

A Novel Paradigm to Investigate the Role of the Retrosplenial Cortex in Spatial Learning and Memory in Freely Behaving Mice

Ob du wirklich richtig stehst, siehst du, wenn das Licht angeht

Dissertation der Fakultät für Biologie
der Ludwig-Maximilians-Universität München



Annette Loidolt

München, 2024

Supervision:

Diese Dissertation wurde angefertigt unter der Leitung von Prof. Tobias Bonhoeffer im Bereich des Max Planck Institute for Biological Intelligence an der Ludwig-Maximilians-Universität München.

Thesis committee:

Erstgutachter: Prof. Tobias Bonhoeffer

Zweitgutachterin: Prof. Laura Busse

Date of submission:

Tag der Abgabe: 10th June 2024

Date of disputation:

Tag der mündlichen Prüfung: 24th April 2025

Erklärung

Ich versichere hiermit an Eides statt, dass meine Dissertation selbstständig und ohne unerlaubte Hilfsmittel angefertigt worden ist.

Die vorliegende Dissertation wurde weder ganz, noch teilweise bei einer anderen Prüfungskommission vorgelegt.

Ich habe noch zu keinem früheren Zeitpunkt versucht, eine Dissertation einzureichen oder an einer Doktorprüfung teilzunehmen.

München, 24. April 2025

Annette Loidolt

TABLE OF CONTENTS

ABSTRACT	IX
1. INTRODUCTION	1
1.1 Navigation and spatial cognition	1
1.2 Spatial memory tasks	3
1.2.1 Behaviour as a tool to study spatial learning and memory	3
1.2.2 Radial arm maze	4
1.2.3 Morris water maze	4
1.2.4 Barnes circular maze	6
1.2.5 Virtual reality mazes	7
1.2.6 Visual cues, landmarks, and geometric information	7
1.2.7 Summary	8
1.3 Neural mechanisms underlying spatial cognition	9
1.3.1 Cognitive map	9
1.3.2 Place cells	10
1.3.3 Head direction cells	11
1.3.4 Grid cells	12
1.3.5 Boundary vector cells	12
1.3.6 Summary	13
1.4 Retrosplenial cortex	15
1.4.1 Anatomy and connectivity	15
1.4.2 Proposed function	17
1.4.2.1 Navigating spaces: egocentric vs. allocentric strategy	17
1.4.2.2 Egocentric coordinate frame	17
1.4.2.3 Allocentric coordinate frame	19
1.4.2.4 Transformation of reference frames	20
1.4.2.5 RSC's role in spatial memory	23
1.5 Aim of this study	25
2. METHODS	27
2.1 Animals	27
2.2 Cohorts and experimental timeline	27
2.3 Surgery	28
2.4 Spatial memory task	30
2.4.1 Arena and reward ports	30
2.4.2 Arena cues	31
2.4.3 Liquid reward	32

2.4.4	Position tracking	32
2.4.5	Arena ceiling	33
2.4.6	Single-trial structure	33
2.4.7	Error trials	33
2.4.8	Session structure	33
2.4.9	Reward port selection.....	34
2.4.10	Trigger placement.....	35
2.4.11	Habituation	37
2.4.12	Pre-training	37
2.4.13	Spatial memory training	37
2.4.14	Dropout criteria	38
2.5	Food deprivation	38
2.6	Chemogenetic inactivation.....	38
2.7	<i>In vivo</i> two-photon Ca²⁺ imaging	38
2.8	Histology and immunohistochemistry	39
2.9	Data analysis	39
2.9.1	Analysis of the Ca ²⁺ imaging dataset.....	39
2.9.2	Analysis of the behavioural dataset.....	40
3.	RESULTS	43
3.1	Spatial memory task performance	43
3.1.1	Mice form spatial memories	43
3.1.2	Low and stable error trial rate	45
3.1.3	Quantifying behavioural task metrics in 3D.....	46
3.1.4	Task engagement.....	47
3.1.5	Error angles differ between visible and hidden trigger trials	52
3.1.6	Summary.....	57
3.2	Chemogenetic inhibition of RSC activity	59
3.2.1	DREADD/CNO-mediated reduction in RSC cell activity.....	59
3.2.2	Mice form spatial memories under CNO conditions	61
3.2.3	Summary.....	66
3.3	Comparing memory-guided and visually-guided navigation	67
3.3.1	Introduction of visible-fixed trigger zones	67
3.3.2	Comparable behavioural metrics in both cohorts	67
3.3.3	Mice only use memory-guided navigation in the absence of visual information	71
3.3.4	Summary.....	73
3.4	Manipulation of arena cues and trigger placement	75
3.4.1	Controlling for baseline performance.....	76
3.4.2	Mice form allocentric spatial memories	77
3.4.3	Summary.....	80

3.5	Mice form spatial memories within a single day	81
3.6	Summary.....	85
4.	DISCUSSION	87
4.1	Do mice form spatial memories?.....	89
4.1.1	In search of a learning curve.....	89
4.1.2	Identifying task engagement in freely behaving animals	92
4.1.3	Quantifying navigation strategies.....	93
4.1.4	Vision vs. memory.....	95
4.2	How fast do mice form spatial memories?	97
4.2.1	Evidence for within-day learning	97
4.2.2	Task benefits	98
4.3	Are spatial memories allocentric?	101
4.3.1	Task design prevents path integration	101
4.3.2	Experimental evidence for allocentric memory.....	102
4.3.3	Why are allocentric memories not anchored to arena cues?.....	103
4.4	Is RSC involved in spatial memory formation?	105
4.4.1	Experimental evidence for chemogenetic RSC inactivation	105
4.4.2	Potential involvement of RSC in spatial memory formation	105
4.4.3	Remaining open questions and next steps	107
4.5	Conclusion.....	109
5.	BIBLIOGRAPHY	111
6.	SUPPLEMENTARY FIGURES	127
	List of Figures.....	135
	Abbreviations	137
	Acknowledgements	139

Abstract

Spatial cognition is vital for the survival of many species, enabling mice and other animals to navigate their environment, locate food resources, and escape from threats. Meanwhile, spatial learning and memory allow an individual to acquire, retain, and recall spatial information. The retrosplenial cortex (RSC) has recently gained attention for its role in spatial cognition and memory, integrating sensory and spatial signals. Despite extensive research on the neural circuits involved in spatial learning through employment of various spatial memory tasks, the specific contributions of the rodent RSC to spatial memory are not fully understood. This gap in knowledge emphasises the need for innovative behavioural paradigms to address previous spatial memory tasks' limitations. In this thesis, I present a novel behavioural paradigm designed to explore the role of the RSC to spatial learning and memory.

The task I developed investigates mechanisms of spatial memory formation, retrieval, and reversal learning in rodents, wherein mice learn the locations of hidden trigger zones within a circular arena. This spatial memory task stands out for its ability to track mice, their behaviour, and task performance in a three-dimensional environment without any physical constraints on head or body movement. This allows for naturalistic animal behaviour and the simultaneous use of portable neuronal recording devices during training. In this study, I identified key task metrics for quantifying behaviour in the spatial memory task, including error angles and active task engagement, to assess learning and task performance. Using these parameters, I demonstrate that the goal-directed navigational strategies employed by mice vary significantly whether they are guided by vision or memory. Specifically, mice resort to memory-guided navigation exclusively in the absence of visual landmark cues. Through manipulations of the external landmarks within the animals' environment during spatial memory training, I demonstrate that mice develop a comprehensive cognitive map of the arena and the surrounding space, using allocentric navigation strategies to locate previously learnt goal locations. The results of this study further show that mice can form robust allocentric spatial memories quickly and efficiently, even under conditions where RSC activity is inhibited. Chemogenetic experiments show that inactivating RSC neuronal activity does not impede the formation or retrieval of spatial memories in this task.

Taken together, these findings highlight the effectiveness of this task for investigating allocentric spatial learning and memory in mice. Characterized by rapid acquisition of task proficiency, high trial rates, reliable memory formation and recall, brief pre-training periods, and unrestricted, natural mouse behaviour, this paradigm opens up new avenues for advanced research into the neural substrates of spatial learning and memory with broad applications in RSC research and beyond.

1. Introduction

1.1 Navigation and spatial cognition

Spatial learning and memory are crucial factors when it comes to an individual's survival — whether that be the ability to explore uncharted places while retaining a sense of direction, to re-locate previously encountered food stocks, or to take the fastest route to the nearest shelter when escaping a predator. Consider the example of a mouse, which, along with its mischief, resides in the London Underground — a 'tube mouse'. Such a tube mouse can see the food crumbs scattered on the platform, maybe even smell them from afar. To a certain degree, the mouse can use its taste receptors to investigate whether those crumbs are in fact edible and safe to consume. Yet, it is the mouse's memory of specific locations, such as the underground pizza vendor where humans leave food scraps and pizza crusts near the bin, that ensures the mouse a semi-reliable food resource. This mouse is also capable of seeing and hearing a human approaching, and perceiving it as a threat. It might even feel the vibration of the ground caused by the nearing footsteps. However, ultimately, it is the knowledge of where the nearest shelter is located and the learnt escape route to safety that will rescue that mouse from being caught or stepped on.

As children, from an early age onwards, we are taught that there are five senses with which we can perceive the world and our surrounding: sight, hearing, taste, smell, and touch. These senses are universally recognized as essential tools for survival across species. However, we often tend to overlook what else is needed in order to survive in the world we live in. One often overlooked aspect of survival is spatial cognition — the ability of an individual to explore and navigate its environment effectively using various sensory and spatial cues. For many species, including the tube mouse, the ability to identify and relocate food sources, avoid predators, and safely navigate through complex environments is as vital as the traditional five senses.

This intricate interaction of sensory input, navigation, and spatial memory is not unique to mice. It is a shared necessity among various species, each displaying remarkable spatial abilities tailored to their peculiar environment. Grey squirrels, for example, can remember the location of their food caches (Jacobs & Liman, 1991) and can relocate them using both visual landmarks and olfactory cues (McQuade et al., 1986). Western scrub-jays, a bird species known for its food caching capabilities, have been shown to re-hide food caches to reduce theft by conspecifics (Clayton et al., 2007), while octopuses avoid re-visiting the same foraging areas that they have previously depleted of resources (Forsythe & Hanlon, 1997). The Egyptian fruit bat has been shown to take very straight and consistent long-range flights to return to the same individual tree night after night, and successfully returns to this location when displaced more than 80 km away from its familiar fruit tree using echolocation to navigate its three-dimensional environment (Tsoar et al., 2011).

The importance of spatial cognition, as well as spatial learning and memory — the cognitive processes that enable an individual to acquire, retain, and recall information about its spatial environment and its spatial location and orientation within it — becomes apparent in humans when these abilities begin to deteriorate, for example, as seen in Alzheimer's disease patients. Such patients struggle to locate objects in their own home or have difficulties in their everyday life navigating previously familiar routes or using public transport, pointing to the fundamental nature of spatial cognition in our daily lives. The effects of spatial

disorientation and navigation deficits in Alzheimer's disease patients are so apparent that they have been suggested as a cognitive marker for preclinical Alzheimer's disease in humans (Coughlan et al., 2018).

Spatial cognition is generally obtained through exploratory behaviour, an instinctive and widely preserved behaviour in all animal species, as well as humans, and spatial memory evolved in many different species. Such behaviours are an expression of an innate need to gather spatial information critical for survival (Berlyne, 1950; Thinus-Blanc, 1996). Despite this universal behaviour and the extensive research around spatial cognition, learning and memory, the underlying neural computations are still a topic of ongoing research. For instance, despite having identified many key brain regions involved in spatial cognition, such as the hippocampus, we still do not understand how these regions interact with other areas, such as the retrosplenial cortex (RSC). The RSC has recently gained attention through its involvement in spatial navigation and memory, the exact nature of which is still unknown. Even though there exists a wide array of behavioural tasks and mazes to investigate spatial learning and memory in both rodents and humans, they all come with their limitations in investigating the neural mechanisms underlying spatial learning and memory in a naturalistic setting.

1.2 Spatial memory tasks

1.2.1 Behaviour as a tool to study spatial learning and memory

The advancement of technology in recent years has significantly transformed the landscape of neuroscience, particularly in the realm of behavioural quantification. With the advancement of animal tracking technology, deep learning and computer vision, researchers now possess the capability to automate the quantification of animal behaviour with unprecedented precision, opening up new avenues for understanding how animals interact with their environments (Pereira, Shaevitz, & Murthy, 2020).

Quantitative descriptions of behaviour, crucially, begin with tracking movements, an endeavour that historically has been time consuming and labour intensive due to manual annotations of video material. However, the emergence of animal pose estimation frameworks, such as DeepLabCut (Mathis et al., 2018), SLEAP (Pereira, Tabris, et al., 2020), and DeepPoseKit (Graving et al., 2019), has significantly mitigated these challenges. By providing detailed analyses of how animals navigate spatial memory tasks, modern tools for behavioural quantification are deepening our understanding of the brain-behaviour nexus. These advancements, therefore, position behavioural description and quantification as an indispensable component of spatial neuroscience research.

However, while the tracking and quantification of animal behaviour keeps improving, the need for new behavioural paradigms grows as the complexity of the cognitive processes we aim to understand, such as spatial cognition and learning, becomes more and more apparent. The following sections will provide an overview of the most commonly used tasks and in rodent research: the radial arm maze, the Morris water maze, and the Barnes circular maze.

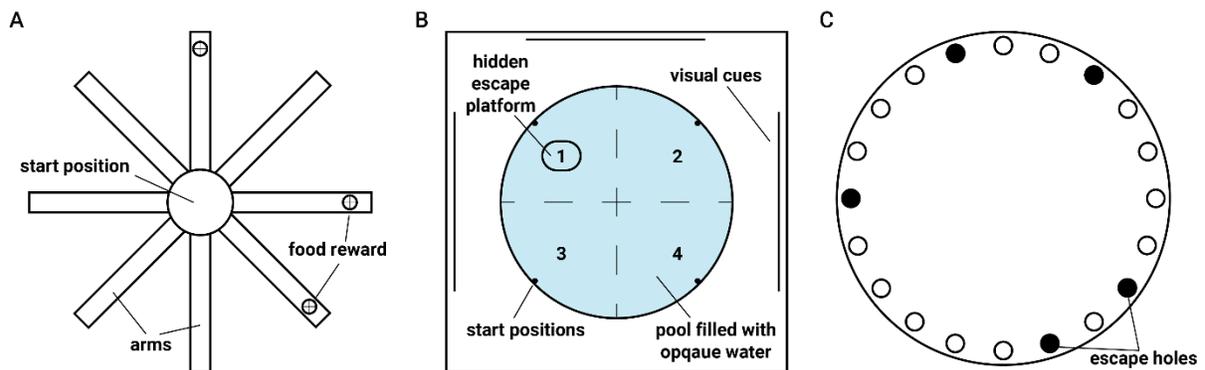


Figure 1.1 Different mazes to evaluate spatial memory

Schematic representations of the 8-arm radial maze (A), the Morris water maze (B), and the Barnes circular maze (C). Adapted from (Paul et al., 2009)

1.2.2 Radial arm maze

The radial arm maze was originally developed by Olton and Samuelson (1976) for rats to study spatial working memory, which is the ability to retain spatial information over a short period of time. As the name implies, the maze consists of eight linear routes (or arms) equidistantly spaced and radiating from a shared centre point that connects the arms. At the end of each arm (or, at the end of a few selected arms) a reward in form of a food pellet is placed in such a way that it is not visible to the animal from the centre of the maze or from the neighbouring arms (e.g. by hiding the pellet in a container or burrowed in the sand) (Figure 1.1 A). Food-deprived animals are thus motivated to locate and visit only arms that contain a food reward. For spatial orientation purposes, the maze is typically surrounded by distal visual cues, used by the animal to identify the baited arms (Morellini, 2013; Paul et al., 2009).

In its original form, the radial arm maze was intended to study spatial working memory, whereby a food pellet is placed at the end of each of the eight arms. The animal starts a training session from the centre of the maze and is allowed to explore freely the arena and all arms until all eight pellets have been collected. Optimal performance of this version of the radial arm maze would be achieved when the animal is visiting each arm only once, collecting all rewards in the least amount of time, while each re-visit of a maze arm is considered an error. Olton and Samuelson theorized that animals generate a 'mental list', an internal record of previously explored and thus unbaited arms in order to successfully master the maze.

Subsequent experiments showed that animals rely predominantly on extra-maze visual cues to navigate the maze and to recognize arms, as a drop in performance was observed when the position of those landmarks was either modified or their visibility reduced (Mazmanian & Roberts, 1983; Suzuki et al., 1980). However, animals can use a strategy by which they consistently visit the arm left (or, right) of the currently visited arm, going in a clockwise (or anticlockwise) direction and thus eliminating the need to recognize and remember specific arms, contradicting the idea of a 'mental list'. This can be prevented by retaining the animal within a given arm for a set amount of time with the help of mobile doors (Dubreuil et al., 2003; Olton et al., 1977). In another version of the radial arm maze, intended to study spatial long-term memory, food pellets are only located in a subset, most often in only three or four, of the eight arms, and over the course of several training days, animals learn to remember the location of the baited arms (Cooper & Mizumori, 1999; He et al., 2002). In other studies, the number of arms were changed to modify the difficulty levels (Cole & Chappell-Stephenson, 2003; Lenck-Santini et al., 2001; Olton et al., 1977). In all of these versions of the task, however, experimental feasibility is reduced by the possibility of olfactory cues being left behind animals in already visited arms, thereby helping the animal to avoid these arms. In addition, this task comes with the limitation that arms are static and thus only allow for straight, pre-defined trajectories so that animals do not need to remember specific routes (Morellini, 2013; Paul et al., 2009).

1.2.3 Morris water maze

The Morris water maze has become one of the most popular behavioural paradigms for studying spatial learning and memory in rodents. It gained popularity due in large part to its dependency on hippocampal function (Bannerman et al., 1999; Morris, 1984). Originally invented by Richard Morris as an alternative

to the radial arm maze, this task was specifically designed to study the role of visual landmarks for spatial orientation and reference (Morris, 1981).

The maze consists of a large, circular pool filled with opaque water. Within this pool, an escape platform is placed in one of the quadrants of the arena, that is slightly submerged a few centimetres below the opaque water surface, making it invisible to the animal. During the acquisition period of the task, the animal is released into the pool from various starting positions at the wall of the arena, triggering an immediate escape response. Motivated by the instinctive desire to escape the water, animals are required to find the hidden platform within the pool (Figure 1.1 B). Animals undergo several trials (4-6) per day for a training period of about 5-10 days, during which the animal learns to remember the location of the platform using the surrounding visual cues for spatial orientation. Memory is subsequently tested in a single probe trial, in which the platform is either moved to a different quadrant of the arena or removed completely, allowing researchers to investigate the spatial learning strategy used by the animal to locate the platform (Morellini, 2013; Paul et al., 2009).

Performance in the Morris water maze is generally evaluated by measuring escape latency—the time required to find the platform within a single trial. Since the platform's diameter (around 10 - 15 cm) is significantly smaller than that of the pool (around 2 m), the chance of randomly encountering the platform is low, necessitating the use of spatial strategies and landmarks to find the fastest route to the platform. Additionally, commonly used measures of performance include the distance travelled to reach the platform, swimming velocity, time spent in the quadrant containing the platform, and the mean proximity to the former platform location during the probe trial when the platform location has been moved to a different arena quadrant. The latter has been identified as the most sensitive measure for assessing water maze performance (Maei et al., 2009).

The Morris water maze was initially designed to stimulate the use of landmarks surrounding the pool and their spatial configuration to form a global, viewpoint-invariant (that is, north-south/east-west) understanding of the maze, which allows animals to locate the hidden platform independent of their starting position in the pool. The behavioural task, however, has since been adapted in various studies to explore different learning and navigation strategies. For instance, marking the platform with a flag, instead of providing extra-maze cues, prompts animals to follow this proximal visual guide in a simple stimulus-response, whereby these animals neither need to apply any spatial strategy nor acquire spatial memory (see (Eichenbaum et al., 1999) for a more detailed explanation of the stimulus-response). If being consistently released into the pool from the same starting point, animals could use their memory of distance and direction to the platform upon removal of the platform flag during the probe trial allowing researchers to assess spatial memory independently of direct visual guides. This variation of the task is sometimes used at the end of an experimental timeline to evaluate whether any previously observed change in behaviour is attributable to spatial memory deficits or is caused by alterations in other functions that can impact task performance (Czajkowski et al., 2014). Similarly, by performing the probe trial in darkness, or by removing or rotating the extra-maze cues, animals learn to remember a sequence of movements needed to reach the platform (Moghaddam & Bures, 1996), that is, they need to remember both direction and distance from the starting point to the platform - a strategy called path integration.

The precision exhibited by animals in finding the platform suggests that animals are capable to gain a global understanding of their environment with the aid of distal visual cues and their relative configuration, making this task ideal to study spatial learning and memory (Paul et al., 2009). However, the use of water immersion as an aversive stimulus introduces stress as a potential confounding factor, with performance differences also attributable to varying susceptibility to stress (Francis et al. 1995) and maze diameter (Van Dam et al., 2006) among mouse strains (further factors influencing task performance are reviewed in (D’Hooge & De Deyn, 2001)). Additionally, animals must learn to overcome the initial instinct to stay close to the walls, circling along the edge of the pool (a behaviour termed thigmotaxis) and to associate the platform with escape from the water (Morellini, 2013). Despite these challenges, the Morris water maze offers significant advantages, including rapid learning (potentially spurred by fear of drowning) and the elimination of the need for water or food deprivation, thereby reducing the overall duration of experiments. The variability in experimental design and use of visual cues and landmarks allows for flexible adaptation to specific research goals, making the Morris water maze a versatile tool for studying spatial learning and memory (Morellini, 2013; Paul et al., 2009).

1.2.4 Barnes circular maze

The Barnes maze is a spatial memory task designed by Carol Barnes (1979) to assess spatial learning and memory in rodents without the use of water, avoiding the stress associated with swimming, and thereby offering an alternative to the Morris water maze. The maze consists of an elevated, circular platform with 18-20 holes around its perimeter, with some of these holes leading to an escape box located underneath the maze via transparent plastic tunnels not visible from the platform (Figure 1.1 C). For spatial orientation, visual cues are placed outside the maze. The task exploits rodents' aversive response to open, brightly lit, or noisy areas to motivate them to seek shelter in one of the escape holes (Barnes, 1979).

During the testing phase, animals are placed in the centre of the platform, and an aversive stimulus, such as a bright light or a loud noise, is being presented to encourage the animal to seek shelter in the escape box. The primary measures of performance to assess the animal's ability to remember the location of the escape holes include the latency to find the correct escape hole (that is, one that is connected to the shelter box) and the number of errors, which are defined as visits to incorrect holes (that is, holes that do not lead to the shelter box) (Paul et al., 2009).

One of the most significant advantages of this behavioural paradigm is the reduction of animal stress, as the task does not involve swimming or immersion in water, and instead leverages the natural behaviour of rodents to explore and seek underground shelter. Furthermore, unlike tasks that require food or water deprivation, the Barnes maze relies on the animal's natural aversion to open, exposed spaces to motivate the search for the escape holes. However, a notable disadvantage of the Barnes maze is that learning can be slower or even absent in some cases. This slower learning curve has been attributed to the lack of highly stressful stimuli, which might not sufficiently motivate the animals to escape, leading to more exploratory behaviour than targeted escape responses (Sunyer et al., 2007). Nevertheless, the Barnes maze offers a valuable behavioural paradigm to assess spatial learning and memory in rodents, particularly when naturalistic behaviour and minimising stress is a priority (Paul et al., 2009).

1.2.5 Virtual reality mazes

The integration of head-fixed rodents and virtual reality (VR) mazes into neuroscience research has provided a novel approach to studying spatial learning and memory, allowing for precise control over environmental and visual stimuli, and the simultaneous recording of neural activity, for example using two-photon microscope setups. This setup often involves the animal running on a treadmill or navigating an air-floating Styrofoam ball, which, while effective for experimental purposes, keeping the animal's head still and centred below the microscope, represents a departure from their natural locomotion and exploration behaviours (Kislin et al., 2014; Nashaat et al., 2016).

One of the primary challenges with this approach is the lack of natural proprioceptive and vestibular feedback. Proprioception, the sense of body position and movement, and the vestibular system, which contributes to balance and spatial orientation, play critical roles in how animals and humans perceive and navigate their environment (Aghajan et al., 2015; Keshavarzi et al., 2022; Wallace et al., 2002). The typical setup of head-fixed VR systems significantly limits or altogether eliminates these types of sensory feedback as the animal's head is fixed in place, and movement is simulated through the VR environment. The lack of natural proprioceptive and vestibular feedback in VR setups has prompted researchers to develop innovative methods to these challenges. For example, Voigts and Harnett (2020) have developed a rotating headpost that allows for simultaneous two-photon imaging in freely moving and rotating mice to better mimic natural locomotion patterns. This setup could be theoretically fitted with a screen to display a VR environment, thereby augmenting common VR mazes with the additional sense of angular head rotation in the horizontal plane (see also (Aronov & Tank, 2014; Chen et al., 2018)).

Despite these limitations, the use of head-fixed rodents in VR mazes offers unique advantages. The ability to record neural activity in real-time while an animal navigates a controlled, virtual environment allows researchers to dissect the neural mechanisms underlying spatial navigation and memory with a level of detail that is difficult to achieve in more naturalistic, freely-behaving settings. While this has opened new avenues for understanding how specific brain regions and neural circuits contribute to spatial cognition, as technology and methodologies continue to evolve, the focus has shifted towards portable recording techniques, such as miniaturised two- and three-photon microscopes (Klioutchnikov et al., 2023; Zong et al., 2021; Zong et al., 2017) and electrophysiological recording probes (Steinmetz et al., 2021), that allow for freely-moving behaviour and the resulting proprioceptive and vestibular processing, as well as for simultaneous recordings of neural activity.

1.2.6 Visual cues, landmarks, and geometric information

Spatial learning in rodents has traditionally focused on their ability to use distal landmarks to navigate and find goals within environments, such as mazes. These distal cues, landmarks outside the maze that the animal cannot interact with directly, contrast with proximal cues, which are elements within (e.g. the platform flag in the Morris water maze) or at the boundary (e.g. arena walls) of the maze. Traditionally, it has been thought that reliance on proximal cues supports stimulus-response strategies rather than the development of a global understanding of the environment, leading to a design preference in spatial tasks that minimizes the presence of proximal cues (Morellini, 2013).

However, the perspective that navigation and spatial memory are primarily supported by distal, extra-maze landmarks has been challenged (Knierim & Hamilton, 2011). Proximal cues, especially those relating to the boundaries of an environment, play a crucial role in navigation, a principle that applies to both animals and humans (Knierim & Hamilton, 2011; Morellini, 2013). This insight has led to a re-evaluation of the importance of environmental geometry in spatial learning. Despite the focus on distal visual cues as spatial orientation aids in the radial arm maze (Olton & Samuelson, 1976), the Morris water maze (Morris, 1984), and Barnes circular maze (Barnes, 1979), the boundaries of these mazes serve as proximal cues, demonstrating that even in tasks designed to emphasize the use of distal cues, proximal cues and the geometric information they provide are inevitable and significant (Sun et al., 2024). Even mazes that do not comprise a physical boundary, such as the elevated platform in the Barnes circular maze, still provide geometric information. Support for this claim comes from studies showing that rodents can navigate an environment and perform place learning based solely on the geometric information provided by the shape of the environment, in the absence of any distinct visual landmarks (Cheng, 1986; Fellini & Morellini, 2011; Fellini et al., 2006; Maurer & Derivaz, 2000; Ramos, 2000). In the case of the Morris water maze task, it could be further shown that mice can learn the location of the hidden platform when the maze was located in a room with a strengthened asymmetric shape (Law et al., 2003), emphasizing the role of environmental geometry and boundary information in the design of spatial memory tasks.

1.2.7 Summary

In concluding this exploration of spatial memory tasks used within rodent neuroscience research, it is clear that each of the discussed mazes — the radial arm maze, the Morris water maze, and the Barnes circular maze — offers unique insights into the mechanisms of navigation and spatial learning, yet also presents its own set of limitations (Paul et al., 2009). The radial arm maze, with its emphasis on spatial working memory, limits search strategies through static and inflexible routes. The Morris water maze, despite its widespread use and valuable contributions to our understanding of spatial learning, relies on a strong aversive stimulus that could affect the animal's performance in ways that may not always reflect purely cognitive deficits, and prevents the application of simultaneous neural recording techniques sensitive to water immersion. The Barnes circular maze, designed to provide a less stressful alternative, may not sufficiently motivate animals, leading to delayed learning or even failure to acquire the desired escape behaviour. Meanwhile, most VR mazes come with the prerequisite of head-fixation preventing naturalistic behaviour, such as rearing and other 3D movements, and limited vestibular processing.

These observations underscore the complexity of designing experiments to study spatial learning and memory, and highlight the ongoing need for innovation in behavioural assessment tools. As we move forward, the development of new recording devices and behavioural paradigms that can overcome the limitations of current mazes is crucial. The quest for more refined and ecologically valid experimental setups is driven by the recognition that learning and memory are influenced by a confluence of factors, including motivational states, sensory inputs, and environmental contexts. Future research tools will likely aim to create more balanced and integrative approaches that can dissect these complex interactions as researchers strive to unravel the intricacies of the brain's navigational and spatial memory systems.

1.3 Neural mechanisms underlying spatial cognition

1.3.1 Cognitive map

The concept of a cognitive map, foundational to our current understanding of navigation and spatial memory, traces its roots back to Edward Tolman's ground-breaking work in 1948. Tolman challenged the prevailing view by proposing that animals navigating mazes do not merely rely on associations between external stimuli (e.g. visual cues) and behavioural responses (e.g. swimming/running towards the goal location). Instead, he suggested that through training, animals develop the ability to infer shortcuts and devise new search strategies, indicative of an internal, global representation or "cognitive map" of their environment (Tolman, 1948). This notion posited that rodents, and by extension other animals, could form global representations of space to efficiently locate goal positions within mazes and their environment, a radical departure from the stimulus-response explanation dominant at the time.

Despite initial scepticism, the discovery of "place cells" in the hippocampus of freely moving rats by John O'Keefe and Jonathan Dostrovsky in 1971 provided empirical support for Tolman's theory. Their work identified specific neurons in the CA1 and CA3 regions of the hippocampus that activated when an animal was in or moved through particular locations in an environment, effectively serving as neural correlates of spatial cognition and memory (O'Keefe & Dostrovsky, 1971). This discovery was not only pivotal in validating the cognitive map hypothesis but also laid the groundwork for further explorations into the neural mechanisms underpinning spatial cognition and navigation. The significance of these findings was further underscored by the awarding of the Nobel Prize in Physiology or Medicine in 2014 to John O'Keefe, May-Britt and Edvard Moser. This accolade was in recognition of their discoveries related to place cells and grid cells, respectively, which has since propelled a wealth of research into the complex neural architectures supporting navigation and spatial memory (Burgess, 2014; Sun et al., 2024). The following section will cover the four fundamental neuronal response properties underlying spatial cognition and the cognitive map theory: place cells, head direction cells, grid cells, and boundary vector cells.

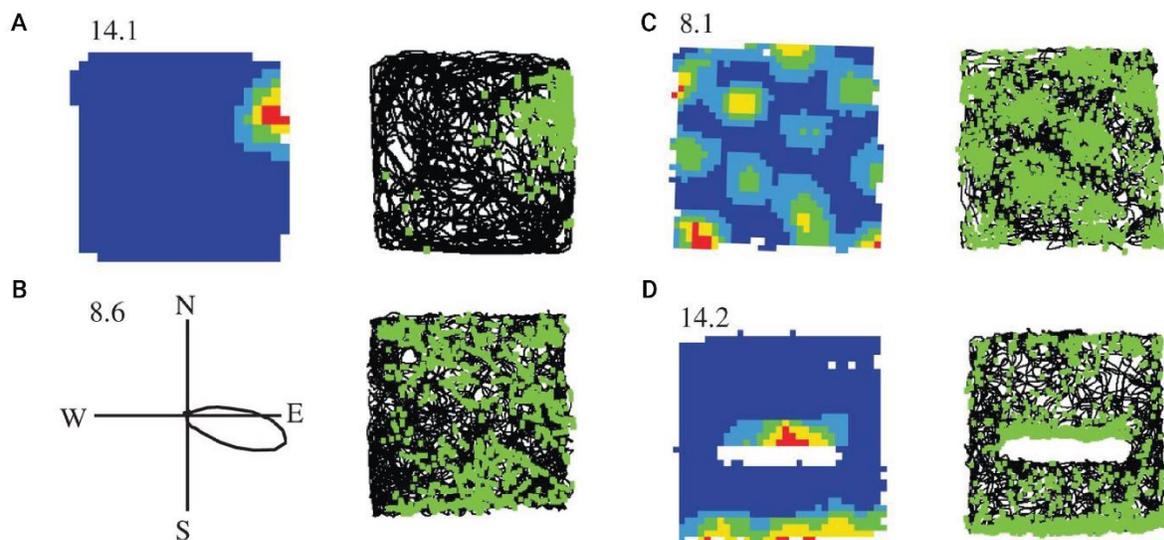


Figure 1.2 The four fundamental neuronal response properties of allocentric spatial coding

An example cell of a place cell (A), head direction cell (B), grid cell (C), and boundary vector cell (D). Firing rates are shown as a function of location (*left, in A, C, and D*) or head direction (*left, in B*), as well as the trajectory of the animal (*black line*) and the locations within the arena at which spikes were recorded (*green*).

Adapted with permission from (Hartley et al., 2014) using data provided by Sarah Stewart and Colin Lever.

1.3.2 Place cells

Place cells were first discovered within the CA1 and CA3 regions of the hippocampus by John O'Keefe and Jonathan Dostrovsky in 1971 during open-field recording experiments in freely moving rats. Their research showed that place cells are specialized neurons characterized by their unique response to spatial environments. Place cells become active whenever an animal enters or occupies a specific location in a given environment (i.e. a place field), and remain silent when the animal is outside this designated place field, coding only for a specific part of the environment (O'Keefe and Dostrovsky, 1971) (Figure 1.2 A).

The establishment of place cell firing patterns occurs rapidly upon an animal's first exposure to a new environment and remains stable over several days (Lever et al., 2002; Thompson & Best, 1990). While the firing pattern of an individual place cell is specific to a given environment, a notable feature of place cells is their ability to "remap", whereby a place cell might code for a specific location in a given environment, but is inactive (or active with a different firing pattern) in another (Muller & Kubie, 1987; O'Keefe & Conway, 1978). This remapping can occur swiftly between different environments (Muller & Kubie, 1987), while place cell firing patterns and place fields are robust to minor environmental changes, such as the removal or absence of a few visual cues (O'Keefe & Conway, 1978; Quirk et al., 1990). Remapping has also been shown to be a result of a change in task requirements (Markus et al., 1995), or in goal location

(Dupret et al., 2010; Ormond & O'Keefe, 2022). Place cell firing patterns also show trajectory dependence, developing preferred firing along particular paths when those trajectories are consistently followed in a constrained setting, such as radial arm or T-shaped mazes (McNaughton, Barnes, and O'Keefe, 1983; Wood et al., 2000), underscoring the dynamic nature of place cells in encoding spatial information. For a more detailed summary of hippocampal place cell features, please refer to the review paper by Chersi and Burgess (2015). Studies using the Morris water maze further clarified the essential connection between hippocampal place cells and spatial memory, demonstrating a significant decline in performance after hippocampal lesions (Morris et al., 1982).

Place cells have not only been identified in rodents, but also in a variety of other mammals, including bats (Ulanovsky & Moss, 2007) and humans (Ekstrom et al., 2003), indicating their fundamental role in navigation and spatial learning across species. Moreover, the firing fields of place cells adapt to 3D spaces, as observed in flying bats (Yartsev & Ulanovsky, 2013) and climbing rats (Grieves et al., 2020).

The discovery of place cells and their representation of spatial location not only highlighted their role in navigation and spatial memory, but also emphasised the need for complete navigational systems within the brain – incorporating not only location, but also directionality and distance.

1.3.3 Head direction cells

Head direction cells are specialized neurons that encode an animal's directional orientation relative to its environment. These cells were initially identified in the pre-subiculum of freely moving rats (Ranck, 1985; Taube et al., 1990a), but subsequent research has found head direction cells distributed throughout the rodent Papez' circuit, a neural pathway in the brain crucial for emotional processing and memory functions. Head direction cells have since been reported in several locations, including thalamic nuclei (Mizumori & Williams, 1993; Taube, 1995), mammillary bodies (Stackman & Taube, 1998), and entorhinal cortex (Sargolini et al., 2006), indicating a widespread neural network for directional sensing.

The defining characteristic of head direction cells is their activation in response to the animal's head facing a particular direction, regardless of the animal's specific location within a given environment. These cells code for the animal's directional heading in the horizontal plane (or, yaw) and are tuned to a narrow range of head directions, each cell having a preferred firing direction (Figure 1.2 B). The orientation of these preferred firing directions is strongly influenced by distal visual cues and landmarks, demonstrating the cells' reliance on external spatial information to guide directional sensing (Chersi & Burgess, 2015). An intriguing aspect of head direction cells is their coordinated activity across different environments. If two cells share a preferred firing direction in one environment, they will maintain this alignment in another environment, despite any changes in the absolute firing direction. This phenomenon indicates that head direction cells operate as a network, with their preferred directions remaining at a fixed angle relative to one another, thereby preserving their spatial relationship across various settings (Taube, 1995; Taube et al., 1990a, 1990b).

The presence of head direction cells has been confirmed in multiple animal species beyond rodents, including non-human primates (Robertson et al., 1999) and bats (Finkelstein et al., 2015). Notably, in bats,

the head direction system operates in 3D, similar to 3D place cells observed in bats (Yartsev & Ulanovsky, 2013).

Head direction cells constitute a fundamental component of the brain's navigation system, providing a neural basis for the internal compass that guides directional orientation (further summarised in Taube, 2007).

1.3.4 Grid cells

The defining characteristic of grid cells is their unique firing pattern: each grid cell comprises multiple firing fields arranged in an equilateral triangular lattice that spans the entire environment the animal is placed in (Figure 1.2 C) (Fyhn et al., 2004). Notably, while the firing patterns of grid cells nearby each other share orientation and scale, they differ in spatial offset. This offset, along with their grid-like arrangement allows for a precise and scalable representation of space (Chersi & Burgess, 2015).

Grid cells were first observed in the medial entorhinal cortex (mEC) by researchers in the Moser lab (Fyhn et al., 2004). This region of the brain receives projections from the pre-subiculum, where head direction cells are found (Ranck, 1985; Taube et al., 1990a), and projects directly to the dorsal hippocampus, thereby positioned as a critical hub in the spatial navigation network. Subsequent research has identified grid cells in the pre- and para-subiculum (Boccaro et al., 2010), further emphasizing their widespread role in spatial cognition. One of the intriguing aspects of grid cells is the variation in the spatial scale of their grid firing patterns across the dorsoventral axis of the mEC, with grids becoming progressively larger from dorsal to more ventral locations (Hafting et al., 2005). Thus, unlike place cells, which can undergo complete remapping of firing patterns in different environments, grid cells maintain their characteristic firing patterns across various environments, adjusting only in terms of their spatial relationship to the surroundings (Fyhn et al., 2007). The adaptability of grid cells to changes in the environment's size and shape has been documented, indicating their role in dynamically encoding spatial information (Barry et al., 2007; Stensola et al., 2012).

Furthermore, the presence of grid cells has been confirmed in multiple species, including rats (Fyhn et al., 2004), mice (Fyhn et al., 2008), bats (Yartsev et al., 2011), and humans (Doeller et al., 2010), highlighting their fundamental contribution to spatial navigation across species.

A crucial distinction between grid cells and place cells lies in the nature of their spatial representations. Grid cells provide a metric for space that encompasses non-local information, therefore theoretically allowing for the computation of distance between two points in a given environment with the help of a translational vector between them (Bush et al., 2015) (see Bicanski & Burgess, 2020 for a more detailed review of neuronal vector coding). This feature positions grid cells as integral components of the brain's navigation system, offering a geometric framework for understanding large-scale navigation and support to the cognitive mapping of space.

1.3.5 Boundary vector cells

Boundary Vector Cells (BVCs) are neurons specifically tuned to barriers at certain distances directions, exhibiting more precise tuning at closer ranges (Figure 1.2 D), and represent a crucial addition to the

spatial navigation network discussed above, particularly in their interaction with place cells (Chersi & Burgess, 2015). Place cell coding has been shown to be significantly influenced by the presence of extended boundaries within an environment. O'Keefe and Burgess (1996) observed that the firing of place cells is not only location-specific but also highly sensitive to the environmental geometry, including the proximity and orientation of boundaries like arena and maze walls. Using rectangular environments with varying dimensions, they noted that the peak firing location of a place cell typically remains constant relative to the nearest boundaries. Furthermore, many firing fields appeared elongated along the environmental axes, suggesting a spatial encoding that integrates boundary proximity. From these observations, they proposed the concept of BVCs. Their model suggests that BVCs provide critical inputs to place cells, enabling the encoding of spatial location in relation to environmental boundaries (Bicanski & Burgess, 2020; Burgess, 2008). Since then, the presence of BVCs has been identified in both the subiculum (Lever et al., 2009) and the mEC (Solstad et al., 2008), indicating their widespread involvement in the spatial navigation network (further reviewed in (Barry et al., 2006).

1.3.6 Summary

The identification and characterization of place cells, head direction cells, and grid cells, as well as BVCs, have outlined a detailed map of the brain's internal GPS system – the cognitive map (Tolman, 1948). The discovery of place cells within the hippocampus laid the foundational stone for this field, revealing how specific areas of the brain are tuned to the physical landscape of our environment. By encoding the distance and direction to boundaries, BVCs allow for a sophisticated mapping of the environment that goes beyond the specific location to include contextual geometric information. The subsequent identification of head direction cells and grid cells expanded our comprehension of the brain's spatial navigation system, highlighting a complex, yet elegantly coordinated, neural network that encodes positional data, boundary cues, and directional orientation.

Building on this rich base of knowledge, it is important to note that the cognitive map only covers the viewpoint-invariant (or, allocentric) perspective on space, and not the perspective of the animal itself, which is viewpoint-variant (or, egocentric), constantly changing as the animal explores its environment. Therefore, the open questions concerning our understanding of how the brain computes sensory and spatial information, how different reference frames are represented in the brain, and how spatial memory is stored and retrieved, has led research to look at other brain regions involved in navigation and spatial memory, such as the retrosplenial cortex (RSC).

