THE UNIVERSITY
of EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

The Black Box Effect:
Investigating the Role of Retroactive Interference on Hippocampal Memory Mechanisms

Daisy Arkell

Doctor of Philosophy
Deanery of Biomedical Sciences
College of Medicine and Veterinary Medicine
University of Edinburgh
2018

Abstract

Recent studies have shown that decreasing sensory stimulation after learning can enhance memory retention in humans. Amnesic patients and healthy controls expressed significantly better memory for both passages of prose and spatial landmarks when learning was followed by a short period filled with restful wake, rather than an unrelated distractor task (Dewar et al, 2010; Craig et al, 2016). This enhancement was suggested to arise from decreases in memory interference processes. These findings suggest that interference from ongoing sensory stimulation could have a much larger impact on memory and everyday life than previously thought. The aim of this thesis was therefore to investigate the role of retroactive interference in hippocampal-dependent memory consolidation, and to explore the neural mechanisms behind this episodic memory enhancement.

To this end, I tested the effects of reducing different types of interference after spatial learning on memory retention in rats. A spatial memory task was used that required no reward, instead using the animal's natural tendency to detect and explore novelty. To exploit this behaviour experimentally to test memory retention, I used the novel object location (nOL) recognition task. My protocol consisted of a single training trial, during which animals could explore two copies of the same novel object placed in an open field arena. Memory for the object locations was then tested 6h or 24h later, when animals were returned to the arena in which now one of these objects was moved to a novel location. Animals that preferred to explore the object at the novel location expressed memory for the location the objects occupied during the training trial. The role of interference on object location memory was assessed by exposing the animals to different, highly-familiar stimuli (i.e., dark or normally lit holding box, home cage, or cagemate in a holding box) during the 1 h period directly following the training trial. We used gentle handling to prevent rats from falling asleep during this period. I found that animals expressed robust nOL memory when exposed to a dark familiar holding box after learning, but not when they were exposed to their home cage, replicating the memory enhancement effect following reduction of visual stimulation seen in humans. Further experiments sought to isolate what aspects of the dark holding box promoted memory retention as compared to the home cage. To this end, after learning, animals were put into their home cages with their cage mates, which was placed in either an enclosed normally lit (white light) box (WB), or an enclosed dark (red light) box (RB). Neither group expressed memory, suggesting that the black box effect was dependent on animals

being socially isolated. Exposure to the WB when alone also prevented the expression of nOL memory. Yet, animals exposed to the RB without cage mates expressed object location memory, establishing that the black holding box effect was dependent on animals being socially isolated and with reduced visual stimulation. These results suggested that interference not only stems from new learning, but can occur simply when exposed to either highly familiar social or visual stimuli.

Object location memory is known to depend on the hippocampus. The activity of pyramidal neurons within the hippocampus (place cells) represents the location of an animal within its environment. This activity is context-dependent, and has been shown to be modulated by the manipulation of objects within these environments (Deshmukh et al, 2013; Burke et al, 2011). Therefore, to explore the neural mechanisms underpinning the 'black box effect', I recorded place cells in the dorsal CA1 of rats. I first focused on the spatially-selective firing of place cells to study whether post-learning stimulation could affect the spatial stability of place cell firing within a novel environment, thereby causing memory interference. Animals explored a novel environment for 10 min, after which they spent 3 h awake in either the WB or RB. Then, 6 h after the initial exposure, animals explored the same environment again. Analysis of place cell firing indicated that whilst the overall firing and spatial properties of place cells were not different between groups, the stability of place fields between the initial and repeated exposures was significantly enhanced in the dark (RB) box group. Therefore, reducing visual stimulation after learning promoted place field stability, consistent with the behavioural results.

To determine whether these changes in place field stability correlated to the strength of object location memory, a third set of experiments investigated the influence of objects on place field expression during a nOL behavioural task. As seen previously, implanted rats expressed object location memory for 6 h when exposed to the RB, but not the WB, after learning. In contrast to these findings, no differences in the firing and spatial properties of place fields both over and between sessions were found between the WB and RB groups. The introduction, movement and removal of objects, however, did affect various measures of place field stability and synchronicity. The apparent object-place field relationship was investigated further, and results suggested that place fields were more likely to be expressed

away from objects during the probe trial if the animal had significant memory for the object locations.

Overall, the results reported in my thesis show that long-term memory formation, in terms of behavioural as well a subset of electrophysiological measures, benefits from reduced sensory stimulation after learning. These findings highlight that even low levels of sensory stimulation can have a drastic impact on spatial memory and correlated neural activity. This has important implications for experimental design, as well as life outside of the laboratory.

Lay summary

Whilst the forgetting of irrelevant information is an important part of everyday life, it can be infuriating when a piece of information you learned just minutes ago has now gone from your memory - even more so when it appears that receiving a phone call or talking to an acquaintance has sped this process up. This disruption of memory caused by incoming information is termed retroactive interference. Much research has covered the topic of both pathological and healthy memory loss due to retroactive interference in humans. It has been shown previously that if you sit alone in a quiet dark room (wakeful rest) after learning a passage of prose you will be able to remember the stories much better than if you had to complete a spot the difference task during this time after learning. This suggests that the reduction of retroactive interference directly after learning can increase memory retention in the long-term. This thesis aims to correlate some of these findings into rats, allowing for a greater depth of investigation into these forgetting mechanisms.

To this end, a memory task was used that required rats to learn the locations of objects. Rats that were put into a black box by themselves after learning these locations, showed significantly better memory for the locations when tested 6 hours later, compared to rats that were put back into their home cage in the light with their cage mates. This replicated findings in humans showing that rats could also benefit from wakeful rest by being put into a black box, i.e. the 'black box effect'. Further experiments found that rats had to be by themselves and in the dark for this benefit to occur. If rats were by themselves in the light, or if they were with their cage mates in the dark, they could not remember the object locations. These results suggested that memory could be disrupted by being able to see the inside of the box or by interacting with cage mates.

The hippocampus is a brain area important for navigation and conscious memory. The hippocampus contains cells which tell you where you are within the environment, by becoming active only when you are in a specific place within in the environment. . These 'place cells' are thought to underlie the memories associated with location. In this study it is required for animals to remember object locations in a memory task. Therefore I recorded place cells in rats to see if their activity is affected by the black box effect. I found that the activity of a percentage of place cells was affected by the black box effect, suggesting that the 'black box effect' can be observed at the single cell level.

<u>Acknowledgements</u>

Declaration

I declare that all presented work is my own unless stated otherwise; and I composed this thesis myself. The work in this thesis has not been submitted for any other degree or professional qualification.

Daisy Arkell

Chapter 1: Data collection for the following novel object location behavioural tasks was performed by students under my supervision. Experiment 1 was partly carried out by Ahnhilli Zhuparris, Chan Chin Hang, Erika Zeigyte and Maria Lopez Quiroga during an undergraduate honours project. Experiment 4 was partly carried out by Emily Sweet during a summer placement in the lab. Emily Sweet also manually scored sampling session videos for this experiment.

Chapter 2 and 3: Perfusions were performed by Dr Anna Smith if the animals had been used in subsequent experiments outside of this thesis. The majority of histology nissl staining was performed by Jane Tulloch and Dr Rosie Jackson.

Pre- and post-clustering MATLAB scripts were written by Steven Huang and edited by both myself and Dr Elizabeth Allison. Population correlation analysis scripts were written by Dr Antonis Asiminas. Scripts performing bootstrap analysis and analysing place field movement were written by myself.

List of abbreviations

A	Amygdala
ACC	Anterior cingulate cortex
ACh	Acetylcholine
AMPA-R	α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptor
ANOVA	Analysis of variance
AP	Anterior-posterior axis
B1	First baseline session
B2	Second baseline session
B3	Third baseline session
b/s	Bits per spike
C1	Context 1
C2	Context 2
Ca^{2+}	Calcium ion
CA1	Cornu ammonis 1
CA2	Cornu ammonis 2
CA3	Cornu ammonis 3
CaMKII	Ca^{2+}/calmodulin-dependent protein kinase II
Cs	Collateral sulcus
d	Discrimination index
d[0-180]	Discrimination index of 3 minute probe trial
DV	Dorso-ventral axis
e-LTP	early long-term potentiation
EC	Entorhinal cortex
fMRI	Functional magnetic resonance imaging
FR	Firing rate
H	Hippocampal formation
HPC	Hippocampus
IsoD	Isolation distance
LEC	Lateral Entorhinal cortex

LFP	Local field potential
LTM	Long term memory
LTP	Long-term potentiation
l-LTP	Late long-term potentiation
MAP-K	Mitogen-activated protein kinase
ML	Medio-lateral axis
MMN	Medial mammillary nucleus
mPFC	medial prefrontal cortex
Na^+	Sodium ion
NMDA-R	N-methyl-D-aspartate receptor
nOL	Novel object location
P	Probe session
p190 RhoGAP	p190 Rho GTPase-activating protein
PCo	Population coordination
PKA	Protein kinase A
PKC	Protein kinase C
PR	Perirhinal cortex
PSD	Post synaptic density
RB	Red box
REM sleep	Rapid eye movement sleep
RI	Retroactive interference
S	Sampling session
SEM	Standard error of the mean
SI	Spatial information content
STM	Short term memory
SWS	Slow wave sleep
vmPFC	Ventro medial prefrontal cortex
WB	White box

Table of Contents

Chapter 1: General Introduction

Whilst the forgetting of irrelevant information is an important part of everyday life, it can be infuriating when a piece of information you learned just minutes ago has now gone from your memory - even more so when it appears that receiving a phone call or talking to an acquaintance has sped this process up. This disruption of memory caused by incoming information is termed retroactive interference. Much research has covered the topic of both pathological and healthy memory loss due to retroactive interference in humans. It has been shown that the reduction of such interference directly after learning can increase memory retention in the long-term. This thesis aims to correlate some of these findings into non-humans, allowing for a greater depth of investigation into these forgetting mechanisms. In this introduction I will cover the role of the hippocampus in memory in both humans and non-humans; the implications of retroactive interference on the forgetting of hippocampal memories in humans; the processes thought to underlie both synaptic and systems consolidation of such memories; and place cells – a possible neural correlate of spatial memory. I will finish by stating the overall aims and hypotheses of this thesis.

1.1 Hippocampal memory in humans and non-humans

Up until the midpoint of the twentieth century it was believed that memory was distributed throughout the brain with no particular locus. This all changed in 1957 with the report of patient H.M. H.M. had such severe epileptic seizures that he underwent bilateral medial temporal lobe resection in an attempt to control them (Scoville and Milner, 1957). This lesion included the hippocampus and surrounding structures, such as the amygdala and parahippocampal gyrus (Corkin et al., 1997), highlighted in Figure 1.1. Although higher capacities were left intact, H.M. suffered from profound forgetfulness for declarative memories some years before the surgery (retrograde amnesia), and for all declarative memories after it (anterograde amnesia) (Gabrieli et al., 1988; Penfield and Milner, 1958; Scoville and Milner, 1957). This severe inability to form new declarative memories helped to identify the medial temporal lobe (MTL) as an important locus for memory in the human brain.
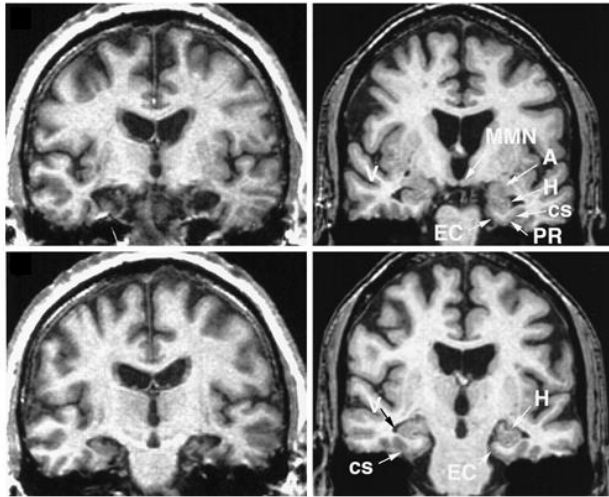
Figure 1.1: A magnetic resonance imaging scan of the brain of H.M (left) and a healthy age-matched control (right). H.M. had a bilateral medial temporal lobe lesion, affecting the hippocampal formation (H), entorhinal cortex (EC), perirhinal cortex (PR), amygdala (A), collateral sulcus (cs) and medial mammillary nucleus (MMN). This lesion produced profound anterograde and retrograde amnesia. This figure is adapted from Corkin et al., 1997.

Through years of thorough neuropsychological testing of H.M., and patients with similar MTL damage, the specificity of hippocampal-dependent memory began to be elucidated. Whilst these patients couldn't consciously acquire any new memories, motor skills could be learned, albeit with no recollection of the learning event (Brooks and Baddeley, 1976; Cohen and Squire, 1980). This showed the separation of two important forms of memory: declarative and procedural. The acquisition of procedural memory – unconsciously learned skilled-based information – clearly did not require the MTL, as this type of memory remained intact in amnesic patients. The acquisition of declarative memory – conscious memory containing information about facts and events – on the other hand did require the MTL. It is also important to note that, whilst patients with MTL damage could not consciously remember any memories minutes after learning, if the information learned could be held in short-term or working memory with no distractions then it could be rehearsed and remembered. However as soon as attention was placed elsewhere this memory was gone (Milner, 1959, 2005; Sidman et al., 1968). The MTL is therefore clearly involved in the acquisition of long-term declarative memory, with short-term or working memory and attention being spared. Studies into patients with more focused damage to just the hippocampus implicated this structure in specific types of long-term declarative memory. The acquisition of memories rooted in contexts, i.e. episodic and spatial memory, appeared to be selectively impaired, whereas the acquisition of general knowledge and memories based around fact, i.e. semantic memory, appeared to be spared (Tulving and Markowitsch, 1998; Vargha-Khadem et al., 1997). Although there has been much debate surrounding exactly what type of memory is hippocampal dependent in humans (Manns et al., 2003), studies into non-humans have provided more conclusive results.

As the MTL structure is highly conserved between species (Figure 1.2), animal models have been used to replicate these findings and investigate the impact of more focused lesions of just the hippocampus, rather than the other MTL structures. This profound anterograde amnesia for long-term spatial memory, coupled with a complete sparing of procedural memory and relative sparing of short-term memory, has been replicated in various animal models of amnesia with varying degrees of success (Murray and Wise, 2010). When using delays of 40 minutes in between learning and recall it was found that monkeys with hippocampal lesions were able to perform as well as controls on novel object recognition and delayed non-matching to location tasks (Murray and Mishkin, 1998). These were both tasks thought to require the hippocampus. However the distinction between short-term and long-term memory is important, as delays of at least 1-3 hour are thought to be needed for the testing of long-term memory (Grecksch and Matthies, 1980; Izquierdo and Medina, 1997). The distinction between purely spatial memory and spatially-related memory is also imperative. Studies both in monkeys and rodents have shown that if memory tests requiring the acquisition and subsequent recollection of purely spatial memory are selectively used, then anterograde amnesia becomes apparent (Aggleton et al., 1986; Clark et al., 2005; Hampton et al., 2004; Lavenex et al., 2006; Murray and Mishkin, 1998; Murray and Wise, 2010). These findings therefore suggest that the consolidation of long-term episodic and spatial memory is hippocampal-dependent.
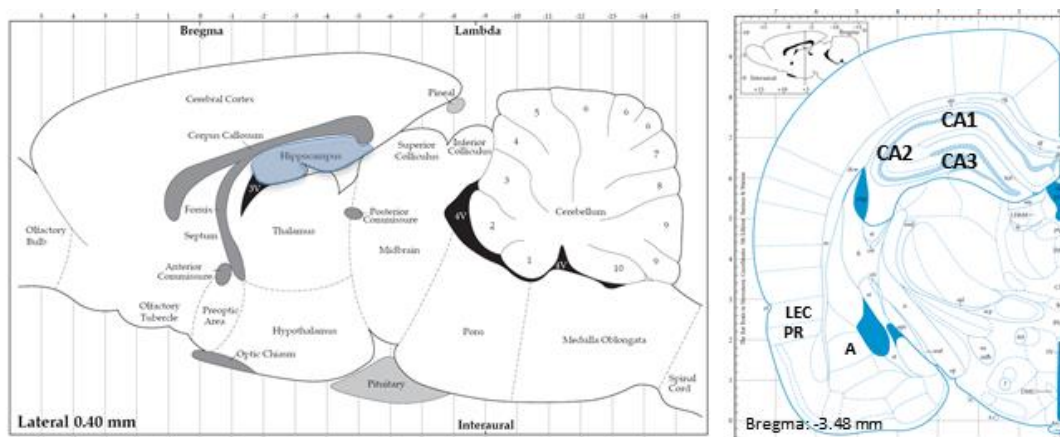
Figure 1.2: A sagittal section of a rat's brain highlighting the location of the hippocampus in blue (left). A coronal section of a rat's brain highlighting the location of the hippocampal formation, including the hippocampal subfields. Other parts of the medial temporal lobe (MTL) are also labelled: perirhinal cortex (PR), lateral entorhinal cortex (LEC) and amygdala (A). MTL structures are highly conserved between species giving relevance to non-human studies of this part of the brain. Figure adapted from Paxinos and Watson, (2004).

1.2 The hippocampus, forgetting and the role of retroactive interference

The conscious acquisition and consolidation of memories is dependent on the hippocampus. However in the healthy population there are often times where memories are consciously learned but become unable to be recalled soon after. There are two main theories behind this non-pathological forgetting, one involving the active decay of memories that have been consolidated and stored, and one involving the interference of memory consolidation or retrieval. The former is thought to occur almost selectively during sleep, systemically removing weaker memories with only those sufficiently 'protected' from active decay processes persisting (Hardt et al., 2013). Interference, on the other hand, occurs when new information disrupts consolidation or retrieval of an old memory (retroactive), or occurs when previously learned information disrupts encoding or memory of new information (proactive). Whilst the theory of active decay is significant and merits further discussion, only the theory of retroactive interference will be covered in this thesis.

1.2.1 Forgetting as a function of time

Among the first to address the issue of forgetting was Ebbinhaus, who in 1885 published his 'forgetting curve', recreated in Figure 1.3. Ebbinhaus taught himself lists of nonsense syllables and subsequently attempted to recall them after increasingly long intervals. He plotted a curve to show how his retention of this learned material decreased as the delay between learning and recall got longer. Forgetting therefore appeared to occur as a function of time. However soon after, Bigham highlighted a problem with the passive decay theory Ebbinhaus had suggested, arguing that the time between encoding and recall of the learned material was seldom void of incoming information. This problem was investigated some years later by Muller and Pilzecker (1900), who proposed that rather than being a function of time, forgetting was driven by the everyday tasks and subsequent processing of new information that occupied this delay between learning and recall.
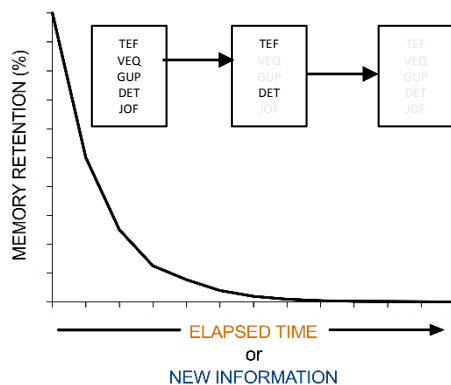
Figure 1.3: Ebbinhaus's forgetting curve plotted memory retention of nonsense syllables over the elapsed time between learning and recall. (orange). Muller and Pilzecker hypothesised that rather than a function of time, forgetting occurred due to the presence of new information between learning and recall (blue).

1.2.2 Forgetting due to retroactive interference

Muller and Pilzecker created the basic methodology for the study of retroactive interference. To investigate this proposed mechanism of forgetting, participants were presented with nonsense syllable pair lists and asked to recall one of the pair when the other was presented to them. The material to be learned was therefore similar to the lists of nonsense syllables used by Ebbinhaus. However in this case the delay period in between the encoding and recall of these syllables was either filled with the learning of another nonsense syllable list or unfilled and comparably void of incoming information. As predicted by Muller and Pilzecker's proposed mechanisms of forgetting, the filled delay period led to a lower recall performance than the unfilled delay period, and had more of a detrimental effect if it directly followed learning.

From this the retroactive interference theory of forgetting was born. This proposed that new learning interfered with old learning in a time-dependent manner, with interference more likely to occur directly after the original material was learned. However it was unclear whether interference occurred due to the similarity of the material to be learned (i.e. they were both lists of nonsense syllables) or whether any new learning would cause this interference effect. The experiment was therefore repeated; this time with the filled delay period consisting of viewing landscape paintings and subsequently describing them to the experimenter in great detail. The same pattern of recollection was seen, with the filled delay leading to a lower recall performance than the unfilled condition. Muller and Pilzecker hypothesised that interference occurred due to mental exertion after learning, rather than the similarity of the new and old learned material.

This question was touched upon by future studies using delay conditions containing material of varying similarity, however many of these studies disagreed with Muller and Pilzecker entirely. Both McGeoch and McDonald (1931) and Dey (1969) argued that material encountered during the delay period had to be similar to the learned material to cause retroactive interference. When the delay was filled with non-similar information participants performed better than when the information was the same as the learned material (Dey, 1969; McGeoch and McDonald, 1931). For example, McGeoch and McDonald implemented a task requiring the learning of 10 adjectives followed either by the learning of interpolated stimuli of varying degrees of similarity, or a 'rest' condition involving reading and choosing

their favourite jokes from a book. The interpolated stimuli included synonyms, antonyms, unrelated adjectives, syllables and numbers. Recall after the delay improved as the similarity of stimuli decreased, with the 'rest' condition showing the best levels of recollection overall. These studies dismissed the existence of interference caused by non-similar material, i.e. diversion retroactive interference (diversion-RI). Although others proposed that both similar and non-similar material could lead to retroactive interference (Skaggs, 1925), diversion-RI was largely ignored until the start of the next century.

The important thing to note is that neither McGeoch and McDonald nor Dey included an unfilled delay condition. Even when a rest interval was used, participants were asked to read or talk to the experimenter. This is in contrast to the true unfilled delay used by Muller and Pilzecker as a control to their diversion interference condition. This suggests that the rest interval could have been acting as a filled delay period, especially since the memory task involved participants learning words, and both reading and speaking are verbal tasks. Another issue that was not addressed by either Muller and Pilzecker or McGeoch, McDonald and Dey's subsequent studies into familiarity was that of whether interference only occurred when participants were required to learn something, as had been the case with the interpolated stimuli used previously, such as the syllable list and the picture task. As much of everyday life is full of the incidental learning of unrelated material, or even just the processing of irrelevant stimuli, this distinction was an important one to make. Therefore for interference to play a common role in everyday forgetting, true diversion-RI had to be a plausible mechanism.

With this in mind, Dewar et al. investigated the possibility that retroactive interference does not require the learning of meaningful information to disrupt the future recollection of memories (Dewar et al., 2007). To this end, participants were required to learn a list of 15 verbally presented nouns followed by a delay period filled with one of various tasks. Of these tasks, two required learning and future recall (listening to a radio recording or watching visual clips of scenes around the University campus and subsequent follow up questions), two contained meaningful material that were not required to be learned (spot the difference or maths problems), and one contained no meaningful material and no learning (a tone detection task). Of the two tasks that required recall, one was verbal and therefore similar to the material to be learned, and one was non-verbal and therefore dissimilar to the learned

material. A true unfilled delay period was also used as a control condition, with participants left by themselves in a quiet, darkened room to rest for the duration (termed wakeful rest).

As predicted the unfilled wakeful rest condition exhibited the best noun recall after the delay, as had been found 100 years earlier by Muller and Pilzecker. Surprisingly all of the other conditions produced the same levels of retroactive interference, with no significant differences between the percentages of correctly recalled nouns. This highlights many findings. Firstly, when the task requires intentional learning, the similarity of the task to the learned material does not affect interference, in contrast to previous findings. This discrepancy could arise from the levels of similarity of the task to the material learned, as although both required listening and recalling verbal content, the material recalled was purely a list of nouns, whereas the task required listening to a radio show and extracting information. It has been suggested that the task has to be almost identical to the learned material to cause interference. This was previously shown by Robinson, who found that when learning a list of eight different four digit numbers interference was much greater when followed by the learning of eight more four digit numbers compared to the multiplication of four digit numbers or the learning of a string of 32 numbers (Robinson, 1920). Therefore, unless highly similar, all material appeared to have a comparable detrimental effect. It should be noted that proactive interference also exists, where previously learned information disrupts encoding or memory of new information, however it appears that this type of interference relies solely on the similarity of the material (Dewar et al., 2007). This suggests that similarity retroactive interference might be comparable to proactive interference and possibly caused by different mechanisms to that of diversion-RI.

Another finding from Dewar et al. was that the delay task did not need to be intentionally learned to cause retroactive interference. Importantly, the delay task also did not need to contain any meaningful new information as in the case of the tone-detection task, which required mental effort but no unintentional learning of novel or informative material. It was also confirmed that a true unfilled delay does reduce the effects of retroactive interference. These findings all agree with Muller and Pilzecker's initial hypothesis of retroactive interference, where interference occurred due to mental exertion after learning rather than the similarity of the new and old learned material. Overall this study showed that diversion-RI is a real phenomenon and therefore a plausible candidate for everyday forgetting. They

also highlighted that the removal of interfering stimuli through the unfilled condition of wakeful rest can lead to the highest levels of retention, suggesting that wakeful rest enhances memory.

Whilst these results indicated that consolidation mechanisms were modulated by interference, the delay periods used between acquisition and test were short. Therefore the proposal that interference affects long-term memory retention processes was clarified in a further experiment (Dewar et al., 2012). Participants listened to passages of prose followed immediately by either a 10-minute filled delay containing a visual spot-the-difference task or a 10-minute unfilled delay in a dark and quiet room. 15-30 minutes later participants were asked to recall the prose, followed by a further surprise recall test 7 days later. There were striking improvements in recall both 15-30 minutes and 7 days after the initial learning session for participants in the unfilled delay condition. This indicated that the unfilled delay immediately after encoding improved memory in the long-term.

Whilst the rehearsal of learned material during the unfilled delay period was unlikely to underlie the increased recall in this condition, this notion remained to be directly tested. Therefore an experiment was carried out where, instead of participants learning known words, they were asked to learn and subsequently recall a list of non-rehearsable non-word items (Dewar et al., 2014). As with the long-term retention experiments described previously, the delay lasted 10 minutes and was either filled with a spot-the-difference task or unfilled. This was followed by surprise recall tests both 15-30 minutes and 7 days after the initial learning. As had been seen previously, the unfilled condition boosted memory retention at both time points compared to the filled condition, indicating that intentional rehearsal was not necessary for the memory enhancing effects of reduced interference.

Further experiments sought to test the effects of an unfilled condition of wakeful rest on spatial memory. Instead of learning prose or lists of words, participants learned a route through a computer-based virtual reality environment containing eight decision points with a distinctive landmark positioned at each (Craig et al., 2015). Immediately after the route was learned to a 100% criterion participants underwent either a filled (spot-the-difference task) or unfilled (wakeful rest) delay condition of 10-minutes. Spatial memory was tested after 10-20 minutes, followed by a surprise test 7 days later. These tests included associative memory

for landmark-direction associations (i.e. which way did you turn at a particular landmark); a cognitive map test where on the presentation of one landmark participants had to give the compass direction to another specified landmark; a temporal landmark order test; and a route memory test. No effect of wakeful rest was seen in the route memory test, as there appeared to be a ceiling effect due to relative easiness of the task. No effect was seen in the cognitive map test either, however this appeared to be due to a floor effect due to the relative difficulty of the task. Wakeful rest did enhance both associative and temporal order memory at both 10-20 minutes and 7 days post-route learning. This experiment was repeated using a much larger virtual environment and longer routes (Craig et al., 2016). The cognitive map test was repeated, however this time participants were put back into the virtual environment next to a landmark and asked to rotate to face another landmark within the environment. Therefore participants were given the surrounding contextual information as well as the landmark position. The floor effect was removed and participants were able to complete the task in both conditions, with more accurate performance after wakeful rest condition compared to the filled delay condition. This suggested that wakeful rest was assisting the consolidation of spatial memories into an internal representation of the environment, allowing the participant to more accurately visualize short-cuts and routes never taken before.

1.2.3 Increased susceptibility to diversion retroactive interference in amnesic patients

In parallel to the research described above, the same laboratories started to apply the concept of diversion-RI to patients with severe anterograde amnesia, either caused by mild cognitive impairment (MCI) or focal brain damage (Cowan et al., 2004; Della Sala et al., 2005). Focal brain damage was caused by a stroke or major head injury, and through computerised tomography (CT) or magnetic resonance imaging (MRI) scans it was found that the damage was restricted to the frontal lobe, thalamus, or parietal lobe. These patients therefore appeared to have intact temporal lobes, unlike the studies mentioned previously where damage was restricted to the MTL. It was ensured all patients had intact short-term memory and verbal reasoning, whilst presenting with anterograde amnesia that had had an abrupt onset after brain damage. Patients with MCI had no signs of any focal lesions (examined using CT or MRI scans) and presented with anterograde amnesia in the absence of any other deficits such as a decrease in verbal reasoning or short-term memory loss. Therefore whilst it was possible that the MTL were affected in MCI and focal brain damage patients, it

appeared that this was not the locus of their memory deficits. It was suggested that rather than being unable to encode new information, as appeared to be the case with patients such as H.M., damage to non-temporal structures could interfere with the filtering of incoming information, overloading the temporal structures with information and leading to faulty consolidation mechanisms. This would manifest itself as an increased susceptibility to diversion-RI, and therefore give patients a significant benefit from the unfilled condition of wakeful rest.

To test this hypothesis, both amnesic patients and healthy age-matched controls were read a list of words to learn which they immediately recalled, followed by a either a delay condition filled with various psychometric tests, none of which required learning and none of which contained material similar to the word list, or an unfilled condition of wakeful rest spent alone in a quiet, dark room (Cowan et al., 2004). These delay conditions lasted 10 minutes, after which participants were asked to recall as many words as possible. There were minimal differences between the patient and control groups for the immediate recall with both able to perform the required task. After the delay period both benefitted from the unfilled condition, recalling more words than after the filled condition, replicating the study by Dewer et al. However the differences between the two groups were striking. Whilst the control group was able to recall on average 46% of the words after the filled delay, only half of the patients could recall any words at all (on average 14%) with the other patients unable to recall even the existence of the test. After the unfilled condition the control group was able to recall on average 74% of the words, showing a significant improvement. This improvement was also seen in the amnesiac group, with an overall increase in retention (on average 49%) and with one patient who had previously been unable to remember anything about the testing procedure now able to remember a proportion of the learned words.

This experiment was repeated, however instead of a word list participants had to listen to a passage of prose, followed by a delay of an hour. Participants were also unaware that they would have to recall the prose after the delay (Cowan et al., 2004). This was to minimise the possibility that participants were constantly repeating the prose and therefore using short-term memory to recall the material. In line with the previous experiment, amnesic patients could not remember the majority of the prose after the filled delay condition (7%), however recollection improved dramatically after the unfilled wakeful rest condition (79%). This

suggested that reducing diversion-RI worked remarkably well to enhance the encoding of memory in amnesic patients. In contrast the control group was able to recall 79% of the content even after the filled delay. Although an improvement was shown in the control group after the unfilled condition (89%) as expected, this indicated that a ceiling effect was present, i.e. the task was too easy. Similar improvements were seen in amnesic patients in future studies, such as when a 10-minute filled delay condition containing a tone-detection task was compared with a 10-minute unfilled condition (Dewar et al., 2010), and when patients with mild cognitive impairment were compared to a healthy age-matched control group (Della Sala et al., 2005), indicating the effect was consistent and reproducible. Interestingly, when Cowan et al. (2004) tested another group containing patients with amnesia due to temporal lobe damage, these patients showed no memory after the filled or unfilled condition. Overall these results suggested that patients with non-temporal amnesia are highly susceptible to diversion-RI as hypothesised, and that the benefit of reducing such interference through wakeful rest requires an intact MTL.

Further studies sought to investigate the time-dependent nature of diversion-RI in amnesic patients. Muller and Piltzecker had briefly explored this idea in 1900 in healthy participants, showing that memories were more susceptible to retroactive interference 17 seconds after learning compared to retroactive interference 6 minutes after learning, with 28% of material recalled compared to 49% of material, respectively. This effect was replicated by Skaggs in 1925, whereby participants were asked to remember the locations of five chess pieces on a chessboard, followed by algebra calculations either directly after learning or at different time-points during a five minute delay. Locations were less likely to be remembered if the task directly followed the initial learning of the locations (Skaggs, 1925). These studies suggested that as consolidation mechanisms progressed, memories became less affected by interference. Therefore if amnesic patients were highly susceptible to diversion-RI due to interference with consolidation mechanisms, then this time-dependent nature of retroactive interference should be even clearer in the amnesic patient group.

Dewar et al. explored this effect in amnesic patients with mild cognitive impairment that spared the temporal lobes (Dewar et al., 2009). Patients and age-matched healthy controls were asked to learn a list of 15 words followed by immediate recall of the same list. Following this an interference task was used with various onset delays. The task consisted of verbally

naming pictures whilst ignoring the word that was overlaid on the picture. This required high levels of mental effort but no new learning. This task occurred either immediately after the initial material was learned (early interference), or was delayed by 3 minutes (middle interference) or 6 minutes (late interference). An unfilled wakeful rest delay condition was also used that lasted the whole 9 minutes between learning and recollection of the material. It was found that all participants recalled the most after the wakeful rest delay condition as expected. All participants also had improved recollection when the interference was delayed by 6 minutes compared to directly after learning, indicating that retroactive interference mechanisms are time-dependent. This improvement in recollection was incredibly striking in the amnesiac group, with no patients able to recall any words in the early interference condition, but with recollection comparable to that seen after 9 minutes of wakeful rest in the late interference condition. This implied that amnesic patients are indeed more susceptible to diversion-RI mechanisms due to interference with consolidation mechanisms. As both the early and late interference groups would have interrupted the rehearsal of words within short-term memory this also suggested that retroactive interference mechanisms disrupt long-term memory consolidation processes.

As had been shown in the healthy population, a further experiment sought to clarify whether the benefit of reducing interference in amnesic patients lasted in the long-term over a 7 day period (Alber et al., 2014), as would be expected if long-term consolidation mechanisms were being modulated. This experiment was based on the study outlined previously, investigating long-term processes in the healthy population (Dewar et al., 2012). As had been found in the healthy population, there were striking improvements in recall in the amnesia group both at 15-30 minutes and 7 days after the initial learning session, indicating that the unfilled delay improved memory in the long-term. This improvement in long-term memory in amnesic patients was remarkable given their normal inability to recall memories even a few minutes after acquisition.

1.2.4 Diversion retroactive interference in non-humans

Overall these studies show clearly and consistently that the levels of interference immediately after the learning of hippocampal-dependent memories can affect how much is remembered in the long-term. Reducing these levels of interference post-learning, via wakeful rest, has been shown to benefit amnesic patients with both acute brain damage and MCI as well as healthy controls. This highlights the importance of these findings, both as developing a treatment for memory loss as well as providing a tool for understanding more about the underlying mechanisms of this effect in the healthy population. The hippocampal-dependent and time-dependent nature of these mechanisms, coupled with the fact that memories appear to be modulated in the long-term without the need for constant rehearsal of the material, suggests that retroactive interference modulates the mechanisms underlying hippocampal memory consolidation. Exactly how interference affects these post-learning processes is unknown. Animal studies have tried to replicate the findings of Muller and Pilzecker in an effort to explore the mechanisms behind diversion-RI at a deeper level, however none were successful as these studies focused on movement as an interfering stimuli. Rats underwent spatial learning tasks followed by 1h, 3h or 5h on a rotating drum. A control group underwent no forced exercise. It was hypothesised that forced movement would act as an interfering stimulus in the same way that learning lists of words or numbers interfered with memory in humans. Surprisingly for the experimenters, forced exercise was minimally better than the control group (Corey, 1931) and in one experiment actually enhanced the rat's spatial memory (Gray, 1937), suggesting that retroactive interference was not replicable in non-humans. Although an obvious confounding factor now, at the time it was unknown that exercise could improve spatial learning and memory (Erickson et al., 2011; O'Callaghan et al., 2007; Shih et al., 2013; Wang and Wang, 2016). Therefore whilst studies have attempted to replicate diversion-RI in non-humans, none have been successful, in all probability due to the protocols used. If the effects of reduced diversion-RI could be replicated in rodents then the underlying mechanisms of such effect could be understood in much greater detail. It is possible that retroactive interference is disrupting consolidation at either the cellular level or at a systems level. Here I will review the concept of memory consolidation and the neural machinery involved.

1.3 The hippocampus and memory consolidation

Memory consolidation is the stabilisation of memory after learning, a process that happens progressively over time. Muller and Pilzecker proposed this phenomenon (Konsolidierung) in the late twentieth century, after experiments showing memory for list learning improved greatly over the first few minutes post-acquisition (Muller and Pilzecker, 1900; Lechner et al 1999). Further experiments have shown different types of mechanisms underlying consolidation seeming to work on two different timescales: fast structural processes working at the cellular level, followed by slower processes working at the systems level. For this report they will be termed 'synaptic consolidation' and 'systems consolidation' respectively.

1.3.1 Synaptic consolidation

Synaptic consolidation refers to the molecular changes that occur at the synapse and that lead to the structural modifications underpinning long-term memory. This is highly conserved throughout species and memory types, having been shown in Aplysia (Castellucci et al., 1970; Lin and Glanzman, 1994), rodents (Lipp and Wolfer, 1998) and humans (Dudai and Morris, 2000; Goelet et al., 1986). The memory trace is thought to occur through changes in synaptic weight or synaptic composition (Redondo and Morris, 2011), with long-lasting traces representing memories that persist.

Protein synthesis is necessary for synaptic consolidation as the products of gene expression are required for synaptic remodelling and growth (Davis and Squire, 1984; Hernandez and Abel, 2008). Within the hippocampus, glutamatergic signalling initiates this process, leading to *de novo* protein synthesis and alterations in synaptic morphology. The activation of α-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid receptors (AMPA-Rs) on the post-synaptic density (PSD) causes an influx of sodium ions (Na+), which leads to subsequent membrane depolarisation and removal of the N-methyl-D-aspartate-receptor (NMDA-R) magnesium block. Given the binding of glutamate and D-Serine to the NMDA-R, this allows influx of calcium ions (Ca^{2+}), which further depolarises the membrane, and, more importantly, activates numerous signalling pathways. These comprise kinases and GTPases such as protein kinase A (PKA), mitogen-activated protein (MAP) kinase, and p190 Rho GTPase-activating protein (p190 RhoGAP) (Lamprecht and LeDoux, 2004; Lattal and Abel, 2001). These then initiate cytoskeletal alterations and modulate receptor expression and trafficking, eventually leading to spinogenesis and synaptogenesis (Lamprecht and LeDoux,

2004). Activity-dependent activation of protein kinase C (PKC) and Ca^{2+} /calmodulin-dependent protein kinase II (CaMKII) lead to the phosphorylation of specific GluA1 sequences, driving AMPA-R insertion from extrasynaptic pools. This is followed by small G-protein-mediated trafficking of AMPA-Rs and the actin polymerisation required to support these increases in AMPA-Rs (Derkach et al., 2007). Activation of PKC and PKA also phosphorylate the GluN1 subunit, promoting endoplasmic reticulum export and trafficking of NMDA-Rs to the plasma membrane (Sans et al., 2003). Modifications of both can lead to increased synaptic potentiation, which corresponds to the strength of long-term memory expression (Migues et al., 2010).

These outlined molecular changes appear to occur during a specific time window initiated by memory acquisition and lasting on a time-scale of minutes to hours. During this time memories are thought to exist in a short-term form, which is labile and sensitive to anything that can inhibit these synaptic processes, such as protein synthesis inhibitors. Only after synaptic consolidation processes are complete does the memory become long-term, existing in a stable state that is resistant to blockade. This is true for synaptic changes induced by learning as well as experimentally induced synaptic changes termed long-term potentiation (LTP). The end of synaptic consolidation has therefore been defined operationally as the time point after learning (or LTP induction) when protein synthesis inhibitors can no longer affect the formation of memory (or LTP). Early-LTP consolidates to late-LTP and short-term memory consolidates to long-term memory.

Although the relationship between hippocampal LTP and synaptic memory consolidation is still under scrutiny, much evidence points towards LTP being a candidate for synaptic consolidation processes, and much of what we know about synaptic consolidation comes from studies investigating LTP mechanisms (Martin et al., 2000). When infused into the hippocampus of rats after training on an inhibitory avoidance task, PKC inhibitors disrupted long-term memory consolidation on a similar time-scale to that shown when using PKC inhibitors to disrupt late-LTP both *in vitro* and *in vivo* (Colley et al., 1990; Huang et al., 1992; Jerusalinsky et al., 1994). Similar results have been shown between late-LTP and long-term memory for both PKA and MAP kinase activity during contextual fear conditioning (Schafe et al., 1999), and for cAMP response element binding protein (CREB) during a Morris Water Maze spatial memory task (Guzowski and McGaugh, 1997). Similar studies have found

increases in such proteins in the hours following spatial learning tasks (Porte et al., 2008). It should be noted that there are studies that have successfully eliminated late-LTP *in vivo* whilst maintaining spatial long-term memory (Meiri et al., 1998), suggesting that mechanisms underlying late-LTP might not fully explain the synaptic consolidation of spatial memory. However given the levels of similarities found between the two processes, LTP still remains a likely candidate for synaptic consolidation of the hippocampal memory trace.

1.3.2 Synaptic consolidation and diversion retroactive interference

The specific time scale of synaptic consolidation (minutes to hours after initial learning) is in line with the time-dependent nature of the benefit of wakeful rest (minutes after initial learning) shown by Dewar et al. (2009). Retroactive interference could therefore decrease memory retention by disrupting the synaptic consolidation processes outlined above. It is possible that new learning could interrupt the synaptic consolidation of previously learned material by competing for plasticity-related proteins such as protein kinase or activity-regulated cytoskeletal-associated protein (Arc). Studies investigating the effects of memory of two similar behaviour tasks have shown that memory traces can compete for Arc if the levels of Arc are not abundant (Martínez et al., 2012). Rats were given weak context-dependent inhibitory-avoidance (IA) task training, producing a memory that under normal circumstances would not be fully consolidated into long-term memory. After this training rats explored a novel open field for 5 minutes, a length of time that would normally lead to long-term memory of the new environment. If the novel open field task directly followed the weak IA task, the weak IA task showed enhanced memory in the long-term and the open field task showed decreased memory in the long-term. If the order of the tasks was reversed and the weak IA task directly followed the novel open field task the same effects were seen. This was dependent on the levels of Arc present, and the effects were only present when a weak IA protocol was used. If a stong IA protocol was used that would normally be consolidated into long-term memory the memory for the novel open field was not affected. This indicated that memory traces were competing for proteins required for synaptic plasticity, and suggests that similarity-retroactive interference caused by the learning of new information is dependent on synaptic plasticity mechanisms. Therefore it is possible that these mechanisms could also underlie diversion-RI if unrelated sensory information is unconsciously learned.

1.3.3 Systems Consolidation

Systems consolidation is widely defined as the reorganisation of hippocampal memories over time and is thought to occur on the order of days to years, starting directly after memory acquisition during periods of quiet wakefulness or sleep (Ben-Yakov et al., 2014; Dudai et al., 2015). It is believed that during this process (or processes) the memory trace is redistributed throughout the neocortex (Osada et al., 2008). Therefore, as time goes on memories become consolidated on a systems level and depend less and less on the hippocampus (Dudai and Morris, 2000). What systems consolidation achieves, and exactly how these processes occur, has been debated for many years. Each key hypothesis will be briefly outlined in the following section.

1.3.4 The standard model of systems consolidation

The standard model of consolidation, partly conceived in the 1970's by Marr and expanded upon in the 1980's by Squire et al., outlined the basic concepts of a prolonged process of memory consolidation (Davis and Squire, 1984; Marr, 1971). This model emerged from the pioneering studies on amnesic patients with MTL damage. Interestingly although these amnesic patients could not form any new declarative memories, some declarative memories could be recalled. For example, H.M. could not consciously form any new memories, however he could remember memories occurring three years or more before his surgery. The older the memory (and further away from the surgery or lesion), the more likely it was to be unaffected and remembered (Scoville and Milner, 1957). It therefore appeared that this lack of retrograde amnesia was temporally graded. Consequently, the main principle of the standard consolidation theory was that memories initially stored in the hippocampus are slowly and gradually transferred to other brain regions, until a point comes when the hippocampus is no longer required for either storage or retrieval of these memories (Marr, 1971; Squire et al., 2014). Interestingly, some thought that part of this systems consolidation process could take many years, with the hippocampus still required as an index of sorts for long-term storage (Teyler and DiScenna, 1985). This index was thought to bind extra-hippocampal sites, allowing the retrieval of highly dispersed memories until the long process of systems consolidation is complete.

1.3.5 The multiple trace theory of systems consolidation

The extent to which old memories depend on the hippocampus has become a highly contentious topic, both in human and animal studies. The standard model of consolidation would predict that over years all memories would become independent of the hippocampus. However the temporally graded nature of retrograde amnesia was found to be incredibly variable, both in amnesic patients and in experimentally induced amnesia in animal models (Nadel and Moscovitch, 1997). In some instances it appeared that all declarative memories were lost, without any gradient at all (Bolhuis et al., 1994; Cermak and O'Connor, 1983; Damasio et al., 1985). It is far-fetched to argue that these decades-old memories were not yet fully consolidated. A second model was therefore proposed – the multiple trace theory (Nadel and Moscovitch, 1997). This theory shares some principles with the standard model of consolidation, such as the hippocampus rapidly encoding all consciously attended information into a highly and sparsely distributed ensemble of hippocampal neurons. However, instead of gradually transferring these ensembles to extra-hippocampal sites over time, every time part of a memory is re-experienced a new trace is formed overlapping the former. Older memories are therefore associated with a greater number of traces distributed more thoroughly throughout the hippocampus meaning the temporal gradient of episodic memories seen in some amnesic patients could be explained by extent of hippocampal damage.

Another problem that was left unexplained by the standard model was the relative sparing of semantic information over episodic memories seen in amnesic patients (Kopelman and Kapur, 2001; Nadel and Moscovitch, 1997). Episodic memory is rooted in contextual information, with the what, where and when acting as a "spatial scaffold" for different elements of the memory. The context of memory acquisition is what makes episodic memory so vivid and rich in detail. It appeared to be this type of memory that was so selectively impaired in patients with hippocampal damage, with most patients and animal models showing no ability to recall episodic memories no matter how old. Semantic memory, on the other hand, is gist-like and not attached to a place or time. This can include common knowledge gained throughout life and arbitrary facts learned. This memory, which lacks detail and vividness, tended to be spared, especially if the damage was unilateral or confined to small portions of the hippocampus (Kopelman and Kapur, 2001). Whilst the existence of truly episodic memory in animals in contentious (Suddendorf and Corballis, 2007) and

difficult to test (although much weight is now being put upon episodic memory tests in non-humans (Fellini and Morellini, 2013; Panoz-Brown et al., 2018; Veyrac et al., 2015; Zhou and Crystal, 2011)), spatial memory is thought to be episodic-like and easy to test through contextually-dependent behavioural experiments. As with human studies, animal studies also showed this distinction in retrograde amnesia. Purely spatial long-term memory showed no temporal gradient in a number of studies (Bolhuis et al., 1994; Clark et al., 2005; Liang et al., 1994; Mumby et al., 1999; Squire, 1992; Sutherland et al., 2001; Winocur, 1990) showing that the hippocampus is essential for retrieval of spatial memories regardless of memory age. Experiments testing non-spatial memories did show a temporal gradient (Kim et al., 1995; Ramos, 1998; Takehara et al., 2002). If all memories become independent of the hippocampus over time, as proposed by the standard model of consolidation, these results would not be expected. The multiple trace theory accounted for this discrepancy through the explanation that factual (semantic) information is present in many different contexts and memories, and therefore many different traces, allowing for its easy extraction from these episodes. When taken away from any contextual framework, these memories would still exist outside of the hippocampus due to the many overlapping non-contextual parts of the memory trace. In contrast memories that relied solely on a contextual framework would be lost.

Whilst this theory offered explanations and reasoning behind variations in the temporal gradient and type of memory affected by retrograde amnesia seen in both amnesic patients and animal studies, other inconsistencies with the multiple trace theory became apparent. The temporal gradient of retroactive interference was not dictated by the extent of hippocampal damage as predicted (Kopelman et al., 1989). It has been suggested that this could be explained by problems in the types of memory test used (Nadel et al., 2007). H.M. and similar patients had presented with decades old episodic memories that had originally appeared untouched and recallable, whereas more recent episodic memories appeared to be lost, suggesting a temporal gradient of episodic memories. However more tightly controlled studies indicated that this lack of contextual detail seen in recent memories was also lacking in remote memories (Corkin, 2002). This suggested that as long as memories remain detailed and rooted in a specific context, they depend on the hippocampus. Therefore with improper testing, truly episodic memories spanning back years might appear intact. However when episodic and semantic information is properly separated only the latter

should remain (Winocur and Moscovitch, 2011). This is supported by fMRI imaging studies in humans where participants were asked to recall episodic memories that were two weeks old (recent) and episodic memories that were ten years old (remote). Results showed that both of these autobiographical memories containing vivid and rich detail engaged the hippocampus, regardless of age (Bonnici et al., 2012). However, this same study showed that, unlike recent memories, remote memories were well represented in the ventro medial prefrontal cortex (vmPFC) as well as the hippocampus. A subsequent experiment showed that when testing and comparing the exact same two week old memory as recalled in the original study, the representation of this memory in the vmPFC became much more apparent two years later (Bonnici and Maguire, 2018). This suggests that episodic memory traces do become more dependent on the cortex over time, even if the hippocampus is required for the retrieval of such memories.

1.3.6 The schema theory of systems consolidation

An update of the multiple-trace theory suggests that, rather than the formation of multiple traces upon the re-experiencing of memories, every time a trace is fully or partly reactivated in the presence of some novelty, it is updated. This update incorporates the new information, strengthening the overall memory trace. This idea that memory traces can transiently return to a labile state even when fully consolidated has been termed reconsolidation (Nadel et al, 2007). Reconsolidation is thought to allow prior knowledge, i.e. consolidated memory traces, to influence the rate of further systems consolidation, as new information can be assimilated into a framework of pre-existing memory traces without the need for weeks of systems consolidation mechanisms. Reconsolidation and the subsequent assimilation of knowledge into these schemas has been shown in a hippocampal-dependent paired-association task (Tse et al., 2007). Rats had to memorise six flavour-place associations within a familiar arena so that when they were cued in the start box with a particular flavour they would know which location (sand well) within the arena was rewarded. After six weeks of training involving thirteen of these sessions, each with six different flavour-place associations rewarded three times, sham lesion animals appeared to learn where each flavour was located within the arena, making less incorrect choices before going to the correct sand well. Animals with hippocampal lesions did not make these associations, and the number of incorrect choices made did not improve. It was thought that if animals with an intact hippocampus could form a schema for the task-rules and event arena over these multiple sessions, then when two

consolidated flavour-place associations were changed to two novel flavour-place associations these new associations would be consolidated at a much quicker rate due to their rapid assimilation into the pre-existing schema. This was found to be the case, with the sham lesion animals recalling these new associations 24h later even after only one training session. Interestingly it appeared that 48h after the one training session this rapidly acquired memory was no longer hippocampal-dependent, as hippocampal lesions at this time-point did not affect memory recall. This appeared to disagree with the multiple trace theory, rather agreeing with the standard model of consolidation in part, as memories became less dependent on the hippocampus over time. However schema-based consolidation occurs at a much more rapid rate than ever predicted by the standard model. Further studies into the schema hypothesis of systems consolidation showed upregulation of immediate early genes in the medial prefrontal cortex directly after learning of the two novel flavour-place associations, indicating parallel encoding of memory traces in the hippocampus and cortex (Tse et al., 2011). This is in line with the theory that these new memories are rapidly assimilated into pre-existing cortical schemas and suggests that synaptic and systems consolidation mechanisms can exist on the same time-scales.

Whilst the schema hypothesis of systems consolidation appears to show that spatial memories become hippocampal-independent once consolidation is completed this is in direct conflict with studies showing that the hippocampus is required for episodic memory retrieval (Bonnici and Maguire, 2018). Bonnici and Maguire have offered an explanation behind this discrepancy. They propose that over time the hippocampus does not store memories, mirroring the standard model and schema hypothesis. It is, however, required for the retrieval of spatially anchored (hippocampal-dependent) memories, serving to re-contextualise the trace in a coherent episode. Without the hippocampus, only the semantic information within the trace can be retrieved. This requirement of the hippocampus during the recall of remote spatial memories has also been shown in animal studies where the recall of such memories resulted in activation of the hippocampus, irrespective of the age of the memory (Broadbent et al., 2006; Schlesiger et al., 2013). It therefore could be suggested that these whilst animals can recall information consolidated into schemas without a hippocampus, the information available to the animal is more semantic in nature, with schemas being more generalised and gist-like (Lewis and Durrant, 2011).

1.3.7 Sharp wave ripples and systems consolidation

Regardless of whether episodic memories ever become truly independent of the hippocampus, a defining feature of systems consolidation appears to be the coordinated communication between the hippocampal and cortical areas. For many years systems consolidation has largely been associated with sleep states, with studies showing that rapid eye movement (REM) sleep facilitates the consolidation of procedural memories whereas slow wave sleep (SWS) facilitates the consolidation of declarative memories (Fowler et al., 1973; Plihal and Born, 1997, 1999). This SWS-mediated facilitation of declarative memories appears to be dependent on low levels of acetylcholine (ACh) (Gais and Born, 2004) and the subsequent reactivation of the hippocampal memory trace (Wilson and McNaughton, 1994). Although first thought to occur solely during sleep, systems consolidation is now thought to occur during all post-learning 'offline' states characterised with low levels of ACh, such as quiet wakefulness (Mednick et al., 2011). It is in these periods of quiescence that interactions between hippocampal-cortical structures appear most coordinated (Siapas and Wilson, 1998).

Network events in the hippocampus are characterised by sharp wave ripples. These transient events comprise of negative potentials, termed sharp waves, in the CA1 stratum radiatum (comprised of the Schaeffer collaterals and commissural fibres of pyramidal cells) overlaid with fast frequency oscillations of around 200Hz, termed ripples, in the CA1 stratum pyramidale (comprised of the pyramidal cell bodies) (Buzsáki et al., 1983; Ylinen et al., 1995). Sharp waves are thought to be generated within the CA3 (Buzsáki, 2015), quickly bringing about fast ripple oscillations within the CA1, with sharp wave ripple events simultaneously occurring throughout the hippocampus and MTL structures (Chrobak and Buzsáki, 1996). The synchronous nature and high firing frequency of ripples is optimal for inducing synaptic plasticity mechanisms in downstream targets, such as cortical neurons (Logothetis et al., 2012), suggesting that hippocampal traces could be intrinsically reactivated and transferred to cortical sites through these network oscillations (Buzsáki, 1996; Ylinen et al., 1995). It has indeed been shown that the incidence of hippocampal sharp wave ripples during sleep increases after the animal has learned a hippocampal-dependent place-association reward task, and that this increase is proportional with improvements in behavioural performance (Ramadan et al., 2009). This post-learning increase in ripples has also been shown in humans (Axmacher et al., 2008). Subsequent studies in rats have shown that the selective disruption

of hippocampal sharp wave ripples during post-learning rest can impair spatial learning of a radial arm maze task (Ego-Stengel and Wilson, 2010; Girardeau et al., 2009). These results suggest that sharp wave ripples during periods of quiescence are necessary for the systems consolidation of hippocampal-dependent memory.

1.3.8 Cortical oscillations and systems consolidation

Cortical network events are characterised by slow frequency oscillations (1-4Hz) and spindle oscillations (7-14Hz) (Siapas and Wilson, 1998). Slow cortical oscillations are thought to temporally couple cortical spindles and hippocampal sharp wave ripples (Clemens et al., 2007), suppressing both during the hyperpolarised down-state of the slow oscillation (Mölle et al., 2006) and producing a temporal window of excitation allowing the increased activity of spindles and ripples during the depolarised up-state of the slow wave (Battaglia et al., 2004). This coordination is suggested to facilitate the reactivation of memory traces in the hippocampus and cortex simultaneously, allowing the transfer of information between the two brain areas (Sirota et al., 2003). Spindle activity has been implicated in the consolidation and integration of novel words into existing knowledge in humans (Clemens et al., 2007), suggesting that cortical spindles are required for schema-like systems consolidation. Boosting slow cortical oscillations via transcranial direct current stimulation or with in-phase auditory stimulation has also been shown to enhance cortical spindle activity, increasing coordination between slow waves and spindles and leading to enhanced declarative memory consolidation in both healthy participants and patients with mild cognitive impairment (Ladenbauer et al., 2016, 2017; Marshall et al., 2006; Ngo et al., 2013). This suggests that the both temporal coordination of hippocampal network oscillations and the temporal coordination of cortical network oscillations are important for the systems consolidation of hippocampal-dependent memories. An elegant study by Maingret et al. showed the significance of the coupling between these two areas (Maingret et al., 2016). Rats were exposed to two identical objects in adjacent corners of an open field for either 3 minutes or 20 minutes. 24 hours later, rats were exposed to the same open field, however one of the objects had moved to a new location. Rats that had explored the original object locations for 20 minutes showed a preference for the novel location and therefore expressed long-term spatial memory; whereas the rats that had only explored the objects for 3 minutes showed no preference, indicating they had no memory of the original object locations. It was found that, during the hour after exploration of the original object locations, the joint occurrence

of hippocampal and cortical oscillations was selectively increased in rats that expressed consolidated memory. A closed-loop stimulation protocol was then used in the hour following learning to increase this temporal coupling seen between sharp wave ripples in the hippocampus and slow oscillations and spindles in the cortex. A control condition used the same stimulation, but the slow wave stimulation was delayed so that the pairing of slow waves and sharp wave ripples far exceeded the timing that would be seen endogenously. Selectively increased hippocampal-cortical coupling enhanced memory consolidation so that 3 minutes of encoding now led to the expression of spatial memory 24h later. This was not seen in the delayed stimulation condition. These results highlight the importance of the fine-tuned coordination of hippocampal sharp wave ripples and cortical spindles and slow waves in the systems consolidation of hippocampal-dependent memory. It should be noted that that this fine-tuned coordination appears to be important in the hour directly following learning, indicating that long-term memory retention can benefit from a relatively short period of systems consolidation processes.

1.3.9 Systems consolidation and diversion retroactive interference

It is possible that retroactive interference leads to a decrease in memory retention through disrupting systems consolidation processes. New learning or incoming sensory information during this consolidation period could interfere with the communication between the hippocampus and the cortex. This incoming information could affect the joint occurrence of hippocampal and cortical oscillations, dampening the coordination between the two structures. Therefore an unfilled delay without incoming information would comparatively enhance this hippocampal-cortical communication, enhancing systems consolidation and leading to increased memory retention for the learned material.

1.4 Neural Correlates of Hippocampal-Dependent Memory

Memories take time to be stabilised into long-term representations in the brain. Patients such as H.M. could only remember items if they were constantly repeated in their short-term memory. Memories were forgotten as soon as the repetition of the items ceased. It has also been shown that memories are most susceptible to retroactive interference in the minutes directly after learning. The effects of retroactive interference in humans are diminished 10-minutes post-learning. The mechanisms underlying both synaptic and systems consolidation are also time-dependent. There are at least two states that a memory trace can occupy: one that is sensitive to amnesiac agents, such as protein synthesis inhibitors or retroactive interference, and decays quickly over a period of hours; and one that, once synaptic and systems consolidation mechanisms are complete, becomes resistant to such blockers and lasts for months if not years (unless reconsolidation occurs through the reactivation of the memory trace). Another way to investigate these memory mechanisms, and the mechanisms that underlie retroactive interference processes, is to identify the neural correlates of memory. The most plausible and striking of these neural correlates is outlined in the next section.

1.4.1 Place Cells and the 'Spatial Map'

Edward Tolman suggested the idea that animals possess an internal representation of the environment – a cognitive spatial map – in 1948. His experiment showed that rats could navigate a maze using short cuts, i.e. routes that the rat had never taken before, to reach different goal locations. This was also true when part of a travelled route was blocked, showing plasticity in spatial navigation (Tolman, 1948). As the rat had previously been allowed to spontaneously explore the maze, it was hypothesised that the rat was building an internal map of the environment during this non-goal orientated exploration time. This map could be subsequently used to efficiently navigate between different goal locations.

This theory was not seriously entertained until the discovery of a possible cellular correlate of this spatial map over twenty years later. As predicted by studies of amnesic patients and lesion studies in animals, these cells were located in the hippocampus, in both the CA1 and CA3 regions (O'Keefe, 1979). These so-called place cells were first recorded by O'Keefe and Dostrovsky in 1971. A place cell is a pyramidal neuron that dramatically increases in firing rate when the head of the rat is in a specific location within the environment, termed the

place field, with individual cells firing in different locations of the same environment (O'Keefe and Dostrovsky, 1971). This increase in firing is so large that in most areas of the environment the cell does not fire at all, or at least fires less than 0.1Hz, but in the firing field of the place cell, firing rates can reach over 20Hz. This large signal-to-noise ratio is highlighted in Figure 1.4, where the firing rate map of an individual place cell is shown. On the left is the path of the rat through the environment over a 10-minute session, and on the right is a heat-map (the rate map) of the location-specific firing of this cell. Blue indicates the cell was not firing, red indicates the maximum firing rate recorded, in this case 15 Hz. It is clear that this particular place cell only fires in one specific location within the open field. It should be noted that place cells can express multiple firing fields at once, although this is usually restricted by the shape and size of the recording environment, with less than 10% of cells expressing more than one field when recorded in a cylinder less than 1m in diameter (Muller and Kubie, 1987).

This discovery of place cells, coupled with the knowledge that the hippocampus was required for spatial learning and memory, led to the cognitive map theory (O'Keefe and Nadel, 1978). As firing fields of different place cells appeared to be fairly evenly distributed throughout the explored environment (Muller and Kubie, 1987) it was suggested that the firing fields of these cells could cover the entire environment giving the animal information about where it was in space at any given time, much like a coordinate position on a map.
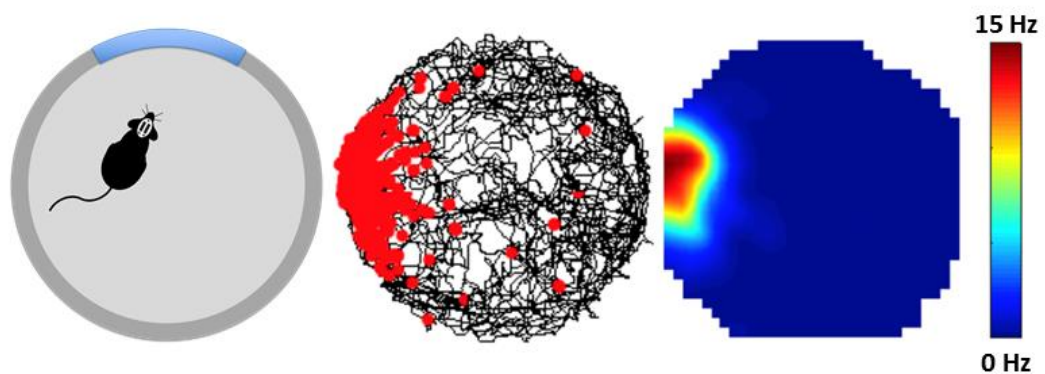
Figure 1.4: Left: Illustration of place cell recording in an open field with a cue card at north. Middle: Black lines represent the path of the rat through the environment over the 10-minute recording session. Red dots represent the locations within the environment where this specific cell fired. Right: The rate map of this cells firing in the environment. Red colours indicate peak firing, in this case 15Hz, and blue colours represent little to no firing. The hot colours therefore represent the firing field of this place cell.

Interestingly, unlike other areas of the brain, such as the sensory and motor cortices (Dräger, 1975; Patel et al., 2014), place cells were found to fire in a non-topographical manner. This means that adjacent place cells do not necessarily express firing fields that are next to each other in the environment, and there is no preservation of the relationship between pairs of place cells and the locations of their firing fields between different environments (Dombeck et al., 2010; Kubie and Muller, 1991). It is populations of place cells firing together that encode where the animal is in space. Any given cell may have more than one field in the environment, especially if the environment is large, and each cell fires in many environments. Therefore the population activity of place cells must be decoded to know which environment and where within this environment the animal is. Whilst more recent studies propose that the idea of the hippocampus acting as a spatial map is too reductionist (Eichenbaum et al., 1999), as place cells encode many different aspects of an environment or episode other than location (Gothard et al., 1996; Kraus et al., 2013; Wood et al., 1999), there is no argument that the hippocampus is required for spatial memory and that place cell firing can represent the animal's location within an environment (Brown et al., 1998).

1.4.2 Place Field Remapping

Although other brain areas also have cells that appear to fire in a location specific manner, such as the entorhinal cortex and the subiculum (Sharp, 1997, 1997), the defining feature of hippocampal place fields is that a unique population of place cells fires in every different context the animal encounters, producing a 'spatial map' that is specific to that particular environment (Muller and Kubie, 1987). Individual place cells therefore express place fields in different locations in different environments and can be completely silent in one environment whilst having a strong field in another (Kubie and Ranck, 1982) due to at least half of place cells appearing to be inactive in an environment at any given time (Thompson and Best, 1989). These changes in firing between environments are termed remapping.

It is unknown exactly what features of the environment determine whether place fields remap. O'Keefe and Conway showed that small changes to the environment, such as the removal of one cue out of many, did not lead to remapping, with large changes to the environment being required to elicit globally different patterns of place cell activity (O'Keefe and Conway, 1978). Muller and Kubie showed that when the shape of the environment was changed from circular to rectangular the population of place cells showed global remapping,

with unique but overlapping populations of neurons firing in each environment (Muller and Kubie, 1987). This global remapping was also seen when animals explored cylinders of different colours (Kentros et al., 1998), or environments with changing colours and odours (Anderson and Jeffery, 2003). On the other hand, another study showed that if only the colour of a prominent cue card changed then half of the animals underwent this remapping of their place cells but half did not (Bostock et al., 1991). This suggested that half of the animals deemed the environment to be the same and half considered themselves to be in a different environment altogether. It has also been shown that the distinction of place cell representations between two geometrically different contexts can require repeated exposures, remapping occurring gradually over time (Lever et al., 2002). The extent to which place cells remap therefore appears to be dependent on the animal's own perception of how different two environments are from each other. It has recently been suggested that variability in the extent of remapping between contexts could also be based on the levels of attention the animal is paying to their surrounding environment, with minimal changes in environmental cues leading to remapping if the animal was not sufficiently attentive to its environment during its first exploration (Kentros et al., 2004; Monaco et al., 2014). Examples of global remapping are shown in Figure 1.5.
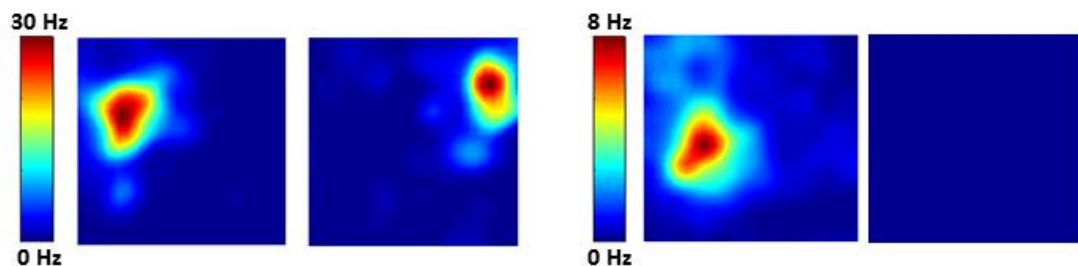
Figure 1.5: Two examples of global remapping of place fields. Each square represents a different box surrounded by very different contextual cues (such as different colours of curtains surrounding the box and different cues attached to these curtains). Left: The cell expresses a place field in a different location. Right: The place cell no longer expresses a field in this environment

As well as global remapping, where place cells change their firing locations, another type of remapping has been observed termed rate remapping, shown in Figure 1.6. Rate remapping occurs when place cells express fields in the same location, but these fields have significantly different firing rates between environments (Leutgeb et al., 2005b). In one study place cells globally remapped if the animal explored the same box within geographically different rooms, and rate remapped if animals explored different boxes located at the same place

within the same room (Leutgeb et al., 2005b). Another study used different contexts, a familiar square and a familiar circular context, that gradually morphed from one to the other (Leutgeb et al., 2005a) to show high levels of rate remapping. The remapping occurred gradually, smoothly transitioning from rates associated with one context to rates associated with the other. This was in direct contrast to a study published in the same year, which showed that there was a distinct point where animals perceived one shape to be different from the other leading to global rather than rate remapping. Remapping did not occur in the gradually morphing contexts until this point was reached (Wills et al., 2005). It should be noted that the latter study recorded from CA1 place cells whilst the former study was recorded from both CA1 and CA3 place cells, and it appears that rate remapping occurs at higher rates in CA3 place fields (Leutgeb et al., 2005a). These studies suggest that rate remapping appears to encode smaller changes in the environment, informing the animal that something has changed but that they might not necessarily be in a different location geographically.
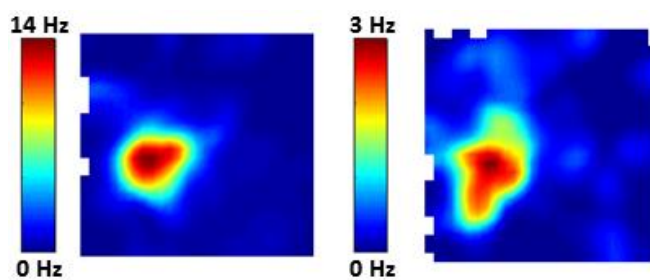
Figure 1.6: A place cell exhibiting rate remapping between two exposures to the same environment. The field is expressed in the same location, however it fires at a significantly lower rate.

What is important to remember is that when animals are put back into what they deem to be the same environment and allowed to explore, the same population of place cells fire in the same locations at similar firing rates, shown in Figure 1.7. This specific pattern of neuronal firing can only occur in this particular environment, indicating that the animal has 'memory' of the features that make up this particular context. These neuronal representations are known to last for days (Muller and Kubie, 1987), even months (Thompson and Best, 1990), implying that this long-term stability is the recollection of the original representation created when the animal first encountered the environment. This suggests that long-term place field stability could be construed as a neuronal correlate of long-term spatial memory for the environment. The following section therefore outlines the parallels between place cell stability and the mechanisms thought to underlie spatial memory consolidation.
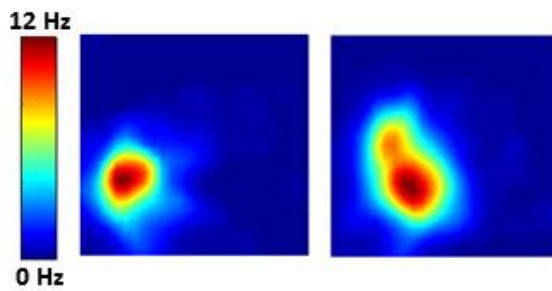
Figure 1.7: A place cell exhibiting both global (location) and rate place field stability. The field is expressed in the same location in the environment at the same rate.

