Reactivation content is important for consolidation of spatial memory

by

Igor Gridchyn

August, 2018

A thesis presented to the
Graduate School
of the
Institute of Science and Technology Austria, Klosterneuburg, Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

The dissertation of Igor Gridchyn, titled "Reactivation content is important for consolidation of spatial memory", is approved by:

Supervisor: Prof. Jozsef Csicsvari, IST Austria, Klosterneuburg, Austria

Signature: _____

Committee Member: Prof. Peter Jonas, IST Austria, Klosterneuburg, Austria

Signature: _____

Committee Member: Prof. Alessandro Treves, SISSA, Trieste, Italy

Signature: _____

Exam Chair: Prof. Maximilian Jösch, IST Austria, Klosterneuburg, Austria

Signature: _____

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

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

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

Signature: _____

Igor Gridchyn

August 27, 2018

Abstract

The hippocampus is a key brain region for spatial memory and navigation and is needed at all stages of memory, including encoding, consolidation, and recall. Hippocampal place cells selectively discharge at specific locations of the environment to form a cognitive map of the space. During the rest period and sleep following spatial navigation and/or learning, the waking activity of the place cells is reactivated within high synchrony events. This reactivation is thought to be important for memory consolidation and stabilization of the spatial representations. The aim of my thesis was to directly test whether the reactivation content encoded in firing patterns of place cells is important for consolidation of spatial memories. In particular, I aimed to test whether, in cases when multiple spatial memory traces are acquired during learning, the specific disruption of the reactivation of a subset of these memories leads to the selective disruption of the corresponding memory traces or through memory interference the other learned memories are disrupted as well.

In this thesis, using a modified cheeseboard paradigm and a closed-loop recording setup with feedback optogenetic stimulation, I examined how the disruption of the reactivation of specific spiking patterns affects consolidation of the corresponding memory traces. To obtain multiple distinctive memories, animals had to perform a spatial task in two distinct cheeseboard environments and the reactivation of spiking patterns associated with one of the environments (target) was disrupted after learning during four hours rest period using a real-time decoding method. This real-time decoding method was capable of selectively affecting the firing rates and cofiring correlations of the target environment-encoding cells. The selective disruption led to behavioural impairment in the memory tests after the rest periods in the target environment but not in the other undisrupted control environment. In addition, the map of the target environment was less stable in the impaired memory tests compared to the learning session before than the map of the control environment. However, when the animal relearned the task, the same map recurred in the target environment that was present during learning before the disruption.

Altogether my work demonstrated that the reactivation content is important: assembly-related disruption of reactivation can lead to a selective memory impairment and

deficiency in map stability. These findings indeed suggest that reactivated assembly patterns reflect processes associated with the consolidation of memory traces.

Acknowledgments

Many people have supported and assisted me throughout all the stages of my PhD and I would like to express my gratitude to them.

First of all, I would like to thank my PhD supervisor Prof. Jozsef Csicsvari for being a wonderful mentor. Jozsef's sharp critical and analytical thinking helped us to devise a proper experimental paradigm, as well as obtain plenty of valuable results from the data analyses. Futhermore, the generous funding provided by Jozsef for my project gave me valuable freedom in exploring different aspects of the project. Completion of my project would not be possible without all the time that we had spent discussing, troubleshooting and improving the project..

Secondly, I would like to thank Prof. Alessandro Treves, my external PhD committee member, and Prof. Peter Jonas, my internal PhD committee member, for their participation in the qualifying exam and in the PhD thesis defense process and for providing comments and suggestions on how to improve my PhD thesis despite their busy schedule.

I would also like to thank an outstanding scientist, Dr. Philipp Schönenberger, who was a postdoctoral researcher in the Csicsvari lab during the early years of my PhD. Given my previous lack of experimental experience and analysis of neural data, Philipp's guidance and attention to detail as well as his comprehensible and intelligible manner of teaching were crucial in setting me on the path to becoming an independent researcher in the many aspects of systems neuroscience.

All members of the Csicsvari lab, having an expertise in diverse fields, have in different ways helped me to improve various aspects of my PhD project, such as experimental design, microdrive assembly, surgery, behavioural training, electrophysiological recordings and data analysis. I want to thank all my colleagues from the Csicsvari Lab: Yosman BapatDhar, Peter Baracskay, Karel Blahna, Charlotte Boccara, Karola Käfer, Michael LoBianco, Alessia Manganaro, Michele Nardin, Joseph O'Neill, Damáris Ketinó Rangel Guerrero, Federico Stella, Jago Wallenschuss, and Haibing Xu.

And last but not least, I would like to thank my parents, Lena and Yuri, and my wife, Vera, for their immense support and an endless patience throughout my PhD.

1 INTRODUCTION 12

1.1 ANATOMICAL ORGANIZATION OF THE HIPPOCAMPUS 12
1.2 PLACE CODING IN THE HIPPOCAMPUS 14
1.2.1 Place cell and the cognitive map 14
1.3 PLACE CELL AND REMAPPING 15
1.4 PLACE CELLS DURING SPATIAL MEMORY TASKS 15
1.5 LOCAL FIELD POTENTIAL IN HIPPOCAMPUS 16
1.6 BEHAVIOURAL STUDIES DEMONSTRATING THE ROLE OF THE HIPPOCAMPUS IN SPATIAL MEMORY 17
1.7 REPLAY OF ACTIVE WAKING NEURONAL ACTIVITY PATTERNS IN REST 19
1.7.1 Hippocampal replay in cognitive functions 21
1.7.2 Disruption of SWR 21
1.8 CLOSED-LOOP SETUPS 23
1.9 RELEVANT STATISTICAL MODELS 24
1.10 AIM OF THE STUDY 26

2 METHODS 28

2.1 ANIMALS, MICRODRIVE CONSTRUCTION AND SURGICAL PROCEDURE 28
2.2 TETRODES ADJUSTING AND RECORDING 30
2.3 BEHAVIOURAL PARADIGM 31
2.4 TRAINING PROCEDURES 34
2.5 POSITION TRACKING SYSTEM AND SPIKE SORTING 35
2.6 CALCULATION OF RATE MAPS 36
2.7 REAL-TIME SIGNAL PROCESSING 36
2.7.1 Original decoding method and its optimizations 38
2.7.2 Closed loop delay 44
2.7.3 Decoding performance 46
2.7.4 Dissociation of two environment reactivations 49
2.8 STATISTICAL ANALYSES 50

3 RESULTS 51

3.1 INFLUENCE OF OPTOGENETIC LIGHT APPLICATION ON THE FIRING OF CA1 CELLS. 51
3.2 INFLUENCE OF ASSEMBLY-RELATED REACTIVATION DISRUPTION ON THE REACTIVATION PATTERNS 53
3.2.1 Firing rates 53
3.3 ASSESSMENT OF BEHAVIOURAL PERFORMANCE ON THE CHEESEBOARD MAZE 62
3.4 INFLUENCE OF OPTOGENETIC LIGHT APPLICATION MEDIATED RATE CHANGES ON BEHAVIOURAL PERFORMANCE 69
3.5 INFLUENCE OF THE REACTIVATION DISRUPTION ON PLACE MAP STABILITY. 70
3.6 INFLUENCE OF THE REACTIVATION DISRUPTION ON PLACE MAP GOAL REMAPPING 76

4 DISCUSSION 80

4.1 ROLE OF REACTIVATION IN MEMORY CONSOLIDATION 80
4.2 ADVANCEMENTS IN THE CLOSED-LOOP METHODOLOGY 85

5 Future work 88

6 BIBLIOGRAPHY 90

List of Abbreviations

1D	one-dimensional
2D	two-dimensional
AAV	adeno-associated virus
ALS	amyotrophic lateral sclerosis
ANOVA	analysis of variance
AP	anterior-posterior
CA1	Cornu Ammonis 1
HSE	high synchrony event
KDE	kernel density estimate
LIA	large amplitude irregular activity
LFP	local field potential
ML	medio-lateral
MUA	multiunit activity
PFS	place field similarity
RAM	random access memory
SD	standard deviation
SIMD	single instruction multiple data
MFB	medial forebrain bundle
OS	operating system
CPU	central processing unit
GB	gigabyte

1 INTRODUCTION

1.1 Anatomical organization of the hippocampus

Although the anatomical organization of the hippocampus has been extensively studied, there is still no full agreement about the nomenclature. Designations used in my thesis follow those of Amaral and Lavanex (2007). The hippocampal formation consists of the **dentate gyrus** (DG), the **hippocampus proper** (containing **CA1, CA2, CA3** subregions) and **subiculum, presubiculum, parasubiculum** and the **entorhinal cortex (EC)**. In rats, the hippocampal formation looks like a banana or resembles a seahorse, from the Latin name of which the structure borrowed its name. It emerges at the posterior end of the septal nuclei, curves over the thalamus and continues in the ventrolateral direction to end in the temporal lobe. This longitudinal axis of the hippocampal formation is referred to as the **septotemporal axis**.

The excitatory connections of the hippocampus are mostly unidirectional in contrast to the neocortex. The excitatory projections connecting different subregions of the hippocampus include those from the DG to CA3, called the **mossy fibres** and from CA3 to CA1 referred to as the **Schaffer collaterals**. These projections together with stellate cells projections from layer II of the EC via perforant path to the DG and CA3 form the **tri-synaptic path (Fig. 1.1)**. Projections from layer III of EC to CA1 are called **temporoammonic pathway**. There is a strong **recurrent connectivity** within the CA3 region, which suggested its role in **pattern completion**, the process by which an entire memory is retrieved from partial cues (McNaughton & Morris, 1987; Edmund T. Rolls, 2007, 2013). However, recent study has shown that the connectivity patterns within the CA3 region are not as dense as had been thought before (Guzman, Schlogl, Frotscher, & Jonas, 2016).

Most of the **sensory information** reaches the hippocampus through the EC, which also is the main relay for the information flow back from the hippocampus to the neocortical areas. Layer V of the EC is innervated by CA1 and subiculum outputs.

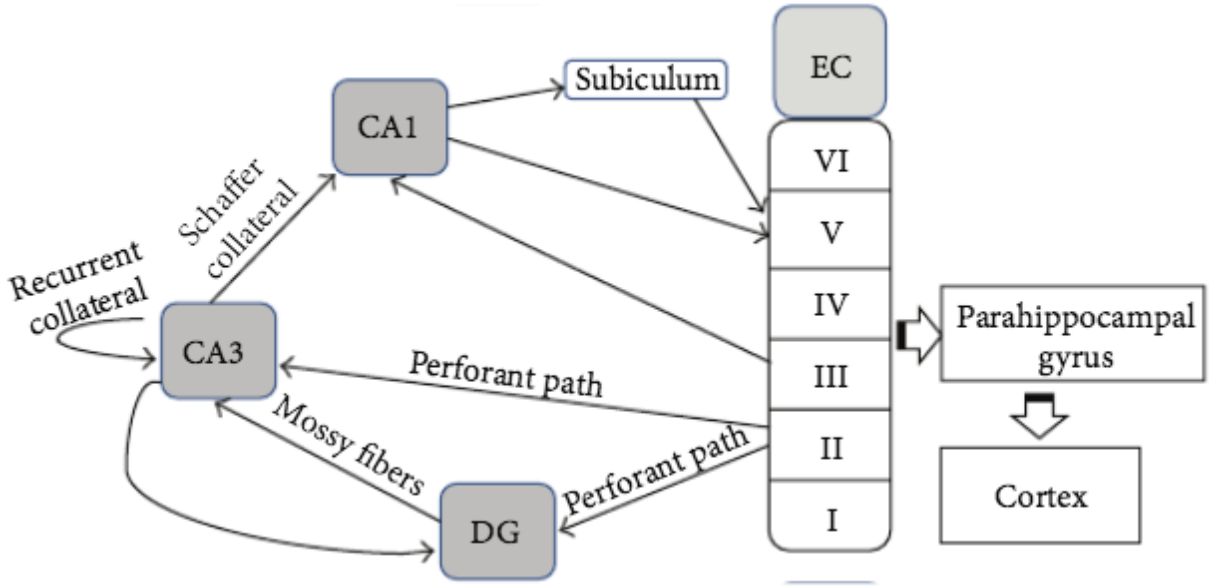

Fig. 1.1 | Schematic illustration of the anatomy of the hippocampal network

Figure adapted from (Yau, Li, & So, 2015)

Principal cell type of the hippocampus proper is the **pyramidal cell**. The soma of these cells is located in the pyramidal cell layer and their basal dendrite trees - in stratum oriens, while apical dendrite trees extend from the stratum radiatum to stratum lacunosum-moleculare. By contrast, in the dentate gyrus granule cells are the excitatory principal cells with their somas forming the stratum granulosum while their dendrites are located in in the stratum moleculare. Dentate gyrus also includes the hilus with excitatory mossy cells.

The **interneuron** populations of the hippocampus are diverse (Freund & Buzsaki, 1996) with at least 16 types of CA1 GABAergic interneurons listed in (Somogyi & Klausberger, 2005) and at least 21 types listed in (Klausberger & Somogyi, 2008), which vary in location of their cell bodies, dendritic trees, projections, protein expression profiles, target domains of pyramidal cells and preferred firing phase during oscillations.

1.2 Place coding in the hippocampus

1.2.1 Place cell and the cognitive map

The fundamental breakthrough in the hippocampal research happened due to the discovery that the behavioural correlate of hippocampal principal cells is related to space. These cells firing in relation to the location of the animal are called **place cells** (O'Keefe & Dostrovsky, 1971a). In the original study, activity of pyramidal units in the hippocampus was recorded while the animals traversed an open space environment. Eight of the 76 recorded units were reported to fire in the particular part of the test environment. Subsequently it was proposed that these spatially-selective cells were a substrate for the allocentric cognitive map of the environment, proposed earlier by Tollman (O'Keefe & Nadel, 1978; Tolman, 1948) and that the position of the rat can be decoded from the activity of these neurons given that sufficient number of them is recorded simultaneously. This hypothesis was later supported by an accurate reconstruction of the animal's position from the simultaneous recording from place cell firing (Wilson & McNaughton, 1993). Multiple studies that followed were able to identify hippocampal spatially selective neurons in non-rodent species, including birds, primates and humans (Ekstrom et al., 2003; Georges-Francois, Rolls, & Robertson, 1999; E.T. Rolls, 1999; E.T. Rolls & O'Mara, 1995; Siegel, Nitz, & Bingman, 2006). In addition to place cells, in other brain regions other spatially-selective cell types have been identified, such as border cells/boundary vector cells in the subiculum (Lever, Burton, Jeewajee, O'Keefe, & Burgess, 2009) and in the medial entorhinal cortex (MEC) (Solstad, Boccara, Kropff, Moser, & Moser, 2008), grid cells in the MEC and pre- and para-subiculum (Boccara et al., 2010; Hafting, Fyhn, Molden, Moser, & Moser, 2005) and head-direction cells in the anterodorsal thalamus, presubicumum and MEC (Taube, 1995, p. 19; Taube, Muller, & Ranck, 1990)

1.3 Place cell and remapping

Following the discovery of place cells, multiple studies went on to describe their properties. One of their important properties is stability. Place cells are stable over long periods of time (Muller, Kubie, & Ranck, 1987) if the environment is stable and no learning occurs. This stability suggests that the place cells are a proper substrate for spatial memory. This view has been questioned by imaging studies claiming that over days place representations may alter (Hainmueller & Bartos, 2018). However it is unclear whether stable recordings could be accomplished in these studies or imaging preparation itself may have led to instability. However, if any aspects of the environment, such as visual cues or locations of objects or food are changed to a significant degree, the firing patterns of the place cells may change (Muller & Kubie, 1987; Thompson & Best, 1989). This capacity is commonly referred to as **remapping**. However, moderate changes to the size of the environment can only lead to a rescaling of the place fields (O'Keefe & Burgess, 1996). Remapping with a change in location is often called 'global remapping' in contrast to 'rate remapping' in the case when only firing rate changes while the location of the place field is stable (S. Leutgeb et al., 2005). Rate remapping occurs when distinct sensory experiences, like sensory features of the enviroment, task-contingent demands and the representation of temporal delay (Eichenbaum, 2004; O'Keefe & Nadel, 1978; Pastalkova, Itskov, Amarasingham, & Buzsaki, 2008) have to be discriminated in the same environment and place cells encode spatial and non-spatial information.

1.4 Place cells during spatial memory tasks

Place cells can exhibit behavioural modulation of the firing rates during memory tasks. In spatial working memory tasks place cells remap their place fields to represent the goal locations. In particular, in the water maze task, the place fields are more frequently observed around the hidden platform (Hollup, Molden, Donnett, Moser, & Moser, 2001). Similarly, in the other tasks involving learning of goal-location, such as reward locations on

the cheeseboard maze, similar goal remapping can be observed constituting an increase in representation of the goal location in the place fields (Dupret, O'Neill, Pleydell-Bouverie, & Csicsvari, 2010). In another task, a conditioned place preference, in which the animal has to reach an unmarked goal location before a reward was received, a secondary place field emerged around the goal location (Hok et al., 2007).

1.5 Local field potential in Hippocampus

Local field potential (LFP) is an electrophysiological signal representing the summed electric current flowing from multiple nearby neurons within a small volume of nervous tissue.

Hippocampus exhibits a rich variety of characteristic oscillatory patterns (Berger, 1929) that vary depending on the behavioural and cognitive state of an animal. Hippocampal theta-band oscillations (5-10 Hz) occur during voluntary movements such as running, head-scanning and sniffing (Vanderwolf, 1969). Place cell firing and theta oscillations are related in several ways: place-related firing tends to occur during theta oscillations. Moreover, theta firing phase of place cells is related to the animal's position relative to the place field of the place cell, a process referred to as the **theta phase precession** (O'Keefe & Recce, 1993).

In the absence of theta oscillations, the large amplitude irregular activity (LIA), as it was named by Vanderwolf (1969), appears. One of the most prominent oscillatory patterns during the LIA are the sharp wave/ripples (SWR), which represent a transient (50-100 ms) fast (150-250 Hz) oscillatory events accompanied by high synchrony bursts of the hippocampal population (Buzsaki, 1986; Buzsaki, Leung, & Vanderwolf, 1983; Csicsvari, Hirase, Czurko, Mamiya, & Buzsaki, 1999b, 1999a). SWRs occur mostly during slow wave sleep, waking immobility, consummatory behaviour, grooming or in pauses in the exploratory behaviour (Buzsaki, Horvath, Urioste, Hetke, & Wise, 1992; Csicsvari et al., 1999a; Foster & Wilson, 2006; O'Neill, Senior, & Csicsvari, 2006). SWRs are thought to arise from the activity in the recurrently connected CA3 region, as the EC inputs are mostly silent during the initiation of the SWR (Chrobak & Buzsaki, 1994; Mizuseki, Sirota, Pastalkova, & Buzsáki, 2009; Sullivan et al., 2011). High recurrent connectivity of CA3 region promotes the

high synchrony bursts after a spontaneous stochastic simultaneous firing of a sufficient number of initiator cells (Schlingloff, Káli, Freund, Hájos, & Gulyás, 2014). Firing rate and synchrony increase in both CA3 and CA1 regions with higher but later emerging synchrony in the CA1 (Csicsvari, Hirase, Mamiya, & Buzsaki, 2000).

1.6 Behavioural studies demonstrating the role of the hippocampus in spatial memory

Multiple behavioural studies suggested hippocampal involvement in spatial memory function (Aggleton & Brown, 1999; Burgess, Maguire, & O'Keefe, 2002; Eichenbaum, 2004; Eichenbaum, Stewart, & Morris, 1990). Historically, the first important conclusions about brain function were made based on observations that were consequences of brain injuries or anomalies. For example, involvement of a specific brain area in speech processing was deduced by Broca in 1861 based on observation of patients with speech impairment that had focalized lesions in the brain. Similarly, the first conclusions about hippocampal functions were made based upon observations of a famous epilepsy patient H.M. (Scoville & Milner, 1957). H.M. had a substantial part of his medial temporal lobe bilaterally removed in an attempt to stop his epileptic seizures. After surgery the seizures stopped. However, the patient developed anterograde amnesia: he was not able to form new memories about events occurring in his personal life, as well as transient retrograde amnesia, absence of the memory of events that happened shortly before the surgery. However, his ability to form other types of memories, for example, a procedural memory while learning a motor task or working memory and language abilities, were preserved. Later it was discovered that Alzheimer's disease-related neurodegeneration, which to a large extent affects the hippocampal formation, causes similar memory deficits (Braak, Braak, Yilmazer, & Bohl, 1996).

In general, inducing structural or functional damage to a specific brain area allows studying the role of these brain areas given a proper behavioural and/or physiological readout. Obviously, deliberate brain damage to humans may not be justified for research purposes. However, with the development of experimental neuroscience it became apparent that **rodent** models may serve as an adequate model for studying multiple brain areas, including

hippocampus (Harper, 2010). Rodents can perform multiple spatial navigation and learning tasks (Robert J. Sutherland & Hamilton, 2004). In addition, mice offer a wide variety of genetically modified lines, allowing to target specific cell types in specific brain areas (Harper, 2010), while rats can perform more complicated behaviour as well as allow recording of larger number of units simultaneously. Rodents will easily engage in spatial navigation and learning tasks to look for food, when they are sufficiently food-deprived (Robert J. Sutherland & Hamilton, 2004).

First studies that linked hippocampal function to behavioural performance in rodents induced hippocampal lesions and observing animal behaviour in spatial navigation tasks. One of the prominent tasks is **the Morris water maze** developed by Morris et al. (1982). In this task, the animal has to find a hidden escape platform, which is located below the water surface level in a cylindrical tank filled with the murky water. The task relies on the rodents' natural aversion to swimming in the water. Animals start the search from the edge of the tank and, after finding the hidden platform, are allowed to stay on the platform for a limited time. In the test phase of the task the platform is removed, and the animal is allowed to swim in the tank for a limited time. In this case, the time that animal spends swimming around the former location of the platform indicates how well it remembers the location. It was shown that animals with lesioned hippocampus are impaired in the learning trials and subsequent probe performance (Clark, Broadbent, & Squire, 2005; Logue, Paylor, & Wehner, 1997; R. G. Morris et al., 1982; R.J. Sutherland, Whishaw, & Kolb, 1983). Dorsal hippocampal lesions have more pronounced effects that ventral hippocampal lesions (E. Moser, Moser, & Andersen, 1993; M. B. Moser, Moser, Forrest, Andersen, & Morris, 1995). Impairment is also achieved through disruption of inputs by perforant pathway or entorhinal cortex. There even exist human analogs of this task to assess memory in traumatic brain injury patients (Astur, Taylor, Mamelak, Philpott, & Sutherland, 2002; Goodrich-Hunsaker & Hopkins, 2010, p. 2; Skelton, Ross, Nerad, & Livingstone, 2006, p. 200).

Experiments in the water do not work well in combination with electrophysiological recordings. Therefore, the 'dry' version of the Morris water maze task was developed to be used along with electrophysiological recording techniques (Barnes, 1979, p. 197; Gilbert, Kesner, & DeCoteau, 1998; Kesner, Farnsworth, & DiMattia, 1989). One of the dry versions

of Morris water maze is the **cheeseboard task**, in which animal has to locate a food hidden in one of the multiple holes in the maze platform. Some studies have adopted the original regular square design by Kesner (Pillay, Kellaway, & Kotwal, 2008, p. 20) while others utilized the random pattern (Yoshida, Goto, & Watanabe, 2001). Example of cheeseboard layout and behavioural tracking is shown in **Fig. 1.2.** Similarly to the original Morris water maze task, in the probe session the animal has to search for food on the maze, but the food location is not baited, and the time spent around the goal location indicates the strength of the memory about the food location. A possible disadvantage compared to the original task is that over multiple days of training animals may become habituated to the absence of the food in the probe sessions and stop moving, unlike to the water maze where the animal is forced to continuously swim. To ensure that the animal has to rely on spatial navigation to locate the food the smell of the food is masked by spreading food dust with the same smell all over the board, and board is rotated between trials to avoid usage of local cues. The cheeseboard maze is less stressful than the Morris water maze or the Barnes maze. Hippocampal-lesioned animals are impaired in the probe performance in the cheeseboard task (Kesner et al., 1989). In the control version of the cheeseboard maze, visual cues are added to the goal locations (Dupret et al., 2010).

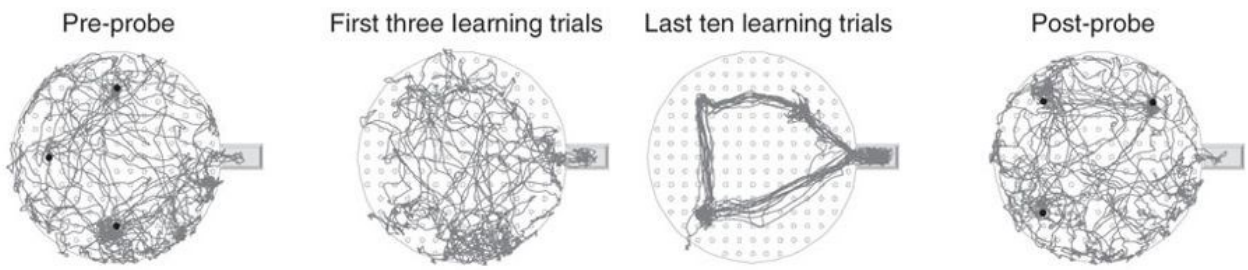

Fig. 1.2 | Cheeseboard layout and behavioural tracking example

An example of the layout of a circular cheeseboard environment with food wells located on a grid. Behavioural tracking of a control animal is shown for three stages of the cheeseboard paradigm: pre-probe, in which the memory for the old goal locations is tested, beginning and end of learning session, in which the emergence of stereotypical trajectories is seen, and the post-probe in which the memory for the new foals is tested. Figure adapted from (Dupret et al., 2010)

1.7 Replay of active waking neuronal activity patterns in rest

Soon after the discovery of the place cells it was observed that place cells fire in sleep as well, however, their firing is no longer correlated with the location of the animal (Wilson & McNaughton, 1994). However, the sleep firing patterns of place cells were discovered to be closely related to the previous spatial firing in active waking periods. This phenomenon was first demonstrated by Wilson and McNaughton(1994) by showing that the high cofiring correlations of place cells with overlapping place fields persisted in sleep after the running sessions, while in a pre-sleep session prior to exploration these cofiring correlations were low. Subsequently it was shown that the firing sequences of the place cells generated by an animal moving in the stereotypical trajectories (linear tracks, circular mazes) are reactivated during the subsequent sleep (Louie & Wilson, 2001; Skaggs & McNaughton, 1996). This latter finding suggests that frequently repeated movement trajectories of the animal are also reactivated in addition to behaviourally relevant places (O'Neill, Pleydell-Bouverie, Dupret, & Csicsvari, 2010). Reactivated patterns tend to occur during SWRs patterns both when expressing single places or trajectories. The reactivation phenomenon is illustrated in the **Fig. 1.3**.

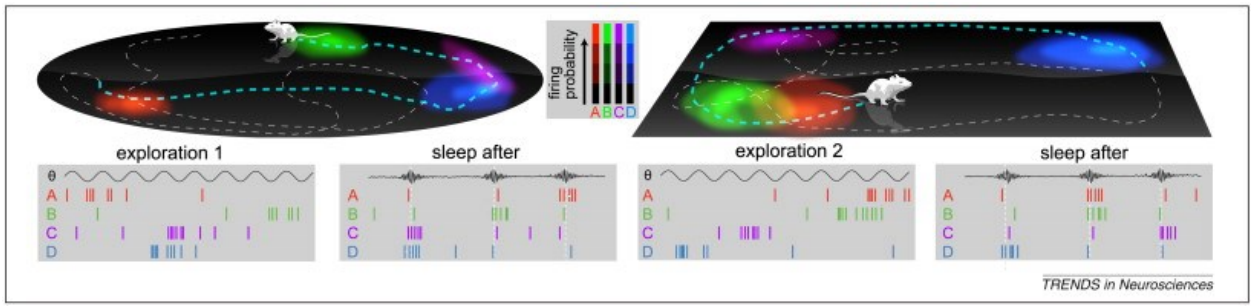

Fig. 1.3 | Reactivation of waking patterns of place cell firing

Schematic illustration of the space-related spiking patterns during waking are their reactivation during the sleep. Top panels illustrated locations of place fields of several place cells in two different environments. Bottom panels illustrated raster plots of firing of these cells during exploration of the environments and in the following rest periods. Cells that fire together during exploration, also fire together in sleep, but different pairs of cells may fire together in different environments as well as in the rest periods following exploration of these environments. Figure adapted from (O'Neill, Senior, Allen, Huxter, & Csicsvari, 2008).

Awake SWRs occurring during exploratory periods exhibit similar oscillatory and synchrony properties (O'Neill et al., 2006) to those observed during sleep. However, they might play a different role. Replay during these SWRs represents the current location or behavioural trajectories originating from the current location of the animal (Foster & Wilson, 2006; O'Neill et al., 2006). These waking SWRs have been linked to memory recall or decision

making (Jadhav, Kemere, German, & Frank, 2012a) and replay during these SWRs can code for the future trajectory of the animal in the goal-related behaviour (Pfeiffer & Foster, 2013).

