Adaptive hierarchical representations in the hippocampus

by

Heloisa Santos Costa Chiossi

January, 2024

A thesis submitted to the Graduate School of the
Institute of Science and Technology Austria
in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

Committee in charge:

Prof. Dr. Edouard Hannezo, Chair
Prof. Dr. Jozsef Csicsvari
Prof. Dr. Gašper Tkačik
Prof. Dr. Magdalena Sauvage

The thesis of Heloisa Santos Costa Chiossi, titled *Adaptive hierarchical representations in the hippocampus*, is approved by:

Supervisor: Prof. Dr. Jozsef L. Csicsvari, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Prof. Dr. Gašper Tkačik, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Prof. Dr. Magdalena Sauvage, Leibniz Institut für Neurobiologie, Magdeburg, Germany

Signature: _____

Defence Chair: Prof. Dr. Edouard Hannezo, ISTA, Klosterneuburg, Austria

Signature: _____

Signed page is on file

ISTA Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Heloisa Santos Costa Chiossi

January, 2024

Signed page is on file

Abstract

The hippocampus is central to memory formation, storage and retrieval over many timescales. Neurons in this brain area are highly selective to spatial position as well as to many other variables of the environment. It is believed that the selectivity patterns of hippocampal neurons reflect the structure of tasks an animal performs. However, especially at timescales longer than a few minutes or hours it is not fully known how these representations evolve, nor how they map to behaviour in the process. In this thesis, I monitored the evolution of hippocampal representations in a novel spatial-associative memory task for rats. Reward locations were associated with global sensory cues (i.e. context); animals had to remember the associations and dig for food in those locations only. I used in vivo electrophysiology to record the activity of the hippocampus dorsal CA1 neurons during the learning period of a few days.

I report here a novel and simple method to classify behaviour performance to account for individual variability in learning speed and spurious performance unrelated to true task rule learning. Using this classification I was then able to investigate neural responses on different stages of learning matched across animals. On the first day of learning, I observed a fast formation of single-cell selectivity to task variables which remained stable over days. I also observed that reward tuning was not a single process but dependent on task-related cognitive load. At the population level, a linear decoding approach revealed a hierarchy in the representation of task variables that changed with learning. In the high-dimensional space of population activity, the representation of contexts was specific to each position in the maze, and could thus be better decoded if the position was known. The decoding of position did not improve with knowledge of other variables. As learning progressed, the hippocampal code underwent a reorganisation of high-variance directions in population activity, identified by principal component analysis. I found that dominant dimensions started carrying increasing amounts of information about task context specifically at those positions where it mattered for task performance. When I contrasted this with variables less relevant to task performance (e.g. movement direction), I did not observe differences in decoding quality over positions nor a reduction of dimensionality with learning.

Overall, the largest changes in CA1 neural response with task learning happened in a matter of a few trials; over days, changes undetectable in single-cell statistics were responsible for re-structuring the hierarchy of neural representations at the population level; these changes were task-specific and reflected different stages of learning. This indicates that complex task learning may involve different magnitudes of response modulation in CA1, which happen at specific time scales linked to behaviour.

Acknowledgments

I would like to acknowledge all of those who helped me over the years, directly or indirectly, to complete this PhD. First of all, I thank Prof. Jozsef Csicsvari for giving me the opportunity to pursue my PhD in systems neuroscience in his research group. I thank Prof. Magdalena Sauvage and Prof. Gašper Tkačik, for taking the time to be part of my committee and Prof. Edouard Hannezo for accepting to be my defense chair.

On the academic side, I especially thank those who believed I could do computational work and mentored me along the way: Prof. Gašper Tkačik, Dr. Michele Nardin and Dr. Wiktor Młynarski. I have learned immensely from them and they are the reason I am still passionate about neuroscience. I thank Jago Wallenschus for all the technical support over the years and the mushroom-picking excursions. I thank Rebecca Morse, which did the neural recordings of one of the animals in this thesis, under my supervision. I also thank my colleagues, Yosman Bapatdhar, Lars Bollman, Predrag Živadinović, Andrea Cumpelik and all others, for the discussions, support and company in the lab, office, conferences and beyond.

I should also acknowledge all the rats used in this project and in all neuroscience projects that I reference here. Without them, our understanding of the brain would still be extremely limited. They also provided me with lessons about animal behaviour that I could not quantify or demonstrate in this thesis.

On a personal side, I thank all the friends who have been by my side over the years, who have struggled with me but also who made life fun. I thank Nikola Čanigová for always being there for me, from the very beginning. I thank Daniel Boocock for all the knowledge, the endless patience, support and friendship throughout this entire journey. I thank Kasumi Kishi for all the hours studying physics together, the deep discussions ranging all topics and the adventures abroad. I thank Nikola Konstantinov for cherishing the little things in any occasion, regardless of all the stressful upcoming deadlines. I also thank those who showed up later along the way: Nick Machnik and Dan Grober, whom I can always count on for having a great day. Galien Grosjean, for all the tea breaks and special snacks that powered long days of work. Jacob Lindloff, for bringing so many joyful days into the challenging final year of this work and for all the late night conversations. All the climbing partners, Steve Baratt, Kiki Wagner, Simon Mayer, Micha, Ondro, Saška and so many others, for making my weekends in Austria the best part of my week. All the DnD friends, who could transform a difficult workday into the funniest and greatest of evenings. And, of course, my family, who supported my daunting idea of moving to Austria without knowing anything about the country or IST; and who cared for me all these years despite the long distance. And a special thanks to my grandmother, Ernestina, who is 96 and fulfilled her promise of staying alive to see me become a doctor.

This project received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 665385.

About the Author

Heloisa Chiossi started her scientific career in 2012 at the University of Campinas in Brazil, where she was granted a Pharmacist degree. During her studies, she worked on two neuroscience-related projects under the supervision of Prof. Fabio Papes. The first project investigated the topographical organisation of Periaqueductal Grey responses to different olfactory stimuli. The second project established primary neuronal cell cultures from olfactory tissues of adult mice; these were used for testing the potential of novel pharmacological agents in supporting the differentiation of neural progenitors. In 2014, she spent one year at the University of Bristol in the United Kingdom studying Neuroscience. During this period, she worked under the supervision of Prof. Ute Leonards, investigating human perception of movement. She also worked in the laboratory of Prof. Matt Jones, where she learned techniques such as *in vivo* electrophysiology of behaving rats and the analysis of rat vocalisations. She wrote her final degree thesis, entitled "Brain oscillations, cognition and drug development", under the supervision of Prof. Matt Jones (Univ. of Bristol) and Wanda Almeida (Univ. of Campinas). In 2017, she joined ISTA for her PhD and, in 2018, she affiliated with the group of Prof. Jozsef Csicsivari. Since then, she has been studying the neural basis of learning and associative memory formation in the hippocampus.

List of Publications and Presentations

Conference posters

Chiossi, H. S. C., Csicsvari, J. (2020) Testing generalisation using a linear digging maze. FENS Forum, online.

Chiossi, H. S. C., Csicsvari, J. (2021) Neural signatures of contextual learning strategies. Austrian Neuroscience Association Meeting, Salzburg, Austria.

Chiossi, H. S. C., Csicsvari, J. (2022) Neural signatures of contextual learning strategies. FENS Forum, Paris, France.

Chiossi, H. S. C., Csicsvari, J. (2023) Hierarchical task variable encoding in the hippocampus. Replay @ CUBRIC, Cardiff, UK.

Chiossi, H. S. C., Nardin, M., Tkacik, G., Csicsvari, J. (2023). Behaviourally relevant reorganisation of the hippocampal population code over multiple days of learning. Bernstein Conference, Berlin, Germany.

Presentations

Chiossi, H. S. C., Csicsvari, J. (2021) A novel behavioural task for testing generalisation of memories. IST Neuro Data Talk, Klosterneuburg, Austria.

Chiossi, H. S. C., Csicsvari, J. (2021) Neural signatures of contextual learning strategies. IST Neuro Data Talk, online.

Chiossi, H. S. C., Csicsvari, J. (2023). Hierarchical variable representation in the hippocampus. MITT-ANA Meeting, Budapest, Hungary.

Table of Contents

1. Introduction

The brain can be understood as an information processing network that grants animals the ability to generate complex behaviours in response to internal and external inputs. Whilst sensory and motor areas can be seen as the two ends of this network, processing layers that lie in between them also play a role in shaping behaviour. These layers have an important feature: their inputs and outputs are within the network itself, so interpreting their representations of the external world is a central challenge in neuroscience. The dynamics of these representations and the computations performed on them are believed to underlie learning and memory processes.

In this thesis, I was interested in the hippocampus, one of the brain areas which is described neither as sensory nor motor, but involved in cognitive processing. I focused on how it represents the external world over the timescale of days, as animals learn associations between space and sensory cues. Here I describe previous research into hippocampal anatomy, physiology and function. I also highlight previous approaches to studying learning and cognition in rodents as well as theoretical models of learning and hippocampal dynamics. Lastly, I discuss the unanswered questions in the field I targeted in the present study.

Behaviour tasks to study cognition

To understand the tuning of individual neurons and brain areas to the external world, they have to be monitored during well-controlled behavioural tasks. Human studies of behaviour are largely based on self-reports in tasks tailored to specific human abilities. A common criticism of animal studies of cognition is that animals do not possess the same abilities as humans or that these abilities cannot be reliably assessed. However, animal models allow the use of a broader range of techniques – such as calcium imaging and in vivo electrophysiology – and the observation of the activity of individual neurons as they behave. This is only possible in humans in exceptional cases. As a result, many works in the past century focused on finding systematic ways to measure cognitive and cognitive-like abilities in animal species. In the field of memory research in rodents, a large set of behavioural paradigms are now well established and have elucidated mechanisms of memory formation, consolidation and retrieval.

One of the first well-described mazes for rodents was developed by the American psychologist Edward C. Tolman and it took the form of a relatively complex labyrinth (Figure 1.1a (Tolman & Honzik, 1930). Without any cues inside, the navigation of this maze relies on the animal's capability of remembering the position of rewards and the path to them. It also contains many bifurcation points, where hypotheses about the decision-making can be tested, especially if parts are altered, e.g. by blocking passages and making forced choices.

Decades after the Tolman maze was invented, Professor Richard Morris invented another navigation and memory-based maze in a much less complex form. It involved a pool of turbid water – so that animals could not see inside – and a hidden platform just below water level. Animals were placed in it to swim until they found the platform (R. G. M. Morris, 1981). This paradigm has the advantage of not relying on the animal's hunger or thirst as motivation. Many manipulations of the brain have been studied with this maze, e.g. lesions and pharmacological interventions, but the presence of water precluded its use with certain techniques, such as electrophysiology. A dry version of the maze was later invented for that purpose, called a cheeseboard maze. In this case, a large surface is used, containing small, evenly spaced holes for hiding food (Kesner et al., 1991).

The Morris Water Maze has led to findings about the role of neurotransmitters and receptors in the different stages of memory formation and recall (R. K. McNamara & Skelton, 1993). The cheeseboard maze (Figure 1.1b) allowed for hippocampus findings in the field of sleep and memory replay, through the tracking of the activity of single cells (Dupret et al., 2010; Gridchyn et al., 2020). Other mazes containing a T-junction (Figure 1.1c) were pivotal for the findings in terms of decision-making during navigation (Johnson & Redish, 2007; Singer et al., 2013; Wood et al., 2000). More recently, virtual mazes together with calcium-imaging allowed for the live recording of large populations of neurons in environments that can be easily and quickly manipulated. This enabled, for example, the study of statistical and topological properties of neural populations in the hippocampus (J. S. Lee et al., 2020; Nieh et al., 2021).

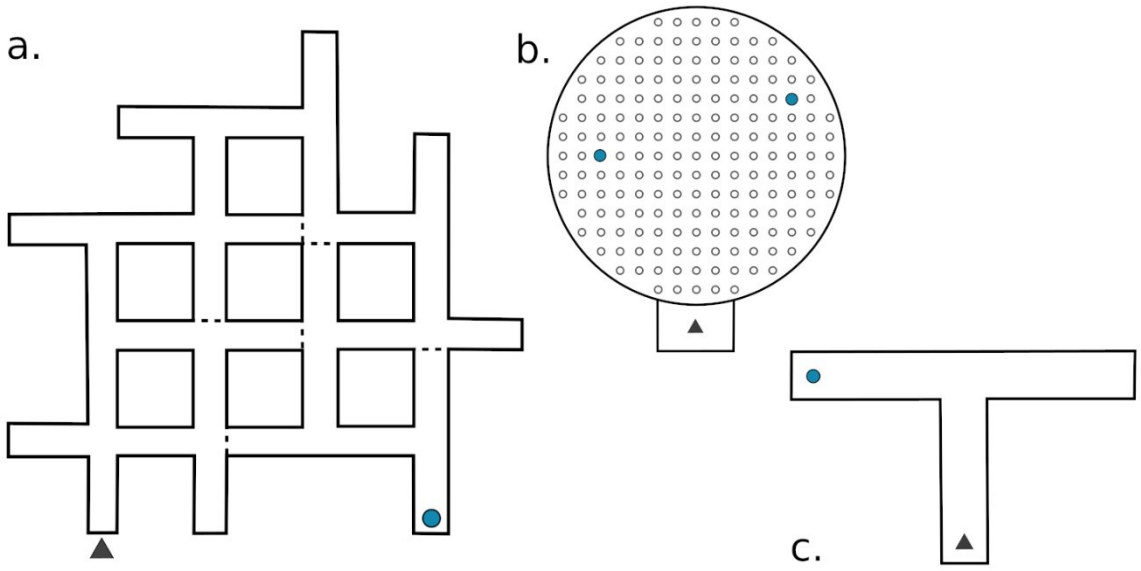

Figure 1.1. Examples of mazes for rats. In all mazes, the start position of the animal is indicated by a triangle and example reward locations are marked by blue circles. a) Tolman maze (Image adapted from Tolman & Honzik, 1930). In this maze the animal has to navigate through several junctions to reach the reward. Dotted lines indicate movable barriers. The experimenter can change the position of the barriers to investigate if the animal remembers the absolute position of the reward or only a sequence of turns associated with one specific configuration of the barriers. b) Cheeseboard maze. Each small circle indicates a small hole in the maze. Animals cannot see which hole contains food until they stand above it. Over time, animals learn to navigate directly from the start box to the reward positions. One or multiple reward positions can be used. c) A T-maze. This maze has a single decision point, at the intesection. The animals start at the bottom arm and need to decide to turn left or right. Many paradigms can be done in this maze, e.g. an alternation task where at each trial reward is placed at the opposite side from the previous trial.

The field of memory research grew a lot also when studying fear. Different from the tasks described above, fear memory paradigms can induce stable memory formation with single exposure events. For example, by inducing a footshock in an environment, future exposures to the same environment lead animals to freeze. One important finding from these experiments was that the fear memory is context-specific and hippocampus-dependent so that a different environment will not lead to the same reaction (Frankland et al., 1998). The concept of "context", however, changes between studies and therefore claims about hippocampal response to context also vary. In the simplest case, each context is defined as a separate environment, and in that case, the hippocampal code can be entirely different between contexts (S. Leutgeb et al., 2004). In other cases, the manipulation is within the same physical space, and what changes are either distal or local sensory cues of the environment (Kim & Lee, 2011; Piterkin et al., 2008). Lastly, context may define more abstract features in an otherwise unchanging environment, let it be a temporary cue (Vaidya et al., 2023; Zemla et al., 2022) or the rule by which the animal needs to use for decision-making, which needs to be inferred without any cue (Kaefer et al., 2020; Wood et al., 2000). In all those instances, it has been shown that hippocampal neurons encode contextual information, albeit in different ways; this shows different levels of sensitivity of the code to various changes in the environment.

One of the great challenges of the approaches mentioned above is finding the right balance between simple and complex behaviour paradigms. Simple paradigms involving a reduced number of variables are easier to interpret and can yield valuable insight into the basic mechanisms of the network. However, they are often not representative of naturalistic behaviour. This implies that conclusions drawn from such studies might not portray the dynamics of the network from when it simultaneously processes multiple sources of information. Complex tasks, on the other hand, can be hard to interpret, since many variables are likely to be highly correlated and their individual contributions to the neural code hard to distinguish. Therefore, there is still a demand for developing behavioural tasks that manipulate sensory and behavioural variables in different amounts so that comparisons between neural activity in each case can be made.

The brain areas involved in cognition

Cognitive tasks engage various brain areas of the mammalian brain, including neocortical areas such as the prefrontal (Euston et al., 2012; Kaefer et al., 2020; Passecker et al., 2019) and parietal cortices (Raposo et al., 2014), as well as periarchicortex areas such as the entorhinal (Eichenbaum & Lipton, 2008; Fyhn et al., 2004) and retrosplenial cortices (Nelson et al., 2014; Smith et al., 2012). The cortex, and especially the neocortex, is enlarged in humans and primates, making it easy to access through invasive and non-invasive research techniques. As a result, previous research has been able to elucidate its role in abilities such as decision-making (Euston et al., 2012; Raposo et al., 2014; Sul et al., 2010), multimodal sensory integration (Alexander & Nitz, 2015; Andersen & Buneo, 2002) and attention (Benchenane et al., 2011).

Located ventrally to the neocortex in mammals – and analogously in other vertebrates – is the hippocampus (Figure 1.2) (Knierim, 2015). The first insights on its cognitive functions date back to a report in 1957, when a patient with lesions to this brain area developed retrograde amnesia (Scoville & Milner, 1957). From then until the present day, advances in experimental techniques allowed discoveries related to its structure, function and interactions with other brain areas. It is now widely accepted that the hippocampus plays a central role in coordinating cognitive processes related to navigation, learning and memory.

During memory formation, sensory information arriving at the cortex is relayed to the hippocampus via its subregion Dentate Gyrus (DG). The DG is densely packed with principal neurons and therefore can assign a different pattern of activity to each experience an animal has (pattern separation). The DG connects to the subregion CA3, where recurrent connections allow for pattern completion: matching features of the current experience to previous knowledge, even in the presence of incomplete information. The CA3 then connects to the subregion CA1, which can integrate information from CA3 together with cortical inputs. The interaction between CA1 and cortical areas such as the prefrontal cortex is thought to be central to memory consolidation (Klinzing et al., 2019; Squire et al., 1984).

Memory consolidation is the phenomenon by which memories become less labile. Information storage about recent experience relies strongly on the hippocampus but over time it is spread over the cortex, as has been confirmed by lesions and other manipulation studies

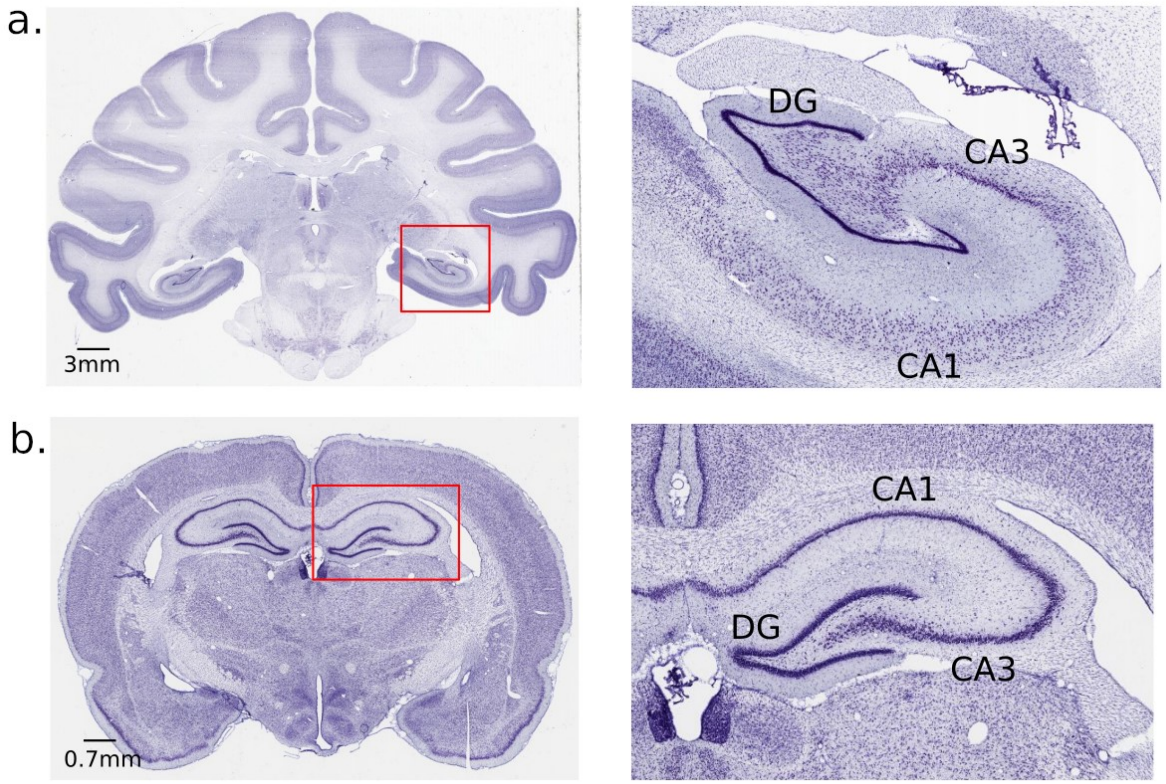

Figure 1.2. Hippocampal anatomy in primates and rodents. Coronal brain sections from the monkey species Macaca Mulatta (a) and rat Rattus Novergicus (b) where the hippocampal region is highlighted by a red square (left). A zoom-in of the hippocampus is shown on the right. Cell bodies were marked using Nissl staining. The Dentate Gyrus (DG), CA3 and CA1 subregions of the hippocampus are also indicated. Images adapted from brainmaps.org (Mikula et al., 2007).

(Frankland & Bontempi, 2005; Nadel & Moscovitch, 1997; Squire, 1992; Tse et al., 2007). This information transfer is believed to happen during sleep and rest periods, through synchronisation between the hippocampus and the cortex. This is when the replay of awake hippocampal activity at compressed timescales is thought to support plasticity between those brain areas. Disrupting sleep therefore has a strong effect on long-term memory consolidation (Klinzing et al., 2019). How much memory retrieval relies on the hippocampus in the long term has been shown to depend on several factors, including time since encoding and memory content (Atucha et al., 2021; Frankland & Bontempi, 2005).

The hippocampal representations of behaviour

The hippocampus is also known to play a role during behaviour, due to its navigation-related activity. Especially in the dorsal portions of the CA3 and CA1 subregions, a subset of neurons – referred to as "place cells" – display a strong modulation of firing rate by the animal's position in the environment (O'Keefe & Dostrovsky, 1971). The activity of those cells not only indicates the animal's current location, but in specific situations can also inform about past or planned future trajectories (O'Neill et al., 2010; Stachenfeld et al., 2017). This uniqueness of firing at each location is thought to "index" experiences to specific points in space as well as aid decision-making brain areas to choose appropriate behaviour outputs given the current state of the animal.

Neurons in the dorsal CA1 do not respond solely to an animal's position. First discovered in the postsubiculum area of the hippocampal formation, another known modulator of hippocampal activity is head direction, which is driven by the anterior thalamus (Goodridge & Taube, 1997; McNaughton et al., 1983; Taube, 1995; Taube et al., 1990; Viejo & Peyrache, 2020). Dorsal CA1 is also known to carry information about goals and rewards (Dupret et al., 2010; Gauthier & Tank, 2018; Hollup et al., 2001; L. Zhang et al., 2022), through dopaminergic modulation from regions such as the Ventral Tegmental Area (Mamad et al., 2017). More interestingly, CA1 neurons can also respond to a complex set of environmental variables (mixed selectivity). These variables can be physical barriers, sensory cues and changes in environment shape and size (J. K. Leutgeb et al., 2005; R. U. Muller & Kubie, 1987; Shapiro et al., 1997; Wiener et al., 1989). Theoretical studies show that mixed selectivity neurons grant coding advantages to the hippocampus (Rigotti et al., 2013).

A result of mixed selectivity is a modulation of the spatial code by the context the animal is in. For example, neurons can change their preferred firing position entirely between two environments, a phenomenon called global remapping (S. Leutgeb et al., 2004). Other types of change in sensory inputs (e.g. visual, tactile or auditory) can lead instead to a modulation of the maximal firing rate at the neuron's preferred location, which is referred to as rate remapping (J. K. Leutgeb et al., 2005). Even when maintaining all sensory cues, the firing of CA1 neurons can be modulated by the animal's upcoming decision to turn right or left (Wood et al., 2000). Some studies suggest that time itself is also encoded as a variable (Kraus et al., 2013; MacDonald et al., 2011). Even when position and sensory cues are maintained, if the time towards the goal is

progressing and the animal is not resting or distracted, a clear sequence of firing related to the elapsed time can be observed (Kraus et al., 2013). There is evidence that time creates correlations between encoded memories; that means that not only overlapping sensory features of experiences dictate the similarity in their representation but also their proximity in time (Ziv et al., 2013).

The representational feature of the dorsal hippocampus suggests it has a complementary role to the cortex during behaviour and that the relationship between those brain areas goes beyond memory consolidation. Neurons in the prefrontal cortex and other cortical areas show a more general representation of the task than the hippocampus, which includes even more complex mixed selectivity neurons (Rigotti et al., 2013), representation of task rules (Wallis et al., 2001) and reward contingencies (Riceberg & Shapiro, 2012). Cortical areas coordinate with the hippocampus in many ways: the synchrony of the prefrontal cortex to the hippocampal theta rhythm in working-memory tasks is related to cognitive flexibility (Benchenane et al., 2010; Zielinski et al., 2019); the entorhinal cortex grid cells are the main input to the hippocampal place cells (Fyhn et al., 2004); the posterior parietal cortex and retrosplenial cortex integrate visuospatial and proprioceptive inputs to support ego and allocentric navigation (Alexander & Nitz, 2015; Whitlock et al., 2008); the retrosplenial cortex complements the hippocampus in contextual coding (Smith et al., 2012). Moreover, memory consolidation and neural activity during task performance are not totally unrelated. Given the increased participation of the cortex in memory storage over long periods of time, many also believe that performing learned tasks requires proportionally less of the hippocampus over time (Frankland & Bontempi, 2005). However, the dynamics of hippocampal representations of behaviour variables in the long term remain unclear.

The timescales of learning

Learning is a process that involves brain changes in many timescales. On a cellular level, learning-induced plasticity can happen on a scale of seconds to minutes (Zucker & Regehr, 2002) and can be long-lasting (Bliss & Collingridge, 1993; Dudek & Bear, 1992). Short-term plasticity involves changes in the probability of neurotransmitter release or the insertion of receptors in the synaptic membrane, whilst long-term plasticity involves mechanisms such as de novo protein synthesis. At the behavioural level, learning takes minutes, hours or even days, so other neural processes must involve these timescales. On what concerns this thesis, when I refer to learning I imply a behaviour process that leads to the formation of a memory at the network level, and this memory can then be used for improving behavioural responses.

On the fastest side of task learning is working memory, where changes are mostly transitory (Miller et al., 2018) and fear conditioning, where changes are quick but long-lasting (Izquierdo et al., 2016), supported by fear-related inputs of the amygdala to the hippocampus. Also episodic memories – or episodic-like in the case of rodents and other animals – are acquired in a single event. An episodic memory can be understood as the cohesive memory of a specific event, including the time (when), the context or location (where) and the actual event (what).

The concept was defined for humans, but animals have also shown similar types of memory (Babb & Crystal, 2006; Clayton & Dickinson, 1998; Tulving, 1972).

On the other side of the learning spectrum is learning through repetition. Instead of acquiring a detailed memory of a single event, we often learn rules and associations through repetitive exposure to events that share a structure or feature that guides our behaviour. One concrete example is associative learning; in experimental paradigms, this means that the subject needs to find a correct pair – of two images or between a sensory cue and a location – to be able to receive a reward. If not explicitly told what the pairings are, many attempts of possible pairings are necessary to learn which ones lead to reward. The difficulty of the task can arise mainly from two sources: the number of features one can pair and the timing between the association and reward. If this timing is too large, the causal relationship between the task and reward might not be perceived.

In rodents, associative tasks are commonly combined with working memory paradigms. A stimulus is provided to the animal at the beginning of each trial and stimulus identity defines the appropriate choice to maximise reward. Stimulus presentation and choice periods are separate in time and the animal needs to store information about the recent stimulus during this interval. Studies using these paradigms often do not target the period of learning the associations themselves, but the working-memory periods in trials after they are learned. They have allowed for discoveries on the role of both the hippocampus and the prefrontal cortex in sustaining memory of the stimulus throughout the delay period before decision-making (Yoon et al., 2008). Associative memory paradigms that do not require working memory have also been used to show the different roles of the hippocampus in recognition memory and familiarity (Sauvage et al., 2008).

At longer time scales are associations in complex task structures, referred to as schemas. In one reference rodent study in the field (Tse et al., 2007), rats were required to associate smells and reward locations in an open arena. The learning of associations took the animals many training sessions. However, once the initial pairings were known, new odour-location pairings were quickly learned. The experimenters believe that understanding the concept that odours and locations are associated – the schema – defines the slow part of the learning curve. When new odours are introduced, the animals already expect a reward location to be associated with them; so within a single session, they are capable of discovering it. This study shows that 24 hours after initial learning hippocampal lesions have a detrimental effect on memory, but not 24 hours after the new set of odours are learned. They conclude that once the animal has a schema for the task, consolidation of new associations is much faster than at first learning.

Models of learning

Many theories have been devised to explain how animals structure their knowledge as they learn a task. One set of theories assumes decisions are based on previously experienced action-value pairs (model-free), while others assume the animal builds a mental model of the entire task and acts based on it (model-based). In model-free models, decision-making is quick,

since the animal can simply choose the action that leads to the highest outcome. The issue, however, is that any changes to the environment require recomputing all action-value pairs, which is inefficient if the environment is dynamic. In model-based learning, the animals are more flexible, as they are able to choose the best action from those never used before. The drawback is that building an internal model containing all possible action paths and making a decision over them is costly, as it requires a lot of experience.

One model that lies between those two categories is that of successor representations (Stachenfeld et al., 2017). In this model, decision at any time point is made by trying to maximise the expected value of future states, weighing by their proximity in time. The probability of future states is computed separately from the reward probability at each state. This means that it can be learned quickly, as in the model-free case, but it can readily respond to dynamic environments, as in model-based. Successor representations predict several phenomena, for example, the backward shift of hippocampal place fields in 1D environments with experience (I. Lee et al., 2004). If extended to contextual learning, one could assign different expected values to each state depending on context. Matching this idea, one experimental result shows that the probability distribution of hippocampal place fields in a linear track is specific between the start position and the reward position of each given context (Zemla et al., 2022).

Other theoretical models accounting directly for contextual learning in the hippocampus include Bayesian models with latent state inference (Fuhs & Touretzky, 2007; Gershman et al., 2010) and reinforcement learning models (Hasselmo & Howard Eichenbaum, 2005; Redish et al., 2007). For example, Fuhs and Touretzky (Fuhs & Touretzky, 2007) attempt to unify different definitions of context through general neural mechanisms. They define context as a set of experiences close in time that share sensory inputs and behaviours. In the brain, contexts contain activity patterns of groups of neurons ("states") that are likely to transition between each other, whilst the probability of transition between states from different contexts is low. Contexts in this model form "latent attractors", which I mention in the next section. Using Bayesian inference in a Hidden Markov Model they can determine when new contexts will be separated in terms of neural activity; this allows predictions about the different levels of hippocampal remapping observed in experiments. Also in agreement with experiments, it can explain the distinction between "context learning", the slow learning that more than one context exists through multiple experiences, and "context inference", the fast identification of the current context once learning has already happened.

Neural population and dimensionality

In the learning models above, neural mechanisms for information storage and processing cannot always be assigned to changes in individual neurons but are better understood at the neural population level. Responses of individual neurons to external variables are often measured via the correlation of firing rate to those variables. This can give reliable predictions in sensory cortices but can be hard to interpret in higher cognitive areas. Methods to study population dynamics have been developed to target this issue, especially in the motor cortex –

due to the rise of brain-computer interfaces for prosthetics. In this field, not only the dynamics of the population in multiple timescales are relevant, but also the identification of parameters stable over days and between subjects. For instance, the concept of a neural manifold was created to describe the low-dimensional space that constrains the neural representation or neural trajectory, and that can be identified regardless of trial-to-trial variability and noise (Gallego et al., 2017; Nieh et al., 2021). This implies a redundancy of the neural code regarding behaviour tasks since a few behavioural variables are represented by thousands or millions of neurons.

The topic of dimensionality of the neural code has received a lot of attention in recent years. Research on decision-making tasks shows that the hippocampal activity maps to a low-dimensional manifold which is similar among animals and has a geometry relevant to the task (Bernardi et al., 2020; Nieh et al., 2021). However, these studies do not track how the geometry changes with time. It is known that the majority of hippocampal place cells change their preferred firing position as the days go by (Ziv et al., 2013), which is referred to as representational drift. This drift is not fully orthogonal to the representational manifold and therefore will affect its geometry, posing a question of how the manifold is used by the animal during learning. Despite the representational drift, animals are capable of learning and spatial information does not decrease over time (Vaidya et al., 2023). An insight into this comes from a study showing that natural constraints to the drift allow for linear decoding of information in downstream areas (Rule et al., 2020). They also show that plasticity mechanisms are sufficient to compensate for the cumulative effect of drift over time. Although the brain is a highly non-linear system, it operates under the constraints of time, energy and number of neurons; a linearly decodable output is the simplest and arguably one of the most efficient codes to be read out by downstream areas (Karpas et al., 2019).

Not only manifolds have been studied at the neural population level, but also normative theories have been used to explain and predict hippocampal population activity. Many models of memory are based on attractor dynamics, in which each memory itself can be understood as an attractor in the high dimensional space of neural activity. Attractors are stable subsets of all possible neural activity patterns, to which the dynamics of the population are driven. These theories can explain pattern completion in CA3 (Treves & Rolls, 1994) and the remapping between environments by CA1 place cells (Battaglia & Treves, 1998; Wills et al., 2005). Attractors grant the network robustness to noise if they are lower-dimensional than the neural activity space (Khona & Fiete, 2022). Attractor network models also allow for predictions in terms of the limit capacity for memory storage, which explains the narrow tuning of hippocampal place fields as well as other features of hippocampal representations (Battaglia & Treves, 1998). Apart from limit capacity, other constraints of the hippocampal network have been investigated using information theoretical approaches (Treves et al., 1996).

2. Aims of the study

In the introduction, I highlighted some of the extensive research from the past decades on hippocampal physiology and function, and some of the questions about this brain area that remain unanswered. Hippocampal representations of the environment have been mostly studied in the context of exploratory behaviours or after learning of goal-oriented behaviours. Less attention has been given to the dynamics of these representations during the learning period itself; it is also not clear how representations are related to animal performance over extended periods of time. Therefore, the main goal of this thesis was to understand how relationships between multiple variables are represented in the hippocampus as the animal learns and how they can support behaviour. I also looked for general mechanisms that can be employed by the brain to shape neural representations according to task demands. To achieve my goal, my specific aims were the following:

1. Compare the statistics of CA1 single-cell tuning throughout different associative learning stages
2. Understand the relationship between the CA1 population code structure and learning

To accomplish these aims, I focused on a learning period of days, as is often how associative learning happens in naturalistic behaviour. In the upcoming chapters, I targeted these aims by the following approaches:

(1) I designed a behavioural task for rats involving spatial and non-spatial learning over many days: animals had to remember the association between global visual-tactile cues and the location of food to be dug in a linear maze. I also developed an analysis method that accounts for individual variability in learning. It is based on a hierarchical clustering algorithm and it classifies training sessions into different learning stages based on animal task-solving strategies.

(2) I electrophysiologically recorded the activity of hippocampus dorsal CA1 neurons as the animals learned. I investigated the cell firing correlation of pyramidal neurons and interneurons to single or multiple task variables over time. I also investigated how task demands influenced the coding of rewards.

(3) I investigated how encoding of information in the population evolves with learning in ways not directly predictable by single-cell changes. I decoded task variables from neural population activity at different learning stages: this allowed me to compare the decoding of task-relevant and irrelevant variables over time, and measure how dependencies between the representation of different variables could affect the overall population code.

3. Methods

M1. Experimental methods

Animals

All the results shown here come from 5 Wild-type, Long Evans Rats of 9 - 13 weeks of age and weighing 300-350g. Animals were caged with littermates before implantation surgery and isolated afterwards. They were kept in a dedicated animal room under a 12 hour light/dark cycle. All animal procedures were carried out in accordance with the Austrian federal law for experiments with live animals, under the project licence number BMBWF-66.018/0018-V/3b/2019.

Tetrode microdrive

Tetrode microdrives were built using a 3D printed plastic case, metal cannulae and four-channel electrodes (tetrodes) made of 4 x 12µm tungsten wires, which could be moved by rotating a small screw at the top of the case. In this thesis either 16 or 32-tetrode microdrives were used. Sixteen tetrodes were placed in the hippocampus and, if a larger microdrive was used, the other half was placed on the retrosplenial cortex, as described below.

