that these tasks can be solved by other, much simpler strategies, such as keeping track of the statistics about the frequency or the average value of each stimulus (Rudy & O’Reilly, 1999; Rudy & Sutherland, 1995). Furthermore, these kinds of tasks often don’t allow for

manipulations to test inference or other types of generalisations, emphasising the need for controlled investigation of structural learning, and its subsequent utilisation to guide behaviour. In particular, a suitable task should use minimal, clearly defined cues, balanced across contingencies in a way that it can only be solved using structural learning.

The latter part has proven particularly difficult – many tasks that meet the first two criteria are so complex that they are too difficult for rodents to learn in a timely manner. One way to make paradigms more amenable to rodent research is the use of olfactory cues. Mice readily attend to odour cues and can detect extremely fast and subtle changes in the structure of odours (Ackels et al., 2021).

Given the sensory salience of odours to rodents, multiple studies have used odours to investigate the role of the hippocampus in relational learning. For example, a 1993 study by Bunsey & Eichenbaum demonstrated that rats can associate specific odour pairs with rewards, and that the ability to distinguish them from different pairs with the same components is dependent on hippocampal circuits (Bunsey & Eichenbaum, 1993). Other studies have shown that the hippocampal formation is needed to distinguish between different mixtures of the same two odours, but are not needed to follow an odour gradient (Eichenbaum et al., 1986; Eichenbaum et al., 1988; Otto et al.,1991) and that using interlocking chains of odour pairs, rats can perform transitive inference judgements (Bunsey & Eichenbaum, 1996; Dusek & Eichenbaum, 1997; Dusek & Eichenbaum, 1998). These studies were pivotal in placing the hippocampus identifying the hippocampus as a critical region for learning relationships between neutral stimuli. However, due to methodological limitations, they could not provide strong hypotheses of how this learning was implemented within hippocampal circuits.

Additionally, since the aim of these studies was not the comparison of the specific neural responses elicited by individual odour stimuli or combinations

thereof, the tasks were not designed in a way that ensured interpretability of single cell signals. For example, they might not ensure that each cue appeared equally often (to control for familiarity effects) or that cues were associated with the same average likelihood of reward (to control for development of preference). In this study, we aimed to design a task that builds on the principles established by these earlier studies, specifically by using odour cues and combining them across a delay to form reward-predictive contingencies.

In this chapter, I will describe the olfactory paired-associates task that we designed to explicitly probe the acquisition and use of structural learning in mice. Crucially, this task can be learnt quickly in head-fixed mice, and once learnt, allows probing the flexible use of the task structure.

Due to the volatility of odours and the ability of rodents to distinguish even odour mixtures with very similar proportions, behavioural tasks with odour cues often are at risk of odour contamination which can lead to inadvertent presentation of "combinatorial odour cues" (Uchida & Mainen, 2003). Therefore, I will further describe the steps we have taken to control for these possibilities.

3.2 Results

To assess structural learning, we designed a novel olfactory paired-associates task. In this task, water-restricted mice were head-fixed on a setup where they were presented with a sequence of two odours separated in time. On the setup, the mouse had access to a waterspout that delivered rewards based on task events. From the right, a stream of air was introduced at a constant rate. The air was removed from the left side by a tube connected to a vacuum pump, thus creating a continuous flow of air across the mouse's nose (**Figure 3.1a**). Into this continued airflow, different odours were injected through a high-accuracy odour delivery manifold, resulting in temporally defined bouts of odourised air that were used as cues (**Figure 3.1b**).

In the task, three possible odours (A, B and C) were presented in pairs, thus leading to 6 possible sequences, three of which were rewarded ('go' trials: AB, BC and CA) and three unrewarded ('no-go' trials: AC, CB and BA, **Figure 3.1c**). The odours were separated from each other by a delay of at least 5s, and the contingencies of reward were counterbalanced such that each cue alone provided no information about upcoming reward, as each odour was equally frequent in 'go' and 'no-go' trials.

For example, in this task design, the sequence CA was rewarded while BA was not rewarded – the meaning of A therefore was ambiguous, unless combined with the previous odour. Furthermore, the task was also designed such that mice could not use a perceptual configuration such as 'AB' to solve the task: B after A was rewarded, while A after B was not rewarded. Thus, mice could neither use an elemental (one-cue based) nor a configural (cue-combination based) strategy and had to retain both the identity of the odour cues as well as their positions within the sequence to behave optimally (**Figure 3.1d**).

On each trial, a randomly selected odour pair was presented to the mouse, and licking behaviour was recorded. Water reward was given 1.5s after the second

odour ('response window'), but the water delivery was not dependent on the licking behaviour of the mouse. Therefore, every 'go' trial was followed by reward and learning of the task was measured as the emergence of anticipatory licking in the response window on 'go'-trials, and a lack of anticipatory licking in 'no-go' trials (**Figure 3.1e**).

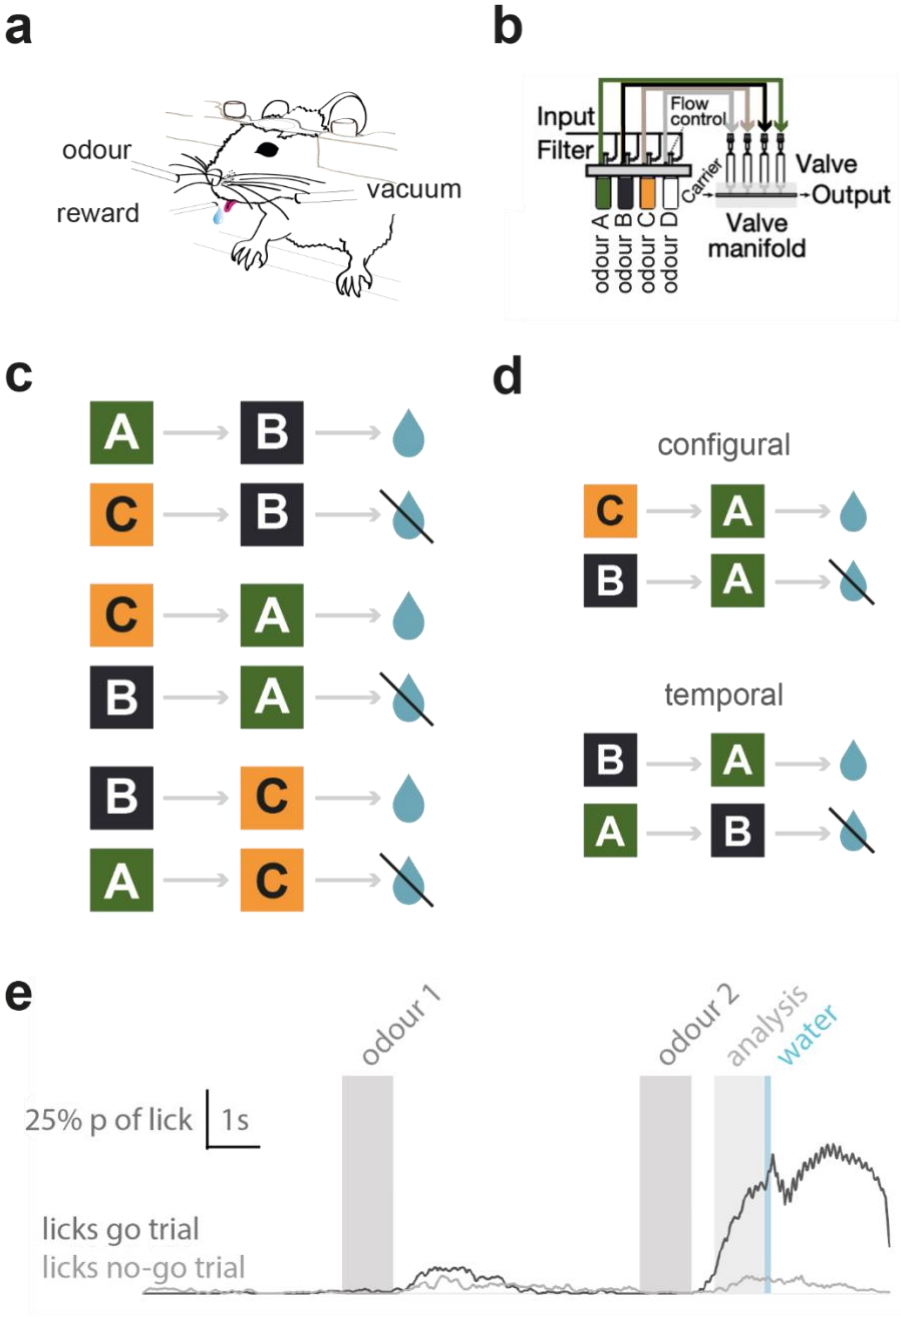

Figure 3.1: Schematic of behavioural setup and task structure. a. A head-fixation system allows for delivery and removal of odour cues as well as reward delivery in a consistent position relative to the animal. **b**. A high accuracy olfactory system for odour delivery injects odourised air into a constant stream of clean air and delivers that to the animal. **c**. Design of the go/no-go task. Each pair of odour cues is separated by a delay and is either rewarded or not rewarded. **d**. The design is counterbalanced so that only retaining both the identity of an odour and its temporal position within a sequence can predict reward. **e**. Trial structure overlaid with example licking behaviour from a well-trained mouse: predictive licking in the analysis window between the delivery of the second odour and the reward in go-trials indicates anticipation of the correct outcome.

Mice learned to perform this task stably within the space of 8 days of training (**Figure 3.2a;** n = 13; repeated measures ANOVA: $F_{(9,45)}$ = 11.53, p = $3.9x10^{-9}$). While all mice reached >70% accuracy within the 10-day training period, some did learn the task markedly quicker than others (**Figure 3.2b**): some mice reached the learning threshold within only a few days while others took over a week.

Mice would undergo several stages of training (outlined in Chapter 2) in order to encourage predictive licking in the response window before reward delivery. When encountering the full task with all 6 contingencies, mice initially behaved suboptimally, often exhibiting marked anticipatory licking before 'no go' trials, but also marked licking after the first odour of each pair (**Figure 3.2c**, top). Over training, mice gradually refrained from licking during the delay period between odours, and well-trained mice would typically initiate licking only after the second odour (**Figure 3.2c**, bottom).

We quantified this learning two ways: first, as a marked increase in the proportion of correct trials from the first to the last session (**Figure 3.2d**; Day 1: p correct = 0.52 +/- 0.2; Day 10: p correct = 0.78 +/- 0.18; n=13, paired t-test $t_{(12)}$ = 10.83, p = $7.0x10^{-8}$). Secondly, we used a metric from signal detection theory called behavioural d-prime (d') which quantifies the ability to predict the identity of a trial as 'go' or 'no go' based solely on licking behaviour. Consistent with the higher proportion of correct trials, d' increased markedly between first and last sessions (**Figure 3.2e;** day 1: d' = 0.03 +/- 0.26; day 10: d' = 1.17 +/- 0.46; n = 13, paired t-test $t_{(12)}$ = 8.34, p = $2.0x10^{-6}$).

In summary, we designed an olfactory structural learning task that mice can rapidly learn.

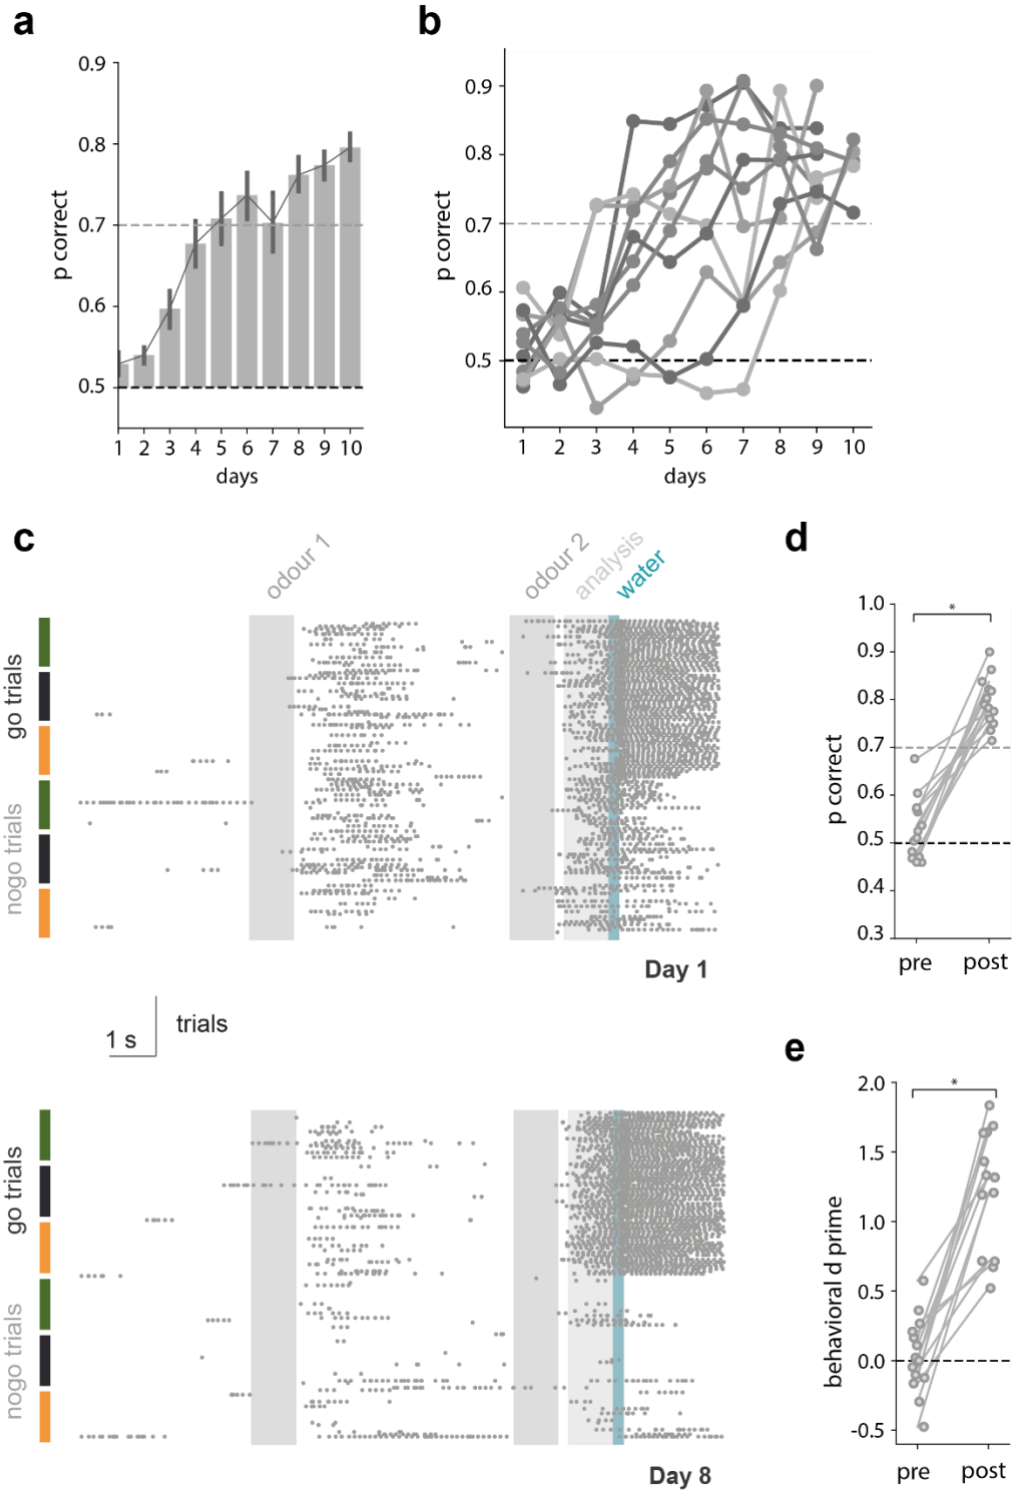

Figure 3.2: Mice can perform structural learning task. a. Mice reach accurate performance over a short timescale, responding correctly to more than 70% of trials in 10 days (n = 13) **b**. Performance over training days for individual mice. **c**. Example licking behaviour (dark grey dots) in a well-trained mouse before and after learning (top and bottom panel, respectively). The coloured bars correspond to trial types with a shared first odour (green: A, black: B, orange: C) **d**. Performance before and after learning (day 1 and day 10, respectively; n = 13) **e**. Discriminability index d' (day 1 and day 10 respectively, n = 13)

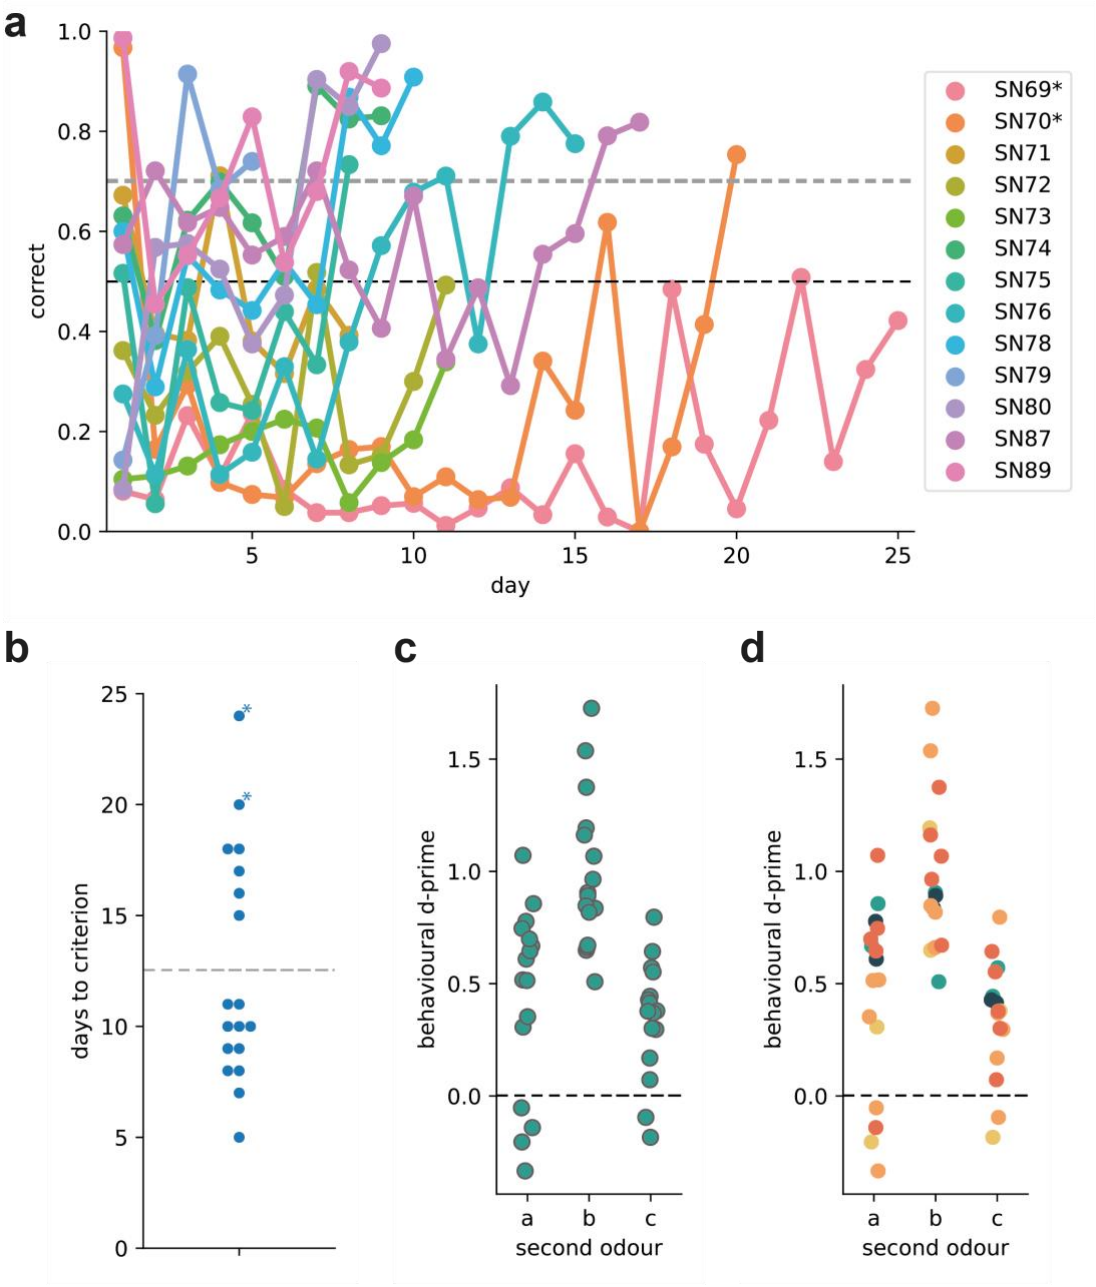

Figure 3.3: Variance across mice within task and shaping stages. a. Performance of mice across shaping stage in which only rewarded pairs are presented (n = 13). Mice marked with an asterisk were partially shaped on a protocol with conditional reward delivery. **b**. Mice took on average 13 days to reach the threshold of predictive licking on >70% of trials (n = 17). In both a and b, Mice marked with an asterisk were subject to multiple changes in the training protocol due to hardware problems. Performance data is unavailable for some mice and they are thus only included in panel b. **c.** Discriminability index d' split by identity of the second odour (e.g. a = b->a and c->a trials; n = 17). **d**. Discriminability index d' split by identity of the second odour and coloured according to training cohort (n = 17).

However, there is variability between mice both in the speed of task acquisition as well as in the level of performance (**Fig. 3.3**). For instance, the individual learning curves during the shaping stage, i.e. the stage in which only rewarded trials are presented, show large variance in the amount of predictive licking on the first day of shaping. Notably, these mice have previously not experienced cued rewards. In the first session, mice clearly employ different strategies: while some mice lick predictively 100% of the time and thus likely lick quite continuously, other mice only achieve <10% correct trials, therefore likely employing a strategy of responsive licking (**Fig. 3.3a**).

In consequence, some mice need many sessions to reach the performance criterion of 70% correct trials, while others graduate to the full task within as little as 5 sessions (**Fig. 3.3b**). Some mice (indicated with an asterisk) were partially trained on a protocol in which reward was only delivered conditionally, i.e. only when the mice licked predictively. Since this change did not improve their learning, they were returned to a Pavlovian reward delivery protocol after a week.

Analysing the performance of mice on trials split by second odour shows that there is variability not only in the time course of learning, but also in what is learnt. Specifically, splitting the discriminability index d' by the second odour reveals that on average, mice tend to perform better on trials where b is the second odour (**Fig. 3.3c**). Since over the course of the project, multiple changes were made to the setup, we next compared whether this bias was isolated to a single cohort of mice (**Fig. 3.3d**). We found that this was not the case: individual variability exists within several cohorts. However, in some cohorts (blue & turquoise) all mice demonstrate balanced discriminability, thus suggesting that unbiased performance in the task is possible.

Accurate delivery of odour cues can be challenging, often impeding the interpretability of behaviour in olfactory tasks. We aimed to mitigate this using highly controllable olfactory delivery system (Schaefer lab (Ackels et al., 2021)).

However, we still wanted to ensure that mice were indeed using the odour cues in the way we intended.

In a first instance, we checked whether circumstantial cues, such as the alterations in air pressure resulting from valve openings and closings, or the audible cues (clicks) associated with valve openings may be contributing to behaviour. We therefore had expert mice undergo sessions of the task where we completely removed the air flow, leaving only auditory valve clicks as cues. In these sessions, mouse behaviour dropped significantly when compared to sessions with normal air flow, indicating that the valve clicks alone are not sufficient to allow the mice to anticipate reward (**Figure 3.3a**; n = 5; repeated measures ANOVA: $F_{(2,8)} = 33.01$, $p = 1.3 \times 10^{-4}$; post-hoc Tukey [Baseline vs. No Airflow] $p = 2.8 \times 10^{-4}$, post-hoc Tukey [No Airflow vs. Recovery] $p = 2.8 \times 10^{-4}$).

We then removed the odour vials, replacing them with empty containers while keeping a constant flow of clean air. This manipulation left both auditory clicks as well as changes in air pressure as available cues. As in the previous manipulation, this change made the behavioural performance drop to chance level (**Figure 3.3b**; n = 4; repeated measures ANOVA: $F_{(2,6)} = 9.77$, $p = 0.01$, post-hoc Tukey [Baseline vs. No Odour] $p = 4 \times 10^{-3}$, post-hoc Tukey [No Odour vs. Recovery] $p = 0.02$), showing that mice required odour identity to behave optimally.

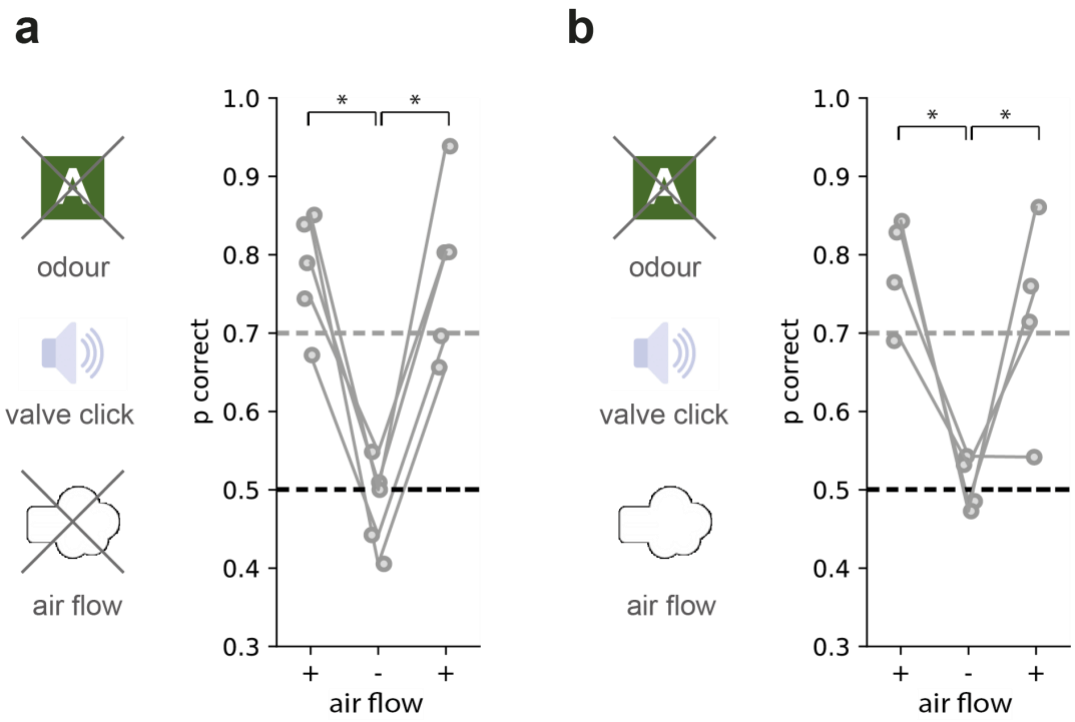

Figure 3.4: Mice use sequence of odours to solve the task. a. Mice cannot perform the task with auditory cues (valve clicks) only (n = 5) **b**. Mice cannot perform the task with only valve clicks and changes in air flow as cues. (n = 4)

We next wanted to evaluate the possibility of mice merging each odour pair into a single percept ('configural cue'), using small concentrations of residual odour that might still be present at the end of the delay period. For example, the sequence of odour A followed by odour B might be encoded as "low/residual [A] + high [B]", while the reverse sequence could be encoded as "low/residual [B] + high [A]". As a first step, we examined the dynamics of each odour pulse using photoionization detection (PID) recordings. These recordings showed that our delivery method reliably produced a rapid increase in odour levels followed by a slower decay (**Figure 3.4b**, orange curve). This temporal pattern indicated that both the odour plume and the associated changes in air flow returned to baseline after approximately 1500ms, well before the end of the delay period and the delivery of the second odour cue.

To rule out the possibility that odour levels below the detection threshold of our measurement equipment influenced mouse behaviour, we conducted a control experiment. In this experiment, we interleaved a new trial type with the standard trials in the task: in these trials, the first odour cue was replaced with a diluted odour cue (1:10000) which was delivered at the same time as the first odour in the normal trials but persisted throughout the entire delay period (schematic in **Figure 3.4a**). We selected this dilution level because it consistently yielded PID readings above baseline for the entire duration of odour delivery and the delay period (**Figure 3.4b**, grey curve and inset). Thus, in these trials, mice lacked a usual first odour cue but instead were presented with a defined, PID-detectable, yet low-concentration residual odour throughout the entire delay period at a level that was consistently higher than our baseline recordings. If mice were indeed using residual odour traces to build a 'configural cue' as previously described, they should exhibit high accuracy performance in trials with the diluted odour cue. Contrary to this hypothesis, their performance significantly declined when encountering these trials (**Figure 3.4c and d**; n = 3; paired t-test: $t_{(2)} = 5.58$, p = 0.03), indicating that mice did not rely on residual odour and require a full-concentration first odour cue to perform well on the task.

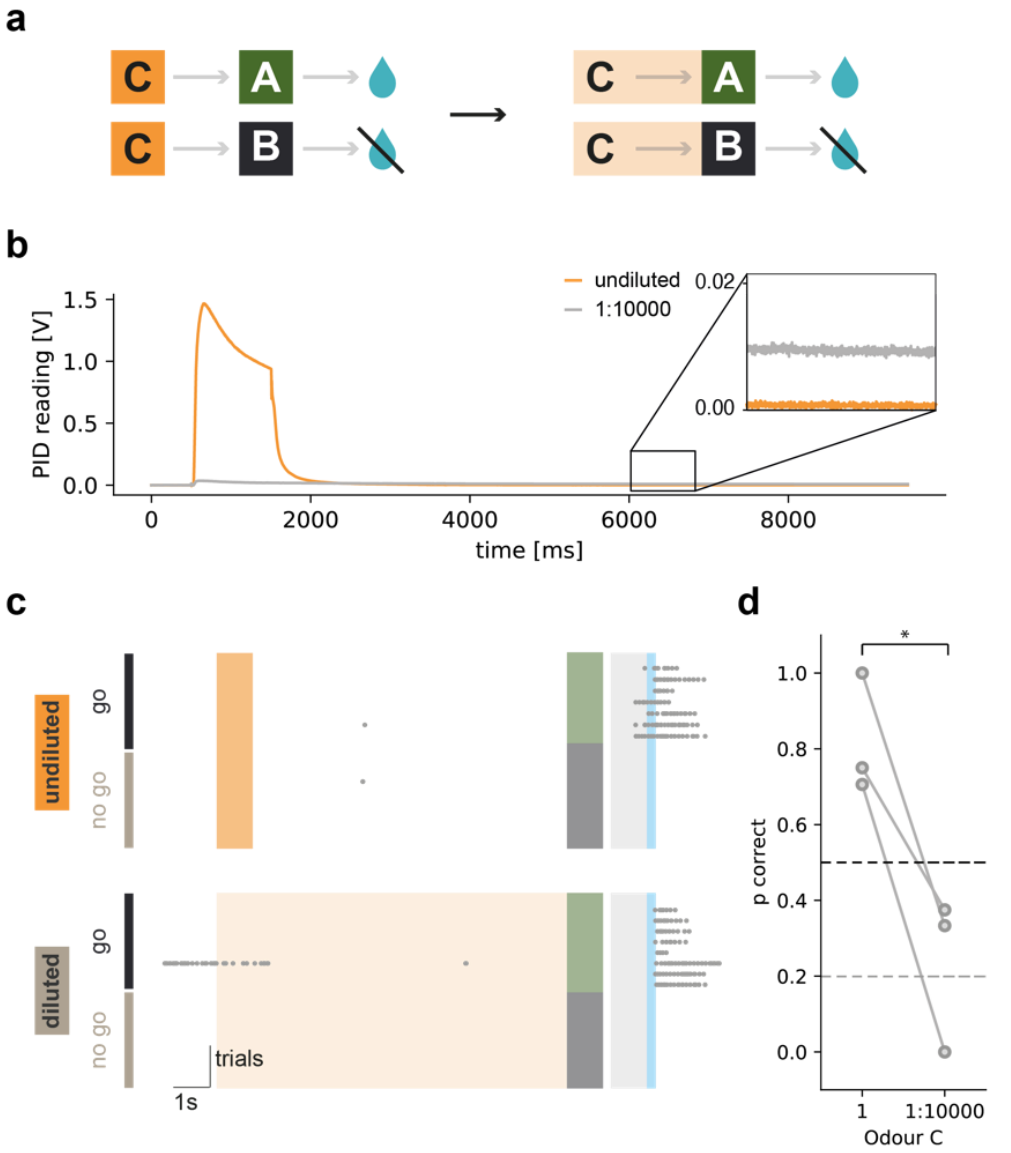

Figure 3.5: Mice do not use residual odour traces for a configural representation. **a**. Schematic of trials to control for the mice using residual odour traces for a configural representation (e.g. a representation of "95% A + 5% C"). In those trials (right) a defined, detectable, but low concentration residual odour was presented during the entire delay period. **b**. Example odour traces of undiluted odour and diluted odour titrated to produce a continuous Photoionization device (PID) reading across the entire delay period (dilution 1:10000). The continuous reading is higher than baseline at the onset of the second cue (see inset) such that if mice were using a configural strategy, they could solve the task. **c**. Licking behaviour during a session with randomly interleaved normal trials (undiluted odour) and trials involving a diluted first odour cue (model of residual odour). On the top, the orange bar represents undiluted presentation of odour C. The light orange shading below is the diluted odour C presented across the delay. The green and black bar represents the possible second odours (green: A, rewarded; black: B, unrewarded) **d**. Mice cannot perform the task using residual odour traces (n = 3)

The above controls suggest that the mice were indeed perceiving the odours as separate temporally concise cues. However, as mentioned in the introduction, it has been shown that mice can keep track of various statistics and average the value of cues over many trials. We therefore wanted to ensure that our subjects were indeed using both the order and the identity of the cues to solve the task, and not e.g. a complex combination of elemental bias towards certain odours together with a rolling average of past rewards.