1.7.1 Hippocampal replay in cognitive functions

This repeated patterns of activity during the SWRs within the short time window has been suggested to promote the potentiation of synapses and might strengthen associations within the hippocampus (Buzsaki, 1986). They are thought to facilitate the transfer of reactivated memory traces to other cortical areas (Buzsaki, 1989; McClelland, McNaughton, & O'Reilly, 1995; Ólafsdóttir, Carpenter, & Barry, 2016). Indeed, hippocampal SWR activity recruits EC (Chrobak & Buzsaki, 1996) and other brain areas (Logothetis et al., 2012). However, recent studies suggested that hippocampus is not the sole initiator of the replay of spatial memories, as medial entorhinal cell assemblies can replay previously visited trajectories of the animal independently of the hippocampus (O'Neill, Boccara, Stella, Schoenenberger, & Csicsvari, 2017).

A correlation relationship between replay and behavioural performance was shown by Dupret et al. (2010). In this study, amount of reactivation of the goal location in the sleep following goal learning on the cheeseboard predicted the performance of the animal in the post-probe session after the sleep.

1.7.2 Disruption of SWR

To demonstrate the importance of SWR-associated replay in cognitive functions, multiple experiments with selective blockade of SWR have been performed (Ego-Stengel & Wilson, 2010; Girardeau, Benchenane, Wiener, Buzsaki, & Zugaro, 2009; Jadhav et al., 2012a; van de Ven, Trouche, McNamara, Allen, & Dupret, 2016). In such studies, a disruption of activity outside of the SWR events was normally applied as a control manipulation.

In the work of Girardeau et al. (2009) animals were trained to learn to locate food in a radial eight-arms maze task. In the subsequent sleep, SWR events were selectively disrupted using electrical stimulation. This procedure was performed across multiple days during which a behavioural impairment was observed in the form of a delayed (8 days instead of 5 in control) emergence of higher than chance performance on the maze as well as lower than control performance starting from day nine. In another study, the animal learned to perform two identical spatial navigation tasks on the same day with sleep sessions following learning of each of the two tasks (Ego-Stengel & Wilson, 2010). During one of these sleep sessions, SWR disruption was performed, As a result, a behaviour was impaired in the task, after which the SWRs were disrupted, but not in the control task.

Role of awake SWRs was also studied by disrupting them during learning of a W-track task where the animal has to sequentially visit each arm of the maze in both directions (Jadhav et al., 2012a). Similarly to the earlier studies, a behavioural impairment was achieved by this manipulation already on the second day and was maintained until the end of the experiment on the day 8. However, in this case only working but not a reference memory impairment was seen, suggesting a role of these SWRs in working memory and spatial decision making. SWR disruption was also performed in rabbits during the inter-trial intervals of the trace eyeblink conditioning task and similarly led to decreased learning performance (Nokia, Mikkonen, Penttonen, & Wikgren, 2012).

Physiological deficits as a consequence of SWR disruption were reported in relation to the stability of place fields (van de Ven et al., 2016). Disruption of SWR activity after exploration of the novel environments prevented the subsequent reinstatement of hippocampal assembly patterns of these environments when the animal was placed back to the same environment, while no such effect was observed for familiar environments. This indicates that the role of SWR is time-limited and once a place field stabilised, it cannot be disrupted during the relay of their patterns. Moreover, this study used optogenetics to disrupt the activity patterns, which excludes the unspecific effect of electrical stimulation unrelated to SWR. Note however that another study (Kovács et al., 2016) observed stable place field of a novel environment when the optogenetic disruption of SWRs was applied subsequently to the first exposure of the novel environment. In this case place fields were equally stable

when SWRs were disrupted or a control sleep disruption was applied when the optogenetic signal was time-shifted relative to the SWR time. It is possible the SWR independent optogenetic inhibition may cause some degree of reduced place field stability which can lead to altered assemblies observed in the (van de Ven et al., 2016) study.

1.8 Closed-loop setups

Closed loop feedback-based experimental setups allow performing manipulations of brain activity conditioned on either the animal's behaviour (location, speed, etc.) or brain activity itself (population events, oscillatory events, individual spikes). Such setups required a real-time signal processing software capable of, depending on the setup type, real-time bandpass filtering, spike detection, classification, etc. Another characteristic part of the closed-loop setup is its feedback component, such as electrical stimulation device, laser setup, light or sound controllers. These components are modulated by the real-time processing software in response to the relevant events.

A **brain-computer interface (BCI)** is a connection between some kind of brain signal (EEG, LFP, spikes) and an external device (manipulator, screen, etc.). Among important medical application of BCI are antiepileptic stimulations (Gluckman, Nguyen, Weinstein, & Schiff, 2001; Psatta, 1983), non-muscular (e.g. by brain signal only) bidirectional (ability to express pre-configured sentences, alarms and binary responses) communication and control (e.g. of robotic devices) in amyotrophic lateral sclerosis (ALS) patients (Escolano, Ramos Murguialday, Matuz, Birbaumer, & Minguez, 2010; Lazarou, Nikolopoulos, Petrantonakis, Kompatsiaris, & Tsolaki, 2018). According to the given description of closed-loop setups, BCIs can be a component of the closed-loop experimental setup and as such can be used to probe hippocampal function.

As discussed above, SWR disruption (Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Jadhav et al., 2012a) provided an example of how closed-loop stimulation was applied in the hippocampus. Not only a specific type of an oscillatory event, but a specific phase of an oscillation can be targeted by a closed-loop feedback, as was the case in the study of Siegle

and Wilson (2014), in which behavioural performance in the encoding segment of the maze was enhanced by stimulation triggered by the peak of theta oscillation, while behavioural performance in the retrieval segment of the maze was enhanced by stimulation triggered by the trough of the theta oscillation. Targeting single cells spikes as the trigger was performed in Lavilleon et al. (2015), where stimulation of the medial forebrain bundle (MFB) triggered by a single cell with large amplitude induced a place preference in the area, where the trigger cell had place field.

1.9 Relevant statistical models

A fundamental task in neuroscience is to understand how neural ensembles represent information and population decoding is a useful method to extract information from neuronal populations. Not only such modelling can shed a light on the neural code, it can also enable the design of sophisticated experimental setups (Deng, Liu, Karlsson, Frank, & Eden, 2016; Grosenick, Marshel, & Deisseroth, 2015) as well as lead to the development of neural prosthetics and BCIs to restore motor function in patients with neurological damage (Chapin, 2004; Hochberg et al., 2012; Schwartz, Cui, Weber, & Moran, 2006). Out of many models that have been developed to describe population activity (Brown, Frank, Tang, Quirk, & Wilson, 1998; Kloosterman, Layton, Chen, & Wilson, 2014; Paninski, Pillow, & Lewi, 2007; Sanger, 2003; Truccolo, Eden, Fellows, Donoghue, & Brown, 2005), I will outline only those which are relevant for my thesis.

One of the earliest models attempted to describe the spiking statistics of a single neuron. Many of models of population activity rely on the assumption of Poisson spiking of hippocampal neurons given a fixed stimulus. It means that for a given fixed time window and stimulus (e.g. location for place cells), number of spikes emitted by a place cell under a given stimulus will follow a Poisson distribution:

$$P\ (n \text{ spikes during } \Delta t) = e^{-r\Delta t}\frac{(r\Delta t)^n}{n!} \quad (1)$$

where r is the mean firing rate of a neuron and Δt is a time interval.

Next important class of statistical models that have found a wide application in the encoding/decoding models for hippocampal data are the Bayesian decoders. In the context of spiking and stimulus, Bayes rule allows to calculate the conditional distribution of stimulus given activity of a neuron based on conditional distributions of activity of a neuron given stimulus and prior distribution of stimulus. For example, if $\boldsymbol{x}$ is a position of an animal and $\boldsymbol{n}$ is a number of spikes emitted by a neuron, then the conditional distribution of the position given the number of spikes is given by:

$$\boldsymbol{P(x|n)} = \frac{\boldsymbol{P(n|x)P(x)}}{\boldsymbol{P(n)}} \quad (2)$$

where $P(x)$ is a prior position distribution and $P(n)$ is a probability to observe n spikes at any position.

An important assumption used in Bayesian models for population vectors is an assumption of independence of spiking of individual cells, i.e. that given M neurons probability to observe these neurons emitting n_i spikes for $i = 1 \dots M$ in a given time window equals to the product of probabilities of independently observing every neuron emitting n_i spikes in this time window:

$$P(n_1, n_2 \dots n_M|x) = \prod_{i=1}^{M} P(n_i|x) \quad (3)$$

where M is number of neurons, n_i is the number of spikes emitted by i-th neuron, and x is a stimulus.

This assumption is not entirely true for the place cells, however, accounting for pairwise relations between spiking probabilities makes the models more complicated, computationally unfeasible for many real-time applications while providing little gain in coding accuracy.

Finally, I would like to point to two classes of models based on their input: first class of models describe relationship between sorted spikes / population vectors and stimulus (Brown et al., 1998; Paninski et al., 2007; Sanger, 2003; Truccolo et al., 2005), while the second class of models describes relationship between unsorted spikes/wave shapes and stimulus (Kloosterman et al. 2014). The main difference between the two lies in the fact that additional operations on the raw stimulus may lead to the loss of information. In the case of

spike sorting, this loss takes the form of discarded spikes due to inability to assign them to one of the defined putative units, as well as unavoidable errors in the spike sorting.

1.10 Aim of the study

Multiple studies have demonstrated the role of the hippocampus in navigation and in all stages of spatial memory formation including encoding, consolidation and recall (Aggleton & Brown, 1999; Burgess et al., 2002; Eichenbaum, 2004; R. G. M. Morris, Schenk, Tweedie, & Jarrard, 1990). Hippocampal place cells and the cognitive map formed by them are involved both in navigation and spatial memory (O'Keefe & Nadel, 1978; Riedel et al., 1999; Squire, 1992). It has been suggested that the hippocampus is involved in memory consolidation through reactivating waking memory traces during SWRs that occur in inactive rest periods (Buzsaki, 1989; McClelland et al., 1995; Skaggs & McNaughton, 1996). Recent work supports this hypothesis relating to the role of SWRs in memory consolidation. It was shown that SWRs during rest periods following learning are required for the successful recall of both spatial memories and the reinstatement of previously established novel spatial maps (Ego-Stengel & Wilson, 2010, p.; Girardeau et al., 2009; van de Ven et al., 2016). These suggest that SWR and SWR-associated reactivation are involved in the consolidation of spatial memories and the stabilisation of newly formed spatial maps. However, these works could not differentiate whether reactivation itself or other factors, such as highly synchronised activity during SWRs, are important for stabilization of spatial memories and maps. Another work showed that the reactivation content predicts the subsequent memory recall performance of the animals learning goal locations: the more frequently a given goal location was reactivated during SWRs, the better the animal remembered that location (Dupret et al., 2010). This provides a correlative evidence for the role of reactivation in the consolidation of goal-related spatial memories. Direct evidence for the role of reactivation in the consolidation of spatial memories still lacks however. Accordingly, the aim of my thesis was to directly test whether the reactivation content encoded in SWR firing patterns is important for consolidation. In particular, I aimed to test whether, in cases when multiple spatial memory traces are acquired during learning, the specific disruption of the reactivation of a subset of these memories leads to the selective disruption of the

corresponding memory traces or through memory interference of the other learned memories are disrupted as well.

To meet my research aim, I introduced a modified version of the cheeseboard maze task, in which animal had to perform the spatial goal learning task in two cheeseboard environments in parallel. I have also designed and developed algorithms and software for the efficient real-time decoding and the closed-loop disruption of neural patterns based on recently devised cluster-less decoding method to address this question. In these experiments, I performed multichannel extracellular recordings to record simultaneous activity of many single units and decoded in real-time this activity in order to optogenetically disrupt high synchrony events (HSE) associated with a specific spatial map representing one the cheeseboards. With these manipulations I aimed at selectively disrupting the recall of the spatial memory related to one of the cheeseboards and disrupt the stability of spatial representation of this cheeseboard. The specific aims were as follows:

(1) Develop a method that is capable of disrupting HSEs with specific reactivation content related to hippocampal spatial maps and verify the disruption efficiency of the method.

(2) Test whether the content of the spiking patterns within HSEs is important for the consolidation of the corresponding memories. More specifically, determine whether the selective optogenetic disruption of HSEs according to their reactivated content can cause specific memory impairment in recalling the corresponding memories.

(3) Test whether the maps-specific HSE disruption leads to the selective destabilisation of maps whose reactivation was disrupted.

2 METHODS

2.1 Animals, microdrive construction and surgical procedure

Seven adult male rats (Long Evans) were used as subjects for experiments. They were housed with an *ad libitum* access to water at all times in groups of three and individually after the virus injection surgery. A 12-hour light / 12-hour dark cycle was maintained with all recording, training and signal tuning conducted during the light phase. The rats weighed 250-350g at the time of surgery. All procedures involving experimental animals were carried out in accordance with Austrian federal law (1974) under a project license approved by the Federal Science Ministry (license number BMWFW-66.018/0015-WF/V3b/2014, permitted from 31.08.2014 till 31.08.2019).

Three animals were used to obtain the data for the hyperparameter tuning, algorithm testing and validation, and testing and debugging the real-time decoding method. With the remaining four animals the full experiments were performed and the data was used for the final analyses. Microdrives holding twenty four individually movable tetrodes and four individually movable octrodes (i.e. 8-wire electrodes) were used for neural recording. The tetrodes and octrodes were constructed from four and eight individual tungsten wires correspondingly (H-Formvar insulation with Butyral bond coat, California Fine Wire, Grover Beach, CA, USA) 12 µm in diameter in tetrodes and 10 µm diameter in octrodes, twisted then heated in order to bind them into a single bundle. Prior to the implantation, tetrode tips were plated with gold to reduce their impedance to 200-300 kΩ. Tetrodes and octrodes were individually inserted into the cannuli, which were individually mounted to shuttles. The Shuttles were mounted to individual screws, which allowed the tetrodes to be independently lowered by turning the screws (250 µm per full screw turn). Four 200 mm per 0.48 NA optic fibre stubs (Doric Lenses) located in the centre of the tetrode array were used to apply laser light directly to the dorsal CA1 area, two in each hemisphere.

I expressed Archaerhodopsin (ArchT) under the control of the CaMKII promoter in the CA1 region of the hippocampus using an adeno-associated virus (AAV) obtained from the Penn

Vector Core facility. An example of histology illustrating the expression of the virus is shown in **Fig. 2.1**

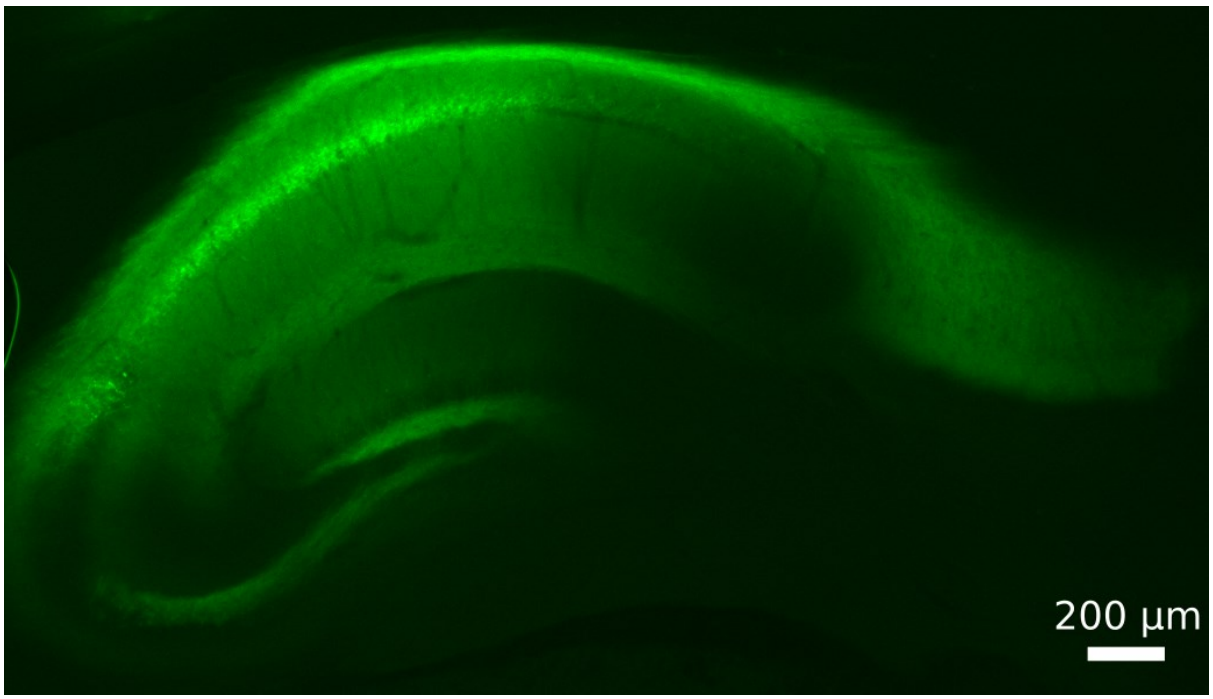

Fig. 2.1 | Example of virus expression

ArchT-YFP expression in the dorsal CA1 area of a virus-injected rat. Sagittal section.

I have chosen the ArchT for my experiments because of its fast kinetics of photocurrent rise and fall (15–85% photocurrent rise time of 7.4 ± 5.5 ms, mean ± SD; 85–15% fall time 18.4 ± 9.0 ms) as well as its fast recovery kinetics (85–15% recovery time 45.2 ± 34.3 ms) (Han et al., 2011).

Virus injection and microdrive implantation surgeries were conducted separately because virus expression took about three weeks. During both surgeries, all rats were anaesthetized with isoflurane (0.5%-3%), oxygen (1-2 l/min), and an initial dose of buprenorphine (0.1 mg/kg). The animals were then placed on a stereotaxic frame with the head levelled between bregma and lambda. During the injection surgery, two narrow diagonal craniotomies 3 mm long were centred at the dorsal hippocampal region CA1 (3.5 mm AP and 2.5 mm ML) after which virus was injected at four sites into the dorsal CA1 region bilaterally with the following coordinates: site 1: 3.2 AP, ±2.0 ML; site 2: 3.87 AP, ±2.67 ML; site 3: 4.53 AP, ±3.33 ML; site 4: 5.2 AP, ±4.0 ML. Before the injection, the virus solution was

diluted 1:24 with physiological NaCl solution. The virus was loaded into a calibrated glass capillary with a tip pulled to an inner diameter of ~12 µm. At each injection site, after capillary insertion and a wait period of 3 min., 300 nl diluted virus solution was pressure injected over 4 min. following a wait period of 3 min. After the injections, the craniotomies were covered with bone wax, and the skin wound was stitched. During the implantation surgery, the craniotomy was centred at the dorsal hippocampus (4.3 mm AP and 3.4 mm ML), and the electrodes were implanted into the deep layer of the neocortex at the depth of 1.4 mm. The implantation coordinates of every tetrode, optic fibres and craniotomy edges are depicted in the **Fig 2.2**. Two stainless-steel screws were positioned above the cerebellum to serve as a ground. Six additional anchor screws were used in order to permanently attach the microdrive assembly to the skull. The electrodes and the microdrive apparatus were paraffin wax coated and daubed with dental acrylic to encase the electrode-microdrive assembly and anchor it to the screws in the skull.

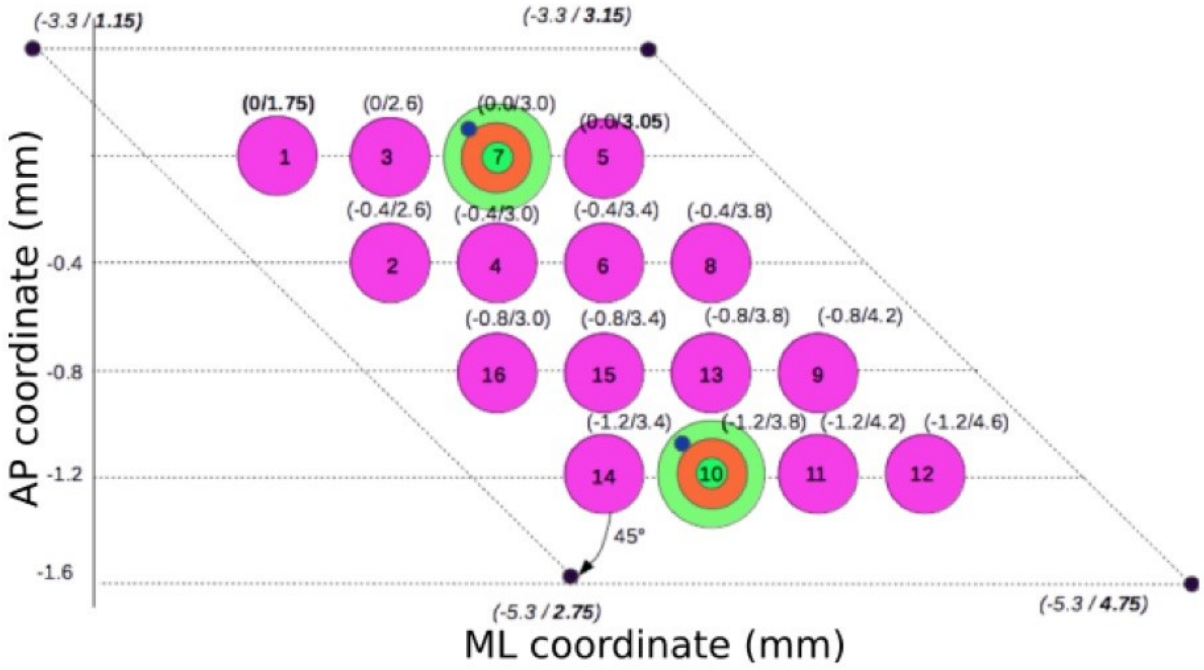

Fig. 2.2 | Coordinates of the electrode array, craniotomy and optic fibres.

The dotted parallelogram designates the craniotomy, the purple circles show the tetrodes locations, while the green circles the optic fibre locations. Coordinates (AP, ML) of the bundle centres: (-4.3, ±3.4).

2.2 Tetrodes adjusting and recording

Wide-band (0.4 Hz-9 kHz) recordings were taken, and the amplified local field potential and multiple-unit activity were continuously digitized at 24 kHz using 128-channel data acquisition system (Axona Ltd., St. Albans, UK). During the selective high synchrony disruption sessions, the signal was then processed in real-time by the custom software, which performed spike detection, feature extraction, decoding of the signal, and the trigger of the light pulse. Two red LEDs mounted on the preamplifier headstage were used to track the location of the animal. In order to simultaneously move 24 tetrodes and 4 octrodes into the pyramidal layer of CA1, tetrodes were adjusted over two weeks while using local field potential signatures as guidance. Adjusting was complicated by the initial instability of electrode locations caused by a large number of tetrodes pushing the brain tissue and by swelling, which can be present for several days following the surgery. During the surgery, the tetrodes were lowered into the deep layers of the neocortex - area, characterized by spiking activity that is less dense than in the hippocampal pyramidal cell layer and shows spindles and delta waves during sleep. After tetrodes pass through the neocortex, they enter the corpus callosum where no spiking activity is observed. When tetrodes reach the upper layers of the CA1 region, sharp waves can begin to be seen as upwards deflections in the LFP in the stratum oriens and then reverse near the cell body layer. In addition in the stratum oriens, ripple oscillations (150-250Hz) appear overlaid on sharp waves. While sharp waves increase in the magnitude and then decrease near the pyramidal layer as ripple amplitude increases and reach their maximum amplitude as action potentials begin to appear on the signal. At this stage, tuning was based on the cell yield with movement limited to ¼ of turn per tetrode per day to avoid delayed movement accumulation. Experiments start when all or a sufficient number of tetrodes and octrodes (at least 16 out of 28) show signs of optimal amplitude unit activity. It is often necessary to perform small adjustments at the end of experiment days due to continued movement of tetrodes. Such fine adjustments were also performed at the beginning of the recording day to maximize the yield after which a delay of one hour was made before the recording started.

2.3 Behavioural paradigm

The cheeseboard apparatus is used to test hippocampal spatial memory function (Kesner et al., 1989). Our lab previously used the cheeseboard to test goal remapping of CA1 place cells (Dupret, O'Neill, & Csicsvari, 2013; Dupret et al., 2010) and medial entorhinal grid cells (Boccara et al., in preparation) as well as long-term dynamics of reactivation after fast learning in the CA1 region (Baracskai et al., in preparation). To be able to test selective disruption of memory consolidation, I expanded this paradigm to two distinct cheeseboard environments with different external visual cues and start boxes at different relative locations to the cheeseboards. The cheeseboards I used are 120 cm in diameter and have semi-spherical holes with a radius of 1.25 cm, whose centres are located in the square grid with a spacing of 8 cm. The cheeseboards were covered with pellet dust to hide the smell of a pellet at the goal location. Between sessions, the cheeseboards were rotated randomly and independently at either 90, 180 or 270 degrees angle to avoid animals using local cues to ensure the hippocampal dependence of the task.

A typical cheeseboard task involves two dimensional trajectories and, based on the preliminary data I have collected from the paradigm with 3 goals in each of the two environments, reliably memorizing multiple (3-4) goal locations in two environments would require about 15 trials to reach an asymptotic behaviour, which conflicts with my goal to prevent the memory consolidation and start disruption as soon as possible. Therefore I have used a single goal in every cheeseboard to speed up the learning process. In addition, after learning animals tend to take simple linear trajectories from the start box to the goal location which helped to simplify the encoding model and improve the decoding performance In addition, I have added 20 cm 'bridge' between start box and cheeseboard to have slightly longer trajectories so that animals would develop higher speed before reaching the goal and cover larger proportion of the environment, thus generating place cell population activity that is easier to decode. To meet my research goals, I have performed experiments in a paradigm, in which animals had to learn a single goal location on both cheeseboards, i.e. they had to reach the goal location after leaving the start box and return to the start box immediately after consuming the goal. The goal locations of the two cheeseboards were different and changed daily. In each day, the updated goal locations were learned in interleaved learning sessions in the two environments in batches of five trials to reduce the recency effect. After the learning sessions, while the animal was

sleeping, my real-time decoding system (see **Section 2.7**) was detecting and disrupting reactivation of one of the environments which I will call **'target'** environment, whereas reactivation of the other environment, which I will call **'control'** environment was allowed to proceed. Which of the environments served as a control and which served as a target was decided randomly before the beginning of the experiment, as was the order in which learning sessions were performed in the two environments. The detailed experimental paradigm described below is illustrated in the **Fig. 2.3**. First, a 20 min. sleeping session with regular 1 Hz, 100 ms laser pulses was recorded to quantify the cell light responses for the post-hoc analysis.

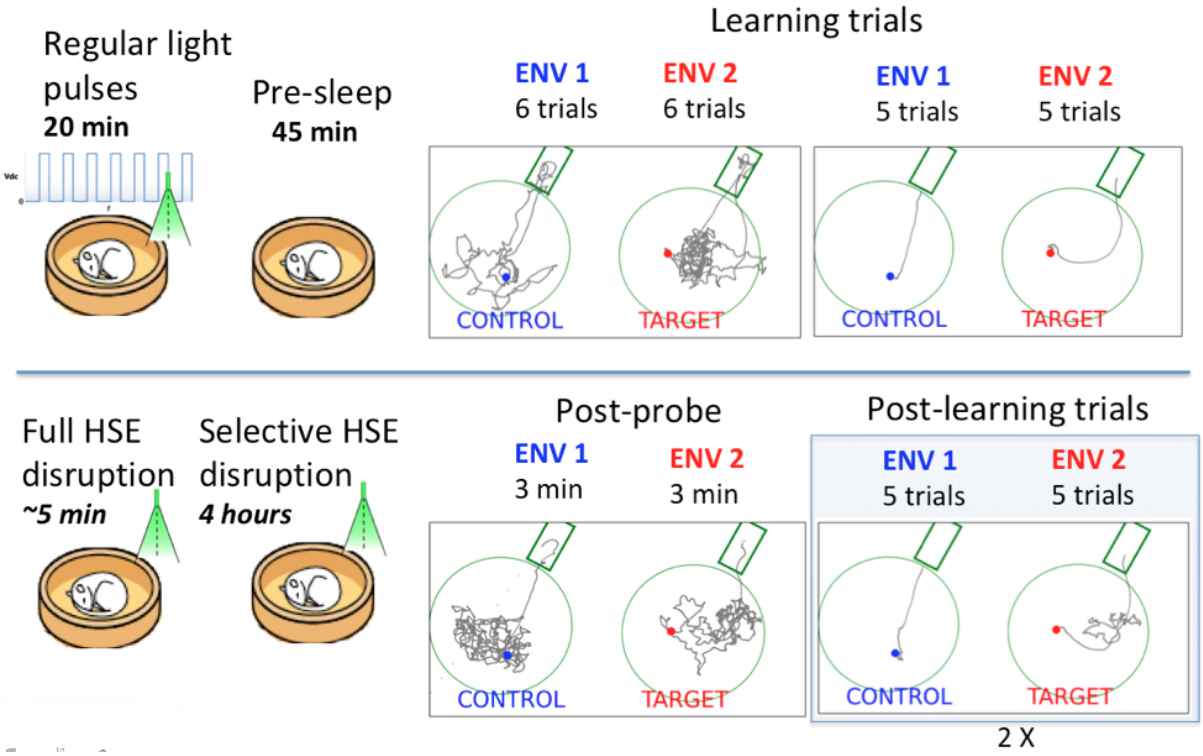

Fig. 2.3 | Experimental paradigm

Illustration of the experimental paradigm (time order left to right, top to down). The orange box depicts the sleeping box, green light triangle mark sessions where light pulses were triggered.