Microdrive implantation surgery

The surgery protocol implemented here has been used in previous research projects (Kaefer et al., 2020; Xu et al., 2019). The entire procedure was performed with animals under isoflurane anaesthesia. Animals received rehydration with a saline/glucose solution mix every 2 hours during the surgery. Analgesia was achieved with Metamizol (Novalgin) for quick initial action and buprenorphine for long-term action over the entire duration of the surgery. Medication dosage was calculated according to the animal's body weight and internal guidelines based on the Austrian animal law.

After shaving the hair over the animal's head, the skin surface was covered in iodine solution and a scalpel was used to make a straight cut that exposed the surface of the skull. With the help of a stereotaxic device, the craniotomy coordinates over the hippocampus were marked. The hippocampus coordinates used were the following (from Bregma, positive ML value indicates the right hemisphere):

- Animal ID jc233 - AP: -2.5 to -4.5 and ML: -0.75 to -3.75
- Animal ID jc243 - AP: -2.7 to -4.6 and ML: -2.0 to -4.3 (15^{o} angle)
- Animal ID jc250 - AP: -2.5 to -4.5 and ML: 1.0 to 4.0
- Animal ID jc253 - AP: -2.5 to -4.5 and ML: 1.0 to 4.0
- Animal ID jc259 - AP: -2.5 to -4.5 and ML: 1.2 to 4.2

Animal jc243 had the microdrive implanted at a 15° angle from the midline, all others had it implanted at a straight angle (perpendicular to the skull). Two animals (jc233 and jc243) also received a second craniotomy over the Retrosplenial cortex:

- Animal ID jc233 - Craniotomy 1 AP: -1.5 to -5.25 and ML: -0.25 to -0.75
  Craniotomy 2 AP: -5.5 to -7.5 and ML: -0.25 to -2.25
- Animal ID jc243 - AP: -2.7 to -6.3 and ML: -0.6 to -1.5 (15° angle)

A dental drill was used for making a small hole on the skull following these coordinates. Using forceps, the dura mater was then removed so that the tetrodes could be lowered in the centre of the craniotomy with the help of the stereotaxic device. During the surgery tetrodes were lowered 1.0 to 1.3 mm into the brain; the remaining depth to reach the CA1 stratum pyramidale was achieved by tuning tetrode depth on the days following surgery. The tetrodes were then covered in a wax-oil mixture (3.5:1) to ensure smooth vertical movement, followed by a layer of bone cement connecting the body of the microdrive to the surface of the skull. After surgery, animals were allowed to recover for at least 7 days before the start of behaviour training.

Maze design

The maze consisted of a raised S-shaped platform, 10 cm wide and 360 cm long. The shape was formed by 5 arms of equal length (80 cm, each overlapping 10cm at the corners) and 90° angles between them. It was surrounded by 30 cm-high transparent walls along the track and a square box on either end, each 20 cm wide surrounded by high black walls and a door to the maze. Each maze arm contained 7 holes 10cm apart (15 cm from the start box). Holes were 5cm wide and 5cm deep and were filled with playground sand. All sand was burned in a lab oven before use, to avoid contamination. Extra tiles could be added to the maze surface, changing its texture and colour and covering subsets of the holes. In this experiment, the possible tiles (determining context) were two: a black, hard plastic tile with equidistant grooves of around 1mm; and a white, EVA-foam sheet tile with a smooth surface. Fixed distal cues were hung around the room to facilitate distinguishing the overall position of the maze.

Habituation and training

Before surgery, a plastic well similar to the wells/holes in the maze was placed inside the animal's home cage, containing cornflakes covered in sand. This was used to habituate animals to digging for food. Initially, it contained a large number of flakes and over the days the ratio of flakes vs. sand was reduced down to a few flakes in the entire well. Once the animals recovered from surgery, they entered a food restriction protocol, aiming to keep them 85-90% of their original weight. The sand well was removed from the cage and they were then habituated to digging in the maze itself.

During pre-training, no contextual cues (floor tiles) were put on the maze. The training started by making only two out of the 35 possible wells available to the animal. Each would

contain a mix of cornflakes and sand, and the flakes would be visible on the surface. The number of flakes was progressively reduced until only a few were available; they were also made progressively less visible until they were buried around halfway deep into the sand. Once animals consistently dug the buried reward, a task-like scenario was introduced, where more wells were available, but not all contained food. Animals had to remember which well contained the food (consistent within a day) and after obtaining two pieces of food, they were expected to run to the end of the maze opposite to the start position. At the end box, they received an extra sucrose pellet. The electrophysiological recordings started once animals were capable of consistently performing this task without any interference from the experimenter.

Behavioural task

Recording days started with a 40 minute sleep session, followed by 40 trials of the task. Context (surface tiling of the maze) was changed every 5 trials so one day consisted of 8 blocks of 5 trials. The starting position (left/right box) was defined by a pseudo-random algorithm that made sure that both boxes were used in equal amounts in either context on a given day and not more than 3 trials in a row, as longer sequences delayed learning. After the 40 trials, animals had another 40 minute sleep session, followed by 4 probe trials.

On a single trial, animals had to use the colour and texture of the maze as cues to decide where one of the rewards would be located (context-dependent, referred to as A or B). A second reward would be found at a fixed location, equal in both contexts, so independent of the cues (context-independent). A total of 8 sand-filled wells were available at each context, 4 of which overlapped between contexts in terms of position (the three rewarded wells plus a control well). Animals were only allowed to dig two wells in a trial; after the second choice was made all other wells were covered and made unavailable. Nevertheless, animals were allowed to go back and forth on the track and dig the chosen wells as many times as they wanted, for up to 5 minutes. Only one corn flake was made available at each reward well, to make sure animals would be motivated to run 40 trials a day. Animals terminate a trial by running to the box at the opposite end from which they started, where they would find a sucrose pellet if the two holes dug were correct. In case of an incorrect choice, there was no reward at the end box. The door to the end box was open only once the animal made two choices. If the animal did not complete the task in 5 minutes the trial was also terminated and no further reward was provided in the trial.

On the first day of training the corn flake was made visible and animals were allowed to dig as many holes as they wanted for the first 2 trial blocks; in the following 2 blocks, flakes were hidden superficially and only two choices were allowed. This protocol was necessary because omitting the food from the first trial demotivated the animals. For the rest of the first day and also the remaining days the food was placed halfway deep into the sand-filled wells and only two choices were allowed. A trial was considered correct when the animal chose to dig the two correct rewards for the current context.

The probe trials consisted of one trial of each category (defined by start side and context), each of 5 minutes. In those trials, the food was placed at the very bottom of the reward well,

making it very difficult for the animal to find it. This allowed us to measure their motivation and confidence about the location of the food. During probe trials, animals were allowed to dig as many holes as they wanted. Trials were terminated after 5 minutes or if the animal ran to the end box.

As a control for this task, I made sure that no residual odours could guide the animal's choice, neither odours from the food nor the animals themselves. Corn flour was mixed in small amounts in the sand and the sand from all the wells was frequently re-mixed together and distributed across all wells again to spread the smells evenly. The maze was also regularly cleaned with alcohol. To test for smell preferences, control experiments were done by hiding cornflakes in random locations and checking for behaviour. Animals never dug such locations and preferred the memorised reward wells.

Electrophysiological recordings and position tracking

In vivo extracellular electrophysiological recordings were made using an Axona Ltd. recording apparatus. The implanted microdrive in the animal's head was plugged into the recording system via a headstage connected to a tether. The system included a pre-amplification of the analogue signal directly at the animal's head, then a second step of amplification and digitization of the signal before relaying it to the computer. The neural signal was recorded at a 24kHz rate. For position tracking, the animal's headstage contained a pair of LEDs, one on each side. The LEDs were of different sizes, to allow identification of the left and right sides of the animal. Using an overhead camera and low light conditions, the system captured the position of the LEDs throughout the task, at a 50Hz rate.

M2. Data Processing

Spike extraction and spike sorting

Spike extraction and automated spike sorting of neuronal spikes were performed using the MountainSort algorithm (Chung et al., 2017). The output from the algorithm was manually curated to remove artefacts, clusters violating the refractory period, clusters with non-biological waveshapes and those not recorded for the entirety of a session. I also manually checked for clusters that needed to be merged or split. Those generally were the result of firing bursts or tetrode drift, processes not fully handled by the algorithm.

Linearization of tracking data

The 2D spatial data of animal movement in the S-shaped maze was transformed to 1D coordinates by a series of steps. First, the pixel position of the maze boundaries were annotated and each of the 5 maze arms was defined as a polygon. Distortions due to the camera position

were corrected, making all arms exactly 80cm. In the next step, the polygons were concatenated by the short axis according to their order to the maze. Since corners overlap 10cm between maze arms in the real maze shape, this was also the case when calculating the coordinates. The short axis was then collapsed and for all analysis, I used only the linear coordinate along the long axis. The maze contained 5 arms of 80cm each with 4 overlapping regions of 10cm each; the end result was linear position data ranging from 0 to 360cm. To compensate for missing tracking data, missing data of less than 3 seconds was replaced by a linear interpolation between the position before and after the missing frames. Position values were also smoothed by using the mean tracking values around each frame, in a window of 5 frames. All off-maze tracking was labelled as such.

M3. Analysis Methods

M3.1 Behaviour

Behavioural performance and error types

Behaviour in a trial was considered correct if animals dug exactly the two rewarded wells of the current context and ran to the box at the opposite end of the maze from where they started. This information was annotated by the experimenter. In case the animal did not dig correctly, the errors were classified as:

1. *Wrong context:* the animal dug the reward of the opposite context to the current trial.
2. *Incomplete:* the animal did not dig at all or dug only one well throughout the maximal duration of the trial (5 minutes). This was usually a sign of demotivation.
3. *Random dig:* the animal dug a hole that was not rewarded in either context.
4. *Other:* the trial was incorrect for a reason not described by the options above.

Reversals

A behaviour was considered a reversal if animals turned around and ran in the direction they originally came from. Trials always ended at the opposite end from where they started, so each reversal was paired with a second one towards the main trial direction. I counted each pair as a single reversal. To quantify them, I first calculated the discrete derivative of the linearised position to determine the movement direction in each frame. Denote with l_t the linearized position at time t, then the discrete derivative is defined as:

$$d_t = l_t - l_{t-1}$$

Left and right movement were determined by negative and positive derivatives, respectively. A zero derivative indicated no movement. A reversal event was defined if two consecutive bins had opposite signs. The total number of reversals in a trial was defined as the sum of all reversal events divided by two (pairing).

Hierarchical clustering of behavioural strategies

To cluster the recording sessions in terms of behaviour, I determined the likely strategies the animals could be using to solve the task on each given trial:

1. Eager: the animal dug whichever 2 of the 3 reward wells came first, depending on the side they started, and regardless of the contextual cues.
2. Win-stay-lose-shift: the animal dug the exact same well they dug in the previous trial if that was a successful trial, and a different combination (overlap of one is allowed) if the previous trial was incorrect. For the first trial of the day, it is not possible to assign evidence to this strategy, given that there is no previous trial for comparison.
3. Contextual: the animal was aware that one well was always rewarded (context-independent) and knew the association between the contextual cues and the second reward. Therefore, the animal dug only the two rewarded wells of the trial.

Each trial was assigned a 1 or 0 for each strategy, depending if the behaviour was evidence for the strategy or not. For a single trial, there could be evidence for more than one strategy. One example of this evidence assignment would be the following: assume the order of rewards starting from the left side is Indep - A - B; if the animal dug Indep and B on the previous trial successfully, and now dug Indep and A when running from the left, the assignment on this trial would be: Strategy1 = 1, Strategy2 = 0 and Strategy3 = 1.

For an entire session, I averaged assignments over all trials, obtaining the mean evidence for each strategy. These values determined the features used in the hierarchical clustering, which worked as follows:

Initially, each point represents a cluster of size 1 each; clusters are merged recursively following the Ward variance minimization criterion. In detail, denote a cluster as $X = \{x_1, \ldots, x_k\}$, and compute its mean $E(X) = \frac{1}{k}\sum_{i=1}^{k} x_i$ and its variance $V(x) = \frac{1}{k}\sum_{i=1}^{k} (x_i - E(X))^2$. At each step, the algorithm computes the variance of each possible merged pair of clusters $V(X_i \cup X_j)$, and selects the pair to merge that yields the lowest increase in total variance. This procedure is computationally demanding; a simplified algorithm with linear complexity in the number of points has been proposed by (Müllner, 2011) and is implemented in the Scipy library cluster.hierarchy.linkage which was used here for analysis. This process is iterated over until all points are clustered together, and that is the root of the tree. Once the tree is built, different cut-off points in the hierarchy yield different numbers of clusters, as discussed in the main text of Chapter 4.

Refrain behaviour

A behaviour was labelled “refrain” when animals successfully held back from eagerly digging the incorrect reward well. More specifically, when the incorrect contextual reward well

was positioned before the correct one, given the trial running direction. To quantify this, I measured the amount of time spent in the first passage over the incorrect well. Time was calculated as the difference between entry and exit times within a radius of 10 cm from the centre of the well.

Digging and non-digging periods

Digging periods were estimated based on animal position and behavioural observations made by the experimenter. A digging period was determined by frames when the animal was within a 10 cm radius from a well dug and no running movement was detected (speed < 3m/s). This means that the time included both digging and reward consumption if they happened at the same place. I termed this "time around reward". This also implies that "trial time excluding digging" includes immobility periods, but away from the dug rewards.

M3.2 Single cells

Firing rate maps

Ratemaps were calculated per cell and category. A category was defined as the combination of trial context and starting side. I only considered periods of movement (>3cm/s) in which the animal was moving in the preferred direction of the trial (e.g. moving left if the trial started on the right start box), to control for the direction selectivity of place fields in linear environments. The maze was divided into 4 cm bins; for each bin, the rate was calculated as the sum of all spikes that happened in that position divided by the occupancy (time spent in that bin). I then smoothed the place fields by convolving each rate map with a Gaussian kernel (sigma=2).

Place cell classification

A cell was defined as a stable place cell if it fulfilled a set of parameters:
(1) Mean firing rate above 0.25 Hz and below 6 Hz. Low firing rates yield unreliable rate maps, and too high firing rates indicate that the clustered unit is most likely to be a putative basket cell, not a pyramidal neuron.
(2) Maximum sparsity of 0.5. If we denote r the rate map vector and $\bar{r}$ the mean firing rate over all positions, the sparsity was calculated as follows (Skaggs et al., 1996):

$$sparsity = (\bar{r})^2 / \overline{r^2}$$

(3) Minimum spatial information of 0.8 between firing rate and position. If we denote $\boldsymbol{r}$ the rate map vector and p the number of position bins, the spatial information was calculated as follows (Skaggs et al., 1992):

$$SI = \sum (r * log_2 r / \bar{r} * p)$$

(4) Minimum stability of 0.8. Stability calculation is defined in the next section.

Place field stability

Place field stability describes how correlated is the activity of a neuron in space across different passages of the same location. Therefore, I first determined the first trial in each category in which the place field appeared. Given that in each session there were only 10 trials per category, stability could only be reliably calculated if the place field appeared before the 5th trial in the category. The trials after appearance were then separated into two random halves, for which I calculated the mean rate maps and their correlation. I repeated this procedure 5 times. The stability score was calculated as the mean correlation over all repetitions.

Firing rate gain

The average firing rate gain per position in each trial category was defined as the mean of the normalised rate map over all pyramidal neurons. If we denote r_i the rate map vector for each neuron i and n the number of pyramidal neurons, it was calculated as follows:

$$gain = \left(\sum_i r_i / \overline{r_i} \right) / n$$

Identification of place fields

To detect place field region(s) in a rate map, I identified consecutive spatial bins with firing rate above two standard deviations from the mean firing rate. In a single rate map, it was sometimes possible to identify more than one region that fulfilled this criterion, indicating the presence of multiple place fields. The centre of each place field was then defined as the point of zero derivative inside such regions and the edges as the position on either side of this centre where the firing rate returns to less than one standard deviation above the mean firing in the entire map. The distance between the two edges determined the place field width. The width I report here was calculated as an average over the field widths calculated in the trial-by-trial rate maps; it was not calculated directly on the average rate map to avoid overestimation due to the place field shift over trials.

The shift of the place fields depends on the direction of movement of the animal (see main Figure 5.4a), so the same place field has a slightly shifted centre in the mean rate map from left vs right trials, if present in both. To account for that and be able to compare the place field across categories, I matched the place fields identified in each trial category by considering fields to be the same if the peak of a field in one category fell within the edges of the field in the other category, even if their centre position did not perfectly match. If fields did not match between rate maps, they were considered as different fields.

Remapping

Two types of remapping have been considered, global and rate remapping, as previously described in (S. Leutgeb et al. 2004; J. K. Leutgeb et al. 2005). Global remapping measures the extent to which place fields completely reorganize from one condition to another, whereas rate remapping measures the extent to which place fields that remain in the same position change their peak firing rate. In the current study, the conditions I compared were the trial categories, defined by the direction and context of each trial.

To quantify global remapping between pairs of categories, for each cell I calculated the Pearson correlation between the average firing rate map in each category. I made sure to keep one variable fixed for this comparison – e.g. if I compared contexts, I only compared maps of the same movement direction. I then calculated the distribution of correlations between all pairs of trials within a single category. Low correlations across categories indicated remapping. To establish if each single cell significantly remapped across categories, I compared its correlation across categories c with the distribution of within-category correlation coefficients. Specifically, I used a z-score test, as follows: denote with μ, σ the mean and standard deviation of the within-category distribution. I then computed the z-scored coefficient as $z = \frac{c - \mu}{\sigma}$, and set a threshold of $z > 1.96$ (i.e., $p < 0.05$) to consider a cell significantly remapped.

For the rate remapping, I calculated a separate score for each place field within a rate map. The identification of place fields is described in the previous section. The remapping score was calculated as the difference in mean place field firing in each condition divided by their sum. If we denote $\overline{r_x}$ and $\overline{r_y}$ the mean firing rates of a field in categories *x* and *y*, the score is calculated as follows:

$$score = (\overline{r_x} - \overline{r_y}) \,/\, (\overline{r_x} + \overline{r_y})$$

As in the global remapping case, I compared the score distribution within and across categories to define a significance threshold. I considered that a cell rate remapped only if it did not globally remap and at least one of its fields had a significant rate remapping score. The reason for the separation of fields was that many place cells displayed multiple place fields and it was not clear *a priori* that all fields would behave the same when the categories changed. Moreover, opposing changes in different fields would lead to a low score and therefore conclusion that no rate remapping occurred when it actually did, but locally.

In-field peak rate difference

Rate remapping scores measured changes in mean firing between conditions but did not give information about its magnitude. Therefore, for each field region, I also compared the magnitude of remapping as the ratio between the peak firing rate inside the Gaussian-smoothed field in each condition.

Place cell reward firing

A cell classified as a place cell in a trial category was considered "active at a reward" if its average firing in the bin nearest the reward was at least 2 standard deviations above the average firing in the entire rate map. As a control, I shifted the rate map of each neuron independently and by a random amount and then calculated the same measure again. I repeat this procedure 100 times to calculate the expected percentage of place field in any location if they did not show any preference for specific location. This is denoted as "shuffle" in Figure 5.6b.

Place field distance to reward

For each place field I calculated the distance between the centre of the place field and the centre of each of the 3 reward locations. These distances were used to determine which is the closest reward to each place field. Given that reward positions are not evenly distributed throughout the maze, I also calculated the distribution of place field distances if place field centres were uniformly distributed in the maze, assuming the same reward positions as in the data. This procedure was repeated 1000 times, to have an estimated distribution of distances for each reward in the uniform place fields case. This distribution could then be compared to that of the data; I used F-statistics to compare the variances of the two distributions.

Place field shift calculation

Since many of the place cells recorded displayed more than one place field, the shift analysis was performed on each place field and trial category separately. For the shift calculation, the rate map was first calculated per trial and the peaks were identified in the same fashion as for the mean rate map, but with a looser threshold (1.5 std from the all-day mean). If a trial contained a peak within the place field region, it was assigned to that place field. The shift slope was then calculated using a linear regression between place field peak position and trial relative number (in the category). Trials in which there was no peak were not included as data points. For a more robust fit, a leave-one-out procedure was used, where I left one data point at a time and re-calculated the regression. The shift was defined as the regression with the lowest mean-square error on the entire data. Place fields with a peak at the edge of the maze (<10cm from the ends) were excluded from this analysis, since often the shift led them "out" of the maze and the calculation was not reliable.

Place field shift correlation to time and speed

To understand if place field shift was simply a result of changes in speed over trials, I calculated the partial correlation between the shift slope and speed using a linear mixed effects model. The model is defined as

$$s = Xv + Yt + Zv{:}t + \varepsilon$$

Where s in the shift slope, v is the velocity, t is the trial number, $v{:}t$ is the interaction between both and ε refers to the residuals. X, Y and Z refer to the coefficients for each regressor,

estimated by the model. I used the implementation from the Python library Statsmodels function *mixedlm*.

Interneuron firing rate trends

For each interneuron I assessed whether the mean firing rate significantly changed as a function of trial number. Using the instantaneous rate over time instead of the mean per trial yielded equivalent results. For identifying the change in firing rate in time, I tested two types of regression: linear and sigmoidal, for detecting slow or abrupt changes in rate. I used Bayesian Information Criterion (BIC) to select the best fit. BIC is used for model selection as it calculates the trade-off between the quality of the fit and the number of parameters in a model. Given a dataset x with n datapoints and a model M, parametrised by a number k of parameters θ, the BIC is calculated as follows:

$$BIC = k\,ln(n) - 2ln(p(x|\theta, M))$$

Using this value, the criteria for determining the trends in the firing rate were the following:

1. Sigmoid (step up/down): if the BIC was smaller for sigmoid than linear fit, the transition was sharp (angle parameter for the sigmoid >0.5) and the difference between upper and lower bounds of the sigmoid were larger than the trial-by-trial fluctuations.
2. Linear (up/down): if the Bayesian Information Criterion was smaller for linear fit than sigmoid fit and the slope of the linear fit was significantly different from zero (p>0.05)
3. No trend: If none of the above criteria were met.

Interneuron selectivity to task variables

To understand if the mean firing rate of an interneuron in a given trial was correlated to the category of the trial I first subtracted the trends previously calculated, and analysed the residuals. Then for each neuron, I performed a two-factor ANOVA where the dependent variable was the residual firing rate in a given trial and the two independent variables were start position and context. I defined a neuron as selective to one or both of these variables if p>0.05 for the given variable or their interaction, respectively.

M3.3 Neural population

Population vector analysis with PCA and t-SNE

I divided the entire maze (360 cm) into 9 spatial bins. The edges of each bin were slightly shifted between animals to make sure that reward locations were not in between bins. All bins were kept at 40cm wide, except at the maze edges, which were 35 and 45 cm if necessary for this adjustment. For each trial and position, I calculated the mean firing rate of each neuron; the activity of all neurons formed a population vector of length n, where n is the number of pyramidal neurons; unless otherwise stated, these analyses included only putative excitatory

neurons. The Principal Component Analysis (Bishop, 2006) or t-SNE (Maaten & Hinton, 2008) methods were then performed using all 360 vectors (40 trials x 9 positions) in a given session. Only periods of movement (>3cm/s) in the preferred direction of the trial were considered for the calculation (e.g. moving left if the trial started on the right start box).

Linear decoding with Support Vector Machines

I used Support Vector Machines (SVM) to decode task variables (e.g. Trial direction or context) from the average population vectors from each position and trial (as described in the previous section) or their projections in a subset of the Principal Components. An SVM finds the best high-dimensional plane that separates the different categories, by finding the plane with the largest distance to points from any category (Bishop, 2006). For estimating decoding accuracy, I used a bootstrapping method. I sampled, with replacement, 360 vectors at each repetition (n=100 repetitions), then 80% of these vectors were used for training the decoder and the remaining 20% for accuracy testing. I used the same bootstrap samples on both decoders as well as for the shuffled label condition. The two decoding strategies employed were the following:

1. Global decoder: a standard SVM which received as labels the variable of interest associated to each vector. The output was the best hyperplane that separates the labels, by reducing the mean square error in the classification of the vectors in the training set.
2. Conditional decoder: each vector was assigned two labels, one for the conditional variable and one for the variable of interest. I subdivided the vectors into groups determined by their conditional variable. For each group, I trained an SVM for classifying the variable of interest, determining a different hyperplane at each value of the conditional variable.
3. Shuffle decoders: as a control, I calculated the average accuracy of both decoders above when I randomly shuffled the labels for the variable of interest, for the same vectors used in each bootstrap sample. The “shuffle” decoder in all plots refers to this manipulation unless explicitly mentioned that the shuffle was done to the conditional variable instead.

The decoding error was calculated as the proportion of matches between the predicted labels and the real labels of each vector in the test set. To control for the different numbers of neurons in each section I also looked at days with large numbers of pyramidal neurons (>80) and divided the data into non-overlapping parts of 10, 20, 40 or 80 neurons. In each subsampled set, I calculated the decoding accuracy for different numbers of Principal Components, to compare the effect of cell identity and number in this measure.

Population vector similarity

Population vector similarity was calculated creating a tensor of dimensions n x p x t where n was the number of neurons, p the number of position bins and t the number of trials. For each fixed t, n x p contains as rows the mean rate map of each neuron in trial t (4cm spatial

bins, 90 bins in total). For the similarity calculated over trials (Figure 6.1b-d), for each position *p*, I calculated the cosine similarity between the vector containing the activity of all neurons in that position for all possible pairs of trial numbers *t* and then averaged the similarity over all positions, obtaining the average similarity between pairs of trials. For the similarity between positions (as in Figure 6.7a-c), I took an average over the last 3 correct trials of a given trial category, obtaining a *n* x *p* matrix for each category. The cosine similarity was calculated between the columns of these matrices.

M4. Statistics

This is a general description of the statistical tests used throughout this thesis. The statistical test used in each analysis is described in the corresponding figure or text. Most comparisons in this thesis concern changes between behaviour clusters. Clusters contained an unequal number of sessions, and different animals could be more or less represented in each cluster. To control for this, in parametric tests, I mostly used two-factor ANOVA controlling for the factor "animal" and checking the effect of "behaviour cluster". Tukey HSD was used as a post-hoc test in these instances. If Repeated Measures ANOVA (parametric) or Wilcoxon test were used (in the non-parametric case), I first obtained the mean value of the statistic for each animal in each behaviour group, then calculated the test statistic in a paired manner. For the Wilcoxon test, I used Holm-Sidak correction for multiple comparisons. For log-normal data, I calculated parametric tests on log-transformed data. For values bound between 0 and 1, such as correlation measures, I first used a Fisher Z-transform and then a Z-test to calculate significance levels.

4. Animal behaviour

In this chapter, I introduce a novel behavioural paradigm for testing space-context association memory in rats. In the following chapters, I come back to this paradigm to investigate hippocampal neural function at the single-cell and population level during associative learning. The hippocampus has been known for many decades to encode information about an individual's spatial position (O'Keefe & Dostrovsky, 1971), time (Kraus et al., 2013; MacDonald et al., 2011) and contextual cues (R. Muller et al., 1987). My aim was to address both the spatial and contextual selectivity of the hippocampus simultaneously and to understand how they are represented as the animal learns their relevance in a goal-oriented task.

Most rodent studies of hippocampal function involve acquiring data from animals already pre-trained in a task, where the only novelty is a novel memory item, or position (Komorowski et al., 2009; Xu et al., 2019). In this study, my goal was to monitor learning for an extended time, so I looked for a paradigm that involved learning over multiple days. This is interesting because it requires memory processes to counteract homeostatic mechanisms of sleep (Grosmark et al., 2012; Vyazovskiy et al., 2009) that could favour forgetting the task, especially in periods when it has not yet been fully learned.

Contextual and associative learning have been largely explored in the field of fear memory research, but contextual-fear memory formation is known to involve different timescales and neural mechanisms than the formation of other types of memories (Phillips & LeDoux, 1992). Associative learning has also been studied in paradigms pairing objects, flavours or odours. These show the importance of the hippocampus for the memory of stimulus pairs but they usually do not include a spatial feature to the task or study the learning process itself (Bunsey & Eichenbaum, 1996; Komorowski et al., 2009; Sauvage et al., 2008). Based on these observations, I defined a set of requirements for the task I wanted to employ:

1. A goal-oriented spatial task
2. Contexts defined by changes within a single physical space
3. Context as the determinant of reward position
4. Fixed trajectories between contexts
5. A task hard enough to require multiple days to be learned
6. Reduced working-memory component, i.e. contextual cues should be accessible at all times

Very recently, a study has been published using a paradigm that partially matches these requirements, using head-fixed animals and odour cues as context, albeit in a working-memory setting (Zemla et al., 2022). Nevertheless, I did not find in the literature a pre-existing behavioural task in freely moving rodents that fulfilled all the requirements; with this purpose, I designed a novel maze and behavioural task. I also performed in vivo tetrode recordings of the

rat dorsal CA1 area of the hippocampus throughout the entire learning period (see schema in Figure 4.1a, an example of the electrode tracks reaching CA1 can be seen in Figure 4.1b).

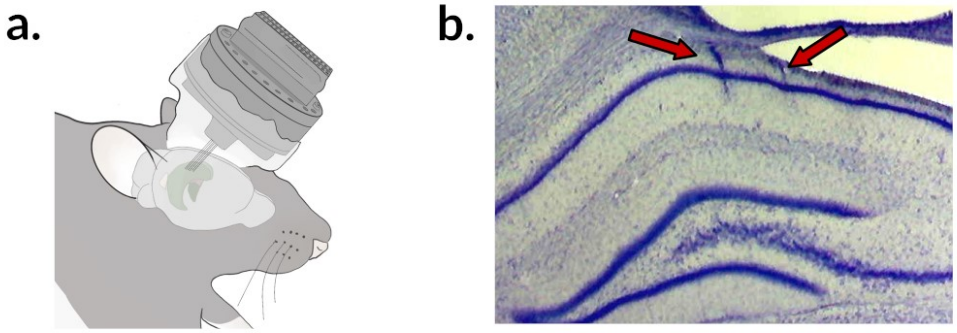

Figure 4.1. In vivo electrophysiology. a) Schematic of a tetrode microdrive implant targeted to the dorsal CA1 area of the hippocampus. Image credits to Sofia Taveira. b) Example of tetrode tracks in the pyramidal layer of the dorsal CA1. Cell bodies were stained with cresyl violet (Nissl staining). Tracks are marked with arrows.

Maze and task design

The details of the maze and task are explained in the Methods section. In short, I designed an S-shaped linear track containing multiple wells levelled with the track, which were filled with sand (Figure 4.2a). A subset of these wells could be covered, and the remaining were available for the animal to dig. Different contexts could be set by changing the colour and texture of the maze floor, whilst all other features and positions remained the same. In the paradigm I implemented, the animal was presented with two contexts. Half the wells available were matched between contexts (Figure 4.2b). At each trial the animal had to retrieve two pieces of food, one of which was always at the same location, independent of the context; the other whose location could be inferred by context but otherwise not indicated by any positional cue. The reward positions were the same every day for each animal, but different between animals. The structure of each training day is illustrated in Figure 4.2c. Animals ran a fixed amount of trials, independent of their performance.

One of the great challenges of using 2D mazes is that changing the reward position leads the animal to choose a new trajectory and, as a result, the maze coverage is not the same between conditions (Gridchyn et al., 2020). The S-maze design forced the animals to a fixed trajectory, making positions comparable between contexts. The behaviour of animals familiar with the task only differed with respect to which positions animals stopped to dig in each context. Trials were equally likely to start on either end of the maze (left/right), to ensure that the order of contextual rewards in the 1D environment did not lead to behavioural bias.

I chose a digging task to probe spatial memory despite the fixed running trajectory: a trial was considered correct if the animal only dug the two correct food locations. After the animal made its two choices, all other wells were covered to prevent any further exploration. Digging added an extra feature to this paradigm: different from running and exploration, it seems to be performed only when individuals are highly motivated; although I do not address it in detail in the current work, I observed that when animals were more confident about reward

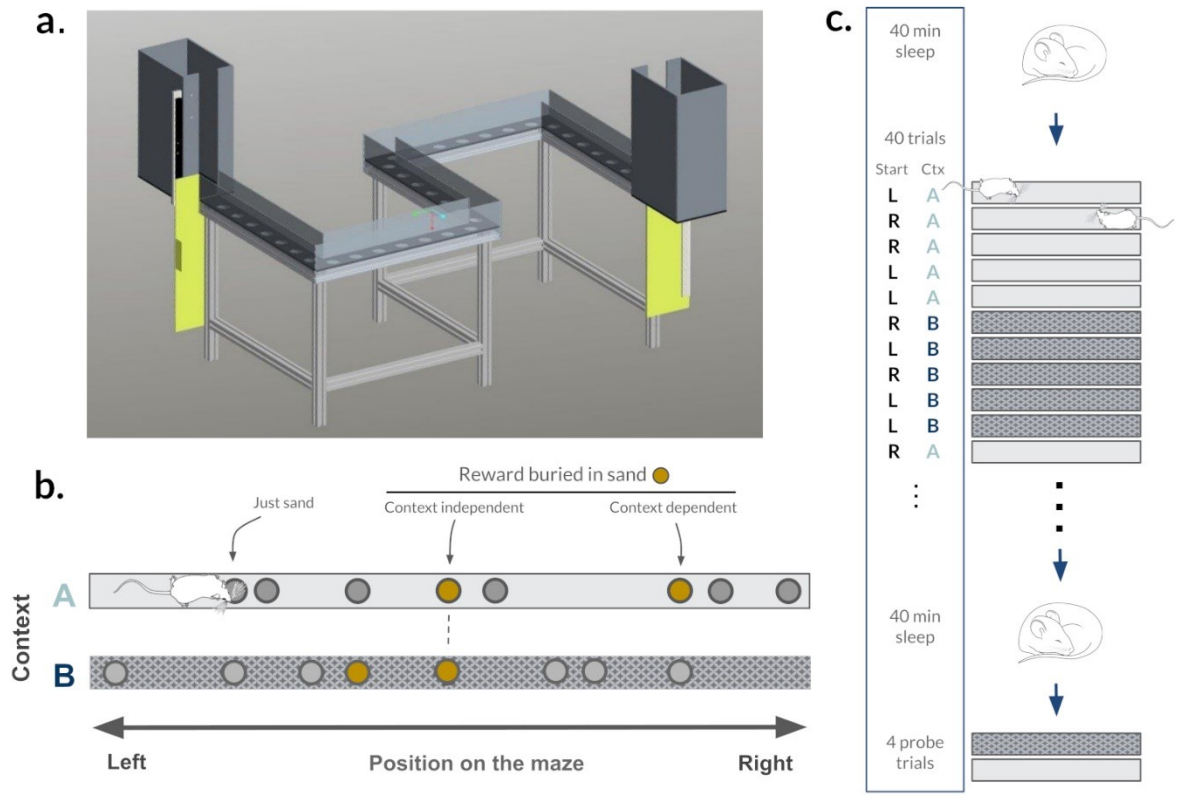

Figure 4.2. Experimental design. a) Schematic of the S-shape digging maze design. Image credits to Todor Asenov. b) Linear representation of the maze and task. Circles represent sand wells. Example reward locations are represented here in yellow, but are not actually marked by any cue in the real maze. Reward locations differ between different animals but are fixed across learning days. c) Structure of a single recording day.

location – i.e. once they had a set strategy for the task – they were more likely to engage in digging and also to insist on it for longer. Although the food was placed at a similar depth in every trial, digging time was very variable; this had to do with how the food was often pushed around together with the sand as the animal dug, making it easier or harder for it to be found. This uncertainty of reward delivery has also been seen in other paradigms (Passecker et al., 2019; Tryon et al., 2017), albeit in a form controlled by the experimenter. In our case, it was determined by the digging variability, but also affected the animals' confidence during the task.

Animals seemed to recognise changes in sensory cues from first exposure, much before they used these cues to guide decision-making. In trials right after a context switch, animals generally halted for a moment as they entered the maze. This behaviour was not observed in trials where no switch happened and therefore may indicate attention to cue change. Although only an empirical observation in terms of behaviour, neural data presented in the next chapters supports the idea that the animals are aware of context cues very early during learning.

Behavioural performance

Overall, achieving a performance of 80% or more required up to 4 days of training, and learning speed varied between animals (Figure 4.3a). Already on the first day, all animals quickly learned which wells could be rewarded and avoided digging all others almost entirely (Figure

4.3b). This is consistent with behaviour in other tasks such as the cheeseboard maze, where animals quickly learn to navigate directly to the correct reward locations (Kesner et al., 1991).

Despite the recognition of contextual cues and memory for reward locations not posing a challenge to the animals, the process of creating an association between them was much slower. This was expected, especially in this scenario where one of the reward locations did not depend on context. Most of the errors in this task could be attributed to digging a reward location in the context it was not rewarded ('Wrong context', Figure 4.3b). In those cases, animals seemed to dig whichever of the "rewardable" wells came first in their trajectory. Digging of other, never-rewarded wells, happened mostly on the first day of learning ('Random dig', Figure 4.3b), usually near locations rewarded on the pre-training days. On the 5th day, most animals showed a slight decrease in performance, but no increase in wrong associations; their performance generally resulted from disengagement from digging, likely due to tiredness ('Incomplete', Figure 4.3b).