1.4 Retrosplenial cortex

1.4.1 Anatomy and connectivity

The RSC operates as a central node in the brain's network for processing spatial and sensorimotor information. Its extensive and diverse connectivity patterns enable it to function as a hub, facilitating communication between various brain regions involved in visual processing, locomotion, navigation, and spatial memory.

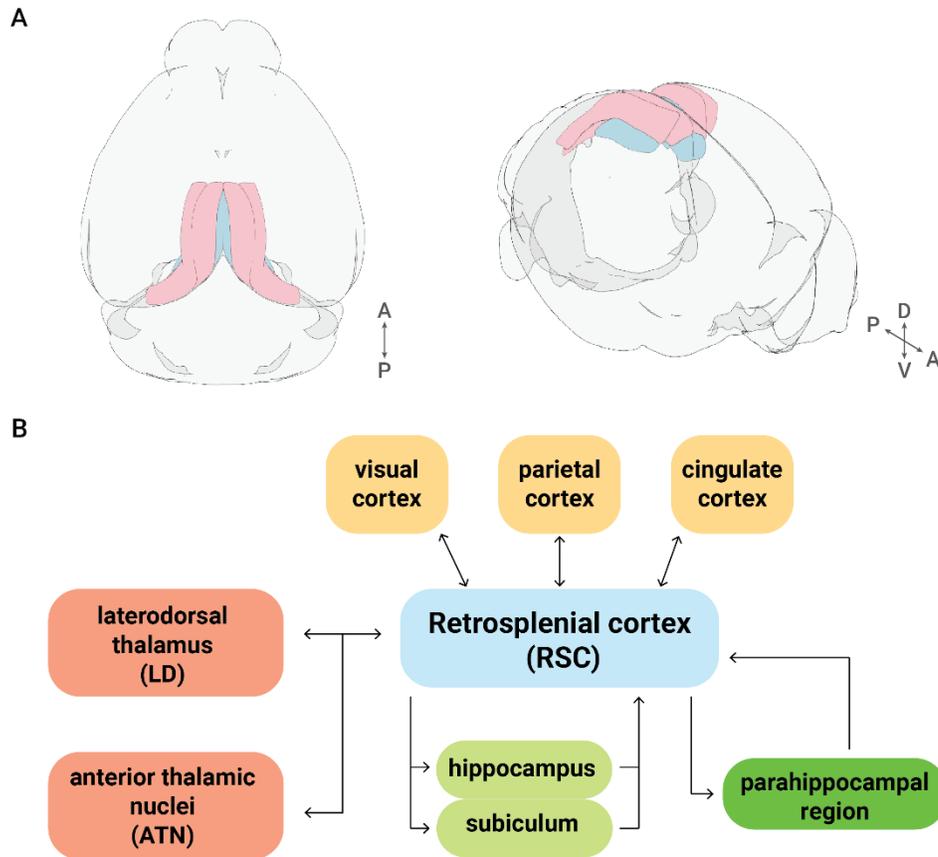


Figure 1.3 RSC anatomy and connectivity

A. Anatomical position of granular (*light blue*) and dysgranular (*light pink*) RSC within the rodent brain from the top (*left*) and the side (*right*). Schematics were generated using (Claudi et al., 2021). B. Diagram depicting gross RSC connectivity with cortical (*yellow*), thalamic (*orange*), hippocampal and subicular (*light green*), as well as parahippocampal regions (*dark green*). Adapted from (Mitchell et al., 2018).

Anatomically, in both humans and non-human primates, the RSC is situated posteriorly to the corpus callosum and aligns with Brodmann's areas 29 and 30, which together with areas 23 and 31 form part of the posterior cingulate cortex (Chrastil, 2018; Kobayashi & Amaral, 2000, 2003). This contrasts with other species such as rodents, which do not have a direct anatomical counterpart to areas 23 and 31. Instead, the entire region of area 29 and 30 is designated as RSC, spanning over half the length of the rodent cerebrum. In addition, RSC is located more dorsally and reaches the brain surface (Figure 1.3 A), making it one of the largest cortical regions in rodents (Vogt & Peters, 1981) (for a more detailed comparison of RSC anatomy across rodents and humans see Mitchell et al., 2018; Vann et al., 2009).

The RSC exhibits a dense network of unilateral and reciprocal connections to both cortical and subcortical regions in the rodent brain, integrating information from various brain regions (Figure 1.3 B). In addition, RSC's subregions are highly interconnected with each other in both primates and rodents (Kobayashi & Amaral, 2003), underscoring RSC's role as an association cortex or a neural hub, (Alexander et al., 2023). Cortically, the RSC is connected to primary and secondary visual cortices (Van Groen & Wyss, 1992a), parietal cortex (Van Groen & Wyss, 1992a), cingulate cortex (Jones & Witter, 2007), and secondary motor cortex (M2) (Yamawaki et al., 2016). It also receives unidirectional inputs from the CA1 region of the hippocampus (Chen et al., 1994b; Miyashita & Rockland, 2007) and from the subiculum (Honda & Ishizuka, 2015; Sugar et al., 2011; Wyss & Van Groen, 1992) providing an indirect link between these two regions. In addition, the RSC connects to the postrhinal cortex and mEC (Czajkowski et al., 2013; Jones et al., 2005), as well as to the general parahippocampal region. Subcortically, the RSC shares reciprocal connections with the anterior (ATN) and laterodorsal (LD) thalamic nuclei (Figure 1.3 B). The LD, in particular, receives inputs from the subiculum, visual association cortex, and lateral mamillary bodies, whereas the ATN receives predominantly inputs from the hippocampal formation and subicular complex (Van Groen & Wyss, 1990; Van Groen & Wyss, 1992a, 1992b, 2003), highlighting the RSC's integration within the broader spatial navigation circuitry. The granular region of the RSC is particularly notable for its reciprocal connections with sites containing head direction cells, such as the ATN, LD and the post-subiculum. Conversely, the dysgranular cortex exhibits stronger interconnectivity with visual areas (Vann & Aggleton, 2005), illustrating the RSC's multifaceted role in integrating spatial and visual information.

The RSC aligns with the classical description of an association cortex given its diverse and extensive connectivity, representing a critical hub by linking numerous cortical and subcortical regions (Alexander et al., 2023). The focus on navigation and spatial cognition with respect to RSC function stems from its anatomical position, acting as a bridge between neocortical regions involved in sensorimotor processing (e.g. visual and motor cortex, as well as parietal cortex), thalamic areas coding for head orientation (e.g. ATN and LD), and the extended hippocampal formation (e.g. hippocampus, entorhinal cortex, and subiculum), which is essential for spatial memory and navigation. This strategic location grants the RSC access to a variety of spatial and sensorimotor information, which is why the RSC has been suggested to potentially mediate between these diverse spatial inputs. Given RSC's anatomy and connectivity, it is widely believed that the RSC's function is more consistent with multiple sensorimotor and spatial processes, rather than any isolated function.

1.4.2 Proposed function

The RSC has attracted research interest due to its extensive connectivity and central positioning within the brain's sensorimotor and navigational circuits. Yet, varying theories surrounding the RSC's function remain. In particular, it is still unclear whether the RSC has a single, generalized role, or whether the RSC fulfils multiple, specialised functions through distinct sub-circuits that share a common modulation by RSC neuron ensembles (Alexander et al., 2023).

1.4.2.1 Navigating spaces: egocentric vs. allocentric strategy

The concept of space is multifaceted and encompasses at least two distinct dimensions: personal (or corporal) space and external space. Personal space includes the location of bodily stimuli, awareness of limb positions, and other sensorimotor feedback that informs an individual about their immediate physical context. External space, on the other hand, extends beyond the individual's body to the environment surrounding them. These dimensions of spatial awareness are crucial for navigating and interacting with the world effectively. The information gathered from both personal and external spaces is synthesized and utilized through two primary processing strategies: egocentric and allocentric (Burgess, 2006, 2008).

Egocentric strategies are based on information provided by bodily cues within a specific sensory domain, be it the retina in visual perception or whisker fields in rodents, or in relation to the individual's actions, such as turning left during a navigation task. This approach anchors the perception of space to the individual's current position and orientation, making it a subjective reference frame that moves with the observer, independent of external spatial cues.

Allocentric strategies, in contrast, offer a viewpoint-invariant encoding of environmental positions. This method relies on reference frames defined by external features, such as landmarks or boundaries, rather than the observer's sensory perception or immediate actions. Allocentric processing allows for a more global, objective understanding of space that is independent of the observer's current location or orientation.

The RSC has been suggested to play a role in translating between these processing strategies and spatial coordinate systems. This implies that the RSC could have a broader, more generalized function by transforming multiple reference frames. By mediating between egocentric and allocentric strategies, the RSC may facilitate comprehensive integration of spatial information, allowing for flexible navigation and interaction with both the corporal and external aspects of space. This integration is fundamental to the animal's ability to adapt and orient within their environment.

1.4.2.2 Egocentric coordinate frame

The egocentric coordinate frame is a system of spatial understanding and navigation that is based on the individual's own body as the central point of reference. This frame is fundamentally shaped by bodily movements and sensory processing, and is largely independent of external spatial cues, positioning all objects and locations in direct relation to the subject's position in space (Burgess, 2006). Examples of egocentric processing of space include the traditional clockwise scheme that is used to define spatial

positions around an individual and the conventional directions – left/right and forward/backward (Klatzky, 1998). The RSC has been shown to process a range of egocentric information – self-motion, egocentric distance and orientation to boundaries, as well as visual information (Alexander et al., 2023).

One of the key aspects of egocentric processing in the RSC involves the sensitivity to self-motion. Research has shown that a significant number of RSC neurons modulate their firing rate in response to linear and angular movement speeds of the animal (Alexander & Nitz, 2015; Chen et al., 1994a; Cho & Sharp, 2001), or its head (Keshavarzi et al., 2022). Based on these findings, the RSC has been attributed a role in path integration – a process where the animal tracks its own movements and displacement to calculate its position in relation to its starting point (sometimes also referred to as dead reckoning). Support for this theory comes from RSC lesion and inactivation studies demonstrating spatial deficits in darkness, where path integration is crucial for navigation (Cooper et al., 2001; Cooper & Mizumori, 1999, 2001; Elduayen & Save, 2014). Additionally, direct projections from RSC to secondary motor cortex (M2) in mice form a circuit that connects hippocampal networks, involved in navigation and spatial memory, to neocortical networks, concerning sensorimotor integration and motor control, via the RSC (Yamawaki et al., 2016).

Further support for egocentric processing in the RSC comes from egocentric boundary vector cells (EBCs) (Alexander et al., 2020; Van Wijngaarden et al., 2020). These cells become active whenever an animal is positioned at a specific distance *and direction* to an environmental boundary (e.g. the wall of an arena or maze). EBCs within the RSC have receptive fields that cover the full spectrum of orientation and distance, playing a pivotal role in mapping environmental shapes and navigating spaces. They persist in novel environments, however EBC firing ceases when physical boundaries of the environment are removed. The ability of EBCs to persist in the absence of visual or tactile (whisker-related) information suggest that EBCs can be modulated by different sensory inputs, further supporting the hypothesis that the RSC's role in spatial cognition could be to learn to associate prominent sensory inputs (such as boundaries) with spatial locations (Alexander et al., 2023; Byrne et al., 2007).

The RSC is also instrumental in integrating movement and position signals with visual information present in the environment. Experiments have shown that visual stimuli, especially those that simulate or amplify optic flow, can influence the processing of linear and angular speeds in both enhancing and suppressing manners, highlighting the RSC's role in visual processing (Clancy et al., 2019; Hennestad et al., 2021; Keshavarzi et al., 2022; Powell et al., 2020). Moreover, certain areas within the RSC display a retinotopic organization, indicating a sensitivity to the egocentric position of visual stimuli (Campbell et al., 2021; Chang et al., 2020; Fischer et al., 2020; Mao et al., 2017; Mao et al., 2020; Minderer et al., 2019; Powell et al., 2020; Zhuang et al., 2017). Such retinotopic properties are critical when it comes to landmark processing and spatial position coding during spatial memory tasks. Additionally, studies on the connections from the RSC to the visual cortex have uncovered a specific tuning to previously learnt visual stimuli, whereby these re-entrant connections played a significant part in shaping the visual cortex's responses (Makino & Komiyama, 2015). This mechanism suggests that the RSC not only processes visual and spatial information but also contributes to the learning and memory of visual stimuli.

1.4.2.3 Allocentric coordinate frame

The allocentric coordinate frame is characterised by its reliance on spatial cues and the spatial arrangement of these landmarks. When using the allocentric strategy, subjects relate their current and goal position by using landmarks in their environment as reference points, just like humans do when using a map to navigate a new city. Targets, and space in general, are thus established through coordinates that are independent of the observer's position or orientation (for example, north/south and east/west). This coordinate frame contrasts with the egocentric perspective, where spatial information is processed relative to the individual's position (Burgess, 2006). The allocentric frame is critical for understanding the broader spatial context, enabling navigation over larger distances and facilitating memory of places beyond immediate sensory perception (Benhamou & Poucet, 1995) (Klatzky, 1998). As discussed earlier (1.3 Neural mechanisms underlying spatial cognition), place cells, head direction cells, and grid cells combine to form a cognitive map that is viewpoint-invariant (i.e. allocentric). The RSC itself possesses sensitivity to allocentric space by coding for both head direction and position (Alexander et al., 2023).

Several studies have identified a subset of RSC neurons that are tuned to the heading direction of the animal relative to the environment and the layout of landmarks (Angelaki et al., 2020; Chen et al., 1994b, 1994c; Cho & Sharp, 2001; Lozano et al., 2017). Interestingly, in comparison to classical head direction cells observed in the subicular complex, which have sharp and narrowly tuned preferred directions, bi-directionally tuned head direction cells with opposing tuning directions were identified in the RSC. These bi-directional head direction firing fields emerged when animals were exposed to a multi-compartment environment with bi-directionally symmetrical landmarks, showing that head direction tuning is anchored to prominent visual landmarks. The presence of both classical and bi-directional head direction cells within the RSC indicates that RSC might translate between local and global coordinate systems (Jacob et al., 2017).

There exist neurons in the rodent RSC that code for allocentric spatial position, which have been recorded using head-fixed two-photon calcium imaging and VR mazes (Fischer et al., 2020; Mao et al., 2017; Mao et al., 2020; Mao et al., 2018). These neurons exhibit firing patterns similar to those of hippocampal place cells, displaying narrowly tuned firing fields that form a sparse, orthogonal coding system correlated with spatial location. Remarkably, RSC 'place cell' activity demonstrates resilience to changes in the environment and partial remapping, reminiscent of neurons in the hippocampal CA1 region (Mao et al., 2017). These spatial representations develop over time and are significantly impacted by hippocampal damage, which led to the hypothesis that the hippocampus might provide the RSC with an index-like, context-specific representation of space which might facilitate the coordinated retrieval of memories (Mao et al., 2018) (see also 1.4.2.4 Transformation of reference frames). Moreover, it could be shown that RSC neurons that track the animal's position along a VR track are driven by visual stimuli and active movement, as decoupling of vision and locomotion impaired spatial position coding in the RSC, indicating that RSC 'place cell' activity is anchored to specific reference points in the environment (Mao et al., 2020). While these head-fixed experiments demonstrate location-related firing patterns in RSC that closely resemble hippocampal place fields, this similarity is less clear in freely moving rodents (Miller et al., 2020; Van Wijngaarden et al., 2020). This difference may result from the distinct dynamics of calcium imaging (in head-fixed animals) versus extracellular electrophysiological recordings (in freely moving animals), and

the absence or presence of vestibular information, respectively, during spatial learning. Additionally, a commonly found confounder in linear VR tracks is that head-fixed animals can only approach locations from a certain angle, consistently receiving the identical visual input (Alexander et al., 2023). As a consequence, the animal is confronted with a viewpoint-dependent scene at any given location, which does not compare with freely behaving tasks (or real world observations) where animals can take different approach angles to the same spatial location. Further head-fixed, VR track experiments will be needed to confirm that RSC place coding is genuinely tied to an allocentric frame of reference. This could be tested by rotating distal visual cues and landmarks. Thus, despite evidence supporting an allocentric framework in the RSC, the question still remains about the conditions under which the RSC encodes spatial location within specific viewpoint-dependent or viewpoint-independent frames.

1.4.2.4 Transformation of reference frames

To summarise, two spatial perspectives – the egocentric (viewpoint-variant), which varies with the observer's viewpoint, and the allocentric (viewpoint-invariant), which remains constant regardless of the observer's position – are believed to involve RSC activity. However, the RSC is not the exclusive seat of any of these reference frames, and its role goes beyond the mere encoding of location, orientation, and direction in the service of navigation. As both egocentric and allocentric information coding takes place in the RSC, the RSC has been instead suggested as a transitioning point between these spatial reference frames (Alexander et al., 2023). An influential theory called the BBB model, named after the authors of this study, Byrne, Becker, and Burgess (2007), proposes that the RSC facilitates the transformation between the egocentric and allocentric coordinate systems, taking into account the anatomical and connectivity network surrounding the RSC (Bicanski & Burgess, 2018; Clark et al., 2018) (further reviewed in Bicanski & Burgess, 2020; Burgess, 2008).

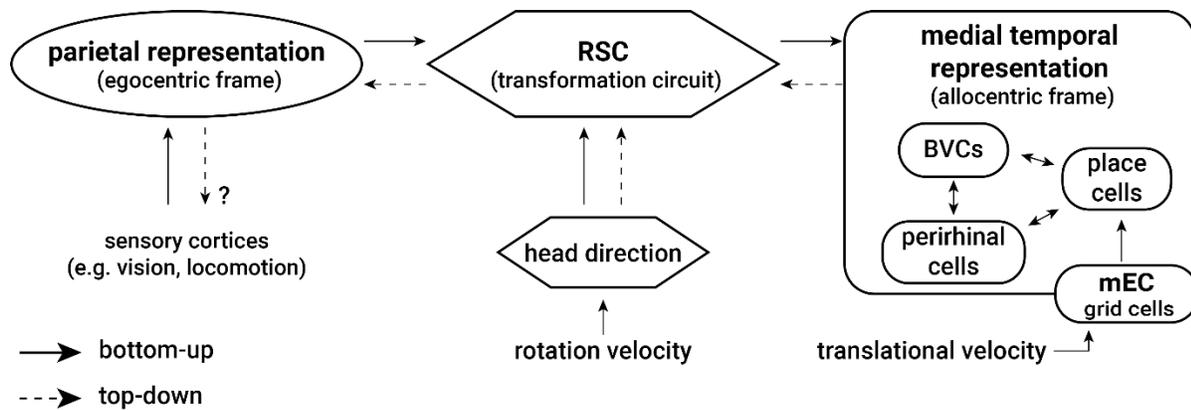


Figure 1.4 The BBB model

The BBB model proposing the RSC as the centre of reference frame transformation between egocentric and allocentric perspectives. *Solid arrows* indicate the bottom-up flow of information (from perception to memory), and *dashed arrows* the top-down flow of information (from memory to imagery and recall). *BVCs* (*boundary vector cells*), *mEC* (*medial entorhinal cortex*).

Adapted from (Byrne et al., 2007) and (Bicanski & Burgess, 2018).

More specifically, the BBB model (Figure 1.4) illustrates how during spatial perception (bottom-up flow of information), sensory inputs (such as vision and locomotion) are further processed in parietal areas to form an egocentric representation of the environment. The RSC in turn uses heading direction (originating from areas such as the subiculum (Taube et al., 1990b) and thalamus (Taube, 1995)), information about head rotation and velocity (Keshavarzi et al., 2022), as well as the parietal egocentric reference frame (Minderer et al., 2019) to map allocentric BVCs, ultimately performing a transformation from egocentric to allocentric space coding. These BVCs (found in the subiculum, (Lever et al., 2009)), subsequently support the memory formation of the allocentric reference frame in the medial temporal region, along with place cells coding for location (in the CA1 and CA3 region of the hippocampus), and perirhinal neurons representing boundary identities (e.g. texture, colour, etc. of boundaries), as suggested by the model (Bicanski & Burgess, 2018). In addition, the reciprocal connections between the entorhinal cortex and the hippocampus might allow the place cell and grid cell representations to combine both motion-related inputs (to grid cells) and environmental sensory information (the BVC inputs to place cells), which allow for the determination and characterization of the animal's current location (for more grid cell involvement in the BBB model see (Burgess et al., 2007), further reviewed in (Barry et al., 2007; Burgess, 2008; O'Keefe & Burgess, 2005)). Conversely, the top-down process enables the reconstruction of egocentric, orientation-specific representation of space from allocentric memory through the RSC.

In other words, the egocentric perception of the local sensory environment, corresponding to a specific body-centred point of view, is transformed into an allocentric map, that is viewpoint-independent, facilitating spatial learning and long-term storage of spatial memories in the medial temporal lobes. On the other hand, the RSC transforms allocentric space into corresponding egocentric perspectives, allowing

for the visualization and mental navigation of space (also referred to as imagery) and for recall of spatial memories.

Empirical evidence supporting the BBB model comes from the RSC's unique neuroanatomy and connectivity patterns (see 1.4.1 Anatomy and connectivity). The RSC maintains connections with both the parietal regions and visual cortex, which are characterized by egocentric mappings of body position and visual fields, as well as with the broader hippocampal formation, recognized for its allocentric spatial representations of location and orientation. Further support comes from EBCs, which represent an intermediate state from purely egocentric sensory input and allocentric heading direction, to allocentric BVCs. These cells have been predicted by the BBB model and have now been reported in the RSC (Alexander et al., 2020; Van Wijngaarden et al., 2020), as well as in the postrhinal cortex, parasubiculum, mEC (Gofman et al., 2019) and in the posterior parietal cortex (Alexander et al., 2022).

Moreover, neurons in the RSC exhibit spatial receptive fields that are sensitive to combinations of egocentric and allocentric features, indicating the RSC's capacity for interfacing across spatial reference frames. For example, the RSC has been shown to encode conjunctions of current route position, position in a larger environment, as well as left/right turns of rats navigating a W-shaped track (Alexander & Nitz, 2015). Some RSC neurons have been shown to exhibit periodic activation patterns that repeat across route segments that have the same shape along a plus-shaped maze, while other neurons exhibited single-cycle periodicity over the entire route, thereby providing a framework for encoding sub-route positions relative to the whole route (Alexander & Nitz, 2017). The RSC has also been shown to differentially encode environmental context (Miller et al., 2020). Chinzorig et al. (2019) observed decreased RSC ensemble activity when animals were prevented from using visual and locomotion cues, and Mao et al. (2018) demonstrated that the spatial activity sequence of RSC neurons is anchored to visual landmarks and driven by locomotion-gated optic flow, highlighting the RSC's role in integrating visual inputs, locomotion, and contextual cues to generate internal representations of space.

Additionally, RSC inactivation leads to angular drift of head direction neurons in relation to visual landmarks, underscoring the RSC's involvement in anchoring allocentric viewpoints to sensory inputs during navigation. This could be shown both in computational modelling (Bicanski & Burgess, 2016) and experimental studies, where the stability of head direction coding in the rat anterodorsal thalamus was significantly reduced following neurotoxic RSC lesions (Clark et al., 2010). Similarly, it could be shown that landmark representations in the RSC of mice navigating a linear VR track are the result of local integration of visual, motor and spatial information as the animal learns the spatial arrangement of environmental features (Fischer et al., 2020). Several RSC inactivation and lesion studies using a wide variety of different spatial navigation and memory tasks, requiring either allocentric or egocentric navigational strategies, produced a range of spatial behavioural deficits (Cooper et al., 2001; Cooper & Mizumori, 1999; Harker & Ian, 2002; Harker & Whishaw, 2004a; Hindley et al., 2014b; Keene & Bucci, 2009; Lukoyanov & Lukoyanova, 2006; Nelson, Hindley, et al., 2015; Nelson, Powell, et al., 2015; Parron & Save, 2004; Pothuizen et al., 2008; St-Laurent et al., 2009; Sutherland et al., 1988; Van Groen et al., 2004; Vann & Aggleton, 2002, 2004, 2005; Vann et al., 2003; Whishaw et al., 2001), although spatial deficits in RSC are generally more subtle compared to spatial impairments caused by hippocampal damage (further summarised in (Aggleton, 2010; Harker & Whishaw, 2004b)).

Although these research outcomes support the BBB model's suggested function of the RSC in enabling spatial reference frame transformations, the conjunctive encoding of both egocentric and allocentric representation, seen as a key indicator of the RSC's unique role in bridging spatial perspectives, has also been reported in several other brain regions. For instance, egocentric bearing and distance, in conjunction with allocentric head direction coding has been reported in the rat postrhinal cortex (LaChance et al., 2019) and posterior parietal cortex (Wilber et al., 2014), and in the mouse post-subiculum (Peyrache et al., 2017) and hippocampal CA1 (Jercog et al., 2019). In the mouse mEC, neurons code for multiple parameters of head and eye movements along with allocentric space information encoding, which could support the integration of visual features into an allocentric representation of space (Mallory et al., 2021). Thus, the RSC is likely not the single place of reference frame transformation, but could represent an addition to a wider network responsible to translate between egocentric and allocentric perspectives.

1.4.2.5 RSC's role in spatial memory

Another theory that goes hand-in-hand with the BBB model is that the RSC plays a part when it comes to the consolidation of hippocampus-dependent spatial episodic memories (Mitchell et al., 2018). According to this hypothesis, the RSC itself is therefore not implicated in the initial learning of spatial relationships but, instead, in accessing previously learnt spatial knowledge to perform tasks or integrating new information into existing memory representations. This concept is based on two other foundational theories in the field: Marr's proposition that (spatial) memories established in the hippocampus undergo gradual consolidation within the neocortex (Marr et al., 1991), and the notion that spatial learning constructs not only specific task-related memories but also a broader cognitive schema within which these memories are embedded (Ghosh & Gilboa, 2014; Morris, 2006). For instance, in a water maze task, an animal's accelerated learning of a new platform location can be attributed to its prior knowledge of the maze's layout, requiring only minimal updates to its spatial schema (further summarized by (Mitchell et al., 2018)).

Support for these ideas comes from studies employing chronic *in vivo* two-photon imaging and immediate early gene (IEG) mapping (a proxy for neuronal activity), showing that specific RSC neuronal ensembles are reactivated after mice are trained on the Morris water maze. Notably, RSC ensemble activity was heightened when animals relied on distal visual cues to navigate to the previously learnt location of the hidden platform, supporting the hypothesis of RSC's involvement in integrating allocentric strategies with sensory inputs during spatial learning (Czajkowski et al., 2014). Other studies also found increased IEG expression in RSC after radial arm maze acquisition (Amin et al., 2006; Czajkowski et al., 2020; Pothuizen et al., 2009; Vann et al., 2000). Experiments employing optogenetic reactivation of RSC ensembles could evoke previously learnt fear responses (Cowansage et al., 2014), highlighting RSC's involvement in recalling spatial memories. While RSC neuronal ensembles have also been identified in other studies (Czajkowski et al., 2020), it could be shown that these ensembles representing space, and sometimes even goal locations, develop with learning (Miller et al., 2019; Vedder et al., 2016), stabilize with learning, persist over time and that the retention of memory correlates with ensemble stability (Milczarek et al., 2018), underscoring a mechanism for memory consolidation within the RSC. Furthermore, studies tracking dendritic spine dynamics in the RSC have found that spine turnover prior to learning can predict future learning and memory performance (Frank et al. 2018). Lastly, it could be shown that optogenetic reactivation of RSC ensemble neurons during sleep can produce a recent memory with features normally

observed in consolidated remote memories, including higher engagement of neocortical areas during retrieval, contextual generalization, and decreased hippocampal dependence (de Sousa et al., 2019).

Together, these findings underscore the RSC's critical function in spatial learning and memory, next to its capacity to bridge egocentric and allocentric information. However, its precise role in spatial memory acquisition and recall has not been elucidated yet.

1.5 Aim of this study

The RSC integrates both sensory and spatial signals and plays a central role in navigation, as well as spatial learning and memory. To investigate RSC's role during spatial learning, memory retrieval and reversal learning in freely behaving animals, I aimed to establish a novel behavioural paradigm, during which mice acquire spatial memory, recall it later in training and perform reversal learning. I set out to obtain a comprehensive dataset of behavioural tracking, which would allow for a detailed analysis of task performance. I performed chemogenetic experiments to establish RSC's role within this task, and a series of landmark manipulations to establish the relative egocentric vs. allocentric nature of the acquired spatial memory.

2. Methods

2.1 Animals

All experiments were conducted in accordance with the institutional guidelines of the Max Planck Society and the local government ethics committee (Beratende Ethikkommission, Regierung von Oberbayern).

In total, 22 adult mice (14 female, 8 male; two to four months of age at the start of the experiment) were bred in-house at the MPI of Biological Intelligence to genetically express the calcium indicator GCaMP6s in excitatory neurons (B6;DBA-Tg(tetO-GCaMP6s) 2Niell/J23, JAX stock #024742; back-crossed for at least seven generations to C57Bl/6NRj) and crossed with B6.Cg-Tg(Camk2a-tTA)1Mmay/DboJ54 (JAX stock #007004; maintained on a mixed background of C57BL/6NRj and C57BL/6J). Mice were kept in standard type II cages (530 cm²) on a reversed 12-hour light/dark cycle. Whenever possible, mice of the same gender were co-housed (two to four mice per cage, depending on litter size and mouse availability), or single-housed if they showed signs of aggression towards co-housed animals. Cages were filled with wooden chips as bedding, and wooden shavings as nesting material. Each cage was equipped with a little red plastic house serving as both shelter and nest for mice, a running wheel, and a red handling tube. Before the start of the behavioural training and food deprivation, mice were fed *ad libitum* with standardized food pellets. Mice had a constant water supply at any given time point in their cages.

2.2 Cohorts and experimental timeline

In this study, a total of two cohorts were trained: 14 mice were used for the chemogenetic experiments (9 female and 5 male animals), and 8 mice (5 female and 3 male animals) made up the behavioural cohort. Animals of both cohorts were trained for a maximum of 42 consecutive days under food restriction to learn the location of three hidden trigger locations: SM1, SM2, and SM3 (spatial memory 1, 2, and 3).

Animals in the behavioural cohort (Figure 2.1 A) underwent a surgery, in which a head bar was implanted (see Methods – 2.3 Surgery). Following the surgery, animals were allowed to recover (two weeks) before the habituation to handling by the experimenter and to the behavioural arena started (1 week) (see Methods – 2.4.11 Habituation), after which access to food was restricted and the behavioural pre-training period started (see Methods – 2.4.12 Pre-training). During the spatial memory training (see Methods – 2.4.13 Spatial memory training), mice learn consecutively the location of three hidden-fixed trigger locations (see Methods – 2.4.8 Session structure). Animals of the behavioural cohort were additionally trained and tested on a set of behavioural training protocols, where either the arena cues or trigger locations were rotated, or animals were presented with a new, randomly placed, hidden trigger location for five consecutive days.

For the chemogenetic experiments (Figure 2.1 B), a cranial window was implanted and inhibitory DREADDs were expressed in the RSC (see Methods – 2.3 Surgery). A Ca²⁺-imaging session (see Methods – 2.7 *In vivo* two-photon Ca²⁺ imaging) was scheduled before the start of the behavioural pre- and spatial memory training. A subset of animals received CNO injections before the imaging session and the start of the training session.

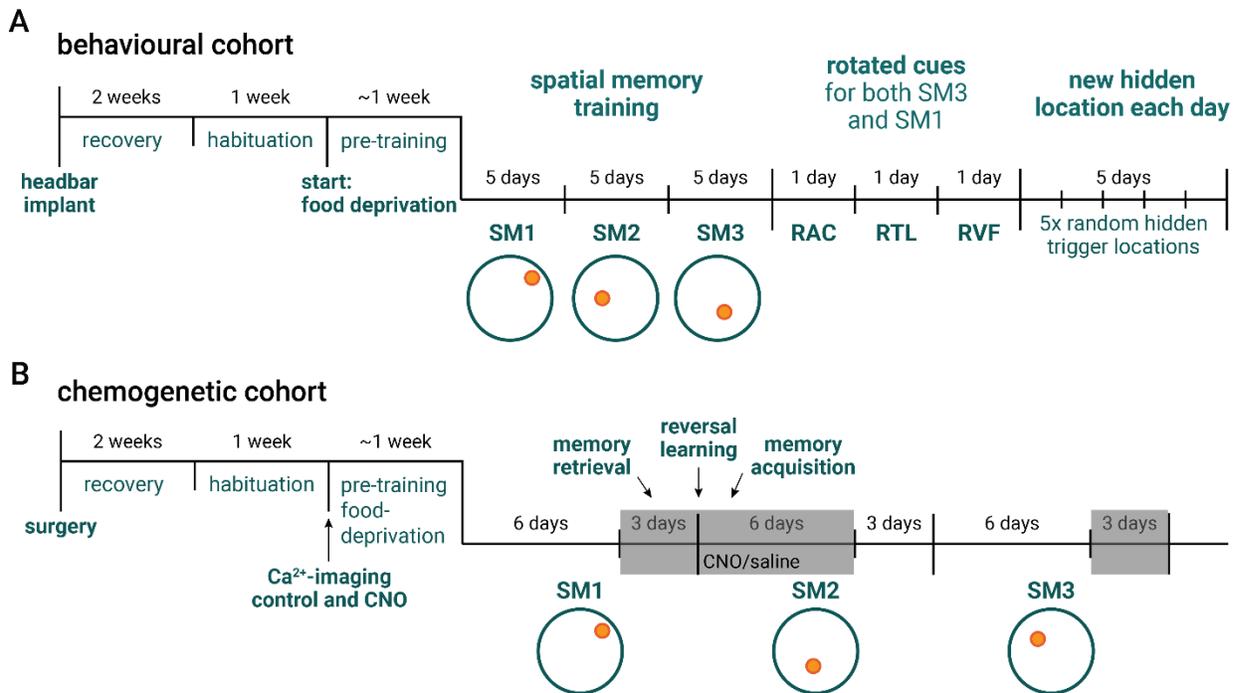


Figure 2.1 Experimental timeline

Experimental timelines for the behavioural (A) and chemogenetic cohort (B), depicting the time point of surgery, recovery and habituation period, start of the food restriction and the behavioural training (pre- and spatial memory training), during which mice learn the location of three hidden-fixed trigger locations (SM1, SM2, and SM3, for spatial memory 1, 2, and 3). A. Animals of the behavioural cohort were additionally trained on a set of experiments where either the arena cues (RAC), trigger locations (RTL), or the visible-fixed trigger locations (RVF) were rotated. Animals were also presented with a new hidden-fixed trigger location for five consecutive days. B. Animals of the chemogenetic cohort underwent a Ca^{2+} -imaging session prior to the start of the behavioural training. On selected days, animals received either intraperitoneal saline or CNO injections (*grey shaded days*) before the behavioural training started to investigate the effect of chemogenetic inhibition of RSC activity on memory acquisition, retrieval, and on reversal learning.

2.3 Surgery

For surgical procedures, mice were anaesthetised with a mixture of fentanyl (0.05 mg kg^{-1} body weight), midazolam (5.0 mg kg^{-1} body weight) and medetomidine (0.5 mg kg^{-1} body weight) in saline, injected intraperitoneally. An analgesic (carprofen, 5.0 mg kg^{-1} body weight) and dexamethasone ($2 \mu\text{g kg}^{-1}$ body weight) were injected subcutaneously after anaesthesia onset, but prior to the start of the surgery. Throughout the surgery, mice were kept at a constant body temperature of $38 \text{ }^{\circ}\text{C}$ and ophthalmic ointment (IsoptoMax/Bepanthen) was applied onto the eyes to protect the cornea from drying out. The

depth of anaesthesia was monitored by observation of the breathing rate and confirmed by regular testing for the absence of a pedal reflex.

The animal was put in a stereotaxic apparatus (Neurostar). After the skin above the skull was disinfected and a topical analgesic (lidocaine (Aspen Pharma), 0.2 mg ml⁻¹) was applied on and underneath the scalp, a rostrocaudal incision was made and the skin was removed above the midline. The skin surrounding the exposed area was adhered to the skull using Histoacryl (B. Braun Surgical). The skull was subsequently dried and scraped with a scalpel to improve later adherence of the head bar to the bone.

For animals of the behavioural cohort, a custom-designed, titanium head bar (approx. 3.0 cm x 1.0 cm) was centred and fixed onto the skull with cyanoacrylate glue (Ultra Gel Matic, Pattex). Dental cement (Paladur) was used to fixate the head bar in place and to cover the remaining exposed skull areas.

In animals of the DREADD cohort, the exposed bregma and lambda positions were used as landmarks to locate the RSC region (see Figure 2.2). A circular craniotomy (4.0 mm in diameter) was performed centred over RSC (at ML 0 mm and approximately 1.0 mm caudal to bregma) using a dental drill.

For chemogenetic inactivation of RSC activity, a viral vector mixture of AAV2/1:CamKII0.4-Cre(2.1·10¹¹ GC ml⁻¹) and AAV2/9:hSyn-DIOhM4D(Gi)-mCherry (2.3·10¹² GC ml⁻¹) was bilaterally injected at six locations at two different depths to target both granular and dysgranular RSC, resulting in a total of twelve injection sites (coordinates: AP -1.58 mm, -2.46 mm and -3.16 mm to bregma, ML +/- 0.3 mm, +/- 0.3 mm, +/- 0.5 mm, respectively, at DV depths of 0.6/1.0 mm, 0.5/1.0 mm, 0.75/1.24 mm, respectively; 120 nl of virus solution per injection site at an injection rate of 40 nl min⁻¹; Nanoject, Neurostar) (Figure 2.2). Each viral vector injection was flanked by a 3-5 min pre and post-injection period. The craniotomy was subsequently covered with a circular cover glass (4.0 mm), which was fixated first with cyanoacrylate glue (Ultra Gel Matic, Pattex) and then secured with dental acrylic (Paladur). A head bar was centred rostral to the craniotomy and fixed onto the skull with cyanoacrylate glue (Ultra Gel Matic, Pattex). Dental cement (Paladur) was used to fixate the head bar in place and to cover the remaining exposed skull areas.

Animals of both cohorts received a subcutaneous injection of an antagonist mixture of naloxone (1.2 mg kg⁻¹ body weight), flumazenil (0.5 mg kg⁻¹ body weight) and atipamezole (2.5 mg kg⁻¹ body weight) in saline at the end of the surgery. Mice were placed under a heat lamp during the recovery period. Animals received postoperative analgesia (carprofen, 5.0 mg kg⁻¹ body weight, injected subcutaneously) for at least two subsequent days following the surgery. In a subset of mice, a second surgery was performed prior to the Ca²⁺-imaging session to remove small patches of re-grown bone under the cranial window.

All animals of the chemogenetic cohort underwent this surgery protocol prior to the Ca²⁺-imaging session and the behavioural training protocol: cranial window placement for visual access for Ca²⁺-imaging; intracranial virus injection for inhibitory DREADD expression in RSC; head bar implantation for head-fixation during the Ca²⁺-imaging session and as a baseplate for head markers for the position tracking during the behavioural training. For animals of the behavioural cohort, solely a head bar was implanted allowing for the identification of individual mice and serving as a baseplate to which the head markers for the position tracking during the behavioural tracking can be attached to.

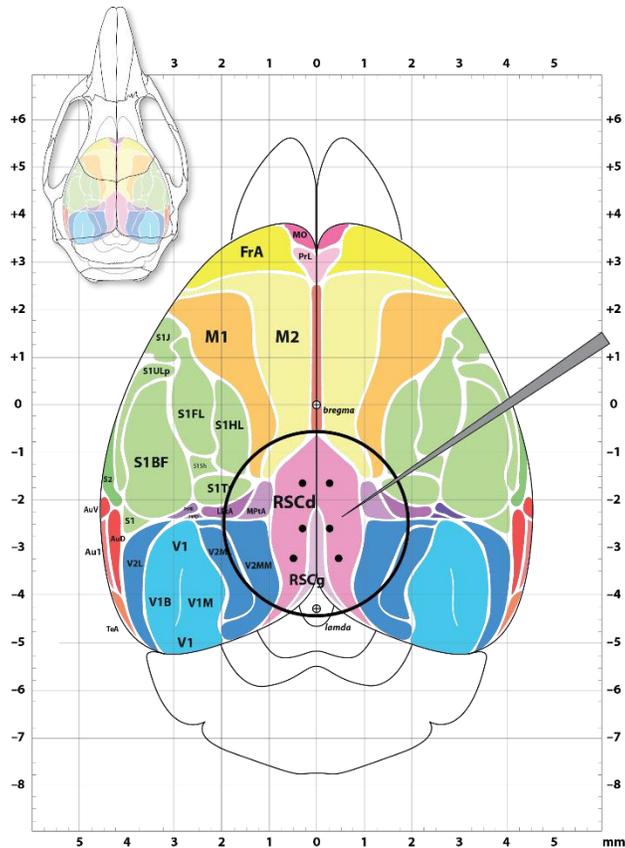


Figure 2.2 Cranial window placement and virus injection sites

Schematic depicting anatomical location of mouse dysgranular (RSCd) and granular (RSCg) retrosplenial cortex. *Black circle* depicts the position of the cranial window (4.0 mm in diameter) in relation to bregma, used as an anatomical landmark. *Black dots* show the approximate location of the virus injection sites at -1.58 mm, -2.46 mm and -3.16 caudal to bregma along the anterior-posterior axis, and ± 0.3 mm, ± 0.3 mm and ± 0.5 mm along the medial-lateral axis, respectively. The virus solution was injected at different cortical depths (0.6/1.0 mm, 0.5/1.0 mm and 0.75 /1.25 mm, respectively starting anterior and moving posteriorly) to target both dysgranular (RSCd) and granular (RSCg) RSC. Adapted from (Watson et al., 2011)

2.4 Spatial memory task

2.4.1 Arena and reward ports

I developed a spatial memory task, in which mice learn to remember hidden trigger locations. This task requires a setup consisting of a round arena (80 cm in diameter) surrounded by 50 cm high walls (Figure 2.3). There are eight equidistant reward ports (Sanworks) embedded in the arena wall, where each port

comprises a light gate, an LED light, and a cannula to deliver a liquid reward, held together by a 3D-printed housing. The light gate in each port ensures the accurate detection of nose pokes, while the LED light serves as a visual signal for the mice, indicating the availability of a reward at this given port. Outside the arena, cannulas are connected to pinch valves via a rubber tubing system and a syringe holding the reward liquid to allow for the precise and controlled delivery of liquid rewards to the mice. The light gate, the LED, as well as the pinch valves are controlled by a system (Bpod, Sanswork) that is interlocked with the position tracking software (OptiTrack, NaturalPoint, Inc.) via a closed-loop.

