

Drift and stabilization of cortical response selectivity

Dissertation

for the award of the degree

Doctor rerum naturalium

of the Georg-August-Universität Göttingen

within the doctoral program

Theoretical and Computational Neuroscience (PTCN)

of the Georg-August University School of Science (GAUSS)

submitted by

Alexander Schmidt

from Karl-Marx-Stadt, now Chemnitz

Göttingen 2020

Thesis Committee

Prof. Dr. Fred Wolf

Campus Institute for Dynamics of Biological Networks and
Max Planck Institute for Dynamics und Self-Organization

Prof. Dr. Siegrid Löwel

Bernstein Focus Neurotechnology (BFNT) and
Johann-Friedrich-Blumenbach Institute for Zoology and Anthropology
Dept. of Systems Neuroscience

Prof. Dr. Jörg Enderlein

III. Physical Institute
Georg-August-Universität Göttingen

Members of the examination board

Referee: Prof. Dr. Fred Wolf

Campus Institute for Dynamics of Biological Networks and
Max Planck Institute for Dynamics und Self-Organization

2nd Referee: Prof. Dr. Siegrid Löwel

Bernstein Focus Neurotechnology (BFNT) and
Johann-Friedrich-Blumenbach Institute for Zoology and Anthropology
Dept. of Systems Neuroscience

Further members of the examination board

Prof. Dr. Jörg Enderlein

III. Institute of Physics
Georg-August-Universität Göttingen

Prof. Dr. Florentin Wörgötter

III. Institute of Physics
Georg-August Universität Göttingen

Prof. Dr. Nils Brose

Dept. of Molecular Neurobiology
Max Planck Institute for Experimental Medicine

Prof. Dr. Alexander Ecker

Institute of Computer Science
University of Göttingen

Date of the oral examination: 23rd November 2020

CONTENTS

| | |
|---|-----------|
| Contents | i |
| 1 Introduction | 1 |
| 1.1 Dynamics of receptive fields | 2 |
| 1.2 The hippocampus | 3 |
| 1.3 Place coding within the hippocampus | 6 |
| 1.4 Virtual reality | 13 |
| 1.5 Overview of this work | 13 |
| 2 Methods | 17 |
| 2.1 Content | 17 |
| 2.2 Experimental setup | 18 |
| 2.3 Motion correction | 20 |
| 2.4 Neuron detection | 21 |
| 2.5 Long-term tracking of neurons | 26 |
| 2.6 Place field detection | 33 |
| 2.7 Artificial activity data | 37 |
| 2.8 Discussion | 37 |
| 3 Long-term recordings of hippocampal activity | 41 |
| 3.1 Content | 41 |
| 3.2 Testing of the place cell detection algorithm | 41 |
| 3.3 A complete set of neurons across all sessions | 46 |
| 3.4 A vast data set | 55 |
| 3.5 Discussion | 58 |
| 4 Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus | 61 |
| 4.1 Content | 61 |
| 4.2 Citation and original contribution | 61 |
| 5 Turnover dynamics of the spatial map under baseline conditions | 85 |

CONTENTS

| | | |
|----------|--|------------|
| 5.1 | Content | 85 |
| 5.2 | Short-term dynamics of place representation | 86 |
| 5.3 | Turnover of the spatial map in CA1 | 88 |
| 5.4 | Population dynamics across different levels | 91 |
| 5.5 | Time in the hippocampus | 104 |
| 5.6 | Discussion | 106 |
| 6 | Influences on network dynamics | 109 |
| 6.1 | Content | 109 |
| 6.2 | The impact of different locations | 110 |
| 6.3 | Individual contributions to observed population dynamics | 113 |
| 6.4 | Discussion | 120 |
| 7 | Summary and discussion | 123 |
| 7.1 | Data / Methods | 123 |
| 7.2 | Dynamics of the spatial representations in CA1 | 124 |
| 7.3 | Overall effect / Future works | 129 |
| A | Constrained non-negative matrix factorization (CNMF) | 131 |
| B | Nested sampling | 137 |
| | REFERENCES | 141 |

Chapter One

Introduction

Everyday life poses a multitude of challenges towards living beings, which in many cases has guided evolution towards the development of a brain, with regions responsible for specific tasks. The hippocampus, being part of the cerebral cortex, is one of these regions which has enjoyed wide attention throughout the last decades due to its widely accepted central role in memory^{1,2} and spatial navigation^{3,4}.

Research on its role in the formation of new memories has been sparked by reports of the patient Henry Molaison in 1954⁵, who had his hippocampi surgically removed in an attempt to relieve epileptic seizures and was subsequently reported to suffer from severe anterograde and partial retrograde amnesia. While the observed anterograde amnesia - the disability to form new memories - suggests a central role of the hippocampus in memory formation, the partial retrograde amnesia suggests that other regions, independent of the hippocampus must play a role in retention of past memories. This independence from the hippocampus has been found to arise from a consolidation process, in which information is gradually transferred from the hippocampus to downstream regions over a period of several days or weeks^{1,2,6}, which requires the hippocampus to maintain a stable representation of the memory.

The biological basis to this process is reported to happen on a synaptic level^{7,8}, where sequences of coactivated neurons are reactivated to strengthen their synaptic connections^{9,10} via Hebbian learning¹¹. However, hippocampal synapses, especially within the CA1 region, are reported to be in a highly dynamic state with almost complete turnover happening on a time scale of a few days^{12,13}. Synaptic connectivity is commonly assumed to have a central role for orchestrated activity and thus the reactivation of neuronal sequences. Memories within the hippocampus are therefore in a transient state, balanced between stability required by the process of consolidation, and turnover, driven by the underlying synaptic dynamics, possibly erasing relevant connections.

Given these circumstances, how can the hippocampus reliably maintain relevant information over the time scales of consolidation, being able to recall short term memories with high precision over several days?

Until recent years reports of dynamics and long term stability of neuronal function on a

time scale of days or weeks have been largely anecdotal and focussed on individual neurons^{3,14}. This possibly introduced a bias towards highly stable individual units. Advances in imaging and analysis techniques throughout the last decade now allow for a description of the dynamics encountered in unchanging environments over days, weeks, or even months for hundreds to thousands of neurons¹⁵⁻¹⁸ and have revealed a high degree of turnover on a population level, even when experimental conditions are held constant. Especially the CA1 region of the hippocampus has received increased attention due to its function as an output layer of the hippocampal loop integrating information from different parts of the hippocampal region (see Sec. 1.2).

These studies have examined the role of the hippocampus in place coding via its famous place cells, which were discovered in the 1970s by John O'Keefe^{3,19}. The representation of place by a cognitive map⁴ (see Sec. 1.3) changes over time and becomes increasingly dissimilar with the initial state^{15,18,20}. An understanding of the underlying statistics and mechanisms, however, is still missing.

In this thesis, I set out to study the statistics of turnover dynamics in CA1 place field coding: Which effects contribute to the observed dynamics of the place map? Are there traces of stability to be found? What is the contribution of subpopulations or individual neurons to the dynamics? And are timescales in agreement with reported synaptic turnover dynamics?

I will examine these questions on the basis of a set of mouse data from the deep CA1 subregion of the hippocampus, by means of a large scale data-analysis of longitudinally recorded neural activity.

1.1 Dynamics of receptive fields

This work will deal with the description of dynamics within the neural code of hippocampal CA1 neurons, to which this chapter will provide an introduction. I will first introduce the basic concepts of functional dynamics (Sec. 1.1.1) and receptive fields (Sec. 1.1.2) and then delve into the hippocampal structure, its neurons and their role in place coding as well as an overview of the current knowledge about place field dynamics.

1.1.1 The function of dynamics

Cortical regions of the brain are in a constant state of balance, in which they have to account for different needs²¹ requiring plasticity (for learning and adaptation) as well as stability (for reliable recall of learned behavior). Especially in some hippocampal regions one can observe a high degree of network dynamics accompanied by a high rate of turnover in synapses¹². These are realized by changes in the activity as well as in the functional role of neurons which is discussed in further detail in Ch. 5.

Recent studies have suggested that this turnover is not only a necessary trade-off to allow for adaptation and learning, but that gradual changes in the neural code give a means to efficiently

timestamp memories and thus provide a way to integrate the perception of long time scales and episodic memory¹⁶. The notion that gradual changes allow for time inference is further supported by a work showing that rat memory works in a *How long ago*, rather than *When-fashion*²².

However, these dynamics are difficult to realize for information encoded on an individual neuron level. This suggests an important role of population-level coding, through which stimulus features can be reliably encoded while allowing response properties of single units to be highly dynamic. Furthermore, it motivates the approach of this study to approach the dynamics from a population level.

1.1.2 Receptive fields

Information in the brain is transmitted and processed via generation and reception of electrical signals - so called *action potentials* (APs), which are transported in between neurons connected via synapses. In a large number of cortical regions, the rate of AP elicitation is reported to be impacted by sensory stimuli in a stereotypical way: neurons have preferred stimuli, for which their average rate of AP generation - the firing rate response - is maximal, declining with increasing difference from it. This region of increased response is termed the *receptive field* of a neuron.

Common examples are motor neurons in the motor cortex, responding to the directionality of arm movements, neurons in the auditory cortex responding to frequencies or pitches, neurons of the visual cortex responding to the orientation of contours²³ or neurons in the hippocampus responding to environmental cues or location itself³. The discovery of such spiking patterns in the visual cortex in the 1950s by Hubel and Wiese²³ sparked still ongoing research on how such preferences arise and how such function is integrated in cortical networks²⁴.

In many cases, a receptive field can be well approximated by a projection onto a single input dimension such as sound frequency²⁵, orientation of visual gratings²³, or, in the case of the hippocampus, places³. However, even though the classification of neurons into units selective to a unique stimulus is a tempting approach to shed light onto the functionality of the brain, reality is not as simple. This can be observed especially in the hippocampus²⁶. While for many neurons a receptive field to a certain type of stimulus (e.g. location) can be detected, at the same time those neurons' activity is modulated by different types of stimuli^{26,27}, resulting in a joint representation of several external or internal variables. A simple example of this is the run-direction-¹⁵ or behavioural-state-dependent response to locations of several neurons²⁸.

1.2 The hippocampus

While the first detailed descriptions of the hippocampal anatomy has been provided by Cajal y Ramon and his student Lorente de No in the early 1900s²⁹, its function remained elusive until several decades later. Firstly, the case of the hippocampal removal of patient H.M. in 1954⁵

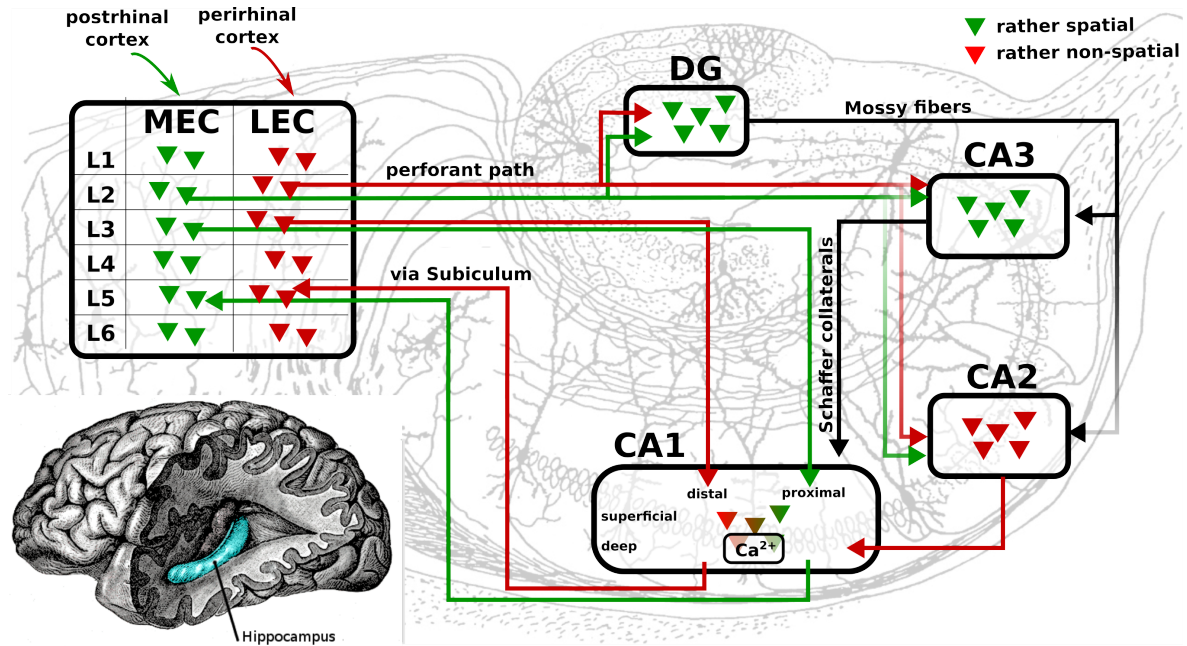


Figure 1.1: Anatomy and differing functions across the hippocampal subregions; lower left: sketch of the human brain with the hippocampus highlighted, from *Henry Gray (1918), Anatomy of the Human Body*; large figure: cross-section of the hippocampus (background from *Santiago Ramón y Cajal (1911) [1909] Histologie du Système nerveux de l'Homme et des Vertébrés*); subregions and major synaptic pathways sketched on top with activity of mainly spatial (green) and mainly non-spatial (red) content highlighted; display loosely inspired by Igarashi, 2014³⁰

followed by severe anterograde and partial retrograde amnesia sparked research which displayed the role of the hippocampus in memory formation, retention, and consolidation¹ with increasing certainty. Secondly, the seminal work by O'Keefe and his student Dostrovksy in 1971¹⁹ found neurons of the hippocampus to respond selectively to parts of the experimental environment, opening a wide field of research on place coding neurons, eventually earning John O'Keefe, as well as May-Britt and Edvard I. Moser the *Nobel Prize in Physiology or Medicine* in 2014.