Second, a **pre-sleep** session of 45 minutes was recorded for subsequent analysis to establish differences in pre- and post- sleep sessions as well as for the calculation of some of the parameters required for the real-time analysis. The animal was sleeping in the cylindrical box with high walls in the centre of one of the cheeseboards. The animals must not have been able to associate their sleep location with one of the environments, because they were

placed in the sleeping box outside of both cheeseboard environments and then the box was placed in the centre of one of the cheeseboards. After the pre-sleep session, the animals were placed for a short time in the home cage, i.e. the cage in which the animal was housed, after which four blocks of the learning sessions were recorded, the first two consisting of six, and the remaining two consisting of five learning trials. At the beginning of each learning session the animal was placed in the start box, and recording was started. Each trial started with opening the door, after which animal entered the cheeseboard, located and consumed the goal, and returned to the start box. Next, animals were placed for a couple of minutes into the home cage, after which they were placed in the sleeping box. For a short (5-7 min) duration all synchrony events were disrupted while the calculation of the model required for decoding the environments was completed. This was followed by a selective disruption session that lasted for four hours and was recorded in two sessions of two hours with only a one minute break between them. After the recording of the sleep, the animal was placed in the home cage for five minutes. Next, the three minute long **post-probe** sessions in each environment were recorded. During the post-probe sessions, animals had to search for food in the cheeseboard environment while the goal location was not baited. The start box door was closed as soon as the animal had left the start box and was opened after the three minutes of exploration. Environments have been presented in the same order as during learning. After the post-probe session, more learning sessions were recorded during which animal had to perform two blocks of five trials in each of the cheeseboard environments (referred to as **'post-learning**'). Later I refer to the last block of post-learning trials (i.e. 5 trials in each environment) as the **end of post-learning.**

2.4 Training procedures

Following the recovery period after surgery, the animals were exposed to the maze over the period of 14 days. Animals were placed on food restriction (>85% of initial weight with a 10 g gain each week), once they passed the 7 days recovery period and exceeded the pre-surgery weight.

To train the animals for the double cheeseboard task, the animals were first habituated to the cheeseboards, sleep box and the start box in both environments. During the habituation phase, the food pellets were randomly and abundantly scattered over the cheeseboards. After an animal started collecting the pellets, the training was gradually transformed to the form of the final task: instead of a free exploration of the cheeseboard, the animal would be moved in and out of the start box, and the baited area of the cheeseboard was gradually shrunk until only one hole was left baited. In addition, the animal was habituated to switch between the two environments and learn two different food locations in the two environments. During the whole training period, the animal was also habituated to the regular light pulses. No post-probe sessions were performed during the training stage to avoid animals habituating to the absence of food after the sleep.

2.5 Position tracking system and spike sorting

The motion of the rat on the maze was tracked by following the trajectory of two LEDs attached to the head stage. A video camera (Dinion Camera, Bosch) was mounted on the ceiling above the recording environment and connected to the dacqTrack video tracking system of the Axona recording system (Axona Ltd, St. Albans, UK). The tracking signal was recorded together with the electrophysiological data on the same data file and read offline to timestamp every tracking sample.

The tetrodes were lowered into the CA1 pyramidal cell layer of the hippocampus, which is densely populated with neuronal cell bodies and action potentials (spikes) of many cells can be recorded simultaneously from a single tetrode location in this layer. Based on their waveforms, a spike sorting method was developed to separate the action potentials of different neurons. Spikes were extracted offline from the digitally filtered (0.8-9 kHz) signal using the same software that was used for online decoding. Action potentials with a power of 6.5 SD from the baseline mean were selected, and spike features were then extracted using principal component analyses. During the subsequent offline unit analysis, spikes from putative individual neurons were segregated into single units by using manual clustering software. Only units with clear cluster boundaries and refractory period in the

autocorrelation were used for further analysis. CA1 pyramidal neurons and interneurons were discriminated by their autocorrelation, firing rate and wave form as previously described (Csicsvari et al., 1999b). I have further confirmed the quality of cluster separation by using the isolation distance based on Mahalanobis distance (Harris, Henze, Csicsvari, Hirase, & Buzsaki, 2000) to ensure that calculated spike clusters did not overlap during the course of recording. In this way, I was able to identify 2578 CA1 pyramidal units.

2.6 Calculation of rate maps

Two-dimensional (2D) place-rate maps were calculated by subdividing the environments on a grid containing 6x6 cm bins. A 6x6 cm square was overlaid on this grid, for each spike of a given cell, centred on the position of the animal when the spike was emitted. Each bin was then incremented by the degree to which this square overlapped with it. The same procedure was then performed with the animal position tracking data to produce a map of occupancy. Both the occupancy (time spent in each bin) and spike matrices were then smoothed with a Gaussian kernel function (SD=10 cm), and divided to produce a place rate map. Many of the pyramidal cells spiked mostly in one of the environments, which allowed decoding the environment reliably from few spikes.

2.7 Real-time signal processing

One of the biggest technical challenges of my research project was to accomplish online spike detection and assembly detection, which required high-speed spike and high synchrony events (HSE) detection and environment decoding. A schematic representation of the closed-loop setup for selective disruption of HSE that I have implemented is depicted in the **Fig 2.4**.

The signal was acquired through the Axona recording system in the packages of 3 data samples (or the last 0.125 ms). As soon as a data package arrived at the recording computer, the software gained access to and processed this data package. Initial processing included

spike detection, feature extraction was described above (2.5 Position tracking system and spike sorting). Next high synchrony events (HSE) were detected.

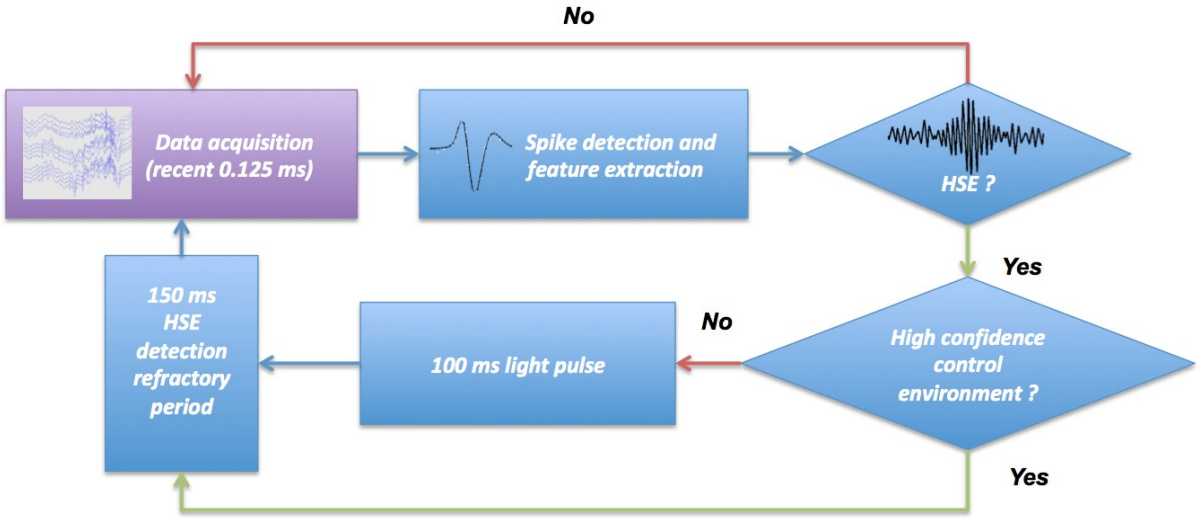

Fig.2.4 | Scheme of the closed loop setup for selective disruption of HSEs

Purple rectangle designates processing that happens in the Axona system, the blue components represent processing that happens in the custom-developed software. Rhombi designate the decision points.

HSEs were detected using multiunit activity (MUA) in the time windows of fixed length (20 ms). Based on the pre-sleep recording, the mean number of spikes in all tetrodes in the 20 ms time windows was calculated. This estimate was then used as a baseline for HSE detection, which was triggered when the number of spikes in a time window exceeded a threshold of 3.5 times the baseline mean. This threshold was then adaptively recalculated every minute based on the actual number of detected HSEs, so that the effective HSE detection rate would be near 1 Hz. This simple method nevertheless allowed reliable detection the high synchrony events, which often coincide with the sharp-wave ripple times. The beginning of the HSE was then adjusted to the time of the first spike in the HSE window. After the HSE detection, real-time decoding (see below) was used to determine, which of the two environments is encoded in the HSE window.

If the HSE was detected, decoding of the environment was performed (see decoding algorithm next). If the HSE with high fidelity encoded a control environment**,** the event was allowed to proceed (i.e. the event was not disrupted optogenetically). Otherwise, if no confident reactivation of the control environment was detected, the light pulse of 100 ms was triggered for the purpose of disrupting the reactivation. In both cases, **detection**

refractory period of 150 ms was applied to avoid multiple detections of the same HSE as well as spiking 'rebound' after the light pulse ceases. Even though most of the previous reactivation disruption studies have performed SWR detection instead of the HSE detection, I have chosen the latter because of the fact that the synchrony increase tends to occur before the onset of ripples (Csicsvari et al., 1999b) gives this MUA-based method advantage of earlier detection over the ripple-power based HSE detection methods. However, to compare the HSE and SWR detection in my data, I have post-hoc detected when the deviations in the ripple oscillations (150-250 Hz) power have exceeded the 4 SD threshold from the mean and compared the obtained times to the real-time HSE detection times. First, a total of 112485 SWR were detected in contrast to 301012 HSE events (total over 21 sessions). The larger number of HSE events could be explained by the enforced HSE detection rate of 1 Hz. Secondly, 92.8% of SWR were detected as HSE and 36.3% of HSE were detected as SWR. Finally and most importantly, HSEs were detected on average 17.3 ms earlier than the SWR.

Because the study was based on small number of animals (n=4), I have quantified the SWR properties, namely the occurrence rate, ripple frequency and the number of ripple oscillations per SWR and compared them to previously reported values. The reported SWR occurrence rates vary from 0.02 to 3 Hz (Ylinen et al., 1995) and can be modulated by novel experience (Eschenko, Ramadan, Molle, Born, & Sara, 2008; Foster & Wilson, 2006) and learning (Kudrimoti, Barnes, & McNaughton, 1999). I found that the mean SWR occurrence rates per animal (0.35 Hz, 0.41 Hz, 0.34 Hz and 0.38 Hz) to be in agreement with previously reported values. Similarly, mean ripple frequency (180.1±15.7, 178.89±16.0, 181.3±16.1, 179.5±15.9) and mean number of cycles (6.3±1.68, 7.17±2.46 6.87±1.93, 7.02±2.12) were in agreement with previously reported ranges of 150-250 Hz and 5-15 cycles correspondingly (Buzsaki et al., 1992; Chrobak & Buzsaki, 1996; Ylinen et al., 1995).

2.7.1 Original decoding method and its optimizations

Most of the recently reported tools for real-time signal processing (Deng et al., 2016; Kloosterman et al., 2014) were not fast enough to disrupt an on-going HSEs, because the short HSE would be over before the necessary calculation would have been performed. To

reach as short delays as possible, I have integrated a detection and decoding software I developed with the Axona recording system to gain access to the signal as soon as it reached the RAM of the recording computer. In addition, I have utilized multiple optimizations including integer arithmetic and SIMD (Single Instruction, Multiple Data) instructions to speed up the signal filtering, which was the slowest part of the signal processing, constituting around 35% of all data processing time. To decrease the delay even more, the analogue hardware signal filtering can be performed. Moreover, 1D models based on linearized trajectories were used instead of 2D models and multiple optimizations were introduced to the original decoding method to make the real-time decoding as efficient as possible.

The real-time decoding was based on the method presented in (Kloosterman et al., 2014). The method relies on the Kernel Density Estimation (KDE) procedure, which is a non-parametric way to estimate a probability density function of a random variable:

$$f(x) = \frac{1}{n}\sum_{i=1}^{n} K\left(\frac{x-x(i)}{h}\right) \quad (4)$$

where $x_1, \ldots, x_n$ is the data sample, h is the kernel width and K is the kernel density estimator, which often is a Gaussian kernel:

$$K(x) = exp\left(-\frac{1}{2}x^2\right) \quad (5)$$

The main idea behind this method is that it allows to bypass the spike sorting step, which is required for unit-based decoding, and decode the position based on the KDE of the bivariate distribution of spike features and stimulus. In the notations below, a denotes spikes features vector and x denotes stimulus, which in my case was a linearised position, N denotes the number of spikes used for KDE and R denotes number of position samples. The bivariate distribution of spike features and location $p(a, x)$, prior location distribution $p(x)$ and distribution of spiking locations $\pi(x)$ are estimated using KDE:

$$p(\boldsymbol{a}, \boldsymbol{x}) = \frac{1}{N}\sum_{n=1}^{N} K_{H_{ax}}\left(\begin{bmatrix}\boldsymbol{a}\\ \boldsymbol{x}\end{bmatrix} - \begin{bmatrix}\tilde{\boldsymbol{a}}_n\\ \tilde{\boldsymbol{x}}_n\end{bmatrix}\right) \quad (6)$$

$$p(\boldsymbol{x}) = \frac{1}{N}\sum_{n=1}^{N} K_{H_x}(\boldsymbol{x} - \tilde{\boldsymbol{x}}_n) \quad (7)$$

$$\pi(\boldsymbol{x}) = \frac{1}{R}\sum_{r=1}^{R} K_{H_x}(\boldsymbol{x} - \tilde{\boldsymbol{x}}_r) \quad (8)$$

Here H_x and H_{ax} are the bandwidth matrices for stimulus only and for spike features and stimulus correspondingly, which are discussed in (Kloosterman et al., 2014). Given these probability estimations, the position encoded in the population activity can be then estimated from the set of new spikes using Bayes formula, Poisson spiking assumption and conditional probabilities of position given the features of spikes in the decoding window:

$$P(\boldsymbol{x_t}|\boldsymbol{a_{1:m}}) = \frac{P(\boldsymbol{a_{1:m}}|\boldsymbol{x_t})P(\boldsymbol{x_t})}{P(\boldsymbol{a_{1:m}})} \tag{9}$$

Where $P(a_{1:m})$ is a normalizing constant such that the posterior $P(x_t\,|a_{1:m})$ is a proper probability distribution, $P(x_t)$ is a prior position distribution, which I defined as a uniform for simplicity, and the conditional probability of spikes features given the position $P(a_{1:n}|x)$ is computed as

$$P(\boldsymbol{a_{1:n}}|\boldsymbol{x}) = (\Delta t)^n \left[\prod_{i=1}^{n} \lambda(\boldsymbol{a_i}, \boldsymbol{x})\right] \left[e^{-\Delta t \lambda(\boldsymbol{x})}\right] \tag{10}$$

where n is a number of spikes in the decoding window Δt, and the generalized rate function $\lambda(a, x)$ and marginal rate function $\lambda(x)$ are defined as:

$$\lambda(\boldsymbol{a}, \boldsymbol{x}) = \frac{spikecount(\boldsymbol{a}, \boldsymbol{x})}{occupancy(\boldsymbol{x})} = \frac{N}{T}\frac{p(\boldsymbol{a}, \boldsymbol{x})}{\pi(\boldsymbol{x})} = \mu\frac{p(\boldsymbol{a}, \boldsymbol{x})}{\pi(\boldsymbol{x})} \tag{11}$$

$$\lambda(\boldsymbol{x}) = \mu\frac{p(\boldsymbol{x})}{\pi(\boldsymbol{x})} \tag{12}$$

This method not only allowed for more accurate decoding, because no information is lost in the clustering process, but also to design experiments, in which decoding has to be used before the clustering can be performed. The workings of the method are depicted in the **Fig 2.5.**

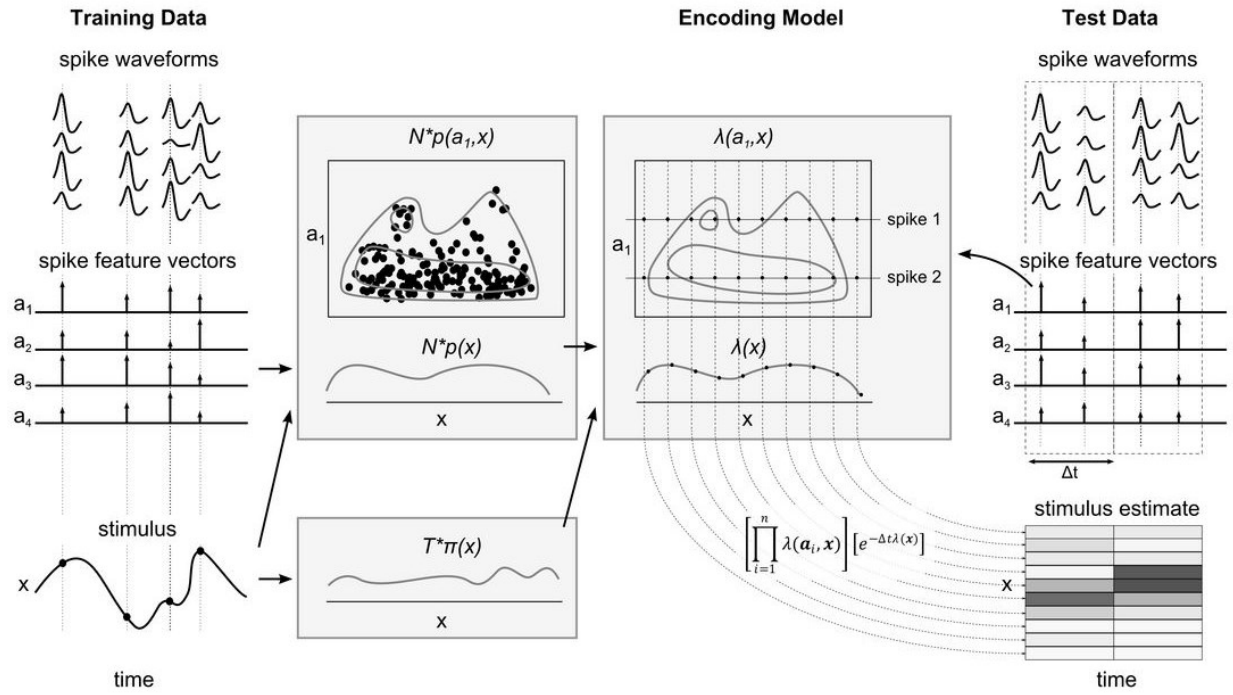

Fig. 2.5 | Illustration of the original decoding method

During the encoding step, spike feature vectors are extracted from spike waveforms. Then, an estimation of the bivariate distribution of spike features and stimulus as well as estimations of stimulus distribution and spike location distribution are calculated based on the extracted spike features and stimulus data. During the decoding step, a conditional probability distribution of stimulus given the spike features are calculated for every new spike and then combined according to (11) to produce the final stimulus estimate. Figure adapted from (Kloosterman et al., 2014)

Given a sequence of spike wave shapes occurring in an HSE across different electrodes, the decoding model outputs a likelihood map, which describes the likelihood of every possible position being encoded in the given spike sequence. For example of such output of decoding an HSE spike pattern during sleep see **Fig. 2.6,** in which the model was trained to decode a linearised position in the two cheeseboard environments. The most likely encoded position is located in the environment on the right, which also means that the environment on the right and not on the left is encoded in the given HSE spike pattern. In my paradigm I was not only interested which environment was encoded in the given HSE spike pattern, but also how confident the algorithm is about its decision about the encoded environment. This **confidence** can be defined as a difference of two maximum likelihoods in the two environments - the higher its absolute value, the more reliable the result of the decoding **(Fig. 2.11)**. In my paradigm, I aimed at disrupting all the events except for the **top most confident 20%, i.e.** ones that preferentially encoded the control environment. The

confidence threshold was initially set based on the distribution of confidences in the pre-sleep HSEs. Later, this threshold was adaptively adjusted every three minutes based on the actual percentage of disrupted events to enforce the overall disruption rate of 80%. The threshold selection in pre-sleep and confidences of disrupted and control events are shown in **Fig. 2.7**.

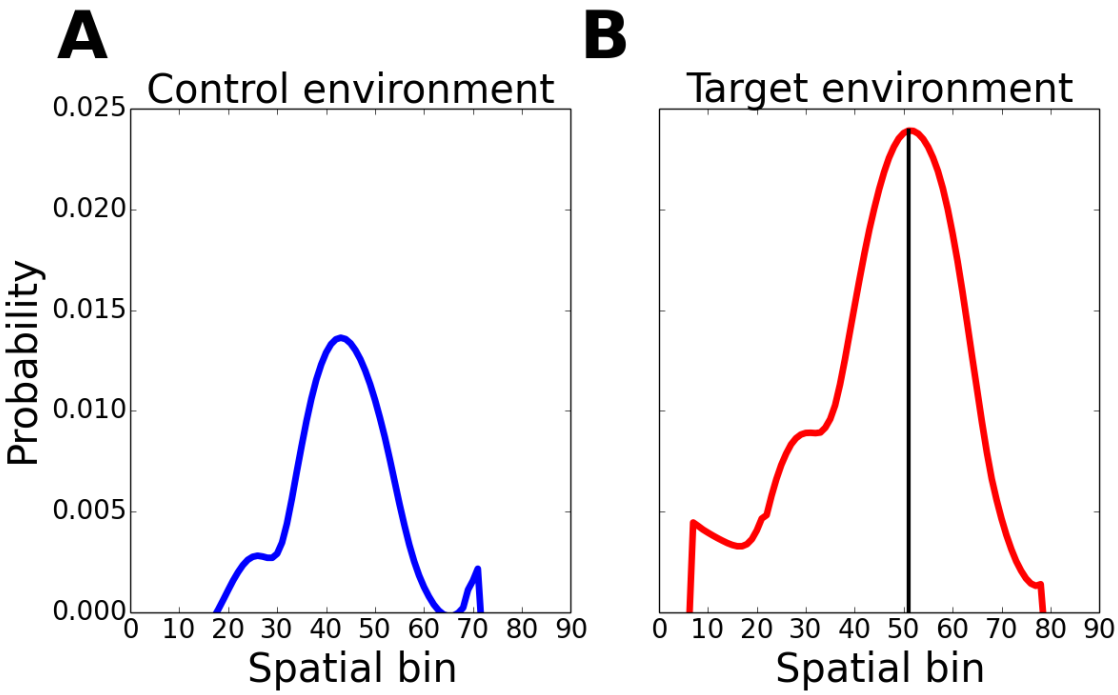

Fig. 2.6 | Example of the decoding output

Probability distribution of the linearised position encoded in the HSE spike pattern based on the model constructed from spiking and linearised positions the animal moves from the start box to the goal and back. The most likely position is designated with the black vertical line and is located in the target environment. Probability of this location is 0.024 and probability of the most likely location in the control environment is 0.014, so the confidence of decoding is the natural logarithm of their ratio: ln (0.024 / 0.014) = 0.54.

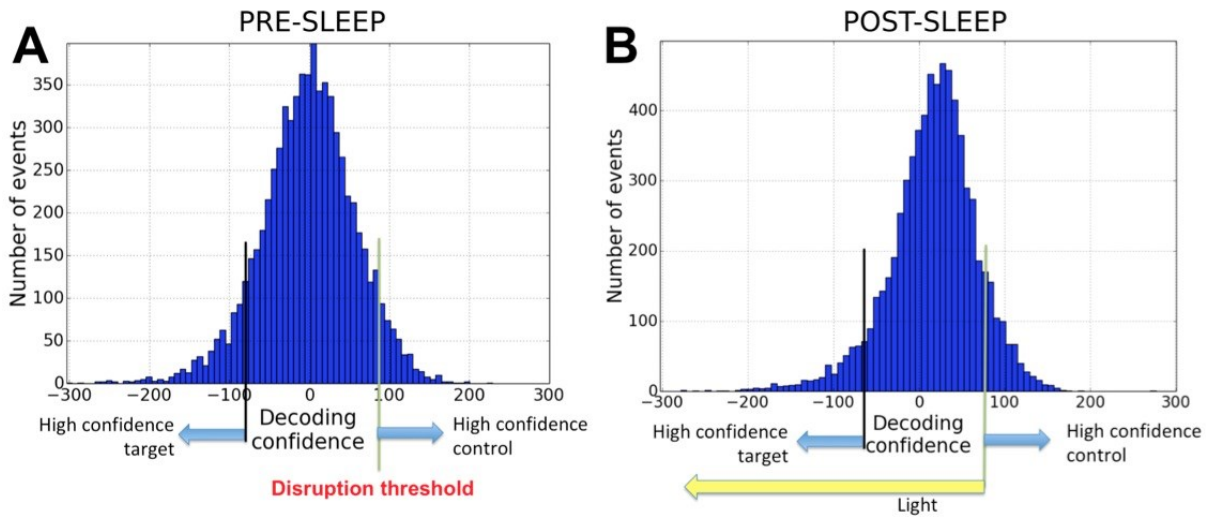

Fig. 2.7 | Illustration of disruption confidence threshold illustration

(A) Distribution of confidences (difference of the two maximum likelihoods in the two environments) of the pre-sleep HSEs and a disruption threshold, selected at the 80% level of the distribution. (B) Distribution of the confidences of the post-sleep HSEs and the events that were disrupted according to a selected threshold.

However, in its original form, the decoding part of this method is prohibitively computationally expensive, because during decoding it requires to compute the KDE for each incoming spike in the decoding window. Therefore, I have introduced multiple optimization procedures to adapt this method for the real-time HSE disruption.

Substitution for real-time KDE

First of all, given similar distribution of spike features in the learning and in the subsequent sleep signals, there was no need to calculate KDE for every new spike, as every new spike will be similar enough to one of the spikes observed before, for which KDE can be pre-computed. Moreover, it was not necessary to pre-compute KDE for all spikes collected during the learning period; around 500 **'pivot'** spikes were enough to gain speed without losing decoding accuracy. These pivot spikes are chosen from all available spikes by an iterative exclusion procedure, according to which, at every iteration, a spike with the largest number of neighbours in a given radius was chosen as a pivot, and all its neighbours were excluded. And if a set of pivot spikes was available, the most similar one to the given new spike can be efficiently found by the search in the **k-d tree** data structure. A k-d tree is a space partitioning data structure for organizing points in a multidimensional space and allows to perform a search in $O(log\ n)$ time. The dimensionality in case of spike features

equals $2T$, where T is the number of channels in a group, i.e. 2 principal components for a wave shape on every channel).

Linearisation

The 2D trajectories of the animals in the cheeseboard environments were linearised by projecting them on the line containing the centre of the start box and the goal location. This allowed operating in the 1D space and save time on expensive matrix operations. However, it must be noted that the algorithm and software are fast enough for 2D decoding as well.

Continuous decoding

To ensure as fast decoding upon the HSE trigger as possible, the software continuously maintained the likelihood map decoded from the most recent time window by means of running average. This allows minimizing the amount of computation immediately after the HSE detection, i.e. at the most delay-critical time point.

Fixed number of spikes

Finally, when this method was applied to decode sleep activity, strategies with fixed time window or a fixed number of spikes can be used. I have opted for the latter since this strategy allowed compensating for the fact that the firing rates scale within the HSE and that the HSE detection can be triggered at different stages of the event whereas a fixed (e.g. 100 ms) time window would contain highly variable number of spikes and could delay the decision and event detection. So in my setup, the decoding of the HSE was always done from a fixed number of the last available spikes. The exact number of spikes was defined as a median number of spikes observed in consecutive 100 ms time windows during learning.

2.7.2 Closed loop delay

Minimizing the processing delay in the closed-loop setups is crucial for many experimental paradigms including the one presented in this thesis. The short processing delay enables the disruption of a larger proportion of reactivation events that encode the target environment and, consequently, leads to more pronounced physiological and behavioural effects.