To this end, we used logistic regression models to estimate the probability of an event occurring – in this case, the mouse licking – based on a set of independent predictor variables (task events, **Figure 3.5a**). By providing the models with access to different sets of task events, we built four models representing distinct strategies.

Our **structural model (S)** incorporates both the identity and temporal order of odour cues within each trial, with six parameters corresponding to the six possible trial types (AB, AC, BA, BC, CA, CB). The **elemental model (E)**, by contrast, only considers the presence or absence of each odour, disregarding their order. For instance, an AB trial and a CA trial would both be recorded as "A present," while a BC trial would be noted as "A absent." The **configural model (C)** predicts responses based on odour combinations but does not differentiate temporal order; AB and BA are treated equivalently as "AB both present." Finally, the **history model (H)** incorporates the choice and outcome of the previous trial as predictor variables (**Figure 3.5b**).

To evaluate how well each logistic regression model predicts mouse behaviour, we then calculated the Area Under the Curve (AUC) using a Receiver Operating Characteristic (ROC) curve. Each model generated a probability between 0 and 1 for whether a mouse will lick on a given trial, based on the predictor variables described above.

To assess model performance, we then wanted to determine how well these probabilities distinguish between trials where licking does and does not occur.

A ROC curve is constructed by varying the decision threshold for classifying a trial as a "lick." At each threshold, the true positive rate (TPR, the proportion of correctly predicted licking trials) is computed and compared to the false positive rate (FPR, the proportion of non-licking trials incorrectly classified as licking). Plotting TPR against FPR across all thresholds produces the ROC curve, and the AUC is calculated as the area beneath this curve.

AUC values range from 0.5, indicating chance-level performance, to 1.0, representing a perfect model. Higher AUC scores indicate better predictive performance, as the model more accurately distinguishes between licking and non-licking trials. By comparing AUC scores across models, we can infer which strategy best captures mouse behaviour.

For instance, if the history model (H) achieved a high AUC, this would suggest that mice rely on past trial outcomes when deciding to lick. Conversely, a high AUC for the structural model (S) would indicate that mice use both odour identity and temporal order to guide their behaviour. If a model yields a lower AUC, it suggests that the factors it considers are less relevant in predicting licking responses.

This approach allowed us to quantitatively compare different models and gain insights into the strategies mice use to learn and make decisions in the task. We found that the structural model best predicted the behaviour of the mice (**Figure 3.5c,** n = 9; repeated measures ANOVA: $F_{(3,27)} = 15.12$, $p = 6.0 \times 10^{-6}$; post-hoc Tukey [S vs C] $p = 2.3 \times 10^{-4}$, post-hoc Tukey [S vs E] $p = 2.4 \times 10^{-5}$, post-hoc Tukey [S vs H] $p = 1.1 \times 10^{-4}$).

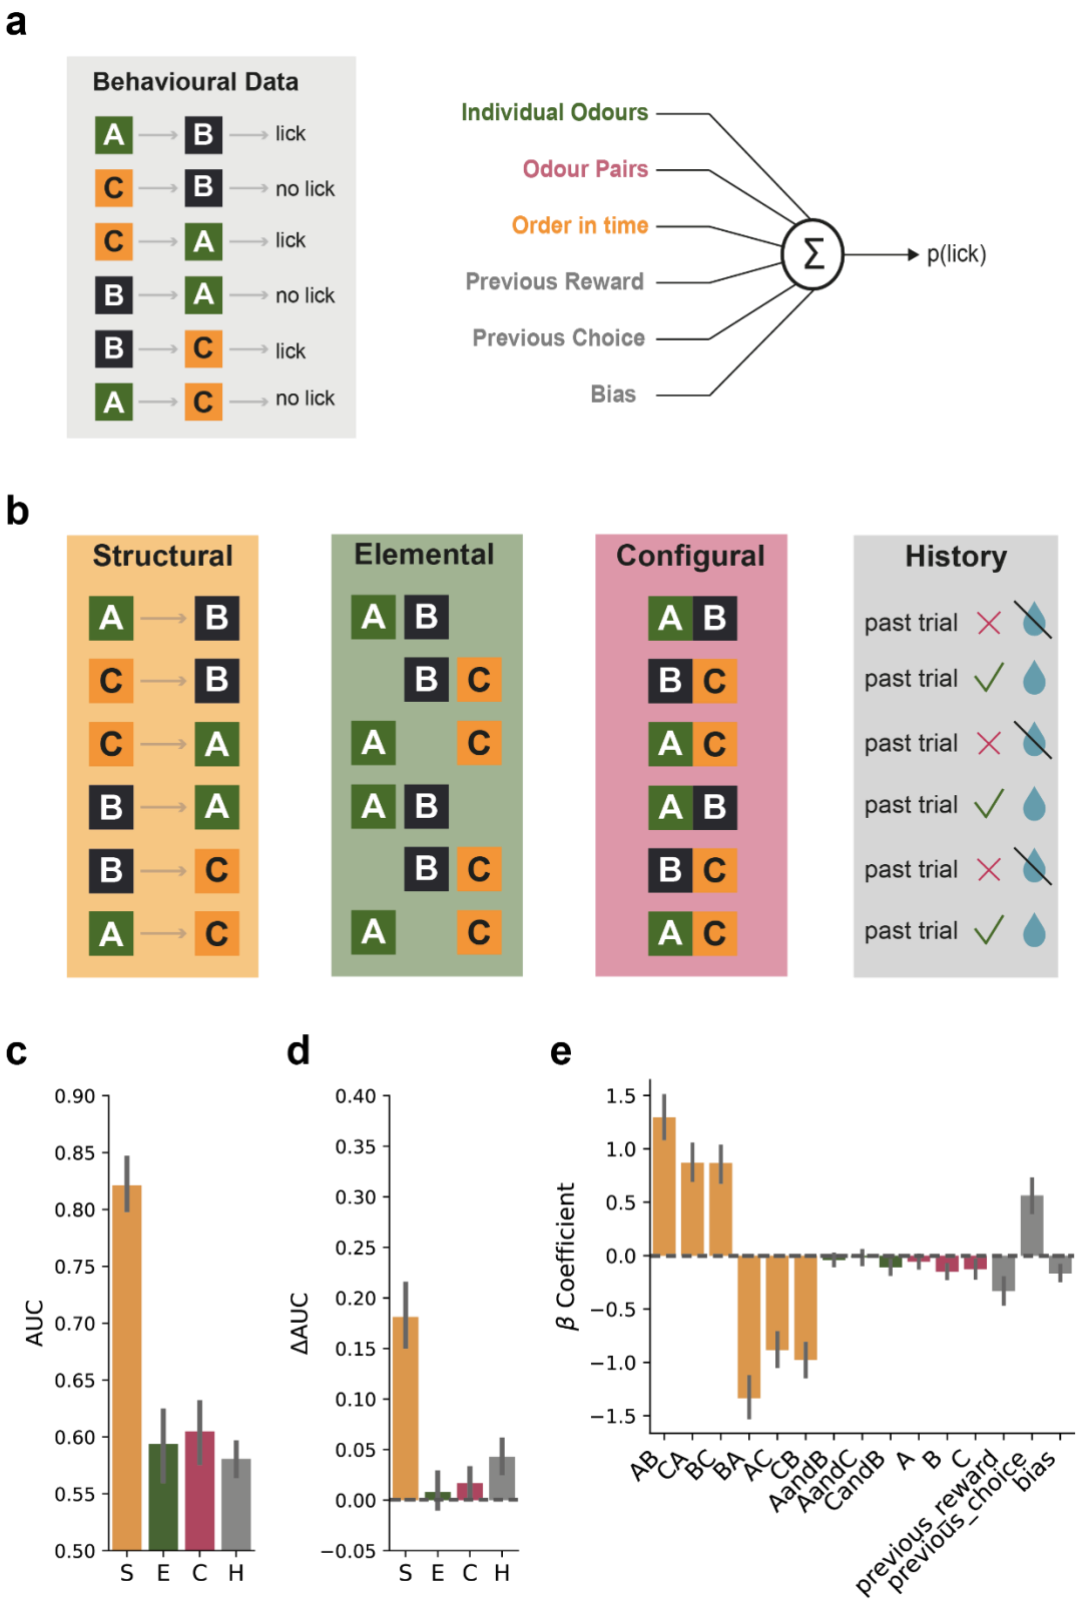

Figure 3.6: Regression analysis shows that mice use the order of cues to solve the task. **a.** Logistic regression analysis uses different task events such as cues and previous outcomes to predict the probability of the mouse licking. **b.** Giving the regression analysis access to different information, we built four different models: structural (has access to both cues and their order in time), elemental (only has access to individual cues, no order), configural (cue

combination, but not order) and finally a model with access only to previous choice and outcomes. **c.** Area under the curve (AUC) for each of the models (1: perfect accuracy, 0.5: chance). The structural model best predicts licking behaviour (n = 9) **d.** ΔAUC is the change in model fit if all models but one are used to predict mouse behaviour. The biggest change in accuracy is observed in the regression that uses all information except the order of cues (n = 9) **e.** Beta coefficients for the predictor variable of each model. Positive coefficients are positively correlated with licking behaviour, while negative coefficients suggest an inverse relation.

To further investigate whether a combination of our models could outperform the structural model, we next calculated delta AUC (ΔAUC). This measure quantifies the contribution of specific predictor variables by comparing model performance with and without them. Specifically, in a first step, predictions are generated using a full model, incorporating all available predictor variables across the different strategies. We then systematically removed individual predictors or groups of predictors – such as trial history or odour identity – and recalculated the AUC.

The difference between the AUC of the full model and the AUC of the reduced model provides the ΔAUC and is a measure of how much predictive power is lost when a given predictor is excluded. By examining the ΔAUC, we could assess whether certain aspects of task structure, odour identity, or trial history contribute significantly to the model's ability to predict licking behaviour. A large drop in AUC when removing a predictor suggests that the mice rely heavily on that information, whereas a small or negligible ΔAUC indicates that the predictor may not be critical for decision-making.

Using this technique, we again find that the structural model (predicting based off temporally ordered odour pairs) has the biggest contribution to the overall accuracy of classification (**Figure 3.5d,** n = 9; repeated measures ANOVA: $F_{(3,27)} = 14.47$, $p = 8.0 \times 10^{-6}$; post-hoc Tukey [S vs C] $p = 3.2 \times 10^{-5}$, post-hoc Tukey [S vs E] $p = 1.0 \times 10^{-4}$, post-hoc Tukey [S vs H] $p = 5.1 \times 10^{-4}$), indicating that the mouse behaviour is indeed best predicted by structural information and not a combination of elemental, configural or historic information.

This was further corroborated when looking at the beta coefficients for each predictor variable in the combined model. Beta coefficients describe the relation between each variable and the predicted outcome. Positive values denote a positive correlation, while negative coefficients suggest an inverse relation. Here as well, the variables that are most strongly influencing the model prediction are the ones containing structural information (**Figure 3.5e,** n = 9; repeated measures ANOVA: $F_{(3,27)} = 19.39$, $p = 4.7 \times 10^{-25}$; planned comparisons p-values: see Appendix 1).

Taken together, this regression analysis suggests that mouse behaviour is best explained by odour pairs in order, suggesting that mice used a structural route to solve our task.

3.3 Discussion

In this chapter, I have presented a novel task paradigm that requires mice to remember an odour cue over a delay, and then either lick or withhold licking depending on the identity of a second cue (**Figure 3.1**). Importantly, contingencies in this task are balanced in a way that the values of all individual odours are ambiguous (i.e. as often rewarded as not), and mice can only solve the task by combining the two cues across the delay.

Our behavioural data shows that mice can learn this task over the space of less than 10 days, developing a robust anticipatory lick response to the rewarded pairs of odours, and withholding their licks for unrewarded pairs (**Figure 3.2**).

To ensure that the mice are indeed using a structural learning strategy, we conducted both experimental controls as well as a regression analysis.

In control experiments with only circumstantial clues (valve clicks or changes in air pressure, **Figure 3.3**) mice did not exhibit accurate anticipatory lick responses, indicating that the odour cues are necessary for performance. In a second experiment, we investigated the possibility of mice circumventing the need for structural learning by using a combination cue made up of residual odour from the first odour at the time of the second odour delivery (**Figure 3.4**). Here as well, mice were not able to perform the task, consistent with a behavioural requirement for strong, temporally distinct odour cues.

To computationally examine strategies that mice might use during the task, we fit different logistic regression models to the behavioural data. We found that in all cases, a model using structural information predicted animal behaviour better than both individual models using elemental or configural cues or past outcomes, or even a model using a combination of all of these variables (**Figure 3.5**).

Overall, we show that our novel task requires mice to use structural information and combine odours across the delay, and that mice successfully learn about the reward contingencies associated to specific temporal combinations of odours.

3.3.1 Comparison to similar tasks from the literature

As previously mentioned, the parallel scientific discoveries of hippocampus as the centre of autobiographical memory in humans and as an essential component of rodent navigation has sparked many theories attempting to unify the two functions into one global theory of hippocampus, and their investigation has produced many tasks that show parallels to our paradigm.

A cued T-maze for example also requires subjects to retain a cue over a period of time (while traversing the stem of the maze) in order to behave correctly at the choice point (turn left or right). However, in the usual setup, there is no further information needed outside of the initial cue and therefore animals could in theory prepare their action even before the choice point, making it difficult to interpret neural signals as they could both represent the past cue identity as well as the future associated action.

The delayed-match-to-sample (DMS) task, widely used in the study of memory, circumvents this issue by only presenting some of the necessary information at the start of the trial: in the initial phase, the monkey is shown an object, and has to retain its identity over a delay phase. After the delay, the monkey is shown two objects, and is asked to select the familiar object. As the position of the objects (right or left) is necessary to plan the correct action, in this task it is easier to tell apart neural signatures of memory (recall of target object) from the planning of a future action. Yet, this task has other limitations as it could conceivably be solved by a recency effect instead of a recall of the specific object. To exclude this strategy, an alternative approach was to change the target the subjects were

meant to select from the sample object to the novel object. In this version, called delayed-non-match-to-sample (DNMS) task, a recency-based strategy is not as easy to implement, but still does not consistently require the subject to disambiguate a stimulus that is equally likely to lead to reward as not, as for example the non-sample object might be completely novel or have an average associative value that is higher or lower than that of the sampled object (Kangas et al., 2011).

The tasks developed by Rudy and Sutherland to investigate their theory of configural learning are carefully designed to avoid potential biases caused by differences in associative value of specific clues (Rudy & Sutherland, 1995; Sutherland & Rudy, 1989). For example, in the "positive patterning" paradigm, reward is given when two stimuli are presented together (AB+) but not when they occur individually (A-, B-). Thus, the values of both cues are equivalent, as is the frequency of their presentation. Other versions of these configural tasks, such as negative patterning (A+, B+, AB-) and biconditional discrimination (AX+, AY-, BX-, BY+) also contain similar sets of counterbalanced cues. However, this only holds true if the configural (combined) stimuli are indeed encoded as a combination of the elemental cues, and not as a separate representation of e.g. 'A&B'. Interestingly, later studies using the same type of paradigms found the necessity of hippocampus to be a function of the spatial distance between the individual elements to be included in the configural cue (Albasser et al., 2013), or required the use of an entire context as one "cue" (Riaz et al., 2017), suggesting that the separation of cues in space or time rather than the combinatorial nature of cues is the factor that makes a given task hippocampus-dependent.

Our task design attempts to circumvent all of the limitations mentioned above: mice can only make the decision on how to act at the delivery of the second odour, since both cues are necessary to correctly anticipate rewards. All individual cues are carefully counterbalanced so that neither familiarity nor

recency effects can be used to solve the task. Finally, a delay of at least 5s ensures that the individual odour cues are too far separated to be encoded as one configural odour mixture, an attribute we specifically controlled for in the experiments shown in **Figure 3.4**.

By designing the task in a way that A before B holds a different meaning to B before A, we also introduce directionality into the cue space, thereby meeting the criteria of "structural learning" as defined by Aggleton (Aggleton et al., 2007). They argue that hippocampus is required to bind several features into a "cue array", but only if there is a specific relationship between the elements (X to the left of Y, B before A, etc.).

In recent years, interest in the type of learning required to understand structure and relations between different objects has gained much attention, leading to the development of a wealth of further tasks, from spatial (Frank et al., 2000; Wood et al., 2000) to non-spatial (MacDonald et al., 2013; Pastalkova et al., 2008). Most of these tasks however exhibit similar drawbacks as described above and thus they can offer only limited insight into the neural representations that underpin the role of hippocampus in these behaviours. With the design of our task, we are well placed to provide new insight into what information is represented in hippocampal populations during structural learning, offering a new piece in the puzzle of the function of hippocampus as a whole.

3.3.2 Technical considerations of the task design

We chose to use odour cues in this task since previous studies have shown olfactory cues to be very salient to mice (Liu et al., 2014; Taxidis et al., 2020; Yun et al., 2023).

However, while providing many advantages for fast learning of complex tasks, olfactory cues present with many unique challenges when compared to visual or

auditory cues. The spread of odours is difficult to predict and equally difficult to measure, both in qualitative as well as quantitative ways, especially when it comes to comparing odour levels across different components. Furthermore, research has shown that mice are extraordinarily good at noticing even small differences in odour composition (Uchida & Mainen, 2003).

To limit the influence of these issues, we have taken multiple measures in both the design of our setup as well as the design of the task. First, we built a high-accuracy odour system adapted from the Schaefer lab (Ackels et al., 2021). Crucially, in this system, clean air is constantly introduced into the setup and used air is sucked away from the setup via a vacuum pump, thus creating a constant flow of air in which odorants are then injected, forming temporally distinct odour plumes with a sharp rise and fall (see **Figure 2.3**). We regularly controlled that the amplitude and dynamics of each odour remained the same across learning by measuring the concentration of each odour with a Photoionization Detector (PID). Since the PID readings are not comparable between odorants due to the different ionisation energies required for different molecules, we further used an airflow sensor to ensure that the changes in airflow due to the injection of odourised air into the airflow were similar in shape and amplitude. Lastly, we conducted multiple control experiments to ensure that mice were indeed using the odour cues as intended (**Figure 3.4 and 3.5**).

Another point of consideration is the use of anticipatory licking as a readout for learning. Since the rewards are delivered in a Pavlovian manner, i.e. independent of the mouse's behaviour, anticipatory licking was not necessary for maximising rewards. Therefore, it is possible that we are missing some of the learning since mice might correctly predict the rewarded outcome, but not display any anticipatory licking.

Another important aspect of our task is that licking outside the reward availability window is not punished. As a result, mice tend to lick both during the inter-trial

interval and in the delay between the two odour cues. Focusing on the latter, the bottom panel in **Figure 3.2c** suggests that anecdotally, the example mouse seems to lick more frequently in response to the first cue when it appears in a rewarded pairing. At first glance, this behaviour seems puzzling since our task design ensures that any given first odour is equally likely to be followed by a reward or no reward.

However, because trials are presented pseudo-randomly, effectively simulating draws without replacement, the probability of encountering a rewarded trial increases following repeated instances of unrewarded trials. At the extremes of this distribution, a mouse could infer that after six consecutive unrewarded trials, the next trial must be rewarded. Even in shorter sequences of trials with the same outcome, this pattern might still influence licking behaviour.

In our logistic regression model, we currently account for trial history only in terms of the outcome of the preceding trial. However, it may be worth exploring whether mice consider trials further back in their history to assess the likely outcome of a trial.

In our current analysis, we use anticipatory licking as a binary metric, where a single lick between the second odour and the reward is counted as a hit. However, it may be informative to consider alternative measures, such as lick rate. A single lick could occur by chance and may not necessarily indicate learning, whereas repeated licking that ramps up towards reward delivery could reflect a stronger association with the stimuli. Additionally, lick rate might serve as a proxy for the animal's certainty, with higher rates indicating greater confidence in the expected outcome. Investigating how this measure varies across task difficulty could provide insights into how mice adapt their behaviour in more challenging paradigms, potentially revealing more nuanced aspects of learning and decision-making.

Lastly, mice in our task are placed on a wheel in order to reduce the stress of head-fixation. In our current setup, we are not quantifying their running nor analysing movement by other means (such as a camera). We can therefore not exclude effects of running and other movement on learning and task performance.

3.2.3 Considerations on the regression analysis

While the data from our regression models helped to convince us that mice are solving the task using both the odour identity and the temporal structure of their presentation, it is by no means a comprehensive strategy to capture the richness of mouse behaviour. To capture a fuller picture of how mice navigate the task structure, several improvements could be made.

Right now, we make individual models for each mouse, calculate the metrics for model performance and then average those metrics over all mice. An alternative approach could be to train a single model on the entire dataset while including mouse ID as a parameter. This might enable us to capture both shared behavioural patterns and individual differences in a more structure way. It would furthermore provide a metric to estimate the individual-level deviations from the group-level behaviour, pointing us more clearly to mice who are outliers in their use of task parameters and enabling us to test specifically whether some mice are more affected by e.g. trial history than others.

Another way to capture more variance in mouse behaviour would be to change the parameters each of our models use. For example, we could refine the way odour identity and order are represented: while the structural model currently considers trial types as distinct categories (AB, AC, BA, etc.), it could be enhanced by incorporating parameters that separately encode "A as first odour" or "B as second odour." This would allow the model to generalise across different

trial types and capture potential biases in how mice respond to specific odours depending on their position within the sequence.

Another improvement could involve incorporating parameters related to the timing and magnitude of licking responses. Currently, the models treat licking as a binary outcome, but introducing a lick rate parameter could provide a more nuanced measure of learning and certainty.

Lastly, as mentioned above, the history model, which currently only accounts for the outcome and choice on the previous trial, could be expanded to include a weighted influence of multiple past trials. This would allow us to test whether mice integrate trial history over longer sequences rather than responding solely to the most recent outcome. This would furthermore allow us to add an interaction term between odour identity and trial history, which could reveal whether mice adjust their responses differently depending on past experiences with specific odours. By refining these parameters, the models could more accurately reflect the complexity of mouse decision-making and learning processes.

Finally, while logistic regression modes are a useful tool for modelling the influence of many factors on binary outcomes like licking behaviour, they are not the only method that can be used. Alternative approaches such as linear mixed models (LMMs) account for hierarchical structure in the data, such as repeated measurements with individual mice. Furthermore, by including random effects, these models allow for individual variation in the tendency to lick, such as is described in **Figure 3.3a**, thus providing a more flexible framework for capturing individual differences.

Another option would be to use Bayesian hierarchical modelling, which offers a probabilistic approach that takes into account prior information about expected behavioural patterns. However, these approaches, while more powerful,

sacrifice interpretability by adding non-linear interactions and more parameters into the mix. Ultimately, the choice of modelling strategy must strike a balance between flexibility, interpretability and the ability to account for nuance and variability within the behaviour.

4 Flexible adaptation to changing task structure reveals a role for hippocampus

4.1 Introduction

In the previous chapter, I presented a novel task that required mice to learn both about the identity of two odour cues as well as their position in time relative to each other. We found that mice could rapidly learn this and showed data from several control experiments as well as logistic regressions to demonstrate that mice were indeed using a structural learning strategy to solve the task.

A hallmark of structural learning is that, instead of learning about items and cues individually, relations between cues are retained and organised into schemas, contexts, or cognitive maps (Behrens et al., 2018). In this, this type of learning differs conceptually from the associative learning. While the latter can be modelled by classic reinforcement learning (RL) algorithms (Bari et al., 2019; Sutton & Barto, 1998) and is associated with dopamine signals in the ventral tegmental area (Howe et al., 2013; Schultz et al., 1997), this is not the case for structural learning.

One hypothesised key advantage of structural learning over reinforcement learning is that, once learnt, relationships between cues can be expressed in a flexible manner and should allow for quick adaptation to new experiences that share the same underlying structure or follow the same rules. A classic and still poignant example of this type of inference was described by Tolman in 1948: rats that experience a maze many times learn not only about the trajectories they take, but also about trajectories they might take in the future. This becomes apparent when the layout of the maze changes and the rats take shortcuts that

weren't previously available to them, using what they have learnt about the spatial layout to connect two places within the maze in a new way (Tolman, 1948).

When 30 years after, place cells forming a map of the spatial environment were discovered in the hippocampus of rats (O'Keefe & Nadel, 1979), this offered a neural substrate for the latent learning in the absence of clear rewards and placed the hippocampus firmly at the centre of the proposed structural learning network.

Many computational studies have since shown that a models based on a structural learning framework can not only reproduce the types of responses recorded in navigation tasks, such as place cells, border cells and many others (Whittington et al., 2022), but also can potentially explain findings from contextual fear conditioning experiments (Gershman et al., 2015) as well as attentional set-shifting tasks (Niv, 2019).

These models have made much headway in demonstrating that structural learning principles offer a unifying framework to describe hippocampal activity across various domains, from spatial navigation to memory encoding. However, only few studies have attempted to experimentally test these hypotheses due to the difficulty in training mice to solve these kinds of structural problems.

We have already shown that the learning in our task satisfies the criteria for structural learning since the cues are temporally separated, individually ambiguous and only become meaningful if combined in a specific order in time. In our next steps, we will test whether mice that are proficient at the task can adapt to changes in task structure and cue identity, testing whether the hypothesised degree of abstraction necessary for structural learning incurs higher degrees of behavioural flexibility.

Furthermore, we will investigate whether hippocampal circuits are necessary for this task by inactivating parts of CA1 during the task with optogenetic manipulations.

4.2 Results

4.2.1 Mice can adapt to changes in task structure

In a first step, we wanted to test how flexibly mice can express the behaviour learnt in the initial version of a task. To test this, we first systematically altered the timing of the task, such that the delay period between the two odour cues grew from 5s to 10s, 20s and eventually to 30s (**Figure 4.1a**).

As would be expected for a higher memory load, the accuracy of behaviour decayed slightly as the delay became longer (**Figure 4.1b and c**; n = 10; repeated measures ANOVA: $F_{(3, 30)} = 4.78$, $p = 7.7 \times 10^{-3}$), but nevertheless, mice remained substantially above chance across all delay periods (5s delay: p correct = 0.84 +/- 0.06; 10s delay: p correct = 0.84 +/- 0.10; 20s delay: p correct = 0.79 +/- 0.12; 30s delay: p correct = 0.68 +/- 0.11). Notably, mice refrained from licking throughout the delay period even at long delays (**Figure 4.1d**). Thus, once mice have learnt the task, they could adapt to changes in cue timing and still correctly anticipate reward.

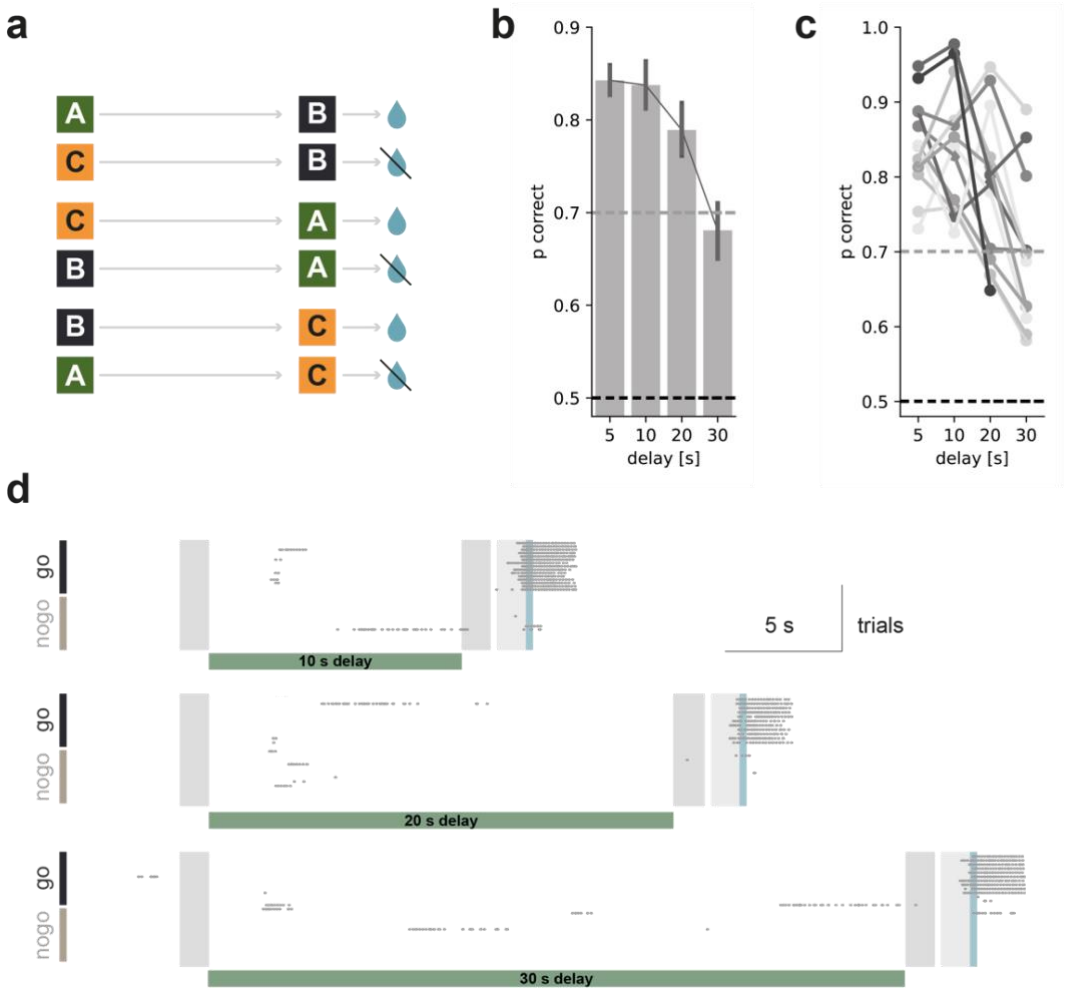

Figure 4.1: Mice can adapt to temporal changes of task structure. a. Different delays serve as an example for changes in the temporal structure of the task: instead of 5s, the delay between the two odour cues is changed to 10s, 20s, or 30s respectively. **b.** Mice can perform the task even at 30s delay (n = 10) **c**. Data from individual mice trained on paradigms with 5s, 10s, 20s and 30s delay **d.** Example licking behaviour for 'go' and 'no go' trials in the 10s, 20s and 30s paradigms respectively.

Since behavioural accuracy decreased with longer delays, we wanted to ensure that mice were still using the full structural information of the task, instead of resorting to a different strategy.

We therefore returned to the logistic regression analysis described in Chapter 3. As before, we split predictor variables into four individual models (**see Figure 3.5b**): a structural model (S) including variables for each odour pair, taking into account their order; an elemental model (E) using only presence or absence of each individual odour cue; a configural model (C) using presence or absence of each odour configuration (making no distinction between AC and CA); and finally a model using past choice and past reward as predictors (H).

Comparing between these individual models, we found that at all delay times, the structural model consistently best predicted licking behaviour (**Figure 4.2a and b**; mixed-design ANOVA, effect of model type: $F_{(3,27)} = 58.90$, $p = 5.7 \times 10^{-12}$; effect of delay: n.s; interaction: n.s).

As a second approach to understand the strategies used by mice, we employed a Feature Subtraction approach. For this, we used a combined model using all predictor variables (from models S, E, C and H) and compared its performance to that of a model where one set of variables was removed. This difference in performance was quantified by the delta AUC (ΔAUC) which therefore represents the unique variance attributable to each group of predictors. As an example, the ΔAUC score for the structural model S describes the change in performance between a complete model and a model using only the predictor variables from E, C and H (i.e. with the structural predictors removed).