1.2.1 Anatomy of the hippocampus

The hippocampus is a structure of the limbic system located beneath the cerebral cortex. In mammals exists one hippocampus on either side of the brain. It has a unique shape, similar across all mammalian species, to which it owes its naming: the resemblance to the sea horse monster in Greek mythology provides the name hippocampus, while its major subregions, CA1 to 4 are named after the ram's horns of Amun from Egyptian mythology: Cornu Ammonis.

It shows an intricate structure with a vast range of interconnectivity between its regions, which goes far beyond of the scope of this work to explain in full detail. I will merely highlight the major distinctions between different regions of the hippocampal formation and the major

pathways being of interest for the interpretation of input to the CA1 region.

Regions of the hippocampus are interconnected to form a unique system for processing information, mainly projecting from the entorhinal cortex (EC), but also from the amygdala and further subcortical regions³¹. Information is mainly forwarded through excitatory connections with a diverse network of interneurons, providing local within-region inhibition (but strong inter-regional inhibition between CA2 and CA3^{32,33}). The following description will focus on the excitatory connectivity between regions.

Pathways of connectivity through the hippocampus Hippocampal connectivity is mostly arranged unidirectionally, transversal to the longitudinal axis³⁴ through the different regions³¹, Fig. 1.1. Its function along the dorso-ventral axis ($\hat{=}$ along the longitudinal axis and orthogonal to the cross-section displayed in Fig. 1.1) changes: the dorsal side is mainly responsible for spatial navigation and memory, while the ventral side processes motivational and emotional behavior³⁵, with higher fractions of place cells and information content in the dorsal part³⁶.

Granule cells from the first region of the hippocampus, the dentate gyrus (DG), receive input from the upper layers 2 and 3 (L2/3) of EC via axons of the perforant path. DG projects further to the CA3 via axons called *mossy fibers* and CA3 then projects to CA1, via connections commonly called *Schaffer collaterals*. This three-step connectivity from EC via DG and CA3 to CA1 has been termed the trisynaptic circuit³¹ and represents a major pathway of information flow through the hippocampus. CA1 then projects to both the subiculum, which integrates information from CA1 and direct connections from the EC layer 3 to provide hippocampal output to the deep layers of the EC (L5/6), and directly to L5/6 of the EC³¹.

In addition to the trisynaptic circuit, CA1 receives major input from EC L3^{33,37} and is part of a recently suggested alternative trisynaptic loop: A range of powerful connections starting from EC L2 directly and indirectly via DG project to region CA2, which then projects to CA1^{32,33}.

Anatomical differences between the neighboring CA3 and CA2 regions are rather gradual along the proximo-distal axis in CA3, with connections from the dentate gyrus becoming less, while input from EC L2 becomes stronger³⁸. At the same time, associational connections between pyramidal neurons become stronger³⁹, while the cells themselves maintain similar properties³¹. It was found, however, that interactions between CA3 and CA2 are dominated by mutual inhibition^{32,33}

Specifics of the CA1 region CA-regions within the hippocampus are structured into different strata, each with a dedicated role. Cell bodies of pyramidal cells are densely packed in the *stratum pyramidale* with their dendrites extending into the other strata, providing specific connections from different pathways³¹. Within-region connectivity is highly diverse: while the CA3 region shows a highly interconnected, recurrent network, recurrent connections within CA1 can barely be found and appear to be mostly directed towards inhibitory interneurons³¹, suggesting distinct functional roles of the two regions.

The CA1 region has been found to be further compartmented into a superficial and a deep layer, based on electrophysiological properties like their modulation of theta-phase locking during REM sleep⁴⁰, but also functional differences. The Schaffer collaterals from the classical trisynaptic loop are found to project to either layer, while the alternative trisynaptic loop with input from CA2 projects exclusively into the deep layer³³.

A further distinction has been found along the proximo-distal axis of CA1, where differences in connectivity suggest different functional behavior: direct connections from the MEC preferentially project to the proximal part of CA1, while afferents from the LEC reach primarily towards the distal parts^{41–43}. An according gradient with decreasing spatial information and theta-phase modulation was reported along the proximo-distal axis⁴⁴. Differences in connectivity of MEC and LEC towards the deep and superficial layers remain yet to be determined.

Information on hippocampal subregions and their connectivity, as well as suggested major functional roles is summarized in Fig. 1.1.

1.3 Place coding within the hippocampus

The two major pathways of input arriving at the hippocampus carry spatial (via the postrhinal and medial entorhinal cortex) and non-spatial (via the perirhinal and lateral entorhinal cortex) information³¹. The different inputs allow for the formation of what has been termed a *cognitive map*⁴ - an internal representation of the external environment, expressed by the activity of the different areas of the hippocampus^{45,46}. This map is not only engaged during exploration through the environment, but also during recall of events that happened within this environment or action planning⁴⁷, implicating an important role of the hippocampus and its spatial representation for memory.

This map is encoded in neural activity on a population level, as has been shown in works using a Naive Bayesian decoder, which is able to reliably predict the location of a mouse from measured neural activity of the CA1 region^{15,16}. At the same time, a large number of neurons is found to display place coding on an individual level with according receptive fields. On both the population and the individual level, changes over time and in response to changes in the environment have been identified and characterized in a number of works.

1.3.1 Formation of the spatial map

Location specific activity can be found in all sub regions of the hippocampus and also the input- and output region, the EC, but shows differences in its exact functionality^{36,48–51}. For the representation of an environment, a unique set of neurons is recruited to be active, while a vast number of neurons remain silent^{52,53}, with differences in code-sparsity between the different regions. However, the population of active neurons is not a set of predefined units responsible for place coding, but changes gradually over time^{15,16} and drastically over different

environments^{53,54}. Of the active units, differing fractions of hippocampal neurons are reported to have a single or few place fields per unit³. Neurons of the entorhinal cortex, in contrast, display more abstract firing maps leading to these neurons termed grid-, head direction-, or boundary cells^{55,56}.

Several studies have reported these fields to be highly dependent on the environment: If an animal is exposed to two different enclosures, the activity of individual neurons in most of the hippocampal regions is completely uncorrelated as long as these enclosures are perceived as significantly different^{54,57,58}, and thus form an orthogonal representation (see *global remapping* below). These maps are quickly established upon first exposure on a timescale of seconds^{59–61}, but require further time and exploration of the environment to stabilize^{60,62}.

Even then, place fields are no fixed entities: Throughout exploration, they change shape and position^{61,63,64} to shift in a direction opposite to the movement direction. This could be explained by a simple model of Hebbian learning within a network in which place units coding for similar places are preferentially connected⁶³ and could be reproduced by a successor representation model⁶⁵. Place field shape changes are usually reverted after a period of rest and return to a mainly symmetrical shape⁶⁴. Furthermore, a place field in some locations does not necessarily cause a response at all instances when an animal crosses given location, expressed in reported activation reliability values as low as 10%¹⁸.

According to their dynamic nature, place fields can adapt to the shape of an environment. They align along barriers⁶⁶, scale with the size of the environment^{66,67} and adapt the geometrical shape, e.g. when a box is changed from a square to a circle^{50,57}.

Lastly, place fields are not uniquely shaped by a location, but integrate all kind of information at different magnitudes. It was found that hippocampal neurons integrate time, running distance and location with individual differences in how informative each of these parameters are, ranging from equal contributions of all factors to a single impacting factor only⁶⁸. A variety of other sensory inputs contribute to the formation of place fields, such as odors⁶⁹, colors^{69,70} or visual objects⁷¹. Furthermore, a study using a virtual environment could successfully disentangle effects of movement and visual cues to show how these effects are integrated in most neurons in a non-linear way⁷². Accordingly, selectivity and stability of place coding is strongly influenced by attentional demands⁷³, with poor environments or tasks resulting in poorer spatial representations and active, in contrast to passive, locomotion improving performance²⁸

1.3.2 Types of remapping

The balance between stability and variation of the neuronal code varies over different cortical areas. Most sensory, but also motor areas are reported to have a high degree of stability^{74–76}, especially after learning^{77–79}, while areas of the hippocampus show a lower degree of stability and higher degree of variability^{15,16,18} befitting its transient role in short term memory consolidation^{1,2}.

Differences can not only be found on a cellular level, but also on a synaptic level, where reported spine turnover rates range from $\approx 1\%$ per day in visual cortex⁸⁰ up to 15% per day in somatosensory cortex⁸¹, with the hippocampal CA1 region lying at around 5%¹². A kinetic model of spine turnover in cortical regions found that observed dynamics in neocortical regions^{82–84} could best be realized by a considerable fraction $> 50\%$ of long-term stable spines. This is supported by a theoretical work suggesting that a backbone of strong, stable synapses could reproduce observed cortical dynamics²¹. Data from the hippocampal region CA1, however, was found to be best modelled by a vanishing fraction of stable spines, suggesting all synapses to be subject to a network-wide turnover and leading to an estimated time scale of 3-6 weeks over which $\approx 100\%$ synapses are renewed¹².

How do these synaptic dynamics influence the representation of a constant environment in the hippocampus? Imaging studies have allowed the recording of several hundreds to thousands of neurons from CA1 over extended periods of up to months. This has enabled the analysis of baseline dynamics in the place representation, identifying substantial changes over according time-scales of synaptic turnover^{15–17}. These studies have found a gradual turnover in the spatial representation, causing a consistently decreasing similarity over the complete time measured. At the same time, the existence of highly stable neurons, maintaining their function over up to several months has been reported in earlier studies, using electrophysiology^{14,85}, challenging the notion of a complete turnover over weeks.

Similar analyses of long-term dynamics in other hippocampal regions are yet to be conducted, but studies on a scale of few days have revealed differing degrees of stability in other hippocampal areas^{18,20,86,87}, further detailed below. A region's temporal stability, together with its response towards changes in the environment allow for conjectures about its supposed role in the hippocampal place coding. CA1 integrates the output of all of these regions, suggesting that understanding of the role of its upstream regions helps to apprehend CA1 functionality.

Before focusing on the description of separate regions, I will briefly outline the different responses hippocampal place maps were reported to show in a wide range of studies. A recent review has categorized these into three major types of responses⁸⁸: global remapping, partial remapping and rate remapping, all of which depend on a wide variety of factors and are usually described with respect to two (or more) different environments, Fig. 1.2.

Global remapping The largest change in hippocampal activity is observed during global remapping: Neuronal activity within two distinct environments becomes completely uncorrelated⁵⁴, with different active populations⁵³ and unrelated spatial firing in neurons active in both^{54,58,66,89}. The perception of *significant difference* thereby is strongly impacted by learning⁵⁷, or the degree of difference^{58,66,90}, where the geometrical shape seems to have a bigger impact than different materials, colors or odors⁵⁷

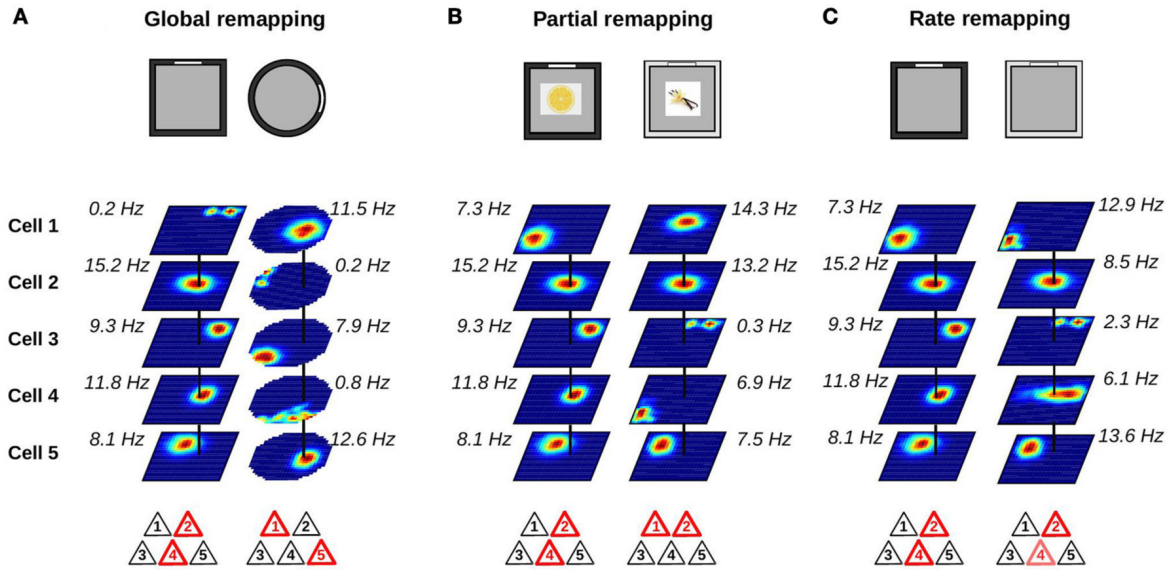


Figure 1.2: Hippocampal responses to changes in the environment range from complete orthogonalization of maps, to single neurons changing their rate response; from Latuske 2018,⁸⁸; **(A)** significant changes in the environment (such as major geometrical changes, top) induce global remapping: recruitment of a statistically independent population of active neurons (sketch in bottom) with uncorrelated place fields between the two environments (colored firing maps indicate localized average firing rate within the environment from low (blue) to high (red) values, maximum firing rate indicated next to plots); **(B)** changes of a subset of environmental details (such as odors and colors) sometimes lead to partial remapping: some neurons maintain their original response properties and some relocate their firing fields or become silent; **(C)** minor changes (such as the color of some object or wall) can lead to rate remapping, only, where the same neurons compose the spatial map with maintained place fields, but change the maximum response firing rate

Partial remapping During partial remapping, some neurons maintain their spatial firing properties, while others apparently randomly change their activity and spatial content. This behavior has been observed in an experiment, where different combinations of colors and odors were presented to rats, showing how neurons responded with place field remapping or stability, and highlighting, how different features are encoded in the activity of a neuron at the same time⁶⁹. In other experiments, the movement of environmental cues leads to partial relocation of place fields, indicating the accordance with certain aspects of the environment^{66,70}.