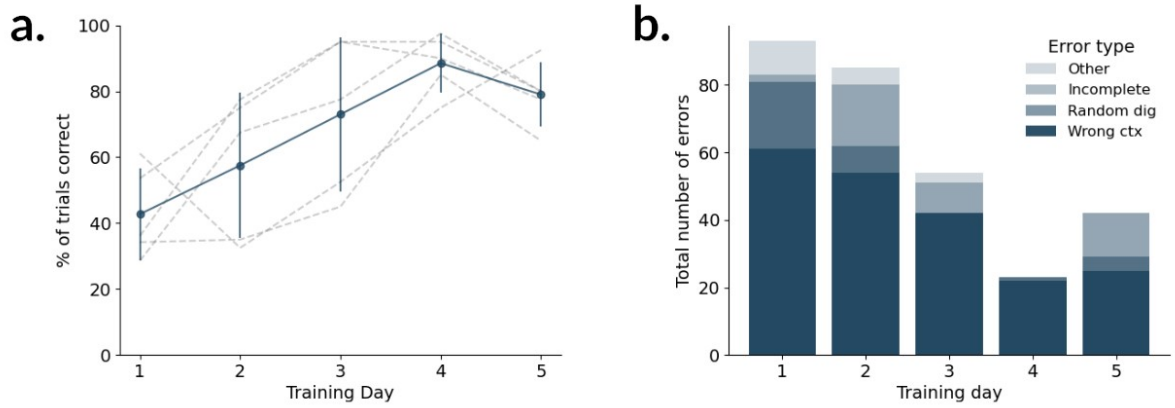

Figure 4.3. Behavioural performance. a) Average performance (number of correct trials / total number of trials) of all animals over training days (in blue) and of each animal separately (light grey). Error bars indicate 1 standard deviation. b) Number of errors summed for all animals as a function of training day and error type.

Individual variability

To better understand the behaviour of each individual, I compared their performance on different trial categories (Figure 4.4). The categories were defined by the contextual cues presented to the animal (A or B) and the side of the maze the trial started (L or R), as both variables seemed to affect behaviour. Whilst for some animals performance increased equally in all trial categories over days, others displayed biased behaviours; this was reflected in high performance in one or two categories in the initial days of learning. I then questioned what features of the behaviour determined these differences in performance.

One relevant feature of the task was that animals were allowed to run back and forth at will, which I refer to as "reversals" (Figure 4.5a). This was introduced as a result of testing different training regimes in pilot experiments. When reverse movement was forbidden, animals were incapable or very slow at learning the task. Even when animals dug correctly, they often did not find food if they did not insist. As a consequence, animals often gave up and moved on to the next reward well. Motivated individuals, however, were likely to return to their initial

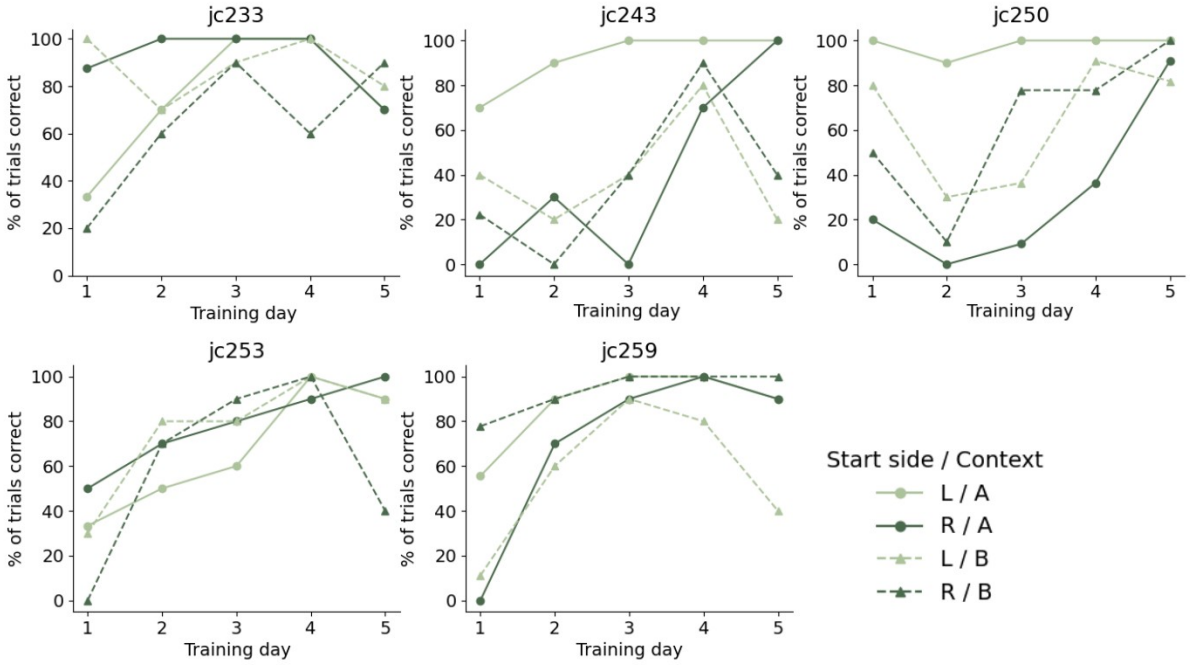

Figure 4.4. Individual performance per trial category. Performance per category was calculated as the number of correct trials in the category divided by the total number of trials in the category. Trial category was defined by which side of the maze the animal started (L or R) and which contextual cues are present (A or B). Figure titles starting indicate the ID of each animal.

choice and dig again, succeeding at finding the food. In case the first choice was wrong, going back also functioned as a confirmation of the incorrect choice. Therefore, reverse movement serves as compensation for task uncertainty as well as an indication of an animal's confidence about reward location. Even though animals seemed to reverse more on later learning days, this difference was not significant (Figure 4.5b). Animals were also more likely to reverse on the initial trial blocks of the day – each block corresponds to 5 consecutive trials of the same context (Figure 4.5d). The number of reversals in a single trial was highly dependent on the animal, but not on their stage of learning (Figure 4.5c). Even though reversals were necessary for learning as a whole and happened in more than half the trials, days with more reversal trials did not directly correlate with performance scores (Figure 4.5e). This confirmed that allowing reversals did not influence the original design of the task.

Behavioural clustering

Averaging behavioural measures between animals per training day (session) yielded inconclusive results, due to their differences in learning speed. Average task performance was also insufficient to describe behaviour, as the same overall performance score could result from various performances in each trial category, as previously mentioned. To tackle this issue, I applied a hierarchical clustering algorithm to classify sessions based on the different task-solving strategies an animal could have employed on each day. In this algorithm, I used three strategies that seemed to most likely explain the observed behaviour:

1. "Eager" strategy: animals dig the first two "rewardable" holes on their path, independent of context;

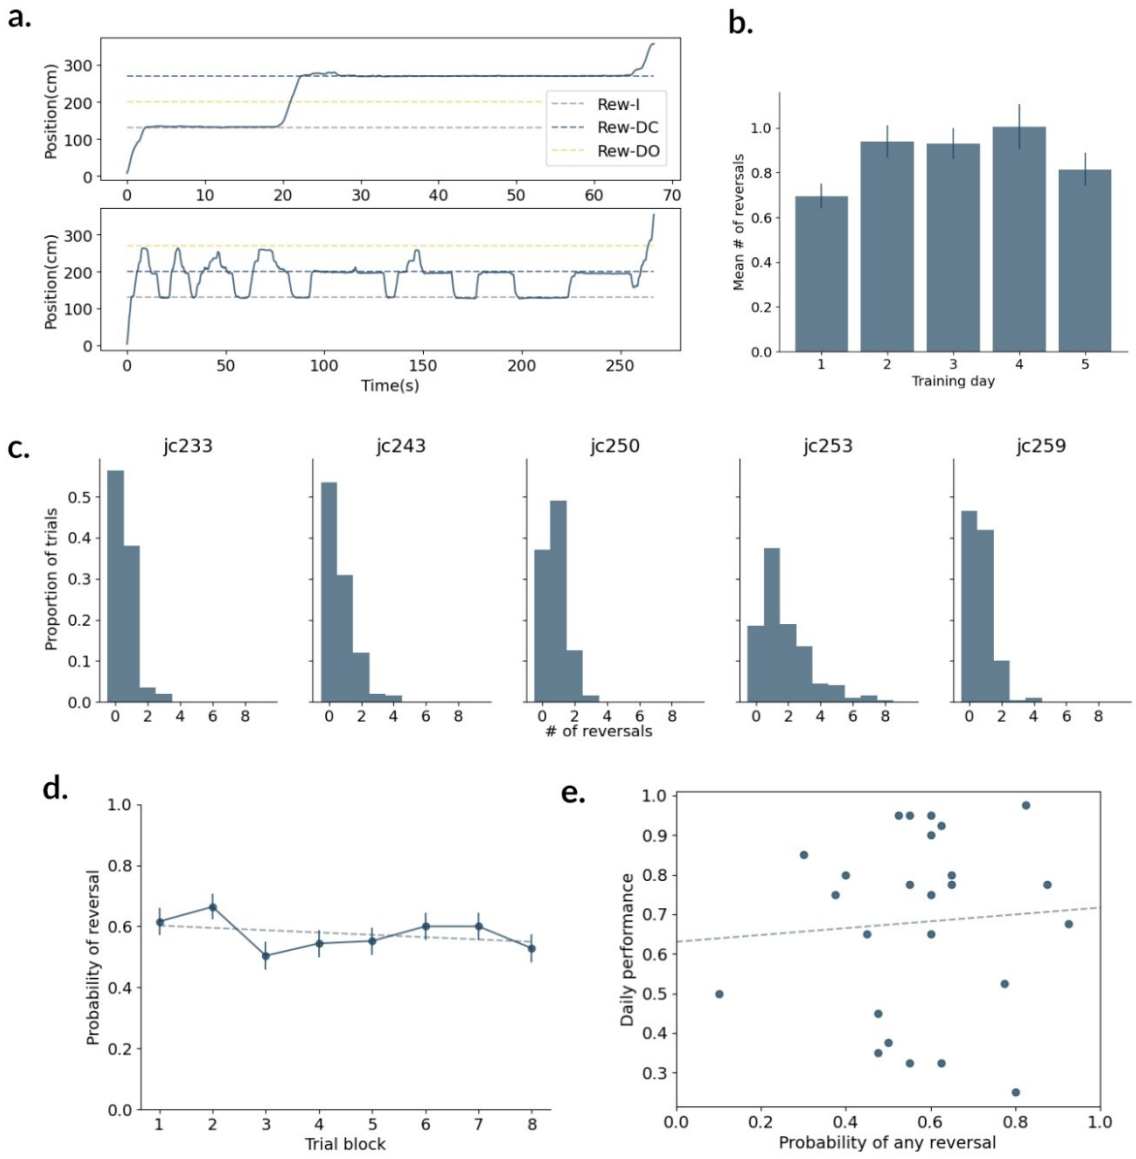

Figure 4.5. Reversals a) Example animal trajectories without (top) and with (bottom) reversal behaviour. Note that the animal mostly stops at reward positions (dashed lines). Rew-I: context-independent reward, Rew-DC: context-dependent reward of the current trial context, Rew-DO: context-dependent reward of the opposite context. b) Average number of reversals per training day. Bars indicate the SEM. Differences are n.s. ($p>0.9$ for all pairs, Wilcoxon test with Holm-Sidak correction). c) Same as (b), but per animal. Animal jc253 has a larger tendency to reversals than the others. d) Probability of an animal doing one or more reversals in a trial as a function of trial block within a given day. The trend is not significant ($r=-0.35$, $p=0.39$, linear regression). e) Relationship between probability of reversal and daily performance ($r=0.069$ and $p=0.743$, linear regression).

2. Win-stay-lose-shift (WSLS): animals dig the same holes as in the previous trial if that was a successful trial. Following an error trial they make a different choice, but independent of contextual cues;

3. Contextual strategy: animals use the cues to make a decision, as expected from the task design.

Each trial was assigned 1 or 0 for each strategy, depending on whether the wells dug were evidence for their use. The same choice of wells could sometimes be evidence for more than one strategy, and in such cases, both are assigned a 1. The average evidence for each strategy in each session was used as the dimensions for the clustering. Details on the algorithm

are described in the Methods section. The resulting hierarchical tree contained two main clusters (Figure 4.6). The first cluster contained sessions in which animals primarily used strategies 1 or 2, which do not require integrating contextual cues into the decision (Naive). This includes the first day of training for most animals, up to the third day for some of them. The second cluster contained sessions where there was evidence for contextual strategy use. Given the tree structure of the algorithm, this could be subdivided into two groups: sessions in which this strategy was used but animals still made a considerable amount of mistakes (Intermediate); and those of high performance on the contextual task (Expert).

The results of the clustering were robust to the addition of strategies. The use of "demotivated" (digging less than two holes) as a strategy did not greatly change the hierarchy. Adding training day as a variable also did not affect the main two clusters either, although it subdivided the Intermediate cluster into training sessions before and after Expert sessions. Since later analysis did not indicate many differences between the Intermediate subgroups, I did not include "training day" in the clustering.

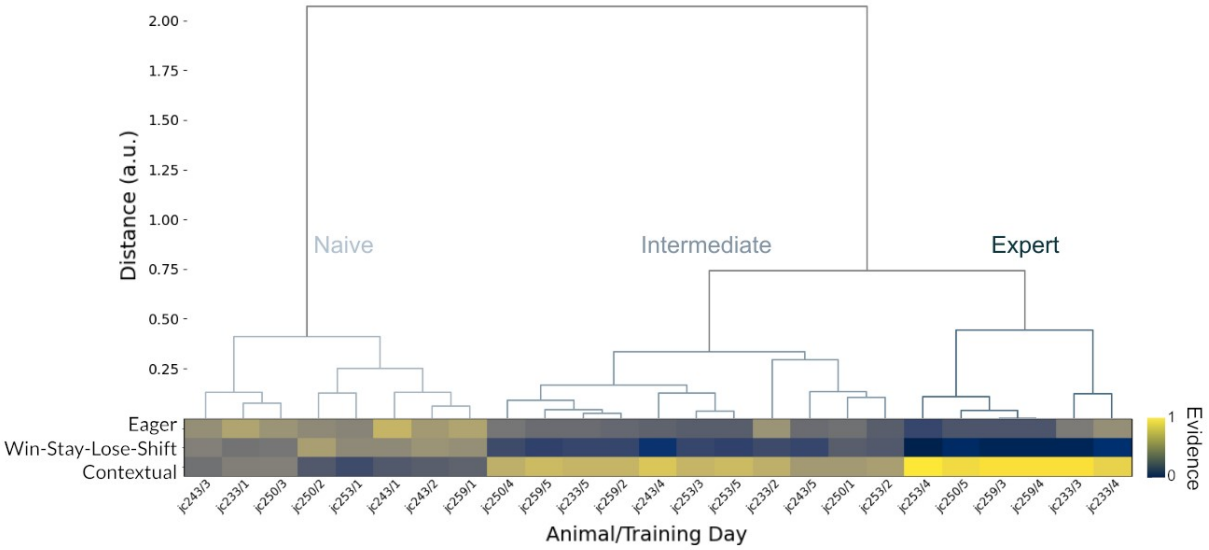

Figure 4.6. Hierarchical clustering of behavioural sessions. Top: Hierarchical clustering of training sessions based on the accumulated evidence of each behavioural strategy. For details on the distance metric please see the Methods section. Bottom: heatmap of evidence for each behavioural strategy in each training day.

Behaviour readouts of animal strategy

To better understand the dynamics of the learning, in Figure 4.7 I show how the evidence for each strategy changed over days. It is important to note that, even when animals employed the contextual strategy, evidence for multiple strategies was expected in certain trials. This expected value depended on the order of the rewards on the maze (i.e. when the context-independent reward was placed in between the contextual ones, 50% of correct trials would give evidence to an eager strategy, whilst only 25% otherwise). I could observe that 3 animals switched to a contextual strategy already on the second day of training; the remaining two animals achieved similar behaviour only around the 4th day. This agrees with the results shown

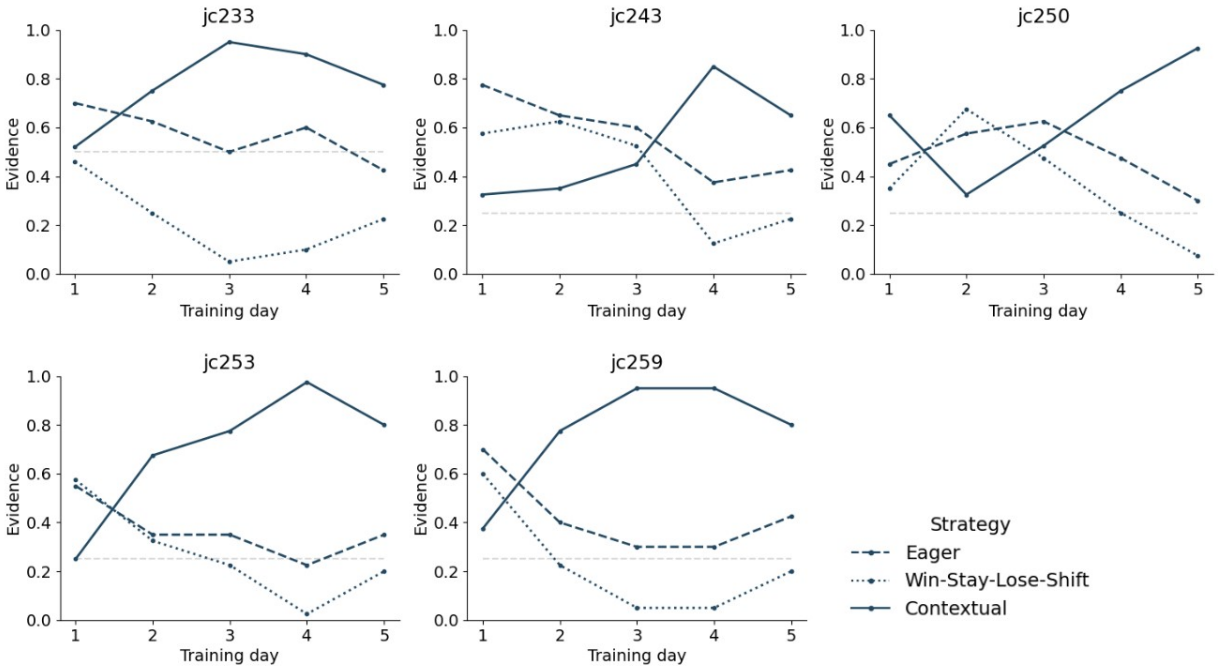

Figure 4.7. Evidence for behavioural strategies. Mean evidence for each strategy, per animal and training day. The same trial can be evidence for more than one strategy so values do not have to sum to one. The dotted grey line indicates expected evidence for the "Eager" strategy when the animal is actually using the correct "Contextual" strategy, as it differs between animals depending on reward locations.

in Figure 4.4, in which I show individual differences in learning each trial category. Both results support the notion that studying learning as a simple function of time can be misleading, as it might mask learning timescales inherent to each individual. For the remainder of this thesis, I looked at variables of interest both in terms of training day and behavioural clusters and presented the grouping that most explained each of them.

I first compared behavioural clusters in terms of task performance. An inherent result of the cluster separation is that sessions in the Naive cluster show high evidence for an "eager" strategy. That means that animals were less likely to refrain from digging as they passed by the reward of the opposite context (Figure 4.8a). When they did succeed in refraining, however, they spent a lot more time around the location before deciding to move on; in the Expert sessions, animals ran past those positions with little sign of uncertainty (Figure 4.8b). In Intermediate sessions, this behaviour was neither of the extremes. In terms of reversals, there was no significant difference between behavioural clusters (Figure 4.8c).

Learning within each session

I also questioned if performance improved within a single training day. I divided the behaviour sessions into 8 training blocks of 5 trials, the rate at which context changed. Figure 4.9a shows that Expert sessions already started with better performance, and the positive slope of performance in time indicates that it improved as the day progressed. This increase is also true for Intermediate but not for Naive sessions. This implies that errors in the Intermediate and

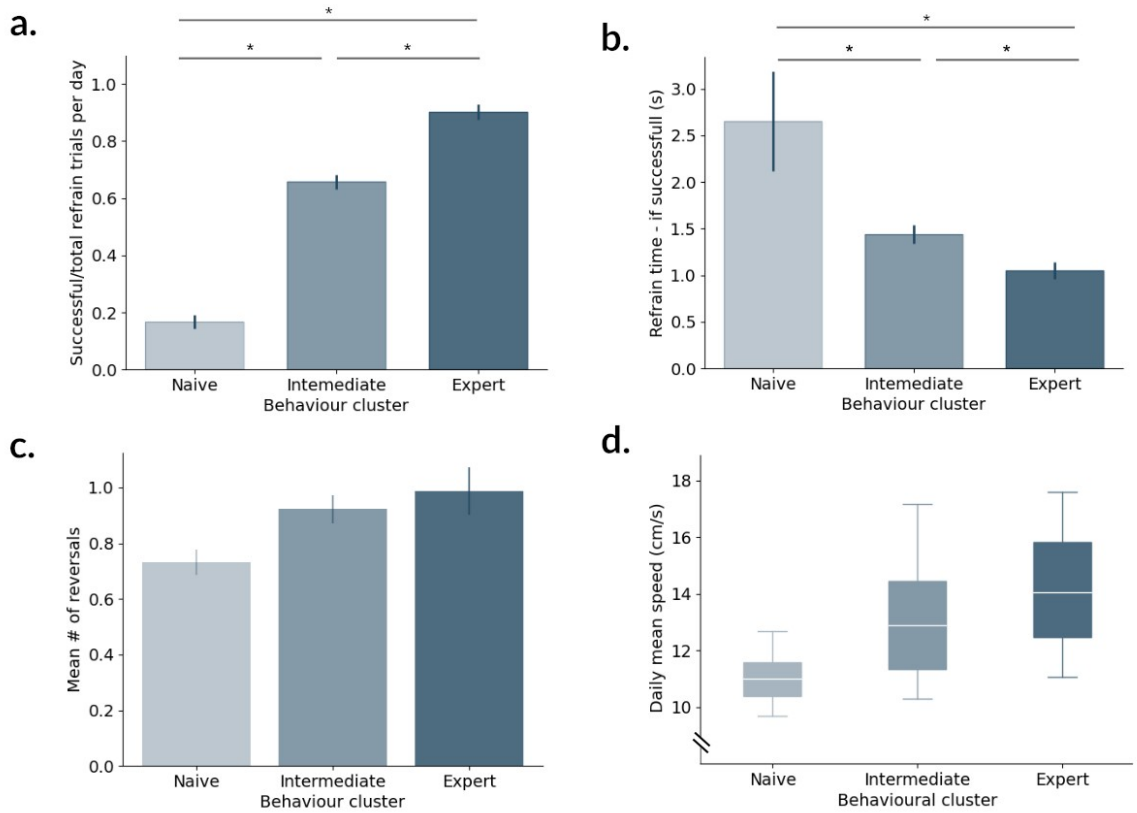

Figure 4.8. Differences between behavioural clusters. a) Proportion of trials that the animal successfully refrained from digging when refrain was necessary, per behavioural cluster (p=3.74e-5 two-factor ANOVA. Post hoc Tukey HSD N vs. I p=5.30e-12, N vs. E p=3.08e-14, I vs. E p=5.67e-6. Calculations on log-transformed data). b) Time, in seconds, spent around the incorrect reward in successful refrain trials; i.e. not digging. (p=1.07e-5 two-factor ANOVA, Post hoc Tukey HSD N vs. I p=0.021, N vs. E p=9.19e-6, I vs. E p=9.38e-4. Calculations on log-transformed data) c) Mean number of reversals per trial, for each behavioural cluster. Differences are n.s. (all p>0.75, Wilcoxon test with Holm-Sidak correction) d) Animal mean speed, excluding immobility periods (< 3cm/s). Differences are n.s. (p=0.324, Two-factor ANOVA). Error bars in panels a-c indicate the SEM.

Expert sessions were more likely to happen at the beginning of the day. Also at this timescale, reversals were not a direct correlate to performance. I was also interested in how likely animals were to switch strategies in the middle of the day, as switching is a sign of cognitive flexibility (Kaefer et al., 2020). In accordance with our expectation regarding the behaviour clusters, I found high evidence for the context strategy already in the first trial block of Expert sessions. In Intermediate sessions, there was non-negligible evidence for all strategies at the beginning of the day, but as the day progressed the contextual strategy became dominant. In the case of Naive sessions, other strategies were preferred and there was no clear sign of strategy learning throughout the day (Figure 4.9b).

Time measurements of behaviour

Once I assessed that the hierarchical clustering outputs a relevant grouping of the training sessions, I investigated other behavioural measures suggestive of learning of the space-context associations. Similarly to the refraining behaviour in Figure 4.8a-b, for most measures, there was a gradual change from Naive to Expert sessions. For example, animals ran slower in naive sessions (Figure 4.8d), which resulted in longer trials, when it comes to movement periods

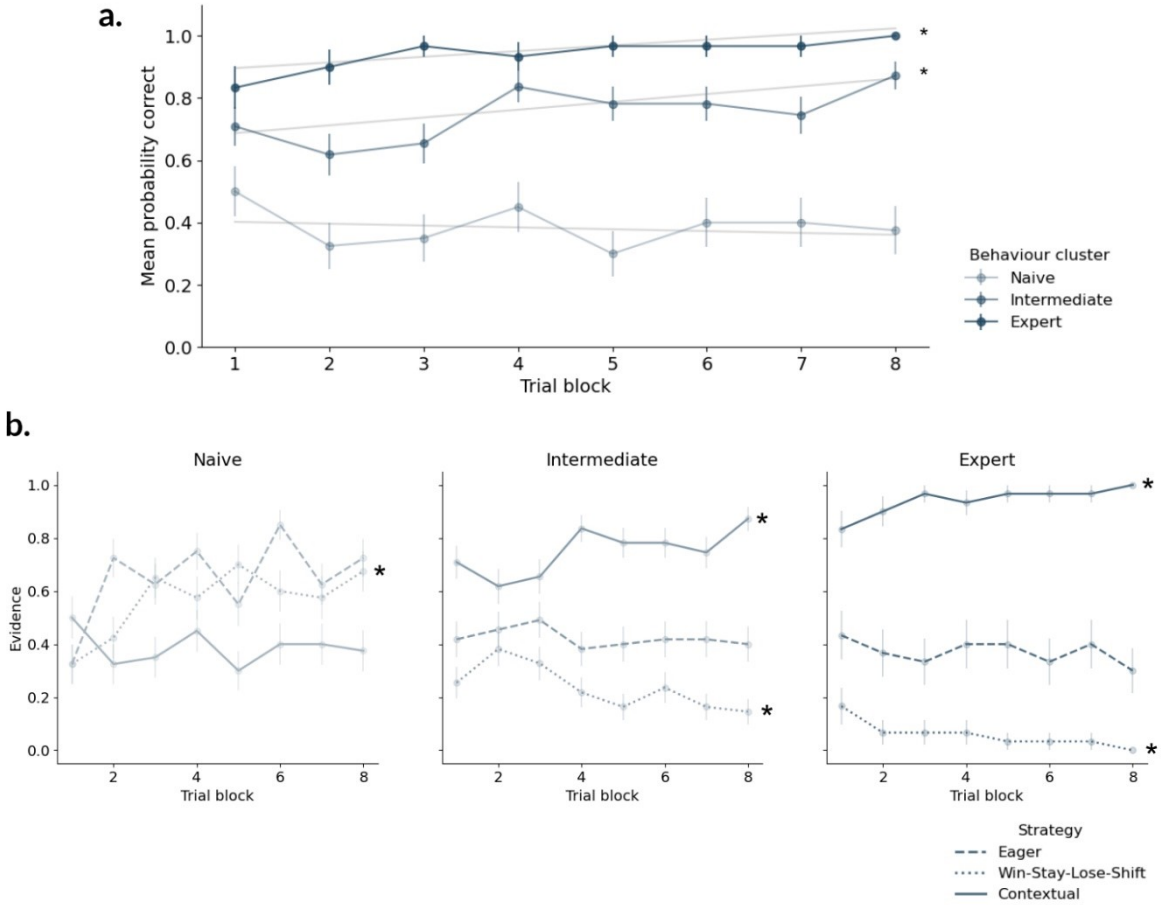

Figure 4.9. Evidence of learning within a day. a) Probability of performing a correct trial within each trial block, averaged over all sessions in a behavioural cluster. Grey lines indicate the result of linear regression. (Linear regression: Naive r=-0.05 p=0.685, Intermediate r=0.22 p=0.037, Expert r=0.36 p=0.011) b) Mean evidence, per trial block, for each of the behavioural strategies (Linear Regression $p<0.05$ for all strategies marked with a *, for all others $p>0.05$). Error bars in all plots indicate the SEM.

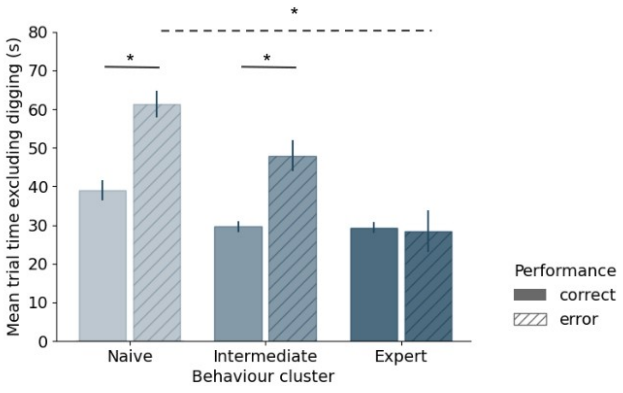

Figure 4.10. Trial time in correct and error trials. Comparison of trial duration, in seconds, between the different behavioural clusters. Here digging periods are excluded, since those are extremely variable within training days. For difference between correct and error, * indicates $p<0.05$ in paired t-tests. For comparison between behaviour clusters, two-factor ANOVA indicates $p<0.05$ only between error trials. Post hoc Tukey HSD results p=0.316 Nvs.I, p=0.023 N vs. E, p=0.232 I vs E. Error bars indicate the SEM.

(Figure 4.10). This slowness might be a sign of higher exploration at this phase. Error trials were also longer for the Naive and Intermediate sessions, indicating animals explored more when they made a mistake, whilst in Expert sessions they quickly finished the trial after errors (Figure 4.10).

Immobility and digging periods were much more variable in duration and not directly related to the animal's speed. This seems inherent to the digging itself; a rat could spend between 3 seconds and more than a minute digging a location on a correct trial (Figure 4.11a, top). As previously mentioned, the reason is that digging is uncertain in nature. Nevertheless, on incorrect trials, if animals mistakenly dug the reward location of the opposite context, they spent a lot less time on it (Figure 4.11a, bottom). This is a sign that such "wrong digging" was not treated equally to the correct ones. This effect was more pronounced in Expert sessions, when wrong digs were quickly noticed by the animals and the time spent on those positions was much lower than for Naive sessions (Figure 4.11b).

In terms of the proportion of time allocated to the reward region, not surprisingly animals spent more time around the context-independent and correct contextual rewards in correct trials than in error trials. This is because after making two choices, all other sand wells

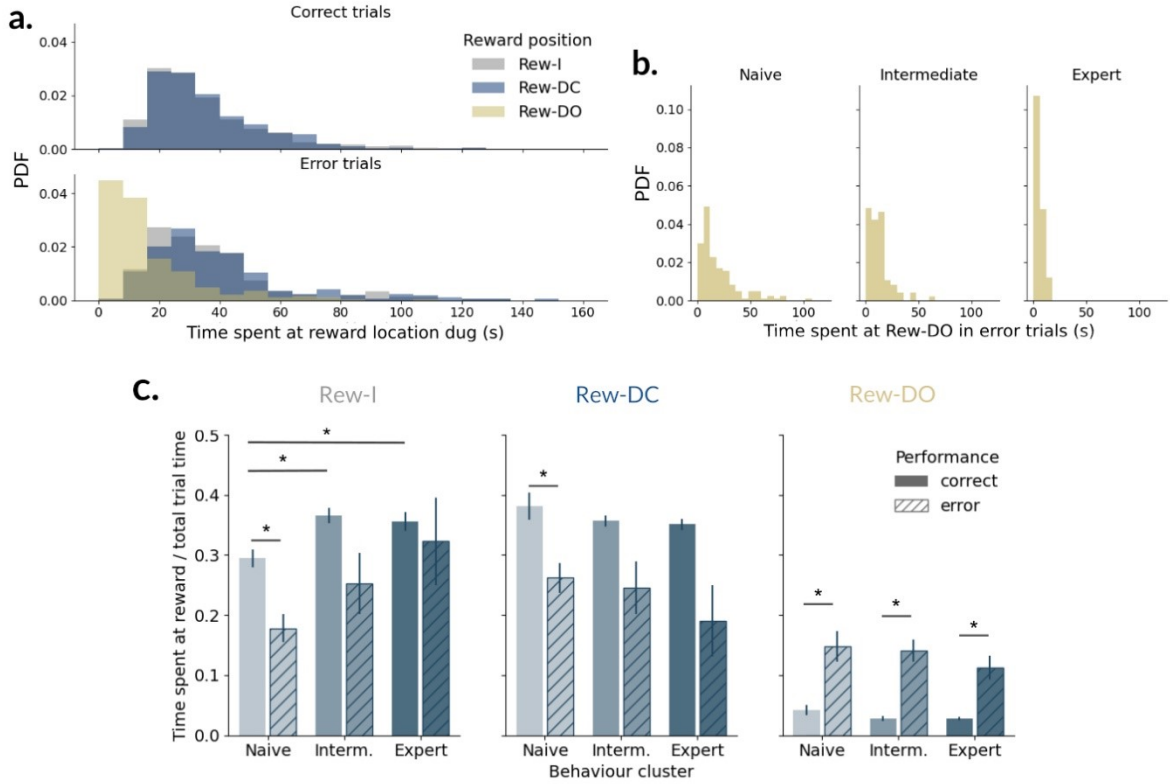

Figure 4.11. Time spent on reward locations. a) Probability Density of time spent over each reward position at each trial, if it was dug. Top: correct trials, Bottom: error trials of the "wrong context" type. Rew-I: Context-independent reward, Rew-DC: Context-dependent reward of the current trial, Rew-DO: Context-dependent reward of the opposite context. b) Same as error trials in (a) but separated between behaviour clusters. c) Proportion of trial duration spent on each reward region, regardless if the animal dug. I first calculated the mean duration over all trials in each session, then I calculated the mean and s.e.m. over all sessions in each cluster. Error bars indicate the SEM. For difference between correct and error, * indicates $p<0.05$ in paired t-tests. For comparison between behaviour clusters, $p<0.05$ only between correct trials in Rew-I (two-factor ANOVA). Post hoc Tukey HSD results p=0.003 Nvs.I, p=0.027 N vs. E, p=0.889 I vs E.

were made unavailable; if an animal made a mistake it means that at least one of the correct reward wells was not accessible and therefore they could not spend time digging it during an error trial. Analogously, they spent more time around the reward of the opposite context on error trials. Only in Expert sessions they spent equal amounts of time around the independent reward in both correct and error trials (Figure 4.11c), indicating that they probably dug this well in both types of trials.

Discussion

Together, these results show that rats were capable of learning this context-space associative task within 3-4 days. In this paradigm, classifying behaviour sessions in terms of animal task-solving strategy yielded better predictions of behavioural variables than simple performance measures. As the animals learned the desired task strategy, they also became faster and spent less time exploring reward locations not associated with the current context. They spent less time deliberating over and made fewer attempts to dig incorrect locations; in case of an incorrect dig attempt, they also spent less time digging. As a note, I performed animal tracking by acquiring solely x/y coordinates of the animal's head over time. For a more fine-grained investigation of behaviour in this maze, especially of digging behaviour, it would be beneficial to acquire full video footage associated with machine learning algorithms for movement classification (e.g. DeepLabCut).

To create this maze and behavioural paradigm I was inspired by a conjunction of previous studies in spatial and contextual memory, which I summarise here. Much of the role of the hippocampus in contextual associative memory has been discovered in fear memory research. Animals show fear-related behaviours such as freezing primarily in contexts (e.g. a chamber) where they received a nociceptive stimulus in the past, a process that requires the hippocampus (Maren et al., 2013; Phillips & LeDoux, 1992). It is well established that fear memory formation also involves specific neural mechanisms, e.g. via the amygdala (Phillips & LeDoux, 1992), which allow associative learning to happen within a single exposure to the contextual cue.

Other types of associative tasks require much longer timescales of learning (Komorowski et al., 2009). For example, animals take several days to learn odour-pair associations, a time comparable to my observations in the S-maze task (Bunsey & Eichenbaum, 1996). In this study, animals were taught to associate pairs of odours embedded in cups of sand or other digging media. After learning, if exposed to a sample odour followed by two options, animals could correctly dig at the odour originally associated with the sample. Lesions to the hippocampus impair animals' ability to choose the correct pairing (Bunsey & Eichenbaum, 1996). Similarly, if odours are associated with specific digging media, the hippocampus is necessary for animals' recollection of the pairing, but not for the memory of each odour or media themselves (Sauvage et al., 2008).