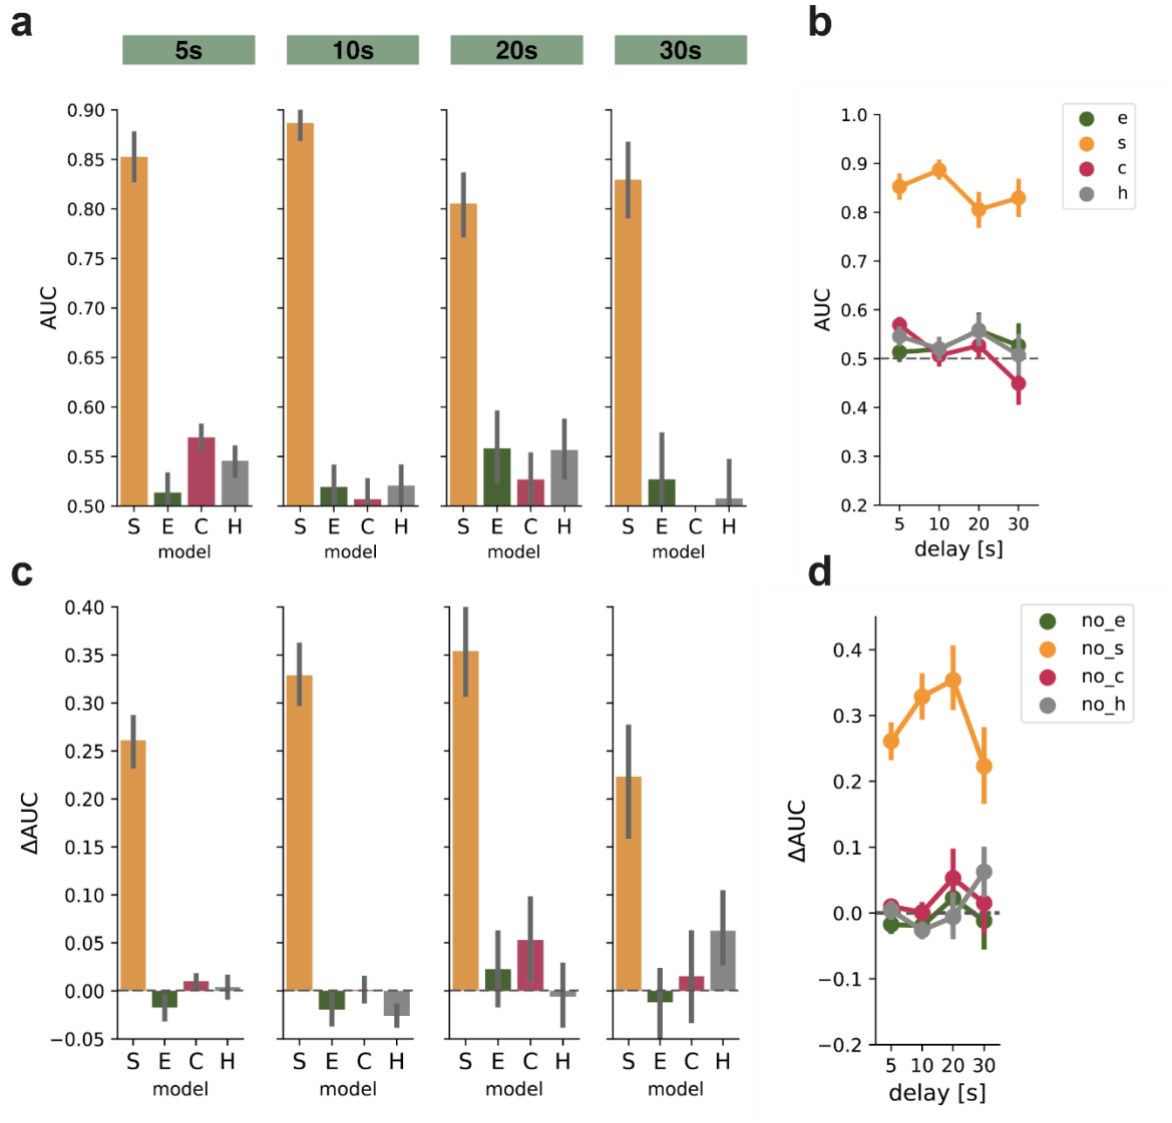

Figure 4.2: Regression analysis shows that mice solve task using the order of cues even at long delays. a. Area under the curve for single models using different predictor (S: structural, E: elemental, C: configural, H: trial/outcome history) for different delays. **b.** Summary of the performance of each model across different delays. Note that the structural model consistently best predicts licking behaviour. **c.** Change in performance of a global regression model when predictors from specific models are removed (ΔAUC) for each delay. **d**. Summary of the change in model accuracy when either structural, elemental, configural or historical information is removed.

Consistent with the results from the single models, removal of the structural predictors has the biggest impact on the performance of the regression model (**Figure 4.2c and d**; mixed-design ANOVA: effect of model type: $F_{(3,28)} = 40.01$, $p = 2.9 \times 10^{-10}$; effect of delay: $F_{(3, 84)} = 3.07$, $p = 0.03$; interaction: n.s.).

Taken together, this suggests that despite the decrease in the behavioural accuracy at long delays, mice still used a structural strategy to solve the task throughout.

Despite this confirmation, we wanted to explore the decrease in performance with longer delays further, examining the types of errors mice made in sessions with longer delays. Within a go/no-go task, there are two possible errors: missed trials (rewarded trials in which mice failed to lick) and false alarm trials (unrewarded trials in which mice licked incorrectly). Therefore, the decrease in correct trials at longer delays could either be caused by a higher proportion of false alarm responses, a higher proportion of missed trials, or a combination of both (**Figure 4.3a**).

To examine this, we compared the proportion of these four outcomes across the different delays (**Figure 4.3b, c-f**). Consistent with the finding that mice perform above chance level even in the 30s paradigm, we observed that the proportion of correct trials was consistently higher than that of error trials of either kind. Interestingly, with increasing delays the proportion of missed trials consistently rose, accompanied by a matching decrease in hit trials (n = 13, Pearson Correlation Coefficient for p miss : $r_{(11)} = 0.38$, $p = 0.01$). The rate of false alarm trials however showed no clear correlation to the delay time, remaining between 5 – 10% for all delays (n = 13; Pearson Correlation Coefficient for p false Alarm $r_{(11)} = 8 \times 10^{-4}$, $p = 0.44$).

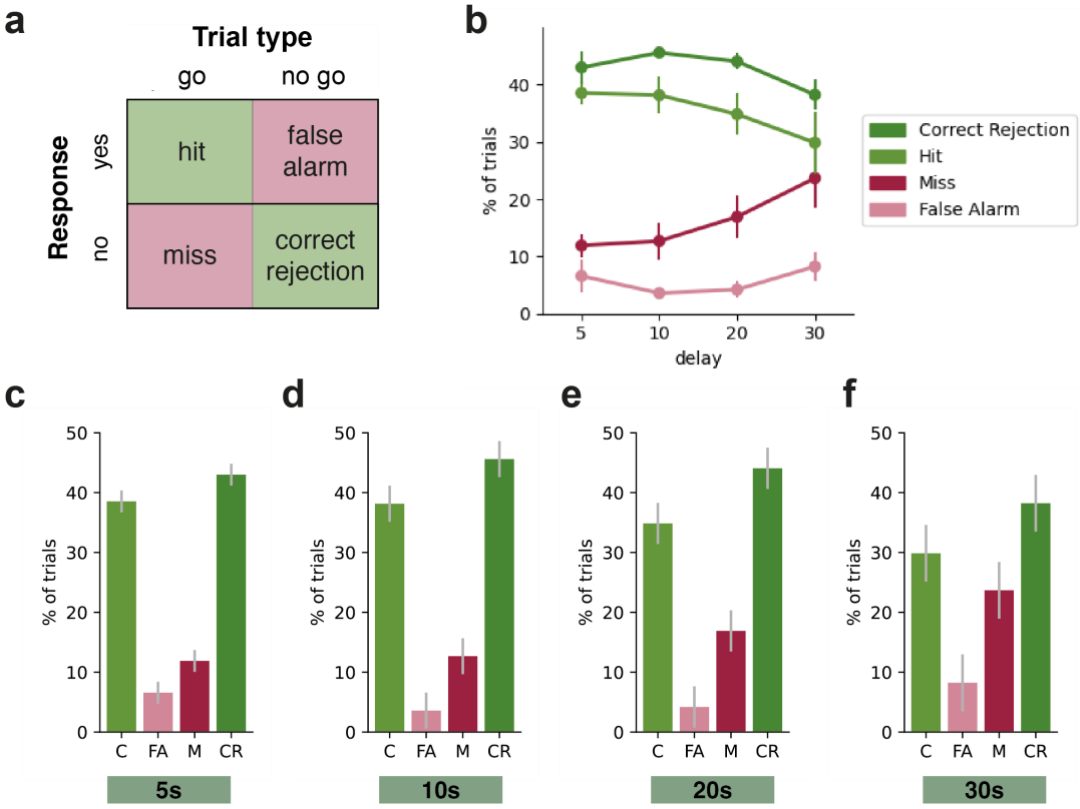

Figure 4.3: Mice miss more rewarded trials at longer delays. a. Within our task, four possible outcomes are possible for a trial, depending on the trial type (go/no-go) and the response (lick/no lick). **b.** Rate of each possible outcome for task paradigms with different delays (summary) **c - f** Occurrence of each outcome in 5s, 10s, 20s, and 30s task, respectively.

All mice first learnt the task with a 5s delay (as shown in **Chapter 3**) and were only moved onto longer delays once they had become proficient. In the same way, only mice who stably performed the task with a 10s delay would be presented with the 20s paradigm, leading to mice reaching the maximal delay of 30s after a variable amount of training days.

Since the motivation for this manipulation was to test whether structural learning conferred greater behavioural flexibility, we next compared the training times and performance on day 1 of training across the different delay paradigms.

Traces in **Figure 4.4b** show that the individual learning trajectories vary across mice: while some mice took as long as 9 days to reach a performance level significantly above chance in the 20s task, others were able to perform at above 80% accuracy from day 1.

On average, the time needed to acquire the task at new delays decreased in tendency over the learning period (**Figure 4.4a**). While it took 15 days for all tested mice to reach good performance of the 10s paradigm, it only took 7 days to acquire the 30s task. Furthermore, the performance on day 1 of training in the 10s and 20s paradigm (i.e. the first day mice ever experienced that particular paradigm) lay higher than performance on day 1 of the 5s paradigm (**Figure 4.4c**). However, when looking at the individual learning curves (**Figure 4.4b**), it is notable that different mice seem to react differently to a change in time structure: some mice (e.g. SN74 or SN87 for the 10s task) maintain a high level of correct trials from the first day, while others (SN91) fall to chance level for a few sessions before recovering performance.

Taken together, these results suggest that knowledge of prior task paradigms may transfer to paradigms with different temporal structures to some mice, providing a learning advantage.

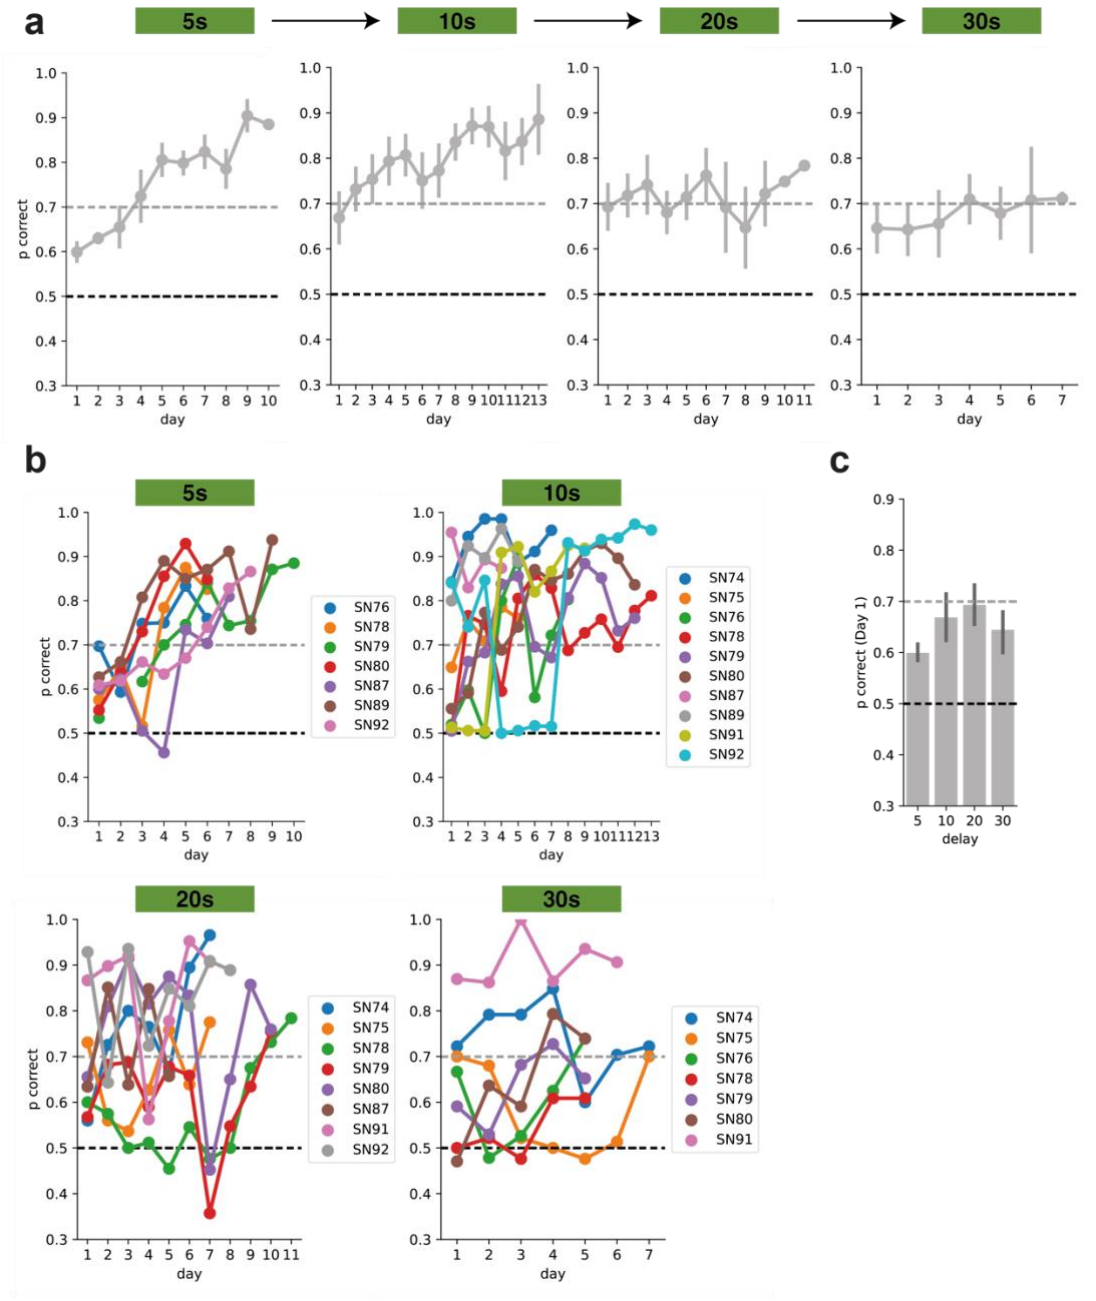

Figure 4.4: Time to learn new delays gets shorter over learning. a. Average performance for the task at different delays. Note that once mice perform above 70% consistently, they are moved to the next stage. **b.** Learning curves for individual mice. Days with manipulations (such as "No airflow" control experiments (cf. Fig 3.4) are omitted in these plots. **c.** Average performance on Day 1 of training in the task at different delays

To further examine the evidence for knowledge transfer from one task paradigm to another, we compared the learning curves of mice that underwent different delay paradigms sequentially (first encountering the task with a 5s delay, then moving to 10s when proficient; schematic in **Figure 4.5a**, left) to the learning curves of mice whose first experience with the task was the 10s paradigm (without any training on or experience of the 5s task, **Figure 4.5b**, left). Thus, in this experiment, we directly compared learning of the 10s task with or without previous experience of the 5s task.

While both cohorts reached a high level of task performance by day 9 of training, the average performance of mice with previous experience of the 5s task exceeded 70% from day 2 onward, while naïve mice only reached this level after 8 days of training (**Figure 4.5c**). Even before the mice reached criterion, the difference between experienced and naïve mice was evident, which we quantified by comparing the average performance across day 1-3 of training between these groups (n = 9 and 3 for experienced and naïve group, respectively; two-sided t-test $t_{(8.54)} = 3.73$ $p = 0.01x10^{-2}$).

In summary, we demonstrate that mice trained on our paired-associates task with a 5s delay between cues can adapt to changes in the task's delay structure, doing so more efficiently than mice with no prior experience of the task structure.

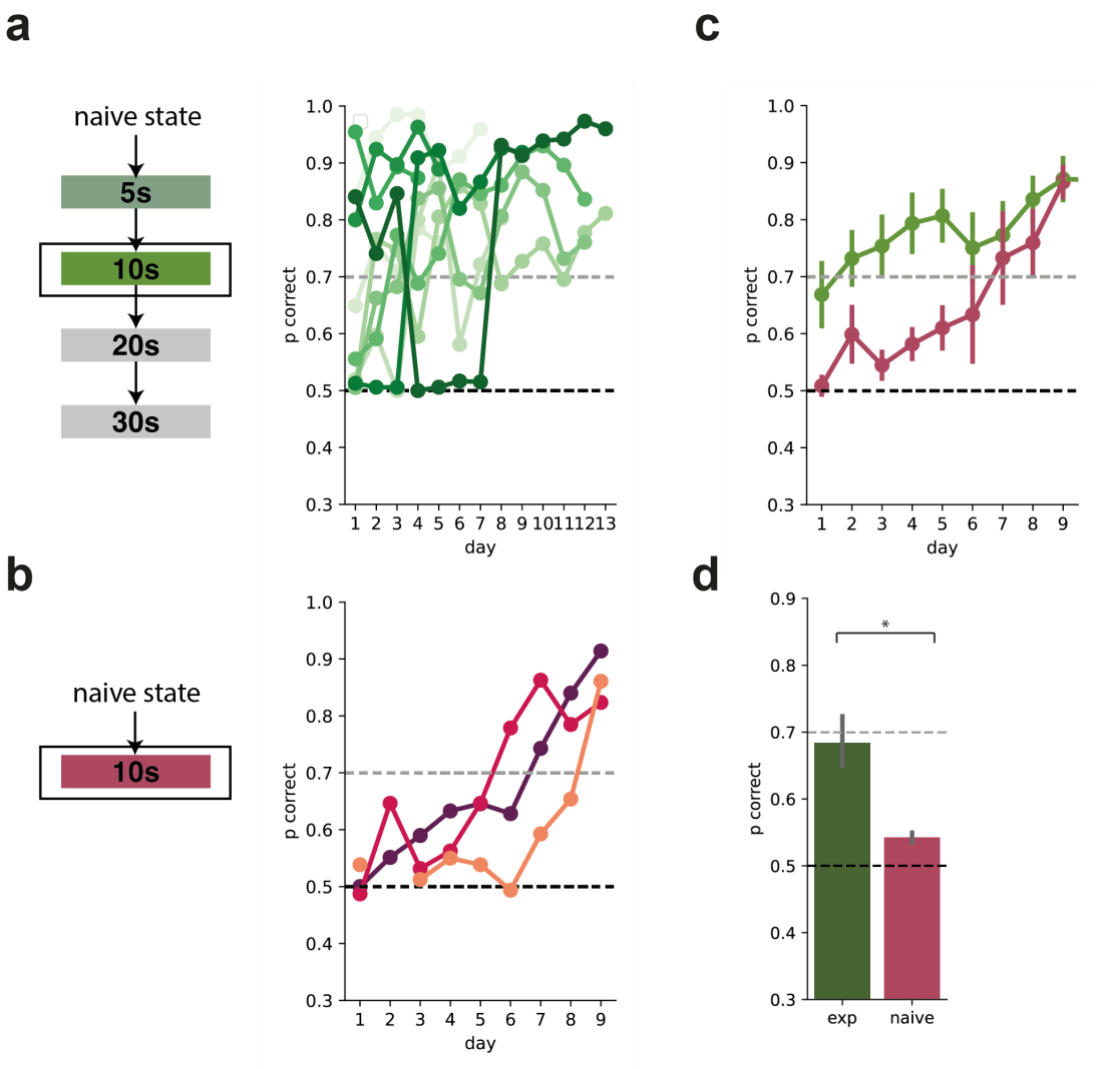

Figure 4.5: Mice with experience of the 5s task paradigm perform better in 10s task. **a**. Individual learning curves of mice training on a task paradigm with 10s delay, after having mastered the basic version of the task with 5s delay. **b**. Learning curves of mice that have no prior experience of the task before encountering the 10s paradigm. **c**. Comparison of average performance across learning between experienced mice (green) and naïve mic (red). **d**. Average performance on the first days of learning (day 1-3) in experienced and naïve mice (n = 9 and 3 for experienced and naïve group, respectively)

4.2.2 Mice can incorporate a new odour into existing task structure

In addition to allowing for flexible adaptation to changes in task architecture, structural learning is also proposed to allow generalisation to novel experiences that share the same underlying structure. To test for this type of behavioural flexibility, we changed the sensory experience of the task while leaving the temporal structure constant.

We achieved this by replacing odour C with a novel odour D (**Figure 4.6a**). Specifically, we introduced 4 new odour pairs (AD, DA, BD, DB), with only 2 trial types remaining familiar from the basic version of the task (AB and BA). Despite this new cue configuration, the underlying structure (two odours separated by 5s, then outcome) remained the same, allowing us to observe how quickly mice could incorporate a new sensory cue into a familiar structure of rules.

We used mice who had learnt the task with the standard set of odours (ABC) and were performing at a high level (**Figure 4.6b**, day -1). When we introduced odour D into the task (day 0), the accuracy of anticipatory licking dropped initially as expected. However, task performance rapidly returned to very high levels (**Figure 4.6b;** $n = 6$, repeated measures ANOVA: $F_{(5, 25)} = 2.66$, $p = 0.04$). In some instances, this recovery even took place within a single session, after the mouse had only very limited exposure to each pair of odours (**Figure 4.6c**). In contrast to learning a new temporal structure (**Figure 4.3**) or learning an entirely new task (see **Chapter 3**), this recovery of previous performance was rapid, suggesting that prior experience with a given task structure facilitates faster learning of new cue combinations that share the same structure.

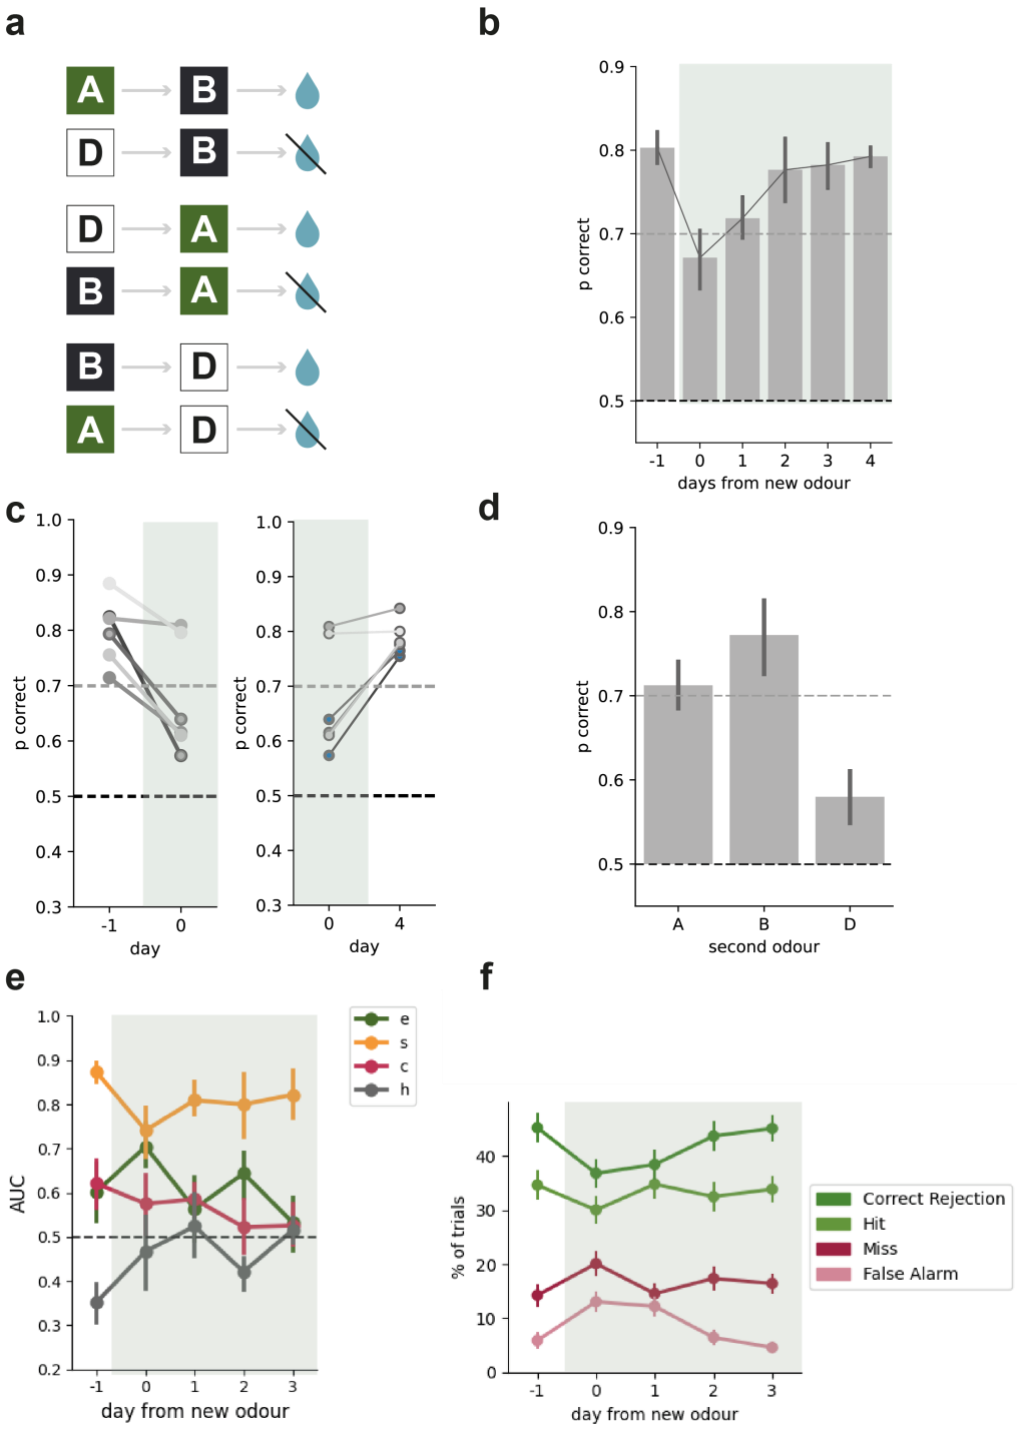

Figure 4.6: Mice can rapidly integrate new odour into task structure. a. After mice have fully learnt the task, odour C is replaced by a novel odour D **b.** Mice rapidly adjust and recover performance (n = 6) **c.** Most subjects experience initial difficulty manifesting in a drop in performance on the day the new odour is introduced (day 0) compared to the previous day (day -1) but recover performance within 1 – 5 sessions (day 5) (n = 6) **d.** Performance on the day of the switch (day 0) grouped by second odour. Performance on trials with the new odour (D) as second odour are most affected**. e.** After a brief rise in the predictive power of both the configural and the elemental model in the days following the switch, the regression shows that after 3 days, the mice return once more to using mostly structural information to solve the task. **f.** This manipulation leads only to a minimal rise both in missed rewarded trials as well as incorrect lick responses to unrewarded odour pairs.

To better understand the effect of replacing one of the odour stimuli, we next looked at the performance on the day of the odour switch (day 0), split according to the second odour in the pair (**Figure 4.6d**). As would be expected, on trials where the new odour D is the second odour, mice perform poorly since both trial types within that group are unfamiliar. In trials where A or B are the second odour, mice perform above criterion, but notably do better on trials where B is the second odour. Overall, this result supports our notion that the mice are able to use their knowledge of the structure of the task to adapt to the exchanging of a familiar cue for a novel one.

This was supported by the results of a logistic regression with the same sets of predictor variables as previously described: even on the day of the odour switch, the structural model remained the most predictive of mouse behaviour (**Figure 4.6e,** mixed-design ANOVA, effect of model: $F_{(4,27)} = 34.40$, $p = 5.4x10^{-5}$; effect of day: n.s; interaction: n.s). On the day when the novel odour was first introduced, the increased predictive power of elemental features suggests that there might be an initial phase of bias while mice learn where the new cue fits into the task structure.

Despite a slight rise in miss and false alarm trials on the switch day, there was no statistically significant effect in the types of errors mice made (**Figure 4.6f;** n = 6, repeated measures ANOVA: $F_{(4, 20)} = 1.09$, $p = 0.38$), suggesting that across the 5 days, these response types did not change in a consistent way across mice.

4.2.3 Ventral Hippocampus is required for solving the task at longer delays

Our results so far suggest that mice indeed solve the task by learning about both the task structure as well as individual cues, enabling them to adapt flexibly to changes to cues as well as to the timing of their delivery.

As mentioned in the introduction to this chapter (section 4.1), learning and use of task structure have been proposed to be dependent on the hippocampus, supported by both experimental results as well as theoretical studies. Therefore, we wanted to test how disrupting neural activity in this area would affect behavioural accuracy in our task.

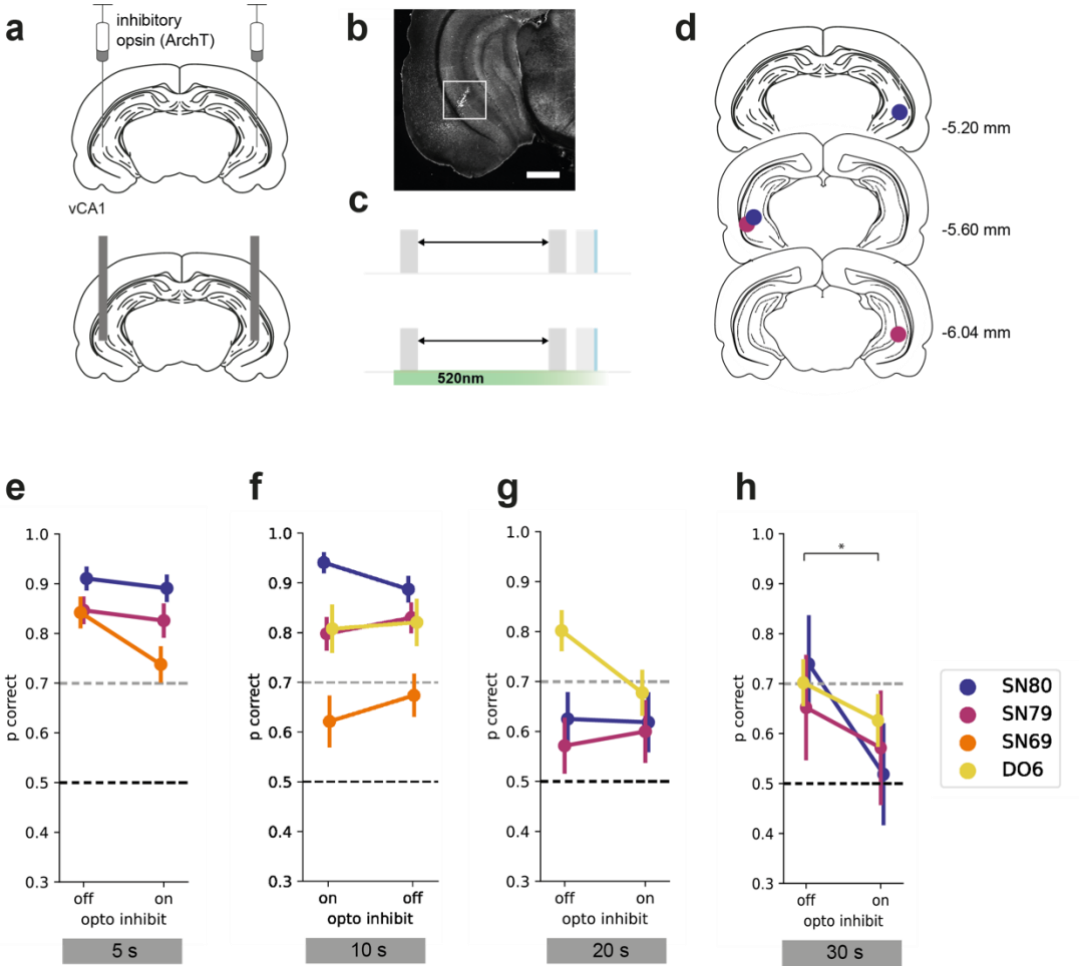

Figure 4.7: Optogenetic inactivation of ventral CA1 impairs behaviour at long delays. **a.** Schematic of bilateral optogenetic inhibition of ventral CA1 region of hippocampus (vCA1): stereotaxic injection of a flexed inhibitory opsin ArchT combined with a CamKii-cre virus (top) and fibre placement (bottom) **b.** Example histology image showing fibre placement, scale bar = 1mm. Box indicates region of interest **c.** Schematic of optogenetic inhibition experiment: on trials with laser stimulation, the laser pulse starts at 500ms before the first cue and fades out after water delivery **d – g**. Bilateral optogenetic inactivation of vCA1 has no significant effect on task performance at delays of 5 – 20s, but in a 30s delay paradigm, laser stimulation impairs performance significantly (n = 3), indicating a role for vCA1 for structural learning with long delays.

To test this hypothesis, we bilaterally injected a mixture of *AAV2/1-FLEX-ArchT-tdTomato* and *AAV2/1-CamKII-Cre* into the CA1 area of ventral hippocampus (vCA1) and implanted optic fibres above the injection site for directed laser light delivery. This enabled us to inhibit principal neurons in vCA1 selectively and reversibly via administration of 520nm illumination (**Figure 4.7a and b**). Due to unforeseen circumstances, we were only able to recover the histological data of two out of the four animals used in this manipulation. The two brains varied in viral expression (data not shown) but both injection site as well as fiber placement were comparable (**Figure 4.7d**).

We tested the effect of this optogenetic inhibition on the behaviour of well-trained mice in versions of the task with either 5s, 10s, 20s or 30s delay. We designed the stimulus such that light was automatically delivered for the duration of an entire trial, from 500ms before the onset of the first cue until after the outcome (**Figure 4.7c**). In addition, light delivery was turned off with a "ramp" to minimise rebound spiking (Chuong et al., 2014). In each mouse, we compared sessions with light delivery to control sessions with no light delivery.