Rate remapping Lastly, there is rate remapping, in which neurons maintain their spatial firing across changes in the environment, but change the intensity of the response - the maximum firing rate, depending on different stimuli. This has been observed in an experiment in which rats had to decide whether to dig for a reward or not, depending on different odors presented to them⁹¹, and also in further studies, using other non-spatial stimuli^{92,93}.

1.3.3 Place coding across different regions

Entorhinal cortex In the main input region to the hippocampus, the entorhinal cortex³¹, cells with a firing modulation in periodic, grid-like patterns are reported^{34,49,55,56,94,95}, named grid cells. They commonly appear in layer 2, and less often in layers 3,5 and 6⁵⁵. Indicating their different afferent inputs, the medial (MEC) and lateral (LEC) entorhinal cortex show different behavior: The MEC shows strong spatial modulation and functional stability, even across slightly different environments^{34,50,94}, indicating a strong, spatial signal and potentially providing a reference frame for its downstream regions. In contrast, the LEC displays little spatial modulation and appears to rather code for context⁹⁶ and object-places within a reference frame^{49,97–99}. It displays strong temporal modulations, previously shown to allow inferring the passage of time over a scale of seconds to hours²⁷.

Other than grid cells, the EC is also home to head-direction cells⁵⁵, mostly observed in layers 3,5 and 6 and border or boundary cells^{100,101}.

Dentate gyrus Most input to the hippocampus arrives at the dentate gyrus, where its granule cells (GC), displaying single or multiple place fields can be found⁵⁰. The region is reported to show stable spatial activity over days and even across visually different environments¹⁸. In an experiment with a morph-box, that can be turned from a square into a circle enclosure, incrementally, it was found that the DG sensibly reacts to gradual changes in the environment with strong rate remapping of its place fields, allowing it to discriminate between even the slightest changes to the box shape⁵⁰. While large changes in the rates could be observed, a global remapping, where the population of active neurons is randomly redrawn, could not be found. These findings indicate an important role of the dentate gyrus for the discrimination of geometrical environment shapes, providing a simplistic, but stable representation of the global environment¹⁸.

CA3 The second bundle of connections of the trisynaptic loop, the mossy fibers project from the dentate gyrus to the CA3 region. The regions recurrent circuitry³¹ has led to the proposal of it acting as an autoassociative system, efficiently completing fragmented patterns^{102–104}. In accordance with this, CA3 is found to provide a stable spatial representation over days^{20,86}. Along the proximodistal axis of CA3 leading into CA2, studies have identified a gradient in neuronal connectivity, specifically a decreased number of inputs from the dentate gyrus, accompanied by an increased recurrent connectivity. This comes with a decrease in representational stability towards the CA3-CA2 border⁸⁷ and an increase in pattern completion functionality¹⁰⁵.

In further accordance with the function of pattern completion, experiments alternating between different environments found CA3 to represent boxes of similar shapes by the same code^{50,53,58,106}. Larger differences resulted in global remapping^{50,53,58,107}, suggesting a switching between pattern completion and pattern separation depending on environment similarity^{58,106}.

After an initially slow phase of formation⁵⁸, CA3 therefore provides a reliable representation of space over a scale of several days, with the ability to distinguish between similar and significantly different environments.

CA2 Following the proximo-distal axis along CA3, increasing differences in connectivity go along with differences in function. This leads into the region CA2, which is reported to play a major role in social behavior^{108–110}. Its place fields are highly dynamic⁸⁷, displaying quickly decaying activity correlations reaching a minimum after only ~ 18h and show higher similarity between different environments, than within the same environment at different time points lying only few hours apart⁸⁶. Studies suppressing CA2 activity suggest the region to play a minor role in place representation^{108,111}.

CA1 The CA1 region integrates input from upstream regions via the different pathways of the trisynaptic circuit and the alternative trisynaptic circuit as outlined above. Output from CA1 is either forwarded to the subiculum, or directly handed over to the deeper layers of EC³¹, underlining the important role of its dynamics for further processing. The place code in CA1 is reported to display some intermediate level of temporal stability, with a higher turnover than found in CA3^{20,86}, but less than in CA2⁸⁶. Due to its central role, CA1 has been the target of studies analyzing its long-term dynamics and found turnover of both, the active population and the spatial code on a time scale of several days to weeks^{15–17}.

Changes in the environment in both visual cues or geometric shape are registered by CA1 with swift changes in its spatial code⁵⁸. These responses are stronger than responses in CA3 when environmental changes are small^{53,58,107} and CA1 also shows global remapping to larger changes. In contrast to CA3, representations of different environments do not fully decorrelate¹⁰⁷ and show a significant overlap of active neurons when similar landmarks are present^{53,58}. In another study, however, varying arrangements of cues could trigger global remapping in CA1 without an according response in CA3¹⁰⁶, highlighting the different functionalities of the two regions and how CA1 is highly discriminative to slight changes in the environment.

The CA1 region is not homogeneous but is comparted into a superficial and a deep layer, with barely any interconnections³¹ and different neuronal and functional properties⁴⁰. The deep layer is found to have a higher place cell fraction and a study examining differences in the spatial code between the two layers could identify the superficial layer to show less temporal variability in its spatial code than the deep layer¹¹². This possibly originates from the connections of CA2 uniquely targeting the deep layer³³.

Further studies, analyzing differences in spatial coding along the proximo-distal axis of CA1 found the proximal part to display higher spatial information content, reflecting its direct input from the MEC, while the distal parts, primarily driven by afferents from the LEC, showed less spatial modulation and more place fields per neuron^{30,44}.

1.3.4 When, What and Where encoded in CA1

Summarizing above information, CA1 is found to integrate all kind of information from spatial and non-spatial context, with footprints of the role of upstream regions visible in different subregions, which is correlated with according connectivities. The classical trisynaptic circuit via the dentate gyrus and CA3 together with direct input from the MEC provides a temporally stable spatial reference system^{18,20}. It is supported by a report of synapses from CA3 to CA1 neurons which strengthen during experience and following periods of sleep¹⁰. However, temporal stability of the CA3 region is found to gradually decrease along the proximo-distal axis of CA3^{18,87}.

The alternative trisynaptic loop^{33,86}, providing input via DG and CA2 together with direct input from the LEC, provides non-spatial, temporally highly variable activity^{27,86}. These different types of input are integrated differently along the proximo-distal axis of CA1, as well as the transversal axes, where two distinct layers exist. The deep layer in general shows higher temporal variability¹¹² and the impact of non-spatial input is increased towards the distal part of CA1.

The interplay of these different inputs provide an explanation for the sensibly context-dependent spatial modulations with respect to the more stable autoassociative system of CA3^{106,107} and via the integration of non-spatial features offers an explanation to activity correlations maintained across different environments^{53,107}. A recent work suggested that its ongoing temporal dynamics could even serve a function in temporal coding, providing time stamps to events by gradual but systematic turnover in spatial coding over days and weeks¹⁶, similar to what has been found on shorter timescales in the LEC²⁷.

These functions suggest that, thanks to the combination of stability and dynamics in CA1, event information of *When, What and Where* can be jointly encoded in CA1²⁰.

1.3.5 Modulations of the place map by environmental salience

Starting from the very first reports of place cell existence, in which the impact of common but also exotic stimuli such as a rubber crocodile or sandpaper on neuronal activity has been researched¹¹³ it was clear that locations are not the only influencing factors to the observed neuronal activity. A wide range of works has since focussed on influences other than animal location on hippocampal neuronal activity and the allocentric cognitive map³ formed in the hippocampus^{97,112,114–125}.

Studies on the impact of visual salience on place maps consistently report visual cues such as landmarks to result in an overrepresentation of according locations in the cognitive map^{3,97,121,122,124,125}). Similar works on the impact of rewarded locations also report according overrepresentations in CA1^{62,112,116–118,120,126,127}, even though this impact appears to be less reliable and dependent on other parameters. A study by Gauthier and Tank¹²⁰ suggests, based on reports from other works, a reward overrepresentation to be dependent on (1) the location not being additionally marked by visual cues¹¹⁷ and (2) frequent shifts of the reward location within the environment¹¹⁸. Furthermore, reward overrepresentation might occur especially during learning^{62,112,117,128}. The study even provides evidence for a separate, small population of

reward coding cells in CA1 (1 – 5%), which reliably code for reward locations across different environments¹²⁰.

Environmental salience is not only found to have impact on the place map density, but also on single neuron stability, coming with an increase in place field stability at landmarks¹²¹ and rewards^{112,118}. This goes along with differences in the superficial and deep layer of CA1, with the latter providing a representation more predictive of animal behavior¹¹².

1.4 Virtual reality

Virtual reality (VR) systems have become increasingly popular for experiments with rodents, as they allow for a tightly controlled environment and enable manipulations that would be difficult or impossible to achieve within a real environment. Furthermore, it allows keeping the animal in place for head-fixated measurements enabling two-photon imaging of neuronal activity. Since the first successful setup of such a VR system¹²⁹, the impact of a virtual instead of a real environment has been examined in a number of studies^{72,130–132}. A key factor to a successful setup of a VR has been suggested to be a panoramic display screen stimulating a major part of the wide visual field of rodents. It was conjectured that animals did perceive smaller monitors, covering parts of their visual field, only, as objects within an environment and not as the environment itself¹²⁹. Toroidal or cylindrical monitors covering 300° up to 360° of the visual field have therefore become popular^{130–132}, but experiments with large 2D monitors have also been shown to succeed^{62,133,134}.

The second crucial part to a successful VR setup is the treadmill¹²⁹ allowing rodents to move despite their head-fixation, and the coupling of the VR with the information from movement to allow for a natural navigation. Animals in such a setup thus have access to information from vision and movement, which in most cases is reported to be enough to reproduce hippocampal activity similar to real-world-like experiences^{72,130,131}. However, studies found that considerably less neurons were engaged in VR experiments¹³⁰ and that vision and movement combine in a non-linear fashion to produce place fields, with minor and major deviations from realistic stimuli or even conflicting information leading to a reduced number and reliability of place cells^{72,130}.

1.5 Overview of this work

This thesis comprises a large scale data analysis of mechanisms during hippocampal place map formation and maintenance. It builds on a number of previous works that have determined different dynamics across the hippocampal fields, especially the CA1 region. The work identifies different processes contributing to the overall turnover observed in the spatial map over time and thereby contributes to the understanding of how the relevant functions of memory and navigation are not only encoded, but also maintained in the hippocampus. The thesis is organized as follows:

Ch. 2 introduces the experimental setup as well as the methods for neuron detection, neuron registration across different sessions and place field detection applied to imaging data. I introduce

a previously published neuron detection method using constrained non-negative matrix factorization (CNMF), which is applied in this thesis by the publicly available toolbox¹³⁵ and used in a novel way to identify the fate of undetected neurons. Detected neurons across different sessions have to be registered to one another to allow tracking of identities over long times. Probabilities of neuron footprints belonging together are computed based on a probabilistic model, adapted with major changes from an earlier work¹⁷. Matching is then performed by an algorithm inspired by above toolbox¹³⁵. In the last part I introduce a new method for place field detection based on a Bayesian hierarchical model¹³⁶, using a method commonly applied in astrophysics for exhaustive exploration of the parameter phase space: nested sampling¹³⁷. It characterizes neurons as place or non-place cells based on the Bayes factor (ratio of model evidences) and provides posterior distributions for place field parameters. Furthermore, it allows for the detection of multiple place fields per neuron and introduces a way for the individual characterization of trials in which place fields are active or inactive.