1.4.3 Place Field Stability and Spatial Memory

The possibility that long-term stability of place fields requires synaptic consolidation processes has been experimentally tested in a number of studies. Late-NMDA-R-dependent-LTP has been shown to be essential for long-term spatial memory (Tsien et al., 1996). It has also been shown that, whilst place cells can still express stable fields for an hour when NMDA-Rs are inhibited, this stability was not apparent when rats were recorded in the same environment 6h later (Kentros et al., 1998). This NMDA-R-independent short-term stability has also been shown in a mouse model where CA1 NMDA-Rs were selectively knocked out. Although spatial specificity was significantly decreased in the knock-out mice compared to the control mice, place field stability was unaffected over a 2.5h period (McHugh et al., 1996a). These results suggest that NMDA-Rs are required for long-term place field stability but not stability in the short-term, implying that long-term place field stability is akin to late-LTP and long-term spatial memory.

These similarities between long-term place field stability and late-LTP mechanisms have also been found in studies using protein synthesis blockers. New protein synthesis is required for late-LTP (Frey et al., 1988; Huang et al., 1996) and for the long-term stability of spatial memories (Ozawa et al., 2017). When mice were injected with anisomycin, a protein synthesis inhibitor, directly after exploring a novel environment, place field stability remained intact for an hour but fields were not stable 6h or 24h later (Agnihotri et al., 2004). If mice were injected after exploration of a familiar environment the stability of place fields was not affected, indicating that protein synthesis was not required for the recollection of the original place cell representation of the environment if already consolidated. Transgenic R(AB) mice with reduced forebrain protein kinase A (PKA) levels also showed place field stability at 1h but not 24h, in contrast to the wild-type control mice that expressed both 1h and 24h place field stability (Rotenberg et al., 2000). Overall these studies suggest that the

long-term stability of place fields requires the same synaptic consolidation mechanisms that are needed for long-term spatial memory.

Just as there are parallels with long-term place field stability, spatial memory and synaptic consolidation, place cells appear to be intrinsically linked to systems consolidation mechanisms. Many theories of systems consolidation agree that over time memories become less dependent on the hippocampus with information gradually transferred to the cortex through repeated reactivation of these memory traces. This reactivation is thought to occur during sleep or quiet wake when levels of ACh are low, and to be intrinsically driven by coordinated network oscillations in both the hippocampus and cortex. It is thought to be hippocampal sharp wave ripples that 'replay' the hippocampal memory trace to the cortex, a phenomenon that has been found in many place cell studies.

Animals will pass through different place fields when running along the length of a linear track. The sequential firing of place cells is therefore governed by the specific order in which place fields are encountered by the animal. This is also true for routes taken through an open field, however studies into the sequential firing of place cells have preferentially used linear tracks as this sequential order is fixed and cannot change throughout the experiment. The sequential firing of place cells appears to be synchronised to intrinsic hippocampal theta rhythms present when the animal is moving (Wallenstein and Hasselmo, 1997). Place cells can fire in the trough or the peak of the theta cycle. Where in the theta cycle the cell fires depends on the animal's location with that cell's field, i.e. when a rat enters the place field the cell fires at a later phase of the theta cycle than when the same cell fires when the rat is leaving the place field (shown in Figure 1.8). Therefore there is spatial information encoded in the timing of spikes, and this information is dependent on theta rhythms. This is termed theta phase precession. Place fields can overlap, so the animal can be travelling through multiple place fields at once, all firing in specific parts of the theta cycle depending on where the animal is on its path through these firing fields. This leads to place fields encountered later in the rat's trajectory firing in the later theta phases, whilst place fields encountered earlier along the linear track fire in the earlier phases of theta (O'Keefe and Recce, 1993).
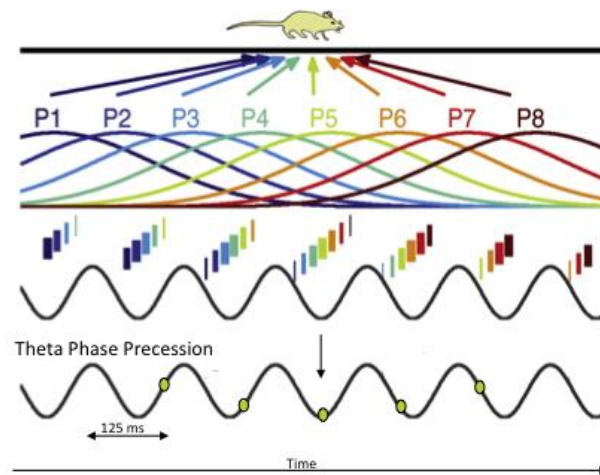
Figure 1.8: Theta phase precession on a linear track. As the rat moves through the different place fields (P1-8) the corresponding place cells fire. This place cell firing is depicted by rectangles that are colour coordinated with their corresponding place fields. The width of the rectangle depicts the firing intensity of the place cell. As the rat enters the place field the corresponding place cell fires late in the theta phase (theta oscillation is depicted in black). The place cell then fires earlier in the theta phase as the rat traverses through the cell's firing field. This is highlighted for P5 (bottom), showing where in the theta phase this cell fires (green oblong) depending on the location of the rat. Place fields overlap so many cells fire in every cycle of theta. The timing of place cell firing within the theta phase therefore encodes spatial information. Figure adapted from Buzsáki, 2010.

These sequences of place field firing are replayed during sharp wave ripple events in the hippocampus, representing a replication of the trajectory the animal took along the linear track or through an environment (Skaggs et al., 1996). Sequences are compressed to the extent that the entire linear track, or environment, is replayed in a single ripple event (Davidson et al., 2009; Diba and Buzsáki, 2007; Foster and Wilson, 2006). This place cell replay occurs both in the awake state during brief pauses in exploration of the environment (Foster and Wilson, 2006), and remotely in a sleep or quiet rest state following spatial learning. Awake place cell replay can occur in reverse, suggesting that it is a form of memory retrieval available to guide the animal on further navigational decisions (Jadhav et al., 2012; Takahashi, 2015), as well as in a forward direction ('preplay') (Diba and Buzsáki, 2007) indicating it can also be used for the planning of future trajectories (Silva et al., 2015). Replay of place cells during sleep or quiescence on the other hand occurs almost exclusively in the same direction as experienced during exploration (Lee and Wilson, 2002; Skaggs et al., 1996). It is this form of place cell replay during sleep or quiescence, occurring only during hippocampal sharp waves in the presence of low levels of ACh, which is thought to underlie the systems consolidation of spatial memories. Just as sharp wave ripples have been shown to be important in the systems consolidation of spatial memory (Ego-Stengel and Wilson,

2010; Girardeau et al., 2009), the replay of place cell firing during these oscillatory events also appears to be imperative to these consolidation processes. It has been shown that the same synchronous patterns of place cell firing, termed cell assemblies, recorded during the learning phase of a reward-place association spatial memory task were reactivated in association with sharp wave ripples during the subsequent rest session. The strength of this cell assembly reactivation during sharp wave ripples predicted the rats' memory performance of the task (Dupret et al., 2010a), indicating that the 'replay' of place fields during sharp wave ripples led to the consolidation of this spatial memory.

Overall the mechanisms underlying the consolidation of spatial memory appear to be closely paralleled with the mechanisms underlying the long-term stability of place field expression.

1.4.4 The relation of place cells to episodic memory

Whilst much of the place cell literature focuses on the spatial nature of place cell firing, it has been shown by numerous studies that place cells also encode non-spatial information, (Igarashi et al., 2014; Sakurai, 1996; Segal et al., 1972; Wood et al., 2000; Young et al., 1994). In some cases a specific subset of place cells encoded non-spatial information, such as odours in a non-matching to sample task (Wood et al., 1999). In this experiment rats were tasked with digging in sand wells for rewards only if the odour of the sand was different to the odour encountered previously. The location of the sand well, order of the odour sequence, and match/non-match contingencies were in a pseudo-random order. A subset of cells fired at a consistent spatial location regardless of odour, others fired at a specific location only when specific odours were encountered there, but importantly another subset of cells fired in response to specific odours regardless of location. In other cases place cells encoded an association between a non-spatial feature of the environment or task and the place in which this occurred. This has been shown for odour-place associations (Komorowski et al., 2009a), where sand wells were rewarded only if the odour correctly matched the associated context, and the association between object identity and object location in a novel object location based task (Manns and Eichenbaum, 2009). Place cells have even been shown to encode the passing of time (Kraus et al., 2013; MacDonald et al., 2011; Manns et al., 2007; Pastalkova et al., 2008) suggesting a method for representing when specific events occurred during a task. Whilst this non-spatial encoding is partly in disagreement with the original 'spatial map' hypothesis, where it was proposed the hippocampal code was primarily spatial, it does lend

weight to the idea that place cells could be underlying any type of hippocampal memory rooted in a contextual framework i.e. episodic memories. Episodic memories are associated with the 'what', 'where' and 'when' of an event, all features that have been shown to be encoded by place cell populations. This, coupled with the fact that place cells have been recorded not only in in rodents, but also in bats (Yartsev and Ulanovsky, 2013), monkeys (Hori et al., 2005) and humans (Ekstrom et al., 2003; Miller et al., 2013), suggests that place cells could be important for more than just simple navigational purposes. Therefore findings from place cell recordings in rodents during spatial memory tasks have the potential to elucidate mechanisms underlying episodic memory consolidation in humans.

1.5 Introduction conclusion and aims of this thesis

This introduction has shown that the consolidation and recall of episodic and spatial memories are dependent on the hippocampus. Consolidation is a time-dependent process involving a period of susceptibility to new learning and incoming information. Studies investigating interference during this consolidation period, i.e. the delay between learning and recall, have shown that two types of retroactive interference (RI) exist, one that relies on the similarity of the information present during the delay (similarity-RI) and one that does not require the information to be similar in any way (diversion-RI). Although the original concept of diversion-RI was contentious, well-designed experiments with true control conditions have definitively proven its existence. From these studies it appears post-learning material does not need to be consciously learned, novel or similar to the memory being consolidated to interfere with such consolidation mechanisms. This body of evidence therefore suggests that diversion-RI could play a role in everyday forgetting. At present it is known that both a spot-the-difference task and a tone detection task lead to diversion-RI. This is hypothesised to be through the mental exertion required to perform these tasks. However it is unknown whether diversion-RI exists without such mental exertion. Therefore the exact type of stimuli that lead to diversion-RI requires further investigation. Although the type of stimuli that affect these consolidation mechanisms is presently unclear, reducing the majority of sensory stimulation through wakeful rest has been shown to consistently

enhance the retention of long-term hippocampal-dependent memories in humans. Wakeful rest in the studies with human subjects consists of being alone in a dark room, suggesting that reducing visual input or reducing social interaction could be underlying the enhanced memory retention shown in these control conditions. The first aim of this thesis is therefore to test whether reduced interference also enhances memory in rats. This was tested using a novel object-location spatial memory task, using varying types of non-task based interfering stimuli immediately after the encoding period. This included modulating the levels of social interaction and visual input individually. This set of experiments allowed exploration into whether diversion-RI requires mental exertion, and answers questions regarding the importance of visual input and social interaction in wakeful rest.

There are distinct similarities between the virtual reality route-learning task used by Craig et al. whilst investigating the effect of wakeful rest on humans, and the navigation of a maze using short-cuts shown in rats by Tolman (Craig et al., 2016; Tolman, 1948). Both required the learning of an internal representation of the environment and the subsequent extraction of spatial information from this consolidated memory. The possibility of such an internal representation, termed the 'spatial map' by Tolman, led to the discovery of a possible neural correlate of spatial memory within the hippocampus. The firing of place cells can signal where an animal is in its environment, and show "memory" for familiar environments through place field stability. This "memory" requires the same processes as synaptic consolidation, and shares many parallels with the mechanisms thought to be behind systems consolidation. It has therefore been suggested that processes underlying place cell "memory", such as place field stability, also underlie long-term spatial memory. If place cells are required for spatial memory tasks, such as the navigation task used by Tolman, then it could be hypothesised that the properties of these cells also underlie the memory enhancement of spatial memories seen after wakeful rest in humans. Although place cells have largely been investigated in rodents, a number of studies have now shown the existence of such cells in humans, indicating translational possibilities. The second aim of this thesis is therefore to investigate whether the stability of place cell activity is also affected by diversion-RI, and the benefit of wakeful rest on the consolidation of spatial memory of a novel environment, as measured by the stability of place fields.

The benefits of wakeful rest in humans have been shown not only in spatial navigation tasks but also in tasks requiring the consolidation of episodic memory, such as the recollection of passages of prose. Whilst episodic memory is more difficult to test in non-humans, it is possible to test place-associations requiring more than simple A to B navigation. Place cells are known to encode more than just location, and numerous studies have shown the encoding of place-associations through place field firing properties. Therefore the final aim of this thesis is to investigate whether place field properties associated with memory for object locations in the novel object location task differ between the wakeful rest and filled delay conditions.

The overarching aims of this thesis are therefore to further our understanding of how reducing retroactive interference through wakeful rest can enhance hippocampal-dependent memories, both behaviourally and mechanistically. The specific hypotheses and predicted outcomes for each experiment are included in each of the three experimental chapters.

Chapter 2: The Black Box Effect: Enhancement of Object Location Memory in Rats

2.1 Introduction

The phenomenon of consolidation (*Konsolidierung*), first described in 1900 by Müller and Pilzecker, showed that over time memories become stronger. These pioneering studies also underlined the importance of incoming information during this consolidation period. If the time following learning (consolidation period) was void of new learning, as is the case, for example, during sleep, memories were more likely to be remembered; if this period, however, was filled with new learning, as often is the case in real-life, memories were more likely to be forgotten. New memories became less sensitive to new learning the longer the time between original learning and the onset of new learning, probably due to the completion of the consolidation of these memories. The memory impairment caused by new learning was termed retroactive interference (RI). Many subsequent studies suggested that the impairment caused by RI was much greater if the new learning, or stimulus, was highly similar to the learned material (similarity-RI) (Skaggs, 1925). A number of these studies even concluded that the material had to be similar for RI to occur (McGeoch and McDonald, 1931; Dey, 1969). Although the idea of diversion-RI – interference caused by non-similar stimuli via mental effort – was suggested by Müller and Pilzecker in their seminal work, this possible source of interference remained a highly contentious issue for many years.

Diversion-RI was largely forgotten until the publication of two studies (Cowan et al., 2004; Della Sala et al., 2005) investigating RI in amnesic patients and healthy controls. These studies hypothesised that if the period after learning was filled with an unrelated distractor task diversion-RI would impair memory consolidation, as had been suggested previously. However, if the post-learning period was unfilled and contained no new learning, consolidation mechanisms were not impaired, and could even be enhanced. The latter had been shown numerous times in sleep studies (Stickgold, 2005; Tamminen et al., 2010; Wamsley et al., 2010), however it was unknown whether memory benefited from sleep because it reduced the amount of new incoming sensory information due to thalamic inhibition (i.e. reducing interference) or through other sleep-dependent mechanisms, such as sleep-dependent memory replay. Cowen and Della Sala also proposed that amnesic patients could be more susceptible to diversion-RI, which would in turn exacerbate any apparent memory problems. It was therefore hypothesised that reducing new learning not only would enhance the memory of healthy participants (through a post-learning unfilled condition), but would also benefit amnesiac patients to a much greater extent.

To test these hypotheses, both amnesic patients and healthy controls learned lists of words and stories immediately followed by either a filled condition (an unrelated distractor task) or an unfilled condition for 10 minutes. Recall of words and stories was greater after the unfilled condition in both groups, confirming that memory could be enhanced when post-learning interference levels were minimal. They also found that memory enhancement was greater in the amnesic patients than the control subjects, suggesting that RI caused by mundane events in daily life may be contributing to the memory deficits seen in these patients. These results were corroborated by a number of further studies, showing that the participants were not rehearsing the learned information or sleeping during the unfilled (wakeful rest) condition (Dewar et al., 2014; Craig et al., 2016; Dewar et al., 2012). This suggests that sleep is not a requirement for memory enhancement produced by wakeful rest.

Some of these studies showed little or no memory enhancement in the healthy participant group due to a ceiling effect (Dewar et al., 2009, 2010). Further studies therefore sought to increase the difficultly of the task. Landmarks and routes were learned during virtual navigation tasks to investigate how well these memories could be integrated into a spatial map of the virtual environment (Craig et al., 2016). To test how well participants could extract spatial information from their internal representation of the virtual environment, they were asked to point to learned landmarks from random locations within the environment. Participants in the wakeful rest condition performed significantly better than those in the distractor-task condition. This suggests that wakeful rest may enhance the formation of a cognitive map. These studies in healthy participants therefore highlighted the benefit of wakeful rest on higher cognitive processes, such as the integration of spatial memories. These brief periods of reduced incoming information have been shown to enhance episodic and spatial memories for up to 7 days (Dewar et al., 2012; Alber et al., 2014). Although this effect is clearly robust, not much is known about the importance of individual elements of wakeful rest, nor how these affect the underlying consolidation processes. This chapter therefore focuses on the role of retroactive interference on memory and consolidation mechanisms in rats.

Human studies show that wakeful rest promotes the integration of spatial memories into the cognitive map. To replicate these findings in rats, a spatial memory task was chosen that relies on spontaneous learning of object locations –the novel object location task (nOL). Both

the spatial navigation task in the human studies and this nOL task used non-guided exploration of the environment, and neither task was motivated by rewards. In addition, they both require the consolidation of long-term memory. In the nOL task there are four phases – habituation, sampling, delay, and probe (Figure 2.1). During the habituation phase, animals explore an empty open field for a number of days, reducing anxiety and forming memory of the context. During subsequent sampling, animals spontaneously explore two novel objects placed in the familiar open field, after which animals return to their home cage for a memory retention period. After this delay period, rats are put back into the open field with one of the objects in its original place, and the other moved to a novel location. Due to the innate preference of rats to explore novelty, if the animals remember the initial object locations, i.e., the places they occupied during sampling, they will explore the object at the novel location more than the one at the old location (Dix and Aggleton, 1999). This, therefore, produces a behavioural read-out of spatial memory.

To replicate the initial human studies, another phase was added directly after sampling – a post-sampling phase in which the level and type of interference was manipulated in each condition. The specific conditions used in the post sampling phase can be seen in Table 2.1. In the human studies, there were many differences between the wakeful rest condition and the control (RI) condition. During wakeful rest subjects were asked to rest quietly, with closed eyes, in a darkened room. Sleep and the ability to rehearse recall content were controlled for by a post-experiment questionnaire and non-recallable verbal content (Dewar et al., 2014). During the interference condition, subjects were asked to participate in non-recall related tasks such as spot the difference or tone detection. These required sustained attention, but not of information in the same modality as the recall task. Therefore, not only was one condition in darkness and the other in light, one required interaction with either a computer or experimenter and continued attention. It could therefore be predicted that the wakeful rest condition only works to enhance memory when compared to a condition with comparably high levels of interference. Unlike in the human studies, the post-sampling conditions in our experiments sought to increase interference in a passive way - through visual input or social interaction - rather than participation in tasks requiring constant attention. This allowed the comparison of conditions with varying levels of interference with a low interference condition that mimicked wakeful rest.
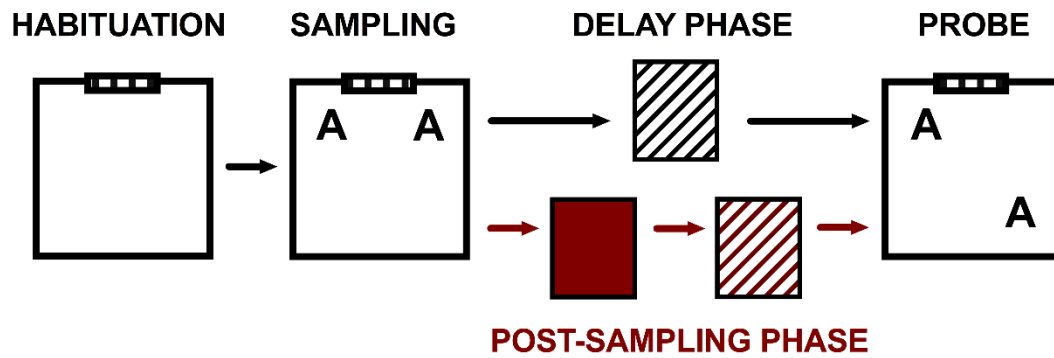
Figure 2.1. The standard novel object location protocol with a delay phase between sampling and probe trials (outlined in black). Animals usually spend the duration of this delay phase in their home cage. The post-sampling phase in the following experiments (outlined in red) aimed to change the amount of interference animals were exposed to during this delay.

If our prediction is true, and memory enhancement is evident only when a condition of low interference is compared to a condition of very high interference, we would expect to see differences between the two conditions only when the latter has high amounts of both social and visual interference. This would suggest a passive protective role from the high levels of RI. If enhancement is seen when the low interference condition is compared to a condition with just social or just visual interference, this could suggest active enhancement of memory, possibly through consolidation processes. These conditions also allowed specific differences to be isolated, addressing the importance of both darkness levels and social isolation independently.

The post-sampling phase was at least one hour long, unlike the 10 minute wakeful-rest in human studies. The rationale for using a 1 h post sampling phase in my experiments is as follows. Synaptic consolidation processes in rats and mice have been shown to last at least one hour after encoding, determined by a period of susceptibility to blockers such as protein synthesis inhibitors. As it was hypothesised that wakeful rest enhances memory through protecting or enhancing these consolidation processes, it was important to reduce interference for the duration of this labile period. Although this extended period of wakeful rest might not be necessary, we wanted to increase the probability of memory enhancement effects.

Although the delays used in the following experiments (6h and 24h) were not as long as those used in the human experiments (7d), both 6h and 24h nOL have previously been associated with long-term spatial memory (Ozawa et al., 2011). As the expression of memory was required after only one sampling exposure (multiple sampling sessions would lead to multiple consolidation periods) longer delays were not feasible. However, when results showed that 6h nOL memory was expressed after one sampling exposure the delay was increased to 24h to ensure that the effects seen were possible at longer delays.

The first experiment aimed to replicate the human effect of wakeful rest on memory retention in rats. As with the human studies it was hypothesised that rats would benefit from a period of reduced sensory stimulation after learning compared to a period of high interference. If this were correct, rats placed into a dark, quiet box with no social interaction after spatial learning (the black box condition) would express significant nOL memory, whereas rats placed directly into their home cage (home cage condition) would not.

The next experiment aimed to replicate the "black box effect" using a more refined protocol. Putting rats into a box with a red light rather than darkness would allow the experimenter to better see the animal, leading to more precise gentle handling procedures to ensure rats were kept awake. This protocol also tested a longer duration of nOL memory (24h). If this refined protocol was effective, you would expect animals in the red light box condition to express significant 24h nOL memory. This experiment also aimed to investigate the role of social stimuli on the black box effect. Previous studies always ensured participants were alone during restful wake, so it could be hypothesised that social isolation alone is sufficient to produce the memory enhancing black box effect. If reducing social stimuli does eliminate the effect of interference, you would expect animals separated individually into holding cages to express memory, even in the presence of other interfering stimuli such as visual input from a comparatively busy lab environment.

Many unrelated distractor tasks used in wakeful rest studies were visual tasks, such as spot-the-difference. It is reasonable to suggest that the visual input from the complex and changing visual scene of the lab environment could also act as a significant source of interference. The next experiment therefore aimed to explore whether a familiar and controlled visual environment was enough to prevent memory enhancement, even when

socially isolated. Animals were placed into boxes identical to those with the red light, with the only difference that the colour of the light was white. This acted to reduce all social stimuli and non-controlled visual input. It was hypothesised that whilst the complex visual scene of the lab could cause RI, a familiar and highly controlled visual environment would not because levels of new learning would be minimal. If this were correct animals in both conditions would express nOL memory. If in fact it was darkness that was required to produce memory enhancement the red box group would express memory, but the white light box group would not.

The previous experiments would indicate whether the reduction of visual input, both novel and familiar, was a requirement for memory enhancement. The last experiment therefore aimed to test whether social isolation was a requirement for the memory enhancement of the black box effect. Animals were placed into the red and white light boxes, but this time in their home cages with their cage mates. If reduction of novel visual interference from the lab was required but social isolation was not one would expect to see memory in both groups. If darkness was required to produce memory enhancement but social isolation was not one would expect to see memory only in the red box home cage condition.

The following experiments will therefore elucidate whether the memory enhancing effects of wakeful rest can be replicated in rats, and whether social isolation, lack of novel visual stimulation and lack of familiar visual stimulation or light are required for these effects to be seen.

2.2 Methods

2.2.1 Animals

Sixty-four male Lister Hooded rats were obtained from Charles River laboratories. During the experimental procedures these animals were aged between 3-8 months. All animals were housed in groups of three or four to prevent social isolation, and all cages had tubes and chewing blocks for enrichment. Animals were kept on a 12-hour light/dark cycle, with training and testing always performed in the light phase of the cycle (training always in the first half of the light cycle). To maintain high-levels of spontaneous exploration animals were kept at 90-95% free-feeding body weight, i.e., animals were kept on a weight maintenance feeding regime. Animals were given 25-30g standard lab chow each per day, and free access to water. All procedures complied with the UK Animals (Scientific Procedures) Act (1986) and the European Communities Council Directive of November 24, 1986 (86/609/EEC). All animal experiments were carried out in compliance with protocols approved by the University of Edinburgh Animal Welfare and Ethical Review Board (AWERB), and under a UK Home Office Project License.

2.2.2 Pre-experiment habituation

Before animals entered into experiments, they underwent a minimum of 5 days of handling procedures. This was to reduce anxiety levels in animals, and to habituate them to human contact and gentle handling procedures used for sleep deprivation. Handling involved gently picking up and holding the animals, allowing them to be carried without any need for restraint. During this time, in groups, animals were also put into large arenas containing many different novel objects, such as metal whisks, ceramic egg cups and wooden 3D shapes, which were only used in this phase of the experiment. This served to familiarize the animals to the presence of novel objects, enhancing exploration in the upcoming behavioural experiments. These objects were not used in the behavioural experiments.

2.2.3 Experimental Design

Each experiment used sixteen animals: eight in each group for experiments 1, 2 and 3; and sixteen in each condition for experiment 4. Experiments 1, 2 and 3 used a between-groups design whereas experiment 4 used a within-groups design. Experimental numbers and conditions are shown in Table 2.1.

	Condition 1	Condition 2
Experiment 1	Home Cage (n=8)	Black Box (n=8)
Experiment 2	Holding Cage (n=8)	Dark (red) Box (n=8)
Experiment 3	Light (white) Box (n=8)	Dark (red) Box (n=8)
Experiment 4	Home Cage in Light (white) Box (n=16)	Home Cage in Dark (red) Box (n=16)

Table 2.1: List of experimental conditions

2.2.4 Equipment

The open field used in all experiments measured 65 x 65 cm, with wooden white walls 60 cm high and a wooden white floor. All experiments used a striped black and white cue card (30 x 20 cm) on the North wall, and a variety of 2D and 3D cues either attached to the top of the walls or just outside the open field in clear view. Cues were arranged in an asymmetric fashion, with one wall always devoid of cues. The floor of open field was covered with the same type of bedding used in the home cages of the animals. Before each trial this bedding was disturbed to ensure no scent trails remained, and faeces removed. The luminance on the floor of the open field was measured using a light meter, and the lights were dimmed such that the open field was at 20±1 lumens. Objects were selected of similar height or width (approximately 10 x 10cm), but of varying textures, colours and shapes. Objects were made of non-porous and easily washable material such as ceramic, glass and metal. It was ensured that no objects had faces or pictures of animals that could have elicited an innate preference or anxiety response. Objects were fixed to clear glass bases (7 x 9 cm), which could be screwed into the floor of the open field for stability during exploration. Before each sampling and probe trial, objects were cleaned thoroughly with alcohol disinfectant wipes to remove any residual scents. Behaviour was recorded using an overhead camera, through Blackmagic video capture software (Blackmagic Media Express version 3.3.1.).

Figure 2.2. Top left: the home cages containing cage mates used in Experiment 1 and 4. Top middle: the black box used in Experiment 1. Top right: the holding cage used in Experiment 2. Bottom left: the dark (red) and light (white) boxes containing individual plastic cages. The dark (red) box was used in Experiment 2 and 3. The light (white) box was used in Experiment 3. Bottom right: the plastic cages were removed from these boxes and replaced with home cages for Experiment 4.

2.2.5 Behavioural testing

2.2.5.1 Habituation

Habituation consisted of two parts: 3 days of exposure just to the post-sampling condition that the animal would experience (1h per day); then 4 days of free exploration of the open field (for 5 or 10 min) followed directly with exposure to the post-sampling condition (1h or 3h). 7 days of exposure for the post-sampling condition ensured that any effects seen were not due to novelty. During these habituation trials, animals were able to sleep in the post-sampling condition, to avoid a sleep deprivation effect. For open field habituation sessions, animals were placed into the open field with their snout facing the corner, ensuring that during the four trials each corner was used once as a starting position for each animal.

2.2.5.2 Sampling

Sampling was carried out 24 h after the last habituation session, and consisted of placing the rat into the open field with two copies of the same novel object for either 5 or 20 minutes (see Table 2.2). Objects were placed in north (NW and NE) or south (SW and SE) positions for experiments 1 and 4, and diagonally opposite corners (NW and SE or SW and NE) for experiments 2 and 3. An example of this is shown in Figure 2.1. These objects were centred 23 cm from the corners of the open field (16.25 cm from each corner wall).

2.2.5.3 Post-sampling condition

After sampling animals were put into one of the post-sampling conditions outlined in Table 2.1 and shown in Figure 2.2. Animals were kept awake for the duration of these post-sampling conditions via gentle handling. This involved gently picking up the animal two inches above the ground, and then placing them down again when the animal started to fall asleep. The method has been outlined previously in sleep deprivation studies, and found to cause minimal amounts of stress (Colavito et al., 2013).

For the first experiment animals were either placed back into their home-cage with their cage-mates or into a black box. The home cage measured 60 x 44 x 30cm, and was half plastic, half metal bars. It was located on a trolley (see Figure 2.2 – top left), in a room adjacent to the testing room, so the test rat would have access to the rich visual scene within the lab environment. Cage mates remained in the home cage during the post sampling period, so the animals were also exposed to social stimuli. The black box measured (38 x 32 x 24 cm), and was entirely covered with a black plastic lid such that when inside the box the animals were in the dark (see Figure 2.2 – top middle). It did not contain any other rats. The black box condition aimed to replicate the restful wake condition used in human studies, where participants were by themselves in a dark room. The home cage condition aimed to be a high interference condition with both visual and social stimuli. These together would provide comparisons between minimal interference and high interference conditions, respectively.

The second experiment aimed to replicate the conditions of the black box, but this time refining the protocol allowing the experimenter to view whether the rat was asleep in the box. To this end, a box was used with a red light. Red light is visible to humans but not to rats (Szél and Röhlich, 1992) so whilst the experimenter could view the rat via a camera, the rat experienced darkness. Each red box (RB) compartment measured 56 x 53 x 51cm, and

contained fans for airflow; cameras connected to screens outside of the boxes; low levels of lighting (30±5 lumens); and plastic cages (42 x 27 x 40cm). Red plastic film covered the lighting panel in these boxes (see Figure 2.2 – bottom left). Another aim of this experiment was to investigate the role of social stimuli in the high interference condition of the first experiment. Would reducing social stimuli eliminate the effect of interference even when visual stimuli were still present? To this end, the second post-sampling condition was a holding cage (see Figure 2.2 – top right). The individual holding cage was similar to the home cage, but smaller in dimensions (54 x 40 x 22cm), and did not contain any other animals. It was located in a holding rack in the same room as the home cage, so as with the home cage the rat would have access to the visual scene within the lab environment. Both the RB condition and holding cage condition had minimal interference from social stimuli. These aimed to compare a minimal interference condition with a visual interference condition, testing whether it was sufficient to remove social interaction in order to get the memory enhancing 'black box effect'.

It appeared that removing social interaction was not sufficient to produce the 'black box effect' (as shown in Figure 2.4). The third experiment therefore aimed to reduce visual interference present in the holding cage condition whilst also maintaining the reduction in social interaction. This explored whether a familiar and controlled visual environment was enough to prevent memory enhancement, even when the animal was socially isolated. The RB condition was used again as the minimal interference condition. To provide a comparable condition where the only difference was visible light, a light (white) box condition was used. The white box (WB) was exactly the same dimensions as the RB, and had the same low levels of lighting. The WB and RB were in different compartments of the same cabinet (shown in Figure 2.2 – bottom left). The only difference between the compartments was that the red plastic film covering the lighting panel in the red boxes was not present in the white boxes. If the exposure to novel or complex scenes was preventing memory enhancement in the second experiment, you would expect to see memory in both the WB and RB conditions. If however a highly familiar stimulus was enough to interfere with memory enhancement you would expect to only see memory in the RB condition.

As social isolation was not enough to drive memory enhancement, it could be argued that it was not required at all. When comparing the home cage condition to the black box condition

in the first experiment, had the home cage had been in the dark would this condition have also produced memory enhancement? To this end, a final experiment sought to test whether it was sufficient to remove visual stimulation in order to get the memory enhancing 'black box effect' or if in fact social isolation was necessary for this effect. The plastic cages were removed from the white and red boxes, and replaced with the home cage of the animal (see Figure 2.2 – bottom right). Therefore one post-sampling condition was the home cage in the WB and one was the home cage in the RB. If lack of visual stimulation was sufficient then the RB home cage condition should show memory and the WB home cage condition should not.

2.2.5.4 Delay Period

After the post-sampling period of experiments 1-3 rats were returned to their home cages. As the home cages were used in the last experiment, prior to this experiment rats were habituated to 'lab-cages', identical to their home cages. Animals were held in these 'lab-cages' after the post-sampling phase. Both home cages and lab cages were located on a trolley in a room adjacent to the testing room. This room had higher light levels than the testing room and animals were with cage mates, i.e. not socially isolated. The first and last experiments were testing 6h memory, so the animals remained this room for the entire 5h post-sampling condition delay. The second and third experiments were testing memory after a 24 h retention interval, so after an hour in this room the cages were transferred back to the colony. This ensured the animals were kept to their 12h light-dark cycle. Animals were then transferred back to this room the following day, 1h before the probe trial commenced.

2.2.5.5 Probe Trial

After a 6 or 24-hour delay, which included the post-sampling condition exposure directly after sampling, rats were placed back into the open field for 3 minutes. In the probe trial one of the objects remained in the same location as in the sampling phase, and the other was moved to a novel location. Each animal was placed into the corner that would never have an object near it during both sampling and probe trials (e.g. sampling positions: SW SE, probe positions: SW NE, rat start position: NW). Object positions were fully counterbalanced within and between groups.

	Habituation to open field	Sampling	Condition Exposure	Delay
Experiment 1	4 x 5m	5m	1h	6h
Experiment 2	4 x 10m	20m	1h	24h
Experiment 3	4 x 10m	20m	3h	24h
Experiment 4	4 x 5m	5m	1h	6h

Table 2.2: Parameters used for each experiment

2.2.6 Analysis and statistics

Videos of the sampling and probe trial were manually scored using a bespoke software (zScore). This involved replaying the videos and depressing a key every time the animal explored a given object. Exploration involved the animal facing the object with its snout from either the ground or on top of the object. Exploration on top of the object, but not facing it (i.e. looking around the environment), was not included. This produced a file with timestamps of exploration for each object, which was then processed further with a custom software (zChop) to calculate the discrimination index at each 10-second accumulative bin. Sampling was also scored to ensure animals had no innate location preference, and to assess whether groups expressed comparable exploratory activity. For the sampling phase only the absolute exploration times were compared, i.e., no preference index was computed. All testing and scoring was completed blind to object novelty and group.

The discrimination index *d* is a ratio that indicates the animal's preference for either the familiar or novel object location, with the total exploration time factored into the following equation:

$$discrimination\ index = \frac{exploration\ time\ of\ novel\ object\ location\ -\ exploration\ time\ of\ familiar\ object\ location}{total\ object\ exploration\ time}$$

If the discrimination index is 0, the animal expressed no preference for either the familiar or novel object location. If the discrimination index is negative the animal expressed a preference for the familiar object location; if the discrimination index is positive the animal expressed a preference the novel object location. Therefore, only positive discrimination indices indicate expression of spatial novelty exploration. Over the course of each 3-minute

probe trial this index was calculated in accumulative 10-second time bins. From this a curve was plotted to show how the discrimination index ratio changed throughout the probe session. Rats were included if they explored object for greater than 20 seconds in total in both the sampling and probe sessions; no animals were excluded as all reached this criterion.

Data analysis was performed with Microsoft Excel and GraphPad Prism (version 7.0, Graphpad, USA). The discrimination index calculated across the whole 3 min probe session (d[0-180]) was used for further statistical analysis to compare between groups or conditions. All datasets were examined for normality using the Shapiro-Wilk normality test. If the data were normally distributed ($p < 0.05$), two-tailed one-sample t-tests were then used to compare the discrimination indices with the hypothetical mean of 0 (no preference). Depending on whether the experimental design was within or between groups, a two-tailed paired or unpaired t-test (respectively) was used to calculate differences in discrimination indices between the two groups or conditions. An F test was used to compare the variances between groups or conditions. Two way ANOVAs were used to test for effects of object position or condition on sampling exploration time. Overall differences in sampling and probe exploration time between conditions were also analysed. If the data was not normally distributed the appropriate non-parametric equivalents were carried out: a Wilcoxon signed-rank test and either a Wilcoxon matched-pairs signed-rank test or Mann-Whitney test.

2.3 Results

2.3.1 Experiment 1: Home cage vs. Black Box

	Home Cage	Black Box
"Interesting" Visual Stimulation	✓	✗
"Habituated" Visual Stimulation	✓	✗
Social Interaction	✓	✗
Sleep	?	✗

Table 2.3: The differences between the two conditions in experiment 1.

This experiment sought to test whether the long-term memory enhancing effects of reduced environmental interference previously shown in humans can also be found in rats. The two post-sampling conditions are comparing high levels of interference (home cage condition) with minimal levels of interference (black box condition), summarised in Table 2.3. The former was situated in a room adjacent to the testing room, giving the rat access to the rich visual scene within the lab environment. Although the rats were habituated to being within the lab environment, the experimenter was moving and various aspects of the environment could change making the visual stimulation 'interesting', i.e., to some extent unfamiliar.. The home cage also contained cage mates and therefore social interaction. Animals in the black box condition on the other hand were in complete darkness and social isolation, as the box was completely covered and in a separate dark room, and contained no cage-mates.

Panels A and B of Figure 2.3 show the discrimination index for all animals. Panel A gives an overview of the entire 3-minute probe trial in 10-second cumulative bins. This indicates that animals exposed to the black box condition for 1 h directly after learning explored the novel object location more than the old location for the duration of the probe trial. Animals exposed to the home cage had no preference for the novel location at any time point. This difference is highlighted in panel B, showing that across the whole 3 minute probe time point, animals in the black box condition had a significant preference for the novel object location (comparison between discrimination index and chance: $t=3.787$, $df=7$, $p=0.0068$) whereas animals in the home cage condition did not (comparison between discrimination index and chance: $t=0.3083$, $df=7$, $p=0.7668$). This difference cannot be explained by differences between the groups in total object exploration times for either the sample or probe phases, as there were no significant differences in either measure between the groups (panel D - sampling: $t=0.8902$, $df=1,14$, $p=0.3884$; panel E - probe: $t=0.05349$, $df=1,14$, $p=0.9581$).
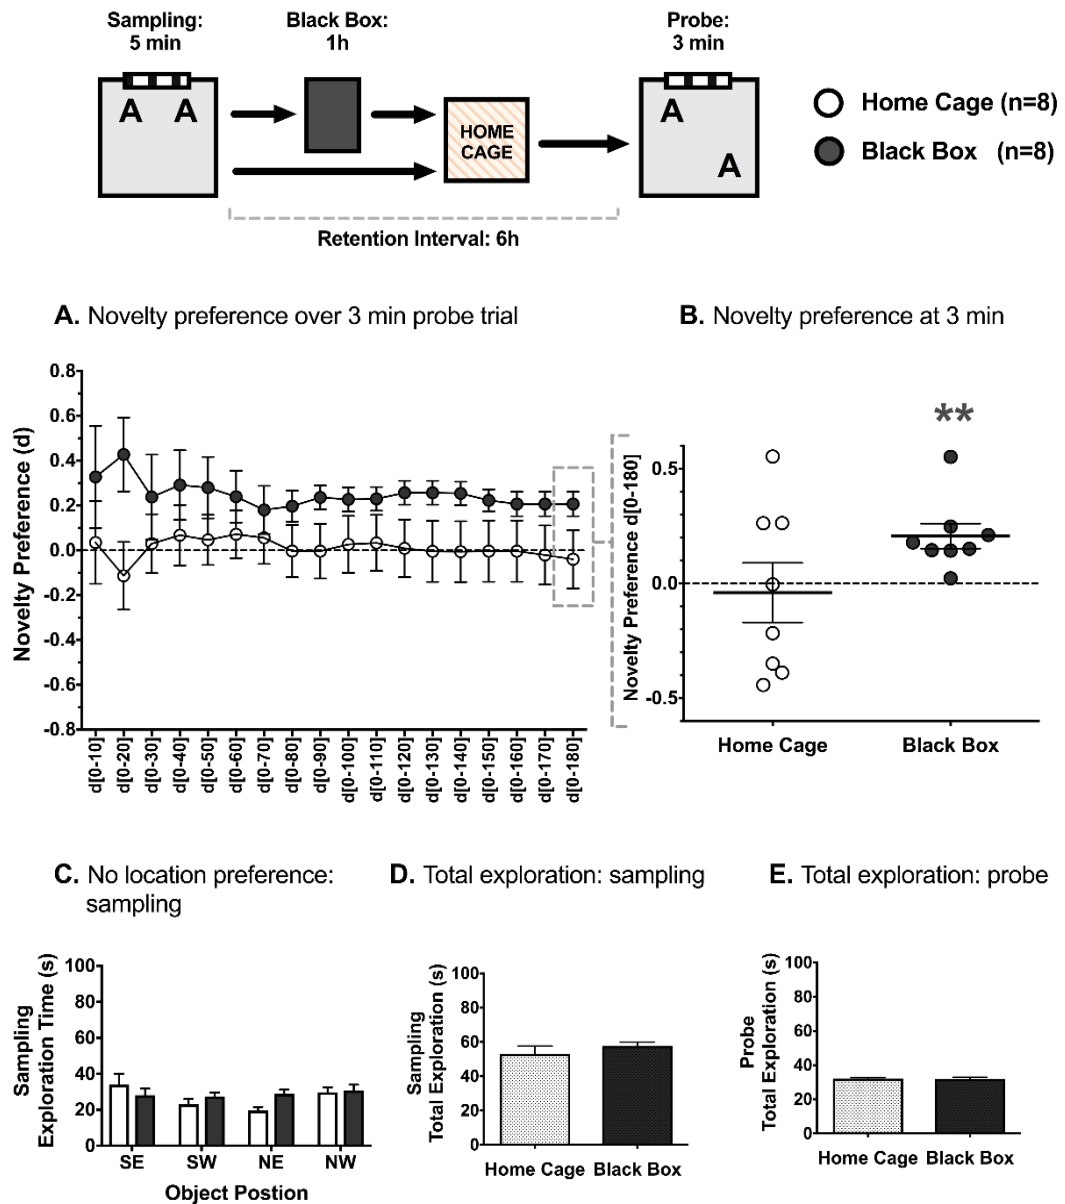
Figure 2.3: Reducing interference after learning promotes retention of 6h long-term object location memory. A) Accumulative novelty preference over the entire 3 minute probe trial B) The 3 minute mark is highlighted, showing significant object location memory for the rats exposed to the black box after learning. ** One-sample t-test: t=3.787, df=7, p=0.0068. C) No location preference shown in sampling for either group. D and E) Total exploration time for sampling and probe trials not different between groups. Error bars represent SEM.

There were no object location preferences for either group during sampling [no main effect of location: $F_{(3,24)}=2.085$, p=0.1288; no interaction between location and group: $F_{(3,24)}=1.805$, p=0.1731] (panel C). The variances between the two groups were significantly different [F=5.68, DFn=DfD=7, p=0.0355], so an unpaired t-test with a Welch's correction was used. Surprisingly this showed no significant difference in discrimination index at the 3-minute time-point between the two groups (t=1.749, df=9.391, p=0.1127).

Overall these results indicate that rats put back into the home cage immediately the learning phase of his object-location task did not show significant memory for object locations 24 h later, but that those exposed to the black box for the first hour after encoding (before being placed in their home cage), did. These data are indicative of memory enhancement in this reduced interference condition, and consistent with results of studies in human subjects described earlier, although the difference between groups did not reach significance.

2.3.2 Experiment 2: Holding cage in the light vs. dark (red) box

	Holding Cage in Light	Dark (red) Box
"Interesting" Visual Stimulation	✓	✗
"Habituated" Visual Stimulation	✓	✗
Social Interaction	✗	✗
Sleep	✗	✗

Table 2.4: The differences between the two conditions in experiment 3

The next experiment had two aims. Firstly it sought to refine the experimental protocol, allowing the experimenter to better determine when to use gentle handling procedures and ensure the rat was not sleeping. To do this a dark (red) box (RB) was used that contained a red light instead of being in darkness. This aimed to be an identical condition to that of the black box used previously, with no interference from social interaction or visual stimulation. A longer delay was also used (24h instead of 6h) to determine whether the effect seen in the previous experiment could last for a much longer duration.

The second aim was to test whether social interaction was a critical variable in the 'black box effect' seen in the first experiment. To do this a holding cage was used as the second condition. This cage was similar to the home cage, and kept in the same room, ensuring that the 'interesting' (unfamiliar) visual stimulation remained. The one difference was that animals were socially isolated, minimising effects from social interaction. This task therefore determined whether social isolation was enough to produce the 'black box effect' seen previously, without the need for reducing visual stimulation (summarised in Table 2.4).

At all time points of the 3 minute probe trial, animals exposed to the holding cage in the light appeared to have no preference for the novel object location, shown in Panel A Figure 2.4. Moreover, at some intervals in this trial these animals appeared to have a preference for the old object location (negative discrimination index). Animals exposed to the RB on the other hand had a preference for the novel object location throughout the probe trial. This difference is also displayed in panel B, showing the significant preference for exploring the novel location in the RB group at 3 minutes (comparison between discrimination index and chance: t=6.02, df=7, p=0.0005), compared to no preference in the holding cage group (comparison between discrimination index and chance: t=0.05257, df=7, p=0.9595). This difference between groups is significant (t=3.841, df=14, p=0.0018), and cannot be explained by variances in exploration times of the objects in sampling (panel D: t=1.029, df=14, p=0.3210) or probe trials (panel E: t=1.089, df=14, p=0.2945). There was however an almost significant effect of object position [$F_{(3,24)}$=2.588, p=0.0765] and a significant interaction between position and group [$F_{(3,24)}$=3.044, p=0.0483] on sampling exploration time (panel C), with positions in the south being explored more than those in the north. As positions were counterbalanced in a way that ensured SE and SW positions were both the old and the novel location within groups, this difference should not have affected the overall results.

These results confirm that the red box can replicate the black box condition, as both showed significant nOL memory, and that the memory enhancing effect can be seen when longer delays are used. The results also imply that social isolation alone is not sufficient to enhance spatial memory as the holding cage group did not express nOL memory. This suggests that the visual scene within the lab environment is preventing the formation or expression of memory, even though the animal is socially isolated.

Figure 2.4: Reducing interference from social interaction and visual input after learning promotes retention of 24h long-term object location memory. Reducing interference from social interaction alone does not. A) Accumulative novelty preference over the entire 3 minute probe trial B) The 3 minute mark is highlighted, showing no memory with social isolation alone, but significant memory when both social interaction and visual input are reduced. ***One-sample t-test: t=6.020, df=7, p=0.0005. Unpaired t-test: t=3.841 d=14 ,p=0.0018. C) No location preference shown in sampling for either group. D and E) Total exploration time for sampling and probe trials not different between groups. Error bars represent SEM.

2.3.3 Experiment 3: Light (white) box vs. dark (red) box.

	Light (white) Box	Dark (red) Box
"Interesting" Visual Stimulation	✘	✘
"Habituated" Visual Stimulation	✓	✘
Social Interaction	✘	✘
Sleep	✘	✘

Table 2.5: The differences between the two conditions in experiment 4.

Social isolation in the light, in the holding cage, was not enough to enhance spatial memory. The holding cage in the second experiment was in a holding rack in a lab environment in which rats may have experienced a complex visual scene made up of furniture, apparatus and computers. As the home cages containing the animal's cage mates were also kept in this room, there could have been the sounds of these cage-mates moving or vocalising as well as their smells serving as other interfering stimuli. This final task sought to establish whether the visual (or other) interference from the goings-on in a holding room had prevented memory enhancement in the "light holding cage" condition of Experiment 2, or whether it was merely visual interference from light. To test this, animals were placed individually into a plastic cage which was within one of two identical boxes: one box was illuminated by a white light (light condition) and one with a red light (dark condition). The latter condition had been previously used in Experiment 2 as the low interference condition. The differences between these conditions are summarised in Table 2.5.

As shown in Figure 2.5, panel A, the animals exposed to the red box (RB) once again showed preference for the novel object location throughout most of the probe trial, this preference becoming more apparent by the end of the trial. In contrast, those exposed to the white box (WB) showed no preference throughout the trial. Again, this is highlighted in panel B, showing that at the 3 minute time point the animals in the RB group show significant preference compared to chance levels (comparison between discrimination index and chance: $t=2.637$, $df=17$ $p=0.0336$), and animals in the WB group do not (comparison between discrimination index and chance: $t=0.01761$, $df=7$, $p=0.9864$). Although the difference at this time point is not significant ($t=1.870$, $df=14$, $p=0.0826$), a power analysis test shows that with 2 more animals in each group (power – 0.5) or 11 animals in each group (power 0.8), this would be significant ($p<0.05$).

A. Novelty preference over 3 min probe trial

B. Novelty preference at 3 min

C. No location preference: sampling

D. Total exploration: sampling

E. Total exploration: probe

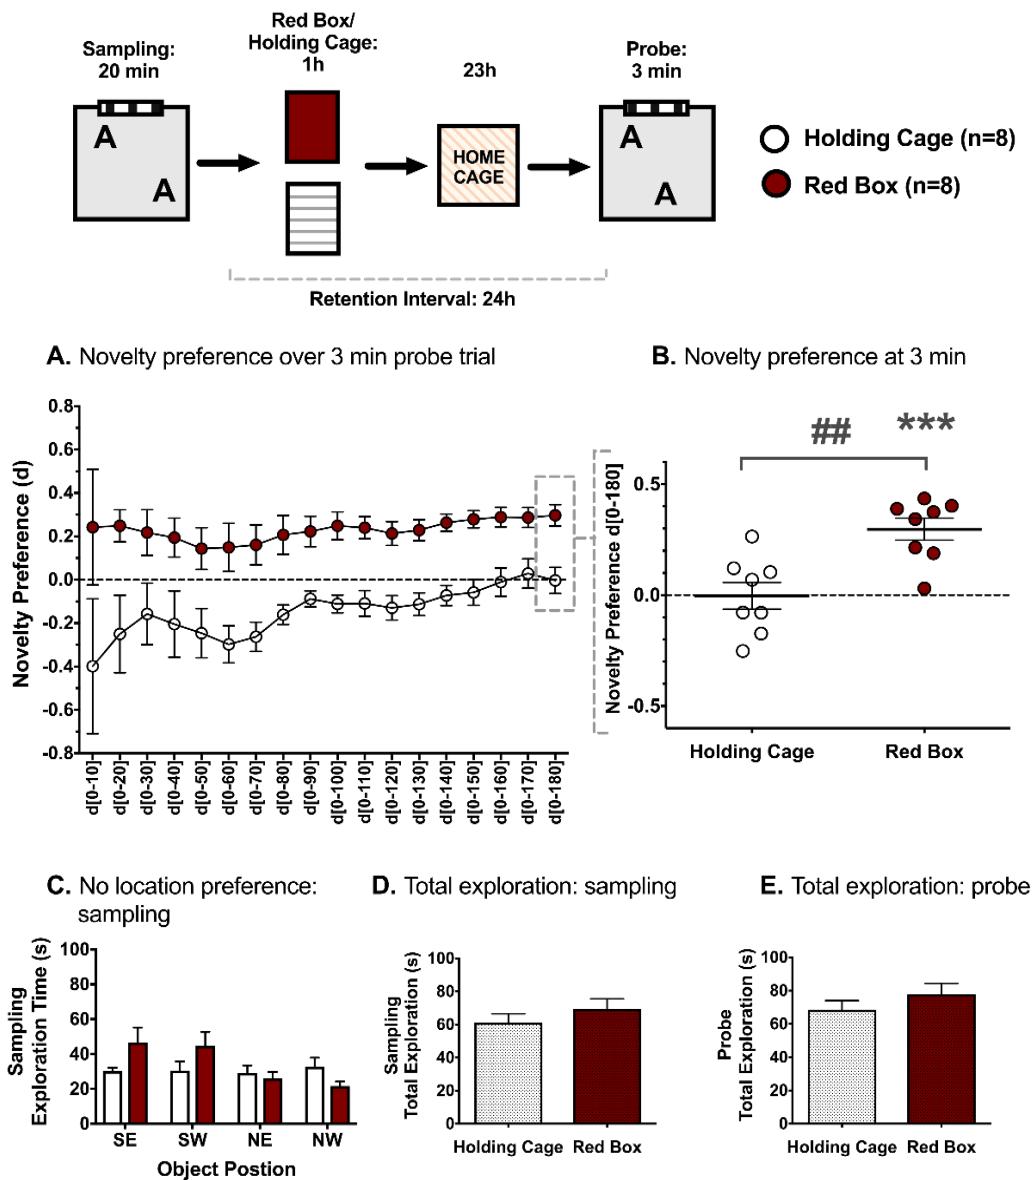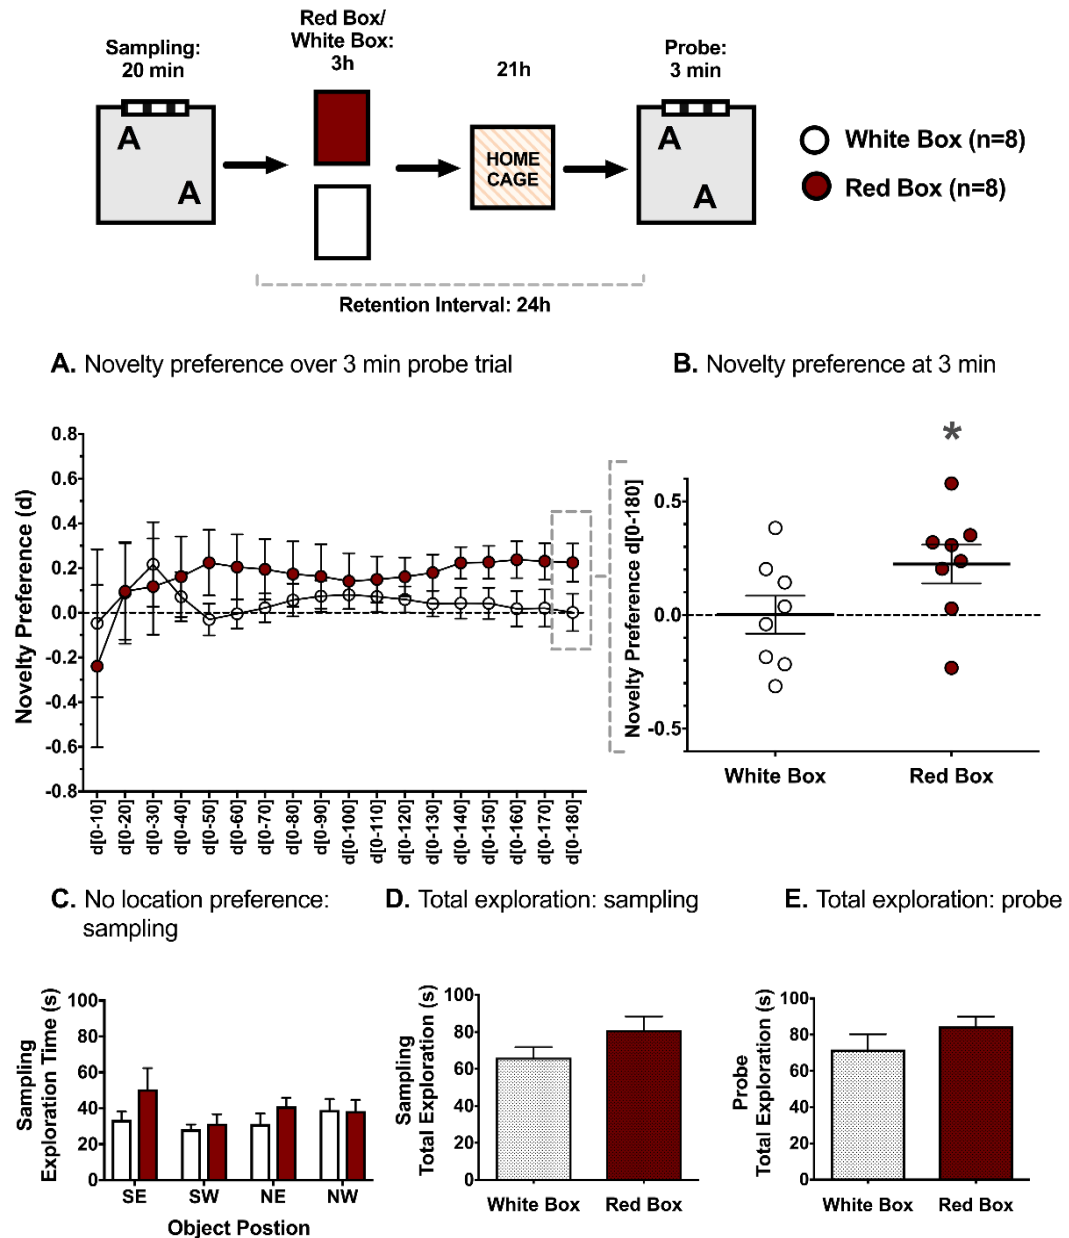**Figure 2.5: Reducing interference from social interaction and light after learning promotes retention of 24h long-term object location memory. Reducing interference from social interaction and familiar visual stimulation does not.** A) Accumulative novelty preference over the entire 3 minute probe trial B) The 3 minute mark is highlighted, showing significant 24h memory in the red box group, but not the white box group. *One-sample t-test: t=2.637, df=17 p=0.0336. C) No location preference shown in sampling for either group. D and E) Total exploration time for sampling and probe trials not different between groups. Error bars represent SEM.

Again, there was no significant difference in exploration time between the two groups in either sampling (panel D: t=1.588, df=14, p=0.1346) or probe trials (panel E: t=1.273, df=14, p=0.2238), and no object location preferences for either group during sampling [no main effect of location: $F_{(3,24)}=1.326$, p=0.2891; no interaction between location and group: $F_{(3,24)}=0.7601$, p=0.5275] (panel C). This shows that visual stimulation from a highly habituated and unchanging environment, or even light itself, is enough to prevent memory enhancement, even when socially isolated.

2.3.4 Experiment 4: Home cage in the light (white light) vs. home cage in the dark (red light)

	Home Cage in Light	Home Cage in Dark
"Interesting" Visual Stimulation	✓	✗
"Habituated" Visual Stimulation	✓	✗
Social Interaction	✓	✓
Sleep	✗	✗

Table 2.6: The differences between the two conditions in experiment 2.

Experiment 3 showed that minimising visual stimuli is required for the memory enhancement of the 'black box effect'. Experiment 2 showed that social isolation alone was not enough to drive this effect. This suggests that social isolation might not be a requirement for memory enhancement, and rather the only requirement of the 'black box effect' is darkness. If the latter is true then if the home cage had been kept in relative darkness during the first experiment, there would not have been differences between the home cage and black box conditions and both would have expressed memory. The next task therefore sought to ascertain whether the memory enhancing effect of the black box condition observed in Experiment 1 required the animals to be socially isolated in the dark or if the effect remained when in the presence of cage mates.

To this end, the home cages of the animal were used in both conditions, however in one condition the home cage was contained within the white box (WB) and in the other it was contained within the red box (RB). The latter condition was therefore the same as both the black box in Experiment 1 and the RB in Experiments 2 and 3, apart from the presence of cage mates leading to social interaction. The only difference between the two conditions was the presence of light (summarised in Table 2.6). If social isolation is not a requirement for the 'black box effect' the RB condition should produce memory and the WB condition should not.
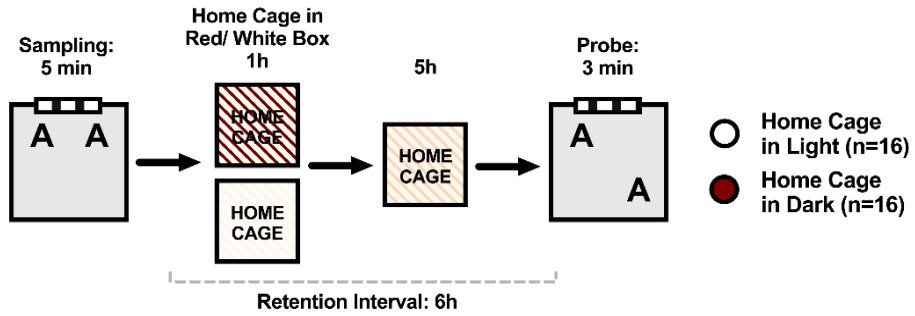
A. Novelty preference over 3 min probe trial   **B.** Novelty preference at 3 min
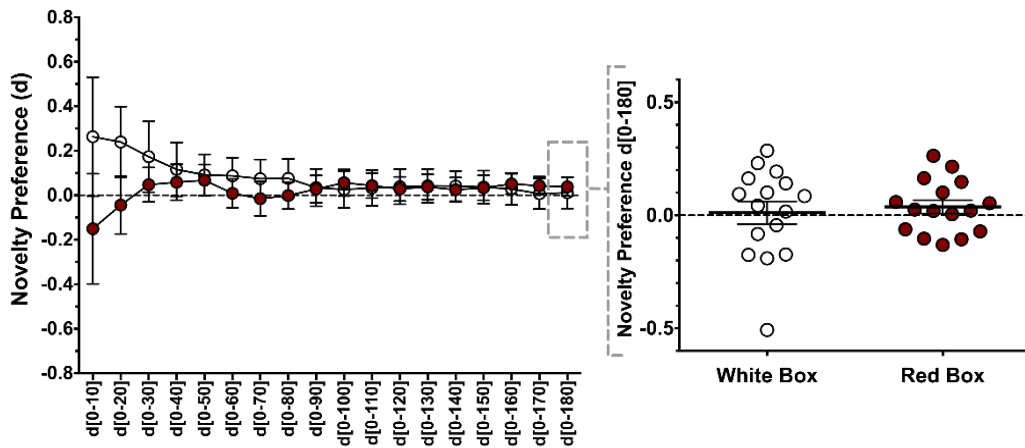
C. No location preference: sampling   **D.** Total exploration: sampling   **E.** Total exploration: probe
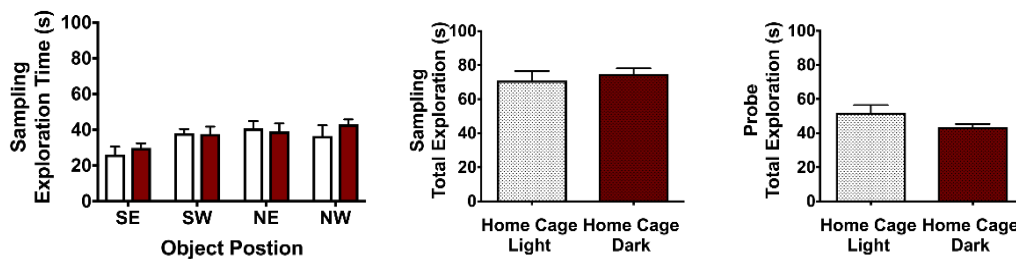
Figure 2.6: 6h long-term object location memory is not enhanced by darkness alone. Social interaction in the home cage disrupts memory enhancement A) Accumulative novelty preference over the entire 3 minute probe trial B) The 3 minute mark is highlighted, showing no memory for either group put in the home cage after learning. C) No location preference shown in sampling for either group. D and E) Total exploration time for sampling and probe trials not different between groups. Error bars represent SEM.

Animals did not appear to have a preference for the novel object location at any point across the probe trial, regardless of whether they were exposed to the home cage in the light or in the dark during the post sampling period, shown in Panel A Figure 2.6. This lack of memory expression is highlighted in panel B, showing no significant preference for the novel object location for either group (comparison between discrimination index and chance: home cage in light: t=0.2232, df=15, p=0.8264; home cage in dark: t=1.272, df=15, p=0.2228). There was also no significant difference in the discrimination index between groups at the 3 minute time point (t=0.4707, df=15, p=0.6447). There were no differences in exploration times for the sampling (panel D – sampling: t=0.7613, df=15, p=0.4583) and probe trials (panel E – probe: t=1.723, df=15, p=0.1055). Although there was an effect of position on sampling exploration time [F(3,28)=3.203, p=0.0384], there was no interaction between position and group [F(3,28)=0.6084, p=0.6084] (panel C). As object position was counterbalanced between groups, this position effect should not have affected the behavioural results. Overall this suggests that animals are required to be socially isolated for the "black box effect" to be seen.

2.3.5 Results overview

These results show that the benefit of restful wake seen in humans can be reproduced in rats. Overall these behavioural tasks show that this 'black box effect' requires both social isolation and darkness: one or the other is not sufficient to enhance spatial memory. Animals must be alone and without visual input during the post-sampling period for this effect to work, even if the visual input is from a highly familiar and 'uninteresting' environment.

2.4 Discussion

2.4.1 Experiment 1: Home cage vs. Black Box

The black box condition - emulating the dark room used in human wakeful rest studies - produced a significant expression of spatial memory. The "control" condition, which in this case was to place the animals back into their home cage, did not. Although there was no significant difference in discrimination scores between the black box condition and the home cage condition, there was significantly more variation of scores in the home cage condition. If an animal has no memory one would expect it to explore old and new object locations for similar lengths of time (d[180]=0). However, in a population some animals might have individual preferences, leading to a spread of d[0-180] values. Due to counterbalancing, one would still expect the average d value for the entire population to be normally distributed around zero. If a population of animals expresses memory, one would expect the majority of animals to explore the new object more than the old, ignoring individual preferences. This would lead to d[0-180] values skewed in the positive direction with less variance over the population, as seen in the present experiment. Under conditions of reduced visual stimulation after learning, humans show enhanced long-term memory (Dewar et al., 2012) and the integration of new information into a spatial map (Craig et al., 2016). The results of the present experiment therefore replicate those found in humans, suggesting that the 'Black Box effect' is conserved across species. As shown in previous human studies, the distractor task did not have to be similar to the material learnt (for example learning new object locations) to disrupt memory consolidation. In the case of the present experiment, the home cage condition provided sufficient interference. This RI could have stemmed from social interaction with cage mates. RI could also have been caused by visual stimuli from within the cage or the surrounding lab environment. In all of these instances, new information could have been encoded, disrupting the consolidation of the previously encoded nOL memory. As with restful wake in humans, the black box condition protected the animals from these RI mechanisms, leading to the expression of nOL memory.

Previous studies using a nOL protocol have shown that a 20 min sampling phase is required to produce consistent expression of 6h nOL memory, and that when the delay is extended to 24h this length of sampling phase only leads to novelty preference in the first minute of exploration (Ozawa et al., 2011). Other studies producing reliable 24h nOL memory expression over the entire 3 min probe trial used multiple sampling sessions over hours (5

min x 5 (Hardt et al., 2010)) or over days (7 min/ day for 3 days (Gaskin et al., 2009)). This indicates that the black and red boxes are indeed enhancing memory expression over what would normally be expected in the protocols used.

2.4.2 Experiment 2: Holding cage in the light vs. dark (red) box

The next experiment explored the possibility of RI in the absence of social interaction. Was social isolation enough to produce the 'black box effect' shown previously or were more factors at play? This time a red box acted as the dark, isolated box condition, allowing the experimenter to ensure the rats were not sleeping. Again this dark box group expressed memory, whereas the socially isolated group in the light (the holding cage condition) did not. Although animals were socially isolated in the holding cage, new visual information from the lab environment could still be encoded, suggesting that visual information is important to RI mechanisms. There were, however, other animals in close proximity to the holding cage. It could have been the interaction with these cage mates, albeit from a distance, that produced RI, possibly through odours or vocalisations.

2.4.3 Experiment 3: Light (white) box vs. dark (red) box.

The third experiment sought to reduce novel visual stimulation, whilst in social isolation away from other possible social cues (WB condition). Surprisingly, even this plain, unchanging and highly habituated environment led to no expression of OL memory. Only the dark box group showed nOL memory expression. This demonstrates that the black box effect is evident even when compared to a control group with minimal interference levels, a factor not yet tested in humans. This also suggests that interference not only stems from new learning or mental exertion but can simply occur due to either a highly habituated visual stimulus, or light itself. This goes against the currently held view of possible sources for interference.

2.4.5 Experiment 4: Home cage in the light (white light) vs. home cage in the dark (red light)

The final experiment highlighted the impact of social interaction on memory consolidation. As expected, animals that were exposed to the home cage in the light condition expressed no nOL memory. Those that were exposed to the home cage in the dark (RB) condition also expressed no nOL memory. These animals had no visual stimuli for the hour after spatial learning. The only stimulus that could have led to the encoding of new information, and therefore RI, was social interaction with cage mates. This suggests that RI produced by social interaction is enough to impair memory consolidation. It also demonstrates that both social

isolation and darkness are required for the black box effect to be seen - something assumed but not previously known.

The fact that both social isolation and darkness are required for the black box effect to emerge suggests that visual information is just as important as social interaction to RI mechanisms. Although no distractor tasks were used, these two experiments show that stimuli of the same primary modality as the nOL task (vision) and of different modality of the nOL task (social interaction) both produce enough interference individually to disrupt memory consolidation.