In experiments with 5s and 10s delays between odour cues, bilateral inhibition of vCA1 had no effect on the performance of expert mice (**Figure 4.7e and f**; 5s: n= 3; p correct $_{opto}$ = 0.82 +/- 0.08, p correct $_{control}$ = 0.87 +/- 0.04; 10s: n = 4; p correct $_{opto}$ = 0.79 +/- 0.13, p correct $_{control}$ = 0.80 +/- 0.09). At 20s delay, there was a marked decrease of behavioural accuracy in one animal, but since the other two subjects performed poorly even under control conditions, this data did not allow for conclusions about a general effect (**Figure 4.7g**; 20s: n = 3; p correct $_{opto}$ = 0.63 +/- 0.04, p correct $_{control}$ = 0.67 +/- 0.12).

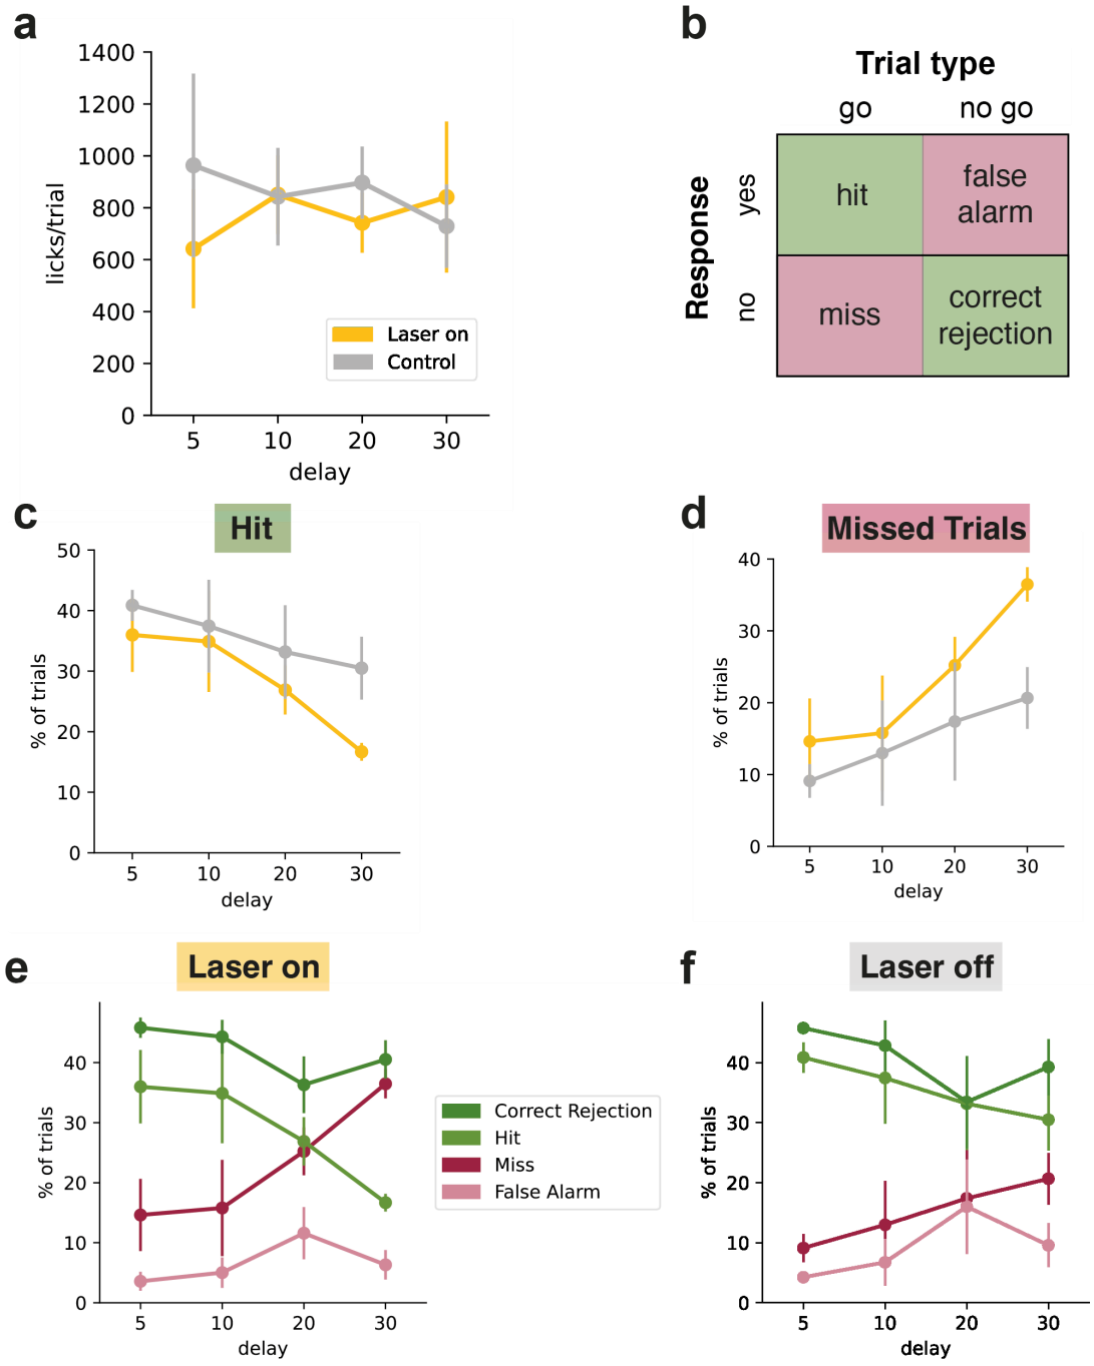

Figure 4.8: Optogenetic inactivation of vCA1 leads to increased number of missed trials at long delays. a. The average number of licks per trial does not change between trials with and without laser stimulation **b.** Possible outcomes in our task are either Hit or Miss for a rewarded trial, or Correct Rejection and False Alarm for unrewarded trials **c.** The frequency of Hit trials per delay time for both trials with optogenetic stimulation (green) as well as control trials (grey) **d.** The frequency of Miss trials in tasks with different delays, green: laser on, grey: control trials **e**. Percentage of trial outcomes in trials with laser stimulation. **f.** Percentage of trial outcomes in trials without laser stimulation.

However, performance accuracy in a task with 30s delay was significantly impaired in sessions with laser stimulation (**Figure 4.7h;** 30s: n = 3; p correct $_{opto}$ = 0.53 +/- 0.05, p correct $_{control}$ = 0.70 +/- 0.04; paired t-test: $t_{(2)} = 8.58$, p = 0.01) and fell to almost chance levels.

To examine what change in behaviour underlay this marked decrease in task performance, we next looked more closely at the licking responses in trials with and without laser stimulation. Since previous studies have proposed behavioural inhibition to be one of the functions of the ventral hippocampus (Gray & McNaughton, 2003), we compared the amount of licking per trial between the sessions with optogenetic inhibition and control sessions (**Figure 4.8a;** two-way repeated measures ANOVA, laser: n.s; delay: n.s., interaction: n.s.). We found that at all delays, mice licked equally frequently under both laser and control conditions.

Then, we returned to the response analysis (**Figure 4.8b**) to determine what types of responses formed the basis of the decrease in behavioural accuracy in laser-stimulated sessions with long delays.

We found that optogenetic inhibition of vCA1 led to a steeper decrease of Hit trials with longer delays (**Figure 4.8c,** repeated measures ANOVA, $F_{(1, 3)} = 40.00$, $p = 7x10^{-3}$) and a matching increase in Miss trials (**Figure 4.8d**). For example, while mice missed 18% +/- 2.4% of rewarded trials in the 30s task under control conditions, this figure rose to 35% +/- 1.2% of rewarded trials when vCA1 was inactivated. In contrast, there was no difference in False Alarm responses between sessions with optogenetic stimulation (**Figure 4.8e,** repeated measures ANOVA: $F_{(1, 3)} = 7.16$, p = 0.08) and control sessions (**Figure 4.8f**).

4.3 Discussion

In this chapter, we wanted to test whether mice could use the knowledge of the task even when cues or their timing were changed, suggesting a level of abstract learning about the task structure that has been postulated by previous studies.

First, we showed that mice proficient at our task could adapt to changes in the underlying temporal structure such as increasing the delay from 5s up to 30s (**Figure 4.1**). Importantly, even though the accuracy of their responses decreased slightly with longer delays, our logistic regression models suggest that mice still used structural information to solve the task (**Figure 4.2**). We next looked at the types of errors mice made across these different task paradigms, and found that with increasing delays, the number of missed trials (rewarded trials where mice failed to respond) increased, whereas the number of false alarm trials (unrewarded trials where mice incorrectly responded) showed no significant correlation with delay length (**Figure 4.3**).

We further show that mice adapt to changes in the timings of cues in a shorter time as compared to the time they require for learning the original task (**Figure 4.4**) and that their behavioural accuracy in early learning is higher than in mice with no prior experience (**Figure 4.5**), indicating that some of the knowledge can be transferred across tasks.

To further test behavioural flexibility in our task, we then changed some of the odour pairs, replacing one odour cue with a new odour that the mice had no prior experience of. We found that mice adapted to this new cue rapidly, recovering their previous levels of accuracy by the third session, and often before that (**Figure 4.6**). We used regression to investigate which variables best predicted mouse behaviour and found that once again, licking was best predicted by a structural model using both odour identity and order. Lastly, we examined the types of responses and found that in the sessions following the odour switch, there is no specific change in errors.

Finally, we investigated whether the hippocampus was required for this proficient performance in our task. To do this, we bilaterally expressed an inhibitory opsin in ventral CA1 (vCA1) and stimulated it by delivering laser light through optic fibres while mice were performing the task (**Figure 4.7**). Inactivation of vCA1 only showed a significant effect on behaviour in the paradigm with 30s delay between odour cues.

This effect was not due to a change in the amount of licking. Using the same response analysis as before, we showed that the drop in behavioural accuracy during optogenetic inactivation was due to a large decrease in correct responses to rewarded trials but was not related to a higher amount of false alarm trials (**Figure 4.8**).

4.3.1 Adaptation to changes in task structure suggests some learning across conditions

To test how flexibly mice were able to express their learnt understanding of the task, we chose to successively change the delay between cues from 5s to 10s, 20s, and finally 30s.

While at first glance, this might not seem like a change requiring an abstract understanding of task structure, it warrants a closer look: to successfully recover performance after the delay changes from 5s to 10s, the animal either has to learn 6 new associations with this new time course or transfer the previously learnt reward contingencies to a longer trial time. The latter strategy is proposed to be more efficient and therefore allowing a faster recovery of performance after a change in temporal structure. However, considering proposed mechanisms of time keeping in the brain, such as e.g. time cells (MacDonald et al., 2013; Eichenbaum, 2014), the neural mechanism of this kind of transfer learning would require a generalisable representation of the task.

In our data, we find hints that mice with previous task experience indeed adapt faster to manipulations of the delay time, both by reaching stable accurate performance earlier and by displaying a higher baseline of behaviour even upon first exposure (**Figure 4.4 and 4.5**).

An important consideration in interpreting these results is that the key comparison is not between our task and a simpler one in which single odours predict reward. While such a task might not require structural learning, repeated exposure to different cue-reward pairings with varying odours or delays would still establish a common task framework, even if it would be a very simple one.

Instead, the appropriate control would be a task in which no common structure exists across successive learning experiences. For example, if mice were first trained on our original task – where reward is predicted by an odour pair with a 5s delay – and then switched to a task where only the second odour predicts reward, or where a distractor odour is presented in the delay, we would not expect prior experience to facilitate learning. In such cases, the underlying task structure would be fundamentally different, preventing transfer of previously learnt contingencies.

However, since we did not perform a control experiment with a non-structural task with similar cues, we cannot exclude that the improved performance on day 1 might be due to alternative explanations. Since the exposure to multiple task paradigms necessarily means more total days of training, this improved learning might also be explained by increased attention to odour cues, or simply higher habituation to head fixation, handling and water restriction.

4.3.2 Behavioural performance after new cue introduction

Our second test to assess behavioural flexibility in our task was to change the cue contingencies. We did this by exchanging one of the three odour cues (odour C) with a new odour (Odour D) that the mice had no prior experience of. In consequence, 4 out of 6 trial types contained novel odour pairs with unfamiliar associations, half of which were rewarded and the other half unrewarded.

We found that mice rapidly adapted to this change, all recovering their baseline levels of behavioural performance by the third session with the new set of odours (**Figure 4.6**). Notably, some mice responded correctly to the new set of odours even within the first day of exposure, thereby requiring less than 10 trials of each type to reconsolidate their behaviour.

This was a stark difference from the time required for learning the original associations (see Chapter 3, **Figure 3.2**), suggesting that maybe abstraction from the basic task allowed for inference to fill the gaps. Theoretical accounts have proposed that a representation built following the tenets of structural learning should allow the modular assembly of task representations, adding and removing cues without having to re-learn the underlying structure (Whittington et al., 2020), offering a computational mechanism by which this type of fast learning might occur.

However, without neuronal data, we cannot make any definitive statements on the role of generalised representations in this behaviour, since several other hypotheses might explain the fast adaption to the new odour. For one, it is possible that despite our best efforts to pick odours that are dissimilar from each other (Pashkovski et al., 2020), odour D was more similar to odour C than to both A and B, therefore making it a problem of pattern completion rather than generalisation. Another possible strategy might be to infer the meaning of the new odour by a process of elimination – if both A and B are still present, the new

odour must therefore be replacing C. Note however that this type of conditional learning still requires an understanding that there are only 3 odours combined into pairs within each trial.

A further limitation of this manipulation is the fact that there is only a short window of time before the mice adapt to the new odour, and a repeated replacing of odour cues would be expected to yield different results since the overall exposure to each cue would be imbalanced.

In the future, it would be interesting to record single cell activity during this task, gaining insight into how odour D is represented in the neuronal population, and whether the neural data supports the hypothesis of a generalised task representation able to flexibly include new odour cues.

4.3.3 Role of vCA1 in bridging delay time

To assess the extent to which hippocampal circuits are necessary for structural learning, we optogenetically inactivated neurons in vCA1 in mice performing the task (**Figure 4.7**) and found that this manipulation markedly impaired task performance. This finding ties in with results of sequence learning studies in rats (Fortin et al., 2002; Kesner et al., 2002) as well as configural learning tasks (Sanderson et al., 2006) and strengthens the hypothesis that the hippocampus may indeed be necessary for this type of learning.

However, this impairment was seen in experiments targeting ventral CA1, in contrast to the above studies which were targeting more dorsal areas of the hippocampus. In the literature, dorsal and ventral hippocampus have been proposed to be functionally distinct due to their differences in gene expression and anatomical projection patterns (Fanselow & Dong, 2010; M. Moser & Moser, 1998), but what exactly the respective role of dorsal and ventral hippocampus might be is not yet fully understood. This is partly due to the general difficulty of

comparing results acquired through different behavioural paradigms, but also due to a bias in anatomical targets between the fields: the navigation field has historically focused on dorsal regions (Bittner et al., 2017; M. B. Moser et al., 1995; O'Keefe & Nadel, 1979) whereas fear conditioning studies have more frequently targeted ventral regions (Kjelstrup et al., 2002; Maren & Holt, 2004; Twining et al., 2020). Without comparative data from more dorsal areas, it is too early to say whether the results of our optogenetic inactivation experiments support the theory of a specific role for ventral hippocampus in structural learning. In the rest of this work, we will focus on the contribution of ventral hippocampus to our task, but an exciting future direction is to explore the different contributions of dorsal and ventral circuits in more detail.

A second interesting finding from our optogenetic manipulations is that inactivation of vCA1 affects behavioural performance only in task paradigms where the delay exceeds 10s. This implies that at short delays, hippocampus might not be required, and other circuits seem to support both the encoding of the stimuli as well as their relationships to each other. Candidate regions for this type of short time-scale binding of cues could be the piriform cortex as well as prefrontal regions (Franks et al., 2011; Liu et al., 2014). Hippocampal circuits only seem to be essential to the task performance at delays longer than 10s. This result resembles the findings that hippocampal lesions have larger effects in fear conditioning paradigms including a delay period (trace interval) before the aversive experience (Quinn et al., 2002), or spatial studies showing that a Delayed-Match-To-Place (DMTP) paradigm requires hippocampus only at longer delays between cues (Spellman et al., 2015).

Interestingly, the deficits in behavioural performance in sessions of the 30s task with inactivation manifest specifically in a larger proportion of missed trials, while the proportion of false alarm responses remains unchanged. This supports our hypothesis that the underlying mechanism is due to the specific role of vCA1

in processing temporal sequences and structural information, rather than a non-specific effect.

One possible explanation for this pattern is that vCA1 inactivation may impair the ability to maintain or retrieve the memory of the first odour over longer delays. Rather than responding incorrectly at chance level, mice may adopt a more conservative decision-making strategy, resulting in a higher rate of omitted responses. If this is the case, there might be an interesting gradual disruption of response confidence with the increase of difficulty (i.e. longer delay times).

Analysing lick rates across these different tasks could provide further insight, e.g. comparing the lick rate in hit trials without laser stimulation to those with laser stimulation in the edge-case of the 20s delay task. If the lick rate is lower under optogenetic stimulation, that might suggest a reduced confidence, especially since the overall licks per trial stay stable. If, on the other hand, lick rates remain similar, this would suggest a failure in recall rather than a shift in decision thresholds.

Investigating such confidence-related behaviours could help refine our understanding of how vCA1 contributes to maintaining and integrating temporally separated cues to guide behaviour.

Taken together, these results suggest that the hippocampus is not needed for this structural learning task per se, as the mice can perform the 10s paradigm well despite optogenetic inhibition. Hippocampal activity only becomes necessary when the task-relevant cues are temporally distant from each other, as they are in the 30s paradigm. This is not a new idea in the field – in the literature, this has been proposed multiple times over the last decades and several potential mechanisms for this binding of cues across spatial and/or temporal distance have been suggested. However, my data show this very clearly for the first time within the same mice and the same experimental paradigm.

To further dissect the role of vCA1 in bridging temporal gaps between stimuli, a critical next step would be to test whether brief, precisely timed optogenetic inhibition within the delay period is sufficient to impair performance. If vCA1 activity is specifically necessary for maintaining stimulus representations across time, then even a transient perturbation – such as a short laser pulse midway through the delay – should be disruptive, potentially yielding similar deficits to continuous inhibition. This experiment would help differentiate whether the observed effects are due to the cumulative duration of inhibition or whether vCA1 is specifically required during critical moments of the delay. Additionally, selectively inhibiting vCA1 only during the first or second odour presentation could clarify whether the hippocampus primarily contributes to encoding, maintenance, or retrieval of stimulus representations. These experiments would allow for more precise mapping of when vCA1 activity is essential within the trial and help refine our mechanistic understanding of its involvement in structural learning.

4.3.4 Technical considerations

The finding that there is no effect of optogenetic inactivation of vCA1 on task performance offers evidence that the laser stimulation by itself is not affecting the perception of stimuli, general motivation, or reward-seeking behaviours (such as anticipatory licking). The latter is pertinent especially since some studies have associated the ventral part of hippocampus with behavioural inhibition (Gray & McNaughton, 2003). We therefore looked specifically at the amount of licking across laser-stimulated and control trials at all delay lengths (**Figure 4.8a**) and found that the average amount of licking per trial was not affected by the length of the delay nor by the presence or absence of laser stimulation.

It would have been valuable to relate the effects of optogenetic inhibition in individual animals to the specific fibre placement and extent of viral expression, especially since we see some variability in effect especially during the 20s task. However, since we were only able to recover two brains, this is not possible with the data available. In future experiments, it would be interesting examine this relationship more thoroughly, using a larger dataset to assess the influence of fibre placement and expression patterns more systematically.

Another consideration concerns the length of laser stimulation: since we are inhibiting throughout the entire trial, the length of the delay time is directly correlated with the length of laser stimulation. This means that in the 30s delay paradigm, the laser stimulation lasts for more than twice the time than in the 10s paradigm.

A considerable limitation of prolonged optogenetic inhibition is the potential for non-specific effects beyond neural silencing. Extended laser stimulation can lead to local tissue heating (Owen et al., 2019), which may alter neural excitability and inadvertently affect behaviour. Additionally, prolonged suppression of vCA1 might induce compensatory network changes (de Jong et al., 2023), including shifts in whole-brain activity or engagement of alternative circuits, making it difficult to isolate the specific contribution of vCA1 to task performance. Thus, we cannot exclude the possibility of our results being a consequence of an unspecific effect related either to the laser pulse itself, or the imbalance in the whole brain activity as a result of inactivating vCA1 neurons for a long time.

The finding that mice specifically miss rewarded trials when vCA1 is inactivated while False Alarm trials remain more stable hint at a more specific effect underlying the decrease in behavioural accuracy, but to fully address this concern, control experiments would be needed. To mitigate these concerns, control conditions such as delivering the same total laser power in a temporally

scrambled manner or stimulating an unrelated brain region would be useful in ruling out non-specific effects, as well as the more specific optogenetic manipulations suggested in the previous paragraph.

To further investigate how exactly vCA1 is required for our task, it would furthermore be of great value to use shorter laser pulses to inhibit neuronal activity specifically during the delay or during the presentation of the first or second odour cue, an idea which I describe in more detail in **Chapter 6.4.1**.

In addition, in order to disentangle the effects of the difficulty level (as conveyed by the absolute delay length) and meta-effects on learning, it would be advantageous to counterbalance the order of tasks. Instead of all mice learning the tasks in order of difficulty level, counterbalancing the training schedule such that some mice would have been exposed to the 30s delay paradigm before the 10s version might resolve some of the ambiguity of our results. For example, by comparing performance at the most difficult level (30s delay) across mice with varying levels of prior experience, we could separate the influence of difficulty versus the total amount of training time better. Furthermore, it would allow to counter the unpreventable "selection bias" by which when mice are removed from the experiment early, only a subset of mice reaches the 30s task stage.

Within the scope of this project, I only had the ability to test the involvement of vCA1 when mice adapted to changes in temporal structure. A future goal would therefore be to investigate whether hippocampal circuits are required for our other structural manipulation: the updating of the task structure to include a novel odour D. Since the novelty of this odour cue is the defining feature of the experiment, it is difficult to implement a within-mouse control for optogenetic manipulations, such as the comparison between trials with and without laser in the delay-time manipulations. Therefore, these experiments would likely need larger cohorts of mice to allow for sufficient statistical power to compare effects

between mice that experience a new cue with or without optogenetic inactivation of CA1.

5 vCA1 population encodes cues and context, but not reward

5.1 Introduction

In the previous chapter, I presented data from multiple experiments to probe generalisation to novel cues in the same temporal structure, and the ability to adapt to changes in task structure while observing the same cues. We found that, after initial learning, mice could rapidly adapt to manipulations of cue value and identity, suggesting flexible use of previously learnt relational structures. To assess the extent to which hippocampal circuits are necessary for such learning we optogenetically inactivated neurons in vCA1 in mice performing the task and found that this manipulation markedly impaired task performance in task paradigms where the delay exceeds 10s.

This result ties in with previous studies suggesting that the function of hippocampus is to bind distinct experiences across space and time, thus forming the basis of both spatial navigation and episodic memory (Aggleton et al., 2007; Eichenbaum & Cohen, 2014; Milivojevic & Doeller, 2013). There is however no consensus yet how this computation might be achieved in hippocampus on a neuronal level.

Following from observations about spatial remapping in CA1 populations in animals navigating distinct environments, one theory posits that hippocampus might act as a multiplex of many overlaying contexts, simultaneously representing the spatial environment as well as distinct perceptual, cognitive, and behavioural events that occur within it. Indeed, it has been shown that neurons in hippocampus robustly encode sensory cues (auditory (Kamiński et al., 2017), visual (Aronov et al., 2017) or olfactory (Wood et al., 1999)), and that

the presence or absence of such cues leads to changes in the representation of spatial environments (Anderson & Jeffery, 2003).

This kind of multiplexing of information can be used to disambiguate between reoccurring cues on the basis of their surrounding context. For example, in a spatial non-match-to-sample task in which rats had to approach cups in variable locations and only dig for a food reward when the odour did not match the previous cup, neurons recorded from CA1 did collectively encode both position, odour identity as well as match or non-match (Wood et al., 1999). Interestingly, the representations of these three task dimensions weren't fully separate: some recorded neurons fired to a specific odour in a preferred location, therefore multiplexing separate information streams not only across a population, but within a single neuron.

The data supporting the idea of the hippocampus as a locus of overlaying contextual maps is largely derived from experiments where cues and events are separated in space, but it has been suggested that the function of hippocampus in binding distinct events might even be more general. Multiple studies found that population activity in hippocampus changed gradually across time, thereby not only encoding events within their spatial context, but also in a temporal context that can span minutes to seconds (Manns et al., 2007; Mau et al., 2018).

Moreover, several studies have found that, analogous to place cells that fire when an animal passes a specific position in space, hippocampus also contains so-called "time cells" that fire at a specific point in time (Itskov et al., 2011; Pastalkova et al., 2008), with the activity of all time cells tiling continuous time in the same way that place cells form are thought to form a cognitive map (Eichenbaum, 2014).

In summary, the hippocampus has been shown to encode multisensory representations of the environment, and displays sequential activity tiling both

space and, more importantly for our task, time, and is therefore placed perfectly to bind separate events across temporal delays and thus solve structural learning problems.

Indeed, the discovery of “splitter cells” supports the idea that neurons in hippocampus can perform such integration. These neurons, discovered in rodents performing an alternating T-maze task, distinguish between (“split”) trajectories through the same segment of space depending on recent past, upcoming choice, or inferences about the state of the environment (Wood et al., 2000; Frank et al., 2000), and can therefore be interpreted as a “context-specific place cell”.

And even in tasks without a clear spatial dimension, equivalent “context-specific time cells” have been identified (MacDonald et al., 2013; Taxidis et al., 2020). In their 2020 study, Taxidis et al. found that a significant proportion of the hippocampal neurons recorded in mice performing a delayed non-match to sample (DNMTS) task fired not only at a specific time in the trial but did so only for one of the two cues. The population of these neurons, termed “odour-time” cells by the authors, encoded both the time elapsed since the first cue as well as the identity of that cue.

This provides us with a clear hypothesis of how the identity of the first odour might be maintained across the delay in our task: if our findings align with the above studies, we can expect to find a representation of the cue at the time of the cue presentation which is then followed by a cue-specific ensemble of neurons that maintains this information until the second cue is delivered.

For the representation of the second cue within hippocampal populations, there are however multiple hypotheses: Information about the first and the second odour might be maintained in two separate populations of cue-time cells which can be combined downstream to predict a given outcome. Alternatively, the first

cue could act as a “contextual” cue, activating a different set of neurons not only to encode the delay, but also subsequent cues. In that case, we could expect the neuronal representations of “A after B” and “A after C” to separate in our recordings, analogous to the splitter cells discussed earlier (Ainge, Tamosiunaite, et al., 2007; Duvelle et al., 2023; Wood et al., 2000).

So far, I have focused on the possible representation of cue structure in our task based on previous findings from the literature. There is however another dimension to our task: In addition to binding two odour cues across a delay, the task requires mice to associate these cue configurations to a specific outcome (rewarded or unrewarded), thereby making the predicted value of the respective odour pair a further task-relevant feature that might be encoded in hippocampus.

Studies from the spatial domain have shown that rewarded locations are associated with a higher density of place cells, suggesting that highly salient events such as rewards are prioritised in hippocampal representation (Hollup et al., 2001; Jarzebowski et al., 2022). A representation of the external and internal experience of reward at the time of its delivery is however not sufficient to explain the anticipatory licking behaviour we have shown in mice – for this, a neural response predicting the outcome before the reward is delivered is required.

Whether this prediction of upcoming reward is indeed encoded in neural activity in hippocampus is so far not conclusively demonstrated. Analogous to a pattern first observed in striatal dopamine neurons (van der Meer et al., 2010), multiple studies found single cells in hippocampus ramping up their firing up to salient behavioural events, some of which are associated to reward (Jarzebowski et al., 2022; Wee et al., 2024). Secondly, in line with the idea of hippocampus as a generator of context-specific sequences, a recent study reported wide-spread remapping in hippocampal CA1 place cells after moving the reward location on a linear track (Sosa et al., 2024).

Both individual neurons, ramping to the expected reward, as well as ensembles of neurons tiling the time leading up to a reward might be a possible underlying mechanism for hippocampus to support the kind of prediction required for anticipatory licking. Whether either type of reported neural response can be related to a general representation of reward as opposed to a value-free learnt sequence of likely state transitions is not yet clear, since the tasks used in the abovementioned studies do not allow a disambiguation of the state within the task sequence and the value associated to the state.

In the following chapter, I will present calcium data recorded from CA1 in mice performing our task. To test the hypotheses stated above, we identified neurons representing task-relevant variables such as individual cues, cue combinations as well as upcoming rewards on the single cell level. In a next step, we then investigate the encoding of these variables on a population level via SVM decoders.

5.2 Results

To gain insight into which part of the computation underlying the solving of our task might be performed by vCA1, we recorded calcium activity using the genetically encoded calcium indicator GCaMP7f expressed in principal neurons in vCA1 through a GRIN lens (∅ 1mm, **Figure 5.1 a** and **b**) during the performance of the 5s task.

Using a UCLA V4 miniscope, we recorded neural activity in expert mice. With the partially automated pipeline (Minian pipeline, (Dong et al., 2022)) described in **Chapter 2**, we then refined the areas of interest (ROIs, see **Figure 5.1c** and **d** for raw and processed images side by side). The pipeline further subtracted background noise, corrected for movement and performed iterative dimensionality reduction to extract calcium traces from each ROI that act as an indicator of neuronal activity.

We aligned the data to the behavioural data using the shared TTL pulse to accurately combine the two data streams, and then matched each behavioural event to a corresponding time point in the calcium data.

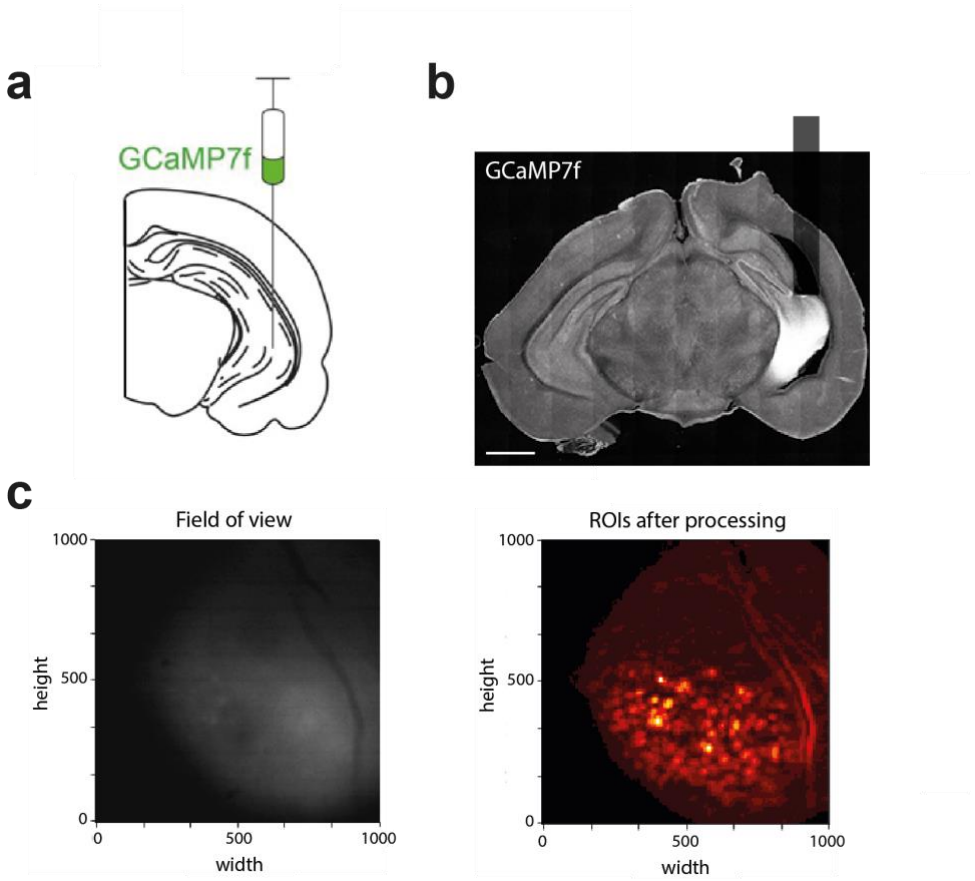

Figure 5.1: Calcium imaging in vCA1 through GRIN lens. a. Schematic of stereotaxic injection to express GCaMP7f unilaterally in ventral CA1. **b.** Example coronal image showing GFP signal from GCaMP7f and position of the GRIN lens for miniscope imaging. Scale bar = 1mm. **c.** Example field of view through the GRIN lens with V4 miniscope **d.** Identified regions of interest (ROIs, maximal projection of all neurons identified during session)

5.2.1 Selectivity of individual neurons in vCA1

Using the single cell calcium data, we first wanted to test whether neurons in CA1 showed differential activity related to task events such as odour cues, cue positions, or trial outcomes. To quantify the extent of neural tuning, we calculated a selectivity index (SI) as the difference in average neural activity over a time window between all trials in which the event is present and the activity in trials which is it not.

For example, to test whether a given neuron is selective for Odour A in the first position, we would take into the account the activity of that neuron in that specific time window (i.e. during the first cue) in all trials. The SI was calculated as the average activity during that time window in all trials in which A was the first cue, subtracted from the average activity in all other trials. To be able to compare SI values across neurons with different baseline activity, we expressed selectivity as a sigma value (σ), defined as the distance of the true SI value from the mean of a distribution of SI values derived from shuffled data from the same cell.