In addition to the algorithmic changes described above, an important update to the Axona system allowed to gain access to stream of data at finer time resolution: instead of default packages of 300 samples, which spans 12.5 ms of signal, in my setup data was available in the batches of 3 packages, which spans only 125 µs of signal thus allowing for faster spike detection and a proper response. It is worth mentioning that the developed algorithm and software are generic with respect to the variables of the distribution and can be used in many different applications that require efficient decoding of any continuous or discrete variables in a cluster-less decoding method. Another advantage of the software is that it is cross-platform and can be with little effort integrated with recording systems other than Axona.

I have defined a processing delay of my setup as a time between the last sample of a spike that triggered a synchrony detection and light pulse, and a significant deflection in the LFP caused by the light, which can be calculated from a spike-triggered average, centred on the last timestamp of the spike that triggered the synchrony detection and the light pulse. According to this measure, a significant change in the LFP occured on average already **1.04 ms ± 0.09 ms (mean ± SD, n=21 sessions)** after the last timestamp of the HSE-triggering spike. It is true, however, that the network response dynamics of some place cells was slower than the LFP response. This was due to the fact that the AAV with CaMKII promoter used resulted in a leaky expression and was expressed in inhibitory interneurons as well (Schoenenberger, O'Neill, & Csicsvari, 2016) (see also **chapter 3.1**). As a result, many pyramidal cells increased their firing due to disinhibition, and these cells responded to the light stimulation with up to a 10-20 ms delay (**Fig. 2.8).** However, those cells that were inhibited suppressed their firing within 1-2 ms after the light onset.

For my experiments, I have used a recording computer with the Intel Core i5 CPU and 8 GB of RAM with Windows 7 operating system (OS) installed. Even though I have used the Windows-based recording system in my setup, because it was a part of already established protocols in the lab, I recommend using Linux-based recording software to reach shorter OS-associated delays. More powerful recording computer could also help gain shorter delays. In addition, a parallel version of spike detection code performs slightly faster than the non-

parallel code on the Linux computer with Intel Core i7 CPU with 6 cores and 12 processing threads, according to my tests. Therefore, using the CPU with similar or larger number of cores can yield an additional gain in processing speed.

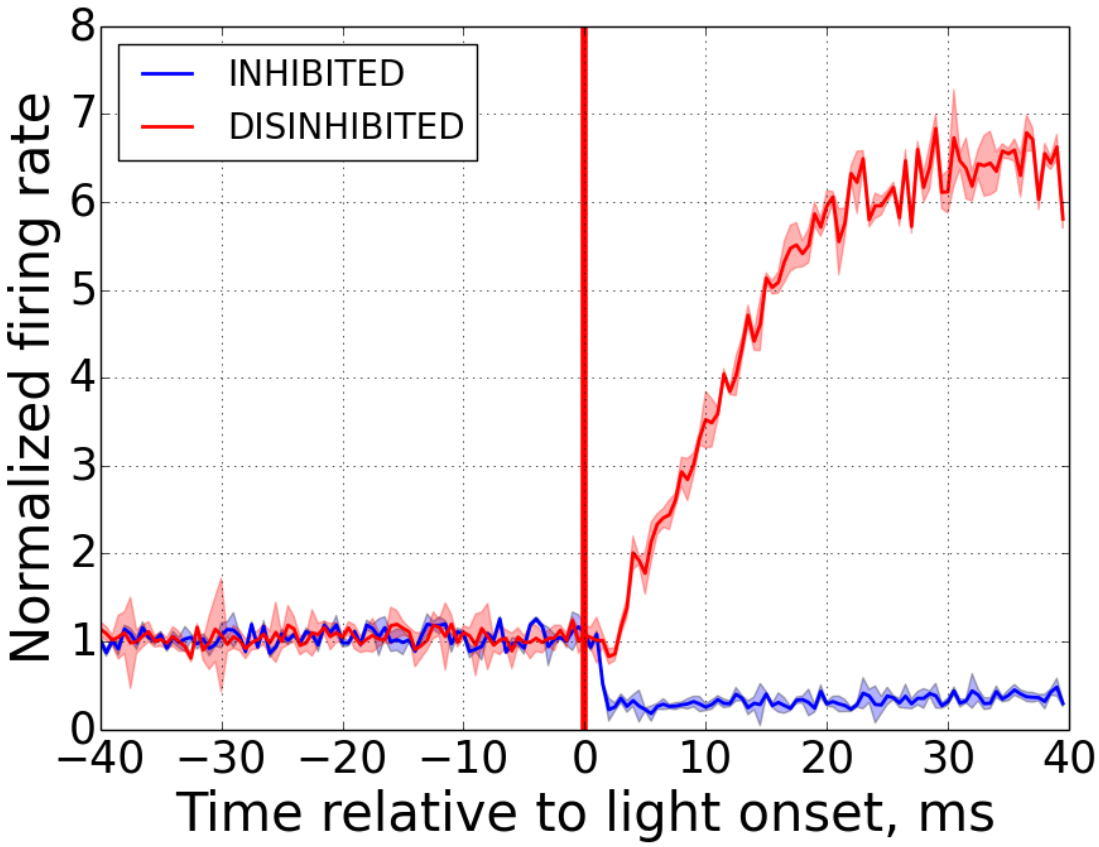

Fig 2.8 | Mean response of firing rate of inhibited and disinhibited cells to the light pulse.

The shaded areas around the response curves show the standard error of the mean. Bin size = 0.5 ms. N=457 inhibited cells, N=732 disinhibited cells. Firing rate was normalized by the mean firing rate in the 100 ms window before the light pulse.

2.7.3 Decoding performance

The described method has multiple hyper-parameters, both present in the original method such as the bandwidth of the kernels for the density estimation, as well as those derived from its optimization such as a number of pivot spikes. These hyperparameters were tuned offline, based on data obtained from pilot animals using the stochastic gradient descent

method with the cost function including both decoding accuracy and decoding error. These hyperparameters need to be retuned if the software will need to be used in a different context.

I have used two measures to evaluate the performance of the decoding model: median decoding error describes how different a decoded position was from the real one, and binary decoding accuracy shows how often the method correctly decodes one of the two environments. The latter is more important for my experimental paradigm, as my goal was to detect which environment was being reactivated or was encoded in the HSE. Performance of the decoding was evaluated using cross-validation: the model was built using 100 ms windows with odd numbers, while evaluation was performed using the 100 ms windows with even numbers. However, windows in the second half of the learning were used only to build and evaluate the model, as the maps are more stable towards the end of learning. First, **Fig 2.9** shows an example of the decoding during the second half of the learning. The median error of position encoding was **9.8 cm** with a 95% CI of [3.6; 17.1] as measured across all recordings.

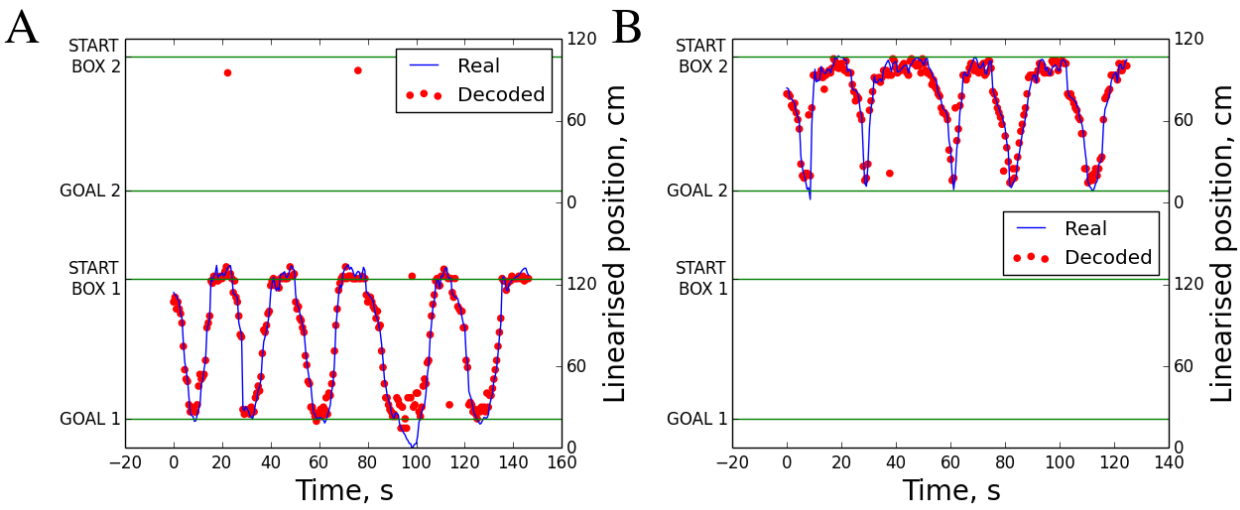

Fig. 2.9 | Real and decoded position during a learning session

A 1D model based on linearized trajectory was used, the position is plotted versus time with real position showed by the blue line and red dots representing a decoded position for every even-numbered 100 ms window. (A) Decoding in the control environment. (B) Decoding in the target environment.

In addition, a quantified decoding accuracy, which reflects the proportion of the decoding windows in which the control environment was encoded and the animal was present in the control environment. Decoding accuracy increased with the confidence of decoding. The 'leakage' of the reactivation of the target environment in the control events that were classified as reactivating the control environments and did not trigger the light pulse. In general, the accuracy of decoding increases with confidence with the least 20% confident time windows containing most of the errors (**Fig 2.10)**.

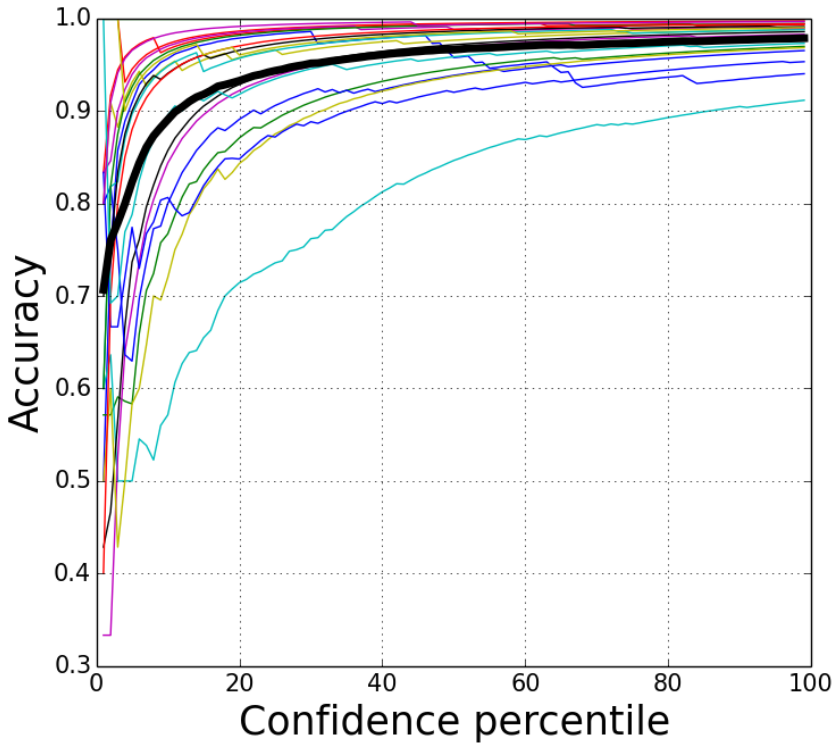

Fig. 2.10 | Dependence of decoding accuracy on the confidence

Each line represents a separate recording day accuracy profile, and a bold line represents the mean. Confidence was measured as the difference of the logarithms of maximum Bayesian encoding probabilities of the two environments, which were rank ordered.

In order to minimize the 'leakage' of the target environment reactivation during the disruption through these being erroneously classified as control environment reactivation, we detected target environment only as the top 80% confident events. On average, the decoding accuracy at 80% was **97.5% ±2.2% (mean ± SD)**. Thus, the method copes very well with the task of binary classification.

Of course, performance of the model depends on the number of cells recorded. This dependence is summarized in **Fig. 2.11** with two measures, decoding accuracy (at 80% confidence) and median error, plotted against the median number of spikes in the 100 ms window for every recording session. However, the correlations of these measures with the number of spikes were not significant (R = 0.37, P = 0.105 for the decoding accuracy at 80% confidence, and R = -0.27, P = 0.244 for the median error).

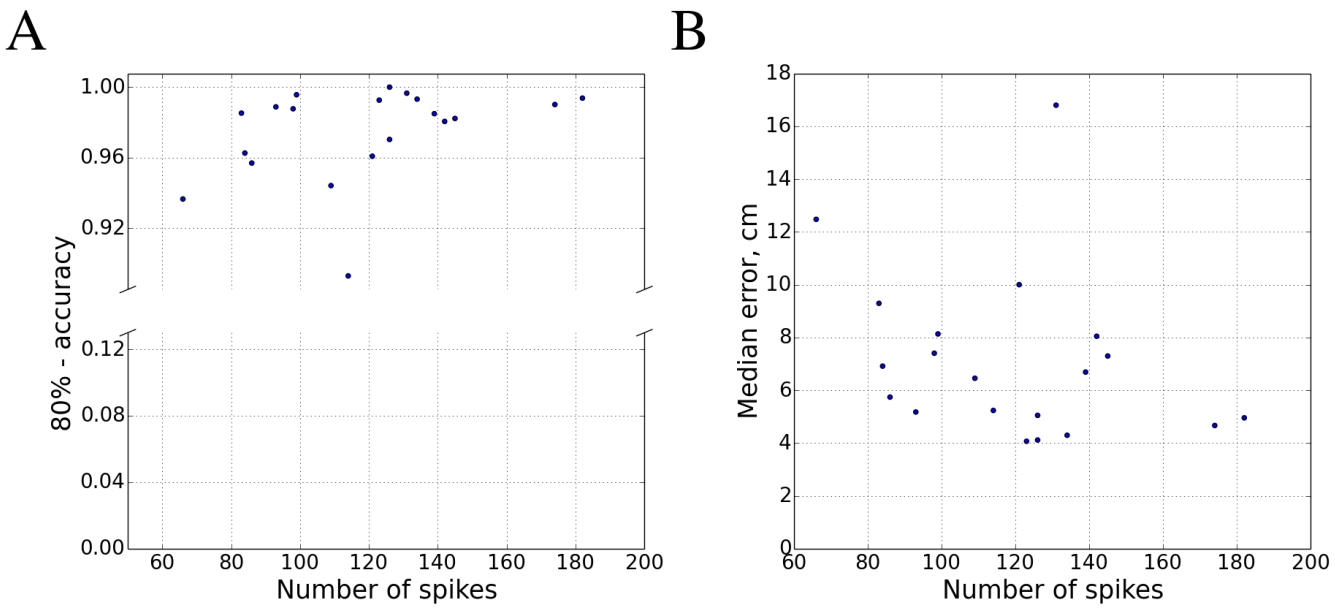

FIg. 2.11 | Dependence of decoding performance measures on the yield.

(A) Decoding accuracy at 80% confidence level versus the median number of spikes in the 100 ms window during the waking activity. (B) Median error of position decoding versus the median number of spikes in the 100ms during the recordings.

2.7.4 Dissociation of two environment reactivations

The method and experimental paradigm rely on the assumption that the reactivation of the two environments happens independently. This assumption itself is a separate research problem that has not been studied so far. To illustrate the presence of two reactivated assemblies, I have analysed the sleep cofiring correlations matrices of cells. Cofiring correlation was measured between pairs of cells by calculating the Pearson correlation of their instantaneous firing rates during the HSE events. **Fig. 2.12** illustrates such correlation matrices for two control recording days, in which cells have been ordered according to the

hierarchical clustering. It is clear that the hierarchical clustering revealed two groups of cells tend to cofire within their group while firing less with less in the other group.

.

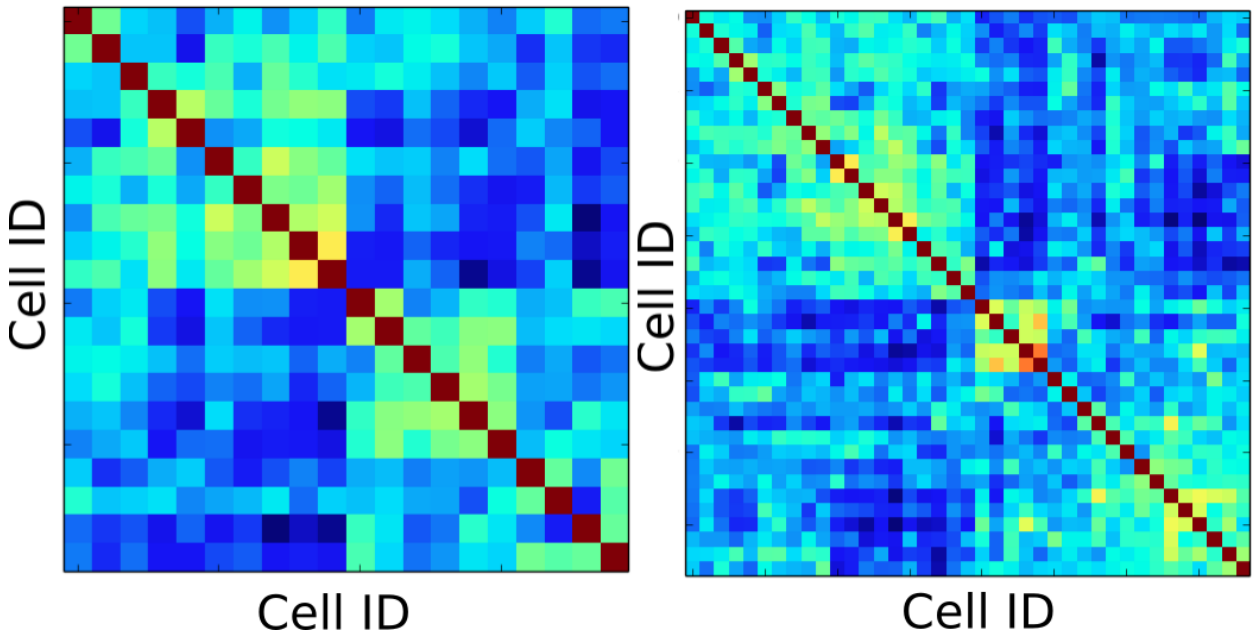

Fig. 2.12 | Sleep cofiring correlation matrices

Two examples of cofiring correlation matrices. Cofiring correlations were calculated during the HSE windows (N = 11473 windows and N = 20 cells for the left matrix, N = 12523 windows and N = 39 cells for right matrix).

2.8 Statistical analyses

The statistical analyses were performed using the C++-based software as well as Python scripts. Mann-Whitney U test was used for non-parametric comparisons of two data samples. Bootstrapping was used to define confidence intervals and significance of the behavioural scores. The post-hoc multiple testing correction using the Holm-Bonferroni method was used in cases of multiple testing. Fisher's r-to-z transform was used to compare the Pearson correlations. Kolmogorov-Smirnoff test was used to compare distributions.

3 RESULTS

3.1 Influence of optogenetic light application on the firing of CA1 cells.

First, I quantified to what degree hippocampal assemblies were affected by the light application. For example, imperfect hippocampus coverage during viral injection as well as non-optimal positioning of the optic fibres may have influenced a small subset of the hippocampal CA1 cells. To quantify how cells were affected by the light, I have calculated the light response of each cell and looked at the distribution of these responses. The cell light response was measured as the natural logarithm of the ratio of firing rate in the light-on period and the baseline rate before the light. Light-on periods started at 10 ms after the light was triggered and lasted until the end of the light pulse while the baseline rate was measured in the 90 ms window preceding the light. These responses were calculated from the 20 minute regular light pulses session **(Fig. 2.1)**, where 100 ms pulses were delivered at regular intervals of 1 s while the animal was rested in the sleep box. As illustrated by examples of the light pulse peristimulus time histograms in **Fig. 3.1** of two pyramidal cells and one interneuron, a subset of pyramidal cells and interneurons were inhibited by the light while some pyramidal cells increased their firing due to disinhibition. Previously using halorhodopsin and CaMKII promoter we obtained similar results (Schoenenberger et al., 2016). Here I confirmed that the leaky expression of the AAV was not specific to halorhodopsin, but ArchT was similarly expressed both on pyramidal cells and interneurons.

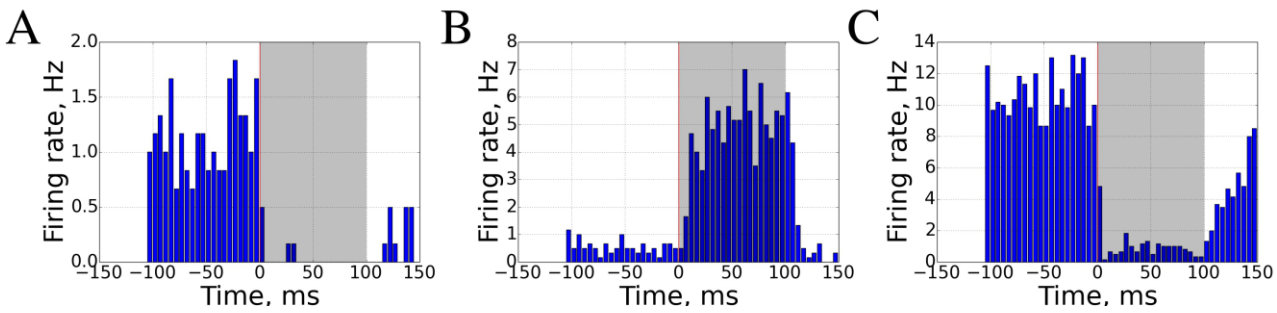

Fig. 3.1 | Examples of cell light responses

(A) Inhibited pyramidal cell light response (B) Disinhibited pyramidal cell light response(C) Inhibited interneuron light response. The vertical red line represents a start of the pulse; the shaded area represents the period of a light pulse. The histogram was calculated based on 20 minutes regular (at a 1Hz rate) 100 ms light pulses (n=1200 pulses) session. Bin size = 5 ms.

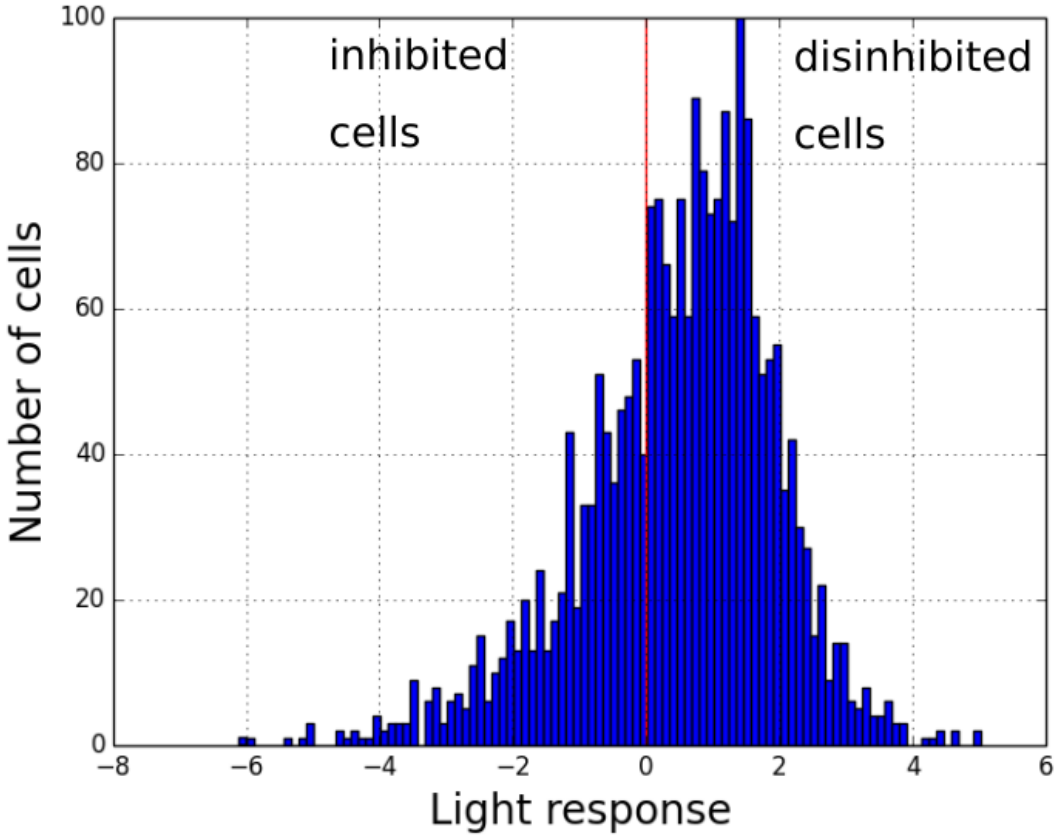

FiIg. 3.2 | Histogram of cell light responses

Total number of cells N = 2578. Disinhibited cells have a positive light response, inhibited – negative. The light response value was calculated as a natural logarithm of the ratio of the firing rates in the light-on period and the baseline period before the light.

Fig. 3.2 describes a distribution of light responses of all recorded pyramidal cells and interneurons. Interestingly, I found that the large proportion of pyramidal cells were excited by the light because of disinhibition. This explanation is also supported by the fact that all

the interneurons that I have recorded were strongly inhibited by the light **(Fig 3.1 C)**. Therefore, I will further refer to the pyramidal cells that have a positive light response, i.e. increased their firing rates during the light pulse, as **'disinhibited**'. The dominance of disinhibition did not impede with the original goal of the study to perform the disruption of consolidation, as disinhibition also disrupts the reactivation patterns because it excites also those cells that preferentially fire in the control environment, as I will demonstrate later in **chapter 3.4**.

I have chosen the following thresholds to define the categories of inhibited and disinhibited cells: a cell was labelled as **inhibited** if its rate dropped at least **25%** as a result of light application and it was labelled as **disinhibited** if its rate increased at least **two-fold** as a result of the light application.

3.2 Influence of assembly-related reactivation disruption on the reactivation patterns

3.2.1 Firing rates

While there is no ground truth data to estimate how accurate the decoding algorithm performed in sleep, analysis of the population activity in the disrupted and control events before and after the decision point of the algorithm can reveal which kinds of ensembles were disrupted and which were not. **Fig. 3.3** contains an example of the raster plots for four following events from the same recording day: a single learning trial in the control and target environments from the start box to the goal and back **(Fig.3.3 A, B)** and two HSE raster plots during subsequent sleep, one HSE, in which the algorithm detected the control environment and which was not disrupted **(Fig. 3.3 C)** and one disrupted HSE **(Fig. 3.3 D).** First of all, as can be seen from the learning trial raster plots, cells preferably fire in the corresponding environment during run and different cells fire as the location of the animal changes. Secondly, the field in the sleep raster plots shows that synchrony has been detected before or soon after the ripples onset. And most importantly, the firing during the HSE in which the light was applied loses selectivity toward the target environment after the

light onset, while the selectivity of the firing in the control HSE persists after the assembly-classification point.

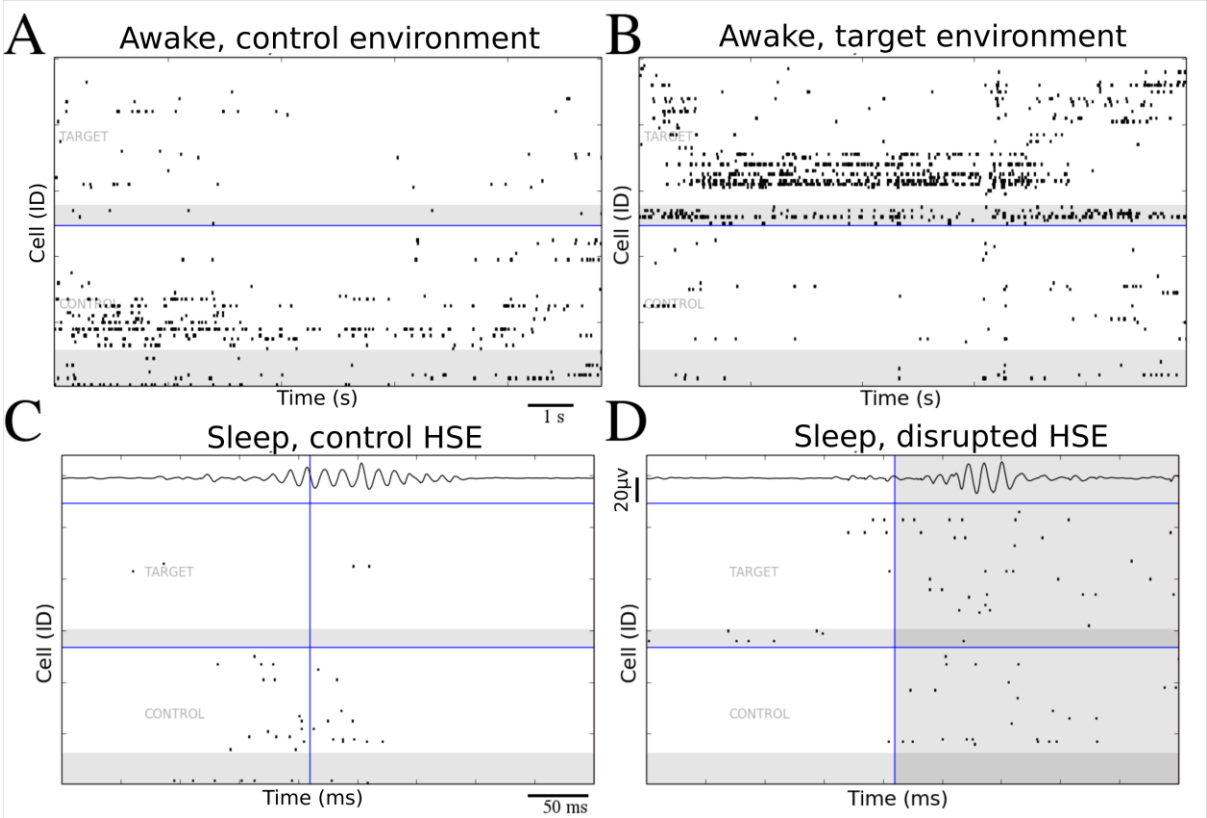

Fig. 3.3 | A representative example of raster plots of learning trials in control and target environments and control and disrupted HSE

(A) Run in the control environment. (B) Run in the target environment. (C) Control HSE with reactivation of the control environment. (D) Disrupted HSE with reactivation of the target environment. Cells were grouped into those with higher firing rate in control (bottom section of the raster plots) and target environments (top section on the waking raster plots and middle section in the sleep raster plots). Inhibited cell rasters are plotted on the grey background, while disinhibited cells are plotted on the white background (light grey after light onset). Rasters for same cells are plotted in the waking and sleep cases, however the ordering is different. The cells in the waking rasters are ordered by the location of the peak of their place field, separately so for inhibited and disinhibited cells, while cells in the sleep rasters are sorted by their light response. The top part of the sleep raster plots shows the LFP in the frequency band of 150-250 Hz around the HSE. The vertical blue line illustrates the time point, at which the HSE was detected, and a classification decision was made.