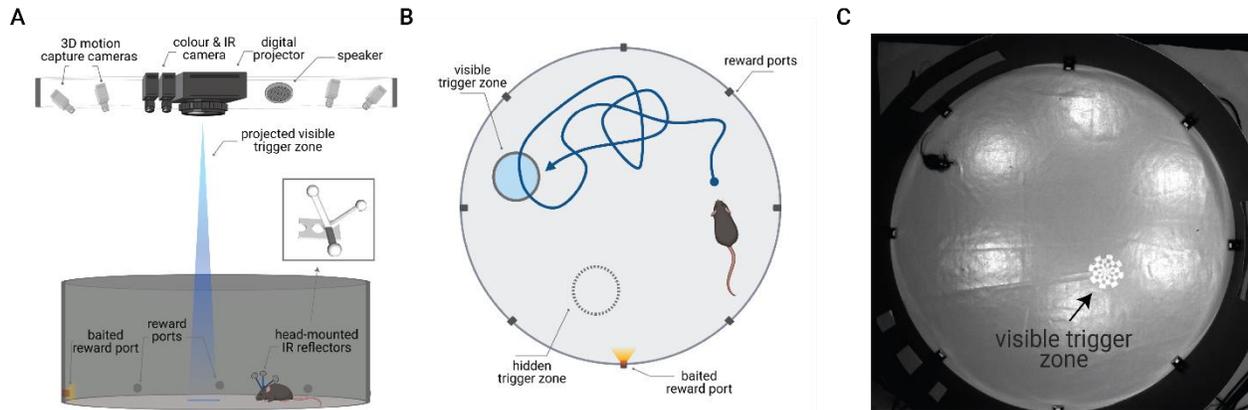


Figure 2.3 Behavioural setup

A. Schematic of the behavioural training setup for the spatial memory task. Four to six infrared motion capture cameras (OptiTrack System, NaturalPoint Inc.) track the animal's location in 3D with the help of head-mounted IR reflectors. A digital projector mounted to the ceiling is used to shine visible trigger zones (8 cm in diameter) into the arena (80 cm in diameter). Two cameras – a colour and an IR camera – to record the training sessions, as well as a speaker are mounted to the ceiling. Reward ports are inserted into the arena wall, and the LED light of each port can be turned on individually to indicate the availability of reward at this port to the mouse. B. Top view schematic of a mouse in the training arena showing an example trajectory (*dark blue*), a visible (*blue*) and a hidden (*grey, dashed*) trigger zone. Eight reward ports are distributed equidistantly. C. Screenshot of a training session recorded with an infrared camera showing a mouse in training and a visible trigger zone.

2.4.2 Arena cues

For visual guidance and spatial orientation, mice are provided with six visual cues in the arena during behavioural training - two blank (white) cue cards, two cue cards with horizontal stripes, and two cue cards with vertical stripes (Figure 2.4). These cues are provided at both proximal (one of each cue card type attached to the inside of the arena wall) and distal (one of each cue card type suspended from the arena ceiling) distances to the arena centre.

During specific training conditions and stages (see Figure 2.1 A – rotated cues), arena cues were rotated 120° degrees clockwise (Figure 2.4).

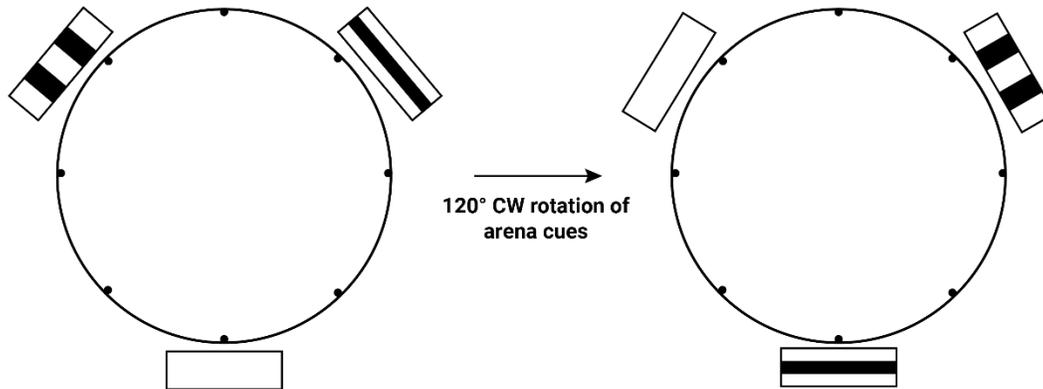


Figure 2.4 Arena cue rotation

Schematic depicting the clockwise (CW) rotation (120° degree) of the arena cues. Both proximal (attached to the inside of the arena wall) and distal (suspended from the arena ceiling) cue cards were rotated in the same direction. For simplicity, only the distal arena cues are shown.

2.4.3 Liquid reward

Throughout the spatial memory task, upon completion of a correct trial, mice were rewarded with either sweetened, infant formula soy milk (SMA Wysoy) or a diluted high-caloric, unflavoured, soy-based drink (Fresubin). Mice of the same cohort received the same type of liquid reward drink throughout the entire training protocol. The choice of sweetened soy milk and the high-caloric soy-based drink as liquid rewards was based on a previous study that showed the effectiveness of these reward types in motivating mice (Goltstein et al., 2018). By consistently providing the same type of liquid reward to mice in training, I aimed to minimise any potential confounding factors and ensure that any observed differences in behaviour were solely due to the experimental manipulations.

2.4.4 Position tracking

Before the start of each training session, a custom-designed and 3D-printed, three-armed IR reflector (see Figure 2.3 A – head marker inset) weighing less than 1.0 g was attached to the animal's head-mounted baseplate with an easily removable, mouldable adhesive gum (patafix, Uhu). These head markers reflect infrared light emitted by four to six motion capture cameras (OptiTrack, NaturalPoint, Inc.), allowing for the tracking of the mouse's position and rotation in a three-dimensional space. This tracking system has an error rate of less than 1 mm ensuring reliable and accurate collection of mouse coordinates in 3D during experimental sessions.

2.4.5 Arena ceiling

Motion capture cameras (six cameras for the chemogenetic cohort, four cameras for the behavioural cohort, for mouse position and rotation tracking), a projector (calibrated to fit the arena's size and projection zone to display visible trigger zones into the arena), a speaker (to aid the learning process and task progression with auditory cues), a colour as well as an IR camera (to record the training session) are attached to the arena ceiling (Figure 2.3).

2.4.6 Single-trial structure

In the arena, visible trigger zones are displayed with the help of a ceiling-mounted projector (Figure 2.3 B-C). The motion capture cameras detect when the mouse moves through these visible trigger zones, which triggers a sound emitted from the speaker. This short (300 ms long) auditory cue indicates to the mouse that it has successfully crossed the visible (or hidden, i.e. not displayed by the projector) trigger zone. Position tracking is interlinked with a system controlling the reward ports. When the mouse passes a visible or hidden trigger zone, the LED light of one of the eight reward ports turns on to indicate the baited reward port. During the pre-training phase, mice have 30 seconds to find the baited reward port, while they have 15 seconds in the spatial memory training phase. A nose poke, detected by the IR light gate, confirms the choice of the correct reward port, a further short (300 ms long) sound is played, and approximately 5 μ l of liquid reward are delivered via the cannula. Immediately after reward delivery, a new trial with the next trigger location starts.

2.4.7 Error trials

There are three types of error trials: 'wrong port' (WP) error, 'trigger, no port' (TNP) error, 'no trigger' (NT) error. In a WP trial, the mouse fails to request the reward from the correct, baited reward port with the LED light on after having traversed the trigger zone, but instead a nose poke is detected in an unbaited, LED-off reward port. This mistake is indicated to the mouse by a short (300 ms long) noise sound before the next trial starts. The purpose of emitting a noise sound for 300 milliseconds after an error trial is to provide a clear indication to the mouse that it has made a mistake and to help reinforce the association between the LED light at the correct, baited port and the reward availability. Additionally, starting the next trial immediately after an error trial allows for continuous training and learning without unnecessary delays. In TNP trials, the mouse is able to locate the visible or hidden trigger zone, however, fails to retrieve the reward from the baited port within a response time of 15 seconds (or 30 seconds in the pre-training phase). The reward sound is not played and the next trial starts right away. The last error trial type - NT trial - can only be made during the pre-training phase where all trigger zones are still visible and when the mouse is not able to locate the trigger zone in a time window of 2 minutes. In this case, a new trial starts with a new visible trigger zone.

2.4.8 Session structure

Mice were trained once per day in a 20-minute training session. For the chemogenetic cohort, a training session was structured in a three-trial structure, whereby every third trial had a trigger that was fixed in location but invisible to the mouse (i.e. hidden-fixed trigger). These trials were interleaved with randomly placed triggers, made visible to the mouse with the help of a light projection (i.e. visible-random trigger)

(Figure 2.5 A). For the behavioural cohort, a four-trial structure was adopted. The second trial of a given training session had a trigger that was both fixed in location and visible to the mouse (i.e. visible-fixed trigger), which was repeated in every fourth trial following this one. Starting with the fourth trial of this training session, every fourth trial had a hidden-fixed trigger. In the remaining, interleaving trials the trigger zone was visible to the mouse and placed randomly in the arena (i.e. visible-random trigger) (Figure 2.5 B).

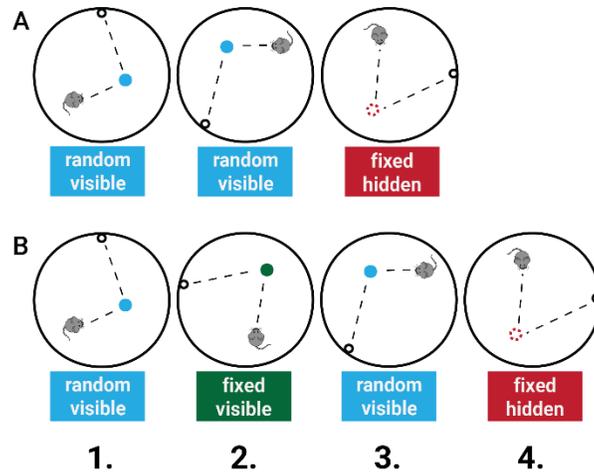


Figure 2.5 Session structure for the chemogenetic and behavioural cohort

Schematic depicting the trial types, trial sequence and order for the two cohorts. A. For the chemogenetic cohort, in every third trial, the trigger location was fixed in position and invisible (hidden) to the mouse (hidden-fixed). The trigger locations in the remaining trials were visible to the mouse and randomly placed (visible-random). B. For the behavioural cohort, every fourth trial starting with the second trial had a visible-fixed trigger, and every fourth trial starting with the fourth trial had a hidden-fixed trigger. The remaining, trigger locations were visible and randomly placed.

2.4.9 Reward port selection

The rewarded port in a given trial is chosen in such a way that mice need to travel at least 40 cm (i.e. the arena radius) to reach the baited port. First, the distances between the trigger location (visible or hidden) and all eight reward ports are calculated. The two ports whose distance to the trigger location is closest to the arena radius (40 cm) are selected, of which one is randomly chosen as the current trial's baited reward port whose LED light is then turned on (Figure 2.6). This means, that for a given fixed trigger position, such is the case for hidden-fixed and visible-fixed triggers, only two out of the eight reward ports are potential baited reward ports.

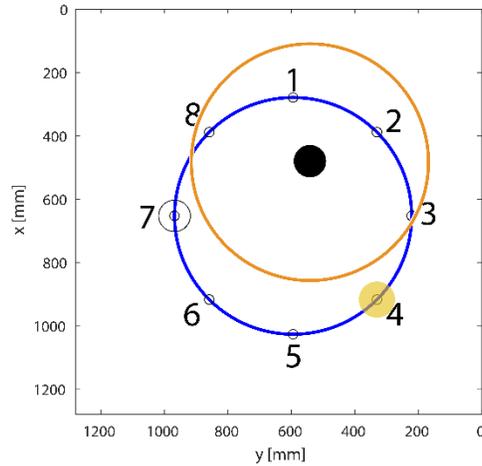


Figure 2.6 Reward port selection

Schematic depicting the logic behind the reward port selection. The *blue circle* depicts the arena boundaries with the embedded reward ports labelled 1-8. The *black, filled dot* depicts the trigger position of an example trial. After disregarding all ports with a distance shorter than the arena radius (40 cm) to the trigger location (all ports that lie within the *orange circle*), the two reward ports closest to the trigger location are selected as potential reward ports (depicted by the *circle* around port 4 and 7), of which one of them is randomly selected as the baited reward port of this given trial (depicted by the *yellow, filled circle* around port 4).

2.4.10 Trigger placement

For the placement of the visible-fixed and hidden-fixed triggers, the arena was divided into tertiles. Each of the three hidden-fixed trigger locations of the three spatial memory training phases SM1, SM2, and SM3, (and visible-fixed trigger locations in the behavioural cohort) were placed in a different tertile, whereby each trigger location in each tertile had a different distance to the arena wall: 12 cm, 28 cm and 20 cm for SM1, SM2 and SM3 (Figure 2.7 A). For the behavioural cohort, a visible-fixed and hidden-fixed trigger could not be placed in the same arena tertile for a given training phase (Figure 2.7 B).

To eliminate the possibility that mice would leave olfactory cues for subsequent animals, the arena floor was wiped after each training session with a mild, skin-friendly cleanser, and the trigger placement pattern was rotated for each animal individually by a certain rotation factor (an example trigger placement with a rotation factor of 60° degrees is shown 2.7).

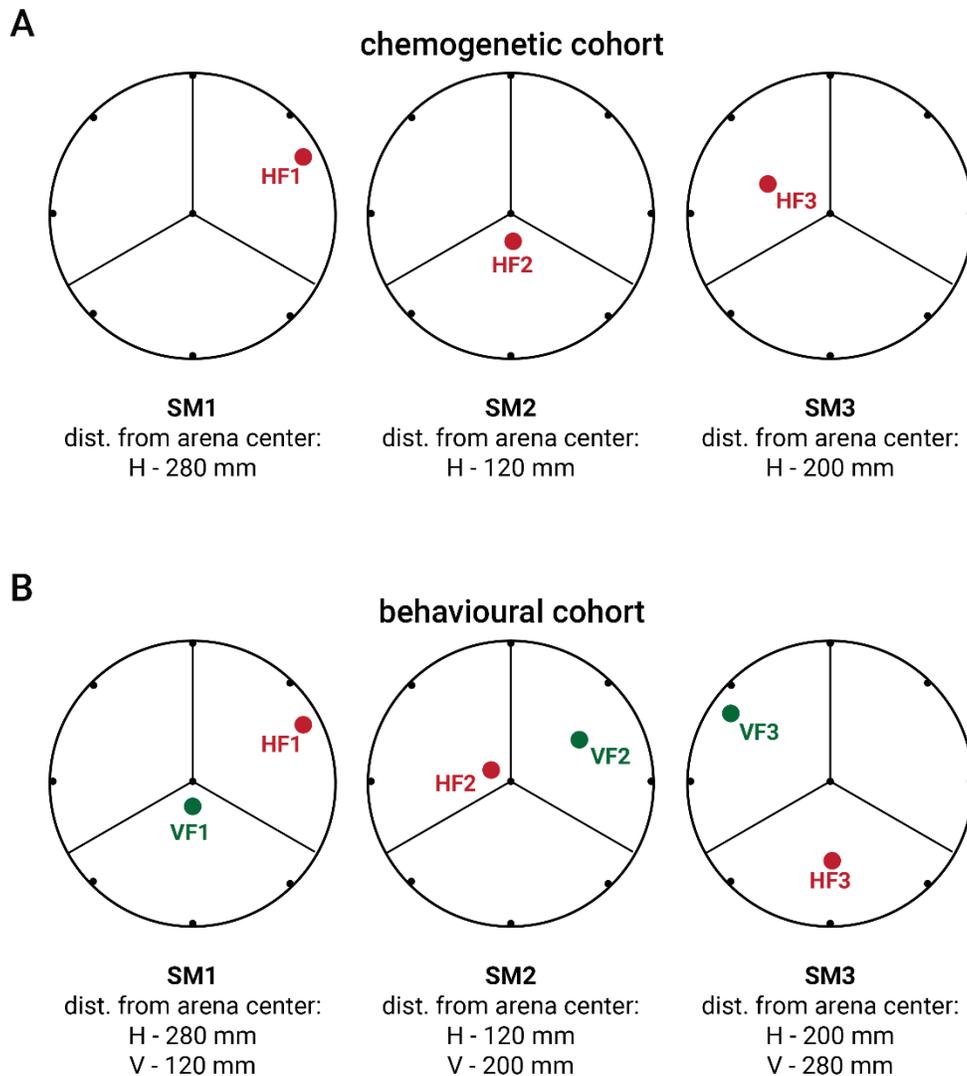


Figure 2.7 Placement of fixed trigger locations

Schematic depicting the placement of the fixed trigger locations in the arena for the three spatial memory training phases (SM1, SM2 and SM3) with an example rotation factor of 60° degrees, calculated as the angle between a norm vector [0,1] and a vector spanning from the arena center to the first hidden-fixed trigger (HF1). A. Placement of the hidden-fixed trigger locations HF1, HF2 and HF3 for animals of the chemogenetic cohort. Each hidden-fixed trigger location has a different distance to the arena centre (120 mm, 200 mm, or 280 mm). B. Placement of the hidden-fixed (HF) and visible-fixed (VF) trigger locations for animals of the behavioural cohort. Each hidden-fixed and visible-fixed trigger location has a different distance to the arena centre (120 mm, 200 mm, or 280 mm), while for each spatial memory training phase (SM1, SM2, and SM3), the hidden-fixed and visible-fixed trigger cannot be placed in the same arena tertile.

2.4.11 Habituation

Animals of both cohorts underwent a two-week long recovery period after the surgery. Subsequently, mice were habituated to handling by the experimenter and to the arena over a time course of approximately one week (Figure 2.1). Each habituation session lasted at least ten minutes during which a mouse was held in the experimenter's gloved hands and then placed into the arena. For the chemogenetic cohort, animals were additionally habituated to head-fixation underneath the two-photon microscope setup. Animals were first placed on the Styrofoam ball, then accustomed to short head fixation by holding the head bar manually for a few seconds, before they were then placed into the head-fixation system of the imaging setup for five to ten minutes (for 3-5 consecutive days).

2.4.12 Pre-training

In the first week of behavioural training and with food restriction onset, mice underwent a pre-training phase (Figure 2.1). This pre-training phase is designed to familiarize the mice with the task structure and the reward ports. It habituates animals to the following task specifics: passing through visible trigger locations, familiarization to auditory cues (upon trigger entry and correct/incorrect port selection), familiarization with light cues (LED light of baited port), and familiarization with operating the reward ports (nose-poke to release the liquid reward). To allow animals to learn the association between the trigger zone and the corresponding LED light indicating the baited reward port, the pre-training phase is split into two protocols: eight-port training (4-10 days) and one-port training (1-2 days).

During the eight-port training, initially, all ports are baited, meaning that the LED lights of all eight reward ports turn on after the first trigger has been traversed. Thus, the mouse can choose from which port to retrieve a reward. Once the mouse has visited a port, this port will not be active in the next trial. Therefore, with each new trial and reward retrieval, one reward port less is baited with reward and the chance to randomly, independently of the light find a correct, baited reward port decreases. Once all eight reward ports have been correctly visited and a reward has been retrieved, the training starts anew and all eight ports are baited in the next trial. Mice spend at least four days training on the eight-port training.

Once a mouse has successfully completed at least 35 correct trials that resulted in reward delivery and consumption in a given training session, the mouse moves on to one-port training. In this training stage, only the LED light of a single baited reward port (see reward port selection) turns on after a successful trigger zone visit, which at this stage are all visible. Mice were trained for one to two days on this training stage.

2.4.13 Spatial memory training

After the pre-training phase, mice move on to the spatial memory training phase (Figure 2.1), where every third or fourth trial has a hidden-fixed (invisible to the mouse) trigger zone for the chemogenetic and behavioural cohort, respectively (Figure 2.5). The location of these hidden-fixed trigger zones must be learnt and remembered by the animals. All mice were trained on three different hidden-fixed trigger zones: SM1, SM2, and SM3.

2.4.14 Dropout criteria

All mice that dropped below 35 trials per daily training session following the pre-training phase, and thus did not show stable behavioural baseline performance, were excluded in the analysis. In the chemogenetic cohort, one mouse was excluded from the analysis, while none of the animals from the behavioural cohort met the dropout criterium.

2.5 Food deprivation

For food restriction, a reference baseline weight was defined as the average weight of an animal during the one-week-long habituation period. During this time mice were weighed daily and had constant *ad libitum* access to food. With the start of the pre-training, the food restriction was started, whereby mice were maintained at 85-90% of their reference weight. Each day after the behavioural training session, mice received an individually measured amount of food (chow) in addition to the soy milk received during the behavioural task. Food restriction lasted for the duration of the behavioural training protocol, but no longer than 42 consecutive days.

2.6 Chemogenetic inactivation

For chemogenetic inactivation experiments, animals of the chemogenetic cohort were randomly assigned to either the control or experimental group. All animals of both groups received an intraperitoneal injection of clozapine-*N*-oxide (CNO; 10 mg kg⁻¹ bodyweight) 45 minutes prior to the Ca²⁺-imaging session to measure the effect of CNO-mediated inactivation of RSC cell activity. During behavioural training, animals of the control group received saline injections intraperitoneally 45 minutes before the start of the behavioural training, while mice of the experimental group received intraperitoneal injections of CNO (10 mg kg⁻¹ bodyweight) 45 minutes before start of the behavioural training. There were in total 12 days on which animals received injections prior to behavioural training (Figure 2.1 B).

2.7 *In vivo* two-photon Ca²⁺ imaging

During the imaging session of the chemogenetic cohort prior to the start of the behavioural training (Figure 2.1 B), animals were securely positioned in a head-fixed arrangement on top of an air-floating Styrofoam ball, enabling the mice to move and run beneath the two-photon imaging system (Thorlabs Bergamo II). The imaging setup featured a pulsed femtosecond Ti:Sapphire laser (Spectra Physics MaiTai DeepSee laser), tuned to 960 nm for calcium imaging or 1040 nm for structural imaging. Setup components included resonant and galvo scanning mirrors, a 16x NA 0.8 immersion objective (Nikon), and a piezoelectric stepper for multiplane imaging (4 planes 20µm apart were imaged in each recording session). The photon collection pathway incorporated a 720/25 nm short-pass filter followed by a dichroic beam-splitter, facilitating simultaneous detection of green and red light using two GaAsP photomultiplier tubes (PMTs; Hamamatsu) equipped with either a 500-550 nm or a 572-642 nm bandpass filter. Control of the system was executed through ScanImage 4.27 (Pologruto et al., 2003). Ultrasound gel, which was diluted with water (3:1) and centrifuged to eliminate air bubbles, was applied on top of the cranial window, and the objective was submerged in the gel. The imaging window was shielded from external light using strips of tape, and all recordings were conducted in the dark.

2.8 Histology and immunohistochemistry

Subsequent to the last training session, animals of the chemogenetic cohort were anaesthetised and transcardially perfused, first, with saline containing lidocaine and heparin (5 mg ml⁻¹ and 2.8 mg ml⁻¹, respectively) and, second, with 4% paraformaldehyde (PFA) in PBS. Brains were post-fixed in 4% PFA in PBS for at least 72 hours, and subsequently in 30% sucrose for at least 72 hours, at 4 °C. Using a microtome, brains were sectioned into 50 µm coronal slices. Slices were stained for mCherry (chicken anti-mCherry [1:1000], followed by anti-chicken Alexa 568) and mounted with medium containing DAPI.

Fluorescent signals from expressed hM4D(Gi)-mCherry, immunohistochemically amplified signals (Alexa 568), as well as DAPI stain signals were acquired with a ZEISS Axio Scan.Z1 slide scanner microscope. Using a confocal microscope (Leica Stellaris, 10x NA 0.4 air objective), four optical sections spaced 3 µm were acquired with a resolution of 550 x 550 µm. The different colour channels were acquired using excitation lasers for DAPI (excitation at 405 nm, emission at 410-419 nm) and Alexa 568 (excitation at 578 nm, emission at 603 nm) (Figure S6).

2.9 Data analysis

2.9.1 Analysis of the Ca²⁺ imaging dataset

Head-fixed locomotion

As the running speed of the head-fixed animals on the Styrofoam ball during the control and CNO recording sessions was measured in arbitrary values, I set the threshold for detection of running behaviour above the 30th percentile of tracked locomotion across all sessions and mice. Subsequently, all frames were excluded fell below this threshold when calculating the mean running speed in *Matlab* (R2020b).

Ca²⁺ imaging dataset

Suite2p was used for frame registration and ROI extraction to obtain the raw fluorescence traces for each cell for each recording session separately. Custom-written *Matlab* software was used for subsequent semi-manual matching of cells between the control and CNO recording sessions. Cells that could not be matched were excluded from subsequent analysis in *Python*. The fluorescence baseline was defined as the 8th percentile of each cell's raw fluorescence trace, excluding the initial dark frames at the onset of each recording. Df/f was defined as

$$\frac{\text{raw fluorescence trace} - \text{baseline}}{\text{baseline}}$$

Subsequently, the mean df/f for each cell and each recording session was calculated. A modulation index was calculated per cell and recording session, which was defined as

$$\text{modulation index} = \frac{df/f \text{ cno} - df/f \text{ control}}{df/f \text{ cno} + df/f \text{ control}}$$

Thus, a negative modulation index indicates that the cell was inhibited during the CNO recording session, whereas a positive modulation index points toward excitation. To assess whether overall cell activity was affected by the intraperitoneal CNO injections, an inhibition vs. excitation ratio was obtained by dividing

the number of cells with a negative modulation index (inhibition) with the number of cells that had a positive modulation index (excitation)

$$ratio = \frac{inhibition}{excitation}$$

This ratio was used to assess inactivation levels during the CNO recording, where a higher ratio points towards high levels of cell inactivation, and a ratio close to 1 indicates a small if any change in RSC cell activity.

2.9.2 Analysis of the behavioural dataset

The following analysis was performed in *Matlab* (R2020b).

Occupancy maps

The occupancy maps were generated by dividing the mouse's trajectory to the hidden-fixed trigger location into spatial bins (15-by-15 bins) and counting the number of occurrences in each bin. The resulting raw map was then smoothed using a two-dimensional Gaussian filter (filter size=6, standard deviation=0.5) to reduce noise and enhance visual clarity, resulting in a final visual representation of the mouse's occupancy in different regions of the behavioural arena. In the chemogenetic cohort, the entire trajectory to the hidden-fixed trigger location, starting from the previous trial's end point to trigger entry, was used to produce the occupancy map. For the behavioural cohort, the trajectory to the hidden-fixed trigger location, starting from the start search point, was used.

Active task engagement – start search points

The start search point (see Results – 3.1.4 Task engagement) along trajectories to the trigger location was defined as the last time point where the animal reaches the 5th percentile of running velocity (of a given trial) before entering the trigger location. The start search point for the reward port was defined as the lowest running velocity point within a time window 0.5 seconds after trigger entry.

Error angles

The global error angle was defined as the angle between the vector spanning from the trial's start point (for trajectories leading to the trigger position: either the previous trial's reward point or the start search point; for trajectories leading to the port: either the trigger position or the start search point) to the current position of the mouse within the arena, and the vector spanning from the trial's start point to the goal position (either the trigger location or the reward port) (see Figure 3.8 A - *left*).

The instantaneous error angle was defined as the angle between a vector spanning from the animal's current position to the target position (for trajectories leading to the trigger: the trigger position; for trajectories leading to the port: the reward port), and a vector spanning from the animal's previous position to the animal's position in the current frame, where the animal's position is recorded using a framerate of 240Hz (see Figure 3.8 A - *right*).

Angular head velocity

The angular head velocity was calculated as the change in azimuth angle (horizontal plane, or yaw) over a period of 200 milliseconds and normalised to angular velocity per second.

Ethograms

For the session ethogram, mouse head position and orientation in 3D (in x, y, and z), angular head velocity (ahv), running speed, running acceleration, were plotted over time, along with the task events: visible or hidden trigger zone entry, reward port visit, and reward delivery. Position, rotation, and ahv values were smoothed over 500 milliseconds, and running velocity and acceleration were smoothed over 213 milliseconds using a centred simple moving average.

For the trial ethograms, mouse running velocity, mouse distance to the trigger zone and baited reward port, and shortest Euclidean distance to the trigger zone and baited reward port were plotted over time. Velocity was smoothed over 100 milliseconds using a centred simple moving average.

Path surplus

Path surplus was defined as the fraction of distance travelled by the mouse from the trial's start point to trigger entry of the Euclidean distance from trial start to trigger entry.

Statistical analysis

In this study, statistical analyses were performed using the Wilcoxon rank sum test to assess the significance of differences between groups. The Wilcoxon rank sum test was chosen due to its suitability for non-normally distributed data or small sample sizes. Assumptions associated with this test include the independence of observations and the ability to rank the data. In instances where multiple comparisons were conducted, a Bonferroni correction was applied to control for Type I errors. Statistical significance was defined as p values below 0.05 (*), 0.01 (**), and 0.001 (***), if not stated otherwise.

3. Results

3.1 Spatial memory task performance

To study spatial learning and memory in mice, I developed a spatial memory task. In this behavioural paradigm, freely moving and behaving animals learn, over the course of several weeks and training sessions, to remember the location of three consecutive hidden triggers, which are circular zones located within a larger behavioural arena (Figure 2.3). The location of these hidden triggers is not indicated to the mice, so that the animal must memorise its location to locate and visit it again. This is made possible, as the location of the hidden trigger zones within the arena does not change in a given training period; the mice have to learn the location of hidden-fixed triggers. When the mouse successfully enters a hidden trigger zone, an LED light turns on at one of the eight reward ports embedded in the arena wall, indicating to the mouse which port is baited with reward in this trial. After the mouse has consumed the reward, the next trial begins. In every third trial, the trigger zone is invisible and fixed in position (hidden-fixed trigger trial). These trials alternate with visible-random trigger trials. In visible-random trigger trials, a trigger zone is projected in a random location into the arena with the help of a ceiling mounted projector. Mice must follow the light projection to find the visible-random trigger. In correct trials, mice receive a small amount of sweetened soy milk at the reward port with the LED light turned on.

3.1.1 Mice form spatial memories

I first tested if mice could follow the task structure, and whether they could learn to remember the hidden-fixed trigger locations. To test this, I trained a cohort of 14 animals over the course of 37 days on three different hidden-fixed trigger locations: SM1, SM2, and SM3 (i.e., spatial memory 1, 2, and 3) (Figure 2.1 B). Each of these hidden-fixed triggers was placed in a different tertile of the arena and had a different distance to the arena centre and wall (see Methods – 2.4.10 Trigger placement; and Figure 2.7 A). The hidden-fixed trigger zones were placed at different distances from the arena wall to avoid thigmotaxis as a potential search strategy across the three different hidden-fixed trigger locations.

All mice, except one, which met the dropout criteria (see Methods – 2.4.14 Dropout criteria) and was subsequently excluded from further analysis, were able to learn the task structure and to recall the location of the hidden-fixed trigger zones. Mice underwent a habituation phase (see Methods – 2.4.11 Habituation), after which mice showed continuous task improvement over the course of the pre-training phase (see Methods – 2.4.12 Pre-training), increasing their correct trial count (i.e. trials that resulted in the mouse receiving a reward at the end of a trial) in a 20-minute training session from 4.5 ± 0.7 correct trials on the first day of training to 124.7 ± 35.7 correct trials on the last day of pre-training (Figure 3.1 A).

After mice became familiar with the trial structure and the reward delivery mechanism and logic, they moved on to the spatial memory training phase (see Methods – 2.4.13 Spatial memory training). During this training stage, mice need to remember the location of the hidden-fixed trigger zone in every third trial, which is no longer projected into the arena. Mice were trained daily on the task, whereby a training session lasted 20 minutes. To quantify task performance, I used the trial count of correct trials per training session (133.5 ± 36.1 of correct trials per session during spatial memory training) (Figure 3.1 A) and the time needed for the mice to locate the trigger zone in hidden-fixed and visible-random trigger trials

separately (7.7 ± 6.3 seconds to visible-random triggers, 8.4 ± 11.3 seconds to hidden-fixed triggers during spatial memory training) (Figure 3.1 B). While a small drop in performance was observed during the first days of a given training phase when mice were introduced to a new hidden-fixed trigger location (Figure 3.1 A, on the first day of training on SM1, SM2, and SM3), I found that performance levels remained largely stable during the spatial memory training (Figure 3.1, Figure S1, Figure S2). Stable behavioural performance was observed both during memory recall, when mice learnt the location of a new hidden-fixed trigger, and during reversal learning, when the hidden-fixed trigger zone switched from SM1 to SM2, and from SM2 to SM3.

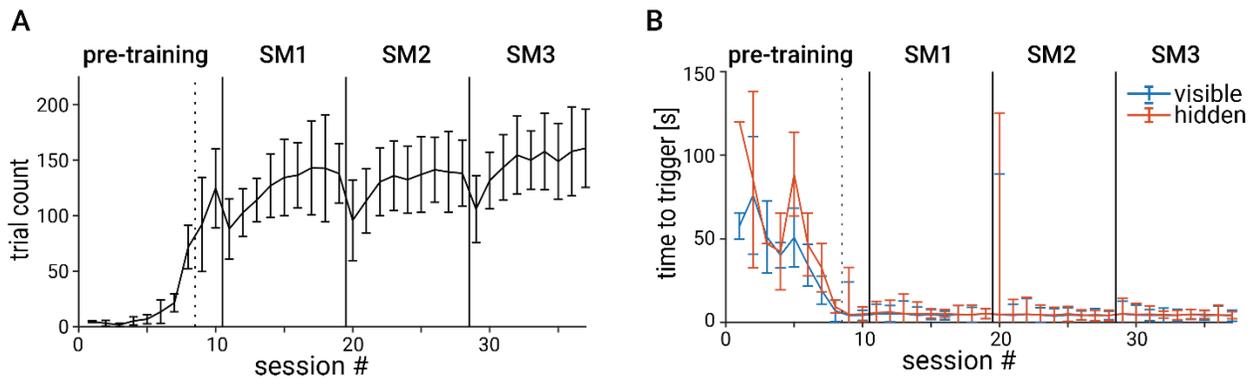


Figure 3.1. Task performance

A. Median (\pm std) trial count throughout pre-training and the spatial memory training (three hidden-fixed trigger zones: SM1, SM2, and SM3) (*black vertical lines*). $n=13$ mice. B. Median (\pm std) time to visible-random (*blue*) and hidden-fixed (*red*) trigger zones. $n=13$ mice.

In addition, I plotted the trajectories of hidden-fixed trigger trials per mouse and training session and colour-coded the resulting occupancy maps (Figure 3.2). I observed that during training on a particular hidden-fixed trigger zone (SM1, SM2, or SM3) occupancy hotspots formed in and around this zone (Figure 3.2, second row – last day of training). Interestingly, during reversal learning, when the hidden-fixed trigger zone moved from one location in the arena to another (SM1 to SM2, and SM2 to SM3), I noticed an occupancy hotspot at the previous hidden-fixed trigger zone. For example, as shown in Figure 3.2, on the first day of training on SM2, a hotspot formed in the upper right quadrant of the arena, which is where the former SM1 hidden-fixed trigger used to be located. Similarly, on the first day of training on SM3, an occupancy hotspot can be observed in the former SM2 hidden-fixed trigger zone in the lower left quadrant of the arena. These results suggest that mice were able to recall previously learnt spatial memory (location of the previous hidden-fixed trigger zone).

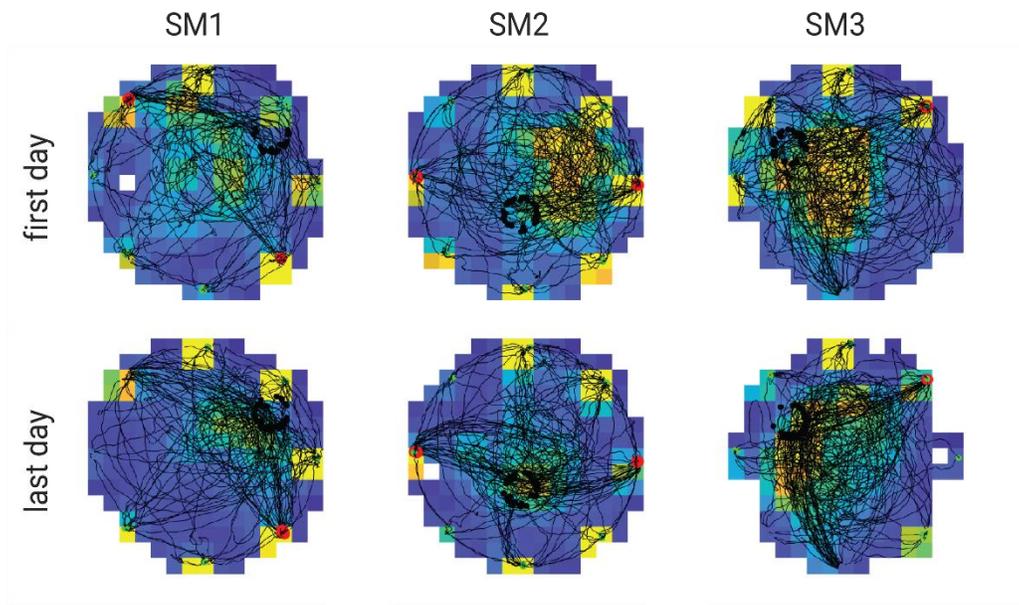


Figure 3.2 Example occupancy maps

Example occupancy plots and trajectories (*black*) of a single mouse for the first (*upper row*) and last (*lower row*) day of training for the three hidden-fixed trigger zones: SM1, SM2, and SM3. Trial start and end points shown as *green* and *red* dots, respectively.

I show that mice are capable of learning and remembering the location of three hidden-fixed triggers - SM1, SM2, and SM3. In the spatial memory task that I developed, mice were able to acquire spatial memories, recall them during subsequent training sessions, and perform reversal learning.

3.1.2 Low and stable error trial rate

In each trial, the mouse has to locate the trigger zone (which can be visible or hidden) and retrieve the reward from the correct reward port (i.e. the port with the LED light turned on). Mice can make any of the following three mistakes in this process, in which case, the trial will be classified as an error trial: visiting the wrong port (wrong port - WP), finding the trigger but not visiting any port subsequently (trigger-no port - TNP), and not finding the trigger (no trigger - NT) (see Methods – 2.4.7 Error trials).

At the beginning of the pre-training session, I observed a moderate number of NT error trials, as the mice had not yet learnt the general task structure. However, as mice learnt to traverse the trigger zones, the number of TNP error trials increased, as animals still needed to learn the association between finding the triggers and the subsequent reward to be retrieved at baited ports. Finally, in the last few training sessions of the pre-training phase, the number of WP errors rose as the mice needed to learn in a final step that only the port with the LED light turned on was baited with reward. To speed up this learning process, mice were trained for two days on the one-port training protocol (see Methods – 2.4.12 Pre-training) before the start of the spatial memory training. At the beginning of the spatial memory training, the error trial rate was low and remained stable throughout the rest of the training protocol (Figure 3.3).

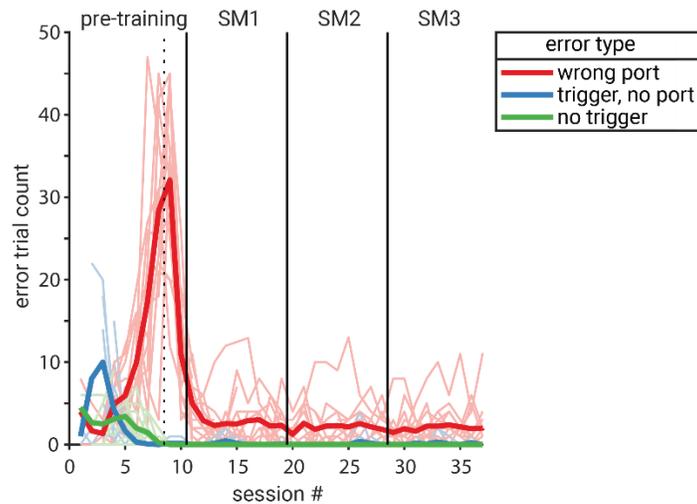


Figure 3.3 Error trials

Error trial count (*median in bold*) for the three types of error trials: wrong reward port (wrong port, *red*); trigger found, but not reward port (trigger, no port, *blue*); and no trigger found (no trigger, *green*). n=13 mice.

The three types of error trials - wrong port (WP), trigger-no port (TNP), and no trigger (NT) error trials - highlight the evolution of task structure learning, which is particularly highlighted during the pre-training phase. On the other hand, during spatial memory training, the error trial rate is low and remains stable throughout training. To focus on trials where mice have successfully recalled the memory of the hidden-fixed trigger location, all error trials were excluded from subsequent analysis.

3.1.3 Quantifying behavioural task metrics in 3D

The setup built for the spatial memory task allows for the precise measure and analysis of the behaviour of the mice in a 3D space. Mice carry head markers during training, which are attached to a baseplate cemented to the mouse head. These infrared reflectors are detected by 3D motion capture cameras allowing for the determination of the animal's position and behaviour in a 3D space.

To gain a better overview of the animal's behaviour during training, I plotted several task metrics in a training session ethogram: position coordinates in 3D, head rotation, angular head velocity, running velocity and acceleration, as well as task events (Figure 3.4). The inset in Figure 3.4 shows the first three trials of an example training session. In the first and second trial the mouse successfully traverses the visible-random trigger zones and retrieves the reward from the baited ports. In the third trial, the mouse locates the hidden-fixed trigger zone, however, requests the reward from a port that is not baited. This wrong port (WP) error trial can be read out from the session ethogram by the detection of the hidden trigger traversal, the port event and the subsequent *absence* of a reward event (Figure 3.4 inset, lower

plot depicting the task events). It is also noteworthy that the running speed increases as mice approach the trigger locations (Figure 3.4 inset, velocity).