This finding is unexpected as it has been shown in previous human studies that the more similar the learning and distractor task, the less participants can recall (McGeoch and McDonald, 1931; Dey, 1969). Similar tasks and modalities are more likely to encode into the same neuronal networks, increasing overlap and the potential for interference mechanisms (Martínez et al., 2014). It would therefore be hypothesised that a learning task and distractor task of the same modality would produce maximum levels of RI; whereas tasks of different modalities could still produce RI, but on much lower levels. It could be suggested that a floor effect is present in our task. If the difficulty of the task was decreased, visual stimuli could have more of an impact than that of social interaction. On the other hand, it could be argued that visual input of the lab environment is not the same as visually encoding a spatial map. Both Dewar et al. (2007) and Robinson et al. (1920) found that unless highly similar, all material appeared to have a comparable detrimental effect. For example, participants had to learn 15 verbally presented nouns followed by various filled delay conditions. Listening to a radio recording with subsequent follow up questions produced the same amount of interference as spot-the-difference problem, even though verbally presented nouns and a radio show would be predicted to be more similar and therefore cause a larger detrimental effect. It could also be argued that nOL may not be just a visual task. Although the spatial aspect of the task should be visual as the cues used were not tactile, the rats still explored the objects by sniffing and whisking. Other modalities could therefore have contributed to the spatial memory.

A further experiment could be carried out using two different modalities, odour and vision, to investigate these interactions. Rats would be trained on an odour discrimination task, with four post-learning conditions: a dark box with no odour; a light box with no odour; a dark box with a strong but non-aversive and highly habituated odour; and a light box with the same odour as dark box. If interference is entirely modality-dependent, both the dark and light boxes without odour would show significant memory; those with odour would express no memory. If interference is entirely visual or light dependent, only the dark boxes would have memory. If interference occurs through an interaction of modalities, the dark, non-odour box would express the strongest memory; the light, odour box wouldn't express memory; and the order of the two other conditions would depend on the hierarchy of vision and odour. It could be suggested that an interaction is the most likely outcome. Even within the same modality, there can be varying levels of interference, depending on either absolute task similarities or required mental effort (McGeoch and McDonald, 1931). On the other hand, it could be that modality is not important, and it is darkness that is required to produce the enhancement effect.

Another interesting finding was that a highly habituated visual stimulus, or light itself, appeared to lead to RI mechanisms. RI is postulated to occur when encoding of new information disrupts the consolidation of previously encoded material. This process therefore requires new learning to disrupt the encoding of object location spatial memory. Most of the post-learning conditions support this definition, as change (and therefore learning) was possible between different exposures. The home cage conditions could have led to new learning, as social interactions were always occurring, some of potentially high valence. The holding cage condition could also theoretically have led to new learning, as the experimenter could be seen in the lab environment, providing a potential distraction. As the home cage group had no memory, even when in the dark, it is highly likely that this need for social isolation stems from the high levels of interference that social interactions can produce. The need for darkness however, is more challenging to explain. The WB should have already been encoded during previous habituation exposures. The animals were by themselves in a controlled environment, with nothing changing between exposures. Although one could argue that the gentle handing procedure qualifies as a novel stimulus, this occurs in both the dark and light boxes equally, so would create the same amount of encoding, and therefore interference, in both groups. This effect also persists if the animals

are not naïve to the experimental procedure, where gentle handling would not be novel. This leads to two questions: is there encoding in the WB; and if there is no new learning, how is memory consolidation disrupted in this group? For the former question further experiments need to be conducted. As there is no memory expression in the WB one would assume that there would be no synaptic consolidation. Therefore, if animals were sacrificed directly after the post-sampling condition and hippocampal immediate early gene levels analysed, one would expect to find increased levels in both conditions if there was encoding in the WB. If no encoding occurred, only the RB would have increased levels. However, only the latter would lead to conclusive results.

If these experiments do suggest that there is no new learning during the post sampling phase, it could be that an alternative interpretation is required. Importantly, the results show that even when the post sampling period conditions have very minimal interference, memory is expressed in the dark box group but not the control group. The OL memory learned is therefore too weak to normally be consolidated. This WB group does not need an explanation of increased interference mechanisms. Instead, the baseline threshold for memory consolidation could be too high, preventing the completion of either cellular or systems consolidation in the white box. This suggests that exposure to the black box could lower this threshold, allowing cellular and systems consolidation mechanisms to occur. It is possible that if the task were made easier, both groups would express memory (summarised in Figure 2.7).
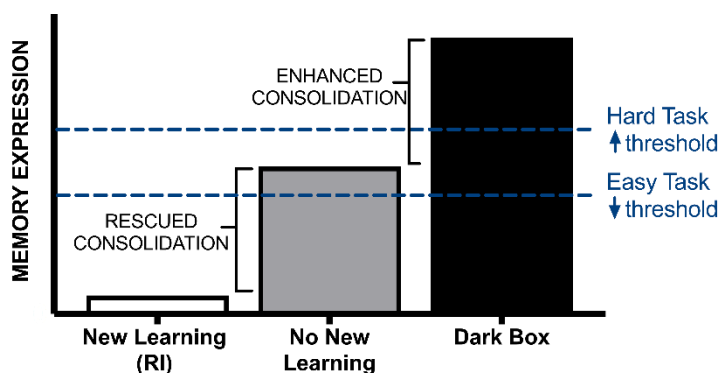
Figure 2.7. Hypothesised differences in memory expression. Different difficulty levels of memory task (blue dotted lines) can highlight different types of consolidation mechanisms: rescued or enhanced.

This has been shown previously in human interference studies, where varying levels of difficulty of spatial memory tasks led to both ceiling and floor effects within the same virtual route-learning experiment (Craig et al., 2015). It could be hypothesised that pathological

problems with memory encoding or retention, such as those seen in patients with mild cognitive impairment or focal brain injury, could lead to increases in consolidation thresholds. This would manifest as a benefit from wakeful rest in patients, even when healthy controls show a ceiling effect. This is in fact the case, having been shown previously on more than one occasion (Cowan et al., 2004; Della Sala et al., 2005). Results from Cowan's research also suggested that amnesiac patients needed an intact temporal lobe, including the hippocampus, to obtain any benefit from wakeful rest. This indicates that hippocampal consolidation mechanisms are somehow required for this benefit to occur. The question should therefore be: what mechanisms are occurring in the black box to enhance memory over the control baseline? The requirement for social isolation can be explained via traditional RI mechanisms, as new learning is possible with social interaction. The need for reduced visual input cannot be explained via traditional RI mechanisms, as a highly familiar visual input should not lead to new learning. Therefore this need for darkness is an attractive candidate for mechanisms underlying memory enhancement.

An important factor in consolidation mechanisms is the exact time-period after learning. Previous studies of RI have highlighted a temporal gradient, where memories are less susceptible to interference as time goes on (Britt, 1935; Dewar et al., 2007). It could therefore be hypothesised that the requirement for social isolation would be time-dependent, becoming less effective as time passed. Studies have shown that in humans, exposure to interference after 10 minutes of wakeful rest does not impair memory enhancement (Dewar et al., 2012). This could be exposing a temporal gradient on the order of minutes, in contrast to the temporal gradient of consolidation, which is thought to be on the order of hours. A time-dependence of 10 minutes would therefore imply that traditional RI mechanisms could underlie the memory enhancement seen. Alternatively, if memory enhancement persisted even if 10 minutes of wakeful rest was preceded by 30 minutes of exposure to interference it could suggest that wakeful rest is simply required at some point during the proposed consolidation window. This would imply that other non-traditional RI mechanisms were contributing to the memory enhancement. Future studies could repeat the black box experiments, but delay the post-learning conditions by 30 minutes. This would ensure that the wakeful rest condition still occurred within the consolidation window, but effects of rescue from retroactive interference would be minimised.

As previously mentioned, the post-sampling delay window after the nOL task is on the order of minutes to hours, so should encompass the labile period for synaptic consolidation mechanisms. This means that the enhanced consolidation potentially produced by the black box could be due to synapse-related proteins and pathways. One hypothesis could be that the black box enhances these mechanisms over baseline levels. However, it seems more plausible that the black box stops new synaptic consolidation, in line with the current retroactive interference dogma. Lower levels of acetylcholine (ACh) are produced in wakeful rest conditions (Marrosu et al., 1995). The opportunistic theory of consolidation outlines the role that these low levels of ACh could play in reducing interfering synaptic consolidation. Low levels of ACh are hypothesised to switch the brain into 'consolidation mode', stopping new incoming information being encoded at the synapse (Rasch et al., 2006). This consolidation mode is seen during slow wave sleep, but also during wakeful rest (i.e. in the black box). This explains how wakeful rest could protect synaptic consolidation mechanisms from RI caused by new learning. It does not however explain how the black box could enhance consolidation mechanisms over baseline levels when no new learning was apparent during the consolidation window.

Both the present experimental results and previous human studies suggest that wakeful rest leads to enhancement of the formation of a cognitive map and the integration of spatial memories into this map (Craig et al., 2016). This associative memory consolidation is thought to benefit greatly from the reactivation of hippocampal memory traces during systems consolidation (Carr et al., 2011). The reactivation of these encoded associations occurs without external input, appearing to be driven by the hippocampus itself (Alvarez and Squire, 1994). This reactivation serves to progressively strengthen the memory trace and distribute or transfer it throughout the cortex (Sutherland and McNaughton, 2000). In rodents this phenomenon has been termed place field replay. Place cells fire in a spatially specific manner, forming firing fields as the animal explores the environment. It is these fields that are thought to underlie the spatial map, which is used as an anchor for spatial memory. Studies have shown that the sequential replay of these place fields occurs on a compressed timescale during 200Hz oscillations in the CA1 pyramidal layer (Ylinen et al., 1995). These transient patterns of oscillatory neuronal firing, termed sharp wave ripples, allow for the coordination of firing between the hippocampus and cortical structures, such as the prefrontal cortex (Sirota et al., 2003). This coordination is thought to link these sparsely related structures,

leading to a memory trace that becomes less dependent on the hippocampus over time. Many studies have now shown that the magnitude of these neuronal reactivations can predict subsequent memory strength in both humans (Deuker et al., 2013; Oudiette et al., 2013; Staresina et al., 2013) and rodents (Dupret et al., 2010b; Ramadan et al., 2009). Selectively blocking these sharp wave ripples directly after learning can also impair or inhibit memory consolidation (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; Jadhav et al., 2012). These studies suggest a highly important role of neuronal reactivation during sharp wave ripples in spatial memory consolidation. Although originally associated only with slow wave sleep (Wilson and McNaughton, 1994), sharp wave ripples are now known to occur during the awake state in both rodents and humans (Kudrimoti et al., 1999; Tambini et al., 2010; Deuker et al., 2013). Notably, these occur most prominently during quiescence, i.e. during wakeful rest. As previously mentioned, wakeful rest is associated with low cholinergic activation. As well as reducing hippocampal synaptic plasticity, this decrease in ACh levels is characterised by the presence of sharp wave ripples (Buzsáki, 1989) and communication between the hippocampus and the cortex (Hasselmo, 1999). This implies that wakeful rest leads to an increase in systems consolidation mechanisms through the heightened presence of sharp wave ripples in the hippocampus. This in turn would enhance communication between the hippocampus and cortical areas, strengthening memory traces and enhancing the expression of nOL memory. Results from the WB versus RB conditions suggest that both social isolation and darkness are required for the enhancement of hippocampal memory systems consolidation. This implies that both social interaction and non-novel visual stimuli could impair this offline reactivation of newly learnt hippocampal memory traces.

2.6 Conclusions

The 'Black Box effect' is a phenomenon conserved between humans and rodents. Reduction of interference directly after spatial learning consistently enhances long-term hippocampal memories, but requires both social isolation and lack of visual stimulation. It is hypothesised that social interaction or a familiar visual stimulus alone is enough to retroactively interfere with systems consolidation processes by impairing sharp wave ripples and hippocampal place field replay.

Chapter 3: Reducing Sensory Stimulation after Spatial Learning Promotes Place Field Stability

3.1 Introduction

Whilst the novel object location (nOL) behavioural task is known to be hippocampal dependent, the exact role of the hippocampus in this task and in spatial memory in general is still unresolved (Lisman et al., 2017). Place cells have been proposed to be the best neural correlate for spatial memory since their discovery almost half a century ago (O'Keefe, 1976). It has been suggested that these cells provide the neuronal framework required for the consolidation and recall of memories associated with spatial locations, such as nOL memory. As the levels of interference after spatial learning have been shown to modulate the strength of such memories it is important to investigate whether interfering mechanisms can also modulate the properties of place cells. For example, it is possible that visual interference could affect the underlying stability of the so-called spatial map through changes in place field stability. Investigating how place cell properties change over time would indicate how the spatial memory of a new environment consolidates in the long-term. This would then allow the exploration into how reducing retroactive interference could affect spatial memory consolidation processes.

Place cells fire in a specific location within an environment, known as the cells place field. In the absence of any changes in the environment, or changes in task demands, a given cell typically fires in the same field across different sessions. These fields are often linked to cues within the environment and have been shown to consistently rotate with the cue if it is moved to a different wall within the environment. Place cells fire in a context-dependent manner, expressing fields in different locations when exposed to different environments. The formation of these fields occurs through the direct exploration of the environment (Rowland et al, 2011). When an animal is first introduced to a novel environment these cells slowly increase their firing rates as the animal repeatedly travels through the firing field of the cell (Mehta et al., 2000). It appears that whilst firing rates increase over the first novel exposure, over time average firing rates decrease as the environment becomes more familiar. When place cells were recorded as rats foraged in two different environments, one familiar and one novel, it appeared that average firing rates in the familiar environment were much lower than that of the novel environment (Nitz and McNaughton, 2004). Another study comparing place cell firing rates in a 'W track' spatial task showed that whilst the average firing rate of the overall population was less in familiar environments, as shown before, the in-field firing rate of certain place cells actually increased with familiarity. These place cells

71

initially had high levels of spatial tuning during the first exposure to the novel environment, meaning place cells with high levels of spatial information were selectively amplified and those with low levels of spatial information were selectively silenced (Karlsson and Frank, 2008). Therefore, after repeated exposures to the same novel environment the active population of place cells had enhanced in-field firing rates and spatial tuning which correlated with an increase in task performance. However not all studies appear to show this pattern associated with a familiar environment, where average place cell firing rates decrease and in-field firing rates and spatial information increase.

One such study involving random foraging during repeated exposures to an initially novel open-field showed that neither place cell firing rate or spatial information changed over six repeated 10 minute exposures within one day (Bett et al., 2013). Another similar study showed that three repeated 10 minute exposures to a novel environment also did not lead to changes in average firing rate or spatial information (Brandon et al., 2014). Whilst these inconsistencies between studies could be due to the shape of the environment or the nature of the rats experience (i.e. a task or spontaneous exploration) it is likely that the number of exposures used in these studies was not enough to express the low firing rates and high levels of spatial information associated with familiar environments. This is supported by further findings from the study by Brandon. Average place cell firing rates and spatial information were compared between the three exposures to the novel environment and one exposure to a very familiar environment. Even the last of the three exposures to the novel environment had significnatly different firing rates and spatial information compared to that of the very familiar environment. This in line with the differences found between novel and familiar environments by Nitz and McNaughton and Karlsson and Frank. Overall this suggests that novel environments are associated with high average firing rates, low peak in-field firing rates and low levels of spatial information, whereas familiar environments are associated with low average firing rates, high peak in-field firing rates and high levels of spatial information.

It is well known that place fields gradually stabilise over time and with repeated exposures to a novel environment. When an animal enters a new environment a unique population of place cells fire, each with a field representing an area of the environment (Muller and Kubie, 1987). Whilst initially these cells do not consistently fire when the rat moves through the place field, the expression of most of these fields stabilises over the first few minutes of the

exposure to this novel environment (Frank et al., 2004; Wilson and McNaughton, 1994). There is a small population that appears to be stabilised instantly, consistently firing when the rat traverses through the place field from the very start of the animals exploration of the environment (Frank et al., 2004; Ven et al., 2016). This within-session stabilization process is not dependent on NMDA-Rs or new protein synthesis, as blocking these processes does not affect the creation of new place fields (Kentros et al., 1998; Agnihotri et al., 2004). However, stabilisation of place fields over a longer time course appears to require both NMDA-R activation and protein synthesis. Specifically, if NMDA-Rs are blocked or protein synthesis inhibited during an animal's first exposure to a novel environment, when animals return to the same environment within an hour of the first exposure, cells typically fire in the same locations as before (i.e. they appear to be stabilised). This is suggested to be akin to early-LTP mechanisms and will not persist over long periods of time. However, if the same animals are returned to the same environment 6 h or 24 h later, the place fields remap, indicating that the long term stabilisation of place field does require the activation of NMDA-R and new protein synthesis (Kentros et al., 1998; Agnihotri et al., 2004). This is consistent with suggestions that long-term place field stabilisation is analogous to late-LTP and long-term spatial memory (Dragoi et al., 2003). In line with these suggestions, it has been found that the success of a spatial task can be associated with the long-term stability of place fields (Barnes et al., 1997; Kentros et al., 2004) in an NMDA-R dependent manner (Dupret et al., 2010a). Kentros trained mice to seek a particular location within an environment when an aversive loud noise sounded. This location was not marked out or obvious to the mice, so the surrounding spatial cues had to be used to perform the task correctly. It was found that the stability of the animal's place fields correlated directly to the animal's performance of the task. Mice had to express stable place fields to perform the task correctly. Dupret used a cheeseboard task where rats had to learn and recall three hidden food-rewarded locations. Over the process of learning the task place fields tended to stabilise more at goal locations, with place field stability directly predicting the animal's performance during a probe trial 2h later and during another probe trial the subsequent day. Both the stabilisation of place fields and the animal's performance were inhibited when NMDA-Rs were blocked, suggesting that the ability of the animal to learn this spatial task was dependent on both place field stability and NMDA-Rs.

These studies show large differences between place field stability in novel and familiar environments. Novel environments are associated with the gradual stabilisation of place fields within the exposure. However the expression of these fields is not stable between exposures to the environment over long-term periods. Familiar environments on the other hand are associated with stable place fields both within the exposure and between exposures over long-term periods. These familiar attributes have been associated with the success of performance in spatial tasks.

Although the population of cells expressing place fields within a given environment can remain very stable over long periods of time (Thomson and Best, 1990), the variability of individual place field firing within this can be extreme. Even when almost identical in speed and path, the number of action potentials fired in a single trajectory of the rat through a place field can vary from robust firing to completely silent (Fenton and Muller, 1998). This suggests that individual place fields signalling an absolute coordinate would be very prone to errors. It is much more likely that the entire population of place cells that fires within a given environment is used together to decode which environment and where in this environment an animal is. Therefore, whilst a spatial 'map' per se is contentious, the idea that place cells can signal where an animal is located within an environment, and that this is built into a temporal framework, is advocated in many studies and reviews on the subject (Eichenbaum, 2017; Kentros, 2006; Schiller et al., 2015). The firing of place cells with fields expressed near each other in an environment allows populations of cells (assemblies) to fire in quick succession, even though they are physically far apart in the CA1. This co-activation of the same patterns of neurons (assembly patterns) within specific time frames is thought to convey where the animal is with the greater precision needed for the internal representation of space (Wilson and McNaughton, 1993). These time frames can represent oscillatory activity, especially at the frequencies of theta and gamma due to their input in the coordination of place cell firing (Harris et al., 2003; Buzsáki, 2010), or the movement of the animal through space in real-time (Fenton and Muller, 1998). As with place field stability, this coordinated firing is suggested to be dependent upon NMDA-R mediated synaptic plasticity (McHugh et al., 1996b). A novel environment is therefore associated with less synchronous and coordinated firing of cell assemblies than a familiar environment.

The first aim of the experiment described in this Chapter was therefore to investigate the properties of place cells in rats during the spontaneous exploration of a novel open-field environment, and then during the following re-exposures at time-points associated with long-term spatial memory (6h and 24h). It is hypothesised that average place cell firing rates will be highest, and peak in-field firing rates and spatial information will be lowest during the first novel exposure. As found previously it could be predicted that average place cell firing rates would decrease and peak in-field firing rates and spatial information would increase over the three exposures as the environment becomes more familiar. However, as the protocol contains only three exposures spaced out over 24 hours it is more likely that the environment would not yet be familiar to the animal, and therefore firing rates and spatial information would not change between these three exposures.

As has been shown before in numerous studies, it would be expected that place field stability would be lowest between the 0h and 6h sessions as the environment is novel. The gradual stability of place field firing would increase within the first exposure and more so with repeated exposures to the same environment. As the 24 h exposure is the third time the animal will have explored the same environment, it is more likely that place cells will express stable firing fields between the 6 h and the 24 h time points. As the long-term stability of place fields is known to depend on NMDA-R activation and new protein synthesis, this stability at 6 h and 24 h should depend upon the length of the initial exposure. Therefore, whilst place field stability should increase over sessions, it is unclear whether 10 minutes of exploration is enough to produce a place field map that is stable 6h later. Given that nOL memory was not expressed at 6h unless interference was greatly reduced, it could be hypothesised that the population of place cells would also not express stable fields between the initial and the 6h exposures.

The second aim of this experiment was to investigate how the black box effect shown in the previous experiment could change these place cell dynamics. To this end, animals were placed into either the light (white) box (WB) or the dark (red) box (RB) after the initial exposure to a novel environment. This sought to replicate the post-sampling condition in Experiment 3 of the previous chapter, which showed that social isolation and lack of stimulation were required to enhance the expression of spatial memory, and that a familiar visual stimulus appeared to retroactively interfere with memory consolidation.

As the novel object location task required the animal to have familiarity towards the old object location compared to the novel object location, and the reduction of interference after learning enhanced this level of familiarity, it could be predicted that at the 6h exposure animals in the red box condition will treat the exposure as familiar, whereas animals in the white box condition will treat the exposure as novel. This leads to the hypothesis that average place cell firing rates will be lower whilst peak in-field firing rates and spatial information will be higher in the red box condition compared to the white box condition during the 6h exposure. It has also been shown that spatial coherence can predict performance on a spatial memory task (Kentros et al., 2004), suggesting that coherence will also be enhanced in the red box condition. As it is predicted that 3 exposures will not be enough to familiarise the animal with the environment, these predictions could also be hypothesised at the 24h time-point.

As the reduction of social and visual interference appeared to enhance spatial memory at both the 6h and 24h time-point, and it is predicted that the 6h exposure will act as a familiar environment only in the red box condition, it could be hypothesised that long-term place field stability would also be enhanced in the red box condition, compared to that of the white box. As stated previously, although individual place fields might not always be stable within and between sessions of exploration, when stability of all cells is averaged together you would still expect to see higher place field stability overall. In relation to place field stability it could also be hypothesised that the co-activation of cell assemblies would be increased between sessions in the RB condition.

The following experiment will therefore elucidate how place field properties change as a novel environment becomes more familiar. It will also explore whether reducing interference after the first exposure to this novel environment can modulate these properties, in the same way that reducing interference enhanced the expression of nOL memory previously.

3.2 Methods

3.2.1 Animals

Eight male Lister Hooded rats were obtained from Charles River laboratories. At the time of surgery these animals weighed between 300-380g, and during the experimental procedures these animals were aged between 3-11 months. Before surgery all animals were housed in groups of four. After surgery, animals were individually housed to prevent damage to the implanted drives. All cages had tubes and chew blocks for enrichment. Animals were kept on a 12-hour light/dark cycle, with recording always performed in the light phase of the cycle. Before surgery animals were granted ab lid food and water. Post-surgery, once animals had recovered to pre-surgery weights and a minimum of 7 days after surgery, animals were put on food-maintenance and kept at 90-95% free-feeding body weight. Animals were therefore given 25-30g standard lab chow each per day, and free access to water. All procedures complied with the UK Animals (Scientific Procedures) Act (1986) and the European Communities Council Directive of November 24, 1986 (86/609/EEC). All animal experiments were carried out in compliance with protocols approved by the University of Edinburgh Animal Welfare and Ethical Review Board (AWERB), and under a UK Home Office Project License.

3.2.2 Pre-surgery habituation

Before surgery, animals underwent a minimum of 5 days of handling procedures. This served to reduce anxiety levels in animals, and to habituate them to human contact and gentle handling procedures used for sleep deprivation. Handling involved gently picking up and holding the animals, allowing them to be carried without any need for restraint. During this time, in groups, animals were put into large arenas containing many different novel objects. This served to reducing freezing and enhance exploration in the upcoming recording experiments. Animals were also transported to the recording room for a minimum of 3 days before surgery. Here, they were habituated to the holding room environment, and to eating whilst sat on the experimenter's lap - a necessity for plugging the animals into the recording equipment post-surgery.

3.2.3 Recording Device

Microdrives, based on a moveable tripod design (Kubie, 1984), were made and implanted unilaterally into the CA1 of the hippocampus. Each drive contained a bundle of 8 tetrodes, which could be advanced further into the brain using supporting screws. These tetrodes were used for recording both single unit activity and local field potentials, simultaneously, from 32 possible channels. As shown in Figure 3.1, the main drive components consisted of mill-max connectors to attach the rat to the recording equipment; drive 'feet' to attach the drive to the animal's skull; three supporting screws in a tri-pod arrangement; and a 3D-printed plastic Microdrive base.

To create the Microdrive base, two strips of mill-max (Mill-max Mfg. Corp, NY), each 9x2 pins in length, were inserted into the 3D-printed plastic and superglued in place. Three pins were removed from the mill-max for the tetrode bundle and two ground-wires. The three supporting screws (Precision Technology Supplies, UK) were inserted firmly into pre-printed holes. These acted as the drive mechanism, each turn lowering the whole base, and therefore the tetrode bundles, further into the brain. One full turn equates to an advancement of 310μm, although in practice drives were lowered on average $1/8^{th}$ of a turn at a time, equating to 38μm advancement.

Each tetrode was made of four lengths of HLM coated 17μm 90% platinum 10% iridium wire (California Fine Wire, CA), tightly twisted together into a bundle and gently heated to affix the lengths together. Eight of these tetrodes were then threaded through a stainless steel inner cannula (21 Gauge Hypodermic Tube, Small Parts Inc., FL) for added support and protection. This cannula was attached to the Microdrive base via a mill-max pin soldered onto the base of the cannula. Each end of tetrode wire was stripped of Teflon insulation and wrapped around an individual mill-max pin, producing high levels of conductance between the tetrodes and the mill-max connector. These wire-to-pin connections were secured using highly conductive sliver paint (Electrolube, UK), and bubble-tested to confirm conductivity. Two 10cm lengths of wire (Vishay Precision Group, Germany), each attached at one end to a mill-max pin, were inserted into the base. These acted as ground wires, reducing interference from non-neuronal electromagnetic sources during recording. To prevent any loss of wire-to-pin connections, the entire base was coated with a thin layer of spray acrylic (Electrolube, UK), followed by a thicker layer of nail varnish. A stainless-steel outer-cannula (17 Gauge

Hypodermic Tube, Small Parts Inc., FL) was threaded over the inner cannula to protect the electrode bundle both during and post-surgery. Once implanted, this outer-cannula rested on the surface of the brain. Three feet, made from Ampehnol (Amphenol Ltd, UK) coated in a thin layer of dental cement (Simplex Rapid acrylic denture polymer, Associated Dental Products Ltd. UK), were threaded onto the end of the supporting screws. The layer of dental cement helped to affix the feet to the animal's skull during surgery.

Either on the day of, or the night before, surgery, the end of the tetrode bundle was cut. This left 3mm of clean tetrode protruding from the end of the cannula. The tip of every electrode in the tetrode bundle was then cleaned and gold-plated (Non-Cyanide Gold Plating Solution, Neurlynx, MT) to reduce and normalise the impedance of each wire. The goal impedance was 120-150kΩ. Impedances were tested using an analogue impendence metre.
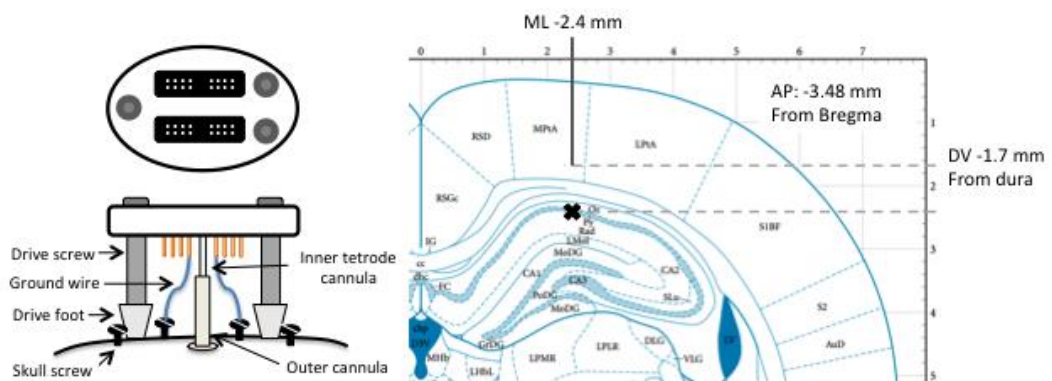
Figure 3.1: Left: Schematic diagram of implanted Microdrive resting on the animal's skull with key components labelled. Both top and side views are shown. Once the electrode was lowered into the brain and ground wires were connected, dental cement was layered around the drive feet and skull screws. This attached the drive to the skull and protected the outer cannula and ground wires. Right: Coordinates of electrode placement during surgery. Dark grey line represents the electrode. The top dashed grey line represents the depth of implantation during surgery. After the animal has fully recovered from surgery the electrode is slowly lowered (via the drive screws) in incremental steps until the electrode is within the CA1 pyramidal cell layer (marked with an x).

3.2.4 Surgery

Rats were anaesthetised with isofluorane gas (Abbott Laboratories, IL) for both induction and maintenance of anaesthesia throughout the surgery. The animal's head was shaved and disinfected before surgery began. At the start of surgery rats were subcutaneously injected with 0.08ml/kg bodyweight small animal Rimadyl (Pfitzer, UK), which acted as a long-lasting

anti-inflammatory pain analgesic. Rats were also subcutaneously injected with 2.5ml of an isotonic saline and glucose solution at the start of surgery and at any point during the surgery where hydration needed to be actively maintained. The animal's body temperature was maintained throughout surgery using a thermostatic heat blanket and their eyes were protected using hydrating eye-gel (Viscotears, TX). The animal's head was fixed into position using a stereotaxic frame (Kopf, CA) with non-traumatic ear bars. A nose cone was fitted to ensure optimal delivery of the anaesthetic and breathing was constantly monitored throughout. Drapes were used to ensure a sterile area for surgery was maintained. A first incision was made down the midline to expose the skull. Once enough of the skull was visible and bregma and lambda were clear and parallel, CA1 coordinates were calculated from bregma (AP -3.5mm, ML -2.4mm) and marked on the skull. Six holes were drilled, skull screws (Fine Science Tools, Germany) were inserted, and a thin layer of dental cement was spread over the exposed skull. This provided a strong connection to the skull for the drive to be attached to. A bigger hole of about 1mm was drilled at the CA1 coordinate and the electrodes were lowered in 1.7mm below the dura, once it had been pierced. The final coordinates of the implanted drive are shown in Figure 3.1. The outer cannula was lowered so it rested on the skull or dura and sterile Vaseline was used to fill the hole containing the electrodes to create a seal. The two ground wires were connected to skull screws using silver paint and dental cement was used to connect the feet of the drive to the skull screws and base of the skull. Once the drive was securely attached to the skull it was surrounded with electrical tape to protect the exposed ground wires. Animals were then removed from the stereotaxic frame and placed into a recovery cage kept at $30^{O}C$ and monitored for at least an hour, until fully conscious. Animals were given free access to food for at least a week after surgery. Pre-surgery body weight had to be reached before screening and recording commenced.

3.2.5 Equipment

Single-unit activity and local field potentials were recorded using a 32-channel Axona USB system (Axona Ltd., St. Albans, UK). The animal was connected to this system via millmax connectors on a headstage, where the signal recorded was also amplified. This headstage amplifier connected to a commutator and a pre-amplifier. After amplification, the signal was bandpass filtered between 300-7000Hz before being processed and recorded by specialist software (DACQ software: Axona Ltd., St. Albans, UK). A boom arm was attached to the

commutator. This was both for the comfort of the rat, as it reduced the overall weight on the head of the animal, and to improve the signal, as it prevented wires from becoming tangled. The dark (red) and light (white) boxes from the previous experiment were used (Chapter 2). The plastic cages from within these were used for rest sessions. Each cage contained bedding specific for an individual rat. This allowed animals to become highly habituated to their specific cage, enhancing the time spent asleep.

3.2.6 Recording environments

Six different recording environments were used. They were of similar sizes ranging from 63 cm to 76 cm, and on average 69 cm in diameter. Each environment represented a different context, and had different coloured walls, and different coloured and textured floors. Two environmental contexts were square and four were cylindrical. Each context also had different distal cues, including curtains to change the size and colour of the surrounding environment, and many different 3D cues hung on these curtains and on furniture throughout the room. These differences were in place to ensure animals could easily discriminate between the different environmental contexts, enhancing place field remapping between them. As the contexts were always in the same room, animals were transported to each context in an opaque and covered bucket. This served to decrease visual cues of the environment surrounding the context, giving the illusion of separate rooms.

3.2.7 Recording Procedure

Microdrives were connected to the recording system via two amplifiers at the end of a flexible recording cable. To do this, animals were placed on the experimenter's lap and given chocolate cereal to eat. This served as a distraction, allowing the experimenter to plug the amplifiers into the mill-max connector of the drive. This connection was bound with electrical tape to prevent disconnection during the experiment. Animals were then placed into a bucket with high walls whilst the reference channels and gains were selected and optimised via DACQ recording software (Axona systems Ltd). The gains equated to a user-defined threshold, where any single unit activity spike crossing this threshold was captured in a 1ms waveform-window, 0.2ms before and 0.8ms after the spike peak, and digitally time stamped. A channel was also chosen for the recording of local field potentials. This optimisation process was only carried out at the start of each experiment – gains were not changed during a two or three day protocol. This allowed for different sessions to be analysed together. Each of the two amplifiers had an infra-red LED attached. These allowed the animal's position to

be continuously tracked through an infra-red sensitive CCTV camera mounted above the open-field. This position information was combined with the time-stamped 1ms waveform-windows, allowing the recording and analysis of spatially modulated cells. Animals were then removed from the bucket and placed into the open field, and the recording was initiated. Chocolate cereal was thrown into the open field when animals ceased exploration, and to help guide the animal to explore the entirety of the environment. Rats were removed from the open field and placed into the bucket in-between recording sessions. The open field was cleaned using soapy water before every session.

3.2.8 Screening

To ascertain if and how much a drive needed to be advanced, each rat was 'screened' in an open box whilst foraging for chocolate cereal. This open box (50cm in diameter) was surrounded by a black curtain, shielding the animals from any cues that were available from the recording environments described above. Screening sessions lasted as long as was required to obtain full coverage of the open field. This was typically around 10 minutes. Both single unit activity and local field potentials were recorded and used to identify how close the electrode tips were to the pyramidal cell layer. Typical pyramidal cell waveforms and theta oscillations indicated that the electrodes were in the correct location. Absence of these signals, and the presence of many characteristic interneuron waveforms, indicated that the electrode tips were above the pyramidal cell layer. In this case the drive would be advanced. Once the majority of electrodes were thought to be in the correct location in the cell layer, rats would begin the experimental protocol. Screening protocols typically lasted a week, therefore serving as habituation to the recording procedure, minimising any animal stress during the experiment. After at least four of each animal's screening sessions, animals were also habituated to both the light (white) and dark (red) boxes for an hour.

3.2.9 Experimental Design and Protocol

Three different experimental designs were used, as the task evolved and improved. For each task animals were brought down from the animal house to the lab before recording each day. Animals were kept in an adjacent room to the recording room before and after recording protocols in their own cages. At the end of the recording day animals were transferred back to the animal house.

The first design (n=1), outlined in Figure 3.2 (top), involved recording neuronal activity during three 10 min exposures to two different novel contexts, over a period of three days. At the start of each protocol animals were transferred to the recording room in a black bucket containing bedding. This prevented them using the route into the room as a cue. This was especially important as each context was visually very different, but physically was within the same room. Animals were also held in this bucket whilst the gains were set, as had happened in screening. Debris was removed and soapy water was used to clean the environment before every exposure to remove possible odour cues. When placed gently in the context animals randomly foraged for cereal for the 10 min duration of the exposure.

Directly after the first (C1 0h) novel exposure to the first context animals were unplugged from the recording equipment and transferred to either the white or red box in the black bucket. Here they spent 3 h being kept awake via gentle handing (as used previously and described in Chapter 1). After the 3 h animals were transferred back to their home cage (animals were housed individually to prevent damage to the head stages). After 3 h in the home cage, where they were allowed to sleep, animals were carried back into the recording room in the bucket, plugged in and gently placed into the same environment as before (C1 6h exposure). After 10 minutes of random foraging animals were unplugged and transferred back to their home cage in the black bucket.

24 h after the initial exposure to the novel environment animals were transported into the recording room in the black bucket, plugged into the recording equipment and allowed to randomly forage for 10 min in the same environment as before (C1 24h). After this animals were held in the black bucket while the context was changed. This included the recording box that animal explored as well as the curtains and surrounding cues. Once the second environment and surrounding context were set up, animals were gently placed into this second novel environment for 10 min (C2 0h). The same procedure that was used for the first novel context was repeated, although if animals had been in the white box during the first run of the protocol they would be put in the red box during the second run and vice versa.
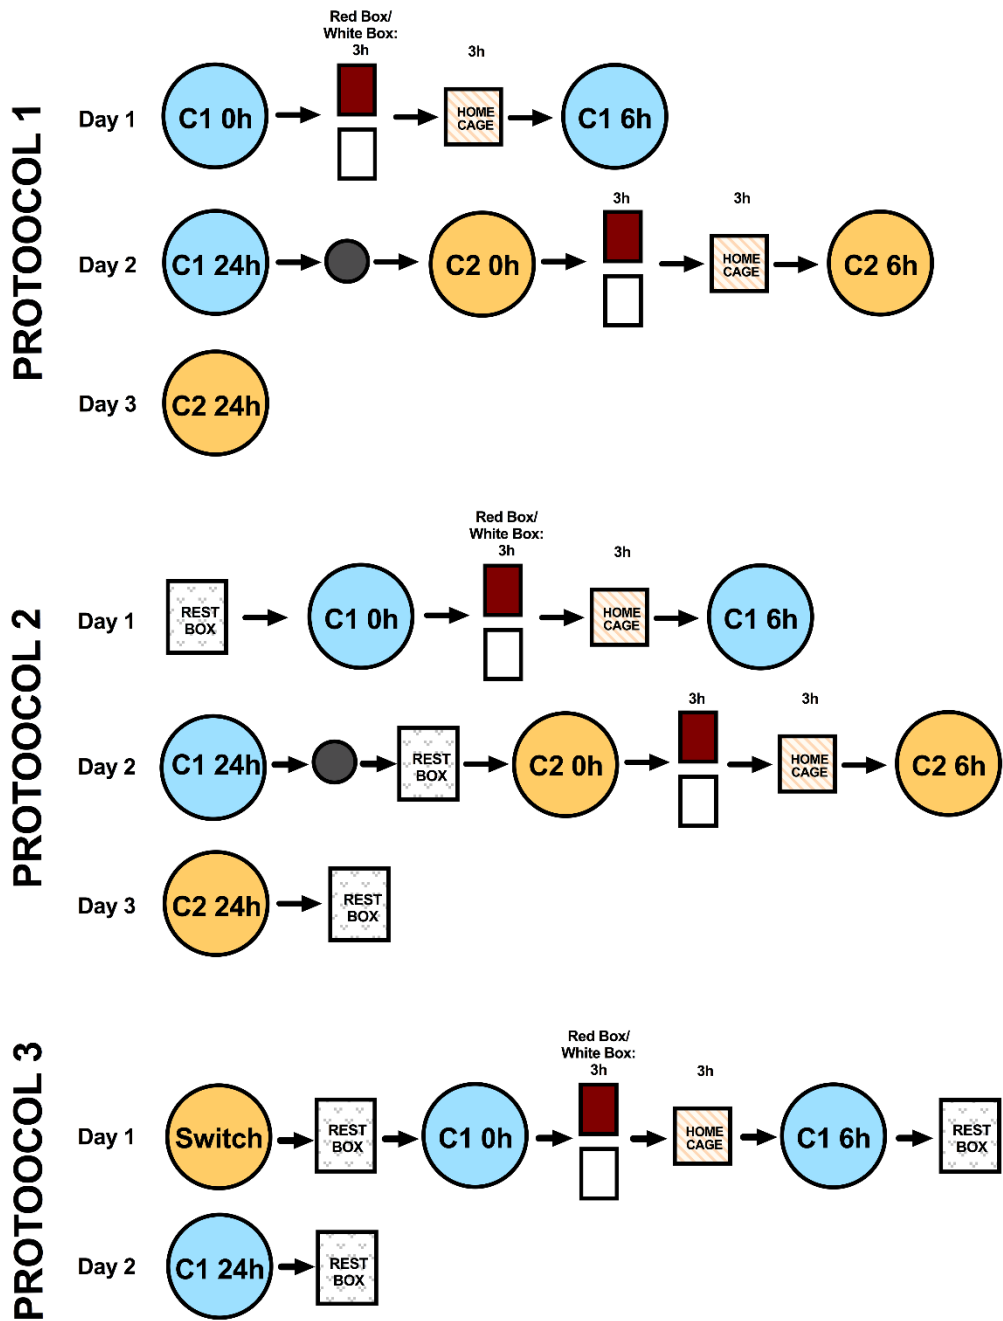
Figure 3.2: The three different protocols used throughout the experiment. The larger different colour circles represent different novel environments. The smaller black circle represents the black holding bucket. Unless specified above, every exposure was 10 minutes long.

The second design (n=2), shown in Figure 3.2 (middle), was exactly the same protocol as above, however rest-recording sessions were added before the initial novel exposures (C1 0h and C2 0h) and after the final exposure (C2 24h) to the environments. Whilst the environment and surrounding context were switched animals were held in the bucket as before. The rest session was therefore recorded after the switch occurred (but before the animals had experienced the switch). Recording cells in a sleep or rest state allowed for detection of neuronal replay. If a cell fired in one session, it was likely to fire during this rest state, albeit at a much lower firing rate. Over recording sessions cells could cease firing for two reasons: instability and remapping of the place fields; or electrodes drifting away from previously recorded cells. If the former was true, cells should still be detectable during this rest state, allowing the cells to be included in the analysis. On the other hand, if the latter were true, cells would not be detected and therefore removed from further analysis.

The final design (n=5), summarised in Figure 3.2 (bottom), aimed to enhance the remapping between the two novel contexts to allow animals to be run through the protocol more than once, increasing the number of cells included in analysis. To ensure the highest levels of remapping possible, radically different environments were used with very different surrounding cues. This included using different shaped environments (cylindrical and square) and using different parts of the recording room, closed off with different coloured curtains. Another aim of the final design was to reduce possible interaction effects from one session to the next. To do this, the gap between exposures to the two different contexts (C1 24h and C2 0h) was increased from 3-5 minutes (plus 10 minutes of rest-recording in the second protocol) to at least 5 days, increasing the time between the light (white) and dark (red) box conditions. An extra exposure was added before the 0 hour, novel exposure to the second context. This allowed the comparison of place fields between the two different contexts, to ensure that there was adequate place field remapping. An extra rest-recording was also added after 6h exposures to increase the number of cells that could be included. Animals were exposed to four contexts in this protocol, rather than the two contexts used in the previous protocols (i.e. the sequence shown in protocol 3 was repeated 4 times, with an interval of 1 week between sessions).

3.2.10 Analysis

3.2.10.1 Signal clustering and output

Following recording by DACQ software (Axona Ltd., St. Albans, UK) the signals were semi-automatically sorted into clusters using the KlusterKwik clustering algorithm (Kadir et al., 2014). This algorithm used various features such as waveform energy, amplitude, width, principle component, and time of peak. These clusters were visualised using Klusters (Hazan et al., 2006) to allow for the manual detection and deletion of noise clusters. Incorrect clustering could also be amended at this stage. A further MATLAB script then calculated properties of these individual clusters, producing a number of output parameters on a cell-by-cell, session-by-session basis. Cell waveform and theta modulation, via an autocorrelogram, were also visualised.

3.2.10.2 Rate Map Analysis

The MATLAB script also created rate maps for individual sessions. Rate maps are effectively heat maps of the spatial firing rate distributions of a cell. They take into account where the animal has been and how much each cell fired in each location of this environment. The following rate map analysis was based on previous methods (Leutgeb et al., 2007). The environment explored by the animal was divided into 2.5cm x 2.5cm bins. To ensure all bins were this size, irrespective of the actual environment size or distance from the camera, a pixel ratio was calculated for each environment and used to convert camera pixels into centimetres. The total number of spikes that fired in a given bin was divided by the length of time the animal spent in that location, therefore assigning a firing rate to each bin. This was then smoothed with a Gaussian filter (sigma 15) centred on each bin, giving more weight to the spikes closest to the centre of the bin (x). The equation for this smoothing is as follows:

$$g(x) = \exp\left(\frac{-x^2}{2}\right)$$

The animal had to be within 5cm of a given bin for a minimum of 100ms for that bin location to be included in the analysis (a minimum dwell time). These bin-specific firing rates were plotted in a heat map, showing where the preferred firing location of a cell was in a given environment.

A number of parameters required place fields to be identified from these rate maps. Place fields were defined as 'islands' containing 10 contiguous bins. Any bins with a firing rate less than 2.5 times the average bottom 50% firing rate (the median firing rate was calculated for the entire rate map, and any bins with a firing rate under this value were averaged to produce this average bottom 50% value) or less than 1Hz were excluded. This was to remove sub-field firing. Inbuilt MATLAB features for 'island' properties, such as area, max firing rate, mean firing rate and the weighted centroid, were then calculated for the top four firing place fields per rate map. The weighted centroid of the place field was used, as supposed to the actual centroid, because pixel intensities within the 'island' were included as weights in the calculation of the centre. The overall place field properties analysed were the number of fields and the average area of all the fields

3.2.10.3 Firing Rate

Firing rate (FR) of place cells was measured in two different ways. Firstly, the average FR for the entire session was calculated by dividing the total number of spikes in a session by the length of the session. This measure did not take into account whether the firing was place field specific. The next measure of FR used a measure from the identified place fields – the peak firing rate of the highest firing place field in a rate map.

3.2.10.4 Spatial parameters:

Spatial information (SI) is a measure of the information content (in bits) conveyed by a single spike (Skaggs et al., 1992), and how well this information can predict the animal's location. It was calculated using the following equation, where i identifies the bin in the rate map, Pi is the probability that this bin is occupied by the animal (dwell time of the animal in this bin/ total recording time), Ri is the mean firing rate of this bin and R is the mean firing rate of the entire rate map:

$$\Sigma i Pi \left(\frac{Ri}{R}\right) \ln \left(\frac{Ri}{R}\right)$$

Spatial sparsity equates to the percentage of the environment in which a cell fires. A very high percentage indicates that cell fires indiscriminately in a large proportion of the context, whereas a lower percentage indicates the cell fires preferably in one location of the

environment. This was calculated using the following equation, which uses the same values as described previously for SI:

$$\frac{(R)^2}{(R^2)} = \frac{(\Sigma PiRi)^2}{(\Sigma PiRi^2)}$$

Spatial selectivity also measures the selectivity of cell firing, but as a function of firing rate rather than the location within the environment. This used the peak firing rate of the rate map (FR_{peak}) and the firing rate for entire session ($FR_{session}$). A high level of spatial selectivity indicates that the peak cell firing rate is much higher than the sub-field firing rates.

$$FR_{peak} / FR_{session}$$

Spatial coherence measures how firing rates in one bin can predict the firing rates in neighbouring bins, calculated as a percentage. Coherence is calculated by returning the z transform of the correlation of the firing rates within a bin with the firing rates of the nearest 8 neighbouring bins. A high level of spatial coherence represents a more consistent pattern of place field firing.

3.2.10.5 Rate map Stability

The Pearsons correlation coefficient was calculated on a cell-by-cell basis for pairs of sessions. Two rate maps, one for each compared session, were centred using the smaller of the maps and scaled to equal sizes. This was to ensure either that the two different contexts were comparable in size and properly aligned, or that the two exposures to the same context were properly aligned to correct for any small movements in context position that could occur. Once these were reshaped to one-dimensional rate maps, all bins where both datasets had sufficient coverage were compared via a Pearson correlation. This therefore produced a correlation coefficient of the two rate-maps on a bin-by-bin basis.

Rate map stability was analysed on a rat-by-rat basis. The median rate map correlation across cells for a given rat was taken as the absolute value of stability and compared across conditions. A percentage stability value was also calculated by dividing all these absolute correlation values in two conditions, stable or unstable, and dividing the number of 'stable'

cells by the total number of cells and multiplying by 100. The higher this percentage was the more stable the cells were for a given rat.

A type of bootstrap analysis was used to calculate this cut-off value of stability. All cells from all rats for a specific session were loaded into a matrix and numbered, and those from another specific session were loaded into a different matrix and numbered. A random number generator selected one rate map from the first matrix and another from the second matrix, and the Pearson correlation was then calculated as described above. Overall, this compared 10,000 random rate map pairs, producing a shuffled dataset with a normal distribution. Based on a p value of 0.05, the 95% value was taken as a value of significant stability.

Whilst the percentage stability analysed levels of place field remapping, another type of remapping looked at more subtle changes in cell firing between two sessions. Instead of the place field remapping to a different preferred firing place, the place field changes the intensity of its firing for two different contexts (Leutgeb et al., 2005b). This means that place fields can be spatially stable over sessions (i.e. fire in the same location), whilst information about the environment or task is still encoded by changes in firing rate (Allen et al., 2012). Rate remapping was calculated by dividing the mean session firing rate (number of spikes/length of session) from the less active session by the more active session. Multiplied by 100, this gave a percentage of rate overlap – how much the firing rates of a cell changed from one session to the next. Only cells that fired in both sessions and were stable in location, i.e. had a higher rate map correlation than the cut-off value of stability calculated using bootstrap analysis, were included in analysis.

3.2.10.6 Cell Assembly Analysis

Population coordination (PCo) was calculated to analyse the correlation of neuronal cell assembly activity, and how this synchronous firing of cell populations changes with repeated exposures to the same environment (as described in Neymotin et al., 2017). To do so, cell assembly firing was quantified. Raster plots of individual cell firing were split into the 10-minute behaviour sessions and activity vectors were created for each rat, session and cell. Activity vectors counted the frequency of spikes within each time bin for the 10-minute session. Different time bins were used to represent different types of neuronal activity: 25ms

represented gamma oscillations and has been suggested as the optimal timeframe for analysing cell assembly activity (Harris et al., 2003); 125ms represented a band of theta oscillations (Buzsáki, 2010); and 1s represented the time taken for rats to pass through a place field in real-time (Fenton and Muller, 1998). The activity vector of each cell was then correlated with every other simultaneously recorded cell's activity vector within a session, for that rat, to create a PCorr vector for that session and rat. These values ranged between -1 and 1, with a higher PCorr indicating more coordinated firing of the recorded cells within a session (i.e. cell assemblies). Kendall's tau non-parametric rank correlation was used as this measure can handle many (0,0) pair correlations produced by sparse activity vectors when cells fire infrequently (Kendall, 1938). The equation is as follows, with pairs being concordantly paired (nc) or discordantly paired (nd):

$$\tau = \frac{n_c - n_d}{(1/2) \cdot n \cdot (n-1)},$$

PCo was calculated by correlating each PCorr vector from one session with a PCorr vector from another session on a rat-by-rat basis. This shows how the coordinated neuronal assembly firing during one session relates to another, i.e. are the same cell assemblies being activated in different sessions.

3.2.10.7 Cell-inclusion

The cell-inclusion parameters had three steps: is it a pyramidal cell; is it spatially tuned; and is it firing enough to be included in analysis. To ascertain whether it was a pyramidal cell the cell waveform spike width had to be greater than 250µs, as anything under this was deemed to be either noise or an interneuron. A visual step was also included to ensure no noise was included. Anything with a clear refractory period, and no obvious contamination from non-neuronal signals was included.

Correlations compared different sessions within the same cell. Not every session included had to reach criterion for both active and spatial property cut-offs. One of the sessions being compared had to be both spatially tuned (spatial information greater than 0.5b/s) and active (mean firing rate greater than 0.15Hz but less than 6Hz). The other session did not have to meet these criteria, as we wanted to include pyramidal cells that remapped between sessions (in which case they might only have fired sufficiently, or been spatially tuned in one of the

two sessions being compared). Either this other session, or a neighbouring session where the rat continued to be plugged in, needed to be active (FR>0.15Hz <6Hz) but not spatial (SI<0.5b/s). In the experimental designs containing sleep sessions, even if the other session or surround sessions were not active enough, sessions could still be included if the cell fired during a neighbouring sleep session (i.e. where the rat had not been unplugged between sleep sessions and pair sessions). To be seen as firing in a sleep session the FR had to be greater than 0.01Hz. When single sessions were analysed (i.e. to assess FR), cells were included if at any point they had been included in a pair-wise analysis. However, if spatial information was being analysed, cells were only included if they had been active (FR>0.15<6Hz). This was because spatial properties can be distorted when the cell is firing with very low rates.

Two output parameters were based on cluster isolation quality, using the same features to define clusters as before (waveform energy, amplitude, width, principle component, and time of peak). Isolation distance (IsoD) is how well defined each cluster is from every other cluster. If a cluster has a high isolation distance, there is little to no overlap between that cluster and all other clusters recorded. This increases the certainty that this cluster has been accurately defined as a single cell. L ratio is how well defined each cluster is from the cluster containing noise signals. A lower ratio indicates that the cluster is well separated from the noise recorded, increasing the certainty that this cluster is does not contain any noise spikes. IsoD and L ratio were used as further cell-inclusion criteria, splitting cells into 5 different categories: excellent (IsoD >20; L ratio <0.1); very good (IsoD >15; L ratio <0.5); good (IsoD >10; L ratio <2); acceptable (IsoD >7; L ratio <4); and unacceptable (IsoD <7; L ratio >4). These parameters were based on curves plotting the relationship between the Iso D and L ratio. Unacceptable cells were excluded from all analysis. When significant differences between the WB and RB conditions were obtained the analysis was repeated using the four different categories of cell quality. This ensured poorly isolated cells did not drive results.

3.2.11 Statistics

Before all statistical analysis, normality was tested using a Shapiro-Wilk normality test. If the data was not normally distributed non-parametric versions of tests were used. This included a Kruskal-Wallis test in place of a one-way ANOVA, and a Mann-Whitney U test in place of an unpaired t-test. Two-way repeated measures ANOVAs were initially used to ascertain

whether there was any difference over the different exposures to the novel environment for the entire dataset, and whether there was an interaction between the time-point and condition. Further post hoc comparisons using Bonferroni corrections were used to test for significance between the different time-points, between the conditions, or between time-points for each condition depending on the type of significant found. When comparing conditions at the level of the cell unpaired t-tests, or Mann-Whitney U tests, were used. Analysis was performed at the cell level for shuffle data analysis, as well as comparing IsoD and L ratio between the two conditions.

If the data showed trends towards significance an a priori power analysis test was performed to calculate the required number of animals with a given power of 0.5 and 0.8.

Whilst many studies have grouped together data from all animals used and analysed data at the cell level, this can lead to type I errors due to high statistical power. Therefore, all statistics were carried out at the rat level. Due to the within-subject design of the study, analysing data at the rat level also allowed me to directly compare individual rats in both conditions.

3.2.12 Perfusion and Histology

At the end of the experiment animals were anaesthetised with isofluorane gas (Abbott Laboratories, IL) until unresponsive and injected intraperitoneally with a lethal dose of a Euthatal, a sodium pentobarbital agent (0.7ml, Merial Animal Health Ltd., UK). Once blink and tail-pinch reflexes ceased, animals were perfused with 0.9% saline followed by 4% formalin. This acted to fix the tissues for further histology. Brains were removed and stored in 4% formalin for at least a week, and then flash-frozen and sliced using a cryostat-microtome to give 40μm sections. These sections were mounted onto polysine slides (Thermo Scientific, UK), stained with 0.1% Cresyl Violet to label Nissl substance in the cytoplasm of neurons and coverslipped in DPX (Sigma-Aldrich, UK). Once dry, sections were viewed under 2.5x magnification with a light microscope to identify where the electrode track had been located, and imaged using Image-Pro Plus (Media Cybernetics, USA).

3.3 Results

3.3.1 Histology

Electrode tracks were identified in all animals in the dorsal CA1 of the hippocampus. The coordinates of the electrode tracks ranged between AP: -3.36mm to -3.0mm and ML -3mm to -1.8mm, with average coordinates of AP: -3.24 ML: 2.54. Images of these tracks in each animal can be seen in Figure 3.3.
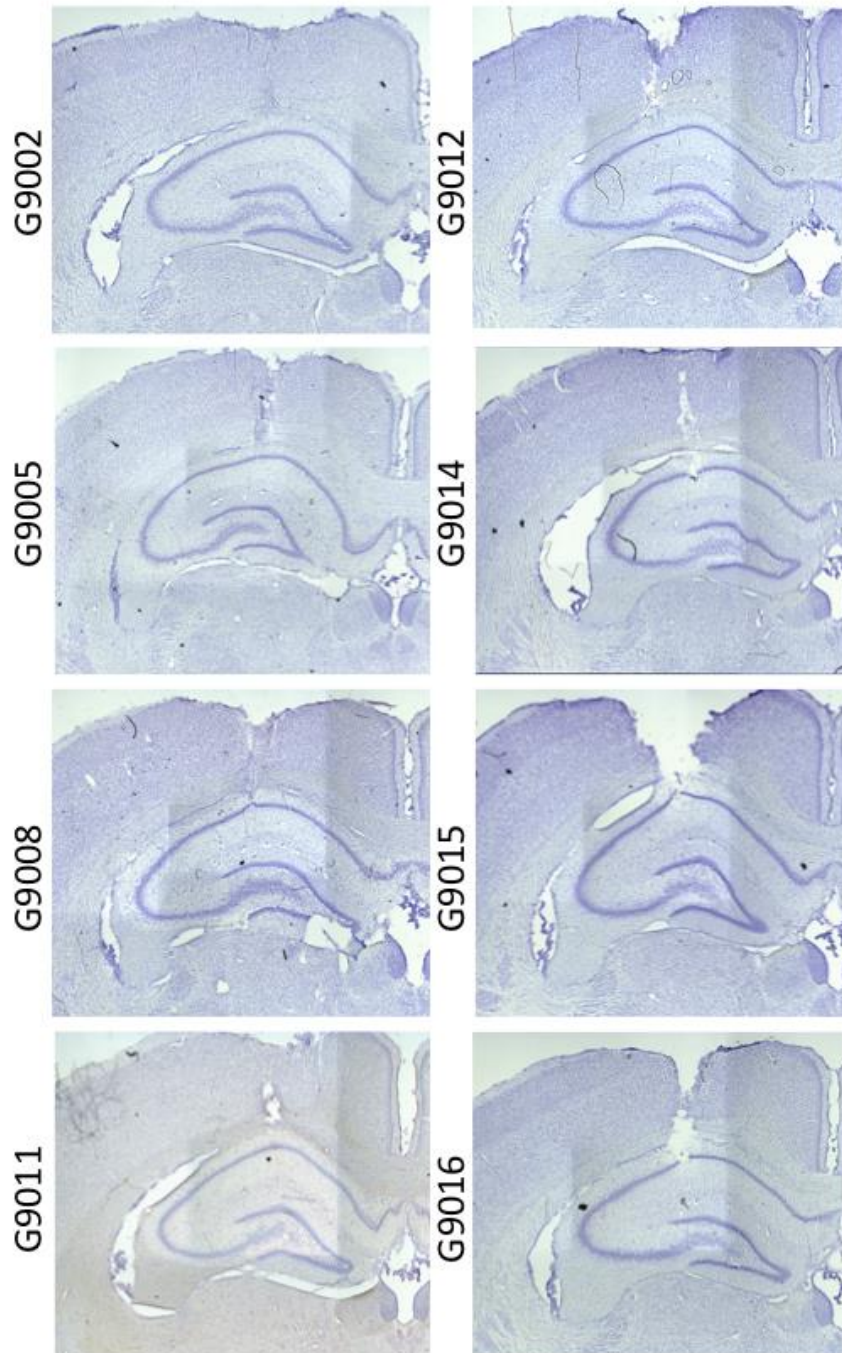
Figure 3.3: Histology showing electrode track positions for all animals used.

3.3.2 Place Cell Properties

A primary aim of this experiment was to investigate various place cell properties across the three exposures to the novel environment. The first exposure was completely novel, whilst the second and third were re-exposures to the same context. The effect of manipulating interference levels after the first exposure was explored by investigating the difference in various place cell properties between the two experimental conditions, WB vs RB. These experimental conditions refer to whether the animal was exposed to the white box (WB) or red box (RB) after the 0h recording sessions. As this was a within subjects design, all animals were tested in both conditions.

3.3.3 Place Cell Firing Properties

Given the evidence that average place cell firing is lower in a familiar environment compared to a novel environment, whereas peak in-field firing appears to be higher (Brandon et al., 2014; Karlsson and Frank, 2008) two different aspects of firing rate (FR) were analysed: average place cell firing rate and peak in-field firing rate. The average place cell firing rate was approximately 1Hz and the peak in-field firing rate of place cells was around 9Hz (Figure 3.4).

Average place cell firing rates significantly changed over the repeated exposures [main effect of time-point: $F_{(2,14)}=8.476$, $p=0.0039$]. Unexpectedly, the novel 0h exposure did not show the highest place cell firing rates, and instead post hoc comparisons using Bonferroni corrections showed that there was a significant increase between the 0h and 6h exposures ($p=0.0043$). There was however a significant decrease between the 6h and 24 exposures as expected ($p=0.0303$). The peak in-field firing rate on the other hand did not significantly change over the three repeated exposures [main effect of time-point: $F_{(2,14)}=1.266$, $p=0.3122$], in contrast to previous findings.
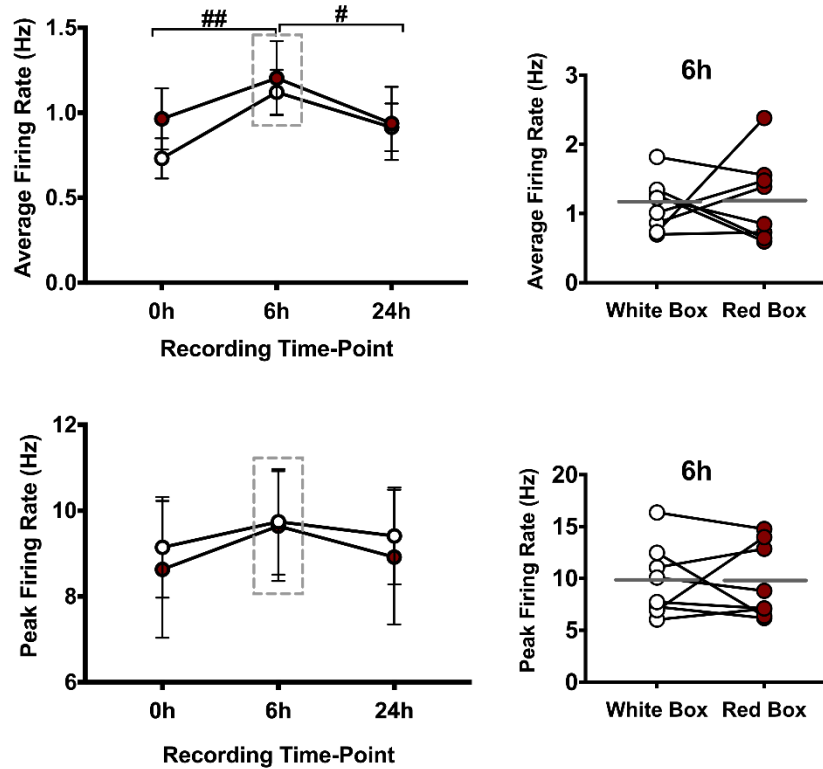
Figure 3.4: The average place cell firing rate (top left) changed significantly over repeated exposures to the same novel environment, however firing rate did not decrease over time-points as expected, instead increasing at the 6h time-point before returning to novel exposure (0h) levels at 24h. Peak in-field right rate (bottom left) in-field firing rate of place cells did not change over repeated exposures. There were no significant differences in average place cell or peak in-field firing rates between conditions. White and red dots refer to the Mean of the respective conditions with Error bars represent SEM. Mean Values for individual rats at the 6h time point are shown to the right, with grey lines representing the average firing rate for each condition. Post hoc comparisons with bonferroni corrections: # p>0.05 ## p>0.01

There were no differences in average place cell firing rate [$F_{(1,7)}=0.2225$, $p=0.6515$] or peak in-field firing rate [$F_{(1,7)}=0.07285$, $p=0.7950$] between the white and red box conditions, and no interactions between time-point or condition [average FR: $F_{(2,14)}=0.7744$, $p=0.4797$; peak in-field FR: $F_{(2,14)}=0.1248$, $p=0.8837$]. No differences between conditions were expected at the novel exposure (0h) as both had undergone exactly the same conditions by this stage. You would expect to see the most sizeable effect during the exposure directly after the white or red boxes (6h) however again this showed no difference between condition for peak FR ($t=0.07467$, $df=7$, $p=0.9426$) or average in-field FR ($t=0.1004$, $df=7$, $p=0.9229$). The two types of firing rates at 6 hours are highlighted in the two panels to the right of Figure 3.4. These results show that reducing interference between the first and second sessions does not alter the in-field firing rate or overall firing rate of place cells.

3.3.4 Spatial Properties

To determine whether the spatial firing properties of place cells were affected by either repeated exposure to a novel environment or changing interference levels, spatial information, spatial sparsity, spatial selectivity and spatial coherence were analysed.

Spatial Information

Spatial information conveys how well the information content of spikes can predict the location of the animal, with higher values carrying more information per spike. Spatial information appeared to increase by approximately 10% from the first exposure of the novel environment to the last, shown in Figure 3.5. This effect of time-point on spatial information was significant [main effect of time point: $F_{(2,14)}=8.262$, $p=0.0043$], indicating that the spatial tuning of place fields does increase over repeated exposures. Further post hoc comparisons using Bonferroni corrections revealed that there were significant differences between the 0h and 6h time-points ($p=0.0077$) and the 0h and 24h time-points ($p=0.0141$) but not the 6h to 24h time-points ($p>0.9999$). This represents an increase in spatial information from the first session to the last two sessions.
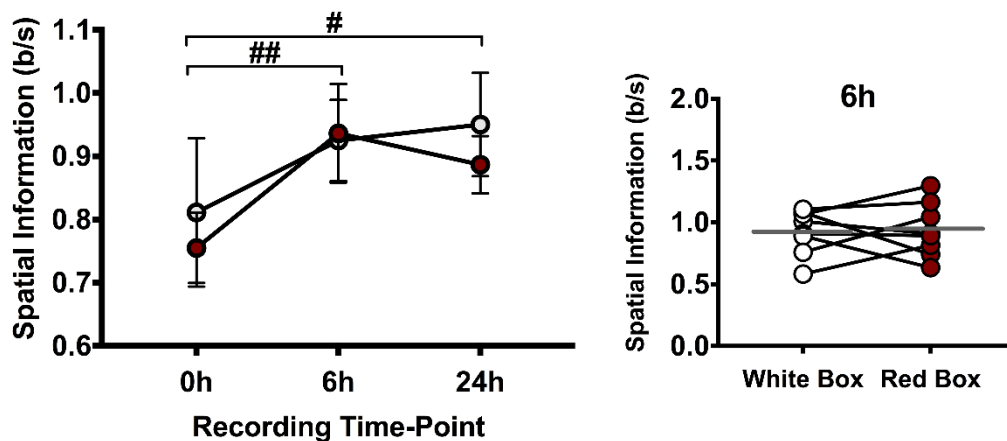
Figure 3.5: Spatial information (bits/spike) increases over repeated exposures to a novel environment, with significant differences between the first exposure and the second and third exposures. There are no differences between white and red box conditions. White and red dots refer to the Mean of the respective conditions with Error bars representing SEM. Grey lines represent mean spatial information at the 6h (right) time-points. Post hoc comparisons with bonferroni corrections: # p<0.05   ## p<0.01

There was no significant difference of spatial information between the white and red box conditions [$F_{(1,7)}=0.2585$, $p=0.6268$] and no significant interaction between time-point and condition [$F_{(2,14)}=0.2844$, $p=0.7567$]. Overall, spatial information increases over time, demonstrating a refinement in spatial representation with repeated exposures to the same environment. However, there was no effect on spatial information of reducing interference after the initial novel exposure.
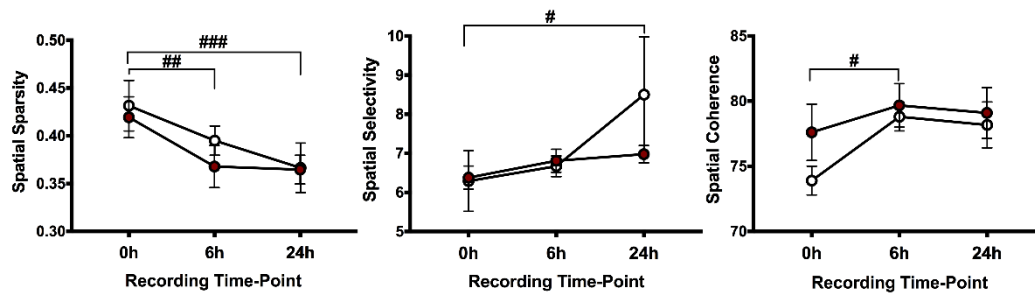
Figure 3.6: Every spatial property measured improved over repeated exposures to the novel environment, but these improvements did not differ between conditions. Spatial sparsity decreases (left) between the first exposure and second and third exposures, whereas spatial selectivity (middle) increases between the first and last exposures and coherence (right) increases between the first and second exposures. White and red dots refer to the Mean of the respective conditions with Error bars representing SEM. Post hoc comparisons with bonferroni corrections: # $p<0.05$ ## $p<0.01$ ### $p<0.001$

Spatial Sparsity

Spatial sparsity is a measure of how selectively a cell fires within an environment. The lower this measure is, the more sparsely it fires. Spatial sparsity decreased by approximately 10% over the three different exposures to novel environment, shown in Figure 3.6 (left). This change in sparsity over time was significant [main effect of time-point: $F_{(2,14)}=18.73$, $p=0.0001$], suggesting that cells fired more selectively within the environment as the environment became more familiar. Post hoc comparisons using Bonferroni corrections revealed that there were significant differences between 0h to 6h time-points ($p=0.0020$) and 0h to 24h time-points ($p=0.0001$), but not 6h to 24h time-points ($p=0.4202$). This suggests that, as with spatial information, there was a significant improvement in the spatial firing of cells between the first session and last two sessions.