Digging has also been introduced in a spatial task studying memory schemas (Tse et al., 2007). Rats learned the association of a food flavour with a location in an open arena where a

reward could be dug out. They show that the initial learning of flavour-location pairings also requires many days, although not the learning of new pairings afterwards. The authors claim that, in the first stage of the training, the animal needs to learn a "schema" of the task: the rule that flavours are associated with reward positions. Later, when exposed to a new flavour, they already expect that a flavour determines a reward location, and therefore learn much faster. The study shows that lesions of the hippocampus affect task learning (Tse et al., 2007). Altogether, these studies that the behavioural paradigm I introduced here is a good candidate for studying the hippocampal dynamics during associative learning.

I was also interested in the spatial aspect of the task. In spatial memory studies, contexts are often defined as separate environments, as in contextual fear paradigms. A previous study by my research group shows that contextual memories of this type can be manipulated separately during consolidation (Gridchyn et al., 2020). Hippocampal responses have also been studied in the case of morphing environments, where the animal is exposed to a square arena that is gradually shifted into a circular one (J. K. Leutgeb et al., 2005). Nevertheless, neither study addresses the learning of context process per se, nor do most studies in this field, which motivated the current work.

Other mazes for rodents that share parallels with my task design include the radial 8-arm maze (Xu et al., 2019) and some conveyor belt paradigms, the latter usually employed in head-fixed settings (Danielson et al., 2016; Vaidya et al., 2023). Similarly to the radial 8-arm maze, where animals have to make a decision about which arm to explore, my design has precisely 8 possible digging sites at any trial, which the animals can choose to dig. Conveyor-belt paradigms often include tactile cues to define spatial locations or to define context (Danielson et al., 2016), as I have used in this case.

The results presented in this chapter show that the task I developed imposed an interesting challenge for rats, which nonetheless they are capable of learning. Animals seemed to quickly familiarise themselves with the maze and learn which positions were ever rewarded. This is in accordance with previous studies of goal-directed spatial learning. The S-maze shares many features with the classic "Morris water maze" and its dry version, the "cheeseboard" maze, albeit in 1D. In those mazes, animals learn the most direct path to the safe platform or reward in a few trials (Kesner et al., 1991; R. Morris, 1984). In my paradigm, I saw a similar phenomenon in which animals moved directly towards their believed rewards, with very little interruption in their movement. Similarly, in the 2D mazes animals explored more in earlier trials (longer trajectories) and in the S-maze, they spent longer time digging at incorrect locations.

Contrary to fast spatial learning, associative learning happened at a slower timescale, highlighting individual differences not noticeable otherwise. Studies often ignore or actively "average out" this individual variability, even when it is large (Zemla et al., 2022). In rodents, it has been shown that individual variability is regulated by external (Körholz et al., 2018) and neural factors (Akiti et al., 2022). In humans, individual differences in learning, memory and decision-making have been extensively studied (Badre et al., 2012; Bors & MacLeod, 1996; Kirchhoff & Buckner, 2006). A review study on human cognition suggests that differences in

behaviour can be linked to specific variations in the brain structure between individuals (Kanai & Rees, 2011).

Investigating decision-making in rodents was not a goal of the current study. Nevertheless, I hypothesised that learning differences between animals emerged from the task-solving strategies they employed. I could not exhaust all possible strategies but pinpointed those most likely to explain the observed behaviour. The hierarchical clustering based on task-solving strategies was agnostic to any predetermined grouping; the evidence for each strategy was related to the wells dug by each animal rather than the experimenter's view of the behaviour. I showed that many behavioural measures not used for clustering are common between sessions in the same cluster. This could also explain the robustness of the clustering to adding/removing strategies from the algorithm. In Chapters 5 and 6, I show how these behaviour clusters were distinct in terms of neural activity.

Other methods for classifying rodent behaviour in terms of task-solving strategy have been described for the Morris water maze, using machine-learning algorithms (Gehring et al., 2015; Graziano et al., 2003). In these examples, the algorithm classifies trajectory patterns on a trial-by-trial basis or even parts of trials. Whilst their approach is focused on studying the differences between behavioural patterns per se, my approach is mostly focused on classifying sessions to investigate other features of the behaviour and neural activity. I am not aware of any studies using hierarchical clustering for this type of behavioural classification. This is a very simple approach that can be easily adapted to other behavioural paradigms, by choosing strategies most suitable in each case. I believe that this type of behaviour clustering can yield a relevant grouping of training sessions and help adjust other analysis to different individual learning speeds.

Lastly, the maze and behavioural paradigm I present here are not only suitable for the analysis I show in the following chapters but could also be used in future studies of other brain areas. For example, the retrosplenial cortex is known to respond to contextual cues (Nelson et al., 2014), as well as being important in spatial tasks, regulating egocentric versus allocentric navigation, among other functions (Alexander & Nitz, 2015; Vann et al., 2009). Therefore it is likely to play an important role in this task. The posterior parietal cortex and prefrontal cortex are known to be involved in decision-making (Diehl & Redish, 2023; Raposo et al., 2014; Sul et al., 2010) and associative learning (Lipton et al., 1999), making them other interesting areas for investigation. The relationship between sleep and associative learning has not been much explored in rodent studies; given that animals learn over multiple days, mechanisms of memory consolidation are probably also central to learning this task (Frankland & Bontempi, 2005; Ramadan et al., 2009; Talamini et al., 2008).

5. Single cells

In Chapter 4 I introduced a novel experimental paradigm that allowed me to study associative learning between space and visual/tactile contexts. The goal of the paradigm was to elicit neural responses in the hippocampus to multiple task variables known for engaging this brain area (Gener et al., 2013; R. U. Muller & Kubie, 1987; O'Keefe & Dostrovsky, 1971; Wible et al., 1986; Wiener et al., 1989) and investigate these responses over many days of learning. In the current chapter, I will look into how individual hippocampal cells behave during this task.

The extent to which CA1 neurons are tuned to spatial variables is highly related to their position in the proximo-distal axis of the hippocampus, which determines their inputs (Beer et al., 2018). Moreover, there has been accumulating evidence that hippocampal place cells do not respond solely to an animal's position, but also to environmental cues and their relationships (Shapiro et al., 1997; Wiener et al., 1989). Mixed-selectivity – neural tuning to multiple variables – has been shown to bring several coding benefits for a neural population (Rigotti et al., 2013). Some argue that place cell activity does not represent space at all, but rather the animal's predictions or experiences (Stachenfeld et al., 2017; Tanaka et al., 2018). This is supported by the fact that the tuning of neurons to non-spatial variables seems more stable than their place code (Tanaka et al., 2018). In addition, there have been reports showing that non-place cells in this brain area also carry information about space (Stefanini et al., 2020) and reward (Gauthier & Tank, 2018; Wood et al., 1999; L. Zhang et al., 2022).

Apart from these properties observed during free exploration or in familiar tasks, in the process of learning the hippocampus is known to adapt its place cell firing. It can discriminate different contexts (Markus et al., 1995; Mizumori et al., 1999; Smith & Mizumori, 2006; Wood et al., 2000; Zhao et al., 2020) and overrepresent reward locations (Dupret et al., 2010; Hollup et al., 2001). These phenomena have been mostly investigated in a time scale of dozens of trials, but less is known about their evolution over long periods of time. Previous recordings of hippocampal activity during multiple days use a free-exploration paradigm (Ziv et al., 2013) or focus only on the subpopulation of neurons with stable place fields over days (Vaidya et al., 2023), precluding the investigation of changes in the firing resulting from learning. Both studies report that only a small percentage of place cells maintain their place field position over days, whilst the majority acquire their field on roughly a daily basis (Vaidya et al., 2023; Ziv et al., 2013). Another study shows that CA1 neurons accumulate place fields around rewards also over days, and the strength and stability of this modulation are dependent on the cell depth in the pyramidal layer (Danielson et al., 2016). Regarding behaviour, one report shows that neurons from high-performing animals accumulate place fields in linear tracks selectively between the start position and reward for each context (Zemla et al., 2022), but it does not investigate the timescale of this accumulation during the process of learning.

Overall, not many other studies address the encoding of fixed rewards and sensory cues in a spatial task over long periods of time, as the environment and the reward locations become familiar. It is not clear how the animals' improving performance over consecutive days is related to changes in the hippocampal encoding in the same period. Therefore I aimed here at investigating the responses of CA1 neurons over various timescales of learning, from a few trials to a few days.

Place cells

I simultaneously recorded dozens of cells (n=1725 putative units in total, mean of 69 ± 25 units per day per animal) from the dorsal CA1 of 5 rats across 5 days of learning each, a period that was sufficient for all animals to achieve a performance of at least 80% of trials correct. The recordings were optimised for renewing the neuronal population recorded on each day; generally, the number of cells recorded on earlier training days was smaller (n=26 units on the worst session from all recordings) than that on later training days (n=109 units on the best session from all recordings), as the electrodes were moved to the centre of the hippocampal pyramidal layer over the days. On average, 80 ± 10% (mean ± STD) of cells recorded in a given day were classified as putative pyramidal cells, whilst the remaining were labelled putative interneurons (mostly fast-firing basket cells).

To classify putative pyramidal neurons as place cells, I measured spatial information and sparsity of each rate map (Skaggs et al., 1996), as well as rate map stability. Trials were divided into categories determined by the combination of the side of the maze where the trial started (Left or Right) and the context (A or B). For the rate map calculation, I used only the periods where the animal ran forward, given their starting side. On a given day, 40 ± 15% of pyramidal cells were classified as stable place cells in at least one of the trial categories (Figure 5.1a), which is in agreement with previous reports on place cells (Schoenenberger et al., 2016), although most of those studies do not separate rate maps into categories. Only 17 ± 13% of place cells displayed stable place fields in all trial categories. This is in agreement with a previous report on a similar context-space association paradigm (Zemla et al., 2022).

In addition to the stable cells, 36 ± 10% of cells were classified as unstable place cells, meaning their map was stable over at least two consecutive trials but not over the entire recording day (Figure 5.1b). The proportions of stable and unstable cells were not significantly different between animals, training days or behaviour clusters, although there seemed to be a trend for a higher number of stable place cells on sessions with better performance and an opposite trend for unstable ones. These trends could be a result of learning or a result of (on average) more cells being recorded from central areas of the CA1 pyramidal layer in expert sessions. A study using a similar task design claims that the proportion of pyramidal cells active as place cells is steady over learning (Vaidya et al., 2023). The same study shows that, over time, there is an increase of cells with long-lasting place fields, matching the trend I observe. They also report earlier emergence of place fields on later days, for a subgroup of place cells. I do not

observe this phenomenon when looking at the entire population (Figure 5.1c) and their subgroup categorisation is not possible in the case of my dataset.

I also observe no change in the width of place fields with experience over days (Figure 5.1d), in agreement with a recent study (Geva et al., 2023). Sharpening of the place fields via synaptic plasticity mechanisms as a result of learning has also been reported (McHugh et al., 1996) but only in the scale of seconds.

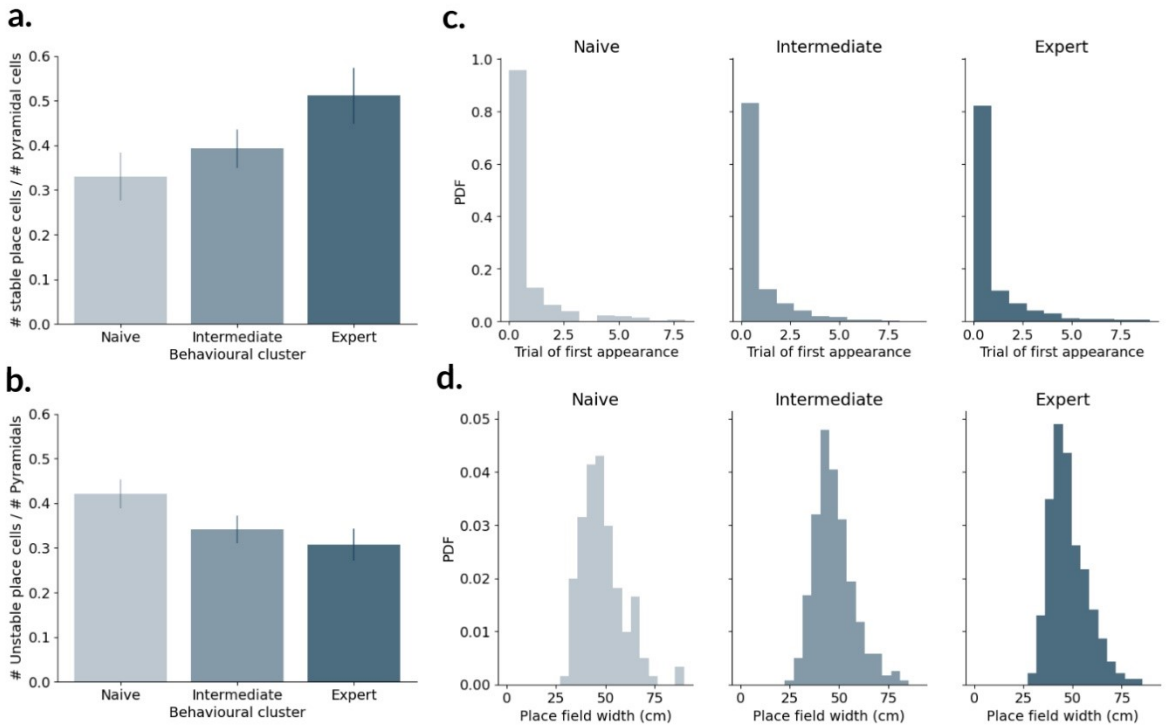

Figure 5.1 Pyramidal cell features. a) Proportion of stable place cells in the pyramidal population. n=25 sessions. Differences are n.s. (One-sided Wilcoxon test with Holm-Sidak correction. p=0.23 N vs I; p=0.18 N vs E; p=0.23 I vs E). b) Proportion of unstable place cells in the pyramidal population. Differences are n.s. (One-sided Wilcoxon test with Holm-Sidak correction. p=0.09 N vs I; p=0.34 N vs E; p=0.56 I vs E). c) First trial of place field appearance in the category the place field is present. Differences are n.s. (Log-rank test with Holm-Sidak correction p=0.97 N vs I; p=0.34 N vs E; p=0.34 I vs E). d) Distribution of place field average size in each behavioural cluster. n=2716 place fields. Differences are n.s. (pval = 0.28. Two-Factor ANOVA after log-transformation of the data). Error bars in panels a-b indicate SEM.

Place field identification

I then detected individual place fields within the rate map of each cell, in the categories the cell was classified as a place cell. Whilst most place cells had only one place field (90±5%) in which they fire above 2 standard deviations of their regular firing, a small proportion of place cells (9±5%) displayed two place fields within a single rate map (example in Figure 5.2). These proportions were not significantly different between behaviour groups ($p>0.33$ for all pairs, Wilcoxon test with Holm-Sidak correction). Some cells also show smaller (< 2 STD) firing peaks at regions that match the place field in a different category. Although not directly addressed in most publications, the presence of multiple place fields has been reported in many studies (Fenton et al., 2008; Schoenenberger et al., 2016; Vaidya et al., 2023). It has also been related to the size of environments, with place cells acquiring more place fields as the environment becomes larger (Fenton et al., 2008; Rich et al., 2014).

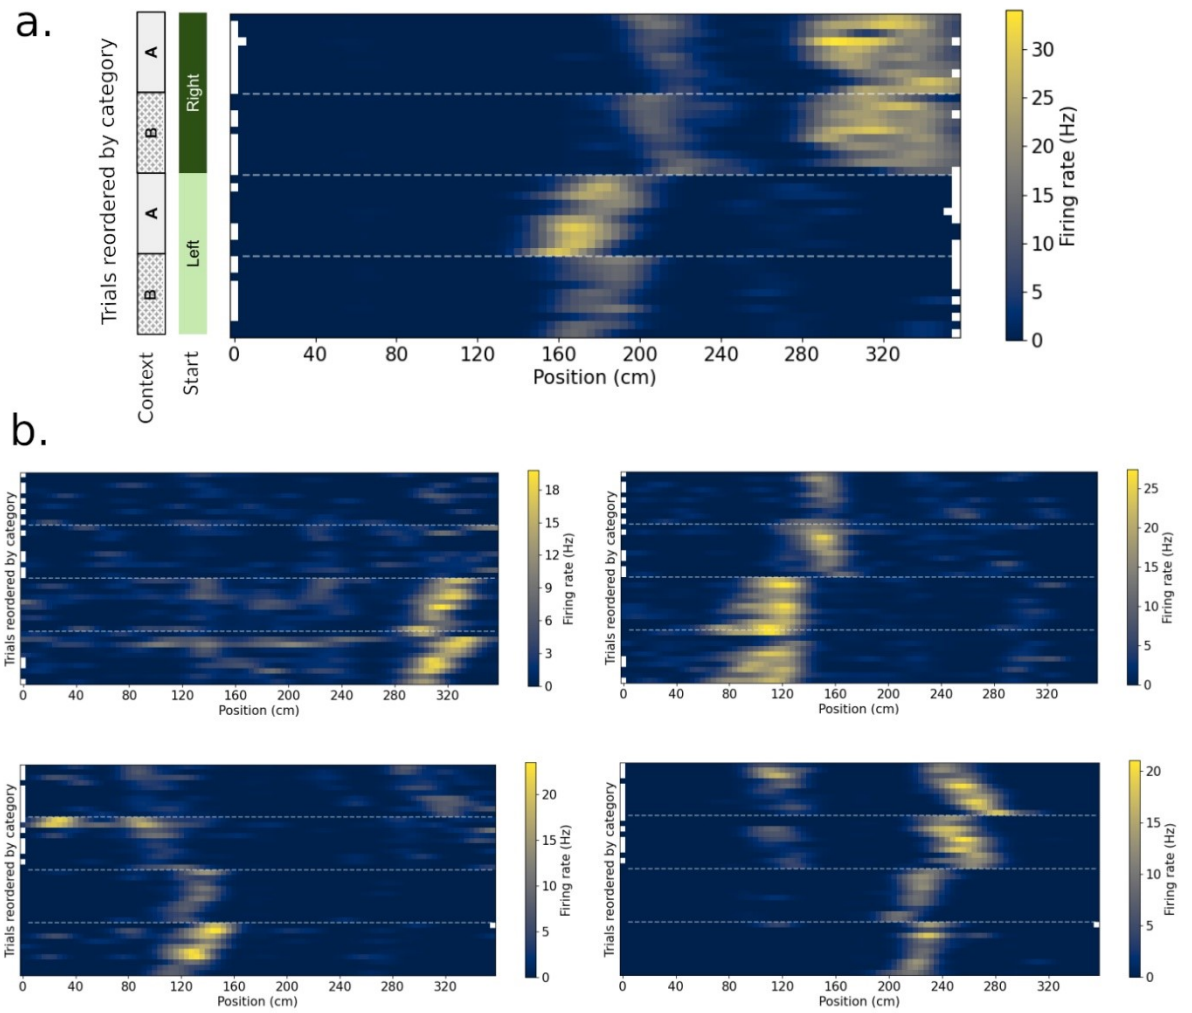

Figure 5.2. Examples of place cell firing rate maps. a) Example place cell with context-specific place fields. Each row represents a trial, plotted separated by their categories (during behaviour they are intermixed). Within each category, trials were ordered from top to bottom. "Start" refers to the side of the maze where the trial started, which defines the main running direction of that trial which is used for rate map calculation (e.g. a right start means the map was calculated only when the animal was running towards the left). This cell shows global remapping between left and right runs in both contexts. It also shows rate remapping between contexts A and B in runs starting on the left side. b) Other 4 example pyramidal neurons classified as place cells, from different animals and training sessions. The trials were ordered the same way as in panel (a).

Place field remapping

Considering that most place cells did not seem to respond equally in all trial categories, I was interested in how they remapped when I changed one task variable at a time. For a fixed context, cell activity was highly modulated by trial direction, which resulted in a high proportion of cells globally remapping, as well as some rate remapping (Figure 5.3a). This is in accordance with previous literature on rodents moving in 1D environments (McNaughton et al., 1983). For a fixed movement direction, most cells did not remap between contexts, although a small percentage showed rate remapping (Figure 5.3b). A recent report in a similar behaviour paradigm describes global rather than rate remapping in response to context change (Zemla et al., 2022), but my results agree with previous studies that show rate remapping in response to

environment and task cue manipulations (Allen et al., 2012; S. Leutgeb et al., 2005). The difference between studies might be due to different recording methods; the first study estimates rate from calcium transients instead of actual spike counts, which could lead to lower correlations – therefore classification as global remapping – and less sensitivity to small changes in firing rate, missing rate remapping cases.

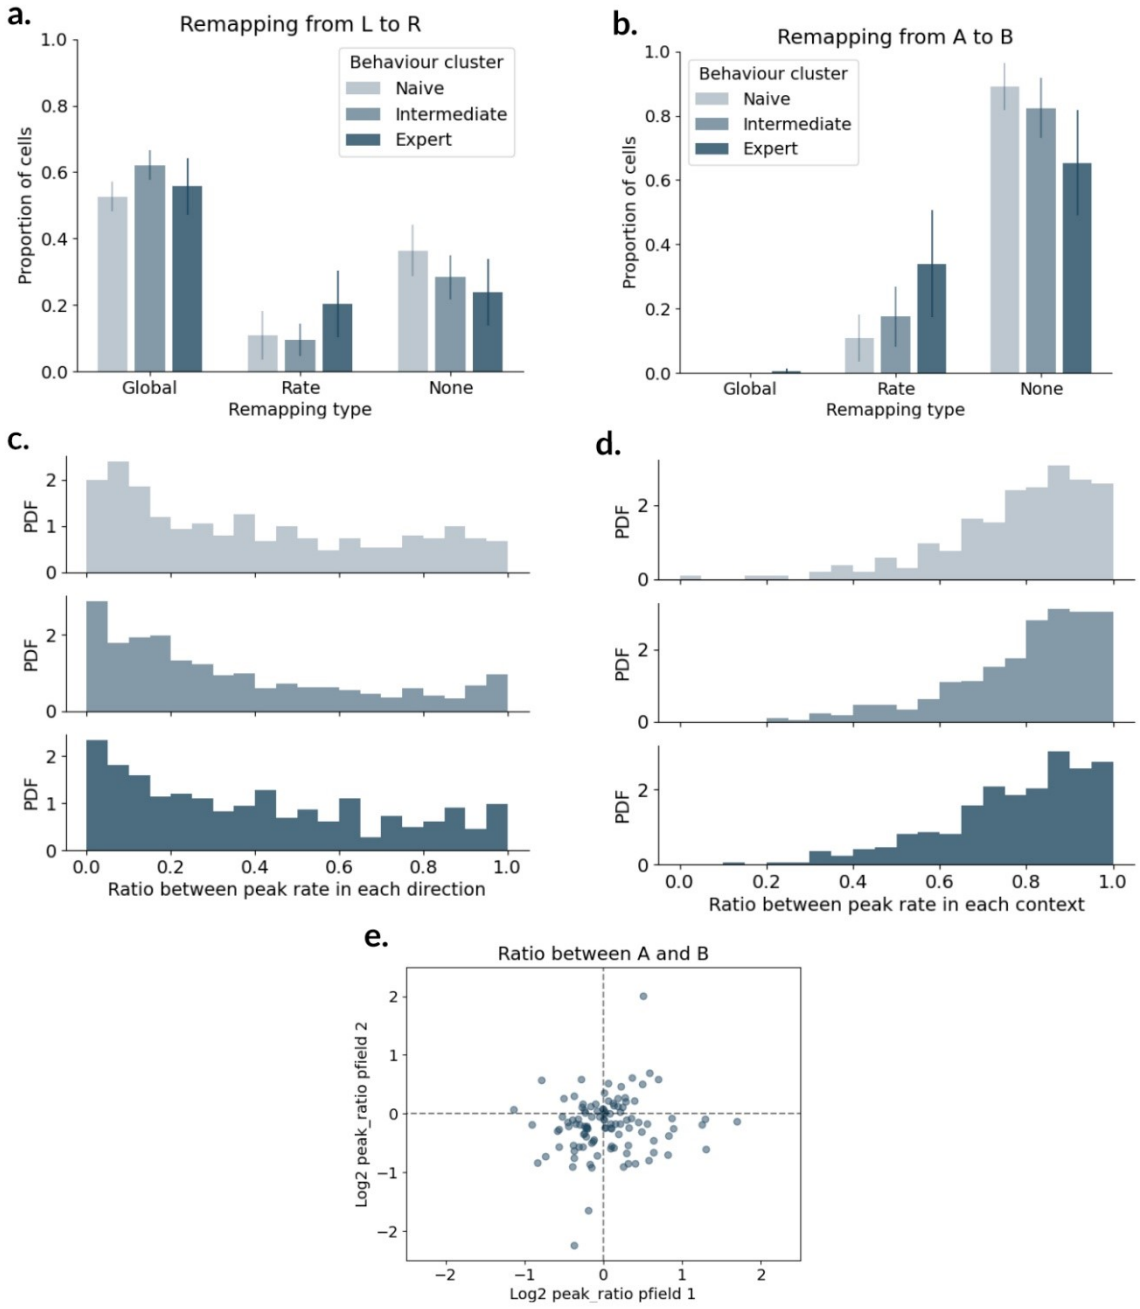

Figure 5.3. Place cell remapping. a) Proportion of place cells that remap between running directions. n=25 sessions. All differences between behaviour clusters in the same remapping category are n.s. (Wilcoxon test with Bonferroni correction, all p>0.5). b) Remapping between contexts. All differences between behaviour clusters in the same remapping category are n.s. (Wilcoxon test with Bonferroni correction, all p>0.15). c) Distribution of ratios between the peak firing rate (lowest/highest) in each direction within the same context, for each place field. d) Same as (c) but between contexts within the same running direction. e) Log from the same values as calculated in (d), for pairs of place fields in the same rate map. Pearson R=0.28, p=7.83e-13. In panel a-b, error bars indicate the SEM.

For a better comparison of rate remapping magnitude, I quantified the ratio between the in-field peak firing rate in each category (smallest/highest peak rate). I did this calculation as long as a place field had been detected in one of the categories in a pair; this allowed the detection of local rate remapping events in globally remapping cells (e.g. when in addition a second field emerged). A ratio close to zero indicates global remapping of the field, whilst other values indicate different strengths of rate remapping. On average, the difference was greater between firing rates of opposite trial directions than between contexts (Figure 5.3c-d), confirming again a stronger effect of directionality than context. Since this analysis was run per place field, I also checked if the remapping effect on two place fields within the same rate map was correlated. My finding was that remapping is specific to each place field, not to each place cell (Figure 5.3e). An example can be seen in Figure 5.2. My initial expectation was that stronger single-cell selectivity to context would emerge over days, as animals become better at associating context to different behaviours, despite no previous work showing it directly. However, there was no immediate difference between behaviour clusters in terms of remapping (Figures 5.3a-d). This indicates that learning this task might be attributed to other changes in the hippocampus or other brain areas.

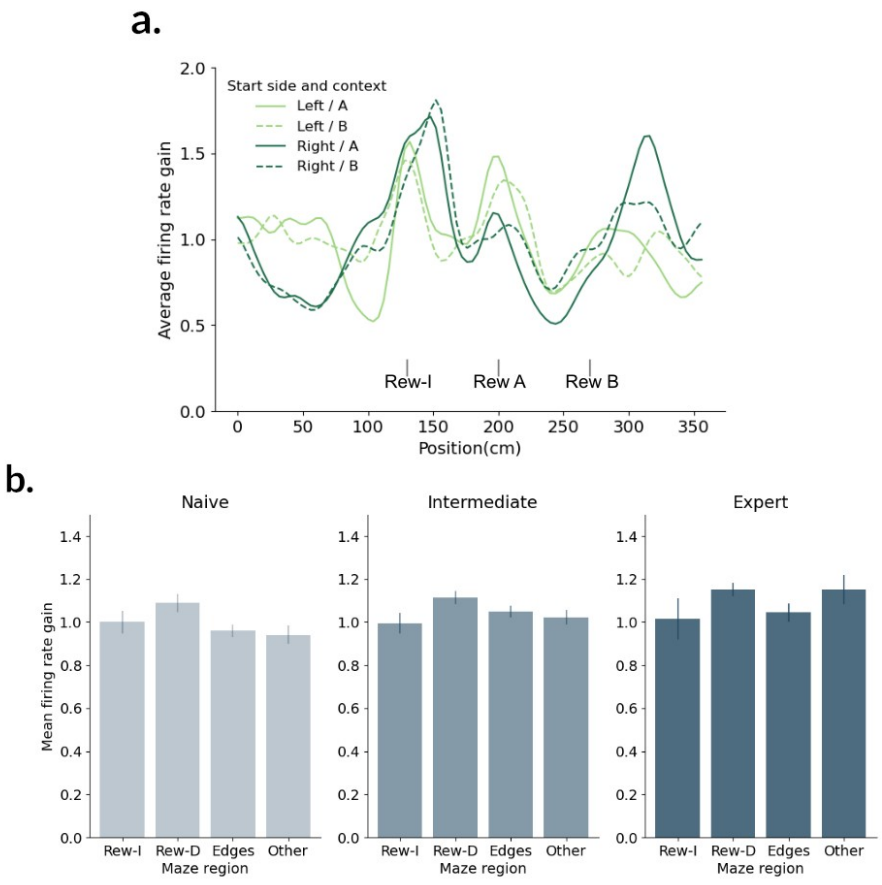

Figure 5.4. Firing rate gain. a) Example of gain around reward. b) Gain per behavioural cluster at different positions. Rew-I = Context Independent Reward, Rew-D = Context Dependent Reward (A and B). Differences between positions are n.s. (Kruskal-Wallis test. Naive p=0.11, Intermediate p=0.06, Expert p=0.14). Error bars indicate the SEM.

Firing rate gain

Next, I checked if the summed activity of all pyramidal cells was dependent on the animal position across the maze. For this I used a measure of firing rate gain, which is the sum of the mean firing rate of all cells per position bin, divided by the average mean rate over all bins, filtered for movement periods (movement velocity >3cm/s). An initial inspection seemed to indicate a higher gain around all three reward locations in all trial categories (Figure 5.4a). However, the strength of this effect was quite variable between recording days and, on average, there was no increase in firing rate gain at reward locations (Figure 5.4b). The variability between the days could have been due to a different proportion of reward-coding neurons in the recorded pyramidal population. According to a recent study, only a small subset of place cells (~20%) shows signs of overrepresentation of reward (Vaidya et al., 2023), which correspond to very few cells in some of my recording days. In that study, this is the same pool of neurons that have a fixed place field over many days.

Despite the mean absolute gain not changing over learning, the difference in gain between contexts did; in one context the activity was below average and in the other, it was above. As the animal learned, the difference became larger at context-dependent rewards (Rew-

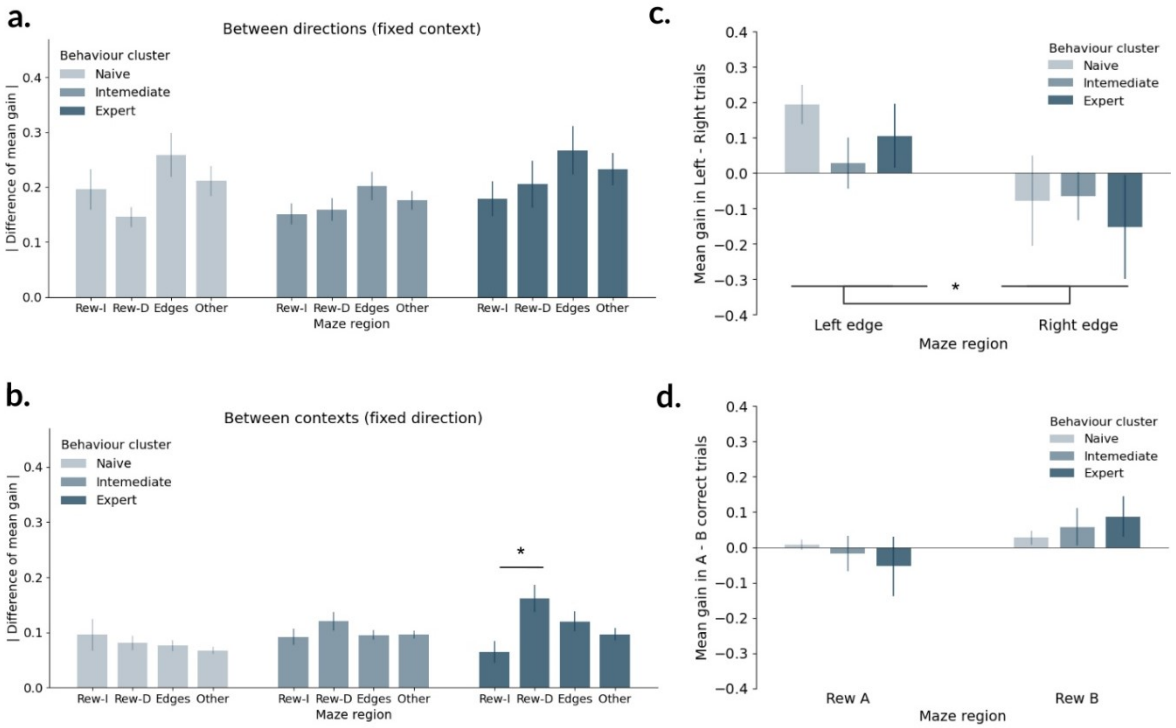

Figure 5.5. Difference in gain between categories. a) Absolute difference in gain between running directions, at specific regions of the maze. Differences between positions are n.s. (Kruskal-Wallis test. Naive p=0.26, Intermediate p=0.59, Expert p=0.70) b) Absolute difference in gain between contexts at specific regions of the maze (Kruskal-Wallis test. Naive p=0.72, Intermediate p=0.76, Expert p=0.02. Dunn's post-hoc in Expert group with Bonferroni correction p=0.03 RI vs RD; p=0.34 RI vs E; p>0.99 RI vs O; p>0.99 RD vs E; p=0.12 RD vs O; p>0.99 E vs O). c) Mean gain in left versus right trials is directional, so the gain is higher when the edge is where the trial starts (Paired t-test, right versus left edge, n=25 sessions, p=0.01). d) Mean gain in A versus B context trials at the context-dependent reward positions. Only correct trials included. There is a trend for higher gain when the position is not rewarded and the animal refrained from digging, but the difference between the two rewards is n.s. (Paired t-test, A vs B reward regions, n=25 sessions, p=0.09). Error bars in all panels indicate SEM.

D) but not in the context-independent case (Rew-I) (Figure 5.5b). When observing only correct trials, there was a trend for higher gain when a context-dependent reward was *not* rewarded (Figure 5.5d), especially in expert sessions. A similar gain difference was also observed between trial directions at the edges of the maze (Figure 5.5a). Those positions are related to entry or exit of the maze, and therefore distinct behaviours: at entry animals are attentive and are likely to be deciding on where to go; at exit, they are focused on receiving a reward at the end box. It is not surprising that there is a difference in neural activity between the two, although this could have been other forms rather than firing rate gain. The gain was higher at entry rather than at the exit of the maze, but no clear trend was visible with learning (Figure 5.5c).

Reward coding

As the training progressed, I expected the accumulation of place fields around reward locations (Dupret et al., 2010; Hollup et al., 2001). This phenomenon, also referred to as goal remapping, is relatively fast, occurring within a few trials of exposure to new rewards. In the case of my task, rewarded locations did not change across days, so I was interested in whether the overrepresentation of goals was restricted to the novelty period or maintained throughout the multiple days necessary for learning.

I focused on reward coding specifically by place cells. For each place field, I measured the distance of the place field to the nearest context-dependent reward at the start and end of the day (Figure 5.6a). I noticed a narrower distribution of place fields around the position of the two context-dependent rewards in expert sessions. This was calculated by comparing the distribution of place field distances to that of a simulated population (with evenly distributed place fields and the same reward locations). I was also interested in when the accumulation happens. As mentioned in Chapter 4, the behavioural cluster 'Intermediate' included sessions in days before the animals achieved asymptotic performance, but sometimes also after, if performance declined on a given day. I extracted from this cluster only the sessions that happened exactly one day before sessions in the 'expert' cluster. What I observed was that the accumulation of place fields at the contextual rewards was already present in these sessions from the beginning of the day, but not in the other sessions classified as intermediate. The same calculation using distance to all three rewards does not show this phenomenon (data not shown).

I then chose a different approach to confirm these results. For each reward position I calculated the proportion of place cells that fired above 2 standard deviations of their rate map average firing rate. One surprising finding was that, despite no change in population gain, the number of place fields at the context-independent reward reduces over learning, which is not the case for other rewards or non-rewarded positions (Figure 5.6b, mean=13% for non-rewarded positions). To my knowledge, there is no previous report of this reduction. Cortical consolidation is likely to play an important role in this task, and it would predict a lesser role of the hippocampus in encoding fixed information such as the context-independent reward. This could be then enacted in the form of less hippocampal resources being allocated to encoding this location. When I compared the width of place fields between reward locations, I observed no

difference at any stage of learning (data not shown). A sharpening of place fields would imply a maintenance of information acuity with less active cells. A different explanation would be that many non-place cells also cover the context-independent reward position since the population firing rate gain was not lower there.