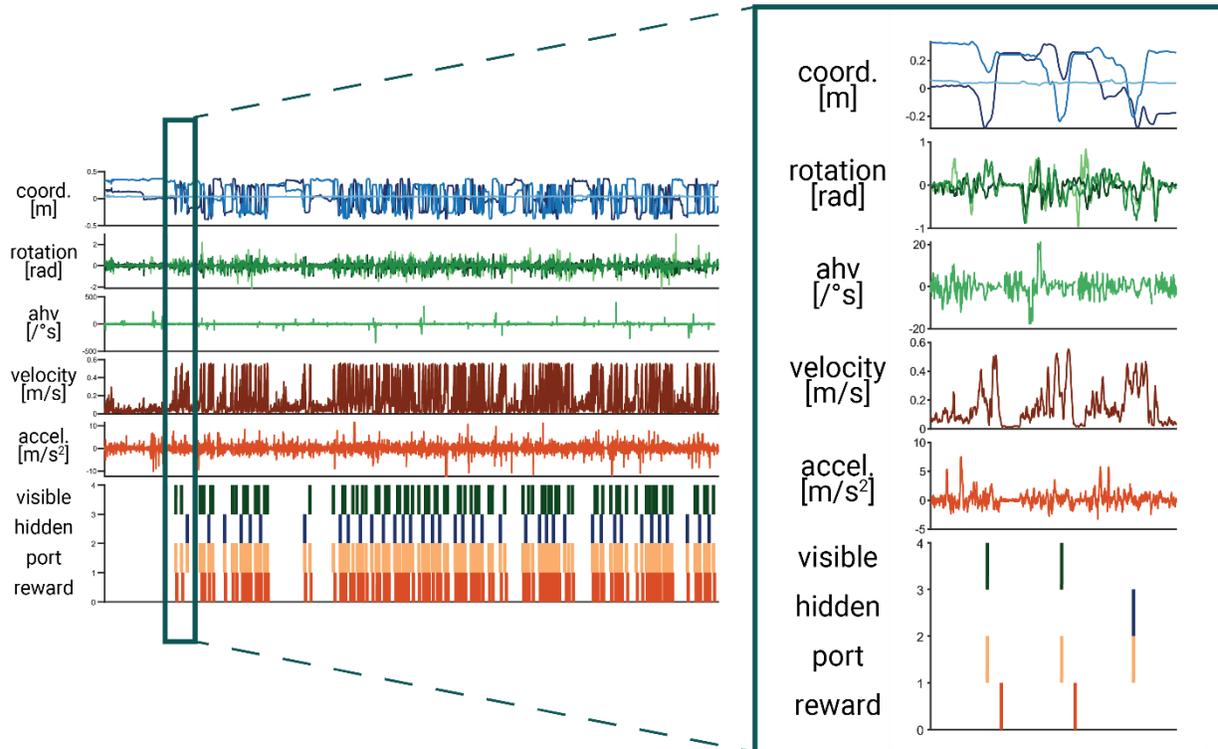


Figure 3.4 Example session ethogram

Ethogram of an example training session (duration of an entire session is 20 minutes, zoom-in shows a 38-second snippet) showing mouse head position and orientation in 3D (in x, y, and z), angular head velocity (ahv), running speed, running acceleration, and the task events: visible or hidden trigger zone entry, reward port visit, and reward delivery over time.

3.1.4 Task engagement

The mice trained on the spatial memory task are food-deprived starting with the onset of the pre-training period and thus motivated to perform the task by the prospect of receiving liquid food rewards after a successful trial. However, the design of the behavioural paradigm requires a certain level of intrinsic motivation as it also allows for mice to spend as much time as desired sitting still, not moving around in the arena to locate the trigger zones. The task design implements no punishment for this disengagement in the task other than mice having less time to find trigger locations and baited reward ports, thereby reducing the amount of reward received in a given training session. I noticed that the mice spent considerable time grooming or consuming the reward, as observed in periods with near-zero velocity

values and unchanging x, y, and z coordinates most often observed immediately after reward delivery (Figure 3.4). It is therefore left to the mice to decide when to start the next trial and to actively search for the current trial's trigger location. Task metrics such as time to trigger, trial duration, and time to port (Figure 3.1 B, Figure S1) are measured from the last trial's port entry and are therefore likely to include times when the animal is not actively engaged in the task but instead is grooming or consuming the reward. To address this issue and to account for this confounder in task timings and durations, I aimed to answer the question: When do mice start actively looking for the next trigger location?

To address this question, I plotted ethograms of individual trials (of variable lengths) to see if I could determine the time point at which mice start a new trial by actively looking for the trigger location. Specifically, I looked at running velocity, mouse distance as well as Euclidean distance to the trigger and baited reward port, and trial time (Figure 3.5).

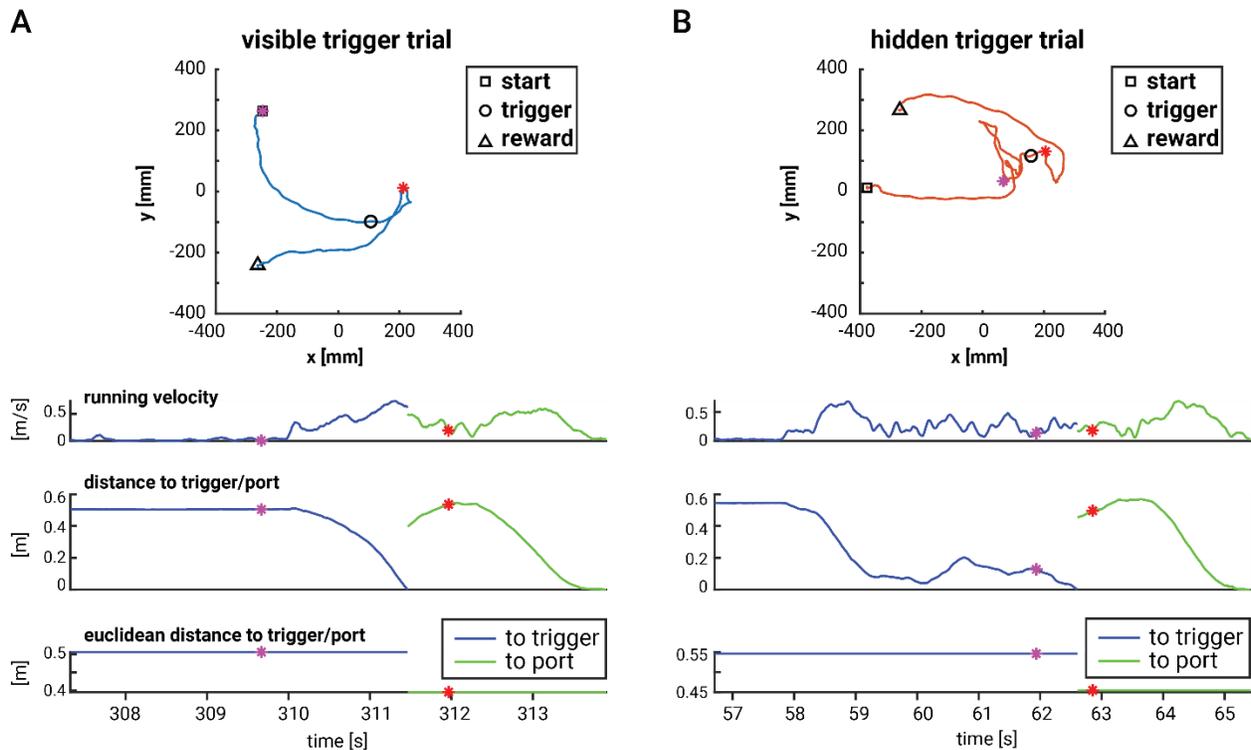


Figure 3.5 Example trial ethograms

Ethogram of two example trials with either a visible-random (A) or hidden-fixed (B) trigger location. (*Top*) Mouse trajectory starting at the last trial's reward port (start, *square*), through trigger zone traversal (trigger, *circle*; in A, trajectory to the visible trigger in *blue*; B, trajectory to the hidden trigger in *red*), and ending at the given trial's baited reward port (reward, *triangle*). (*Bottom*) Mouse running velocity, mouse distance to the trigger zone and baited reward port, and shortest Euclidean distance to the trigger zone and baited reward port plotted against time. Values corresponding to the trajectory to the trigger zone are shown in *dark blue*, and to the reward port in *green*. The *pink* and *red* asterisks denote the start search time points for trajectories leading to the trigger zone and the reward port, respectively. The examples shown here depict correct trials that resulted in reward delivery at the end of the trial.

First, I find that animals spent a large fraction of the trial duration consuming the reward and/or grooming and were thus not actively engaged in the next trial. This is evident in the example trial ethograms by the low running velocity, with values close to zero for large fractions of the trial. This disengagement in the task can be further seen by the distance between mouse position and the trigger, which does not decrease for the first seconds of the trial. The start of the active search for the trigger location becomes visible when the mouse speeds up (increased velocity) and the distance to the trigger starts decreasing. I thus defined the start of active task engagement as the last time point where the mouse reaches the 5th percentile of running velocity (of a given trial) before entering the trigger location. In the example trial ethogram in Figure 3.5, this active start search point lies at a time point when the mouse is still positioned

at the start point, i.e., the previously rewarded port, as the mouse had not yet moved (Figure 3.5 A – top panel, pink asterisk in square). By employing this 'start search' point as the beginning of task metrics, such as trial-duration or time-to-trigger, more reliable measures of active task engagement can be achieved (Figure 3.6).

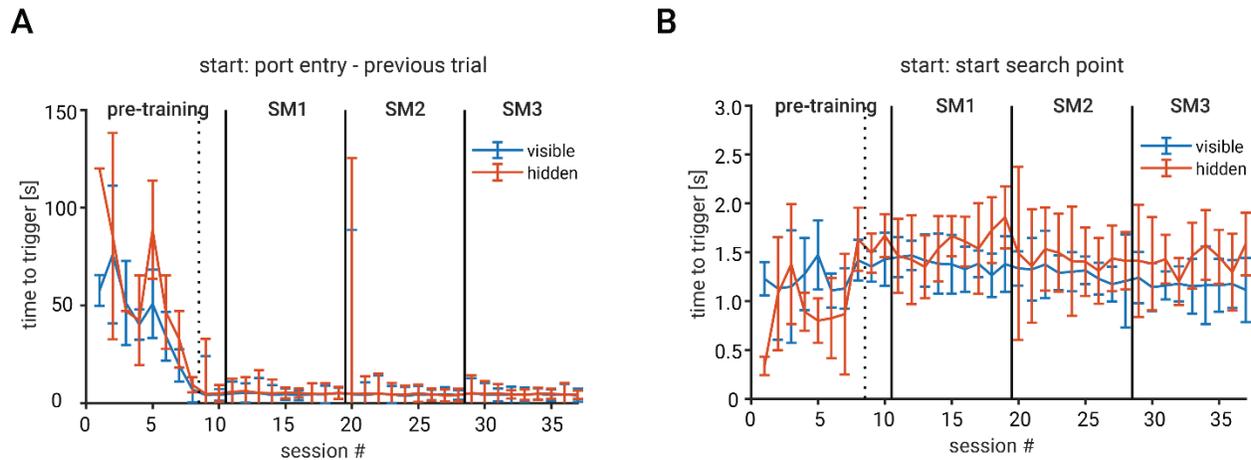


Figure 3.6 Time to trigger

Median (\pm std) time to visible-random (blue) and hidden-fixed (red) trigger zones measured from the time point of port entry in the previous trial (A) or from the start search point of active task engagement (B). $n=13$ mice.

Second, I find that mice tend to overshoot the trigger location. This overshooting was most often observed in hidden-fixed trigger trials (Figure 3.5 B – top panel) and can be explained by the absence of visual feedback as to whether the hidden trigger location has been traversed yet. Consequently, the time spent overshooting the trigger location is included in time measurements, such as time-to-port. To address this second confounder, I defined the time point of the active search for the port as the lowest velocity point within a time window of 500 milliseconds after trigger entry (Figure 3.5 B – red asterisk). To test the observation that mice predominantly overshoot the trigger location in hidden-fixed trigger trials, I measured the distance from the centre of the trigger location to the start search point along the trajectory to the port for both visible-random and hidden-fixed trigger trials separately (Figure 3.7 A-C). I find this distance to be significantly larger in hidden-fixed trigger trials compared to trajectories from visible-random trigger zones (first day of training: $p < 0.0005$ ***; second day of training: $p < 0.0005$ ***; Wilcoxon rank sum test, Bonferroni correction; Figure 3.7 C).

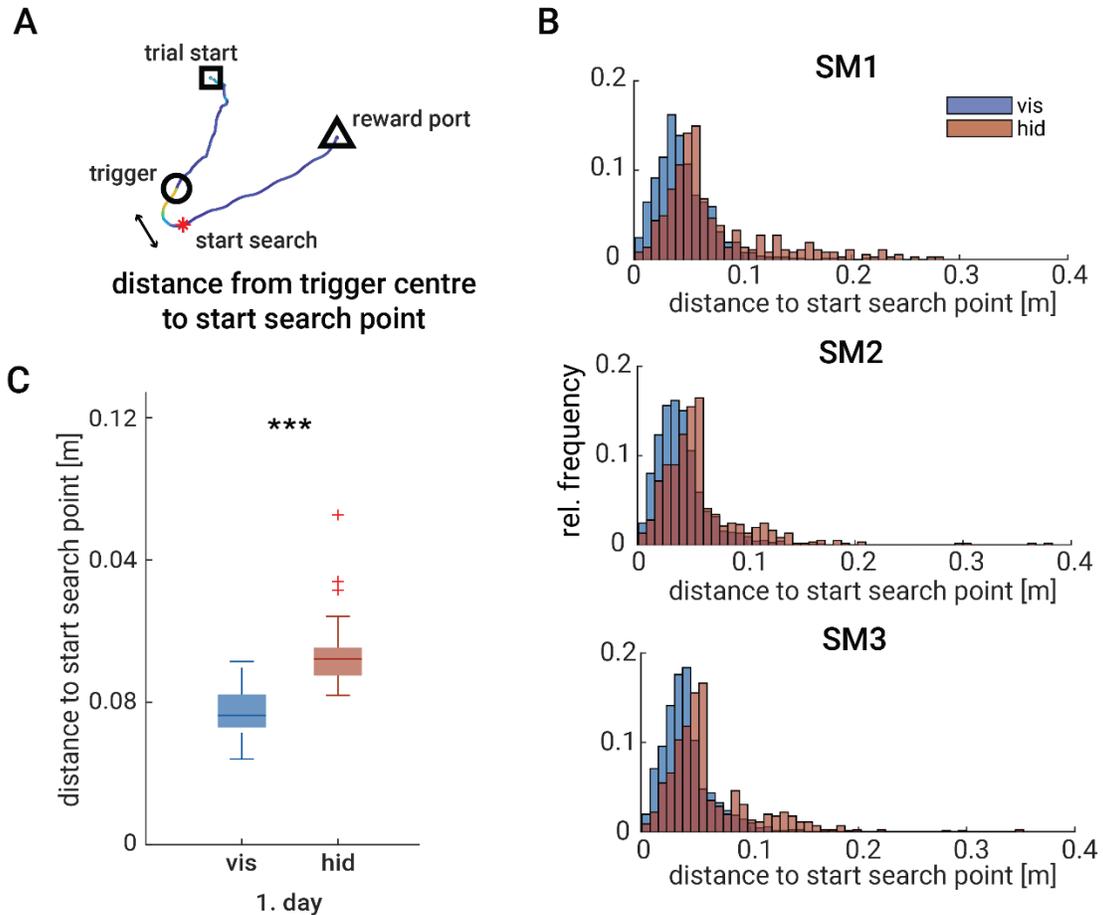


Figure 3.7 Trigger precision

A. Schematic depicting an example trajectory from trial start (*square*), to trigger (*circle*), and reward port (*triangle*), and the measurement of the distance from trigger centre to start search point (*red asterisk*). B. Distribution of the individual trials' distances measured from the trigger centre to the start search point for visible-random (*blue*) and hidden-fixed (*red*) trigger trials for the three spatial memory training stages: SM1, SM2, and SM3. C. Boxplot comparison of the distance from trigger centre to start search point for visible-random (*blue*) and hidden-fixed (*red*) trigger trials for the 1st training day of all spatial memory training stages. *Wilcoxon rank sum test, Bonferroni correction. *** $p < 0.0005$. $n=13$ mice.*

To summarise, I have identified 'start search' points along trajectories of individual trials that serve as time points to measure the start of active task engagement to find either the trigger location or the baited reward port. These points are necessary to determine the duration and timing of active engagement in the task. By identifying these time points, I can eliminate confounding factors, such as time spent consuming rewards or grooming, and found a difference between visible-random and hidden-fixed trigger trials by measuring trigger precision.

3.1.5 Error angles differ between visible and hidden trigger trials

Next, I set out to quantify and characterise the trajectories leading to the visible-random and hidden-fixed trigger zones in more detail. Mice use their vision in order to locate visible-random triggers, while they must rely on their memory to find hidden-fixed triggers. I wanted to examine if the difference between these two strategies – vision and memory – was reflected in the trajectories by identifying their respective error angles.

I decided to focus on two types of error angles - a global and an instantaneous error angle (Figure 3.8). The global error angle is defined as the angle that spans between the vector from the trial start point (i.e. the start search point) to the trigger location and the vector from the start point to the current position of the mouse. This global error angle thus measures the angle between the position of the mouse and the trigger location in relation to where the mouse started the trial, considering mouse location but not orientation, or heading direction, of the animal (Figure 3.8 A, *left*). In contrast, the instantaneous error angle is measured as the angle between the current position of the mouse and the trigger location in relation to the animal's position at the previous time point. The instantaneous error angle thus provides information about the heading direction of the mouse, but not its general position in relation to start and goal point (Figure 3.8 A, *right*).

These two error angles thus provide different information about trajectory efficiency, local searches, and heading direction. While both these error angles measure the deviance from a straight trajectory from the start point to the trigger, the global error angle provides a measure of the overall understanding of the mouse as to where the trigger location is, whereas the instantaneous error angle detects smaller deviations from a straight trajectory right away. This implies, that if a mouse were to start a trial going in the wrong direction that leads away from the trigger (goal) position, the global error angle would be high throughout this trial. In contrast, the instantaneous error would start decreasing the moment the mouse is correcting its path, facing towards the trigger location. Similarly, during a local search in which the mouse is close to the hidden-fixed trigger location, but requires a number of smaller left and right turns to eventually enter the invisible trigger zone, the global error angle would not be able to detect these small changes in heading direction while the instantaneous error angle would remain high throughout this focal search (Figure 3.8 B).

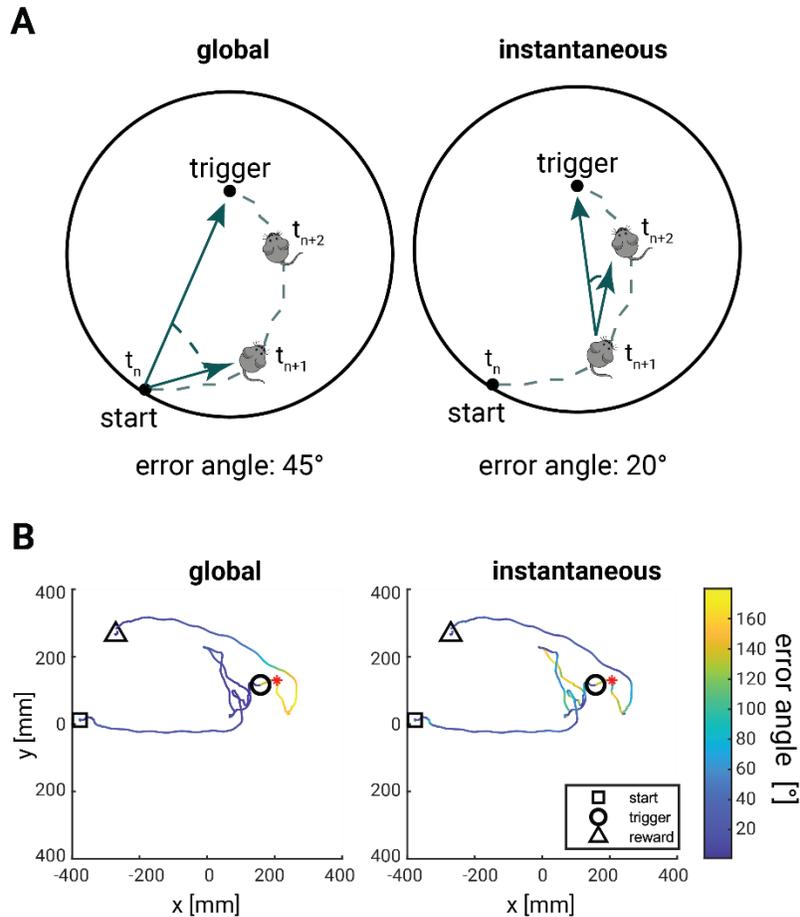


Figure 3.8 Error angle definition and example

A. Schematic of the global (*left*) and instantaneous (*right*) error angle calculation. B. Example trajectory of a single trial of an example mouse. Colour-coded global (*left*) and instantaneous (*right*) error angle. *The example shown here depicts a correct trial that resulted in reward delivery at the end of a hidden-fixed trial.*

I calculated the mean values of both the global and instantaneous error angle from the start search point to trigger entry for visible-random and hidden-fixed trigger trajectories, separately. I find higher global error angles in visible-random trigger trial trajectories ($p < 0.0005$ *** on the first day of training; $p < 0.0005$ *** on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice; Figure 3.9 A, *left*) and higher instantaneous error angles in trajectories leading to hidden-fixed triggers ($p < 0.0005$ *** on the first day of training; $p < 0.0005$ *** on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice; Figure 3.9 A, *right*).

As a control, I calculated the mean error angle of trajectories leading to the reward port, as measured from the start search point to port entry. In all trial types, hidden-fixed or visible-random, mice need to use their vision to navigate to the indicated reward port, which is why I expected the error angles along the trajectory to the reward port to be similar across hidden-fixed and visible-random trigger trials.

Indeed, I do not find a significant difference in instantaneous error angle between visible-random and hidden-fixed trigger trials ($p = 0.32$ on the first day of training; $p = 0.33$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice; Figure 3.9 B, *lower right distribution*). Interestingly, I find significantly higher global error angle trajectories leading to the port in visible-random trials ($p < 0.0005$ *** on the first day of training; $p < 0.0005$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice; Figure 3.9 B, *lower left distribution*). I hypothesised that the higher global error angle could be explained by mice running a curved, instead of a straight, trajectory to visible-random triggers (Figure 3.11). Mice can adjust their path along their way to the visible-random trigger as they receive visual feedback, and subsequently continue this curved trajectory *en route* to the reward port. This is in contrast to the high instantaneous error angles leading to the hidden-fixed trigger, where mice often perform a local search near the hidden trigger location until they eventually enter the trigger zone (Figure 3.11; and see Figure 3.8 B for an example for a local search near the hidden-fixed trigger zone).

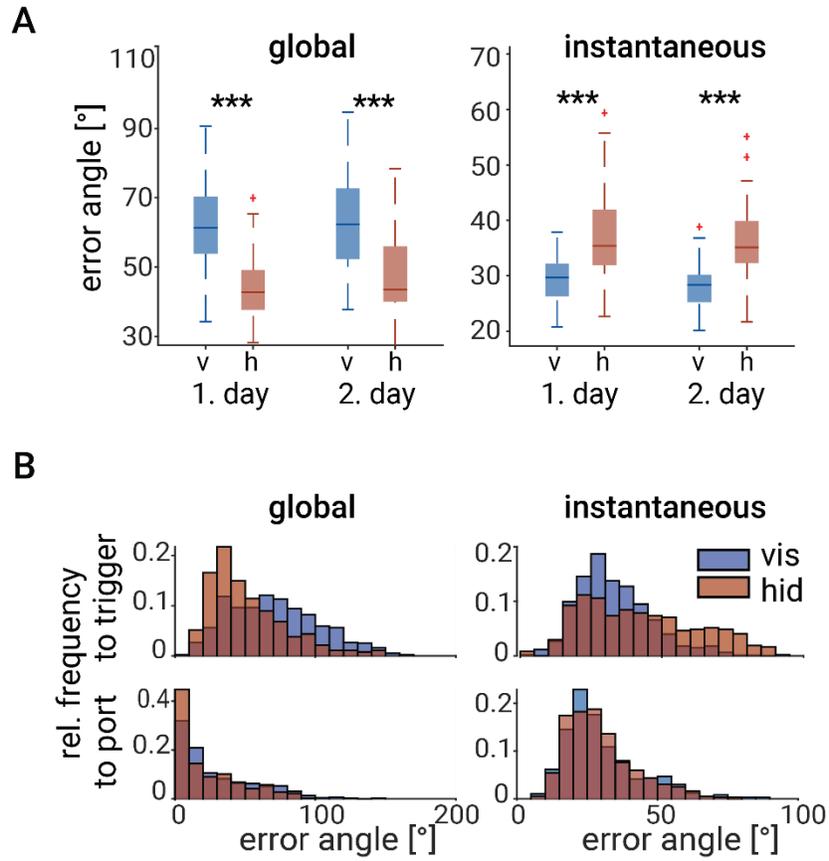


Figure 3.9 Error angles of trajectories leading to triggers and ports

A. Boxplot comparing error angles (global, *left*; instantaneous, *right*) between visible-random (*blue*) and hidden-fixed (*red*) trigger trials for the first and second day of new hidden trigger presentations. *Wilcoxon rank sum test, Bonferroni correction, *** $p < 0.0005$. $n=13$ mice.* B. Distribution of global (*left*) and instantaneous (*right*) error angles of trajectories leading to the trigger zone (*upper*) and to the port (*lower*) for an example training day (first day of SM1 training). *Blue, visible-random trigger trajectories; red, hidden-fixed trigger trajectories. $n=13$ mice.*

This hypothesis is supported by a significantly higher mean angular head velocity measured from the start search point to trigger entry in trials with a visible-random trigger zone compared to hidden-fixed trigger trials ($p < 0.0005$ *** on the first day of training; $p < 0.0005$ *** on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice. Figure 3.10, *left*), a difference I found to be significantly different along the trajectory to the port on the second, but not on the first day of training ($p = 0.087$ on the first day of training; $p = 0.005$ ** on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice. Figure 3.10 A, *right*). Again, these differences in angular head velocity could be explained by continuously curved trajectories leading to the visible-random trigger and ultimately to the reward port (but see, Figure 3.10 right panel. There is no significant difference in angular head velocity on the first day of training between visible-random and hidden-fixed trigger trials).

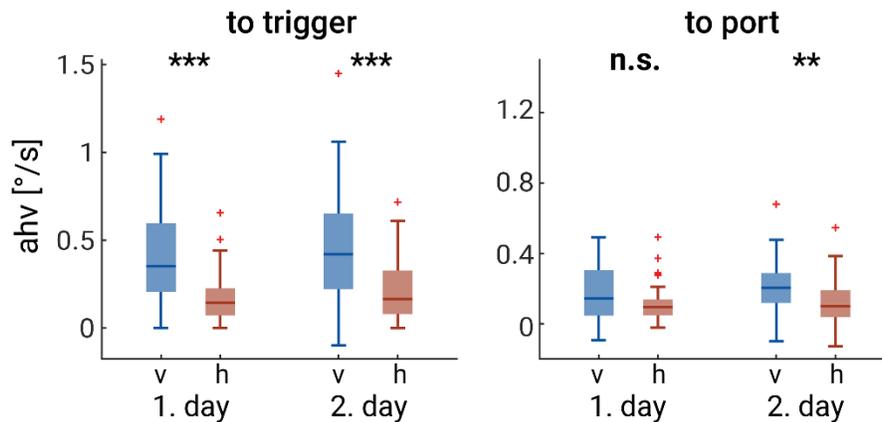


Figure 3.10 Angular head velocity

Mean angular head velocity (ahv) measured from the start search point to trigger (*left*) and to port (*right*) for visible-random (*blue*) and hidden-fixed (*red*) trigger trials. Wilcoxon rank sum test, Bonferroni correction, *** $p < 0.0005$, ** $p < 0.005$. $n=13$ mice.

In summary, I find high global error angles in trajectories leading to visible-random triggers, and high instantaneous error angles to hidden-fixed triggers. This difference in error angles can be explained by mice running curved trajectories to visible-random triggers (a curved trajectory that continues to the reward port), and by the local search when mice were searching for the hidden-fixed trigger (Figure 3.8 B). Overall, I show that error angles work well as task metrics for capturing the difference between these two trajectory types. Error angles characterise the trajectories to visible-random and hidden-fixed trigger zones better than time measurements, such as time-to-trigger and time-to-port (Figure S3). Here, I only observe a significant difference between visible-random and hidden-fixed trigger trials on the second day of training during the trigger search only (to trigger: $p = 0.09$ on the first day of training; $p = 0.02$ * on the second day of training; to port: $p = 0.88$ on the first of training; $p = 0.65$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction, $n=13$ mice. Figure S3).

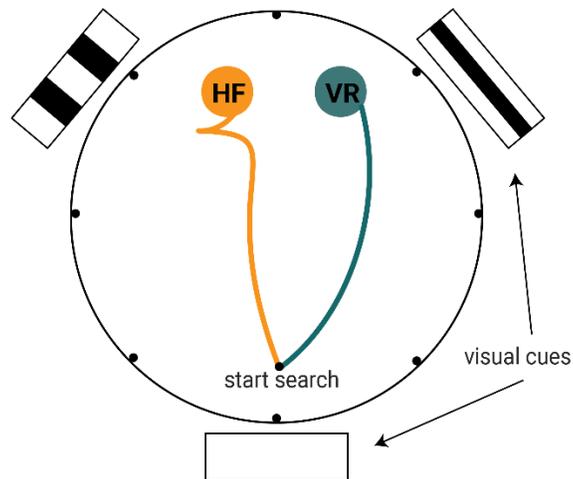


Figure 3.11 Schematic of error angle hypothesis

Schematic of the different trajectories to memory-guided hidden-fixed (HF, *orange*) and visible-random (VR, *blue*) trigger zones.

3.1.6 Summary

I have designed and built a setup that enables us to track the position and rotations of mice in a 3D space, allowing us to measure various parameters related to their performance on a spatial memory task I developed. During this task, mice learn the location of hidden-fixed trigger zones and are able to recall their position within the arena on subsequent training days, as well as perform reversal learning. Overall, mice show consistent performance on the task, with a small number of error trials. I have identified specific points in time during the trial when mice start displaying active task engagement, as well as error angles that correspond to their search for visible-random and hidden-fixed trigger zones. I found that trajectories to visible-random triggers tend to have larger global error angles, while those to hidden-fixed triggers have larger instantaneous error angles. These metrics allow us to quantify and distinguish between curved trajectories to visible-random triggers and local searches near hidden-fixed triggers.

3.2 Chemogenetic inhibition of RSC activity

In a next step, I wanted to investigate RSC's role in the spatial memory task. Specifically, I asked if animals could still learn and acquire spatial memories when RSC activity is chemogenetically inactivated (Figure 2.1 B). I expressed inhibitory DREADDs in excitatory neurons in both granular and dysgranular RSC in all animals of the chemogenetic cohort (n=13 mice).

3.2.1 DREADD/CNO-mediated reduction in RSC cell activity

To test that CNO injections led to DREADD-mediated inactivation of RSC cell activity, I imaged calcium activity in 12 out of 13 mice prior to the start of the behavioural training. I recorded RSC neuronal activity in two settings: under control conditions (no injection) and after intraperitoneal CNO injections. To compare neuronal activity under these two conditions, I controlled that the same region was imaged in the two imaging sessions by re-finding the same neurons using the blood vessel pattern as landmarks for orientation (Figure 3.12).

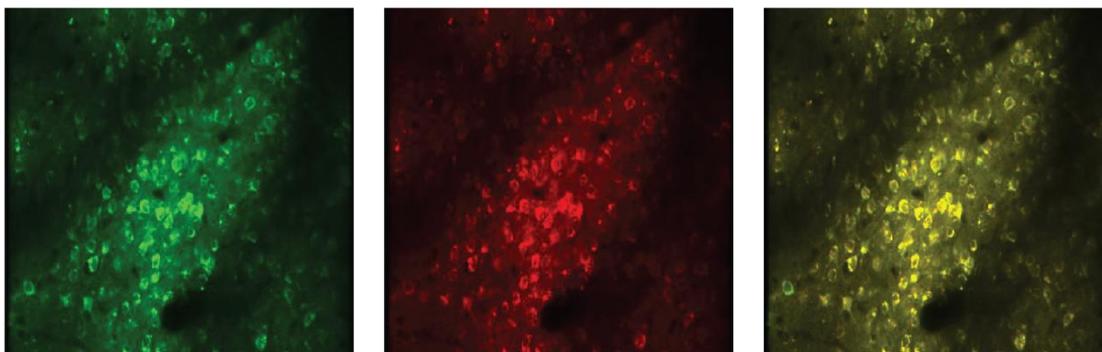


Figure 3.12 Example imaging planes

Two-photon imaging field of view of an example mouse. Animals were Ca^{2+} imaged under both control (left column, *green*) and CNO (middle column, *red*) conditions. The blood vessel pattern was used to re-find the same FOV and cells across the two imaging sessions, during each a z-stack consisting of four planes (only plane is shown here) was obtained. The overlay of the two recordings is shown on the *right*.

To control for locomotion-induced, rather than DREADD/CNO-mediated, changes in cell activity I compared the running behaviour of the head-fixed mice on the Styrofoam ball during the control and CNO recording sessions. Running behaviour was indistinguishable between the control and CNO recording sessions and the mean running speed per mouse compared between these two conditions did not significantly differ ($p = 0.98$, Wilcoxon rank sum test. $n=12$ mice. Figure 3.13). Therefore, locomotion and running behaviour were excluded as confounding factors in neural activity changes.

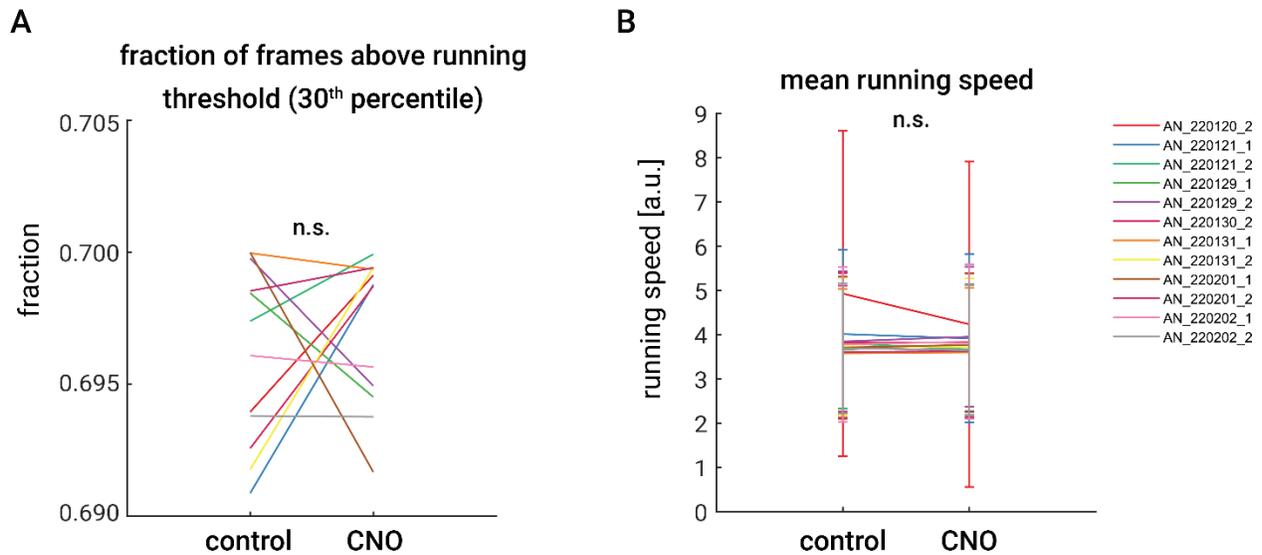


Figure 3.13 Locomotion during Ca²⁺-imaging session

A. Fraction of frames that lie above the defined running threshold (30th percentile) per mouse and recording session (control vs. CNO) and were thus included in the calculation of mean running speed. $P=0.54$, Wilcoxon rank sum test. $n=12$ mice. B. Mean running speed per mouse on the Styrofoam ball during the control and CNO Ca²⁺-imaging session (arbitrary units). $P=0.98$, Wilcoxon rank sum test. $n=12$ mice.

For ROI extraction Suite2p was used, whereby the red channel, detecting mCherry co-expressed with inhibitory DREADDs, was used as a structural marker for image registration and motion correction, while the green channel detecting GcAMP6s signal served as an indicator for neuronal activity.

Whenever possible, I semi-automatically matched the recorded cells from the control imaging session with those of the CNO imaging session using custom-written software to investigate neuronal activity change on a single-cell basis (11 out of 12 mice). The mean df/f values were calculated for each cell and recording session individually and compared across the two recording sessions (Figure 3.14 A). A modulation index (see Methods – 2.8.1 Analysis of the Ca²⁺ imaging dataset) was used to determine if a cell showed higher levels of excitation or inhibition during CNO recording conditions. Using this modulation index, I calculated the excitation vs. inhibition ratio, which I found to be higher when comparing cell activity during control and CNO conditions (Figure 3.14 A, B). As a control, I compared the first with the second half of either only the control recording session, or only the CNO recording session, where I observed lower inhibition vs. excitation ratios (Figure 3.14 C), pointing towards an overall reduction of cell activity following i.p. CNO injections and thus successful DREADD/CNO-mediated, chemogenetic inactivation of RSC activity.

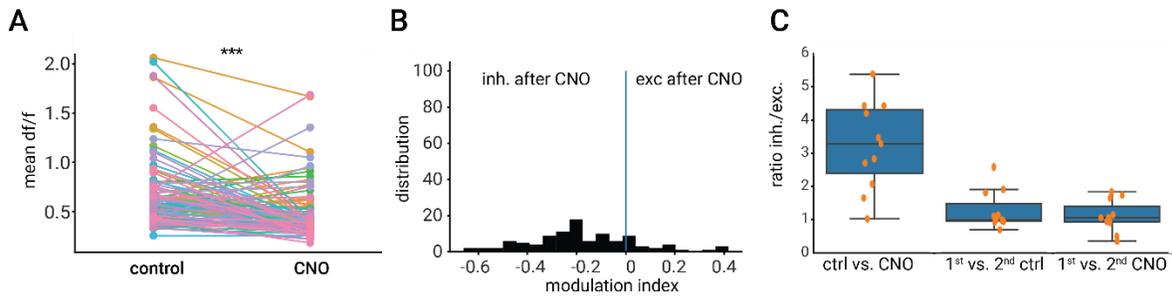


Figure 3.14 Chemogenetic inactivation of RSC activity

A. Mean df/f values of all cells recorded and cell-matched across the control and CNO Ca^{2+} imaging session of an example mouse. $p < 0.001$ ***, *Wilcoxon rank sum test*. $n=104$ cells. B. Distribution of the modulation indices of all detected and cell-matched cells from the control and CNO Ca^{2+} imaging session of the example mouse from A. Negative and positive modulation index values indicate inhibition and excitation of cell activity in the CNO recording session compared to the control recording, respectively. $n=104$ cells. C. Inhibition/excitation ratio for the control vs. CNO recording session (*left*), 1st half vs. 2nd half of the control session (*centre*), and 1st half vs. 2nd half of the CNO recording session (*right*). $n=11$ mice.

3.2.2 Mice form spatial memories under CNO conditions

After the Ca^{2+} imaging session, mice started the pre-training and subsequent spatial memory training. A cohort of 13 animals was split into two groups: a control (7 mice) and experimental (6 mice) group, which received intraperitoneal saline or CNO injections prior to selected training sessions, respectively. The training days on which animals received injections were chosen in such a way that the effect of RSC inactivation could be observed during memory acquisition, memory retrieval and during reversal learning (Figure 2.1 B).

To compare overall task performance between the control and CNO groups, I counted correct and incorrect trials per training session, as well as the time it took mice to locate the hidden-fixed trigger zones. I find comparable behaviour between the two group without significant differences in these behavioural metrics (Figure 3.15 A-E).

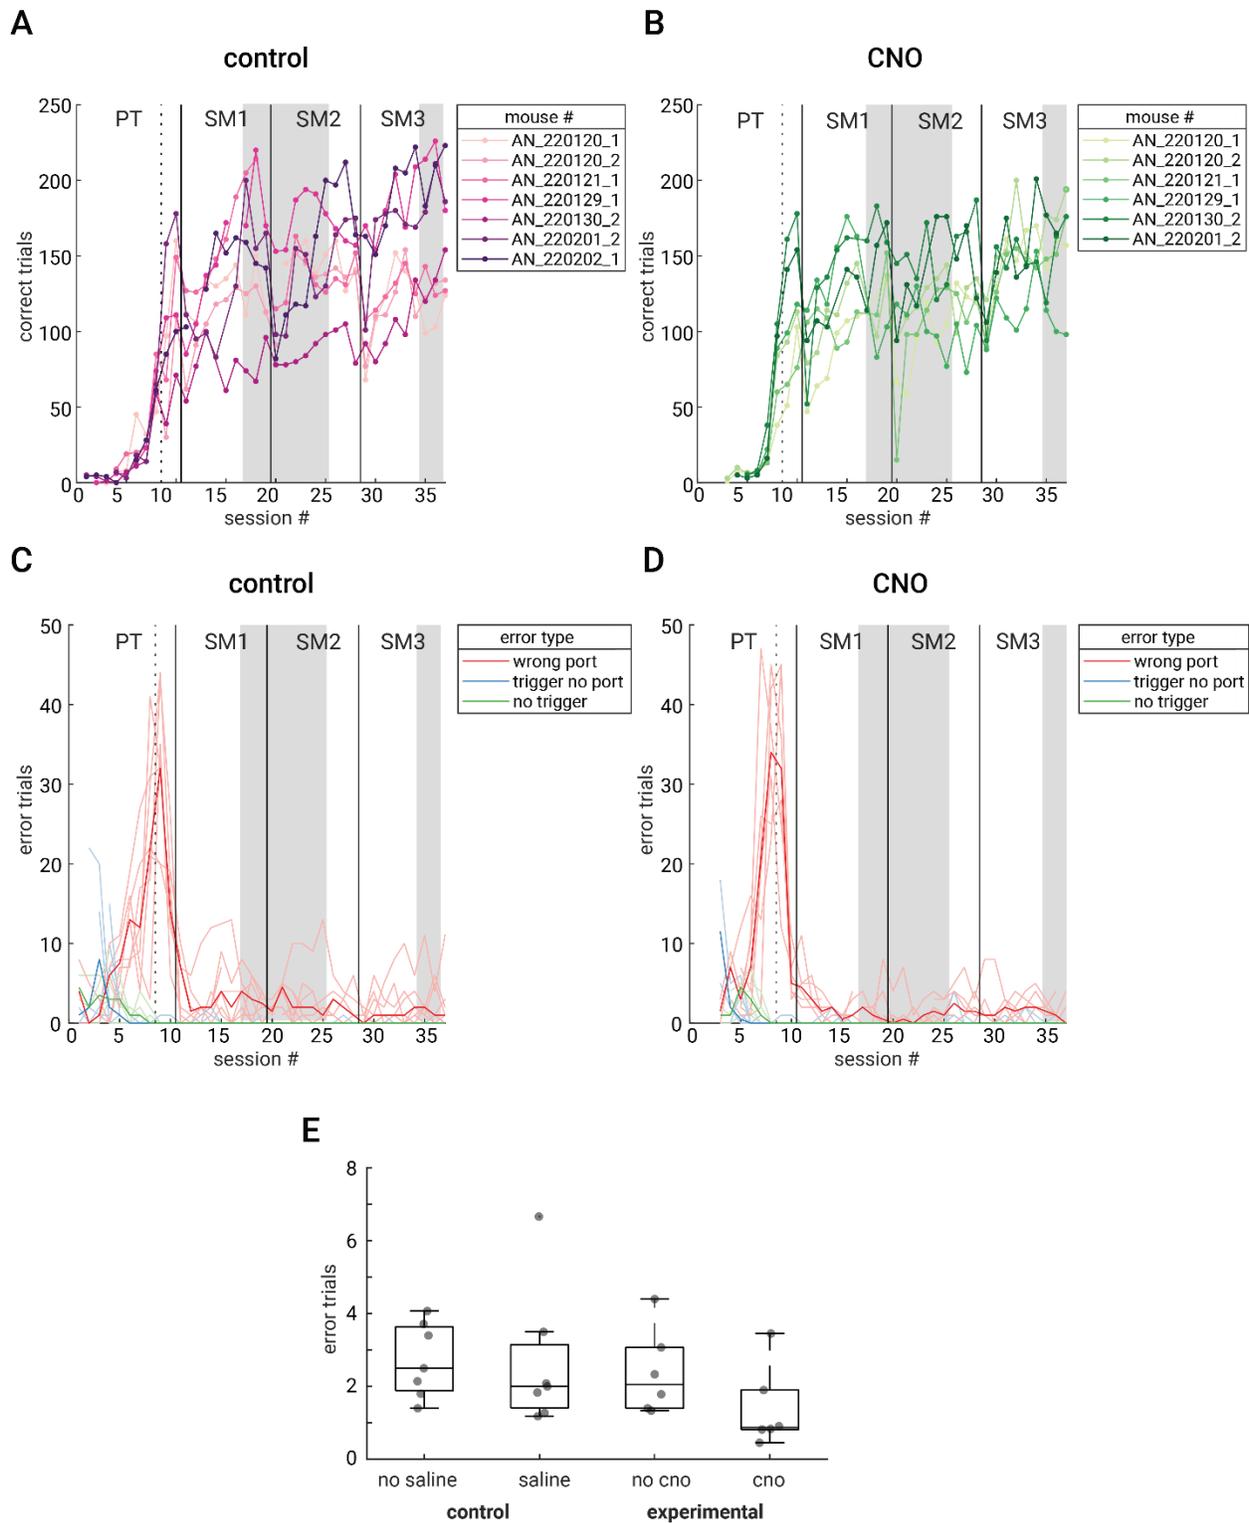


Figure 3.15 Task performance under control and CNO conditions

Correct trial count per session for animals of the control (A) and CNO (B) group. Error trial count per session (*median in bold*), separated for error trial type (wrong port, red;

trigger, no port, *blue*; no trigger, *green*), for animals of the control (C) and CNO (D) group. Error trial count for animals of the control and experimental group separated by days on which they received saline/CNO injections prior to the training session (E).