The following analysis presents an evidence for the selective disruption of population activity at the single cell level. In order to compare the firing responses of cells, I selected those cells that preferentially fired in either the control or the target environments. Thus, the '**control cell**' group includes cells that have higher firing rate in the control than in the target environment and vice versa for the '**target cells**'. Then, I compared the firing of these cells before and after the assembly detection in the disrupted HSEs **(Fig. 3.4 A-C)**. In addition, I also examined firing before and after the assembly detection in the HSEs that were classified (offline during the analysis) as reactivating the target environment with high confidence (top 20%) **(Fig 3.4 D-F)**. Finally, I have compared the firing before and after the assembly detection in the undisrupted HSEs, i.e. those that were with high confidence classified as reactivating the control environment **(Fig, 3.4 G-I)**. In addition, I have performed this analysis separately for disinhibited and inhibited cells. The Mann-Whitney U test was used to compare the firing rates. Below I present my results and conclusions of firing rate comparisons in the three groups of events:

- In the control HSEs the firing rates of control environment-encoding cells were significantly higher than the firing rates of target environment-encoding cells before the assembly detection for disinhibited cells ($p < 5*10^{-5}$), inhibited cells ($p < 0.006$) and all cells ($p < 6*10^{-12}$) as well as after the assembly detection ($p < 0.04$ for inhibited cells, $p < 0.0008$ for disinhibited cells, $p < 8*10^{-6}$ for all cells). First, this serves as an evidence for the assumption that the real-time decoder had correctly identified the events reactivating the control environment because cells that preferentially fired in the control environment also fired stronger in the decoding as compared to target environment coding cells in cases when the light was not triggered. Second, the significant difference in firing rates of control and target cells after the assembly detection demonstrates that the reactivation is proceeding normally.

- In case of the events, confidently (80%) classified as reactivating the target environment, the firing rates of target environment-encoding cells were significantly higher than those of the control environment-encoding cells before the assembly detection for disinhibited cells ($P < 3*10^{-5}$), inhibited cells ($P < 0.02$) and all cells ($P <$

$6*10^{-13}$). However, after the assembly detection, which coincides with light application in case of disrupted HSEs, the rates were significantly different for all cells ($P < 0.04$) but not for inhibited ($P > 0.81$) and disinhibited ($P > 0.32$) cells. Similarly to the non-disrupted HSEs, this serves as an evidence for the assumption that the decoder is also capable of correctly identifying events that reactivate the target environment with high confidence. Moreover inhibited cells of both the control and target environment-encoding cells suppressed their firing while disinhibited cells increased it.

- Finally, in the case of all disrupted events, the firing rates of target environment-encoding cells were significantly higher than those of the control environment-encoding cells before the assembly detection for disinhibited cells ($P < 0.009$) and all cells ($P < 8*10^{-6}$) but not for inhibited cells ($P > 0.23$). Of course, all disrupted events also contained events that may have encoded the control environment but could not be confidently encoded. After the light onset, there was no significant difference in disinhibited ($P > 0.79$), inhibited ($P > 0.67$) and all ($P > 0.48$) cells.

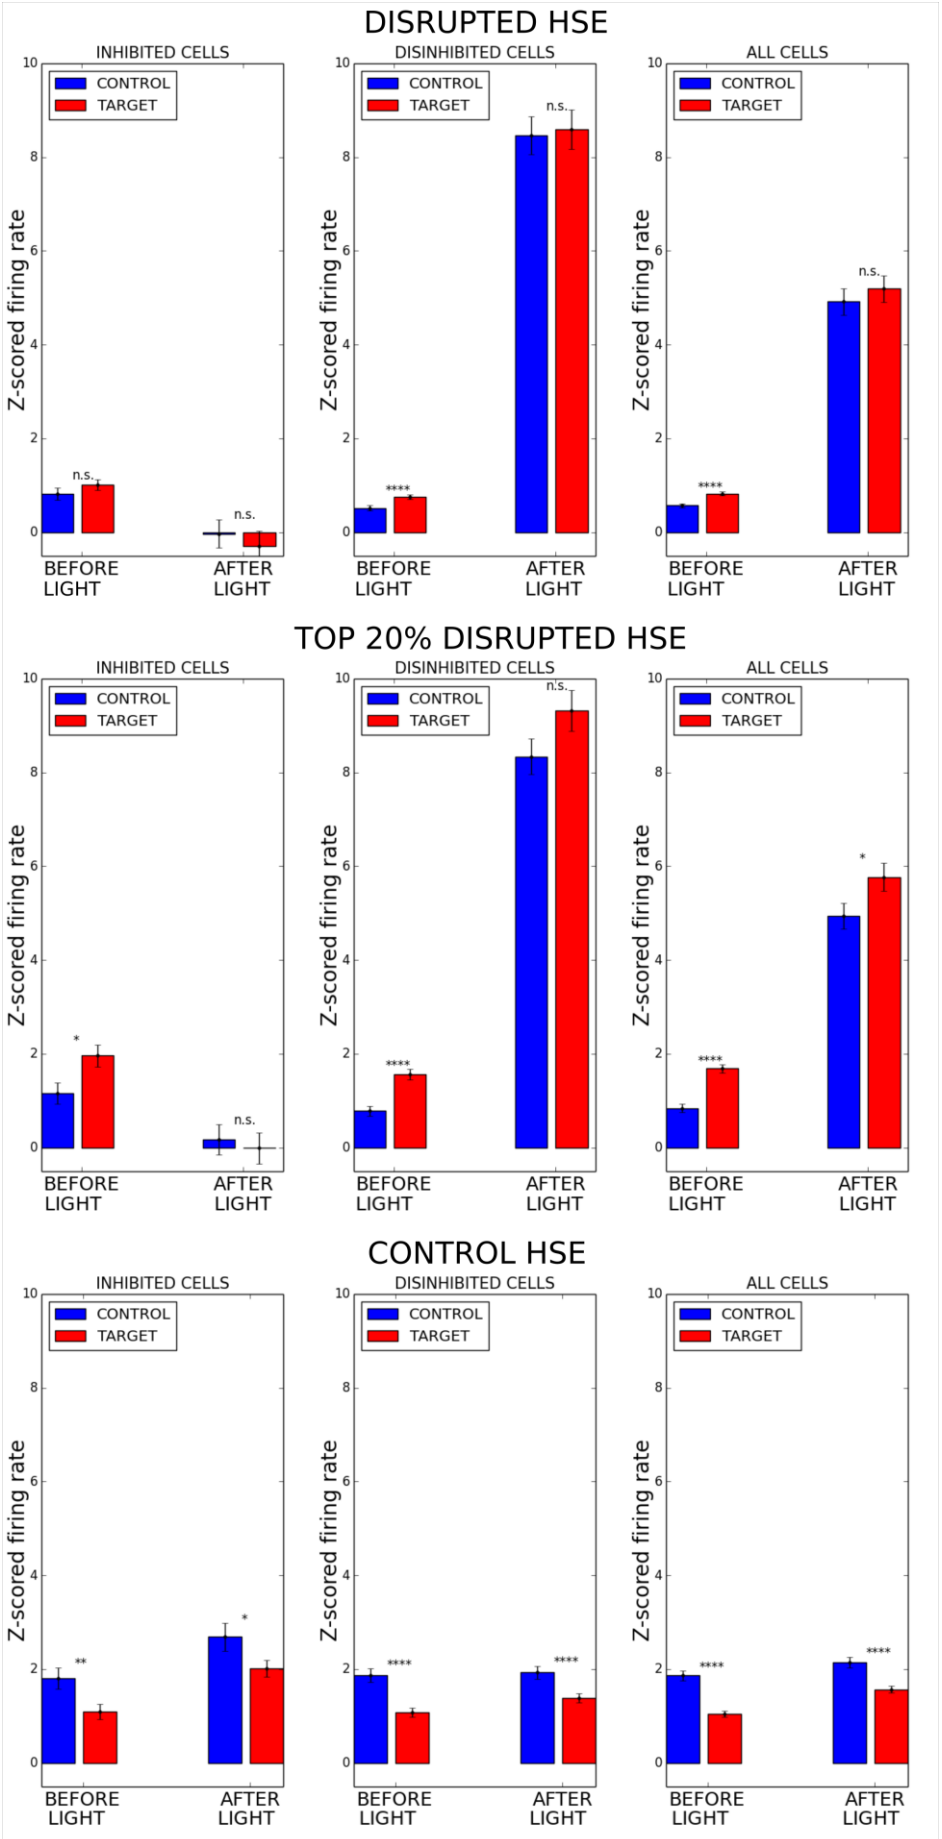
DISRUPTED HSE
INHIBITED CELLS
DISINHIBITED CELLS
ALL CELLS
CONTROL
TARGET
Z-scored firing rate
n.s.
****
BEFORE LIGHT
AFTER LIGHT
TOP 20% DISRUPTED HSE
*
**
CONTROL HSE

Fig. 3.4 | Firing rates of control and target cells during control and inhibited events

Firing rates of control and target environment-encoding cells. Control cells were defined as cells that had higher firing rate in the control than in the target environment, target cells were defined as cells that had higher firing rate in the target than in the control environment. (A-C) Firing rates in all the disrupted HSEs: inhibited cells (A), disinhibited cells (B), and all cells (C). (D-F) Firing rates in the events, classified as confidently (80%) reactivating the target environment: inhibited cells (D), disinhibited cells (E), all cells (F). (G-I) Firing rates in the control HSEs (not disrupted), the ones reactivating a control environment with high confidence: inhibited cells (G), disinhibited cells (H), all cells (I). Two left bars in each figure show the firing rates in the 100 ms window before the light, while the right two bars in each figure show the firing rates in the 100 ms window starting 20 ms after the light trigger. This calculation delay is explained by the time, which was required for the disinhibition to reach its peak. Same control (N = 422 for all cells, N = 88 for inhibited cells and N = 154 for disinhibited cells) and target (N = 492 for all cells, N = 98 for inhibited cells and N = 160 for disinhibited cells) environment-specific cells have been used to compare rates in all three categories of HSE. Firing rates of cells were normalised (Z-scored) according to their baseline mean and SD.

To summarize the results represented in the **Fig. 3.4**, the firing rates of control and target environment-encoding cells during and after the assembly detection windows demonstrated that the real-time algorithm was able to correctly decode reactivation of control or target environment and predicted the persisting reactivation of the control environment. The light pulses disrupted the reactivation of the target environment activity, while reactivation of the control environment activity persisted.

Below I present three additional measures that support the conclusions of the correct real-time decoding as well as disrupted consolidation of the target environment. The categories of control / target cells in these figures comprise the same cells as in the **Fig. 3.4.** First, **Fig. 3.5** demonstrates the correlations of mean firing rates of control and target cells in the disrupted and control events 100 ms before the assembly detection and 100 ms after the assembly detection. The rates were highly correlated in the control events while correlations were much weaker in the inhibited events. The correlation of mean rates of the control environment-encoding cells in the control HSEs was higher than the correlations of mean rates in the disrupted HSEs for both control ($P < 6*10^{-17}$, z-test) and for target ($P < 3*10^{-11}$) environment-encoding cells as well as the correlation of mean rates of the target environment-encoding cells in the control HSEs ($P < 0.04$). Secondly, I have analysed the cofiring of control and target cell pairs in the 100 ms before the assembly detection and 100 ms after the assembly detection in the control and in the disrupted events conditioned on the pre-sleep cofiring correlations (**Fig. 3.6**). Cofiring as before was measured as the Pearson

correlation of the instantaneous firing rates (i.e. spike counts). The cofiring correlations of control cell pairs in the control events were significantly higher than the cofiring correlations of target cell pairs in the control events ($P < 7*10^{-16}$, z-test) as well as than the cofiring correlations of target cells pairs in the target events ($P < 7*10^{-16}$, z-test). The high correlations of the cofiring of target cells in the control events in **Fig.3.4** and **Fig 3.3** indicate according to the firing rate analysis that they cell pairs simultaneously suppress their firing in this case.

Finally, when quantifying reactivation, a commonly used measure is a **cofiring correlation** of waking and sleep firing patterns (Wilson & McNaughton, 1994). **Fig. 3.7** represents the correlations of cofiring in control and target environments and cofiring in control and disrupted HSEs calculated separately for control and target cells. The correlations of cofiring of control environment-encoding cell pairs measured during learning in the control environment and compared with that in the 100 ms time window before and after the assembly detection were equally high and similar ($P > 0.99$, Z-test). The correlations of cofiring of target-encoding cell pairs in the target environment with their cofiring in the 100 ms window after the assembly detection was lower during the light-on period in the disrupted HSEs than before the assembly detection ($P < 4*10^{-12}$, z-test). Therefore, reactivation of target environment is disrupted, while the reactivation of the control environment remained intact.

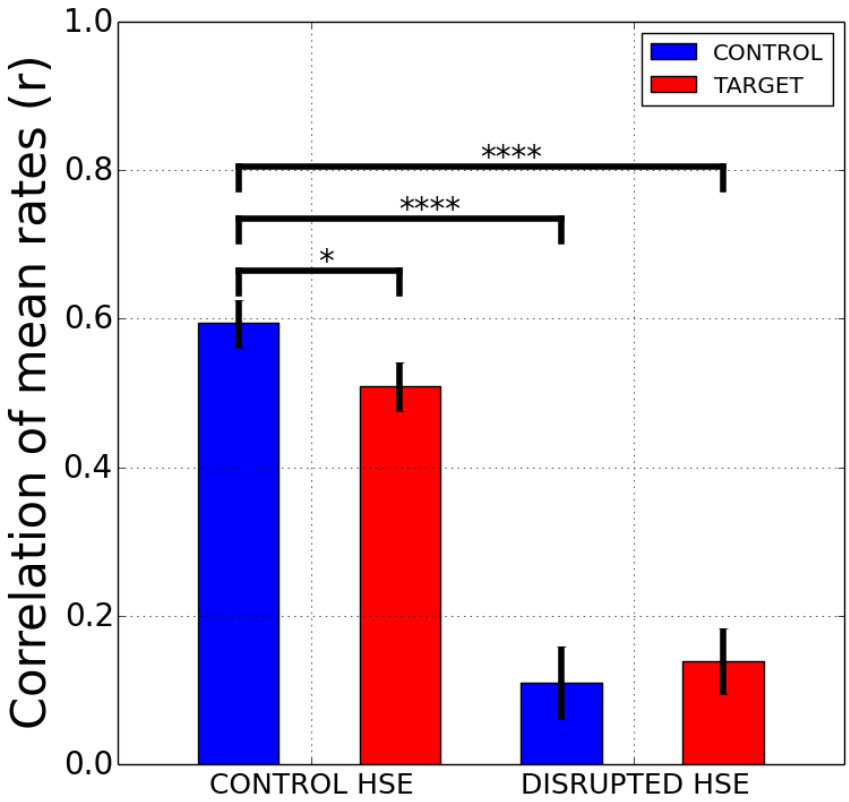

Fig. 3.5 | Correlation coefficient of mean firing rate of control and target cells in the control and inhibited events

Mean firing rates were calculated separately for control (N = 422) and target (N = 492) environment-encoding cells in the 100 ms time windows before the assembly detection and 100 ms after the assembly detection for and averaged across sessions.

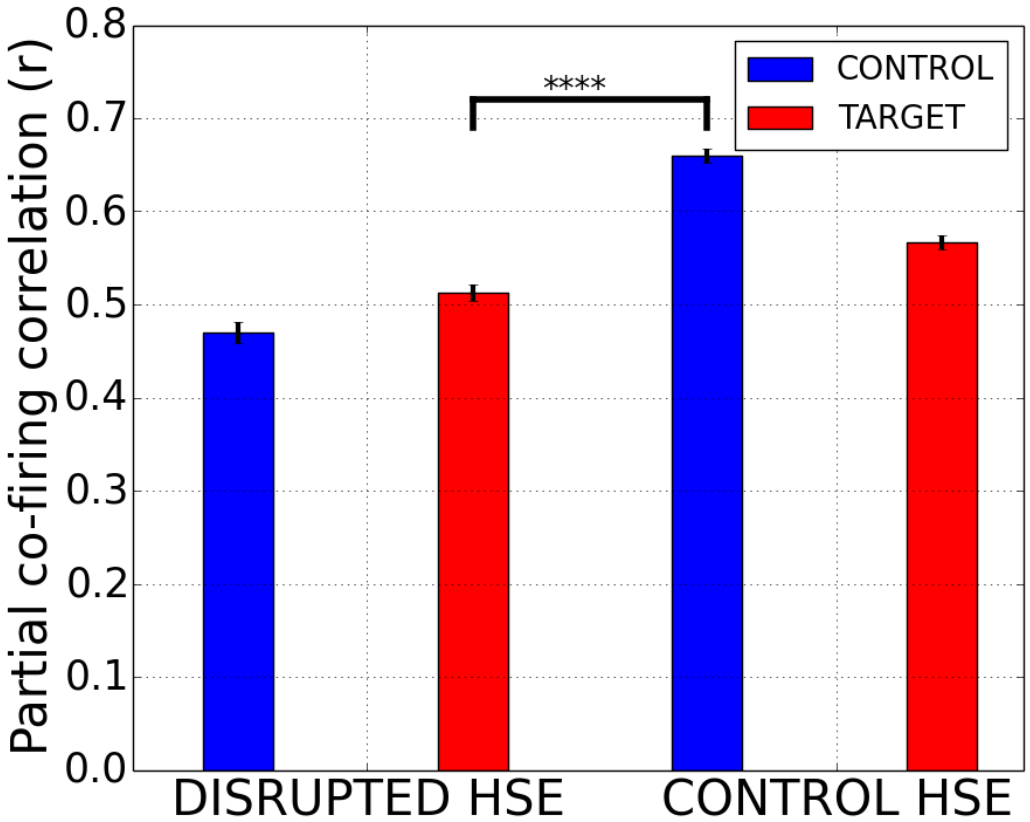

Fig. 3.6 | Partial correlations of cofiring of control and target cells in the control HSEs and disrupted HSEs.

Partial correlations of cofirings of control (N = 4953 cell pairs) and target (N = 7603 cell pairs) environment-encoding cells in 100 ms windows before and after assembly detections controlled for correlations at the beginning of learning. These were calculated for control HSEs and disrupted HSEs that confidently classified as reactivating the target environment.

The results presented in this section suggests that the real-time decoder was able to classify the environments that were being reactivated in the detected HSEs and that the application of light preferentially disrupted the reactivation of the target environment, while the reactivation of the control environment remained intact in the HSEs for the confidently encoded events.

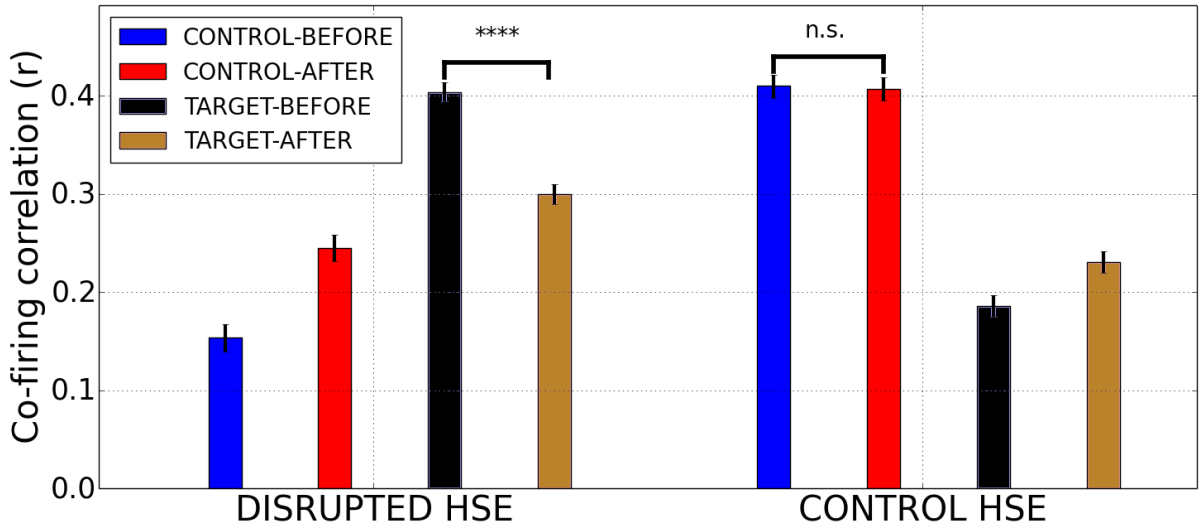

Fig. 3.7 | Cofiring correlations of the target environment-specific and control-environment specific cells between the end of learning and the activity before and after HSE detection in control and disrupted events

A reactivation of end of learning waking pattern was calculated separately for control (N = 4953 pairs) and target (N = 7603 pairs) cells pairs in the 100 ms windows before and after the assembly detection, separately for control HSEs and disrupted HSEs, confidently classified as reactivating the target environment.

3.3 Assessment of behavioural performance on the cheeseboard maze.

Behavioural trajectories of the rats (N = 4) were acquired while they performed a spatial memory task using the double cheeseboard paradigm (number of sessions N = 21). Animals had to remember a single food location in each of the two cheeseboard environments, and these locations changed daily so that each day the animal had to learn new goal locations on the cheeseboards (see Methods). Therefore, in the initial learning trials, animals had to learn these locations. An asymptotic performance was normally reached on the 2nd or 3rd trial of the initial learning trials in each of the environments and an asymptotic trial trajectory would normally constitute a straight line from the start box to the goal and back (**Fig. 3.8 B**). After four hours of sleep, during which all the HSEs that did not encode the control environment with high confidence were disrupted, animals were allowed to spend 3 minutes in each of the environments during the post-probe session **(Fig. 3.8 C)**, during which the goal location was not baited. The order of presentation of the environments was the same as during the learning sessions. After the post-probe sessions, the post-learning

sessions were recorded, in which the food was provided at the same location. This provided an additional behavioural performance measure to check if animal went straight to the goal in the first trial of this session as well as it allowed to obtain rate maps of the cells in during more stereotyped behaviour **(Fig. 3.8 D)**.

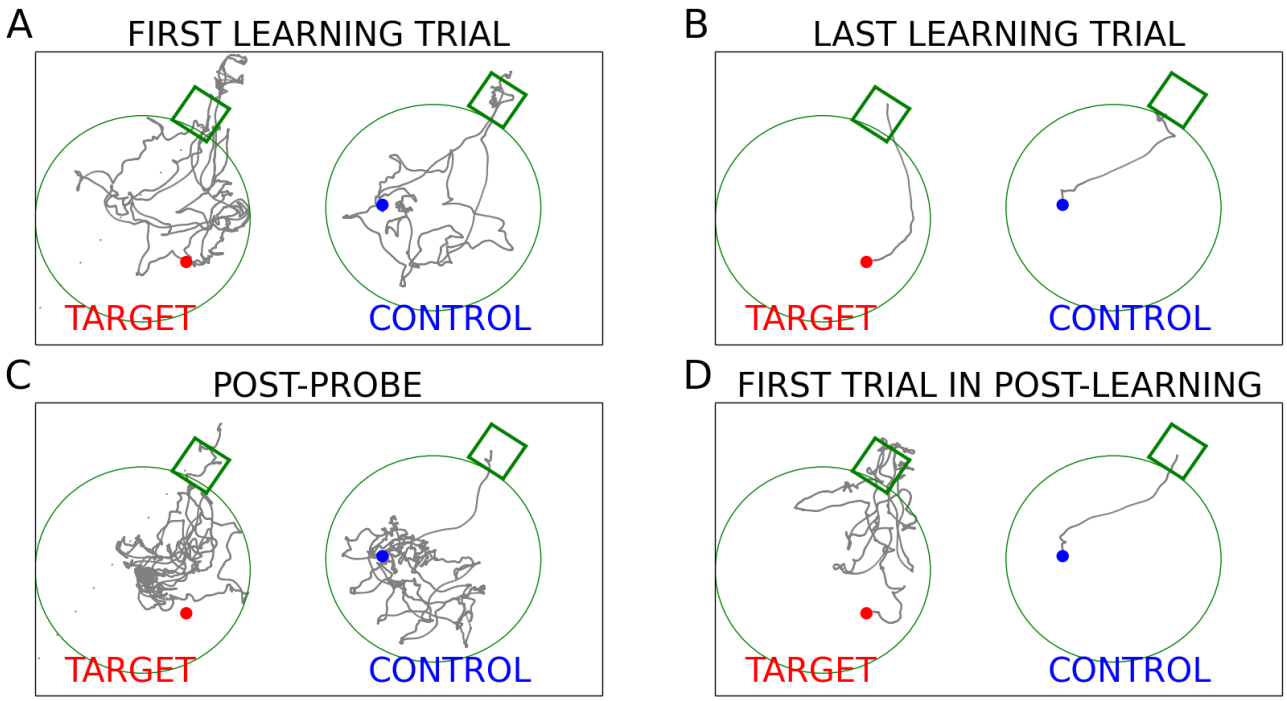

Fig. 3.8 | Behavioural trajectories in different stages of the paradigm

Example of tracked trajectories during different stages of learning and during the post-probe. The target environment is on the left, control - on the right. Goal locations are designated by the dots. Tracking during learning is shown only until the time when the goal was reached.

In each recording day, the behavioural performance of the animals was recorded, and the following behavioural performance measures were analysed in both control and target environments:

1. **Number of crossings** of the goal area (area within 10cm from the centre of the rewarded well)
2. **Dwell time –** proportion of post-probe time spent in the goal area
3. **Excess first trial path length** to reach the goal in the post-learning session (difference between the length of the actual trajectory taken by the animal to reach the goal and the shortest possible path, i.e. a straight line between start box centre and goal location)

The measures are illustrated in the **Fig. 3.9.**

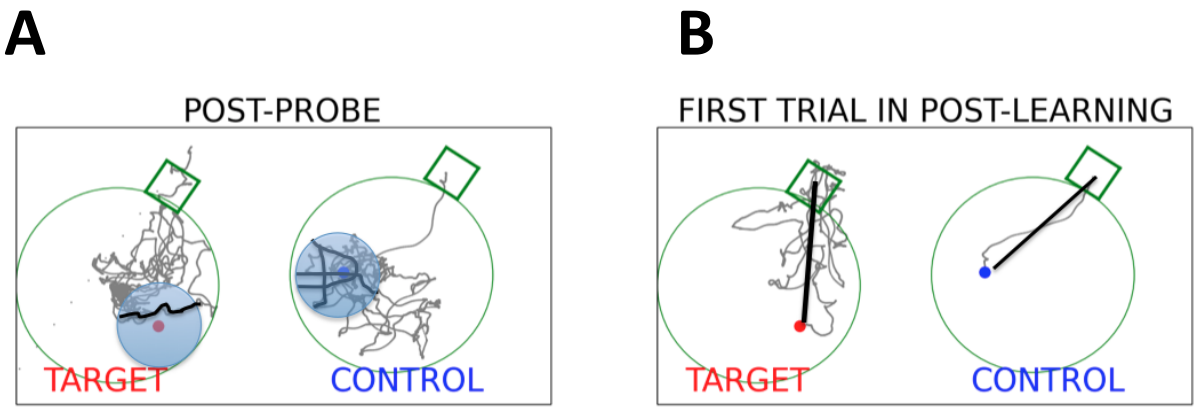

Fig. 3.9 | Illustration of the behavioural measures

(A) Illustration of how dwell time and a number of crossing measures were calculated. Blue semi-transparent circles centred at the goal location represent the goal area of a 10 cm radius circle centred around the goal well. This was used for quantification of goal area crossings (bold segments of trajectories) and total time spent in the goal area.