Above introduced methods are tested and applied in Ch. 3. The place field detection algorithm is applied to a large set of artificially generated neuron activity, providing parameter-dependent estimates of detection quality by using classification parameters of recall, precision and false positives. The detection of neurons from imaging data and their registration across different sessions is further elaborated by an analysis of resulting neuron parameters and neuron movement across different sessions. Data from multiple recording sessions allows further investigating the fate of undetected neurons in each session by inferring their anatomical position. I show how this can be used to confirm silence in most cases, but also how it results in a 20 – 30% increase of active neurons per session. Finally, I present resulting data for six different mice used for further analysis in Ch. 5, 6.

Ch. 4 is a published work, which investigates the mechanisms of place map formation during initial exposure of an animal to a novel environment in a virtual reality setup⁶². It focuses on the mechanisms underlying the expression of overrepresentations of salient features such as a landmark and an unmarked reward location in the hippocampal spatial map, where previous works have merely reported on the existence of such increased densities near salient features^{3,114,116,117,120}. The chapter uses data of neurons tracked across consecutive sessions to identify position-dependent rates of turnover and stabilization of place fields throughout the course of learning and finds how overrepresentations are generated by locally increased rates of stabilization. Differences in the encoding of positions corresponding to landmarks and rewards are identified from different time courses at which overrepresentations are created. Furthermore, experiments using randomized reward delivery or random presentation of landmarks identify how the increased density in the hippocampal map is composed of neurons with different response properties: neurons coding for the expectation or reception of a visual or reward stimulus and neurons coding for the place itself. A further experiment finds how the overrepresentation of landmarks vanishes in Shank2-deficient mice, while the reward remains overrepresented, indi-

cating different underlying mechanisms for how the two cues are integrated in the hippocampal spatial map.

The following chapter 5 explores the dynamics of the spatial representation under baseline conditions, which is after mice are extensively familiarized with a simple, non-changing environment similar to previous studies^{15,16,18}. What are the underlying mechanisms on a network level contributing to the turnover of the spatial map over time? Are there distinct contributions which can be described independently? And what are common features across different animals? Firstly, the trial-based characterization of place field activity allows me to identify within session turnover of the place map. I then analyze changes in the role of a neuron between sessions to characterize short-term and long-term turnover on three different levels defined by neuron activity, existence of a place field and place field location. On two levels - neuron activity and place field location - long term effects reaching over several dozens of recording sessions can be identified, indicating traces of long-term stability in the spatial map despite stronger short-term turnover effects in the maintenance of a place coding than reported in upstream regions CA3 and MEC. While I find qualitatively similar results across different mice, there are strong differences in quantities of neuron and place field recurrence across animals. I analyze how neuron parameters and intermittent contributions to the spatial map are indicative of subsequent place field stability and how silence may play a role in preserving neuron function. Finally, I characterize the dynamics depending on different axes of time along which hippocampal dynamics may operate. I am not able to find conclusive evidence of any time-axis being the defining factor, which further elaborated on in the discussion.

The final chapter 6 departs from a network-level description of the dynamics to elucidate the role of place field location and individual contributions of neurons to the observed turnover dynamics. While the previous chapter was able to capture averaged effects, this one asks whether stability and other effects described on a network level can also be identified on a single neuron level. I analyze how place field parameters and transition probabilities from and to place coding sessions are modulated by the place field location to find how in almost all mice only the edges of the linear track lead to significant modulations of these properties. I then define parameters for the individual stability of a neuron in activation, expression of a place field and place field location, each, and how these behave on a short time scale of few sessions. The analysis allows me to describe how observed dynamics are not carried by a specific subpopulation of neurons, but by an overall turnover on a network level. I can identify characteristic time scales of reliable contributions of individual neurons to the overall network activity, which can then be compared to timescales from synaptic turnover and dynamics of upstream regions in the discussion.

The last chapter 7 summarizes the results of this thesis and discusses how observed dynamics on a network and individual level might indicate footprints of the dynamics of input regions upstream to CA1, or the effect of underlying synaptic turnover, or both. Finally, I lay out a number of promising directions for future analyses on hippocampal dynamics and how the role of CA1

with its particular dynamics can be further elaborated.

Chapter Two

Methods

2.1 Content

Calcium imaging has become a popular method for recording neuronal activity. It enables the recording of hundreds, up to several thousands of neurons at the same time, and provides visual information about neuron positions, helpful for the registration of the neuron identities across different recording sessions. Advancements in measurement techniques themselves have gone hand in hand with the development of analytical methods for automated extraction of neuronal components from recorded images. Especially throughout the last years, a number of toolboxes for neuron detection have been published^{135,138,139}, aiming to provide a ready-to-use analytic pipeline.

In this work, analyzing long-term dynamics not only on a network-, but also on an individual neuron level, it is important to have a high performance, and good understanding of neuron detection and identification. Further, the specific focus on spatial coding highlights the importance of place field detection from neuronal activity.

This chapter will introduce the experimental and analytic methods applied in this work to prepare the data for further analysis. The first sections will introduce the experimental setup and data collection (Sec. 2.2) as well as preparation of data for neuron extraction (Sec. 2.3). Sec. 2.4 will present the workings of the neuron identification via the toolbox CaImAn¹³⁵ based on the CNMF method^{140,141} and briefly discuss parameters obtained for component classification. Registration of neuron identities based on a previously published model¹⁷, but with some major changes applied in this work is presented in Sec. 2.5. Finally, I will present an algorithm for place field detection, using a hierarchical Bayesian model¹³⁶ (Sec. 2.6), which introduces major changes towards more commonly used place field detection methods, such that it characterizes place fields using a Bayesian approach, obtains posterior distributions of field parameters, allows the detection of multiple place fields per neuron and the identification of trials within a single session, in which the field is active.

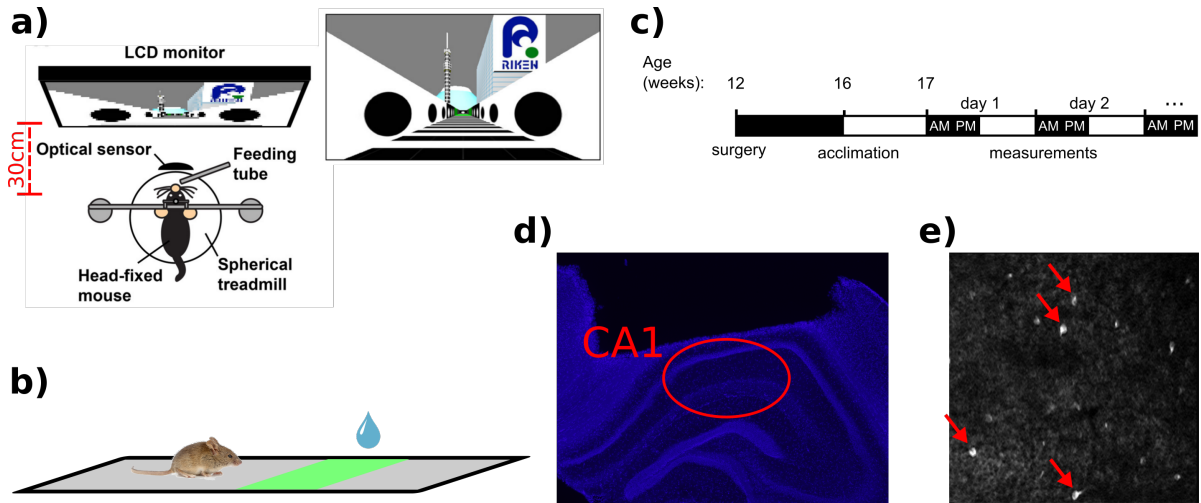


Figure 2.1: Experimental setup for data used in Ch. 5 & 6 and examples from the recording region; **(a)** left: setup with the mouse headfixed on top of a styrofoam spherical treadmill in front of an LCD monitor; right: example of the virtual environment as seen by the mouse; from [134]; **(b)** schematic of the virtual linear track with rewarded location at the green zone; **(c)** time schedule of experiments: surgery is performed after in ~ 3 month old mice, allowing for a 4 week recovery period followed by acclimation to the experimental setup. Mice are held in a reversed day-night rhythm, such that recording sessions at AM and PM (~ 4 h apart) are performed during the night cycle of the animal, 1-2 times per day; **(d)** coronal section of the hippocampus with highlighted CA1 region; top part is the removed cortical area to allow for implantation of an imaging window (see text); **(e)** example of a single imaging frame from Ca^{2+} -imaging. High intensities show high density of Ca^{2+} -ions, corresponding to neurons emitting an action potential (a few indicated by red arrows)

2.2 Experimental setup

Data used within this study has been obtained at RIKEN Brain Science Institute, by experimenters for the data in Ch. 4 detailed in the according chapter and for the data in the other chapters by Kotaro Mizuta and Yukiko Sekine. The data was made available for this work within the scope of a collaboration between Kyoto University and the Max Planck Society. All experiments were conducted in accordance with institutional guidelines and protocols approved by the RIKEN and Kyoto University Animal Experiments Committees. Procedures of mouse-line preparation, surgery, VR setup and imaging are explained in detail in the methods section of Ch. 4, but I will briefly introduce the main points here and highlight differences in the setup between the content presented in Ch. 4 and the other chapters.

Mice and Surgery Adult male Thy1-G-CaMP7 mice expressing the calcium indicator G-CaMP7^{142,143} and DsRed2 in hippocampal pyramidal neurons were used for the experiments. Mice were at least 12 weeks old at the beginning of the surgery and housed for a ~ 4 week recovery period after surgery. Surgery consisted of the placement of screws within the skull to

provide anchors for a head plate and the implantation of an imaging window (stainless steel ring) to be able to optically access the hippocampus, by removing part of the overlying cortex (for details see Ch. 4, Methods Details)

Virtual environment A virtual environment (VR) was constructed from a freely rotating styrofoam ball (20cm \varnothing) in front of a wide screen LCD monitor (30cm distance), covering 81° horizontal and 51° vertical visual field of a mouse placed on top of the ball, sufficiently large to cover the majority of the mouse's binocular and monocular visual fields¹⁴⁴. The remainder of the environment was covered by a black curtain. Throughout recording sessions, mice were headfixed in this position via the implanted head plate, which was screwed into a rigid cross bar and posts. Movements of the ball were registered with an optical mouse and transmitted to the computer controlling the VR, to move the animal position accordingly. Position and velocity of the mouse were registered and stored in 20ms intervals, with frames registered to according images recorded by the microscope. Water rewards (5 μ l/reward) were delivered from microdispenser unit positioned directly in front of the animals mouth. Reward events were triggered by the animal entering a reward zone within the VR.

Data of Ch. 3, 5, 6 is from a VR different from the one described in Ch. 4 (which is described in detail in Ch. 4⁶², Methods Details) and described as follows¹³⁴. The track segment was 120cm long, measured as the number of rotations of the ball required to move from one end of the track to the other, multiplied by the circumference of the ball. The width of the track was set at 10cm, although the mouse moved only one- dimensionally along the midline of the track with its view angle fixed toward the direction of movement. The track was divided into two segments and the walls of each segment were textured as follows: black circles on white background for 0-40cm; white circles on black background for 40-120cm. The floor was textured with black and white stripes aligned perpendicular to the direction of movement. The space above the track was gray. Three large objects (*tower*, *mountain* and *building*) were placed outside the track so that the mouse running the track could use them as global cues. The target zone was defined at 60-80cm from the origin and denoted with green panels on the walls and floor as local cues. When mice reached the end of the track, movement within the VR was inhibited for 4 seconds, while the mouse was still able to run on the treadmill. At the end of the stopping period, the mouse position within the VR was teleported back to the beginning of the track to start a new trial. A reward was delivered after a delay period of 0-2secs (depending on recording session, see Sec. 3.4) after entrance to the reward location, if the mouse was still within the reward location after the delay period. Licking times were not registered.

Experiment Before the start of recording sessions, mice were habituated to handling from the experimenter and the setup by being placed in the experimental setup, with the LCD deactivated. Recording sessions started after a few days of habituation and from the very first encounter of the mouse with the VR. Data presented in Ch 5, 6 is from six different mice, starting earliest

from session 15, to avoid recording initial effects (further discussion in Sec. 5.3.1). Throughout the sessions used for the analysis, the VR was set up according to above description. For some mice, the prior sessions included relocation of the reward position, in which case the last change to the environment was at least 3 sessions before the first session included into the analysis. Sessions were recorded 1-2 times during weekdays per day during the morning and the afternoon, with same-day sessions spaced ~ 4 h apart and sessions between consecutive days ~ 20 -24h apart. Weekends resulted in 68-72h intervals between two sessions.