In addition, in the position known for holding reward but not at the current category (Context-dependent reward of the opposite context: Rew-DO) there was the highest

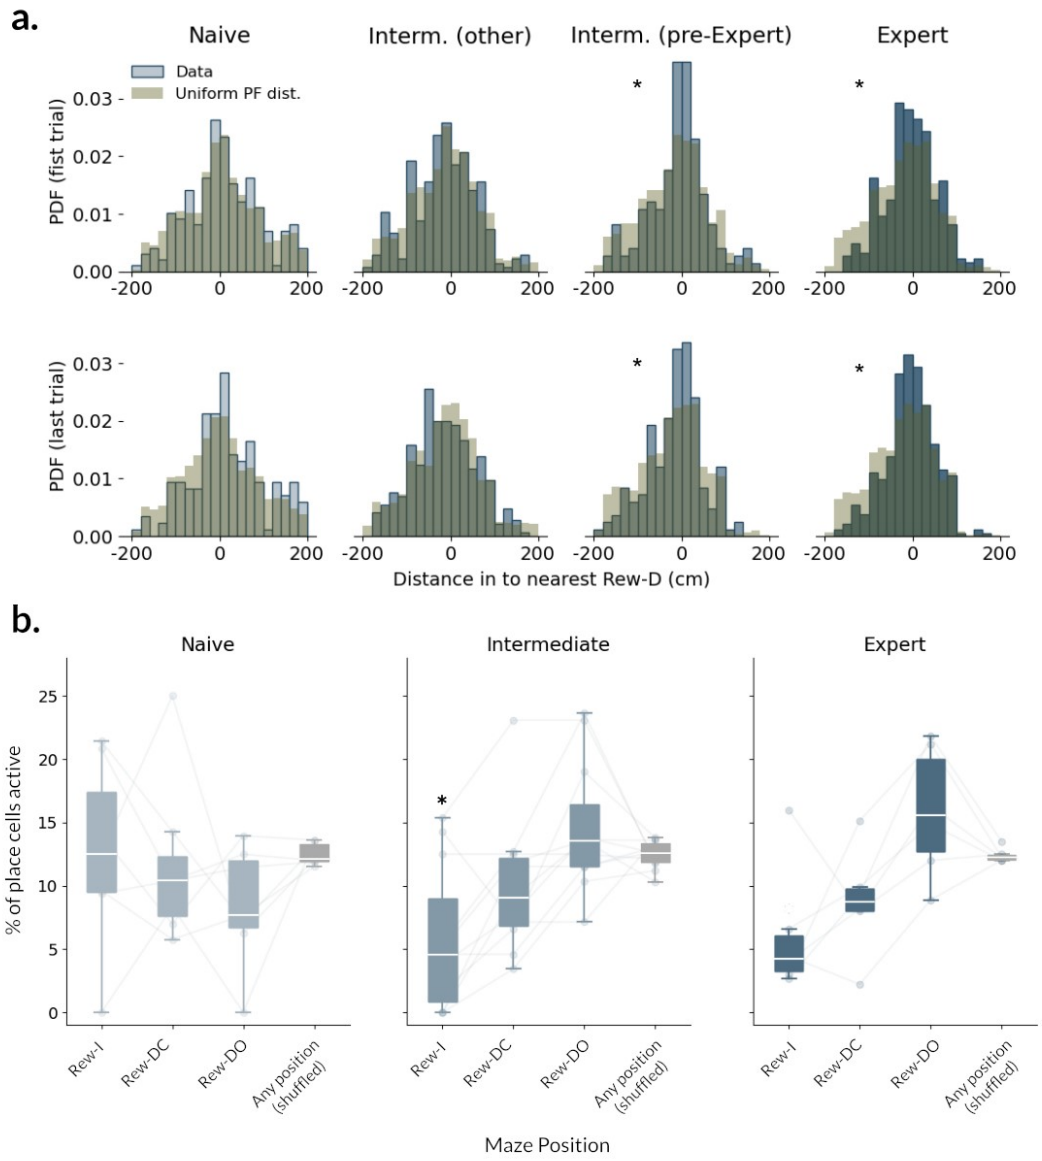

Figure 5.6. Reward coding during learning. a) Concentration of place fields around the context-dependent rewards (Rew-D) depending on behavioural cluster, at the beginning and end of the day. The Intermediate cluster was separated between sessions immediately preceding Expert sessions and other sessions. I also plotted the expected distribution of distances if the place fields (PF) were uniformly distributed in the maze and rewards were located the same as in the experiment (One-sided F-test between uniform distribution and data. First trial: p=0.27 N, p=0.17 I (other), p=0.02 I (pre), p=5e-6 E; Last trial: p=0.37 N, p=0.09 I (other), p=0.01 I (pre), p=2e-6 E) b) Percentage of all place fields that fired significantly around different reward positions, for each behavioural cluster (2 standard deviation above its mean firing per position). In the shuffled condition, the place field of each cell was independently shifted by a random amount and the percentages were then calculated. The process was repeated 100 times for each recording session. Rew-I: Context-independent reward. Rew-DC: Context-dependent reward of the current context. Rew-DO: Context-dependent reward of the opposite context (Wilcoxon test with Holm-Sidak correction comparing each reward position versus the shuffle. Naive: p>0.29 for all rewards, Intermediate: p=0.04 Rew-I, p=0.37 Rew-DC, p=0.37 Rew-DO. Expert: p>0.06 for all rewards). Note that the low number of sessions in the Expert group makes the statistics unreliable in this case, so given more samples Rew-I is likely to be truly lower than the shuffle.

accumulation of place fields in intermediate and expert sessions, as in the firing rate gain results. For the rodents, refraining from digging at this location seems to be the hardest behaviour to be learned in this task. This indicates that reward expectation is not the only driver of place field accumulation but in general positions of behavioural relevance or high cognitive load. Place field accumulation around landmarks has also been shown (Geiller et al., 2017; R. U. Muller & Kubie, 1987), but I am not aware of previous studies showing it at positions determined by behaviour (refrain vs dig) and after days of learning.

Place field shift

Another phenomenon was consistently present across learning days: for a given place cell, as the day progressed the centre of the place fields in each category progressively shifted towards the opposite direction from the running – i.e. the neurons fired earlier at each trial

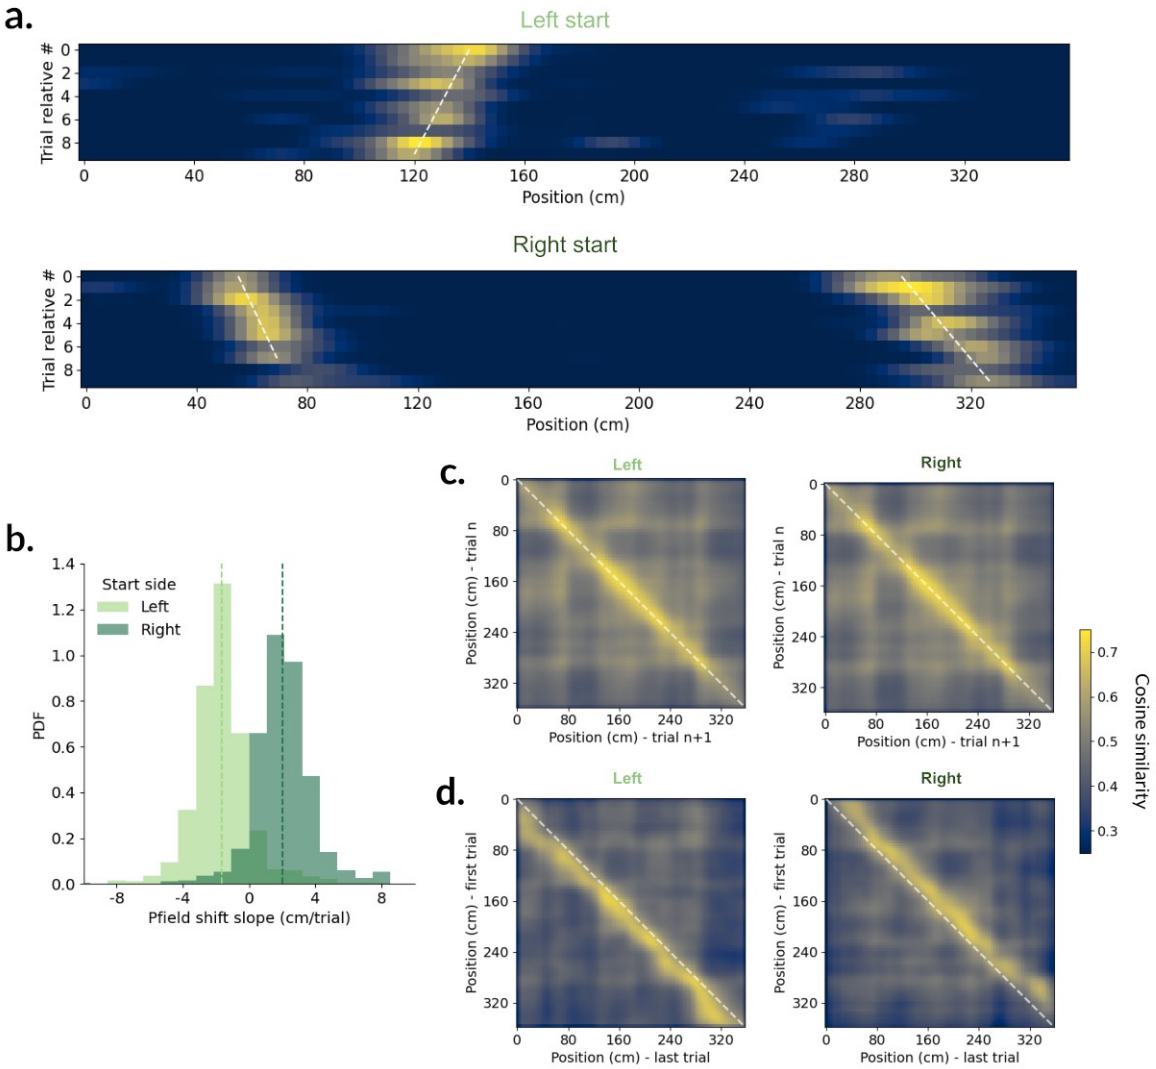

Figure 5.7. Place field shift. a) Two example place fields from different cells, one in left start trials (top) and the other from right trials (bottom). Brighter colours indicate higher firing rate. b) Relative shift slope on left vs right trials, for all place fields. c-d) Population vector similarity across positions (only pyramidal neurons), averaged over all animals. c) Similarity between consecutive trials. d) Similarity between the first and last trial of the day, showing a shift towards the side the animal is coming from.

(Figure 5.7a). Over all place fields, this shift was mostly non-zero (Figure 5.7b) and independent of the distance of the place field centre to the nearest reward (Figure 5.8c). This has been previously reported (I. Lee et al., 2004), but others have also shown the opposite shift, towards the rewards (Xu et al., 2019).

I then defined population vectors containing the mean firing rate of each neuron at a given position and trial. Between neighbouring trials the population vector for a given position was highest correlated with itself (Figure 5.7c), but between the beginning and end of the session, the shift becomes evident (Figure 5.7d). This showed that the effect involved most of the population. These measures were calculated using all sessions, as there was no significant difference in shift when I separated behavioural clusters (data not shown).

The place field shift effect could have been caused by changes in speed as the day progressed since previous reports indicate speed as a strong modulator of the firing of place cells (McNaughton et al., 1983). Indeed, animal speed increased throughout each day. However, this effect was less pronounced in later days (Figure 5.8b), whilst the shift effect remained the same across days (Figure 5.8a). To confirm that speed was not the sole driver of the shift, I calculated the partial correlation of shift with speed and trial number (as a proxy for time). The results

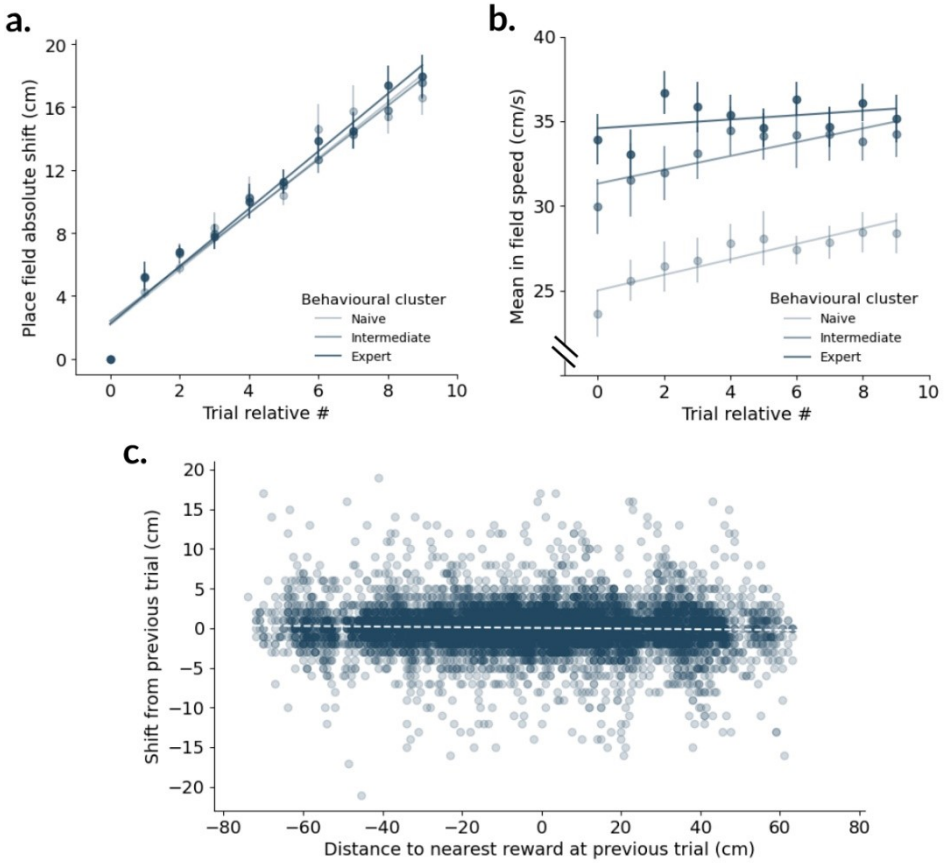

Figure 5.8. Place field shift correlation to speed and time. a) Absolute shift of the place field as trials progress, for each behavioural cluster and the corresponding linear fit. Since place fields are category specific and trials for different categories are interleaved during training, trials numbers are its relative number within the category. b) Mean speed within the place field at each trial, over all place fields, for each behavioural cluster, and corresponding linear fits. c) Neither shift magnitude (in cm) of the place field centre nor its directionality are related to the distance to the rewards in the maze.

showed a non-zero contribution of speed to place field shift, but also a strong independent effect of trial number (linear mixed effects model: coef=1.45 and $p<1.5e^{-4}$ for speed, coef=2.74 and $p<3.5e^{-20}$ for trial number and coef=2.76 and $p<2.4e^{-5}$ for their interaction). If I shuffled the place field centre position between the trials the correlations were abolished.

Place field shift backwards has been suggested to underlie behavioural time scale plasticity (Priestley et al., 2022). In that report, however, the shift happens between the first and second exposure to the environment and is not observed in subsequent trials. Another hypothesis, which is in stronger agreement with my observations, is that shift reflects predictive coding. If the place fields do not represent the current position but the expected future position, then the place field peak is indeed expected to shift backwards and be modulated by speed (Battaglia et al., 2004; Chen et al., 2013). Theoretical accounts of predictive coding also foresee earlier firing as the animal becomes more certain of the upcoming experience (Stachenfeld et al., 2017).

Interneurons

The analyses above were performed on pyramidal neurons and, more specifically, place cells. Nevertheless, I was also interested in the changes in interneuron firing during this task. My first observation was that many interneurons showed a monotonic change in mean firing rate as the day progressed. For some, this change was linear, whilst others seemed to suddenly increase or decrease their mean firing once during the day, between two stable means (linear

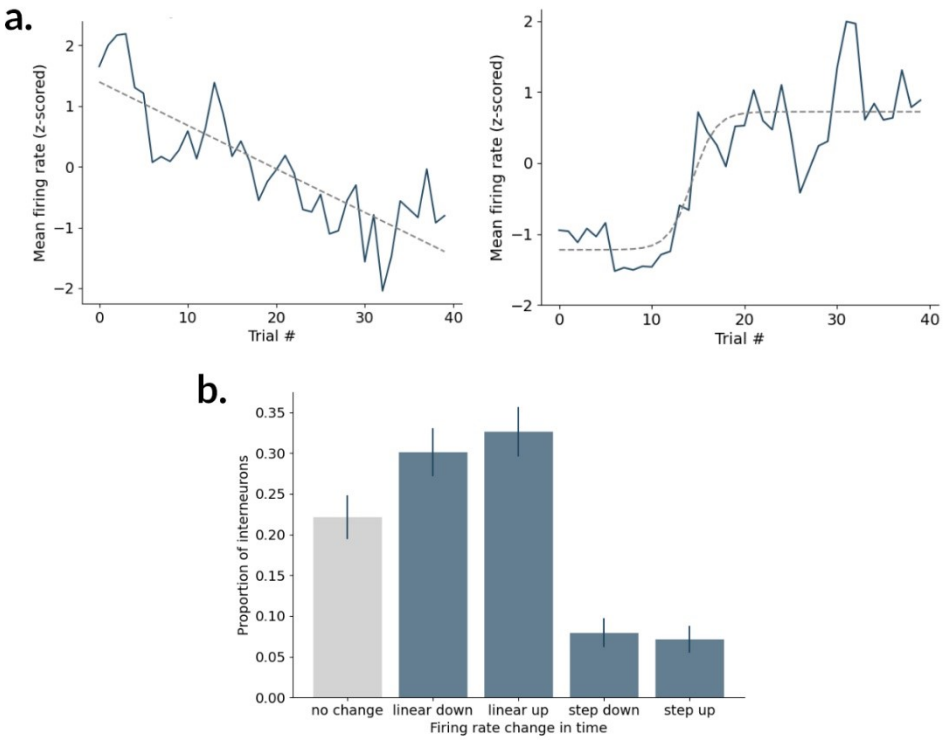

Figure 5.9. Interneuron firing rate trends in time. a) Example of interneurons with firing rate drift over time, either in a linear trend (left) or a step-like change (right). b) Proportion of interneurons with each type of shift. Error bars are the standard deviation from a binomial distribution using Wilson Score Interval. (n=185).

and step, respectively, Figure 5.9a). In total, around 63% of total interneurons recorded showed a linear trend, 15% a step change and the remainder had no significant change in mean rate (Figure 5.9b). The linear trend has been previously described in the literature and associated with learning (Dupret et al., 2013). The small proportion of 'step change' interneurons in my dataset, however, prevents me from making conclusions about their function. Further studies aiming at recording from the interneuron layers of the dorsal CA1 are necessary for characterising them better.

On top of their general trend, interneurons exhibited fluctuations on a trial-by-trial basis. Interneuron selectivity to spatial features of the environment has been previously reported (Ego-Stengel & Wilson, 2007), so I wondered if these fluctuations were related to specific task variables. In Figure 5.10a I show the mean firing rate over trials of an example interneuron after detrending (subtraction of the line or step fit calculated above, see Methods). In this case, the interneuron seems to be much more active in trials of context B, starting on the left side. When I measured this selectivity for all recorded interneurons, around 39% were selective to one of the two variables (trial start side or context) or both (Figure 5.10b). Given that a great portion of the inhibition is local, the selectivity of principal cells to specific task variables is likely to drive different interneurons at each task condition. I did not observe any difference between behavioural clusters in terms of selectivity to different trial directions or contexts (Figure 5.10c).

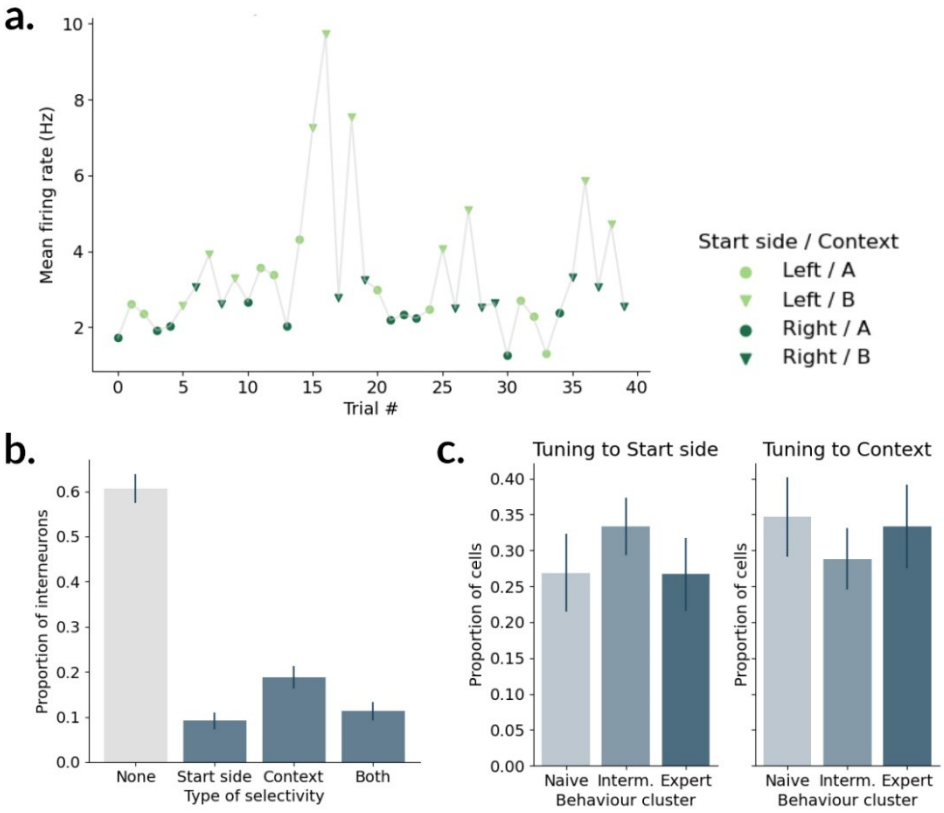

Figure 5.10. Interneuron tuning to task variables. a) Selectivity to different task variables from an example neuron. b) Proportion of interneurons selective to different task variables over the entire recording, determined by correlation with each variable using two-way ANOVA. Error bars are the standard deviation from a binomial distribution using Wilson Score Interval. c) Same as the two middle columns in (b) but separated by behavioural cluster. In panel b-c, error bars indicate the SEM.

Discussion

These single-cell results recapitulate a lot of different findings from the hippocampus literature. I showed that, in the task I designed, measures such as the number of place cells, the width of place fields and the magnitude of place field shift were stable throughout learning days. Place cell activity was not the same in all reward positions but depended on their behavioural significance, a difference that emerged only once animals became experts in the task. I also showed that the firing rate of many interneurons monotonically changed over time and that they could be tuned to different task variables. Here I discuss how these findings match and expand what has been previously reported about learning-related changes in hippocampus dorsal CA1.

My data showed that the response of a large percentage of neurons, both pyramidal and interneurons, was modulated by task variables such as context and direction of movement as well as their combinations. Nevertheless, since many cells did not remap between conditions or only rate remapped, the representation of different contexts was not fully dissimilar. In agreement with this, it has been shown that overlapping features of the environment drive overlapping representations in hippocampal CA1, as opposed to the separate subpopulations of neurons encoding each context/environment in CA3 (S. Leutgeb et al., 2004). It is also known that the response of these CA1 place fields to changes in the environment is slower than in CA3 (I. Lee et al., 2004).

The number of neurons tuned to space (place cells) and context remained constant over days, largely unaffected by familiarity. In a novel environment, dopamine signals during behaviour are thought to promote learning, via the facilitation of long-term potentiation of CA1 synapses and increased assembly reactivation during sleep (Li et al., 2003; C. G. McNamara et al., 2014). This dopaminergic effect is likely absent or reduced after the first day of my paradigm, and therefore the drivers of learning and the stability I observe in this case are less understood. This might be related to the increased homogeneity of inputs to CA1 after repetitions when compared to exposure to a new environment (Cohen et al., 2017). A consistent input could explain the stable quality of place and context coding in the hippocampus with familiarity, despite the reduced role of dopamine and the gradual change in the active ensembles (Mankin et al., 2012).

I also show that the encoding of reward positions by place cells was dependent on their behavioural significance, or cognitive load. At initial learning, novelty and reward signals from other brain areas were likely to have a strong influence on hippocampal activity. This would explain why all reward regions were covered by a similar number of place cells; also the slightly higher activity around the reward positions the animal actually received a reward in a given context. Later in learning, rather the opposite was observed, with the position requiring refraining behaviour – and therefore at the category in which it was not rewarded – being the most represented. I show this both in terms of place cells and in the firing rate gain of the population. This was an indication that the main drivers of CA1 activity changed over days and that during learning the cell tuning to reward became more refined.

Previous research has shown that different reward contingencies modulate hippocampal activity (H. Lee et al., 2012), with more activity expected when reward is less certain, at least in humans (Vanni-Mercier et al., 2009). It is also known that dopaminergic activity is higher in more uncertain states (Fiorillo et al., 2003). In the paradigm I show here, probability of reward at context-dependent positions is 50% if the animal is agnostic to context, but learning the associations showed that it is deterministic. Although uncertainty signals could explain the neural activity earlier in learning, motivational state or task-specific movement at each reward position might be better explanations later on (Kennedy & Shapiro, 2009; Wiener et al., 1989). The position unrewarded in a category can be seen as a punishment; the animal was limited to two choices per trial and, if they dug there, they got one reward less as a consequence. Therefore the animal is likely to be in a different mental state when refraining and digging. I am not aware of studies of hippocampal associative learning that investigate neural tuning when the association determines very different behaviours at the same position. Our unlikely observation of stronger firing during refrain opens venues for investigation of the mechanisms controlling place cell recruitment beyond reward signals.

A similar phenomenon to that of reward is observed at the edges of the maze, where the animal starts and terminates the trials. The same position shows a stronger overall population activity (gain) at the start of the trial. I argue that this is the point where the animal integrates contextual cues and makes a decision on where to dig next, which can be considered a higher cognitive load than running towards the reward at the end box once the trial is over. This hypothesis is in accordance with previous studies that show higher hippocampal activity and theta power during decision-making periods (Belchior et al., 2014).

I also reported place field shifts in position, previously hypothesised to contribute to predictive coding and learning (Battaglia et al., 2004; Chen et al., 2013; Stachenfeld et al., 2017). In my data, the shift happens every day, in the same timescale from other reports: within dozens of trials within a single day. What previous studies do not show and I observed here is that the shift is present over long periods of learning, even when the animal is familiar with the task. Future research with different techniques is necessary to track the same neurons over the entire learning period, to understand if the shift is continuous for stable cells or if there is a "reset" of place field position during sleep (Vaidya et al., 2023). This would give insight if place field shift is a sign of predictive learning over many days or if other, slower processes are at play.

Overall, the statistics of hippocampal single-cell tuning was not greatly altered over days despite the animal's increasing performance. These performance changes are likely to be a result of an increasing engagement of cognitive cortical areas over time. Goal-oriented navigation based on memory is known for relying on the hippocampus in both novel and familiar environments, whilst cortical engagement is more prominent in the latter, at least in humans (Patai et al., 2019). A sign of change in cortical contribution over learning is the reduction in the behavioural effect of CA1 lesions versus an increase in that of prefrontal areas during long-term recall of associations in rodents (Tse et al., 2007). In the next chapter, I also investigate how undetected changes in the single cell level can accumulate and translate into significant

differences in the population activity, which can then be used by downstream cortical areas for better task solving.

The data reported here aimed at understanding changes in the neural code with learning. For this purpose, I separated learning stages in behavioural clusters and focused most of our statistical analysis on comparing these clusters. Each cluster contained a different number of training sessions, and therefore different animals were more or less represented in each. I made an effort to select statistical methods that control for these factors; those have the drawback of reducing the statistical power usually obtained by pooling together neurons from all animals, as is customary in the field. As a result, many of the trends observed here were “non-significant”, but can still be of relevance and give insight into learning-related neural changes in this particular behavioural task.

6. Neural population

In rodents, it has been shown that the emergence of place fields in CA1 is a random Poisson process related to a cell's inputs and its intrinsic excitability (J. S. Lee et al., 2020). As a consequence, the activity of a single cell at a given location does not provide direct information about reward or other environmental variables. According to Lee et al. (J. S. Lee et al., 2020), rather than deterministically recruiting reward-coding neurons, reward signals lead to an upward modulation of the excitability of all hippocampal neurons specifically at those locations. This results in an overall increase in the number of fields emerging around the reward location, which makes information about rewards available only at the population level. If a variable – e.g. context – affects a place cell's inputs instead of its excitability, it would alter instead the maze locations where place fields are likely to emerge. Similarly, downstream areas would only be able to determine context from a set of simultaneously active neurons and not from the presence or absence of a field in a single cell. These findings highlight the relevance of studying hippocampal representations at the population level.

A previous study in primates shows that, given enough neurons, linear classifiers can be used to decode any arbitrary variable encoded in the hippocampal population (Bernardi et al., 2020). Linear decodability at the population level does not require neurons to be exclusively selective to single variables, but it is equally effective if they have mixed selectivity (Bernardi et al., 2020; Kaufman et al., 2022). This divisibility of the activity patterns in any arbitrary set would suggest they are randomly spread in the neural activity space. Contrary to this intuition, however, is the observation of a non-random organisation of the neural activity that facilitates their generalisation to novel situations (Bernardi et al., 2020) and a structure in the population activity that reflects task-imposed associations between variables (Gulli et al., 2020) (for an opposing view on hippocampal generalisation see (Kaefer et al., 2020; Samborska et al., 2022)). The processes by which these structured representations are formed during learning remain to be elucidated.

Task-variables can be considered decodable from the neural population activity if the activity patterns corresponding to each value of this variable are in some way separated in the neural activity space. If the decodability is linear, this means that the separating boundary is a line or high-dimensional plane. The value of the task-variable can be then predicted (i.e. decoded) when a novel activity pattern is observed, based on which side of the plane this activity lies. If any arbitrary division of the neural activity is linearly decodable (i.e. can be separated by a plane), it also implies that the hippocampal code is high-dimensional, i.e. not collapsed into too few patterns of activity (Bernardi et al., 2020). However, others argue that behaviour constrains the neural activity to a low-dimensional manifold (Nieh et al., 2021). These results are not necessarily contradictory; a non-trivial organisation of the neural code could be the result of patterns of activity lying in a non-linear, low-dimensional manifold embedded in a high-

dimensional space. Different definitions of dimensionality and higher dimensionality estimates yielded from linear methods often hinder the comparison between studies (Altan et al., 2021). Relative dimensionality measures are nevertheless a valuable tool for understanding the dynamics of the neural code during behaviour.

In the current chapter, I aimed to investigate the simultaneous representation of task variables in the rat hippocampus during learning; I followed the population activity as variables became gradually associated in behaviour due to task demands. The goal was to identify if these associations were also present in the hippocampal code, what was the structure of the representation at different stages of learning and how fast it emerged in relation to behaviour.

Principal Component Analysis

My first approach to visualising the population activity in lower dimensions was to perform Principal Component Analysis (PCA). From a mathematical perspective, the significance of dimension in this method is straightforward, since each dimension (PC) consists of a weighted sum of the activity of all neurons, is orthogonal to the other dimensions and the order of the principal components is determined by the percentage of variance of the full data that they explain. To calculate the PCA, I obtained a population vector containing the mean firing rate of

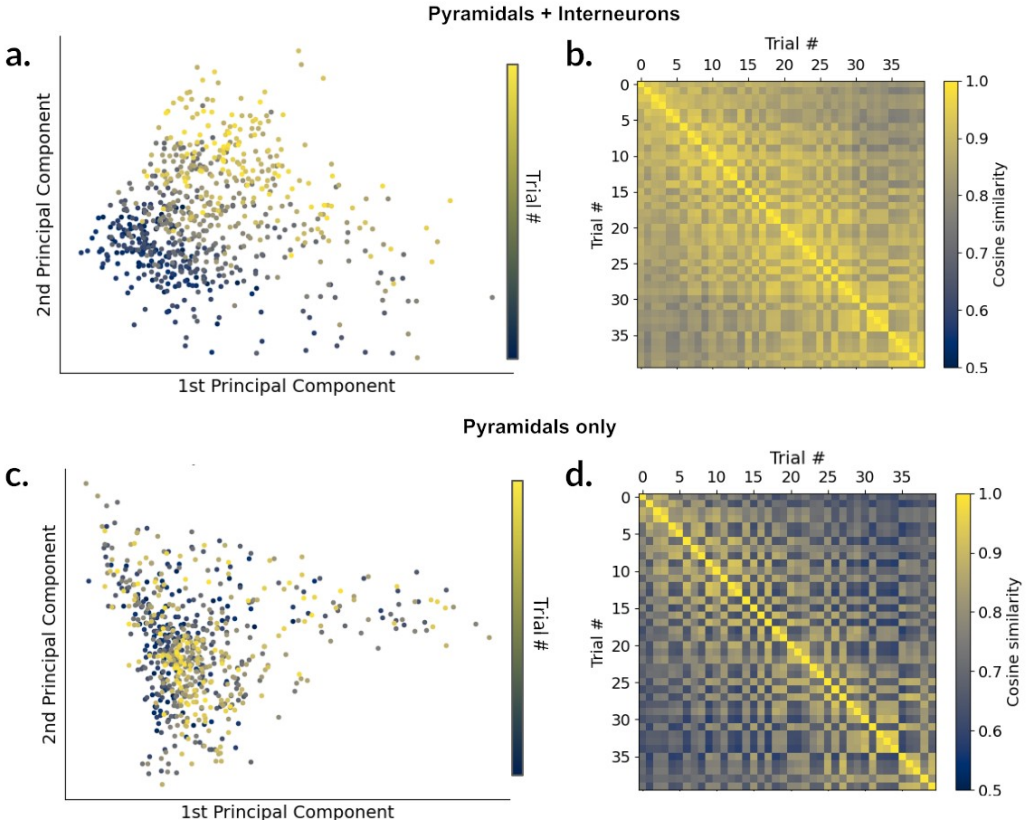

Figure 6.1. Population vectors in time. Example data for a single recording session a) Representation of projection of population vectors calculated per position and trial, projected into the first two principal components. Population vectors were calculated using all neurons in the population. Colour indicates trial numbers. b) Cosine similarity between population vectors in each trial, calculated per position and averaged over all positions. Population vectors were calculated using all neurons in the population. c/d) Same as a/b but calculated using only pyramidal neurons.

the recorded neurons for each maze position in each trial; an n-dimensional vector where n is the total number of neurons in the dataset. The PCA was then calculated using all the vectors in a given session.