Using these σ values, we identified neurons that showed a strong selectivity (< 3) to either a specific odour, an odour pair, or an outcome. To quantify these different selectivities, we calculated SIs for each neuron at four different time windows within each trial: during the first cue, the delay, the second cue, or in the response window before reward delivery (**Figure 5.2a**).

In total, we found that 18% of neurons recorded in all mice and sessions (43 out of 239 neurons from 3 mice and 7 sessions) showed selective activity to trial events in at least one of the examined time windows using our conservative criteria (**Figure 5.2b**).

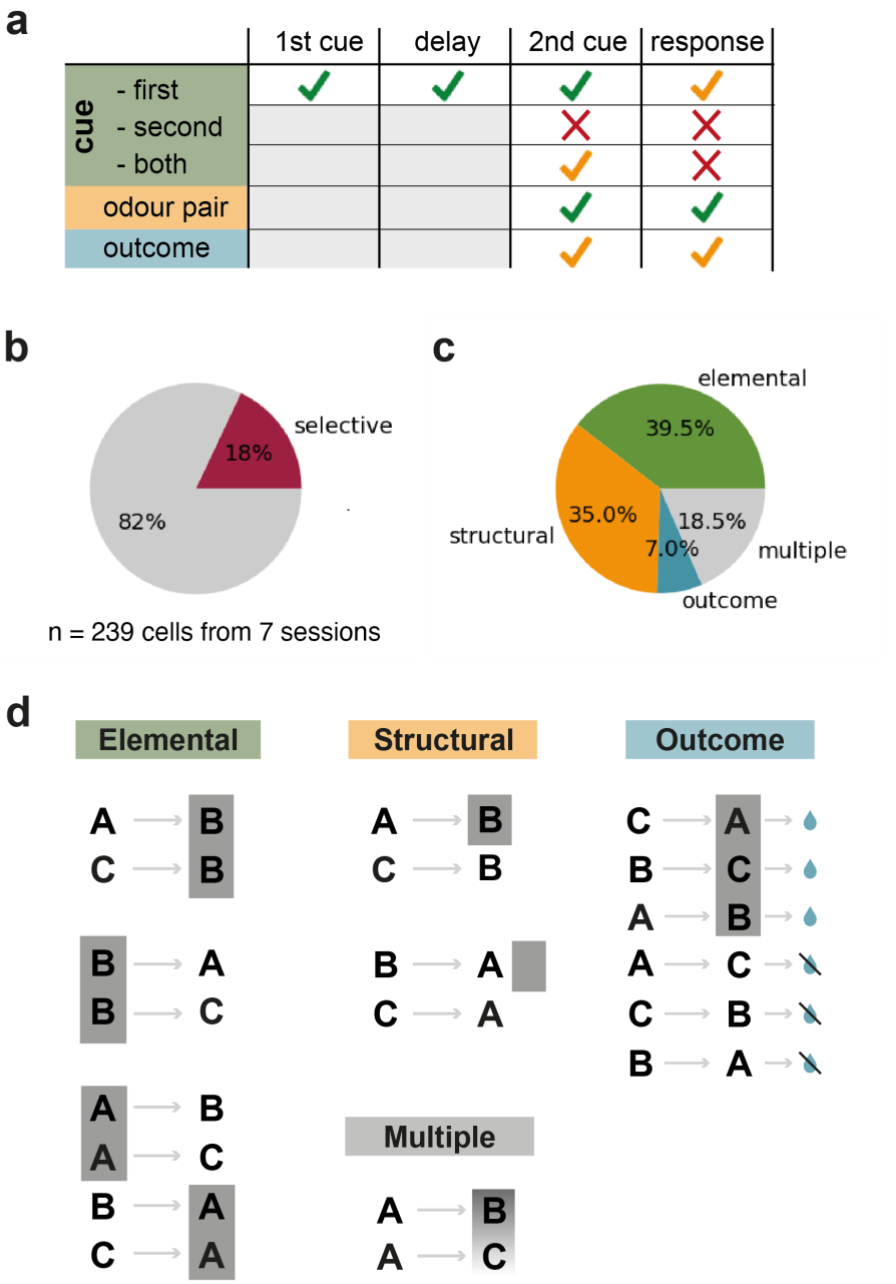

Figure 5.2: Individual vCA1 neurons selectively respond to structural task events.
a. We quantified the selectivity of neurons to different task events (cues, odour pairs, outcomes) at different time points within the trial (first cue, delay, second cue, response window), and found neurons for most but not all combinations of time bin and task event. **b**. Proportion of neurons selective for task events within the total population of recorded cells. **c**. Proportion of neurons within the population of selective cells showing elemental, structural or outcome selectivity, or a combination of those. **d.** Schematic of response types of different patterns. Elemental: response depends on only one odour. Structural: response depends on odour pair. Outcome: response depends on expected trial outcome. Multiple: respond to multiple trial types, but stronger to some than to others.

We then grouped these selective neurons according to the type of selectivity: “elemental” selectivity patterns were those responding to only one odour cue (e.g. “Odour A” selective); “structurally” selective neurons responded to a specific odour only in cases where a specific odour had preceded it (e.g. “A->B” selective); outcome selective neurons were those that responded to all trials that shared the same outcome (i.e. rewarded or not rewarded); lastly, some neurons showed a mixed selectivity for elemental and structural events (e.g. respond weakly to “CA” but strongly to “BA”) (**Figure 5.2d**).

We quantified the occurrence of these types and found that similar proportions of neurons showed an elemental and structural selectivity pattern (39.5% and 35%, respectively). 7% were selective for outcomes, and 18.5% showed mixed selectivity to more than one category (**Figure 5.2c**).

Together, the data shows that about a fifth of the neurons we recorded selectively responded to at least one task event. Most of these neurons were selective to individual odour cues or cue combinations, with only few responding to the expected trial outcome. Interestingly, we found example neurons for almost all combinations of cue and time window, apart from one: we found no example of neurons that were selective to a specific odour in second position (e.g. selective to Odour B, irrespective of whether A or C came before).

To further characterise the responses, we quantified which events were most frequent within each of these categories of selectivity, and we furthermore examined the time courses of these selective responses.

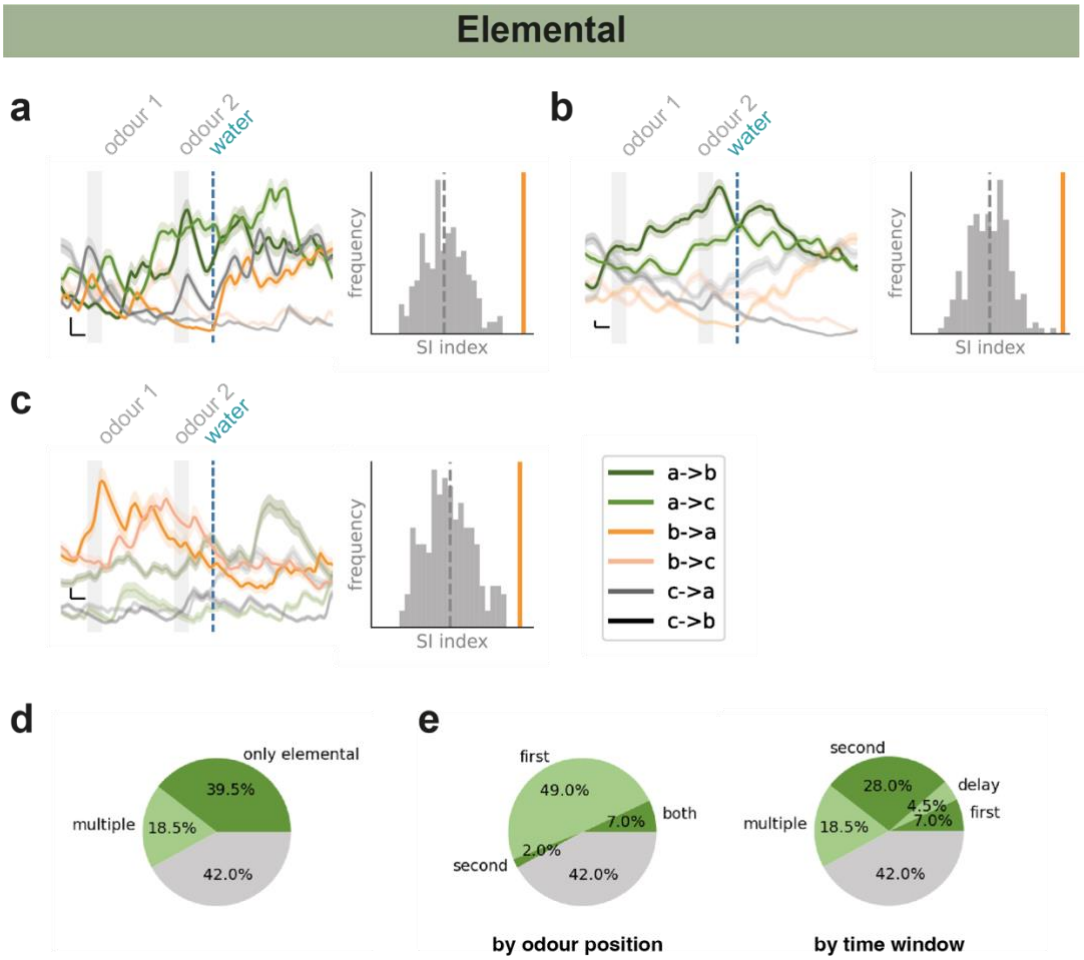

Figure 5.3: Individual vCA1 neurons display selectivity for elemental features of task. a. – c. Average calcium signal from example neuron split across trial types (left) and the selectivity index (SI, orange line) plotted against the distribution of SI indices from shuffled data from the same neuron (grey, right panel). **a**. Example neuron selective for Odour A irrespective of position. **b**. Example neuron preferably responding to Odour A in position 1, but also responsive to Odour B in position 2. **c.** Example neuron selective for Odour B only in position 1. **d.** Proportion of neurons that have any elemental selectivity: a total of 58% of selective neurons respond to specific odours. **e.** Separated by odour position (left): most respond to the first odour cue (49%). Only few neurons respond to a given odour irrespective of position (7%), and even fewer respond to second cues only (2%). Split by time window (right): Most neurons respond with a slow time course, the peak of their selectivity occurring at the second cue (28%) or spanning several time windows (18.5%), with fewer displaying peak selectivity during the first cue or the delay.

5.2.1.1 Elemental responses

First, we looked at those neurons that showed selectivity to "elemental" components (i.e. individual odours). As shown above, 39.5% of selective neurons were exclusively selective for one specific odour (e.g. Odour A). We further found that all of the neurons that showed mixed selectivity (18.5%) had some elemental features, bringing the total proportion of neurons responding to specific odours to 58% (**Figure 5.3d**).

Within this population, some cells responded equally to a specific odour irrespective of which position it appeared in (**Figure 5.3a**, example neurons selective for Odour A as both first and second cue); others responded more strongly to an odour in the first position, but still showed a response to the same odour in the second position (**Figure 5.3b**, example neuron selective for Odour A as first cue); finally, we identified neurons that only responded to an odour when it appeared a specific position (**Figure 5.3c**, example neuron selective for odour B only as first cue).

We next looked at the distribution of these features across the whole population of elemental (i.e. context-independent) neurons and found that most responded to their preferred odour most strongly (or at all) when it appeared as the first cue (49% of all selective neurons). Only a fraction of that number (7%) responded to their preferred odour irrespective of position, and even fewer (2 %) responded to an odour only in the second position (**Figure 5.3e, left**).

Interestingly, despite this bias for first-odour selectivity within the population, only a small number of neurons showed strong selectivity at the time of odour presentation (7%, **Figure 5.3e, right**). Similarly, few neurons (4.5%) were selectively active during the delay. The bulk of the neurons displaying selectivity for specific activity (as illustrated by our examples) either showed the highest difference in activity during the time of the second odour delivery (28%) or a

combination of multiple time windows (18.5%), with their activity most often spanning both delay, second cue delivery and the response window.

Taken together, we found that more than half of the selective neurons (58%) responded selectively to specific odour cues. While some of these neurons responded to their preferred odour irrespective of its position, others fired only for their preferred odour in a specific position, most often in the first position. Interestingly, the majority of elemental responses were delayed with respect to the cue they were selective for, thus maintaining information about an odour after its delivery.

5.2.1.2 Structural responses

We next looked at all cells that were selective for odour pairs, i.e. “structurally” selective neurons. In addition to the 35% of selective neurons that were uniquely responding to specific odour pairs, another 18.5% of neurons responded to odour pairs, but also showed a second response pattern (**Figure 5.4d**).

This pattern was a type of rate-remapping: all of the neurons that showed selective firing in two categories were selective for a specific odour generally, but responded most strongly when this odour appeared as the second odour in a specific pair. For example, a neuron might be elementally responsive to Odour B (regardless of its position in the trial), yet might fire especially strongly to Odour B when it appears after Odour A, therefore having both an elemental and a structural selectivity.

Similarly to the previously described elemental neurons, structurally selective neurons also displayed several distinct temporal dynamics. In some, firing was constricted to the time around the second cue delivery (**Figure 5.4a**) or the outcome (**Figure 5.4c**), while other neurons showed ramping activity from the cue delivery to the outcome (**Figure 5.4b**).

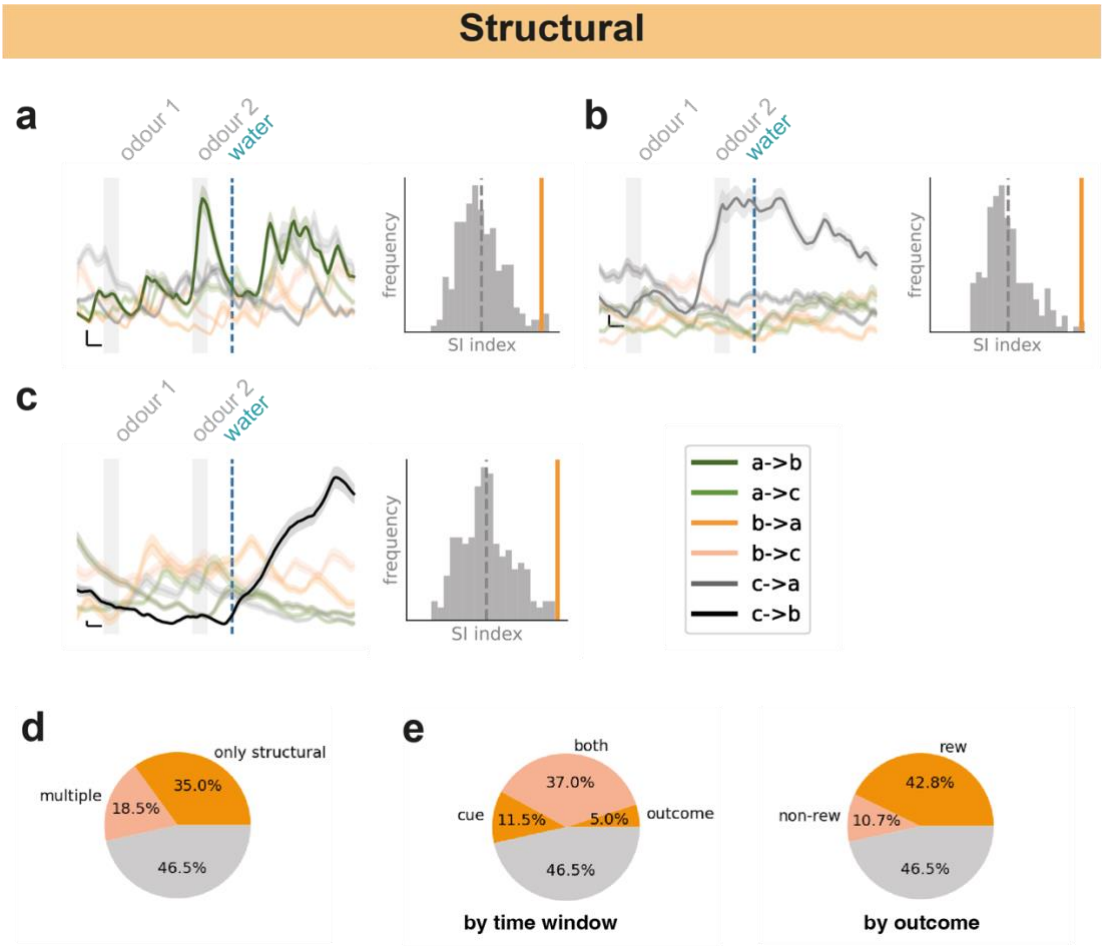

Figure 5.4: Individual vCA1 neurons are selective for specific trial types. a. – c. Average calcium signal from example neuron split across trial types (left) and the selectivity index (SI, orange line) plotted against the distribution of SI indices from shuffled data from the same neuron (grey, right panel). **a**. Example neuron selective for AC (Odour C only when preceded by A) **b**. Example neuron responding to CA (Odour A only when preceded by Odour C) **c.** Example neuron selective for CB (Odour B only when preceded by Odour C) **d.** Proportion of neurons that show structural selectivity: a total of 53.5% of selective neurons respond to odour pairs. **e.** Left: Peak selectivity separated by time window; Most responses have a slow time course, indicated by their peak firing spanning both the cue delivery and the outcome period (37%), although some neurons show temporally constrained firing at the cue or outcome (5% and 11.5%, respectively). Right: Peak selectivity separated by trial outcome: most neurons that show trial-type specific firing are selective for rewarded odour pairs (42.8% of all selective neurons).

Quantifying these different temporal patterns, we found that the latter type in which selective firing occurred throughout cue presentation until reward delivery or omission was by far the most common (37% of selective neurons, **Figure 5.4e**, left) while selective firing at the time of the second cue or the outcome was less common (11% and 5% of selective neurons, respectively). Furthermore, when the population of structurally selective (i.e. context-dependent) neurons was split by whether their preferred trial type was rewarded or unrewarded (**Figure 5.4e**, right), we found that a majority of neurons (80% of structurally selective neurons, 42.8% of all selective neurons) showed selective firing to rewarded pairs, whereas only a minority (20% of structurally selective neurons, 10.7% of all selective neurons) responded preferentially to unrewarded odour pairs.

In summary, we found that roughly half of the selective neurons (53.5%) displayed contextually modulated firing, i.e. structural responses. This type of firing was only found for odours in second position (since the first odour represents the context) and was more likely for rewarded odour pairs. Structurally selective neurons most often showed sustained firing from the delivery of the second cue up to the outcome delivery.

5.2.1.3 Outcome responses

Lastly, a subset of neurons showed selective activity for trial outcome (e.g. all rewarded odour pairs, **Figure 5.5a**). Interestingly, all outcome-selective neurons showed at least some selective firing in the time window between the second cue and the outcome and therefore might not solely be reflections of the shared sensory experience of rewarded (or unrewarded) trials. These outcome-selective neurons were however rare, accounting only for 7% of selective cells (**Figure 5.5b**).

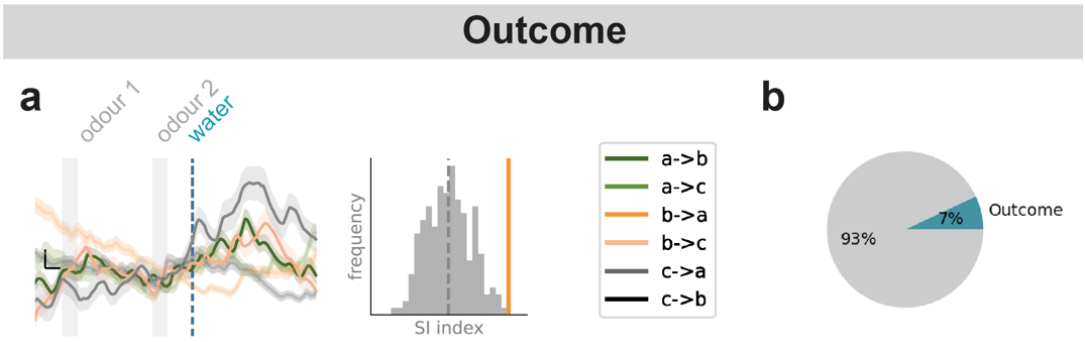

Figure 5.5: Only few vCA1 neurons display activity based on reward contingency **a.** Average calcium signal from an example neuron selective for all rewarded trial types, split across trial types (left) and the selectivity index (SI, orange line) plotted against the distribution of SI indices from shuffled data from the same cell (grey, right panel). **b.** Proportion of selective neurons that show outcome selectivity

5.2.1.4 Summary

In summary, we find that neurons in vCA1 respond to various task events, showing selective activation both to single odours as well as odour combinations, while the expected outcome of the trial is less robustly encoded on the single cell level. The majority of neurons display an activity pattern of broad temporal ramps, often peaking several seconds after the event they are selective to, suggesting that information about past odour cues might be encoded in the population across the delay and even at the time of the next cue

5.2.2 Population representations

Given these findings on the single cell level, we next wanted to look at the encoding of task-relevant information encoded at the population level. To investigate this, we used a series of linear classifiers that utilised the calcium activity of the entire population of recorded neurons to predict trial events. Specifically, we trained an ensemble of linear support vector machines (SVMs) to discriminate between different trial types based on the activity of all recorded neurons the time of the second cue, i.e. the time point that the animal has to make the decision to lick or to withhold licking (**Figure 5.6a**). To do this, we constructed different pseudo-populations by matching trials of the same type from different animals in a randomised way (**see Chapter 2**). We did this by pooling neurons from different mice and sessions while ensuring that trial types are matched in a way that avoids introducing spurious correlations. By randomly pairing trials across neurons (e.g., matching the first AB trial of one neuron with the fifth AB trial of another), we aimed to prevent artificial correlations that could arise simply from the structure of the dataset rather than true neural coding.

The key advantage of this strategy is that it allows to build a large, representative dataset for decoding while ensuring that any structure detected by the SVM is robust and not an artefact of a particular trial alignment. Additionally, because each pseudo-population includes all neurons but with different trial assignments, it automatically gives a measure of how consistent decoding results are across different random samplings.

5.2.2.1 Encoding of trial types

As a first step, we tested whether a linear classifier trained on this dataset could learn to correctly discriminate between trials with different odour pairs. We trained the classifier on 75% of labelled data (i.e. the population vector of the trial together with the true odour pair) and then tested on the remaining 25% of data (**Figure 5.6b**), and found that the decoder could assign the correct trial type

significantly above chance level (**Figure 5.6c**, n = 100, classifier accuracy = 0.31 +/- 0.09; shuffle accuracy = 0.14 +/- 0.08, paired t-test $t_{(99)} = 13.91$, $p = 4.1 \times 10^{-31}$).

We then checked how many neurons were necessary in order to achieve this level of accuracy by training classifiers on subsets of our data with randomly selected neurons. We found that classifiers using the population activity of as low as 20 neurons averaged above-chance performance, and that the accuracy of the classifiers plateaued as they included more than 100 neurons, indicating some redundancy in the network (**Figure 5.6d**).

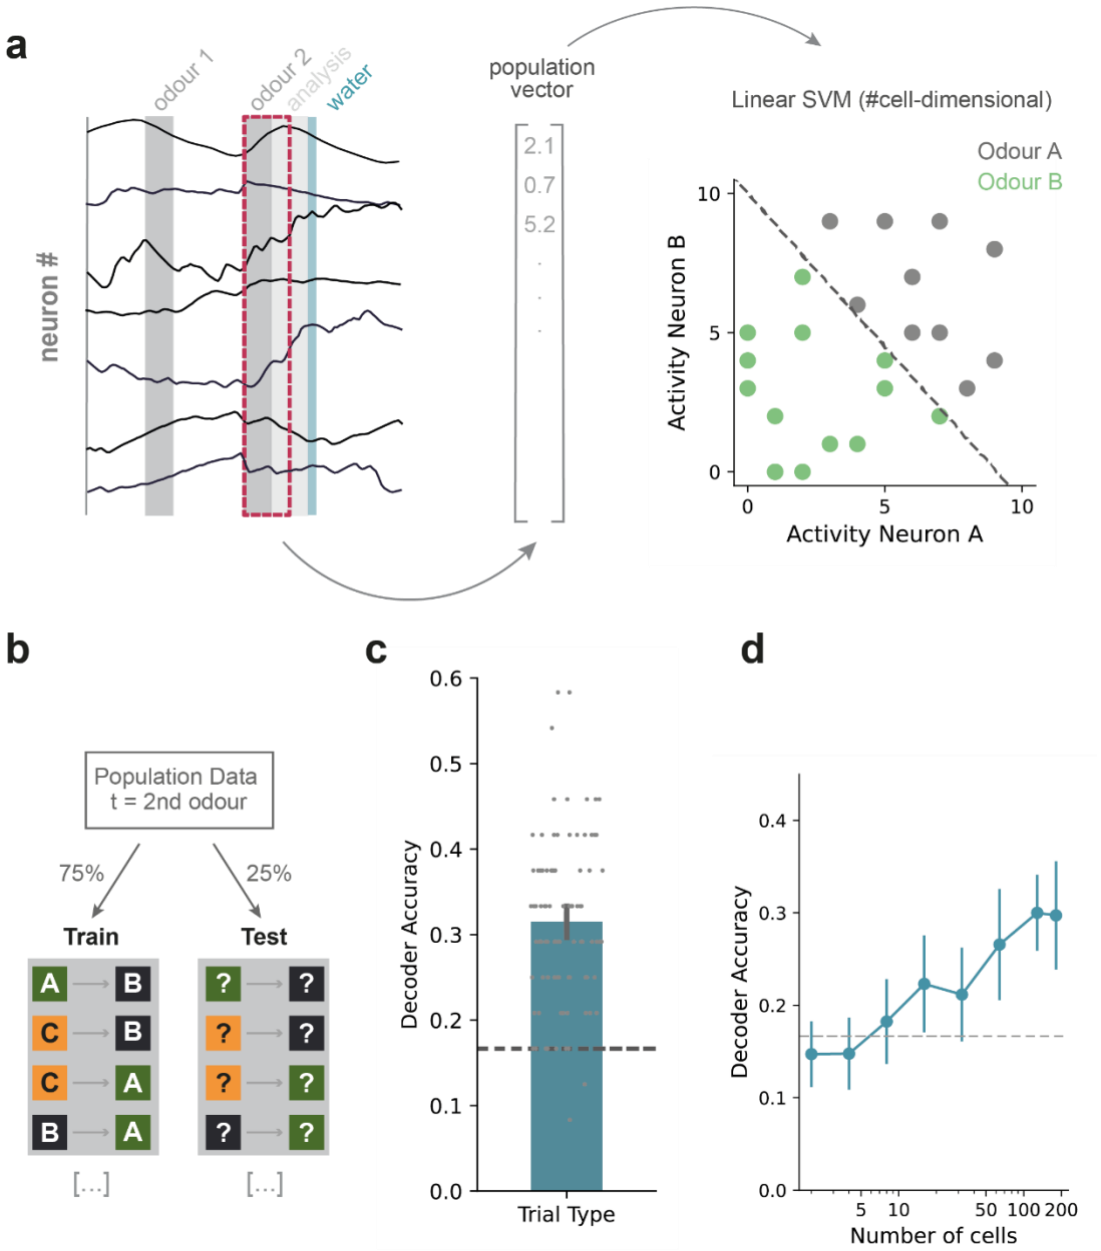

Figure 5.6: vCA1 population activity encodes trial types. a. Schematic of workflow for decoding analysis. The population vector is the average activity of each recorded neuron at a specific time window (in this case: around the presentation of the second cue). Linear classifiers are then trained to separate the population activity based on specific trial events, e.g. the identity of the first cue. **b.** Schematic of split between training data (labelled) and test data (unlabelled). **c.** Accuracy of decoders (n = 100) **d.** Above-chance decoder accuracy requires a dataset including at least 25 neurons. Accuracy plateaus after 100 neurons.

As mentioned above, these classifiers were using the population activity at the time of the second cue delivery. We therefore hypothesised that the identity of the second cue should be strongly represented in the population at this time, both due to the recency as well as due to the importance of this information for decision making.

But to successfully predict outcomes in our task, the identity and order of both cues needs to be retained. We therefore further wanted to check whether the identity of the first cue was encoded in the population activity at this later time point.

5.2.2.2 Encoding of first and second cue

To do this, we used the same dataset as before, combining all recorded neurons into one population. This time, instead of labelling the data according to the odour pair, we only included information about the first or second odour, respectively (**Figure 5.7 a and b**). For example, in order to test for the encoding of the second odour cue, we would train the classifier on 75% of our neural data labelled with the second odour present in each trial, and then ask the classifier to predict the labels for the remaining 25% of the data.

We found that both classifier types (predicting the first and the second odour of seen in the trial, respectively) performed above chance, which we simulated by training decoders on shuffled data. As we expected, decoders trained to identify the second odour cue from the neural data did so with above chance accuracy (**Figure 5.7c**; right, n = 100, classifier accuracy = 0.52 +/- 0.11, shuffle accuracy = 0.33 +/- 0.11, paired t-test $t_{(99)} = 12.20$, $p = 6.5 \times 10^{-26}$). Surprisingly however, the classifiers trained to predict the first odour from the neural activity at the time of the second cue performed as good and even slightly better (**Figure 5.7c**; left, n = 100, classifier accuracy = 0.54 +/- 0.09, shuffle accuracy = 0.31 +/- 0.11, paired t-test $t_{(99)} = 14.05$, $p = 1.4 \times 10^{-31}$), indicating a multi-second maintenance of this information within the network.

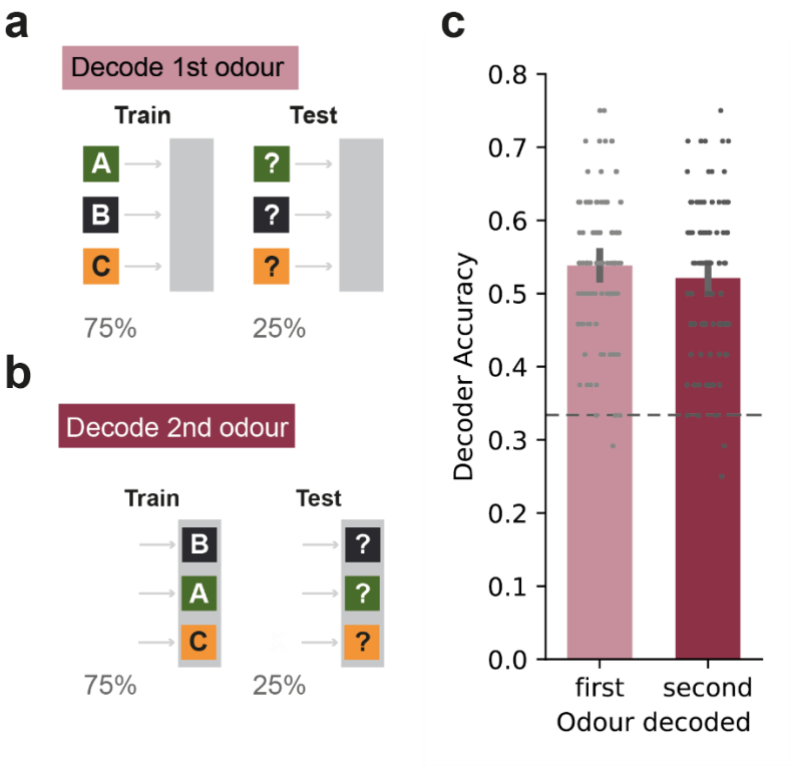

Figure 5.7: Population activity at the time of the second cue encodes both first and second odour. a.-b. Schematic of datasets used to train and test decoders for the first odour of the pair (**a**) and the second odour of the pair (**b**). **c**. Accuracy of decoders trained to classify by first odour (light red, n = 100) or by second odour (dark red, n = 100)

5.2.2.3 Influence of contextual encoding

Given this finding, we next wanted to test to what extent this information might be a result of context-specific encoding. For example, the "first odour" decoders might simply have learnt that the very specific activity pattern related to an AB trial and the pattern associate to an AC trial belong into the same category "Odour A first", without the two patterns sharing much in terms of neural similarity.

To test this hypothesis, we trained linear classifiers on data containing only one of the two trial types with a given attribute, and then tested the accuracy of the predictions in the other. For example, the training set might only contain AB trials, while the test set contained only AC trials, and so on for the other possible first odours.

As expected, we found that the accuracy of these decoders was lower than that of decoders trained on all possible trial types. However, this drop in predictive performance was less than expected and both "first odour" and "second odour" classifiers trained on partial data still correctly classified trials at above chance level (First odour: n = 100, accuracy = 0.43 +/- 0.6 shuffle accuracy = 0.32 +/- 0.07, significantly greater than chance, paired t-test $t_{(99)} = 1.60$, $p = 9.2x10^{-24}$; Second odour: n = 100, accuracy = 0.44 +/- 0.6, shuffle accuracy = 0.33 +/- 0.08, paired t-test $t_{(99)} = 11.13$, $p = 3.3x10^{-22}$).

This suggests that consistent with the types of responses seen at the single cell level, some of the information within the vCA1 network at the time of the second cue is encoding the elemental identity of both first and second cue.