During recording periods, mice were held on a water restriction schedule and checked daily to ensure that they maintained at least 85% of their preoperative body weight and exhibited no signs of abnormal behavior. Mice were held in standard, non-enriched homecages in single breeding after surgery to avoid damage of the imaging window by other mice and were subject to a reversed day-night rhythm, such that experiments performed during the day occurred during the night-cycle of the animals. Animals were consistently treated by the same experimenter. A session consisted of 10 minutes of imaging, while the mouse was allowed to freely behave within the VR, and usually completed several runs (trials) through the linear track.

Imaging Imaging was performed using a 2-photon microscope mounted on top of the head-fixed mouse (further details in Ch. 4 Method Details). The microscope was focussed at a depth approximately $150\mu\text{m}$ from the hippocampal surface, imaging from excitatory, pyramidal neurons in the deep layer of CA1 in an intermediate proximo-distal location. A field of view (FoV) of $532 \times 532\mu\text{m}$ was recorded by 512×512 pixels at 15Hz frames per second, resulting in a total of 8989 frames for each session (~ 10 min). Prior to each recording session, the signal from DsRed2 was used to align the FoV to a previous session, using blood vessels and characteristic neuron arrangements as reference.

2.3 Motion correction

The headfixed setup of the mouse during sessions allows to keep the FoV stable and in position throughout a session. Some movement on a micrometer scale, however remains and has to be corrected for. It comes in two flavors: rigid motion, in which frames as a whole are shifted in x-y-direction with respect to a reference and non-rigid motion, which often comes in wavelike patterns during animal movement on a timescale shorter than the scanning frequency, as the FoV is scanned row-wise, from top to bottom.

After some denoising via a median filter with a $3 \times 3 \times 3$ -pixel neighborhood to reduce noise-contributions to the following methods, I apply an algorithm provided by Juan-Florez Weidinger to correct for the two kinds of motion.

Rigid motion correction via maximum correlation For two frames $I_1(x, y)$ and $I_2(x, y)$ at time points t_1 and t_2 , pixel intensity values are used to compute the cross-covariance $c'_{1,2}(\Delta x, \Delta y)$,

with $(\Delta x, \Delta y)$ the relative shift of the frames. This can be implemented efficiently, using that the cross-covariance of two images equals their convolution, with indices flipped in one image ($x \leftarrow -x, y \leftarrow -y$):

$$(2.1) \quad \begin{aligned} c'_{1,2}(\Delta x, \Delta y) &= \int_{-\infty}^{\infty} \tilde{g}_1(x, y) \tilde{g}_2(x + \Delta x, y + \Delta y) dx dy \\ &= \int_{-\infty}^{\infty} \tilde{g}_1(-x, -y) \tilde{g}_2(\Delta x - x, \Delta y - y) dx dy \end{aligned}$$

$$(2.2) \quad (s_x, s_y) = \max_{\Delta x, \Delta y} c'_{1,2}(\Delta x, \Delta y)$$

starting from the definition of the cross-correlation in Eq. 2.1 and denoting \tilde{g} as the mean-subtracted images g . The rigid motion between the two frames (s_x, s_y) is reversed via an image shift, using nearest neighbor interpolation.

Non-rigid motion correction via optical flow estimation Non-rigid motion is captured in terms of optical flow - a vector field describing the localized net-intensity shift in between two compared images. The applied algorithm for motion correction (implemented based on <https://xcorr.net/2014/08/02/non-rigid-deformation-for-calcium-imaging-frame-alignment/>, opened on 24.April 2020) is based on the Lucas-Kanade method¹⁴⁵ and assumes the flow field along a small batch of scanned rows of the 2-photon microscope to be approximately constant. It calculates the flow field from temporal and spatial intensity gradients between the compared images, using least squares to converge towards an optimal solution. The obtained flow field is then used to reverse the non-rigid motion, using an image shift with linear interpolation, allowing for sub-pixel registration accuracy.

2.4 Neuron detection

Throughout the development of calcium imaging, a variety of methods have been developed for automated neuron detection, ranging from segmentation based on pixel-correlation¹⁴⁶ over principle component analysis combined with an independent component analysis (PCA-ICA)¹⁴⁷ to variants of a non-negative matrix factorization algorithm (NMF¹⁴⁰). A variant of the latter, introducing a number of constraints, giving it the name constrained non-negative matrix factorization (CNMF) has been implemented in the toolbox CaImAn and has shown to produce reliable results, on par with human labeler performance¹³⁵. Throughout this work I am using this toolbox and therefore have dedicated the following section to provide an overview of its function and the parameters obtained from it, important for later component classification.

2.4.1 Detecting neuron candidates

Neuron candidates are extracted from one session at a time, using the online implementation OnACID of the CNMF-implementation of the publicly available toolbox CaImAn¹³⁵ on motion

corrected Ca^{2+} -image stacks. It searches for an optimal factorization of the recorded image stack $\mathbf{Y} \in \mathcal{R}^{d \times T}$ into

$$(2.3) \quad \mathbf{Y} = \mathbf{A} \cdot \mathbf{C} + \mathbf{b} \cdot \mathbf{f} + \varepsilon,$$

with $d = x \times y$ px the number of pixels in a single frame and T the number of frames per session. K is the number of neurons with locally constrained spatial footprints $\mathbf{A} \in \mathcal{R}^{d \times K}$ and temporal dynamics $\mathbf{C} \in \mathcal{R}^{K \times T}$. K_b is the number of background components, represented by the corresponding footprints $\mathbf{b} \in \mathcal{R}^{d \times K_b}$ and temporal dynamics $\mathbf{f} \in \mathcal{R}^{K_b \times T}$. A pixelwise independent process $\varepsilon(t)$ captures noise of various sources. Temporal components $c(t)$ (rows of \mathbf{C}) of neuron candidates are interpreted as an autoregressive process of order $k = 1$ with parameter γ_k to account for calcium indicator dynamics

$$(2.4) \quad c(t) = \sum_{k=1}^p \gamma_k c(t-k) + s(t).$$

$s(t)$ is the neuron activity deconvolved from calcium indicator dynamics, giving an estimate of spiking probability at time t , which can be thresholded to obtain spike numbers (see below). Demixing of temporal components $c(t)$ is achieved through sparsity constraints on both, $s(t)$ and the individual spatial components a_j . The algorithm is initialized by a greedy method searching for locally correlated activity in \mathbf{Y} as a first guess of neuron candidates and optimizes those by alternating updates of temporal and spatial components.

The CNMF algorithm and its methods as implemented in CaImAn¹³⁵ are explained in further detail in Appendix A.

2.4.2 Inferring neuron footprints from other sessions

Identities of components of each session s are registered with according components from other sessions (see Sec. 2.5.2 below), to allow for an identification of neurons, whose footprints were not detected in some of them. For neurons, which were not detected in session s , shape and position are inferred from the pixelwise sum of spatial footprints of the temporally closest sessions before and after in which this neuron appears, each aligned towards the reference position of session s , see Sec. 2.5.1. For each session, a CNMF run consisting of alternating temporal and spatial updates is initialized with the whole set of neuron candidates, consisting of such session-wise inferred footprints additional to already detected neurons. Initial guesses of the temporal component are provided as gaussian noise (inferred footprints) or the previously obtained calcium trace (previously identified neurons).

2.4.3 Evaluation of neuron candidates

The CNMF algorithm introduces a number of constraints, which enforce spatial and temporal components to follow assumptions, such as locally constrained cell body shapes, temporal traces

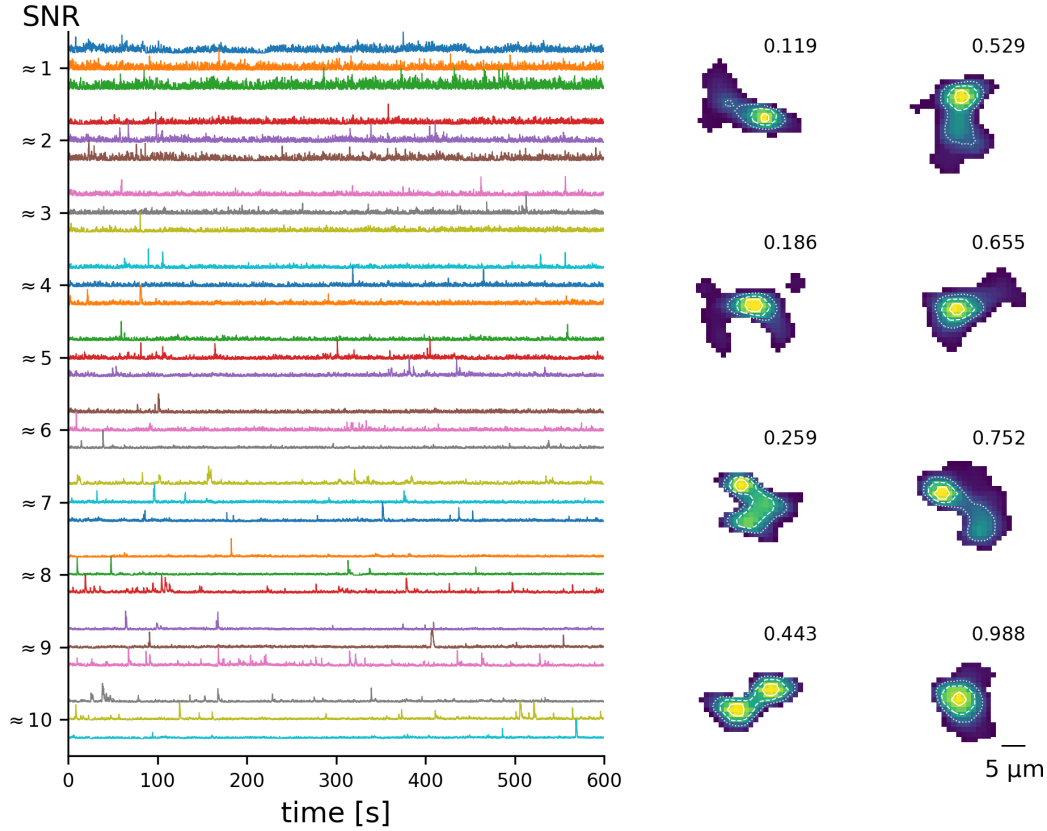


Figure 2.2: Examples of identified neuronal components at different evaluation scores; **(a)** temporal traces at differing SNR levels; **(b)** spatial footprint for all non-zero entries of a_j (colourscale), with contours at 0.3 (dotted), 0.6 (dashed) and 0.9 (solid) of maximum intensity

to include calcium indicator dynamics and uncorrelated activity between neurons, see App. A. Nevertheless, a number of quality assessments are performed by CaImAn, which I will use in this work as well and briefly introduce throughout the following lines:

Neuron shapes are evaluated by a convolutional neural network (CNN) classifier, provided in a pretrained version by CaImAn¹³⁵. Training data was obtained from manually annotated sets of neurons from hippocampal and other cortical areas. For each spatial footprint it provides a scalar value $\in [0, 1]$ describing how closely it resembles a typical neuron shape, with 1 the best match, Fig. 2.2(b). A threshold θ_{CNN} is used for classification.

Consistency with raw data for each spatial footprint is evaluated, using Eq. 2.3. Removing components a_i , c_i of neuron i from \mathbf{A} and \mathbf{C} results in $\mathbf{A}_{\setminus i}$, $\mathbf{C}_{\setminus i}$ and $\mathbf{Y}_i = \mathbf{Y} - \mathbf{A}_{\setminus i}\mathbf{C}_{\setminus i} - \mathbf{bf}$, with the latter being the raw data with the contributions of all background components and neurons, except for i removed. Intervals of a few 100ms around the first N_p events of neuron i are identified, in which it does not coactivate with neighbouring neurons and a proxy spatial footprint

$a_i^{(Y)}$ is constructed from the average over Y_i for the union of all time intervals. The spatial correlation c_r of a_i (the footprint obtained from CNMF) and the footprint obtained from raw data $a_i^{(Y)}$ is computed on a small neighborhood around the centroid of a_i and serves as a test value for this component. If it is below a threshold value θ_r , it is rejected.

Signal to noise ratio (SNR) for all neurons is computed to estimate how well the maximum amplitudes of each temporal component can be explained by baseline fluctuations, Fig. 2.2(a). It is implemented via the calculation of a pseudo-z-score

$$(2.5) \quad z_i(t) = \frac{\tilde{c}_i(t) - \text{Baseline}(\tilde{\mathbf{c}}_i)}{\text{Noise}(\tilde{\mathbf{c}}_i)}$$

with $\tilde{c}_i(t) = c_i(t) + r_i(t)$ the non-denoised calcium trace of neuron i , where \mathbf{r}_i is the residual from Eq. 2.3 at the position of component i . $\text{Baseline}(\tilde{\mathbf{c}}_i) = \text{Baseline}_c$ is obtained from a detrended $\tilde{\mathbf{c}}_i$, by sorting intensity values and finding the position at the center of the largest plateau (range of values, over which the intensity increases slowly), assuming that most of the temporal trace is given by noise fluctuations, interrupted by sparse, but large excursion during an action potential. Noise_c is obtained from values of \tilde{c}_i beneath Baseline_c , as the standard deviation of a half-normal distribution:

$$(2.6) \quad \text{Noise}_c = \frac{\text{std}(\tilde{c}_i(t_n))}{\sqrt{1 - \frac{2}{\pi}}}, \quad t_n = \{t : \tilde{c}_i \leq \text{Baseline}_c\}$$

The probability $p_i(t)$ of an amplitude of the temporal trace at time t is calculated as an average over N_s consecutive timesteps, corresponding to the characteristic decay time of the calcium indicator and provides the SNR via

$$(2.7) \quad p_i(t) = \left(\prod_{j=1}^{N_s-1} \Phi(-z_i(t+j)) \right)^{\frac{1}{N_s}}$$

$$(2.8) \quad \text{SNR}_i = -\Phi^{-1}(\min_t p_i(t)),$$

where Φ is the cumulative density function of the standard normal distribution. The SNR therefore describes how many standard deviations above noise level the maximum amplitude of each component is and can be thresholded with θ_{SNR} .