I further plotted occupancy maps to investigate memory acquisition, recall and reversal learning more qualitatively. Occupancy maps look comparable between animals of the control and CNO group (Figure 3.16 A-B). As observed and described previously (Figure 3.2), hotspots form where the hidden-fixed trigger location is placed (memory acquisition; last days of training in Figure 3.16). These hotspots initially remain in the same place when the hidden-fixed trigger location changes (memory recall; first days of training on SM2 and SM3 in Figure 3.16), however form anew in the place where the new hidden-fixed trigger location is located (reversal learning; last days of training on SM2 and SM3 in Figure 3.16).

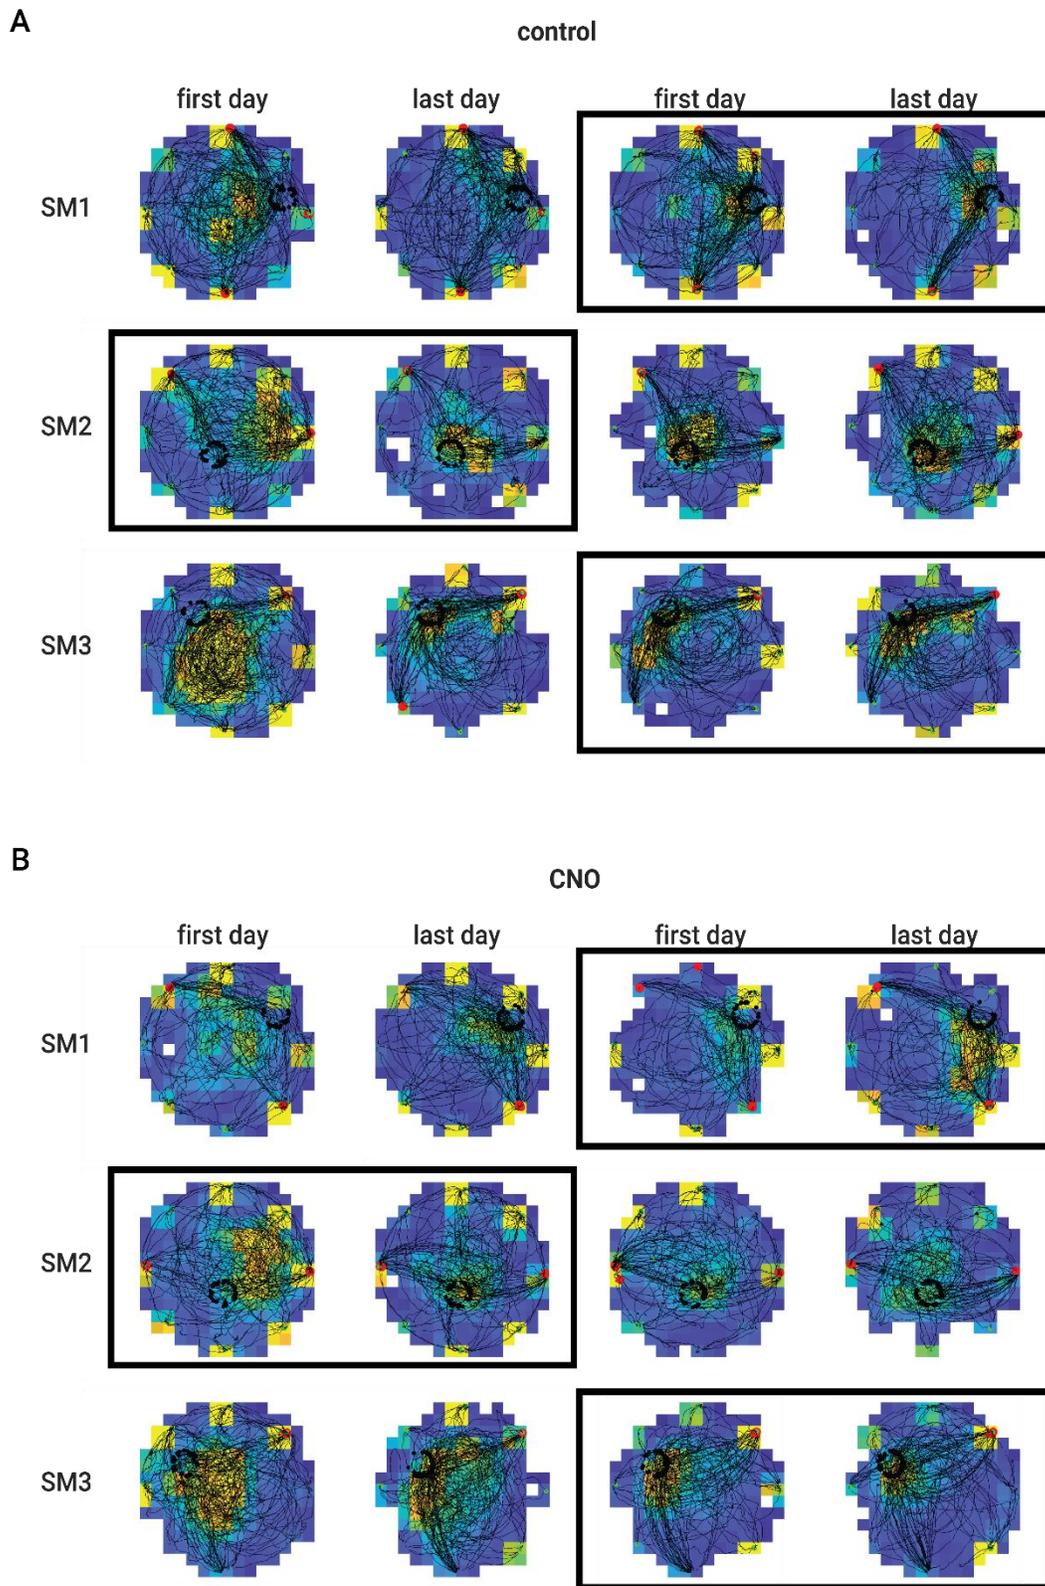


Figure 3.16 Occupancy map examples of a control and experimental mouse

Example occupancy plots and trajectories (*black*) of a mouse from the control (A) and a mouse from the CNO (B) group. First and last days of training for the different training stages and the three hidden-fixed trigger zones (SM1, SM2, and SM3). First and last days of training on which animals received saline (control group, A) or CNO (experimental group, B) i.p. injections are shown in a *black box*. Trial start and end points shown as *green* and *red* dots, respectively.

To investigate whether animals change their behaviour and/or running trajectories when re-locating the hidden-fixed trigger locations when RSC is inactivated, I calculated the mean global and mean instantaneous error angle along the trajectory to hidden-fixed and visible-random triggers (measured from the start search point). No significant differences in either instantaneous (for trajectories leading to hidden-fixed triggers: $p = 0.70$ on the first day of training; $p = 0.70$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction; Figure 3.17 A; for trajectories leading to visible-random triggers: $p = 0.40$ on the first day of training; $p = 0.24$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction) or global error angle (for trajectories leading to hidden-fixed triggers: $p = 0.90$ on the first day of training; $p = 0.30$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction; Figure 3.17 B; for trajectories leading to visible-random triggers: $p = 0.18$ on the first day of training; $p = 0.09$ on the second day of training; Wilcoxon rank sum test, Bonferroni correction. $n=7$ mice for the control group, $n=6$ mice for the CNO group) between animals of the control and CNO group was detected.

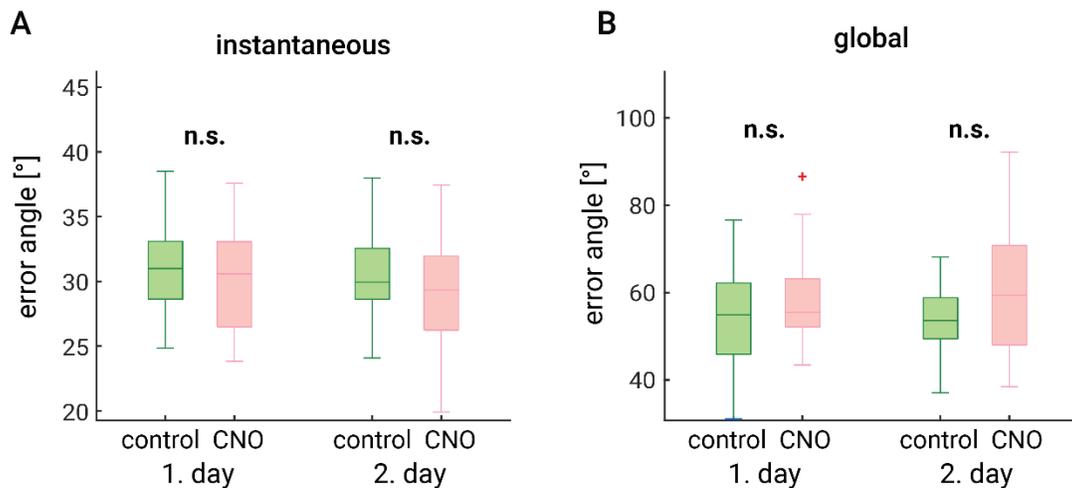


Figure 3.17 Error angles under control and CNO conditions

Boxplots comparing the instantaneous (A) and global (B) error angle of trajectories leading hidden-fixed triggers between animals of the control (green) and CNO (pink) group for the first and second day of training on SM1, SM2, and SM3. *Wilcoxon rank sum test, Bonferroni correction. $n=13$ mice.*

3.2.3 Summary

In summary, I conclude that chemogenetic inhibition was successful as I observe a reduction in RSC cell activity under CNO conditions, compared to control conditions. However, there were no behavioural differences in task performance (correct/incorrect trial count per training session and instantaneous/global error angle to hidden-fixed or visible-random triggers) comparing animals that received CNO or saline injections prior to behavioural training, indicating that mice can still perform the spatial memory task when RSC activity is inactivated.

3.3 Comparing memory-guided and visually-guided navigation

Next, I aimed to compare task performance between vision-guided (visible trigger) and memory-guided (hidden trigger) trials. In the previous cohort, animals were only able to use vision to find the randomly placed visible trigger zones, whereas they had to rely on memory to find the hidden-fixed trigger zones, since no visual indication of the hidden trigger location was provided. Therefore, I wondered whether these two conditions – vision-guided and memory-guided – could lead to different navigation strategies, potentially also explaining the higher global error angles in visible-random trigger trials and higher instantaneous error angles in hidden-fixed trigger trials observed in the previous cohort. However, in order to compare these two navigation strategies, there needs to be a condition under which animals can choose between vision-guided and memory-guided navigation.

3.3.1 Introduction of visible-fixed trigger zones

To test this, I trained a new cohort of eight mice (five female and three male mice) on the spatial memory task. As this cohort was intended for a purely behavioural assessment, only a head-mounted baseplate was implanted during surgery, which is necessary to attach the head-mounted IR reflectors for position tracking. I trained the animals for five days each on three different hidden trigger locations that were fixed in position - SM1, SM2, and SM3 (Figure 2.1 A). To test whether memory-guided navigation can improve task performance also in visually-guided trials, I introduced a new trial type in addition to hidden-fixed and visible-random trigger trials. This new trial type involved a visible, fixed trigger zone allowing mice to rely on either vision or memory to locate it. That was made possible as this trigger zone was projected into the arena, and its location did not change for a given training period. I modified the training session from a three-trial structure to a four-trial structure, alternating between a visible-fixed and hidden-fixed trigger zone for every other trial. These fixed-trigger trials were interleaved with visible-random trigger trials (Figure 2.5 B).

3.3.2 Comparable behavioural metrics in both cohorts

I first determined whether the mice had learnt to perform the task and could remember the location of hidden-fixed trigger zones. To assess this, I quantified the number of correct and error trials during the pre-training and spatial memory training periods (Figure 3.18). I then performed a qualitative analysis and compared these values with the trial counts from the previous, chemogenetic cohort, where similar task performance was observed. In both cohorts, animals displayed increased error trials during the pre-training phase, which decreased during the one-port training of the pre-training and remained low through the spatial memory training on SM1, SM2, and SM3 (Figure 3.18 C, D).

Furthermore, animals from both cohorts showed stable behavioural performance throughout the spatial memory training, with a mean trial count of 130 ± 32 trials in the behaviour cohort, and 153 ± 43 trials for animals of the chemogenetic cohort (Figure 3.18 A-B). The mean error trial count throughout spatial memory training was below 5 per session with a similar standard deviation observed in both cohorts (5.0 ± 3.9 error trials in the behavioural cohort, and 3.9 ± 3.6 error trials in the chemogenetic cohort) (Figure 3.18 C-D).

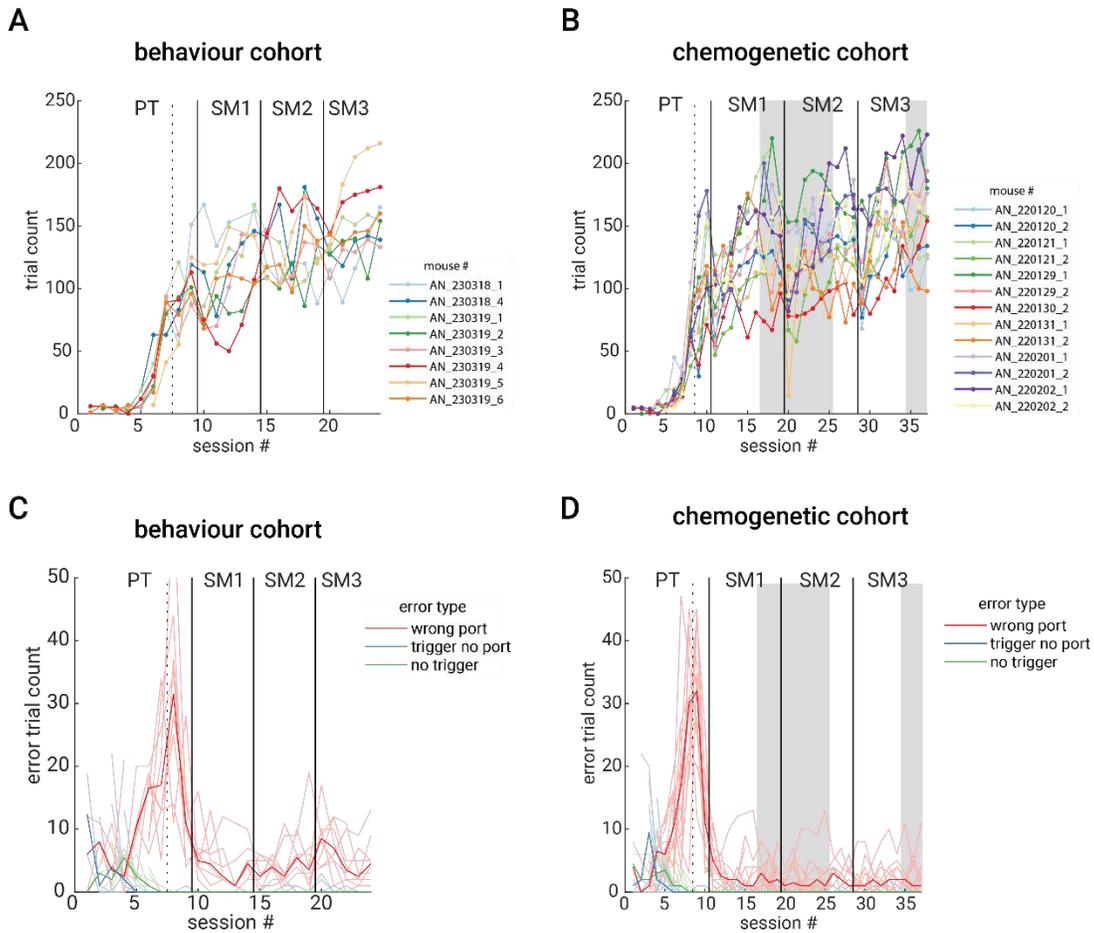


Figure 3.18 Trial count comparison of the behavioural and chemogenetic cohort

Number of correct trials per mouse and training session throughout pre-training (PT) and spatial memory training (SM1, SM2, and SM3) for the behavioural (A) and chemogenetic (B) cohort. Error trial count, separated by type of error (wrong-port in *red*, trigger-no-port in *blue*, and no-trigger in *green*), shown for the pre-training (PT) and spatial memory training (SM1, SM2, and SM3) phases for the behavioural (C) and chemogenetic (D) cohort. Median values are shown in darker colours, whereas individual mouse values are shown in fainter colours. $N=8$ mice for the behaviour cohort, $n=13$ mice for the chemogenetic cohort.

Mice were able to acquire spatial memories, recall them during later training sessions, and also learn to reverse previously learnt behaviour, as demonstrated by the occupancy maps of individual training sessions (Figure 3.19). Occupancy hotspots that formed near and around the hidden-fixed trigger zone during training on a particular hidden-fixed trigger zone provide evidence of this phenomenon, i.e. memory. On the first day of training on a new hidden-fixed trigger location, I observed an occupancy hotspot in the area where the previous hidden-fixed trigger zone used to be (i.e., memory recall). That is, an occupancy hotspot is observed on day 1 of SM2 and SM3 training at the location where the hidden-

fixed trigger used to be during SM1 and SM2 training, respectively. Additionally, new hotspots were observed when animals were trained longer on new hidden-fixed trigger locations (i.e., reversal learning). These new hotspots became apparent on day 3, 4 and 5 of a given training stage. This reversal learning process was repeated twice: first by changing the hidden trigger zone from SM1 to SM2, and then from SM2 to SM3.

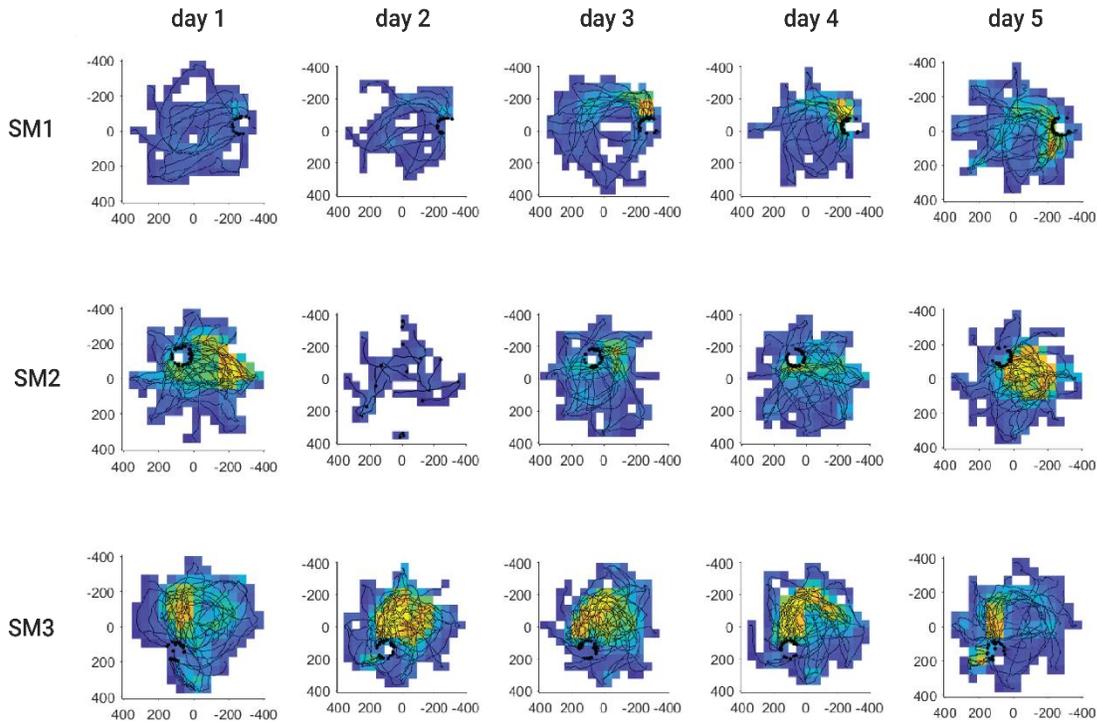


Figure 3.19 Example occupancy maps of the spatial memory training

Occupancy maps calculated from hidden-fixed trigger trials only for each training session (day 1-5) for the three spatial memory training stages (SM1, SM2, and SM3) of an example mouse (AN_230319_3). *Yellow* areas show higher occupancy.

Over the course of the five-day training sessions per SM1, SM2 and SM3, mice showed continuous improvement in task performance. I assessed this task improvement by individually analysing the trial duration for the three trial types – visible-random, visible-fixed, and hidden-fixed (Figure 3.20). I find a more prominent difference in trial duration between visible and hidden trigger trials on the first day of training, compared to the last day (examples are shown for SM1 and SM3 in Figure 3.20). Furthermore, the difference in trial duration between visible and hidden trigger trials is more prominent in SM1 training than for SM3 on both the first and last day of training (SM1: vr vs. hf: $p = 0.003$ * on the first day of training,

$p = 0.003$ * on the last day of training; vf vs. hf: $p = 0.003$ * on the first day of training, $p = 0.003$ on the last day of training; SM3: vr vs. hf: $p = 0.0002$ ** on the first day of training, $p = 0.02$ on the last day of training; vf vs. hf: $p = 0.0002$ ** on the first day of training, $p = 0.13$ on the last day of training; Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice), showing that mice exhibit task improvement with each new hidden-fixed trigger location they are trained on. Additionally, I looked at more direct, time-independent task metrics such as path surplus, which measures the fraction of the distance travelled by the mouse in relation to the Euclidean distance from the trial start point to the point of trigger entry. This new task metric, path surplus, supports previous results (SM1: vr vs. hf: $p = 0.0002$ ** on the first day of training, $p = 0.004$ * on the last day of training; vf vs. hf: $p = 0.0002$ ** on the first day of training, $p = 0.002$ * on the last day of training; SM3: vr vs. hf: $p = 0.0003$ ** on the first day of training, $p = 0.08$ on the last day of training; vf vs. hf: $p = 0.002$ * on the first day of training, $p = 0.13$ on the last day of training; Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice), where mice improve task performance both over the course of the five training days, as well with each training on a new hidden-fixed trigger position (Figure S4). However, no significant differences were observed between visible-random and visible-fixed trigger trials.

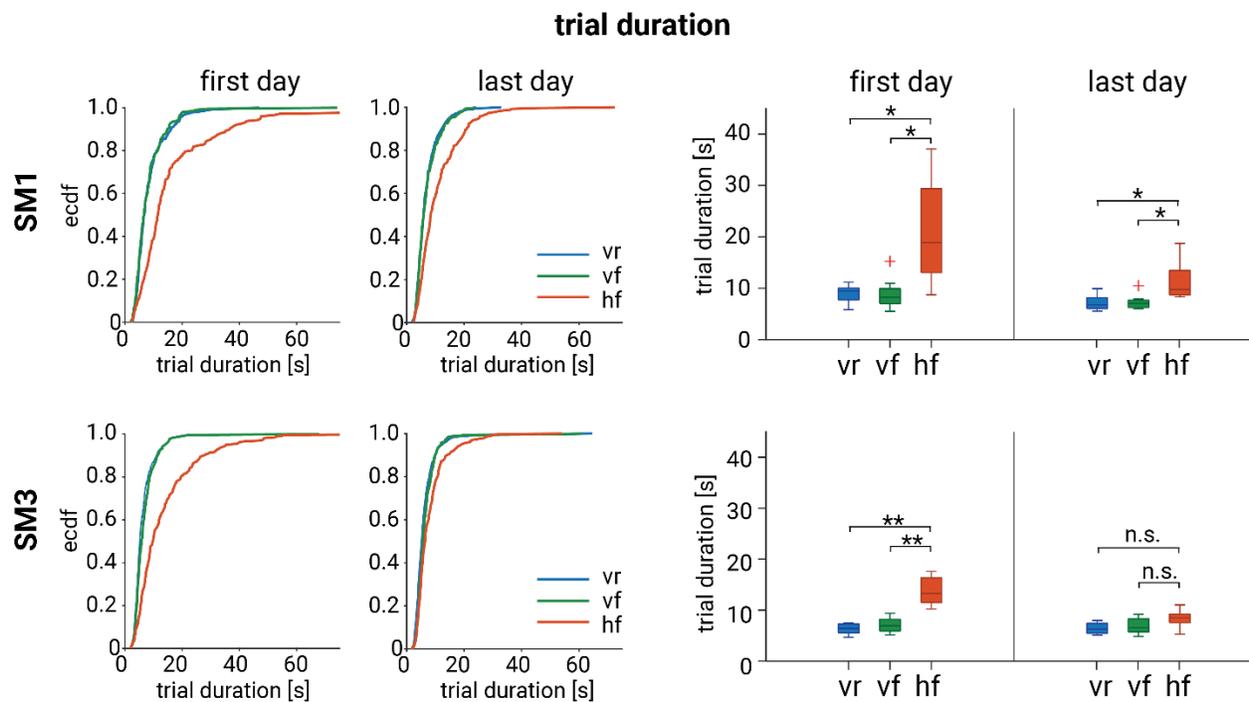


Figure 3.20 Trial duration as a metric for task performance

Trial duration, as measured from trial start to port entry, for the first and last (5th) day of training shown for SM1 (upper row) and SM3 (lower row) as empirical cumulative distribution function (ecdf) (left) and boxplot (right), Wilcoxon rank sum test, Bonferroni correction, ** $p < 0.0016$, * $p < 0.008$. $n = 8$ mice.

3.3.3 Mice only use memory-guided navigation in the absence of visual information

To characterise the trajectories leading to the visible-random, visible-fixed, and hidden-fixed trigger locations, I determined if the error angle results from the chemogenetic cohort could be reproduced in the behavioural cohort. And, if so, if trajectories to the visible-fixed trigger zone display a greater resemblance to vision-guided trajectories (visible-random trigger trials) or memory-guided trajectories (hidden-fixed trigger trials).

These findings are consistent with previous results, showing a higher mean instantaneous error angle, as measured from the start search point to trigger entry, during hidden-fixed trigger trials compared to visible-random and visible-fixed trigger trajectory types (vr vs. hf: $p = 0.00007$ *** (day 1), $p = 0.0003$ ** (day 2), $p = 0.0009$ ** (day 5); vf vs. hf: $p = 0.002$ * (day 1), $p = 0.03$ (day 2), $p = 0.002$ * (day 5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice) (Figure 3.21 A-B). These results suggest that mice rely mainly on vision to locate the visible-fixed trigger locations, as there is no significant difference between visible-random and visible-fixed trigger trajectories. Similar to the chemogenetic cohort, I did not find any significant differences in the mean instantaneous error angle along the trajectory from the start search point to the reward port entry across all the different trial types (Figure 3.21 B).

I further find a higher mean global error angle in visible trigger trajectories when comparing trajectories from the start search point to trigger entry (Figure 3.21 C-D). This difference between hidden-fixed and visible trigger trials is significant (vr vs. hf: $p = 0.00006$ *** (day 1), $p = 0.0003$ ** (day 2); vf vs. hf: $p = 0.00201$ * (day 1), $p = 0.002$ * (day 2); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice), but not on the last day of training (vr vs. hf: $p = 0.017$ (day 5); vf vs. hf: $p = 0.2$ (day 5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice); Figure 3.21 D). In addition, there is no significant difference between visible-random and visible-fixed trigger trajectories on any training day. This further supports the hypothesis that mice use their vision to locate the visible-fixed trigger zone. As a control, I obtained the mean global error angle from the start search point to the port entry (Figure 3.21 D). I found that there were significantly higher mean global error angles along the trajectory leading to the port in visible than in hidden trigger trials on the first days of training (vr vs. hf: $p = 0.002$ * (day 1), $p = 0.02$ (day 2), $p = 0.19$ (day 5); vf vs. hf: $p = 0.002$ * (day 1), $p = 0.14$ (day 2), $p = 0.04$ (day 5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice; Figure 3.21 D), which is consistent with the results obtained from the chemogenetic cohort. As suggested earlier, this phenomenon could be explained by continued curved trajectories that pass through the visible trigger and extend to the reward port.

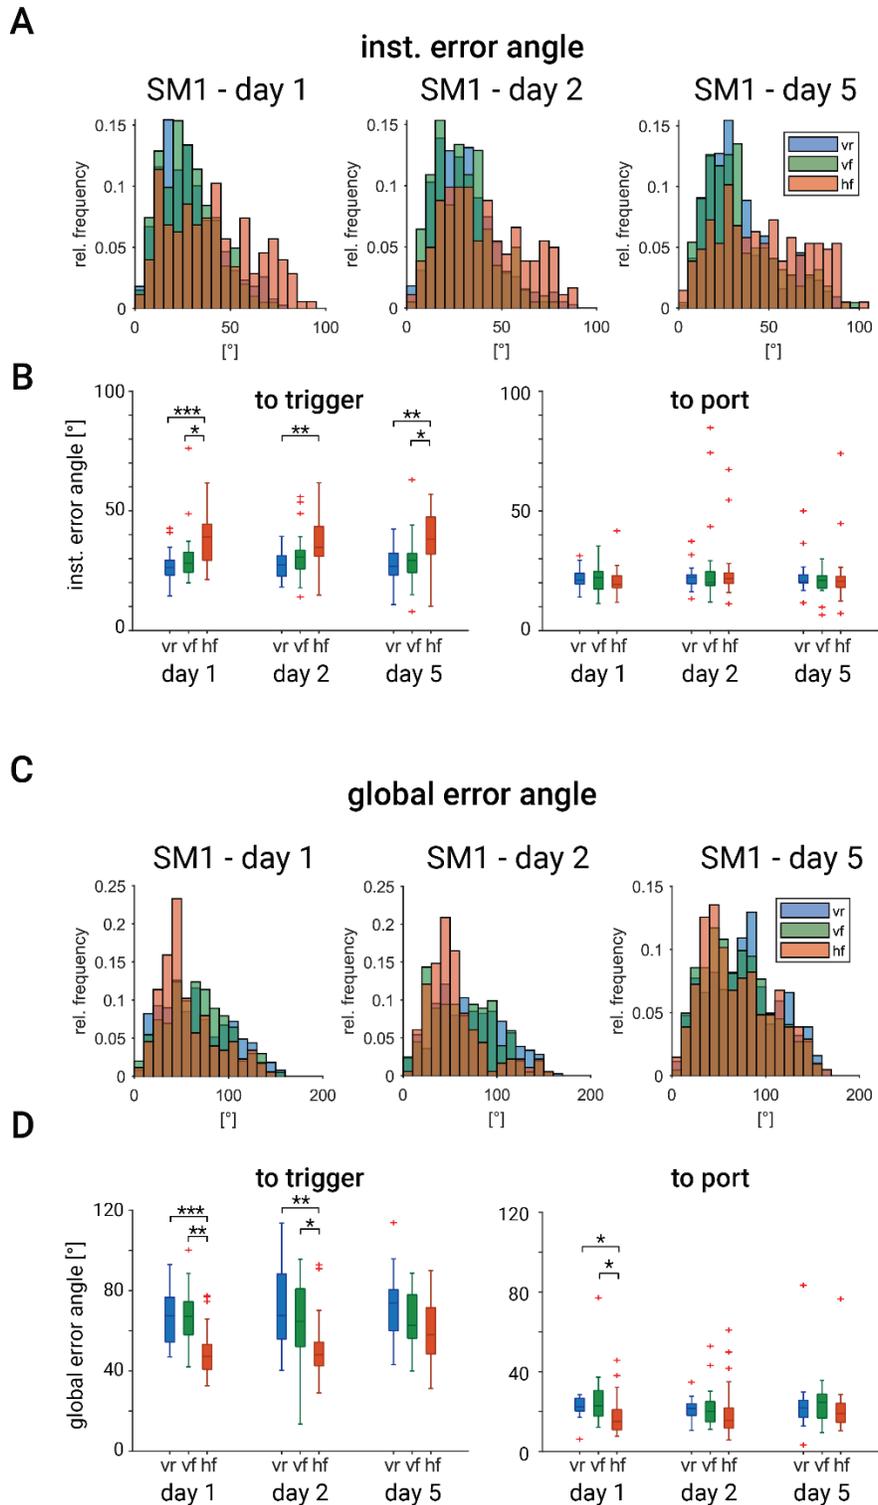


Figure 3.21 Error angles of different trial types

A. Distribution of the mean instantaneous error angle, as measured from the start search point to trigger entry, for visible-random (vr, blue), visible-fixed (vf, green), and hidden-fixed (hf, red) trigger trials. B. Boxplot comparison of the mean instantaneous

error angle for visible-random, visible-fixed, and hidden-fixed trigger trials, as measured along the trajectory to the trigger (start search point to trigger entry) (*left*) and to the port (start search point to port entry) (*right*). C. Distribution of the mean global error angle measure from start search point to trigger entry for the three trial types. D. Boxplot comparison of the mean global error angle for the different trial types, for both the trajectory to the trigger (*left*) and to the port (*right*), measured from the start search point. *Wilcoxon rank sum test, Bonferroni correction, *** $p < 0.00011$, ** $p < 0.0011$, * $p < 0.0055$. $n = 8$ mice.*

3.3.4 Summary

I conducted a set of experiments to compare the task performance of mice when relying on either vision (visible trigger trials) or memory (hidden trigger trials). To do so, I trained a new cohort of mice on the spatial memory task, allowing them to rely on either vision or memory by introducing a new trial type with a visible, fixed trigger location. I found that mice were able to remember the location of the hidden-fixed trigger zones and performed similarly to the previous cohort. I used error angles to characterise and quantify trajectories and were able to reproduce previously observed error angle patterns. I discovered that mice exhibited higher instantaneous error angles during hidden-fixed trigger trials and higher global error angles during visible-random and visible-fixed trigger trajectories. Since there was no significant difference between the visible-random and visible-fixed trigger trajectories, I hypothesised that mice relied primarily on vision to locate the visible-fixed trigger locations and that there was no memory component involved in these trials, indicating that mice only use memory-guided navigation strategies when they cannot use visual information.

3.4 Manipulation of arena cues and trigger placement

My next objective was to investigate the visual information animals rely on when using vision-guided navigation strategies. Specifically, I set out to test whether the animal's memory of the learnt trigger locations is anchored to the arena cues. This would imply that the spatial memories the mice had formed during task acquisition are allocentric, and not egocentric.

Therefore, to test this hypothesis, I performed a set of manipulations. Initially, I established a stable behavioural performance on hidden-fixed trigger location learning during SM3 and SM1. The subsequent experiments involved physically rotating the arena cues or reprogramming the trigger projection zones to create a "rotated" perception:

1. Physical rotation of arena cues

My aim was to determine if spatial memories and, consequently, the occupancy hotspot surrounding the hidden trigger location, would rotate in response to the rotation of the arena cues (Figure 3.22 C). I achieved this by interchanging each arena cue card with the adjacent one, resulting in a 120-degree clockwise rotation of arena cues (Figure 2.4). The rest of the behavioural arena, as well as the relationship between distal and proximal arena cues, remained unaltered.

2. Programmed rotation of all fixed trigger locations (visible-fixed and hidden-fixed)

To determine whether mice utilised other global landmarks, beyond the arena cues purposely provided to them, to form a cognitive map of the arena, arena cues were kept in place (just as they were during the spatial memory training, Figure 3.22 B) while hidden-fixed and visible-fixed trigger projection zones were reprogrammed so that they underwent a 120-degree counter-clockwise rotation (Figure 3.22 D).

3. Rotation of the visible-fixed trigger location only

I recognised the potential for animals to use the visible-fixed trigger location as a point of reference to locate the hidden-fixed trigger location. Notably, these two fixed trigger locations maintain a consistent distance and orientation (120 degrees counter-clockwise rotation from the visible-fixed to the hidden-fixed trigger location), even in the case of rotated arena cues. Given the four-trial sequence - vr-vf-vr-hf - with the visible-fixed trigger being initially presented, mice could potentially utilise this single trial to orient themselves, constructing a new cognitive map of their surroundings independent of the arena cues. To account for this possibility, I rotated the arena cues 120 degrees clockwise (as I did in step 1, rotation of arena cues, Figure 3.22 C). In addition, I reprogrammed the visible-fixed trigger location such that it rotates 240 degrees counter-clockwise, while retaining the hidden-fixed trigger's original position (Figure 3.22 E).

By manipulating the arena cues and trigger locations, I aimed to determine if spatial memories are influenced by the visual environment or if they are independent of it. This allowed me to investigate whether the mice relied on their internal perspective or on external cues for navigation and memory recall.

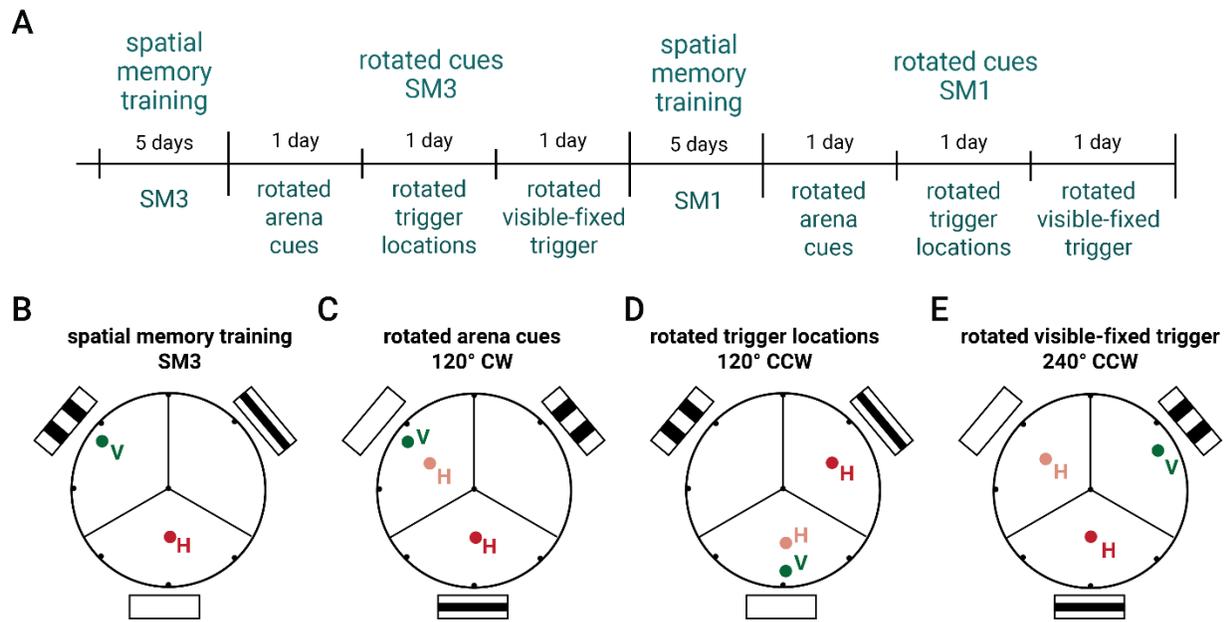


Figure 3.22 Timeline and logic of arena cue and trigger placement manipulations

A. Experimental timeline of the arena cue and trigger placement manipulations. Stable behavioural task performance was established before the three manipulations – rotated arena cues (1), rotated trigger locations (2), and rotated visible-fixed trigger (3) - were tested. This set of experiments was performed twice: once for SM3 (*left column*), and subsequently for SM1 (*right column*) trigger placements. B. Original placement of hidden-fixed and visible-fixed trigger locations for SM3 (*left*) and SM1 (*right*). C. Arena cues and trigger positions following the rotation of the arena cues (1). The *pink dot* indicates the position of the hidden-fixed trigger location if mice anchored their memory to the arena cues. D. Same as in C, but after re-programming and rotation of the fixed trigger locations. The *pink dot* indicates the location of the hidden-fixed trigger if mice anchored their memory to other allocentric cues, besides the arena cues. E. Same as in C and D, but after rotation of the visible-fixed trigger location. The *pink dot* indicates the position where mice would look for the hidden-fixed trigger if they remembered the position of the hidden-fixed trigger in relation to the visible-fixed trigger location. CW – clockwise, CCW – counter-clockwise.