There were no differences between the two conditions [$F_{(1,7)}=0.3575$, $p=0.5687$] and no interaction between time-point and condition [$F_{(2,14)}=0.3515$, $p=0.7097$]. This implies that the apparent improvement in the spatial firing of cells was not modulated by a reduction in interference after the novel exposure to the environment.

Spatial selectivity

Spatial selectivity is a measure of within field versus out-of-field firing rates. A higher level of selectivity would indicate that a place cell had limited place fields with high firing rates and an overall low level of background firing. This measure increased slowly but significantly over the three exposures [main effect of time-point: $F_{(2,14)}=5.245$, $p=0.0199$], shown in Figure 3.6 (middle). This suggested a higher signal-to-noise ratio of place field firing as the environment became more familiar, a trend seen previously in spatial tasks (Karlsson and Frank, 2008). Post hoc comparisons using Bonferroni corrections showed that there was a significant difference between 0h and 24h time-points ($p=0.0214$). However there were no differences between adjacent time-points (0h to 6h: >0.9999; 6h to 24h: $p=0.1258$), suggesting this increase in spatial selectivity requires more than two exposures to the novel environment.

Again, there was no difference between conditions [$F_{(1,7)}=0.301$, $p=0.6003$] nor interaction between time-point and condition [$F_{(2,14)}=1.494$, $p=0.2581$]. This follows the pattern of spatial properties described so far, with both conditions showing an improvement in spatial selectivity from the first, novel exposure, to the last exposure, but with no variation due to reduced interference.

Spatial Coherence

Spatial coherence is a measure of place field "smoothness". Higher percentages of coherence represent more consistent firing patterns of place fields. This measure increased significantly over the three exposures [main effect of time-point: $F_{(2,14)}=4.492$, $p=0.0311$], indicating that the consistency of place field firing also improved over repeated exposures to the environment Figure 3.6 (right). Further post hoc comparisons using Bonferroni corrections showed that there was only a significant difference between the 0h and 6h time-points ($p=0.0422$), with 6h to 24h ($p>0.9999$) and 0h to 24h ($p=0.1082$) comparisons showing no significant differences.

Again, there was no difference between conditions [$F_{(1,7)}=0.6861$), $p=0.4348$] or interaction between time-point and condition [$F_{(2,14)}=1.572$, $p=0.2421$], indicating that the reduction of interference after the first exposure did not affect the consistency of place field firing patterns.

Overall, the spatial properties of cells and place fields improved from the first novel session to the last two sessions in both conditions. Reducing interference between these sessions did not significantly impact on any spatial properties measured.

3.3.5 Place Field Properties

Exposure to environments can affect both the number and size of place fields (Fenton et al., 2008). Repeated exposures to the same environment have also been associated with decreases in place field size (Brandon et al., 2014). The number and size of place fields was therefore calculated for the RB and WB conditions over the three repeated exposures to the novel environment.
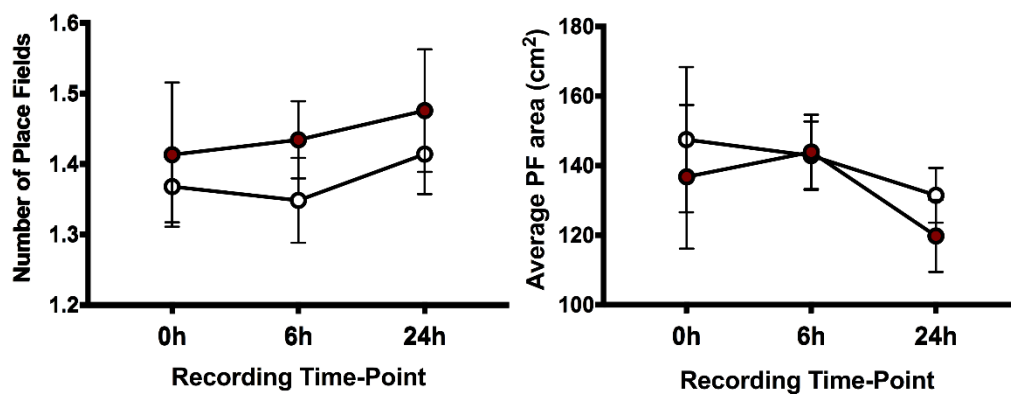
Figure 3.7: The number (left) and size (right) of place fields did not increase over repeated exposures to the novel environment or differ between conditions. White and red dots refer to the Mean of the respective conditions with Error bars representing SEM.

Number of Place Fields

The mean number of place fields per cell in each session was between 1 and 2 for both conditions, as shown in Figure 3.7 (left). This number did not significantly change between sessions [main effect of time-point: $F_{(2,14)}=0.8052$, $p=0.4667$], indicating that cells did not express different numbers of place fields as the environment became more familiar. The majority of place cells expressed only one field, with just over a quarter of cells having two or more place fields, and around half of those having three or more fields (data not shown). As with place field number, the percentage of cells with multiple fields did not change significantly over repeated exposures [>1 PF: $F_{(2,14)}=0.774$, $p=0.4799$; >2 PFs: $F_{(2,14)}=1.55$, $p=0.2465$].

There were also no differences between the WB and RB conditions for number of fields [$F_{(1,7)}=0.4531$, $p=0.5225$] or the percentage of multiple fields [>1 PF: $F_{(1,7)}=0.7494$, $p=0.4153$; >2 PFs: $F_{(1,7)}=0.1827$, $p=0.6819$] and no interactions between time-point and condition [Number of PF: $F_{(2,14)}=0.1471$, $p=0.8645$; >1 PF: $F_{(2,14)}=0.1456$, $p=0.8658$; >2 PFs: $F_{(2,14)}=0.6141$, $p=0.5551$].

Place Field Size

The average area of the place fields recorded also did not change between exposures [main effect of time-point: $F_{(2,14)}=1.747$, $p=0.2101$] suggesting that the size of place fields does not depend on how familiar the environment is. There was also no difference in place field area between the two conditions [$F_{(1,7)}=0.1735$, $p=0.6895$] or interaction between time-point and condition [$F_{(2,14)}=0.1396$, $p=0.8709$]. This is shown in Figure 3.7 (right). This implies that the level of interference after the first exposure to a novel environment does not alter the number or size of place fields expressed in subsequent exposures. Unlike the spatial properties of place cells, these results also indicate that the number and size of place fields expressed was not affected by repeated exposures to the same environment.

3.3.6 Place Field Stability

Another main aim of this experiment was to determine how well place fields remained stable over sessions. For all measures of stability analysed, the 0h to 6h and 6h to 24h and 0h to 24h time-points were compared. It was hypothesised that the RB condition could enhance place field stability between the 0h and 6h time-points, as this condition was shown to enhance spatial memory at 6h. Enhancing stability between the 6h and 24h time-points would depend on how many exposures the animal would need to produce stable place fields. Enhancing stability between 0h and 24h would depend upon the consistency of the stability of place fields over all three sessions, as low levels of remapping between the first and second, and the second and third exposures could appear as much higher levels of remapping between the first and the last exposure. To test the stability of place fields over these exposures three measures were analysed: median correlation of rate-maps; the percentage of stable cells; and rate-remapping

100
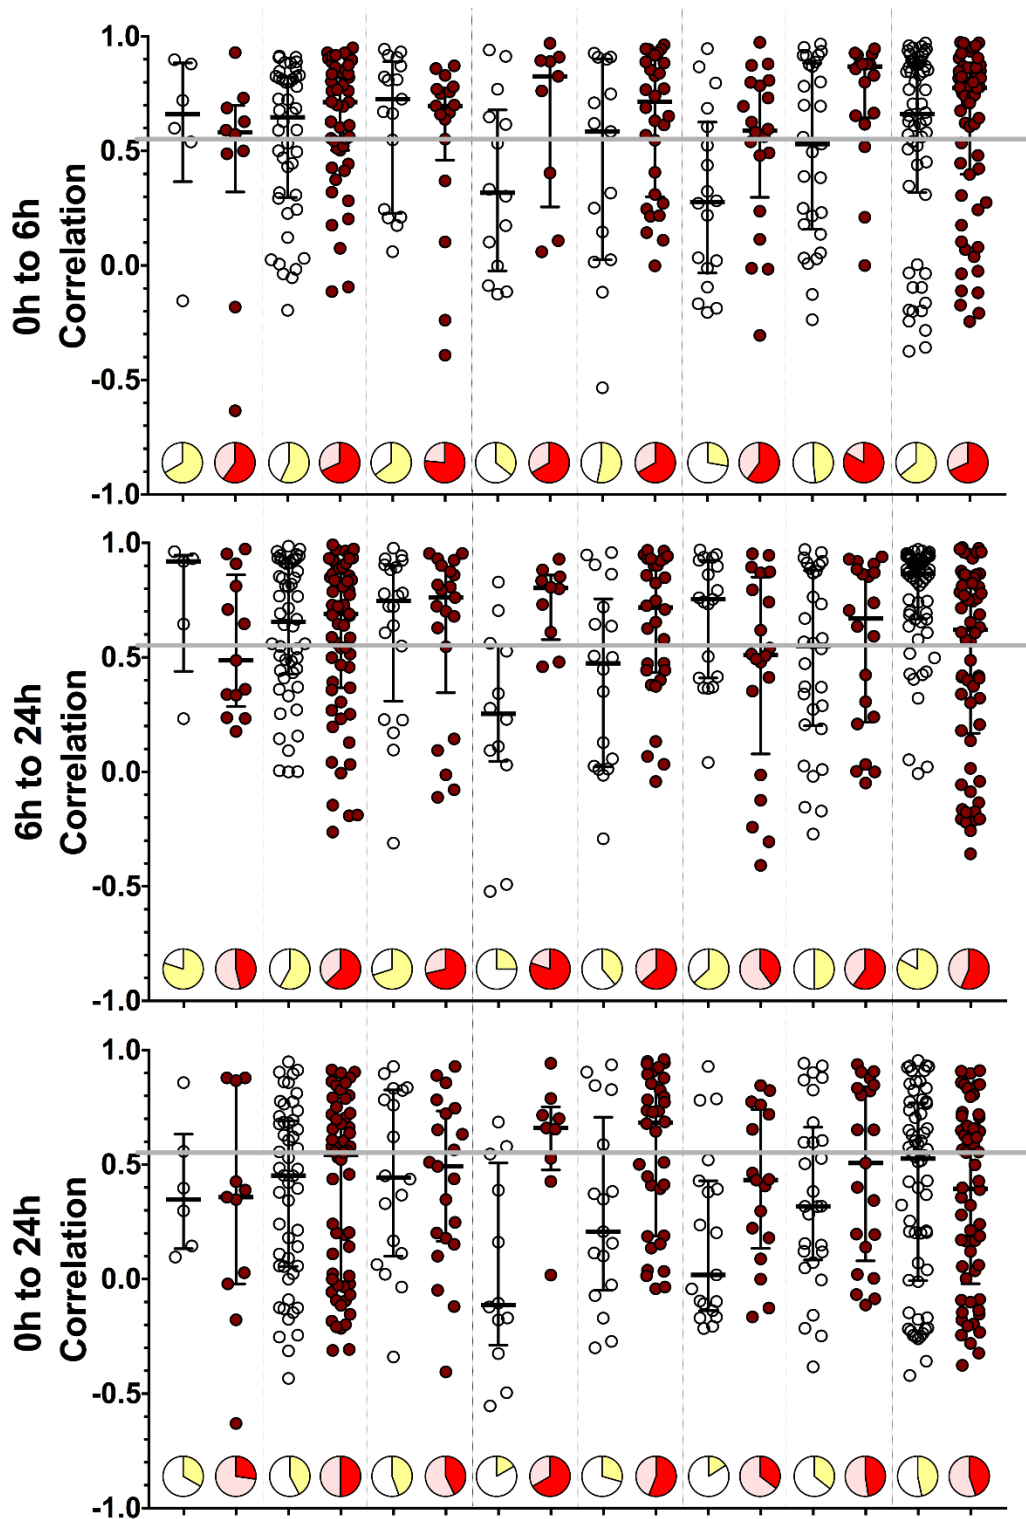
Figure 3.8: Pearson correlation values for individual cells are plotted for correlations between 0h and 6h (top), 6h and 24h (middle) and 0h and 24h (bottom). Overlaid box and whisker plots represent the median correlation values and interquartile ranges for each rat in each white box or red box condition. Correlations of 1 represent highly similar rate maps with stable place fields. This place field stability is also shown as a percentage of stable cells in pie charts. Every pie chart represents the percentage of stable (dark colours) and unstable (light colours) cells for each rat in each white box or red box condition. This measure of stability was calculated using a bootstrap cut-off, shown as a grey line intersecting the correlation values at 0.55 correlation. All values above this cut-off were 'stable' and all values below were 'unstable'.

Median Correlation

The first measure of stability analysed was the median correlation of place field rate maps. If the two rate maps were similar, this implies that the field-specific spatial firing of the place cell was maintained between the two sessions. This is known to be an NMDA receptor and protein synthesis-dependent phenomenon, suggesting that a higher correlation pertains to the better expression of spatial memory. All of the correlation values for the individual cells recorded are plotted in Figure 3.8, split into rats and conditions with the median correlation values overlaid as box and whisker plots.
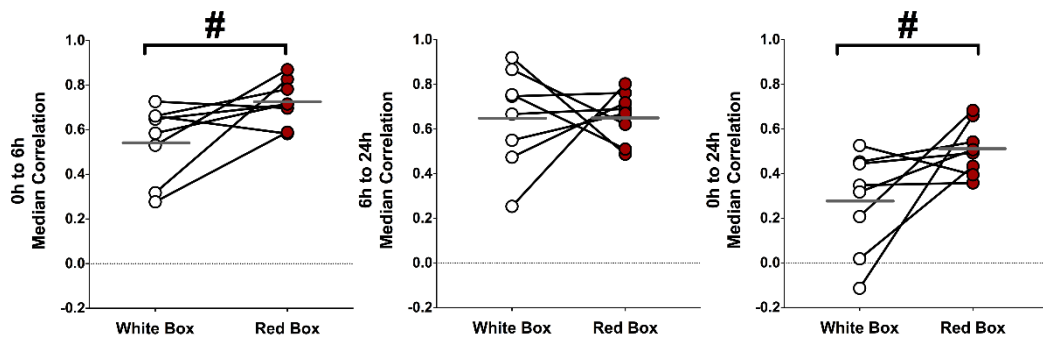
Figure 3.9: The median correlation of place cell rate maps was significantly higher in the red box condition compared to the white box condition when comparing the 0h and 6h exposures (left). There are no differences in correlation when comparing the 6h and 24h exposures (middle). The median correlation was also significantly higher in the red box condition when comparing the 0h and 24h exposures (right). Grey lines represent average values. Post hoc comparisons with bonferroni corrections: # p<0.05

As has been shown previously by a number of studies, median correlations changed greatly over the three exposures [main effect of time-point comparison: $F_{(2,14)}=22.54$, $p<0.0001$]. Further tests also revealed a significant interaction between time-point comparison and condition [$F_{(2,14)}=4.512$, $p=0.0307$]. The WB condition showed a slight gradual increase in correlation over the three exposures, as the correlation from 6h to 24h was minimally higher than the correlation from 0h to 6h, however this was not significant (post hoc comparisons with bonferroni corrections: $p=0.2602$). Whilst rate map correlations between adjacent sessions were comparably high in the WB condition, rate map correlations between the first and the last session were significantly lower (0h to 6h and 0h to 24h comparisons: $p=0.0007$; 6h to 24h and 0h to 24h comparisons: $p<0.0001$). Contrastingly the RB condition showed a much higher correlation between 0h and 6h with which then decreased minimally over time to stability levels seen in the WB condition between 6h and 24h. Rate map correlations

between the 0h and the 24h session were significantly lower than correlations of the 0h and 6h sessions (post hoc comparisons with bonferroni corrections: p=0.0058), indicating that correlations decreased over time in both conditions.

As shown in Figure 3.9, the average median correlation of the white box condition between 0h and 6h was 0.55, compared to a much higher correlation of 0.72 for the red box condition. These correlations were significantly higher in the red box condition compared to that of the white (post hoc comparisons with bonferroni corrections: p=0.0256), demonstrating that reducing interference after the initial exposure to a novel environment led to enhanced rate map correlation. In contrast, the correlation values between the 6h and 24h time points were almost identical for the white box (0.65), and red box (0.66) conditions (post hoc comparisons with bonferroni corrections: p>0.9999). This suggests that whilst one 10 minute exposure is not enough to produce place field stability at 6h if the animal is exposed to the WB after the initial exposure, two 10 minute exposures (0h and 6h) are enough to produce place field stability at 24h. Interestingly, when comparing the first and last exposures (0h to 24h) the correlation values were 0.28 for the WB condition and 0.51 for the RB condition. This increase in correlation in the RB condition was also significant (post hoc comparisons with bonferroni corrections: p=0.0028). As the RB condition showed high levels of stability between 6h and 24h and 0h and 24h, this suggests that the place fields fired consistently throughout all three exposures. Higher values of correlation are thought to represent increased place field stability and spatial memory. This suggests that reducing interference after spatial learning enhances spatial memory by increasing place field stability. In contrast, the WB condition showed high levels of stability between 6h and 24h, but not between 0h and 24h. This indicates that whilst place fields were stabilised after the second exposure, this pattern of firing was very different from the initial exposure. This suggests that not only does reducing interference enhance stability between the first and second exposures, the map remains stable over a 24h period. It also implies that if the map is unstable between the first and second exposures at least one more exposure is needed to achieve that stability.

As shown in Figure 3.8, the correlation values of individual cells appear to be split into two clusters, within each rat and condition. It would seem that these two clusters represent stable and unstable cells. Bootstrap analysis was used to produce a stable cell median correlation cut-off, statistically splitting the cells into these two "stable" and "unstable". This value can be seen as the dividing line in Figure 3.8.

Figure 3.10 is plotting the correlation data for the 0h to 6h correlations in a different way by grouping all the cells from all the rats in the white box condition together and all the cells from all the rats in the red box condition together to show the distribution of these correlations. This highlights how both the white box and red box conditions have distributions that are shifted drastically to the right compared to the shuffled data set at this time-point. The average correlation of the shuffled data was 0.01 - significantly lower than that of the white box (correlation 0.50, U=117908, p=<0.0001) and red box (correlation 0.61, U=76695, p=<0.0001) condition cells. Consistent with the "rat-level" analysis described above (Figure 3.10), the average correlation was significantly increased in the red box condition compared to that of the white, this time at a cell level (U=18257, p=0.0123). The 95th percentile value of the shuffled data represents the significant bootstrap value (p=0.05), in this case 0.55.

Figure 3.11 shows the difference between shuffled and actual correlations at the 6h to 24h time-point. Again the shuffled data correlation (0.1) was significantly lower than that of the white box (correlation 0.61, U=209828, p<0.0001) and red box (correlation 0.56, U=313010, p<0.0001) condition cells. Mirroring the previous results which analysed correlations on a rat by rat basis, there was no difference between the white box and red box conditions at this time-point (U=22812, p=0.1240). The bootstrap value for this data set was also 0.55.
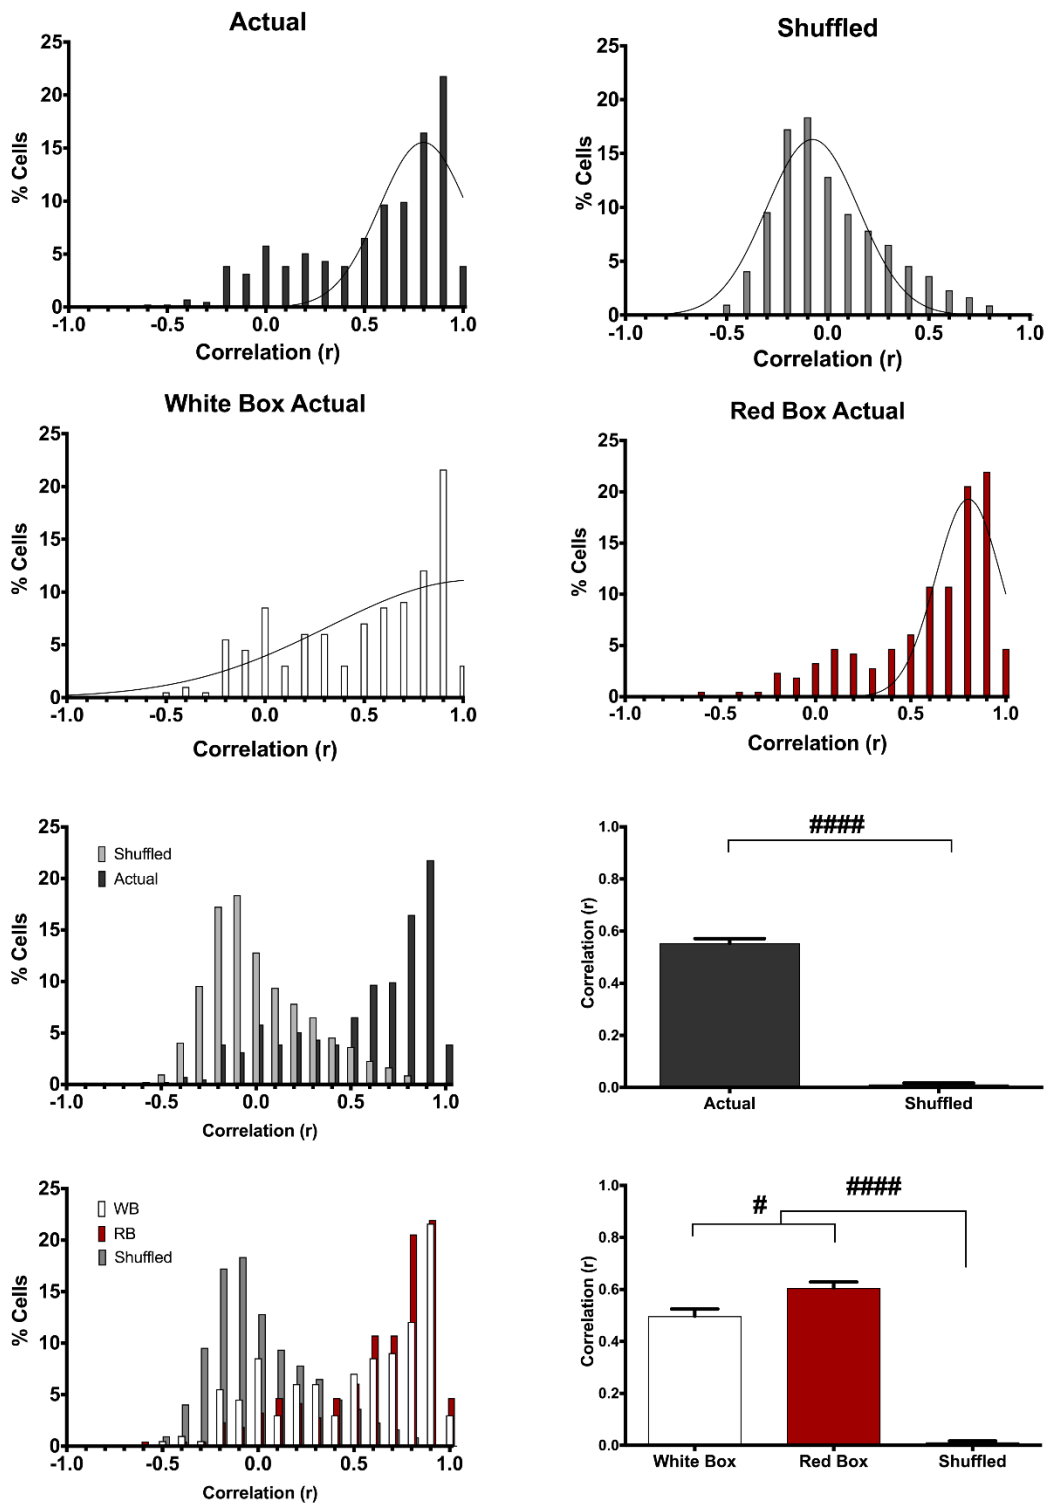
Figure 3.10: These values represent correlations of the 0h and 6h exposures. Frequency distributions of correlation values from cells from the white box condition and red box condition were shifted to the right, with cells having significantly higher correlation values than those of the shuffled data set. There was also a significant difference between correlations of cells from the white box and red box conditions. Error bars represent the SEM. Mann-Whitney U test: # p<0.05   #### p<0.0001
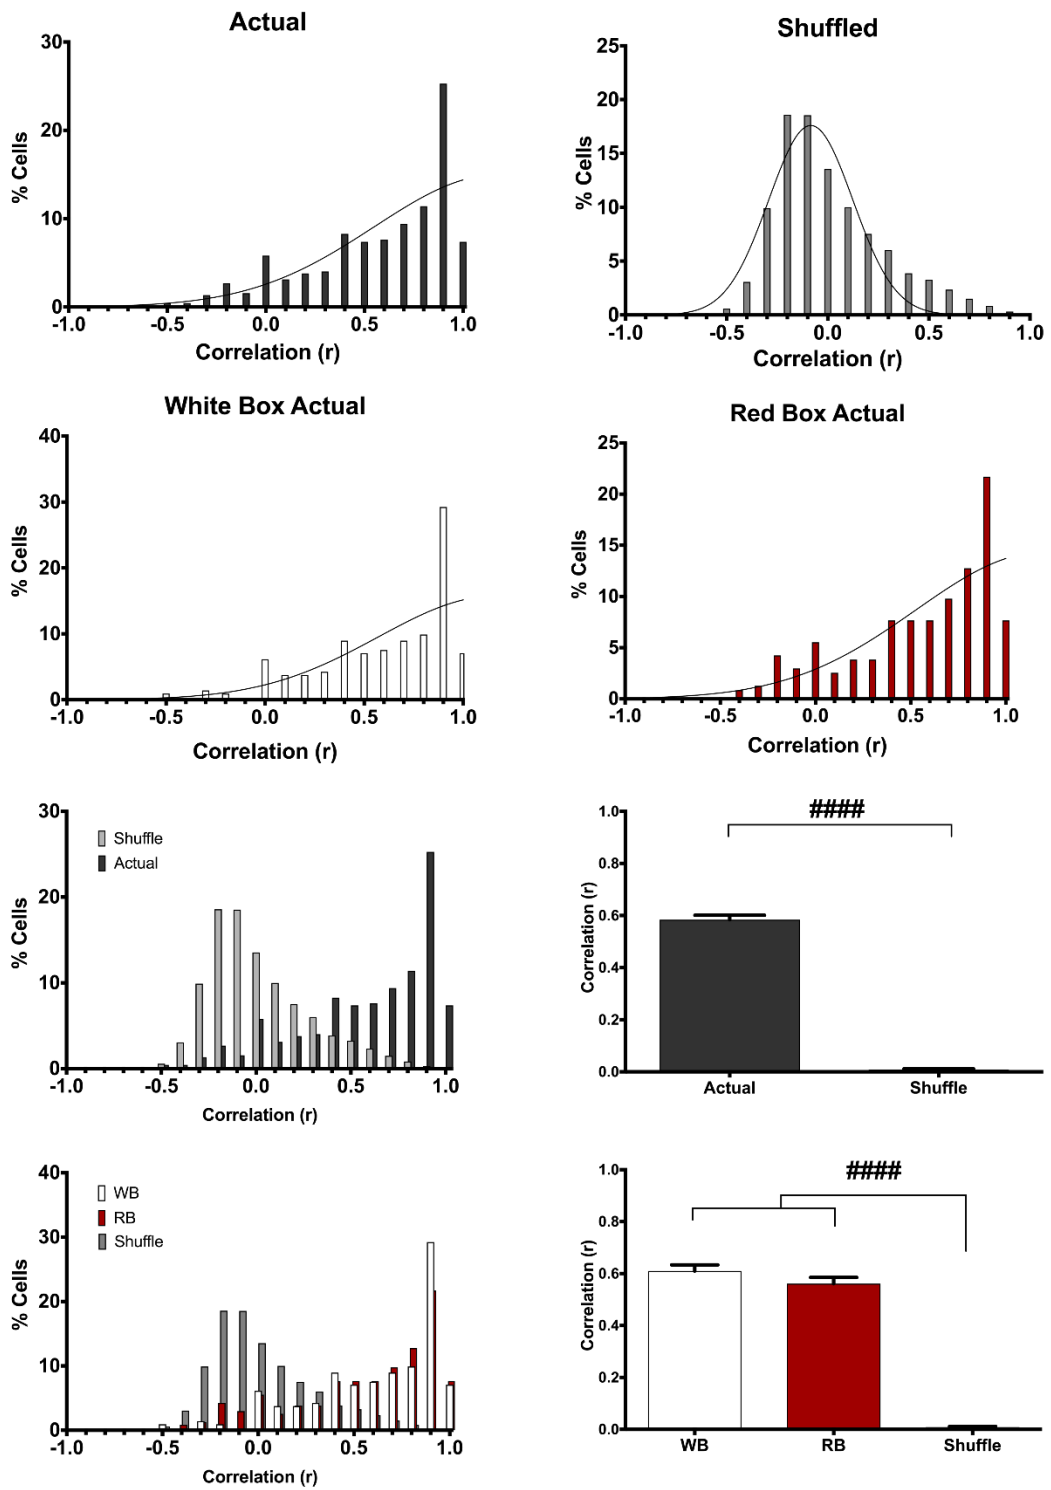
Figure 3.11: These values represent correlations of the 6h and 24h exposures. Frequency distributions of correlation values from cells from the white box condition and red box condition were shifted to the right, with cells having significantly higher correlation values than those of the shuffled data set. There was no difference between correlations of cells from the white box and red box conditions. Error bars represent the SEM. Mann-Whitney U test: #### p<0.0001
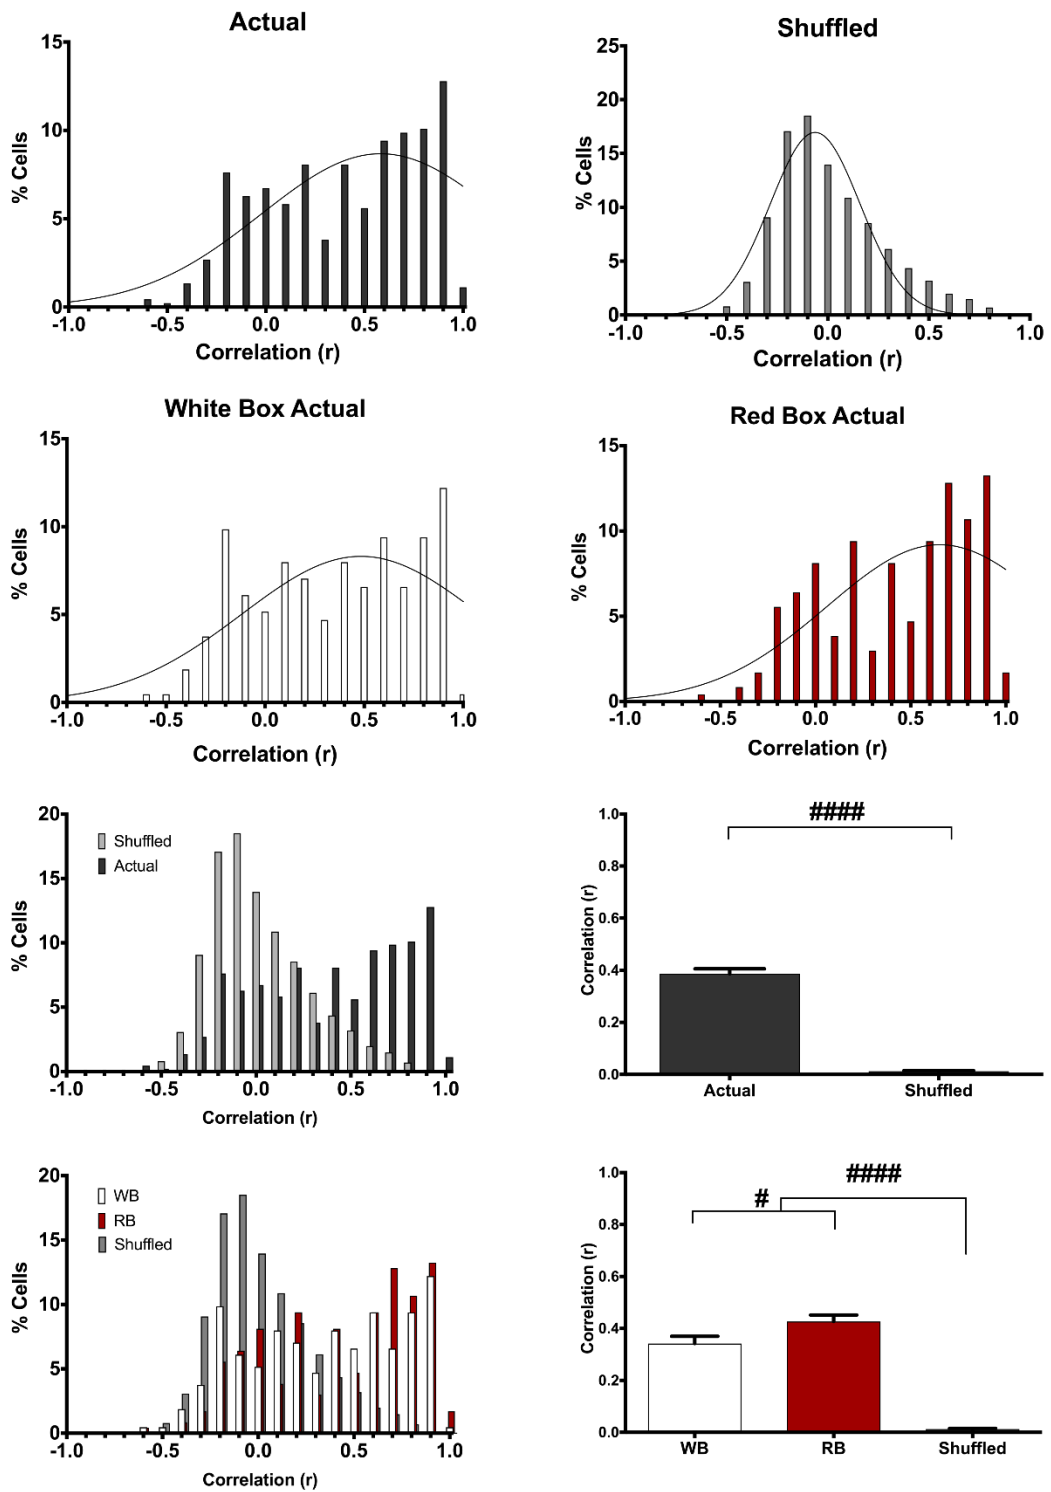
Figure 3.12: These values represent correlations of the 0h and 24h exposures. Frequency distributions of correlation values from cells from the white box condition and red box condition were shifted to the right, with cells having significantly higher correlation values than those of the shuffled data set. There was also a significant difference between correlations of cells from the white box and red box conditions. Error bars represent the SEM. Mann-Whitney U test: # p<0.05   #### p<0.0001

The distribution of median correlations between the first and last exposures are shown in Figure 3.12, again showing that the shuffled data correlation (0.1) was significantly lower than that of the white box (correlation: 0.34, U=565718, p<0.0001) and red box (correlation: 0.43, U=465940, p<0.0001) condition cells. Consistent with the previous rat-by-rat analysis, the average correlation of rate maps at the cell level was significantly increased in the red box condition compared to the white box condition (U=21874, p=0.0255). The bootstrap value for this data set was slightly lower, at 0.53. As the bootstrap values for each of the three time-point comparisons were minimally different (0h to 6h: 0.55; 6h to 24h: 0.55; 0h to 24h: 0.53) the overall bootstrap value of 0.55 was used for analysis. However it should be noted that using different bootstrap values for the different time-point comparisons did not change the outcome of the results.

Percentage Stability

Based on the bootstrap values for stable vs unstable place cells calculated in the analyses described above, the correlation values for all cells were split into stable (>0.55) and unstable (<0.55), producing a percentage of spatially stable cells for each rat and condition. These are shown in Figure 3.8 as pie charts, and summarised in Figure 3.9. Examples of the rate maps of stable and unstable cells are shown in Figure 3.15, highlighting the differences in stability seen between the conditions for the 0h to 6h comparison.
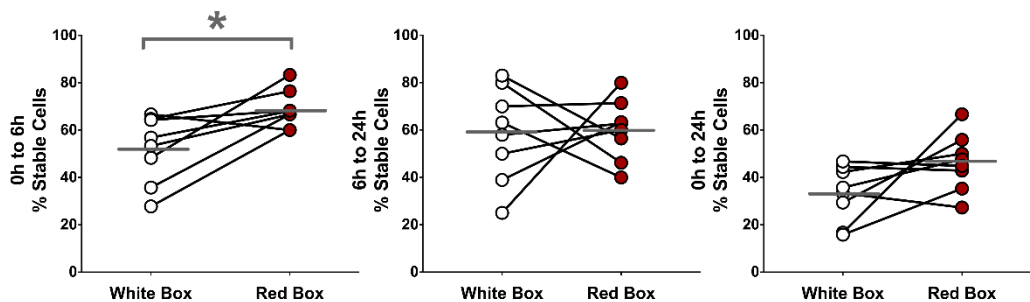
Figure 3.13: The percentage of stable cells was significantly higher in the red box condition compared to the white box condition when comparing the 0h and 6h exposures (left). There are no differences in stability when comparing the 6h and 24h exposures (middle). Although not significant, there are trends towards increases in percentage stability in the red box condition when comparing the 0h and 24h exposures (right). Grey lines represent average values. Paired t-test: * p<0.05

The percentage of cells with stable place fields significantly changed over the three exposures to the novel environment [main effect of time-point comparison: $F_{(2,14)}=38.57$, $p<0.0001$]. Although there was no main effect of condition [$F_{(1,7)}=2.446$, $p=0.1618$] and the interaction

between time-point comparison and condition was not significant [$F_{(2,14)}=2.631$, $p=0.1071$], the same trends in stability were seen as with median correlation. Cells were least stable between the 0h and 24h exposures for both conditions (post hoc comparisons with bonferroni corrections: 0h to 6h and 0h to 24h $p<0.0001$; 0h to 24h and 6h to 24h $p<0.0001$).

Although there were no interactions between time-point comparison and condition, further t-tests did indicate that some differences might be apparent between the two groups. As shown in Figure 3.9, when the 0h and 6h time-points were compared, the red box condition exhibited a significantly higher percentage of stable cells compared to that of the white box condition ($t=3.13$, $df=7$, $p=0.0166$). On average, 52.2% of cells were stable in the white box condition, compared to 68.7% in the red box condition. These results using a different measure of place field stability mirror our previous findings and show that exposure to the red box can increase the proportion of place cells that are stable between the first and second exposures to a novel environment.

As found with the median correlation of cells, the percentage stability between the 6h and 24h time-points was not significantly different between conditions ($t=0.1433$, $df=7$, $p=0.8901$), with almost identical averages of stable cells: 58.5% in the white box condition compared to 60.0% in the red box condition. When the 0h and 24h time-points were compared the red box condition had a higher percentage of cells with stable place fields (46.3%) than the white box condition (33.1%), although again this was only a trend and not significant ($t=2.015$, $df=7$, $p=0.0838$), with further power analyses indicating 2 more rats (power=0.5, n=10) or 10 more rats (power=0.8, n=18) being needed to show significance. This pattern of results again suggests that stability of place fields is more consistent over time in the red box condition.

Rate Remapping

The final measure of place field stability analysed was rate remapping. Rate remapping is where the location of place field firing remains stable (i.e. the place field rate maps are well correlated) but the firing rate of place fields changes over different exposures to an environment. This is thought to give the animal information about subtle changes in the environment that would not necessarily lead to global remapping of place fields. Rate remapping was calculated by dividing the mean session firing rate (number of spikes/length of session) from the less active session by the more active session. Multiplied by 100, this gave a percentage of rate overlap – how much the firing rates of a cell changed from one session to the next. This analysis was only completed if the cell fired in both sessions and if the correlation between the two sessions was over 0.55. This ensured the location of the place field remained stable.
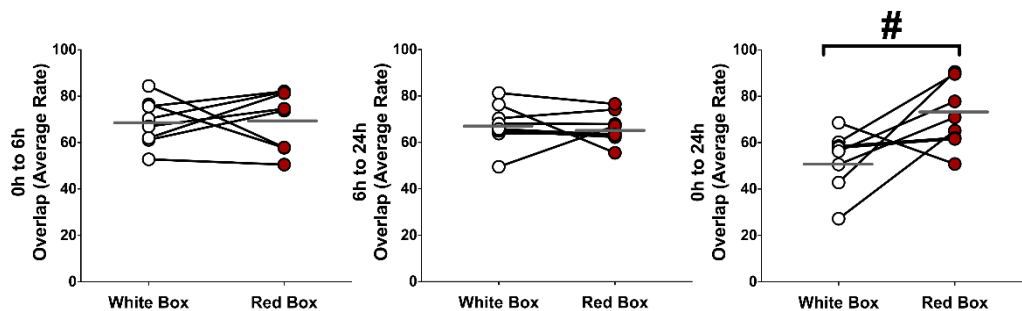
Figure 3.14: There were significantly higher levels of rate remapping between the 0h and 24h exposures (right) in the white box condition compared to the red box condition. Rate remapping between 0h to 6h (left) and 6h to 24h (middle) did not differ between conditions. Grey lines represent the average rate of remapping. Post hoc comparisons with bonferroni corrections: # $p<0.05$

When conditions were analysed together there was no significant effect of time-point comparison on the levels of rate remapping calculated [main effect of time-point comparison: $F_{(2,14)}=1.647$, $p=0.2278$]. This suggests that repeated exposures to a novel environment did not affect this measure of stability. However, although there was no significant effect of condition on rate remapping [$F_{(1,7)}=3.042$, $p=0.1246$], there was an almost significant interaction between time-point comparison and condition [$F_{(2,14)}=3.475$, $p=0.0595$]. Post hoc comparisons with bonferroni corrections showed that whilst the red box condition showed no significant differences between any time-point comparisons ($p>0.9999$), the white box condition showed significant differences between the 0h to 6h and

0h to 24h time-point comparisons (p=0.0423). As cells expressing fields with stable locations showed on average 68.73% and 67.53% overlap in firing rates between 0h to 6h and 6h to 24h comparisons, respectively, the decrease in overlap seen at 0h to 24h suggests that cells are variably stable in firing rate over the exposures. This is also apparent when comparing the conditions, as only the 0h to 24h time-point comparison was significantly different between the white box and red box conditions (post hoc comparisons with bonferroni corrections: p=0.0188), shown in Figure 3.14. It should be noted that there were no differences in average place cell firing rate between the two groups, implying that these differences are not due to overall changes in place cell firing rate.

Overall these results suggest that the cells recorded in the red box condition express fields that are more stable in location and in firing rate throughout the three exposures to the novel environment. Significantly fewer cells recorded in the white box condition express fields that are stable in location, and those that are stable in this way show significant increases in rate remapping between the first and last exposures to the environment. This indicates that even the "stable" population of cells is more unstable when higher levels of visual interference follow spatial learning.

Figure 3.15: Next page. Representative place cell rate maps showing varying levels of stability between the 0h and 6h exposures. Cells at the top show the highest correlations, whereas cells at the bottom have the lowest correlations. These cells are split into stable and remapped. G number (eg. G9002) represents different rats used. Each rate map shows the area of the environment explored by the animal within a single exposure. This is why some are circular and some are square. The firing rate above each rate map represents the peak in field firing rate of the rate map, i.e. the red part of the rate map. The blue part of the rate map represents no or minimal amounts of place cell firing in those locations in the environment.

Representative cells of White Box

52% Cells Stable

	16 Hz	29 Hz
G9012 IsoD: 89 SI 1.32 b/s **Corr: 0.93**		

	15 Hz	16 Hz
G9015 IsoD: 14 SI 1.27 b/s **Corr: 0.88**		

	9 Hz	15 Hz
G9005 IsoD: 20 SI 1.14 b/s **Corr: 0.79**		

	11 Hz	17 Hz
G9011 IsoD: 27 SI 1.27 b/s **Corr: 0.65**		

	2 Hz	9Hz
G9016 IsoD: 16 SI 1.15 b/s **Corr: 0.57**		

48% Cells Remapped

	16 Hz	15 Hz
G9016 IsoD: 22 SI 0.89 b/s **Corr: 0.44**		

	28 Hz	22 Hz
G9015 IsoD: 25 SI 1.28 b/s **Corr: 0.25**		

	3 Hz	11 Hz
G9008 IsoD: 23 SI 0.88 b/s **Corr:0.17**		

	17 Hz	20 Hz
G9014 IsoD: 11 SI 0.58 b/s **Corr: 0.02**		

	6 Hz	10 Hz
G9002 IsoD: 40 SI 0.99 b/s **Corr: -0.15**		

Representative cells of Red Box

68% Cells Stable

	21 Hz	23 Hz
G9005 IsoD: 20 SI 1.31 b/s **Corr: 0.93**		

	12 Hz	20 Hz
G9014 IsoD: 20 SI 0.90 b/s **Corr: 0.87**		

	20 Hz	21 Hz
G9016 IsoD: 48 SI 1.77 b/s **Corr: 0.81**		

	29 Hz	22 Hz
G9002 IsoD: 16 SI 1.94 b/s **Corr: 0.73**		

	11 Hz	11 Hz
G9016 IsoD: 22 SI 1.40 b/s **Corr: 0.71**		

	14 Hz	16 Hz
G9008 IsoD: 16 SI 0.79 b/s **Corr: 0.66**		

	7 Hz	7 Hz
G9014 IsoD: 18 SI 0.97 b/s **Corr: 0.56**		

32% Cells Remapped

	16 Hz	8 Hz
G9005 IsoD: 19 SI 0.67 b/s **Corr: 0.43**		

	9 Hz	7 Hz
G9016 IsoD: 49 SI 0.69 b/s **Corr: 0.18**		

	15 Hz	16 Hz
G9002 IsoD: 33 SI 2.60 b/s **Corr: -0.18**		

112

3.3.7 Poor cluster quality cannot account for the enhanced place field stability following exposure to the red box

To ensure both the enhanced rate map stability between the 0h and 6h time points and the decrease in rate remapping between the 0h and 24h time points in the red box condition were not being driven by poor isolation of clusters, four different isolation distances and l ratios were used of varying quality: acceptable (IsoD >7; L ratio <4); good (IsoD >10; L ratio <2); very good (IsoD >15; L ratio <0.5); and excellent (IsoD >20; L ratio <0.1). These cluster qualities were chosen by visually dissecting points on a scatter plot of isolation distance and l ratio, shown in Figure 3.16. There were no differences in average isolation distance (U=23484, p=0.2214) or l ratio (U=23287, p=0.1720) between the white and red box conditions, shown in Figure 3.15. There were also no differences in the percentages of cells included in each cluster quality between conditions (t=2.747e-10, df=4, p=>0.9999). The exact number of cells included in each band are shown below in Table 3.1. These four different cluster quality bands were used to include or exclude cells used in the analysis of both median correlation and percentage stability. This was implemented for the 0h to 6h time-points for rate map stability and the 0h to 24h time-points for rate mapping, as this is where significant changes were clearly seen.

	White Box Condition	Red Box Condition
Unacceptable (IsoD <7; L ratio >4)	216 (17)	233 (19)
Acceptable (IsoD >7; L ratio <4)	199 (36)	214 (32)
Good (IsoD >10; L ratio <2)	163 (79)	182 (78)
Very Good (IsoD >15; L ratio <0.5)	84 (35)	104 (49)
Excellent (IsoD >20; L ratio <0.1)	49 (49)	55 (55)

Table 3.1: Number of cells included in each band of cluster quality (exact number in individual band).
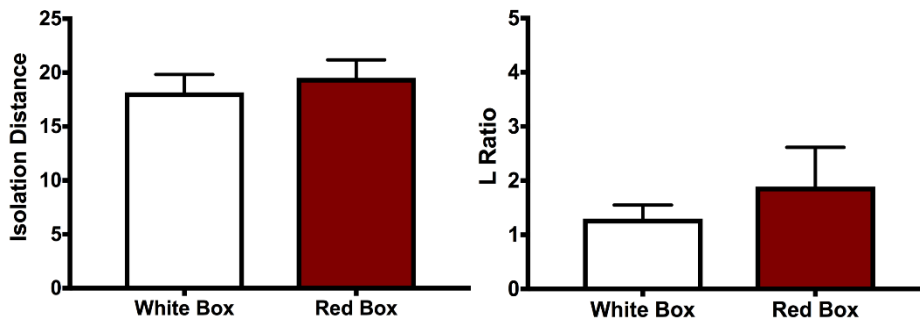
Figure 3.15: Both of the qualities of cluster quality used (Isolation Distance (left) and L Ratio (right)) showed no significant differences between the white and red box conditions. Error bars represent SEM
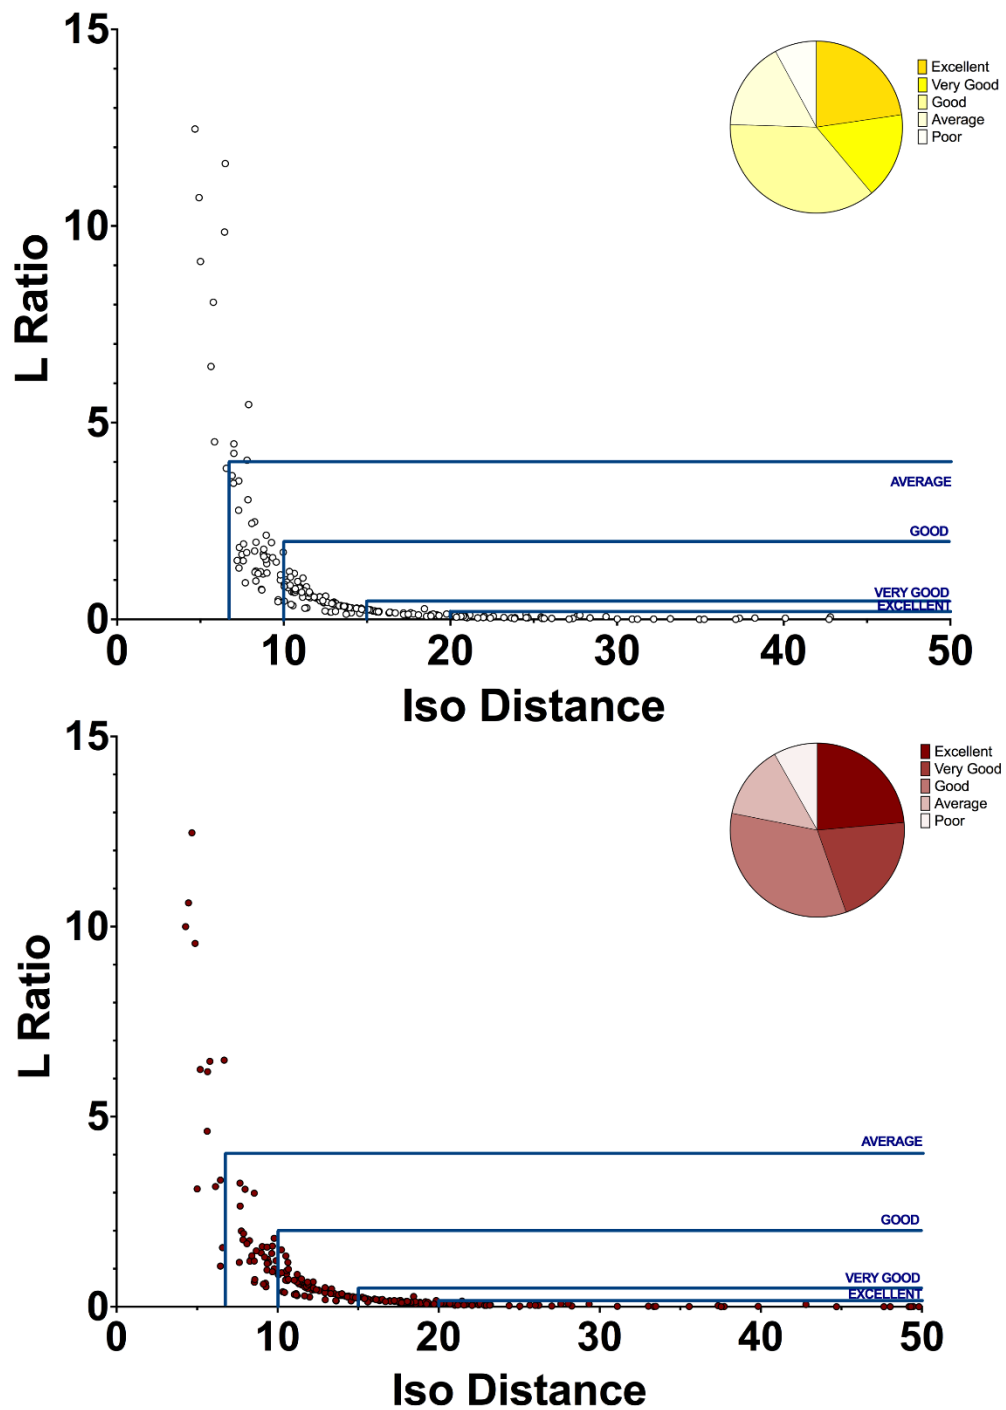
Figure 3.16: Scatter plots showing isolation distance plotted against l ratio for both the white box condition (top) and the red box condition (bottom). The blue dissecting lines represent the four different cluster quality cut-offs chosen: average (Iso D>7, l ratio <4), good (Iso D>10, l ratio <2), very good (Iso D>15, l ratio <0.5) and excellent (Iso D>20, l ratio <0.1). Anything that was below the average cut-off was excluded from analysis. The pie charts show the proportion of cell qualities for each condition.

As shown in Figure 3.17, when all cells that were rated as having an acceptable cluster quality or better were included in the analyses, there was a significant increase in median correlation in the red box condition compared to that of the white (t=3.13, df=7, p=0.0166). However, when the cluster quality criteria were more stringent this enhancement disappeared, with minimum good quality (t=2.216, df=7, p=0.0623), minimum very good quality (t=1.812, df=7, p=0.1129), and minimum excellent quality (t=0.5218, df=7, p=0.6179) showing no significant difference between conditions. The inclusion of good quality clusters showed trends towards significance, and a power analysis showed that either 1 more rat (power=0.5 n=9) or 7 more rats (power=0.8, n=15) or would be needed to achieve significance. This suggests that reducing the number of cells used within the analysis reduced the overall power, even though the statistical analyses were performed at the level of individual rats (rather than cells). However, power analysis shows that 4 (power=0.5, n=12) or 14 (power=0.8, n=22) more rats would be needed to achieve significance for the very good cluster and better quality band; and 108 (power=0.5, n=115) or 225 (power=0.8, n=233) more rats would be needed to achieve significance for the excellent cluster quality band. These are very unrealistic numbers of animals, implying that inclusion of cells with lower cluster quality could be driving the significant results shown.

Interestingly, two very different trends emerge when the data are split into two conditions representing the different experimental protocols used. Described fully in the methods section (Figure 3.2), protocols 1 and 2 (n=3) did not leave any time between the first 24h session and the second 0h session. This also meant that minimal time was given between exposure to the white and red boxes. Protocol 3 (n=5), on the other hand, left at least 5 days between the first and second contexts. More time was therefore left between white and red box exposure as well. Protocols 1 and 2 are subsequently referred to as the old protocols, and protocol 3 is referred to as the new protocol. As shown in Figure 3.17, animals from the old protocols showed no significant difference of correlation at any band of cluster quality, including when only acceptable cells were excluded (acceptable: t=0.3357, df=2, p=0.7690; good: t=0.6533, df=2, p=0.5806; very good: t=1.601, df=2, p=0.2505; excellent: t=0.8399, df=2, p=0.4894). In contrast, animals tested using the new protocol exhibited a consistent significant enhancement in correlation in the red box condition compared to the white box condition. This was true for acceptable (t=3.905, df=4, p=0.0175), good (t=3.447, df=4, p=0.0261), very good (t=4.636, df=4, p=0.0098) and excellent (t=3.202, df=4, p=0.0328)

cluster quality bands. This is what would be expected if cluster quality was not driving the enhancement of rate map correlation.
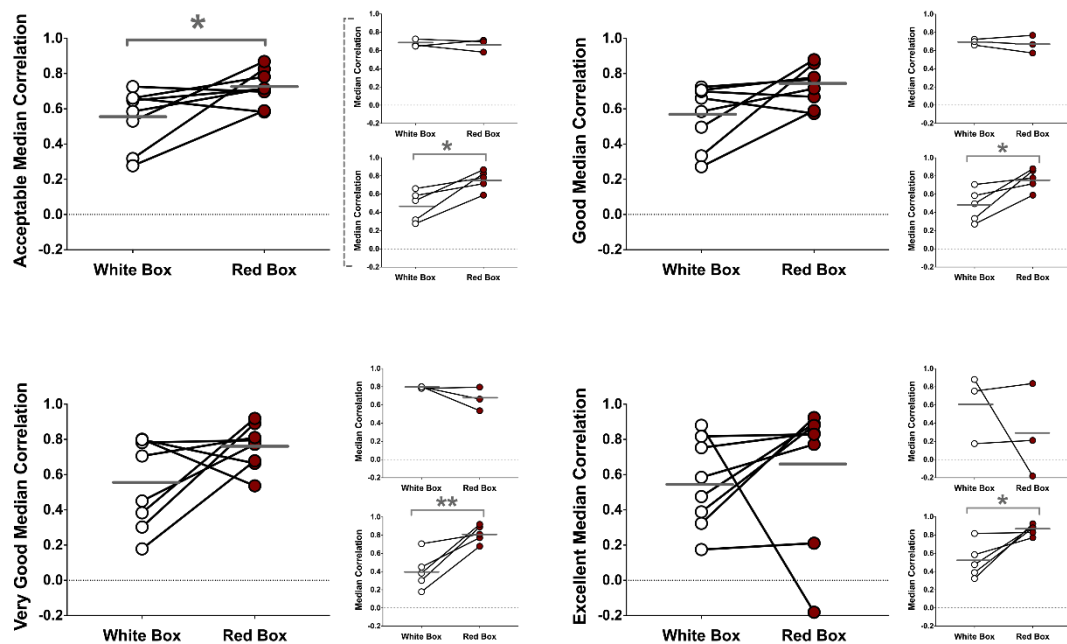
Figure 3.17: The median correlation between 0h and 6h exposures was only significantly different when the lowest band of cluster quality was used (top left). However when animals were split into the two types of protocol used – old and new – animals on the new protocol showed significant differences between red and white box conditions for all bands of cluster quality. Animals on the old protocol showed no differences between conditions regardless of cluster quality band. For each band of cluster quality, the small top graph represents the animals on the old protocol (n=3) and the small bottom graph represents animals on the new protocol (n=5). Grey lines represent the average median correlation for each condition. Paired t-test: * p<0.05 ** p<0.01.

As shown in Figure 3.18, the percentage of stable cells was also significantly increased in the red box condition when compared to the white box condition when taking data from rats tested in all protocols. However this was only true when the acceptable (t=3.13, df=7, p=0.0166) or good (t=2.586, df=7, p=0.0361) bands of cluster quality were used. For very good (t=1.545, df=7, p=0.1662) and excellent (t=1.777, df=7, p=0.1189) bands of quality, there was no significant difference shown between conditions. Power analysis showed that 7 (power=0.5, n=15) or 21 (power=0.8, n=29) more rats would be needed for the very good band of cluster quality; and 4 (power=0.5, n=12) or 14 (power=0.8, n=22) more rats would be needed for the excellent band of cluster quality. Again, this highlights that removing cells could lead to an underpowered analysis. As described for the median correlation analysis, the different protocols also appeared to show different trends in percentages of stability. Animals put through the old protocols exhibited no difference between conditions for any

bands of cluster quality (acceptable: t=0.8998, df=2, p=0.4632; good: t=0.3205, df=2, p=0.7790; very good: t=1.272, df=2, p=0.3313; excellent: t=0.1388, df=2, p=0.9023). On the other hand, animals put through the new protocol showed significant differences between the two conditions for all bands of cluster quality used (acceptable: t=3.775, df=4, p=0.0195; good: t=3.511, df=4, p=0.0247; very good: t=3.489, df=4, p=0.0252; excellent: t=3.373, df=4, p=0.0280). This mirrors the findings for median correlation described above. It appears that cluster quality was not driving the enhancement of median correlation or percentage stability, but protocol design is extremely important for both of these effects to be seen.
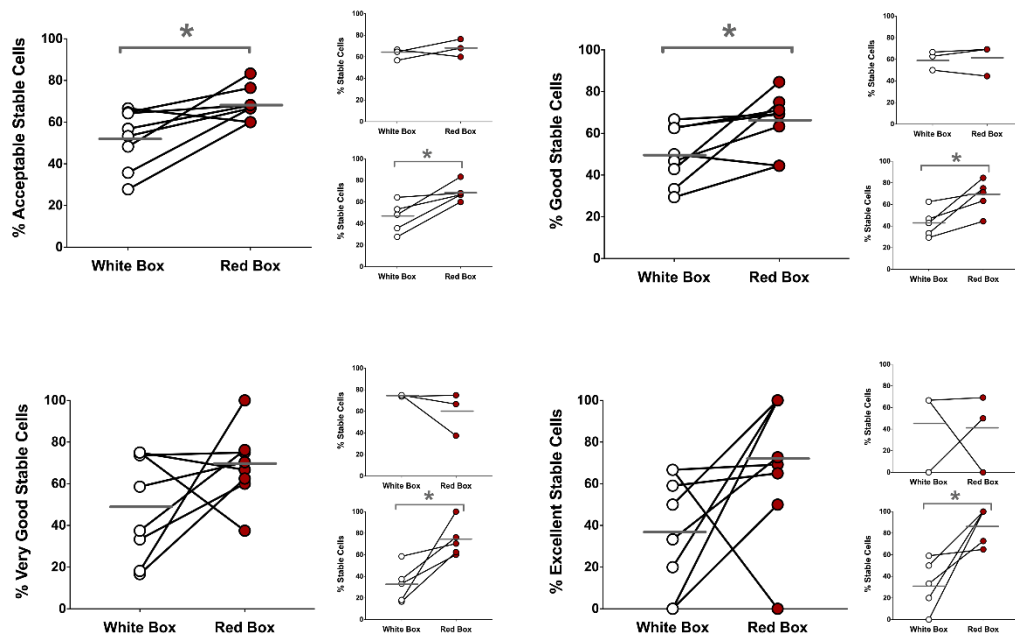
Figure 3.18: The percentage of stable cells between 0h and 6h exposures was only significantly different when the average and good bands of cluster quality were used (top left and top right). However when animals were split into the two types of protocol used – old and new – animals on the new protocol showed significant differences between red and white box conditions for all bands of cluster quality. Animals on the old protocol showed no differences between conditions regardless of cluster quality band. For each band of cluster quality, the small top graph represents the animals on the old protocol (n=3) and the small bottom graph represents animals on the new protocol (n=5). Grey lines represent the average median correlation for each condition. Paired t-test: * p<0.05

Interestingly when a two-way repeated measures ANOVA was run with only the very good cells recorded from only the five rats on the new protocol the main effect of condition on percentage stability was almost significant [F(1,4)=6.146, p=0.0683] and the interaction between time-point comparison and condition was significant [F(2,8)=7.673, p=0.0138]. When only the excellent cells recorded from only the five rats on the new protocol were analysed, the main effect of condition on percentage stability was also almost significant

[F(1,4)=6.483, p=0.0636]. These results suggest that the non-significance of this measure found previously could have been due to both poorly isolated cells and confounding effects from the old protocol.
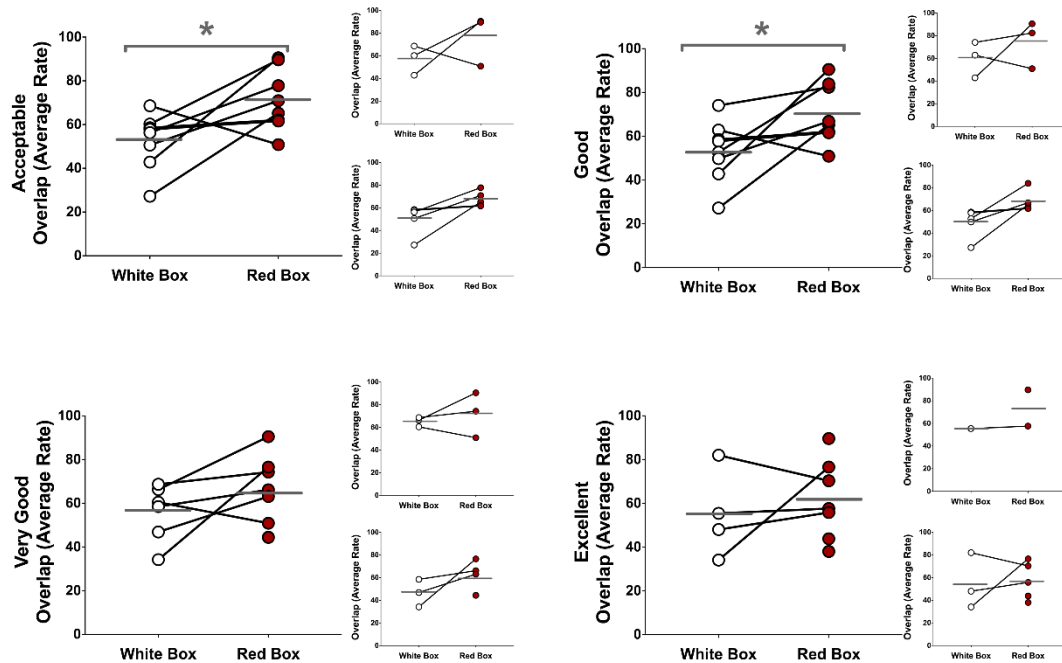
Figure 3.19: The levels of rate remapping between 0h and 24h exposures differed significantly between red and white box conditions when both the average (top left) and good (top right) bands of cluster quality were used. There were also trends towards significant differences when the very good (bottom left) band of cluster quality was used, suggesting poor cluster quality was not driving results. Analysis using higher bands of cluster quality was underpowered due to the lower numbers of "stable" cells available for analysis in both conditions. When old and new protocols were separated there were trends towards significant differences between the conditions when using the average and good bands of cluster quality, however reducing the numbers in each group led to analyses too underpowered to draw conclusions. For each band of cluster quality, the small top graph represents the animals on the old protocol (n=3) and the small bottom graph represents animals on the new protocol (n=5). Grey lines represent the average median correlation for each condition.   Paired t-test: * p<0.05

As shown in Figure 3.19, the amount of rate remapping evident between the 0h and 24h time-points in the white box condition was significantly higher than in the red box condition when taking data from rats tested in both new and old protocols. This was true when the acceptable (t=2.454, df=7, p=0.0439) and good (t=2.413, df=7, p=0.0466) bands of cluster quality were used. However when the very good (t=1.986, df=5, p=0.1038) and excellent (t=0.8854, df=3, p=0.4432) bands of cluster quality were used there was no significant difference between the two conditions. However it should be noted that as cluster quality increased, less cells were included in the analysis. Whilst this was not a problem for median rate map correlation and percentage stability, rate remapping only analyses the "stable" cells

in the population. Therefore a decrease in the number of cells analysed also decreased the number of rats involved in the analysis, due to insufficient numbers of "stable" cells between the 0h and 24h exposures at each band of quality. Only 6 rats had cells in both conditions that could be analysed in the 'very good' band of cluster quality, and only 4 rats had cells in both conditions in the 'excellent' band of cluster quality. Therefore analysis became underpowered. This is especially evident in the very good band of cluster quality, as a power analysis using data from the 6 rats that could be analysed in both conditions showed that with one more rat (power: 0.5, n=7) or six more rats (power: 0.8, n=11) the difference in rate remapping would be significant between conditions. Analysis was also split into new and old protocols as before, however due to a decrease in the number of rats analysed, conclusions could not be reached. It does appear that for both the acceptable and good cluster quality bands that animals from the new protocol were trending towards significance (acceptable: t=2.715, df=4, p=0.0532; good: t=2.672, df=4, p=0.0557) whereas animals from the old protocol were not (acceptable: t=1.013, df=2, p=0.4177; good: t=0.84, df=2, p=0.4893). Overall it appears that the increase in rate remapping between the 0h and 24h time-points in the white box condition was not driven by the quality of clusters included in analysis.

3.3.8 Place Field Remapping

To ensure that place fields were remapping between the different environments (which would indicate that the rat considered each different environment to be novel), place field stability was analysed between the switch environment and the following next novel environment exposure (0h). The switch environment was always the environment used in the previous experiment, i.e. either the familiar environment explored before the first exposure to the next novel environment in the new protocols, or the third exposure to the first novel environment (C1 24h) in the old protocols.

As the white or red box condition could be affecting the rate of subsequent remapping which in turn could affect the stability after the next white or red box condition, conditions are labelled as white box to red box or red box to white box to highlight directionality. As the protocol design appeared to significantly affect place field stability the old and new protocol conditions are also compared.