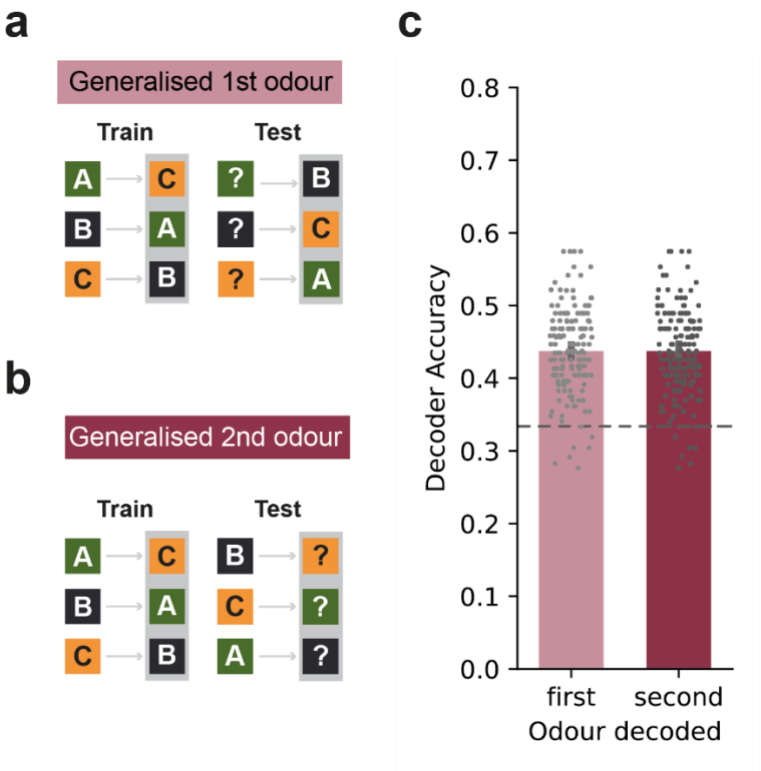

Figure 5.8: Information about odours partially generalises. a.-b. Schematic of datasets used to train and test decoders for the generalisation of the first odour of the pair (**a**) and the second odour of the pair (**b**), respectively. This is achieved by training the decoder on one trial type with the particular characteristic, and then testing on the other (e.g. train decoder on A-B data, then test on A-C data to identify generalised encoding of "A first"). **c**. Accuracy of decoders trained to classify generalised first odour (light red, n = 100) or by second odour (dark red, n = 100)

5.2.2.4 Outcome information in population activity

As shown in Chapter 1, expert mice show through their licking behaviour that they correctly anticipate the outcomes of trials based on the odour pairs, suggesting that information about upcoming rewards is available to them. Looking at individual neurons, we found only few cells showing differential activity depending on trial outcome (see Figure 5.5). We hypothesised however that at the population level, the high number of neurons with structural selectivity could add up to significant encoding of the outcome of a trial.

To investigate this, we first trained decoders to discriminate rewarded from unrewarded trials based on the neural activity at the time of the second cue, using all trial types in both training and testing data (**Figure 5.9a**).

We found that the performance of these decoders was less accurate than the decoders trained to classify neural activity based on odour cues (**Figure 5.9c**; left, n = 100, accuracy = 0.56 +/- 0.11, shuffle = 0.48 +/- 0.09, significantly greater than shuffle: paired t-test, $t_{(99)} = 5.84$, $p = 2.1 \times 10^{-8}$).

Given the low accuracy of outcome decoding, we next wanted to test if any general (i.e. not context-dependent) information about upcoming rewards existed within the population. To do this, we used the same approach as previously shown for odour decoding: we trained the decoders on data containing only one type of rewarded trial, and one type of unrewarded trial, and then tested on a different set of rewarded & unrewarded trials (**Figure 5.9b**). To avoid the decoder using contextual information, we made sure to balance context within each dataset: for example, the training set might contain AB trials (rewarded) and AC trials (unrewarded) which share the first odour cue and therefore should have the same informational content until the delivery of the second cue. The test set would contain all other trial types, balanced such that rewarded and unrewarded trials occur with the same likelihood.

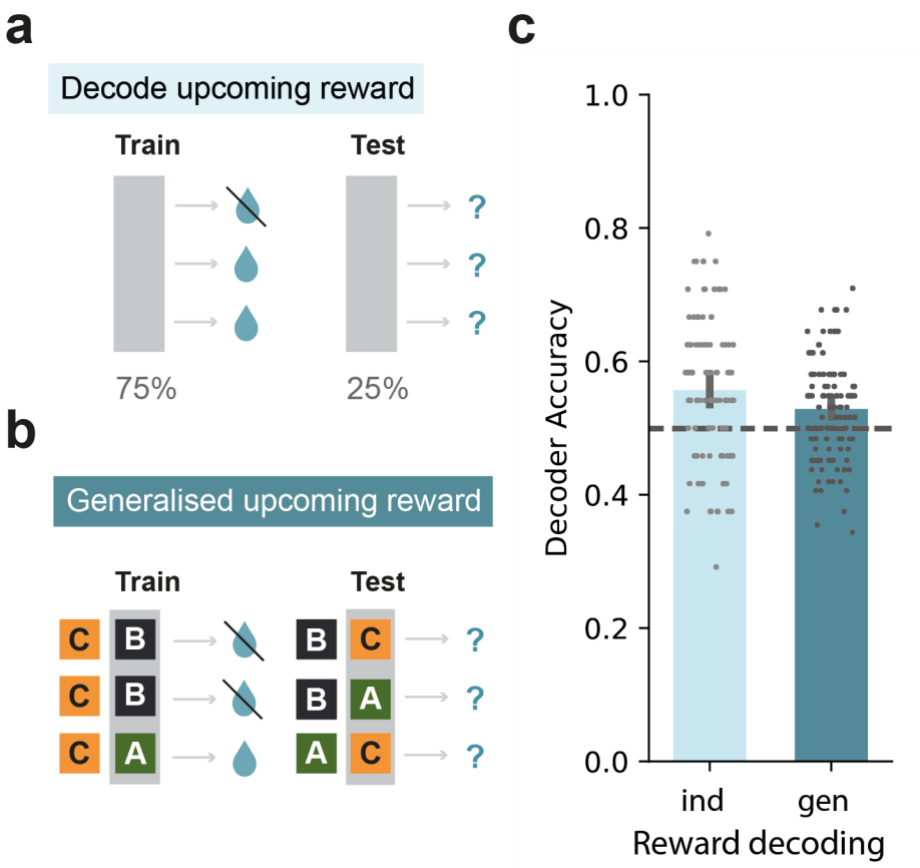

Figure 5.9: Population activity does not reliably encode anticipated outcome **a.** Schematic of data used to train and test classifiers to discriminate between rewarded and unrewarded trials. Data is labelled according to its outcome, and split 75/25 for training and testing. **b.** Schematic of data to test for the generalisation of the outcome across trial types. This is achieved by training the decoder on one type with a particular outcome, and then testing on the other (e.g. train decoder on C-A data, then test on B-C data to identify generalised encoding of "reward"). **c**. Accuracy of decoders trained to classify trials based on outcome. Left: decoder trained on full dataset (all trial types included, n = 100); Right: decoder trained on split data to detect generalised representation of reward (n = 120)

As we expected, the decoding accuracy of these classifiers was lower than the classifiers trained on a balanced dataset, performing only slightly above chance (**Figure 5.9c**, right: n = 120, accuracy = 0.52 +/- 0.07, shuffle = 0.49 +/- 0.03, not significantly greater than shuffle).

In summary, this data shows that expected outcomes are not strongly encoded in the vCA1 population, and that the low level of information about outcomes is due to contextual information rather than a general representation of upcoming reward or non-reward.

5.2.2.5 Comparison of cue and reward encoding in population

To compare the encoding of these different task features (first or second odour cue, odour pairs or outcomes) across decoders, we chose to return to the sigma value (σ) already used for the selectivity analysis of individual cells. This is necessary because "chance" performance corresponds to different values for the different decoders, as chance is a function of the number of classes. For example, a decoder trained on outcomes would be 50% accurate at chance performance since there is only two options (rewarded/non-rewarded), whereas a decoder trained on discriminating between trial types should be 16.6% correct when assigning classes randomly, since there are six different trial types.

We found that sigma values for all decoders trained to discriminate between trials with different task cues (both individual odours as well as odour pairs) were higher than those for outcome decoders (**Figure 5.10;** balanced decoders: ANOVA, $F_{(3)} = 10.84$, $p = 3.0 \times 10^{-6}$; post-hoc Tukey [trial type vs reward] = 3.94, $p = 5.5 \times 10^{-4}$; post-hoc Tukey [cue 1 vs reward] = 5.58, $p = 2.6 \times 10^{-7}$; post-hoc Tukey [cue 2 vs reward] = 3.51, $p = 2.7 \times 10^{-3}$; generalised decoders: ANOVA, $F_{(2)} = 12.85$, $p = 5.0 \times 10^{-6}$; post-hoc Tukey [cue 1 vs reward] = 3.91, $p = 3.3 \times 10^{-4}$; post-hoc Tukey [cue 2 vs reward] = 4.25, $p = 8.4 \times 10^{-5}$). This suggests that the vCA1 population favours encoding of cues and contexts over the encoding of predicted outcomes.

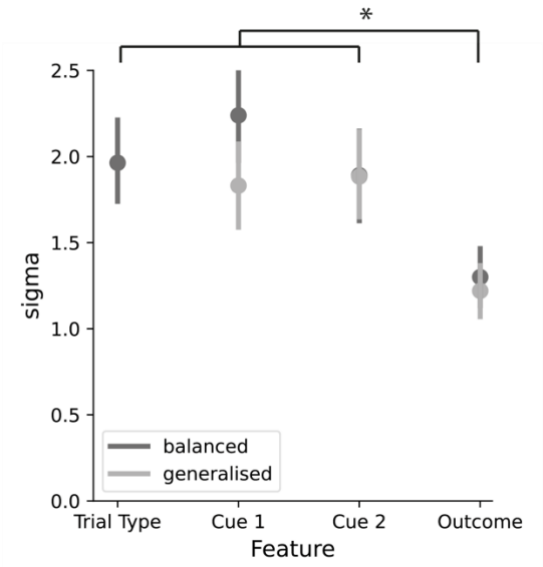

Figure 5.10: vCA1 population encodes cues more robustly than outcomes Sigma values for decoders trained on different task features. Dark grey: decoders trained on balanced data containing all possible trial types. Light grey: decoders trained to discriminate based on generalised information, i.e. trained on one trial type and tested on another. Sigma expresses the distance of the average performance of an ensemble of decoders from chance as a multiple of the ensemble's standard deviation.

Furthermore, we found that the largest difference between decoders trained on balanced data and split data for the same feature was found in the ensembles of "first cue" decoders (σ(Cue 1, balanced) = 2.24 +/- 0.14; σ(Cue 1, generalised) = 1.83 +/- 0.12, paired t-test: T = 2.23, p = 2.6×10^{-2}) which might be a consequence of the decoders classifying trials on the basis of neural data from around the second cue delivery.

5.3 Discussion

In this chapter, I have presented calcium activity recorded from ventral CA1 neurons in mice performing a paired associates odour task with a delay of 5s. I've shown that neurons in vCA1 display selective activity to task events such as cues, cue combinations and outcomes.

I've further shown results from linear classifiers trained on the activity of the population of recorded neurons which suggest that while cue combinations (trial types) as well as individual cues are robustly encoded in the population at the choice point (i.e. the time of the second cue delivery), information about the expected outcome of the trial (reward/no reward) is not significantly encoded in the vCA1 population.

5.3.1 Selectivity for cues and cue combinations in vCA1

Within the population of neurons that show differential firing for trials containing specific cues, we find three different patterns of selective activity.

In line with previous studies (Eichenbaum et al., 1987; Manns et al., 2007; Taxidis et al., 2020), some neurons robustly responded to a specific odour irrespective of its position within the trials (elemental encoding). Interestingly, many of these responses were delayed with regards to the stimulus delivery, thus maintaining information about the cue for several seconds after its delivery.

Additionally, we found neurons that preferentially fired to an odour in a specific position (e.g. "B first"). Most of these neurons fired preferentially in the delay period or around the time of the second cue, recalling the "odour-time" cells described in previous studies (MacDonald et al., 2013; Taxidis et al., 2020).

Together, these neurons with elemental (i.e. context-independent) response patterns made up 39.5% of all cells modulated by task events.

Interestingly, we found very few neurons that fired to an odour in second position in a context-independent manner, i.e. irrespective of which odour was presented before. Instead, the representation of the second odour was strongly contextually modulated: 35% of the neurons displaying selective activation related to task events responded in a contextual manner, e.g. “B after A”. The firing of these neurons therefore is dependent on past experience (the first cue), recalling the splitter cells found in T-maze studies (Ainge et al., 2007; Frank et al., 2000; Wood et al., 2000).

Taken together, the cue-related selective activity in CA1 contains all information necessary to solve our structural learning task: an elemental representation of individual odour cues, maintained across the delay, which can then be combined with the sensory inputs related to the second cue into a structural (context-dependent) response that could then be associated to an outcome.

In future work, it would be interesting to see whether the responses of neurons with task-relevant selectivity can be used as predictors for mouse behaviour, i.e. whether the response of the neurons is more pronounced on trials where the animal responds correctly. The granularity of this analysis would be further aided by recording licking as a continuous rate rather than a binary event, as discussed in **Chapter 3.3.2**.

Analysing lick rate as a continuous measure would allow us to disentangle its potential role as both a motor output and a proxy for decision confidence. If neural activity correlates with licking purely as a motor behaviour, this would likely result in a general alignment between task-relevant firing and licking irrespective of trial outcome. However, if lick rate also reflects the animal’s level of certainty, we might observe stronger neuronal responses on trials where mice lick more vigorously before correct choices. Investigating these relationships in

detail could provide further insight into how hippocampal activity contributes to both decision-making and behavioural execution.

These future additions aside, our finding that single cells in CA1 encode all information necessary for solving our task is further corroborated by the analysis of the activity of the entire population of recorded neurons at the time of the second cue delivery. Using linear classifiers, we found that the identity of the first odour is equally well encoded in the vCA1 population at the time of the second odour as the second odour itself, despite the difference in recency of the actual sensory cue. This strongly suggests that vCA1 neurons maintain information about the first odour across the delay. Our data does not have sufficient granularity to conclusively state whether the neurons firing across the delay support either a “constant firing” type or “population tiling the delay” type of information maintenance. It is furthermore possible that this delay activity is supported by other brain regions, such as prefrontal areas (as reported in (Spellman et al., 2015)).

Interestingly, we found that splitting our training and testing data to test the level of generalisation (i.e. context-independent activity) within the population, we found that at the time of the second cue, both the identity of the first as well as the second odour can be decoded from the population activity. However, the accuracy of this decoding is lower than that of decoder trained on the full data, indicating that some information is lost when the decoder does not have access to contextual encoding. This effect is especially strong for the decoding of the first cue, indicating that at the time of the second cue, a larger part of the information about the preceding cue encoded in the vCA1 population is contextual. This is consistent with our findings on the single cell level: neurons with a structural selectivity (e.g. “A->B”) contain information about Odour A as the first cue, but this representation is contextual and thus not available to the decoders trained on split data.

This finding might indicate that representations of cues in vCA1 become more contextually modulated as the time passes, or alternatively could be seen as evidence that neural activity in vCA1 is shaped by task demand (since a contextual representation of the first odour is necessary to solve the task).

In summary, data from individual neurons as well as the population activity as a whole strongly supports the idea that vCA1 encodes task-relevant features, in an elemental (i.e. context-independent) which is then supplemented by structural (i.e. context-dependent) activity, the latter of which contains latent information about the structure of the task.

An interesting follow-up analysis to link findings on the single-cell level to the population decoders would be to test decoders using only subsets of neurons, specifically those that were previously identified as firing selectively for specific task events, or conversely those that did not show any selectivity in the single cell analyses. If the decoding accuracy remains high when using only the selective neurons, this would suggest that only a subset of cells carry task-relevant information. If, on the other hand, the decoding accuracy remains high even when using only non-selective neurons, that would suggest that the single-cell selectivity measures are too conservative, or that weakly tuned neurons still contribute in a mixed-effects way. Taken together, these two types of decoders could yield deeper insight into how task-relevant information is distributed across the population.

5.3.2 Utility of task-dependent activity in vCA1

It is important to point out that in light of the results of our optogenetic inactivation experiments, activity in vCA1 is likely not required for performance of our task at this 5s delay. We performed these experiments at 5s to ensure imaging quality (implant stability is negative affected by weeks-long training) as well as to maximise statistical power (mice on average perform fewer trials per

session in paradigms with long delays). Thus, while the task-dependent activity in vCA1 recorded in the 5s task reflects all features to solve our task, it is probably not used in this way, at least at this delay.

This is in line with the hippocampus literature, which has shown encoding of task-relevant features in tasks that don't require hippocampus. One such example is the delayed match-to-sample task that was used to identify time cells (MacDonald et al., 2013), another the alternating T-maze task that underlies much of the research into splitter cells (Ainge, Tamosiunaite, et al., 2007; Frank et al., 2000).

We have started imaging mice performing the task at longer delays, but unfortunately due to the increased length of trials and lower number of mice, we have not yet reached sufficient numbers of sessions and cells to allow for statistically meaningful analysis or conclusions.

However, a first look at these recorded neurons from the 30s task suggests that the results presented in this chapter might generalise to the way the task is represented even at time scales where vCA1 is essential to task performance (**Figure 5.10**). We found anecdotal evidence of neurons displaying elemental (**Figure 5.10a**) as well as structural firing patterns (**Figure 5.10b**). We did however also find cells that, over the longer time scale of trials, displayed more complex firing patterns that included periods of time with different selectivities (**Figure 5.10c**). More data is needed to see whether these examples are representative of the vCA1 neurons more generally, and whether the same features are encoded in the population as a whole.

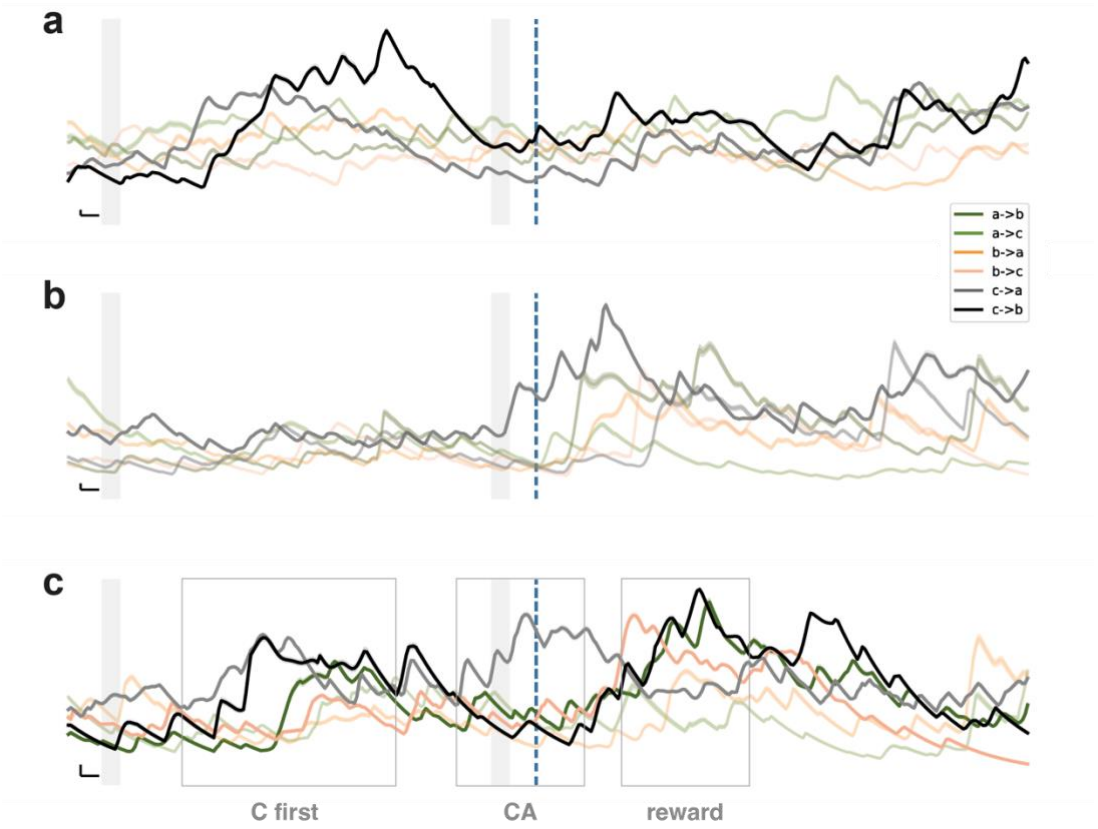

Figure 5.11: First recordings in mice performing the task at longer delays corroborate findings from 5s data
a. Calcium activity recorded from an example neuron displaying elemental selectivity (selective firing in the delay period on trials where C was the first odour) **b**. Example neuron exhibiting structural selectivity (selective firing in the response window only on CA trials). **c.** Example neuron showing a more complex firing pattern that has elements of both elemental, structural and outcome-related firing.

5.3.3 Representation of reward in vCA1

As mentioned in the introduction to this chapter, recent studies have investigated representation of reward in hippocampal populations, putting forward the hypothesis that reward is represented in hippocampus not only as a sensory cue, but a salient feature that might not only reorganise cognitive maps but even be predicted by sequential activity in hippocampal populations (Sosa et al., 2024).

In our data, we identified only few neurons (6% of cells with task-related activity) whose activity was selective for a specific outcome. However, we did find a much higher proportion of “structural” neurons firing selectively for the rewarded

odour pairs than for unrewarded pairs (80% and 20%, respectively), consistent with the previously reported "overrepresentation" of place cells around reward locations (Dupret et al., 2010; Hollup et al., 2001; Jarzebowski et al., 2022). Some of these structural neurons showed activity that spanned the time from cue to reward. They did however not consistently display any "ramp"-like activity as it is seen e.g. in striatal neurons (Howe et al., 2013). Given the temporal dynamics of calcium indicators such as GCaMP7f, this could be a consequence of calcium dynamics. Therefore, a conclusive statement of whether or not neurons in vCA1 display ramping activity might be easier to achieve with more temporally precise recording techniques such as electrophysiology.

Despite the bias towards rewarded pairs in the recorded neurons and the maintenance of trial-type-specific activity throughout the response window up to the outcome, linear SVM decoders trained on the population activity at the time of the second cue (choice point) were unable to reliably predict whether a trial was rewarded.

This suggests that while the encoding of the odour pairs is influenced by the outcome, there is no general representation of value (upcoming reward) in vCA1.

5.3.4 Methodological considerations

In the analysis to identify neurons displaying task-related activity, we focused on four separate time windows that were relevant to our task: the first cue delivery, the delay, the second cue deliver and the time around the outcome.

This method might however lead us to overlook neurons with differential activity around the borders of these events, and to overrepresent neurons whose peak selectivity aligns with these experimenter-chosen time bins. A more objective way of identifying neurons with selectivity towards task events would be to

calculate the selectivity across a sliding time window, thus investigating the occurrence of selective neurons in a more unbiased way.

Furthermore, since selectivity was assessed across multiple neurons, time points, and conditions, multiple comparisons should have been considered in order to estimate the likely proportion of false positives. To this end, the shuffled control data already used to provide a baseline for each individual test could be used: By applying the same selectivity criteria to the shuffled dataset with randomised labels, we could get an estimate of how many neurons would be expected to appear selective in the absence of meaningful signal. If the proportion of selective neurons in the real data significantly exceeds that observed in the shuffled data, this suggests that selectivity is not simply a consequence of random fluctuations. A statistical test such as permutation testing could be used to quantify whether the observed difference in occurrence might arise by chance.

To formally correct for multiple comparisons, different statistical approaches can be considered. One method is false discovery rate (FDR) correction, such as the Benjamini-Hochberg procedure, which controls the proportion of false positives while maintaining statistical power. A more conservative alternative is Bonferroni correction, which adjusts the significance threshold by the number of comparisons, though this can be overly strict in large datasets. Alternatively, permutation-based significance testing, where the analysis is repeated on shuffled data across many iterations, allows for an empirically determined significance threshold. Applying such corrections would ensure that the neuronal selectivity we report is not driven by statistical artefacts but instead reflects genuine task-related activity.

A further concern is the anecdotal observation that the average calcium activity of some cells seems to rise before the event to which the cell is putatively selective for. Since often these events are cues and are thus not predictable, this

implies that either there is a systematic artefact that misaligns the calcium data with respect to the behavioural data, or else suggests that the calcium activity is not in fact a response to the cue at all.

As a first pass, we could use a raster plot of all trials instead of the averaged calcium activity shown in Figures 5.3-5.5, to assess whether these seemingly pre-emptive responses are a consequence of averaging over trials with very varied peaks, or whether it is a consistent phenomenon across many trials.

As a next step, it would be helpful to investigate the potential sources of a systematic time-shift in the data. Looking at our data collection and analysis, there are three points at which a small temporal misalignment might be introduced. Firstly, we record behavioural events and calcium activity on different computers, which means that we have to synchronise their respective clock times. To do this, we send a TTL pulse from the Arduino that delivers the behavioural cues to the miniscope DAQ board. This signal is sent at the time of the first cue and triggers the start of the miniscope recording. However, while the initial TTL pulse ensures a common starting point, any drift in clock speed between the two computers could cause increasing misalignment over time. To assess whether this is a concern, we could embed additional TTL pulses, for example et the end of the session which can be used to apply post-hoc temporal realignment.

A second potential source of the temporal shift might be the median filtering step in the Minian pipeline, which, while effective at reducing salt-and-pepper noise, can also introduce temporal smoothing. If the filter window is too large, it may blur fluorescence transients and shift activity slightly backward in time, creating the appearance of pre-event responses. A simple way to assess this would be to compare raw fluorescence traces with the filtered data and determine whether the effect persists when using a smaller filter window or omitting the filter altogether.

Another possible explanation lies in the temporal deconvolution step of the CNMF algorithm. Since CNMF estimates putative spike trains from calcium fluorescence signals based on an assumed decay model, inaccuracies in this model—such as overly strong regularisation—could lead to small shifts in detected transients. This could cause activity to appear earlier than it truly occurs. Furthermore, because CNMF updates spatial and temporal components iteratively, there is a risk of signal leakage between neighbouring neurons, potentially distorting the timing of extracted activity. To test for this, comparing raw fluorescence traces to CNMF-extracted signals and adjusting temporal constraints within the algorithm would help clarify whether the effect arises from the analysis pipeline rather than the underlying neural activity.

Another way to test whether the temporal shift is a product of the processing of the data with the Minian pipeline would be to use artificial data in the form of a step function. If the filtering steps or the CNMF causes a systematic shift, the output step function would be shifted with respect to the original one.

Additionally, there are methodological consideration when it comes to the SVM decoders used for population analysis. We only trained the linear decoders on the population activity around the second odour. We therefore cannot make any claims about the amount of information encoded in the population activity at any other time within the trial. Importantly, we postulate that given the presence of outcome-selective neurons in our own data and the findings about reward-neurons in the literature, we likely would be able to robustly decode the outcome of the trial if we trained the decoder with the population activity around that time. Our findings thus make no claim about reward representation, but rather about the representation of predicted outcome (i.e. the expectation of reward in the seconds before).

Another consideration is the use of pseudo-populations. While this method is a well-established approach in population decoding analyses (ref, ref), particularly when simultaneous recordings from large populations are not feasible, it has some caveats. For example, if neurons exhibit strong within-session correlations, e.g. slow drifts in activity, disrupting their natural trial-to-trial relationships might slightly underestimate the true decoding performance when compared to a real, simultaneously recorded population. Furthermore, this type of analysis does not allow to investigate differences in encoding between animals. However, given the low count of cells per animal in our recordings, it is not possible to conduct a population-level analysis without combining data across sessions and animals, since a minimum number of 100 cells is required to make these analyses yield meaningful results.

6 Discussion

6.1 Summary

In this thesis, I have investigated what role hippocampal CA1 might play in a structural learning task.

In **Chapter 3**, I described a novel paradigm: an olfactory paired-associates task in which mice are presented with a sequence of two odours separated by a delay and have to learn that certain odour pairs are indicative of reward, while others are unrewarded. Importantly, the contingencies of reward are counterbalanced such that each cue alone provides no information about upcoming reward - both configurally (A after B is rewarded, while A after C is not rewarded), and temporally (A after B is rewarded, while B after A is not rewarded), such that mice can only solve the task by combining the two cues across the delay.

Mice learned to perform this task stably within the space of 8 days of training, and both control experiments as well as the results of logistic regression analysis indicate that mice indeed use a structural learning approach to solve the task.

In **Chapter 4**, I tested how flexibly mice can use these previously learnt structures. I showed that mice were able to maintain high accuracy behaviour when the delay times between the odours were altered, even across very long delays of up to each 30s between each odour cue. Furthermore, I found that mice could rapidly recover performance after introduction of a novel odour into the task, suggesting that previous experience of a given task structure allows faster incorporation of new cue combinations that share the same structure.

Lastly, I investigated whether hippocampus is required for this proficient performance on our task. Using bilateral optogenetic inhibition of vCA1, I found

that inactivation markedly impaired task performance only in paradigms with a 30s delay between the cues.

In **Chapter 5**, I presented calcium activity recorded from vCA1 neurons in well-trained mice performing our task with a delay of 5s between cues. I found robust encoding of task variables, with individual neurons displaying selective activity to events such as individual odours, cue combinations as well as outcomes. This encoding wasn't time-locked to the time of the actual event, but instead many neurons maintained the identity of the first odour cue across the delay or recalled the first odour at the time of the second cue. Furthermore, we found neurons that fired consistent with a context-specific representation, e.g. neurons that fired to odour A, but only if it was preceded by odour B.

I further presented results from linear classifiers trained on the activity of the entire population of recorded neurons which indicate that while cue combinations (trial types) as well as individual cues are robustly encoded in the population at the choice point (i.e. the time of the second cue delivery), information about the expected outcome of the trial (reward/no reward) is not significantly encoded in vCA1.

6.2 Representations of cues and structure in vCA1

When recording calcium activity in neurons in vCA1 from mice performing our task, we found that 18% of neurons showed selective firing to task events, a proportion is in line with findings from studies using similar tasks (MacDonald et al., 2013; Taxidis et al., 2020). Of these selective neurons, 93% were found to be selective for parts of the cue structure, either by showing selectivity for individual odours, odour pairs or a combination of both, indicating that task-relevant information is strongly encoded in the population.

Consistent with previous research into the representation of odours in hippocampus, we found neurons that responded to a given odour irrespective of its position within the trial (Eichenbaum et al., 1987; Manns et al., 2007; Taxidis et al., 2020). This purely elemental type of response however accounted only for 7% of selective neurons.

A larger proportion (49%) of selective cells showed selectivity for specific odours in first position (e.g. “odour A as first cue”). The selective firing of these neurons was most often time-shifted such that the peak of mean activity centred not around the cue delivery but lay in the delay period or even during the delivery of the next. This finding is consistent with the odour-modulated time cells described both in rats (MacDonald et al., 2013) as well as mice (Taxidis et al., 2020). The timing of these cells however was specific to task structure, i.e. we did not find any time-shifted odour cells tiling the seconds after the second cue. This suggests that these neurons represent information both about recent cues, but also about the task structure in which “first cue” is a meaningful position in time that needs to be tracked in order to optimally predict outcomes.

This position-specific cue selectivity is similar to the multiplexed representations that have been demonstrated in studies using spatial navigation tasks, such as landmark vector cells or odour-in-place cells (Deshmukh & Knierim, 2013), except that in our task, the neurons represent a specific cue at a specific position in time rather than in space.

In our task design, the first and second cue carry different meanings: while the first cue can be seen as setting the context (i.e. narrowing the space of possible sequences down to two, only one of which is rewarded), the second cue represents the actual choice point (i.e. the second odour, evaluated within the context of the first, enables the mouse to predict the outcome).

In line with this, we found that the representation of the second cue is strongly contextual: 35% of selective neurons fired differentially in response to one specific odour pair, e.g. “B after A”. The activity of these neurons therefore contains information about past and present: their firing integrates the previous cue with the current read out and can therefore be used to inform future action. This is consistent with the splitter cells reported in T-maze experiments that fire differentially based on past routes or upcoming choices (Ainge, Tamosiunaite, et al., 2007; Wood et al., 2000).

Taken together, within the population of selective neurons in vCA1, we found an elemental representation of odour identity, an encoding of odour-in-position maintained across the delay and context-specific conditional encoding of the second cue. To our knowledge, this is the first time that neurons maintaining a cue across a delay and cells integrating this maintained past with a current cue have been demonstrated in the same task.

Furthermore, we found that the peak activity of the population of selective neurons was biased towards the time of the second cue, making the neural representation in vCA1 at the choice point well suited to support solving the task. We corroborated this hypothesis with decoders trained on the entire population of recorded neurons irrespective of their selectivity and found that, in line with a strong encoding of both context and sensory inputs, the decoders trained on neural activity around the time of the second cue delivery were able to correctly predict the identity of both the second as well as the first cue.

Overall, these findings underscore the importance of hippocampus in the processing of temporal sequences and the encoding of sensory events within task-specific contexts.

These findings raise the question of whether selective neurons constitute a distinct functional cell type within vCA1 or whether their response patterns

simply emerge from different input configurations. Given that CA1 pyramidal neurons receive afferent inputs both at their apical dendrites from entorhinal cortex and closer to the soma from CA3 (Takahashi & Magee, 2009), it is plausible that their selectivity for task features arises from the integration of these inputs rather than from inherent differences in cell identity.

The hippocampus is known to flexibly encode information depending on behavioural context (McKenzie et al., 2016), and task-selective responses may thus reflect dynamic input-driven modulations rather than distinct neuronal subpopulations. This aligns with work showing that the same hippocampal neurons can exhibit different firing properties across tasks (Komorowski et al., 2009), reinforcing the idea that these representations are an emergent feature of network activity rather than being hardwired into specialised cell types.

More broadly, our findings on the single-cell level contribute to the view of the hippocampus as a region that transiently encodes and maintains structured information about the past to facilitate learning. Unlike brain areas that encode predefined task-relevant features, the hippocampus is thought to act as an association machine that retains recent experiences until their relevance becomes clear, allowing post-hoc learning (Wallenstein et al., 1998; Ranganath, 2010; Schapiro et al., 2017).