3.4.1 Controlling for baseline performance

After completing the spatial memory training, I controlled for consistent and stable behavioural task performance. To achieve this, I calculated both the mean and the standard deviation of the trial count of the final three days of training on SM3 prior to the rotation experiments (155.21 ± 27.76 trials per training session). Once stable performance was established (Figure 3.22 B), I proceeded to the subsequent phase, in which I assessed the animal's response to arena cue and trigger location manipulations (Figure 3.22 A). This entailed rotating the arena cues (step 1, Figure 3.22 C), rotating all trigger locations (step 2, Figure

3.22 D), and rotating the arena cues and the visible, (but not the hidden) fixed trigger location (step 3, Figure 3.22 E). This set of experiments was then repeated with a different arena cue - trigger location mapping (SM1). I first confirmed stable baseline performance by training mice on the initial hidden-fixed trigger location, SM1, for four to five days, depending on the animal's individual learning speed and task performance. Again, a consistent performance baseline was confirmed for the last three days of this training stage (185.58 ± 45.61 trials per training session) before repeating the three arena cue and trigger location rotations (steps 1, 2, 3). Each of these steps consisted of a single training session (Figure 3.22 A).

3.4.2 Mice form allocentric spatial memories

When comparing baseline and rotated-arena-cue (RAC, step 1) conditions, I observed no discernible behavioural differences (Figure 3.23 A-B). I plotted the occupancy of individual mice during hidden-fixed trigger trials to determine whether mice searched for the hidden-fixed trigger in the rotated position that maintained its relation to the arena cues (Figure 3.23 C). However, I found that the occupancy hotspot did not shift with the arena cues (120-degree clockwise rotation) and instead remained at the original trigger location. This pattern held true for both repetitions, where mice consistently displayed a higher occupancy at the original SM3 and SM1 hidden-fixed trigger locations, but not at the rotated position. These findings indicate that spatial memories are not anchored spatially to arena cues.

Nevertheless, I did observe altered behaviour during rotated-trigger-location (RTL, step 2) conditions. I compared trajectories to the visible trigger locations (both random and fixed) with trajectories to the hidden-fixed trigger location. In particular, I used the mean instantaneous error angle from the start search point to trigger entry as a proxy for behavioural performance. The instantaneous error angle quantifies smaller left and right turns often observed during local searches, thereby providing a quantifiable proxy for hidden-fixed trigger searches in the incorrect trigger location. I found a significantly higher mean instantaneous error angle in trajectories towards hidden-fixed trigger locations in comparison to both visible-random and visible-fixed trigger locations (SM3: vr vs. hf: $p = 0.0011$ *; vf vs. hf: $p = 0.0003$ **; Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice; Figure 3.23 B, RTL). Although not statistically significant, I similarly observed increased instantaneous error angles during hidden-fixed trigger trajectories in the second repetition (SM1: vr vs. hf: $p = 0.003$ *; vf vs. hf: $p = 0.007$; Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice). These findings align with the occupancy hotspots of individual mice at the 'old', un-rotated trigger location for both the SM3 and SM1 trigger locations (Figure 3.23 C, RTL) and are consistent with an allocentric understanding of the arena and the anchoring of spatial memories to allocentric landmarks within the environment.

I identified a higher mean instantaneous error angle in hidden-fixed trigger trials that did not reach statistical significance in either repetition (SM3: vr vs. hf: $p = 0.0047$; vf vs. hf: $p = 0.0047$; SM1: vr vs. hf: $p = 0.065$; vf vs. hf: $p = 0.1605$; Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice) under conditions where the visible-fixed trigger location rotated (RVF) in accordance with the arena cues, while the hidden-fixed trigger location did not (Figure 3.23 A-C). While this could suggest that the visible-fixed trigger serves as a more important anchor than the arena cues, this difference in instantaneous error angle might be attributed to the confusion of the animals resulting from the previous day's altered placement of the hidden-fixed trigger zone during rotated trigger location (RTL) conditions. This interpretation is further

supported by the absence of an occupancy hotspot at the rotated trigger location (Figure 3.23 C, RVF), excluding the visible-fixed trigger as a primary anchor point for allocentric navigation.

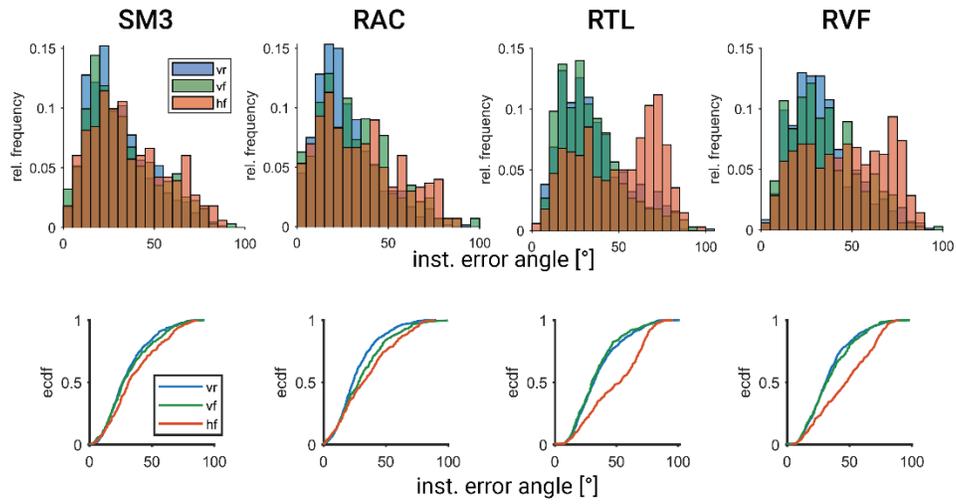
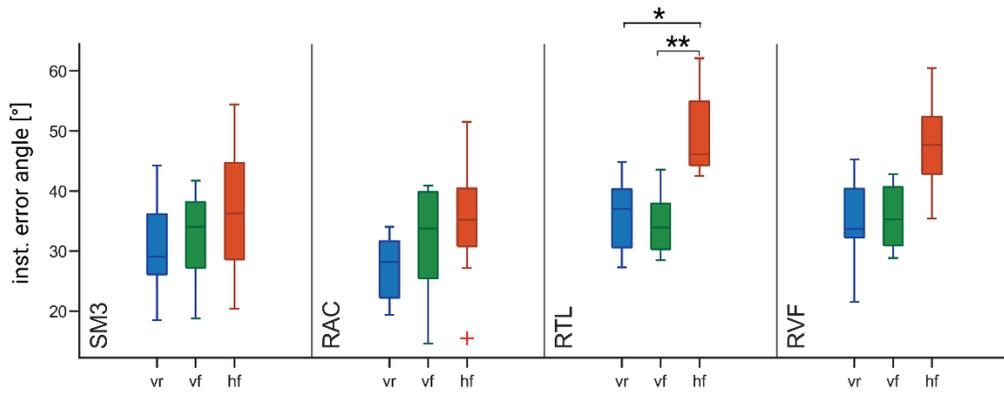
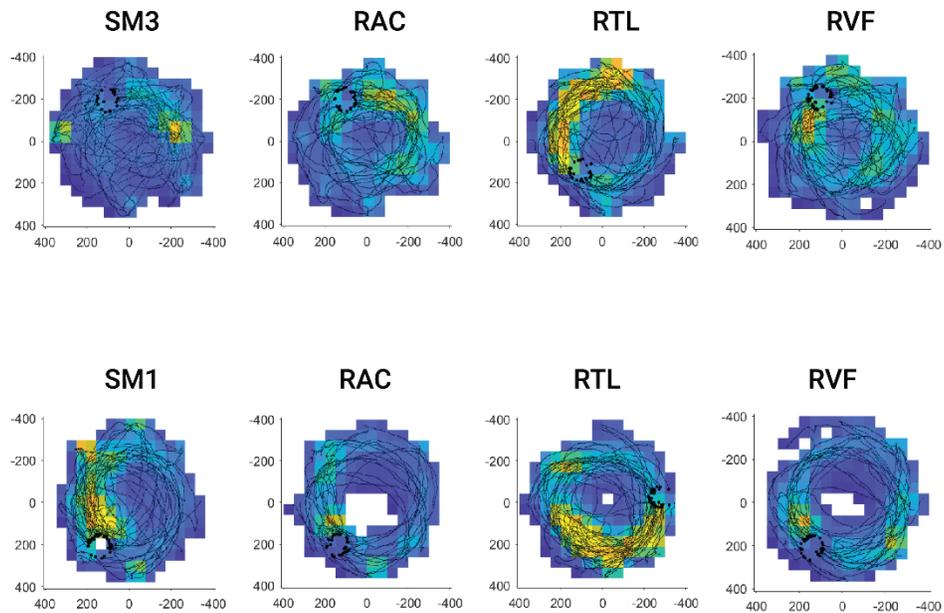
A**B****C**

Figure 3.23 Instantaneous error angle during arena cue and trigger location manipulations

A. Distribution and empirical cumulative distribution function (ecdf) of the mean instantaneous error angle for the three trial types (visible-random in *blue*, visible-fixed in *green*, and hidden-fixed in *red*) during baseline condition (here shown for SM3 training), rotated arena cues (RAC), rotated trigger locations (RTL), and rotated visible-fixed trigger location (RVF) conditions. B. Boxplot comparing the mean instantaneous error angle for all three trial types under baseline (SM3) and three manipulation (RAC, RTL, and RVF) conditions. C. Occupancy maps calculated from the trajectories during hidden-fixed trigger trials of an example mouse during baseline (SM3), RAC, RTL, and RVF conditions. *Wilcoxon rank sum test, Bonferroni correction, ** $p < 0.0008$, * $p < 0.004$. $n = 8$ mice.*

3.4.3 Summary

By rotating either trigger locations and/or arena cues, I could show that previously formed spatial orientations of the arena are not anchored to arena cues, but instead rotate with trigger locations. This supports the idea that spatial learning and memory are allocentric, ruling out egocentric spatial learning and path integration as possible navigation strategies. Instead, I hypothesized that mice rely on other external landmarks in this task to construct a cognitive map of their surrounding environment.

3.5 Mice form spatial memories within a single day

In a final step, I wanted to determine how fast mice can learn the location of a new hidden trigger to determine the speed of memory acquisition and reversal learning. In particular, I wanted to find out if mice can learn the location of a new hidden-fixed trigger location in a single 20-minute training session to investigate their within-day learning abilities.

To do so, each day, the visible-fixed and hidden-fixed trigger locations were randomly placed in the arena in a manner that avoided any overlap between these two fixed trigger locations. The arrangement and relation between these two fixed triggers remained consistent for the duration of each training session. Consequently, over the course of five consecutive days, mice were trained on five different fixed trigger zones – R-SM1, R-SM2, R-SM3, R-SM4 and R-SM5 (Figure 3.24 A).

Mice were able to learn the location of a new hidden-fixed trigger location within the span of a single training session. This is evidenced by occupancy hotspots that formed in each training session in the location where the hidden-fixed trigger location has been positioned on that particular day (Figure 3.24 B). Strikingly, during the five days of training involving random trigger placements, mice displayed an ability to recollect the previous day's hidden trigger location. This recollection was so pronounced that instances of two occupancy hotspots occasionally became evident (Figure 3.24 B, R-SM2 and R-SM5).

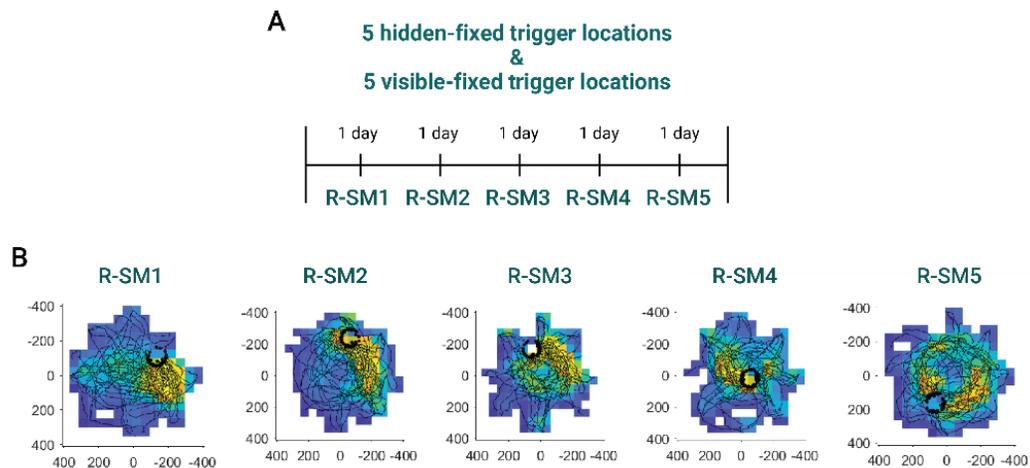


Figure 3.24 Timeline and occupancy maps during random trigger placement

A. Experimental timeline where mice had learnt the location of a new visible-fixed and hidden-fixed trigger location each day for five consecutive training days. B. Occupancy maps calculated from the trajectories of hidden-fixed trigger trials of an example mouse for the five random hidden-fixed trigger locations – R-SM1-5. Note, that in some instances an occupancy hotspot can be seen at the location of the previous day's hidden-trigger position (R-SM2 and R-SM5).

Additionally, the observed patterns of error angles observed previously during the spatial memory training were successfully replicated during this random-trigger-placement training phase. I find higher instantaneous error angle trajectories during hidden-fixed trigger trials compared to visible ones, a trend that was consistently significant from the third day of random trigger training until the fifth day (vr vs. hf: $p = 0.007$ (R-SM1), $p = 0.004$ (R-SM2), $p = 0.001$ * (R-SM3), $p = 0.003$ * (R-SM4), $p = 0.003$ * (R-SM5); vf vs. hf: $p = 0.021$ (R-SM1), $p = 0.004$ (R-SM2), $p = 0.001$ * (R-SM3), $p = 0.0006$ ** (R-SM4), $p = 0.0006$ ** (R-SM5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice; Figure S5 A). In contrast, global error angles exhibited a consistent, although not statistically significant, pattern of being overall lower during trajectories aimed at hidden-fixed triggers, in comparison to trajectories directed towards the visible triggers (vr vs. hf: $p = 0.19$ (R-SM1), $p = 0.32$ (R-SM2), $p = 0.51$ (R-SM3), $p = 0.06$ (R-SM4), $p = 0.03$ (R-SM5); vf vs. hf: $p = 0.19$ (R-SM1), $p = 0.05$ (R-SM2), $p = 0.13$ (R-SM3), $p = 0.16$ (R-SM4), $p = 0.23$ (R-SM5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice; Figure S5 B).

However, while mice were capable to remember a new trigger location each day, overall performance levels did not compare with those achieved during longer training periods, e.g. during spatial memory training when mice were trained on a single hidden trigger location for five consecutive training days. Furthermore, performance appeared to decline over the course of five random-trigger training days. This decline in performance manifested through the gradual increase in time required by the mice to locate the hidden-fixed, but not the visible triggers, as measured from the start search point to trigger entry (vr vs. hf: $p = 0.003$ * (R-SM1), $p = 0.004$ (R-SM2), $p = 0.0003$ ** (R-SM3), $p = 0.004$ (R-SM4), $p = 0.0006$ ** (R-SM5); vf vs. hf: $p = 0.007$ (R-SM1), $p = 0.002$ * (R-SM2), $p = 0.0002$ ** (R-SM3), $p = 0.002$ * (R-SM4), $p = 0.0002$ ** (R-SM5); Wilcoxon rank sum test, Bonferroni correction, $n=8$ mice; Figure 3.25).

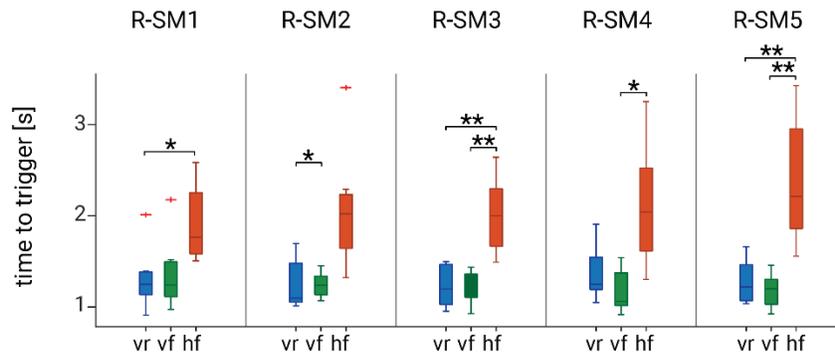


Figure 3.25 Task performance during random trigger training

A. Boxplot comparison of the time to trigger, measured from the start search point to trigger entry, for the three trial types (visible-random in *blue*, visible-fixed in *green*, and hidden-fixed in *red*) for the five days of random trigger placement training (R-SM1-5). *Wilcoxon rank sum test, Bonferroni correction, ** $p < 0.0006$, * $p < 0.003$. $n = 8$ mice.*

This suggests, that while mice can quickly learn the location of a newly hidden-fixed trigger zone in a single training session, their spatial memory may not be as strong or reliable compared to when animals undergo longer training periods. Nevertheless, this indicates the possibility of improving the timeline of spatial memory training by reducing the number of training days required for each hidden-fixed trigger location while increasing the number of reversal learning incidents that can be studied.

3.6 Summary

I designed a spatial memory task and built a behavioural setup that allows for the tracking of mouse position and rotation in an arena, and the quantification of behavioural task metrics in 3D with the help of a multi-camera 3D tracking system. In this spatial memory task, mice learn to memorize the fixed location of hidden trigger zones. Trials with hidden triggers are interleaved with trials with randomly placed, visible trigger zones, where mice need to use vision, instead of memory, to navigate to the trigger location. I find that mice use different strategies to navigate towards visible-random and hidden-fixed trigger zones. They approach hidden-fixed trigger locations with straight-forward (trigger-directed) trajectories and often perform a local search near the hidden trigger zone. This is in contrast to visible-random trigger trials, in which mice typically run curved trajectories towards the trigger zone, as quantified by global and instantaneous error angles. I find no significant differences in behaviour and task performance between control and experimental animals during chemogenetic inactivation of RSC neuronal activity. Upon introduction of a new trial type with a visible trigger zone that is fixed in location, to directly test how performance differs between memory-guided (hidden-fixed trigger zone) and vision-guided (visible-random trigger zone) trials, I find no significant differences to visible-random trigger trials, suggesting that mice prefer to rely on vision to locate the trigger zone when given the option between memory and vision. I show that mice use allocentric navigation strategies, anchoring their global internal map to surrounding, stable landmarks, but not to the arena cues, as demonstrated by a series of arena cue rotation and trigger location manipulations. Finally, I show that mice are capable of learning a new hidden-fixed trigger location within a single training session, even though task performance levels do not reach those of longer training protocols.

4. Discussion

Spatial cognition is a fundamental cognitive function crucial for survival, enabling mice and other species to navigate environments, locate resources, and avoid dangers. Despite extensive research into the neural circuits underpinning spatial cognition and learning through various spatial memory tasks, the precise role of the RSC in spatial memory remains elusive. This knowledge gap underscores the need for innovative behavioural paradigms to address previous spatial memory tasks' limitations. In this thesis, I introduce a novel behavioural paradigm aimed at investigating the contribution of the RSC to spatial learning and memory.

The developed task probes mechanisms of spatial memory formation, retrieval and reversal learning in rodents, wherein mice learn the locations of hidden trigger zones within a circular task arena. The setup allows dense tracking of animal behaviour in a 3D space and the detailed quantification of task performance, whereby key task parameters, such as error angles and active task engagement, were identified for assessing learning and task proficiency. Investigations revealed that mice form spatial memories even under conditions of RSC neuronal inactivation. Chemogenetic experiments were conducted to test RSC's role in acquiring spatial memory in this behavioural paradigm. Surprisingly, the results obtained indicate that RSC inactivation does not hinder task performance across memory acquisition, recall, and reversal learning phases. Furthermore, this work elucidates the circumstances wherein mice favour memory-guided navigation, highlighting a predominant reliance on visual cues over spatial memory when both are available. This underscores the complex interplay between navigational strategies and sensory information, emphasising the need for suitable task metrics for behaviour quantification. Through manipulations of visual cues and trigger placements, I demonstrated that spatial memories formed by mice are allocentric, anchored to stable environmental landmarks other than the arena cues. Moreover, mice exhibit the capability to form spatial memories within a single day.

Collectively, these findings underscore the task's suitability for investigating allocentric spatial learning and memory in mice. With rapid task proficiency acquisition, high trial rates, reliable memory formation and recall, short pre-training periods, and natural, unrestricted mouse behaviour, this paradigm opens new research avenues into the neural substrates of navigation, spatial learning, and memory.

4.1 Do mice form spatial memories?

The spatial memory task outlined in this thesis is well-suited for investigating spatial memory formation, recall, and reversal learning. Nevertheless, before employing this task, it is crucial to ascertain whether mice have indeed acquired spatial knowledge. I demonstrate the formation of spatial memories by identifying several key behavioural metrics, compensating for the absence of a conventional learning curve typically used to assess task proficiency. Additionally, as mice are unrestricted in their movement and behaviour within the arena during training sessions, I have identified the need to exclude trial times of task disengagement and present a possible means to quantify the active task engagement of freely behaving animals in this task. Moreover, I propose error angles as a method for quantifying different navigation strategies, such as those guided by memory or vision.

4.1.1 In search of a learning curve

Generally, in order to ascertain whether an animal has learnt to perform a given behavioural task, it is customary to analyse a learning curve. This learning curve typically takes the form of a sigmoidal function, depicting the gradual improvement of a specific task metric over time, usually measured in trials or training sessions. Depending on the behavioural paradigm, the learning curve is flat at the beginning of the training when animals have not yet acquired any memory, followed by a more or less steep improvement in task performance as learning progresses. The rate of improvement correlates with the speed of learning; hence, faster learning results in a steeper learning curve. In extreme cases of rapid learning, such as one-shot learning, where the animal learns within a single trial, the learning curve manifests as a step function. Once animals have acquired the task, the learning curve flattens out, whereby the resulting plateau indicates consistent and stable task performance.

In the context of the spatial memory task outlined in this thesis, the presence of a learning curve is less apparent. Conventional plotting of task metrics over training sessions (as exemplified in Figure 3.1) does not reveal a discernible learning curve. This absence of a well-defined learning curve poses challenges in defining the time point at which animals have learnt and acquired spatial memory. Consequently, questions arise regarding whether mice have indeed learnt and, if so, how this learning should be determined. In this study, I looked at different trial times, trial counts, path surplus, and occupancy to confirm that mice have formed spatial memories.

Trial times

Commonly used indicators of learning in other spatial memory tasks include trial times (Cain et al., 2006; Gawel et al., 2019; Paul et al., 2009). In the context of the spatial memory task discussed here, key time parameters include the time taken by an animal to locate the trigger location (time-to-trigger), the time taken to collect the reward from the port (time-to-port), and the cumulative trial duration, which represents the sum of these two metrics.

I anticipated to observe a learning curve and task improvement reflected in a gradual decrease in time-to-trigger, as mice are expected to become faster at locating the hidden trigger zone as they memorise its

location within the arena. Conversely, no significant improvement is expected in time-to-port, as this metric mainly reflects the response of the mice to a stimulus – that is, running to the port when the light signals reward availability. Similarly, the anticipated decrease in time to trigger should also manifest in trial duration, given that it encompasses both the time taken to find the trigger and the time needed to move to the reward port.

The development of a learning curve could occur at two stages within the training protocol. First, improvements might be observed over the five consecutive days of training on a specific hidden trigger location (e.g. within SM1 training). Secondly, one might expect a gradual enhancement throughout the spatial memory training, transitioning from SM1 to SM2, and finally to SM3, as animals become more proficient, and thus faster, at learning new hidden trigger zones.

Upon examining the time to trigger, no discernible learning curve is evident throughout the spatial memory training phase, nor within an individual training stage (as shown in Figure 3.1 B, and in Figure S2 B-C). This lack of a learning curve persists despite the expectation of visible task improvement reflected in time-to-trigger, particularly in hidden trigger trials, where mice must recall previously learnt spatial information. Conversely, no apparent learning curve is observed in time-to-port across the training stages. Notably, a subtle improvement is detected in trial duration over individual training days (day 1 to day 5 of a given training stage, such as SM1, or SM3), specifically in hidden trigger trials (Figure 3.20). This improvement provides evidence that mice have indeed learnt the location of the hidden trigger zone, despite the lacking learning curve. Additionally, task improvement and a trial duration reduction are observed throughout the training stages (from SM1 to SM3), albeit not consistently clear (Figure S1 A).

In contrast, most of these task metrics exhibit a learning curve during the pre-training phase, where mice are familiarised with the task structure. In this initial stage, spatial learning is not involved as all trigger zones are visible and indicated to the mouse, prompting for simple stimulus-response behaviour. This task structure learning curve is evidenced by the decrease in trial duration (Figure S1 A), time-to-trigger (Figure 3.1 B), and time-to-port (Figure S1 B) throughout the pre-training phase. Finally, pre-training performance plateaus, reflecting consistent, stable baseline performance prior to the start of the actual spatial memory training (SM1, SM2, and SM3 training). Note, that during task structure learning, improvement in time-to-port is anticipated, as learning in this stage pertains to mice learning to retrieve the reward from the port when the light is illuminated, rather than spatial learning itself.

Together, these findings suggest that spatial learning and memory formation manifest subtler than task structure learning, and require the use of further task parameters to determine spatial learning.

Trial counts

The progression of task structure learning becomes most apparent when examining the count of correct and error trials within a specific timeframe, such as a training session of a pre-determined duration. Using trial counts as a proxy for task performance is valid in this spatial memory task as a daily training session is defined by a set duration (i.e. 20 minutes), rather than terminating a training session once a certain number of correct trials or a specified amount of reward has been collected.

A subtle learning curve is observable within the training stages when analysing the average trial count across mice, as indicated by a small increase in trial count over the training stages (over SM1, SM2, and SM3 training), albeit failing to reach statistical significance (Figure 3.1 A). Although this learning curve is even less discernible when examining the correct trial counts of individual mice (Figure S2a, Figure 3.18 A-B), it offers preliminary evidence that animals have learnt and formed spatial memories.

Similar to previously discussed observations, task structure learning (i.e. improvement over the pre-training phase) is clearly visible. This phenomenon is best illustrated by plotting the different error trial types and their occurrence within each daily training session (Figure 3.3, Figure 3.18 C-D). Initially, a moderate number of trials resulting in timeouts is evident at the onset of the pre-training, indicating that mice had yet to grasp the general task structure and failed to visit the trigger zones. As mice familiarised themselves with traversing the trigger zones, the number of error trials where reward collection timed out increased, reflecting the animals' ongoing learning of the association between finding the triggers and the subsequent reward retrieval at baited ports. Towards the latter part of the pre-training phase, error trials in which animals request a reward from an unbaited port increase, signifying the mice's need to learn that only the port with the LED light turned on holds a reward. Finally, at the end of the pre-training, error trial counts remain low and stable.

However, following pre-training, during spatial memory training, stable counts of both correct and error trials emerge, substantiating that mice have indeed acquired spatial memories but lacking a learning curve.

Path surplus

In search for a task metric that unequivocally quantifies spatial learning in this task, I looked beyond conventional timing metrics and trial counts typically employed in similar studies. In this study, path surplus is defined as the ratio of the distance covered by the mouse to the shortest possible path required to reach the trigger zone. Unlike time-to-trigger, path surplus provides a more direct, time-independent measure, excluding inactive periods where the animal may be stationary.

The analysis of path surplus revealed further evidence for spatial memory acquisition. Across the five-day learning period within a particular training stage, noticeable task improvements were observed. This progression is exemplified by the comparison between the initial and final day of training on stages SM1 and SM3 (depicted in Figure S4). As anticipated, an improvement was also evident across different training stages, where mice exhibited reduced path surplus on the first day of training on SM3 compared to the first training day on SM1 (Figure S4), bolstering the argument for spatial memory acquisition in this task.

Occupancy

In many spatial memory tasks, it is customary to examine the duration animals spend in proximity to or directly at the goal location that needs to be remembered and located (Gawel et al., 2019; Lukoyanov et al., 2005; Paul et al., 2009; Vann et al., 2003).

Given that traditional trial time measurements and trial counts failed to produce a discernible learning curve in my analysis, I turned to occupancy as a qualitative indicator of task proficiency (see Figure 3.2 and Figure 3.19). By plotting occupancy maps, a visual “learning curve” emerged across the training days of the spatial memory training phase. Spatial learning becomes apparent as occupancy hotspots emerge over the training days (days 1-5), shift to new goal locations during reversal learning (SM1 to SM2, and SM2 to SM3), and gradually refine into smaller, more precise occupancy clusters with continued training (later days of a training stage). These observations provide a clear, albeit visually interpreted and challenging-to-quantify, demonstration of spatial memory acquisition within the behavioural paradigm presented in this thesis.

While identifying the precise moment at which mice achieve task proficiency is challenging, the combination of the quantitative assessment of the task parameters discussed, alongside the qualitative assessment of occupancy maps, provides compelling evidence for the establishment of spatial memories by mice in this task.

4.1.2 Identifying task engagement in freely behaving animals

In experimental contexts where animals are head-fixed during spatial memory tasks, and particularly in behavioural paradigms where animals are unrestricted in their movement and behaviour, accounting for periods of task disengagement becomes crucial for accurate behaviour and task metric quantification. When mice are head-fixed and running on an air-floating Styrofoam ball, task disengagement is typically detected through animal immobility, where the animal remains stationary on the ball. This immobility can be quantified by tracking the movement of the Styrofoam ball, serving as a proxy for running behaviour (Fischer et al., 2020). However, in freely behaving tasks, using general immobility as a proxy for task disengagement is less effective, as animals may explore their surroundings while still not actively participating in the task. This complicates the establishment of clear task metric thresholds that delineate active task engagement from disengagement within a trial.

In the described task, mice occasionally move away from the port after reward consumption, possibly to consume their reward or to engage in grooming activities. This behaviour is characterised by minimal movement in the x, y, and z coordinates following reward delivery (see Figure 3.4) and general, low running velocity (see Figure 3.4 and Figure 3.5). These periods of task disengagement can last several (Kolb et al., 1983) seconds, a significant duration considering that some trials themselves last only a few seconds (see Figure S1 A and Figure 3.20), and can sometimes encompass the majority of the trial duration (as illustrated in Figure 3.5 B, where task disengagement accounts for more than half of the trial duration).

To address these prolonged periods of task disengagement, during which mice may be grooming or freely exploring the arena, I utilised start search points to define active task engagement. These time points aim to capture the moment when mice transition from grooming or other non-task-related behaviours to actively searching for the trigger location (see Figure 3.5 A).

Start search points are defined as the last point in time where the animal reaches the 5th percentile of running velocity along the trajectory towards the trigger location. However, the static definition of the

start search threshold can occasionally fail to accurately capture the exact moment of active task engagement. For instance, in Figure 3.5 B, the start search point along the path to the hidden trigger is placed too late in time, thereby capturing the mouse mid-search rather than at the onset of active engagement, which results in the exclusion of a significant portion of a local search phase from the analysis.

Despite the limitations associated with the static definition of the start search point, the necessity for defining active task engagement becomes evident when examining trial time measurements. While the comparison of time-to-trigger measured from the end of the previous trial and from the start search point (Figure 3.6 A-B) does not reveal a discernible learning curve over spatial memory training (SM1, SM2, and SM3), the use of start search points allows for more realistic time measurements of task performance. Particularly in the case of the time-to-trigger (Figure 3.6 B), it becomes evident how quickly (within a few seconds) animals can locate both visible and hidden trigger zones, validating this method of quantifying task engagement.

4.1.3 Quantifying navigation strategies

When analysing and quantifying trajectories, the task metrics previously discussed, such as time measurements or occupancy maps, are inadequate for capturing nuanced differences in trajectories or identifying distinct navigation strategies. For example, time-to-trigger fails to distinguish between trials involving hidden triggers, where mice rely on memory to recall the trigger zone's location, and trials involving visible triggers, where mice follow a light beacon projected into the arena. This inability to distinguish between trial types persists regardless of whether active task engagement is considered (see Figure 3.6). Similarly, occupancy maps are insufficient for quantifying running trajectories.

However, navigation guided by vision (e.g. visible trigger trials) versus memory (e.g. hidden trigger trials) exhibits notable differences. These differences become apparent when assessing the precision with which animals approach visible and hidden trigger zones. The phenomenon of overshooting trigger zones can be quantified by measuring the distance from the trigger zone centre to the start search point of active port search (see Figure 3.7 A), which is higher in hidden trigger trials compared to visible trigger trials (see Figure 3.7 C), thereby highlighting differences in navigation guided by memory or vision. I hypothesised that the absence of visual feedback in hidden trigger trials leads mice to overshoot the trigger zones. Conversely, the visual feedback of the light projection enables mice to target the centre of visible trigger zones more precisely. It is important to note that overshooting occurs in both visible and hidden trigger trials, as exemplified in Figure 3.5 where the distance between the trigger centre (denoted by the circle) and the start search point of the trajectory leading to the port (denoted by the red asterisk) is higher in the visible trigger example (A) compared to the hidden trigger example (B). Albeit, the significant difference in trigger precision suggests distinct navigation strategies in hidden and visible trigger trials, highlighting the importance of quantifying these strategies with appropriate task metrics.

I further hypothesised that mice tend to follow curved trajectories to visible triggers, as the visual feedback permits for path adjustments at any point. In contrast, during hidden trigger trials where animals need to recall memory, mice are more likely to run along straighter paths followed by a local search (see

Figure 3.11 for a schematic depicting the different running trajectories, and Figure 3.5 for example trajectories). To discern these navigation differences more effectively, I employed error angles, as done in previous studies (Cain et al., 2006; Save, 1997; Valerio & Taube, 2012). To quantify trajectories leading to hidden and visible triggers, I defined two types of error angles - global and instantaneous - to capture both larger and smaller deviations from a straight trajectory, respectively (see Figure 3.8 A). If a mouse begins a trial by moving away from the trigger zone, the global error angle remains high until the animal reaches the trigger zone, while the instantaneous error angle detects path corrections faster and decreases as the animal starts facing the trigger zone. The colour-coded trajectory examples of a hidden trigger trial in Figure 3.8 B illustrate how the instantaneous error captures the local search near the hidden trigger zone better than the global error angle, remaining relatively low throughout the local search, highlighting the trajectory differences these two types of error angle unveil.

The hypothesis of mice adopting curved trajectories towards visible triggers and straight trajectories towards hidden triggers is supported by significantly higher global error angles in visible trigger zones and significantly higher instantaneous error angles in hidden trigger trials (see Figure 3.9 and Figure 3.21). Mice can afford to take curved paths when visual feedback is available, allowing for course corrections at any time. Conversely, they are likely to run straighter paths in hidden trigger trials, where reliance on memory recall is necessary. Support for this hypothesis comes from the reproducibility of these error angle patterns across the two cohorts, suggesting that error angles are a suitable measure for characterising trajectories and distinguishing between goal-directed navigation guided by memory or vision.

For control purposes, error angles were also calculated along trajectories leading to the reward port. Here, I hypothesised that error angles should be similar regardless of whether a visible or hidden trigger had to be found previously, as mice must rely on their vision to locate the reward port with the LED light turned on. However, it is worth noting that in hidden-fixed trigger trials, a small memory component remains, as only two reward ports are possible following a hidden-fixed trigger location (see Methods - 2.4.9 Reward port selection). Nonetheless, these ports are on opposing sides of the arena and randomly chosen on a trial-by-trial basis, making it impossible for mice to predict which port will deliver the next reward. While differences in error angle patterns between visible-random and hidden-fixed trigger trials cannot be replicated along the trajectories leading to the reward port, in some instances, significantly higher global errors were also detected in trajectories leading to the port following a visible trigger (see Figure 3.9 B, Figure 3.21).

I speculated that these differences detected in a few instances when measuring the global, but not the instantaneous error angle, could result from a continued curved trajectory from the visible trigger zone to the reward port. To test this hypothesis, I examined whether similar results could be obtained using a comparable, yet different metric, such as angular head velocity. Angular head velocity should be higher in curved trajectories than straight ones, thereby providing a similar metric to the global error angle. Using this metric, I find support for my hypothesis of continued curved trajectories following visible triggers, as once again, significantly higher angular head velocity is measured along the running trajectory to the port following visible triggers only (see Figure 3.10).

In summary, commonly used time measurements are inadequate for quantifying trajectories and navigation strategies used in this task. Error angles prove to be a more nuanced measure of task performance when comparing visible-random and hidden-fixed trigger trajectories than time measurements. For instance, the time taken by mice to locate either the visible-random or hidden-fixed trigger zone from the start search point does not capture this difference in trajectory and behaviour as effectively as error angles (see Figure 3.1 B, Figure S2 B-C). Although I expected to observe a clear difference between visible and hidden trigger trials in terms of time-to-trigger, with the former involving a simple stimulus-response and the latter requiring memory recall, this difference was not evident. It is unlikely that task disengagement influenced these time measurements to the extent that the time taken to find the trigger was prolonged to the point where small differences between visible and hidden trigger trials were no longer statistically detectable. In fact, this is not the case, as there is no difference between trial types when considering the time-to-trigger measured from the start search point (see Figure S3). This leads to the conclusion that error angles provide a superior measure for quantifying navigation strategies compared to conventional time measurements.

4.1.4 Vision vs. memory

Using error angles to quantify navigation strategies, I could show that mice only use memory-guided navigation in the absence of visual information. Support for this claim comes from the comparable performance between visible-random and visible-fixed trigger trials across various metrics, including trial duration (Figure 3.20), error angles (Figure 3.21), and path surplus (Figure S4). Additionally, during arena cue and trigger rotations (Figure 3.23), as well as during random trigger training (Figure 3.25), mice exhibit similar behavioural responses and running trajectories irrespective of whether the visible trigger is random or fixed in position. From these observations, it can be inferred that mice predominantly utilise vision rather than memory to locate the visible-fixed trigger trials.

The introduction of a visible trigger fixed in position within the arena provides mice with the opportunity to memorise its location instead of relying solely on visual cues. However, it is important to acknowledge that it was not experimentally tested whether mice were aware of the spatial stability of the visible-fixed trigger zones. This aspect could be addressed experimentally by making the visible-fixed trigger zone invisible at the end of the training protocol. Should mice retain the location of the visible-fixed trigger zone in their memory, the animals would display increased searching in the visible-fixed trigger zone which would manifest as a hotspot in the occupancy map.

4.2 How fast do mice form spatial memories?

While most behavioural paradigms utilise learning curves to determine the speed of task proficiency, the spatial memory task described here raises questions about the rate at which animals learn and form spatial memories. In this section, I will discuss the experimental evidence for within-day learning in the spatial memory task, explore potential enhancements to the task, and outline the comparative advantages it offers over other commonly used spatial memory tasks.

4.2.1 Evidence for within-day learning

Evidence supporting within-day learning, defined as task proficiency and memory consolidation occurring within a single training session, comes from experiments where mice had to learn a new location each training session, as the location of the hidden-fixed trigger zone changed each day for five consecutive training days (see Figure 3.24 A).

Analysis of the animal's occupancy during hidden trigger trials across these five random trigger locations (Figure 3.24 B) reveals the rapid formation of occupancy hotspots at the new hidden trigger location within a single 20-minute training session. This qualitative evidence suggests memory acquisition within a day. Furthermore, in some instances, an additional occupancy hotspot was observed at the trigger location from the previous training day. The presence of two occupancy hotspots in a single training session visually demonstrates that mice not only recall the previous day's hidden trigger location, indicating recent spatial memory recall, but also acquire a new spatial memory for the current day's hidden trigger location, indicative of reversal learning and spatial memory formation – all within a 20-minute training period.

It is worth noting, however, that performance was worse compared to instances where animals were trained for longer durations on a given trigger location, as evidenced by progressively increasing times to locate the hidden trigger across the five random locations, implying a gradual exhaustion of task performance (see Figure 3.25). This decline in performance can also be quantified by examining the error angles (see Figure S5). While consistently higher instantaneous error angles were observed in hidden trigger trials and consistently higher global error angles in visible trigger trials, significant differences were only evident on a few select days of random trigger training. In summary, comparable error angle patterns were observed as with earlier training periods during the spatial memory training (see Figure 3.21), providing further evidence for spatial memory formation within a single training session and within-day learning. However, these significance patterns were not as consistent as with longer training periods.

When reviewing the experimental evidence pointing towards within-day learning, it is important to consider that this series of experiments was conducted after animals have been trained to learn the location of three prior hidden trigger locations over an extended period during the spatial memory training (that is the hidden trigger placement in SM1, SM2, and SM3). Additionally, animals underwent a series of arena cues and trigger manipulations before starting the random trigger training (see Figure 2.1 A). Consequently, it is plausible that animals may have been disoriented by these preceding experiments, which could potentially account for the overall lower task performance. It would thus be intriguing to

observe how animals perform when presented with a new hidden trigger location each day, starting this random trigger training immediately after the pre-training phase.

However, despite this potential spatial confusion induced by previous experiments, the presence of occupancy hotspots at both the current and previous training day's hidden trigger locations (see Figure 3.24 B) indicates that spatial memory recall, formation, and reversal learning are happening within a single training session during random trigger training. This suggests the possibility of refining the experimental timeline of the spatial memory training by reducing the number of training days allocated to each hidden trigger location (e.g., only two to three days of training on the hidden trigger location in SM1, SM2, and SM3). Such a timeline refinement would not only decrease the overall experimental duration but also increase the number of hidden trigger locations a mouse can learn throughout the training protocol, as well as the number of reversal learning incidents that can be investigated using this task with a single cohort of mice.

4.2.2 Task benefits

Learning the hidden trigger location in this task is not only rapid but also efficient across the entire experimental timeline. Animals require only about one week of pre-training to familiarize themselves with the task structure. Following this short pre-training period, animals quickly achieve stable and reliable task proficiency, managing up to 250 trials within a daily 20-minute training session. This efficiency in forming and reliably recalling spatial memories contrasts with other spatial memory tasks, in which animals typically require considerably longer to learn the goal location (Sunyer et al., 2007), along with additional drawbacks.