Remapping between environments appeared to happen consistently as both median correlation and percentage stability between the switch and 0h exposures were much lower than that of all other comparisons (0h to 6h, 6h to 24h, 0h to 24h). Unlike these other time-point comparisons there were no differences between white box and red box conditions for the switch to 0h correlation (median correlation: t=0.4472, df=11, p=0.6634; percentage stability: t=0.6523, df=11, p=0.5276). This remained true if white and red box conditions were split further into the two different protocols used. However there were differences in both measures between the old and new protocols, with the old protocol showing significantly higher percentage stability (t=8.588, df=11, p<0.0001) and higher median correlations that were trending towards significance (t=2.118, df=11, p=0.0578), as shown in Figure 3.20. Therefore, although both conditions remapped consistently between environments in both protocols, the levels of remapping were significantly higher in the new protocol. This could suggest reasons behind the overall differences in place field stability seen between the two protocols.
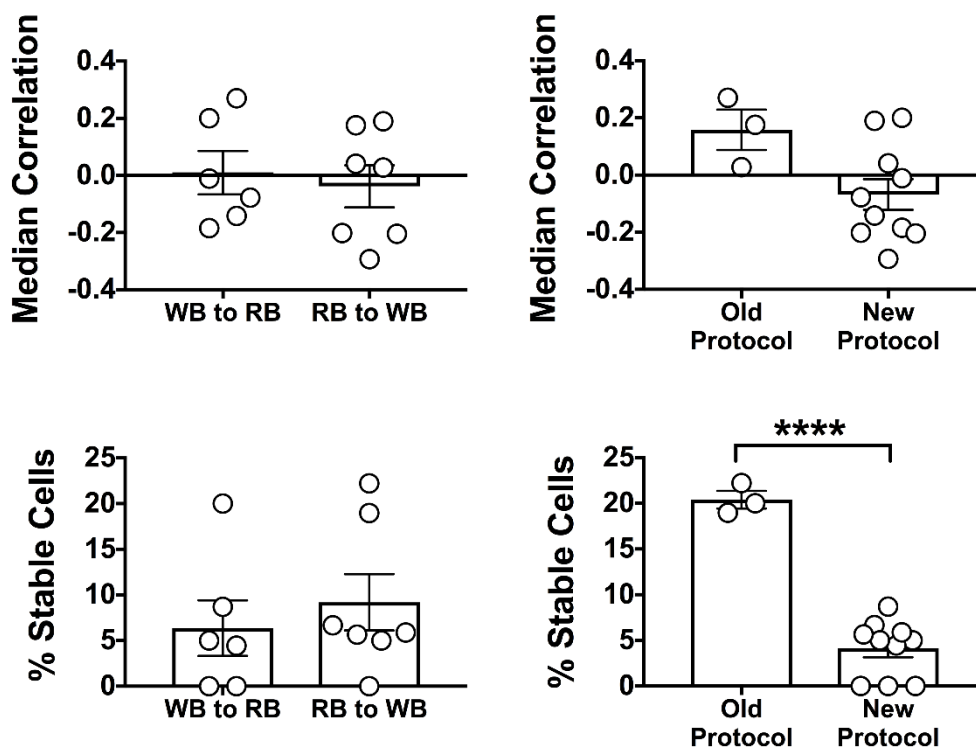
Figure 3.20: Remapping between the switch and 0h exposures did not differ between white and red box conditions. There were almost significant differences between animals on the old and new protocols for median correlation (top right) and significant differences between protocols for the percentage of stable cells (bottom right). This indicates that animals on the new protocol remapped significantly more than those on the old protocol. Error bars represent SEM. Paired t-test: **** p<0.0001.

3.3.9 Place Cell Population Firing – Neuronal Assemblies

The enhancement of place field stability was clear at the level of the place cell. The next step of analysis sought to investigate place cell firing at the population level, exploring whether reduced levels of interference affected the synchronicity of neuronal assembly firing. Only animals from the new protocol (n=5) were included in these analyses. The first step of this analysis calculated the correlation of cell firing between simultaneously recorded cells (PCorr). An activity vector was first calculated for every cell. This contained the number of times this cell had fired during every 25ms of the entire session, indicating high or low frequencies of firing throughout the recording. The activity vector for every cell was then correlated with the activity vector for every other cell to show whether cells in a population were likely to fire, or not to fire, within the same 25ms bins during the whole session, producing PCorr values for every cell pair. High levels of correlation would indicate that these place cells had the same levels of activity, be it variably high (eg. If one cell fired the other was also likely to fire) or low (eg. if one cell was silent the other cell was likely to be silent) levels of activity, throughout the 10 minute exposure. PCorr values for every cell-pair were averaged for every rat and then compared across the three exposures to a given context (0h, 6h and 24h) within conditions and between conditions. This was repeated for two other two time-bins (125ms and 1s).

At the 25ms time-bin there was a decrease in PCorr values over exposures to the environment indicating that the cells became less temporally synchronous across exposures, shown in Figure 3.21 (left). This was trending towards significance [main effect of time-point: $F_{(2,8)}=3.591$, p=0.0771], with 1 (power=0.5, n=6) or 3 (power=0.8, n=8) more animals needed to reach significance. Although further post hoc comparisons using Bonferroni corrections showed that there were no significant differences between 0h to 6h (p=0.2206), 0h to 24h (p=0.1082), or 6h to 24h (p>0.9999) comparisons, it was apparent that there was a large amount of variance in PCorr values at the 6h time-point. Further tests showed that the decrease in PCorr across time-points was only apparent in the WB condition, with results again trending towards significance [$F_{(1.703, 6.814)}=3.823$, p=0.0810]. When individual comparisons were explored in the WB condition, there appeared to be a decrease in coordinated firing of neuronal assemblies between the 0h and 6h exposures (0.006714 to 0.002742) that was trending towards significance (t=2.421, df=4, p=0.0727). A power analysis showed that with 1 more rat (power=0.8, n=6) or with 4 more rats (power=0.8, n=9)

significance would be reached. There was no difference in PCorr between the 6 and 24h exposures (t=0.321, df=4, p=0.7643). This mirrors the patterns of both absolute rate map correlation and the proportion of stable cells in the white box condition. This suggests that when the firing locations of place cells are not stable, the underlying synchronicity of place cells also decreases. It is important to note that the timescales for these phenomena are very different, the former being on a timescale of 10 minutes and the latter being on a timescale of 25ms. This implies that the stability of the place field map could be directly linked to the synchronicity of place cell assembly firing within intrinsic gamma oscillations. As found previously with place field stability, the WB condition also showed a significant decrease in correlation of cell assembly firing between the 0h and 24h time-points (t=2.84, df=4, p=0.0468), suggesting a lack of consistency in the coordinated firing of cell assemblies over time. This change in PCorr over exposures was not seen in the RB condition [$F_{(1.543, 6.17)}$=1.706, p=0.2518], implying that the neuronal assembly synchronicity was more stable between sessions in the gamma frequency band (25ms).
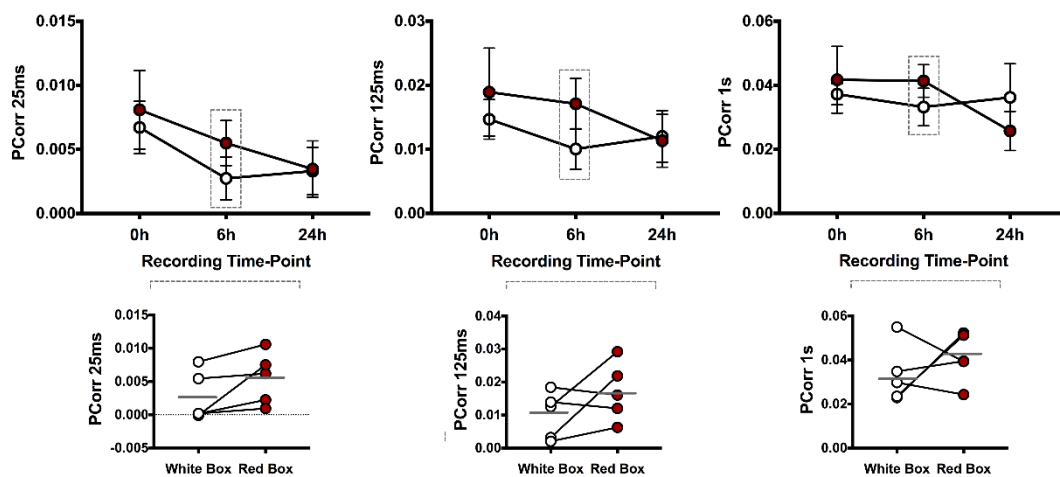
Figure 3.21: Cells became less temporally synchronous over repeated exposures to the novel environment when using the 25ms time-bin for analysis (top left) in both red and white box conditions (the average PCorr value is represented by red and white dots respectively). However only the white box condition showed significant differences between 0h and 24h exposures. Differences between 0h and 6h were also trending towards significance in this condition. Average median PCorr values at the 6h exposure were trending towards being significantly lower in the white box condition compared to the red box condition (bottom left). Neither the 125ms time-bin (middle) nor the 1s time bin (right) showed differences in PCorr values between exposures or conditions.

This difference between conditions is highlighted at the 6h time point (Figure 3.21 – bottom left), where the RB condition appeared to have higher median PCorr value than the WB condition (RB: 0.005495; WB: 0.002742). This trended towards significance (t=2.178, df=4, p=0.0950), with two more rats (power=0.5, n=7) or six more rats (power=0.8, n=11) being needed to reach significance. These results imply that reducing interference after an exposure to a novel environment could increase neuronal assembly synchronicity at the population level as well as place field stability at the individual cell level, however as there was no significant interaction between time-point and condition [F(2,8)=0.4651, p=0.6440] this effect appears to be incredibly underpowered, requiring further testing to draw any conclusions.

The 125ms and 1s time-bins showed no changes in PCorr values over exposures [main effect of time-point: 125ms: F(2,8)=0.9224, p=0.4360; 1s: F(2,8)=1.911, p=0.2097], differences between conditions [125ms: F(1,4)=0.4538, p=0.5375; 1s: F(1,4)=0.006034, p=0.9418] or interaction between exposures and conditions [125ms: F(2,8)=1.225, p=0.3434; 1s: F(2,8)=1.082, p=0.3837]. This suggests that any slight changes seen are specific to time frames associated with intrinsic gamma oscillations. This is shown in Figure 3.21 (125ms – middle; 1s – right).

These PCorr results could suggest that the synchronicity of cells decreases across sessions in the white box condition but not in the red box condition, implying that differences in place field stability might be able to be seen at the level of the cell assembly. However these slight differences in the temporal relationships between simultaneously recorded cells are only apparent on a time frame associated with intrinsic gamma oscillations.

The next step of these analyses sought to investigate the comparison of these PCorr values between sessions further by correlating every PCorr value with every other PCorr value on a rat-by-rat basis. Instead of showing the average synchronicity of cell assemblies within a session for every rat, and whether this average synchronicity changed over sessions, this analysed how the synchronicity of each individual cell pair differed between sessions. High PCo values would suggest that the synchronous cell assembly firing in one session was very similar to that of another session, low values would suggest that the temporal relationships of cell assemblies were changing between sessions. Figure 3.22 (top) shows the range of

PCorr values, ordered from the lowest to highest values of all of the cells during the 6h exposure. This order of cells is used in the 0 and 24h exposure bars to highlight the differences in PCorr of individual cells over the three sessions. This highlights the similarities or differences of individual cell-pair firing synchronicity between sessions. PCo values are given at the top of each bar (r).
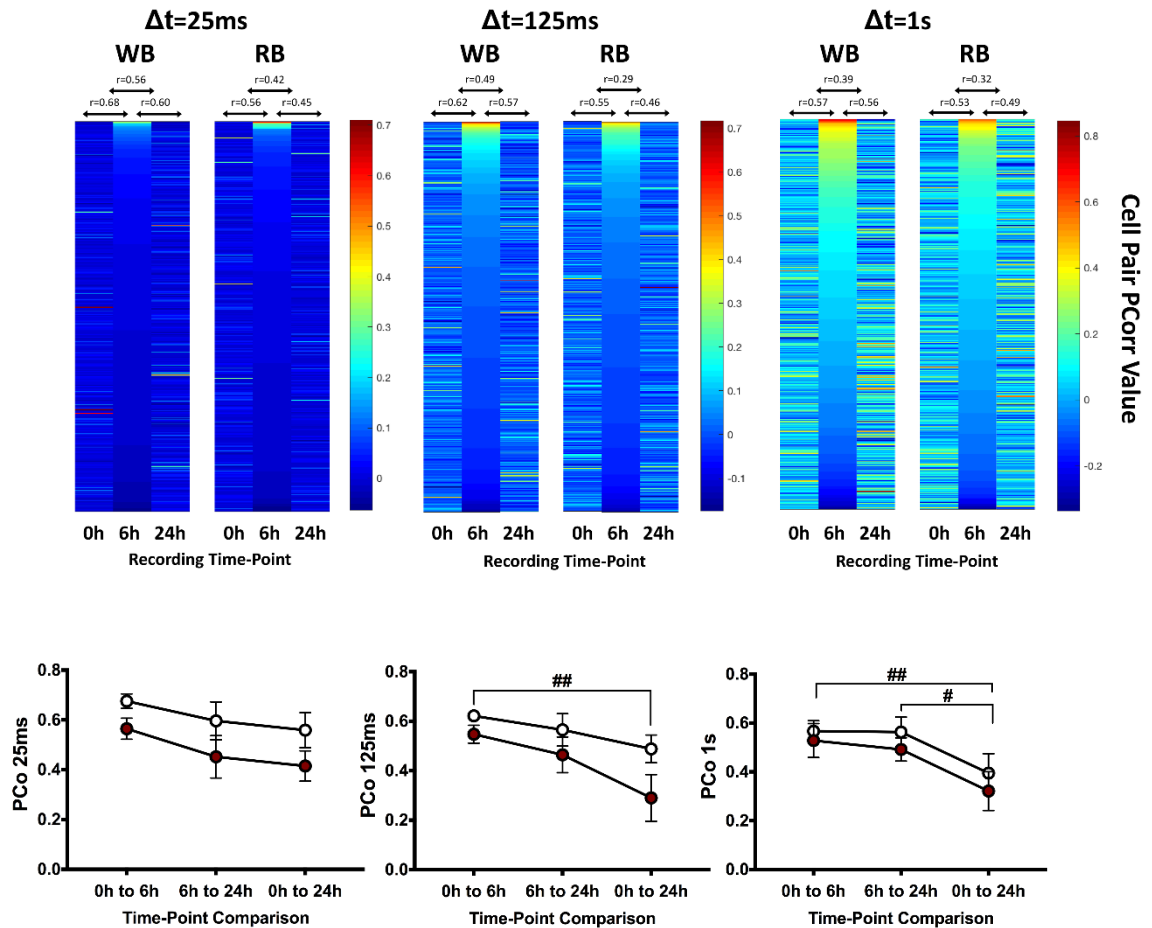
Figure 3.22: The average synchronicity of every individual cell pair decreased between the 0h and 6h exposures and the 0h and 24h exposures. This was true for every time bin used for analysis (25ms – left, 125ms – middle, 1s – right), however changes between time-point comparisons were only significant when using the 125ms and 1s time-bins. There were no significant differences between conditions or interactions between time-point comparison and condition (red dots represent average PCo values for the red box condition, white dots represent average PCo values for the white box condition). For every time-bin used, and each condition there are three bars representing the range of PCorr values at 0h, 6h and 24h. These are down-sampled to 500 values to aid visual representation. These values are ordered from the lowest to highest values of all of the cells during the 6h exposure. This order of cells is used in the 0 and 24h exposure bars to highlight the differences in PCorr of individual cells over the three sessions. The 'r' values above the bars represent the pearson correlations between every cell pair PCorr value, i.e. the PCo values between the two exposures. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: # p<0.05   ## p<0.01

Consistent with trends seen in the PCorr values, when individual Pcorr values (representing cell pairs) for the gamma associated time frame (25ms) were correlated across sessions (PCo) the coordination appeared to decrease over time. As before, this was trending towards significance, but underpowered [main effect of time-point comparison: $F_{(2,8)}=3.344$, $p=0.0880$]. Further post hoc comparisons using Bonferroni corrections showed that there were no significant differences between any time-point comparisons and, unlike the PCorr values themselves, neither the red box or white box condition showed any changes in PCo over time [WB: $F_{(1.986, 7.943)}=0.979$, $p=0.4162$; RB: $F_{(1.722, 6.887)}=3.2$, $p=0.1076$]. There was no main effect of condition [$F_{(1,4)}=2.845$, $p=0.1669$] nor significant interaction between time-point and condition [$F_{(2,8)}=0.06796$, $p=0.9348$] (Figure 3.22 – bottom left). Whilst you might expect PCo values to be higher in the RB condition as PCorr values are increased, this can be explained. If across different exposures certain cells always fire together (high Corr) and at the same time certain cells never fire together (low PCorr) this would produce a high PCo value. However if most cells never fired together (low PCorr) on both exposures this would also produce a high PCo value. The latter could explain why the WB condition does not have lower PCo values.

PCo values associated with the 125ms time-bin significantly decreased over time [main effect of time-point comparison: $F_{(2,8)}=10.79$, $p=0.0054$], indicating that assembly firing of neuronal populations become more desynchronised from the first exposure to the last exposure of a novel environment (Figure 3.22 – middle). Further post hoc comparisons using Bonferroni corrections showed that there were significant differences between the 0h to 6h and 0h to 24h time-point comparisons ($p=0.0054$) and an almost significant difference between the 0h to 24h and 6h to 24h time-point comparisons ($p=0.0556$). This indicated that the cell assembly synchrony during the final exposure to the environment was very different to the first two exposures. The final exposure was 24h or 18h after these first two exposures, indicating that assembly firing of neuronal populations becomes more desynchronised over longer delays. It should be noted that levels of spatial remapping and rate remapping were also highest at the 24h time-point. There was also no interaction between time-point comparison and condition [$F_{(2,8)}=0.9669$, $p=0.4296$].

PCo values associated with the 1s time-bin, shown in Figure 3.22 (right) also significantly decreased over time [main effect of time-point comparison: $F_{(2,8)}=11.88$, $p=0.0040$], mirroring the other time-bins used. Further post hoc comparisons using Bonferroni

corrections showed that there were significant differences between the 0h to 6h and 0h to 24h time-point comparisons (p=0.0065) and significant differences between the 0h to 24h and 6h to 24h time-point comparisons (p=0.0124). This is in line with the results above using the 125ms time-bin for analysis. This indicates that the 125ms and 1s time-bins were most sensitive to changes in synchrony of cell assembly firing between sessions occurring on separate days. Again there was no interaction between time-point comparison and condition [F(2,8)=0.1492, p=0.8713].

For both conditions and within all time-bins the PCo value is lowest at the 0h to 24h time-point comparison, which mirrors the levels of stability seen at this time-point. This suggests that changes in stability on the cellular level can be seen at the population level. Overall, although the 25ms time-bin weakly suggested an increase in neuronal assembly synchronicity in the RB, there was no obvious effect at the population level.

3.3.10 Results Overview

The average firing rate of place cells did change significantly across repeated exposures to a novel environment, however not as expected. Firing rates were higher in the second exposure (6h) in both conditions, decreasing as expected for the third exposure (24h). Peak in-field firing rate did not differ between exposures. All spatial properties measured improved from the first novel exposure to the last two exposures in both the white box and red box conditions. This suggests that over time there is a significant improvement in consistent spatial representation of a novel environment. However, the number and size of the place fields did not change over time with repeated exposures. Surprisingly, reducing interference after the first novel exposure did not significantly impact any of these place field properties during the following two exposures.

The stability of place fields appeared to increase over time between 0h and 6h or 6h and 24h, however there were significant differences between the red and white box conditions. The reduction of interference in the red box condition led to enhancement of median correlation between firing rate maps for each animal, and the proportion of cells for each animal that showed a correlation that significantly exceeded a specific threshold of stability. This enhancement also appeared to be more consistent over time, with rate maps being more similar between the first and last exposures in the red box condition. It was confirmed that

cluster quality was not driving the clear enhancement of stability between the 0h and 6h exposures to the environment. However it was necessary to use a protocol with sufficient time between experimental repeats, suggesting that protocol design is extremely important for the effects of reduced interference to be seen. Interestingly, even when only the "stable" population of cells was analysed, levels of rate remapping significantly increased when comparing the 0h to 24h exposures to the 0h and 6h and the 6h and 24h exposures in the white box condition. It appeared that cluster quality was also not driving these differences in rate remapping seen between the two conditions.

When analysing neuronal assembly firing using a time frame associated with gamma oscillations and hippocampal replay, the RB condition showed slightly more consistency in the synchronous firing of cell assemblies between the different exposures. The WB on the other hand showed decreases in coordinated firing between 0h to 6h and 0h to 24h exposures, again indicating less consistency in cell assembly firing. Although PCo values appeared to decrease over time, indicating much less synchronicity at a population level between the 0h and 24h exposures, there was no obvious difference between WB and RB conditions at this level.

Overall, only properties associated with spatial tuning and stability changed over time, with place fields becoming better spatially tuned and more stable as the environment became more familiar. The reduction of visual interference after exposure to a novel environment significantly enhanced place field stability when measured 6h later. This enhancement also appeared to increase the consistency of place field firing and cell assembly firing between every exposure, with much more similar patterns of firing from the first to last exposures. These results correlate well with our previous behavioural findings, extending the black box effect to the phenomenon of place cell "memory" for a novel environment.

3.4 Discussion

This experiment set out to test the hypothesis that decreased interference following an initial exposure to a novel environment would lead to an enhancement in the firing properties of place cells typically associated with familiar environments, and a decrease in firing properties typically associated with novel environments, when the animal was exposed to the same environment 6h and 24h later.

The first correlates of place cell "memory" assessed were the average place cell and peak in-field firing rates. Previous studies have shown that average place cell firing rates are lower in familiar environments than in novel environments (Brandon et al., 2014; Karlsson and Frank, 2008; Nitz and McNaughton, 2004), whereas peak in-field firing rates appear to be higher in familiar environments then in novel environments (Karlsson and Frank, 2008). However, in our dataset these patterns were not observed. Whilst average place cell firing rate did change significantly over the repeated exposures to the novel environment, it appeared to be highest in the second exposure at 6h. This is surprising, as this exposure wouldn't be deemed the most novel or the most familiar. It is possible that animals had no memory for the first exposure to the environment, however this is incredibly unlikely as over half of place cells expressed stable fields between the first and second exposures. This increase in firing rate therefore requires further testing. Average place cell firing rate did decrease for the final exposure at 24h, however this was only to levels seen in the initial and most novel exposure. Therefore there was no change in average place cell firing rate when comparing the first and the last novel exposures. This was in line with the peak in-field firing rate results, which showed no change across any repeated exposures. It could be argued that after three 10-minute exposures spaced over 24 hours the environment was not fully familiar to the animal. Studies comparing familiar and novel environments tend to familiarise animals over many exploration sessions for periods spanning multiple weeks (Nitz and McNaughton, 2004). One such study compared a familiar environment, in which the animal had explored daily for a week, to three 10-minute exposures to a novel environment (Brandon et al., 2014). They too found significant differences in average place cell firing rate between the familiar and novel environments, even when comparing the last of the exposures to the novel environment (i.e. the most familiar 'novel exposure') to the familiar environment. This suggests that more exposures would be required to see the effects of increasing familiarity on place cell within field firing rates.

Whilst average place cell and peak in-field firing rates remained relatively stable over the three exposures to the novel environment, with no differences seen between the first and last exposures, we found that the fields themselves became more spatially tuned. Spatial information, selectivity and coherence increased, and spatial sparsity decreased between the first and final exposures in both conditions. This indicates that spatial tuning improved rapidly during the initial exposure. These data are consistent with increased spatial tuning with repeated exposures to a novel environment reported previously (Karlsson and Frank, 2008). The fact that all four spatial properties measured increased over the three exposures supports this suggestion. These improvements in spatial information, sparsity, coherence and selectivity relate to more information contained in each neuronal spike, more location specific firing, smoother place fields and a higher signal-to-noise ratio of place field firing respectively. These changes can be thought of as a neural correlate of familiarity with the environment. This indicates that increases in familiarity can be seen over just three exposures to a novel environment.

Although place fields became more spatially tuned over time there did not appear to be any significant changes in place field size. Whilst this is surprising, as other studies have shown decreases in the size of place fields over just three exposures to a novel environment (Brandon et al., 2014), it could be the case that when these exposures are spaced further apart (i.e. 6 or more hours between exposures compared to 5 minutes) more exposures are needed overall to highlight these subtle changes.

In contrast to initial predictions, neither firing rate changes nor spatial properties of place cells were enhanced in the sessions following exposure to the red box compared to exposure to the white box. This is unexpected as both of these place field properties have been suggested to be physiological correlates of familiarity with an environment (Kentros et al., 2004). As the red box is known to enhance spatial nOL memory, you could expect an overall improvement in spatially selective firing of place cells after exposure to this box. However one major difference between the nOL task and the exploration of the novel environment is that the latter only requires foraging and no actual task. The animal could therefore have decreased attention to the novel environment, a factor that has been shown to significantly change the levels of spatial information and firing rate compared to tasks requiring attention of the animal to its surroundings (Kentros et al., 2004). One major question, however, is how

do the spatial properties of place cells improve over repeated exposures spaced 24-hours apart whilst not being enhanced by a condition known to enhance spatial memory? A study where animals were not required to attend to a task or changes in the environment showed that whilst place field stability increased over exposures in the long-term, suggesting that spatial memory had been consolidated, spatial information did not increase (Bett et al., 2013). This implies that the consolidation of spatial memory is not always concomitant with changes in spatial properties of place cells. However, in our study spatial information did increase, even if this didn't correlate with possible memory enhancement. It is possible that any changes in spatial information correlating with memory enhancement would be so subtle that a larger environment would be needed to increase the range of spatial information possible, especially as the changes in place field stability of the cell population did not range from completely stable to completely unstable and were instead much more subtle.

As hypothesised, place fields became more stable over repeated exposures to the same environment, consistent with many different studies. The white box condition showed slightly higher levels of stability between the 6h and 24h exposures to the environment compared to the 0h and 6h. This is in line with previous studies showing that rate map correlations between the first two exposures are much lower than those between the second and third exposures to a novel environment (Bett et al., 2013). This suggests that high levels of stability are reached only after two 10 min exposures to a novel environment, whereas after just one 10 min exposure, the place field "memory" is not as strong. In contrast, although the red box condition showed similar levels of stability as the white box condition between 6h and 24h, stability levels were higher between the 0h and 6h exposures. This significant difference of 0h to 6h place field stability between the two conditions suggests that reducing interference after an animal explores a novel environment for the first time enhances the long-term stability of place fields.

Interestingly, whilst there was no difference between conditions in stability between the 6h and 24h exposures, the red box condition had significantly higher median correlation values between the first and last exposures. This suggests that the increase in stability between the initial novel exposure and the 6h exposure, caused by a reduction in interference, could have consolidated the initial pattern of place field firing at a much quicker rate. This firing pattern could then be recalled much more consistently at both 6h and 24h exposures, leading to

almost half of the cells recorded remaining stable over a 24h period. In contrast, it suggests that the white box condition did not fully consolidate the spatial memory of the first exposure to the novel environment. Although the same levels of stability were seen between the 6h and 24h exposures, the pattern of firing at this last time point was very different to the initial pattern produced during the first exposure, with less than a third of cells recorded remaining stable. This could be due to how consistently the initial 'map' could be recalled over time.

This pattern of place field stability was also seen in levels of rate remapping. Animals in the red box condition showed relatively low levels of rate remapping throughout the experiment. Animals in the white box condition also showed low levels of rate remapping between the 0h and 6h exposures and the 6h and 24h exposures, however levels of rate remapping were significantly higher between 0h and 24h exposures compared to the red box condition. This suggests that even cells that are considered "stable" are not being consistently recalled over time. Overall these differences in place field stability suggest that the reduction of interference in the red box could be enhancing the consolidation of the spatial map.

It has been shown previously that neuronal firing patterns associated with a novel environment appear to be preferentially reactivated during subsequent sharp wave ripples (O'Neill et al., 2008; Ven et al., 2016), implying that the stabilisation of place field firing patterns after the exposure to a novel environment requires consolidation via sharp wave ripple activity. An important future experiment would therefore be to record LFP activity of animals whilst in the red and white light boxes to investigate whether animals in the red box have increased levels of sharp wave ripples after the initial novel exposure compared to animals in the white box. It would be hypothesised that increased levels of place field reactivation during ripple activity could lead to the increased stability of place cell firing seen in the red box condition. Another future experiment directly linked to this proposal would involve directly blocking sharp wave ripple activity in the red light box, possibly via optogenetic stimulation, in an attempt to inhibit the increase of place field stability seen in this condition. If stopping ripple activity blocked the enhancement of place field stability it would suggest that enhanced sharp wave ripples and place field replay do indeed underlie the effects of reduced interference after spatial learning. It is important to note that a variation of this experiment has been previously carried out in mice, and that it showed the opposite to this hypothesis (Kovács et al., 2016). Mice explored a novel environment for 24 minutes, followed by three hours of sleep where sharp wave ripple events were

optogenetically disrupted (or a control condition where disruption was offset by 1.32s) and a final 24-minute exposure to the same novel environment. Mice in both the control and ripple disrupted conditions showed rate map correlations between the first and second exposures of approximately 0.5, suggesting place fields remained stable even when sharp wave ripples were disrupted. Whilst this appears to contradict the hypothesis that blocking sharp wave ripples would disrupt place field stability, there is a possible explanation. Whilst the three-hour delay used could be labelled as long-term memory, this is contentious. Place field studies have shown that consolidation mechanisms, including NMDA-R activation and PKA activity, are required for place field stability at 6 and 24 hours, but not 1-2 hours after the initial novel exposure (Agnihotri et al., 2004; Kentros et al., 1998; Rotenberg et al., 2000). It could therefore be suggested the stability of place fields after 3 hours would not require consolidation mechanisms and ergo sharp wave ripples. Studies investigating the difference between early-LTP and late-LTP support this possibility, as early-LTP has been shown to last for 3-5 hours *in vitro* (Frey et al., 1988). This suggests that late-LTP, a process thought to be akin to the long-term stability of place fields (Dragoi et al., 2003), does not occur until at least the 5 hour time-point. It is also important to note that the study recorded from mice and not rats. Mice have less stable place cell firing than rats and need to pay more attention to their environment, such as in a task, to improve this stability over longer intervals (Kentros et al., 2004). This could suggest that without the requirement of attention, the long-term stabilisation of place fields could take longer in mice. Overall this suggests that repeating this experiment in rats and at longer delays would be needed to reach the conclusion that sharp wave ripples are or are not required for place fields to become stable in the long-term.

Other experiments have shown that sharp wave ripple activity and subsequent stabilisation of place fields can predict an animal's performance on a spatial task (Dupret et al., 2010a). This therefore supports the possibility that reducing interference after spatial learning actively improves consolidation mechanisms through enhancing sharp wave ripples and place field replay, consistent with the nOL behavioural results discussed previously showing that exposure to the red box after spatial learning enhances memory retention.

Surprisingly, when the quality of clusters was taken into account significant differences emerged between the two different types of protocol used. The old protocol that combined two different contexts into a 3-day experiment appeared to not show any enhancement of

place field stability in the red box condition. In contrast, the new protocol that separated each novel context by at least 5 days showed significant levels of enhanced place field stability in the red box condition, even when only the best quality clusters were used in analysis. One possible reason behind this difference could be that cells recorded during the new protocol were only recorded over a period of 24h for each experiment, whereas cells recorded during the old protocol were recorded over a period of 3 days. As the gains were optimised at the start of every experiment, but weren't changed for the duration, cells recorded during the old protocol were more likely to become lost or become noisier by the end of each experiment. However if this were true then you would expect to still see the red box effect when only the best clusters were used in the analysis. As the quality of clusters used did not make a difference it suggests that other factors are at play.

All the rats that underwent the old protocol appeared to have very stable cells throughout, which could be an indicator of problems with remapping between the contexts used. When looking at the remapping values between the contexts for these rats, although they were still low enough to be classed as unstable and remapping (i.e. an average median correlation of 0.16 with an average of 20.4% stable cells) these values were much higher than those of the rats from the new protocol (i.e. an average median correlation of -0.07 with an average of 4.13% stable cells). Whilst the new protocol used both cylindrical and square environments the old protocol only used cylindrical environments. Therefore although the cues surrounding the environment were changed and the colour of the environment itself was different, there were only 5 minutes in between exposures of two cylindrical environments of similar sizes. Although previous studies have shown that global remapping can occur between environments of similar shape and size if the colour of the environment (Kentros et al., 1998) or cue card is changed (Bostock et al., 1991), only one or two cells were recorded simultaneously during these experiments with increased numbers of cells coming from many repeats of the same protocol. This means that the overall percentage of remapping might be hard to define as the few cells recorded are used to infer the effects on the whole population. It should be noted that although Bostock et al. reported that global remapping could occur, it did not occur for every animal in every experiment suggesting that the animal's attention to their environment was an important factor. Other studies have also shown that changing one or two environmental cues does not destabilise the place field map on a global scale (O'Keefe and Conway, 1978). This suggests that the amount of remapping seen could be

dependent on the protocol and environments used. Whilst it seems logical that leaving longer between exposures to two different contexts would increase the levels of remapping, place cells are known to be able to differentiate between a seemingly infinite number of contexts, with firing remaining stable in each individually (Colgin et al., 2008). Therefore you might not expect the previous context to affect how well fields stabilised in the next, even if it was directly after. However, it is likely that the decrease in remapping seen in this sub-group could have masked any subtle differences in place field stability overall as the initial baseline of stability would be changed (20% of cells could already be stable compared to only 4%).

The temporal relationships between simultaneously recorded cells were also assessed. This was to determine how these changed across the three exposures to the novel environment, as well as whether these differed between the red and white box conditions. The first type of analyses explored how synchronous cells were within sessions, and whether this synchronous firing increased or decreased between sessions or between conditions. High PCorr values relate to high levels of synchronicity of cell assembly firing within a specific timeframe, in this case 25ms, 125ms or 1s.

Although the clear changes in place field stability seen previously were not seen when looking into the synchronous firing of cell assemblies at 125ms or 1s time frames, subtle changes in cell assembly firing (PCorr) were seen when using a time frame associated with gamma oscillations and hippocampal replay. As with place field stability, the red box condition showed more consistency in the firing of cell assemblies between the different exposures. The white box condition showed decreases in the coordinated firing of cell assemblies between 0h to 6h and 0h to 24h exposures. This suggests that the simultaneous activation of place cells within this time frame is important for the spatial encoding of a novel environment. This is a logical conclusion, as the activation of multiple place cells to signal an animal's location is much less prone to error than a single cell firing (Wilson and McNaughton, 1994). The importance of cell assembly activation at this 25ms timeframe has been shown before. One study recorded cell assemblies that were activated during exploration of a novel environment using this 25ms time frame. These assemblies were then subsequently recorded during a rest session followed by another exposure to the same novel environment. It was found that the reinstatement of these cell assemblies directly correlated with their

reactivation during sharp wave ripples, and that blocking sharp wave ripples disrupted this cell assembly stability over sessions. This was only true for a novel environment, as it appeared that reactivation during rest and disruption of sharp wave ripples did not affect the stability of cell assemblies in a familiar environment (Ven et al., 2016). This suggests that the consolidation of a novel environment occurs during the hours of rest following the exposure, and that this consolidation occurs at the level of cell assemblies mediated by sharp wave ripples during gamma oscillations. The 25ms time frame for these cell assemblies is thought to be optimal for the synchronicity of intrinsic gamma oscillations with externally modulated place cell spiking and plasticity mechanisms (Harris et al., 2003). As with the stability of place fields, these results appear to support the idea that the red box is enhancing spatial memory through the reactivation of place cells during sharp wave ripples, however as the effect appeared to be very underpowered further testing is required to draw conclusions.

The next set of analyses assessed whether this synchronous firing of cell assemblies remained stable across the different sessions, i.e. were the same cell assemblies co-active in each session or did this change. High PCo values would suggest high levels of similarity in cell assembly activation between sessions, whereas low values would suggest that the temporal relationships between cell assemblies were changing.

When looking at the coordinated firing of these cell assemblies at the population level (PCo) the similarity of cell assembly firing decreased over repeated exposures. 0h to 24h PCo values were significantly lower than 0h to 6h values, which is consistent with the place field stability results. This suggests that more obvious global changes can be seen at the population level, and that the likelihood of coordinated activity enduring over a long 24h period is less. PCo values also appeared low in both conditions, despite significant differences in place field stability seen between the red or white box conditions at the cell level. Whilst this could be seen as surprising, this result could theoretically be expected if the role of the place cell in spatial memory and the contentious issue of the spatial map theory are reconsidered.

Many arguments have been made against the spatial map theory, with much emphasis put upon the other roles of hippocampal place cells. Since the discovery of these cells it has been shown that they can encode much more than just the animals location. A large number of place cells seem to also be controlled by other modalities such as odour, and can be

modulated by motivation, goals, and the task at hand. This includes firing specifically at goal locations (Gothard et al., 1996) or for specific odours (Wood et al, 1999), during different timing delays used in a delayed-match-to-sample task (Hampson et al., 1993) or when waiting for a reward (Kraus et al., 2013) and during specific behaviours of the animal such as grooming or eating (O'Keefe, 1976; Ranck, 1973). It is therefore suggested that rather than encoding different parts of the environment in a solely spatial manner, they are encoding the regularities between different events, such as exposures of an environment or trials of a task, into a temporal framework. A subset of these many different environmental cues, such as the sound of a buzzer or reward of food, could then be used to recall these memories quickly allowing the animal to flexibly solve new tasks or navigate more efficiently through different environments. Coordination of firing at a population level is therefore thought to give the rat information not only about its location in space but also about the temporal relation of any other salient cues present at the time of encoding, and whether these repeat consistently in different episodes.

It could be argued that the place cells recorded during this spontaneous exploration of a novel environment would be acting as "true place fields" and only encoding the animal's location within an environment rather than things such as routes, goals or expectations. This idea was suggested by Eichenbaum (Eichenbaum et al., 1999) when commenting on an experiment by Muller and Kubie (Muller and Kubie, 1987) in which rats spontaneously foraged in a completely controlled open field. Spatial cues would be the only available cues that were consistent enough to encode in a meaningful manner, as every other cue would be randomised. Without the need to encode many different types of cues within the same temporal episode, it is plausible that the coordinated firing of cell assemblies at the population level would not be necessary for encoding of the event. Therefore, whilst you might not expect high levels of coordination in the cell population during spontaneous exploration, this synchronicity at a population level could be required for successful memory encoding and consolidation during a spatial memory behavioural task, such as nOL.

3.5 Conclusions

Reducing visual interference after exploration of a novel environment enhances long-term place field stability. This is supports the idea that the enhancement of nOL spatial memory seen previously is due to increases in sharp wave ripples and place cell replay. It is therefore hypothesised that reducing retroactive interference after spatial learning enhances place field stability in such a way that correlates to the strength of nOL memory.

Chapter 4: Investigating Whether Place Cell Properties Associated with Memory for Object Locations are Enhanced by the Black Box Effect

4.1 Introduction

In the preceding chapters of this thesis, I have shown that reducing interference after spatial learning of a novel object location (nOL) task can lead to enhanced long-term memory expression (Chapter 2) and that reducing interference after spatial learning of a novel environment can lead to increased long-term stability of the underlying place fields (Chapter 3). However, it is not known whether place field properties related directly to memory for object locations are similarly enhanced by reducing interference after the initial exposure to the objects. In the nOL task in Chapter 2, the animals were first habituated to the environment for 4 sessions before being exposed to the objects in the environment. Therefore, unlike in Chapter 3 where I assessed stability of place fields between successive exposures to a novel environment, in the object location task the environment is familiar and I am assessing memory for the locations of objects within an already familiar environment. The purpose of this experiment described in this chapter was therefore to assess whether place cell properties associated with memory for object locations are enhanced by reducing interference after an object sampling period. Any place cell properties that correlate to the strength of memory could be likely to contribute to how object locations are encoded into long-term representations in the hippocampus, or even how they are later retrieved during the probe trial to signal object location familiarity. Such properties could include the stability of place fields, firing rate or the association of place fields with the objects.

As seen in the previous chapter you would expect place field expression to be stable after multiple exposures to the same environment, regardless of whether the animal had been exposed to a condition with reduced interference levels after the first exposure. What could be interesting to investigate is whether these fields would be expressed in the same locations, and therefore remain stable, after the addition of objects to the same environment. Some studies have suggested that global remapping might not occur when objects are introduced or moved as the context has not changed to a great enough extent to be considered a different environment altogether. Instead it has been shown that rather than changes in place field locations to signal that something in the environment has changed there could be variations in place cell firing rates (Larkin et al., 2014). Larkin recorded place cells during a novel object location paradigm using a sampling session, with two novel objects in a familiar environment, as the baseline session. This was then compared to the probe session where one of the two now familiar objects had moved to a novel location. Larkin

found that the movement of one of the two identical objects in a familiar open-field did not lead to any obvious changes in the level of place field stability, but instead led to significantly higher levels of average place cell firing rate (the number of spikes/the total session time) indiscriminately throughout the environment. This was suggested to act as a 'generalised novelty signal' allowing new information to be incorporated into the pre-existing spatial map, as firing rates decreased within the second half of the probe session as the new object location became more familiar. Although not tested, it would be plausible that if the baseline was a session without objects, and this was compared to the sampling session where two novel objects were introduced, this novelty signal would be even more pronounced.

On the other hand, rather than indiscriminate increases in firing rate signalling novelty, it is possible that the place fields will undergo rate remapping. This is where locations of fields remain stable but their firing rate changes as if they were in a different environment. Therefore one rate would be associated with the environment without objects, and one would be associated with the environment with objects, although these rates would not change within sessions. Rate remapping has been implicated in studies where the local context changed but the surrounding environment remained the same (Leutgeb et al., 2005b). Animals were either recorded in the same room but in different open fields (square and circular, or black walls and white walls) or in the same open field in different rooms. Place cells of animals recorded in the latter exhibited global remapping as expected, with correlations between the two rate maps being low. In contrast, place cells of animals in the former exhibited very minor shifts in place field location, but very big differences in the ratios of place cell firing rates. This suggests that the animal knows it is in the same room so global remapping does not occur, but that there are local cues that have changed and therefore need to be updated. As the same context and surrounding environment were used in the novel object location experiment, with the objects serving as changing local landmarks, it could be suggested that rate remapping would occur rather than remapping on a global scale. However, it should be noted that there were no objects in the study by Leutgeb, and also the highest levels of rate remapping were recorded from CA3 cells, whereas in the study by Larkin et al and in Chapter 2 of this thesis, cells were recorded in CA1.

Another study using a conditional discrimination task, where rats had to discriminate between two odours and respond to one or the other depending on which of two contexts

they were in, found increases in firing rate correlating with learning of the task in CA1 and CA3 neurons equally (Komorowski et al., 2009b). Rats were exposed to two identical pots with two different odours and allowed to dig in one for a food reward. In one context one of the odours was rewarded and in a different context the other odour was rewarded, requiring animals to make an association between an odour and a context. Unlike the study by Larkin, Komorowski found that this increase in place cell firing rate was specific to fields in significant locations, i.e. the rewarded odour in a particular context. This suggests that rather than a generalised novelty signal, firing rates near objects could signal learning of the new object locations. Although this study used two different contexts within the same behavioural protocol, unlike the nOL task, it could be suggested that firing rates of place fields could increase near the objects, leading to overrepresentation of these object locations.

On the other hand, other studies have shown that stable place fields can become disrupted upon the addition of novel objects to the environment. One such study trained rats to run around a circular track for a food reward to familiarize themselves with the environment, after which eight unique novel objects were placed at different locations within the track. The location of CA1 place fields significantly changed after this introduction of novel objects, which could suggest that these objects were acting as local landmarks or cues signifying a new environment and therefore leading to remapping on a global scale. The overall number of place fields also increased. After the rats ran twenty laps of the track interacting with the objects, all eight of them were randomly shuffled. Therefore all familiar object locations now had a different but familiar object location. This caused place cells to remap further, indicating that place field expression could be modulated by both the introduction and movement of objects (Burke et al., 2011). In contrast to the previous studies, neither of these novel conditions led to a significant change in place cell firing rates. One of the differences between these studies is the number of objects used. The studies that indicated changes in firing rate only had two different objects present at one time, whereas the study that indicated global remapping had eight different objects appearing or moving simultaneously. It could be suggested that for global remapping to occur high levels of change or novelty are required.

The level of novelty needed for changes in place cell properties to become apparent has been explored in a recent study by Zheng et al., (2016). Three different object tasks were used:

novel object location, novel object, and a combination of the two where both the location and object were novel in the probe session. Zheng found that the in-field firing rates of place cells were significantly increased when the animal was exploring the novel object location specifically during fast gamma oscillations. This was true for both the novel object location task and the novel object-novel object location task, indicating that place field firing increased during location novelty only near the novel object locations. Although it should be noted that these increases in firing rate were much more apparent when both the object and object location were novel, suggesting that if analysis is unable to separate firing rates during fast gamma from firing rates during slow gamma, then these differences might be too small to see.

The first aim of this experiment was therefore to investigate how place cell properties, including the stability of place field expression and place cell firing rate, related to the addition of objects and subsequent changes in object location within a familiar environment. It was hypothesized that place fields would have the highest levels of stability between the two baseline sessions before the introduction of objects in the sampling session, and remain stable when objects were introduced or moved, as the addition of only two objects should not be enough to signal global remapping when the overall context and surrounding environment did not change. However it could be predicted that the average place cell firing rate would increase when objects were added and moved to act as a generalized novelty signal, as seen by Larkin. The peak in-field firing rate of fields might also increase when objects were added and moved, however there might not be sufficient levels of novelty to see such changes.

Interestingly it appears that although a large proportion of cells can remain stable when objects are introduced to the environment, some studies have found that another smaller population of cells are directly modulated by objects, a phenomenon termed partial remapping (Skaggs and McNaughton, 1999). More specifically, fields that are expressed in close proximity to objects appear to be more directly modulated by object movement even though place fields themselves are not more likely to be expressed close to objects in an open field environment (Deshmukh and Knierim, 2013). One such study by Lenck-Santini found that CA1 place fields located in close proximity to objects underwent much higher levels of remapping when objects were rotated than those of more distantly located fields (Lenck-

Santini et al., 2005). Two novel but distinct objects were added to a familiar cylindrical environment which had a large cue card attached to the internal wall. Rats were allowed to explore these objects freely whilst CA1 place cells were recorded for two sessions, acting as baseline recordings. Objects were then moved to positions rotated 90 degrees relative to the cue card. Place cells reacted in two different ways depending on the proximity of their fields to the objects. Fields expressed more than 10 cm away from objects remained relatively stable relative to the spatial location of the cue card (mean correlation 0.60). However, although the mean firing rate of these cells did not change, the in-field firing rate significantly decreased over sessions, showing that whilst these fields were stable in location, their coherence decreased when objects were moved. In contrast fields expressed within 10 cm of the objects were unstable (mean correlation 0.36), either changing location within the environment or becoming silent. From this finding, it could be predicted that place cells would be more likely to be unstable during the nOL experiment if they were in close proximity to the objects, indicating that even if global remapping does not occur, partial remapping could be prevalent in certain areas of the open field.

Another study by Deshmukh and Knierim highlighted how these unstable fields located close to the objects could actually be tracking the objects, suggesting that place cells not only encode the overall context through stable place field representations, but also encode specific object locations through moveable place field representations (Deshmukh and Knierim, 2013). Rats were trained to randomly forage for cereal in an open field without objects and once the environment was familiar two different types of object location manipulation were used. First, four objects were added, and after these object locations became familiar one object was moved to a novel location. In another experiment only one object was used, changing location three times over three different sessions. This study highlighted many different types of place field movement in direct relation to object movement, albeit in a very small number of the overall place cell population. A number of fields in close proximity to the objects were shown to track the object location when the object moved to a different location within the same environment. Some cells even appeared in the old object location, either after the object moved or after the object was removed from the environment entirely, suggesting that a subset of place cells expressed fields that tracked previous object locations.

The second aim of this experiment was therefore to investigate the movement of place fields in relation to the objects present. It was predicted that place fields would undergo the most change between the sample and probe session if in close proximity to the object that moved between these two sessions. For example place fields could be object bound, firing near the object location in the sampling session, as well as near the objects new location in the probe session. It is also possible that place fields would appear near objects when they were introduced during the sampling phase of the protocol and then remain in these locations once one object moved to a novel location, serving to 'trace' the old object location, or appear only in the probe session near the old object location, serving to signal that an object had moved.

Deshmukh and Knierim showed one such place field that traced the old object location and whose activity persisted for more than a day, a pattern seen previously in object location responses in both the lateral entorhinal cortex (Tsao et al, 2011) and the anterior cingulate cortex (Weible et al, 2012). This suggests some indication of long-term object location memory expressed at the level of the individual place cell in the CA1. Other studies have shown that object location memory can also be seen at the level of cell assembly coordination (Manns and Eichenbaum, 2009). Similar to the study by Burke et al (2011) described earlier, rats were trained to run around a circular track for food. Once the track was familiar, objects were added one at a time, with the rat running three laps of the track between every addition. Every time an object was added the objects already on the track were moved to novel locations. The exploration time by rats of each object decreased over the three repeated track runs, only increasing again when the object was moved to a new location for the next set of three track runs. Manns and Eichenbaum showed that the firing of cell assemblies was coordinated to both object location and identity simultaneously. A decrease in the coordination of cell assembly firing coincided with increased exploration of the object, when the object was moved to a new location. This increase in exploration indicated that the animal remembered where the object had been and therefore that the object was now in a novel location. The decrease in cell assembly coordination suggested that the underlying place field population activity also reflected this change.

Although these experiments suggest that changes in place field firing or firing locations could be acting as a neural correlate for long-term object location memory, only two studies linked these changes in place cell properties with an nOL behavioural read out of memory strength. Larkin correlated the increase in place cell firing during the probe trial with an increase in nOL memory expression, and Manns and Eichenbaum correlated the decrease in cell assembly coordination with the increase in object exploration when the object was moved. However neither of these studies had a condition that expressed no nOL memory to act as a control. Therefore it was only assumed that any findings were related to memory strength, but the memory and the changes in neuronal firing were not directly correlated. The third aim of this experiment was therefore to compare any such neural correlates with the novelty preference scores calculated from implanted animals that were recorded during the nOL behavioural task. This would also allow me to directly test the hypothesis that reducing visual interference after spatial learning enhanced nOL memory through underlying place field mechanisms. It could be predicted that animals in the red box condition would detect novelty in the probe trial to a greater extent than animals in the white box condition, and therefore would exhibit higher levels of remapping between the sampling and probe sessions. However previous studies suggest that the movement of only one object might not lead to global remapping of place fields. Therefore I predicted that animals in the red box condition would show higher levels of place field stability in areas not containing objects, which would lead to an overall increase in place field stability compared with the white box condition during the probe trial. I also predicted that a higher proportion of place cells would express fields that 'traced' the object location in the sampling session, remaining near the location of the object during the sampling sessions, even after the object was moved to a new location during the probe trial. This is because 'trace' cells have been previously associated with place cell "memory" (Deshmukh and Knierim, 2013).

The overall aims were therefore to investigate how the addition, movement and removal of novel objects within a familiar environment modulated place cell properties, such as stability or tracing of the object location, and whether any of these properties were altered by the reduction of visual interference after spatial learning, therefore acting as a neural correlate for nOL memory.

4.2 Methods

4.2.1 Animals

Of the eight animals used in the recording experiment described in Chapter 3, four carried onto this experiment. Eight additional male Lister Hooded rats were obtained from Charles River laboratories, weighing between 300-350g before surgery. During the experimental protocol the twelve animals were aged between 3-12 months. Before surgery all animals were housed in groups of four. After surgery, animals were individually housed to prevent damage to the implanted drives. All cages had tubes and chew blocks for enrichment. Animals were kept on a 12-hour light/dark cycle, with testing and recording always performed in the light phase of the cycle. Before surgery animals were granted ab lid food and water. Post-surgery, once animals had recovered to pre-surgery weights and a minimum of 7 days after surgery, animals were put on food-maintenance and kept at 90-95% free-feeding body weight. Animals were therefore given 25-30g standard lab chow each per day, and free access to water. All procedures complied with the UK Animals (Scientific Procedures) Act (1986) and the European Communities Council Directive of November 24, 1986 (86/609/EEC). All animal experiments were carried out in compliance with protocols approved by the University of Edinburgh Animal Welfare and Ethical Review Board (AWERB), and under a UK Home Office Project License.

4.2.2 Surgery

Pre-surgery habituation was implemented as described previously (Chapter 3). This included habituation to random objects in a large arena and habituation to aspects of the recording protocol (such as eating cereal whilst sitting on the experimenter's lap). The same type of Kubie moveable microdrives with 8 tetrodes were implanted into the CA1 of the hippocampus (AP: -3.5mm, ML: -2.4mm, DV: -1.7mm) using the same surgery protocol as described in Chapter 3. Animals were left to recover for at least a week until pre-surgery body weight had been reached before recording procedures commenced.

4.2.3 Environmental Contexts

Two different square contexts were used. These had the same width (65cm) but different heights (60 cm and 40 cm). Contexts used had different coloured walls, and different coloured and textured floors. Each context had different cues surrounding the context, including curtains to change the size and colour of the surrounding environment, and many different 3D cues hung from these curtains or attached to the walls of the context. This

included a black and white striped cue card on the north wall of one context and a purely white cue card of bigger dimensions on the north wall of the other context. As the contexts were always in the same room, animals were transported to each context in a bucket. This served to decrease visual cues of the environment surrounding the context, giving the illusion of separate rooms.

4.2.4 Equipment

A 32-channel Axona USB system (Axona Ltd., St. Albans, UK) connected to a head stage amplifier, commutator and pre-amplifier were used to record single unit activity and local field potentials, and track the animal's location, as previously described in Chapter 3. After amplification, the signal was bandpass filtered between 300-7000Hz before being processed and recorded by specialist software (DACQ software: Axona Ltd., St. Albans, UK).

Objects were selected of similar height or width (approximately 10 x 10cm), but varying textures, colours and shapes. It was ensured that no objects had faces or pictures of animals that could have elicited an innate preference or anxiety response. The reflectiveness of objects was also taken into consideration so the tracking software only recorded the LED lights on the animal's headstage. Objects were affixed to clear glass bases (7 x 9cm), which could be screwed into the floor of the open field for stability during exploration. Before each sampling and probe trial, objects were cleaned thoroughly with alcohol disinfectant spray to remove dirt and any residual scents. The animal's behaviour was recorded using an overhead camera that fed into a DVD recorder. The same dark (red) and light (white) boxes used in the previous experiments were used for the post-sampling period. The plastic cages within these boxes were used as the rest boxes. Animals were recorded in these rest boxes before and after the behavioural trials.

4.2.5 Recording Procedure

The recording procedure was carried out as described previously (Chapter 3). In short, animals were plugged into the head stage amplifier via a millmax connection strengthened with tape, and placed into a holding bucket, while recording parameters were optimised (gains, reference channels etc) in the DACQ software (Axona Ltd., UK). Animals were then gently placed in the open field and allowed to explore. Chocolate cereal was thrown into the open field only when objects were not present. This was to ensure food placement did not bias exploration of the objects, affecting the behavioural memory read-out. This meant that

exploration of the environment did not always reach 100% in the object-containing trials; trials were removed from analysis if any session had less than 85% coverage. This led to one animal being removed from analysis for all four repeats of the protocol, and one rat removed from analysis for one repeat of the protocol.

4.2.6 Habituation

Prior to surgery, the animals were allowed to explore large arenas containing objects, 10 minutes per day for 3 days. This served to acclimatise animals to exploring novel objects by themselves, and to enhance exploration in the upcoming behavioural experiments. These objects were not used in the behavioural sessions. Animals were also habituated to aspects of the recording protocol (such as eating cereal whilst sitting on the experimenter's lap). Following surgery, they received further exposure to the same large arena as before to ensure full habituation to novel objects. Animals were then habituated to the context used in the behavioural sessions for at least 4 days, 10 minutes per day in each context. As two different contexts were used, animals were habituated to each context directly prior to its use, i.e. they were habituated to the first context just before the first two repeats of the experiment, and then the second context just before the second two repeats of the experiment. During these context habituation sessions, the animals were plugged into the recording equipment, so that these sessions also acted as screening sessions for cells. Cell screening procedures were carried out as described in Chapter 3, ensuring that all animals had electrode tips in the CA1 pyramidal cell layer, and that pyramidal cell activity could be detected prior to the onset of testing. After at least four of each animal's context habituation/screening sessions, animals were habituated to both the light (white) and dark (red) boxes for an hour. The animal was unplugged from the recording cable and recordings were not made during the white and red box sessions.

4.2.7 Behavioural testing: Experimental Design

On test days, animals followed a protocol depicted in Figure 4.1. First, the animal was plugged into the recording headstage and cable, and placed into a bucket so that the gains could be optimised. Then they were gently placed in the rest box, and recorded for 10-minutes whilst in quiescence. The recording from this session was used to identify the population of pyramidal cells that were being recorded that day (as not all cells would necessarily be active or express place fields in all 5 behavioural sessions). After this rest session the animal was placed into the holding bucket for 3-5 minutes, after which they were gently placed into a

corner of the open field (one of the two environmental contexts to which they had been habituated) with their snout facing the corner. During this initial 10 min baseline session (B1), there were no objects present, and the rat was allowed to foraging for cereal. Next, they were placed into the holding bucket for 3-5 minutes, allowing time to clean the context. The animal was then placed in the same corner of the same open field as before for a second 10 min baseline session (B2), followed by the holding bucket for another 3-5 minutes. While the animal was in the holding bucket, the context was cleaned and two identical objects were secured in two locations of the context - both were 15 cm from the corner of the box, and located at adjacent corners. The animal was placed back into the context in the same corner as before for a 10 min sampling session (S) and allowed to explore the novel objects, this time without foraging for cereal. After the sampling session the animal was unplugged from the recoding system, and transferred, in the bucket, to either a light (white) or dark (red) box, where they were kept awake via gentle handling for an hour (described in Chapter 2). After an hour, the animal was transferred back to their home cage situated in the holding room of the lab for 5 hours, where they were allowed to sleep. 6 h after the initial sampling session, once the animal's microdrive was plugged back into the headstage amplifier, the animal was placed into the same context and corner with the same two objects as before for a 10 min probe session (P). This time one of the objects had moved to a novel location (one of the previously un-occupied corners of the box, so that the two objects now occupied diagonal corners). If an animal expressed novel object location memory, they would be expected to explore the novel object location significantly more than the old object location during the first three minutes of this probe session. After the probe session the animal was transferred back to the holding bucket for 3-5 minutes whilst the objects were removed and the context was cleaned, followed by being placed back into the same context and corner as before, this time without objects, for a final baseline session (B3). This was followed by one last recording session in the rest box, to check that the same population of cells was being recorded.
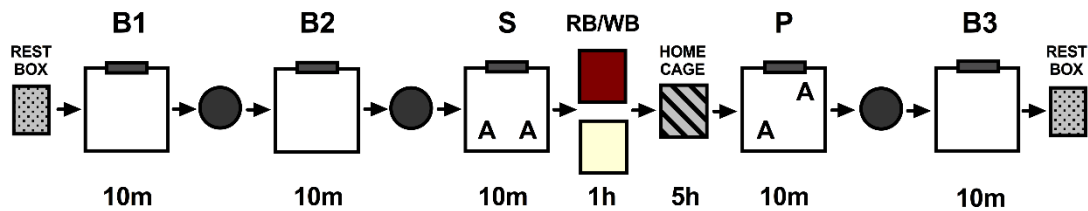
Figure 4.1: Experimental protocol: Animals were recorded during rest sessions (rest box) for 10m before the behavioural protocol commenced. A holding bucket (black circle) was used in between sessions in the open field for 3-5m, to allow time for cleaning and moving objects.

This protocol was repeated four times for each animal using two different contexts. The first and second repeats were conducted in the same context, and used the same sampling object locations, but different objects and a different novel location in the probe session. The third and fourth repeats were conducted in the second context. These used the same sampling object locations (different to the sampling locations used the first and second repeats), but different objects and a different novel location in the probe session. This ensured that the novel probe trial object locations were always completely novel within a context. There was a minimum of three days between repeats using the same context, to decrease the possibility that animals would remember the objects locations used in the previous repeat. When animals moved onto the third and fourth repeats in a different context they underwent at least 4 days of 10-minutes of habituation to the new context before the protocol commenced. This ensured that the context and surrounding environment was familiar and that place fields were stable for the two baseline sessions (B1 and B2) within the protocol. All object locations, post-sampling conditions, and objects were counterbalanced between animals.

This protocol sought to investigate place cell properties, such as stability or firing rates, during the first two baseline sessions without objects, and how this was impacted by the addition, movement and removal of novel objects. This could then be correlated to the discrimination index of probe trial exploration – a behavioural readout of 6h nOL memory.

4.2.8 Analysis and Statistics

4.2.8.1 Behaviour Analysis

Videos of the sampling and probe sessions were manually scored using specially written software (zScore), as described in Chapter 2. Only the first 3 minutes of the probe session were scored. zChop was used to calculate the discrimination index at each 10 second

cumulative bin for the probe session. The discrimination index is a ratio that indicates the animal's preference for either the familiar or novel object location, with the total exploration time factored into the following equation:

$$discrimination\ index = \frac{exploration\ time\ of\ novel\ object\ location\ -\ exploration\ time\ of\ familiar\ object\ location}{total\ object\ exploration\ time}$$

Again, all 10 minutes of the sampling session were scored to ensure animals had no innate location preference, and that neither condition explored significantly more or less than the other. All scoring was completed blind to object novelty and condition.

Data analysis was performed in Microsoft Excel, and further statistical testing and graphing of data was carried out in GraphPad Prism (version 7.0, Graphpad, USA). The final discrimination index bin encompassing the whole three-minute probe trial was analysed, as before, with both one-sample tests within conditions (comparing to chance performance) and paired t-tests between conditions (non-parametric tests were used if appropriate). Two-way ANOVAs were used to test for effects of object position or condition on sampling exploration time. Differences in sampling and probe exploration time between conditions were also analysed.

4.2.8.2 Place Cell Analysis

Place cell analysis was carried out exactly as described in Chapter 3. Once spike sorting was completed, rate maps were produced and various levels of analyses were undertaken for every session. As before, this included firing rate properties (average place cell firing rate over the session and peak in-field firing); spatial properties (spatial information, spatial sparsity, spatial selectivity and spatial coherence); general place field properties (number and area); and stability of place cells (median correlation; shuffle analysis; percentage stability and rate remapping). The same cell-inclusion parameters were used as before. Cells with unacceptable cluster isolation quality (IsoD <7; L ratio >4) were excluded throughout all analysis. Cell assembly analysis was also conducted as described in Chapter 3 to calculate both the PCorr and PCo values for every session, adjacent session and condition. PCorr analysed the synchronicity of cell assembly firing within a session, the average of which could be compared between conditions (red vs white box) as well as between sessions to explore

to the extent of synchronised cell assembly activity in each adjacent session. PCo analysed the correlation of the synchronicity of individual cell assemblies between sessions to explore whether the same cell assemblies were active and synchronous between sessions.

4.2.8.3 Object and place field movement analysis

Various methods were implemented to analyse how place fields move in reaction to object movement. For all of these analyses, the open field was split either into quadrants or into two areas: near an object or not near an object. To do this, the position data for the open field was split into an 8x8 64-bin array, as shown in Figure 4.2. Each quadrant was therefore 32.5 cm x 32.5 cm (sixteen 8.125 x 8.125 cm bins), and the object bases measured 15 cm in diameter, meaning that place fields were at maximum 15.5 cm away from the object within each quadrant. The four central bins of each quadrant were taken as being near an object (highlighted in blue in Figure) with the object centered within these areas. Each 'object-containing area' measured 16.25 x 16.25 cm the place fields had to be situated within 4cm of an object location to be included. The twelve outside bins of each quadrant were not near an object.
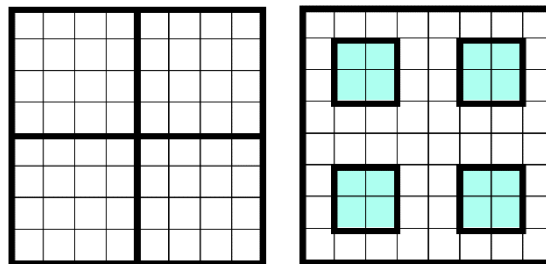
Figure 4.2: The open field split into an 8 x 8 64-bin array, and then grouped into quadrants (left) or near an object (right, highlighted in blue) and not near an object (right, shown in white).

Percentage of Place Fields Near Objects

The simplest analysis split the quadrants into two conditions: containing, or not containing, an object. Each place field was assigned to a bin defined by their weighted centre of mass. The number of place fields in bins within object containing quadrants, and the number in bins within quadrants not containing objects was counted for each rat in each session (sample and probe). A percentage was calculated by dividing the number of place fields in object containing bins by the total number of place fields for each rat and session. This was repeated for the bins in the area near an object compared to bins in the area not near an object. These analyses disregarded any object movement.

Percentage of Place Fields in Relation to Object Movement

The next analysis was very similar, but took into account object movement. Quadrants were labelled depending on object presence during sampling and probe trials: quadrants with no objects, with a stationary object, with a moving object, or with an appearing object. From this a percentage of place fields in bins within each quadrant was calculated for each session, as had been done for the last type of analysis. This was repeated for bins within the areas near or not near objects, with the same labels as the quadrants: no object (not near object area), stationary object, moving object or appearing object. The different conditions are summarised in Figure 4.2. These were calculated for all sessions, as the type of quadrant or object area was consistent throughout all object-containing sessions. This meant that a baseline of place field locations before the introduction of objects (B1 and B2 sessions) could be compared to sessions when objects were introduced. Any changes between baseline sessions and object containing sessions would suggest that place field locations were modulated by object movement.
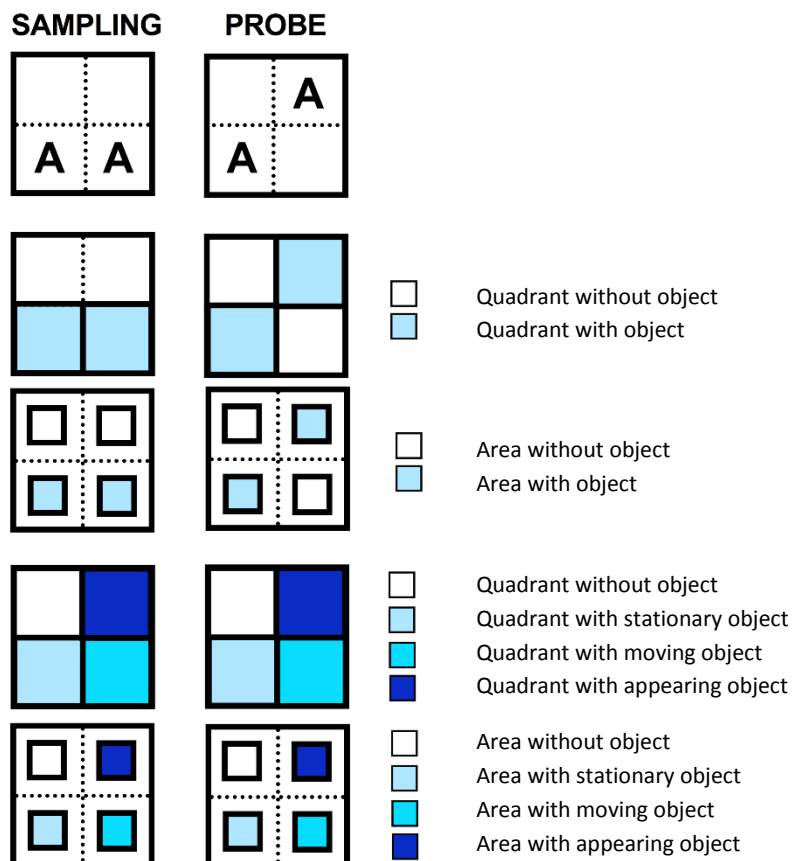
Figure 4.2: How the open field was divided for analyzing the percentage of place fields near object quadrants or near objects (top) or the percentage of place fields in relation to object movement (bottom).

153

Types of Place Field Movement

The next type of analysis sought to calculate in exactly what way the place fields for each cell were changing in regards to object movement. Place cells were again assigned bins within the 64-bin array, acting as coordinates, with a maximum of four fields per cell. These field-containing bins were assigned to a quadrant (no object, stationary object, moving object, appearing object). The number of place fields in each of these quadrants was then compared between sessions where object movement was apparent (B2 to sampling, sampling to probe, and probe to B3). This was repeated for object areas (area without an object, area with a stationary object, area with a moving object and area with an appearing object).

Figure 4.3: Different types of place field movement possible between baseline and sampling sessions. The green dot represents a single place field. The column on the left represents the open field during baseline and the column on the right represents the open field during sampling. 'A' represents novel objects.

These comparisons showed the overall change in place field locations between two sessions as before, but this time these changes were classified as fields appearing (an overall net gain in fields in that quadrant or object area), fields disappearing (an overall net loss of fields in that quadrant or object area) or the number of fields remaining stable (no net change in fields in that quadrant or object area). This avoided tracking an individual place field, an impossibility with this data set. These types of change were compared between quadrants and object areas in the same session and in adjacent sessions, producing a number of possible types of overall place field movement in response to different types of object movement (summarised in Figure 4.3, Figure 4.4 and Figure 4.5).

Baseline to Sampling Place Field Movement

When comparing the second baseline session and sampling session (B2 to S) four types of place field movements of interest were classified (shown in Figure 4.3). Fields could either appear where objects appeared (a net gain in the stationary and moving object quadrants or areas), disappear when objects appeared near them (a net loss in the stationary and moving object quadrants or areas), appear in a location that wasn't near an object (a net gain in the no object and appearing object quadrants or areas) or disappear from a location that wasn't near an object (a net loss in the no object and appearing object quadrants or areas).

Sampling to Probe Place Field Movement

When comparing the sampling and probe sessions, where one object was moved, nine types of place field movements were classified, summarised in Figure 4.4. Fields could be object-bound and follow the moving object (an overall net loss of place fields from the moving object quadrant or area, and an overall net gain of place fields in the appearing object quadrant or area), trace the object that moved (at least one field in the moving object quadrant or area in both sampling and probe trials), appear in the novel object location (a net gain of fields in the appearing object quadrant or area without a net loss in the appearing object quadrant or area), misplace the moved object (a net gain in the moving object quadrant or area), be disrupted by the object moving (a net loss in the moving object quadrant or area without a net gain in the appearing object quadrant or area), be disrupted by the object appearing (a net loss in the appearing object quadrant or area), disappear from a quadrant not containing an object in either sampling or probe trials (a net loss in the no object quadrant or area), appear in a quadrant not containing an object in either sampling or probe trials (a net gain in

155

the no object quadrant or area), or remain stable within the quadrant or area that did not contain an object (at least one place field with no net gain or loss in the no object quadrant or area).

Figure 4.4: Different types of place field movement possible between sampling and probe sessions. The green dot represents a single place field. The column on the left represents the open field during sampling and the column on the right represents the open field during the probe session. 'A' represents novel objects.

Probe to Baseline Place Field Movement

When comparing the probe and final baseline session (P to B3) where both objects disappeared, six types of place field movements were classified, summarised in Figure 4.5.