Before thresholding, components obtained from CaImAn display a wide range of signal-to-noise (SNR) values, ranging from noise only ($\text{SNR} \approx 1$), to components with easily identifiable action potentials ($\text{SNR} \gtrsim 3$, up to > 20), Fig. 2.2(a). Similarly, neurons display the whole range of shapes from barely being similar to a real neuron ($\text{CNN} \approx 0$), up to perfect neuron shapes ($\text{CNN} \approx 1$), Fig. 2.2(b).

2.4.4 Neuron activity

The deconvolved activity \mathbf{S} is obtained by interpreting the calcium traces \mathbf{C} as a superposition of a number of independent decay processes, following point events at spike times of a neuron, whose

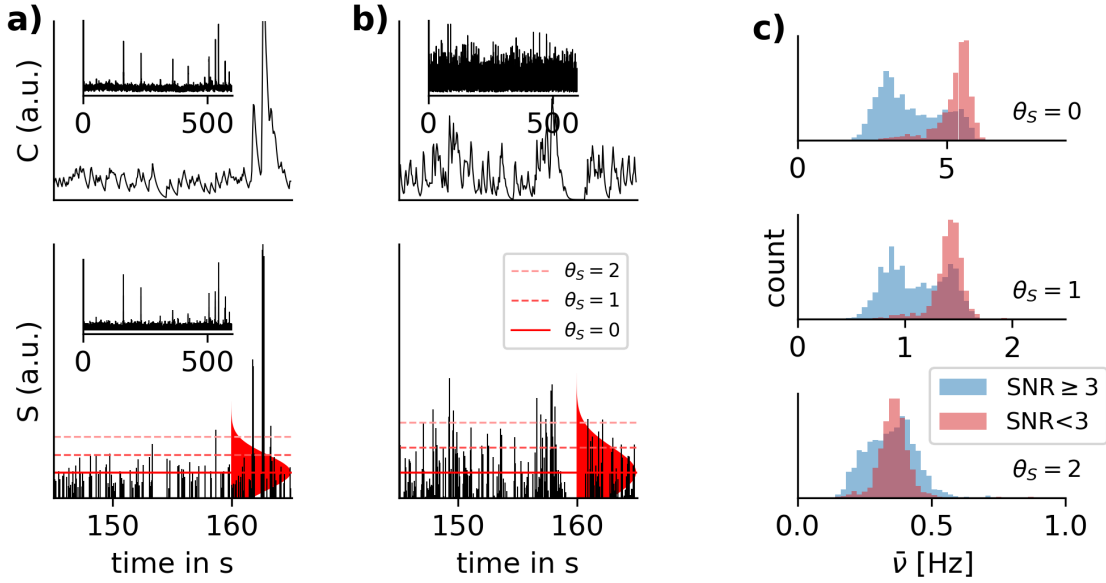


Figure 2.3: Neuron activity is detected from deconvolved signal and is sensitive to thresholding; **(a)** top: denoised calcium trace c_j of a neuron with high SNR extracted by CaImAn, inset: data of complete session; bottom: deconvolved activity s_j of the neuron above; baseline activity (solid red) obtained from median of non-zero values, dashed lines indicate 0, 1 and 2 SD above baseline; assumed normal distribution of noise fluctuations displayed in red; **(b)** same as (a) for a neuron with low SNR; **(c)** firing rate distributions within one session obtained from threshold values $\theta_S = 0, 1 \& 2$ for neurons with $\text{SNR} \geq 3$ (blue) and $\text{SNR} < 3$ (red)

intensity are registered accordingly in \mathbf{S} . The implemented OASIS-algorithm¹⁴⁸ reliably extracts large spikes at correct time points, Fig. 2.3(a), but also attempts to describe noisy patterns by a vast number of small entries in \mathbf{S} , resulting in vivid baseline activity, Fig. 2.3(a-b). An additional option of CaImAn to further refine the sparseness constraint on deconvolved temporal traces s_j to eliminate such patterns does not yield any notable improvement on removing these patterns. Therefore, when calculating the firing rate of a neuron from \mathbf{S} , I use a thresholded version as follows:

I compute the baseline activity of \mathbf{S} as the median over its non-zero values and obtain the variance from non-zero entries below baseline, according to Eq. 2.6 with $c_i \rightarrow s_i$. Assuming a normal distribution of activity values around baseline with said variance, Fig. 2.3(a,b), bottom, red gaussian, a threshold θ_S (red dashed lines in Fig. 2.3(a,b)) for distinction between spikes and noise can be expressed in terms of multiples θ_S of the standard deviation (SD).

I compute spike numbers and times according to

$$(2.9) \quad N_{sp} = \sum_t n_{sp}(t) \quad , \quad n_{sp}(t) = \left\lfloor \frac{s_i(t)}{\text{Baseline}_s + \theta_S \cdot \text{Noise}_s} \right\rfloor ,$$

allowing multiple spikes to occur per time bin, when s_i surpasses the denominator by large amplitudes. This assumes a linear relationship between recorded intensity and underlying spike

number, which holds approximately for low spike numbers^{142,149}.

The inferred spike number displays a dependence on the actual threshold. Firing rates are not only increasing with lower threshold values, but especially neurons, earlier detected to have low SNR values increase their activity due to their heavy fluctuations, Fig. 2.3(c). For the detection of place fields, I use a threshold of $\theta_S = 1$, including some of the assumed noise in S , which should average over all locations, if non-informative. It results in a high event rate of neurons with low SNR values and overall activity of ≈ 1 Hz, Fig. 2.3(c), center. It is worth noting here, that the SNR was computed on the basis of the calcium traces \mathbf{C} and before denoising, while computations here use the denoised, deconvolved \mathbf{S} .

2.5 Long-term tracking of neurons

A central problem to characterizing long-term neuronal dynamics is, to combine separately recorded sessions into a single dataset, by registering spatial footprints across different sessions to one another, describing the underlying neurons' activity at different timepoints over the course of days, weeks or months. The data available within the scope of this work contains a vast number of neurons, densely packed, and recorded in up to 100 sessions over 3 months. This results in a very rich dataset of several 1000s of neurons, but also poses some challenges on the registration of neurons to one another.

2.5.1 Alignment of recordings - a common reference frame

Before each session, the mouse is brought into head-fixation on the experimental setup. This includes readjusting the FoV, which is done manually, guided by anatomical landmarks, such as blood vessels. Still, the FoVs can deviate by several μm between sessions, additional to possible rotations around some axis, or movement in z-direction. Lastly, biological movements like shifts of the brain or brain region, or individual neuron migration can lead to non-linear deviations in the image.

A 2D-projection of the recordings provides the basis for rigid- and non-rigid motion correction between two distinct sessions. For each session, it is calculated from pixelwise summation over all detected spatial footprints, resulting in an image $g(x, y)$ the size of the FoV, see Fig. 2.4(a).

Rigid motion between distinct sessions s_1 and s_2 is obtained using Eq. 2.1,2.2 on the respective projected images from the position of the maximum Pearson Correlation coefficient $C_{s_1, s_2}(\Delta x, \Delta y) = \frac{c'_{s_1, s_2}(\Delta x, \Delta y)}{\sigma_1 \sigma_2}$, Fig. 2.4(b). A few sessions with FoV shifts $> 50px$ are rejected from the analysis, while most shifts remain at distances $< 25\mu\text{m}$; see inset Fig. 2.4(b). In some cases sessions with very bad imaging quality have to be removed as well.

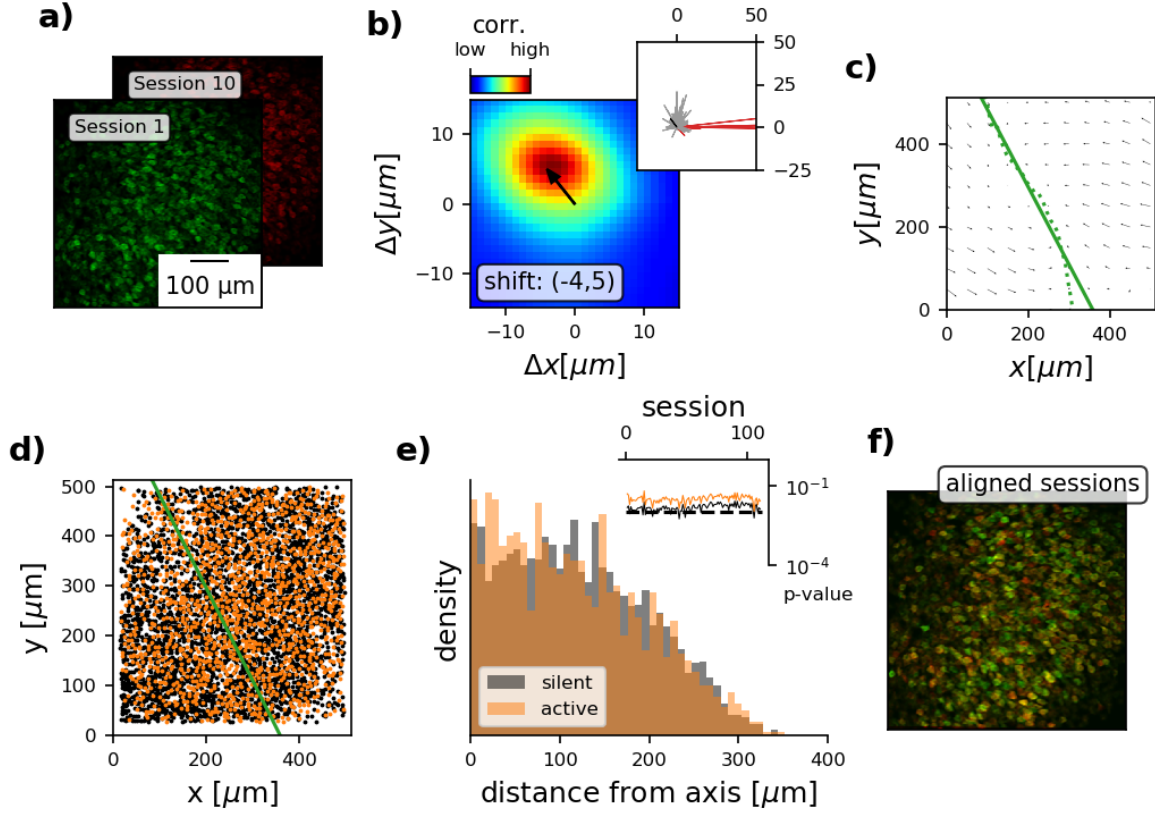


Figure 2.4: Rigid- and non-rigid alignment of session recordings show stability of FoV and provide a common reference frame; **(a)** projection of all footprints onto the FoV of two different sessions of one mouse; **(b)** Pearson correlation C_{s_1, s_2} of images in **(a)**, depending on image shift $(\Delta x, \Delta y)$; vector highlights identified rigid shift between sessions, which is determined from the position of the maximum; inset: shift vectors of all sessions with respect to first session (displayed session in black, rejected ones in red); **(c)** vector field of optical flow between rigid-shift-aligned example sessions, vector sizes are enhanced by a factor 10 for displaying purposes; green line shows the position of vanishing flow (dotted, inferred from coarse-grained flow) and a linear regression (solid), highlighting a probable axis of rotation out of the x-y-plane; **(d)** position of detected (orange) and silent (black) neurons in session s_2 , with rotation axis on top; **(e)** distribution of neuron centroid distances from inferred rotation axis in **(c)** for detected (orange) and silent (grey) neurons; inset: p-values of KS-test, whether distributions of detected (orange) and undetected (black) neurons are drawn from the same distribution as the distribution of all neurons, significance level $p = 0.01$ highlighted (black dashed line); **(f)** overlaid, aligned sessions after rigid-shift correction and reversal of optical flow, (s_1 in green, s_2 in red, thus overlaps appear as yellow)

Non-rigid deformation between two sessions which can not be captured by rigid shifts are detected using a python-implementation of the Farneback algorithm, computing the optical flow between the two projected images. It uses a polynomial approximation of pixel neighborhoods to estimate displacement fields from observed transformations between the two images¹⁵⁰ to provide an estimate of the movement. Single-pixel-estimations of displacements are averaged over a large neighborhood (here 128x128 pixels block-size), to capture general movements between the images on a sub-pixel resolution, while preventing finescale and individual neuron-position adjustments.