An additional, significant advantage of the task outlined here is the separation of the reward location from the target learning site. Unlike other tasks, such as the radial arm maze or those designed as dry versions of the Morris water maze (Morales et al., 2020), where the reward location overlaps with the learnt location within the arena, the task described here is only indirectly reward-guided. Instead, it primarily relies on mice recalling their memory or visually orient themselves on environmental landmarks to find the hidden trigger zone, while rewards are delivered at a separate reward port. This setup allows for a clearer differentiation between neuronal coding for spatial correlates and reward.

Although the radial arm maze also permits natural, unrestricted animal movement and behaviour, its design limits animals to static trajectories due to the unchanging layout of the arms. Furthermore, originally developed to assess spatial working memory, the radial arm maze is less effective for studies of long-term spatial memory retention.

The feasibility of within-day learning of the hidden trigger location in this task is likely due to the brief duration of the trials, which enables a high volume of trials per session. Unlike the Morris water maze, where mice must swim — an activity that can induce significant stress — the task described in this thesis involves running, which is less stressful and more natural for the mice. This, and the possible application of neuronal recording devices during training, makes the spatial memory task a preferable alternative to water-based mazes.

Perhaps the most significant advantage of this task is that it allows for natural, unrestricted mouse behaviour, making it ideal for studying navigation, spatial exploration and spatial learning in a naturalistic setting. In contrast to head-fixed VR tasks that restrict proprioception and head movement, this spatial memory task supports natural head movements and proprioceptive feedback, while allowing for the simultaneous application of portable recording devices.

Taken together, this task is ideally suited for investigating allocentric, long-term spatial learning and memory in rodents. It offers a robust platform for exploring the neuronal mechanisms underlying navigation and spatial learning within the RSC and likely other brain regions within the broader spatial network.

4.3 Are spatial memories allocentric?

Mice trained on the spatial memory task develop a cognitive map of the task arena and the space around them and therefore have a global, spatial understanding of their environment. This implies that the spatial memories they form regarding the hidden trigger zones are allocentric rather than egocentric. Consequently, egocentric navigation strategies, such as path integration, can be ruled out as potential strategies employed by the mice trained in this task.

If spatial memories were egocentric, the location of the hidden trigger would be encoded based on the animal's self-centred perspective. Mice would rely on path integration to locate the hidden-fixed trigger position, recalling sequences of left and right turns as well as forward steps leading to the hidden-fixed trigger. This navigation process would largely disregard visual landmarks since mice do not require a comprehensive global, 'north-south, east-west' understanding of the arena for locating the hidden-fixed trigger. Conversely, allocentric coding involves the use of global landmarks and viewpoints. Mice construct a global map of their surroundings by correlating a set of visual landmarks, enabling them to determine the hidden-fixed trigger's position in space.

In this section, I discuss how the task structure design itself prevents path integration as a possible navigation strategy to locate hidden trigger zones. Additionally, I examine the experimental evidence that points toward allocentric spatial memories. Finally, I explore potential explanations regarding which environmental features could be used by the animals as spatial reference points to anchor their allocentric spatial memories.

4.3.1 Task design prevents path integration

The selection of baited reward ports following a randomly placed visible trigger zone effectively prevents path integration as a workable navigation strategy.

In both cohorts trained on the task in this study, each memory-guided trial with a hidden-fixed trigger zone is followed by a trial with a randomly placed visible trigger zone. The choice of which port is baited after a mouse encounters a trigger zone – regardless of whether this trigger is visible-random, visible-fixed, or hidden-fixed – is determined in a manner that allows only two possible ports to be baited for any given trigger location. These two ports are the ones whose distance to the trigger zone is closest to the arena radius, while maintaining a minimum distance of 40 cm from the trigger location (see Figure 2.6 for a schematic of the reward port selection). This design ensures that for any given trigger location, such as hidden-fixed or visible-fixed triggers, only two reward ports can possibly be baited. However, since each trial with a fixed trigger zone is succeeded by a trial with a visible-random trigger zone (see Figure 2.5), which subsequently triggers the baiting of an equally semi-random reward port, mice commence a hidden trigger trial from a random starting point, i.e. from the reward port from the preceding visible-random trigger trial. This variation in starting points renders the counting of steps in a certain direction ineffective to find the hidden trigger zone, thus making path integration void.

Consequently, the task design itself encourages the formation of a global map of the arena and necessitates the use of allocentric navigation strategies to locate the hidden trigger zone from memory.

4.3.2 Experimental evidence for allocentric memory

Experimental evidence that the spatial memories mice formed during task training are allocentric, stems from a series of arena cue and trigger location manipulations.

Throughout the entire training protocol, mice were provided with both proximal arena cues affixed to the arena wall and distal arena cues suspended from the arena ceiling (Figure 2.4). These arena cues are meant to serve the animals as visual landmarks in an otherwise radially symmetrical environment. Although the ceiling of the arena is not symmetrical due to the placement of cameras and other setup equipment (Figure 2.3 A), the design of the behavioural setting aimed to create radial symmetry throughout the rest of the arena. To assess whether mice use these arena cues as landmarks, the spatial relation of these cues to the learnt hidden trigger location was disrupted, which, it is assumed, would in turn also disrupt their sense of orientation, if mice were indeed to use the arena cues as spatial anchorage points.

To assess and quantify the animal's behaviour in response to the arena cue and trigger manipulations, I used the instantaneous error angle. This measure detects smaller left/right turns and produces increased values in response to local searches in the wrong location, making it the appropriate task metric in this set of experiments. Additionally, occupancy maps were employed to qualitatively identify hotspots where mice predominantly searched for the trigger zone.

During "step 1: 120-degree clockwise rotation of the arena cues" (Figure 3.22 C), the physical location of the hidden-fixed trigger remained unchanged. However, its relation to the arena cues changed. If mice anchored their global memory of the hidden trigger location to the arena cues, animals would search for the trigger zone in a location corresponding to a 120-degree clockwise rotation. However, no occupancy hotspot formed at the rotated location (pink dot in Figure 3.22 C), indicating that spatial memories were not spatially anchored to the arena cues. Unlike scenarios where mice encounter a newly placed hidden trigger location each day, resulting in the formation of two occupancy hotspots indicative of fast reversal learning happening within a single session, no such occurrence was observed here. This suggests that mice did not initially search in the 'rotated' location and, instead, searched for the hidden trigger in the correct location. Furthermore, there was no significant increase in instantaneous error angles on that training day, which could be interpreted as a sign of increased searching in the vicinity of the 'rotated' trigger zone, providing further support that mice do not use the arena cues as spatial reference points.

In "step 2: 120-degree counter-clockwise rotation of the trigger locations" (Figure 3.22 D), the fixed trigger locations were re-programmed, resulting in higher instantaneous error angles in hidden-fixed trigger trials, as well as two occupancy hotspots forming at both the previous and current day's trigger locations. This implies that animals predominantly searched for the hidden trigger in the same physical location as during baseline performance, suggesting that mice anchor their spatial memories to allocentric landmarks other than the arena cues.

In both scenarios – clockwise rotation of the arena cues in step 1 (Figure 3.22 C) and counter-clockwise rotation of the trigger locations in step 2 (Figure 3.22 D) – the spatial relation between the arena cues and the fixed trigger positions is the same. That holds true for both the newly positioned, correct trigger location and the 'old', incorrect trigger location. However, while the physical environment (i.e. the arena cues) rotated in one scenario, the programmed environment (i.e. the projected fixed trigger zones) rotated in the other. These manipulations allow for exploration whether mice form any global, allocentric understanding of their environment, whether these spatial memories are spatially referenced to the environment, and, if so, whether the arena cues are utilised by the animals to spatially reference the learnt hidden trigger location

Finally, to conclusively rule out the possibility of mice using the visible-fixed trigger location as a spatial reference point for locating the hidden-fixed trigger by rapidly adjusting their sense of orientation in accordance with the visible-fixed trigger location, a final experiment was conducted: "step 3: rotation of the arena cues and the visible-fixed trigger location" (Figure 3.22 D). Similar to step 1 (rotation of the arena cues), if mice construct their global map primarily from the arena cues, they would be expected to search for the hidden trigger location at a site also rotated 120 degrees clockwise, maintaining the spatial relation of the arena cues and the hidden-fixed trigger location. Conversely, if mice utilise the visible-fixed trigger as a spatial anchor point, their memory of the hidden-fixed trigger location would undergo a counter-clockwise rotation of 120 degrees, rotating along with the visible-fixed trigger location. Finally, if mice were to use environmental landmarks other than the arena cues as spatial reference points, animals would continue searching at the same location as during baseline performance. Through this experimental approach it is possible to determine whether mice rely on the visible-fixed trigger as their primary reference point for locating the hidden trigger. The absence of an occupancy hotspot in a rotated location – clockwise, or anti-clockwise – demonstrates that neither the arena cues, nor the visible-fixed trigger zone are used by mice as reference points. Despite the higher instantaneous error observed in hidden-fixed trigger trials compared to the visible trigger trials in this rotation experiment, this difference does not reach statistical significance and is likely attributed to the animals' confusion stemming from the preceding rotation step on the previous training day. Here, in step 2, mice indeed re-learnt the location of the hidden trigger zone, as its physical position changed. As a result, animals, utilising allocentric landmark anchoring, had to re-learn the location of the hidden trigger the day before. In addition, as discussed earlier, there exists no experimental evidence that mice are even aware that the visible-fixed trigger location remains stable in position, making it unlikely that mice use the visible-fixed trigger as reference point.

Together, these rotation experiments provide experimental evidence that mice have formed allocentric spatial memories, using stable landmarks in their environment as spatial reference points, excluding the arena cues.

4.3.3 Why are allocentric memories not anchored to arena cues?

Given that mice do not utilise the intentionally provided arena cues as global landmarks, yet form spatial memories through allocentric navigation strategies, there must be other spatial features in the animals'

environment that serve as stable anchor points for the animals' spatial orientation. For these environmental landmarks to serve as effective reference points, they must be spatially stable but can manifest in various sensory forms. Several potential confounders might deter mice from using arena cues for global orientation: potential auditory, visual, and olfactory reference points.

The arena where mice undergo spatial memory task training is not enclosed in a soundproof box. Despite efforts to maintain quiet experimental conditions, various sounds and auditory inputs may reach the mice during training, including sound coming from the PCs and the air-conditioning system in the setup room, noise from other mice or the experimenter, and the click sound of individual reward port valves, to name a few. These auditory stimuli could serve as spatial reference points if distinguishable by the mice.

The arena features both proximal arena cues affixed to the arena wall and distal arena cues suspended from the arena ceiling (Figure 2.4). While the design of the behavioural task arena aimed for radial symmetry, certain visual features, such as cameras and equipment placement on the ceiling, are asymmetrical. Mice may utilise these visual cues on the ceiling for orientation, instead of the arena cues (Qiu et al., 2021). Additionally, the consistent entry point from which mice are released into the arena may serve as a visual orientation guide.

In addition, it is plausible that mice use olfactory cues for spatial orientation during spatial memory task training (Hindley et al., 2014a; Muysers et al., 2024). Although the arena is cleaned between different animals' training sessions to minimise odours, residual olfactory cues and the experimenter's scent could serve as spatial anchor points. To further ensure olfactory variability across training sessions, the physical location within the arena changes for each mouse. This is achieved by rotating the trigger location pattern by a rotation factor that is unique for each animal while maintaining a consistent spatial relationship of fixed trigger locations and their distance to the arena wall (Figure 2.7).

Since it remains uncertain whether mice utilise these sensory cues for spatial orientation – auditory, visual, or olfactory – conducting experiments in a sound-proof and odourless arena with a covered ceiling would provide valuable insights into the spatial reference points used by the animals trained on the spatial memory task. Moreover, releasing animals from different entry points or utilising a covering to prevent prior visually-guided spatial orientation could further elucidate the role of these sensory cues in spatial orientation. However, considering the task's aim of mice acquiring spatial memories through allocentric navigation strategies, the necessity of identifying the exact landmarks used for orientation may be questioned depending on the question being asked in a given experiment.

4.4 Is RSC involved in spatial memory formation?

While the expression of inhibitory DREADDs in the RSC leads to neuronal inactivation under CNO conditions, mice trained on the task exhibited no spatial deficits, neither during the acquisition and retrieval of spatial memories nor during reversal learning. This raises several questions: Is RSC activity essential for spatial learning and/or memory within this specific spatial memory task? Was neuronal activity chemogenetically inactivated? This section discusses the experimental evidence confirming RSC neuronal inactivation, explores potential reasons for the lack of observable behavioural deficits, and considers unanswered questions about RSC function that this behavioural paradigm could address.

4.4.1 Experimental evidence for chemogenetic RSC inactivation

In the chemogenetic experiments, the virus expressing inhibitory DREADDs was administered at twelve injection sites: three injection sites per hemisphere, at two different depths per injection site (see Figure 2.2). These coordinates were chosen to ensure the viral targeting of neurons in both the dorsally located dysgranular and the superficial granular subregions of RSC, as well as most RSC neurons along the rostro-caudal axis (Figure S6). Functional inactivation was confirmed through Ca²⁺ imaging before the onset of the behavioural training (Figure 2.1 B, Figure 3.12), showing a reduction in cellular activity when CNO was administered (Figure 3.14).

However, it is possible that not all neurons in RSC were fully silenced, and residual activity might have been sufficient for memory formation and retrieval. Moreover, functional proof of cellular inactivation under CNO conditions via two-photon imaging was limited to the habituation period before actual spatial memory training began, leaving open the possibility that RSC was not (or, not sufficiently) inactivated during the spatial learning phase.

4.4.2 Potential involvement of RSC in spatial memory formation

Despite the functional evidence supporting chemogenetic inhibition of RSC activity in the habituation phase, no behavioural differences were observed in experimental animals under CNO conditions during spatial memory training. Animals that received CNO injections prior to training sessions displayed similar task performance to control animals in terms of trial count (Figure 3.15 A-B), error trials (Figure 3.15 C-D), arena occupancy (Figure 3.16), and error angles (Figure 3.17). This absence of behavioural differences could be attributed to a potentially incomplete inactivation of all RSC neurons, inactivation occurring during the habituation phase but not during the spatial memory training, or it may suggest that RSC activity is not critical for mastering this spatial memory task. Alternatively, compensatory mechanisms in other brain regions may enable spatial memory formation and retrieval in the context of this behavioural paradigm. Investigating these possibilities requires re-visiting RSC's potential role in navigation and spatial learning.

The RSC is implicated in egocentric spatial information processing, which comprises neurons sensitive to a range of movement-related variables, such as self-motion and linear and angular movement speed (Alexander & Nitz, 2015; Chen et al., 1994c; Cho & Sharp, 2001; Keshavarzi et al., 2022). Consequently, it

is proposed that the RSC contributes to path integration – a navigation process where animals continuously track and integrate changes in linear and angular displacement to self-locate. RSC inactivation or lesioning leads to navigational deficits in conditions requiring path integration, such as darkness (Cooper et al., 2001; Elduayen & Save, 2014), as well as in tasks relying on self-motion cues (Elduayen & Save, 2014; Whishaw et al., 2001). Given that the task described here requires allocentric navigation, relying not on path integration but instead on the spatial configuration of stable landmarks in the environment, this could explain the preserved ability of mice to perform on this task despite RSC inactivation.

Following RSC impairment, consistent behavioural deficits are reported in tasks involving allocentric spatial processing, particularly those requiring visual cues for orientation (Hindley 2014). These tasks include learning the location of the hidden platform in the Morris water maze (Sutherland et al., 1988; Vann & Aggleton, 2002, 2004; Whishaw et al., 2001), the radial arm maze (Keene & Bucci, 2009; Pothuizen et al., 2008; Vann & Aggleton, 2004), and object-in-place discriminations (Parron & Save, 2004). Experimental evidence suggests that the dysgranular RSC is particularly critical for allocentric navigation strategies, as rats with selective dysgranular RSC lesions struggle to use distal visual cues to guide spatial working memory, and instead rely on path integration (Vann & Aggleton, 2005). Together, these findings imply that in the spatial memory task discussed in this thesis, mice should not be able to successfully locate the hidden trigger zone using only the configuration of stable landmark cues without RSC involvement.

Allocentric spatial processing in the RSC also involves head direction and place coding. Several studies have identified RSC neurons that are tuned to the heading direction of the animal relative to the environment and the landmark layout (Angelaki et al., 2020; Chen et al., 1994b, 1994c; Cho & Sharp, 2001; Lozano et al., 2017), and head-direction cell sensitivity to landmarks is reduced following RSC lesions (Clark 2010). However, similar head-direction cells are found in various brain regions, including the pre-subiculum (Ranck, 1985; Taube et al., 1990b), thalamic nuclei (Mizumori & Williams, 1993), mammillary bodies (Stackman & Taube, 1998), and entorhinal cortex (Sargolini et al., 2006), which could substitute for the RSC in encoding allocentric directional information necessary for this spatial memory task. Additionally, place cell-like coding in the rodent RSC has been reported (Fischer et al., 2020; Mao et al., 2017; Mao et al., 2020; Mao et al., 2018; Miller et al., 2019), though it remains unclear if this place coding is inherited from hippocampal regions, such as the CA1 region of the hippocampus, where classical place cells are predominantly found. Since only RSC activity was inhibited in this study, the ongoing hippocampal place cell activity within the broader spatial coding network could explain why animals maintain task performance and process allocentric positional information without RSC contribution.

The concept of distributed processing also applies to spatial memory itself. Rather than being localized to a single brain area, memory typically involves multiple regions, creating a network that allows for redundancy and resilience. This distributed nature ensures that several areas can simultaneously store similar information and compensate for deficits that may arise in one region. Specifically, the RSC is integrated into an extensive network that includes connections to the hippocampus, medial entorhinal

cortex, and various other subcortical and cortical areas, all of which collectively contribute to guiding navigation and spatial memory processing.

The RSC plays a pivotal role in processing both egocentric and allocentric space. However, its role may be more crucial in integrating these two types of navigation rather than merely processing them individually. The BBB model provides an anatomical framework in which the RSC is integral in the transformation of egocentric and allocentric reference frames (Byrne et al., 2007). This specialised function could explain why the spatial deficits usually observed following RSC lesions, inactivation or inhibition are generally less pronounced than those resulting from damage to the hippocampus or anterior thalamic nuclei. Research by Nelson, Hindley, et al. (2015) highlights the selective nature of RSC lesion-induced spatial deficits: in a Morris water maze task where the location of a submerged platform was either determined by the geometric properties of the arena or by prominent visual landmark, RSC-lesioned rats displayed impairments in locating the platform specifically after having been placed passively on the platform several times in preceding trials, but not in conditions after which animals were trained to actively swim to the platform themselves. This suggests that RSC damage primarily disrupts the ability to switch between different spatial frames of reference and viewpoints when animals are navigating to the platform starting from unfamiliar points in the arena. In its current form, the task employed in this study does not require mice to switch from allocentric to egocentric reference frames, or vice versa, as allocentric navigation is sufficient to perform this task and to learn the location of the hidden trigger zone.

Nevertheless, the involvement of RSC in translating between allocentric and egocentric reference frames could be experimentally verified by adapting the current task version. Rotation experiments have demonstrated that mice form allocentric memories anchored to stable environmental landmarks. By modifying the task conditions — such as operating in complete darkness or removing all landmarks that facilitate allocentric navigation (e.g., using an odourless and sound-proof box with a covered ceiling) and ensuring that all trials start from the same location — it is possible to assess how quickly and accurately mice locate a hidden trigger zone under normal conditions and when RSC activity is inhibited through chemogenetic or optogenetic means. Analysing error angles could further help quantify task performance and distinguish between egocentric and allocentric navigation strategies.

Finally, despite thorough analysis, it remains possible that some behavioural differences are not detectable with the current metrics and parameters used in this study. Further research might elucidate subtle variations in behaviour and/or task performance that could provide deeper insights into RSC's role in spatial memory and navigation.

4.4.3 Remaining open questions and next steps

Several open questions about RSC's function remain, many of which could be addressed with the help of the spatial memory task described here.

For many of the experiments described below, it is crucial to monitor RSC activity while mice engage in the task to better understand RSC's role in navigation and learning. This can be achieved by employing portable recording devices, such as miniaturised one- and two-photon microscopes or high-density

electrophysiological recording probes (Aharoni & Hoogland, 2019; Steinmetz et al., 2021; Zong et al., 2021; Zong et al., 2017), which, together with the task described here, allow for the study of spatial learning in freely behaving animals, experiencing unaffected proprioception and vestibular inputs, such as angular head velocity, and natural visual flow. In this context, it would be interesting to see if the dysgranular and granular RSC differ in their neuronal activity and representation of space, as has been suggested before (Pothuizen et al., 2009; Vann & Aggleton, 2005). Additionally, examining the variations in spatial representation across the superficial and deep layers of the RSC could provide further insights into its function.

It is also worthwhile to investigate the relationship between RSC and the hippocampus. Spatial deficits resulting from RSC lesions or inactivation are generally more subtle compared to those arising from hippocampal impairments (Aggleton, 2010; Harker & Whishaw, 2004b; Hayashi & Sato, 2023). As inhibition of hippocampal activity affects spatial learning in the Morris water maze task (Bannerman et al., 1999; Morris, 1984), it can be assumed that inactivating hippocampus would also hinder spatial learning in the case of the task described here. In this context, it is interesting to note the anatomical relationship, where RSC receives more inputs from the CA1 region of the hippocampus than it sends back, whereby the nature of the information being relayed remains unknown (Mitchell et al., 2018). It will thus be important to determine the interaction between these two structures during memory formation, retrieval and reversal learning. This task offers an excellent opportunity to test for spatial representations in the RSC and hippocampus and observe how these change as mice form new memories. Further analysis could focus on the spatial memory neuronal ensembles that form and are reactivated in the RSC during spatial learning (Czajkowski et al., 2014). Investigating their stability over time could clarify the distinct role of the RSC in spatial memory formation, retrieval, and reversal learning.

The RSC has also been suggested to play a role in predictive coding (Alexander et al., 2023). This theory suggests that the brain continuously generates and updates predictions about incoming sensory information, and uses differences between these predictions and actual sensory inputs (called prediction errors) to refine its understanding of the world. This process helps in efficiently processing sensory information by focusing on unexpected or new data that deviates from predictions (Han & Helmchen, 2024). It has been hypothesised that the RSC generates predictions based on allocentric and conceptual knowledge and compares these predictions with incoming perceptual information, such as the rearrangement or manipulation of landmark cues. By employing this task, it could be tested whether this process occurs, and, if so, identify the specific predictions formed in the RSC.

4.5 Conclusion

In this thesis, I have introduced a novel behavioural paradigm designed to investigate the role of the RSC in spatial learning and memory in freely behaving mice.

This spatial memory task is particularly noteworthy for its capability to track mice and their behaviour in a three-dimensional environment without any physical constraints on head or body movements. This setup facilitates dense tracking and detailed analysis of naturalistic animal behaviour, alongside the qualitative assessment of task performance and proficiency.

This task offers several key advantages over traditional behavioural paradigms employed to study spatial learning and memory. First, the task allows for the efficient and reliable acquisition of long-term spatial memories. I show that mice are able to recall spatial memories at later training times, to adapt to new spatial information, and to update their memories of the learnt goal location within a short period of time. Second, the task requires minimal pre-training, making it highly time efficient. Mice can learn multiple spatial locations within the experimental timeline. Animals display within-day learning capabilities, which increases the number of reversal learning events that can be studied with this task. Third, the task produces stable and reproducible results, affirming its reliability for future studies. Finally, unlike many other behavioural paradigms designed to investigate spatial learning and memory, the task described here has the advantage that the learnt goal location does not spatially coincide with the reward location, allowing for a clear differentiation between neuronal coding for spatial correlates and reward in future studies employing portable recording devices.

In this work I highlight the importance of selecting appropriate metrics and task parameters for behavioural quantification in spatial memory tasks. The use of error angles and active task engagement points has enabled me to investigate the navigational strategies employed by mice trained on the spatial memory task. This nuanced analysis has shown that goal-directed navigation guided by vision and memory differs and that mice only use memory-guided navigation in the absence visual information.

Through manipulations of the external landmarks within the animal's environment during spatial memory training, I have shown that mice form a global, cognitive map of the arena and the space around them using allocentric navigation strategies to locate learnt goal locations. The results of this study further demonstrate that mice can form robust allocentric spatial memories quickly and efficiently, even under conditions where RSC activity is inhibited. The chemogenetic experiments I performed reveal that the inactivation of RSC neuronal activity does not hinder the formation or recall of spatial memories in this task.

In conclusion, this thesis introduces a robust new behavioural tool for further exploration of spatial navigation and cognition. The behavioural paradigm developed holds significant promise for future research into the neural mechanisms underlying spatial learning and memory, with many potential applications in RSC research and beyond. Through continued refinement and future application of this task, we can look forward to a richer understanding of the brain's navigation system and its role in memory and learning.

5. Bibliography

- Aggleton, J. P. (2010). Understanding retrosplenial amnesia: Insights from animal studies. *Neuropsychologia* 48(8), 2328-2338
- Aghajan, Z. M., Acharya, L., Moore, J. J., Cushman, J. D., Vuong, C., & Mehta, M. R. (2015). Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. *Nature Neuroscience* 18(1), 121-128
- Aharoni, D., & Hoogland, T. M. (2019). Circuit investigations with open-source miniaturized microscopes: past, present and future. *Frontiers in Cellular Neuroscience* 13, 141
- Alexander, A. S., Carstensen, L. C., Hinman, J. R., Raudies, F., William Chapman, G., & Hasselmo, M. E. (2020). Egocentric boundary vector tuning of the retrosplenial cortex. *Science Advances* 6(8), eaaz2322
- Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature Neuroscience* 18(8), 1143-1151
- Alexander, A. S., & Nitz, D. A. (2017). Spatially Periodic Activation Patterns of Retrosplenial Cortex Encode Route Sub-spaces and Distance Traveled. *Current Biology* 27(11), 1551-1560
- Alexander, A. S., Place, R., Starrett, M. J., Chrastil, E. R., & Nitz, D. A. (2023). Rethinking retrosplenial cortex: Perspectives and predictions. *Neuron* 111(2), 150-175
- Alexander, A. S., Tung, J. C., Chapman, G. W., Conner, A. M., Shelley, L. E., Hasselmo, M. E., & Nitz, D. A. (2022). Adaptive integration of self-motion and goals in posterior parietal cortex. *Cell Reports* 38(10), 110504
- Amin, E., Pearce, J. M., Brown, M. W., & Aggleton, J. P. (2006). Novel temporal configurations of stimuli produce discrete changes in immediate-early gene expression in the rat hippocampus. *European Journal of Neuroscience* 24(9), 2611-2621
- Angelaki, D. E., Ng, J., Abrego, A. M., Cham, H. X., Asproдини, E. K., Dickman, J. D., & Laurens, J. (2020). A gravity-based three-dimensional compass in the mouse brain. *Nature Communications* 11(1), 1855
- Aronov, D., & Tank, D. W. (2014). Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* 84(2), 442-456
- Bannerman, D., Yee, B., Good, M. A., Heupel, M., Iversen, S., & Rawlins, J. (1999). Double dissociation of function within the hippocampus: a comparison of dorsal, ventral, and complete hippocampal cytotoxic lesions. *Behavioral Neuroscience* 113(6), 1170
- Barnes, C. A. (1979). Memory deficits associated with senescence: a neurophysiological and behavioral study in the rat. *Journal of Comparative and Physiological Psychology* 93(1), 74
- Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience* 10(6), 682-684

- Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K., & Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences* 17(1-2), 71-98
- Benhamou, S., & Poucet, B. (1995). A comparative analysis of spatial memory processes. *Behavioural Processes* 35(1-3), 113-126
- Berlyne, D. E. (1950). Novelty and curiosity as determinants of exploratory behaviour. *British Journal of Psychology* 41(1), 68
- Bicanski, A., & Burgess, N. (2016). Environmental anchoring of head direction in a computational model of retrosplenial cortex. *Journal of Neuroscience* 36(46), 11601-11618
- Bicanski, A., & Burgess, N. (2018). A neural-level model of spatial memory and imagery. *eLife* 7, e33752
- Bicanski, A., & Burgess, N. (2020). Neuronal vector coding in spatial cognition. *Nature Reviews Neuroscience* 21(9), 453-470
- Boccaro, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., & Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience* 13(8), 987-994
- Burgess, N. (2006). Spatial memory: how egocentric and allocentric combine. *Trends in Cognitive Sciences* 10(12), 551-557
- Burgess, N. (2008). Spatial Cognition and the Brain. *Annals of the New York Academy of Sciences* 1124(1), 77-97
- Burgess, N. (2014). The 2014 Nobel Prize in Physiology or Medicine: a spatial model for cognitive neuroscience. *Neuron* 84(6), 1120-1125
- Burgess, N., Barry, C., & O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus* 17(9), 801-812
- Bush, D., Barry, C., Manson, D., & Burgess, N. (2015). Using Grid Cells for Navigation. *Neuron* 87(3), 507-520
- Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychological Review* 114(2), 340-375
- Cain, D. P., Humpartzoomian, R., & Boon, F. (2006). Retrosplenial cortex lesions impair water maze strategies learning or spatial place learning depending on prior experience of the rat. *Behavioural Brain Research* 170(2), 316-325
- Campbell, M. G., Attinger, A., Ocko, S. A., Ganguli, S., & Giocomo, L. M. (2021). Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex. *Cell Reports* 36(10), 109669

- Chang, H., Esteves, I. M., Neumann, A. R., Sun, J., Mohajerani, M. H., & McNaughton, B. L. (2020). Coordinated activities of retrosplenial ensembles during resting-state encode spatial landmarks. *Philosophical Transactions of the Royal Society B* 375(1799), 20190228
- Chen, G., King, J. A., Lu, Y., Cacucci, F., & Burgess, N. (2018). Spatial cell firing during virtual navigation of open arenas by head-restrained mice. *eLife* 7, e34789
- Chen, L. L., Lin, L.-H., Barnes, C. A., & McNaughton, B. L. (1994a). Head-direction cells in the rat posterior cortex. *Experimental Brain Research* 101(1), 24-34
- Chen, L. L., Lin, L.-H., Barnes, C. A., & McNaughton, B. L. (1994b). Head-direction cells in the rat posterior cortex - I. anatomical distribution and behavioral modulation. *Experimental Brain Research* 101(1), 8-23
- Chen, L. L., Lin, L.-H., Barnes, C. A., & McNaughton, B. L. (1994c). Head-direction cells in the rat posterior cortex - II. Contributions of visual and ideothetic information to the directional firing. *Experimental Brain Research* 101(1), 24-34
- Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition* 23(2), 149-178
- Chersi, F., & Burgess, N. (2015). The Cognitive Architecture of Spatial Navigation: Hippocampal and Striatal Contributions. *Neuron* 88(1), 64-77
- Chinzorig, C., Nishimaru, H., Matsumoto, J., Takamura, Y., Berthoz, A., Ono, T., & Nishijo, H. (2019). Rat Retrosplenial Cortical Involvement in Wayfinding Using Visual and Locomotor Cues. *Cerebral Cortex* 30(4), 1985-2004
- Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience* 115(1), 3-25
- Chrastil, E. R. (2018). Heterogeneity in human retrosplenial cortex: A review of function and connectivity. *Behavioral Neuroscience* 132(5), 317
- Clancy, K. B., Orsolich, I., & Mrsic-Flogel, T. D. (2019). Locomotion-dependent remapping of distributed cortical networks. *Nature Neuroscience* 22(5), 778-786
- Clark, B. J., Bassett, J. P., Wang, S. S., & Taube, J. S. (2010). Impaired head direction cell representation in the anterodorsal thalamus after lesions of the retrosplenial cortex. *Journal of Neuroscience* 30(15), 5289-5302
- Clark, B. J., Simmons, C. M., Berkowitz, L. E., & Wilber, A. A. (2018). The retrosplenial-parietal network and reference frame coordination for spatial navigation. *Behavioral Neuroscience* 132(5), 416-429
- Claudi, F., Tyson, A. L., Petrucco, L., Margrie, T. W., Portugues, R., & Branco, T. (2021). Visualizing anatomically registered data with brainrender. *eLife* 10, e65751
- Clayton, N. S., Dally, J. M., & Emery, N. J. (2007). Social cognition by food-caching corvids. The western scrub-jay as a natural psychologist. *Philosophical Transactions of the Royal Society B* 362(1480), 507-522

- Cole, M. R., & Chappell-Stephenson, R. (2003). Exploring the limits of spatial memory in rats, using very large mazes. *Animal Learning & Behavior* 31(4), 349-368
- Cooper, B. G., Manka, T. F., & Mizumori, S. J. Y. (2001). Finding your way in the dark: The retrosplenial cortex contributes to spatial memory and navigation without visual cues. *Behavioral Neuroscience* 115(5), 1012-1028
- Cooper, B. G., & Mizumori, S. J. Y. (1999). Retrosplenial cortex inactivation selectivity impairs navigation in darkness. *NeuroReport* 10(3), 625-630
- Cooper, B. G., & Mizumori, S. J. Y. (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. *Journal of Neuroscience* 21(11), 3986-4001
- Coughlan, G., Laczó, J., Hort, J., Minihane, A.-M., & Hornberger, M. (2018). Spatial navigation deficits — overlooked cognitive marker for preclinical Alzheimer disease? *Nature Reviews Neurology* 14(8), 496-506
- Cowansage, Kiriana K., Shuman, T., Dillingham, Blythe C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct Reactivation of a Coherent Neocortical Memory of Context. *Neuron* 84(2), 432-441
- Czajkowski, R., Jayaprakash, B., Wiltgen, B., Rogerson, T., Guzman-Karlsson, M. C., Barth, A. L., Trachtenberg, J. T., & Silva, A. J. (2014). Encoding and storage of spatial information in the retrosplenial cortex. *Proceedings of the National Academy of Sciences* 111(23), 8661-8666
- Czajkowski, R., Sugar, J., Zhang, S.-J., Couey, J. J., Ye, J., & Witter, M. P. (2013). Superficially Projecting Principal Neurons in Layer V of Medial Entorhinal Cortex in the Rat Receive Excitatory Retrosplenial Input. *Journal of Neuroscience* 33(40), 15779-15792
- Czajkowski, R., Zglinicki, B., Rejmak, E., & Konopka, W. (2020). Strategy-Specific Patterns of Arc Expression in the Retrosplenial Cortex and Hippocampus during T-Maze Learning in Rats. *Brain Sciences* 10(11), 854
- D'Hooge, R., & De Deyn, P. P. (2001). Applications of the Morris water maze in the study of learning and memory. *Brain Research Reviews* 36(1), 60-90
- de Sousa, A. F., Cowansage, K. K., Zutshi, I., Cardozo, L. M., Yoo, E. J., Leutgeb, S., & Mayford, M. (2019). Optogenetic reactivation of memory ensembles in the retrosplenial cortex induces systems consolidation. *Proceedings of the National Academy of Sciences* 116(17), 8576-8581
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature* 463(7281), 657-661
- Dubreuil, D., Tixier, C., Dutrieux, G., & Edeline, J.-M. (2003). Does the radial arm maze necessarily test spatial memory? *Neurobiology of Learning and Memory* 79(1), 109-117
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience* 13(8), 995-1002

- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The Hippocampus, Memory, and Place Cells. *Neuron* 23(2), 209-226
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature* 425(6954), 184-188
- Elduayen, C., & Save, E. (2014). The retrosplenial cortex is necessary for path integration in the dark. *Behavioural Brain Research* 272, 303-307
- Fellini, L., & Morellini, F. (2011). Geometric information is required for allothetic navigation in mice. *Behavioural Brain Research* 222(2), 380-384
- Fellini, L., Schachner, M., & Morellini, F. (2006). Adult but not aged C57BL/6 male mice are capable of using geometry for orientation. *Learning & Memory* 13(4), 473-481
- Finkelstein, A., Derdikman, D., Rubin, A., Foerster, J. N., Las, L., & Ulanovsky, N. (2015). Three-dimensional head-direction coding in the bat brain. *Nature* 517(7533), 159-164
- Fischer, L. F., Soto-Albors, R. M., Buck, F., & Harnett, M. T. (2020). Representation of visual landmarks in retrosplenial cortex. *eLife* 9, e51458
- Forsythe, J. W., & Hanlon, R. T. (1997). Foraging and associated behavior by *Octopus cyanea* Gray, 1849 on a coral atoll, French Polynesia. *Journal of Experimental Marine Biology and Ecology* 209(1), 15-31
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446(7132), 190-194
- Fyhn, M., Hafting, T., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Grid cells in mice. *Hippocampus* 18(12), 1230-1238
- 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
- Gawel, K., Gibula, E., Marszalek-Grabska, M., Filarowska, J., & Kotlinska, J. H. (2019). Assessment of spatial learning and memory in the Barnes maze task in rodents—methodological consideration. *Naunyn-Schmiedeberg's Archives of Pharmacology* 392, 1-18
- Ghosh, V. E., & Gilboa, A. (2014). What is a memory schema? A historical perspective on current neuroscience literature. *Neuropsychologia* 53, 104-114
- Gofman, X., Tocker, G., Weiss, S., Boccara, C. N., Lu, L., Moser, M.-B., Moser, E. I., Morris, G., & Derdikman, D. (2019). Dissociation between Postrhinal Cortex and Downstream Parahippocampal Regions in the Representation of Egocentric Boundaries. *Current Biology* 29(16), 2751-2757
- Goltstein, P. M., Reinert, S., Glas, A., Bonhoeffer, T., & Hübener, M. (2018). Food and water restriction lead to differential learning behaviors in a head-fixed two-choice visual discrimination task for mice. *PloS one* 13(9), e0204066

- Graving, J. M., Chae, D., Naik, H., Li, L., Koger, B., Costelloe, B. R., & Couzin, I. D. (2019). DeepPoseKit, a software toolkit for fast and robust animal pose estimation using deep learning. *eLife* 8, e47994
- Grieves, R. M., Jedidi-Ayoub, S., Mishchanchuk, K., Liu, A., Renaudineau, S., & Jeffery, K. J. (2020). The place-cell representation of volumetric space in rats. *Nature Communications* 11(1), 789
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* 436(7052), 801-806
- Han, S., & Helmchen, F. (2024). Behavior-relevant top-down cross-modal predictions in mouse neocortex. *Nature Neuroscience*, 1-11
- Harker, K. T., & Ian, Q. W. (2002). Impaired Spatial Performance in Rats with Retrosplenial Lesions: Importance of the Spatial Problem and the Rat Strain in Identifying Lesion Effects in a Swimming Pool. *Journal of Neuroscience* 22(3), 1155
- Harker, K. T., & Whishaw, I. Q. (2004a). Impaired place navigation in place and matching-to-place swimming pool tasks follows both retrosplenial cortex lesions and cingulum bundle lesions in rats. *Hippocampus* 14(2), 224-231
- Harker, K. T., & Whishaw, I. Q. (2004b). A reaffirmation of the retrosplenial contribution to rodent navigation: reviewing the influences of lesion, strain, and task. *Neuroscience & Biobehavioral Reviews* 28(5), 485-496
- Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2014). Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B* 369(1635), 20120510
- Hayashi, T., & Sato, N. (2023). Contribution of the retrosplenial cortex to route selection in a complex maze. *Neuroscience Research*
- He, J., Yamada, K., Nakajima, A., Kamei, H., & Nabeshima, T. (2002). Learning and memory in two different reward tasks in a radial arm maze in rats. *Behavioural Brain Research* 134(1), 139-148
- Hennestad, E., Witoelar, A., Chambers, A. R., & Vervaeke, K. (2021). Mapping vestibular and visual contributions to angular head velocity tuning in the cortex. *Cell Reports* 37(12), 110134
- Hindley, E. L., Nelson, A. J. D., Aggleton, J. P., & Vann, S. D. (2014a). Dysgranular retrosplenial cortex lesions in rats disrupt cross-modal object recognition. *Learning & Memory* 21(3), 171-179
- Hindley, E. L., Nelson, A. J. D., Aggleton, J. P., & Vann, S. D. (2014b). The rat retrosplenial cortex is required when visual cues are used flexibly to determine location. *Behavioural Brain Research* 263, 98-107
- Honda, Y., & Ishizuka, N. (2015). Topographic distribution of cortical projection cells in the rat subiculum. *Neuroscience Research* 92, 1-20
- Jacob, P.-Y., Casali, G., Spieser, L., Page, H., Overington, D., & Jeffery, K. (2017). An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nature Neuroscience* 20(2), 173-175

- Jacobs, L. F., & Liman, E. R. (1991). Grey squirrels remember the locations of buried nuts. *Animal Behaviour* 41(1), 103-110
- Jercog, P. E., Ahmadian, Y., Woodruff, C., Deb-Sen, R., Abbott, L. F., & Kandel, E. R. (2019). Heading direction with respect to a reference point modulates place-cell activity. *Nature Communications* 10(1), 2333
- Jones, B. F., Groenewegen, H. J., & Witter, M. P. (2005). Intrinsic connections of the cingulate cortex in the rat suggest the existence of multiple functionally segregated networks. *Neuroscience* 133(1), 193-207
- Jones, B. F., & Witter, M. P. (2007). Cingulate cortex projections to the parahippocampal region and hippocampal formation in the rat. *Hippocampus* 17(10), 957-976
- Keene, C. S., & Bucci, D. J. (2009). Damage to the retrosplenial cortex produces specific impairments in spatial working memory. *Neurobiology of Learning and Memory* 91(4), 408-414
- Keshavarzi, S., Bracey, E. F., Faville, R. A., Campagner, D., Tyson, A. L., Lenzi, S. C., Branco, T., & Margrie, T. W. (2022). Multisensory coding of angular head velocity in the retrosplenial cortex. *Neuron* 110(3), 532-543. e539
- Kislin, M., Mugantseva, E., Molotkov, D., Kuleskaya, N., Khirug, S., Kirilkin, I., Pryazhnikov, E., Kolikova, J., Toptunov, D., & Yuryev, M. (2014). Flat-floored air-lifted platform: a new method for combining behavior with microscopy or electrophysiology on awake freely moving rodents. *Journal of Visualized Experiments*(88), e51869
- Klatzky, R. L. (1998). Allocentric and egocentric spatial representations: Definitions, distinctions, and interconnections. In *Spatial cognition: An interdisciplinary approach to representing and processing spatial knowledge* (pp. 1-17). Springer Berlin Heidelberg.
- Klioutchnikov, A., Wallace, D. J., Sawinski, J., Voit, K.-M., Groemping, Y., & Kerr, J. N. (2023). A three-photon head-mounted microscope for imaging all layers of visual cortex in freely moving mice. *Nature Methods* 20(4), 610-616
- Knierim, J. J., & Hamilton, D. A. (2011). Framing spatial cognition: neural representations of proximal and distal frames of reference and their roles in navigation. *Physiological Reviews* 91(4), 1245-1279
- Kobayashi, Y., & Amaral, D. G. (2000). Macaque monkey retrosplenial cortex: I. Three-dimensional and cytoarchitectonic organization. *Journal of Comparative Neurology* 426(3), 339-365
- Kobayashi, Y., & Amaral, D. G. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *Journal of Comparative Neurology* 466(1), 48-79
- Kolb, B., Sutherland, R. J., & Whishaw, I. Q. (1983). A comparison of the contributions of the frontal and parietal association cortex to spatial localization in rats. *Behavioral Neuroscience* 97(1), 13
- LaChance, P. A., Todd, T. P., & Taube, J. S. (2019). A sense of space in postrhinal cortex. *Science* 365(6449), eaax4192