(B) Difference between actual path and direct path representing the shortest distance to the goal. The black lines represent the shortest possible trajectory to reach the goal, that is, a straight line between the edge of the start box and location of the goal.

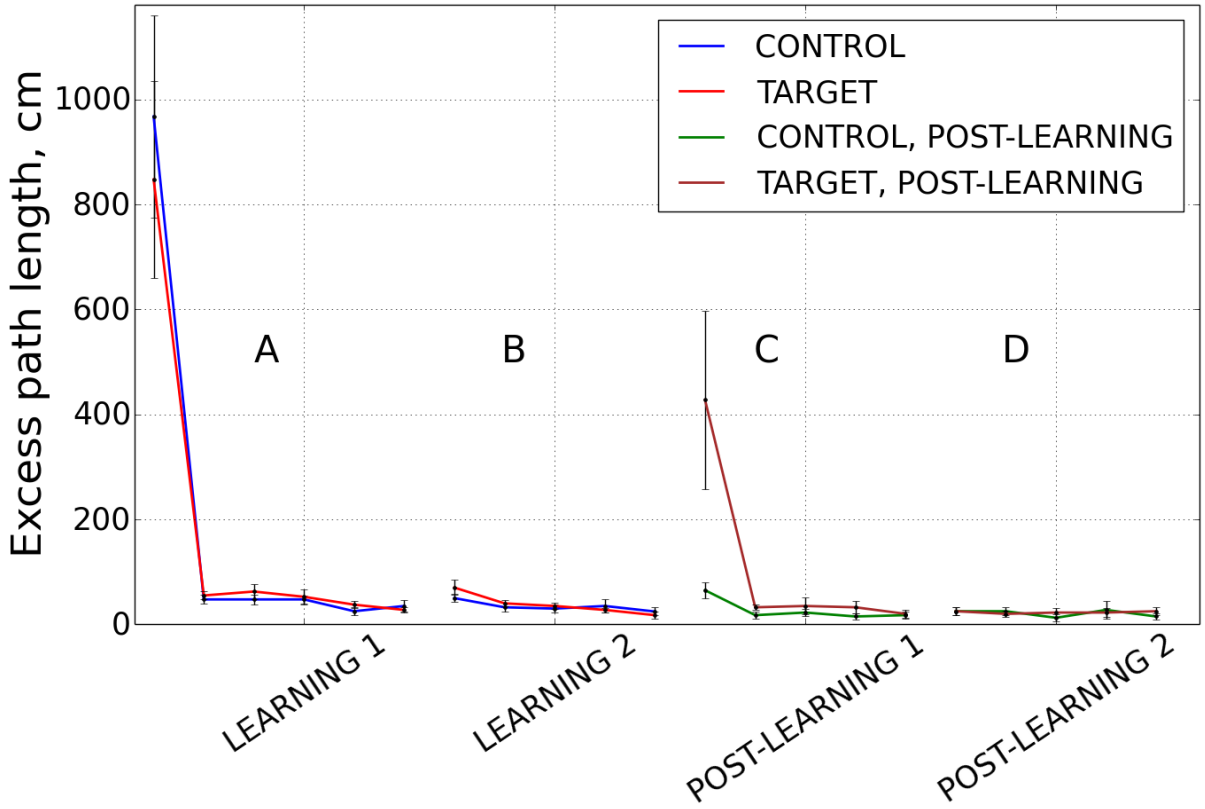

Fig. 3.10 | Average learning curves.

The excess path length measured in each trial (mean ±SEM) was used as a behavioural measure for the learning curves. (A) Learning curve of the first block of learning trials. (B) Learning curve of the second block of the learning trials. (C) Learning curve of the first block of post-learning trials. (D) Learning curve of the second block of the post-learning trials

To understand the dynamics of learning in the control and target environments, I have analysed the learning curves that show the average excess path length for every trial, averaged across all learning sessions. These learning curves are shown in **Fig. 3.10** and are grouped into four groups according to the experimental paradigm: the first block of learning trials **(Fig. 3.10 A)**, the second block of learning trials **(Fig 3.10. B)**, the first block of post-learning trials **(Fig. 3.10 C)** and the second block of post-learning trials **(Fig. 3.10 D)**. This analysis indicated that an asymptotic performance was reached already on the second trial and that there were no differences in the performances in the control and target environment except for the first trial in the post-learning sessions. Also, performance in the target environment in the post-learning sessions was not on the level of the very first trial, but rather between the performance of the very first trial and an asymptotic performance. To further investigate this deficit, I performed an animal-by-animal and trial-by-trial analysis of the performance measures.

To compare performance in control and target environment and account for the day- or animal-specific changes in the behaviour in both control and target environments, the behavioural score was calculated for every recording day based on performance measure values in the control (PM_C) and the target (PM_T) environments as

$$BS = \frac{PM_C - PM_T}{PM_C + PM_T} \quad (13)$$

21 sessions were recorded in total with 6, 4, 5, and 6 sessions per animal. N = 21 sessions were used in the statistical analyses if not specified otherwise. Out of 21 recorded sessions, the animals performed better in control than in the target environments in all measures: in 18 days according to the dwell time **(Fig. 3.11 A)**, in 16 days according to normalized path length **(Fig 3.11 B)** and in 15 days according to the number of crossings with an equal performance in two days **(Fig. 3.11 C)**. The mean (±SEM) values of these measures are shown in **Fig. 3.11**, while the same for each animal as well as the individual values of performance measures in control and target environments can also be seen in **Fig 3.12.**

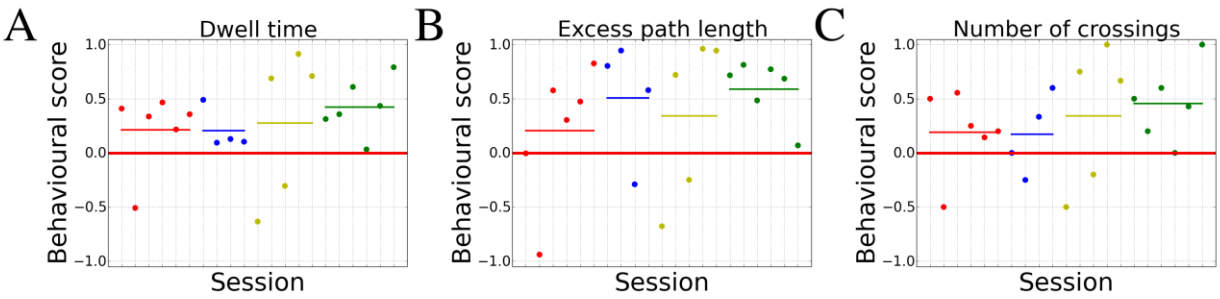

Fig. 3.11 | Behavioural scores of the performance measures.

Positive score value means that performance in the control environment was better than in the target, and vice versa for the negative score value. Each colour represents a different animal. Horizontal lines of a corresponding colour represent the mean score for an animal. (A) Behavioural scores of the dwell time measures. (B) Behavioural scores of the excess first trial path length measure. (C) Behavioural scores of the number of crossings measure.

The median values of all three behaviour measures for control and target environments were statistically different ($P < 0.0001$ for dwell time, $P < 0.00005$ for a number of crossings and $P < 0.02$ for excess path length, Mann-Whitney U test, **Fig 3.12)**.

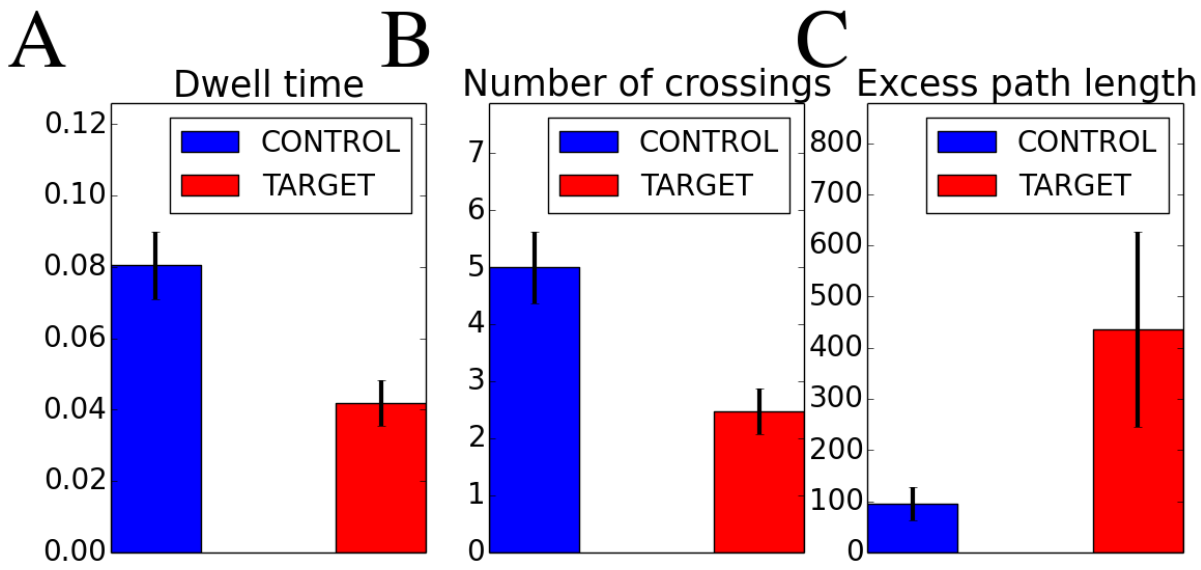

Fig. 3.12 | Behavioural measures in control and target environment

(A) Dwell time – proportion of time spent in the 10 cm area around the goal. (B) Number of crossings of the 10 cm area around the goal. (C) Normalized path length. N = 21 sessions for both control and target environments.

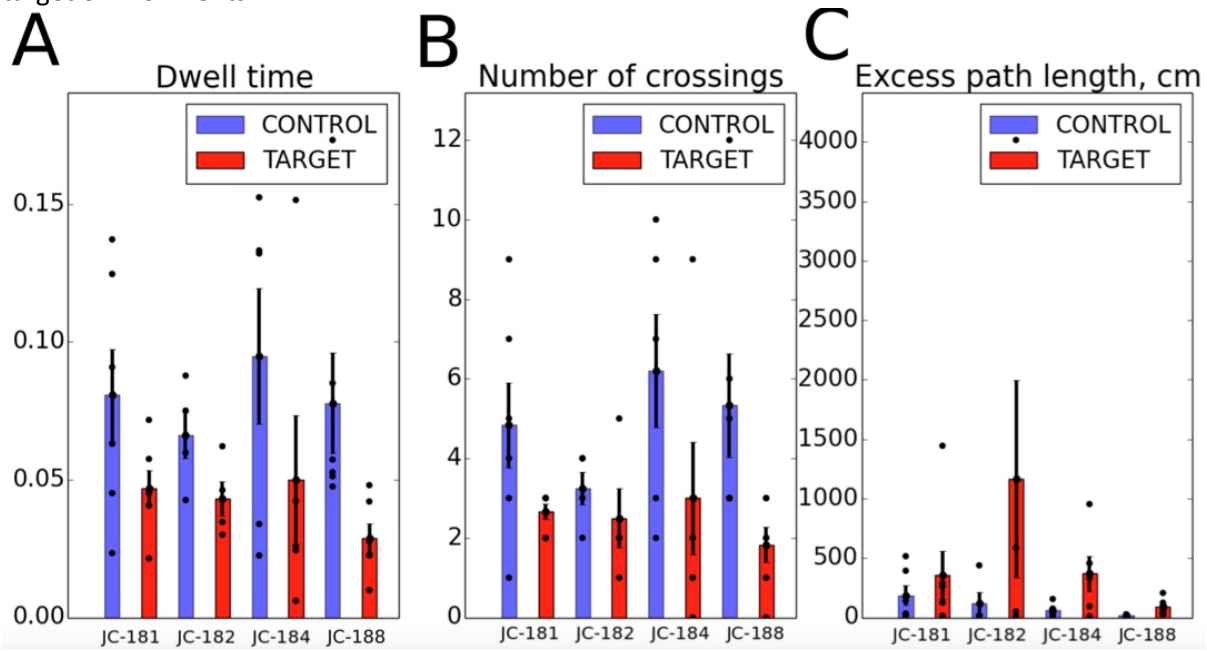

Fig. 3.13 | Animal-wise means of behavioural measures

Individual dots represent session-wise values. (A) Dwell time. (B) Number of crossings. (C) Normalized path length.

Two-way ANOVA analysis showed that there was no significant difference between different animals **(Fig. 3.13)** in any of the performance measures ($P > 0.68$ for dwell time, $P > 0.56$ for number of crossings, $P > 0.29$ for first trial path length, ANOVA) and that there was

significant difference between performance in the target and control environments for all performance measures ($P < 0.004$ for dwell time, $P < 0.004$ for number of crossings, $P < 0.03$ for first trial path length, ANOVA), while no significant interactions were detected between animals and environments ($P > 0.89$ for the dwell time, $P > 0.65$ for the number of crossings, $P > 0.73$ for the excess path length). In addition, all behavioural scores were significantly larger than 0 ($P < 0.002$ for dwell time, $P < 0.0005$ for number of crossings, $P < 0.0006$ for first trial path length, bootstrapping with n =10000 iterations). As the results of all animals were analysed together, it is possible that remapping of place fields of the suppressed cells show significant effect because of the results of a single animal. Therefore, I repeated the bootstrapping analysis by leaving out an animal from the data set and performed the same analysis on the remaining animals. After repeating the analysis by leaving out each animal in turn, I took the least significant (that is, highest) P-values from these calculations. Even in this case, the behaviour scores were significantly larger than 0 for all performance measures (max. $P < 0.03$ for the dwell time, max. $P < 0.03$ for the number of crossings, max. $P < 0.03$ for the excess first trial path length, bootstrapping with n=10000 iterations). A line graph of behavioural scores allows grasping variance in performance across animals and across days (**Fig. 3.14**).

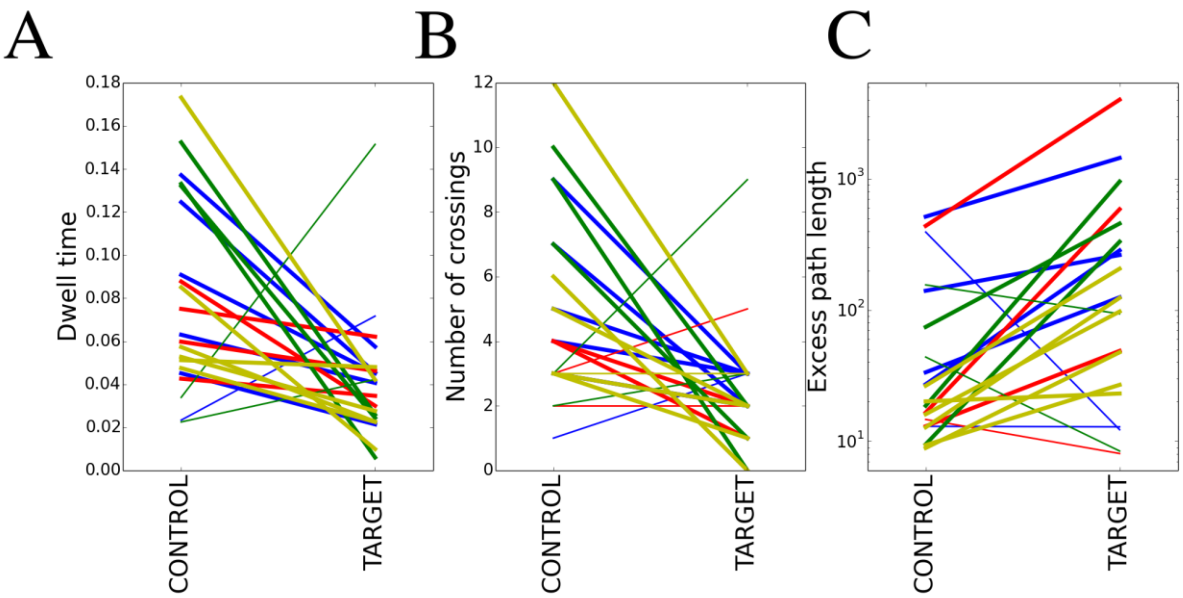

Fig. 3.14 | Line graph of behavioural measures

Each line represents a single recording day, bold lines in the days when the performance in a control environment was better, thin lines - when the performance in the target environment was better.

These results show that my manipulation, which disrupted reactivation without impacting the HSEs that encoded confidently the control environment, led to the behavioural impairment in the post-probe and post-learning sessions in the target but not in the control environments. On the other side, the memory deficit was weak in some instances and in a minority of cases memory performance was stronger in the target environment. Therefore, in the following chapters, I examined the effect of my procedure on the firing of the recorded cells and on the stability of the hippocampal maps.

3.4 Influence of optogenetic light application mediated rate changes on behavioural performance

First I have tested to what degree the behaviour impairment was related to disruption of reactivation by analysing the correlation between behavioural performance and correlation of firing rates before and after assembly detection of control and target environment-specific cells in control and disrupted HSEs correspondingly. I found that these correlations were small but significant for dwell time ($R = 0.123$, $P < 0.003$, **Fig. 3.15 A**) and number of crossings ($R = 0.117$, $P < 0.004$, **Fig 3.15 C**) but not for first trial excess path length ($R = -0.038$, $P > 0.32$, **Fig. 3.15 B**). The weak correlations could be explained by outliers in the behavioural measures, impairment in which was not achieved despite the present physiological effect on the recorded cells.

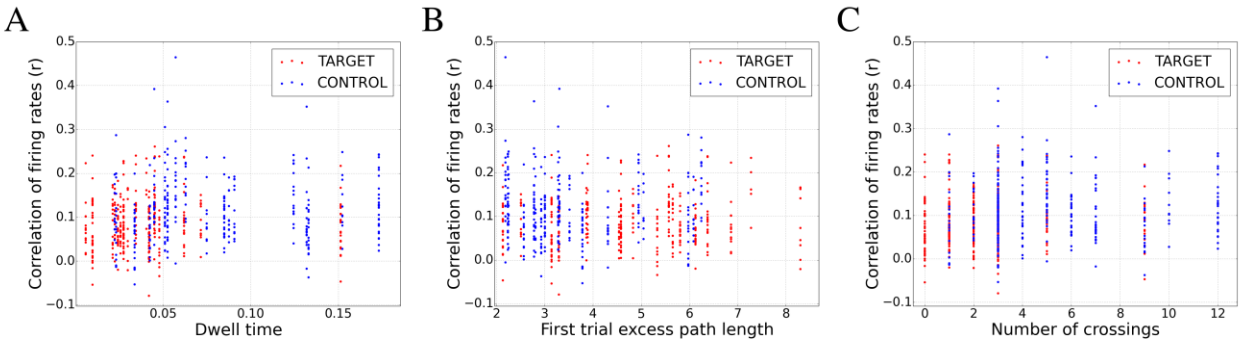

Fig. 3.15 | Correlation of behavioural performance measures and firing rate correlations before and after assembly detection.

Correlations of firing rates of control (N = 422) and target (N = 492) environment-specific cells were calculated in for rates in the 100 ms before and after assembly detection in the control and disrupted HSEs, confidently classified as reactivating the target environment correspondingly. These correlations were then correlated to the three performance measures: dwell time (A), first trial excess path length (B) and number of crossings (C).

3.5 Influence of the reactivation disruption on place map stability.

It was shown earlier that learning new goal locations in a familiar cheeseboard environments causes reorganization of place maps (Dupret et al., 2010) and that optogenetic blockade of SWR leads to disruption of reinstatement of assembly patterns (van de Ven et al., 2016) but not the place fields (Kovács et al., 2016). In addition, blockade of SWR during learning in the cheeseboard maze prevents stabilization and refinement of the hippocampal map (Roux, Hu, Eichler, Stark, & Buzsáki, 2017). Therefore, I examined whether my manipulations selectively prevented the stabilisation of the map of the target environment.

Fig. 3.16 illustrates representative examples of rate maps of four cells in four sessions: end of learning, post-probe, start of post-learning, end of the post-learning. In these examples, the rate map in the target environments was destabilized in the post-probe and start of post-learning sessions and reverts back to the old map at the end of the post-learning session, while the rate map of the control environment was stable across all sessions.

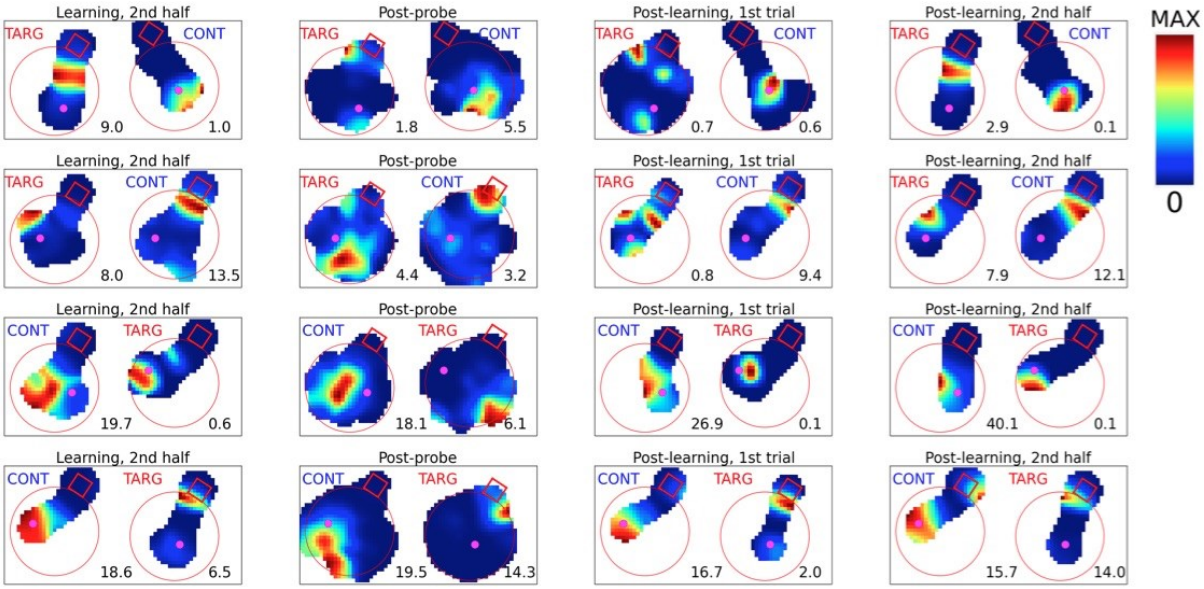

Fig. 3.16 | Rate maps of four cells in four sessions in control and target environments.

The target environment is on the left, control is on the right. The red circle represents the edges of the cheeseboard, the red rectangle represents the start box. Purple dots represent the goal location.

To investigate how the selective disruption of reactivation affected the place coding properties of the cells, I have analysed place field stability, goal-related remapping and the stability of the firing correlations (cofiring) patterns of place cells in both control and target environments. As I aimed to trigger the light in all synchrony events except those in which the control environment was reactivated with high confidence, cells which fire more in the target environment were expected to be affected more by such manipulation. Indeed, I have observed weaker place field stability in the target than for the control environments. This result is demonstrated in the **Fig. 3.17 A-B,** where the similarity of place fields was compared between post-probe and end of learning sessions (A) and start of post-learning and the end of learning sessions (B). A place field similarity (PFS) measure is a correlation of the two rate maps and demonstrates, how similar the spatial firing of cells is in the two

conditions, e.g. two exploration or learning sessions. I have calculated PFS for the cells in the control as well as in the target environment that had valid place fields in control and target environments (i.e. passed the minimal coherence of 0.3 (Skaggs & McNaughton, 1996) and maximum sparsity of 0.5 (Muller & Kubie, 1989) criteria). I have included only cells that have exhibited remapping between the first and second half of the learning (PFS < 0.5) to exclude stable cells that do not encode goal-related memory traces. Also, I have limited the mean firing rate for the cells to the range from 0.5 to 5 Hz, because lower firing rates cells do not produce place fields reliable enough for the comparison, and cells with higher firing rate may have been putative interneurons. Also, I have completed the comparisons for all cells combined as well as separately for inhibited and disinhibited cells, as defined in the previous section. I have performed the post-hoc multiple testing correction using the Holm-Bonferroni method. As the results show, there was a decrease in PFS for the rate maps of the target environment compared to the control environment for all groups of cells in the post-probe vs. end of learning comparisons ($P < 0.005$ for disinhibited, $P < 0.02$ for inhibited and $P < 0.006$ for all cells, Mann-Whitney U Test, Holm-Bonferroni correction for multiple testing; **Fig 3.17 A**) as well as beginning of post-learning (i.e. first trial of post-learning) vs. end of learning comparisons ($P < 0.04$ for disinhibited, $P < 0.04$ for inhibited and $P < 0.03$ for all cells, Mann-Whitney U Test, Holm-Bonferroni correction for multiple testing; **Fig 3.17 B**). Surprisingly there was no difference in the PFS of the end of learning and end of post-learning maps (i.e. last 5 trials of post learning) of the two environments ($P>0.38$, Mann-Whitney U Test; **Fig 3.17 C**). This suggests that hippocampal maps of the target environment, originally disrupted by selective disruption of reactivation, were reinstated towards the end of learning as the animal relearned the goal locations.

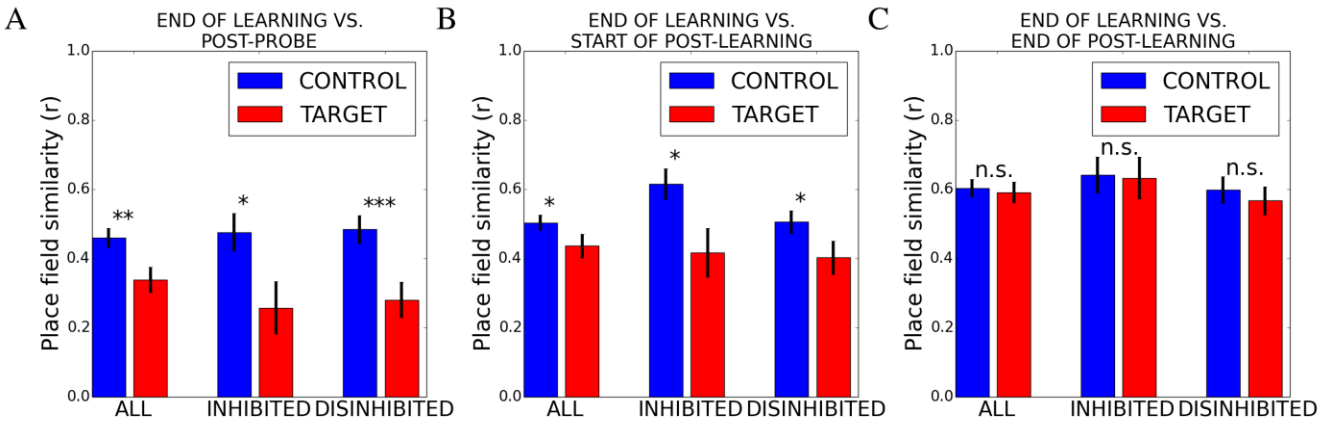

Fig. 3.17 | Place fields similarities of the control and target environment rate maps.

(A) End of learning vs. post-probe. (B) End of learning vs. start of post-learning. (C) End of learning vs. end of post-learning. Same cells were used to calculate the place field similarities sessions: N = 119 for control, N = 145 for target environment in all cells group; N = 28 for control, N = 37 for target in inhibited cells group; N = 61 for control, n=69 for target environment in the disinhibited cell group.