In Figures 6.1 and 6.2 I plotted the projection of each population vector into the first 2 PCs and checked if vectors calculated in similar task conditions (e.g. same trial) lay closer in this low dimensional space. If I included both interneurons and putative pyramidal neurons, I observed a strong correlation of the first two PCs with trial numbers – a proxy for time (Figure 6.1a). This relationship was much weaker if only pyramidal neurons were used in the calculation (Figure 6.1c). The cosine similarity between population vectors over trials confirmed this result (Figures 6.1b/d), which was expected due to the strong time drift of firing rate in half the interneuron population, shown in Chapter 5. Similarly, the presence of place cells and other spatially tuned pyramidal neurons resulted in nearby positions being closely represented in the first PCs (Figure 6.2a). Other variables such as context and direction of movement did not seem

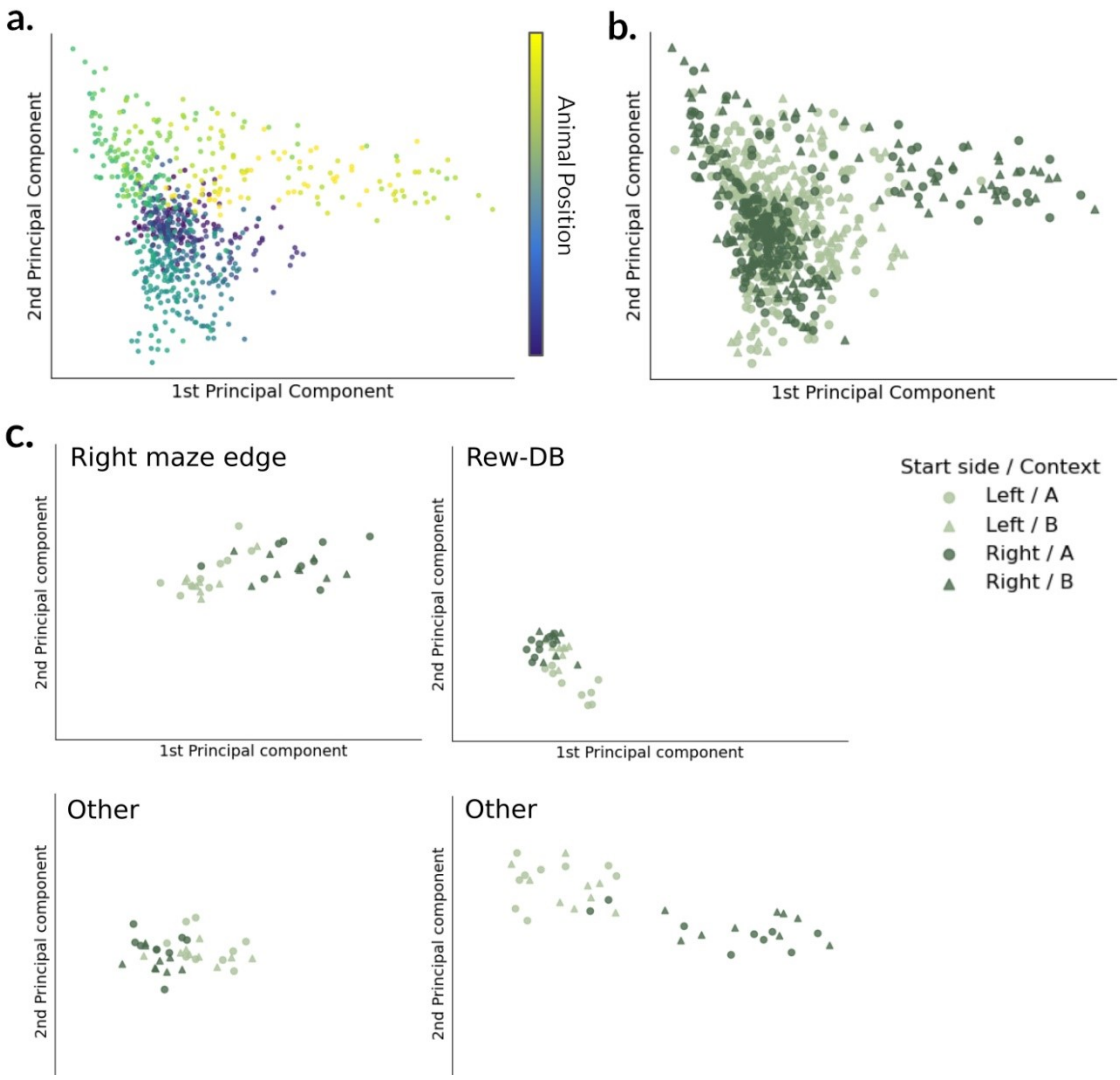

Figure 6.2. Task variable representation in PCA space. Example projection of population vectors calculated per position and trial in one session, using only pyramidal neurons and projected into the first two principal components. a) Color indicates animal position in the maze. b) Colors and markers indicate start side of the trial (i.e. direction of movement) and context. c) Same as (b) but plotting individual positions separately for visualisation. Right maze edge: first position bin at the right edge of the maze. Rew-DB: Context-dependent reward of context B. Other: positions that are neither reward nor edges of the maze.

to be organised within the first principal components at first glance (Figure 6.2b). However, when plotting vectors from a single position, there seemed to be some organisation of at least one of these variables (Figure 6.2c). To investigate the encoding of such task variables, for the remainder of this chapter, I focused on the population vectors containing only putative pyramidal neurons.

Non-linear methods

Visualising principal components is unfortunately restricted to 2 or 3 dimensions at a time, and it is overall a challenge to visually grasp the structure of the data in as few as 4 or 5 dimensions. Yet, these dimensions still hold a lot of the variance of the data (Figure 6.8a-b). Therefore, I recurred to non-linear high-dimensional data visualisation methods to explore other structures in the data that might not be easily observed with PCA. Here I show results from t-

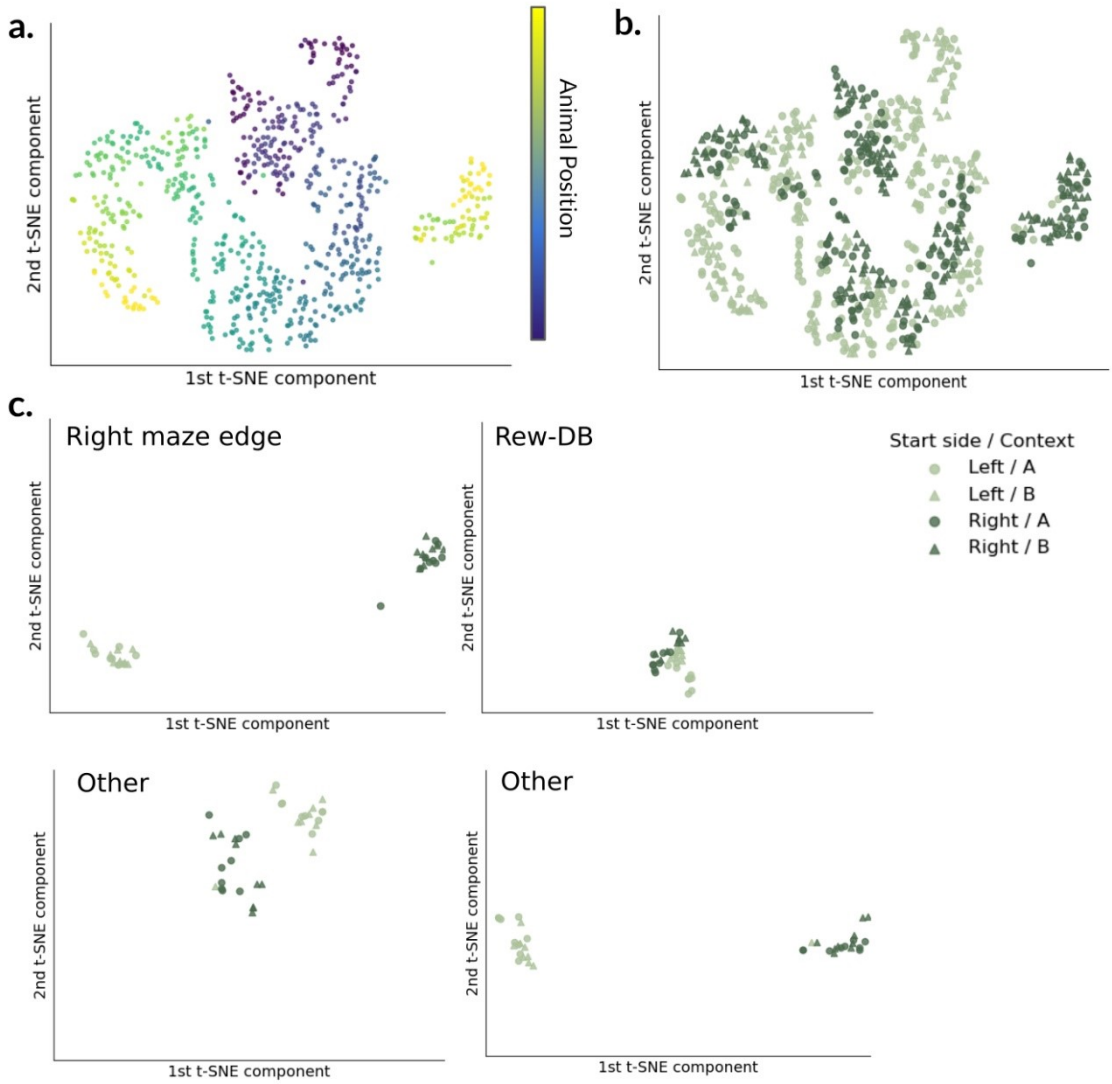

Figure 6.3. Task variable representation in t-SNE. Example projection of population vectors calculated per position and trial in one session, using only pyramidal neurons and embedded in a 2-dimensional t-SNE, perplexity=50. This is the same recording session as in figure 6.2. a) Colour indicates animal position in the maze. b) Colours and markers indicate start side of the trial (i.e. direction of movement) and context. c) Same as (b) but plotting individual positions separately for visualisation. Right maze edge: first position bin at the right edge of the maze. Rew-DB: Context-dependent reward of context B. Other: positions that are neither reward nor edges of the maze.

SNE, as other non-linear methods such as UMAP and Multidimensional Scaling (MDS) yielded qualitatively similar results.

In the 2-dimensional t-SNE, the separation of positions was even clearer than in the first few principal components (Figure 6.3a) and, similarly, I observed only local separation of the other task variables (Figure 6.3b-c). t-SNE tries to balance global versus local distances from higher dimensions when calculating distances in the low-dimensional projection; nearby points in the original data lay close also in the t-SNE projection, giving an intuition on their structure (Maaten & Hinton, 2008). This became clear in Figure 6.3a, where a trajectory-like organisation can be observed following the sequence of positions. This is expected in the hippocampus since there is usually an overlap of active place cells between nearby positions, creating smooth transitions in neural activity space as the animal traverses the environment. In PCA space, this trajectory was also present, but better observed using 3 dimensions.

Using a 3D PCA projection, it was also possible to observe a split of the representation in one of the maze ends, where each side of the split corresponded to the representation of that position in the one running direction. These two subgroups were equidistant to the representation of the other positions. In the t-SNE embedding, this was represented as two extremely distant groups (position coloured yellow in Figure 6.3a), one of which does not seem to be continuous with the other positions. This emerges because, if the data structure has inherently more dimensions than the t-SNE embedding, clusters of points might be placed further than they are in reality. This also suggests that the structure of this representation is probably more than 2-dimensional, which resulted in the distortions of the distances in the case of the 2D embedding.

In addition to the distortions due to dimensionality, in t-SNE the local/global balance of distances is determined by a perplexity parameter that highly influences the projected outcome. t-SNE and UMAP algorithms are also stochastic, which means that every initialization yields a slightly different projection. So although these are powerful tools for visualising high-dimensional data, they are not suitable for quantitative analysis based on the Euclidian distances between points. In this project, t-SNE results served to confirm the qualitative observation obtained from PCA, which could then be used in subsequent analysis for the characterisation of the representational geometry in Euclidean space.

Decoding of task variables

Given the structure observed in the low-dimensional representations of the data – global separation of positions, local separation of direction and context – I looked for quantitative ways to assess this structure. My first step was to check how accurately a simple linear decoder (i.e. a Support Vector Machine, see Methods) could separate different task variables. I denote this a "global" decoder, since it is trained using population vectors from all positions and trials. In Figure 6.5a I show the accuracy for independent classification of these three task variables: trial context, direction, and maze position. For all of them, the classifier accuracy was above chance

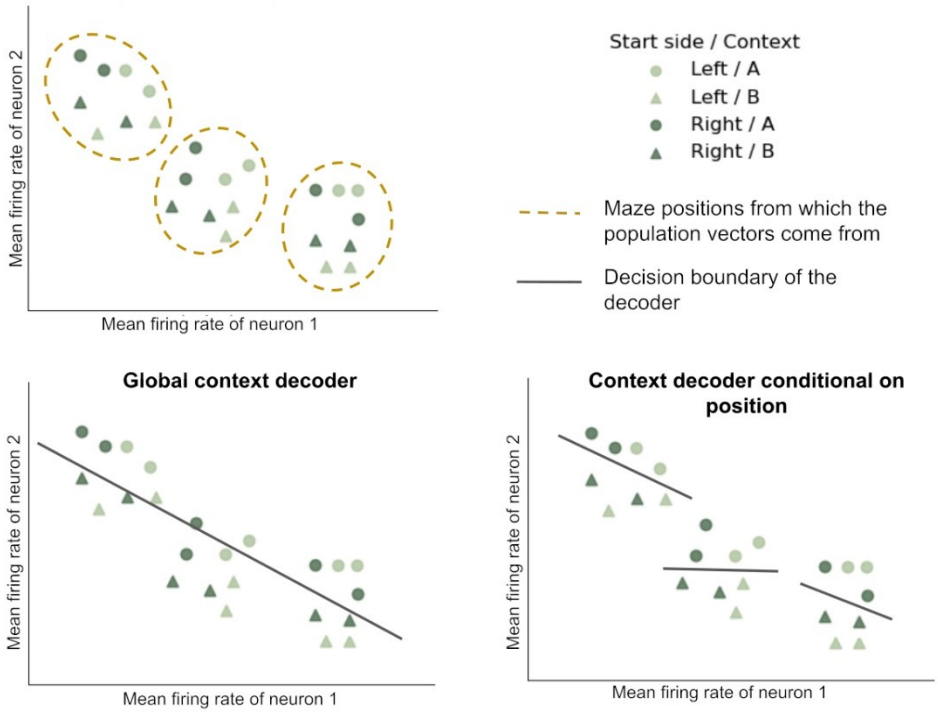

Figure 6.4. Schematic of global vs. conditional decoders. Schematic of the difference between using a global linear decoder (Support Vector Machine) or multiple, one for each conditional variable. In this example the variable decoded is context and the conditional variable is maze position. Both context and position are labels assigned to each population vector and the decoder finds the high-dimensional plane that best separates the labels to be decoded.

levels, with position being the best decoded variable in all behaviour clusters. The accuracy for context was much lower, despite a slight upward trend with learning.

If the local separation of trial direction and context observed in PCA was not a spurious observation it should influence decoding; therefore, I asked whether position-specific decoders for those variables would perform better than a global one. The separating boundary for the variable of interest (the decoded variable) was calculated separately for each label of another variable (the conditional variable), so I called it a conditional decoder. A 2D schematic of the two types of decoders is presented in Figure 6.4. In Figure 6.5b I show a comparison in performance between the two decoders when decoding context alone or conditional on position. I observed a gain in accuracy for the conditional decoder, which indicates that the separating boundary for context was not aligned between positions. To control for overfitting unrelated to the relationship between variables – due to the increased number of parameters – I shuffled the labels for the conditional variable keeping the decoded variable labels intact. If the accuracy of the conditional decoder was simply a result of its number of parameters, it should perform better also in this case. Conversely, if the conditional variable was informative about the decoded variable then the accuracy was expected to drop. I observed no signs of overfitting: shuffled conditional labels yielded worse accuracy than both other decoders.

It is important to note that the decoding accuracy is dependent on the size of the position bins used for calculating the population vectors. Although the trends remain the same, better accuracy is achieved with larger bins, especially for decoding variables such as context and

direction. This might result from averaging out part of the position-related variance when averaging over a larger area of the maze, making differences due to other variables more dominant. For too small spatial bins, the accuracy of context decoding drops close to chance level. Given that the mean place field width is around 48 cm (Chapter 5), if bins are small a lot of the neurons fire equally in nearby bins. In future studies, it might be relevant to use decoding methods which take position as a continuous variable, so that this issue becomes irrelevant.

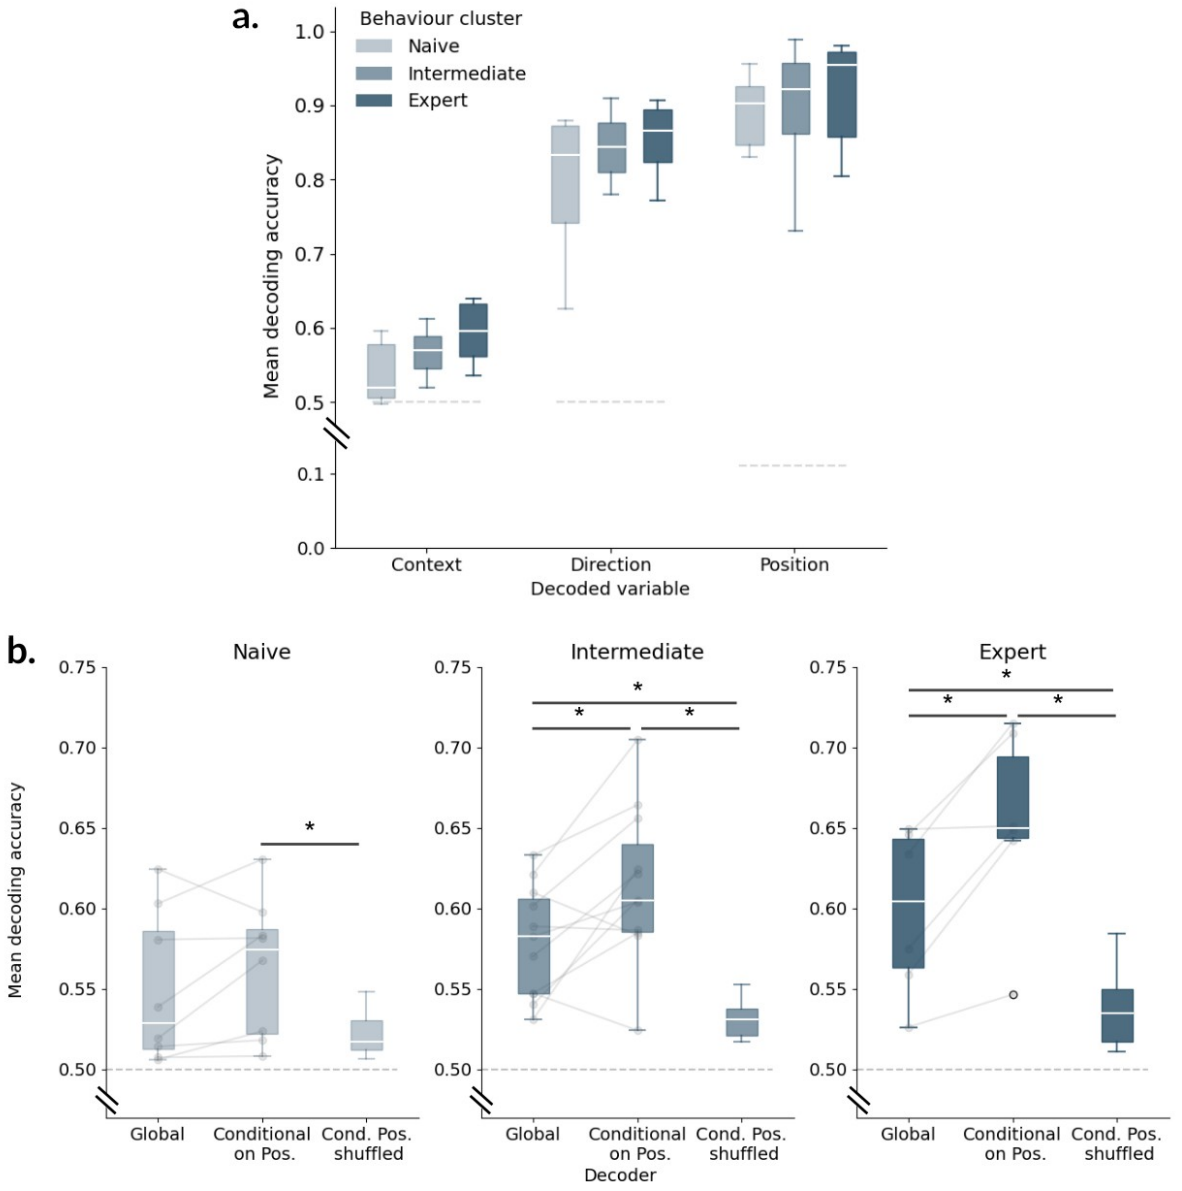

Figure 6.5. Results from linear decoding of tasks variables. a) Decoding accuracy of different task variables using a global decoder, for each behaviour cluster. Accuracy indicates the proportion of predicted labels that match the real labels in the held-out test data set. Positions were divided in discrete bins for this calculation (see Methods). Each boxplot is calculated using one data point per session in the cluster, which represents the mean decoding accuracy for that session, using bootstrapped samples of populations vectors. All differences between behaviour clusters for the same variable are n.s. ($p>0.19$ for all comparisons, Mann-Whitney U with Holm-Sidak correction). b) Comparison of the global decoder of context and the decoding of context conditional on maze position, per behaviour cluster. We also show the accuracy of the conditional decoder in case the conditional labels were shuffled between population vectors. (Paired t-test with Holm-Sidak correction. Naive: Sim vs Cond $p=0.146$, Sim vs Cond(shuffle) $p=0.085$, Cond vs. Cond(shuffle) $p=0.013$; Intermediate: Sim vs Cond $p=0.016$, Sim vs Cond(shuffle) $p=0.001$, Cond vs. Cond(shuffle) $p=0.001$; Expert: Sim vs Cond $p=0.012$, Sim vs Cond(shuffle) $p=0.009$, Cond vs. Cond(shuffle) $p=0.009$.

From the remainder of this chapter, I use 40cm bins, which allow for a clear distinction in performance between the two decoder types.

Decoding on dimension-reduced data

The results from linear decoding were not surprising, given that the number of features (neurons) was much higher than the number of labels to distinguish (2 in the case of context or side, 9 in the case of positions). Considering I observed a separation of variables already in the first few principal components, I then asked whether all dimensions of the data were required to achieve high decoding accuracy. For that, I projected the population vectors into the principal components and showed the decoder only a limited number of PC dimensions at a time. I then tested the decoder performance at all the possible decoded/conditional variable combinations of trial context, direction and maze position.

In Figures 6.6a-d I show the result calculated for the Expert behaviour cluster. As previously shown, decoding of context was better in the conditional case, strongly when conditioned on position but also when conditioned on movement direction (Figures 6.6a-b). It is also important to note that this difference is larger in a "middle-range" of PCs fed to the decoder. One hypothesis is that variability in the first principal components is dominated by differences in position. If I exclude this source of variability by feeding only vectors from the same position, smaller fluctuations in the same PCs caused by context become decodable. The global decoder, on the other hand, would only be capable of extracting contextual information from PCs in which the axis of variability for context is not parallel to that of position, which seems to be at much later PCs. Nevertheless, once the decoders have access to a large enough number of PCs – which is still much less than the total number of dimensions – their performance comes closer again.

A qualitative analysis of the weights assigned to each PC supports the argument above. The global decoder assigned low weights to the initial PCs and its accuracy increase matched the point where later PCs were available and assigned larger weights (Figure 6.7a). Despite the different number of PCs provided in each run, the weights were relatively stable. The conditional decoder, on the other hand, exhibited two regimes: first, early PCs were assigned high weights; once more PCs were available, the weights became distributed. This indicates that although these first PCs provided context information, the middle dimensions added decodability for this variable. This shift also coincided with the point where performance starts to plateau, suggesting that further PCs do not carry much more information about this variable. Weights were considerably different between positions, supporting that the optimal separability between contexts is local (Figure 6.7b). It is important to take into consideration that the different number of samples used for training each parameter versus the number of features provided to the decoder can affect some of the results observed here. Therefore, controlling for such factors and quantifying these values across all sessions is necessary to confirm these results.

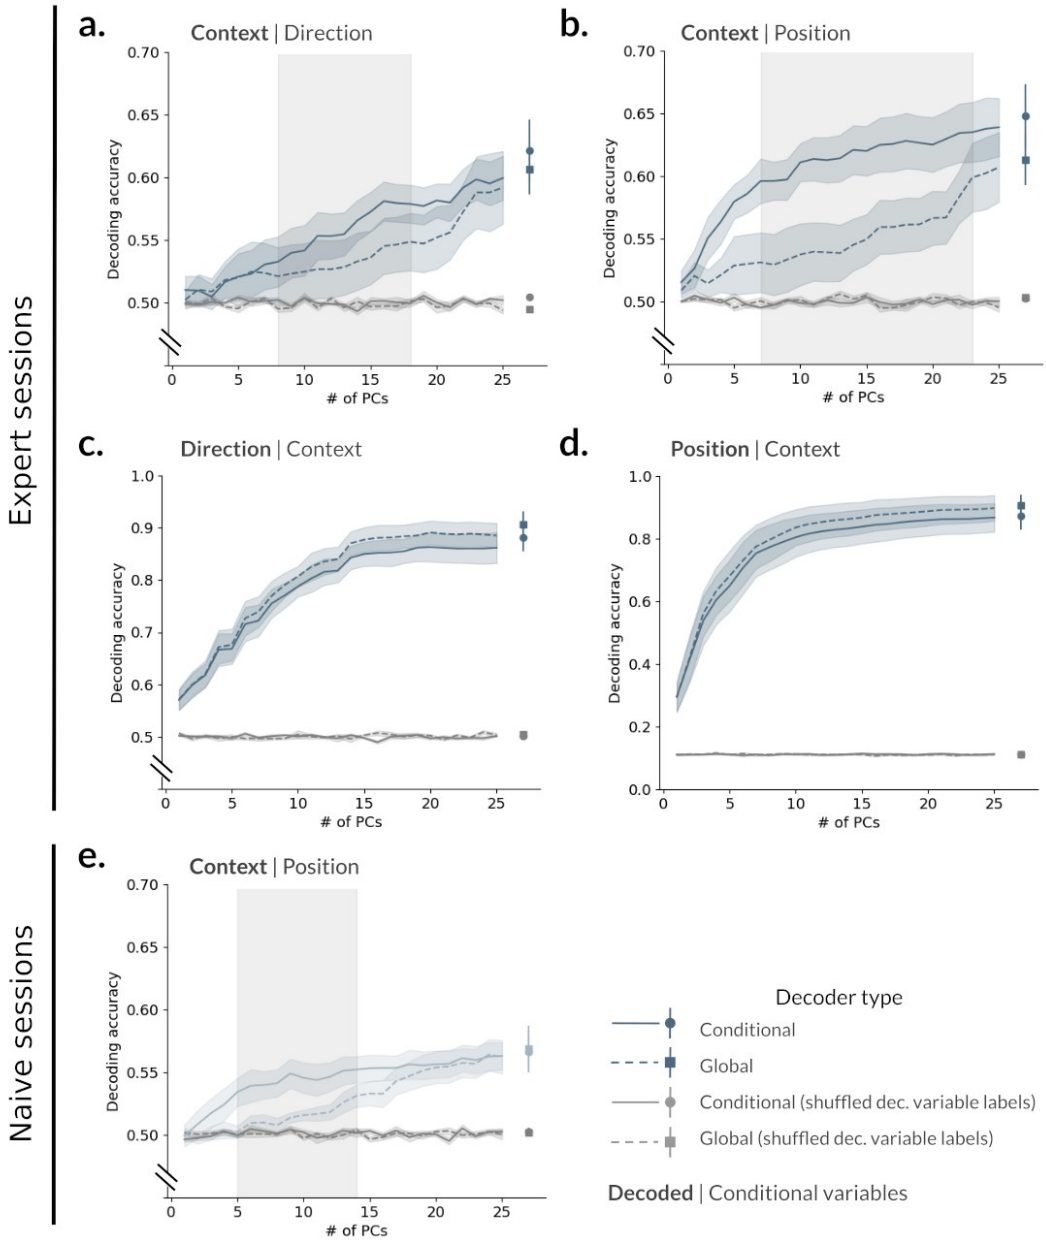

Figure 6.6. Conditional decoder accuracy as a function of principal components. a) Accuracy of decoding context, by the global decoder or conditional on trial direction (accuracy value is the average over both directions). The population vectors were projected in the principal components and increasing number of these projections were fed to the decoder for accuracy calculation (in order of variance explained). Calculated for all sessions in the Expert behaviour cluster. Regions around the lines indicate S.E.M. over bootstrap samples, the shaded range in gray indicates a 95% C.I. for the difference between the global and conditional decoders, paired t-test. Accuracy indicates the proportion of predicted labels that match the real labels in the held-out test data set. Positions were divided in discrete bins for this calculation (see Methods). b) Same as (a) but conditional on maze position. c) Same as (a) but the trial context is the conditional variable, and the decoded variable is trial direction. d) Same as (a) but the trial context is the conditional variable, and the decoded variable is position. e) Same as (b) but calculated for sessions in the Naive behaviour cluster. Note the difference in the y axis between the plots, since the maximal decoding accuracy is highly different between task variables. $p<0.05$ in the shaded regions, paired t-test.

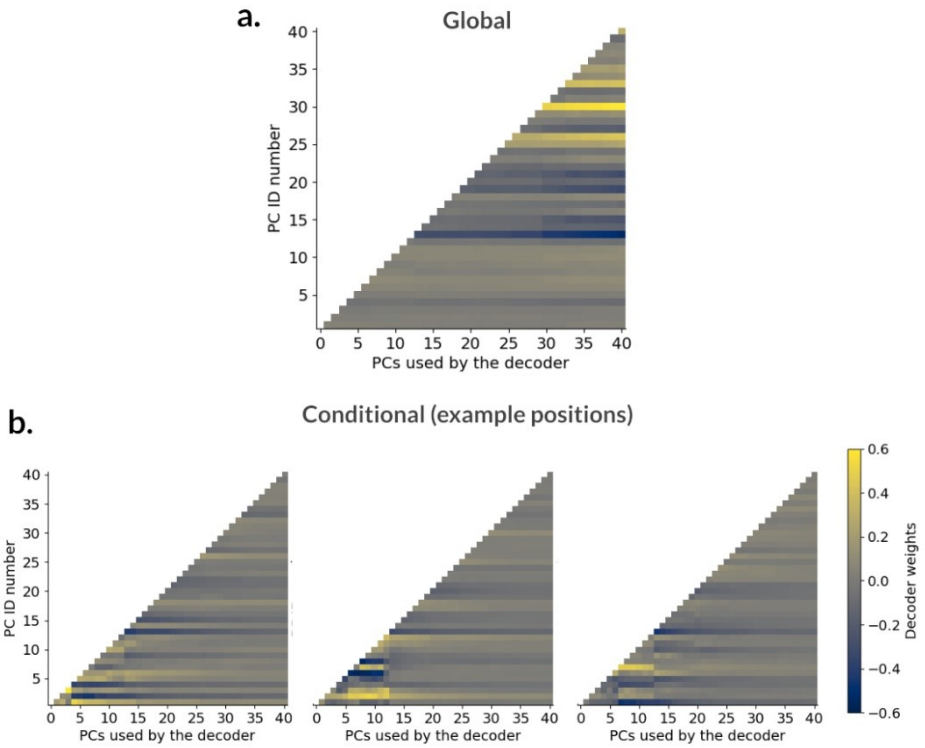

Figure 6.7. Examples of linear decoding weights. a) Example of decoder weights for the global context decoder. Population vectors were projected into a subset of PCs for training the decoder. Each column corresponds to decoding using a different number of PCs (adding cumulatively in order of their variance explained from left to right). The colours indicate the weight each PC received in that case, averaged over 100 bootstrapped repetitions. If the colour in a given row changes as it progresses to the right, it means that the weight associated to that PC changes as more PCs were available to the decoder. PC IDs were ordered from highest to lowest variance explained. b) Same as (a) but for the decoder conditional on position. Weights are specific to each position in the maze, so here the examples from 3 positions are shown.

The hierarchy between variables and learning

To understand the relationship between variables, I then considered the inverse scenario: decoding trial direction or position, conditional on the context label of the population vectors. What I observed was that the accuracy of the two decoders was equivalent (Figures 6.6c-d), an indication that context is not directly informative about the other variables. If we compared trial direction and position, direction decoding was better when conditional on the position at all possible numbers of PCs fed to the decoder, and this did not change with learning (see Figure 6.11). The inverse was not true; position decoding was also equal between decoders regardless of movement direction. This can be understood as a hierarchy between variables: position is on top of the hierarchy, encoded independently from the other variables and driving most of the variance (and therefore the easiest to linearly decode). The code for the direction of movement is then modulated at each position; context sits at the bottom of the hierarchy, dependent on the other two variables. A previous study in a different behavioural paradigm reports a similar hierarchical structure in hippocampal representations (McKenzie et al., 2014).

Variables that are highly decodable using the global decoder are unlikely to benefit from information about less precisely encoded variables (i.e. context) (Figure 6.5a). The results from

place cells in Chapter 5 already hinted at this structure. Position modulates the activity of many cells between almost silent and highly active. Direction and context lead to strong and weak levels of remapping, respectively. What the population analysis captures is how this translates into dimensionality and availability of information to downstream areas, as well as capturing the influence of non-place cells on the code. Given that position decoding accuracy is high and stable over all behaviour clusters, I did not pursue any further position decoding per se but rather the dependency of other variables on it.

Lastly, I compared context decoding in the Expert versus Naive sessions (Figures 6.6b and 3.6e, respectively). The maximum accuracy achieved increased with learning, as well as the difference between decoders at lower numbers of PCs. The difference between decoders could be a sign that context drives more variability as it becomes relevant for behaviour, which would make it better decodable from earlier PCs in the conditional case. However, it is likely that the difference between decoders simply scales with the maximum accuracy. Decoder weights do not provide much insight into this issue, as the phenomena described in Figure 6.7 were the same over all behaviour clusters (data not shown). In the "future directions" section I discuss other methods that could help clarify this hypothesis.

Variance explained and cell numbers

It is important to note that the variance explained by each principal component varies between sessions, as a result of different identities and number of neurons. By definition, the maximum number of principal components is defined by the number of features (i.e. neurons) in the data. Therefore, on days with fewer neurons, there was a steeper relationship between the variance explained and the absolute number of principal components (Figure 6.8a). On the other hand, the variance explained per percentage of all principal components resulted in the opposite scenario; the same percentage of PCs explained more of the variance when calculated in a larger dataset, as it represented more PCs (Figure 6.8b).

To control for this difference between sessions, I measured how much the decoding accuracy per PC changed if I subsampled the population. On days when a large number of

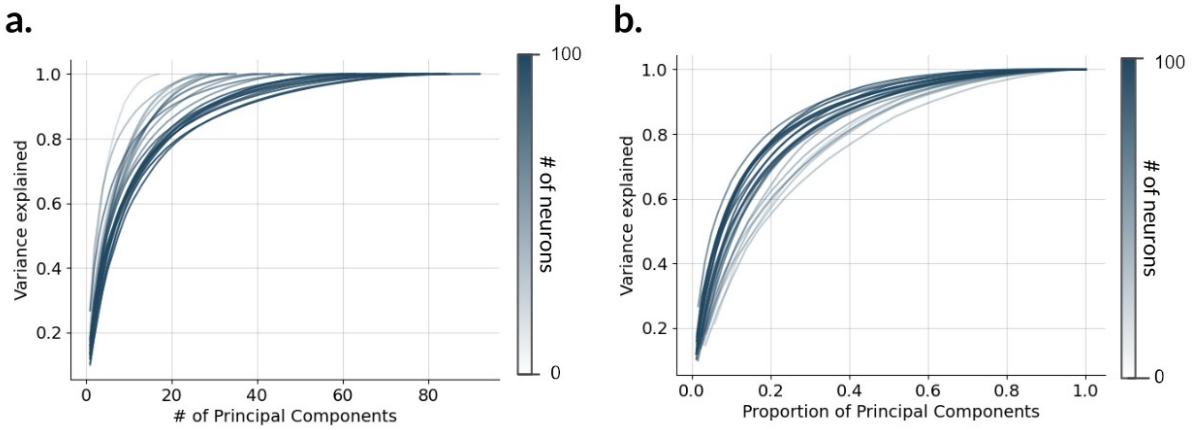

Figure 6.8. Variance explained versus number of neurons recorded. a) Variance explained per principal component for all sessions. Difference colours indicate different animals and opacity represents the number of neuron in the session, with more opaque colour meaning less neurons. b) Same as (a) but per proportion of the principal components.