- Law, J. W., Lee, A. Y., Sun, M., Nikonenko, A. G., Chung, S. K., Dityatev, A., Schachner, M., & Morellini, F. (2003). Decreased anxiety, altered place learning, and increased CA1 basal excitatory synaptic transmission in mice with conditional ablation of the neural cell adhesion molecule L1. *Journal of Neuroscience* 23(32), 10419-10432
- Lenck-Santini, P. P., Save, E., & Poucet, B. (2001). Place-cell firing does not depend on the direction of turn in a Y-maze alternation task. *European Journal of Neuroscience* 13(5), 1055-1058
- Lever, C., Burgess, N., Cacucci, F., Hartley, T., & O'Keefe, J. (2002). What can the hippocampal representation of environmental geometry tell us about Hebbian learning? *Biological Cybernetics* 87, 356-372
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience* 29(31), 9771-9777
- Lozano, Y. R., Page, H., Jacob, P.-Y., Lomi, E., Street, J., & Jeffery, K. (2017). Retrosplenial and postsubicular head direction cells compared during visual landmark discrimination. *Brain and Neuroscience Advances* 1, 2398212817721859
- Lukoyanov, N. V., & Lukoyanova, E. A. (2006). Retrosplenial cortex lesions impair acquisition of active avoidance while sparing fear-based emotional memory. *Behavioural Brain Research* 173(2), 229-236
- Lukoyanov, N. V., Lukoyanova, E. A., Andrade, J. P., & Paula-Barbosa, M. M. (2005). Impaired water maze navigation of Wistar rats with retrosplenial cortex lesions: effect of nonspatial pretraining. *Behavioural Brain Research* 158(1), 175-182
- Maei, H. R., Zaslavsky, K., Teixeira, C. M., & Frankland, P. W. (2009). What is the most sensitive measure of water maze probe test performance? *Frontiers in Integrative Neuroscience* 3, 493
- Makino, H., & Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nature Neuroscience* 18(8), 1116-1122
- Mallory, C. S., Hardcastle, K., Campbell, M. G., Attinger, A., Low, I. I. C., Raymond, J. L., & Giocomo, L. M. (2021). Mouse entorhinal cortex encodes a diverse repertoire of self-motion signals. *Nature Communications* 12(1), 671
- Mao, D., Kandler, S., McNaughton, B. L., & Bonin, V. (2017). Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nature Communications* 8(1), 243
- Mao, D., Molina, L. A., Bonin, V., & McNaughton, B. L. (2020). Vision and Locomotion Combine to Drive Path Integration Sequences in Mouse Retrosplenial Cortex. *Current Biology* 30(9), 1680-1688
- Mao, D., Neumann, A. R., Sun, J., Bonin, V., Mohajerani, M. H., & McNaughton, B. L. (2018). Hippocampus-dependent emergence of spatial sequence coding in retrosplenial cortex. *Proceedings of the National Academy of Sciences* 115(31), 8015-8018

- 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
- Marr, D., Willshaw, D., & McNaughton, B. L. (1991). *Simple memory: a theory for archicortex*. Springer, Birkhäuser Boston.
- Mathis, A., Mamidanna, P., Cury, K. M., Abe, T., Murthy, V. N., Mathis, M. W., & Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nature Neuroscience* 21(9), 1281-1289
- Maurer, R., & Derivaz, V. (2000). Rats in a transparent morris water maze use elemental and configural geometry of landmarks as well as distance to the pool wall. *Spatial Cognition and Computation* 2, 135-156
- Mazmanian, D. S., & Roberts, W. A. (1983). Spatial memory in rats under restricted viewing conditions. *Learning and Motivation* 14(2), 123-139
- McQuade, D. B., Williams, E. H., & Eichenbaum, H. B. (1986). Cues Used for Localizing Food by the Gray Squirrel (*Sciurus carolinensis*). *Ethology* 72(1), 22-30
- Milczarek, M. M., Vann, S. D., & Sengpiel, F. (2018). Spatial memory engram in the mouse retrosplenial cortex. *Current Biology* 28(12), 1975-1980
- Miller, A. M. P., Mau, W., & Smith, D. M. (2019). Retrosplenial Cortical Representations of Space and Future Goal Locations Develop with Learning. *Current Biology* 29(12), 2083-2090
- Miller, A. M. P., Serrichio, A. C., & Smith, D. M. (2020). Dual-Factor Representation of the Environmental Context in the Retrosplenial Cortex. *Cerebral Cortex* 31(5), 2720-2728
- Minderer, M., Brown, K. D., & Harvey, C. D. (2019). The Spatial Structure of Neural Encoding in Mouse Posterior Cortex during Navigation. *Neuron* 102(1), 232-248
- Mitchell, A. S., Czajkowski, R., Zhang, N., Jeffery, K., & Nelson, A. J. D. (2018). Retrosplenial cortex and its role in spatial cognition. *Brain and Neuroscience Advances* 2, 239821281875709
- Miyashita, T., & Rockland, K. S. (2007). GABAergic projections from the hippocampus to the retrosplenial cortex in the rat. *European Journal of Neuroscience* 26(5), 1193-1204
- Mizumori, S. J. Y., & Williams, J. D. (1993). Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience* 13(9), 4015-4028
- Moghaddam, M., & Bures, J. (1996). Contribution of egocentric spatial memory to place navigation of rats in the Morris water maze. *Behavioural Brain Research* 78(2), 121-129
- Morales, L., Tomàs, D. P., Dalmau, J., De la Rocha, J., & Jercog, P. E. (2020). High-throughput task to study memory recall during spatial navigation in rodents. *Frontiers in Behavioral Neuroscience* 14, 64

- Morellini, F. (2013). Spatial memory tasks in rodents: what do they model? *Cell and Tissue Research* 354(1), 273-286
- Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation* 12(2), 239-260
- Morris, R. G. M. (1984). Developments of a water-maze procedure for studying spatial learning in the rat. *Journal of Neuroscience Methods* 11(1), 47-60
- Morris, R. G. M. (2006). Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. *European Journal of Neuroscience* 23(11), 2829-2846
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature* 297(5868), 681-683
- 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
- Muysers, H., Chen, H.-L., Hahn, J., Folschweiller, S., Sigurdsson, T., Sauer, J.-F., & Bartos, M. (2024). A persistent prefrontal reference frame across time and task rules. *Nature Communications* 15(1), 2115
- Nashaat, M. A., Oraby, H., Sachdev, R. N., Winter, Y., & Larkum, M. E. (2016). Air-Track: a real-world floating environment for active sensing in head-fixed mice. *Journal of Neurophysiology* 116(4), 1542-1553
- Nelson, A. J. D., Hindley, E. L., Pearce, J. M., Vann, S. D., & Aggleton, J. P. (2015). The effect of retrosplenial cortex lesions in rats on incidental and active spatial learning. *Frontiers in Behavioral Neuroscience* 9, 11
- Nelson, A. J. D., Powell, A. L., Holmes, J. D., Vann, S. D., & Aggleton, J. P. (2015). What does spatial alternation tell us about retrosplenial cortex function? *Frontiers in Behavioral Neuroscience* 9, 126
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature* 381(6581), 425-428
- O'Keefe, J., & Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15(7), 853-866
- O'Keefe, J., & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research* 31, 573-590
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Research*
- Olton, D. S., Collison, C., & Werz, M. A. (1977). Spatial memory and radial arm maze performance of rats. *Learning and Motivation* 8(3), 289-314

- Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes* 2(2), 97
- Ormond, J., & O'Keefe, J. (2022). Hippocampal place cells have goal-oriented vector fields during navigation. *Nature* 607(7920), 741-746
- Parron, C., & Save, E. (2004). Comparison of the effects of entorhinal and retrosplenial cortical lesions on habituation, reaction to spatial and non-spatial changes during object exploration in the rat. *Neurobiology of Learning and Memory* 82(1), 1-11
- Paul, C.-M., Magda, G., & Abel, S. (2009). Spatial memory: Theoretical basis and comparative review on experimental methods in rodents. *Behavioural Brain Research* 203(2), 151-164
- Paxinos, G., & Franklin, K. B. (2001). *Paxinos and Franklin's the mouse brain in stereotaxic coordinates* (second ed.). Academic Press.
- Pereira, T. D., Shaevitz, J. W., & Murthy, M. (2020). Quantifying behavior to understand the brain. *Nature Neuroscience* 23(12), 1537-1549
- Pereira, T. D., Tabris, N., Li, J., Ravindranath, S., Papadoyannis, E. S., Wang, Z. Y., Turner, D. M., McKenzie-Smith, G., Kocher, S. D., Falkner, A. L., Shaevitz, J. W., & Murthy, M. (2020). SLEAP: Multi-animal pose tracking. *bioRxiv*, 2020.2008.2031.276246
- Peyrache, A., Schieferstein, N., & Buzsáki, G. (2017). Transformation of the head-direction signal into a spatial code. *Nature Communications* 8(1), 1752
- Pologruto, T. A., Sabatini, B. L., & Svoboda, K. (2003). ScanImage: Flexible software for operating laser scanning microscopes. *Biomedical Engineering Online* 2, 1-9
- Pothuizen, H. H. J., Aggleton, J. P., & Vann, S. D. (2008). Do rats with retrosplenial cortex lesions lack direction? *European Journal of Neuroscience* 28(12), 2486-2498
- Pothuizen, H. H. J., Davies, M., Albasser, M. M., Aggleton, J. P., & Vann, S. D. (2009). Granular and dysgranular retrosplenial cortices provide qualitatively different contributions to spatial working memory: evidence from immediate-early gene imaging in rats. *European Journal of Neuroscience* 30(5), 877-888
- Powell, A., Connelly, W. M., Vasalaukaite, A., Nelson, A. J. D., Vann, S. D., Aggleton, J. P., Sengpiel, F., & Ranson, A. (2020). Stable Encoding of Visual Cues in the Mouse Retrosplenial Cortex. *Cerebral Cortex* 30(8), 4424-4437
- Qiu, Y., Zhao, Z., Klindt, D., Kautzky, M., Szatko, K. P., Schaeffel, F., Rifai, K., Franke, K., Busse, L., & Euler, T. (2021). Natural environment statistics in the upper and lower visual field are reflected in mouse retinal specializations. *Current Biology* 31(15), 3233-3247. e3236
- Quirk, G., Muller, R., & Kubie, J. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *Journal of Neuroscience* 10(6), 2008-2017

- Ramos, J. M. (2000). Influence of the shape of the experimental room on spatial learning in rats. *Physiology & Behavior* 70(3-4), 351-357
- Ranck, J. B. (1985). Head direction cells in the deep cell layer of dorsolateral pre-subiculum in freely moving rats. *Electrical activity of the archicortex*
- Robertson, R. G., Rolls, E. T., Georges-François, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. *Hippocampus* 9(3), 206-219
- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., & Moser, E. I. (2006). Conjunctive Representation of Position, Direction, and Velocity in Entorhinal Cortex. *Science* 312(5774), 758-762
- Save, E. (1997). The contribution of visual and inertial mechanisms to navigation in total darkness. *Animal Learning & Behavior* 25(3), 324-334
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science* 322(5909), 1865-1868
- St-Laurent, M., Petrides, M., & Sziklas, V. (2009). Does the cingulate cortex contribute to spatial conditional associative learning in the rat? *Hippocampus* 19(7), 612-622
- Stackman, R. W., & Taube, J. S. (1998). Firing Properties of Rat Lateral Mammillary Single Units: Head Direction, Head Pitch, and Angular Head Velocity. *Journal of Neuroscience* 18(21), 9020-9037
- Steinmetz, N. A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., & Broux, M. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* 372(6539), eabf4588
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature* 492(7427), 72-78
- Sugar, J., Witter, M. P., van Strien, N. M., & Cappaert, N. L. (2011). The retrosplenial cortex: intrinsic connectivity and connections with the (para) hippocampal region in the rat. An interactive connectome. *Frontiers in Neuroinformatics* 5, 7
- Sun, Y., Nitz, D. A., Xu, X., & Giocomo, L. M. (2024). Subicular neurons encode concave and convex geometries. *Nature*, 1-9
- Sunyer, B., Patil, S., Höger, H., & Lubec, G. (2007). Barnes maze, a useful task to assess spatial reference memory in the mice. *Nature Protocol Exchange* 10, 1038/nprot.2007.1390
- Sutherland, R., Whishaw, I., & Kolb, B. (1988). Contributions of cingulate cortex to two forms of spatial learning and memory. *Journal of Neuroscience* 8(6), 1863-1872
- Suzuki, S., Augerinos, G., & Black, A. H. (1980). Stimulus control of spatial behavior on the eight-arm maze in rats. *Learning and Motivation* 11(1), 1-18

- Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience* 15(1), 70-86
- Taube, J. S. (2007). The head direction signal: origins and sensory-motor integration. *Annual Review of Neuroscience* 30, 181-207
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience* 10(2), 420-435
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience* 10(2), 436-447
- Thinus-Blanc, C. (1996). *Animal spatial cognition: Behavioural and brain approach*. World Scientific Publishing Company.
- Thompson, L. T., & Best, P. J. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research* 509(2), 299-308
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review* 55(4), 189
- Tsoar, A., Nathan, R., Bartan, Y., Vyssotski, A., Dell'Omo, G., & Ulanovsky, N. (2011). Large-scale navigational map in a mammal. *Proceedings of the National Academy of Sciences* 108(37), E718-E724
- Ulanovsky, N., & Moss, C. F. (2007). Hippocampal cellular and network activity in freely moving echolocating bats. *Nature Neuroscience* 10(2), 224-233
- Valerio, S., & Taube, J. S. (2012). Path integration: how the head direction signal maintains and corrects spatial orientation. *Nature Neuroscience* 15(10), 1445-1453
- Van Dam, D., Lenders, G., & De Deyn, P. P. (2006). Effect of Morris water maze diameter on visual-spatial learning in different mouse strains. *Neurobiology of Learning and Memory* 85(2), 164-172
- Van Groen, T., Kadish, I., & Wyss, J. M. (2004). Retrosplenial cortex lesions of area Rgb (but not of area Rga) impair spatial learning and memory in the rat. *Behavioural Brain Research* 154(2), 483-491
- Van Groen, T., & Wyss, J. M. (1990). Connections of the retrosplenial granular a cortex in the rat. *Journal of Comparative Neurology* 300(4), 593-606
- Van Groen, T., & Wyss, J. M. (1992a). Connections of the retrosplenial dysgranular cortex in the rat. *Journal of Comparative Neurology* 315(2), 200-216
- Van Groen, T., & Wyss, J. M. (1992b). Projections from the laterodorsal nucleus of the thalamus to the limbic and visual cortices in the rat. *Journal of Comparative Neurology* 324(3), 427-448
- Van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *Journal of Comparative Neurology* 463(3), 249-263

- Van Wijngaarden, J. B., Babl, S. S., & Ito, H. T. (2020). Entorhinal-retrosplenial circuits for allocentric-egocentric transformation of boundary coding. *eLife* 9, e59816
- Vann, S. D., & Aggleton, J. P. (2002). Extensive cytotoxic lesions of the rat retrosplenial cortex reveal consistent deficits on tasks that tax allocentric spatial memory. *Behavioral Neuroscience* 116(1), 85-94
- Vann, S. D., & Aggleton, J. P. (2004). Testing the importance of the retrosplenial guidance system: effects of different sized retrosplenial cortex lesions on heading direction and spatial working memory. *Behavioural Brain Research* 155(1), 97-108
- Vann, S. D., & Aggleton, J. P. (2005). Selective dysgranular retrosplenial cortex lesions in rats disrupt allocentric performance of the radial-arm maze task. *Behavioral Neuroscience* 119(6), 1682-1686
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience* 10(11), 792-802
- Vann, S. D., Brown, M., & Aggleton, J. P. (2000). Fos expression in the rostral thalamic nuclei and associated cortical regions in response to different spatial memory tests. *Neuroscience* 101(4), 983-991
- Vann, S. D., Kristina Wilton, L. A., Muir, J. L., & Aggleton, J. P. (2003). Testing the importance of the caudal retrosplenial cortex for spatial memory in rats. *Behavioural Brain Research* 140(1), 107-118
- Vedder, L. C., Miller, A. M. P., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cerebral Cortex* 27(7), 3713-3723
- Vogt, B. A., & Peters, A. (1981). Form and distribution of neurons in rat cingulate cortex: areas 32, 24, and 29. *Journal of Comparative Neurology* 195(4), 603-625
- Voigts, J., & Harnett, M. T. (2020). Somatic and Dendritic Encoding of Spatial Variables in Retrosplenial Cortex Differs during 2D Navigation. *Neuron* 105(2), 237-245
- Wallace, D. G., Hines, D. J., Pellis, S. M., & Whishaw, I. Q. (2002). Vestibular Information Is Required for Dead Reckoning in the Rat. *Journal of Neuroscience* 22(22), 10009-10017
- Watson, C., Paxinos, G., & Puelles, L. (2011). *The mouse nervous system* (Illustrated ed.). Academic Press.
- Whishaw, I. Q., Maaswinkel, H., Gonzalez, C. L. R., & Kolb, B. (2001). Deficits in allothetic and idiothetic spatial behavior in rats with posterior cingulate cortex lesions. *Behavioural Brain Research* 118(1), 67-76
- Wilber, A. A., Clark, B. J., Forster, T. C., Tatsuno, M., & McNaughton, B. L. (2014). Interaction of egocentric and world-centered reference frames in the rat posterior parietal cortex. *Journal of Neuroscience* 34(16), 5431-5446
- Wyass, J. M., & Van Groen, T. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: A review. *Hippocampus* 2(1), 1-11

- Yamawaki, N., Radulovic, J., & Shepherd, G. M. G. (2016). A corticocortical circuit directly links retrosplenial cortex to M2 in the mouse. *Journal of Neuroscience* 36(36), 9365-9374
- Yartsev, M. M., & Ulanovsky, N. (2013). Representation of Three-Dimensional Space in the Hippocampus of Flying Bats. *Science* 340(6130), 367-372
- Yartsev, M. M., Witter, M. P., & Ulanovsky, N. (2011). Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature* 479(7371), 103-107
- Zhuang, J., Ng, L., Williams, D., Valley, M., Li, Y., Garrett, M., & Waters, J. (2017). An extended retinotopic map of mouse cortex. *eLife* 6, e18372
- Zong, W., Wu, R., Chen, S., Wu, J., Wang, H., Zhao, Z., Chen, G., Tu, R., Wu, D., & Hu, Y. (2021). Miniature two-photon microscopy for enlarged field-of-view, multi-plane and long-term brain imaging. *Nature Methods* 18(1), 46-49
- Zong, W., Wu, R., Li, M., Hu, Y., Li, Y., Li, J., Rong, H., Wu, H., Xu, Y., & Lu, Y. (2017). Fast high-resolution miniature two-photon microscopy for brain imaging in freely behaving mice. *Nature Methods* 14(7), 713-719

6. Supplementary Figures

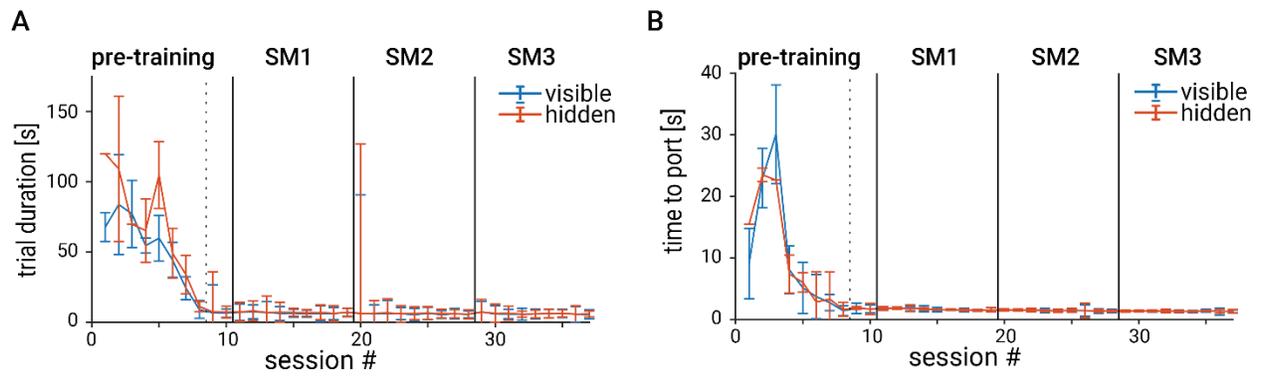


Figure S1 Task performance: trial duration and time-to-port

A. Median (\pm std) trial duration for visible-random (*blue*) and hidden-fixed (*red*) trigger trials. $n=13$ mice. B. Median (\pm std) time to port following the visit of a visible-random (*blue*) or hidden-fixed (*red*) trigger zone. $n=13$ mice.

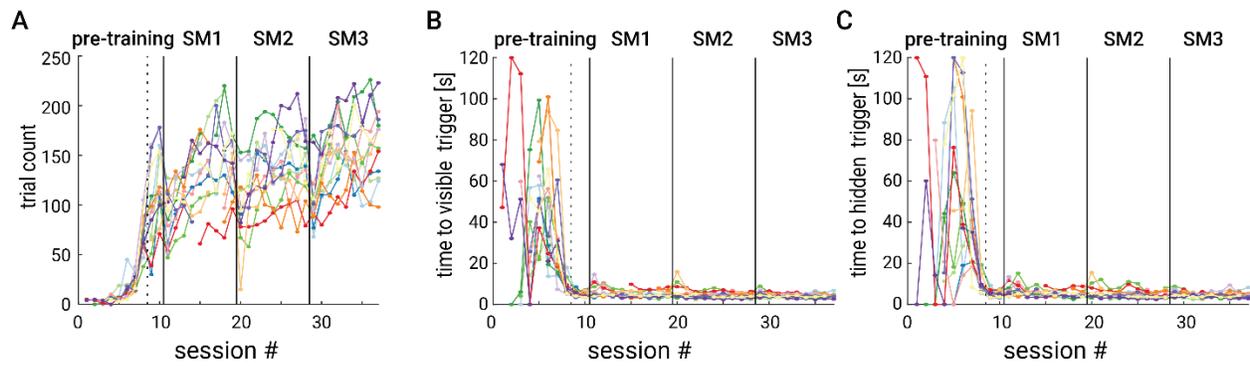


Figure S2 Individual mouse task performance

A. Median trial count throughout pre-training and the spatial memory training (three hidden-fixed trigger zones: SM1, SM2, and SM3) (*black vertical lines*) for each individual mouse. $n=13$ mice. B. Median (\pm std) time to visible-random, and C. Median time to hidden-fixed trigger zones for each individual mouse. $n=13$ mice.

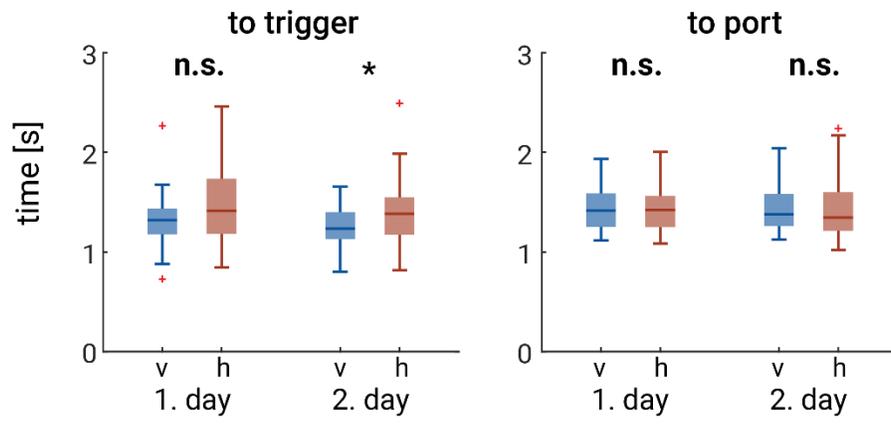


Figure S3 Time from start search point to trigger entry

Time to trigger (*left*) and to port (*right*) measured from the start search point for visible-random (*blue*) and hidden-fixed (*red*) trigger trials. *Wilcoxon rank sum test*, *Bonferroni correction*, * $p < 0.025$. $n=13$ mice.

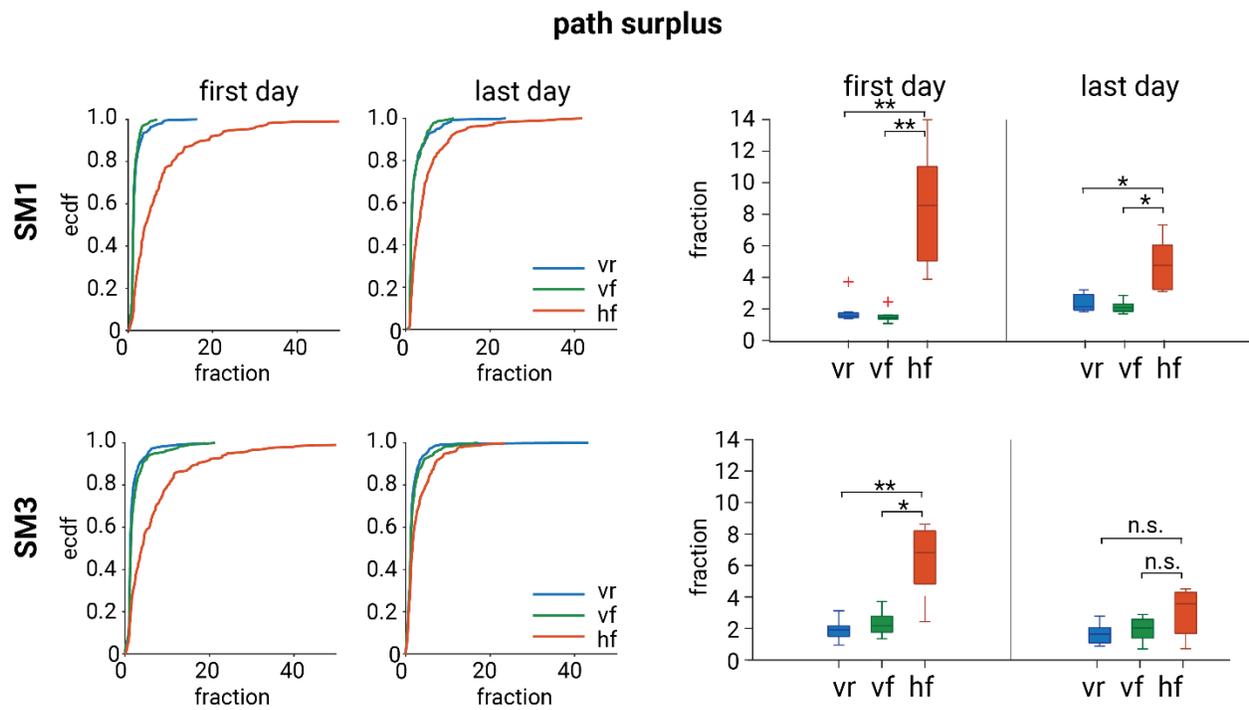


Figure S4 Path surplus as a metric for task performance

Path surplus, calculated as the fraction of actual distance travelled by the mouse to the shortest path possible between trial start point and trigger entry, on the first and last day of training shown for SM1 (*upper row*) and SM3 (*lower row*) as empirical cumulative distribution function (ecdf) (*left*) and boxplot (*right*). Wilcoxon rank sum test, Bonferroni correction, ** $p < 0.0016$, * $p < 0.008$. $n = 8$ mice.

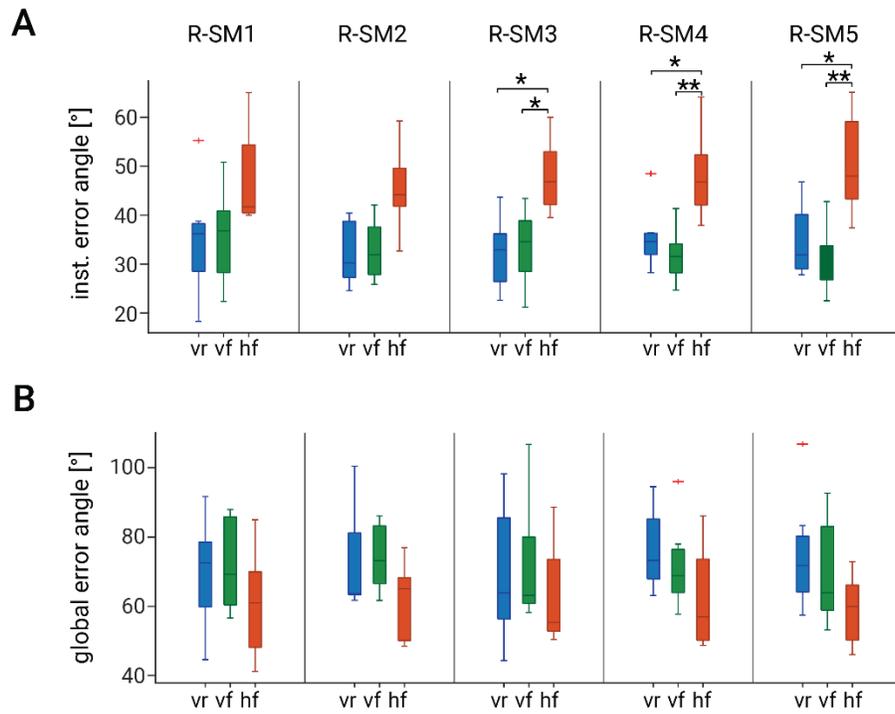
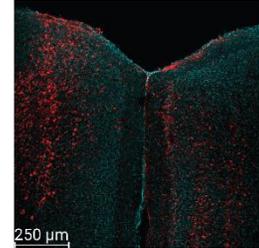
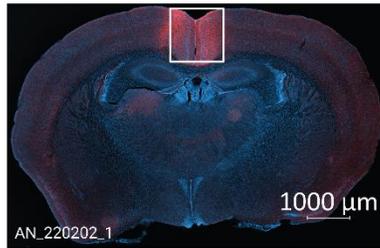
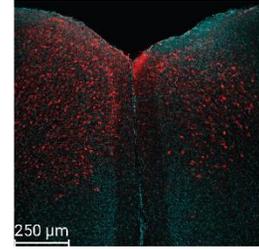
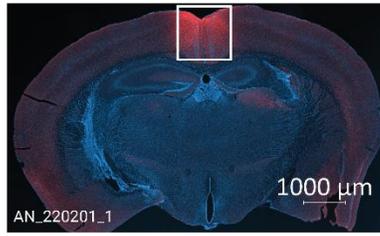
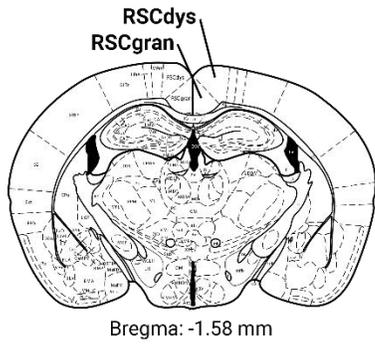


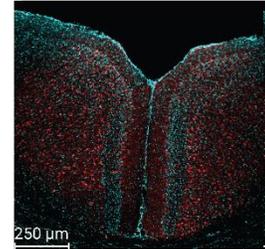
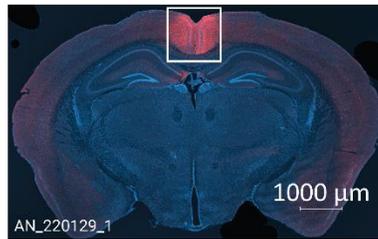
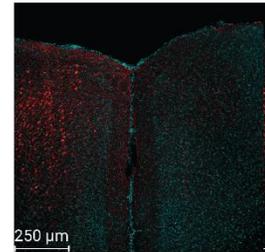
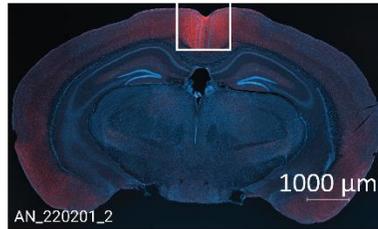
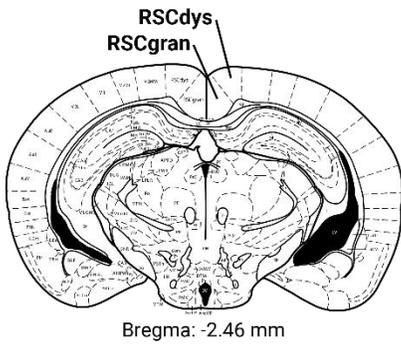
Figure S5 Task performance during random trigger training

A. Boxplot comparison of the mean instantaneous error angle for the three trial types (visible-random in *blue*, visible-fixed in *green*, and hidden-fixed in *red*) for the five days of random trigger placement training (R-SM1-5). B. Same as in A, but for mean global error angles. *Wilcoxon rank sum test, Bonferroni correction, ** $p < 0.0006$, * $p < 0.003$. $n = 8$ mice.*

A



B



C

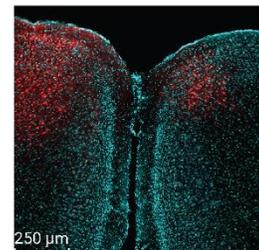
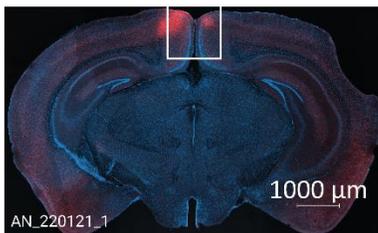
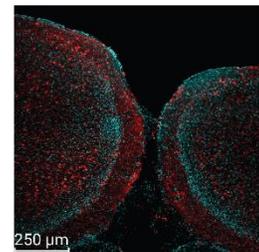
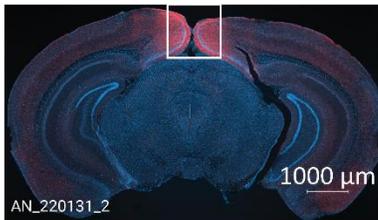
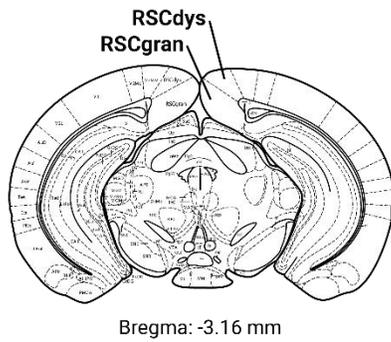


Figure S6 Expression of inhibitory DREADDs in RSC neurons

Coronal slices from six example mice are shown along the rostro-caudal axis, at positions -1.58 mm (A), -2.46 mm (B), and -3.16 mm (C) caudal to bregma, where viral injections were targeted. *Left column*, schematic depicting anatomical location of dysgranular and granular RSC. *Middle column*, epifluorescent images showing expression of hM4D(Gi)-mCherry in RSC neurons in *red* and DAPI staining in *blue*. Zoom-ins on the right are indicated by *white squares*. *Right column*, zoom-in images obtained with a confocal microscope showing expression of hM4D(Gi)-mCherry in RSC neurons in *red* and DAPI staining in *cyan*.

Schematics (*left*) of the coronal sections were adapted from (Paxinos & Franklin, 2001)

List of Figures

Figure 1.1 Different mazes to evaluate spatial memory.....	3
Figure 1.2 The four fundamental neuronal response properties of allocentric spatial coding.....	10
Figure 1.3 RSC anatomy and connectivity	15
Figure 1.4 The BBB model.....	21
Figure 2.1 Experimental timeline.....	28
Figure 2.2 Cranial window placement and virus injection sites	30
Figure 2.3 Behavioural setup	31
Figure 2.4 Arena cue rotation	32
Figure 2.5 Session structure for the chemogenetic and behavioural cohort	34
Figure 2.6 Reward port selection.....	35
Figure 2.7 Placement of fixed trigger locations	36
Figure 3.1. Task performance	44
Figure 3.2 Example occupancy maps	45
Figure 3.3 Error trials	46
Figure 3.4 Example session ethogram	47
Figure 3.5 Example trial ethograms	49
Figure 3.6 Time to trigger	50
Figure 3.7 Trigger precision	51
Figure 3.8 Error angle definition and example	53
Figure 3.9 Error angles of trajectories leading to triggers and ports.....	55
Figure 3.10 Angular head velocity	56
Figure 3.11 Schematic of error angle hypothesis	57
Figure 3.12 Example imaging planes.....	59
Figure 3.13 Locomotion during Ca ²⁺ -imaging session	60
Figure 3.14 Chemogenetic inactivation of RSC activity	61
Figure 3.15 Task performance under control and CNO conditions	62
Figure 3.16 Occupancy map examples of a control and experimental mouse.....	64
Figure 3.17 Error angles under control and CNO conditions.....	65
Figure 3.18 Trial count comparison of the behavioural and chemogenetic cohort	68
Figure 3.19 Example occupancy maps of the spatial memory training.....	69
Figure 3.20 Trial duration as a metric for task performance	70
Figure 3.21 Error angles of different trial types.....	72
Figure 3.22 Timeline and logic of arena cue and trigger placement manipulations	76
Figure 3.23 Instantaneous error angle during arena cue and trigger location manipulations.....	80
Figure 3.24 Timeline and occupancy maps during random trigger placement	81
Figure 3.25 Task performance during random trigger training	83
Figure S1 Task performance: trial duration and time-to-port	127
Figure S2 Individual mouse task performance.....	128
Figure S3 Time from start search point to trigger entry	129

Figure S4 Path surplus as a metric for task performance	130
Figure S5 Task performance during random trigger training	131
Figure S6 Expression of inhibitory DREADDs in RSC neurons	133

Abbreviations

3D	Three-dimensional
AHV	Angular head velocity
ATN	Anterior thalamic nuclei
CNO	clozapine N-oxide
DREADD	Designer receptors exclusively activated by designer drug
FOV	Field of view
hf	Hidden, fixed trigger
IEG	Immediate early genes
IR	Infrared
LD	Laterodorsal thalamic nucleus
mEC	Medial entorhinal cortex
MWM	Morris water maze
NT	'No trigger' error trial
R-SM 1/2/3/4/5	Randomised spatial memory location 1/2/3/4/5
RAC	Rotated arena cues
ROI	Region of interest
RTL	Rotated trigger locations
RVF	Rotated visible, fixed trigger location
RSC	Retrosplenial cortex
SM 1/2/3	Spatial memory location 1/2/3
TNP	'Trigger, no port' error trial
vf	Visible, fixed trigger
VR	Virtual reality
vr	Visible, random trigger
WP	'Wrong port' error trial

Acknowledgements

Completing this PhD project has been both a challenging and enriching experience, and has only been made possible by the unwavering support of many incredible individuals to whom I owe my deepest gratitude. Each of you has contributed to this journey in invaluable ways, and for that, I am deeply thankful.

I am profoundly grateful to all the people who supervised me along the way, providing invaluable guidance and support. **Sandra Reinert** is one of the most exceptional young neuroscientists I know. She embodies all the values necessary to succeed in science and academia, while never losing sight of showing kindness and compassion. She has provided not only exceptional supervision and feedback on my project and this thesis itself, but also rekindled my passion for the project when I got lost in all the nitty-gritty details of planning experiments and in the daily struggles of PhD life. Her lasting influence on me and the project extends far beyond her time in the lab, and I hope many others will benefit from her compassionate guidance. **Drago Guggiana Nilo** played a crucial part in guiding the analysis of this project in the second part of my PhD. He understood the struggles I faced, offering encouragement and kind words when I needed them the most. **Tobias Rose** was instrumental in helping me design this project at the beginning of my PhD.

Special thanks to my Doktorvater, **Tobias Bonhoeffer**, for welcoming me into the vibrant lab community he has created, full of wonderful people. He has been a role model to me with his enthusiasm for technological advances and through his leadership, characterised by fairness and compassion, nurturing the environment that is needed to tackle difficult scientific challenges. His continuous support has been vital, extending beyond daily lab operations and the scope of this project.

I want to thank the members of my thesis advisory committee, **Jonathan Whitlock** and **Christian Leibold**, for their constructive feedback and advice which significantly improved this project.

The lab technicians have been my everyday heroes. **Volker Staiger** has been first and foremost a friend, always lending a supportive ear whenever needed. He not only taught me how to prepare brain samples, but also helped me find whatever I needed in the abyss of the lab cupboards. **Frank Voß** ensured that mice were always available whenever I needed them thanks to his meticulous animal breeding. I also want to thank him for never giving up, no matter how small the pieces I needed soldered together. **Claudia Huber** greatly assisted me with the histology aspects of the project and I also want to thank her for introducing me to the art of home-brewed kombucha. **Max Sperling** provided relentless technical support both in the setup and office. I want to thank him for not laughing when I used unconventional terms in describing the PC related issues I faced, as well as for ensuring that our coffee machine was always up and running. **Dominik Lindner** greatly helped with the histology and was particularly helpful with the last-minute cell counting when I had to prepare for an upcoming TAC meeting.

The administrative support from **Daniela Pade** and **Meike Hack** has been indispensable, helping me to smoothly navigate the often-complex administrative waters.

To all the **BonBros and BonBabes**, both past and present – thank you for all the unforgettable memories from our trips to Montecastelli in Tuscany, our ski and summer retreats, and our daily lunches and

excursions to the far away Mensa. Our daily laughs and your support have made PhD life a lot easier and immensely more enjoyable.

A heartfelt thank you to my **friends** in Vienna, Munich, London, and around the world. Your patience and understanding when scheduling hangouts around my experiments, as well as your continued support despite the logistical challenges and geographical distance, have meant the world to me. I have yet to find a way to inform mice about weekends.

To my **family** – my parents who instilled the values of education and curiosity in me; my mother, who exemplifies lifelong learning like no other; my father who paved the way for my academic pursuits; and my sisters, Sophie and Laura, who make me proud to be a big sister – your support has been foundational in my journey and I couldn't have made without you.

Last, but not least, to my **Peer** and partner. Thank you for supporting me along this journey, even when that meant moving to yet another city in yet another country. I started this journey with you by my side and I am glad that I will end this chapter with you by my side. You helped me out in more ways than I ever thought possible and I can't wait to start this next chapter of our life together. Thank you for always believing in me and making me feel like no mountain is too high.