In addition to the mean, I also compared the distribution of the PFS for all cells. These were significantly different when comparing post-probe rate maps ($P < 0.02$, Kolmogorov-Smirnoff test; **Fig 3.18 A**) and beginning of post-learning rates maps ($P < 0.03$, Kolmogorov-Smirnoff test; **Fig. 3.18 B**) with end of learning rate maps, but not when comparing the end of learning with the end of post-learning ($P > 0.92$, Kolmogorov-Smirnoff test; **Fig 3.18 C**)

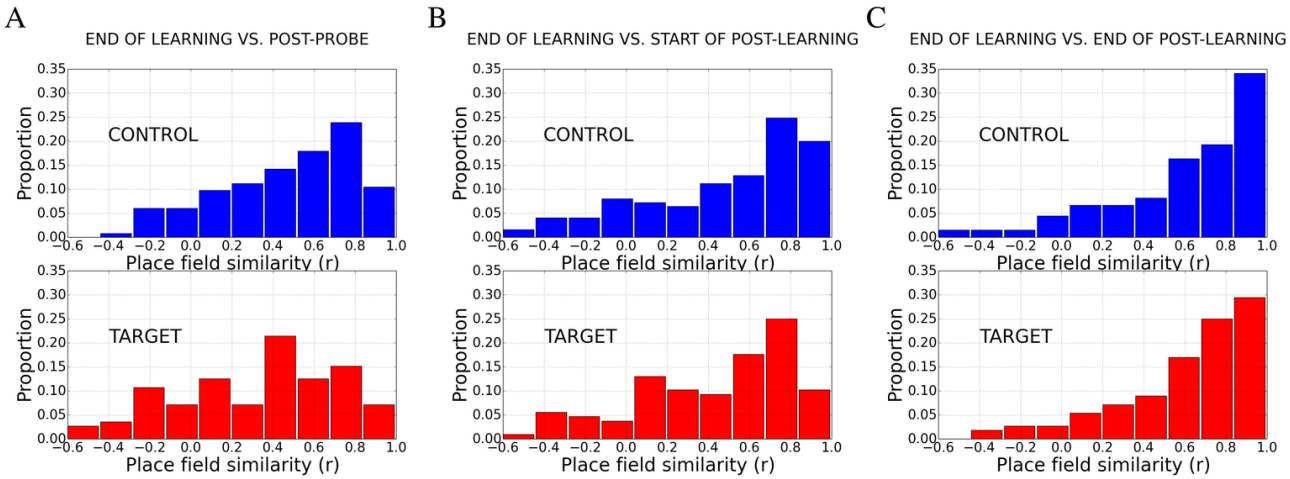

Fig. 3.18 | Distributions of the PFS

(A) End of learning vs. post-probe. (B) End of learning vs. start of new learning. (C) End of learning vs. end of post-learning. Same cells have been used to analyse place field similarities for all combinations of sessions (N = 119 for control, N = 145 for target environment).

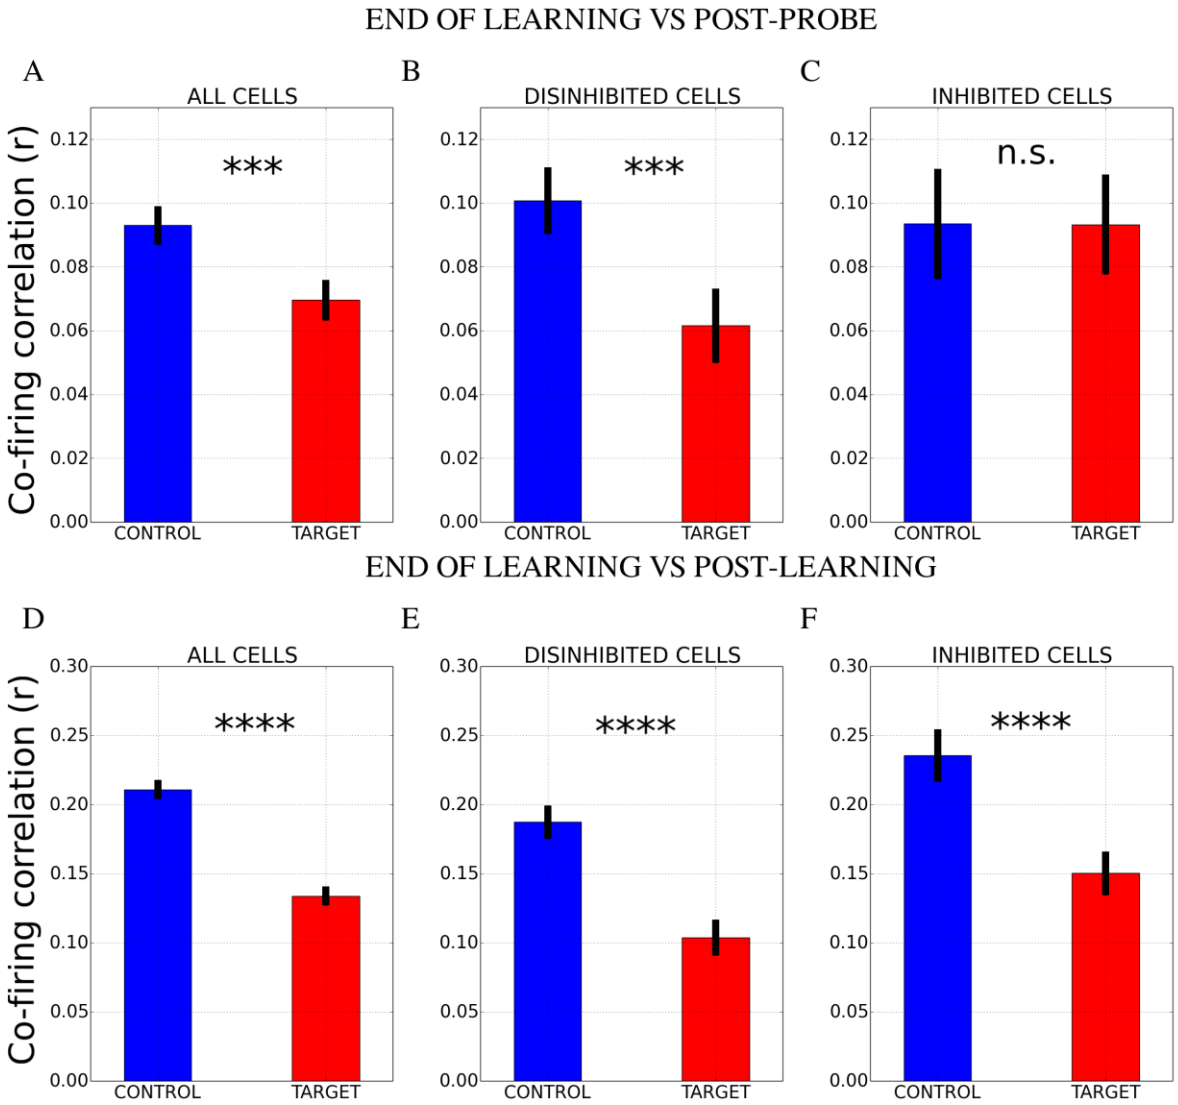

Fig. 3.19 | Cofiring correlations of post probe / end of learning and start of post-learning

End of learning vs. post-probe: (A) all cells, N = 26788 in control / 24785 in target, (B) disinhibited cells, N = 8998 in control / 7294 in target. (C) inhibited cells N = 3313 in control / 4019 in target. End of learning vs. start of post-learning: (D) all cells N = 19629 in control / 20833 in target, (E) disinhibited cells, N = 6355 in control / 5684 in target, (F) inhibited cells, N = 2528 in control / 3832 in target. Only cells that specifically fired stringer in the tested environment were used but these were correlated with all other non-specific cells.

In order to assess to what degree the destabilisation of place fields in the post-probe and at the beginning of the post-learning led the different place-related assembly patterns, I calculated the cofiring patterns of the recording cells by measuring the cofiring of all cell pairs and testing whether this cofiring matrix of all possible cell pairs remained similar after the assembly-related light disruption. Cofiring was measured for every cell pair as the correlation of numbers of spikes emitted by a pair of cells in relevant 100 ms time windows.

I included the same sets of cells for which the PFS analysis was performed, and the similarity of a cofiring matrix of the cell pairs in different condition was assessed by measuring how strongly these patterns were correlated. The cofiring patterns were less similar for target environment-specific cells than for control environment-specific cells in comparison of post-probe to the end of learning ($P < 0.007$ for disinhibited, $P < 0.004$ for all, z-test; **Fig. 3.19 A, B**) as well as comparison of the start of post-learning to the end of learning ($P < 2*^{10-6}$ for disinhibited, $P < 2*^{10-16}$ for all cells, z-test; **Fig. 3.19 D, E**). An exception here was the cofiring of the inhibited cells in the post-learning session vs. end of learning ($P > 0.49$, z-test; **Fig. 3.19 C**) but between the end of learning and post-probe it was significantly different ($P < 0.0004$, z-test; **Fig. 3.19 F**).

3.6 Influence of the reactivation disruption on place map goal remapping

Finally, I compared to what degree place cells encoding the control and target environment moved their place fields towards the goal. As previous work found that change of the goal locations leads to the goal remapping of place cells in which many of them move their place fields closer to the goal (Dupret et al., 2010). Therefore, the selective disruption of reactivation is expected to cause a selective deficit in the goal remapping in the target environment place cells compared to the control environment place cells. To quantify goal remapping, I calculated the distance from the place field peak to the goal location. **Fig. 3.20** quantifies how distance from the peak of the centre of the place fields changed from the beginning of learning to the start of post-learning for place cells that were active either in control and target environment. Again, the cells used for these analyses were the same as used for the PFS analyses. While there was no significant difference in the mean goal distance of place field peak at the beginning of learning when control and target environment place fields were compared ($P > 0.38$, Mann-Whitney U test; **Fig. 3.20**), at the beginning of post-learning the place fields of the cells in the control environment were closer to the goal than in the target environment ($P < 0.002$, Mann-Whitney U test; **Fig. 3.20**). This effect was further confirmed by comparing the distribution of goal-directed place field shifts from the start of learning to the start of the post-learning. This was calculated by subtracting the distance from the peak of the place field in the start of the learning session

from the distance from the peak of the place field in the start of post-learning session. This shows that place fields of 67% of control cells shifted closer to the goal while numbers of cells in the target environment that move their place fields towards and away from the goal were approximately equal. The two distributions were significantly different ($P < 0.0002$, Kolmogorov-Smirnoff test; **Fig. 3.21**). A relatively mild effect on the goal remapping could be explained by the fact the only a proportion of CA1 cells exhibits a goal remapping as well as by possible shift of place field peaks towards to goal due to stereotyped behaviour and not the goal remapping.

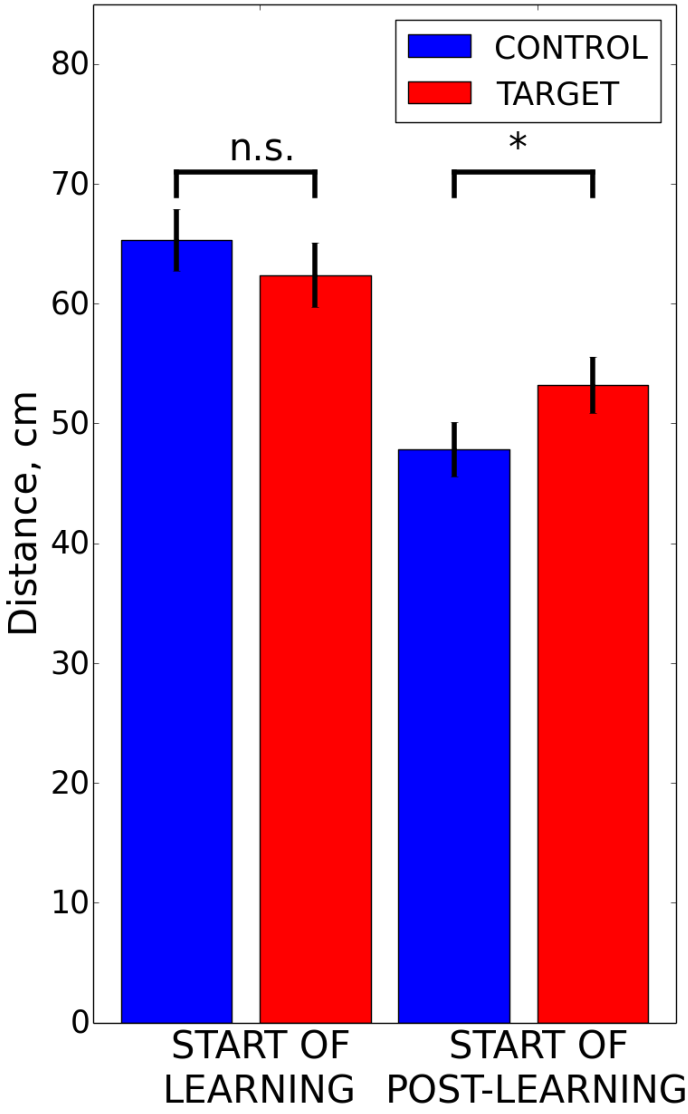

Fig. 3.20 | Distances from place field centre to the goal location

Place fields at the beginning of learning and beginning of post-learning were used for comparison. Only cells with valid place fields, i.e. with min. coherence of 0.5 and maximum sparsity of 0.3 in both sessions were used (N = 146 for control, N = 161 for target environment).

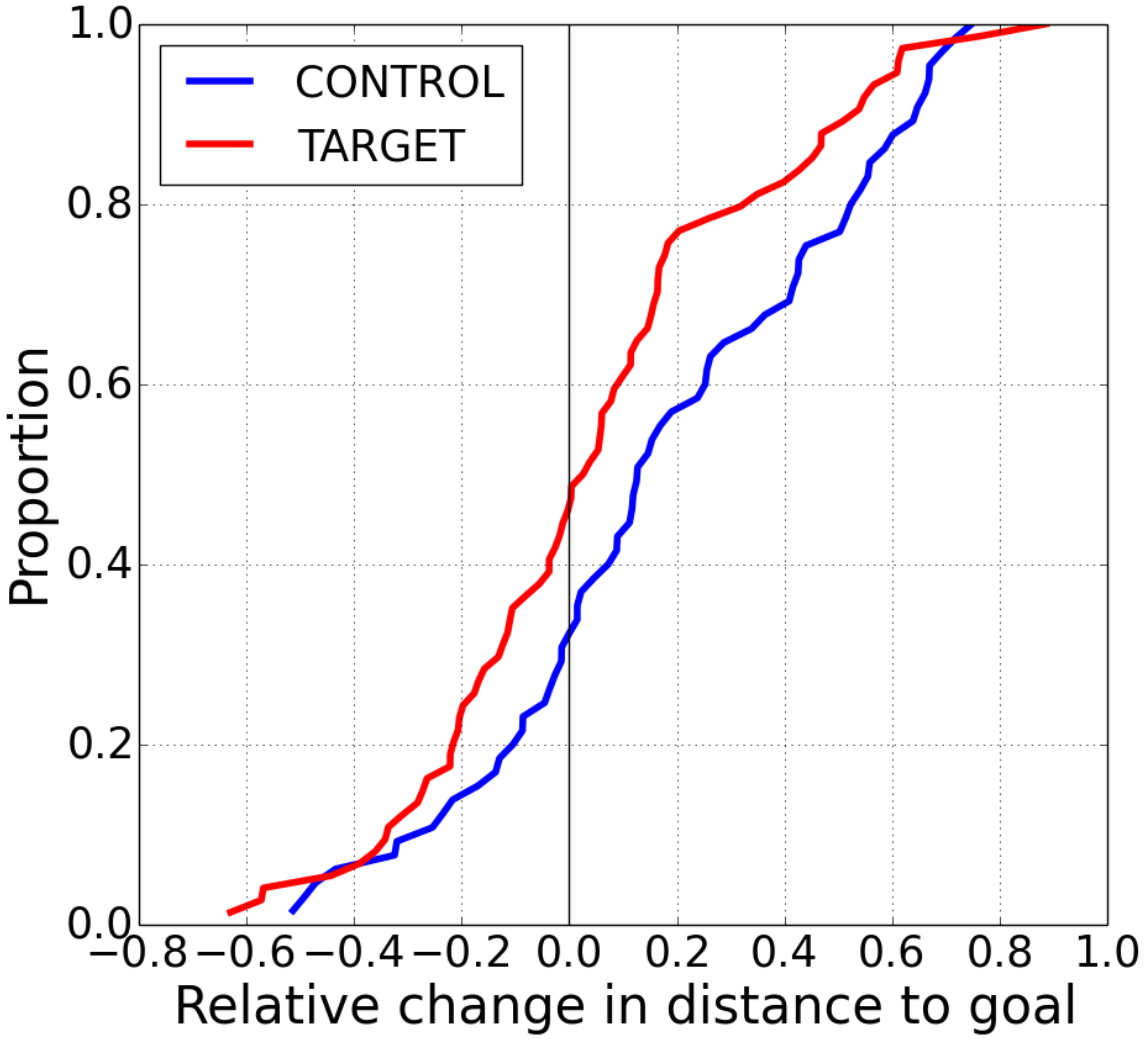

Fig. 3.21 | Cumulative Distribution of place field shifts towards the goal

Differences between the distance from the peak of the place field in the start of the learning session and the start of the post-learning session were normalized by the distance from place field peak to the goal in the start of the learning session.

Variability in the behavioural impairment could be partially explained by the amount to which the maps were selectively destabilized as a result of disrupting reactivation. Indeed, I have analysed correlations between place field similarities of the end of learning and post-probe session in control and target environments and behavioural performance measures described in the previous section: dwell time, number of crossings, first trial path length (**Fig. 3.2**). I have found these correlations to be small but significant ($R = 0.1886$ $P < 0.02$ for dwell time; $R = 0.153$, $P < 0.05$ for number of crossings; $R = -0.178$, $P < 0.02$ for first trial path

length; **Fig. 3.22**). Therefore, the stronger destabilization of place fields was associated with a stronger behavioural deficit.

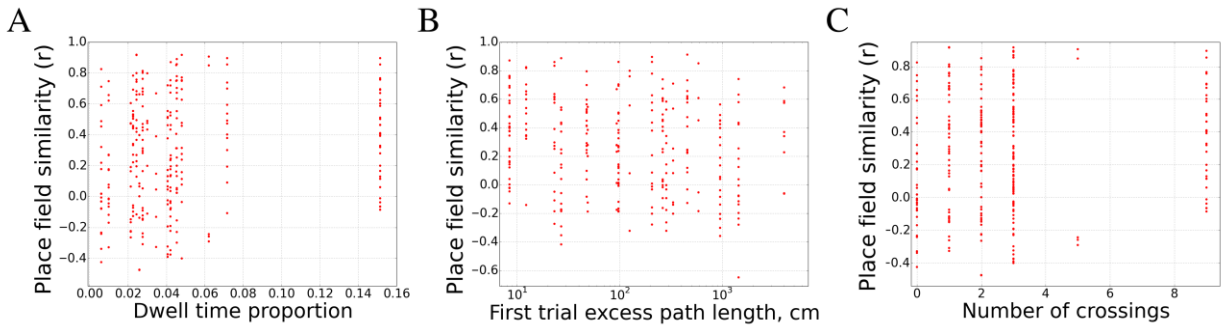

Fig. 3.22 | Correlation between behavioural measures and PFS between the end of learning and post-probe

Every dot represents a single cell. (A) Dwell time versus place field similarities, N = 398; (B) First trial path length versus place field similarities, N = 398; (C) Number of crossings versus place field similarities, N = 398.

4 DISCUSSION

In my thesis, I examined in the hippocampus the role of assembly-specific reactivation in the consolidation of associated memory traces. I have developed a method for the real-time decoding of spike patterns within HSEs and applied it for the classification of cheeseboard environments that were reactivated during the HSEs. In doing so, I preferentially disrupted the reactivation of one of the two cheeseboard environments by means of optogenetics, while the animal was resting and sleeping after learning to collect the food at fixed locations in these two environments. To confirm that the method preferentially affected the spike patterns associated with the target but not the control cheeseboard environments, I have compared post-hoc the unit activity before and after the detection of the HSE content. This analysis indeed demonstrated that my manipulation disrupted more those assemblies that replayed the target environment- than those replaying the control environment. I have then analysed the effect of this assembly-related disruption on the spatial memory performance of the animal and stability of the place maps related to these spatial memory traces. First, I found that my manipulation was able to induce a selective behavioural deficit, i.e. impaired memory about the food location in the target environment, the reactivation of which was preferentially disrupted, but not in the control environment. Second, I have found that the stability of place maps representing the target environment was more disrupted than those of the control environment maps when maps before the optogenetic disruption were compared to those in the post-probe and the first trial of post-learning sessions. In addition, during the first trial of the post-learning session place cells did not represent the goal at target environment as densely as in the control one suggesting that goal biased spatial maps are important for the successful recall of these goals. However, this deficit was transient and the maps established before the optogenetic disruption were reinstated after the learning was reinforced during additional learning trials in the post-learning session.

4.1 Role of reactivation in memory consolidation

Hippocampal place cell assemblies represent the cognitive map of the environment, which is thought to be the substrate for spatial navigation and memory (O'Keefe & Dostrovsky, 1971b; Tolman, 1948). While there is a broad experimental evidence that hippocampus is important for the formation of episodic and spatial memories at all stages of memory processing, i.e. encoding, consolidation and recall (Aggleton & Brown, 1999; Burgess et al., 2002; Eichenbaum, 2004; R. G. M. Morris et al., 1990), it is less clear how stable hippocampal space representations are and what is the role of their stability or flexibility in learning. These representations were long thought to remain stable across many days and unit recordings aiming at recording from the same cells across days confirmed this notion (Kentros et al., 1998; Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002; Muller & Kubie, 1987; O'Keefe & Nadel, 1978). However, some recent studies observed a varying degree of stability of hippocampal maps across days or even within the same day (Hainmueller & Bartos, 2018; Ziv et al., 2013). These studies however all used calcium imaging, which may have caused long-term instability of the recorded maps because of the calcium sensors used or due to the means these were expressed such as viral expression. Conceptually it would be hard to resolve how spatial memories could be successfully recalled with instable maps. As in the Dupret et al. (2010) study, in which spatial maps were reinstated in the recall stages, the spatial maps of the control environments remained relatively stable in my experiments as well during the post-probe and post-learning sessions. Their stability was maintained even after their partial reactivation disruption given that only those reactivation events were not disrupted that were confidently encoded as control environments by the algorithm. Hence as past results, my work indicates that stable maps are required for the successful recall of goal-related spatial memories. I will provide further support for this hypothesis later. The transient destabilization and reinstatement of the CA1 spatial maps could be interpreted in the light of the attractor dynamics in the CA3 region, which is an ability of the neural activity to converge to a set of stable sets from similar neighbouring states (Knierim & Zhang, 2012). In the context of my paradigm, different stable states could represent the spatial locations and trajectories in the two different environments. And if the CA3 maps have not been affected by the optogenetic manipulation, its preserved attractor dynamics could facilitate the reinstatement of the maps in the CA1. However, for this reinstatement to occur, a complementary input from the EC was required. This input was generated in the EC during the post-learning trials in both environments.

Explorations of novel environments or significant changes to the familiar environment lead to global remapping of the location of place fields while partial sensory changes or different behavioural context can trigger rate remapping where the location of the place field does not change but the rate at which a place cell fires in its field changes (Allen, Rawlins, Bannerman, & Csicsvari, 2012; J. K. Leutgeb et al., 2005; Muller & Kubie, 1987). Importantly not only spatial factors but also cognitive ones, such as those related to learning can also trigger place field location remapping in an otherwise familiar environment. When animals learn a fixed goal location, place fields of different place cells accumulate near that location (Hollup et al., 2001), while the changing of goal locations causes further remapping of CA1 place fields to represent the new goal locations (Dupret et al., 2010). It has been further shown that proportion of place cells representing the goals in these newly established maps during learning predicted the memory performance of the animal in recalling the goals. In my case, more place cells represented the goals while recalling the control environment than the target environment, the latter of which is associated with a memory deficit. These further suggest that the reinstatement of the goal-specific map is required for the successful recall of the goal locations.

Several studies have suggested importance of reactivation for memory consolidation (Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; van de Ven et al., 2016). As reactivation happens mostly during sharp-wave ripples (SWR), these SWR events were targeted by closed-loop stimulation to disrupt them and demonstrate their role. As a result, a behavioural impairment was seen after few days of training for the same spatial memories with subsequent SWR disruption. In my setup, I have detected the high synchrony periods instead of ripple oscillations, which allowed to respond to the reactivation onset earlier on, because synchrony has been shown to emerge slightly earlier than ripple oscillations (Csicsvari et al., 1999b, p. 199). The memory deficit demonstrated in these studies was not seen during the first few days of stimulation. However, it is hard to judge the memory disruption role of their manipulations in the earlier stimulation days because even control animals were not able to perform the complex maze learning tasks during the earlier stimulation days. In my study, the disruption of the synchrony events caused the behavioural deficit immediately after the rest sessions. Of course, in my work the animal was familiar with the task in general, it only had to update the memory representation while

using an existing memory schema (Tse et al., 2007). Note, however, that in addition to the behavioural differences, other factors may have also contributed why my method was more efficient and was able to disrupt spatial memories on the same day when they were acquired. These include longer disruption sessions (4 hours in contrast to 1 hour in these studies), and detection of high synchrony events instead of ripple events, which allowed earlier disruption.

In addition to SWR disruption, pharmacological approaches have also been used to disrupt reactivation (Dupret et al., 2010; Silva, Feng, & Foster, 2015). Dupret et al. (2010) showed that NMDA receptor (NMDAR) blockade during learning leads to the impaired reactivation of the goal-specific of place cell activity. Although such goal-related maps were reactivated to some degree, old maps that were present before learning were also reactivated causing a mixed reactivation of old and the updated memory traces. NMDAR blockade in (Dupret et al., 2010) has led to behavioural impairment as well on the cheeseboard maze task; although the animal was able to learn the task during the learning trials, it was not able to recall the learned goal locations in the probe trial. It is unclear however in this work whether the memory interference during reactivation or the blockade of neuronal plasticity itself caused the subsequent memory deficit. In a similar manner, the above-mentioned SWR-blockade experiments may have also caused alterations in learning-associated synaptic plasticity. They used electrical stimulation to disrupt SWRs and electrical stimulation during SWRs has been shown to alter CA3-CA1 synaptic weights (King, Henze, Leinekugel, & Buzsaki, 1999). Therefore, indirectly triggered synaptic plasticity may have been a cause of the memory impairment seen after electrical stimulation-mediated SWR disruption. In my experiments, assembly-specific SWR activity was disrupted through optogenetic means. Although a subgroup of pyramidal cells was disinhibited by my manipulations, it is unlikely that disinhibition itself caused synaptic changes that influenced subsequent memory recall considering that cells encoding the target environment were also disinhibited and accordingly fired during disrupted HSEs. In my manipulation, it is also unlikely that reactivation of interfering memory patterns was a cause of the memory deficit as one might expect from the Dupret et al. (2010) NMDAR blockade results. In my case, the assembly-specific disruption did not inhibit the subsequent network activity, but it triggered the expression of altered assembly patterns. In fact, many of the cells increased their firing after

because of disinhibition causing both target and control environment-encoding cells to fire together. Yet my manipulations lead to the disruption of target environment only without affecting the control environment in the majority of experiments.

As mentioned above, my data supports the hypothesis that the successful reinstatement of the previously established goal-oriented map is needed for the recall of previously learned spatial memories. In the Dupret et al. (2010) study, NMDAR blockade not only disrupted reactivation and recall of previously learned goal locations but destabilized the previously established learning-associated maps which were not reinstated following learning. This further supports the role of goal-oriented maps in goal-related navigation. This finding of the Dupret et al. (2010) study was in fact suggested by earlier work showing that (a) new maps established in the novel environment are not reinstated in subsequent exploration of the same environments if NMDAR were blocked during the formation of novel maps (Kentros et al., 1998) (b) if the escape platform location was shifted under NMDAR blockade in the Morris water maze task, the animal is not able to recall to updated platform location in later days (Muller & Kubie, 1987). The effect of the stability of new maps established following novel environment exploration has been examined through optogenetic SWR blockade. Van de Ven et al. (2016) showed that hippocampal assemblies were destabilized after SWRs were silenced in rest sessions in between consecutive explorations of a novel environment. However, this study did not examine the stability of place fields directly but used an indirect place assembly measure using Eigen projections of the expressed place assembly population vectors. Another study by Kovacs et al. (2016) directly compared the place fields and did not see a destabilization of CA1 place fields as a result of silencing SWRs after exploration of novel environments. Here the fewer number of place cells may have caused a false negative error and in a larger data set smaller but statistically significant differences in place fields may have also been revealed. In my case, the stability of the target environment map was impaired only in the probe and in the first trial of post-learning sessions. In addition, similarly to (van de Ven et al., 2016), I have observed a selective stability impairment in the cells that have remapped between start and end of learning. However, at the end of post-probe sessions, the maps that were present at the end of learning were reinstated, which might indicate that the maps in fact stabilised during the learning session itself and their reactivation was not needed to maintain these representations. However, their reactivation would have been needed for the right map to

be recalled at the beginning of the post-probe session. Reinforcement of food at the same goal locations reinstated the map established before. Previously, Schoenenberger et al. (2016) showed that optogenetically inhibiting place cells during exploration cause lasting remapping of their place fields once the animal returns to the same location. My manipulation was similar but was performed during sleep, and sleep disruption could not completely destabilise the same place fields. These together indicate that activity-dependent plasticity may be more effective during waking exploration to alter synaptic weights, which are relevant for place cell stability.