156

Trace old object location
Field next to appearing object in probe remains in this location after object has been removed

Sampling position trace
Field tracing sampling object location remains when objects are removed

New field near old object locations
Field appears when objects are removed in old object containing location

New field away from old object locations
Field appears when objects are removed in old non-object containing location

Disrupt by object removal
Field next to object in probe trial disappears when object is removed

Lost field away from old object locations
Field disappears when objects are removed in non-object containing location
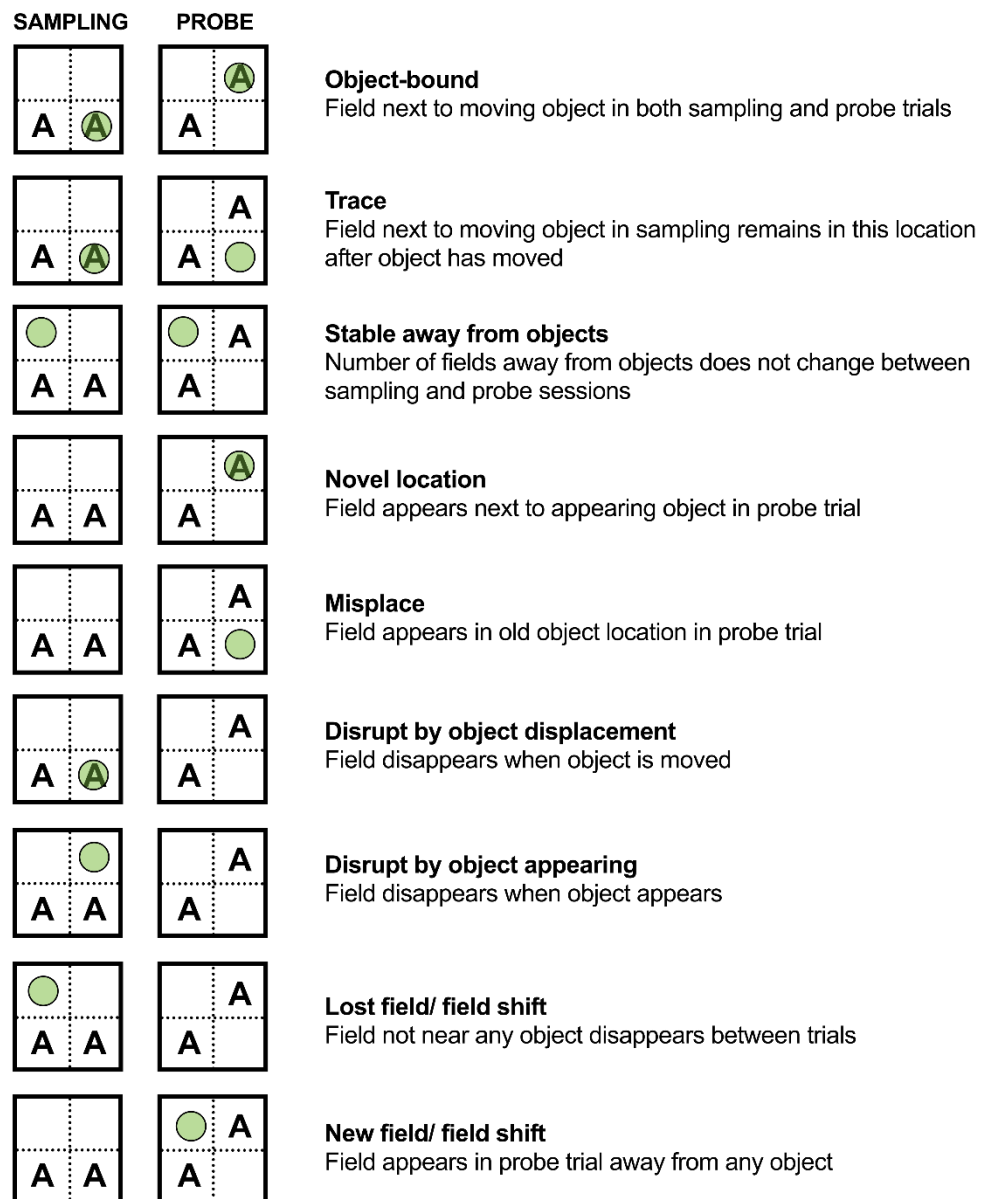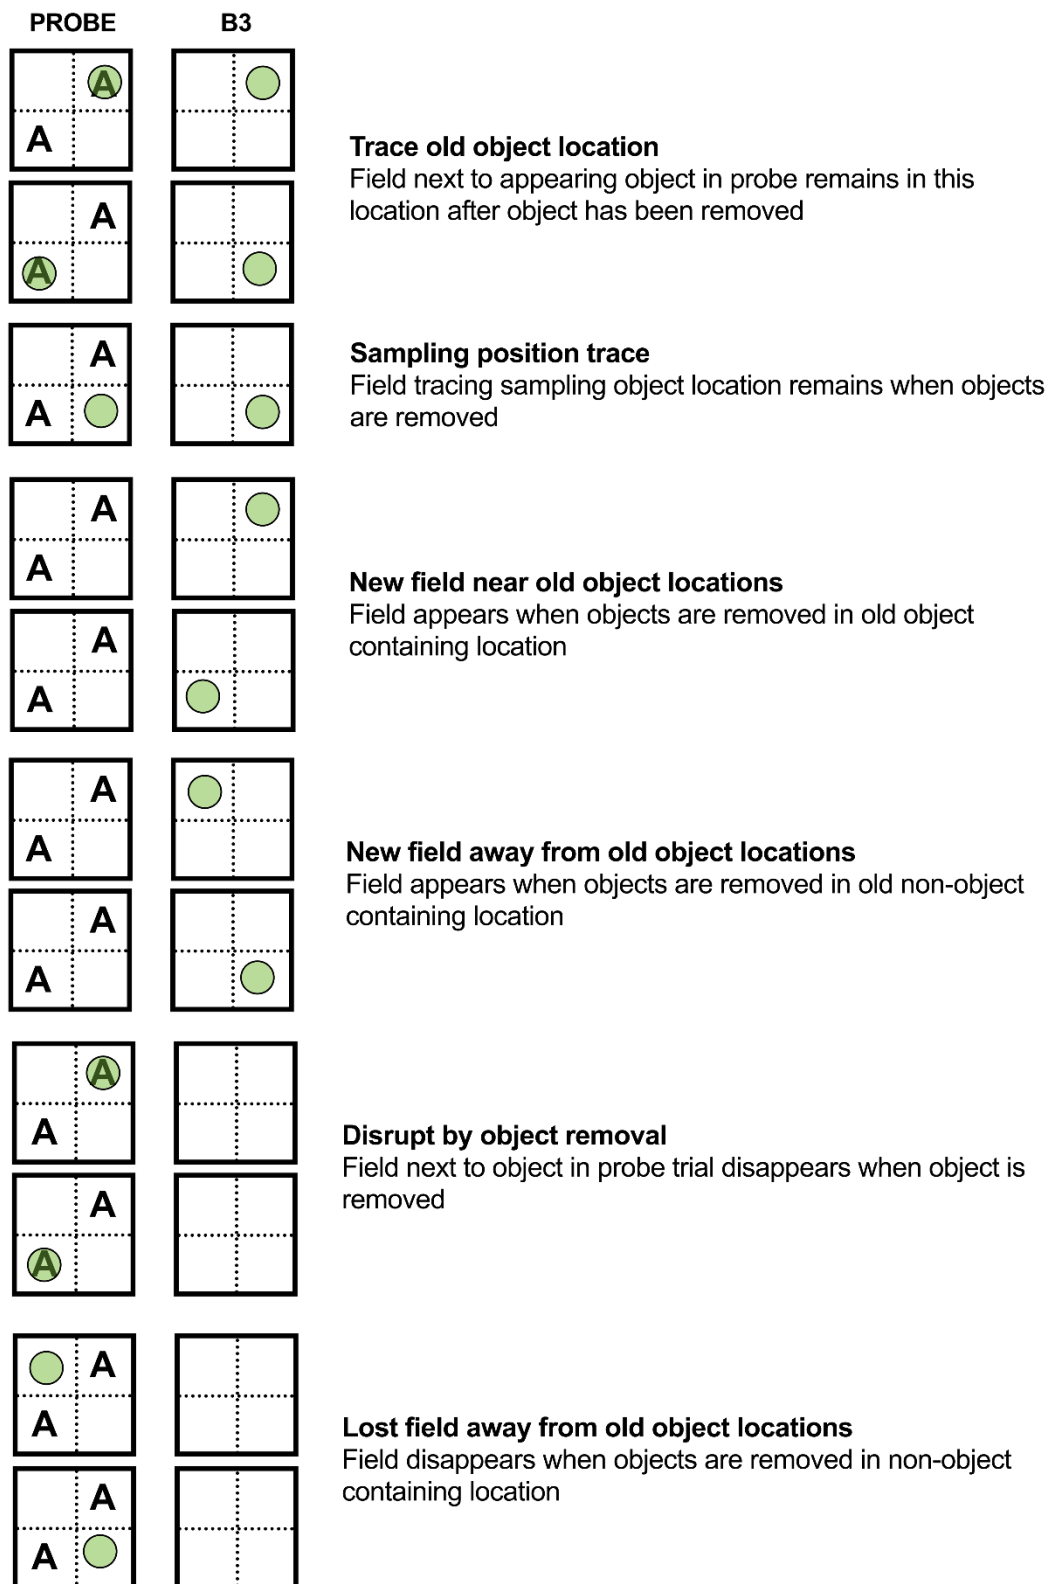
Figure 4.5: Different types of place field movement possible between probe and baseline sessions. The green dot represents a single place field. The column on the left represents the open field during the probe session and the column on the right represents the open field during baseline (B3). 'A' represents novel objects.

157

Fields could either trace where the objects had been (at least one field in the appearing object quadrant or area in both probe and B3 trials, or at least one field in the stationary object quadrant or area in both probe and B3 trials), trace where the object was during sampling (at least one field in the moving object quadrant or area in both probe and B3 trials), appear where the objects were during the probe trial (a net gain in either the appearing or stationary object quadrants or areas), appear away from where the objects were during the probe trial (a net gain in either the no object or moving object quadrants or areas), be disrupted by the removal of objects (a net loss in either the appearing or stationary object quadrants or areas) or disappear from quadrants or areas that didn't contain objects during the probe trial (a net loss in either the no object or moving object quadrants or areas).

These analyses could highlight whether place fields were more likely to appear or disappear in relation to object movement, and whether there were any types of place field movement that appeared more in the red box condition than the white box condition suggesting that the movement of place fields could underlie the strength of nOL memory expression.

4.2.8.4 Statistical Analysis

Statistics were carried out as described in Chapter 2 for behavioural analysis, and Chapter 3 for place cell analysis. Data analysis was performed in Microsoft Excel, and further statistical testing and graphing of data was carried out in GraphPad Prism (version 7.0, Graphpad, USA) and SPSS Statistics (version 22.0, IBM, USA). Where appropriate, two-way repeated measures ANOVAs were used, followed by post hoc multiple comparison tests with Bonferroni corrections to assess significance between time-points or conditions. One-way ANOVAs, as well as two-tailed paired and unpaired t-tests, were also also. Normality was calculated using a Shapiro-Wilk normality test to ensure parametric and non-parametric tests were used correctly. Pearson's correlations were used to correlate nOL memory discrimination indices with place cell measures, on a rat-by-rat, trial-by-trial, and discrimination index vs place cell measure basis.

4.2.9 Perfusion and Histology

Animals were perfused with formalin as described in Chapter 3, and their brains were collected and stored in 4% formalin for at least a week before sectioning. Brains were sliced into 40μm sections using a cryostat-microtome, mounted onto polysine slides and stained with cresyl violet. Once dry, sections were viewed with a light microscope under 2.5x magnification to identify the electrode tracks

4.3 Results

**4.3.1 Novel object location behavioural results**

TI first examined the behavioural performance of the rats to determine whether the behavioural effect shown previously, where reducing visual stimulation after spatial learning (the red box condition) led to memory enhancement in the nOL task, was replicated in this cohort of animals. This was compared to a condition with higher levels of non-novel visual stimulation (the white box condition). In both cases animals were socially isolated during exposure to the red or white box. There were several changes between the procedure used in the behavioural experiments reported in Chapter 2, and those used in the current cohort. First, the current cohort of animals was implanted with microdrives, and had been single housed since the time of surgery. Second, while the initial difference in memory between the white box and red box conditions was observed with a 24h retention test (reported in Chapter 2), in the current cohort I used a 6h retention interval. A shorter retention interval was used to reduce the risk of losing cells between the sample and probe trials, as electrode drift is more likely to occur over longer periods. A between subjects protocol was also used, whereby every rat was tested in both the red box and white box conditions, so cells from the same animal were recorded under both conditions.
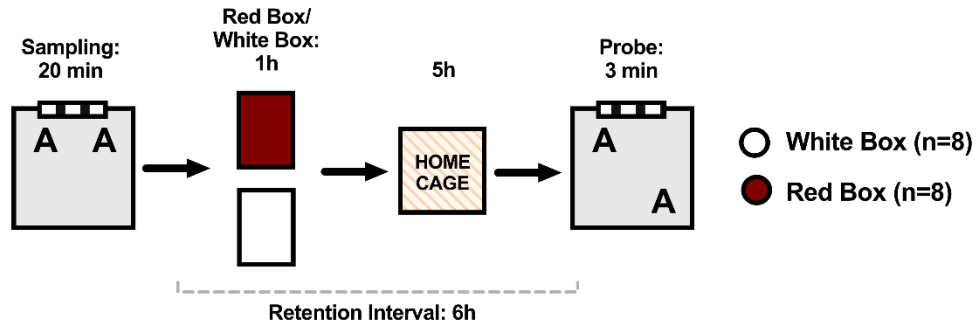
A. Novelty preference over 3 min probe trial
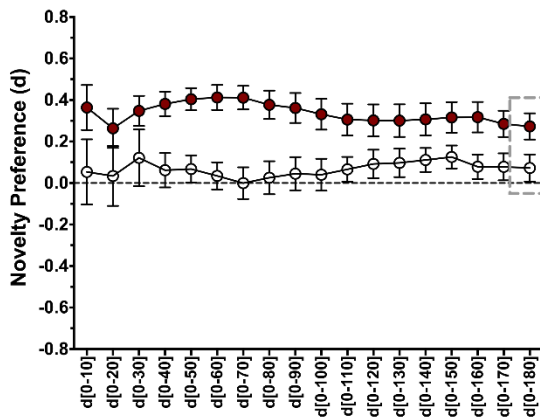
B. Novelty preference at 3 min
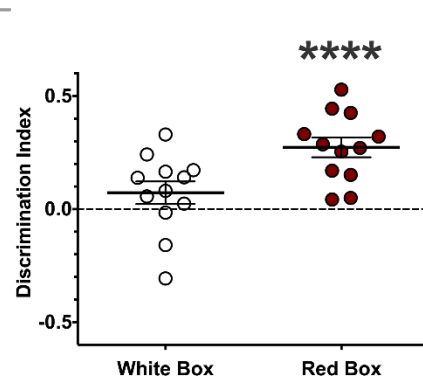
C. No location preference: sampling
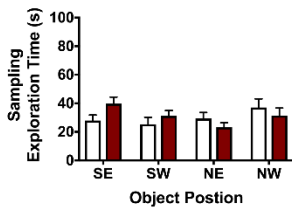
D + E . Total exploration: sampling and probe
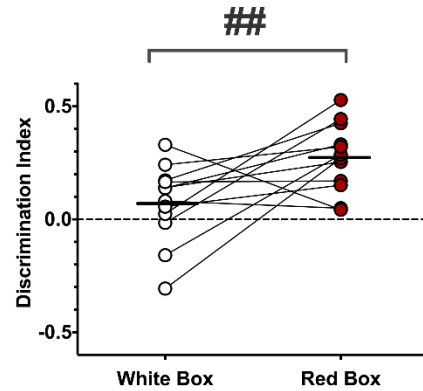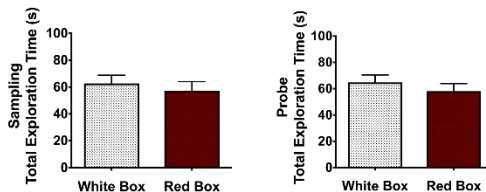
Figure 4.6: Reducing interference from social interaction and light after learning promotes retention of 6h long-term object location memory in implanted rats. Reducing interference from social interaction and familiar visual stimulation does not. A) Accumulative novelty preference over the entire 3 minute probe trial B) The 3 minute mark is highlighted, showing significant 6h memory in the red box condition, but not the white box condition (top) and a significant difference between the top conditions (bottom). C) No location preference shown in sampling for either condition. D and E) Total exploration time for sampling and probe trials not different between conditions. Error bars represent SEM. **** one-sample t-test p<0.0001; ## paired t-test p<0.01.

Panels A and B of Figure 4.6 show the discrimination index for all animals. Animals underwent two trials in each condition. Therefore every data point for one animal is the average of both trials. Panel A is an overview of the entire 3 min probe trial in 10 s cumulative bins. This indicates that when animals were exposed to the red box condition for 1h directly after learning they explored the novel object location more than the old location for the entire duration of the probe trial. In contrast, when the animals were exposed to the white box they showed no preference for the novel object at any time point. This difference is highlighted in panel B (top), showing that across the whole 3 min probe time point animals in the red box condition had significant preference for the novel object location (comparison between discrimination index and chance: t=6.240, df=11, p<0.0001) with a discrimination index value of 0.27. In contrast, the animals in the white box condition showed no preference for the novel object at any time point, with a discrimination index value of 0.07 at three min (comparison between discrimination index and chance: t=1.448, df=11, p=0.1755). This difference between conditions is highlighted in panel B (bottom), where the same animals expressed significantly enhanced memory when in the red box condition compared to when in the white box condition (t=2.701, df=11, p=0.0206). This effect was not driven by differences in total exploration time for sampling (t=1.001, df=11, p=0.3382) or probe trials (t=1.020, df=11, p=0.3296). There were also no preferences for object location for either condition during sampling [no main effect of location: F(3,86)=1.558, p=0.2054; no interaction between location and condition: [F(3,86)=0.1914, p=0.1333], shown in Figure 4.6 (panels D + E). Overall this shows that reducing visual stimulation after spatial learning can enhance 6 h nOL memory. This replicates our previous findings at 24 h, and indicates that implanted animals also benefit from reduced interference.

4.3.2 Place Cell Analysis

Given these results showing that when tested in the red box condition animals expressed nOL memory, while when tested in the white box condition they did not, the next step of this experiment was to analyse firing properties of the place cells recorded during this nOL behavioural protocol. These analyses aimed to investigate how place cells fired in a more familiar environment, how this firing changed with the introduction and movement of objects, and if any of these properties correlated with the expression of nOL memory. Unless otherwise stated, measures did not significantly correlate to the expression of nOL memory.

4.3.3 Histology

Electrode tracks were identified in all of the animals analysed in the dorsal CA1 of the hippocampus. The coordinates of the electrode tracks ranged between AP: -3.6mm to -3.2mm and ML -3mm to -2.4mm, with average coordinates of AP: -3.34 ML: 2.63. Images of these tracks in each animal can be seen in Figure 4.7 below.
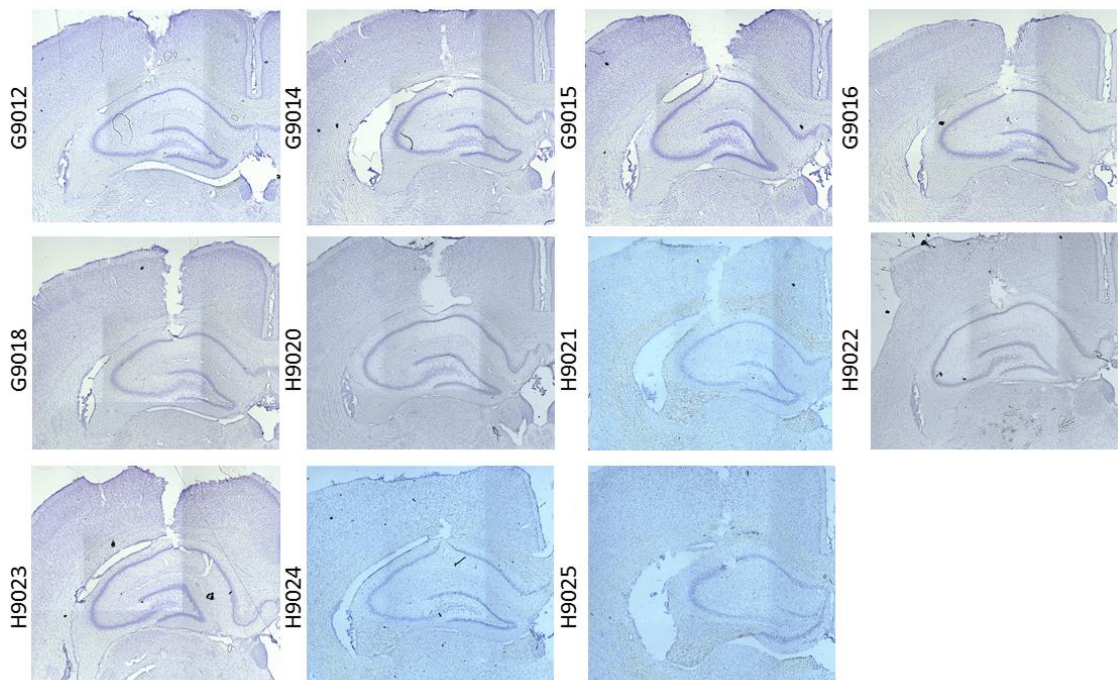
Figure 4.7: Histology showing electrode track positions for all of the eleven animals included in place cell analysis.

162

4.3.4 Cell inclusion: Isolation Distance and L Ratio

Overall, 1073 cells were included for analysis (WB=523; RB=550), with an average number of 48 (WB= 47; RB=49) and median number of 42 (WB=43, RB=39) per animal. There were no differences in cluster quality as measured by the isolation distance (U=152660, p=0.1744) or L ratio (U=153954, p=0.2616) between conditions. There were also no differences in the percentage of cells included in each cluster quality (t=2.98e-10, df=8, p>0.9999), shown in Figure 4.8. The exact numbers of cells included in each band are shown in Table 4.2.
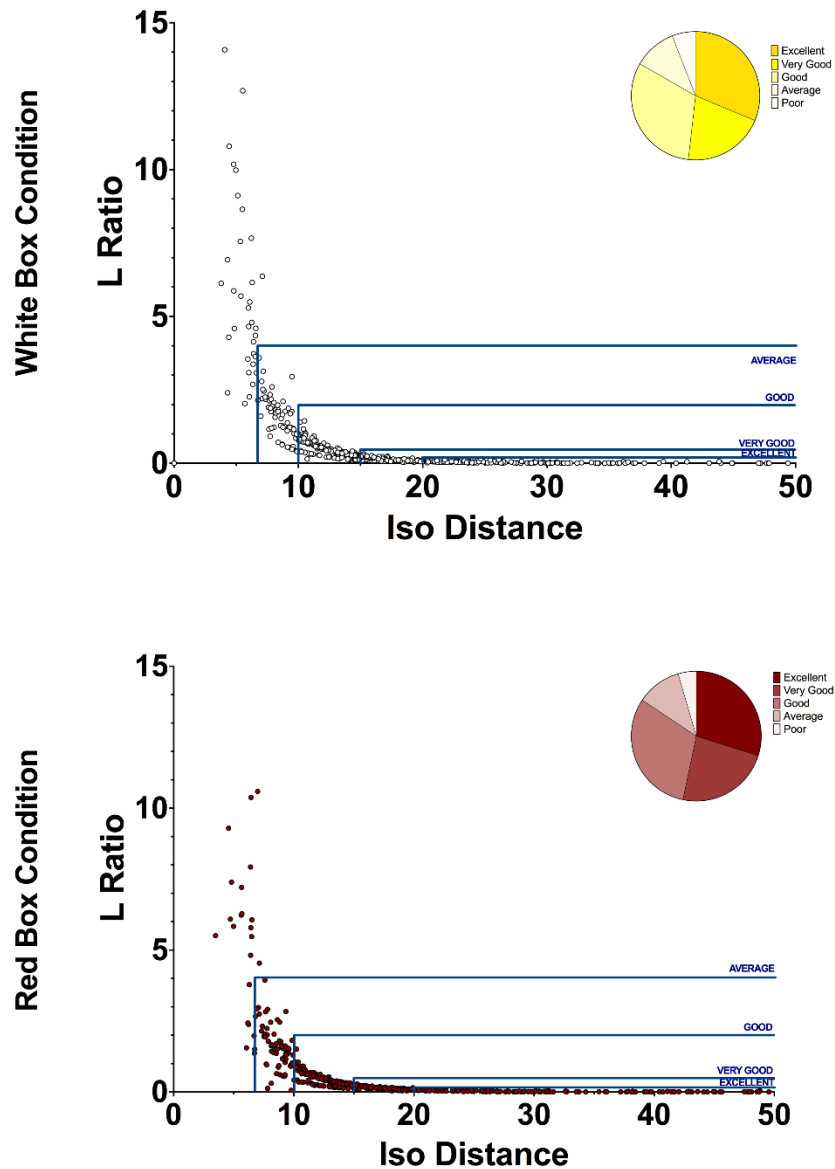
Figure 4.8: Scatter plots showing isolation distance plotted against l ratio for both the white box condition (top) and the red box condition (bottom). The blue dissecting lines represent the four different cluster quality cut-offs chosen: average (Iso D>7, l ratio <4), good (Iso D>10, l ratio <2), very good (Iso D>15, l ratio <0.5) and excellent (Iso D>20, l ratio <0.1). Anything that was below the average cut-off was excluded from analysis. The pie charts show the proportion of cell qualities for each condition.

	White Box Condition	Red Box Condition
Unacceptable (IsoD <7; L ratio >4)	**556** (33)	**576** (26)
Acceptable (IsoD >7; L ratio <4)	**523** (60)	**550** (64)
Good (IsoD >10; L ratio <2)	**463** (175)	**486** (179)
Very Good (IsoD >15; L ratio <0.5)	**288** (114)	**307** (135)
Excellent (IsoD >20; L ratio <0.1)	**174** (174)	**172** (172)

Table 4.2: Number of cells included in each band of cluster quality (exact number in individual band).

4.3.5 Firing Properties

Previous studies have shown that place cell firing rate indiscriminately increases when novelty, such as the addition of objects to a known environment, is detected (Larkin et al, 2014) or increases specifically in fields at learnt locations (Komorowski et al., 2009b). To this end, we investigated both the average place cell firing rates and peak in-field firing rates of cells over all nOL behavioural sessions to analyse how firing rate changes in response to novel objects in a familiar environment, and whether any changes are correlated to the enhancement of memory expression seen only in the red box condition.

Average Place Cell Firing Rate

The mean place cell firing rate changed significantly over the five sessions [$F_{(4,40)}$=4.108, p=0.0070], as shown in Figure 4.9 (left), which could indicate that firing rate increases over time in a familiar environment or that it increases in response to changes in the environment. Sampling, probe and the final baseline session all included manipulations of objects in some way, so this increase over time could be seen as the generalised novelty signal suggested in Larkin et al (2014). However further analysis showed that there were no significant increases in firing rate between adjacent sessions, with only B1 to probe (post hoc comparisons with bonferroni corrections: p=0.0123) and B1 to B3 (post hoc comparisons with bonferroni corrections: p=0.0421) changing significantly, suggesting that this increase in firing was minimal. There were also no differences between the conditions [$F_{(1,10)}$=2.719, p=0.1302] or interactions between session and condition [$F_{(4,40)}$=1.954, p=0.1203]. This suggests that whilst the increased mean firing rate of cells could be in relation to changes of object location within the environment as hypothesised, this difference is not correlated to the strength of nOL memory expression.

Peak In-Field Firing Rate

Unlike the mean place cell firing rate, the peak in-field firing rate did not change significantly over the five different sessions [$F_{(4,10)}=0.2179$, $p=0.9269$], as shown in Figure 4.9 (right). This indicates that the peak firing rate was not modulated by the presence or the movement of objects in a familiar environment. This could suggest that in relation to overall firing rate, the peak in field firing rate actually decreases over sessions. This is in direct contrast to the study by Komorowski et al (2009) which reported increases in peak in field firing rates at important locations. There were also no differences between the red and white box conditions [$F_{(1,10)}=0.243$, $p=0.6327$] and no interaction between the sessions and the conditions [$F_{(4,40)}=0.597$, $p=0.6669$], suggesting that the overrepresentation of object locations by increased place field firing was not underlying the enhancement of nOL memory as hypothesised.
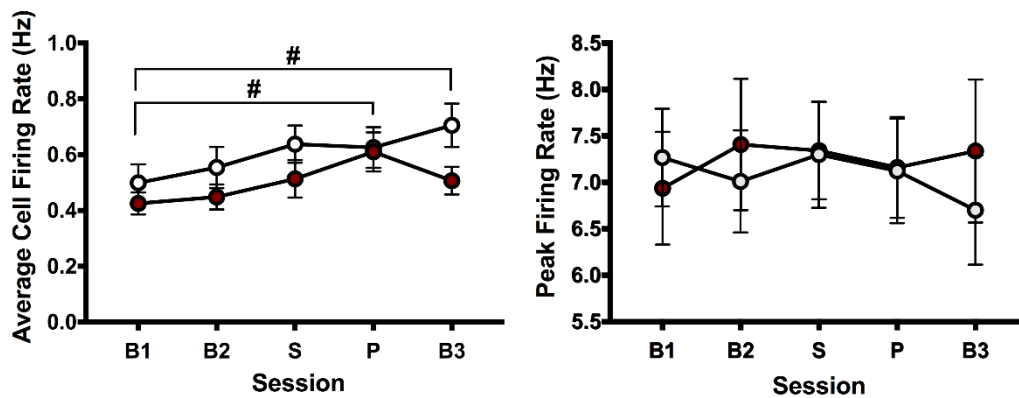
Figure 4.9: As expected, the average place cell firing rate (left) increased over sessions, possibly in relation to object movement occurring in sampling (S), probe (P) and baseline (B3) sessions. Peak in-field firing rate (right) did not change over sessions as had been expected. There were no differences between the red box (red dots) and white box (white dots) conditions, suggesting that firing rate does not correlate to spatial memory as has been shown previously. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: # $p<0.05$.

These results show that although overall firing rates increased as predicted, suggesting that some form of novelty response could be occurring, these were not affected by reducing visual interference after learning and therefore not correlated to the strength of nOL memory. This is in direct contrast to the findings by Larkin, however possibly emphasises the need for a comparative condition with no memory expression to reach accurate conclusions. Interestingly, both mean firing rates and peak in-field firing rates were much lower than those recorded in a completely novel environment during the experiment described in Chapter 3. This highlights that novelty of the environment can lead to much higher levels of place cell firing than the novelty of objects within a familiar environment.

4.3.6 Spatial Properties

It has been shown that whilst spatial properties of place cells are likely to be improved in a familiar environment compared to a novel environment (Brandon et al., 2014), the introduction of objects to a familiar environment, or movement of objects within a familiar environment, does not often lead to changes in spatial information or sparsity (Kyd and Bilkey, 2005), although some studies have reported decreases in spatial coherence upon the movement of objects within a familiar environment (Lenck-Santini et al., 2005). To investigate this further, and determine whether the introduction and movement of objects could alter the spatial firing properties of place cells, various spatial properties were analysed across sessions of the nOL protocol.

Spatial information

Spatial information, a measure of the spatial content carried by place cell firing, showed a significant interaction between session and condition [$F_{(4,40)}=3.987$, $p=0.0082$]. Further post hoc comparisons revealed that although neither group showed any differences in spatial information between adjacent sessions, the red box condition showed significantly lower spatial information carried by place cell firing in the probe session (WB vs RB post hoc comparisons with bonferroni corrections: $p=0.0149$), shown in Figure 4.10 (top left). This suggests that if an animal expresses object location memory during the probe trial then their place cell firing will carry less spatial information. However when nOL memory scores were compared with the average spatial information carried by place cell firing in a given session and for a given rat there was no significant correlation between the two. Overall these results suggest that objects and object location memory can modulate spatial information of place cells, however it is not clear in what way spatial information changes throughout the sessions. It would be interesting to investigate whether spatial information changes differently in different parts of the environment. Although not tested, it is plausible that variations in spatial information could be a result of place cell firing near objects carrying different levels of spatial information than place cell firing away from objects.

Spatial sparsity

Spatial sparsity, a measure of how confined the firing field of a place field is compared to the rest of the environment, changed significantly over sessions for both conditions [main effect of session: $F_{(4,40)}=7.384$, $p=0.0001$]. As shown in Figure 4.10 (bottom left), spatial sparsity

appeared to increase when objects were introduced, decrease when one object was moved, and increase again when they were both removed. This increase in sparsity between baseline and sampling sessions implies that place field firing became more indiscriminate throughout the environment when objects were introduced, a pattern that was nearing significance (B2 to S post hoc comparisons with bonferroni corrections: $p=0.0551$). Further analysis showed the increase between the probe last baseline sessions was significant (P to B3 post hoc comparisons with bonferroni corrections: $p=0.0230$). There were no differences between conditions [$F_{(1,10)}=0.297$, $p=0.5977$] and no interaction between session and condition [$F_{(4,40)}=1.413$, $p=0.2473$] indicating that spatial sparsity did not relate to nOL memory strength.

Spatial Selectivity

Unlike spatial sparsity, spatial selectivity did not change between sessions [main effect of session: $F_{(4,40)}=0.6509$, $p=0.6296$]. This is shown in Figure 4.10 (top right). These results suggest that the confinement of spikes to the firing within a place field is not affected by the presence of objects. There was again no difference between conditions [$F_{(1,10)}=0.001379$, $p=0.9711$], and although the interaction between session and condition was trending towards significance [$F_{(4,40)}=2.181$, $p=0.0886$] no further paired t-tests between conditions were significant for any session comparison. These results indicate that spatial selectivity is not modulated by the strength of nOL memory.

Spatial coherence

Spatial coherence showed clear differences between session for both conditions [main effect of session: $F_{(4,80)}=14.29$, $p=<0.0001$]. This suggests that object manipulation can change the consistency of place field firing. Coherence very clearly decreased when objects were introduced and increased when they were moved and when they were removed, shown in Figure 4.10 (bottom right). This change was most evident between baseline and sampling sessions (B2 to S post hoc comparisons with bonferroni corrections: $p<0.0001$). The difference between sampling and probe sessions was also significant (S to P post hoc comparisons with bonferroni corrections: $p=0.0162$) indicating that moving objects actually increased the spatial coherence of place field firing. The further increase in coherence between the probe and final baseline sessions was significant (P to B3 post hoc comparisons with bonferroni corrections: $p=0.0479$). Whilst results suggest that the presence and

movement of objects can affect place field coherence there were no differences between conditions [$F_{(1,20)}=0.9213$, $p=0.3486$] or interactions between session and condition [$F_{(4,80)}=0.3627$, $p=0.8345$]. As with all the spatial properties analysed, this indicates that the levels of interference post-learning and the strength of nOL memory produced did not modulate the spatial coherence of place field firing.
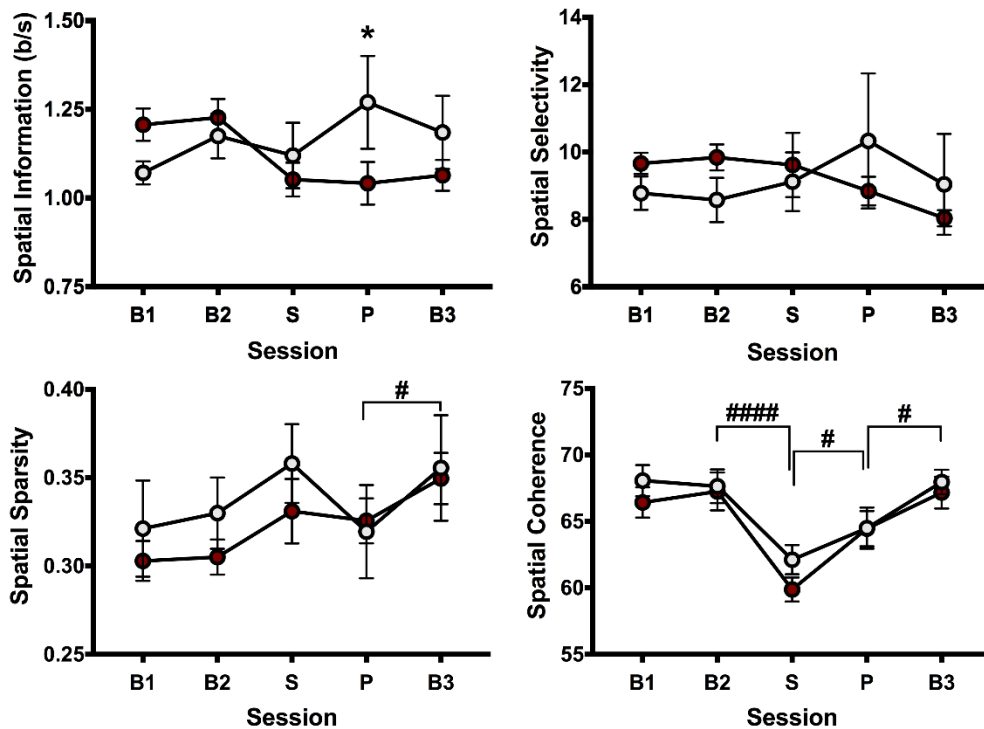
Figure 4.10: Spatial information (top left) and spatial selectivity (top right) did not change significantly over sessions as expected. However spatial information did differ between conditions, with a significant interaction between session and condition emerging as a difference between red box (red dots) and white box (white dots) conditions during the probe session. Spatial sparsity (bottom left) and coherence (bottom right) showed significant changes over sessions, with coherence showing the most evident changes. Spatial sparsity increased upon the introduction of objects (S), decreased when objects were moved (P) and increased when they were removed (B3). Spatial coherence decreased upon the introduction of objects (S), and increased when the were moved (P) and removed (B3). There were no differences in sparsity or coherence between the red and white box conditions. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: * p<0.05 between conditions at that session; # p<0.05 between sessions, #### p<0.0001 between sessions.

As predicted, the spatial property that showed the most consistent changes over sessions was spatial coherence. Coherence decreased with the addition of objects, as expected. However coherence actually increased when objects were moved, in direct contrast to the Lenck-Santini (2005) study showing that coherence decreased with object movement. In the latter study both objects were moved simultenously. As only one object was moved in the present experiment it could be suggested that coherence only decreases when all landmarks within the environment change location at once.

168

4.3.7 Place Field Stability

Median Correlation

The next set of analyses sought to investigate how place field stability was modulated by the introduction, movement or removal of objects. It was hypothesised that stability would be highest in the first two baseline sessions, as animals were familiar to the environment, and that the presence of objects would signal a change in this environment, leading to the destabilisation of the spatial map. Burke et al. (2011) found that the introduction and subsequent movement of objects led to global remapping of the place cell population, however this experiment used 8 different objects rather than the 2 used in the present experiment. Therefore it could be predicted that there would not be enough change in the environment to lead to global remapping.
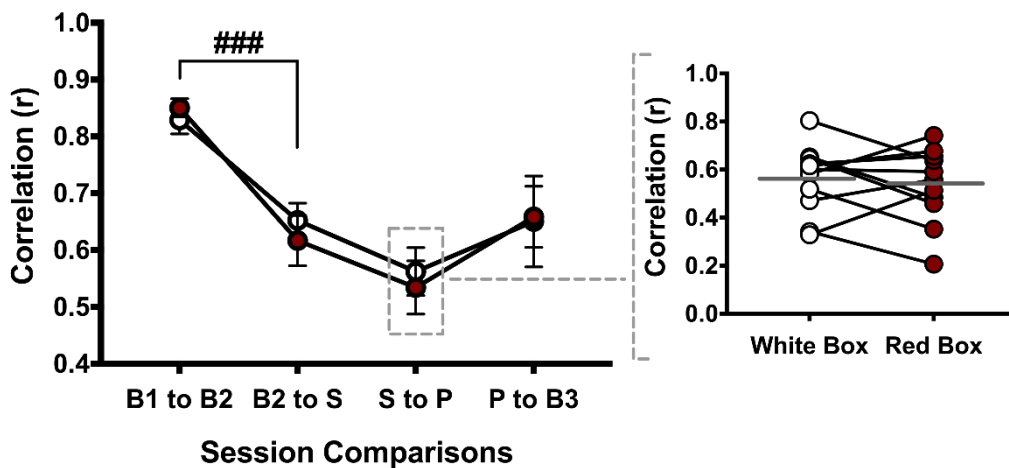
Figure 4.11: Median correlations of rate maps showed significant differences between adjacent sessions. The correlation between rate maps was highest between the two baseline sessions (B1 to B2) as expected. Correlations decreased when objects were added (B2 to S) and when one object was moved (S to P), increasing when both objects were removed (P to B3). This suggests that the addition and movement of objects can lead to partial remapping of place fields, as seen previously. Unexpectedly, there were no differences in correlations between the white box (white dots) and red box (red dots) conditions. This is highlighted for the sampling to probe comparison (right), with grey lines representing the average percentage for each condition. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: ### p<0.001

As expected, the two initial baseline sessions, B1 and B2, had very high correlation values of around 0.85. This dropped significantly to about 0.6 when objects were introduced, and further still when one of these objects was moved. When objects were removed, correlation values increased again to around 0.65, but did not reach previous baseline levels. These clear patterns are shown in Figure 4.11. There was a significant effect of session on correlation values [main effect of session: $F_{(3,30)}=17.17$, p=<0.0001]. Further analysis revealed that

correlation values decreased significantly between the two baseline sessions and the second baseline and sampling session (B1 to B2 and B2 to S post hoc comparisons with bonferroni corrections: p=0.0002) and increased between probe and the last baseline session to levels trending towards significance (S to P and P to B3 post hoc comparisons with bonferroni corrections: p=0.0994). These results show that the addition of objects to a familiar environment can destabilise the spatial map. Interestingly removing objects appeared to increase the stability of the place fields again, but not to the high levels of stability seen at the start of the experiment.

Given our previous findings indicating that reducing visual stimulation after spatial learning promotes place field stability in a novel environment, another important aim of this experiment was to investigate whether place field stability correlated with the strength of nOL memory. It was hypothesised that there would be differences between the red and white box conditions correlating with memory strength. As the red box condition expressed stronger memory for the locations of objects in the sampling trial it was possible that the subsequent movement of one of these objects would destabilise the map to a greater extent. Therefore the red box condition would show less stability between sampling and probe trials. However it was also possible that the red box condition would show greater stability between sampling and probe trials, making encoding object locations onto the underlying spatial map more reliable.

These results show that there was no difference in correlation between the white box and red box conditions for any session comparisons [$F(1,10)=0.08733$, $p=0.7736$] and no interaction between session and condition [$F(3,30)=0.2099$, $p=0.8887$]. The absolute correlation values between the sampling and probe sessions are shown in **Figure 4.12** for every cell, and the similarity of the correlation values of the red and the white box conditions between sampling and probe sessions are highlighted as they represent exposures to the environment before and after the animals experienced the white and red box conditions. These results imply that whilst the addition of objects can destabilise the spatial map, this instability is not related to the strength of memory expression. This also suggests that although a decrease in visual stimulation after a nOL sampling trial enhances spatial memory, it does not enhance the correlation between place field rate maps. It is surprising there was no difference between the two conditions given the two possible hypotheses. However it is

possible that whilst overall stability levels were not different, there could be differences in stability depending on the proximity of the field to objects within the environment.
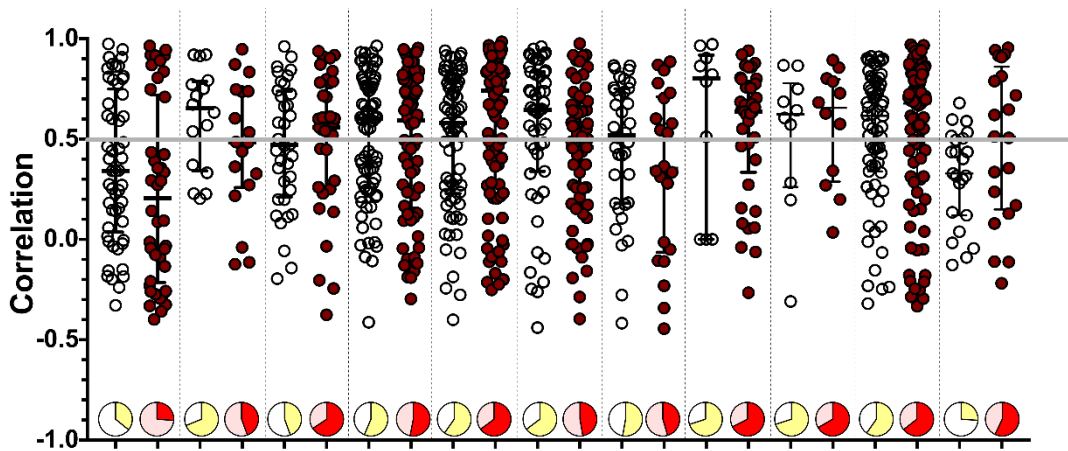

Figure 4.12: Pearson correlation values for individual cells are plotted for correlations between sampling and probe sessions. Overlaid box and whisker plots represent the median correlation values and interquartile ranges for each rat in each white box (white dots) or red box (red dots) condition. Correlations of 1 represent highly similar rate maps with stable place fields. This place field stability is also shown as a percentage of stable cells in pie charts. Every pie chart represents the percentage of stable (dark colours) and unstable (light colours) cells for each rat in each white box or red box condition. This measure of stability was calculated using a bootstrap cut-off, shown as a grey line intersecting the correlation values at 0.50 correlation. All values above this cut-off were 'stable' and all values below were 'unstable'.

Shuffle Analysis

As carried out in the previous chapter, bootstrap analysis was used to produce a stable cell median correlation cut-off, statistically splitting the cells into "stable" or "unstable". The value can be seen as the dividing line in Figure 4.12. This was another way to analyse the median correlation values.

Figure 4.13 is plotting correlation data for every adjacent session comparison (B1 to B2, B2 to S, S to P, P to B3) in a different way by grouping all the cells from all the rats in the white box condition together and all the cells from all the rats in the red box condition together to show the distribution of these correlations. This highlights how both the white box and red box conditions have distributions that are shifted drastically to the right compared to the shuffled data set for all session comparisons. These correlation values were significantly higher in white box condition (B1 to B2: U=163279, p<0.0001; B2 to S: U=478692, p<0.0001; S to P: U=615275, p<0.0001; P to B3: U=307975, p<0.0001) and red box condition (B1 to B2:

U=158170, p<0.0001; B2 to S: U=613700, p<0.0001; S to P: U=798937, p<0.0001; P to B3: U=322390, p<0.0001) compared to the shuffled data set. Consistent with the "rat-level" analysis described above (Figure 4.11), there were no differences between conditions at the cell level for any pairs of sessions (B1 to B2: U=74791, p=0.0817; B2 to S: U=92459, p=0.4071; S to P: U=107531, p=0.5423; P to B3: U=70873, p=0.5709).
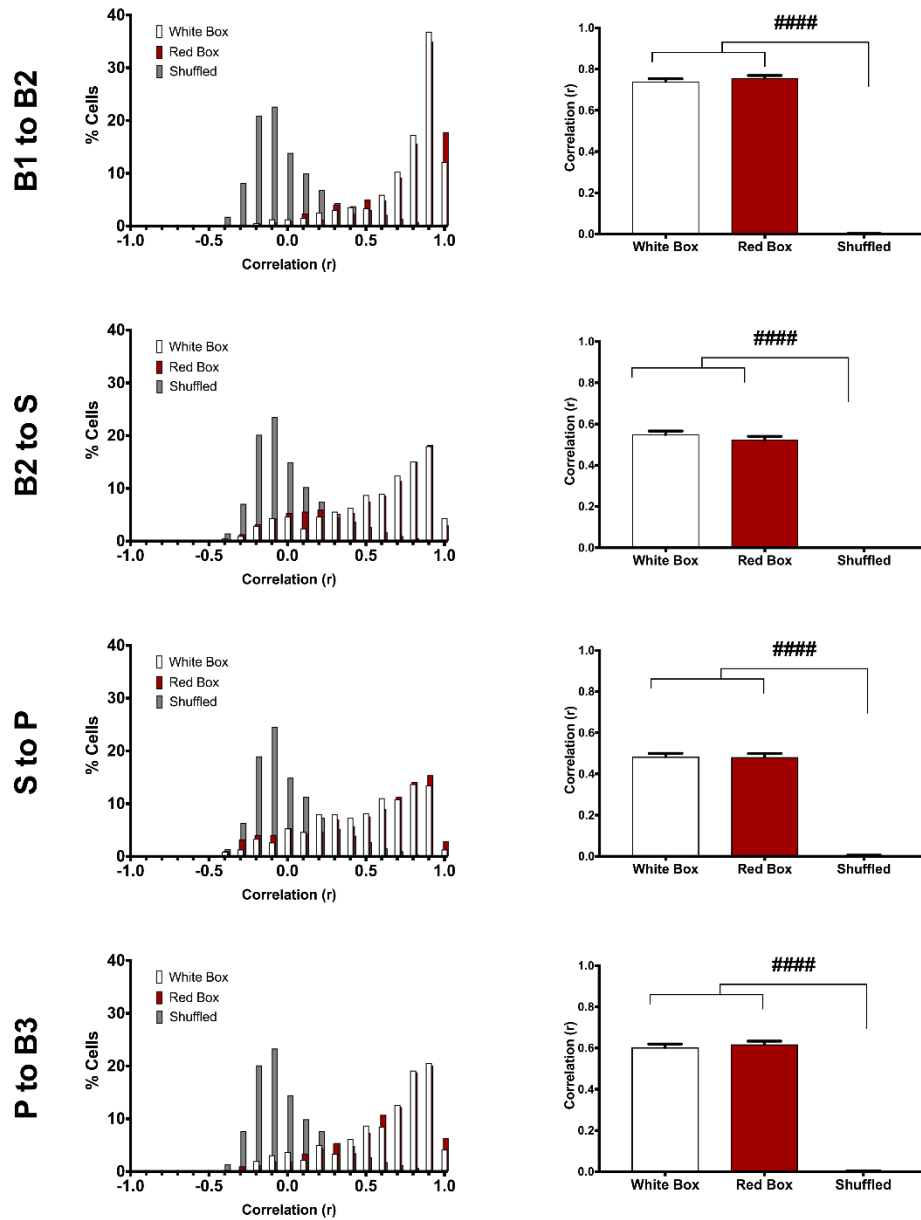
Figure 4.13: These values represent correlations of the first two baseline sessions (B1 to B2 – top), the second baseline session and the sampling session (B2 to S – second to top), the sampling and probe sessions (S to P – second to bottom) and the probe and final baseline sessions (P to B3 – bottom). Frequency distributions of correlation values from cells from the white box condition and red box condition were shifted to the right, with cells having significantly higher correlation values than those of the shuffled data set. There was no significant difference between correlations of cells from the white box and red box conditions. Error bars represent the SEM. Mann-Whitney U: #### p<0.0001

172

The 95[th] percentile values of the shuffled data were as follows: B1 to B2 = 0.53; B2 to S = 0.50; S to P = 0.47; P to B3 = 0.50. As the bootstrap values for the adjacent session comparisons were minimally different the average bootstrap value of 0.50 was used for analysis. However it should be noted that using the corresponding different bootstrap values for the different session comparisons did not change the outcome of the results (data not shown).

Percentage of Stable Cells

Based on the bootstrap values for stable vs unstable place cells calculated in the analyses described above, the correlation values for all cells were split into stable (>0.5) and unstable (<0.5), producing a percentage of spatially stable cells for each rat and condition. These are shown in Figure 4.12 as pie charts for sampling to probe session comparisons, and summarised for all adjacent session comparisons in Figure 4.14.
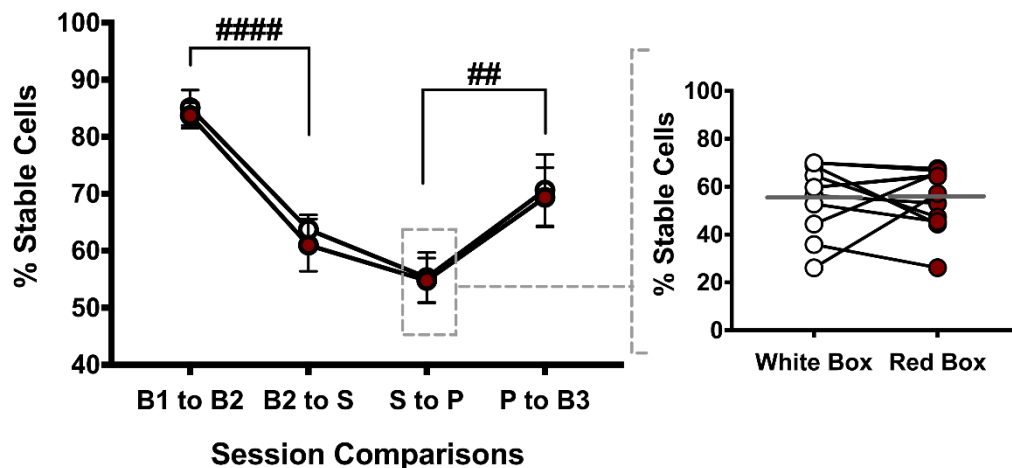
Figure 4.14: The percentage of cells with median rate map correlations over 0.5 showed significant differences between adjacent sessions. This percentage of stable cells was highest between the two baseline sessions (B1 to B2) as expected. Stability decreased when objects were added (B2 to S) and when one object was moved (S to P), increasing when both objects were removed (P to B3). This suggests that the addition and movement of objects can lead to partial remapping of place fields, as seen previously. Unexpectedly, there were no differences in correlations between the white box (white dots) and red box (red dots) conditions. This is highlighted for the sampling to probe comparison (right), with grey lines representing the average percentage for each condition. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: ## $p<0.01$, #### $p<0.0001$.

The percentage of stable cells compared to unstable cells was calculated for every adjacent session. As shown in Figure 4.14, this data almost perfectly mirrored the changes in median correlation shown above. Again, there were high levels of stable cells when comparing initial baseline sessions. This stability decreased with both the addition and movement of objects, and increased again when objects were removed. There was a significant change in

percentage stability over sessions [main effect of session: $F_{(3,30)}=20.84$, $p<0.0001$], with further analysis revealing significant decreases in stability between the two baseline sessions and baseline and sampling sessions (B1 to B2 and B2 to S post hoc comparisons with bonferroni corrections: $p<0.0001$) and significant increases between probe and the last baseline (S to P and P to H3 post hoc comparisons with bonferroni corrections: $p=0.0035$). This indicates that these changes can be seen using various measures of stability. These results also indicate that whilst the spatial map was destabilised when objects were added, this was not global remapping as almost 60% of cells still expressed stable place fields. Therefore partial remapping was occurring.

As with median correlation, there was no difference between the conditions [$F_{(1,10)}=0.4341$, $p=0.5249$] or interaction between session and condition [$F_{(3,30)}=0.03345$, $p=0.9916$]. The similarity in percentage stability between the conditions for the sampling to probe comparison is highlighted in Figure 4.14. Although a clear enhancement of nOL memory expression was seen in animals in the red box condition, the underlying stability of the cells did not appear to mediate this effect.

Rate remapping

A final measure of stability across sessions was used to investigate more subtle changes. Rate remapping is a measure of the ratio of changes in overall firing rates between adjacent sessions. Rate remapping occurs when cells are spatially stable but have different rates associated with different environments (Leutgeb et al., 2005b). Therefore if the ratio of firing rates of "stable" cells significantly changed between sessions it could indicate that the place cells were reacting to the addition of objects to a familiar environment. Only rate maps that had median correlations of over 0.5 (i.e. a "stable" place cell) were used for analysis.

There was a significant change in the levels of rate remapping of "stable" place cells over sessions [main effect of session comparison: $F_{(3,30)}=6.141$, $p=0.0022$], as shown in Figure 4.15. Levels of rate remapping were lowest between the two initial baseline sessions, as expected when recording from a familiar environment (cells therefore had the highest overlap of firing rates). Levels of rate remapping increased when objects were added, increasing further when objects were moved, and decreasing when objects were removed. This indicated that the addition of objects modulated the firing rates of individual cells. These were the same patterns of change seen in the correlation analysis, even though rate

remapping analysed only "stable" cells. However further analysis showed there were no significant changes between adjacent session comparisons. Rate remapping between the two baseline sessions was only significantly less (i.e. significantly more overlap) than rate remapping between the sampling and probe sessions (B1 to B2 and S to P post hoc comparisons with bonferroni corrections: p=0.0012). This implies that although 60% of cells remained spatially stable when objects were added, around 35% of these stable cells underwent increased levels of rate remapping.
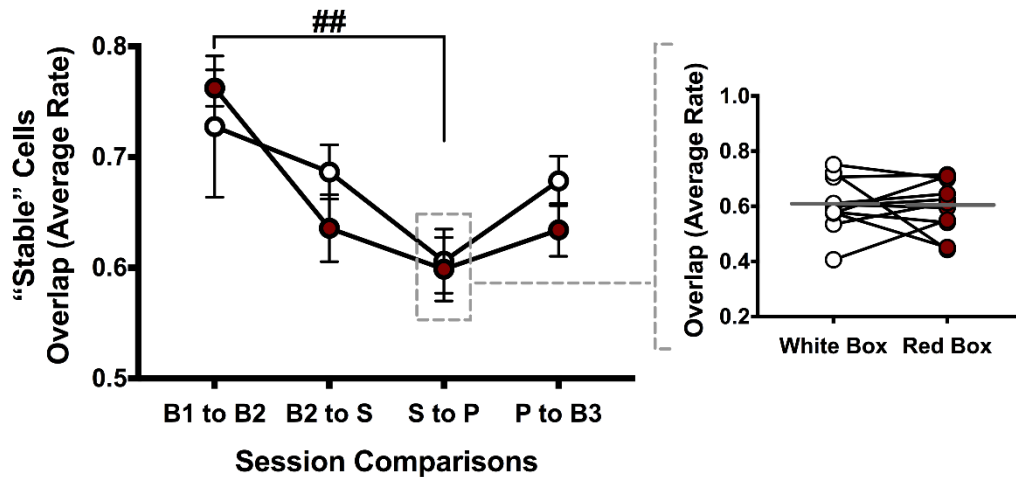
Figure 4.15: Levels of rate remapping changed significantly over sessions. As with the changes in correlation, levels of remapping were lowest between the two baseline sessions (B1 to B2) as expected. Levels of rate remapping increased when objects were added (B2 to S) and when one object was moved (S to P), decreasing again when both objects were removed (P to B3). This suggests that the addition and movement of objects can lead to rate remapping of place fields in the "stable" population of place cells. Unexpectedly, there were no differences in correlations between the white box (white dots) and red box (red dots) conditions. This is highlighted for the sampling to probe comparison (right), with grey lines representing the average percentage for each condition. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: ## p<0.01.

As with median correlation and percentage stability, it was hypothesised that levels of rate remapping would correlate to nOL memory strength. It was possible that there would be higher levels of rate remapping between sampling and probe sessions in the red box condition due to the change in object location being more noticed in this condition. However there were no differences between conditions [F(1,10)=0.6013, p=0.4560] or interaction between session and condition [F(3,30)=0.7775, p=0.5158]. This again suggests that although place cells reacted to objects there was no correlation with this modification and the expression of memory.

4.3.8 Place Cell Assembly Firing

As there were no obvious differences in place field stability at the single cell level, it is possible that the coordinated firing of groups of cells in a population could be underlying the behavioural changes seen. It has been suggested that the integration of object identity and location requires the coordinated firing of populations of place cell assemblies within a session (Manns and Eichenbaum, 2009). It has been shown that when a familiar object was moved to a novel location, the synchronous firing of these cell pairs decreased, possibly signifying a change in this object identity-object location association previously encoded by the animal. Therefore this suggests that although the synchronous firing of cell assemblies might not change within sessions (PCorr) the correlation of these patterns of synchronous firing between sampling and probe sessions (PCo) may decrease due to the discrepancy between object identity and object location. This would suggest that changes in PCo value would also relate to the strength of nOL memories and be significantly different between the two conditions, with stronger memories correlating to smaller PCo values between the sampling and probe sessions. To investigate this further both the correlation of cell firing within (PCorr) and between (PCo) sessions were calculated using three different time-bins.

Synchronous Firing of Cell Assemblies within Sessions (PCorr)

The changes in coordination of cell assembly firing over the five sessions were most striking when using the 25ms time-bin. There were clear patterns in the coordinated firing of cell assemblies in both conditions, as shown in Figure 4.16 (left). The levels of coordinated cell assembly firing remained low but consistent for the two baseline sessions (B1 and B2). Surprisingly, more cells showed coordinated firing with other cells when objects were introduced in the sampling session compared to the adjacent baseline session, which then decreased when one object was moved to a new location in the probe trial. PCorr values changed significantly over the sessions [main effect of session: $F_{(4,40)}=3.702$, $p=0.0118$], and although no adjacent sessions showed significant changes in PCorr values, the difference between the second baseline and sampling session was almost significant (B2 to S post hoc comparisons with bonferroni corrections: $p=0.0813$) and the difference between the first baseline session and the sampling session was significant (B1 to S post hoc comparisons with bonferroni corrections: $p=0.0195$). There were no significant differences between the conditions [$F_{(1,10)}=0.02601$, $p=0.8751$] or interaction between session and condition [$F_{(4,40)}=0.189$, $p=0.9428$].

When looking at results using the 125ms time-bin, these distinctive patterns of cell assembly coordination were also seen, as shown in Figure 4.16 (middle). The correlation of cell assembly firing within a session increased significantly whenever objects were introduced, decreased when they moved and increased again when they were removed. Although there were also significant changes in PCorr values over the sessions [main effect of session: $F_{(4,40)}=2.693$, $p=0.0445$], no session comparisons were significant. This indicates that this pattern is seen most strongly when using the 25ms time-bin. There were were no significant differences between conditions [$F_{(1,10)}=0.01427$, $p=0.9073$] and no interaction between session and condition [$F_{(4,40)}=0.7765$, $p=0.5470$].

When calculating PCorr using a 1s time-bin, this variance caused by session disappeared completely for both conditions [main effect of session: $F_{(4,40)}=1.249$, $p=0.3060$], as shown in Figure 4.16 (right). This indicated that the levels of cell assembly coordination did not change in relation to object movement when looking at a time-bin associated with the real-time movement of the rat through a place field. There were also no differences between conditions [$F_{(1,10)}=0.06035$, $p=0.8109$] and no interaction between condition and session [$F_{(4,40)}=1.054$, $p=0.3918$].
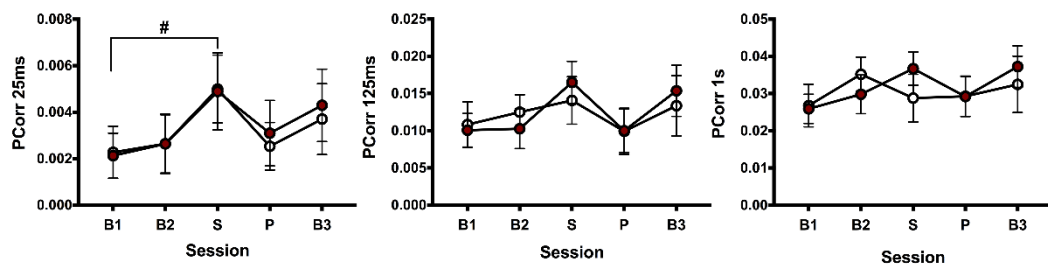
Figure 4.16: The synchronicity of cell assembly firing was lowest within the two baseline sessions (B1 and B2). This synchronicity then increased when objects were added (S), decreased when one object was moved (P) and increased again when objects were removed (B3). This was most apparent for the 25ms time bin (left), although the main effect of session was also significant for the 125ms time bin (middle). This suggests that the addition of novel objects increased the synchronicity of cell assembly firing. There were no differences between the white box (white dots) and red box (red dots) conditions. Error bars represent SEM. Post hoc comparisons with bonferroni corrections:# $p<0.05$

These results suggest that when looking at time-bins associated with internally driven oscillations (gamma – 25ms and theta – 125ms) the coordinated firing of cell assemblies were modulated by the addition and movement of objects within a familiar environment.

Interestingly the addition of objects increased synchronicity of cell pair firing, whereas the movement of objects appeared to decrease this coordination. It is possible that the increase seen in sampling could be a 'novelty' signal, signalling the presence of novel objects. As the objects should be familiar by the probe session the novelty signal decreases. It is also a possibility that such a novelty signal could be related to the encoding of object identity-object location associations within cell assemblies, as suggested by Manns and Eichenbaum. Rather than signalling a decrease in novelty, the decrease in coordination between sampling and probe sessions could relate instead to 'mismatch' signal, i.e. cell assemblies signalling the mismatch between the object identity-object location associations encoded during sampling. Unlike the results from Chapter 3, where recording from a novel environment was analysed, the levels of interference in the hours directly after initial spatial learning did not appear to affect cell assembly firing, and PCorr values were not significantly correlated to nOL memory strength.

The Correlation of Cell Assembly Firing Between Sessions (PCo)

The next step of these analyses sought to investigate the comparison of these PCorr values between sessions further by correlating every PCorr value with every other PCorr value on a rat-by-rat basis. Instead of showing the average synchronicity of cell assemblies within a session for every rat, and whether this average synchronicity changed over sessions, this analysed how the synchronicity of each individual cell pair differed between sessions. High PCo values would suggest that the synchronous cell assembly firing in one session was very similar to that of another session, low values would suggest that the temporal relationships of cell assemblies are changing between sessions. Therefore if specific cell assemblies were signalling a 'mismatch' signal when objects moved locations, one would expect the sampling to probe PCo value to be low. Figure 4.17 (top) shows the range of PCorr values, ordered from the lowest to highest values of all of the cells during the sampling session. This order of cells is used in the bars relating to the other sessions to highlight the differences in PCorr of individual cells over the three sessions. This highlights the similarities or differences of individual cell-pair firing synchronicity between sessions. PCo values are given at the top of each bar.

There were significant changes in PCo values over the different sessions for all time-bins used [main effect of session comparison: 25ms time-bin – $F_{(3,30)}=18.79$, $p<0.0001$; 125ms time-

bin – F(3,30)=18.73, p<0.0001; 1s time-bin – F(3,30)=22.51, p<0.0001], shown in Figure 4.17 (bottom). As predicted the synchronicity of cell assemblies between sessions was highest when there were no objects, and lowest when objects were moved. Coordination of cell assemblies at a population level (PCo) decreased significantly when objects were introduced (B1 to B2 compared to B2 to S: post hoc comparisons with bonferroni corrections: 25ms time-bin: p=0.0001; 125ms time-bin: p=0.0001; 1s time-bin: p<0.0001) and increased when objects were removed (S to P compared to P to B3: post hoc comparisons with bonferroni corrections: 25ms time-bin: p=0.0002; 125ms time-bin: p=0.0016; 1s time-bin: p=0.0081).

These changes could be expected at the 25ms and 125ms time-bins, as any significant differences in PCorr values between sessions would lead to a decrease in the correlation of cell assembly synchronicity (PCo) between the sessions. However, it is unexpected that these changes in PCo values remain when there are no significant changes in PCorr values between sessions, as is seen for the 1s time-bin. As described in Chapter 3 this could be attributed to an increase in variance of PCorr values. The average levels of cell assembly synchronicity (PCorr) would remain stable between sessions, whilst the coordination of individual cell assembly pairs could be very different, i.e. populations of cells could be consistently firing or not firing together depending on the object locations. Therefore it could also be that the populations of cells that are active during the sampling and probe sessions tend to be inactive during the non-object sessions.

Whilst there were clear changes in PCo values between sessions, there were no differences between the red box and white box conditions for any of the time-bins used [25ms time-bin – F(1,10)=0.1567, p=0.7005; 125ms time-bin – F(1,10)=0.0969, p=0.7620; 1s time-bin – F(1,10)=0.03194, p=0.8617] or interactions between session and condition [25ms time-bin – F(3,30)=1.968, p=0.1400; 125ms time-bin – F(3,30)=0.8862, p=0.4594; 1s time-bin – F(3,30)=0.122, p=0.9464]. There were also no significant correlations between PCo values and strength of nOL memory. This is in contention with the predicted results as it was hypothesised that the red box condition would show much lower PCo values between the sampling and probe sessions due to higher levels of 'mismatch' signalling. This suggests that although cell assembly population firing is clearly affected by the addition, movement and removal of objects, this is not underlying the enhanced nOL memory seen in the red box condition.
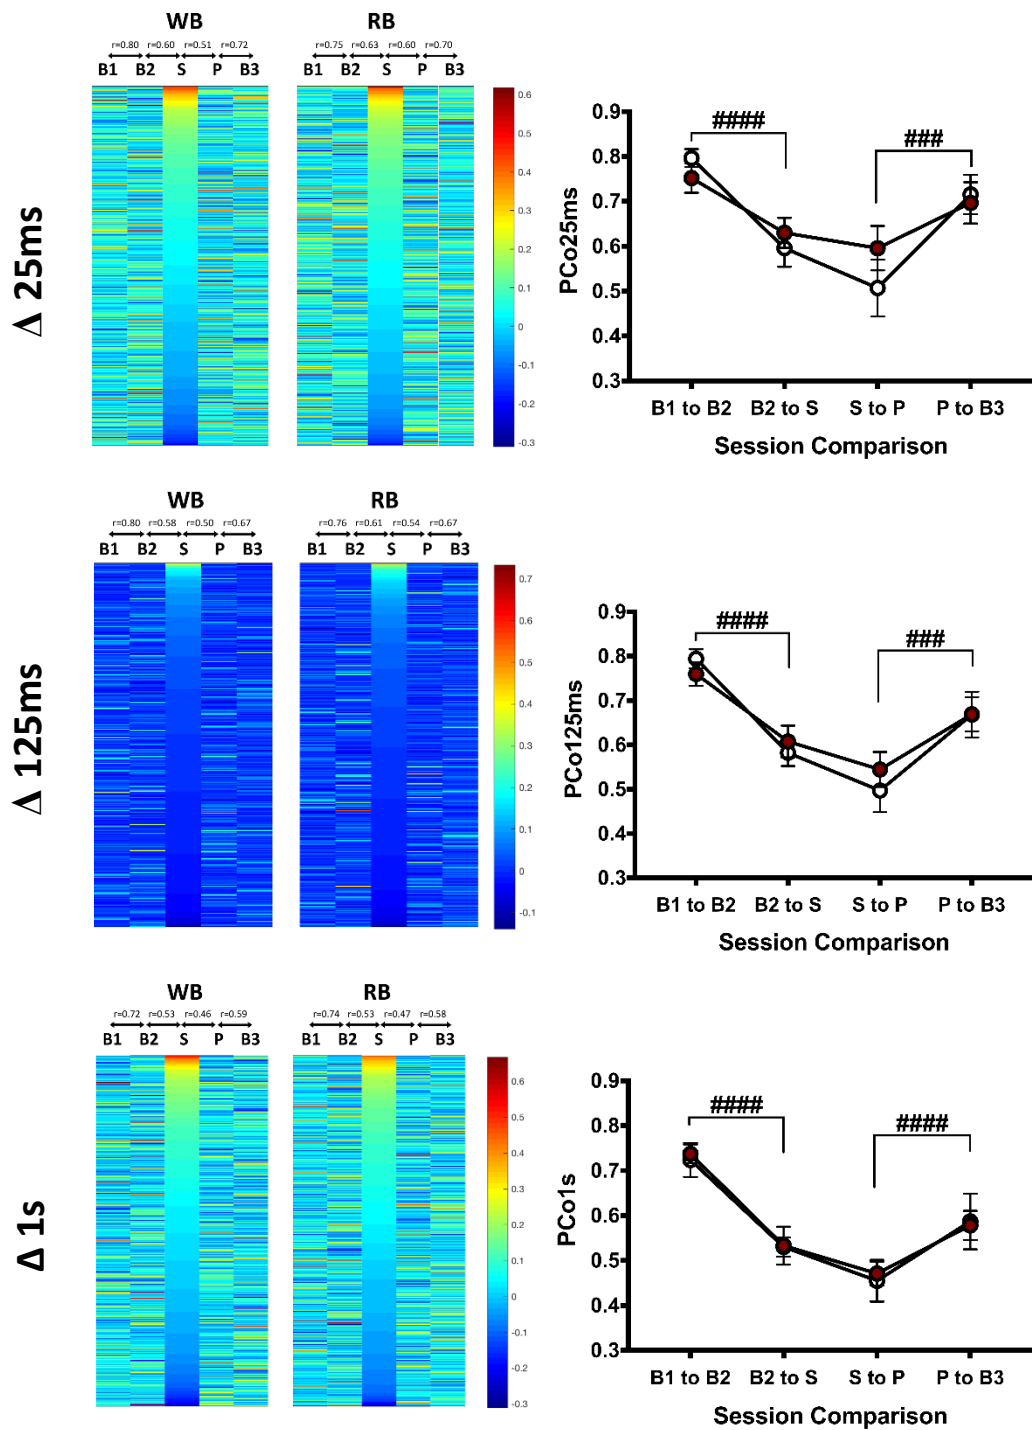
Figure 4.17: The synchronicity of individual cell assembly pairs (PCo) within rats changes significantly over the five sessions, with temporal synchronicity being most correlated between the two baseline sessions (B1 to B2), becoming less correlated when objects were added (B2 to S), when they were moved (S to P), correlation increasing again when they were removed (P to B3). These results mirror the changes in place field stability seen in this experiment. This changes can be seen in the 25ms time frame (bottom left), the 125ms time frame (bottom middle) and 1s time frame (bottom right). There were no differences between rats in the red box (red dots) or white box (white dots) conditions. Error bars represent SEM. The bars at the top represent the range of all the PCorr values from all of the rats, ordered from the lowest to highest values of all of the cells during the sampling session, for all three time frames used. This order of cells is used in the bars relating to the other sessions to highlight the differences in PCorr of individual cells over the three sessions. PCo values are given at the top of the bars. Post hoc comparisons with bonferroni corrections: ### p<0.001 #### p<0.0001.