In many cases, resulting displacement vectors are arranged in a non-arbitrary manner, Fig. 2.4(c), indicating a governing systematic movement between the two sessions. The circular arrangement is indicative of a rotation of the imaging plane. The inhomogeneous length of vectors, however, suggest yet another movement: in many cases, it is possible to find an axis running through the FoV, at which the flow field vanishes (green dashed line), to which it is possible to fit a line (green solid line), Fig. 2.4(c). Such a behavior could originate in a rotation along this axis into z-direction - out of focus of the microscope.

In the case of such a rotation, just a few degrees would result in differences in the focus height of several μm at further distances from the rotation axis, surpassing the radius of a neuron, and possibly the height of the relatively thin deep layer of CA1 and thus losing the data of this neuron.

To test for this possibility, I calculate the centroid position of each detected neuron in the FoV, as well as the position of neurons which were not detected in this session (Fig. 2.4(d)), using the later described registration procedure, Sec. 2.5.2. Neuron centroids of these undetected neurons were obtained as an average over the centroids of temporally adjacent sessions. In the case of a rotation out of focus of whole parts of the FoV, there should be sections of the image at large distances from the rotation axis in which there are statistically more undetected neurons than expected. The null-hypothesis of detected, as well as undetected neurons being drawn from the same distribution as the overall distribution, and not follow some distance dependency from the rotation axis is tested by a KS-test. It compares the distance-dependent distribution of detected and undetected neurons (see Fig. 2.4(e)) to the joint distribution, but can not be rejected at significance level $p = 0.01$ for almost all sessions, Fig. 2.4(e), inset.

A **common reference frame** for all recordings is established by applying the reverse of rigid and non-rigid motion as identified above to the spatial footprints of neuron candidates using the python version of OpenCV's remap function with cubic interpolation, to achieve well aligned sessions, Fig. 2.4(f).

2.5.2 Matching Neurons

Individual neuron activity generally changes unpredictably from measurement to measurement and using it - or any derived measure like the neurons function or information content - for registering two neuron candidates to each other, involves strong assumptions on the continuity of

its function and thereby imposes an unwanted bias. Spatial information such as the position and shape, however, refers to the underlying anatomy, which - in the adult brain - changes at most slowly. It has therefore been used in various works to assign identified footprints to one another by methods, that require differing levels of refinement, depending on the properties of the region that was imaged^{17,135}. Sparse arrangement of neurons as found in many cortical regions can be matched with less effort, by assigning nearest neighbors to one another. Other regions, such as the hippocampal CA1-region analyzed here, in which cell bodies are densely packed in the *stratum pyramidale*³¹, however, pose additional challenges, as cells - due to the point spread function of the microscope as well as the focus length in z-direction - can appear as overlapping. In a dense region, a registration based on the distance of a pair of neurons, therefore can lead to a whole number of competing candidates for matching, of which only one should be the underlying real match.

2.5.2.1 A model of matching probabilities

An approach by Sheintuch et al.¹⁷ introduced a matching procedure based not only on the distance d , but also on the spatial correlation of neuron footprints c_{fp} . The following lines will lead through the algorithm applied by me, with changes towards the originally proposed algorithm pointed out in the later discussion, Sec. 2.8.

In a common reference frame (see Sec. 2.5.1) of temporally adjacent sessions s_1 and s_2 , for each neuron from s_1 , the distance d of its center of mass to all neurons in s_2 is calculated, identifying the ones closer than some threshold (here $12\mu\text{m}$) as candidates for being the same neuron. For each of these pairs, the maximum Pearson correlation coefficient c_{fp} for tightly cropped versions of footprints is calculated (Eq. 2.1), allowing for arbitrary shifts of integer pixel values between them.

Values of d and c_{fp} are registered into a two-dimensional histogram, along with the information whether the pairs were nearest neighbours (NN), or not (nNN), Fig. 2.5(a). The projection of histogram entries to the axis of centroid distances d reveals a bimodal distribution, with the modes composed mainly of nearest neighbours (low distance, high correlation) or non-nearest neighbours (high distance, low correlation), c.f. Fig. 2.5(a), left. The projection towards the axis of footprint correlations c_{fp} , however, shows heavily overlapping distributions of NN and nNN, with NN contributing a major part of the well-correlated footprints, Fig. 2.5(a), top.

A one-dimensional model describing distance- or correlation-distributions can be obtained by fitting a function to either of these projections, composed of a weighted sum of a NN-distribution f_{NN} and a nNN distribution f_{nNN} , where function types are chosen to represent the data appropriately, see Fig. 2.5(a) left and top. The contribution of the NN distribution to the total distribution now assigns a probability $p^{(m)}$ to each distance- or footprint-correlation-value,

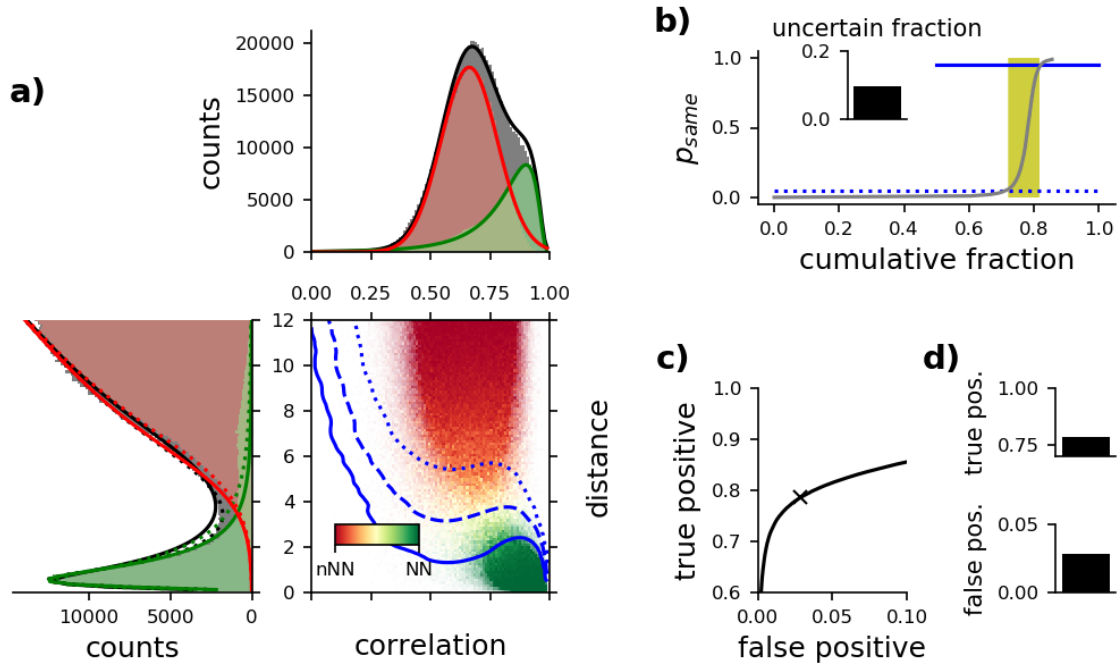


Figure 2.5: Neuron registration is based on a probabilistic model¹⁷ using centroid distance and footprint correlation; **(a)** top: empirical distribution of maximal spatial footprint correlation c_{fp} for neurons of two compared sessions closer than $12\mu\text{m}$ of each other, green: nearest neighbors (NN), red: non-nearest neighbors (nNN), grey: all values; solid lines are best fit of gaussian (nNN) and reverse log-normal (NN); left: empirical distribution of centroid distances with colors as before; dotted lines are fits to NN and nNN data separately, solid lines are functions obtained from a fit to the joint histogram; bottom right: pixelwise ratio of NN pairs to total number of detected pairs from only NN (green) to only nNN (red), linear fadeout for pair counts < 100 ; isolines of registration probabilities 0.05 (blue, dotted), 0.5 (dashed), 0.95 (solid) assigned by the model; **(b)** cumulative fraction of matching probabilities over all registered pairs ($d < 12\mu\text{m}$) for two dimensional model; probability thresholds marked as in (a); region of uncertain registration marked in yellow, range of uncertain region displayed in inset; **(c)** pseudo-RoC for joint model, estimated from the fraction of nearest neighbors that were matched; performance at decision threshold $P = 0.5$ marked by \times ; **(d)** top: fraction of true positives (of NN-groundtruth), at $P = 0.5$; bottom: fraction of false positives (of nNN-groundtruth) at decision level $P = 0.5$; data obtained over 112 sessions of a single mouse, comparing consecutive sessions, only

describing how probable it is that two footprints belong to the same neuron:

$$(2.10) \quad p^{(m)}(d) = \frac{f_{NN}(d)}{f_{NN}(d) + f_{nNN}(d)} \quad , \quad p^{(m)}(c_{fp}) = \frac{f_{NN}(c_{fp})}{f_{NN}(c_{fp}) + f_{nNN}(c_{fp})}$$

I choose a lognormal (NN) and the product of a linear- and a sigmoid-function (nNN) to approximate the distributions of distances, following the originally proposed model¹⁷. The shapes of the distributions of spatial footprint correlations are modelled by a reversed lognormal function (NN, $x \leftarrow 1 - x$) and a gaussian (nNN).

NN and nNN distributions for either one-dimensional model show differing degrees of overlap, corresponding to an area of uncertainty, whether neuron pairs should be matched to one another, or not, Fig. 2.5(b). A two dimensional model, using information available from both dimensions, allows for a clearer separation of NN and nNN distributions, Fig. 2.5(a), bottom right. Fits as above can be performed for each row or column of the histogram, providing a distribution of c_{fp} -values of $p^{(m)}(d) \cdot f_{NN}(c_{fp}|d) + (1 - p^{(m)}(d)) \cdot f_{nNN}(c_{fp}|d)$ for each distance, or vice versa, with the weights between NN and nNN distribution estimated from the one dimensional models. This process results in a model of NN, nNN, and overall distributions, described by a handful of function parameters, only.

The probability of two footprints belonging to the same neuron now follows from the definition of conditional probabilities:

$$(2.11) \quad p^{(m)}(d, c_{fp}) = \frac{p^{(m)}(d) \cdot f_{NN}(c_{fp}|d)}{p^{(m)}(d) \cdot f_{NN}(c_{fp}|d) + (1 - p^{(m)}(d)) \cdot f_{nNN}(c_{fp}|d)}$$

Either of the resulting probabilities Eq. 2.10 (one-dimensional model), 2.11 (two-dimensional model) can now be used as an input to an algorithm solving the assignment problem of neuron registration. Throughout the remainder of this work, I apply the latter, 2-dimensional model, resulting in probabilities as displayed in Fig. 2.6.

When building the model, I only include neuron candidates, passing a threshold of their signal-to-noise ratio $\theta_{SNR} = 3$ and of their raw-data coherence $\theta_r > 0.5$. To avoid a bias from registered values of neurons from extremely dense, or extremely sparse regions, I estimate the neuron density with a gaussian kernel density estimator and only include neurons within the 5 to 95 percentile of the values evaluated at neuron positions when registering distances and footprint correlations to the distributions.

In lack of groundtruth for the underlying data, Sheintuch et al.¹⁷ proposed an assessment of the performance, using nearest neighbours as a ground truth for footprints belonging together. In a dataset of dense neuron bodies, in which several neurons are usually not detected in either of two sessions, this definition of a ground truth is sure to come with a number of wrongly classified *true matches* (e.g. distant NN in Fig. 2.5(b)) and therefore has to be taken with a grain of salt. However, it gives a first idea of how well the algorithm is performing, compared to the *naive* approach of assigning nearest neighbours, RoC in Fig. 2.5(c). Applying a registration threshold of $P = 0.5$ (only neurons with larger registration probability are assigned), it shows a low rate of

false positives, Fig. 2.5(d), and an intermediate rate of true positives, as it rejects less probable candidates, which are nearest neighbors to one another, Fig. 2.5(d).

2.5.2.2 Neuron registration - an assignment problem

Given registration probabilities for each pair of spatial footprints from two different sessions, pairs can now be assigned to one another. The simplest algorithm would identify the partner with maximum matching probability for each footprint and match them together. This however ignores that some footprints may compete for the best match, especially in the case of high neuron density, or have several matching candidates with low probability only, e.g. when neurons are not detected in every session. These kind of problems can be dealt with by methods solving the assignment problem, of which I choose the Hungarian algorithm¹⁵¹ - a combinatorial optimization algorithm.

It is often formulated in terms of minimizing the cost of assigning workers to tasks, given individual cost attributions to different tasks via a matrix, where each row is a worker and each column is a task. In this case, I choose the cost to be $1 - p^{(m)}$, thus 0 when footprints have 100% certainty to belong to one another and 1 when there is no chance of them being the same neuron, according to the model built above:

$$(2.12) \quad \Omega_{s_1, s_2} = \begin{pmatrix} & \vdots & & \vdots & \\ \cdots & 1 - p_{i,j}^{(m)} & \cdots & 1 - p_{i,k}^{(m)} & \cdots \\ & \vdots & & \vdots & \end{pmatrix}$$

In the cost matrix Ω_{s_1, s_2} , each row stands for a neuron candidate from the reference session s_1 , while the columns are neuron candidates from the second session s_2 . Different numbers of footprint candidates in the two sessions are compensated for by adding additional *proxy*-neurons, that match to any other with a cost 1. The algorithm finds the set of matches minimizing the total costs by iterating through a number of well-defined steps, assigning most certain matches first and iteratively approaches less certain pairs. Applying a threshold to the individual costs of final matches ($p^* = p^{(m)}$ of according footprints) allows to reject matches with a low, or vanishing matching probability p^* (here < 0.5) and thereby accounting for neurons that are not detected in one of the sessions.