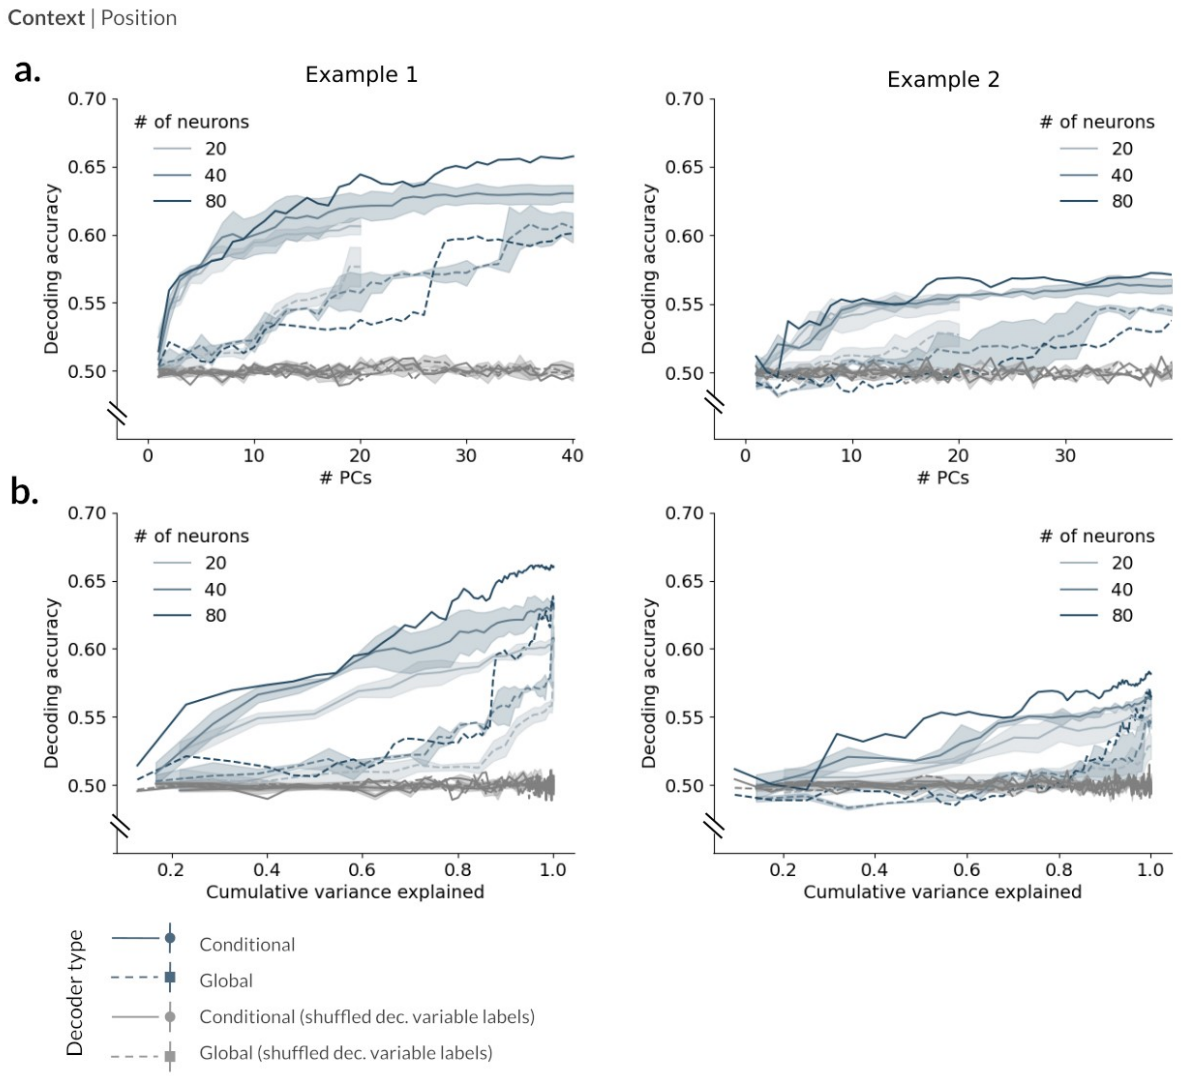

Figure 6.9. The effect of subsampling on linear decoding. a) Decoding accuracy of context conditional on position from two example sessions of >80 pyramidal neurons recorded, when the PCA was calculated from a different sized pool of neurons. Error regions calculated from non-overlapping sets of neurons. b) Same as (a) but per cumulative proportion of variance explained.

pyramidal neurons were recorded (>80 neurons), it was possible to compare the accuracy per PC obtained from the analysis on non-overlapping subsets of neurons. My results show that the decoding accuracy per PC was robust to the sample size, despite the variance explained per PC changing between samples (Figure 6.9). Therefore, sessions could be compared based on the number of principal components used for decoding. As decoder accuracy generally plateaued around 15-20 PCs, I included all sessions with more than 20 pyramidal neurons in PCA calculations (only one session excluded) and show results up to 25 PCs for better visualisation.

Decoding at different positions

The performance of the decoders conditional on position shown until here was calculated as an average over all positions in the maze. I then asked if this performance was equal between positions or specific to locations related to behaviour. In the case of decoding context, I observed a large difference in decoding accuracy between positions; most of the separation of context

happened at context-dependent rewards (Figure 6.10a). This was not a reward-related effect, since at the context-independent reward this was not observed (Figure 6.10b) and at other positions in the maze the separation was also poor (Figures 6.10c-d). This indicates an association of context and place at the hippocampal level that is related to the task; the information of context is encoded only at the position where context influences decision-making. I also made sure to only use running periods when calculating the population vectors, in order to exclude decoding gains due to differences in behaviour (digging or not) at those locations. As a result of this local effect, learning-related gains were also constrained to the context-dependent reward locations.

The separation of trial direction, on the other hand, was equally accurate at all positions (Figures 6.11). Nevertheless, at the edges of the maze, higher accuracy was achievable by both decoders with a much smaller number of principal components (Figure 6.11c). In Chapter 5 I showed that the neural population was more active at an edge position when the animals entered the maze, compared to when they exited. Therefore, this difference in single-cell activity is likely to be a large source of variance at these positions, bringing information about direction to lower PCs. There was no clear difference in direction decoding between behaviour clusters, suggesting that the representation of trial direction was unaffected by learning.

One major possible confound to the difference in decoding between positions is velocity, which is known to strongly modulate hippocampal activity (McNaughton et al., 1983). Although all population vectors have been calculated during running periods, no difference was made between different velocity profiles. Unfortunately, velocity and behaviour are tightly correlated in this case. When animals were moving towards a reward location, they slowed down, compared to running past it in the context it was not rewarded. This difference in speed was especially true in the expert sessions and milder in the naive sessions, in which they were likely to slow down even if eventually they made the correct decision. The same was true regarding the separation of movement direction at the maze edges. Animals ran faster when exiting the maze in comparison to when entering it, regardless of context. Further control analysis would be necessary to detangle these effects, e.g. by comparing the neural activity in two different velocity profiles at a behaviourally irrelevant maze region.

Population code versus behaviour performance

It is not possible to appropriately train linear decoders to classify a trial in terms of behaviour (correct or error trial), due to the low number of error trials available for training the decoder in most sessions. Instead, I asked whether the information about context in the neural activity was different in correct versus error trials. This was the main aspect of the task the animals had to learn: the association between the contextual cues and reward location. For this, I looked at how the conditional decoder I previously trained had performed on the population vectors from the cross-validation test set, depending on the trial the population vector came from. I focused on the neural activity at the reward locations dug in each trial.

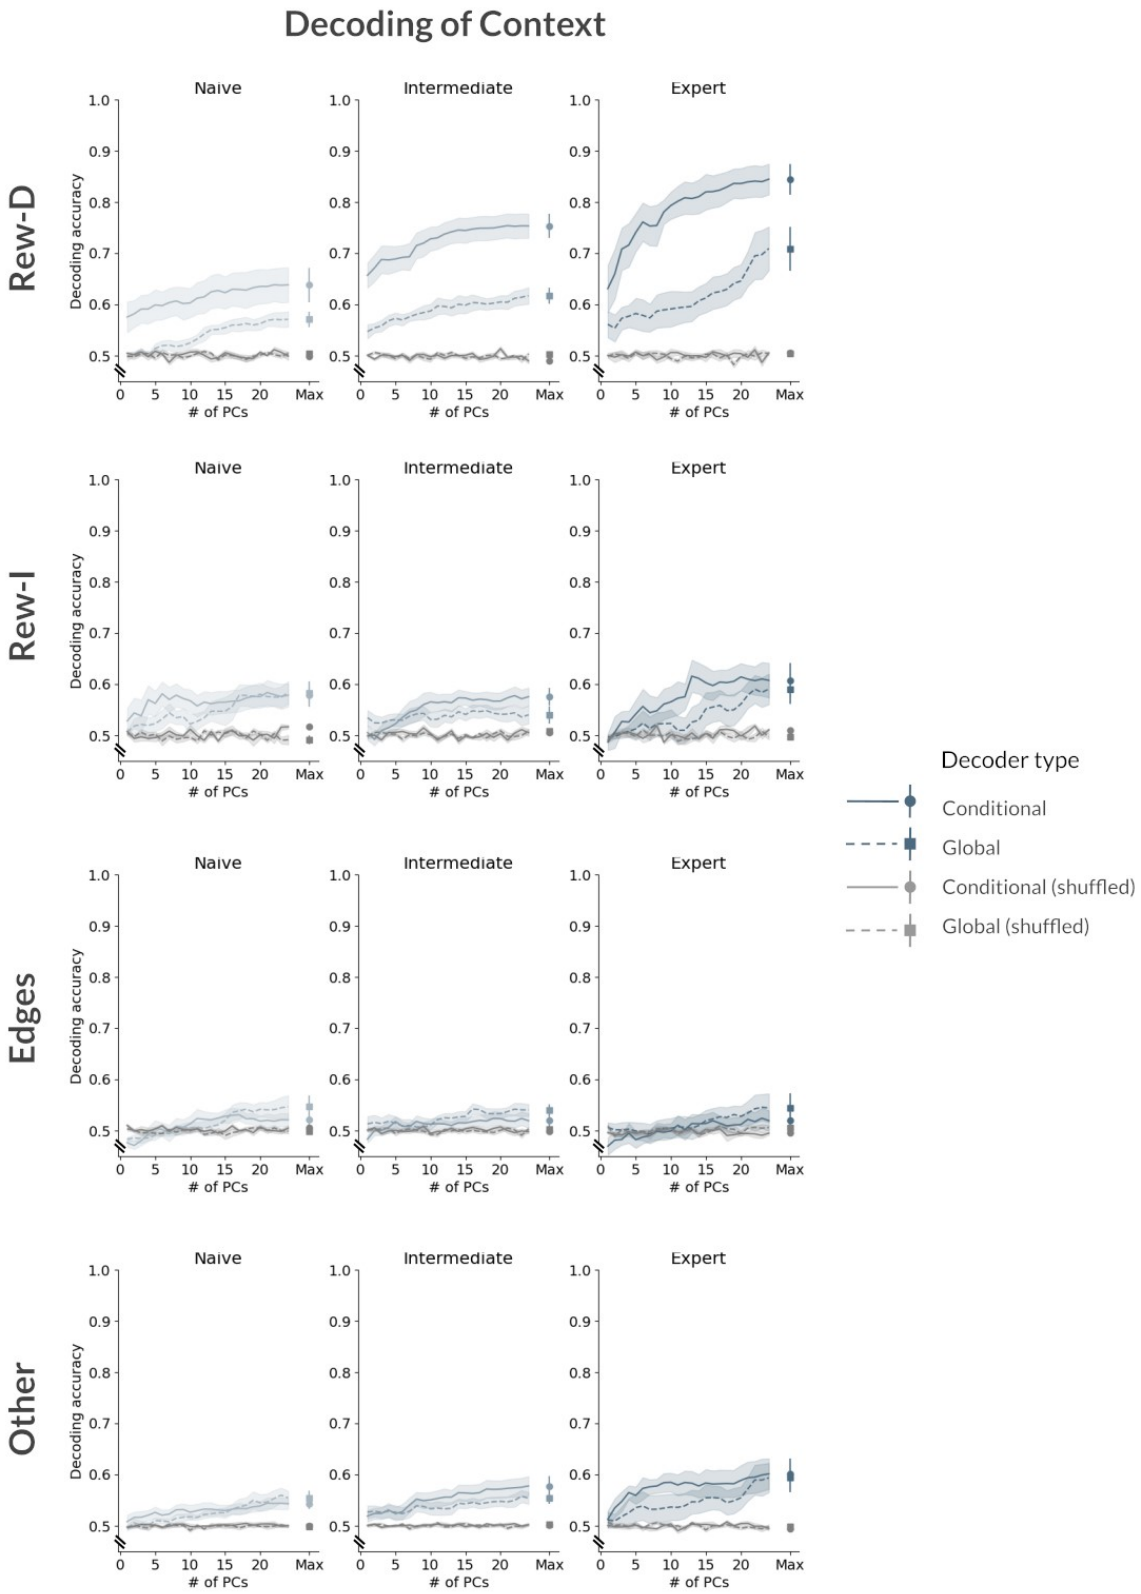

Figure 6.10. Decoding context at different maze positions. Decoding of the variable context using a global decoder or a decoder conditional on position. Since performance is calculated for a held-out test set for each bootstrap sample, performance for population vectors in the test data from different positions can be calculated for both decoders. For each maze position the performance is separated between the different behaviour clusters. Rew-D: Context-dependent rewards, Rew-I: Context-independent rewards.

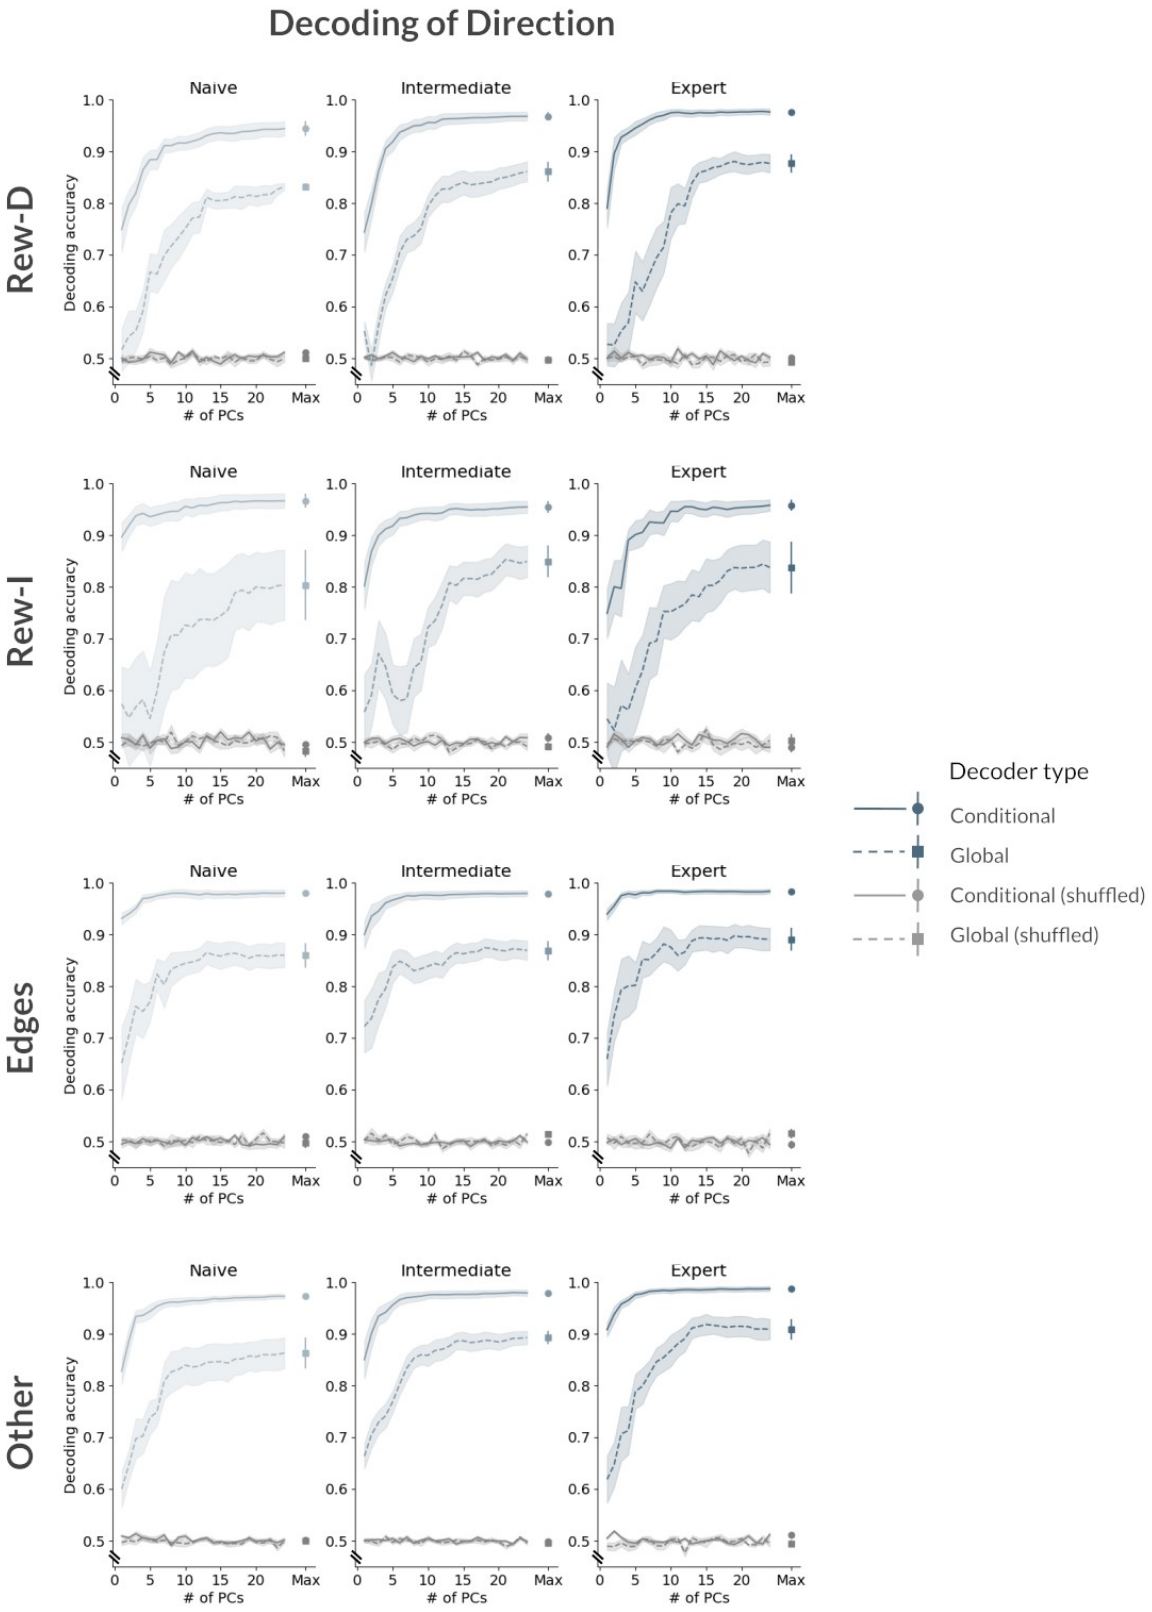

Figure 6.11. Decoding trial direction at different maze positons. Decoding of the variable direction using a global decoder or a decoder conditional on position. Since performance is calculated for a held-out test set for each bootstrap sample, performance for population vectors in the test data from different positions can be calculated for both decoders. For each maze position the performance is separated between the different behaviour clusters. Rew-D: Context-dependent rewards, Rew-I: Context-independent rewards.

At context-dependent rewards, the context was better decoded in correct than in error trials. This difference was small in the naive sessions and became much larger at later learning stages; the accuracy of context decoding in expert sessions was close to 90% during correct trials and at chance level during error trials (Figure 6.12a). This effect was not observed at the context-independent reward (Figure 6.12b), in accordance with the fact that the maximum decoding of context was in general worse at this location. The current analysis cannot exclude the possibility that the main differences in decoding between context-dependent and independent rewards result from the difference in place field coverage between the two (Chapter 5).

The local decoding results showed that the decoding of context is worse at the contextual reward during error trials, so I used a cosine similarity approach to clarify the difference in the representation in these trials. For each trial, I measured the similarity between the population activity vector at the contextual reward dug – even if incorrect – to the average activity at that position over all correct trials. First, I looked at this measure for each correct trial versus the

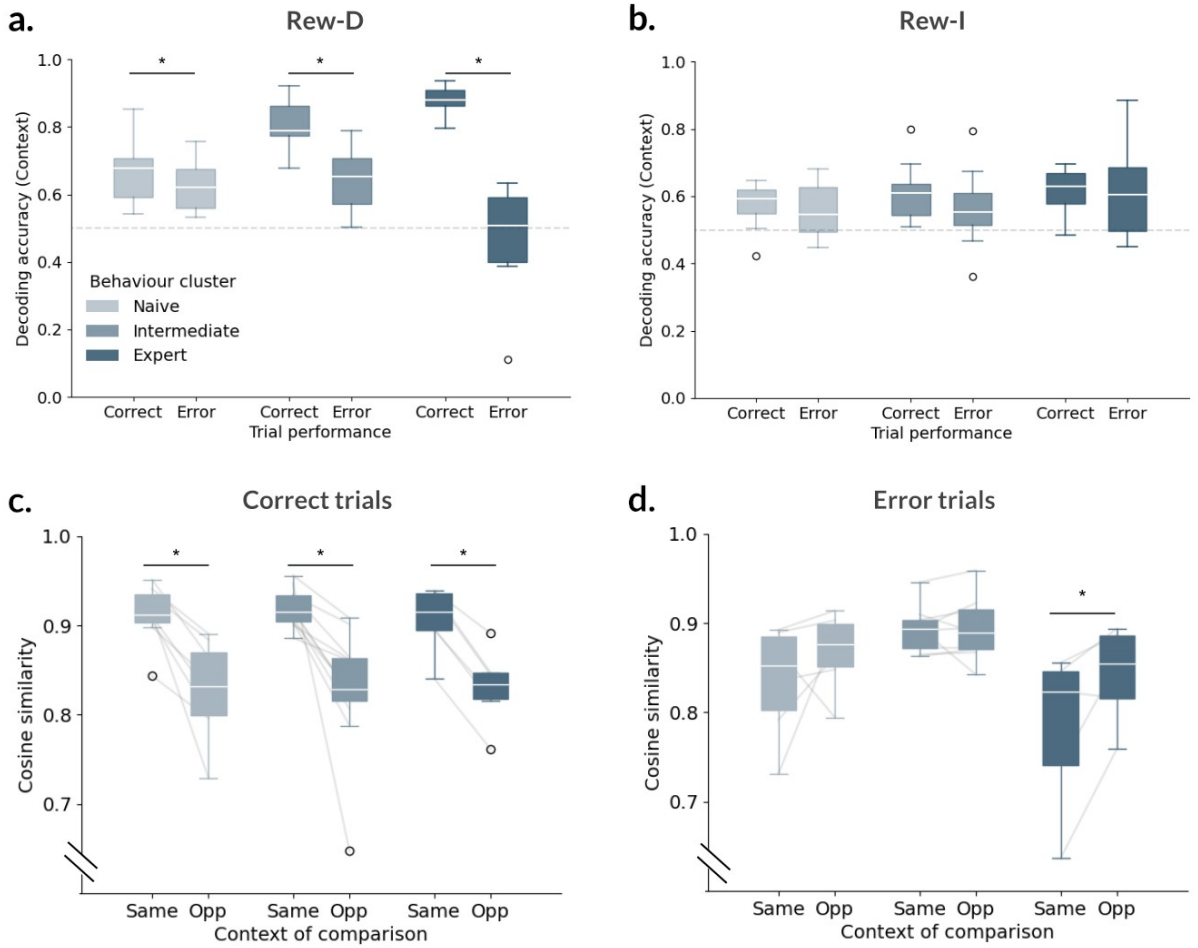

Figure 6.12. Differences between correct and error trials. a-b) Accuracy for conditional decoding of context at reward positions, for held-out population vectors coming from correct or error trials, per behaviour cluster. a) Context-dependent rewards (Correct vs error: Naive p=0.034, Intermediate p=5.19e-5, Expert p=0.003, Paired t-test), b) Context-Independent reward. (All differences between Correct and Error p>0.3, Paired t-test). c-d) Cosine similarity between population vector at the context-reward position dug at a given trial versus the average vector of all correct trials of either the same or opposite context. c) When the position dug was the correct one for the current context – correct trial (Same vs Opposite: Naive p=6.66e-4, Intermediate p=0.001, Expert p=3.47e-5, Paired t-test). Other types of error trials were not considered in this analysis. d) When the position dug was that from the opposite context – error trial (Same vs Opposite: Naive p=0.168, Intermediate p=0.178, Expert p=0.034, Paired t-test).

average as a control measure. Not surprisingly there was high similarity between the neural activity in each trial and the average over all correct trials of the same context (Figure 6.12c). This indicates that the activity was reproducible over trials, when the context and behaviour were the same. The similarity to the average activity at that position in trials of the opposite context was lower; this indicates that the neural activity was different between the two contexts, despite the animal being at the same location.

Contrary to correct trials, during error trials the population activity at the dug position was most similar to that from the opposite context, and increasingly so as the animal learned the task (Figure 6.12d). Here, only errors when animal dug the reward location of the opposite context were considered. Even though population vectors were calculated only using running periods, this result can be due to the fact that behaviour during error trials was overall more similar to that of the opposite context, where the animal also would chose to dig the same location. Together, these results suggest that animals dug the reward of the wrong context not as a result of ambiguous population activity as previously described (Xu et al., 2019), but as an actual encoding of the wrong context.

Reward coding in the population

Despite the difference in context coding between reward types in my data, I asked whether there was any sign of a shared reward code across them. For this, I calculated the similarity between population vectors from pairs of positions, from trials of the same context but different running directions (Figure 6.13a). Since direction is a strong modulator of place cell activity in linear environments, as shown in Chapter 5, there is very little similarity between pairs of positions, even when the same position is considered. Nevertheless, between pairs of rewards, the similarity was higher, even when different reward locations were compared (Figure 6.13b). In the opposite comparison – when the running direction was maintained but the context was different – the correlations between the same positions in the maze were much higher (Figure 6.13c). This was expected given that the effect of context on single-cell firing was much milder (Chapter 5). Nevertheless, when different positions were compared, correlations were higher if both positions were rewarded, regardless of their identity (Figure 6.13d). No significant difference was found between behaviour clusters in both analyses (data not shown).

Reward coding in the hippocampus has been extensively studied and a common reward code has been often reported (Gauthier & Tank, 2018; L. Zhang et al., 2022). In the chapter introduction, I mentioned that a common reward code might simply be the result of a higher propensity of cells to fire at those locations, increasing the chance of the same neurons covering multiple reward locations (J. S. Lee et al., 2020). The fact that the similarity between reward positions is stable over learning is another sign that the reward code emerges quickly and is independent of the associative aspects of the task. One contradictory result to this claim is my observation of different place cell coverage context-dependent and independent rewards, mentioned in Chapter 5. The similarity between reward locations also opens the question if

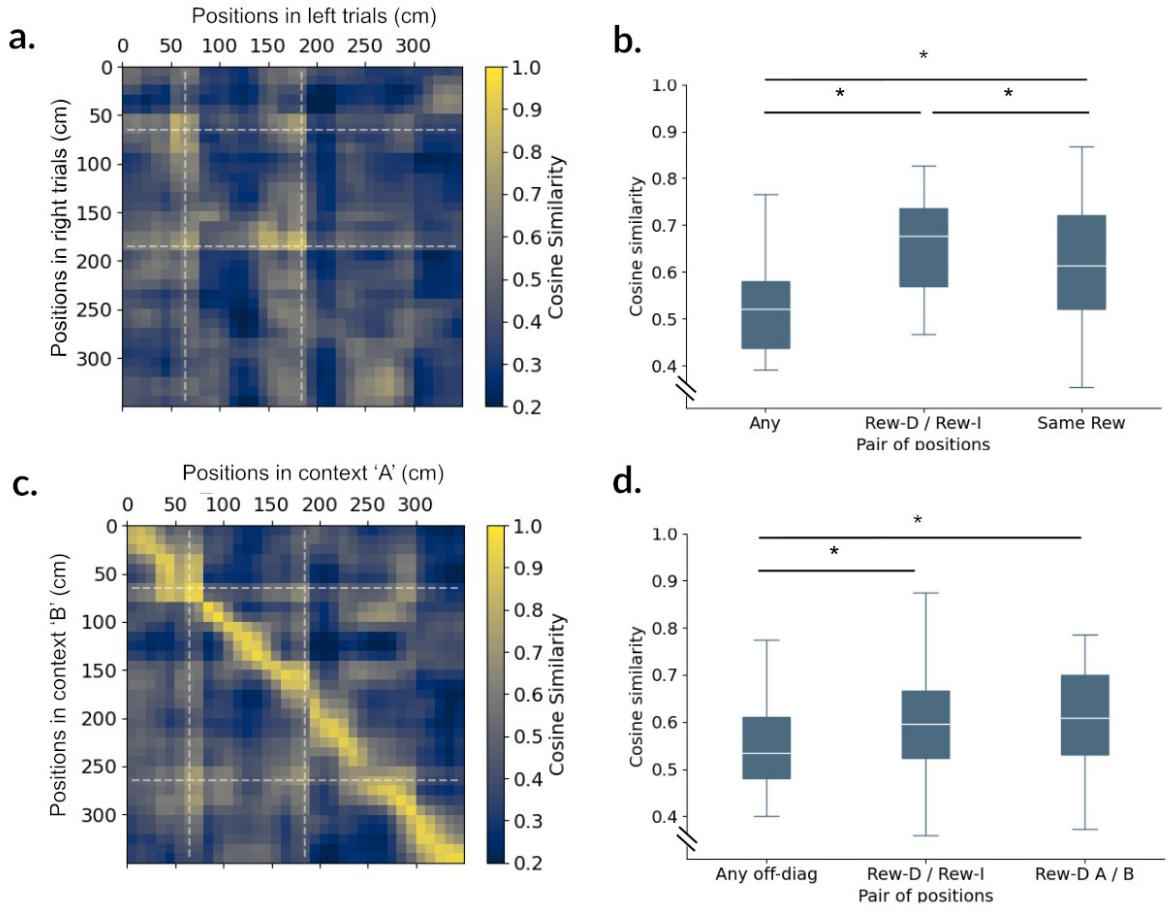

Figure 6.13. Similarity between rewarded positions in the population level. a) Similarity between average population vectors from each position, between the last 3 correct trials of same context but opposite direction. Example from a single session. White lines indicate reward positions. b) Similarity between rewarded positions versus between other pairs of positions when comparing different trial directions, over all sessions Rew-D = Context-dependent reward position, Rew-I = Context-independent reward position. (Repeated Measures ANOVA p=1.68e-20; Paired t-test with Holm-Sidak correction Any to RewD/I p=6.53e-21; Any to Same p=1.12e-9; RewD/I to Same p=3.80e-3). c) Same as (a) but comparing categories of same direction but opposite context. d) Same as (b) but comparing different contexts. Only off-diagonal elements of the matrix were considered, since similarity between the same positions was much higher (Repeated Measures ANOVA p=3.00e-6; Paired t-test with Holm-Sidak correction Any to RewD/I p=6.92 e-5; Any to RewDA/B p=3.47e-7; RewD/I to RewDA/B p=0.11).

linear decoders for certain variables are more aligned between reward positions than at other positions in the maze, which remains to be studied.

Discussion

In this chapter, I have shown that behaviour variables are hierarchically represented in the hippocampus population code and that this representation adapts to meet behaviour demands. Such a hierarchy has been observed in previous studies, both in the spatial domain as well as for various task dimensions, but previous studies did not investigate changes in the representations during learning (McKenzie et al., 2014; H. Zhang et al., 2023). As in the single-cell case, the representation of time and spatial position in all behavioural stages was mostly related to the activity of interneurons and pyramidal cell populations, respectively. Other variables were also encoded in the pyramidal population but were anchored to the position code. As the animal learned, the code adapted so that variables became more decodable at positions

where they were relevant for behaviour; e.g. the variable context was better decoded where it determines the presence of reward. Encoding of context in those locations was also correlated to trial performance, indicating a link between hippocampal representations and behaviour over many days. Lastly, I showed that there was a common code for reward that disregards position information and that did not change with learning.

Linear regressors and classifiers have been extensively used to decode information from the hippocampus and to detangle the contribution of individual neurons to the population code (Bernardi et al., 2020; Levy et al., 2023; Samborska et al., 2022). One recent study monitoring CA1 activity in rodents shows that context can be decoded from patterns of co-activity of hippocampal pyramidal neurons with improved accuracy and stability over long periods of time (Levy et al., 2023). The authors describe the decoding of context – two separate environments – as a result of a small set of anti-coactive cells, including place and non-place cells equally. In my data, I could not identify a clear subgroup of cells that contribute to decoding, as decoding and PCA weights were spread over a large portion of the population. This might have to do with the difference in the concept of context in our study, which leads to much milder changes in cell firing than two distinct environments.

The study above shows that context decoding improves with experience, as I also observed in my data. However, it does not include any structured learning paradigm, which could affect the neural representation of task variables. Here I investigated the process by which behavioural demands shape neural representations, monitoring a continuous learning period of a goal-oriented behaviour. I show that, as the animal learned the association between context and reward location, these two variables became increasingly bound in the population representation. Moreover, this process of adaptation to behavioural demands happened at a timescale that correlates with learning.

I investigated this relationship between variables through conditional decoders. Certain variables were better decoded locally (conditional on position) than globally; this implies that the features (i.e. neurons) representing these variables are not the same across the entire environment. Features used by a global decoder to separate a variable across all locations have to be composed of neurons that are globally active to a certain extent (non-place cells) and also globally responsive to context. In the hippocampus, the population activity at a given position is largely influenced by place cells. If they are also sensitive to context, i.e. through rate remapping, they influence the features representing context at that location. At a different position, the same subset of place cells is silent, so context encoding must also shift to other neurons. In this scenario, mixed selectivity is what makes the code local. Studying the features of this mixed selectivity and how it changes with learning can provide future insights into the structure of the hippocampal population code. Linear decodability of variables on the population level can be achieved with single neurons with linear or non-linear mixed selectivity. The first enables the generalisation of information and robustness to noise whilst the second grants higher flexibility in the number of decodable variables from a single stimulus (Kaufman et al., 2022). The function of a brain area might determine the balance between the two. Population level analysis would

suggest that brain areas such as the prefrontal cortex are more involved in generalisation and the hippocampus in the specificity of each experience (Kaefer et al., 2020; Samborska et al., 2022; Tang et al., 2023), but a link to mixed selectivity has not been made.

In primates, complex task structures that impose novel relationships between variables are reflected in the hippocampal code (Bernardi et al., 2020). My results suggest that this structure also exists in rodents, is built at the timescale of learning and reflects a reorganisation of the contributions of different variables to the code. Here I also argue that task-relevant information might move to high-variance principal components with learning: a sign of migration towards a lower-dimensional representation of task-relevant features. Studies in the motor cortex show a similar type of structure, pointing out that manipulating the dimensionality of representations might be a robust widespread mechanism to interface between neural activity and behaviour (Gallego et al., 2017). Nonetheless, the dynamics and topological changes in hippocampal activity during learning remain to be elucidated. In the "future directions" section of this thesis, I discuss how to target this topic in my data with geometry and dimensionality-related analytical approaches.

7. Future Steps

In this thesis I investigated how the activity of single cells in the hippocampus translates into an effective population code of the external world and how goal-directed learning modulates this code to provide better information of task-relevant variables to downstream brain areas. For this purpose, I used linear methods such as Principal Component Analysis and Support Vector Machines. Although these methods allowed important insights into the data, they limit which aspects of the data can be described. Here I discuss further analysis that would allow me to confirm and expand the results I have presented.

Variability in neural activity versus task variables

A more direct way than linear decoders to measure the relationship between principal components and task variables is by the use of Canonical correlations (CC). CC gives a measure of the variability in each PC that can be accounted for by each task variable; it has been previously used to demonstrate the correlation between the first PCs and position/head direction in a network model of the hippocampus (Recanatesi et al., 2021). In my data, this could be used to confirm the contribution of position and context to the variability of different PCs as the animal learns. It has the advantage of being a continuous measure and can be calculated per PC independently (not cumulative PC as I used for the linear decoder), making any pair of PCs directly comparable.

Moreover, both linear decoders and CC can be used to investigate how variables contribute to the neural activity at different stages of the task. Previous studies have pointed out that the encoding of variables is dynamic, tied to the behaviour in a timely fashion (Bernardi et al., 2020; Samborska et al., 2022). In this thesis I have focused on average activity vectors per position and trial, which hides dynamics related to action and decision-making. Looking at the neural activity around the time of action – e.g. when the animal enters the maze, just before digging, etc. – and at different stages of learning can add insight on contributions of the hippocampal representations to specific behaviours.

The dimensionality of neural representations

Weighting the contributions of task variables to different PCs is only one of many ways to assess the relationship between neural representations and dimensionality. It has been proposed that brain areas balance between shattering dimensionality (SD) and the generalisation capability of the code (Bernardi et al., 2020). These approaches use linear classifiers to determine: (1) how many arbitrary binary divisions of the data are linearly decodable (SD); and (2) how well a decoder trained in a fixed condition performs when the condition changes (generalisation). The authors suggest that, in well trained animals, the hippocampus has high SD and high generalisation score. They also show that during error trials

generalisation is reduced. In my data, it would be possible to investigate the emergence of these features from naive to expert animals, measuring the effect of learning and behaviour on the structure of the neural code.

Despite the term dimensionality, SD can be interpreted very differently than the dimensionality of a manifold, the activity subspace to which population vectors are constrained. SD is focused on the divisibility of the data, whilst manifold analyses focus on its topology. Here I showed that most of the variance explained in my dataset is found in the first ~10-15 principal components, as long as the number of neurons recorded exceeded this value. The neural activity, however, is likely to lie in a non-linear manifold of lower dimension than that estimated by PCA (Altan et al., 2021; Nieh et al., 2021), as it is also upper-bounded by task complexity (Gao et al., 2017).

In the task I proposed, learning was correlated with a reduction in dimensionality in PC space. However, it is not clear if this results from a reduction in dimensionality of the neural manifold or a manifold realignment to the PCs. Changes in topology of hippocampal representations over experience would suggest the first (H. Zhang et al., 2023), but it remains to be tested if this is the case during goal-oriented learning. It is important to note that intrinsic dimensionality measures are not very robust to noise (Altan et al., 2021) and therefore the absolute value of dimensionality extracted by those methods is likely to yield false interpretations. Indeed, it has been shown that noise increases the dimensionality of the hippocampal spatial code (Hazon et al., 2022). Nevertheless, assuming the level of noise is the same across recordings, the relative changes in intrinsic dimensionality during learning can still be used to evaluate the formation of hippocampal representations.