It should be noted that these trends in PCo values also mirror the trends found in other measures of correlation such as the correlation of rate maps and stability of place fields. These results suggest that these changes in place field stability between sessions can be seen at the cell population level as well as the individual cell level. The fact that there are no differences between the WB and RB conditions also suggests that the enhanced expression of nOL memory seen in the RB condition is not expressed at a detectable level within place cells at the assembly level or the cell population level.

4.3.9 Place Field Properties

Whilst the enhanced expression of nOL memory was not detectable at the level of the place cell, it is possible that place field properties, such as size, number or movement in relation to object movement could correlate to the strength of object location memory. It has been previously shown that the size and number of place fields expressed can be modulated by the presence of objects (Burke et al., 2011; Kyd and Bilkey, 2005), with the introduction of objects leading to an increased number of smaller place fields. To investigate whether this occurred during a nOL task, and if this modulation was correlated with memory strength, both the number and the size of place fields were analysed across sessions and between conditions.

<u>Number of Place Fields</u>

As shown in Figure 4.18, the average number of place fields was modulated by object presence, varying significantly across sessions for both conditions [main effect of session: $F(4,40)=16.74$, $p=<0.0001$]. The number of place fields increased when objects were introduced (B2 to S post hoc comparisons with bonferroni corrections: $p<0.0001$) and decreased when they were removed (P to H3 post hoc comparisons with bonferroni corrections: $p=0.0056$). The percentage of place cells with more than one field also varied by about 10% across sessions [main effect of session: >1PF: $F(4,40)=18.6$, $p=<0.0001$; >2PF: $F(4,40)=6.765$, $p=0.0003$], showing the same patterns of increasing with object introduction (B2 to S post hoc comparisons with bonferroni corrections: >1PF: $p<0.0001$; >2PF: $p=0.0009$) and decreasing with object removal (P to H3 post hoc comparisons with bonferroni corrections: >1PF: $p=0.0004$). Overall this indicates that the presence of objects in a familiar environment increases the number of place fields expressed. Interestingly, during baseline sessions B1, B2 and B3, the average number of place fields per cell was around 1.4 and the

fraction of cells with more than one place field was around 30%. These are both very similar to the numbers seen for all sessions in Chapter 3, where again no objects were present.
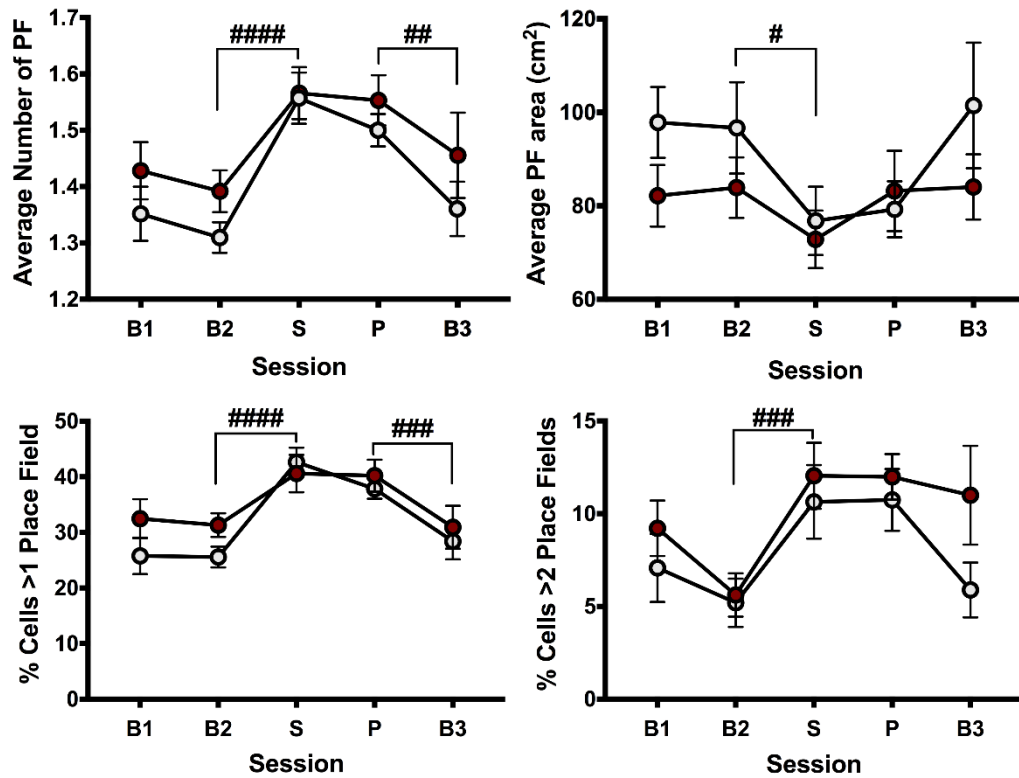
Figure 4.18: The average number of place fields increased when objects were introduced during the sampling session (S), only decreasing when objects were removed during the last baseline session (B3) (top left), as expected. This is mirrored by changes in the number of cells expressing more than one (bottom left) and more than two (bottom right) place fields. As the number of place fields increased, the size of the fields decreased (top right, sampling), and when the number of fields decreased during the last baseline session (B3) the size of the fields increased again. This suggests that place cells were reacting to the presence of objects by expressing a larger number of smaller fields, as has been found before. There were no differences between the red box (red dots) and white box (white dots) conditions, however the size of place fields during the probe session did negatively correlate to novel object location memory strength. Error bars represent SEM. Post hoc comparisons with bonferroni corrections: # $p < 0.05$, ## $p < 0.01$, ### $p < 0.001$, #### $p < 0.0001$.

There were no differences between the white box and red box conditions for the average number of place fields [$F_{(1,10)}=1.873$, $p=0.2011$] or the percentage of cells with multiple place fields [main effect of condition: >1PF: $F_{(1,10)}=1.032$, $p=0.3336$; >2PF: $F_{(1,10)}=2.038$, $p=0.1839$], and no interactions between session and condition [number of fields: $F_{(4,40)}=0.3626$, $p=0.8338$; multiple fields: >1PF: $F_{(4,40)}=0.9759$, $p=0.4315$; >2PF: $F_{(4,40)}=0.6076$, $p=0.6595$]. This suggests that whilst place cells were modulated by the presence of objects, this did not relate to the strength of nOL memory expressed.

Place Field Area

The area of place fields also changed significantly over sessions [main effect of session: $F_{(4,40)}=4.325$, $p=0.0053$], as shown in Figure 4.18. Whilst the number of place fields increased with the introduction of objects, between B2 and sampling, the area of place fields appeared to significantly decrease (B2 to S post hoc comparisons with bonferroni corrections: $p=0.0444$). This suggests that whilst more place fields appear with the introduction of objects, overall these fields are smaller. The baseline place field area during the first and second baseline sessions is a lot lower than that shown in the previous chapter, however it must be noted that the animals had received at least 4 habituation sessions to the environment prior to B1, so the environment was familiar.

There were no differences in place field size between conditions [main effect of condition: $F_{(1,10)}=3.039$, $p=0.1119$] and no interaction between sessions and conditions $F_{(4,40)}=1.25$, $p=0.3055$]. However when the average place field area for every rat for every probe session was correlated with the corresponding nOL memory score there was a significant negative correlation ($r=-0.4805$, $p=0.0011$). There were no significant correlations for the other behavioural sessions (B1, B2, S or B3), indicating that the expression of smaller place fields during the probe session was associated with better nOL memory.

4.3.10 Place Field Movement

The movement of place fields in relation to objects has been shown in a variety of studies. It appears that place fields are least stable when in close proximity to objects (Lenck-Santini et al., 2005), and that some place fields even follow objects when they are moved around an environment or trace them once they are removed (Deshmukh and Knierim, 2013). As the stability of a large population of place cells decreased when objects were introduced to the open field this indicates that place fields could be modulated by the addition, movement, and removal of objects. Although there was no overall difference in stability between the white and red box conditions at a single cell or population level, it is possible that the degree of place field-object interactions occurring could reveal differences between the two conditions.

Percentage of Fields Near Objects

To this end, initial analysis sought to calculate the percentage of place fields near an object, or within an object quadrant, compared to the total number of place fields across all of the cells recorded for a given rat in a given session. This was analysed for both sampling and probe sessions for every rat, and then compared between conditions. As shown in Figure 19 (top), the percentage of place fields within an object quadrant did not change between sessions [main effect of session: $F_{(1,10)}=0.03384$, $p=0.8577$]. For both sessions place fields were just as likely to be within or out of an object quadrant, with on average 49.41% of fields in an object quadrant in sampling and 53.78% in probe in the white box condition; and 53.55% in sampling and 48.51% in probe in the red box condition. There were also no differences between the white and red box conditions [main effect of condition: $F_{(1,10)}=0.0204$, $p=0.8893$] and no interactions between session and conditions [$F_{(1,10)}=2.418$, $p=0.1510$].

When comparing the percentage of place fields near an object, rather than in an object quadrant, the percentage was much less, as expected due to comparative area size, as shown in Figure 19 (bottom). Unlike the quadrant analyses, there was an interaction between session and condition [$F_{(1,10)}=5.177$, $p=0.0462$], indicating that the levels of interference in the hour after spatial learning might have an impact of the subsequent location of place field expression. Further post hoc comparisons revealed there were trends towards an increase in the number of fields near an object in the probe session for the WB condition (S to P post hoc comparisons with bonferroni corrections: $p=0.0529$), with an average of 11.33% of place fields near an object during the sampling session, and 19.03% during the probe session. On the other hand the percentages of place fields near an object in red box condition did not change between sessions, with 14.05% of place fields near an object during the sampling session and 12.22% during the probe session (S to P post hoc comparisons with bonferroni corrections: $p>0.9999$). As expected there were no significant differences between conditions during the sampling session, however post hoc tests revealed an almost significant increase in the percentage of fields near an object in the probe session (WB vs RB post hoc comparisons with bonferroni corrections: $p=0.0886$). When the percentage of place fields near objects in the probe sessions were correlated with nOL behaviour discrimination indices for every individual rat there was an almost significant negative correlation ($r=-0.2892$, $p=0.0600$), suggesting that animals are more likely to express nOL memory if fields are

184

expressed away from object locations in the probe session. There was no apparent correlation between the percentage of place fields near objects in the sampling session and nOL scores.
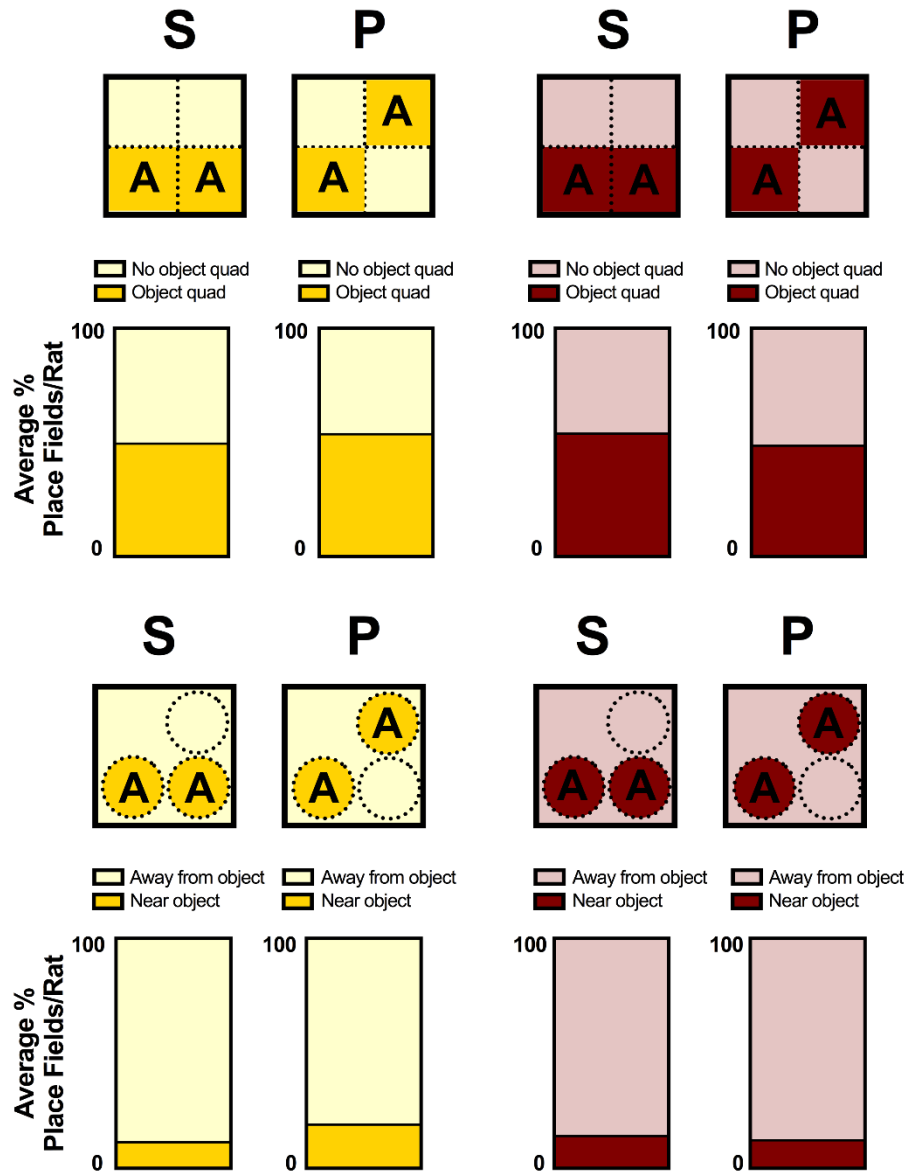
Figure 19: The percentages of place fields within object containing quadrants was not significantly different to the percentage of fields within quadrants that did not contain objects, for either the white box (top left, yellow bars) or the red box (top right, red bars) conditions. There were also no differences in these percentages between sampling and probe sessions, or between the two conditions. The percentages of place fields near an object did not appear to be significantly different between sampling and probe sessions for the red box condition (bottom right), however the white box condition showed trends towards significant increases in the percentage of place fields near objects in the probe session (bottom left). There were also trends towards significance between the number of fields expressed near objects in the white box condition compared to the red box condition during the probe trial. Further analysis revealed the number of fields expressed near objects in the probe session to be significantly negatively correlated to novel object location memory strength.

Interestingly in the condition expressing nOL memory the object-containing bins covered 12.5% of the total open field area. This suggests that overall place fields were not preferentially expressed near objects if the animal remembered the previous object locations from the sampling session, as a similar proportion of place fields (14.05%) were expressed near objects in the red box condition during the probe trial.

Percentage of Place Fields in Relation to Object Movements

Whilst the analyses just described indicate that place fields were preferentially expressed away from objects in the probe trial, the correlation between the percentage of fields and nOL memory scores was not quite significant. However this analysis did not take the actual movement of objects into account. Previous studies have found that place fields are more likely to be affected by the movement of objects when the objects are close to the place field than far away (Lenck-Santini et al., 2005), suggesting that the percentage of place fields in the vicinity of objects might be variable throughout the different sessions. To determine whether place fields were more likely to be expressed near or away from moving objects, object areas were re-labelled as not containing an object (no object); containing a stationary object; containing a moving object; or containing an appearing object, as shown in Figure 4.20. These analyses were repeated for the different object quadrants, but as before there were no differences within or between conditions (data not shown).

Again, the percentage of place fields near each object were calculated for each rat, this time for every session. Differences across sessions could indicate that the expression of fields is more changeable in the vicinity of certain objects. It should be noted that no change in percentage between sessions does not indicate no change in the numbers of place fields, as earlier it was shown that the number of fields increased when objects were introduced. Instead, the number of fields can increase whilst maintaining the same proportions.

A three-way within-subjects ANOVA showed that there was an interaction between session, condition and object [F(2.654), df=12, p=0.003]. Further analysis showed that the percentage of place fields did not change significantly between sessions in the red box condition [main effect of session: $F(4,40)=1.054$, $p=0.3920$] and that although there was an effect of object area [$F(3,30)=580.4$, $p<0.0001$] there was no interaction between object area and session [$F(12,120)=0.8596$, $p=0.5895$]. In the WB condition however, there was an effect of object

186

area [F(3,30)=605.3, p<0.0001] and an interaction between object area and session [F(12,120)=1.945, p=0.0354]. Further tests revealed that this interaction was only apparent between sampling and probe sessions for place fields expressed away from any object locations (S to P area away from object locations: post hoc comparisons with bonferroni corrections: p=0.0182). Therefore there were significantly less place fields expressed away from object locations during the probe trial compared to the sampling trial in the white box condition.
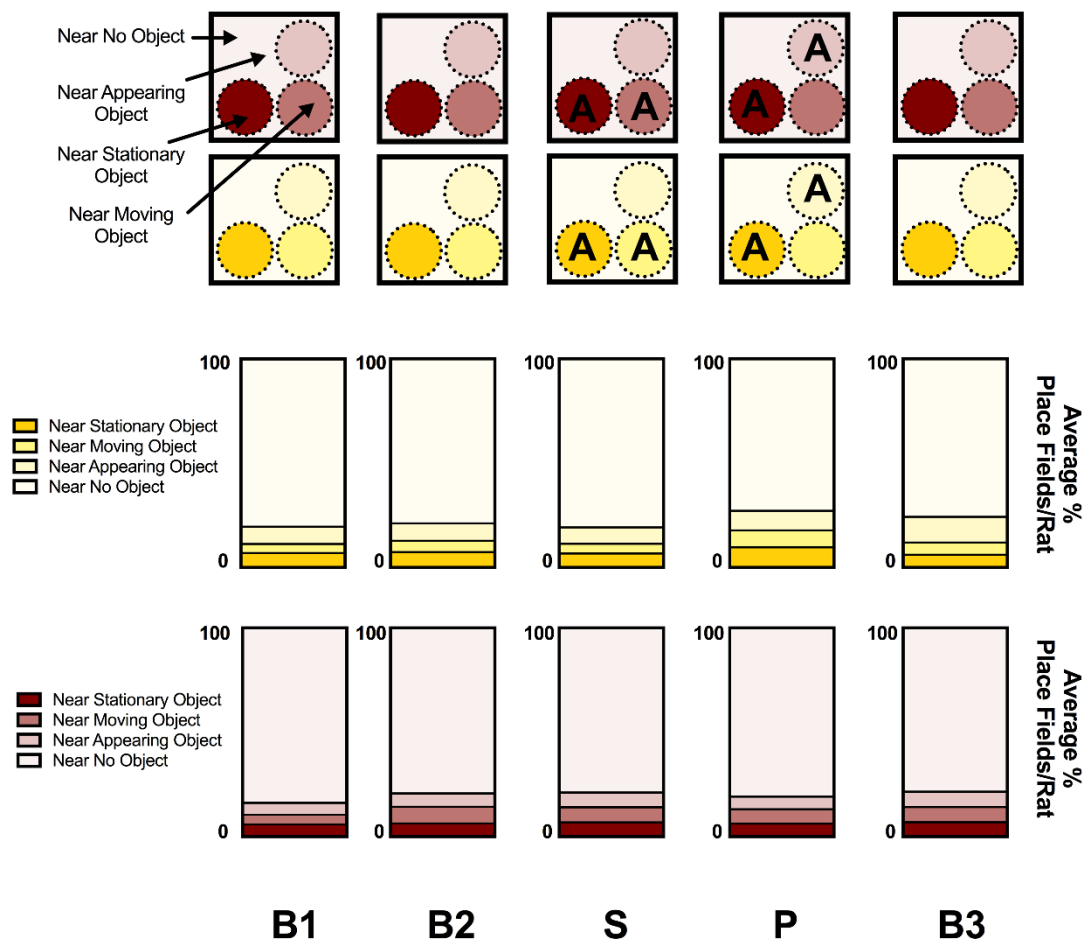
Figure 4.20: The proportions of fields in close proximity to each object in the open field over the five different sessions (outlined at the top). The white box (yellow bars) condition showed a significant interaction between session and the percentage of fields expressed in each type of object movement quadrant over the sessions, with significantly less fields expressed in the area not containing objects during the probe session compared to the sampling session. There was a positive correlation between the number of fields expressed away from objects in the probe session and novel object location memory strength.

Interestingly, when the percentage of fields within object areas was compared to individual nOL memory scores, the percentage of fields within vicinity of the objects during the probe session (i.e. the stationary object and the probe object) negatively correlated to nOL memory scores (stationary object: r= -0.4698, p=0.0015; probe object: r= -0.3458, p=0.0231), whilst the percentage of fields away from any object locations was positively correlated to nOL memory scores (r=0.5196, p=0.0004). These correlations were not apparent in in the two initial baseline sessions (B1 and B2) or the sampling session (S). During the final baseline session there was also a positive correlation between the percentage of fields expressed away from any previous object locations and the nOL memory scores (r=0.3887, p=0.0120). These results indicate that the strength of nOL memory is related to the expression of fields away from objects. Although surprising it is possible that the expression of memory requires stable fields that aren't modulated by being in the immediate vicinity of objects. This is because previous studies have shown that place fields are most likely to be modulated by the introduction or movement of objects when in close proximity to said objects (Lenck-Santini et al., 2005). Stable fields away from objects could produce a stable spatial map on which to encode object locations, whilst fields near objects could form associations with object identity and object location, becoming more disrupted when objects were moved. Therefore the next set of analyses sought to investigate the different types of place field movement that could be associated with objects and object movement.

4.3.11 Types of Place Field Movement

To further elucidate how place fields were changing in relation to objects, the net movement of fields was calculated between sessions on a cell-by-cell, rat-by-rat basis. The fields of a place cell were once again assigned to object areas and the numbers of fields in each area were compared between sessions. As has been found previously by Deshmukh and Knierim (2013), it was suggested that place fields could be object-bound and therefore moving with the object, tracing the old object location when the object moved or was removed, appearing or disappearing with the object, or even appearing where the object used to be. Therefore these different types of movement were analysed for B2 to S; S to P; and P to B3 session comparisons. It was hypothesised that differences between conditions would only be evident in the sampling to probe or probe to final baseline session comparisons, i.e. after exposure to either condition. It should be noted that these analyses were calculated for object quadrant as well as object area, and whilst higher percentages of cells exhibited different types of place field movement within the object quadrants (due to the relative size between

object areas and object quadrants), the proportions remained the same. These quadrant analyses were not shown as no correlations between movement and novel object location (nOL) memory were found, possibly due to covering less specific areas of the environment.

Baseline to Sampling Sessions

When objects were introduced in the sampling session (Figure 4.21), place cells that expressed a net increase of fields near objects were categorised as having field(s) that appeared in a novel object location. Place cells that expressed an overall loss of place fields near objects were categorised as having field(s) disrupted by the appearance of objects.
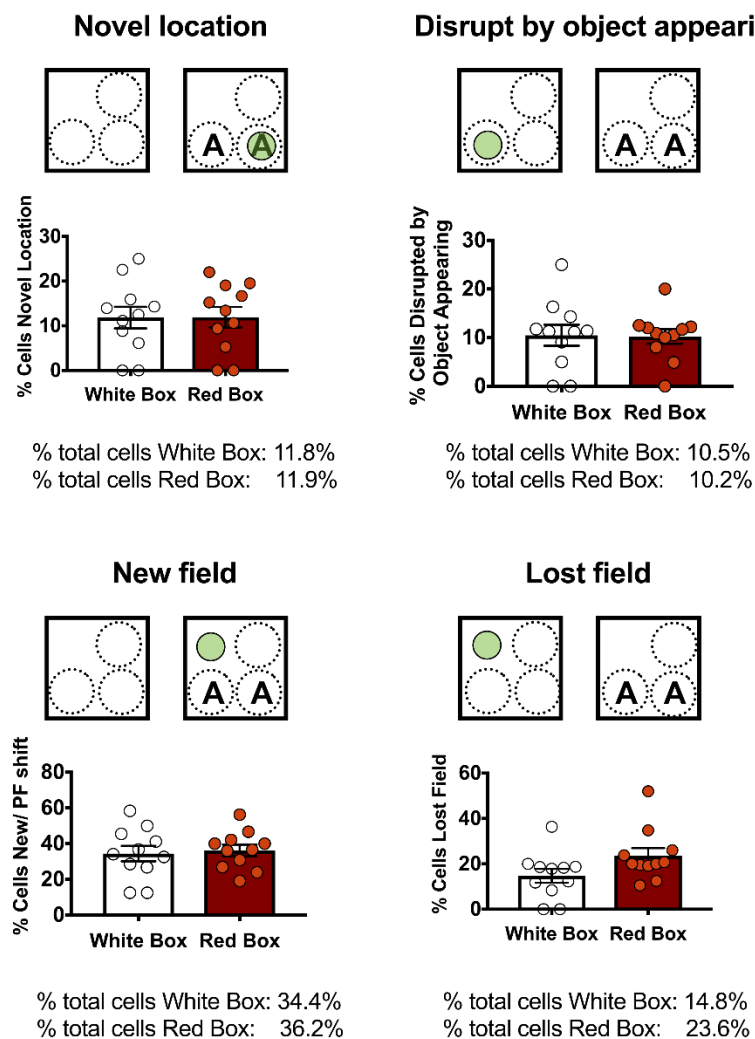
Figure 4.21: Different types of characterized place field movement between the second baseline and sampling sessions (B2 to S). Overall place fields were more likely to appear than disappear when objects were introduced. There were no differences in place field movement between the white and red box conditions. Individual dots represent the percentage of cells exhibiting a type of place field movement for an individual rat. Error bars represent the SEM. % of total cells represents the percentage calculated from all rats combined.

Overall around 12% of place cells expressed fields that appeared in the vicinity of the objects when they were introduced in the sampling session (WB: 11.8%, RB: 11.9%), indicating that a proportion of the place cell population was reacting to the presence of objects. Around 10% of cells expressed fields in the baseline session that then disappeared when objects were introduced near the field in the sampling session (WB:10.5%, RB: 10.2%). This suggests that fields could also be disrupted by the addition of objects. Over a third of place cells expressed new fields away from object locations (WB: 34.4%, RB: 36.2%), which was significantly more than had fields that disappeared away from object locations (WB: 14.8%, RB: 23.6%) [main effect of field change (gain or loss) on percentage of cells: $F_{(1,10)}=22.86$, $p=0.0007$; no effect of condition: $F_{(1,10)}=2.852$, $p=0.1222$; or interaction of field change and condition: $F_{(1,10)}=0.7086$, $p=0.4196$]. This suggests that place fields were more likely to appear than disappear when objects were introduced, a finding in line with the increase in the number of place fields expressed between baseline and sampling sessions, outlined earlier. As expected there were no correlations between place field movement (the percentage of cells expressing fields that either appeared or disappeared in regards to the introduction of objects into the environment) and nOL memory expression.

Sampling to Probe Sessions

There were numerous possibilities for place field movement in relation to objects when one object was moved between the sampling and probe sessions, summarised in Figure 4.22. It was hypothesised that certain types of place field movement between the sampling and probe sessions could correlate to nOL memory strength, such as cells expressing fields that moved with the object (object-bound cells), or cells expressing an increase fields in the location of the old object (misplaced cells), possibly to signify that an object location had changed. The analysis earlier in this chapter showed that the percentage of place fields expressed away from objects in the probe trial positively correlated with the strength of nOL memory, and the percentage of these fields did not change between sampling and probe sessions in the red box condition, but decreased between sampling and probe sessions in the white box condition. Therefore it was also hypothesised that place cells would express a stable number of fields (no net loss or gain of fields) away from object locations if the animal expressed nOL memory. It is possible that stability in the number of fields expressed away from objects could provide a more consistent spatial map for the animal.

The first important place field-object interaction analysed calculated the percentage of object-bound cells, i.e. cells expressing fields that appear to follow the object as it moves. Place cells that expressed an overall loss of fields in vicinity of the moving object area, and an overall gain in fields in the vicinity of the probe object area were categorised as having fields that followed the moving object. Both the white and red box condition exhibited such object-bound fields, however this was only seen in a very small population of cells recorded (WB: 0.4%, RB: 0.9%). A minimally larger population of cells 'traced' the old location of the moving object (WB: 1.4%, RB: 1.9%), i.e. they were near the moving object during sampling, and remained in this location during the probe trial, even though the object had moved away. Place cells expressing 'trace' fields had at least one field in the vicinity of the moving object area in both sampling and probe trials. Neither the percentage of 'object-bound' not 'trace' cells showed any differences between the two conditions, and nOL memory scores did not correlate to the percentage of cells expressing such fields, suggesting that although fields were following or tracing the object this was not in any relation to memory strength.

Object-bound

% total cells White Box: 0.4%
% total cells Red Box: 0.9%

Trace

% total cells White Box: 1.4%
% total cells Red Box: 1.9%

Stable away from objects

% total cells White Box: 29.7%
% total cells Red Box: 31.2%

New field/ field shift

% total cells White Box: 29.4%
% total cells Red Box: 28.8%
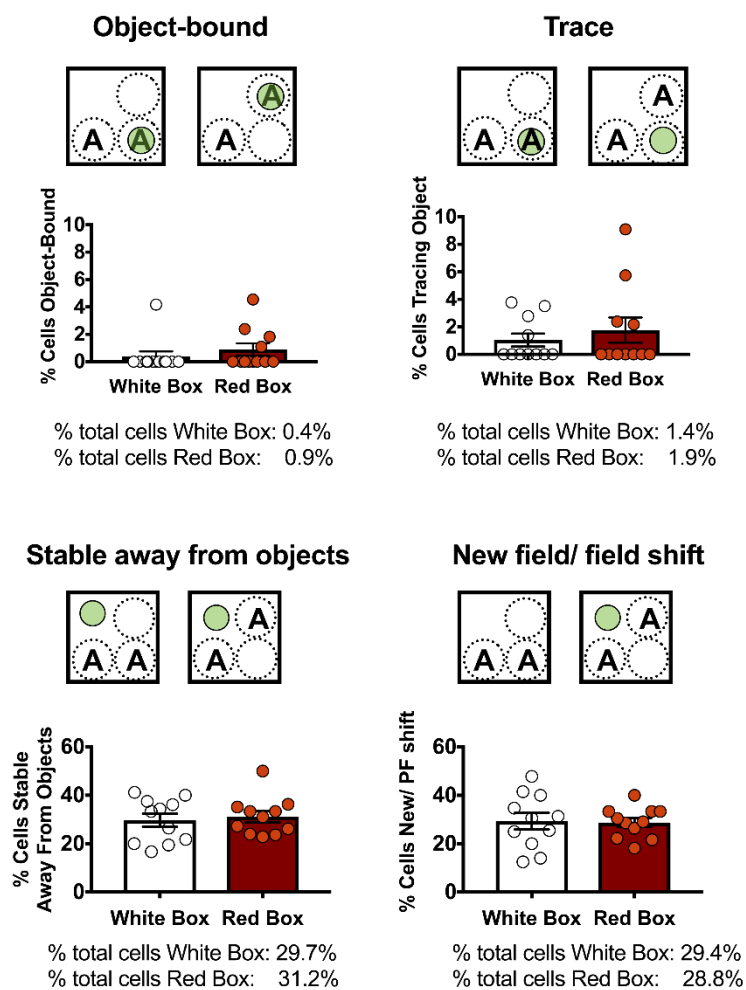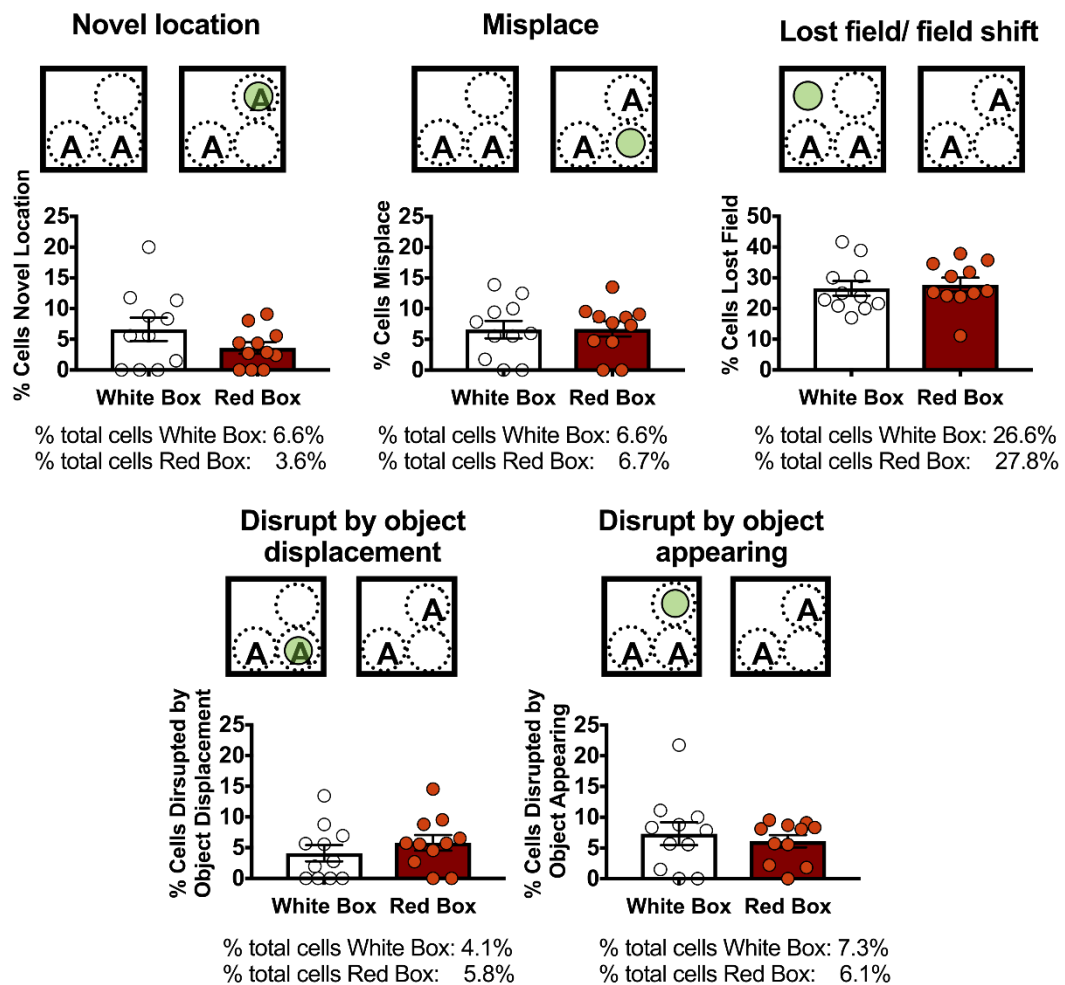
Figure 4.22: Refers to figure from previous page as well as figure above. Different types of characterized place field movement between the sampling and probe sessions. Place fields were just as likely to appear or disappear from the areas near the moving object or appearing object. As has been found previously a small proportion of cells were object-bound, and a slightly larger proportion traced the old object location. There were no differences in place field movement between the white and red box conditions, however the percentage of cells with fields appearing away form object locations significantly correlated to nOL memory strength. Individual dots represent the percentage of cells exhibiting a type of place field movement for an individual rat. Error bars represent the SEM. % of total cells represents the percentage calculated from all rats combined.

As the percentage of place fields away from objects in the probe trial positively correlated with the strength of nOL memory, and the percentage of these fields did not change between sampling and probe sessions in the red box condition, it was suggested that memory could be correlated to a stable number of fields between the sampling and probe sessions. Therefore the percentage of cells with no net loss or gain of fields expressed in the area away from objects was analysed. This area was away from both sampling object locations and probe object locations. Surprisingly, although almost a third of cells expressed a stable number of fields away from objects (WB: 30%, RB: 31%) there were no differences between

the two conditions and no significant correlation with nOL memory. This indicates that whilst the total percentage of fields away from objects can predict nOL memory, this is not in relation to the stability of the number of fields expressed.

If the stability of the number of fields in this area was not correlated with memory strength, it could be possible that memory was instead related to an increase of fields expressed during the probe trial. Therefore the appearance of place fields in relation to object location were analysed. Cells with a net increase in place fields within the area away from any objects between sampling and probe trials were categorised as having new or shifted fields. Similar to the percentage of cells with a stable number of fields, around a third of cells expressed new fields in the probe trial away from any object locations (WB: 29.4%, RB: 28.8%). Although there were no significant differences between conditions (t=0.161, df=10, p=0.8753), when the percentage of cells with new fields was correlated to individual nOL memory scores there was a significant positive correlation (r=0.3148, p=0.0398), indicating that cells were more likely to express new fields away from objects if the animal had nOL memory. This correlation was limited to the fields expressed away from objects. Place cells that had a net increase in fields expressed in the vicinity of the appearing object, but no decrease in fields in the vicinity of the moving object, were categorised as expressing novel object location fields. Place cells with a net increase in fields within the vicinity of the moving object were classified as having misplaced fields. Similar numbers of place cells were categorised as having novel object location fields (WB: 6.6%, RB: 3.6%) and misplaced fields (WB: 6.6%, 6.7%) and no types of place field appearance near old or new object locations showed differences between the white and red box conditions or correlations with nOL memory strength.

The disappearance of place fields was also analysed. Place cells exhibiting a net decrease of fields in the vicinity of the moving object, but no overall increase in the vicinity of the appearing object, were categorised as being disrupted by object displacement. Cells with a net decrease of fields in the vicinity of the appearing object were categorised as being disrupted by object appearance. Cells with a net decrease of fields in the area away from any objects were classified as having lost or shifting field(s). As expected to the relative sizes of the areas, the number of place cells with disappearing fields was largest in the area not containing objects (WB: 26.6%, RB: 27.8%). There were similar proportions of place cells with disappearing fields that had been disrupted by either the object appearing in that location

(WB: 7.3%, RB: 6.1%) or by the object being moved from that location (WB: 4.1%, RB: 5.8%). These results indicated that place fields can be disrupted by both the appearance and movement of objects, however neither showed significant differences between conditions nor were correlated to the strength of nOL memory.

Overall these results suggests that, although a small proportion of cells had fields that moved with the moving object, a larger proportion of cells had fields that appeared or disappeared in relation to object movement. In addition, place fields were just as likely to appear as disappear when objects were moved. Surprisingly, the only correlate of nOL memory strength was the appearance of new place fields away from object locations during the probe trial. It is possible that any other differences or correlations are masked by the length of the probe trial recording, a suggestion that is explored in greater depth within the discussion.

Probe to Baseline Sessions

The final set of analyses looked at place field-object interaction when objects were removed after the probe session. Place cells were classified as expressing fields that traced the old object locations if they had at least one field in the vicinity of the appearing object in both probe and B3 sessions, or at least one field in the vicinity of the stationary object in both probe and B3 sessions. As seen before, when sampling and probe trials were compared, a small proportion of cells had fields that traced the old object location(s) in both conditions (WB: 6.5%, RB: 4.9%). This suggests that a small subset of place cells have fields that aren't disrupted by the removal of objects, even when in close proximity to said objects during the probe trial. A smaller proportion of cells traced the location of the moving object during sampling in both the probe and B3 sessions (WB: 3.4%, RB: 4.2%). Place cells were classified as tracing the sampling location if at least one field was expressed in the moving object quadrant during both probe and B3 sessions. This suggests a small subset of cells were unaffected by the movement and the subsequent removal of objects. There was no difference between conditions for either type of place field stability (trace: t=0.568, df=10, p=0.5825; sampling position trace: t=0.09803, df=10, p=0.9238), however when values for every rat for every session were correlated with the corresponding nOL discrimination indices, there was a significant negative correlation for the percentage of cells that traced an object location between probe and habituation and the strength of nOL memory (r= -0.4795, p=0.0015). This suggested that if animals expressed memory in the probe trial, place fields

would not remain in the location of the objects after their removal. Although surprising, this could indicate if fields are stable and unaffected by object movement when in the close vicinity of objects, less detection of novelty is available to the animal. There were no significant correlations between the proportion of cells tracing the sampling location and nOL memory.

In contrast, a number of cells had disappearing fields between P and B3 sessions. Between 10 and 20% of cells (WB: 18.1%, RB: 10.6%) had fields that were near an object during the probe session, which then disappeared when the objects were removed. This suggested that some cells were modulated by the removal of objects when in close proximity. Interestingly, there were trends towards a significant increase in the proportion of disappearing cells near objects in the white box condition (t=1.953, df=10, p=0.0794), and when the proportion of cells in every session for every rat were correlated with the corresponding nOL discrimination indices there was again a significant negative correlation (r= -0.4691, p=0.0020). This indicates that if animals expressed memory for the novel object location in the probe trial, fields expressed near these new object locations were less likely to disappear upon the removal of said objects. However, as fields were also less likely to trace the object location this suggests that it was in fact the proportion of cells expressing fields near objects throughout the entirety of the probe trial that correlated negatively to the strength of nOL memory, rather than how the place fields reacted to the subsequent removal of objects. A larger proportion of cells had fields that disappeared from the areas not containing objects during the probe trial (WB: 37.2%, RG: 38.0 %). Unlike the proportion of disappearing cells near objects, there were no differences between the two conditions or significant correlations between the proportion of cells disappearing away from old object locations and nOL memory strength.

Around 15% of cells expressed at least one new place field during B3 in locations that had previously contained an object during the probe trial. This was true of both the white box and red box conditions (WB: 16.6%, RB: 12.0%), with no differences between them (t=0.6566, df=10, p=0.5262) or significant correlations with nOL memory strength. The percentage of cells with appearing fields in the old object locations was therefore approximately the same as the percentage of cells with disappearing fields from these areas.
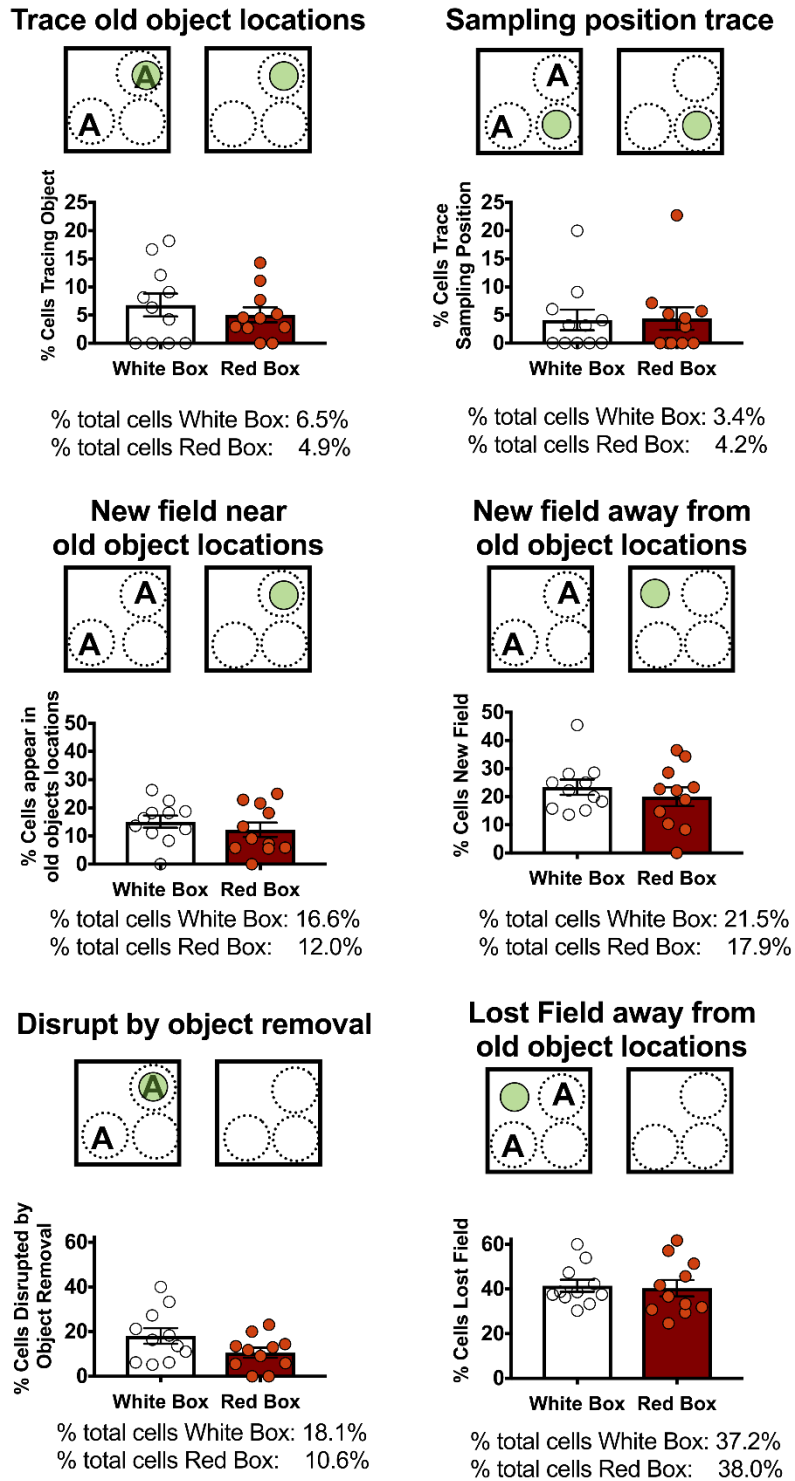
Figure 4.23: Different types of characterized place field movement between the probe and final baseline sessions. A proportion of cells expressed fields that traced the object positions, as predicted. The percentage of both these cells and the percentage of cells disrupted by object removal negatively correlated to nOL memory strength. This was possibly due to these cells expressing fields near objects during the probe session. More place fields disappeared than appeared when objects were removed. Individual dots represent the percentage of cells exhibiting a type of place field movement for an individual rat. Error bars represent the SEM. % of total cells represents the percentage calculated from all rats combined.

Interestingly, a proportion of cells actually expressed a new field away from the previous probe object locations after their removal (WB: 21.5%, RB: 17.9%), although this was much lower than the percentage of cells with disappearing fields away from object locations. This suggests that place cells express less fields when objects are removed, something that was mirrored in the number of place fields analysis. Although there were no differences between conditions in the proportion of cells expressing new fields away from objects (t=0.6514, df=10, p=0.5295), there was an almost significant negative correlation between the proportion of cells and nOL memory strength (r= -0.2881, p=0.0687). Although this was not significant, it does indicate that the positive correlation between nOL memory strength and the increase in the proportion of cells expressing new fields away from objects was specific to the sampling and probe trial comparison.

Overall these results suggest that place field expression can be modulated by object movement in many different ways, in line with previous findings. Interestingly, the only positive correlation between the proportion of place cells and nOL memory strength was for place fields appearing away from objects between sampling and probe trials. This indicates that these analyses were unable to identify changes near objects that could be underlying the detection of novel locations or strength of spatial memory. Examples of place field movement between all sessions are shown in Figure 4.24.

Figure 4.24: (Next page) Examples of place field movement between the five different sessions (H1, H2, S, P and H3). White boxes represent the open field, with 'A' indicating where objects were during the sampling and probe sessions. Each rate map shows the area of the environment explored by the animal within a single exposure. The firing rate is given below each rate map and can be different for every session (i.e. some sessions have very low firing rates). The red part of the rate map represents the peak in field firing rate. The blue part of the rate map represents no or minimal amounts of place cell firing in those locations in the environment. White pixels indicate the rats did not explore this part of the environment. Top: this cell expressed a spatially stable field throughout the five sessions. Second to top: This cell expressed fields that appeared with both the objects in sampling, and then had one field that was object-bound and followed the moving object. This field then disappeared when the objects were removed. Middle: This cell expressed spatially stable fields until one object was moved in the probe trial, where it expressed fields away from objects. Place field firing returned to the location seen in H1 H2 and S when the objects were removed. Second to bottom: This cell expressed a field that was spatially stable until objects were added near the field, where it was disrupted by object appearance. This field only returned to the previous location when objects were removed. Bottom: This field appeared with one of the objects in sampling, was object-bound and followed this object into the probe session, and then traced the old object location when the object was removed in H3.

B1 B2 S P B3
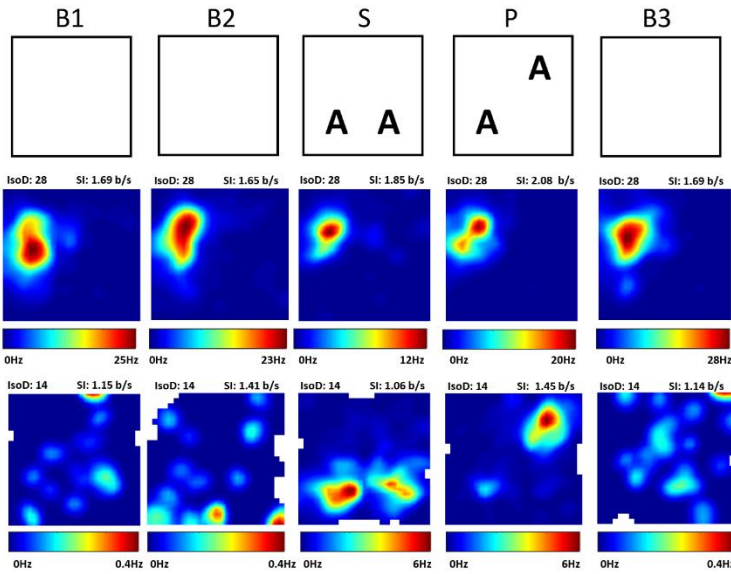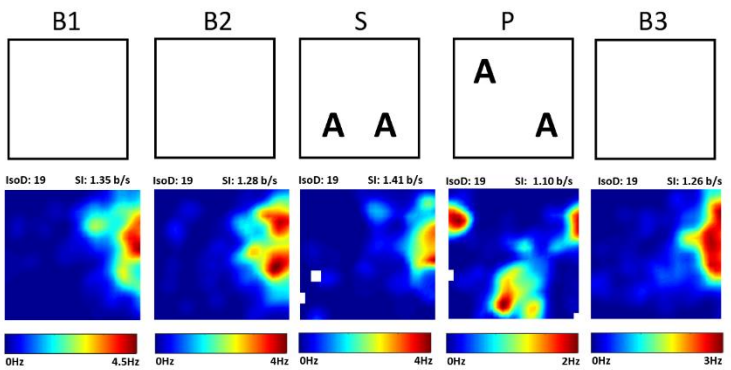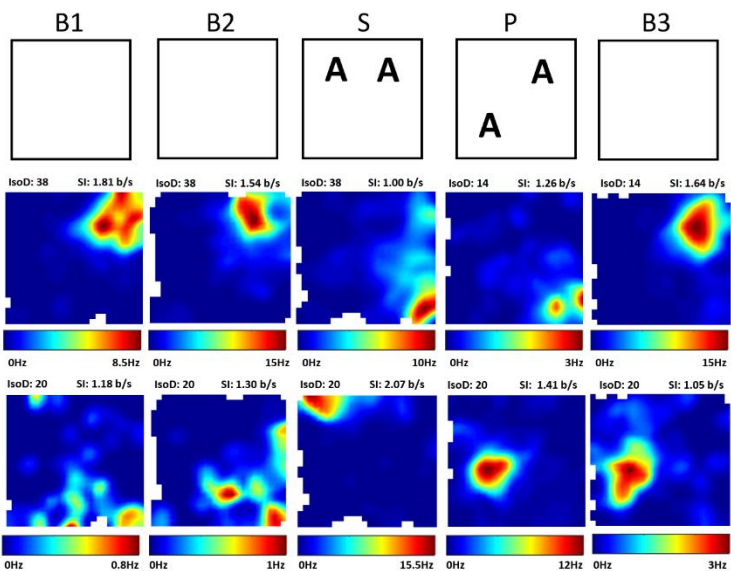
Spatially stable

H2 to S: Appear with object

S to P: Object-bound

P to B3: Disappear when object
        is removed

B1 B2 S P B3

S to P: Appear away from
       objects

B1 B2 S P B3

H2 to S: Disrupt by object
       appearing

H2 to S: Appear with object

S to P: Object-bound

P to B3: Trace object location

198

4.3.12 Results Overview

Reducing visual stimulation after spatial learning enhanced 6h nOL memory in implanted rats, replicating our previous findings at 24h. This allowed us to compare all types of analysis between the two conditions and directly correlate findings from each rat with their strength of nOL memory expression.

Place Cell Properties

The average place cell firing rate increased significantly over sessions, suggesting that firing rate increased with novelty, in line with findings from previous studies (Larkin et al., 2014). However this increase was minimal, suggesting that further experiments or analyses would need to be carried out to draw conclusions. Spatial information carried by place cells in the probe session was significantly lower in the red box condition, indicating that this measure was correlated to the strength of memory. However when spatial information was compared to individual nOL scores there was no significant correlation. Spatial selectivity did not change over sessions. Spatial sparsity and coherence both worsened when objects were introduced. Surprisingly, spatial coherence improved when objects were moved, directly contrasting with the worsening of coherence seen by Lenck-Santini (2005). The addition of objects to a familiar environment led to instability of the spatial map and partial remapping of the place cell population. Median correlation, the percentage of stable cells and rate remapping all showed a decrease upon the introduction of objects to the environment, which decreased further when objects were moved, and increased when objects were removed, indicating these object manipulations can lead to both spatial and rate remapping of place fields. Unexpectedly, these changes did not correlate with the expression of spatial memory. This suggests that place cell stability is only affected by the reduction of interference when the environment itself is novel.

Place cell population correlation analysis

The next analyses investigated whether cell population firing was influenced by the presence of objects, and whether these changes were underlying the differences in nOL memory strength. Time-bins representing gamma (25ms) oscillations showed significantly more cells exhibiting coordinated firing with other cells in a session when objects were introduced, suggesting that cell assemblies were encoding object identity-object location associations. This then decreased when one object was moved to a new location. As expected, these

significant changes led to a decrease in the correlation (PCo) between the sessions when objects were introduced, and an increase when they were removed. Surprisingly although the time-bin representing the animal's movement through a place field (1s) showed no differences in PCorr between sessions, the same changes in PCo values shown in the other time-bins were seen. This could suggest that the populations of cells that are active during the sampling and probe sessions tend to be inactive during the non-object sessions. Once again, there were no differences between the WB and RB conditions.

Place Field Properties

Peak in-field firing rate did not change, suggesting that in relation to overall firing rate, the peak firing rate actually decreased over sessions, in direct contrast to the study by Komorwski et al. (2009). As found previously by Burke (2011), the number and size of place fields were modulated by object presence. Results suggested that whilst more place fields appeared with the introduction of objects, overall these fields were smaller. The number of fields then decreased when objects were removed. Although there were no significant differences between the WB and RB conditions, place field size during the probe trial negatively correlated with nOL memory strength. This was not apparent for any other sessions suggesting that smaller place fields are associated with stronger nOL memory.

The percentage of place fields expressed away from objects positively correlated to the strength of nOL memory in both the probe and final baseline session, and the percentage of place fields expressed near objects during the probe trial negatively correlated to the strength of nOL memory. These results suggested that stronger memory was associated with place fields being located away from objects during the probe trial. Further analysis investigating the movement of place fields corroborated these findings, showing that the strength of nOL memory was positively correlated to place fields appearing away from objects during the probe trial. Surprisingly, stability of the number of fields away from objects did not relate to memory strength. Place cells did express fields that were object-bound, following the object as it moved, and fields that traced the old object locations. Although previous studies have suggested these types of movement are correlates of spatial memory, the percentage of cells expressing fields that were object-bound or that traced the object locations were not significantly different between red and white box conditions and did not correlate to the strength of nOL memory.

Overall these results suggest that whilst object movement can clearly modulate many different place cell properties, the enhanced expression of nOL memory seen in the RB condition is not expressed as a change in place cell stability or at a place cell population level. A number of place field properties do correlate to the strength of nOL memory, although it appears that these underlying changes occur away from object locations.

4.4 Discussion

This experiment set out to test the hypothesis that decreased interference following the learning trial of a spatial novel object location (nOL) memory task would lead to predicted differences in memory-related properties of place cells between red box (decreased interference) and white box (control) conditions, such as place field stability or the association of place fields with specific object locations, when the animal underwent a probe trial 6h later.

Novel Object Location Behaviour

To ensure that any place field properties analysed could be correlated to behavioural results, the first part of the experiment involved testing the animal's nOL memory strength during the first three minutes of the probe trial. As seen previously with unimplanted animals, those that underwent the red box condition directly after the sampling session expressed significant nOL memory at 6h, whereas those that underwent the white box condition expressed no memory. This difference was evident throughout the whole of the 3 minutes of probe trial analysed, and significantly different between the two conditions. This confirms that implanted animals can also benefit from reduced visual interference, as predicted. As the difference between the red and white box conditions had only been shown previously (in Chapter 2) using a 24h delay between sampling and probe sessions and a 3h exposure to the red or white box condition, this experiment also confirmed that this effect could be replicated using different delays (6h) and a shorter period of reduced interference after learning (1h). Importantly, these behavioural results allowed the comparison of place field properties in animals tested under two conditions; one that expressed memory and one that did not, which has not been done before. As not every animal expressed memory in the red box condition, due to the natural spread of the data, this also allowed place field properties to be directly correlated to individual novelty preference scores.

<u>Firing Properties of Place Cells</u>

As changes in firing rate have been implicated previously in multiple types of spatial memory task, the next analysis sought to investigate how the average place cell firing rates and peak in-field firing rates changed over sessions and whether these changes correlated to memory strength. Previous studies have shown that the average place cell firing rate indiscriminately increases when objects were moved during a nOL memory task (Larkin et al, 2014). This is line with our results showing that average firing rate increases over the sessions, as the majority of sessions featured object manipulations of some kind. Although these increases were minimal, and only became apparent when the baseline sessions were compared with the probe session, typically average place cell firing rates decrease as the environment becomes more familiar to the animal (Brandon et al., 2014; Karlsson and Frank, 2008). This suggests that increases in firing rate were due to the movement of objects.

However, unlike the results reported by Larkin et al., where firing rate correlated to object location novelty preference scores, the increase in firing rate seen did not correlate to nOL memory expression, and there was no significant difference between the two conditions. It should be noted that Larkin did not include a condition that had no memory, instead correlating changes in place cell firing with the overall expression of memory shown by all animals used. This appears to highlight the need for studies to include control conditions where memory expression is not expected if properties are hypothesised to correlate to the strength of memory expression. Yet it is unlikely that the increases in firing rate seen by Larkin did not represent the detection of some sort of novelty in the environment, as firing rates decreased as behaviour associated with the novelty of the new object location (i.e. the time spent exploring the new object location vs the old object location) decreased during the probe session.

Although in the present study animals expressed preference for the novel location for the first three minutes of the probe trial, it is unlikely that this preference would remain throughout the entire 10-minute recording session. It could be hypothesised that increases in firing rate would be most evident at the start of the sampling, probe and final baseline sessions, as this is where the addition, movement or removal of objects would be most novel, respectively. Further analysis could aim to calculate average firing rates for the first half of each 10-minute session and compare this to the second half. The discrimination index values

could also be calculated for the entire 10-minute probe trial and correlated to the firing rate minute-by-minute. It would be predicted that as preference for the novel location decreases, so would average firing rate, as had been reported by Larkin. It would also be interesting to analyse these changes in firing rates in respect to location, for example analysing how rates change in the vicinity of objects over the sessions. However as Larkin found that this increase in firing rate was not specific to object locations, it could be hypothesised that no differences would be seen between quadrants or object areas.

Unlike average place cell firing rate, the peak in-field firing rate did not differ significantly between sessions. This is in contrast to the findings of a previous study by Komorowski et al. (2009) which reported increases in peak firing rates at important locations associated with food rewards in specific contexts. These results also appear to be inconsistent with the original hypothesis that peak in-field firing would increase due to object locations specifically becoming overrepresented. However, it is possible that whilst the average peak in-field firing rate did not change, certain place fields would increase in firing rate and certain place fields would decrease, depending on their proximity to the object locations. As with average place cell firing rates, there were no significant differences between the red and white box conditions, however it is possible that if place fields did overrepresent object locations that these could then show differences correlating to the strength of memory expression. Zheng et al. reported this phenomenum of increased place field firing near novel object locations during a nOL behavioural task. The average peak in-field firing rate increased only in fields expressed close to the novel object location, and only during fast gamma oscillations. Firing rate did not increase in fields close to the familiar object location, and did not increase during slow gamma oscillations. These results suggested that these increases in firing rate acted as a novelty detector, aiding to encode novel object locations into the pre-exisiting spatial map of the familiar environment. It would therefore be interesting to investigate whether the red box condition would increase firing rates of place fields near object locations, and whether these increases would also be specific to gamma oscillations.

Spatial Properties of Place Cells

The next possible correlates of nOL memory strength analysed were various spatial properties of place cell activity. Previous studies investigating the modulation of spatial tuning of place cell firing have shown that these properties do not change with the addition

of objects (Kyd and Bilkey, 2005). This suggests that whilst spatial tuning improves as the environment becomes more familiar (Brandon et al., 2014), the increased spatial information carried by place cell firing is not further affected by the addition of objects to the environment.

Although spatial properties of place cells were predicted not to change over sessions, the place cells fired in a more indiscriminate fashion throughout the environment when objects were introduced. As the place cells expressed significantly more fields when objects were introduced, this decrease in spatial sparsity is logical, as the place cell would be firing in more locations throughout the environment. Spatial selectivity, the amount of in-field firing compared to out of field firing, did not significantly change over the sessions and therefore did not appear to be modulated by the presence of objects, as expected.

Although spatial information also appeared not to be consistently modulated by objects, there was a significant interaction between session and condition. This difference in spatial information between the two conditions during the probe session suggests that spatial information can be modulated by the strength of object location memory. This is surprising as previous studies have shown no changes in spatial information when objects were introduced (Kyd and Bilkey, 2005). One possibility is that animals were perceiving the objects as landmarks within the environment rather than non-stationary objects (Scaplen et al., 2014). Animals who had memory for these 'landmarks' in the sampling session would notice the change in their location during the probe trial. As high levels of spatial information have been associated with familiarity of an environment (Brandon et al., 2014), this could indicate that cells recorded from animals in the red box condition detect this change in the environment more readily, making the overall environment appear more novel to the animal. However it is interesting that the addition and removal of objects did not also produce such decreases in spatial information. It should also be noted that spatial information of place cells recorded during the probe session was not directly correlated to the individual nOL memory scores. It is possible that whilst an overall interaction of session and condition can be seen, this effect is not robust enough to be seen when individual nOL scores are used.

The modulation of spatial properties in close proximity to objects has been shown previously by Lenck-Santini et al., (2005), who found that changes in spatial coherence were limited to

the population of cells with fields expressed over 10 cm away from objects. It would therefore be interesting to analyse the spatial properties of place cells that express fields in close proximity to objects compared to those expressing fields further away. It could be predicted that coherence would decrease as object proximity decreased, and that spatial information might change, although possibly in unpredictable ways. Overall the spatial properties of place cells did not appear to be consistently different between red and white box conditions, suggesting that these properties do not underlie the expression of nOL memory.

Place Cell Stability

The introduction of objects into a familiar environment has been shown to affect the stability of place cell firing in three different ways. Place cells can undergo global remapping, where the entire population of place cells change their spatial firing patterns, either by starting or stopping firing, or changing the location of place field firing (Burke et al., 2011). Cells can also undergo rate remapping, where place field locations remain stable but the average place cell firing rate changes within the session; or partial remapping where a proportion of the place cell population destabilises whilst the other proportion remains stable (Lenck-Santini et al., 2005). One study even reported no changes in any measures of place cell stability when objects were introduced to the environment (Kyd and Bilkey, 2005). Therefore the next set of analyses sought to investigate which of these changes occurred within the novel object location paradigm used, and whether these measures of place cell stability correlated to the strength of nOL memory expression.

As predicted, both the median correlation of rate maps and the percentage of stable cells was highest between the two baseline sessions (B1 and B2), indicating that the environment was indeed very familiar to the animals. When objects were introduced to this familiar environment, the median correlation and percentage of stable cells between session B2 and S decreased significantly, and there was further decrease when one object was moved to a novel location (S to P). This suggests that the addition and movement of objects to a familiar environment can destabilise the spatial firing of place cells that previously had very stable fields. It had been hypothesised that the introduction of only two objects, and then subsequent movement of only one of these objects, would not be enough to lead to global remapping, as global remapping only appeared to occur in studies that introduced and

moved many objects simultaneously (Burke et al., 2011). As hypothesised, this destabilisation of the spatial map did not occur throughout the entire place cell population, as around 60% of cells remained spatially stable. Therefore the introduction of objects led to partial remapping, as found previously by Lenck-Santini et al. (2005). This level of stability within a place cell population has been reported before by Scalpen et al., where the 90^O rotation of an 'object floor cue', a pair of small 2D shapes attached to the floor of an open field, led to 57.9% of the place cells spatially remapping. Interestingly the proportion of cells that remapped was less (around 40%) when animals had only been exposed to the environment for less than nine days (Scaplen et al., 2014). This is in line with the amount of habituation to the environment given to animals in the present experiment.

Interestingly, a proportion of cells that were spatially "stable" underwent rate remapping when objects were introduced, moved and removed, suggesting that objects could also disrupt or destabilise the firing ratios of place cells within a familiar environment. However, 60% of these "stable" cells did not undergo rate remapping, indicating that a proportion of the cell population remained stable in both location and firing rate.

Surprisingly there were no significant differences between the red and white box conditions for any measures of stability analysed. This is in contrast to our previous results in Chapter 3, showing that the reduction of interference enhances stability of place cell firing within a novel environment. It was hypothesised that place cell stability would either be higher in the red box condition, due to an increased stability of the underlying spatial map aiding the encoding of novel object locations, or lower in the red box condition, showing that place cells were detecting the change in object location to a greater extent. Although it appeared that neither of these possibilities were correct, it is plausible that place cells in the red box condition could show both of these properties. Lenck-Santini et al., (2005) reported that cells expressing fields close to objects were more likely to undergo spatial remapping, whereas fields away from objects were more likely to show changes in firing rate. Therefore it could be hypothesised that place cells in the red box condition would show enhanced levels of stability away from objects and decreased levels of stability in the vicinity objects. This possibility requires further analysis. An important future experiment would be to repeat this protocol with a sampling session, followed by the red or white box condition, then 6h later by another session with the same objects and object locations. This would allow the analysis

of place field stability across sessions containing objects, without the confounding factor of object movement. It could be predicted that the red box condition would show enhanced stability between these two object containing sessions, compared to the white box condition.

These results also appear to highlight differences in processes between the consolidation of a novel environment and the consolidation of novel object location within a familiar environment. The information derived from the stability of locations might not be enough to encode these complex interactions. It is possible that the object identity-object location associations are actually encoded within the place cells that express "unstable" place fields between baseline and sampling sessions. Although there were obviously no differences between the two conditions in the proportion of "unstable" place cells, it is possible that the properties within this population might have been different between the two conditions.

Synchronous firing of place cell assemblies

The next set of analyses sought to investigate whether the synchronous firing of cell assemblies was influenced by the presence of objects, and whether these changes were underlying the differences in nOL memory strength. It has been shown previously that upon the movement of a familiar object to a novel location, the synchronous firing of cell assembly pairs decreased, possibly to signify a mismatch in the learned object identity-object location associations (Manns and Eichenbaum, 2009). On the other hand Larkin et al. found that the synchronous coordinated firing of cell pairs during sharp wave ripples recorded during the nOL task increased when one object was moved to a novel location, suggesting that coordinated firing was acting as a novelty signal.

Our results mirror those found by Manns and Eichenbaum, showing that the synchronicity of cell assembly firing decreased between the sampling and probe trials. This therefore supports the idea that the synchronicity of cell assembly pairs represents object identity-object location associations within the nOL memory task. Interestingly the synchronicity actually increased between the baseline (B2) and sampling sessions, upon the introduction of novel objects. A possibility for this increase could be that cell assembly pairs are encoding these object identity-object location associations during the sampling session, whereas this association encoding is not needed during exploration of the familiar open field in the baseline session (in line with results in Chapter 3). It should also be noted that these changes

in synchronicity were only seen when using the time frames associated with gamma (25ms) and theta (125ms) intrinsic oscillations. As these changes appeared to be more apparent when using the 25ms time frame, it could be suggested that the object location associations could have been encoded during awake sharp wave ripples. The possible encoding of object location associations during gamma oscillations has been proposed previously by both Larkin et al. (2014) and Zheng et al. (2016). Larkin reported an increase in coordinated synchronous cell assembly firing during awake sharp wave ripples in the probe session that persisted into the subsequent rest session, raising the possibility that that these cell assemblies were selectively replayed and consolidated (Larkin et al., 2014). Zheng reported an increase in place field firing associated with fast gamma oscillations when animals were close to novel object locations. To test the possibility that the increase in cell assembly coordination seen during sampling is relating to the encoding of object-object place associations I would need to correlate cell assembly firing with the oscillatory activity.