In our task, vCA1 neurons encode the relationship between sequentially presented odour cues, enabling mice to use past information at a later decision point. This aligns with evidence that hippocampal representations are shaped by experience-dependent plasticity and play a key role in learning associations that span time (MacDonald et al., 2013; McKenzie et al., 2014). Crucially, since we find no reliable encoding of outcomes, our results highlight that this associative function of hippocampus is not performed in isolation but likely emerges from interactions with regions such as prefrontal cortex and nucleus accumbens, which are implicated in decision-making and reinforcement learning (Durstewitz

et al., 2010; Preston & Eichenbaum, 2013; Gerraty et al., 2014; Liu et al., 2014; LeGates et al., 2018; Zhou et al., 2019).

Taken together, these findings reinforce the notion that the hippocampus is not dedicated to a single cognitive function but rather serves as a flexible, experience-dependent system for structuring information over time.

6.3 Encoding of outcomes in vCA1

In our task, six different trial types can lead to two different outcomes: reward, or no reward. Mice have learnt these associations, as is evident by the high accuracy of their predictive licking.

Yet, within our neural data, we did not find significant encoding of these shared outcomes at the time of this anticipatory behaviour. While we found a small number of individual neurons that displayed selective firing for all trial types with a shared outcome, their mean firing generally peaked around the delivery of the reward, therefore possibly encoding the shared sensory experience of reward rather than an expectation of the outcome. Furthermore, decoders trained on the population data at the time of the second cue performed only slightly above chance when predicting whether a given trial was rewarded or not.

However, within the population of neurons that were selective for specific trial types, 80% of neurons responded to rewarded odour pairs. This is in line with previous research demonstrating the clustering of place fields near reward locations (Dupret et al., 2010; Hollup et al., 2001; Jarzebowski et al., 2022).

Taken together, our data therefore supports the notion that hippocampus preferentially encodes cues and contexts that are related to rewarding outcomes. Whether this is due to the value of these outcomes, to their saliency,

or their association to an action (as opposed to unrewarded trials, which require no further action or attention) remains open. We find only weak evidence of encoding of generalised value (i.e. context-independent shared outcomes) in both the individual neural selectivity as well as in the population as a whole.

This is in contrast to recent studies published on outcome encoding in hippocampus more broadly (Gauthier & Tank, 2018), and vCA1 specifically (Biane et al., 2023), that have found generalised representations of reward that weren't restricted to reward consumption and, in the latter case, even were in anticipation of the predicted outcome. A possible explanation for this discrepancy might lie in the difference in task complexity: in the aforementioned studies, the cues signalling reward were unambiguous. Thus, in contrast to our task, there was no need for contextual differentiation of the neural representation, and therefore the outcomes might be encoded in the same cognitive map, thus looking like a generalised representation of reward, whereas in our task due to the need to disambiguate "B after A" from "B after C" this behavioural requirement might have caused neural representations to diverge.

Recent computational work has indeed suggested that in complex environments with a high number of possible futures, it might be non-optimal to encode possible future trajectories in the same representation as the associated outcomes (Stachenfeld et al., 2017; Whittington et al., 2020). In these frameworks, representations of states and the transitions between them are learnt separately from the values associated with those states, leading thus to a cognitive map of state space that does not encode generalised value. It is suggested that hippocampus might be the neural substrate for this map of state space (Gershman, 2018).

The data I've presented in this thesis supports this view of hippocampal function, suggesting that in tasks that require more complex associations than direct mapping from a cue to an outcome, vCA1 neurons encode information about

cues and contexts with a bias towards rewarded cue pairs, but do not encode generalised representations of expected outcome.

6.4 Future research directions

Investigating the neural mechanisms of encoding and representing relational structures in a task-specific way is crucial for a deeper understanding of how our brains are able to adapt to the constantly changing requirements of the environment and support flexible adaptive behaviour. In this thesis, I have presented data providing more insight into the role of vCA1 neurons in a structural learning task, specifically investigating how neurons in vCA1 encode relationships between ambiguous cues.

However, many new questions arise from this data. Below, I outline several directions for future research that might answer some of these outstanding questions.

6.4.1 Differential HC involvement depending on delay length

In **Chapter 4**, I presented data showing that optogenetic inactivation of vCA1 only affects behavioural performance of our task in paradigms where the delay exceeds 10s. While this suggests that ventral hippocampus is required to bridge longer delays, it is unclear what the exact role of vCA1 is within this process. To investigate this question in more details, I suggest a two-pronged approach.

In a first step, more temporally defined optogenetic manipulations might be able to refine our understanding of which part of the process vCA1 is performing: by inactivating only during the delivery of the first cue, the delay, or the delivery of the second cue, respectively, it would be possible to distinguish between a role for vCA1 primarily in cue representation, memory maintenance across the delay, or integration of past and present cues.

Given our calcium imaging data as well as studies proposing that CA1 fulfils the role of a mismatch detector (Axmacher et al., 2010; Kumaran & Maguire, 2006), I would hypothesise that the last inactivation window (during the second cue delivery) is likely to have the most significant impact on behavioural performance.

In a second step, recording calcium data in the 30s delay paradigm would be a logical next step to investigate the difference in hippocampal involvement between different delay lengths. With this data, it would be possible to compare the neural activity in vCA1 in a 30s paradigm to the data I've shown in this thesis recorded in the 5s task, and determining if there are differences in the encoding of task variables, and how those differences may relate to the difference in hippocampal requirement at these two delays. In light of the vast literature showing that hippocampus encodes rich and context-sensitive representations of task variables in many tasks despite not being required for their performance (Ainge et al., 2007; Biane et al., 2023; Ferbinteanu et al., 2011), it seems likely that the task-specific encoding would not change substantially in the paradigm with a longer delay. This hypothesis is supported by anecdotal data I've presented in **Figure 5.11**.

Lastly, due to the temporal resolution of the calcium indicators used in our recordings and number of recorded neurons per session, our data from the 5s paradigm did not allow us to investigate the encoding of temporal sequence in fine detail. With recording data from the 30s task, it might be possible to compare neural activity to several proposed models for delay activity in hippocampus.

One proposed mechanism for maintaining the identity of a past odour cue over a long time interval is through persistent neural activity across the delay period (**Fig 6.1B**). This type of neural activity has been most thoroughly described in prefrontal areas for example in working memory studies (Fuster & Alexander,

1971; Erlich et al., 2011; Liu et al., 2014; Yang et al., 2014), but examples of this type of delay firing have also been shown in hippocampus and the adjoining subiculum (Deadwyler & Hampson, 2006; Hampson et al., 2011). In this framework, a different set of neurons specific to each odour would become active after the first cue presentation and maintain activity until the second cue, thus allowing an association between the two odours across the delay period.

Another proposed mechanism is that a population of neurons fire in sequence, tiling the delay time between first and second cue (**Fig 6.1C**). This type of chaining of neural activity has specifically been observed in tasks that like ours require the mouse to remember an odour cue throughout a delay period to compare it to a subsequent cue (Macdonald et al., 2013; Taxidis et al., 2020).
such as persistent neural activity of specific neurons, temporal sequences of ensembles of neurons, or even "silent" memory maintenance mechanisms (**Figure 6.1**)

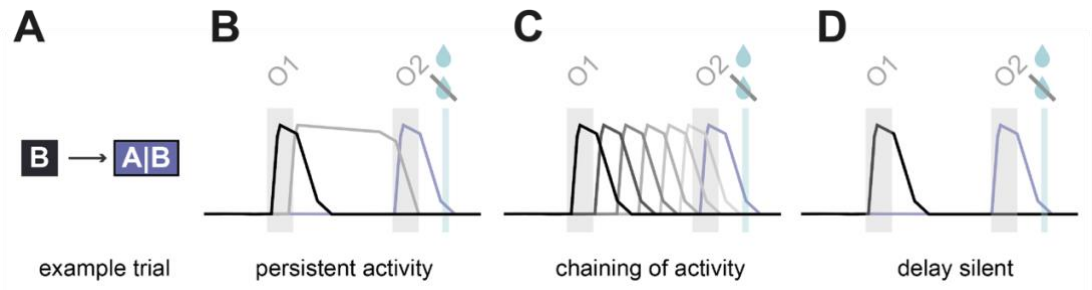

Figure 6.1: Proposed theories of hippocampal activity across delay period.
a. Schematic of example trial **b.** Maintenance of a past odour cue over the delay period through persistent neural activity of single neurons. **c**. Maintenance of a past odour cue by a population of neurons firing in sequence ("tiling"). **d.** "Silent" memory maintenance that is not characterised by specific firing patterns in hippocampus

However, studies showing such activity often use short delays (e.g. 2-5s) and even at those shorter delays the amount of information encoded by the neural population decreases over the delay (Taxidis et al., 2020), suggesting that this mechanism may not be well suited to longer delay paradigms.

Other mechanisms that have been put forward for memory maintenance over a delay period are “silent”, i.e. are not characterized by specific firing patterns in hippocampus itself (**Fig 6.1C**). For example, during trace fear conditioning, there is only limited evidence for delay activity in hippocampus. Despite this, population dynamics before and after conditioning differed significantly (Ahmed et al., 2020). This may be a consequence of memory stored in synaptic weights rather than in neuronal firing rates, a mechanism put forward in theoretical work to enable memory maintenance despite sparse firing (Mongillo et al., 2008). Another possible mechanism underlying “silent” delays may be that the memory is maintained outside the hippocampus, or in the activity between specific regions such as for example between hippocampus and prefrontal areas (Spellman et al., 2015).

6.4.2 Role of HC in generalisation

When exchanging one odour from the task with a novel odour cue that mice had no previous experience of, we found that they adapted to this change rapidly, returning to previous levels of behavioural accuracy in a span of days, and sometimes within one session.

This is evidence supporting theoretical accounts that propose that structural learning might underly abstraction and inference (Whittington et al., 2020). To further investigate the way hippocampal activity supports structural learning, recording the calcium activity of neurons in vCA1 is a very exciting prospect. Specifically, it would be very interesting to compare the neural representation of the novel cue to the representations of the familiar cues.

Results from a combined study in humans and mice performing a transitive inference task showed that hippocampus is involved in constructing representations of novel stimuli and that hippocampus preserves the learned temporal statistics in this representation of a new association (Barron et al.,

2020). However, since the individual cues in that task were unambiguous, in contrast to the individual odours in our task that occur in both rewarded and unrewarded contingencies, the question of how a novel cue might be fit into several new associations is still open.

As mentioned in **Chapter 1**, especially the domains of generalisation and categorisation are not uniquely ascribed to hippocampal circuits, but likely also involve frontal cortices. Specifically, hippocampus-PFC interactions are proposed to be essential to memory integration and generalising knowledge between tasks or stimuli (Gerraty et al., 2014; Rygula et al., 2010; Zeithamova et al., 2012). Thus, it is likely that the integration of a novel cue relies not solely on hippocampus, but on its interaction with frontal cortices.

6.4.3 Projection specificity

This aspect of projection-specific information processing raises another interesting research direction: previous work from our lab has shown how distinct, parallel output circuits in ventral hippocampus perform unique roles in behaviour (Wee et al., 2024).

A further exciting research direction would therefore lie in examining the down-stream targets of the hippocampal representations I characterised in this thesis. From the literature, it is suggested that value assignment might be dopamine mediated and therefore take place in the Nucleus Accumbens (NAc) (Howe et al., 2013). Furthermore, manipulations of the vHC-NAc projection have been shown to drive or suppress reward-seeking behaviours (LeGates et al., 2018; Zhou et al., 2019). Following from this, structurally responsive neurons that represent information about a specific contingency (i.e. odour pair) might project preferentially to NAc.

However, both PFC and OFC have also been implicated in credit assignment in goal-directed behaviours (Takahashi et al., 2011). Therefore, the PFC might also be a possible target that might receive and utilise this structural information. Overall, however, most of the studies relating frontal cortices to credit assignment use tasks requiring some level of inference or generalisation (Walton et al., 2010; Wang et al., 2020; R. C. Wilson et al., 2014).

Given this evidence, investigating the HC-PFC projection might be especially interesting in those sessions where animals have to adapt to changes in the paradigm, whether they are changes to the cues or the temporal structure, since in those sessions and surrounding sessions it might be possible to disambiguate signals related to general credit assignment and update signals in response to changes.

6.4.4 Development of hippocampal representation during learning

In this thesis, I've investigated calcium recordings in vCA1 from expert mice performing our task and shown that neurons this region represent task-relevant variables in both elemental and structural ways. From the literature, we know that hippocampal representations are dynamic, with known phenomena such as remapping and representational drift introducing change into hippocampal representations (Ziv et al., 2013, Sanders et al., 2020).

Therefore, future research might follow the calcium activity in hippocampus throughout learning, to investigate both when neurons in hippocampus start to represent certain features and relate this activity to phases of learning, and further to assess the amount of change the representations undergo both during and after learning. While electrophysiological recordings in spatial paradigms have shown the emergence of task-related representations within seconds (Bittner et al., 2017), results from tasks with similar structural demands as ours

have indicated that in these cases, representations emerge over the course of days, correlated to levels of learning (Taxidis et al., 2020). Given the different layers of task understanding represented in the different responses we find in the recorded neurons in vCA1 (elemental cue representation to structural response), these types of response might emerge at different times over the course of learning.

In summary, in this thesis, I described a new task to investigate structural learning in mice and presented evidence for task-related representations within vCA1 that contain all information to correctly predict outcomes. Understanding the functional contribution of these representations and their role in generalisation and inference as well as investigating the projection targets included in this circuit will be important topics of future investigation.

Bibliography

Ackels, T., Erskine, A., Dasgupta, D., Marin, A. C., Warner, T. P. A., Tootoonian, S., Fukunaga, I., Harris, J. J., & Schaefer, A. T. (2021). Fast odour dynamics are encoded in the olfactory system and guide behaviour. *Nature*, *593*(7860), 558–563. https://doi.org/10.1038/s41586-021-03514-2

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: Common and distinct neural substrates during event construction and elaboration. *Neuropsychologia*, *45*(7), 1363–1377. https://doi.org/10.1016/j.neuropsychologia.2006.10.016

Aggleton, J. P., & Pearce, J. M. (2001). Neural systems underlying episodic memory: insights from animal research. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *356*(1413), 1467–1482. https://doi.org/10.1098/rstb.2001.0946

Aggleton, J. P., Sanderson, D. J., & Pearce, J. M. (2007). Structural learning and the hippocampus. *Hippocampus*, *17*(9), 723–734. https://doi.org/10.1002/hipo.20323

Ahmed, M. S., Priestley, J. B., Castro, A., Stefanini, F., Canales, A. S. S., Balough, E. M., Lavoie, E., Mazzucato, L., Fusi, S., & Losonczy, A. (2020). Hippocampal Network Reorganization Underlies the Formation of a Temporal Association Memory. *Neuron*, *107*(2), 283-291.e6. https://doi.org/10.1016/j.neuron.2020.04.013

Ainge, J. A., van der Meer, M. A. A. van der, Langston, R. F., & Wood, E. R. (2007). Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus*, *17*(10), 988–1002. https://doi.org/10.1002/hipo.20301

Ainge, J. A., Tamosiunaite, M., Woergoetter, F., & Dudchenko, P. A. (2007). Hippocampal CA1 Place Cells Encode Intended Destination on a Maze with Multiple Choice Points. *Journal of Neuroscience*, *27*(36), 9769–9779. https://doi.org/10.1523/jneurosci.2011-07.2007

Akam, T., Rodrigues-Vaz, I., Marcelo, I., Zhang, X., Pereira, M., Oliveira, R. F., Dayan, P., & Costa, R. M. (2021). The Anterior Cingulate Cortex Predicts Future States to Mediate Model-Based Action Selection. *Neuron*, *109*(1), 149-163.e7. https://doi.org/10.1016/j.neuron.2020.10.013

Albasser, M. M., Dumont, J. R., Amin, E., Holmes, J. D., Horne, M. R., Pearce, J. M., & Aggleton, J. P. (2013). Association rules for rat spatial learning: The importance of the hippocampus for binding item identity with item location. *Hippocampus*, *23*(12), 1162–1178. https://doi.org/10.1002/hipo.22154

AlSubaie, R., Wee, R. W., Ritoux, A., Mishchanchuk, K., Passlack, J., Regester, D., & MacAskill, A. F. (2021). Control of parallel hippocampal output pathways by amygdalar long-range inhibition. *eLife*, *10*, e74758. https://doi.org/10.7554/elife.74758

Amaral, D. G., & Witter, M. P. (1989). The three-dimensional organization of the hippocampal formation: A review of anatomical data. *Neuroscience*, *31*(3), 571–591. https://doi.org/10.1016/0306-4522(89)90424-7

Andersen, P., Bliss, T. V. P., & Skrede, K. K. (1971). Lamellar organization of hippocampal excitatory pathways. *Experimental Brain Research*, *13*(2), 222–238. https://doi.org/10.1007/bf00234087

Anderson, M. I., & Jeffery, K. J. (2003). Heterogeneous Modulation of Place Cell Firing by Changes in Context. *The Journal of Neuroscience*, *23*(26), 8827–8835. https://doi.org/10.1523/jneurosci.23-26-08827.2003

Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, *543*(7647), 719–722. https://doi.org/10.1038/nature21692

Axmacher, N., Cohen, M. X., Fell, J., Haupt, S., Dümpelmann, M., Elger, C. E., Schlaepfer, T. E., Lenartz, D., Sturm, V., & Ranganath, C. (2010). Intracranial EEG Correlates of Expectancy and Memory Formation in the Human Hippocampus and Nucleus Accumbens. *Neuron*, *65*(4), 541–549. https://doi.org/10.1016/j.neuron.2010.02.006

Balleine, B. W., Leung, B. K., & Ostlund, S. B. (2011). The orbitofrontal cortex, predicted value, and choice. *Annals of the New York Academy of Sciences*, *1239*(1), 43–50. https://doi.org/10.1111/j.1749-6632.2011.06270.x

Bari, B. A., Grossman, C. D., Lubin, E. E., Rajagopalan, A. E., Cressy, J. I., & Cohen, J. Y. (2019). Stable Representations of Decision Variables for Flexible Behavior. *Neuron*, *103*(5), 922-933.e7. https://doi.org/10.1016/j.neuron.2019.06.001

Barron, H. C., Reeve, H. M., Koolschijn, R. S., Perestenko, P. V., Shpektor, A., Nili, H., Rothaermel, R., Campo-Urriza, N., O'Reilly, J. X., Bannerman, D. M., Behrens, T. E. J., & Dupret, D. (2020). Neuronal Computation Underlying Inferential Reasoning in Humans and Mice. *Cell*, *183*(1), 228-243.e21. https://doi.org/10.1016/j.cell.2020.08.035

Behrens, T. E. J., Muller, T. H., Whittington, J. C. R., Mark, S., Baram, A. B., Stachenfeld, K. L., & Kurth-Nelson, Z. (2018). What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior. *Neuron*, *100*(2), 490–509. https://doi.org/10.1016/j.neuron.2018.10.002

Biane, J. S., Ladow, M. A., Stefanini, F., Boddu, S. P., Fan, A., Hassan, S., Dundar, N., Apodaca-Montano, D. L., Zhou, L. Z., Fayner, V., Woods, N. I., & Kheirbek, M. A. (2023). Neural dynamics underlying associative learning in the dorsal and ventral hippocampus. *Nature Neuroscience*, *26*(5), 798–809. https://doi.org/10.1038/s41593-023-01296-6

Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., & Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*, *357*(6355), 1033–1036. https://doi.org/10.1126/science.aan3846

Blanchard, D. C., & Blanchard, R. J. (1972). Innate and conditioned reactions to threat in rats with amygdaloid lesions. *Journal of Comparative and Physiological Psychology*, *81*(2), 281–290. https://doi.org/10.1037/h0033521

Blanchard, D. C., Blanchard, R. J., Lee, E. M. C., & Fukunaga, K. K. (1977). Movement arrest and the hippocampus. *Physiological Psychology*, *5*(3), 331–335. https://doi.org/10.3758/bf03335340

Bunsey, M. and Eichenbaum, H. (1993) 'Critical role of the Parahippocampal region for paired-associate learning in rats.', *Behavioral Neuroscience*, 107(5), pp. 740–747. doi:10.1037//0735-7044.107.5.740.

Bunsey, M. and Eichenbaum, H. (1996) 'Conservation of hippocampal memory function in rats and humans', *Nature*, 379(6562), pp. 255–257. doi:10.1038/379255a0.

Burgess, N. (2002). The hippocampus, space, and viewpoints in episodic memory. *The Quarterly Journal of Experimental Psychology Section A*, *55*(4), 1057–1080. https://doi.org/10.1080/02724980244000224

Burwell, R. D., & Amaral, D. G. (1998). Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *Journal of Comparative Neurology*, *398*(2), 179–205. https://doi.org/10.1002/(sici)1096-9861(19980824)398:2<179::aid-cne3>3.0.co;2-y

Bussey, T. J., Dias, R., Redhead, E. S., Pearce, J. M., Muir, J. L., & Aggleton, J. P. (2000). Intact negative patterning in rats with fornix or combined perirhinal and postrhinal cortex lesions. *Experimental Brain Research*, *134*(4), 506–519. https://doi.org/10.1007/s002210000481

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

Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the Past and Imagining the Future: A Neural Model of Spatial Memory and Imagery. *Psychological Review*, *114*(2), 340–375. https://doi.org/10.1037/0033-295x.114.2.340

Cenquizca, L. A., & Swanson, L. W. (2007). Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Research Reviews*, *56*(1), 1–26. https://doi.org/10.1016/j.brainresrev.2007.05.002

Chamberlain, S. R., & Sahakian, B. J. (2006). The neuropsychology of mood disorders. *Current Psychiatry Reports*, *8*(6), 458–463. https://doi.org/10.1007/s11920-006-0051-x

Chen, J. L., Margolis, D. J., Stankov, A., Sumanovski, L. T., Schneider, B. L., & Helmchen, F. (2015). Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. *Nature Neuroscience*, *18*(8), 1101–1108. https://doi.org/10.1038/nn.4046

Chuong, A. S., Miri, M. L., Busskamp, V., Matthews, G. A. C., Acker, L. C., Sørensen, A. T., Young, A., Klapoetke, N. C., Henninger, M. A., Kodandaramaiah, S. B., Ogawa, M.,

Ramanlal, S. B., Bandler, R. C., Allen, B. D., Forest, C. R., Chow, B. Y., Han, X., Lin, Y., Tye, K. M., … Boyden, E. S. (2014). Noninvasive optical inhibition with a red-shifted microbial rhodopsin. *Nature Neuroscience*, *17*(8), 1123–1129. https://doi.org/10.1038/nn.3752

Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292), 1464–1468. https://doi.org/10.1126/science.aaf0941

Davidson, T. L., McKernan, M. G., & Jarrard, L. E. (1993). Hippocampal Lesions Do Not Impair Negative Patterning: A Challenge to Configural Association Theory. *Behavioral Neuroscience*, *107*(2), 227–234. https://doi.org/10.1037/0735-7044.107.2.227

Davis, M. (1992). The Role of the Amygdala in Fear and Anxiety. *Annual Review of Neuroscience*, *15*(1), 353–375. https://doi.org/10.1146/annurev.ne.15.030192.002033

Deshmukh, S. S., & Knierim, J. J. (2013). Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus*, *23*(4), 253–267. https://doi.org/10.1002/hipo.22101

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*, *17*(23), 9285–9297. https://doi.org/10.1523/jneurosci.17-23-09285.1997

Diodato, A., Ruinart de Brimont, M., Yim, Y. S., Derian, N., Perrin, S., Pouch, J., … & Fleischmann, A. (2016). Molecular signatures of neural connectivity in the olfactory cortex. *Nature communications*, *7*(1), 12238.

Dong, Z., Mau, W., Feng, Y., Pennington, Z. T., Chen, L., Zaki, Y., Rajan, K., Shuman, T., Aharoni, D., & Cai, D. J. (2022). Minian, an open-source miniscope analysis pipeline. *eLife*, *11*, e70661. https://doi.org/10.7554/elife.70661

Dudchenko, P. A., & Wood, E. R. (2014). *Space,Time and Memory in the Hippocampal Formation*. 253–272. https://doi.org/10.1007/978-3-7091-1292-2_10

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, *13*(8), 995–1002. https://doi.org/10.1038/nn.2599

Durstewitz, D., Vittoz, N. M., Floresco, S. B., & Seamans, J. K. (2010). Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning. *Neuron*, *66*(3), 438–448. https://doi.org/10.1016/j.neuron.2010.03.029

Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations. *Proceedings of the National Academy of Sciences*, *94*(13), 7109–7114. https://doi.org/10.1073/pnas.94.13.7109

Dusek, J. A., & Eichenbaum, H. (1998). The hippocampus and transverse patterning guided by olfactory cues**.** *Behavioral neuroscience*, *112*(4), 762.

Duvelle, É., Grieves, R. M., & van der Meer, M. A. van der. (2023). Temporal context and latent state inference in the hippocampal splitter signal. *eLife*, *12*, e82357. https://doi.org/10.7554/elife.82357

Eichenbaum, H. (2001). The hippocampus and declarative memory: cognitive mechanisms and neural codes. *Behavioural Brain Research*, *127*(1–2), 199–207. https://doi.org/10.1016/s0166-4328(01)00365-5

Eichenbaum, H. (2013). Memory on time. *Trends in Cognitive Sciences*, *17*(2), 81–88. https://doi.org/10.1016/j.tics.2012.12.007

Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, *15*(11), 732–744. https://doi.org/10.1038/nrn3827

Eichenbaum, H. (2017). On the Integration of Space, Time, and Memory. *Neuron*, *95*(5), 1007–1018. https://doi.org/10.1016/j.neuron.2017.06.036

Eichenbaum, H., & Cohen, N. J. (2014). Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*, *83*(4), 764–770. https://doi.org/10.1016/j.neuron.2014.07.032

Eichenbaum, H., Fagan, A. and Cohen, N. (1986) 'Normal olfactory discrimination learning set and facilitation of reversal learning after medial-temporal damage in rats: Implications for an account of preserved learning abilities in amnesia', *The Journal of Neuroscience*, 6(7), pp. 1876–1884. doi:10.1523/jneurosci.06-07-01876.1986.

Eichenbaum, H., Kuperstein, M., Fagan, A., & Nagode, J. (1987). Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *The Journal of Neuroscience*, *7*(3), 716–732. https://doi.org/10.1523/jneurosci.07-03-00716.1987

Eichenbaum, H., Mathews, P. and Cohen, N.J. (1989) 'Further studies of hippocampal representation during odor discrimination learning.', *Behavioral Neuroscience*, 103(6), pp. 1207–1216. doi:10.1037/0735-7044.103.6.1207.

Ekstrom, A. D., & Bookheimer, S. Y. (2007). Spatial and temporal episodic memory retrieval recruit dissociable functional networks in the human brain. *Learning & Memory*, *14*(10), 645–654. https://doi.org/10.1101/lm.575107

Elzakker, M. V., O'Reilly, R. C., & Rudy, J. W. (2003). Transitivity, flexibility, conjunctive representations, and the hippocampus. I. An empirical analysis. *Hippocampus*, *13*(3), 334–340. https://doi.org/10.1002/hipo.10083

Fagan, A., Mathews, P., & Cohen, N. J. (1988). Hippocampal system dysfunction and odor discrimination learning in rats: Impairment of facilitation depending on representational demands. *Behavioral Neuroscience*, *102*(3), 331-339.

Fanselow, M. S. (1990). Factors governing one-trial contextual conditioning. *Animal Learning & Behavior*, *18*(3), 264–270. https://doi.org/10.3758/bf03205285

Fanselow, M. S., & Dong, H.-W. (2010). Are the Dorsal and Ventral Hippocampus Functionally Distinct Structures? *Neuron*, *65*(1), 7–19. https://doi.org/10.1016/j.neuron.2009.11.031

Ferbinteanu, J., Shirvalkar, P., & Shapiro, M. L. (2011). Memory Modulates Journey-Dependent Coding in the Rat Hippocampus. *The Journal of Neuroscience*, *31*(25), 9135–9146. https://doi.org/10.1523/jneurosci.1241-11.2011

Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, *5*(5), 458–462. https://doi.org/10.1038/nn834

Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory Encoding in the Hippocampus and Entorhinal Cortex. *Neuron*, *27*(1), 169–178. https://doi.org/10.1016/s0896-6273(00)00018-0

Franks, K. M., Russo, M. J., Sosulski, D. L., Mulligan, A. A., Siegelbaum, S. A., & Axel, R. (2011). Recurrent Circuitry Dynamically Shapes the Activation of Piriform Cortex. *Neuron*, *72*(1), 49–56. https://doi.org/10.1016/j.neuron.2011.08.020

Funahashi, S., Bruce, C., & Goldman-Rakic, P. (1993). Dorsolateral prefrontal lesions and oculomotor delayed-response performance: evidence for mnemonic "scotomas." *The Journal of Neuroscience*, *13*(4), 1479–1497. https://doi.org/10.1523/jneurosci.13-04-01479.1993

Gauthier, J. L., & Tank, D. W. (2018). A Dedicated Population for Reward Coding in the Hippocampus. *Neuron*, *99*(1), 179-193.e7. https://doi.org/10.1016/j.neuron.2018.06.008

Gerraty, R. T., Davidow, J. Y., Wimmer, G. E., Kahn, I., & Shohamy, D. (2014). Transfer of Learning Relates to Intrinsic Connectivity between Hippocampus, Ventromedial Prefrontal Cortex, and Large-Scale Networks. *The Journal of Neuroscience*, *34*(34), 11297–11303. https://doi.org/10.1523/jneurosci.0185-14.2014

Gershman, S. J. (2018). The Successor Representation: Its Computational Logic and Neural Substrates. *The Journal of Neuroscience*, *38*(33), 7193–7200. https://doi.org/10.1523/jneurosci.0151-18.2018

Gershman, S. J., Norman, K. A., & Niv, Y. (2015). Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, *5*, 43–50. https://doi.org/10.1016/j.cobeha.2015.07.007

Graham, K. S., Barense, M. D., & Lee, A. C. H. (2010). Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, *48*(4), 831–853. https://doi.org/10.1016/j.neuropsychologia.2010.01.001

Gray, J. A., & McNaughton, N. (2003). *The Neuropsychology of Anxiety*. Oxford University Press.

van Groen, T., & Wyss, J. M. (1990a). The connections of presubiculum and parasubiculum in the rat. *Brain research*, *518*(1-2), 227-243.

van Groen, T., & Wyss, J. M. (1990b). Extrinsic projections from area CA1 of the rat hippocampus: Olfactory, cortical, subcortical, and bilateral hippocampal formation projections. *Journal of Comparative Neurology*, *302*(3), 515–528. https://doi.org/10.1002/cne.903020308

van Groen, T., Haren, F. J. van, Witter, M. P., & Groenewegen, H. J. (1986). The organization of the reciprocal connections between the subiculum and the entorhinal cortex in the cat: I. A neuroanatomical tracing study. *Journal of Comparative Neurology*, *250*(4), 485–497. https://doi.org/10.1002/cne.902500407

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., Cox, J., & Svoboda, K. (2014). Procedures for Behavioral Experiments in Head-Fixed Mice. *PLoS ONE*, *9*(2), e88678. https://doi.org/10.1371/journal.pone.0088678

Haber, S. N., Liu, H., Seidlitz, J., & Bullmore, E. (2022). Prefrontal connectomics: from anatomy to human imaging. *Neuropsychopharmacology*, *47*(1), 20–40. https://doi.org/10.1038/s41386-021-01156-6

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801–806. https://doi.org/10.1038/nature03721

Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2014). Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1635), 20120510. https://doi.org/10.1098/rstb.2012.0510

Hassabis, D., Kumaran, D., & Maguire, E. A. (2007). Using Imagination to Understand the Neural Basis of Episodic Memory. *The Journal of Neuroscience*, *27*(52), 14365–14374. https://doi.org/10.1523/jneurosci.4549-07.2007

Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, *104*(5), 1726–1731. https://doi.org/10.1073/pnas.0610561104

Hollup, S. A., Molden, S., Donnett, J. G., Moser, M.-B., & Moser, E. I. (2001). Accumulation of Hippocampal Place Fields at the Goal Location in an Annular Watermaze Task. *The Journal of Neuroscience*, *21*(5), 1635–1644. https://doi.org/10.1523/jneurosci.21-05-01635.2001

Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. M., & Graybiel, A. M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*, *500*(7464), 575–579. https://doi.org/10.1038/nature12475

Itskov, V., Curto, C., Pastalkova, E., & Buzsáki, G. (2011). Cell Assembly Sequences Arising from Spike Threshold Adaptation Keep Track of Time in the Hippocampus.

The Journal of Neuroscience, *31*(8), 2828–2834. https://doi.org/10.1523/jneurosci.3773-10.2011

Jarrard, L. E. (1978). Selective hippocampal lesions: Differential effects on performance by rats of a spatial task with preoperative versus postoperative training. *Journal of Comparative and Physiological Psychology*, *92*(6), 1119–1127. https://doi.org/10.1037/h0077516

Jarzebowski, P., Hay, Y. A., Grewe, B. F., & Paulsen, O. (2022). Different encoding of reward location in dorsal and intermediate hippocampus. *Current Biology*, *32*(4), 834-841.e5. https://doi.org/10.1016/j.cub.2021.12.024

Jay, T. M., & Witter, M. P. (1991). Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, *313*(4), 574–586. https://doi.org/10.1002/cne.903130404

Jezek, K., Henriksen, E. J., Treves, A., Moser, E. I., & Moser, M.-B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. *Nature*, *478*(7368), 246–249. https://doi.org/10.1038/nature10439

de Jong, L. W., Nejad, M. M., Yoon, E., Cheng, S., & Diba, K. (2023). Optogenetics reveals paradoxical network stabilizations in hippocampal CA1 and CA3. *Current Biology*, *33*(9), 1689-1703.