In some of my sessions the behavioural deficit was not pronounced or even reversed. Although I have examined several parameters, I was unable to pinpoint a possible cause, which could have explained why my manipulation did not work in those sessions. However, as my results suggest, the spatial representation was reinstated during later stages of the reinforcing post-learning trials. During the failed days, this reinstatement could happen before or during the post-probe leading to the absence of selective behavioural impairment. Indeed, a correlation between place field similarity and memory performance supports this claim.

4.2 Advancements in the closed-loop methodology

Recent developments in closed-loop experiments allowed to trigger real-time responses to neural events in different brain regions with various applications including antiepileptic stimulations (Gluckman et al., 2001; Psatta, 1983), locomotion restoration after spinal chord injuries (Capogrosso et al., 2016), deep brain stimulation (DBS) for treatment of neurogenic pain (Nguyen et al., 2000) and treatment of Parkinson's disease (Pollak et al., 2002). In hippocampal studies, closed-loop setups have been developed to respond to a wide variety of neural event triggers, including specific phases of theta oscillations (Siegle & Wilson, 2014), spikes of an individual cells (de Lavilléon et al., 2015), ripple-band oscillations (Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Jadhav, Kemere, German, & Frank, 2012b; Roux et al., 2017; van de Ven et al., 2016) and overall activity of a population of neurons (Newman et al., 2015). In contrast to my method, the latter did not demonstrate any real-

time decoding, only spike detection and feedback control of overall firing rate on the desired level were performed. Moreover, developments in optogenetics allowed moving from electrical stimulation to optogenetic stimulation, which excludes the unspecific effect of electrical stimulation.

In hippocampal studies, electrical and optogenetic disruption studies have established a role of SWR in the spatial learning and memory and the stability of the hippocampal cognitive map (Ego-Stengel & Wilson, 2010, p.; Girardeau et al., 2009; Jadhav et al., 2012b; van de Ven et al., 2016). In these experiments, an electrical impulse, in case of electrical stimulation, or a light pulse, in case of optogenetic stimulation, was delivered in response to the onset of the ripple-band (150-250 Hz) oscillations. In my paradigm, I have employed an emerging class of methods, which allow addressing the content of the SWR-associated HSE reactivation event, in my case, in the context of a binary classification of spiking patterns.

Studies of neural firing patterns have been so far largely limited to analyses of the offline-sorted spiking data. Because of this, it was hard to use them for paradigms with dynamic coding, e.g. involving place cell remapping. However, development of cluster-less decoding methods (Chen, Kloosterman, Layton, & Wilson, 2012; Deng et al., 2016; Kloosterman et al., 2014) has allowed to skip the clustering step in the decoding models as well as to achieve a better decoding quality than in the cluster-based methods, because clustering is inevitably associated with information loss in the form of sorting errors and discarded spikes. This class of methods was used, for example, to create a setup capable of real-time detection of forward and reverse replay (Deng et al., 2016). Nevertheless, there were no demonstrations of setups capable of efficiently responding to the content of neural patterns during transient events, like SWRs. In my project, I have developed such a system that can deliver a light pulse based on the outcome of the spiking pattern classification.

An important modification that gave my method an advantage over existing SWR-disruption techniques is the detection of the HSE from the multiple-unit activity. Detecting ripples requires at least two full cycles, by the time of which the beginning segment of an HSE can be progressed. It has been shown that MUA occurs before the onset of the ripples (Csicsvari

et al., 1999b). And even a simple method that compares the total number of spikes in all tetrodes to the baseline allows to detect HSEs earlier than ripple oscillations.

At the time I started my project, there were no reported methods for decoding spiking patterns with latencies below 2 ms. However, recent improvements in performance and modularity of neural signal processing have dramatically decreased the response latencies as well as made easier integration of closed-loop setups (Ciliberti & Kloosterman, 2017; Dutta, Ackermann, & Kemere, 2018; Siegle et al., 2017). Nevertheless, setting up, tuning and debugging real-time performance-critical systems is a challenge and no published work has yet applied these methods to study hippocampal function. Therefore, my work provides the first experimental application of such spike pattern decoding methods in closed-loop setups.

5 Future work

My results revealed that the content of SWRs is important for the consolidation of encoded memory traces. Disruption of memory traces associated with target environment caused a behavioural deficit during the memory tests right after resting period while the animals were able to perform the task in the undisrupted control environment. In addition, this manipulation caused the destabilization of hippocampal place map; albeit this destabilization was only transient and after reinforcing it in the post-learning trials, the maps have returned to the original state. This fact, along with other observations, such as reasons for the failure to produce these deficits in some of the sessions, remain unexplained. Below I suggest future research topics that could elucidate these uncertainties.

According to my results, the destabilization of the hippocampal map of the target environment was only transient, and the representations have reinstated at the end of new learning trials. It is unclear, however, which region was the source of this returned representations. According to the structural and functional connectivity of CA1, the candidate sources of the map include the MEC and CA3 regions. Disrupting reactivation in these regions while detecting HSE in CA1 might help to answer this question. In particular, as the ripples and reactivation originates in CA3, targeting CA3 instead of CA1 could result in more pronounced behavioural effect through affecting larger part of the network as well as earlier disruption. In addition, the physiological effect could be more permanent due to inability to recover the spatial representation in CA1 from the CA3 input.

My research was based on the assumption that reactivation of two environments, whose representation is different enough, takes place in a dissociated manner. This question has not been sufficiently studied, but the demonstrated ability of my method to cause the selective behavioural, as well as physiological impairment in relation to the transient destabilisation of the target maps, hint that this is the case. A detailed analysis of reactivation content with multiple environments could provide more insight into this question and can reveal whether reactivation epochs exist, which incorporate multiple environments. If these exist, one can decipher the role of such mixed reactivation events as well.

My results regarding to map stability for learning goal locations in familiar environments were different from those of SWR disruption experiments following exploration in a novel environment. This may suggest that the network mechanism of the formation and stabilisation of goal-related and novel environment-associated maps are different. To decipher further these differnces, I suggest performing the similar experiments with two novel environments in which in the following exploration only the reactivation of one of the environment is blocked to test whether content specific blockade leads to lasting destabilisation of the blocked map or similar to the cheeseboard disruption will be transient only.

And last but not least, the closed-loop real-time decoding method that I have developed can be, with some effort, adapted for use with different brain areas, stimulus modalities and recording systems. Obviously, the hyper-parameters of the statistical models will need to be established for the specific cases, but this can be easily done, for example, by means of stochastic gradient descent optimization, which I have used in my project.

BIBLIOGRAPHY

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *The Behavioral and Brain Sciences*, *22*(3), 425–444; discussion 444-489.

Allen, K., Rawlins, J. N. P., Bannerman, D. M., & Csicsvari, J. (2012). Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping. *Journal of Neuroscience*, *32*(42), 14752–14766.

Amaral, D. G., & Lavenex, P. (2007). Hippocampal Neuroanatomy. In P. Andersen (Ed.), *The hippocampus book* (pp. 37–109). Oxford University press.

Astur, R. S., Taylor, L. B., Mamelak, A. N., Philpott, L., & Sutherland, R. J. (2002). Humans with hippocampus damage display severe spatial memory impairments in a virtual Morris water task. *Behavioural Brain Research*, *132*(1), 77–84.

Barnes, C. A. (1979). Memory deficits associated with senescence: a neurophysiological and behavioral study in the rat. *J.Comp Physiol Psychol.*, *93*(0021-9940 (Print)), 74–104.

Berger, H. (1929). Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, *87*(1), 527–570. https://doi.org/10.1007/BF01797193

Boccara, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., & Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience*, *13*(8), 987–994. https://doi.org/10.1038/nn.2602

Braak, H., Braak, E., Yilmazer, D., & Bohl, J. (1996). Functional anatomy of human hippocampal formation and related structures. *Journal of Child Neurology*, *11*(4), 265–275.

Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., & 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*(0270-6474 (Print)), 7411–7425.

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*(0896-6273 (Print)), 625–641.

Buzsaki, G. (1986). Hippocampal sharp waves: their origin and significance. *Brain Res.*, *398*(0006–8993), 242–252.

Buzsaki, G. (1989). Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience*, *31*(0306–4522), 551–570.

Buzsaki, G., Horvath, Z., Urioste, R., Hetke, J., & Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science*, *256*(0036–8075), 1025–1027.

Buzsaki, G., Leung, L. W., & Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.*, *287*(0006–8993), 139–171.

Capogrosso, M., Milekovic, T., Borton, D., Wagner, F., Moraud, E. M., Mignardot, J.-B., … Courtine, G. (2016). A brain–spine interface alleviating gait deficits after spinal cord injury in primates. *Nature*, *539*(7628), 284–288. https://doi.org/10.1038/nature20118

Chapin, J. K. (2004). Using multi-neuron population recordings for neural prosthetics. *Nature Neuroscience*, *7*(5), 452–455. https://doi.org/10.1038/nn1234

Chen, Z., Kloosterman, F., Layton, S., & Wilson, M. A. (2012). Transductive neural decoding for unsorted neuronal spikes of rat hippocampus. *Conference Proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, *2012*, 1310–1313. https://doi.org/10.1109/EMBC.2012.6346178

Chrobak, J. J., & Buzsaki, G. (1994). Selective activation of deep layer (V-VI) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat. *J.Neurosci.*, *14*(0270–6474), 6160–6170.

Chrobak, J. J., & Buzsaki, G. (1996). High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *J.Neurosci.*, *16*(0270-6474 (Print)), 3056–3066.

Ciliberti, D., & Kloosterman, F. (2017). Falcon: a highly flexible open-source software for closed-loop neuroscience. *Journal of Neural Engineering*, *14*(4), 045004. https://doi.org/10.1088/1741-2552/aa7526

Clark, R. E., Broadbent, N. J., & Squire, L. R. (2005). Impaired remote spatial memory after hippocampal lesions despite extensive training beginning early in life. *Hippocampus*, *15*(1050-9631 (Print)), 340–346.

Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., & Buzsaki, G. (1999a). Fast network oscillations in the hippocampal CA1 region of the behaving rat. *J.Neurosci.*, *19*(1529–2401), RC20.

Csicsvari, J., Hirase, H., Czurko, A., Mamiya, A., & Buzsaki, G. (1999b). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *J.Neurosci.*, *19*(0270–6474), 274–287.

Csicsvari, J., Hirase, H., Mamiya, A., & Buzsaki, G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron*, *28*(0896–6273), 585–594.

de Lavilléon, G., Lacroix, M. M., Rondi-Reig, L., & Benchenane, K. (2015). Explicit memory creation during sleep demonstrates a causal role of place cells in navigation. *Nature Neuroscience*, *18*(4), 493–495. https://doi.org/10.1038/nn.3970

Deng, X., Liu, D. F., Karlsson, M. P., Frank, L. M., & Eden, U. T. (2016). Rapid classification of hippocampal replay content for real-time applications. *Journal of Neurophysiology*, *116*(5), 2221–2235. https://doi.org/10.1152/jn.00151.2016

Dupret, D., O'Neill, J., & Csicsvari, J. (2013). Dynamic Reconfiguration of Hippocampal Interneuron Circuits during Spatial Learning. *Neuron*, *78*(1), 166–180.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat.Neurosci.*, *13*(1546-1726 (Electronic)), 995–1002.

Dutta, S., Ackermann, E., & Kemere, C. (2018). Analysis of an open source, closed-loop, realtime system for hippocampal sharp-wave ripple disruption. https://doi.org/10.1101/298661

Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, *20*(1098-1063 (Electronic)), 1–10.

Eichenbaum, H. (2004). Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron*, *44*(0896–6273), 109–120.

Eichenbaum, H., Stewart, C., & Morris, R. G. (1990). Hippocampal representation in place learning. *J Neurosci*, *10*(0270-6474 (Print)), 3531–3542.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(1476-4687 (Electronic)), 184–188.

Eschenko, O., Ramadan, W., Molle, M., Born, J., & Sara, S. J. (2008). Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*, *15*(4), 222–228. https://doi.org/10.1101/lm.726008

Escolano, C., Ramos Murguialday, A., Matuz, T., Birbaumer, N., & Minguez, J. (2010). A telepresence robotic system operated with a P300-based brain-computer interface: initial tests with ALS patients. *Conference Proceedings: ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, *2010*, 4476–4480. https://doi.org/10.1109/IEMBS.2010.5626045

Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, (1476-4687 (Electronic)). Retrieved from PM:16474382

Freund, T. F., & Buzsaki, G. (1996). Interneurons of the hippocampus. *Hippocampus*, *6*(1050–9631), 347–470.

Georges-Francois, P., Rolls, E. T., & Robertson, R. G. (1999). Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position or place. *Cereb.Cortex*, *9*(1047-3211 (Print)), 197–212.

Gilbert, P. E., Kesner, R. P., & DeCoteau, W. E. (1998). Memory for spatial location: role of the hippocampus in mediating spatial pattern separation. *Journal of Neuroscience*, *18*(0270-6474 (Print)), 804–810.

Girardeau, G., Benchenane, K., Wiener, S. I., Buzsaki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nat.Neurosci*, *12*(1546-1726 (Electronic)), 1222–1223.

Gluckman, B. J., Nguyen, H., Weinstein, S. L., & Schiff, S. J. (2001). Adaptive electric field control of epileptic seizures. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *21*(2), 590–600.

Goodrich-Hunsaker, N. J., & Hopkins, R. O. (2010). Spatial memory deficits in a virtual radial arm maze in amnesic participants with hippocampal damage. *Behavioral Neuroscience*, *124*(3), 405–413. https://doi.org/10.1037/a0019193

Grosenick, L., Marshel, J. H., & Deisseroth, K. (2015). Closed-loop and activity-guided optogenetic control. *Neuron*, *86*(1), 106–139. https://doi.org/10.1016/j.neuron.2015.03.034

Guzman, S. J., Schlogl, A., Frotscher, M., & Jonas, P. (2016). Synaptic mechanisms of pattern completion in the hippocampal CA3 network. *Science*, *353*(6304), 1117–1123. https://doi.org/10.1126/science.aaf1836

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(1476-4687 (Electronic)), 801–806.

Hainmueller, T., & Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, *558*(7709), 292–296. https://doi.org/10.1038/s41586-018-0191-2

Han, X., Chow, B. Y., Zhou, H., Klapoetke, N. C., Chuong, A., Rajimehr, R., … Boyden, E. S. (2011). A high-light sensitivity optical neural silencer: development and application to

optogenetic control of non-human primate cortex. *Frontiers in Systems Neuroscience*, *5*, 18. https://doi.org/10.3389/fnsys.2011.00018

Harper, A. (2010). Mouse models of neurological disorders—A comparison of heritable and acquired traits. *Biochimica et Biophysica Acta (BBA) - Molecular Basis of Disease*, *1802*(10), 785–795. https://doi.org/10.1016/j.bbadis.2010.05.009

Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H., & Buzsaki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J.Neurophysiol.*, *84*(0022–3077), 401–414.

Hochberg, L. R., Bacher, D., Jarosiewicz, B., Masse, N. Y., Simeral, J. D., Vogel, J., … Donoghue, J. P. (2012). Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, *485*(7398), 372–375. https://doi.org/10.1038/nature11076

Hok, V., Lenck-Santini, P. P., Roux, S., Save, E., Muller, R. U., & Poucet, B. (2007). Goal-related activity in hippocampal place cells. *Journal of Neuroscience*, *27*(1529-2401 (Electronic)), 472–482.

Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. B., & Moser, E. I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J.Neurosci*, *21*(1529-2401 (Electronic)), 1635–1644.

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012a). Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science*, *336*(6087), 1454–1458.

Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012b). Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science*, *336*(6087), 1454–1458.

Kentros, C., Hargreaves, E., Hawkins, R. D., Kandel, E. R., Shapiro, M., & Muller, R. V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science*, *280*(0036-8075 (Print)), 2121–2126.

Kesner, R. P., Farnsworth, G., & DiMattia, B. V. (1989). Double dissociation of egocentric and allocentric space following medial prefrontal and parietal cortex lesions in the rat. *Behavioral Neuroscience*, *103*(5), 956–961.

King, C., Henze, D. A., Leinekugel, X., & Buzsaki, G. (1999). Hebbian modification of a hippocampal population pattern in the rat. *J.Physiol*, *521 Pt 1*(0022–3751), 159–167.

Klausberger, T., & Somogyi, P. (2008). Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science*, *321*(1095-9203 (Electronic)), 53–57.

Kloosterman, F., Layton, S. P., Chen, Z., & Wilson, M. A. (2014). Bayesian decoding using unsorted spikes in the rat hippocampus. *Journal of Neurophysiology*, *111*(1), 217–227. https://doi.org/10.1152/jn.01046.2012

Knierim, J. J., & Zhang, K. (2012). Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System. *Annual Review of Neuroscience*, *35*(1), 267–285. https://doi.org/10.1146/annurev-neuro-062111-150351

Kovács, K. A., O'Neill, J., Schoenenberger, P., Penttonen, M., Ranguel Guerrero, D. K., & 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*(10), e0164675. https://doi.org/10.1371/journal.pone.0164675

Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: effects of behavioral state, experience, and EEG dynamics. *J.Neurosci.*, *19*(1529–2401), 4090–4101.

Lazarou, I., Nikolopoulos, S., Petrantonakis, P. C., Kompatsiaris, I., & Tsolaki, M. (2018). EEG-Based Brain-Computer Interfaces for Communication and Rehabilitation of People with Motor Impairment: A Novel Approach of the 21 st Century. *Frontiers in Human Neuroscience*, *12*, 14. https://doi.org/10.3389/fnhum.2018.00014

Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. A., McNaughton, B. L., … Moser, E. I. (2005). Progressive transformation of hippocampal neuronal representations in "morphed" environments. *Neuron*, *48*(0896-6273 (Print)), 345–358.

Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M. B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, *309*(1095-9203 (Electronic)), 619–623.

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

Lever, C., Wills, T., Cacucci, F., Burgess, N., & O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, *416*(0028-0836 (Print)), 90–94.

Logothetis, N. K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H. C., … Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. *Nature*, *491*(7425), 547–553. https://doi.org/10.1038/nature11618

Logue, S. F., Paylor, R., & Wehner, J. M. (1997). Hippocampal lesions cause learning deficits in inbred mice in the Morris water maze and conditioned-fear task. *Behavioral Neuroscience*, *111*(1), 104–113.

Louie, K., & Wilson, M. A. (2001). Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, *29*(0896–6273), 145–156.

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol.Rev.*, *102*(0033-295X (Print)), 419–457.

McNaughton, B. L., & Morris, R. G. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.*, *10*(10), 408–415.

Mizuseki, K., Sirota, A., Pastalkova, E., & Buzsáki, G. (2009). Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron*, *64*(2), 267–280. https://doi.org/10.1016/j.neuron.2009.08.037

Morris, R. G., Garrud, P., Rawlins, J. N., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, *297*(0028-0836 (Print)), 681–683.

Morris, R. G. M., Schenk, F., Tweedie, F., & Jarrard, L. E. (1990). Ibotenate Lesions of Hippocampus and/or Subiculum: Dissociating Components of Allocentric Spatial Learning. *The European Journal of Neuroscience*, *2*(12), 1016–1028.

Moser, E., Moser, M. B., & Andersen, P. (1993). Spatial learning impairment parallels the magnitude of dorsal hippocampal lesions, but is hardly present following ventral lesions 6. *J.Neurosci.*, *13*(0270-6474 (Print)), 3916–3925.

Moser, M. B., Moser, E. I., Forrest, E., Andersen, P., & Morris, R. G. (1995). Spatial learning with a minislab in the dorsal hippocampus. *Proc Natl Acad Sci U.S.A*, *92*(0027–8424), 9697–9701.

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J.Neurosci.*, *7*(0270-6474 (Print)), 1951–1968.

Muller, R. U., & Kubie, J. L. (1989). The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, *9*(0270-6474 (Print)), 4101–4110.

Muller, R. U., Kubie, J. L., & Ranck, J. B., Jr. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J.Neurosci.*, *7*(0270-6474 (Print)), 1935–1950.

Newman, J. P., Fong, M., Millard, D. C., Whitmire, C. J., Stanley, G. B., & Potter, S. M. (2015). Optogenetic feedback control of neural activity. *ELife*, *4*, e07192. https://doi.org/10.7554/eLife.07192

Nguyen, J. P., Lefaucher, J. P., Le Guerinel, C., Eizenbaum, J. F., Nakano, N., Carpentier, A., … Keravel, Y. (2000). Motor Cortex Stimulation in the Treatment of Central and Neuropathic Pain. *Archives of Medical Research*, *31*(3), 263–265. https://doi.org/10.1016/S0188-4409(00)00078-3

Nokia, M. S., Mikkonen, J. E., Penttonen, M., & Wikgren, J. (2012). Disrupting neural activity related to awake-state sharp wave-ripple complexes prevents hippocampal learning. *Frontiers in Behavioral Neuroscience*, *6*, 84. https://doi.org/10.3389/fnbeh.2012.00084

O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(6581), 425–428. https://doi.org/10.1038/381425a0

O'Keefe, J., & Dostrovsky, J. (1971a). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, *34*(0006–8993), 171–175.

O'Keefe, J., & Dostrovsky, J. (1971b). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.*, *34*(0006–8993), 171–175.

O'Keefe, J., & Nadel, L. (1978). *Hippocampus as a Cognitive Map*. Oxford: Clarindon.

O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(1050–9631), 317–330.

Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nature Neuroscience*, *19*(6), 792–794. https://doi.org/10.1038/nn.4291

O'Neill, J., Boccara, C. N., Stella, F., Schoenenberger, P., & Csicsvari, J. (2017). Superficial layers of the medial entorhinal cortex replay independently of the hippocampus. *Science (New York, N.Y.)*, *355*(6321), 184–188. https://doi.org/10.1126/science.aag2787

O'Neill, J., Pleydell-Bouverie, B., Dupret, D., & Csicsvari, J. (2010). Play it again: reactivation of waking experience and memory. *Trends Neurosci.*, *33*(1878-108X (Electronic)), 220–229.

O'Neill, J., Senior, T., & Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron*, *49*(0896-6273 (Print)), 143–155.

O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nat.Neurosci.*, *11*(1097-6256 (Print)), 209–215.

Paninski, L., Pillow, J., & Lewi, J. (2007). Statistical models for neural encoding, decoding, and optimal stimulus design. *Progress in Brain Research*, *165*, 493–507. https://doi.org/10.1016/S0079-6123(06)65031-0

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*(1095-9203 (Electronic)), 1322–1327.

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, *497*(7447), 74–79.

Pillay, N. S., Kellaway, L. A., & Kotwal, G. J. (2008). Early detection of memory deficits and memory improvement with vaccinia virus complement control protein in an Alzheimer's disease model. *Behavioural Brain Research*, *192*(2), 173–177. https://doi.org/10.1016/j.bbr.2008.03.038

Pollak, P., Fraix, V., Krack, P., Moro, E., Mendes, A., Chabardes, S., … Benabid, A.-L. (2002). Treatment results: Parkinson's disease. *Movement Disorders: Official Journal of the Movement Disorder Society*, *17 Suppl 3*, S75-83.

Psatta, D. M. (1983). Control of chronic experimental focal epilepsy by feedback caudatum stimulations. *Epilepsia*, *24*(4), 444–454.

Riedel, G., Micheau, J., Lam, A. G., Roloff, E. L., Martin, S. J., Bridge, H., … Morris, R. G. (1999). Reversible neural inactivation reveals hippocampal participation in several memory processes. *Nat.Neurosci*, *2*(1097-6256 (Print)), 898–905.

Rolls, Edmund T. (2007). An attractor network in the hippocampus: theory and neurophysiology. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *14*(11), 714–731. https://doi.org/10.1101/lm.631207

Rolls, Edmund T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience*, *7*. https://doi.org/10.3389/fnsys.2013.00074

Rolls, E.T. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, *9*(1050-9631 (Print)), 467–480.

Rolls, E.T., & O'Mara, S. M. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus*, *5*(1050-9631 (Print)), 409–424.

Roux, L., Hu, B., Eichler, R., Stark, E., & Buzsáki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nature Neuroscience*, *20*(6), 845–853. https://doi.org/10.1038/nn.4543

Sanger, T. D. (2003). Neural population codes. *Current Opinion in Neurobiology*, *13*(2), 238–249.

Schlingloff, D., Káli, S., Freund, T. F., Hájos, N., & Gulyás, A. I. (2014). Mechanisms of sharp wave initiation and ripple generation. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *34*(34), 11385–11398. https://doi.org/10.1523/JNEUROSCI.0867-14.2014

Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2016). Activity-dependent plasticity of hippocampal place maps. *Nature Communications*, *7*, 11824. https://doi.org/10.1038/ncomms11824

Schwartz, A. B., Cui, X. T., Weber, D. J., & Moran, D. W. (2006). Brain-Controlled Interfaces: Movement Restoration with Neural Prosthetics. *Neuron*, *52*(1), 205–220. https://doi.org/10.1016/j.neuron.2006.09.019

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J.Neurol.Neurosurg.Psychiatry*, *20*(0022-3050 (Print)), 11–21.

Siegel, J. J., Nitz, D., & Bingman, V. P. (2006). Lateralized functional components of spatial cognition in the avian hippocampal formation: evidence from single-unit recordings in freely moving homing pigeons. *Hippocampus*, *16*(1050-9631 (Print)), 125–140.

Siegle, J. H., López, A. C., Patel, Y. A., Abramov, K., Ohayon, S., & Voigts, J. (2017). Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. *Journal of Neural Engineering*, *14*(4), 045003. https://doi.org/10.1088/1741-2552/aa5eea

Siegle, J. H., & Wilson, M. A. (2014). Enhancement of encoding and retrieval functions through theta phase-specific manipulation of hippocampus. *ELife*, *3*, e03061.

Silva, D., Feng, T., & Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature Neuroscience*, *18*(12), 1772–1779. https://doi.org/10.1038/nn.4151

Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, *271*(0036–8075), 1870–1873.

Skelton, R. W., Ross, S. P., Nerad, L., & Livingstone, S. A. (2006). Human spatial navigation deficits after traumatic brain injury shown in the arena maze, a virtual Morris water maze. *Brain Injury*, *20*(2), 189–203. https://doi.org/10.1080/02699050500456410

Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science (New York, N.Y.)*, *322*(5909), 1865–1868. https://doi.org/10.1126/science.1166466

Somogyi, P., & Klausberger, T. (2005). Defined types of cortical interneurone structure space and spike timing in the hippocampus. *J.Physiol*, *562*(0022-3751 (Print)), 9–26.

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychol.Rev.*, *99*(0033-295X (Print)), 195–231.

Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., & Buzsáki, G. (2011). Relationships between hippocampal sharp waves, ripples, and fast gamma oscillation: influence of dentate and entorhinal cortical activity. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*(23), 8605–8616. https://doi.org/10.1523/JNEUROSCI.0294-11.2011

Sutherland, R.J., Whishaw, I. Q., & Kolb, B. (1983). A behavioural analysis of spatial localization following electrolytic, kainate- or colchicine-induced damage to the hippocampal formation in the rat. *Behav Brain Res.*, *7*(0166-4328 (Print)), 133–153.

Sutherland, Robert J., & Hamilton, D. A. (2004). Rodent spatial navigation: at the crossroads of cognition and movement. *Neuroscience & Biobehavioral Reviews*, *28*(7), 687–697. https://doi.org/10.1016/j.neubiorev.2004.09.012

Taube, J. S. (1995). Place cells recorded in the parasubiculum of freely moving rats. *Hippocampus*, *5*(6), 569–583. https://doi.org/10.1002/hipo.450050608

Taube, J. S., Muller, R. U., & Ranck, J. B., Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, *10*(0270-6474 (Print)), 420–435.

Thompson, L. T., & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *J.Neurosci.*, *9*(0270-6474 (Print)), 2382–2390.

Tolman, E. (1948). Cognitive maps in Rats and men. *Psychol.Rev.*, *55*(4), 189–208.

Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, *93*(2), 1074–1089. https://doi.org/10.1152/jn.00697.2004

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., … Morris, R. G. M. (2007). Schemas and memory consolidation. *Science (New York, N.Y.)*, *316*(5821), 76–82. https://doi.org/10.1126/science.1135935

van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, *92*(5), 968–974. https://doi.org/10.1016/j.neuron.2016.10.020

Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr.Clin.Neurophysiol.*, *26*(0013–4694), 407–418.

Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space 71. *Science*, *261*(0036–8075), 1055–1058.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*(0036–8075), 676–679.

Yau, S., Li, A., & So, K.-F. (2015). Involvement of Adult Hippocampal Neurogenesis in Learning and Forgetting. *Neural Plasticity*, *2015*, 1–13. https://doi.org/10.1155/2015/717958

Ylinen, A., Bragin, A., Nadasdy, Z., Jando, G., Szabo, I., Sik, A., & Buzsaki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *J.Neurosci.*, *15*(0270–6474), 30–46.

Yoshida, M., Goto, K., & Watanabe, S. (2001). Task-dependent strain difference of spatial learning in C57BL/6N and BALB/c mice. *Physiology & Behavior*, *73*(1–2), 37–42.

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