The coordination of cell assembly pairs between sessions was also analysed. As would be expected, the changes seen in PCorr values for the 25ms and 125ms time frames led to a decrease in the correlation (PCo) between the sessions when objects were introduced, and an increase when objects were removed. Surprisingly although the time-bin representing the animal's movement through a place field (1s) showed no differences in PCorr between sessions, this 1s time-bin showed the same patterns of decreasing PCo values seen in the 25ms and 125ms time-bins. This could suggest that the populations of cells that are active during the sampling and probe sessions tend to be inactive during the non-object sessions. However, the findings could also result from high levels of variance of cell assembly synchronicity in the 1s time-frame. For example, the average PCorr value for the 1s time-bin could remain the same, whilst more populations of cells could be consistently firing or not firing together, leading to changes in 1ms time-frame PCo values.

It was predicted that both conditions would show increased synchronous firing of cell assemblies within the sampling session (PCorr), as this coordinated firing has been associated with the encoding of object identity-object location associations (Manns and Eichenbaum, 2009). However only the red box condition would fully consolidate these cell assembly associations. Therefore during the probe trial this synchronous firing would decrease more in the red box condition, to signify a greater mismatch in the object identity-object location

associations learned in the sampling session. It was then predicted that the correlation of this synchronous firing between the sampling and probe sessions (PCo) would be lower in the red box condition, as this synchronous firing of cell assemblies would be very different between sessions (high synchronicity within the sampling session and low synchronicity within the probe session).

Surprisingly neither PCorr nor PCo values showed any differences between the red and white box conditions. Although initially our results might not seem logical, it could be suggested that 'mismatch' cell assemblies would behave in similar ways to the 'mismatch' place cells recorded by Fyhn et al. (2002). Fyhn showed that in an annular water maze task, a subset of CA1 place fields fired either when the escape platform was encountered in a novel location, or when the rat was swimming where it thought the escape platform should be during a probe trial. However these 'mismatch' cells only fired the first few times this discrepancy was encountered, becoming silent for the rest of the session (Fyhn et al., 2002). If mismatch cell assemblies behaved in the same way, they would be desynchronised for the first few encounters of either the familiar object in the novel location (appearing object area) or the familiar location without the object (moved object area), but very quickly the firing of either this cell assembly pair or other cell assembly pairs would become synchronised to encode the new object location configuration. This would mean that any differences in cell assembly synchronicity within the probe trial would be impossible to see when the whole 10-minute probe trial was analysed. It would be interesting to analyse the PCorr and PCo values for the first minute of the probe trial to establish whether this suggestion could be true, with lower levels of synchronicity predicted in the red box condition, correlating with the strength of nOL memory. It would also be imperative to analyse cell assembly firing and synchronicity based on their firing location within the environment. It could be predicted that cell assemblies in the vicinity of objects would behave very differently to those far away from objects, possibly with greater levels of 'mismatch' desynchronised firing either close to locations where objects used to be or close to new object locations.

Place Field Properties

The possible place field correlates of spatial memory analysed were the number and size of fields expressed. Previous studies have shown that whilst the number of fields per cell increases when objects are introduced to a familiar environment by about 10%, the size decreases (Burke et al., 2011; Kyd and Bilkey, 2005), a pattern that was also observed in this data set. This increase in the number of fields when objects were introduced to the environment, and the subsequent decrease in the number of fields on their removal, could indicate that multiple fields are required to encode the presence of multiple objects within an environment. Although more fields might be required to encode different object locations, the number of fields expressed per cell was not correlated with nOL memory strength.

As predicted, the size of place fields decreased when objects were added and increased when they were removed. This suggests that whilst more fields are required to encode multiple object locations, the size of these fields are smaller. This indicates that the incorporation of object locations into memory of a familiar environment is reflected in the properties of place fields. Interestingly, the size of place fields in the probe session was negatively correlated to the strength of nOL memory. This was not seen for any other session. Although the average size of place fields has not previously been correlated to the memory of novel object locations within a familiar environment, it is possible that smaller fields could be associated with better incorporation of object locations into spatial memory. Therefore although a high proportion of cells express multiple fields when objects are introduced, only cells expressing multiple small fields show significant nOL memory. Surprisingly it appeared that the firing of cells recorded in the red box condition carried significantly less spatial information in the probe trial, indicating that these small fields were less spatially tuned. It is possible that place field size could be linked to the association of objects and object location, whereas spatial information could be linked to changes in the overall environment. It would be interesting to investigate spatial information and place field size in an experiment that contained two sampling sessions, i.e. two exposures to the familiar environment containing two identical objects in the same locations during both exposures, meaning the overall environment did not change. It is possible that place cells would express smaller fields whilst carrying higher levels of spatial information in the red box condition if the objects remained in the same locations.

Place field expression and movement in relation to object movement

To investigate whether there were any differences in the expression of place fields in the vicinity of object locations, the percentage of fields near every object as well as the percentage of cells away from any objects were analysed. As Deshmukh and Knierim (2013) had previously reported that fields were not preferentially expressed near objects, it was expected that there would be no differences in the percentage of place fields expressed near objects over the different sessions. Interestingly, this was true for the red box condition but not the white box condition. This could suggest that Deshmukh and Knierim's finding was only true if animals have memory of previous object locations. Although object location memory was not tested by Deshmukh and Knierim, it is plausible that the animals in this experiment did have memory for the object locations.

The percentage of place fields expressed away from objects positively correlated to the strength of nOL memory in both the probe and final baseline session. The percentage of place fields expressed near objects during the probe trial negatively correlated to the strength of nOL memory. These results suggested that stronger memory was associated with place fields being located away from objects during the probe trial. Therefore although the percentage of place fields located away from objects during the probe trial positively correlated with nOL memory, as the red box condition showed no differences between sampling and probe sessions it appeared to signify stability in the percentages of place fields, not change. This analysis calculated the overall percentage of place fields either near or far away from objects on a rat-by-rat basis. Therefore it is possible that although the proportion of place fields did not change between any sessions (including the initial baseline sessions) when the rat showed nOL memory, individual place cells could have been modulating their place field expression differently depending on whether objects were present or moving. This is also suggested in the results from the cell assembly analysis, where although PCorr values did not change significantly over sessions in the 1s time-bin, PCo values decreased upon the addition and movement of objects. Therefore the next analysis was important to detect changes of individual place cells.

The next type of analyses aimed to determine whether the types of place field movement reported by Deshmukh and Knierim were observed in the context of an nOL task, and to explore whether they correlated to the strength of nOL memory expression. Many different

types of place field movement in relation to object movement were found for all session comparisons, mirroring the movements seen in previous studies. It emerged that a subset of place cells expressed fields that appeared in the location of the novel objects during the sampling session. When one object was moved between sampling and probe sessions a very small percentage of fields followed the object (0.5-1%). This place field firing in response to an object irrespective of its spatial location has been shown previously in similar proportions of recorded place fields (3.3%: Deshmukh and Knierim, 2013). Other types of place field movement, such as appearing where the object had moved to, disappearing when the object moved away, or being displaced when the object moved near, all occurred in similar proportions (5-6%), indicating fields could be modulated by object movement within the environment.

As previous analysis had shown that stronger memory was associated with place fields being located away from objects during the probe trial, it was hypothesised that this higher percentage of fields could relate to a higher proportion of cells expressing a stable number of fields away from objects. These cells could lead to a more consistent spatial map for the encoding of novel object locations. Surprisingly, further analysis into the movement of place fields showed that although around 30% of cells expressed a stable number of fields away from objects, this was not correlated to the strength of nOL memory. It should be noted that this analysis calculated cells expressing a stable number of fields, not cells expressing spatially stable fields. Therefore it could still be predicted that cells expressing fields away from object locations would be more spatially stable in the red box condition, as outlined previously.

Misplace cells (cells with fields appearing where the object had been) were also found in this experiment. These cells were reported in the earliest place cell papers (O'Keefe, 1976; Ranck, 1973), although in much higher proportions, with O'Keefe reporting 23% and Ranck 29% of place cells compared to the 6% seen in this study. A very small number of place fields that fired where the object was in the sampling session also continued firing in the same location even after the object was removed. These fields that "trace" object locations have also been shown previously, however again in much higher proportions (Deshmukh and Knierim reported around 12% of cells, compared to around 1.5% found in this current experiment). When both objects were removed around 4% of cells fired close to where the object had been during sampling in both the probe and last baseline sessions, and around 6% traced one

of the probe object locations, again indicating that place fields could trace the object location. It should be noted that much larger percentages of cells expressed fields that traced the object locations when quadrant analysis was used, which required the centre of place fields to be within 15 cm of the object, a distance in line with the study by Deshmukh and Knierim. However the percentages of cells remained in the same proportions regardless of whether quadrant or object area analysis was used, indicating that the method did not change the overall results.

Surprisingly, although it has previously been suggested and even assumed that cells expressing fields that either misplace or trace the object location are showing 'object location memory' (O'Keefe, 1976; Deshmukh and Knierim, 2013), as the activity of these cells can last overnight, there were no differences in the proportions of these cells between red and white box condition and no correlations with nOL discrimination scores. This was true in both sampling to probe, and probe to baseline session comparisons. This suggests that although these misplace and trace cells exist they are not a direct correlate to nOL memory, as previously thought. To my knowledge this is the only experiment that has correlated this type of place field stability with a simultaneous memory read-out. These so-called memory cells have also been shown in the CA3 (Deshmukh and Knierim, 2013), the lateral entorhinal cortex (Tsao et al., 2013) and the anterior cingulate cortex (Weible et al., 2009, 2012), with the latter two showing firing specific to both the context and old object locations that can last multiple days. Like the studies describing trace and misplace firing in CA1 cells, those describing LEC trace cells did not directly correlate this activity to a behavioural read-out. One ACC study did show that the neuronal differences correlated directly with object recognition memory strength, but not with object location memory strength (Weible et al., 2012), possibly suggesting that larger amounts of novelty are needed to record a neural correlate of memory.

It is also possible that whilst a subset of cells clearly traces the object location throughout the duration of the probe trial in the current study, a different subset of 'mismatch' cells could be firing for just the first seconds or minutes needed for the animal to learn the new object location configuration, as reported by Fyhn et al. (2002). It could be this latter subset that directly correlates to object location memory expression. Therefore, repeating the

analyses, but only for the first 30s or 1 min of each trial might yield differences between the red and white box conditions.

In the last baseline trial a small population of cells expressed place fields that traced the location of objects during the probe trial, and another population expressed fields near objects during the probe trial that subsequently disappeared during the baseline session. Both of these types of place field movement negatively correlated to the expression of nOL memory. Although initially surprising, both of these types of place field movement require place fields to be within the vicinity of objects during the probe trial. Analysis comparing sampling to probe sessions showed that the number of cells expressing new fields away from object locations also correlated to nOL memory strength. It was also shown that cells expressing fields near objects during the probe trial are associated with decreased strength of nOL memory. These results suggest that different types of analysis lead to the same conclusions.

It appears that the strength of nOL memory is directly correlated to the location of place fields in the probe trial. The more place fields that appear away from objects, the stronger the expression of memory. This suggests that there are significant changes between the sampling and probe sessions, and that these can correlate to spatial memory. It is possible that these overall changes away from objects actually signify fleeting changes occurring in the vicinity of objects in the first 30s or 1 min of the probe trial. However because the entire 10-minute probe trial was analysed together these 'mismatch' cells cannot be measured as they become silent or move within the environment in reaction to the change in object locations. Therefore although it is interesting that areas away from objects appear to be directly associated with nOL memory strength, this could be a secondary effect to changes occurring near objects. For example, if 'mismatch' cells expressed fields near object locations for the first 30s of the probe trial, and then subsequently remapped elsewhere in the environment, due to the relative proportions of object and no object areas it is likely that these cells would express fields away from object locations for the remainder of the probe session. Therefore a higher proportion of 'mismatch' cells could lead to an increase in place fields appearing away from object locations if the entire 10-minute session is analysed together. As there were no changes in place field movement that directly correlated to the

strength of nOL memory when quadrants were used for analysis, rather than object areas, this is a distinct possibility that requires further testing to investigate conclusively.

4.5 Conclusions

The addition of objects to a familiar environment leads to a partial destabilisation of the spatial map. Object movement can clearly modulate many different place cell properties. As found previously in a number of studies, place fields can trace object locations after object removal, which could be a correlate of object 'memory'. Surprisingly the enhanced expression of nOL memory seen in the red box condition is only expressed at a detectable level in areas not containing objects. It is possible that any expression of nOL memory at either the cell level or the population level in the activity of hippocampal CA1 place cells is only detectable whilst the object locations remain novel. For example, it is possible that whilst object identity-object location associations are encoded at higher rates in the red box condition, via the enhanced replay of synchronous cell assemblies, this is not detectable during the whole duration of the probe trial. Future experiments exploring this possibility are discussed further in the general discussion.

5. Conclusions and Future Directions

The experiments described in this dissertation were driven by previous studies into the benefits of reduced retroactive interference on human episodic memory (Craig et al., 2016; Dewar et al., 2014). Craig and Dewar investigated how the levels of interference in the post-learning period could affect the retention of the learned material. These studies focused on diversion retroactive interference (diversion-RI), which is interference due to material that is dissimilar to the learned material. Through many experiments it was concluded that an unfilled condition termed wakeful rest, where participants would sit alone in a quiet and dark room, was beneficial to memory consolidation; whereas a filled condition containing a task that required some sort of mental exertion, such as spot-the-difference, would lead to decreased memory retention. We sought to replicate these studies in rats to identify what post-learning conditions are required for benefit from wakeful rest, and whether place cell "memory", a possible neuronal analogue of spatial memory, was underlying these memory enhancements.

5.1 What post-learning conditions are required for the benefit of wakeful rest on spatial memory?

Interestingly my data indicate that both social isolation and decreased visual input are required to produce the beneficial effects of wakeful rest in rats (i.e. the black box effect). Social interaction or visual input during the first hour of memory consolidation is enough to prevent the retention of long-term spatial memory. Whilst this confirms that diversion-RI exists in non-human species, and that the reduction of such interference improves memory retention in rats, these results suggest that neither new learning nor 'mental exertion' are necessary for interference to occur. This is in contrast to both Dewar (2007) and Muller and Pilzecker (1900), who assumed that attending to a task was required to cause diversion-RI. Our results suggest that the task used by Dewar wasn't promoting unconscious learning of new material per se, and rather was leading to suboptimal consolidation. Therefore being alone in the dark was allowing consolidation processes to work optimally, possibly through increased communication within the hippocampus or between the hippocampus and cortical structures through sharp wave ripple events.

These results also suggest that minimal amounts of normal everyday activities can be detrimental to the consolidation of new information, which could have a large impact on everyday life. This demonstrates just how detrimental interference could be to patients with damage outside of the temporal lobes. If simply being awake and interacting with someone in a well-lit environment can impede memory consolidation processes, then it is no wonder that patients who are more susceptible to these types of retroactive interference cannot consciously acquire new long term declarative memories. It would be fascinating to repeat these types of studies in both amnesiac patients and healthy controls. I would predict that patients left alone in a light room during the delay period between learning and recall would recall significantly less learned material compared to patients left alone in a dark room. This could also be true for patients left in a dark room with other people to interact with, compared to patients left by themselves in a dark room.

Thus, my behavioural results show that the 'Black Box effect' is a phenomenon conserved between humans and rodents, and further our understanding into the exact post-learning conditions required for benefit from wakeful rest.

5.2 Do place cells show enhanced "memory" when interference is reduced via wakeful rest?

Behavioural novel object location (nOL) memory studies indicated that wakeful rest (red box condition) enhanced long-term spatial memory, whereas increased visual input (white box condition) prevented the expression of long-term spatial memory. Wakeful rest therefore enhanced the consolidation of the object locations experienced during the sampling trial, making the old object location more familiar in the probe trial. Whilst nOL memory involved learning the locations of novel objects within a familiar environment, the first place cell experiment assessed changes in neuronal activity that occurred during the first three exposures to a novel environment (Chapter 2). Although my data suggest that the environment was not fully familiar to the animal by the third exposure, as indicated by no changes in place cell firing rate or number of place fields across the three exposures, place cell "memory" of the environment appeared to be enhanced by wakeful rest.

When properties associated with place cell "memory" were analysed it was shown that animals exposed to the red box condition after the first exposure to the environment

expressed fields that were more stable across the three exposures. The location and rate of fields expressed in the initial exposure remained stable between the first, second and third exposures (6h and 24h later) if the animal was exposed to the red box condition directly after the initial exposure. In contrast, if the animal was exposed to the white box condition then there was significantly less stability between the first and second exposures, and only after the second exposure at 6h were the place fields stable (i.e. between the 6 h and 24 tests). Therefore wakeful rest can modulate place field stability in the CA1. This suggests that the consolidation of the 'spatial map' and the strength of a purely spatial memory, i.e. familiarity of an environment, benefitted from wakeful rest through mechanisms associated with place cell "memory" consolidation.

Surprisingly this enhancement of place field stability was not seen at the cell assembly level. This suggests that a population of individual place cells is sufficient to encode a memory that is purely spatial and with no need to form associations through cell assembly pairs. It has been suggested previously that only in these conditions can place cells act as "true place cells", as very rarely do populations of place cells appear to encode purely spatial information.

5.3 Do place field properties associated with memory for object locations in the nOL task differ between the wakeful rest and filled delay conditions?

Interestingly when directly investigating how place cell properties changed during a nOL behavioural experiment (Chapter 3), properties previously associated with place cell "memory" did not correlate with the expression of nOL memory. Place field stability decreased upon the introduction of objects into the familiar environment and their subsequent movement, however this occurred in both conditions equally. These results imply that the encoding of novel object locations within a familiar environment is very different to the encoding of a novel environment itself, suggesting that object-place associations are encoded differently to purely spatial information. From further analysis into place field movement it appeared that the destabilisation of the spatial map seen was at least in part due to place fields moving in response to the objects appearing, moving and disappearing throughout the behavioural trials. However the only type of place field movement that correlated with nOL memory strength appeared to be the expression of new fields away from object locations during the probe trials. Place cell movement that has

conventionally been associated with spatial memory, such as cells that trace the old object location when that object has been moved, did not correlate to nOL memory strength. These results suggested that whilst place cell stability and familiarity of an environment are relatively easy to measure, as properties are unlikely to drastically change throughout the recording session, place cells signalling small changes in a familiar environment are much harder to record. As has been shown with "mismatch" cells in previous studies (Fyhn et al., 2002), firing that is related to the detection of these changes in the environment disappears rapidly as the animal learns and assimilates this new information. I propose that the best way to investigate whether the animal has encoded and fully consolidated object locations would be to record place cells representing object-place associations during these consolidation periods rather than during the subsequent probe sessions. It is hypothesised that cell assembly pairs encode such object-place associations, especially as our data showed that the synchronous firing of these cell pairs associated with intrinsic gamma oscillations increased when objects were added to the environment.

Therefore, a future experiment that could be enlightening would be to record the reactivation of cell assembly pairs during sharp wave ripples whilst the rat is in either the white or red box. I would predict that cell assembly pairs that were synchronous during the sampling session would be reactivated more than pairs that weren't synchronous. This is proposed in Larkin et al. (2014) where an increase in coordinated synchronous cell assembly firing during awake sharp wave ripples in the probe session persisted into the subsequent rest session, implying that these cell assemblies would be selectively replayed and consolidated. The reactivation of these synchronous cell assembly pairs would therefore represent the consolidation of object identity-object location associations encoded during the sampling session. I would further predict that the reactivation of these synchronous pairs would be greater in the red box condition, directly correlating to nOL memory expression. This would suggest that animals in the red condition undergo enhanced consolidation through the reactivation of cell assembly pairs, and that this enhanced consolidation protects the animals from any further interference that might be present once the animals are put back into their home cage.

It is possible, however, that even with further experiments no cellular representation of object location or even object-place associations would be found in the CA1 of the hippocampus. Therefore alternative hypotheses need to be explored. The entorhinal cortex (EC) is the primary source of cortical input into the hippocampus (van Strien et al., 2009; Witter et al., 2017). The parallel processing model dictates that this input is anatomically segregated into spatial and non-spatial functional roles, with the former being processed in the MEC and the latter in the LEC (Burwell, 2000). This information is then thought to be combined in the hippocampus to form a complete episodic memory. Possible cellular correlates of object memory within the EC have been shown previously in various tasks. When single cells were recorded during a contextual odour discrimination task, neurons in both the MEC and LEC were highly selective for both object and spatial dimensions of a task, with the former prioritizing location information and the latter prioritizing object information (Keene et al., 2016). A subset of LEC neurons have also been shown to fire in close proximity to a novel object when the animal explored said object in a familiar open field (Tsao et al., 2013). Interestingly when this object was removed from the environment, a different subset of LEC neurons were shown to become active at locations where objects had previously been in the environment (i.e. trace cells). Unlike the 'mismatch' cells recorded in the CA1, which fired only in the first few seconds of novelty detection (Fyhn et al., 2002), LEC 'trace' cells fired in a stable fashion for at least the whole recording session. If the object memory was sufficiently strong, stable trace firing could be recorded weeks after the removal of the object was initially detected. Together these neuronal representations of past and present object locations in the LEC could provide two types of input into the hippocampus required for the retrieval of novel object location memory.

Other studies in mice have shown that cells in the ACC can also respond to the "memory" of past object locations, with 'trace' cells firing when one of two familiar objects was removed from an open field. Like the 'trace' cells in the LEC, these ACC cells fired in the same location weeks after the object was removed (Weible et al., 2012). The ACC has previously been implicated in the retrieval of remote spatial memories (Bontempi et al., 1999; Teixeira et al., 2006; Weible et al., 2012). As the LEC has strong bidirectional connections to the ACC (Jones and Witter, 2007), this suggests that this circuit could be important in underlying the storage and retrieval of long-term object location memories.

The lack of distinct place cell correlates of object location memory found in Chapter 4 could therefore be explained and even predicted by the hypothesis that these correlates exist outside of the hippocampus, possibly within the MEC, LEC or ACC. It would therefore be interesting to repeat the experiments described in Chapter 4 whilst recording from the MEC, LEC and ACC. It could be predicted that animals exposed to the red box would exhibit higher selectivity of cell firing for the locations of objects, or even significantly higher 'trace' cell firing at the old object location during the probe trial. This would indicate that the enhancement of object location memory via the 'Black Box Effect' could be seen at a cellular level, just not in the CA1 of the hippocampus as previously hypothesized.

It should be noted that even if the cellular 'trace' of nOL memory is found to be extra-hippocampal, the nOL task itself remains hippocampal-dependent. This can be explained by the role of the hippocampus in episodic memory, where it is thought to act as part of an associative network, grouping memories together in time and space (Eichenbaum et al., 2012). Therefore even if the 'trace' of object location memory is stored outside of the hippocampus, the consolidation and retrieval of such memory would require the spatio-temporal framework created by the hippocampus. The 'Black Box Effect' could therefore be enhancing the possible representations of object location memory outside of the hippocampus, or the hippocampal framework that is thought to link these memories within the context of the nOL task.

5.4 Concluding remarks

This thesis emphasises the importance of taking time after learning to reduce incoming sensory information, allowing consolidation processes to work optimally. It also serves to highlight the impact that experimental design can have on the outcome of a study. What an animal experiences directly after learning can impact both the retention of spatial memory, and mechanisms associated with place cell "memory" consolidation. Therefore it is important to take into consideration and control post-learning stimuli, such as visual stimulation and social interaction, to prevent diversion-RI mechanisms affecting experimental outcome.

References

Aggleton, J.P., Hunt, P.R., and Rawlins, J.N.P. (1986). The effects of hippocampal lesions upon spatial and non-spatial tests of working memory. Behav. Brain Res. *19*, 133–146.

Agnihotri, N.T., Hawkins, R.D., Kandel, E.R., and Kentros, C. (2004). The long-term stability of new hippocampal place fields requires new protein synthesis. Proc. Natl. Acad. Sci. U. S. A. *101*, 3656–3661.

Alber, J., Della Sala, S., and Dewar, M. (2014). Minimizing interference with early consolidation boosts 7-day retention in amnesic patients. Neuropsychology *28*, 667–675.

Allen, K., Rawlins, N.P., Bannerman, D.M., and Csicsvari, J. (2012). Hippocampal place cells can encode multiple trial-dependent features through rate remapping. J. Neurosci. Off. J. Soc. Neurosci. *32*, 14752–14766.

Alvarez, P., and Squire, L.R. (1994). Memory consolidation and the medial temporal lobe: a simple network model. Proc. Natl. Acad. Sci. U. S. A. *91*, 7041–7045.

Anderson, M.I., and Jeffery, K.J. (2003). Heterogeneous Modulation of Place Cell Firing by Changes in Context. J. Neurosci. *23*, 8827–8835.

Axmacher, N., Elger, C.E., and Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. Brain *131*, 1806–1817.

Barnes, C.A., Suster, M.S., Shen, J., and McNaughton, B.L. (1997). Multistability of cognitive maps in the hippocampus of old rats. Nature *388*, 272–275.

Battaglia, F.P., Sutherland, G.R., and McNaughton, B.L. (2004). Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions. Learn. Mem. Cold Spring Harb. N *11*, 697–704.

Ben-Yakov, A., Rubinson, M., and Dudai, Y. (2014). Shifting Gears in Hippocampus: Temporal Dissociation between Familiarity and Novelty Signatures in a Single Event. J. Neurosci. *34*, 12973–12981.

Bett, D., Stevenson, C.H., Shires, K.L., Smith, M.T., Martin, S.J., Dudchenko, P.A., and Wood, E.R. (2013). The postsubiculum and spatial learning: the role of postsubicular synaptic activity and synaptic plasticity in hippocampal place cell, object, and object-location memory. J. Neurosci. Off. J. Soc. Neurosci. *33*, 6928–6943.

Bolhuis, J.J., Stewart, C.A., and Forrest, E.M. (1994). Retrograde amnesia and memory reactivation in rats with ibotenate lesions to the hippocampus or subiculum. Q. J. Exp. Psychol. B *47*, 129–150.

Bonnici, H.M., and Maguire, E.A. (2018). Two years later – Revisiting autobiographical memory representations in vmPFC and hippocampus. Neuropsychologia *110*, 159–169.

Bonnici, H.M., Chadwick, M.J., Lutti, A., Hassabis, D., Weiskopf, N., and Maguire, E.A. (2012). Detecting Representations of Recent and Remote Autobiographical Memories in vmPFC and Hippocampus. J. Neurosci. *32*, 16982–16991.

Bontempi, B., Laurent-Demir, C., Destrade, C., and Jaffard, R. (1999). Time-dependent reorganization of brain circuitry underlying long-term memory storage. Nature 400, 671–675.

Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. Hippocampus *1*, 193–205.

Brandon, M.P., Koenig, J., Leutgeb, J.K., and Leutgeb, S. (2014). New and distinct hippocampal place codes are generated in a new environment during septal inactivation. Neuron *82*, 789–796.

Britt, S.H. (1935). Retroactive inhibition: a review of the literature. Psychol. Bull. *32*, 381–440.

Broadbent, N.J., Squire, L.R., and Clark, R.E. (2006). Reversible hippocampal lesions disrupt water maze performance during both recent and remote memory tests. Learn. Mem. Cold Spring Harb. N *13*, 187–191.

Brooks, D.N., and Baddeley, A.D. (1976). What can amnesic patients learn? Neuropsychologia *14*, 111–122.

Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C., and Wilson, M.A. (1998). A Statistical Paradigm for Neural Spike Train Decoding Applied to Position Prediction from Ensemble Firing Patterns of Rat Hippocampal Place Cells. J. Neurosci. *18*, 7411–7425.

Burke, S.N., Maurer, A.P., Nematollahi, S., Uprety, A.R., Wallace, J.L., and Barnes, C.A. (2011). The Influence of Objects on Place Field Expression and Size in Distal Hippocampal CA1. Hippocampus *21*, 783–801.

Burwell, R.D. (2000). The parahippocampal region: corticocortical connectivity. Ann. N. Y. Acad. Sci. 911, 25–42.

Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. Neuroscience *31*, 551–570.

Buzsáki, G. (1996). The Hippocampo-Neocortical Dialogue. Cereb. Cortex *6*, 81–92.

Buzsáki, G. (2010). Neural Syntax: Cell Assemblies, Synapsembles, and Readers. Neuron *68*, 362–385.

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. Hippocampus *25*, 1073–1188.

Buzsáki, G., Leung, L.W., and Vanderwolf, C.H. (1983). Cellular bases of hippocampal EEG in the behaving rat. Brain Res. *287*, 139–171.

Carr, M.F., Jadhav, S.P., and Frank, L.M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. Nat. Neurosci. *14*, 147–153.

Castellucci, V., Pinsker, H., Kupfermann, I., and Kandel, E.R. (1970). Neuronal mechanisms of habituation and dishabituation of the gill-withdrawal reflex in Aplysia. Science *167*, 1745–1748.

Cermak, L.S., and O'Connor, M. (1983). The anterograde and retrograde retrieval ability of a patient with amnesia due to encephalitis. Neuropsychologia *21*, 213–234.

Chrobak, J.J., and Buzsáki, G. (1996). High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. J. Neurosci. Off. J. Soc. Neurosci. *16*, 3056–3066.

Clark, R.E., Broadbent, N.J., and Squire, L.R. (2005). Hippocampus and Remote Spatial Memory in Rats. Hippocampus *15*, 260–272.

Clemens, Z., Mölle, M., Erőss, L., Barsi, P., Halász, P., and Born, J. (2007). Temporal coupling of parahippocampal ripples, sleep spindles and slow oscillations in humans. Brain *130*, 2868–2878.

Cohen, N.J., and Squire, L.R. (1980). Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that. Science *210*, 207–210.

Colavito, V., Fabene, P.F., Grassi-Zucconi, G., Pifferi, F., Lamberty, Y., Bentivoglio, M., and Bertini, G. (2013). Experimental sleep deprivation as a tool to test memory deficits in rodents. Front. Syst. Neurosci. *7*.

Colgin, L.L., Moser, E.I., and Moser, M.-B. (2008). Understanding memory through hippocampal remapping. Trends Neurosci. *31*, 469–477.

Colley, P.A., Sheu, F.S., and Routtenberg, A. (1990). Inhibition of protein kinase C blocks two components of LTP persistence, leaving initial potentiation intact. J. Neurosci. *10*, 3353–3360.

Corey, S.M. (1931). The relationship between compulsory physical exercise and the ability of the white rat to learn and relearn an elevated skeleton maze. J. Comp. Psychol. *11*, 291–318.

Corkin, S. (2002). What's new with the amnesic patient H.M.? Nat. Rev. Neurosci. *3*, 153–160.

Corkin, S., Amaral, D.G., Gilberto González, R., Johnson, K.A., and Hyman, B.T. (1997). H. M.'s Medial temporal lobe lesion: Findings from magnetic resonance imaging. J. Neurosci. *17*, 3964–3979.

Cowan, N., Beschin, N., and Della Sala, S. (2004). Verbal recall in amnesiacs under conditions of diminished retroactive interference. Brain *127*, 825–834.

Craig, M., Dewar, M., Sala, S.D., and Wolbers, T. (2015). Rest boosts the long-term retention of spatial associative and temporal order information. Hippocampus *25*, 1017–1027.

Craig, M., Dewar, M., Harris, M.A., Sala, S.D., and Wolbers, T. (2016). Wakeful rest promotes the integration of spatial memories into accurate cognitive maps. Hippocampus *26*, 185–193.

Damasio, A.R., Eslinger, P.J., Damasio, H., Van Hoesen, G.W., and Cornell, S. (1985). Multimodal amnesic syndrome following bilateral temporal and basal forebrain damage. Arch. Neurol. *42*, 252–259.

Davidson, T.J., Kloosterman, F., and Wilson, M.A. (2009). Hippocampal replay of extended experience. Neuron *63*, 497–507.

Davis, H.P., and Squire, L.R. (1984). Protein synthesis and memory: A review. Psychol. Bull. *96*, 518–559.

Della Sala, S., Cowan, N., Beschin, N., and Perini, M. (2005). Just lying there, remembering: improving recall of prose in amnesic patients with mild cognitive impairment by minimising interference. Mem. Hove Engl. *13*, 435–440.

Derkach, V.A., Oh, M.C., Guire, E.S., and Soderling, T.R. (2007). Regulatory mechanisms of AMPA receptors in synaptic plasticity. Nat. Rev. Neurosci. *8*, 101–113.

Deshmukh, S.S., and Knierim, J.J. (2013). Influence of local objects on hippocampal representations: landmark vectors and memory. Hippocampus *23*.

Deuker, L., Olligs, J., Fell, J., Kranz, T.A., Mormann, F., Montag, C., Reuter, M., Elger, C.E., and Axmacher, N. (2013). Memory Consolidation by Replay of Stimulus-Specific Neural Activity. J. Neurosci. *33*, 19373–19383.

Dewar, M., Garcia, Y.F., Cowan, N., and Sala, S.D. (2009). Delaying Interference Enhances Memory Consolidation in Amnesic Patients. Neuropsychology *23*, 627–634.

Dewar, M., Della Sala, S., Beschin, N., and Cowan, N. (2010). Profound retroactive interference in anterograde amnesia: What interferes? Neuropsychology *24*, 357–367.

Dewar, M., Alber, J., Butler, C., Cowan, N., and Della Sala, S. (2012). Brief Wakeful Resting Boosts New Memories Over the Long Term. Psychol. Sci. *23*, 955–960.

Dewar, M., Alber, J., Cowan, N., and Sala, S.D. (2014). Boosting Long-Term Memory via Wakeful Rest: Intentional Rehearsal Is Not Necessary, Consolidation Is Sufficient. PLOS ONE *9*, e109542.

Dewar, M.T., Cowan, N., and Della Sala, S. (2007). Forgetting due to retroactive interference: A fusion of early insights into everyday forgetting and recent research on anterograde amnesia. Cortex J. Devoted Study Nerv. Syst. Behav. *43*, 616–634.

Dey, M.K. (1969). Retroactive inhibition as a function of similarity of meaning in free-recall learning. Psychol. Forsch. *33*, 79–84.

Diba, K., and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. Nat. Neurosci. *10*, 1241–1242.

Dix, S.L., and Aggleton, J.P. (1999). Extending the spontaneous preference test of recognition: evidence of object-location and object-context recognition. Behav. Brain Res. *99*, 191–200.

Dombeck, D.A., Harvey, C.D., Tian, L., Looger, L.L., and Tank, D.W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. Nat. Neurosci. *13*, 1433–1440.

Dräger, U.C. (1975). Receptive fields of single cells and topography in mouse visual cortex. J. Comp. Neurol. *160*, 269–289.

Dragoi, G., Harris, K.D., and Buzsáki, G. (2003). Place representation within hippocampal networks is modified by long-term potentiation. Neuron *39*, 843–853.

Dudai, Y., and Morris, R.G.M. (2000). To consolidate or not to consolidate: what are the questions? (Oxford University Press).

Dudai, Y., Karni, A., and Born, J. (2015). The Consolidation and Transformation of Memory. Neuron *88*, 20–32.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010a). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. *13*, 995–1002.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010b). The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. *13*, 995–1002.

Ego-Stengel, V., and Wilson, M.A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. Hippocampus *20*, 1–10.

Eichenbaum, H. (2017). On the Integration of Space, Time, and Memory. Neuron *95*, 1007–1018.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., and Tanila, H. (1999). The hippocampus, memory, and place cells: is it spatial memory or a memory space? Neuron *23*, 209–226.

Eichenbaum, H., Sauvage, M., Fortin, N., Komorowski, R., and Lipton, P. (2012). Towards a functional organization of episodic memory in the medial temporal lobe. Neurosci. Biobehav. Rev. 36, 1597–1608.

Ekstrom, A.D., Kahana, M.J., Caplan, J.B., Fields, T.A., Isham, E.A., Newman, E.L., and Fried, I. (2003). Cellular networks underlying human spatial navigation. Nature *425*, 184–188.

Erickson, K.I., Voss, M.W., Prakash, R.S., Basak, C., Szabo, A., Chaddock, L., Kim, J.S., Heo, S., Alves, H., White, S.M., et al. (2011). Exercise training increases size of hippocampus and improves memory. Proc. Natl. Acad. Sci. *108*, 3017–3022.

Fellini, L., and Morellini, F. (2013). Mice Create What–Where–When Hippocampus-Dependent Memories of Unique Experiences. J. Neurosci. *33*, 1038–1043.

Fenton, A.A., and Muller, R.U. (1998). Place cell discharge is extremely variable during individual passes of the rat through the firing field. Proc. Natl. Acad. Sci. U. S. A. *95*, 3182–3187.

Fenton, A.A., Kao, H.-Y., Neymotin, S.A., Olypher, A., Vayntrub, Y., Lytton, W.W., and Ludvig, N. (2008). Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly-arranged, and expanded place fields in the larger space. J. Neurosci. Off. J. Soc. Neurosci. *28*, 11250–11262.

Foster, D.J., and Wilson, M.A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. Nature *440*, 680–683.

Fowler, M.J., Sullivan, M.J., and Ekstrand, B.R. (1973). Sleep and memory. Science *179*, 302–304.

Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal Plasticity across Multiple Days of Exposure to Novel Environments. J. Neurosci. *24*, 7681–7689.

Frey, U., Krug, M., Reymann, K.G., and Matthies, H. (1988). Anisomycin, an inhibitor of protein synthesis, blocks late phases of LTP phenomena in the hippocampal CA1 region in vitro. Brain Res. *452*, 57–65.

Fyhn, M., Molden, S., Hollup, S., Moser, M.-B., and Moser, E.I. (2002). Hippocampal Neurons Responding to First-Time Dislocation of a Target Object. Neuron *35*, 555–566.

Gabrieli, J.D.E., Cohen, N.J., and Corkin, S. (1988). The impaired learning of semantic knowledge following bilateral medial temporal-lobe resection. Brain Cogn. *7*, 157–177.

Gais, S., and Born, J. (2004). Low acetylcholine during slow-wave sleep is critical for declarative memory consolidation. Proc. Natl. Acad. Sci. U. S. A. *101*, 2140–2144.

Gaskin, S., Gamliel, A., Tardif, M., Cole, E., and Mumby, D.G. (2009). Incidental (unreinforced) and reinforced spatial learning in rats with ventral and dorsal lesions of the hippocampus. Behav. Brain Res. *202*, 64–70.

Girardeau, G., Benchenane, K., Wiener, S.I., Buzsáki, G., and Zugaro, M.B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. Nat. Neurosci. *12*, 1222–1223.

Goelet, P., Castellucci, V.F., Schacher, S., and Kandel, E.R. (1986). The long and the short of long–term memory—a molecular framework. Nature *322*, 419–422.

Gothard, K.M., Skaggs, W.E., and McNaughton, B.L. (1996). Dynamics of Mismatch Correction in the Hippocampal Ensemble Code for Space: Interaction between Path Integration and Environmental Cues. J. Neurosci. *16*, 8027–8040.

Gray, W.L. (1937). The effect of forced activity on the maze performance of white rats. J. Comp. Psychol. *23*, 475–512.

Grecksch, G., and Matthies, H. (1980). Two sensitive periods for the amnesic effect of anisomycin. Pharmacol. Biochem. Behav. *12*, 663–665.

Guzowski, J.F., and McGaugh, J.L. (1997). Antisense oligodeoxynucleotide-mediated disruption of hippocampal cAMP response element binding protein levels impairs consolidation of memory for water maze training. Proc. Natl. Acad. Sci. *94*, 2693–2698.

Hampson, R.E., Heyser, C.J., and Deadwyler, S.A. (1993). Hippocampal cell firing correlates of delayed-match-to-sample performance in the rat. Behav. Neurosci. *107*, 715–739.

Hampton, R.R., Hampstead, B.M., and Murray, E.A. (2004). Selective hippocampal damage in rhesus monkeys impairs spatial memory in an open-field test. Hippocampus *14*, 808–818.

Hardt, O., Migues, P.V., Hastings, M., Wong, J., and Nader, K. (2010). PKMζ maintains 1-day- and 6-day-old long-term object location but not object identity memory in dorsal hippocampus. Hippocampus *20*, 691–695.

Hardt, O., Nader, K., and Nadel, L. (2013). Decay happens: the role of active forgetting in memory. Trends Cogn. Sci. *17*, 111–120.

Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. Nature *424*, 552–556.

Hasselmo, null (1999). Neuromodulation: acetylcholine and memory consolidation. Trends Cogn. Sci. *3*, 351–359.

Hazan, L., Zugaro, M., and Buzsáki, G. (2006). Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. J. Neurosci. Methods *155*, 207–216.

Hernandez, P.J., and Abel, T. (2008). The role of protein synthesis in memory consolidation: Progress amid decades of debate. Neurobiol. Learn. Mem. *89*, 293–311.

Hori, E., Nishio, Y., Kazui, K., Umeno, K., Tabuchi, E., Sasaki, K., Endo, S., Ono, T., and Nishijo, H. (2005). Place-related neural responses in the monkey hippocampal formation in a virtual space. Hippocampus *15*, 991–996.

Huang, Y.-Y., Colley, P.A., and Routtenberg, A. (1992). Postsynaptic then presynaptic protein kinase C activity may be necessary for long-term potentiation. Neuroscience *49*, 819–827.

Huang, Y.Y., Nguyen, P.V., Abel, T., and Kandel, E.R. (1996). Long-lasting forms of synaptic potentiation in the mammalian hippocampus. Learn. Mem. *3*, 74–85.

Igarashi, K.M., Lu, L., Colgin, L.L., Moser, M.-B., and Moser, E.I. (2014). Coordination of entorhinal-hippocampal ensemble activity during associative learning. Nature *510*, 143–147.

Izquierdo, I., and Medina, J.H. (1997). Memory formation: the sequence of biochemical events in the hippocampus and its connection to activity in other brain structures. Neurobiol. Learn. Mem. *68*, 285–316.

Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. Science *336*, 1454–1458.

Jerusalinsky, D., Quillfeldt, J.A., Walz, R., Da Silva, R.C., Medina, J.H., and Izquierdo, I. (1994). Post-training intrahippocampal infusion of protein kinase C inhibitors causes amnesia in rats. Behav. Neural Biol. *61*, 107–109.

Jones, B.F., and Witter, M.P. (2007). Cingulate cortex projections to the parahippocampal region and hippocampal formation in the rat. Hippocampus 17, 957–976.

Kadir, S.N., Goodman, D.F.M., and Harris, K.D. (2014). High-dimensional cluster analysis with the Masked EM Algorithm. Neural Comput. *26*, 2379–2394.

Karlsson, M.P., and Frank, L.M. (2008). Network dynamics underlying the formation of sparse, informative representations in the hippocampus. J. Neurosci. Off. J. Soc. Neurosci. *28*, 14271–14281.

Keene, C.S., Bladon, J., McKenzie, S., Liu, C.D., O'Keefe, J., and Eichenbaum, H. (2016). Complementary Functional Organization of Neuronal Activity Patterns in the Perirhinal, Lateral Entorhinal, and Medial Entorhinal Cortices. J. Neurosci. Off. J. Soc. Neurosci. 36, 3660–3675.

Kendall, M.G. (1938). A New Measure of Rank Correlation. Biometrika *30*, 81–93.

Kentros, C. (2006). Hippocampal place cells: The "where" of episodic memory? Hippocampus *16*, 743–754.

Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of Long-Term Stability of New Hippocampal Place Cell Maps by NMDA Receptor Blockade. Science *280*, 2121–2126.

Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. Neuron *42*, 283–295.

Kim, J.J., Clark, R.E., and Thompson, R.F. (1995). Hippocampectomy impairs the memory of recently, but not remotely, acquired trace eyeblink conditioned responses. Behav. Neurosci. *109*, 195–203.

Komorowski, R.W., Manns, J.R., and Eichenbaum, H. (2009a). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. J. Neurosci. Off. J. Soc. Neurosci. *29*, 9918–9929.

Komorowski, R.W., Manns, J.R., and Eichenbaum, H. (2009b). Robust Conjunctive Item-Place Coding by Hippocampal Neurons Parallels Learning What Happens Where. J. Neurosci. *29*, 9918–9929.

Kopelman, M.D., and Kapur, N. (2001). The loss of episodic memories in retrograde amnesia: single-case and group studies. Philos. Trans. R. Soc. Lond. Ser. B *356*, 1409–1421.

Kopelman, M.D., Wilson, B.A., and Baddeley, A.D. (1989). The autobiographical memory interview: A new assessment of autobiographical and personal semantic memory in amnesic patients. J. Clin. Exp. Neuropsychol. *11*, 724–744.

Kovács, K.A., O'Neill, J., Schoenenberger, P., Penttonen, M., Guerrero, D.K.R., and Csicsvari, J. (2016). Optogenetically Blocking Sharp Wave Ripple Events in Sleep Does Not Interfere with the Formation of Stable Spatial Representation in the CA1 Area of the Hippocampus. PLOS ONE *11*, e0164675.

Kraus, B.J., Robinson, R.J., White, J.A., Eichenbaum, H., and Hasselmo, M.E. (2013). Hippocampal "Time Cells": Time versus Path Integration. Neuron *78*, 1090–1101.

Kubie, J.L. (1984). A driveable bundle of microwires for collecting single-unit data from freely-moving rats. Physiol. Behav. *32*, 115–118.

Kubie, J.L., and Muller, R.U. (1991). Multiple representations in the hippocampus. Hippocampus *1*, 240–242.

Kubie, J.L., and Ranck, J.B. (1982). Tonic and Phasic Firing of Rat Hippocampal Complex-Spike Cells in Three Different Situations: Context and Place. In Conditioning: Representation of Involved Neural Functions, C.D. Woody, ed. (Boston, MA: Springer US), pp. 89–98.

Kudrimoti, H.S., Barnes, C.A., and McNaughton, B.L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. J. Neurosci. Off. J. Soc. Neurosci. *19*, 4090–4101.

Kyd, R.J., and Bilkey, D.K. (2005). Hippocampal Place Cells Show Increased Sensitivity to Changes in the Local Environment Following Prefrontal Cortex Lesions. Cereb. Cortex *15*, 720–731.

Ladenbauer, J., Külzow, N., Passmann, S., Antonenko, D., Grittner, U., Tamm, S., and Flöel, A. (2016). Brain stimulation during an afternoon nap boosts slow oscillatory activity and memory consolidation in older adults. NeuroImage *142*, 311–323.

Ladenbauer, J., Ladenbauer, J., Külzow, N., de Boor, R., Avramova, E., Grittner, U., and Flöel, A. (2017). Promoting Sleep Oscillations and Their Functional Coupling by Transcranial Stimulation Enhances Memory Consolidation in Mild Cognitive Impairment. J. Neurosci. Off. J. Soc. Neurosci. *37*, 7111–7124.

Lamprecht, R., and LeDoux, J. (2004). Structural plasticity and memory. Nat. Rev. Neurosci. *5*, 45–54.

Larkin, M.C., Lykken, C., Tye, L.D., Wickelgren, J.G., and Frank, L.M. (2014). Hippocampal output area CA1 broadcasts a generalized novelty signal during an object-place recognition task. Hippocampus *24*, 773–783.

Lattal, K.M., and Abel, T. (2001). Different Requirements for Protein Synthesis in Acquisition and Extinction of Spatial Preferences and Context-Evoked Fear. J. Neurosci. *21*, 5773–5780.

Lavenex, P.B., Amaral, D.G., and Lavenex, P. (2006). Hippocampal lesion prevents spatial relational learning in adult macaque monkeys. J. Neurosci. Off. J. Soc. Neurosci. *26*, 4546–4558.

Lee, A.K., and Wilson, M.A. (2002). Memory of Sequential Experience in the Hippocampus during Slow Wave Sleep. Neuron *36*, 1183–1194.

Lenck-Santini, P.-P., Rivard, B., Muller, R.U., and Poucet, B. (2005). Study of CA1 place cell activity and exploratory behavior following spatial and nonspatial changes in the environment. Hippocampus *15*, 356–369.

Leutgeb, J.K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C.A., McNaughton, B.L., Moser, M.-B., and Moser, E.I. (2005a). Progressive Transformation of Hippocampal Neuronal Representations in "Morphed" Environments. Neuron *48*, 345–358.

Leutgeb, J.K., Leutgeb, S., Moser, M.-B., and Moser, E.I. (2007). Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. Science *315*, 961–966.

Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., and Moser, M.-B. (2005b). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science *309*, 619–623.

Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. Nature *416*, 90–94.

Lewis, P.A., and Durrant, S.J. (2011). Overlapping memory replay during sleep builds cognitive schemata. Trends Cogn. Sci. *15*, 343–351.

Liang, K.C., Hon, W., Tyan, Y.M., and Liao, W.L. (1994). Involvement of hippocampal NMDA and AMPA receptors in acquisition, formation and retrieval of spatial memory in the Morris water maze. Chin. J. Physiol. *37*, 201–212.

Lin, X.Y., and Glanzman, D.L. (1994). Long-term potentiation of Aplysia sensorimotor synapses in cell culture: regulation by postsynaptic voltage. Proc. Biol. Sci. *255*, 113–118.

Lipp, H.-P., and Wolfer, D.P. (1998). Genetically modified mice and cognition. Curr. Opin. Neurobiol. *8*, 272–280.

Lisman, J., Buzsáki, G., Eichenbaum, H., Nadel, L., Ranganath, C., and Redish, A.D. (2017). Viewpoints: how the hippocampus contributes to memory, navigation and cognition.

Logothetis, N.K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H.C., Besserve, M., and Oeltermann, A. (2012). Hippocampal–cortical interaction during periods of subcortical silence. Nature *491*, 547–553.

MacDonald, C.J., Lepage, K.Q., Eden, U.T., and Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. Neuron *71*, 737–749.

Maingret, N., Girardeau, G., Todorova, R., Goutierre, M., and Zugaro, M. (2016). Hippocampo-cortical coupling mediates memory consolidation during sleep. Nat. Neurosci. *19*, 959–964.

Manns, J.R., and Eichenbaum, H. (2009). A cognitive map for object memory in the hippocampus. Learn. Mem. *16*, 616–624.

Manns, J.R., Hopkins, R.O., and Squire, L.R. (2003). Semantic Memory and the Human Hippocampus. Neuron *38*, 127–133.

Manns, J.R., Howard, M.W., and Eichenbaum, H. (2007). Gradual Changes in Hippocampal Activity Support Remembering the Order of Events. Neuron *56*, 530–540.

Marr, D. (1971). Simple memory: a theory for archicortex. Phil Trans R Soc Lond B *262*, 23–81.

Marrosu, F., Portas, C., Mascia, M.S., Casu, M.A., Fà, M., Giagheddu, M., Imperato, A., and Gessa, G.L. (1995). Microdialysis measurement of cortical and hippocampal acetylcholine release during sleep-wake cycle in freely moving cats. Brain Res. *671*, 329–332.

Marshall, L., Helgadóttir, H., Mölle, M., and Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. Nature *444*, 610–613.

Martin, S.J., Grimwood, P.D., and Morris, R.G. (2000). Synaptic plasticity and memory: an evaluation of the hypothesis. Annu. Rev. Neurosci. *23*, 649–711.

Martínez, M.C., Alen, N., Ballarini, F., Moncada, D., and Viola, H. (2012). Memory traces compete under regimes of limited Arc protein synthesis: Implications for memory interference. Neurobiol. Learn. Mem. *98*, 165–173.

Martínez, M.C., Villar, M.E., Ballarini, F., and Viola, H. (2014). Retroactive interference of object-in-context long-term memory: role of dorsal hippocampus and medial prefrontal cortex. Hippocampus *24*, 1482–1492.

McGeoch, J.A., and McDonald, W.T. (1931). Meaningful relation and retroactive inhibition. Am. J. Psychol. *43*, 579–588.

McHugh, T.J., Blum, K.I., Tsien, J.Z., Tonegawa, S., and Wilson, M.A. (1996a). Impaired Hippocampal Representation of Space in CA1-Specific NMDAR1 Knockout Mice. Cell *87*, 1339–1349.

McHugh, T.J., Blum, K.I., Tsien, J.Z., Tonegawa, S., and Wilson, M.A. (1996b). Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. Cell *87*, 1339–1349.

Mednick, S.C., Cai, D.J., Shuman, T., Anagnostaras, S., and Wixted, J.T. (2011). An opportunistic theory of cellular and systems consolidation. Trends Neurosci. *34*, 504–514.

Mehta, M.R., Quirk, M.C., and Wilson, M.A. (2000). Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields. Neuron *25*, 707–715.

Meiri, N., Sun, M.-K., Segal, Z., and Alkon, D.L. (1998). Memory and long-term potentiation (LTP) dissociated: Normal spatial memory despite CA1 LTP elimination with Kv1.4 antisense. Proc. Natl. Acad. Sci. U. S. A. *95*, 15037–15042.

Migues, P.V., Hardt, O., Wu, D.C., Gamache, K., Sacktor, T.C., Wang, Y.T., and Nader, K. (2010). PKMzeta maintains memories by regulating GluR2-dependent AMPA receptor trafficking. Nat. Neurosci. *13*, 630–634.

Miller, J.F., Neufang, M., Solway, A., Brandt, A., Trippel, M., Mader, I., Hefft, S., Merkow, M., Polyn, S.M., Jacobs, J., et al. (2013). Neural Activity in Human Hippocampal Formation Reveals the Spatial Context of Retrieved Memories. Science *342*, 1111–1114.

Milner, B. (1959). The memory defect in bilateral hippocampal lesions. Psychiatr. Res. Rep. Am. Psychiatr. Assoc. *11*, 43–58.

Milner, B. (2005). The medial temporal-lobe amnesic syndrome. Psychiatr. Clin. North Am. *28*, 599–611, 609.

Mölle, M., Yeshenko, O., Marshall, L., Sara, S.J., and Born, J. (2006). Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. J. Neurophysiol. *96*, 62–70.

Monaco, J.D., Rao, G., Roth, E.D., and Knierim, J.J. (2014). Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. Nat. Neurosci. *17*, 725–731.

Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. Off. J. Soc. Neurosci. *7*, 1951–1968.

Mumby, D.G., Astur, R.S., Weisend, M.P., and Sutherland, R.J. (1999). Retrograde amnesia and selective damage to the hippocampal formation: memory for places and object discriminations. Behav. Brain Res. *106*, 97–107.

Murray, E.A., and Mishkin, M. (1998). Object Recognition and Location Memory in Monkeys with Excitotoxic Lesions of the Amygdala and Hippocampus. J. Neurosci. *18*, 6568–6582.

Murray, E.A., and Wise, S.P. (2010). What, if anything, can monkeys tell us about human amnesia when they can't say anything at all? Neuropsychologia *48*, 2385–2405.

Nadel, L., and Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. Curr. Opin. Neurobiol. *7*, 217–227.

Nadel, L., Winocur, G., Ryan, L., and Moscovitch, M. (2007). Systems consolidation and hippocampus: two views. Debates Neurosci. *1*, 55–66.

Neymotin, S.A., Talbot, Z.N., Jung, J.Q., Fenton, A.A., and Lytton, W.W. (2017). Tracking recurrence of correlation structure in neuronal recordings. J. Neurosci. Methods *275*, 1–9.

Ngo, H.-V.V., Martinetz, T., Born, J., and Mölle, M. (2013). Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. Neuron *78*, 545–553.

Nitz, D., and McNaughton, B. (2004). Differential Modulation of CA1 and Dentate Gyrus Interneurons During Exploration of Novel Environments. J. Neurophysiol. *91*, 863–872.

O'Callaghan, R.M., Ohle, R., and Kelly, Á.M. (2007). The effects of forced exercise on hippocampal plasticity in the rat: A comparison of LTP, spatial- and non-spatial learning. Behav. Brain Res. *176*, 362–366.

O'Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. Exp. Neurol. *51*, 78–109.

O'Keefe, J. (1979). A review of the hippocampal place cells. Prog. Neurobiol. *13*, 419–439.

O'Keefe, J., and Conway, D.H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. Exp. Brain Res. *31*, 573–590.

O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–175.

O'Keefe, J., and Nadel, L. (1978). The hippocampus as a cognitive map (Oxford : New York: Clarendon Press ; Oxford University Press).

O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus *3*, 317–330.

O'Neill, J., Senior, T.J., Allen, K., Huxter, J.R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. Nat. Neurosci. *11*, 209–215.

Osada, T., Adachi, Y., Kimura, H.M., and Miyashita, Y. (2008). Towards understanding of the cortical network underlying associative memory. Philos. Trans. R. Soc. B Biol. Sci. *363*, 2187–2199.

Oudiette, D., Antony, J.W., Creery, J.D., and Paller, K.A. (2013). The Role of Memory Reactivation during Wakefulness and Sleep in Determining Which Memories Endure. J. Neurosci. Off. J. Soc. Neurosci. *33*, 6672–6678.

Ozawa, T., Yamada, K., and Ichitani, Y. (2011). Long-term object location memory in rats: effects of sample phase and delay length in spontaneous place recognition test. Neurosci. Lett. *497*, 37–41.

Ozawa, T., Yamada, K., and Ichitani, Y. (2017). Differential requirements of hippocampal de novo protein and mRNA synthesis in two long-term spatial memory tests: Spontaneous place recognition and delay-interposed radial maze performance in rats. PLOS ONE *12*, e0171629.

Panoz-Brown, D., Iyer, V., Carey, L.M., Sluka, C.M., Rajic, G., Kestenman, J., Gentry, M., Brotheridge, S., Somekh, I., Corbin, H.E., et al. (2018). Replay of Episodic Memories in the Rat. Curr. Biol. *28*, 1628-1634.e7.

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsáki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. Science *321*, 1322–1327.

Patel, G.H., Michael, D.K., and Snyder, L.H. (2014). Topographic organization in the brain: Searching for general principles. Trends Cogn. Sci. *18*, 351–363.

Paxinos, G., and Watson, C. (2004). The Rat Brain in Stereotaxic Coordinates–The New Coronal Set, 5th Edn.

Penfield, W., and Milner, B. (1958). Memory Deficit Produced by Bilateral Lesions in the Hippocampal Zone. AMA Arch. Neurol. Psychiatry *79*, 475–497.

Plihal, W., and Born, J. (1997). Effects of Early and Late Nocturnal Sleep on Declarative and Procedural Memory. J. Cogn. Neurosci. *9*, 534–547.

Plihal, W., and Born, J. (1999). Effects of early and late nocturnal sleep on priming and spatial memory. Psychophysiology *36*, 571–582.

Porte, Y., Buhot, M.C., and Mons, N.E. (2008). Spatial memory in the Morris water maze and activation of cyclic AMP response element-binding (CREB) protein within the mouse hippocampus. Learn. Mem. *15*, 885–894.

Ramadan, W., Eschenko, O., and Sara, S.J. (2009). Hippocampal Sharp Wave/Ripples during Sleep for Consolidation of Associative Memory. PLOS ONE *4*, e6697.

Ramos, J.M. (1998). Retrograde amnesia for spatial information: a dissociation between intra and extramaze cues following hippocampus lesions in rats. Eur. J. Neurosci. *10*, 3295–3301.

Ranck, J.B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. Exp. Neurol. *41*, 461–531.

Rasch, B.H., Born, J., and Gais, S. (2006). Combined blockade of cholinergic receptors shifts the brain from stimulus encoding to memory consolidation. J. Cogn. Neurosci. *18*, 793–802.

Redondo, R.L., and Morris, R.G.M. (2011). Making memories last: the synaptic tagging and capture hypothesis. Nat. Rev. Neurosci. *12*, 17–30.

Rotenberg, E., Abel, T., Hawkins, R.D., K, E.R., and Muller, R.U. (2000). Parallel instabilities of long-term potentiation, place cells, and learning caused by decreased protein kinase A activity. J Neurosci *20*, 8096–8102.

Sakurai, Y. (1996). Hippocampal and neocortical cell assemblies encode memory processes for different types of stimuli in the rat. J. Neurosci. *16*, 2809–2819.

Sans, N., Prybylowski, K., Petralia, R.S., Chang, K., Wang, Y.-X., Racca, C., Vicini, S., and Wenthold, R.J. (2003). NMDA receptor trafficking through an interaction between PDZ proteins and the exocyst complex. Nat. Cell Biol. *5*, 520–530.

Scaplen, K.M., Gulati, A.A., Heimer-McGinn, V.L., and Burwell, R.D. (2014). When do objects become landmarks? Hippocampal place cells respond differently to manipulations of visual cues depending on size, perspective, and experience. Hippocampus *24*, 1287–1299.

Schafe, G.E., Nadel, N.V., Sullivan, G.M., Harris, A., and LeDoux, J.E. (1999). Memory Consolidation for Contextual and Auditory Fear Conditioning Is Dependent on Protein Synthesis, PKA, and MAP Kinase. Learn. Mem. *6*, 97–110.

Schiller, D., Eichenbaum, H., Buffalo, E.A., Davachi, L., Foster, D.J., Leutgeb, S., and Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. J. Neurosci. *35*, 13904–13911.

Schlesiger, M.I., Cressey, J.C., Boublil, B., Koenig, J., Melvin, N.R., Leutgeb, J.K., and Leutgeb, S. (2013). Hippocampal activation during the recall of remote spatial memories in radial maze tasks. Neurobiol. Learn. Mem. *106*, 324–333.

Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psychiatry *20*, 11–21.

Segal, M., Disterhoft, J.F., and Olds, J. (1972). Hippocampal Unit Activity during Classical Aversive and Appetitive Conditioning. Science *175*, 792–794.

Sharp, P.E. (1997). Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: comparison with hippocampal place cells. Behav. Brain Res. *85*, 71–92.

Shih, P.-C., Yang, Y.-R., and Wang, R.-Y. (2013). Effects of Exercise Intensity on Spatial Memory Performance and Hippocampal Synaptic Plasticity in Transient Brain Ischemic Rats. PLOS ONE *8*, e78163.

Siapas, A.G., and Wilson, M.A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. Neuron *21*, 1123–1128.

Sidman, M., Stoddard, L.T., and Mohr, J.P. (1968). Some additional quantitative observations of immediate memory in a patient with bilateral hippocampal lesions. Neuropsychologia *6*, 245–254.

Silva, D., Feng, T., and Foster, D.J. (2015). Trajectory events across hippocampal place cells require previous experience. Nat. Neurosci. *18*, 1772–1779.

Sirota, A., Csicsvari, J., Buhl, D., and Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. Proc. Natl. Acad. Sci. U. S. A. *100*, 2065–2069.

Skaggs, E.B. (1925). Further studies in retroactive inhibition. Psychol. Monogr. *34*, 1–60.

Skaggs, W.E., McNaughton, B.L., Gothard, K.M., and Markus, E.J. (1992). An Information-theoretic Approach to Deciphering the Hippocampal Code. In Proceedings of the 5th International Conference on Neural Information Processing Systems, (San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.), pp. 1030–1037.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus *6*, 149–172.

Squire, L.R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychol. Rev. *99*, 195–231.

Squire, L.R., Cohen, N.J., and Nadel, L. (2014). The Medial Temporal Region and Memory Consolidation: A New Hypothesis. Mem. Consol. Psychobiol. Cogn. 185.

Staresina, B.P., Alink, A., Kriegeskorte, N., and Henson, R.N. (2013). Awake reactivation predicts memory in humans. Proc. Natl. Acad. Sci. *110*, 21159–21164.

Stickgold, R. (2005). Sleep-dependent memory consolidation. Nature *437*, 1272–1278.

van Strien, N.M., Cappaert, N.L.M., and Witter, M.P. (2009). The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. Nat. Rev. Neurosci. 10, 272–282.

Suddendorf, T., and Corballis, M.C. (2007). The evolution of foresight: What is mental time travel, and is it unique to humans? Behav. Brain Sci. *30*, 299–313; discussion 313-351.

Sutherland, G.R., and McNaughton, B. (2000). Memory trace reactivation in hippocampal and neocortical neuronal ensembles. Curr. Opin. Neurobiol. *10*, 180–186.

Sutherland, R.J., Weisend, M.P., Mumby, D., Astur, R.S., Hanlon, F.M., Koerner, A., Thomas, M.J., Wu, Y., Moses, S.N., Cole, C., et al. (2001). Retrograde amnesia after hippocampal damage: recent vs. remote memories in two tasks. Hippocampus *11*, 27–42.

Szél, Á., and Röhlich, P. (1992). Two cone types of rat retina detected by anti-visual pigment antibodies. Exp. Eye Res. *55*, 47–52.

Takahashi, S. (2015). Episodic-like memory trace in awake replay of hippocampal place cell activity sequences. ELife *4*, e08105.

Takehara, K., Kawahara, S., Takatsuki, K., and Kirino, Y. (2002). Time-limited role of the hippocampus in the memory for trace eyeblink conditioning in mice. Brain Res. *951*, 183–190.

Tambini, A., Ketz, N., and Davachi, L. (2010). Enhanced brain correlations during rest are related to memory for recent experiences. Neuron *65*, 280–290.

Tamminen, J., Payne, J.D., Stickgold, R., Wamsley, E.J., and Gaskell, M.G. (2010). Sleep Spindle Activity is Associated with the Integration of New Memories and Existing Knowledge. J. Neurosci. *30*, 14356–14360.

Teixeira, C.M., Pomedli, S.R., Maei, H.R., Kee, N., and Frankland, P.W. (2006). Involvement of the anterior cingulate cortex in the expression of remote spatial memory. J. Neurosci. Off. J. Soc. Neurosci. 26, 7555–7564.

Teyler, T.J., and DiScenna, P. (1985). The role of hippocampus in memory: a hypothesis. Neurosci. Biobehav. Rev. *9*, 377–389.

Thompson, L.T., and Best, P.J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. J. Neurosci. Off. J. Soc. Neurosci. *9*, 2382–2390.

Thompson, L.T., and 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 Res. *509*, 299–308.

Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol. Rev. *55*, 189–208.

Tsao, A., Moser, M.-B., and Moser, E.I. (2013). Traces of Experience in the Lateral Entorhinal Cortex. Curr. Biol. *23*, 399–405.

Tse, D., Langston, R.F., Kakeyama, M., Bethus, I., Spooner, P.A., Wood, E.R., Witter, M.P., and Morris, R.G.M. (2007). Schemas and Memory Consolidation. Science *316*, 76–82.

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., Bito, H., and Morris, R.G.M. (2011). Schema-Dependent Gene Activation and Memory Encoding in Neocortex. Science *333*, 891–895.

Tsien, J.Z., Huerta, P.T., and Tonegawa, S. (1996). The Essential Role of Hippocampal CA1 NMDA Receptor–Dependent Synaptic Plasticity in Spatial Memory. Cell *87*, 1327–1338.

Tulving, E., and Markowitsch, H.J. (1998). Episodic and declarative memory: Role of the hippocampus. Hippocampus *8*, 198–204.

Vargha-Khadem, F., Gadian, D.G., Watkins, K.E., Connelly, A., Van Paesschen, W., and Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. Science *277*, 376–380.

Ven, G.M. van de, Trouche, S., McNamara, C.G., Allen, K., and Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. Neuron *92*, 968–974.

Veyrac, A., Allerborn, M., Gros, A., Michon, F., Raguet, L., Kenney, J., Godinot, F., Thevenet, M., Garcia, S., Messaoudi, B., et al. (2015). Memory of Occasional Events in Rats: Individual Episodic Memory Profiles, Flexibility, and Neural Substrate. J. Neurosci. *35*, 7575–7586.

Wallenstein, G.V., and Hasselmo, M.E. (1997). GABAergic Modulation of Hippocampal Population Activity: Sequence Learning, Place Field Development, and the Phase Precession Effect. J. Neurophysiol. *78*, 393–408.

Wamsley, E.J., Tucker, M.A., Payne, J.D., and Stickgold, R. (2010). A brief nap is beneficial for human route-learning: The role of navigation experience and EEG spectral power. Learn. Mem. *17*, 332–336.

Wang, X.-Q., and Wang, G.-W. (2016). Effects of treadmill exercise intensity on spatial working memory and long-term memory in rats. Life Sci. *149*, 96–103.

Weible, A.P., Rowland, D.C., Pang, R., and Kentros, C. (2009). Neural Correlates of Novel Object and Novel Location Recognition Behavior in the Mouse Anterior Cingulate Cortex. J. Neurophysiol. *102*, 2055–2068.

Weible, A.P., Rowland, D.C., Monaghan, C.K., Wolfgang, N.T., and Kentros, C.G. (2012). Neural Correlates of Long-Term Object Memory in the Mouse Anterior Cingulate Cortex. J. Neurosci. *32*, 5598–5608.

Wills, T.J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor Dynamics in the Hippocampal Representation of the Local Environment. Science *308*, 873–876.

Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. Science *261*, 1055–1058.

Wilson, M.A., and McNaughton, B.L. (1994). Reactivation of hippocampal ensemble memories during sleep. Science *265*, 676–679.

Winocur, G. (1990). Anterograde and retrograde amnesia in rats with dorsal hippocampal or dorsomedial thalamic lesions. Behav. Brain Res. *38*, 145–154.

Winocur, G., and Moscovitch, M. (2011). Memory transformation and systems consolidation. J. Int. Neuropsychol. Soc. JINS *17*, 766–780.

Witter, M.P., Doan, T.P., Jacobsen, B., Nilssen, E.S., and Ohara, S. (2017). Architecture of the Entorhinal Cortex A Review of Entorhinal Anatomy in Rodents with Some Comparative Notes. Front. Syst. Neurosci. 11.

Wood, E.R., Dudchenko, P.A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. Nature *397*, 613–616.

Wood, E.R., Dudchenko, P.A., Robitsek, R.J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron *27*, 623–633.

Yartsev, M.M., and Ulanovsky, N. (2013). Representation of Three-Dimensional Space in the Hippocampus of Flying Bats. Science *340*, 367–372.

Ylinen, A., Bragin, A., Nádasdy, Z., Jandó, G., Szabó, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. J. Neurosci. Off. J. Soc. Neurosci. *15*, 30–46.

Young, B.J., Fox, G.D., and Eichenbaum, H. (1994). Correlates of hippocampal complex-spike cell activity in rats performing a nonspatial radial maze task. J. Neurosci. *14*, 6553–6563.

Zheng, C., Bieri, K.W., Hwaun, E., and Colgin, L.L. (2016). Fast Gamma Rhythms in the Hippocampus Promote Encoding of Novel Object–Place Pairings. ENeuro *3*.

Zhou, W., and Crystal, J.D. (2011). Validation of a rodent model of episodic memory. Anim. Cogn. *14*, 325–340.