Johnston, M., Scarf, D., Wilson, A., Millar, J., Bartonicek, A., & Colombo, M. (2021). The effects of hippocampal and area parahippocampalis lesions on the processing and retention of serial-order behavior, autoshaping, and spatial behavior in pigeons. *Hippocampus*, *31*(3), 261–280. https://doi.org/10.1002/hipo.23287

Jung, M. W., Qin, Y., McNaughton, B. L., & Barnes, C. A. (1998). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cerebral Cortex (New York, N.Y.: 1991)*, *8*(5), 437–450. https://doi.org/10.1093/cercor/8.5.437

Kamiński, J., Sullivan, S., Chung, J. M., Ross, I. B., Mamelak, A. N., & Rutishauser, U. (2017). Persistently active neurons in human medial frontal and medial temporal lobe support working memory. *Nature Neuroscience*, *20*(4), 590–601. https://doi.org/10.1038/nn.4509

Kandel, E. R. (2021). *Principles of Neural Science* (S. Siegelbaum, S. Mack, & J. D. Koester, Eds.; 6th edition). McGraw Hill. https://neurology.mhmedical.com/content.aspx?bookid=3024§ionid=254326759

Kangas, B. D., Berry, M. S., & Branch, M. N. (2011). On the development and mechanics of delayed matching-to-sample performance. *Journal of the Experimental Analysis of Behavior*, *95*(2), 221–236. https://doi.org/10.1901/jeab.2011.95-221

Kerr, K. M., Agster, K. L., Furtak, S. C., & Burwell, R. D. (2007). Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. *Hippocampus*, *17*(9), 697-7

Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The Role of the Hippocampus in Memory for the Temporal Order of a Sequence of Odors. *Behavioral Neuroscience*, *116*(2), 286–290. https://doi.org/10.1037/0735-7044.116.2.286

Kesner, R. P., Hunsaker, M. R., & Ziegler, W. (2011). The role of the dorsal and ventral hippocampus in olfactory working memory. *Neurobiology of learning and memory*, *96*(2), 361-366.

Kjelstrup, K. G., Tuvnes, F. A., Steffenach, H.-A., Murison, R., Moser, E. I., & Moser, M.-B. (2002). Reduced fear expression after lesions of the ventral hippocampus. *Proceedings of the National Academy of Sciences*, *99*(16), 10825–10830. https://doi.org/10.1073/pnas.152112399

Köhler, C. (1985). Intrinsic projections of the retrohippocampal region in the rat brain. I. The subicular complex. *Journal of Comparative Neurology*, *236*(4), 504–522. https://doi.org/10.1002/cne.902360407

Komorowski, R. W., Manns, J. R., Eichenbaum, H. (2009). Robust conjunctive item–place coding by hippocampal neurons parallels learning what happens where. Journal of Neuroscience, 29(31), 9918–9929. https://doi.org/10.1523/JNEUROSCI.1378-09.2009

Kumaran, D., & Maguire, E. A. (2006). An Unexpected Sequence of Events: Mismatch Detection in the Human Hippocampus. *PLoS Biology*, *4*(12), e424. https://doi.org/10.1371/journal.pbio.0040424

Lee, I., & Kesner, R. P. (2004). Encoding versus retrieval of spatial memory: Double dissociation between the dentate gyrus and the perforant path inputs into CA3 in the dorsal hippocampus. *Hippocampus*, *14*(1), 66–76. https://doi.org/10.1002/hipo.10167

LeGates, T. A., Kvarta, M. D., Tooley, J. R., Francis, T. C., Lobo, M. K., Creed, M. C., & Thompson, S. M. (2018). Reward behaviour is regulated by the strength of hippocampus–nucleus accumbens synapses. *Nature*, *564*(7735), 258–262. https://doi.org/10.1038/s41586-018-0740-8

Lehn, H., Steffenach, H.-A., Strien, N. M. van, Veltman, D. J., Witter, M. P., & Haberg, A. K. (2009). A Specific Role of the Human Hippocampus in Recall of Temporal Sequences. *Journal of Neuroscience*, *29*(11), 3475–3484. https://doi.org/10.1523/jneurosci.5370-08.2009

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. *Science*, *315*(5814), 961–966. https://doi.org/10.1126/science.1135801

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary Vector Cells in the Subiculum of the Hippocampal Formation. *The Journal of Neuroscience*, *29*(31), 9771–9777. https://doi.org/10.1523/jneurosci.1319-09.2009

Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental Brain Research*, *133*(1), 23–32. https://doi.org/10.1007/s002210000397

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

Lu, J., Li, C., Singh-Alvarado, J., Zhou, Z. C., Fröhlich, F., Mooney, R., & Wang, F. (2018). MIN1PIPE: A Miniscope 1-Photon-Based Calcium Imaging Signal Extraction Pipeline. *Cell Reports*, *23*(12), 3673–3684. https://doi.org/10.1016/j.celrep.2018.05.062

Luskin, M. B., & Price, J. L. (1983). The topographic organization of associational fibers of the olfactory system in the rat, including centrifugal fibers to the olfactory bulb. *Journal of comparative neurology*, *216*(3), 264-291.

MacDonald, C. J., Carrow, S., Place, R., & Eichenbaum, H. (2013). Distinct Hippocampal Time Cell Sequences Represent Odor Memories in Immobilized Rats. *The Journal of Neuroscience*, *33*(36), 14607–14616. https://doi.org/10.1523/jneurosci.1537-13.2013

Maguire, E. A., & Mullally, S. L. (2013). The Hippocampus: A Manifesto for Change. *Journal of Experimental Psychology: General*, *142*(4), 1180–1189. https://doi.org/10.1037/a0033650

Manns, J. R., Howard, M. W., & Eichenbaum, H. (2007). Gradual Changes in Hippocampal Activity Support Remembering the Order of Events. *Neuron*, *56*(3), 530–540. https://doi.org/10.1016/j.neuron.2007.08.017

Marbach, F., & Zador, A. M. (2017). A self-initiated two-alternative forced choice paradigm for head-fixed mice. *bioRxiv*, 073783. https://doi.org/10.1101/073783

Maren, S. (1996). Synaptic transmission and plasticity in the amygdala. *Molecular Neurobiology*, *13*(1), 1–22. https://doi.org/10.1007/bf02740749

Maren, S. (2005). Synaptic Mechanisms of Associative Memory in the Amygdala. *Neuron*, *47*(6), 783–786. https://doi.org/10.1016/j.neuron.2005.08.009

Maren, S., & Holt, W. G. (2004). Hippocampus and Pavlovian Fear Conditioning in Rats: Muscimol Infusions Into the Ventral, but Not Dorsal, Hippocampus Impair the Acquisition of Conditional Freezing to an Auditory Conditional Stimulus. *Behavioral Neuroscience*, *118*(1), 97–110. https://doi.org/10.1037/0735-7044.118.1.97

Maren, S., Phan, K. L., & Liberzon, I. (2013). The contextual brain: implications for fear conditioning, extinction and psychopathology. *Nature Reviews Neuroscience*, *14*(6), 417–428. https://doi.org/10.1038/nrn3492

Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, *262*(841), 23–81. https://doi.org/10.1098/rstb.1971.0078

Mau, W., Sullivan, D. W., Kinsky, N. R., Hasselmo, M. E., Howard, M. W., & Eichenbaum, H. (2018). The Same Hippocampal CA1 Population Simultaneously Codes Temporal Information over Multiple Timescales. *Current Biology*, *28*(10), 1499-1508.e4. https://doi.org/10.1016/j.cub.2018.03.051

McAlonan, K., & Brown, V. J. (2003). Orbital prefrontal cortex mediates reversal learning and not attentional set shifting in the rat. *Behavioural Brain Research*, *146*(1–2), 97–103. https://doi.org/10.1016/j.bbr.2003.09.019

McDannald, M. A., Saddoris, M. P., Gallagher, M., & Holland, P. C. (2005). Lesions of Orbitofrontal Cortex Impair Rats' Differential Outcome Expectancy Learning But Not Conditioned Stimulus-Potentiated Feeding. *The Journal of Neuroscience*, *25*(18), 4626–4632. https://doi.org/10.1523/jneurosci.5301-04.2005

McDonald, R. J., Murphy, R. A., Guarraci, F. A., Gortler, J. R., White, N. M., & Baker, A. G. (1997). Systematic comparison of the effects of hippocampal and fornix-fimbria lesions on acquisition of three configural discriminations. *Hippocampus*, *7*(4), 371–388. https://doi.org/10.1002/(sici)1098-1063(1997)7:4<371::aid-hipo3>3.0.co;2-m

McKenna, J. T., & Vertes, R. P. (2004). Afferent projections to nucleus reuniens of the thalamus. *Journal of Comparative Neurology*, *480*(2), 115–142. https://doi.org/10.1002/cne.20342

McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Riviere, P. D., Eichenbaum, H. (2014). Hippocampal representation of related and opposing memories develop with distinct time courses. eLife, 3, e01079. https://doi.org/10.7554/eLife.01079

McKenzie, S., Keene, C. S., Farovik, A., Bladon, J., Place, R., Komorowski, R., & Eichenbaum, H. (2016). Representation of memories in the cortical–hippocampal system: Results from the application of population similarity analyses. Neurobiology of learning and memory, 134, 178-191.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the "cognitive map." *Nature Reviews Neuroscience*, *7*(8), 663–678. https://doi.org/10.1038/nrn1932

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, *10*(10), 408–415. https://doi.org/10.1016/0166-2236(87)90011-7

van der Meer, M. A. A., Johnson, A., Schmitzer-Torbert, N. C., & Redish, A. D. (2010). Triple Dissociation of Information Processing in Dorsal Striatum, Ventral Striatum, and Hippocampus on a Learned Spatial Decision Task. *Neuron*, *67*(1), 25–32. https://doi.org/10.1016/j.neuron.2010.06.023

Milivojevic, B., & Doeller, C. F. (2013). Mnemonic Networks in the Hippocampal Formation: From Spatial Maps to Temporal and Conceptual Codes. *Journal of Experimental Psychology: General*, *142*(4), 1231–1241. https://doi.org/10.1037/a0033746

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*(1), 167–202. https://doi.org/10.1146/annurev.neuro.24.1.167

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *The Journal of Neuroscience*, *16*(16), 5154–5167. https://doi.org/10.1523/jneurosci.16-16-05154.1996

Monchi, O., Petrides, M., Petre, V., Worsley, K., & Dagher, A. (2001). Wisconsin Card Sorting Revisited: Distinct Neural Circuits Participating in Different Stages of the Task Identified by Event-Related Functional Magnetic Resonance Imaging. *The Journal of Neuroscience*, *21*(19), 7733–7741. https://doi.org/10.1523/jneurosci.21-19-07733.2001

Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, *297*(5868), 681–683. https://doi.org/10.1038/297681a0

Moser, M. B., Moser, E. I., Forrest, E., Andersen, P., & Morris, R. G. (1995). Spatial learning with a minislab in the dorsal hippocampus. *Proceedings of the National Academy of Sciences*, *92*(21), 9697–9701. https://doi.org/10.1073/pnas.92.21.9697

Moser, M., & Moser, E. I. (1998). Functional differentiation in the hippocampus. *Hippocampus*, *8*(6), 608–619. https://doi.org/10.1002/(sici)1098-1063(1998)8:6<608::aid-hipo3>3.0.co;2-7

Neunuebel, J. P., & Knierim, J. J. (2014). CA3 Retrieves Coherent Representations from Degraded Input: Direct Evidence for CA3 Pattern Completion and Dentate Gyrus Pattern Separation. *Neuron*, *81*(2), 416–427. https://doi.org/10.1016/j.neuron.2013.11.017

Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, *22*(10), 1544–1553. https://doi.org/10.1038/s41593-019-0470-8

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. https://doi.org/10.1016/0006-8993(71)90358-1

O'Keefe, J., & Nadel, L. (1979). Précis of O'Keefe & Nadel's The hippocampus as a cognitive map. *Behavioral and Brain Sciences*, *2*(4), 487–494. https://doi.org/10.1017/s0140525x00063949

Omer, D. B., Las, L., & Ulanovsky, N. (2022). Contextual and pure time coding for self and other in the hippocampus. *Nature Neuroscience*, 1–10. https://doi.org/10.1038/s41593-022-01226-y

Öngür, D., Cullen, T. J., Wolf, D. H., Rohan, M., Barreira, P., Zalesak, M., & Heckers, S. (2006). The Neural Basis of Relational Memory Deficits in Schizophrenia. *Archives of General Psychiatry*, *63*(4), 356–365. https://doi.org/10.1001/archpsyc.63.4.356

Otto, T., Schottler, F., Staubli, U., Eichenbaum, H., & Lynch, G. (1991). Hippocampus and olfactory discrimination learning: effects of entorhinal cortex lesions on olfactory

learning and memory in a successive-cue, go–no-go task. *Behavioral neuroscience*, *105*(1), 111.

Owen, S. F., Liu, M. H., & Kreitzer, A. C. (2019). Thermal constraints on in vivo optogenetic manipulations. Nature neuroscience, 22(7), 1061-1065.

Panoz-Brown, D., Corbin, H. E., Dalecki, S. J., Gentry, M., Brotheridge, S., Sluka, C. M., Wu, J.-E., & Crystal, J. D. (2016). Rats Remember Items in Context Using Episodic Memory. *Current Biology*, *26*(20), 2821–2826. https://doi.org/10.1016/j.cub.2016.08.023

Parker, N. F., Cameron, C. M., Taliaferro, J. P., Lee, J., Choi, J. Y., Davidson, T. J., Daw, N. D., & Witten, I. B. (2016). Reward and choice encoding in terminals of midbrain dopamine neurons depends on striatal target. *Nature Neuroscience*, *19*(6), 845–854. https://doi.org/10.1038/nn.4287

Pashkovski, S. L., Iurilli, G., Brann, D., Chicharro, D., Drummey, K., Franks, K. M., Panzeri, S., & Datta, S. R. (2020). Structure and flexibility in cortical representations of odour space. *Nature*, *583*(7815), 253–258. https://doi.org/10.1038/s41586-020-2451-1

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science*, *321*(5894), 1322–1327. https://doi.org/10.1126/science.1159775

Pavlov, P. I. (1927). Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. *Annals of Neurosciences*, *17*(3), 136–141. https://doi.org/10.5214/ans.0972-7531.1017309

Paxinos, G., & Franklin, K. B. J. (2019). *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates*. Academic Press.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*(6), 532–552. https://doi.org/10.1037/0033-295x.87.6.532

Pennartz, C. M. A., Ito, R., Verschure, P. F. M. J., Battaglia, F. P., & Robbins, T. W. (2011). The hippocampal–striatal axis in learning, prediction and goal-directed behavior. *Trends in Neurosciences*, *34*(10), 548–559. https://doi.org/10.1016/j.tins.2011.08.001

Phillips, R. G., & LeDoux, J. E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behavioral Neuroscience*, *106*(2), 274–285. https://doi.org/10.1037//0735-7044.106.2.274

Quinn, J. J., Oommen, S. S., Morrison, G. E., & Fanselow, M. S. (2002). Post-training excitotoxic lesions of the dorsal hippocampus attenuate forward trace, backward trace, and delay fear conditioning in a temporally specific manner. *Hippocampus*, *12*(4), 495–504. https://doi.org/10.1002/hipo.10029

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

Ranganath, C. (2010). Binding Items and Contexts. *Current Directions in Psychological Science*, *19*(3), 131–137. https://doi.org/10.1177/0963721410368805

Rao, Y. L., Ganaraja, B., Murlimanju, B. V., Joy, T., Krishnamurthy, A., & Agrawal, A. (2022). Hippocampus and its involvement in Alzheimer's disease: a review. *3 Biotech*, *12*(2), 55. https://doi.org/10.1007/s13205-022-03123-4

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory*. Appleton-Century-Crofts.

Riaz, S., Schumacher, A., Sivagurunathan, S., van der Meer, M., & Ito, R. (2017). Ventral, but not dorsal, hippocampus inactivation impairs reward memory expression and retrieval in contexts defined by proximal cues. *Hippocampus*, *27*(7), 822–836. https://doi.org/10.1002/hipo.22734

Rich, E. L., & Shapiro, M. (2009). Rat Prefrontal Cortical Neurons Selectively Code Strategy Switches. *The Journal of Neuroscience*, *29*(22), 7208–7219. https://doi.org/10.1523/jneurosci.6068-08.2009

Root, C. M., Denny, C. A., Hen, R., & Axel, R. (2014). The participation of cortical amygdala in innate, odour-driven behaviour. *Nature*, *515*(7526), 269–273. https://doi.org/10.1038/nature13897

Rudy, J. W., & O'Reilly, R. C. (1999). Contextual Fear Conditioning, Conjunctive Representations, Pattern Completion, and the Hippocampus. *Behavioral Neuroscience*, *113*(5), 867–880. https://doi.org/10.1037/0735-7044.113.5.867

Rudy, J. W., & Sutherland, R. J. (1995). Configural association theory and the hippocampal formation: An appraisal and reconfiguration. *Hippocampus*, *5*(5), 375–389. https://doi.org/10.1002/hipo.450050502

Rygula, R., Walker, S. C., Clarke, H. F., Robbins, T. W., & Roberts, A. C. (2010). Differential Contributions of the Primate Ventrolateral Prefrontal and Orbitofrontal Cortex to Serial Reversal Learning. *The Journal of Neuroscience*, *30*(43), 14552–14559. https://doi.org/10.1523/jneurosci.2631-10.2010

Sánchez-Bellot, C., AlSubaie, R., Mishchanchuk, K., Wee, R. W. S., & MacAskill, A. F. (2022). Two opposing hippocampus to prefrontal cortex pathways for the control of approach and avoidance behaviour. *Nature Communications*, *13*(1), 339. https://doi.org/10.1038/s41467-022-27977-7

Sanders, H., Wilson, M. A., & Gershman, S. J. (2020). Hippocampal remapping as hidden state inference. *eLife*, *9*, e51140. https://doi.org/10.7554/elife.51140

Sanderson, D. J., Pearce, J. M., Kyd, R. J., & Aggleton, J. P. (2006). The importance of the rat hippocampus for learning the structure of visual arrays. *European Journal of Neuroscience*, *24*(6), 1781–1788. https://doi.org/10.1111/j.1460-9568.2006.05035.x

Schmajuk, N. A., & Buhusi, C. V. (1997). Stimulus Configuration, Occasion Setting, and the Hippocampus. *Behavioral Neuroscience*, *111*(2), 235–258. https://doi.org/10.1037/0735-7044.111.2.235

Schoenbaum, G., Setlow, B., Nugent, S. L., Saddoris, M. P., & Gallagher, M. (2003). Lesions of Orbitofrontal Cortex and Basolateral Amygdala Complex Disrupt Acquisition of Odor-Guided Discriminations and Reversals. *Learning & Memory*, *10*(2), 129–140. https://doi.org/10.1101/lm.55203

Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, *275*(5306), 1593–1599. https://doi.org/10.1126/science.275.5306.1593

Scoville, W. B., & Milner, B. (1957). LOSS OF RECENT MEMORY AFTER BILATERAL HIPPOCAMPAL LESIONS. *Journal of Neurology, Neurosurgery & Psychiatry*, *20*(1), 11. https://doi.org/10.1136/jnnp.20.1.11

Sellami, A., Abed, A. S. A., Brayda-Bruno, L., Etchamendy, N., Valério, S., Oulé, M., Pantaléon, L., Lamothe, V., Potier, M., Bernard, K., Jabourian, M., Herry, C., Mons, N., Piazza, P.-V., Eichenbaum, H., & Marighetto, A. (2017). Temporal binding function of dorsal CA1 is critical for declarative memory formation. *Proceedings of the National Academy of Sciences*, *114*(38), 10262–10267. https://doi.org/10.1073/pnas.1619657114

Shepherd, G. M. (Ed.). (2003). *The synaptic organization of the brain*. Oxford university press.

Sosa, M., Plitt, M. H., & Giocomo, L. M. (2024). Hippocampal sequences span experience relative to rewards. *bioRxiv*, 2023.12.27.573490. https://doi.org/10.1101/2023.12.27.573490

Spellman, T., Rigotti, M., Ahmari, S. E., Fusi, S., Gogos, J. A., & Gordon, J. A. (2015). Hippocampal–prefrontal input supports spatial encoding in working memory. *Nature*, *522*(7556), 309–314. https://doi.org/10.1038/nature14445

Squire, L. R., Horst, A. S. van der, McDuff, S. G. R., Frascino, J. C., Hopkins, R. O., & Mauldin, K. N. (2010). Role of the hippocampus in remembering the past and imagining the future. *Proceedings of the National Academy of Sciences*, *107*(44), 19044–19048. https://doi.org/10.1073/pnas.1014391107

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, *20*(11), 1643–1653. https://doi.org/10.1038/nn.4650

Sul, J. H., Kim, H., Huh, N., Lee, D., & Jung, M. W. (2010). Distinct Roles of Rodent Orbitofrontal and Medial Prefrontal Cortex in Decision Making. *Neuron*, *66*(3), 449–460. https://doi.org/10.1016/j.neuron.2010.03.033

Sutherland, R. J., & Rudy, J. W. (1989). Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology*, *17*(2), 129–144. https://doi.org/10.3758/bf03337828

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.

Swanson, L. W., & Cowan, W. M. (1977). An autoradiographic study of the organization of the efferet connections of the hippocampal formation in the rat. *Journal of Comparative Neurology*, *172*(1), 49–84. https://doi.org/10.1002/cne.901720104

Takahashi, H., & Magee, J. C. (2009). Pathway interactions and synaptic plasticity in the dendritic tuft regions of CA1 pyramidal neurons. Neuron, 62(1), 102–111. https://doi.org/10.1016/j.neuron.2009.03.007

Takahashi, Y. K., Roesch, M. R., Wilson, R. C., Toreson, K., O'Donnell, P., Niv, Y., & Schoenbaum, G. (2011). Expectancy-related changes in firing of dopamine neurons depend on orbitofrontal cortex. *Nature Neuroscience*, *14*(12), 1590–1597. https://doi.org/10.1038/nn.2957

Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *The Journal of Neuroscience*, *10*(2), 436–447. https://doi.org/10.1523/jneurosci.10-02-00436.1990

Taxidis, J., Pnevmatikakis, E. A., Dorian, C. C., Mylavarapu, A. L., Arora, J. S., Samadian, K. D., Hoffberg, E. A., & Golshani, P. (2020). Differential Emergence and Stability of Sensory and Temporal Representations in Context-Specific Hippocampal Sequences. *Neuron*, *108*(5), 984-998.e9. https://doi.org/10.1016/j.neuron.2020.08.028

Tennant, S. A., Clark, H., Hawes, I., Tam, W. K., Hua, J., Yang, W., Gerlei, K. Z., Wood, E. R., & Nolan, M. F. (2022). Spatial representation by ramping activity of neurons in the retrohippocampal cortex. *Current Biology*, *32*(20), 4451-4464.e7. https://doi.org/10.1016/j.cub.2022.08.050

Tolman, E. C. (1932). *Purposive behavior in animals and men*. University of California Press.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208. https://doi.org/10.1037/h0061626

Twining, R. C., Lepak, K., Kirry, A. J., & Gilmartin, M. R. (2020). Ventral Hippocampal Input to the Prelimbic Cortex Dissociates the Context from the Cue Association in Trace Fear Memory. *The Journal of Neuroscience*, *40*(16), 3217–3230. https://doi.org/10.1523/jneurosci.1453-19.2020

Uchida, N., & Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nature Neuroscience*, *6*(11), 1224–1229. https://doi.org/10.1038/nn1142

Valero, M., & Prida, L. M. de la. (2018). The hippocampus in depth: a sublayer-specific perspective of entorhinal–hippocampal function. *Current Opinion in Neurobiology*, *52*, 107–114. https://doi.org/10.1016/j.conb.2018.04.013

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Paesschen, W. V., & Mishkin, M. (1997). Differential Effects of Early Hippocampal Pathology on Episodic

and Semantic Memory. *Science*, *277*(5324), 376–380. https://doi.org/10.1126/science.277.5324.376

Vertes, R. P., Hoover, W. B., Szigeti-Buck, K., & Leranth, C. (2007). Nucleus reuniens of the midline thalamus: Link between the medial prefrontal cortex and the hippocampus. *Brain Research Bulletin*, *71*(6), 601–609. https://doi.org/10.1016/j.brainresbull.2006.12.002

Vogel, P., Hahn, J., Duvarci, S., & Sigurdsson, T. (2022). Prefrontal pyramidal neurons are critical for all phases of working memory. *Cell Reports*, *39*(2), 110659. https://doi.org/10.1016/j.celrep.2022.110659

Wallenstein, G. V., Hasselmo, M. E., & Eichenbaum, H. (1998). The hippocampus as an associator of discontiguous events. Trends in neurosciences, 21(8), 317-323.

Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, *411*(6840), 953–956. https://doi.org/10.1038/35082081

Walton, M. E., Behrens, T. E. J., Buckley, M. J., Rudebeck, P. H., & Rushworth, M. F. S. (2010). Separable Learning Systems in the Macaque Brain and the Role of Orbitofrontal Cortex in Contingent Learning. *Neuron*, *65*(6), 927–939. https://doi.org/10.1016/j.neuron.2010.02.027

Wang, F., Schoenbaum, G., & Kahnt, T. (2020). Interactions between human orbitofrontal cortex and hippocampus support model-based inference. *PLoS Biology*, *18*(1), e3000578. https://doi.org/10.1371/journal.pbio.3000578

Wee, R. W. S., Mishchanchuk, K., AlSubaie, R., Church, T. W., Gold, M. G., & MacAskill, A. F. (2024). Internal-state-dependent control of feeding behavior via hippocampal ghrelin signaling. *Neuron*, *112*(2), 288-305.e7. https://doi.org/10.1016/j.neuron.2023.10.016

Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W., & Behrens, T. E. J. (2022). How to build a cognitive map. *Nature Neuroscience*, *25*(10), 1257–1272. https://doi.org/10.1038/s41593-022-01153-y

Whittington, J. C. R., Muller, T. H., Mark, S., Chen, G., Barry, C., Burgess, N., & Behrens, T. E. J. (2020). The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation. *Cell*, *183*(5), 1249-1263.e23. https://doi.org/10.1016/j.cell.2020.10.024

Wikenheiser, A. M., & Schoenbaum, G. (2016). Over the river, through the woods: cognitive maps in the hippocampus and orbitofrontal cortex. *Nature Reviews Neuroscience*, *17*(8), 513–523. https://doi.org/10.1038/nrn.2016.56

Wilson, F. A. W., Scalaidhe, S. P. Ó., & Goldman-Rakic, P. S. (1993). Dissociation of Object and Spatial Processing Domains in Primate Prefrontal Cortex. *Science*, *260*(5116), 1955–1958. https://doi.org/10.1126/science.8316836

Wilson, R. C., Takahashi, Y. K., Schoenbaum, G., & Niv, Y. (2014). Orbitofrontal Cortex as a Cognitive Map of Task Space. *Neuron*, *81*(2), 267–279. https://doi.org/10.1016/j.neuron.2013.11.005

Witter, M. P., Naber, P. A., Haeften, T. van, Machielsen, W. C. M., Rombouts, S. A. R. B., Barkhof, F., Scheltens, P., & Silva, F. H. L. da. (2000). Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus*, *10*(4), 398–410. https://doi.org/10.1002/1098-1063(2000)10:4<398::aid-hipo6>3.0.co;2-k

Wood, E. R., Dudchenko, P. A., & Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, *397*(6720), 613–616. https://doi.org/10.1038/17605

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location. *Neuron*, *27*(3), 623–633. https://doi.org/10.1016/s0896-6273(00)00071-4

Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, *34*(10), 515–525. https://doi.org/10.1016/j.tins.2011.06.006

Yun, M., Hwang, J. Y., & Jung, M. W. (2023). Septotemporal variations in hippocampal value and outcome processing. *Cell Reports*, *42*(2), 112094. https://doi.org/10.1016/j.celrep.2023.112094

Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and Ventral Medial Prefrontal Activation during Retrieval-Mediated Learning Supports Novel Inference. *Neuron*, *75*(1), 168–179. https://doi.org/10.1016/j.neuron.2012.05.010

Zhang, L., Liang, B., Barbera, G., Hawes, S., Zhang, Y., Stump, K., Baum, I., Yang, Y., Li, Y., & Lin, D. (2019). Miniscope GRIN Lens System for Calcium Imaging of Neuronal Activity from Deep Brain Structures in Behaving Animals. *Current Protocols in Neuroscience*, *86*(1), e56. https://doi.org/10.1002/cpns.56

Zhou, Y., Zhu, H., Liu, Z., Chen, X., Su, X., Ma, C., Tian, Z., Huang, B., Yan, E., Liu, X., & Ma, L. (2019). A ventral CA1 to nucleus accumbens core engram circuit mediates conditioned place preference for cocaine. *Nature Neuroscience*, *22*(12), 1986–1999. https://doi.org/10.1038/s41593-019-0524-y

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., Gamal, A. E., & Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, *16*(3), 264–266. https://doi.org/10.1038/nn.3329

Appendix

Appendix 1

Results of planned comparisons: post-hoc Tukey tests for beta coefficients (**Figure 3.5e**)

X	Y	mean(X)	mean(Y)	diff	T	p-tukey
A	AB	-0.057	1.295	-1.351	-6.006	1.70E-06
A	AC	-0.057	-0.885	0.828	3.680	2.51E-02
A	BA	-0.057	-1.335	1.278	5.683	7.93E-06
A	BC	-0.057	0.866	-0.923	-4.103	5.95E-03
A	CA	-0.057	0.868	-0.925	-4.113	5.75E-03
A	CB	-0.057	-0.977	0.920	4.089	6.27E-03
AB	AandB	1.295	-0.041	1.335	5.935	2.40E-06
AB	AandC	1.295	-0.016	1.311	5.825	4.06E-06
AB	B	1.295	-0.151	1.445	6.425	2.17E-07
AB	C	1.295	-0.126	1.421	6.315	3.75E-07
AB	CandB	1.295	-0.110	1.405	6.244	5.34E-07
AB	bias	1.295	-0.167	1.461	6.496	1.52E-07
AB	previous_choice	1.295	0.563	0.731	3.250	n.s.
AB	previous_reward	1.295	-0.332	1.627	7.231	3.37E-09
AC	C	-0.885	-0.126	-0.758	-3.371	n.s.
AC	B	-0.885	-0.151	-0.734	-3.261	n.s.
AC	AandB	-0.885	-0.041	-0.844	-3.751	2.00E-02
AC	AandC	-0.885	-0.016	-0.868	-3.861	1.39E-02
AC	CandB	-0.885	-0.110	-0.774	-3.442	n.s.
AC	previous_reward	-0.885	-0.332	-0.552	-2.455	n.s.
AC	previous_choice	-0.885	0.563	-1.448	-6.436	2.05E-07
AC	bias	-0.885	-0.167	-0.718	-3.190	n.s.
AandB	BA	-0.041	-1.335	1.295	5.754	5.68E-06
AandB	BC	-0.041	0.866	-0.907	-4.032	7.67E-03

AandB	CA	-0.041	0.868	-0.909	-4.041	7.42E-03
AandB	CB	-0.041	-0.977	0.936	4.160	4.85E-03
AandC	BA	-0.016	-1.335	1.319	5.864	3.38E-06
AandC	BC	-0.016	0.866	-0.882	-3.922	1.12E-02
AandC	CA	-0.016	0.868	-0.885	-3.932	1.09E-02
AandC	CB	-0.016	-0.977	0.960	4.269	3.23E-03
B	BA	-0.151	-1.335	1.184	5.265	5.36E-05
B	BC	-0.151	0.866	-1.017	-4.522	1.22E-03
B	CA	-0.151	0.868	-1.019	-4.531	1.18E-03
B	CB	-0.151	-0.977	0.826	3.670	2.59E-02
BA	C	-1.335	-0.126	-1.209	-5.374	3.28E-05
BA	CandB	-1.335	-0.110	-1.225	-5.446	2.37E-05
BA	bias	-1.335	-0.167	-1.168	-5.194	7.34E-05
BA	previous_choice	-1.335	0.563	-1.899	-8.439	4.58E-12
BA	previous_reward	-1.335	-0.332	-1.003	-4.458	1.57E-03
BC	C	0.866	-0.126	0.993	4.412	1.88E-03
BC	CandB	0.866	-0.110	0.977	4.341	2.47E-03
BC	bias	0.866	-0.167	1.033	4.593	9.22E-04
BC	previous_choice	0.866	0.563	0.303	1.347	n.s.
BC	previous_reward	0.866	-0.332	1.199	5.328	4.04E-05
C	CA	-0.126	0.868	-0.995	-4.421	1.81E-03
C	CB	-0.126	-0.977	0.850	3.780	1.82E-02
CA	CandB	0.868	-0.110	0.979	4.350	2.38E-03
CA	bias	0.868	-0.167	1.035	4.602	8.88E-04
CA	previous_choice	0.868	0.563	0.305	1.356	n.s.
CA	previous_reward	0.868	-0.332	1.201	5.337	3.87E-05
CB	CandB	-0.977	-0.110	-0.866	-3.851	1.43E-02
CB	bias	-0.977	-0.167	-0.810	-3.599	3.24E-02
CB	previous_reward	-0.977	-0.332	-0.644	-2.864	n.s.
CB	previous_choice	-0.977	0.563	-1.540	-6.845	2.54E-08