Predictive coding

Hippocampal population activity can be investigated in the framework of predictive coding. Studies of the hippocampus and neural networks argue that navigation can be efficiently solved by a neural population representing the future instead of the current position of the animal (Recanatesi et al., 2021; Stachenfeld et al., 2017). If this is true, it would be interesting to understand the timescale of this process. It could be that the shift towards a future representation happens in a matter of one of a few trials (Priestley et al., 2022) or that it continues to evolve as long as the animal improves at the task. In Chapter 5 I showed a trial-by-trial shift of place fields towards the animal's incoming direction, which on the single cell level did not seem to differ between learning stages. At the population level it is still possible that the representation correlates to the animal's future position with varying lags as the animal learns. In this case, the use of population vectors calculated in time is more suitable than using trial-average vectors, since the future can differ at each passage over a given location. Predictive coding analysis could reveal if better animal performance can be assigned to more distant, precise predictions of task variables.

Representational drift

When referring to long-term dynamics of the hippocampal population, representational drift becomes an important factor (Ziv et al., 2013). Drift seems at first in contradiction with the observation of sleep influence in learning. Neural connections are strengthened for cells co-active during awakeness, which would predict a stable activity of these cells in the following day (Dupret et al., 2010; Gridchyn et al., 2020). While this is true for novel learning, many argue that familiarity reduces the effect of sleep on hippocampal plasticity, so that learning becomes largely relayed to the cortex (Frankland & Bontempi, 2005). Despite this shift, it has been shown that the hippocampus is still important for certain types of behaviour at long timescales (Atucha et al., 2021). Not much is known about the role of hippocampal representations on behaviour over these periods, nor about their dynamics.

One previous study indicates that the encoding of context resists representational drift at the population level, despite variations in the activity of single neurons (Keinath et al., 2022). The study points out that the encoding of context happens in dimensions orthogonal to drift, in a scenario of purely exploratory behaviour. In this thesis I have shown that task demands impose relationships between the encoding of different task variables; it is therefore less clear if context representation would be resistant to drift, considering its representation is connected to space. One hypothesis is that the information about context represented in later PCs is detached from the spatial representation, and therefore could be orthogonal to drift. My behavioural paradigm is a good candidate to study these phenomena, but new recordings aimed at monitoring the activity of the same neurons over time are necessary, e.g. by using calcium imaging or silicon probes.

The neurobiology of digging

One topic that has not been very explored in the literature but is relevant in our dataset is digging behaviour. Like running, digging is an innate motor behaviour that requires a certain level of repetitive action. Despite its use in many experimental paradigms (Sauvage et al., 2008; Tse et al., 2007), not much is known of how this behaviour engages the brain, especially the hippocampus. For example, since the animal is awake and mobile, theta oscillations are expected to dominate hippocampal activity and phenomena such as theta phase precession might also be in place. However, digging seems to require higher levels of motivation than running; this might require neuromodulation from the midbrain or top-down control from cortical areas. Understanding the differences and similarities between running and digging at a neural level could help understand what are the main influences of motor action to hippocampal encoding, in a situation where position is fixed, for example. Our dataset allows for the investigation of some of these questions, but also future studies employing precise tracking of movement and a more constrained digging paradigm are necessary for targeting all aspects of this behaviour.

List of Figures

List of Abbreviations

BIC	Bayesian Information Criterion
DG	Dentate Gyrus
CA1	Cornu Ammonis 1 (subregion of the hippocampus)
CA3	Cornu Ammonis 3 (subregion of the hippocampus)
MDS	Multidimensional Scaling
PC	Principal Component
PCA	Principal Component Analysis
PF	Place Field
PV	Population Vector
Rew-D	Context-dependent reward (any)
Rew-DA	Context-dependent reward of context A
Rew-DB	Context-dependent reward of context B
Rew-DC	Context-dependent reward of the current trial context
Rew-DO	Context-dependent reward of the opposite context
Rew-I	Context-independent reward
SEM	Standard Error of the Mean
t-SNE	t-distributed Stochastic Neighbours Embedding
UMAP	Uniform Manifold Approximation and Projection

References

Akiti, K., Tsutsui-Kimura, I., Xie, Y., Mathis, A., Markowitz, J. E., Anyoha, R., Datta, S. R., Mathis, M. W., Uchida, N., & Watabe-Uchida, M. (2022). Striatal dopamine explains novelty-induced behavioral dynamics and individual variability in threat prediction. *Neuron*, *110*(22), 3789-3804.e9. https://doi.org/10.1016/j.neuron.2022.08.022

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature Neuroscience*, *18*(8), 1143–1151. https://doi.org/10.1038/nn.4058

Allen, K., Rawlins, J. N. P., Bannerman, D. M., & Csicsvari, J. (2012). Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping. *The Journal of Neuroscience*, *32*(42), 14752–14766. https://doi.org/10.1523/JNEUROSCI.6175-11.2012

Altan, E., Solla, S. A., Miller, L. E., & Perreault, E. J. (2021). Estimating the dimensionality of the manifold underlying multi-electrode neural recordings. *PLOS Computational Biology*, *17*(11), e1008591. https://doi.org/10.1371/journal.pcbi.1008591

Andersen, R. A., & Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, *25*, 189–220. https://doi.org/10.1146/annurev.neuro.25.112701.142922

Atucha, E., Ku, S.-P., Lippert, M. T., & Sauvage, M. M. (2021). *Remembering the gist of an event over a lifetime depends on the hippocampus* (p. 2021.04.14.439803). bioRxiv. https://doi.org/10.1101/2021.04.14.439803

Babb, S. J., & Crystal, J. D. (2006). Episodic-like Memory in the Rat. *Current Biology*, *16*(13), 1317–1321. https://doi.org/10.1016/j.cub.2006.05.025

Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral Prefrontal Cortex and Individual Differences in Uncertainty-Driven Exploration. *Neuron*, *73*(3), 595–607. https://doi.org/10.1016/j.neuron.2011.12.025

Battaglia, F. P., Sutherland, G. R., & McNaughton, B. L. (2004). Local Sensory Cues and Place Cell Directionality: Additional Evidence of Prospective Coding in the Hippocampus. *Journal of Neuroscience*, *24*(19), 4541–4550. https://doi.org/10.1523/JNEUROSCI.4896-03.2004

Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, *58*(6), 7738–7753. https://doi.org/10.1103/PhysRevE.58.7738

Beer, Z., Vavra, P., Atucha, E., Rentzing, K., Heinze, H.-J., & Sauvage, M. M. (2018). The memory for time and space differentially engages the proximal and distal parts of the hippocampal subfields CA1 and CA3. *PLOS Biology*, *16*(8), e2006100. https://doi.org/10.1371/journal.pbio.2006100

Belchior, H., Lopes-dos-Santos, V., Tort, A. B. L., & Ribeiro, S. (2014). Increase in hippocampal theta oscillations during spatial decision making. *Hippocampus*, *24*(6), 693–702. https://doi.org/10.1002/hipo.22260

Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P. L., Gioanni, Y., Battaglia, F. P., & Wiener, S. I. (2010). Coherent Theta Oscillations and Reorganization of Spike Timing in the Hippocampal-Prefrontal Network upon Learning. *Neuron*, *66*(6), 921–936. https://doi.org/10.1016/j.neuron.2010.05.013

Benchenane, K., Tiesinga, P. H., & Battaglia, F. P. (2011). Oscillations in the prefrontal cortex: A gateway to memory and attention. *Current Opinion in Neurobiology*, *21*(3), 475–485. https://doi.org/10.1016/j.conb.2011.01.004

Bernardi, S., Benna, M. K., Rigotti, M., Munuera, J., Fusi, S., & Salzman, C. D. (2020). The Geometry of Abstraction in the Hippocampus and Prefrontal Cortex. *Cell*, *183*(4), 954-967.e21. https://doi.org/10.1016/j.cell.2020.09.031

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning* (1st ed.). Springer New York, NY. https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf

Bliss, T. V. P., & Collingridge, G. L. (1993). A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, *361*(6407), Article 6407. https://doi.org/10.1038/361031a0

Bors, D. A., & MacLeod, C. M. (1996). Individual Differences in Memory. In *Memory* (pp. 411–441). Elsevier. https://doi.org/10.1016/B978-012102570-0/50014-8

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

Chen, G., King, J. A., Burgess, N., & O'Keefe, J. (2013). How vision and movement combine in the hippocampal place code. *Proceedings of the National Academy of Sciences*, *110*(1), 378–383. https://doi.org/10.1073/pnas.1215834110

Chung, J. E., Magland, J. F., Barnett, A. H., Tolosa, V. M., Tooker, A. C., Lee, K. Y., Shah, K. G., Felix, S. H., Frank, L. M., & Greengard, L. F. (2017). A Fully Automated Approach to Spike Sorting. *Neuron*, *95*(6), 1381-1394.e6. https://doi.org/10.1016/j.neuron.2017.08.030

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*(6699), Article 6699. https://doi.org/10.1038/26216

Cohen, J. D., Bolstad, M., & Lee, A. K. (2017). Experience-dependent shaping of hippocampal CA1 intracellular activity in novel and familiar environments. *eLife*, *6*, e23040. https://doi.org/10.7554/eLife.23040

Danielson, N. B., Zaremba, J. D., Kaifosh, P., Bowler, J., Ladow, M., & Losonczy, A. (2016). Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1. *Neuron*, *91*(3), 652–665. https://doi.org/10.1016/j.neuron.2016.06.020

Diehl, G. W., & Redish, A. D. (2023). Differential processing of decision information in subregions of rodent medial prefrontal cortex. *eLife*, *12*, e82833. https://doi.org/10.7554/eLife.82833

Dudek, S. M., & Bear, M. F. (1992). Homosynaptic long-term depression in area CA1 of hippocampus and effects of N-methyl-D-aspartate receptor blockade. *Proceedings of the National Academy of Sciences*, *89*(10), 4363–4367. https://doi.org/10.1073/pnas.89.10.4363

Dupret, D., O'Neill, J., & Csicsvari, J. (2013). Dynamic Reconfiguration of Hippocampal Interneuron Circuits during Spatial Learning. *Neuron*, *78*(1), 166–180. https://doi.org/10.1016/j.neuron.2013.01.033

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), Article 8. https://doi.org/10.1038/nn.2599

Ego-Stengel, V., & Wilson, M. A. (2007). Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus*, *17*(2), 161–174. https://doi.org/10.1002/hipo.20253

Eichenbaum, H., & Lipton, P. A. (2008). Towards a functional organization of the medial temporal lobe memory system: Role of the parahippocampal and medial entorhinal cortical areas. *Hippocampus*, *18*(12), 1314–1324. https://doi.org/10.1002/hipo.20500

Euston, D. R., Gruber, A. J., & McNaughton, B. L. (2012). The Role of Medial Prefrontal Cortex in Memory and Decision Making. *Neuron*, *76*(6), 1057–1070. https://doi.org/10.1016/j.neuron.2012.12.002

Fenton, A. A., Kao, H.-Y., Neymotin, S. A., Olypher, A., Vayntrub, Y., Lytton, W. W., & Ludvig, N. (2008). Unmasking the CA1 Ensemble Place Code by Exposures to Small and Large Environments: More Place Cells and Multiple, Irregularly Arranged, and Expanded Place Fields in the Larger Space. *Journal of Neuroscience*, *28*(44), 11250–11262. https://doi.org/10.1523/JNEUROSCI.2862-08.2008

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete Coding of Reward Probability and Uncertainty by Dopamine Neurons. *Science*, *299*(5614), 1898–1902. https://doi.org/10.1126/science.1077349

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*, *6*(2), 119–130. https://doi.org/10.1038/nrn1607

Frankland, P. W., Cestari, V., Filipkowski, R. K., McDonald, R. J., & Silva, A. J. (1998). The dorsal hippocampus is essential for context discrimination but not for contextual conditioning. *Behavioral Neuroscience*, *112*(4), 863–874. https://doi.org/10.1037//0735-7044.112.4.863

Fuhs, M. C., & Touretzky, D. S. (2007). Context Learning in the Rodent Hippocampus. *Neural Computation*, *19*(12), 3173–3215. https://doi.org/10.1162/neco.2007.19.12.3173

Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M.-B. (2004). Spatial Representation in the Entorhinal Cortex. *Science*, *305*(5688), 1258–1264. https://doi.org/10.1126/science.1099901

Gallego, J. A., Perich, M. G., Miller, L. E., & Solla, S. A. (2017). Neural Manifolds for the Control of Movement. *Neuron*, *94*(5), 978–984. https://doi.org/10.1016/j.neuron.2017.05.025

Gao, P., Trautmann, E., Yu, B., Santhanam, G., Ryu, S., Shenoy, K., & Ganguli, S. (2017). *A theory of multineuronal dimensionality, dynamics and measurement* [Preprint]. Neuroscience. https://doi.org/10.1101/214262

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

Gehring, T. V., Luksys, G., Sandi, C., & Vasilaki, E. (2015). Detailed classification of swimming paths in the Morris Water Maze: Multiple strategies within one trial. *Scientific Reports*, *5*(1), 14562. https://doi.org/10.1038/srep14562

Geiller, T., Fattahi, M., Choi, J.-S., & Royer, S. (2017). Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nature Communications*, *8*(1), Article 1. https://doi.org/10.1038/ncomms14531

Gener, T., Perez-Mendez, L., & Sanchez-Vives, M. V. (2013). Tactile modulation of hippocampal place fields. *Hippocampus*, *23*(12), 1453–1462. https://doi.org/10.1002/hipo.22198

Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, *117*(1), 197–209. https://doi.org/10.1037/a0017808

Geva, N., Deitch, D., Rubin, A., & Ziv, Y. (2023). Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron*, *111*(15), 2357-2366.e5. https://doi.org/10.1016/j.neuron.2023.05.005

Goodridge, J. P., & Taube, J. S. (1997). Interaction between the Postsubiculum and Anterior Thalamus in the Generation of Head Direction Cell Activity. *Journal of Neuroscience*, *17*(23), 9315–9330. https://doi.org/10.1523/JNEUROSCI.17-23-09315.1997

Graziano, A., Petrosini, L., & Bartoletti, A. (2003). Automatic recognition of explorative strategies in the Morris water maze. *Journal of Neuroscience Methods*, *130*(1), 33–44. https://doi.org/10.1016/S0165-0270(03)00187-0

Gridchyn, I., Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2020). Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. *Neuron*, *106*(2), 291-300.e6. https://doi.org/10.1016/j.neuron.2020.01.021

Grosmark, A. D., Mizuseki, K., Pastalkova, E., Diba, K., & Buzsáki, G. (2012). REM Sleep Reorganizes Hippocampal Excitability. *Neuron*, *75*(6), 1001–1007. https://doi.org/10.1016/j.neuron.2012.08.015

Gulli, R. A., Duong, L. R., Corrigan, B. W., Doucet, G., Williams, S., Fusi, S., & Martinez-Trujillo, J. C. (2020). Context-dependent representations of objects and space in the primate hippocampus during virtual navigation. *Nature Neuroscience*, *23*(1), 103–112. https://doi.org/10.1038/s41593-019-0548-3

Hasselmo, M. E. & Howard Eichenbaum. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Networks*, *18*(9), 1172–1190. https://doi.org/10.1016/j.neunet.2005.08.007

Hazon, O., Minces, V. H., Tomàs, D. P., Ganguli, S., Schnitzer, M. J., & Jercog, P. E. (2022). Noise correlations in neural ensemble activity limit the accuracy of hippocampal spatial representations. *Nature Communications*, *13*(1), 4276. https://doi.org/10.1038/s41467-022-31254-y

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. *Journal of Neuroscience*, *21*(5), 1635–1644. https://doi.org/10.1523/JNEUROSCI.21-05-01635.2001

Izquierdo, I., Furini, C. R. G., & Myskiw, J. C. (2016). Fear Memory. *Physiological Reviews*, *96*(2), 695–750. https://doi.org/10.1152/physrev.00018.2015

Johnson, A., & Redish, A. D. (2007). Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point. *Journal of Neuroscience*, *27*(45), 12176–12189. https://doi.org/10.1523/JNEUROSCI.3761-07.2007

Kaefer, K., Nardin, M., Blahna, K., & Csicsvari, J. (2020). Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching. *Neuron*, *106*(1), 154-165.e6. https://doi.org/10.1016/j.neuron.2020.01.015

Kanai, R., & Rees, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience*, *12*(4), 231–242. https://doi.org/10.1038/nrn3000

Karpas, E. D., Maoz, O., Kiani, R., & Schneidman, E. (2019). *Strongly correlated spatiotemporal encoding and simple decoding in the prefrontal cortex* (p. 693192). bioRxiv. https://doi.org/10.1101/693192

Kaufman, M. T., Benna, M. K., Rigotti, M., Stefanini, F., Fusi, S., & Churchland, A. K. (2022). The implications of categorical and category-free mixed selectivity on representational geometries. *Current Opinion in Neurobiology*, *77*, 102644. https://doi.org/10.1016/j.conb.2022.102644

Keinath, A. T., Mosser, C.-A., & Brandon, M. P. (2022). The representation of context in mouse hippocampus is preserved despite neural drift. *Nature Communications*, *13*(1), 2415. https://doi.org/10.1038/s41467-022-30198-7

Kennedy, P. J., & Shapiro, M. L. (2009). Motivational states activate distinct hippocampal representations to guide goal-directed behaviors. *Proceedings of the National Academy of Sciences*, *106*(26), 10805–10810. https://doi.org/10.1073/pnas.0903259106

Kesner, R. P., Farnsworth, G., & Kametani, H. (1991). Role of Parietal Cortex and Hippocampus in Representing Spatial Information. *Cerebral Cortex*, *1*(5), 367–373. https://doi.org/10.1093/cercor/1.5.367

Khona, M., & Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, *23*(12), Article 12. https://doi.org/10.1038/s41583-022-00642-0

Kim, J., & Lee, I. (2011). Hippocampus is necessary for spatial discrimination using distal cue-configuration. *Hippocampus*, *21*(6), 609–621. https://doi.org/10.1002/hipo.20784

Kirchhoff, B. A., & Buckner, R. L. (2006). Functional-Anatomic Correlates of Individual Differences in Memory. *Neuron*, *51*(2), 263–274. https://doi.org/10.1016/j.neuron.2006.06.006

Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, *22*(10), Article 10. https://doi.org/10.1038/s41593-019-0467-3

Knierim, J. J. (2015). The hippocampus. *Current Biology*, *25*(23), R1116–R1121. https://doi.org/10.1016/j.cub.2015.10.049

Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust Conjunctive Item–Place Coding by Hippocampal Neurons Parallels Learning What Happens Where. *The Journal of Neuroscience*, *29*(31), 9918–9929. https://doi.org/10.1523/JNEUROSCI.1378-09.2009

Körholz, J. C., Zocher, S., Grzyb, A. N., Morisse, B., Poetzsch, A., Ehret, F., Schmied, C., & Kempermann, G. (2018). Selective increases in inter-individual variability in response to environmental enrichment in female mice. *eLife*, *7*, e35690. https://doi.org/10.7554/eLife.35690

Kraus, B. J., Robinson, R. J., White, J. A., Eichenbaum, H., & Hasselmo, M. E. (2013). Hippocampal "Time Cells": Time versus Path Integration. *Neuron*, *78*(6), 1090–1101. https://doi.org/10.1016/j.neuron.2013.04.015

Lee, H., Ghim, J.-W., Kim, H., Lee, D., & Jung, M. (2012). Hippocampal Neural Correlates for Values of Experienced Events. *Journal of Neuroscience*, *32*(43), 15053–15065. https://doi.org/10.1523/JNEUROSCI.2806-12.2012

Lee, I., Rao, G., & Knierim, J. J. (2004). A Double Dissociation between Hippocampal Subfields. *Neuron*, *42*(5), 803–815. https://doi.org/10.1016/j.neuron.2004.05.010

Lee, J. S., Briguglio, J. J., Cohen, J. D., Romani, S., & Lee, A. K. (2020). The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value. *Cell*, *183*(3), 620-635.e22. https://doi.org/10.1016/j.cell.2020.09.024

Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. A., McNaughton, B. L., Moser, M.-B., & Moser, E. I. (2005). Progressive Transformation of Hippocampal Neuronal Representations in "Morphed" Environments. *Neuron*, *48*(2), 345–358. https://doi.org/10.1016/j.neuron.2005.09.007

Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005). Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. *Science*, *309*(5734), 619–623. https://doi.org/10.1126/science.1114037

Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1. *Science*, *305*(5688), 1295–1298. https://doi.org/10.1126/science.1100265

Levy, E. R. J., Carrillo-Segura, S., Park, E. H., Redman, W. T., Hurtado, J. R., Chung, S., & Fenton, A. A. (2023). *A manifold neural population code for space in hippocampal coactivity dynamics independent of place fields* (p. 2021.07.26.453856). bioRxiv. https://doi.org/10.1101/2021.07.26.453856

Li, S., Cullen, W. K., Anwyl, R., & Rowan, M. J. (2003). Dopamine-dependent facilitation of LTP induction in hippocampal CA1 by exposure to spatial novelty. *Nature Neuroscience*, *6*(5), 526–531. https://doi.org/10.1038/nn1049

Lipton, P. A., Alvarez, P., & Eichenbaum, H. (1999). Crossmodal Associative Memory Representations in Rodent Orbitofrontal Cortex. *Neuron*, *22*(2), 349–359. https://doi.org/10.1016/S0896-6273(00)81095-8

Maaten, L. van der, & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, *9*(86), 2579–2605.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "Time Cells" Bridge the Gap in Memory for Discontiguous Events. *Neuron*, *71*(4), 737–749. https://doi.org/10.1016/j.neuron.2011.07.012

Mamad, O., Stumpp, L., McNamara, H. M., Ramakrishnan, C., Deisseroth, K., Reilly, R. B., & Tsanov, M. (2017). Place field assembly distribution encodes preferred locations. *PLOS Biology*, *15*(9), e2002365. https://doi.org/10.1371/journal.pbio.2002365

Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*, *109*(47), 19462–19467. https://doi.org/10.1073/pnas.1214107109

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

Markus, E. J., Qin, Y. L., Leonard, B., Skaggs, W. E., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, *15*(11), 7079–7094. https://doi.org/10.1523/JNEUROSCI.15-11-07079.1995

McHugh, T. J., Blum, K. I., Tsien, J. Z., Tonegawa, S., & Wilson, M. A. (1996). Impaired Hippocampal Representation of Space in CA1-Specific NMDAR1 Knockout Mice. *Cell*, *87*(7), 1339–1349. https://doi.org/10.1016/S0092-8674(00)81828-0

McKenzie, S., Frank, A. J., Kinsky, N. R., Porter, B., Rivière, P. D., & Eichenbaum, H. (2014). Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. *Neuron*, *83*(1), 202–215. https://doi.org/10.1016/j.neuron.2014.05.019

McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neuroscience*, *17*(12), 1658–1660. https://doi.org/10.1038/nn.3843

McNamara, R. K., & Skelton, R. W. (1993). The neuropharmacological and neurochemical basis of place learning in the Morris water maze. *Brain Research Reviews*, *18*(1), 33–49. https://doi.org/10.1016/0165-0173(93)90006-L

McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, *52*(1), 41–49. https://doi.org/10.1007/BF00237147

Mikula, S., Trotts, I., Stone, J. M., & Jones, E. G. (2007). Internet-enabled high-resolution brain mapping and virtual microscopy. *NeuroImage*, *35*(1), 9–15. https://doi.org/10.1016/j.neuroimage.2006.11.053

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*(2), 463–475. https://doi.org/10.1016/j.neuron.2018.09.023

Mizumori, S. j. y., Ragozzino, K. e., Cooper, B. g., & Leutgeb, S. (1999). Hippocampal Representational Organization and Spatial Context. *Hippocampus*, *9*(4), 444–451. https://doi.org/10.1002/(SICI)1098-1063(1999)9:4<444::AID-HIPO10>3.0.CO;2-Z

Morris, R. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods*, *11*(1), 47–60. https://doi.org/10.1016/0165-0270(84)90007-4

Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, *12*(2), 239–260. https://doi.org/10.1016/0023-9690(81)90020-5

Muller, R., Kubie, J., & Ranck, J. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *The Journal of Neuroscience*, *7*(7), 1935–1950. https://doi.org/10.1523/JNEUROSCI.07-07-01935.1987

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, *7*(7), 1951–1968. https://doi.org/10.1523/JNEUROSCI.07-07-01951.1987

Müllner, D. (2011). *Modern hierarchical, agglomerative clustering algorithms* (arXiv:1109.2378; Version 1). arXiv. https://doi.org/10.48550/arXiv.1109.2378

Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, *7*(2), 217–227. https://doi.org/10.1016/S0959-4388(97)80010-4

Nelson, A. J. D., Hindley, E. L., Haddon, J. E., Vann, S. D., & Aggleton, J. P. (2014). A novel role for the rat retrosplenial cortex in cognitive control. *Learning & Memory*, *21*(2), 90–97. https://doi.org/10.1101/lm.032136.113

Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D., & Tank, D. W. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature*, *595*(7865), 80–84. https://doi.org/10.1038/s41586-021-03652-7

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'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: Reactivation of waking experience and memory. *Trends in Neurosciences*, *33*(5), 220–229. https://doi.org/10.1016/j.tins.2010.01.006

Passecker, J., Mikus, N., Malagon-Vina, H., Anner, P., Dimidschstein, J., Fishell, G., Dorffner, G., & Klausberger, T. (2019). Activity of Prefrontal Neurons Predict Future Choices during Gambling. *Neuron*, *101*(1), 152-164.e7. https://doi.org/10.1016/j.neuron.2018.10.050

Patai, E. Z., Javadi, A.-H., Ozubko, J. D., O'Callaghan, A., Ji, S., Robin, J., Grady, C., Winocur, G., Rosenbaum, R. S., Moscovitch, M., & Spiers, H. J. (2019). Hippocampal and Retrosplenial Goal

Distance Coding After Long-term Consolidation of a Real-World Environment. *Cerebral Cortex*, *29*(6), 2748–2758. https://doi.org/10.1093/cercor/bhz044

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

Piterkin, P., Cole, E., Cossette, M.-P., Gaskin, S., & Mumby, D. G. (2008). A limited role for the hippocampus in the modulation of novel-object preference by contextual cues. *Learning & Memory*, *15*(10), 785–791. https://doi.org/10.1101/lm.1035508

Priestley, J. B., Bowler, J. C., Rolotti, S. V., Fusi, S., & Losonczy, A. (2022). Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences. *Neuron*, *110*(12), 1978-1992.e6. https://doi.org/10.1016/j.neuron.2022.03.026

Ramadan, W., Eschenko, O., & Sara, S. J. (2009). Hippocampal Sharp Wave/Ripples during Sleep for Consolidation of Associative Memory. *PLoS ONE*, *4*(8), e6697. https://doi.org/10.1371/journal.pone.0006697

Raposo, D., Kaufman, M. T., & Churchland, A. K. (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience*, *17*(12), 1784–1792. https://doi.org/10.1038/nn.3865

Recanatesi, S., Farrell, M., Lajoie, G., Deneve, S., Rigotti, M., & Shea-Brown, E. (2021). Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nature Communications*, *12*(1), Article 1. https://doi.org/10.1038/s41467-021-21696-1

Redish, A. D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: Implications for addiction, relapse, and problem gambling. *Psychological Review*, *114*(3), 784–805. https://doi.org/10.1037/0033-295X.114.3.784

Riceberg, J. S., & Shapiro, M. L. (2012). Reward Stability Determines the Contribution of Orbitofrontal Cortex to Adaptive Behavior. *Journal of Neuroscience*, *32*(46), 16402–16409. https://doi.org/10.1523/JNEUROSCI.0776-12.2012

Rich, P. D., Liaw, H.-P., & Lee, A. K. (2014). Large environments reveal the statistical structure governing hippocampal representations. *Science*, *345*(6198), 814–817. https://doi.org/10.1126/science.1255635

Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, *497*(7451), Article 7451. https://doi.org/10.1038/nature12160

Rule, M. E., Loback, A. R., Raman, D. V., Driscoll, L. N., Harvey, C. D., & O'Leary, T. (2020). Stable task information from an unstable neural population. *eLife*, *9*, e51121. https://doi.org/10.7554/eLife.51121

Samborska, V., Butler, J. L., Walton, M. E., Behrens, T. E. J., & Akam, T. (2022). Complementary task representations in hippocampus and prefrontal cortex for generalizing the structure of problems. *Nature Neuroscience*, *25*(10), Article 10. https://doi.org/10.1038/s41593-022-01149-8

Sauvage, M. M., Fortin, N. J., Owens, C. B., Yonelinas, A. P., & Eichenbaum, H. (2008). Recognition memory: Opposite effects of hippocampal damage on recollection and familiarity. *Nature Neuroscience*, *11*(1), Article 1. https://doi.org/10.1038/nn2016

Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2016). Activity-dependent plasticity of hippocampal place maps. *Nature Communications*, *7*(1), Article 1. https://doi.org/10.1038/ncomms11824

Scoville, W. B., & Milner, B. (1957). LOSS OF RECENT MEMORY AFTER BILATERAL HIPPOCAMPAL LESIONS. *Journal of Neurology, Neurosurgery & Psychiatry*, *20*(1), 11–21. https://doi.org/10.1136/jnnp.20.1.11

Shapiro, M. L., Tanila, H., & Eichenbaum, H. (1997). Cues that hippocampal place cells encode: Dynamic and hierarchical representation of local and distal stimuli. *Hippocampus*, *7*(6), 624–642. https://doi.org/10.1002/(SICI)1098-1063(1997)7:6<624::AID-HIPO5>3.0.CO;2-E

Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR Activity Predicts Correct Decisions during the Initial Learning of an Alternation Task. *Neuron*, *77*(6), 1163–1173. https://doi.org/10.1016/j.neuron.2013.01.027

Skaggs, W. E., McNaughton, B. L., Gothard, K. M., & Markus, E. J. (1992). An information-theoretic approach to deciphering the hippocampal code. *Proceedings of the 5th International Conference on Neural Information Processing Systems*, 1030–1037.

Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*(2), 149–172. https://doi.org/10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K

Smith, D. M., Barredo, J., & Mizumori, S. J. Y. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus*, *22*(5), 1121–1133. https://doi.org/10.1002/hipo.20958

Smith, D. M., & Mizumori, S. J. Y. (2006). Learning-related development of context-specific neuronal responses to places and events: The hippocampal role in context processing. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *26*(12), 3154–3163. https://doi.org/10.1523/JNEUROSCI.3234-05.2006

Squire, L. R. (1992). Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, *99*(2), 195–231. https://doi.org/10.1037/0033-295X.99.2.195

Squire, L. R., Cohen, N. J., & Nadel, L. (1984). The Medial Temporal Region and Memory Consolidation: A New Hypothesis. In *Memory Consolidation*. Psychology Press.

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

Stefanini, F., Kushnir, L., Jimenez, J. C., Jennings, J. H., Woods, N. I., Stuber, G. D., Kheirbek, M. A., Hen, R., & Fusi, S. (2020). A Distributed Neural Code in the Dentate Gyrus and in CA1. *Neuron*, *107*(4), 703-716.e4. https://doi.org/10.1016/j.neuron.2020.05.022

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

Talamini, L. M., Nieuwenhuis, I. L. C., Takashima, A., & Jensen, O. (2008). Sleep directly following learning benefits consolidation of spatial associative memory. *Learning & Memory*, *15*(4), 233–237. https://doi.org/10.1101/lm.771608

Tanaka, K. Z., He, H., Tomar, A., Niisato, K., Huang, A. J. Y., & McHugh, T. J. (2018). The hippocampal engram maps experience but not place. *Science*, *361*(6400), 392–397. https://doi.org/10.1126/science.aat5397

Tang, W., Shin, J. D., & Jadhav, S. P. (2023). Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. *Cell Reports*, *42*(3), 112246. https://doi.org/10.1016/j.celrep.2023.112246

Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, *15*(1), 70–86. https://doi.org/10.1523/JNEUROSCI.15-01-00070.1995

Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, *10*(2), 420–435. https://doi.org/10.1523/JNEUROSCI.10-02-00420.1990

Tolman, E. C., & Honzik, C. H. (1930). Introduction and removal of reward, and maze performance in rats. *University of California Publications in Psychology*, *4*, 257–275.

Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374–391. https://doi.org/10.1002/hipo.450040319

Treves, A., Skaggs, W. E., & Barnes, C. A. (1996). How much of the hippocampus can be explained by functional constraints? *Hippocampus*, *6*(6), 666–674. https://doi.org/10.1002/(SICI)1098-1063(1996)6:6<666::AID-HIPO9>3.0.CO;2-E

Tryon, V. L., Penner, M. R., Heide, S. W., King, H. O., Larkin, J., & Mizumori, S. J. Y. (2017). Hippocampal neural activity reflects the economy of choices during goal-directed navigation. *Hippocampus*, *27*(7), 743–758. https://doi.org/10.1002/hipo.22720

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., Witter, M. P., & Morris, R. G. M. (2007). Schemas and Memory Consolidation. *Science*, *316*(5821), 76–82. https://doi.org/10.1126/science.1135935

Tulving, E. (1972). Episodic and semantic memory. In *Organization of memory* (pp. xiii, 423–xiii, 423). Academic Press.

Vaidya, S. P., Chitwood, R. A., & Magee, J. C. (2023). *The formation of an expanding memory representation in the hippocampus* [Preprint]. Neuroscience. https://doi.org/10.1101/2023.02.01.526663

Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, *10*(11), 792–802. https://doi.org/10.1038/nrn2733

Vanni-Mercier, G., Mauguière, F., Isnard, J., & Dreher, J.-C. (2009). The Hippocampus Codes the Uncertainty of Cue–Outcome Associations: An Intracranial Electrophysiological Study in Humans. *Journal of Neuroscience*, *29*(16), 5287–5294. https://doi.org/10.1523/JNEUROSCI.5298-08.2009

Viejo, G., & Peyrache, A. (2020). Precise coupling of the thalamic head-direction system to hippocampal ripples. *Nature Communications*, *11*(1), Article 1. https://doi.org/10.1038/s41467-020-15842-4

Vyazovskiy, V. V., Olcese, U., Lazimy, Y. M., Faraguna, U., Esser, S. K., Williams, J. C., Cirelli, C., & Tononi, G. (2009). Cortical Firing and Sleep Homeostasis. *Neuron*, *63*(6), 865–878. https://doi.org/10.1016/j.neuron.2009.08.024

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

Whitlock, J. R., Sutherland, R. J., Witter, M. P., Moser, M.-B., & Moser, E. I. (2008). Navigating from hippocampus to parietal cortex. *Proceedings of the National Academy of Sciences*, *105*(39), 14755–14762. https://doi.org/10.1073/pnas.0804216105

Wible, C. G., Findling, R. L., Shapiro, M., Lang, E. J., Crane, S., & Olton, D. S. (1986). Mnemonic correlates of unit activity in the hippocampus. *Brain Research*, *399*(1), 97–110. https://doi.org/10.1016/0006-8993(86)90604-9

Wiener, S., Paul, C., & Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *The Journal of Neuroscience*, *9*(8), 2737–2763. https://doi.org/10.1523/JNEUROSCI.09-08-02737.1989

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O'Keefe, J. (2005). Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science*, *308*(5723), 873–876. https://doi.org/10.1126/science.1108905

Wood, E. R., Dudchenko, P. A., & Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, *397*(6720), Article 6720. 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

Xu, H., Baracskay, P., O'Neill, J., & Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze. *Neuron*, *101*(1), 119-132.e4. https://doi.org/10.1016/j.neuron.2018.11.015

Yoon, T., Okada, J., Jung, M. W., & Kim, J. J. (2008). Prefrontal cortex and hippocampus subserve different components of working memory in rats. *Learning & Memory*, *15*(3), 97–105. https://doi.org/10.1101/lm.850808

Zemla, R., Moore, J. J., Hopkins, M. D., & Basu, J. (2022). Task-selective place cells show behaviorally driven dynamics during learning and stability during memory recall. *Cell Reports*, *41*(8), 111700. https://doi.org/10.1016/j.celrep.2022.111700

Zhang, H., Rich, P. D., Lee, A. K., & Sharpee, T. O. (2023). Hippocampal spatial representations exhibit a hyperbolic geometry that expands with experience. *Nature Neuroscience*, *26*(1), 131–139. https://doi.org/10.1038/s41593-022-01212-4

Zhang, L., Prince, S. M., Paulson, A. L., & Singer, A. C. (2022). Goal discrimination in hippocampal nonplace cells when place information is ambiguous. *Proceedings of the National Academy of Sciences*, *119*(11), e2107337119. https://doi.org/10.1073/pnas.2107337119

Zhao, X., Wang, Y., Spruston, N., & Magee, J. C. (2020). Membrane potential dynamics underlying context-dependent sensory responses in the hippocampus. *Nature Neuroscience*, *23*(7), Article 7. https://doi.org/10.1038/s41593-020-0646-2

Zielinski, M. C., Shin, J. D., & Jadhav, S. P. (2019). Coherent Coding of Spatial Position Mediated by Theta Oscillations in the Hippocampus and Prefrontal Cortex. *Journal of Neuroscience*, *39*(23), 4550–4565. https://doi.org/10.1523/JNEUROSCI.0106-19.2019

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

Zucker, R. S., & Regehr, W. G. (2002). Short-Term Synaptic Plasticity. *Annual Review of Physiology*, *64*(1), 355–405. https://doi.org/10.1146/annurev.physiol.64.092501.114547