Figure 2.6: Examples of neuron matches with differing probabilities (p^* from Eq. 2.11 marked next to plots); contours of reference (black) and compared (red) footprints with contours displayed at 0.3 (dotted), 0.6 (dashed) and 0.9 (solid) of maximum intensity

2.5.2.3 Repeated matching

Spatial footprints from a sequence of datasets can be registered to neuron identities by repeatedly applying above method. I follow an approach, inspired by the matching algorithm implemented in CaImAn, but deviating in the update of the reference set (step 4) to require consistency in the shape of neuron footprints and reduce the impact of possibly false matches:

1. **Align** all footprints to a common reference frame, Sec. 2.5.1
2. **Initialize a reference set** A^{ref} of neurons from session 1 footprints
3. **Match** the chronologically next session with footprints A^{new} to A^{ref} , Sec. 2.5.2.2
4. **Update reference set** with new footprints for all neurons j in A^{ref} and A^{new} :

$$(2.13) \quad A_j^{ref} \leftarrow \begin{cases} (1 - \frac{p^*}{2}) \cdot A_j^{ref} + \frac{p^*}{2} \cdot A_j^{new} & , \text{ if match found} \\ A_j^{ref} & , \text{ if no match in } A^{new} \text{ found} \\ A_j^{new} & , \text{ if no match in } A^{ref} \text{ found} \end{cases}$$

5. **Repeat** steps 3. and 4. until the last session

2.6 Place field detection

The spatial content of neuron activity is usually inferred from the firing rate map of the respective cell. In this work a simple model of a tuning curve aims to describe the underlying data by a small number of parameters. A work of Cronin¹³⁶, introducing a hierarchical Bayesian model framework to infer posterior distributions of tuning curve parameters provides the basic idea to the method described here. A number of introduced additions allow for the identification of multiple place fields per cell, as well as to obtain the reliability of place fields, or more specifically the trials on the linear track, in which the place field was activated.

2.6.1 From data to field estimates

The position of the mouse within the linear virtual reality (length 120cm) is partitioned into $n_b = 100$ equally sized bins. Trials are identified as the time intervall between two successive *teleport* events, defined as the timepoint at which the animal position is reset from bin 100 to bin 0. Partial trials at the beginning and end of a session are not considered in the analysis. Mouse velocity is gaussian filtered (SD = 0.2secs) and used to define periods of active exploration by requiring the velocity to be $\geq 1 \frac{cm}{sec}$. The dwelltime at each bin is defined as the time the animal spent at this location during active exploration in the given trial.

Then, for each neuron individually, place field detection follows the subsequent steps:

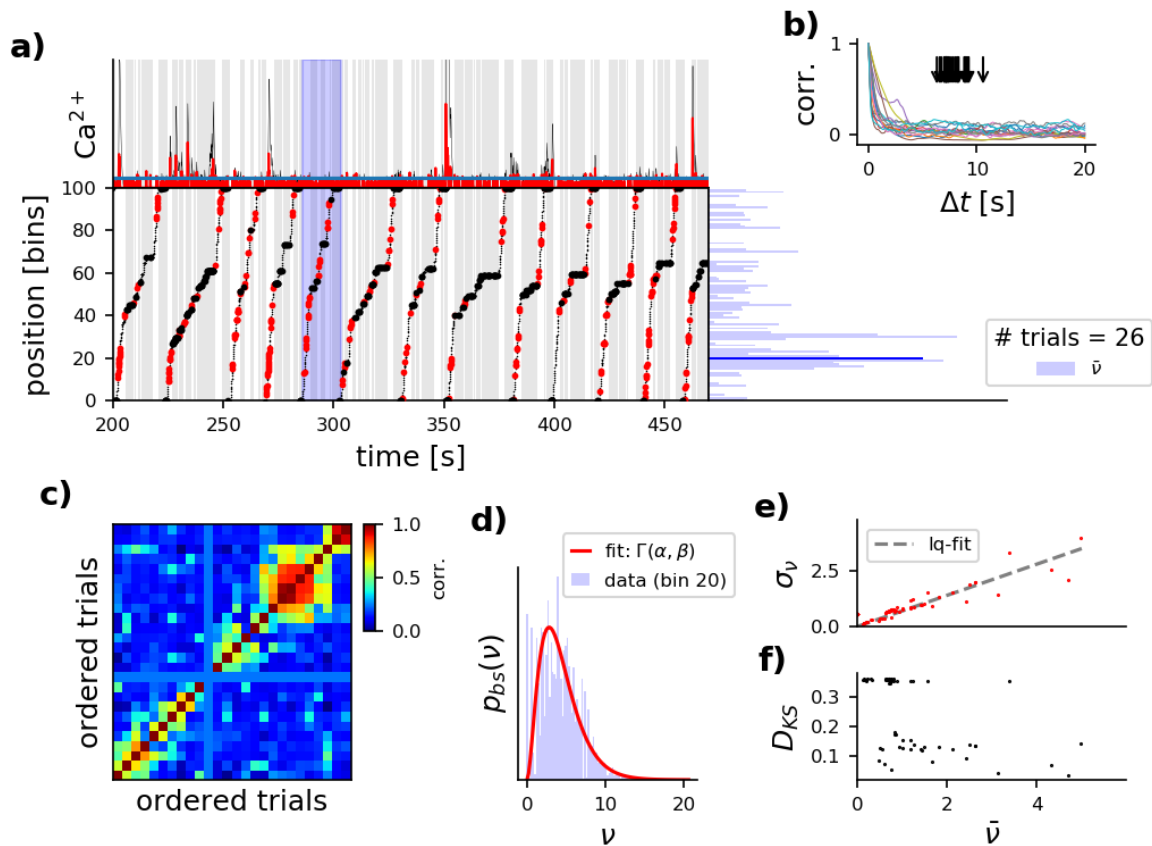


Figure 2.7: Firing maps obtained from animal behaviour and neuronal activity are used to obtain noise statistics and correlated trials for place cell detection; **(a)** top: denoised (black) and deconvolved (red) calcium activity of a single neuron over parts of a session; bottom: animal location within the virtual environment (small black dots) binned into 100 bins of equal length. Calcium activity measurements are superimposed as red and black dots with larger sizes implying larger amplitudes. I distinguish between active (white region, red dots) and inactive (grey region, black dots) periods, inferred from smoothed animal velocity, with the former requiring animal movement $> 1\text{cm/sec}$ within the VR. Calcium activity distributions for each bin (lower right plot, example in (c)) are obtained from trial averaged firing rates (a single trial is when the mouse runs through the VR once, example highlighted in blue). lower right: average (blue bars) calcium activity rate per bin from, same x-scale as (e). Highlighted in dark blue: example bin, displayed in (d); **(b)** Autocorrelation of calcium activity (20 randomly picked neurons) shows statistical independence of trials, with decay times smaller than all trial lengths (black arrows); **(c)** Pearson correlation values between firing rate maps of single trials, ordered according to clustering from a hierarchical tree (see text). Trial clusters with average correlation value > 0.5 are considered in place cell detection; **(d)** distribution of bootstrapped average calcium activities for a single bin, highlighted in (a), bottom right. Fit of a gamma distribution with empirically obtained mean and standard deviation (red); **(e)** standard deviation σ_v shows an approximately linear, but neuron-dependent relationship with mean activity \bar{v} , fit with 0-intercept as dashed line. One datapoint per bin; **(f)** maximum distance of cumulative density of empirical vs fit data of activity distributions as shown in (d)

Neuron firing rate map $\bar{v}_t(x)$ is computed for each trial as the average number of APs per second at each position bin. The deconvolved Ca^{2+} -activity data $s(t)$ is aligned with the position measurement such that each spike event can be assigned to a bin. Spike events are defined as values of $s(t)$ larger than a baseline value (see Sec. 2.4.4, median + 1 SD). APs are defined as the value of $s(t)$ expressed as the integer multiple of $s(t)$ of the above defined threshold, such that multiple spikes per timeframe are possible. Trialwise, APs occurring during active exploration are registered to the current position bin of the animal and the number of APs divided by the dwelltime yields the trial firing rate map $\bar{v}_t(x)$ of trial t for bin x .

Identification of correlated trials via hierarchical

clustering. Pearson correlation coefficients C_{t_1, t_2} between gaussian filtered $\bar{v}_t(x)$ are calculated (SD= 2bins) and used to build a hierarchical tree, based on the distance $d_C = 1 - C_{t_1, t_2}$. Closest trials with respect to d_C are assigned to one another and form a new cluster. Distance between clusters is calculated as the average distance between their trials $d = \sum_{i,j} \frac{d(u[i], v[j])}{N_u N_v}$ (with u , v trials from distinct clusters and N_u , N_v the number of respective trials per cluster) and assigned to one another, until $d_C > 0.5$, see Fig. 2.7(c) for trial correlations sorted according to the hierarchical tree. Clusters containing an absolute of 3 trials or more, and a fraction of 20% of all trials within the session are considered for further analysis.

Detection of place fields using nested sampling.

For each identified trial cluster the joint firing rate map $\bar{v}(x)$ is computed as the location-wise within-cluster AP number divided by within-cluster dwelltime. Firing rate of bins with 0 APs is estimated as $1/\text{dwelltime}(x)$; bins with 0 dwelltime are not considered. Empirical data of within-cluster firing rate maps is modelled by one model with, and one without a place field:

(2.14)

$$\bar{v}(x) = TC(x|A_0, A, \sigma_\theta, \theta) = \begin{cases} A_0 & , \text{ non-place cell, null model} \\ A_0 + A \frac{1}{\sqrt{2\pi\sigma_\theta^2}} \sum_{k=-1}^1 \exp\left(-\frac{(x-\theta+kN_b)^2}{2\sigma_\theta^2}\right) & , \text{ place cell} \end{cases}$$

A_0 is the flat baseline-activity and parameters $(A, \sigma_\theta, \theta)$ describe a gaussian-shaped modulation at some point θ , see Fig. 2.7.




Figure 2.8: Firing map of place fields is modelled as a gaussian modulation on top of a baseline rate; the four parameters of the model are displayed: field location θ , field width, σ , maximum (A) and baseline (A_0) firing rate; location-dependent fluctuations (blue distributions) are obtained empirically, using Eq. 2.16

Binwise activity is estimated to be drawn from a Gamma-distribution with mean value $\bar{v}(x)$ and SD $\sigma_v(x)$, which, together with the tuning-curve model Eq. 2.14 constitutes the hierarchical bayesian model, introduced in an earlier work on inference of orientation tuning¹³⁶.

Parameters of the Gamma-distribution are obtained from the binwise bootstrapped ($N = 10000$) mean firing rate values over the whole set of trials, where $\sigma_v(x)$ is found to be linearly dependent on the mean $\sigma_v(x) = \gamma\bar{v}(x)$ with some factor γ , Fig. 2.7(d).

The likelihood of the model, given the data is determined by a product of binwise Gamma-distributions, which are found to well approximate the distribution of firing rates over trials (see Fig. 2.7(c,e))

$$(2.15) \quad L(\vec{v}) = \prod_{i=0}^{n_b} f_{\Gamma}(v_i, \alpha_i, \beta_i)$$

$$(2.16) \quad f_{\Gamma}(x; \alpha, \beta) = \frac{\beta^{\alpha} x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}, \quad \Gamma(\alpha) = \int_0^{\infty} x^{\alpha-1} e^{-x} dx, \quad \Re(\alpha) > 0$$

$$\alpha = \left(\frac{\bar{v}}{\sigma_v} \right)^2 \approx \gamma^{-2}, \quad \beta = \frac{\bar{v}}{\sigma_v^2} \approx (\gamma^2 \bar{v})^{-1},$$

Estimation of the posterior likelihood is performed using nested sampling^{137,152,153} (for more details on the method, see App. B) with uniform priors

$$A_0 \in [0, 10] \quad , \quad A \in [2, 100] \quad , \quad \sigma_{\theta} \in [2, 20] \quad , \quad \theta \in [0, 100],$$

where θ is treated as a periodic variable.

The ratio of resulting evidences of the two models $\frac{\mathcal{Z}_1}{\mathcal{Z}_0}$ (with \mathcal{Z}_1 the place field model, \mathcal{Z}_0 the null-model) provides the Bayes factor, used as a threshold for distinction between place field and no-place field. Posterior distributions of parameters are stored for further analysis. Multiple place fields within a single neuron are allowed, if their field centers are no closer than 10 bins from one another. In the case of closer place fields, the one with lower Bayes factor is rejected.

2.6.2 Additional measures

Mutual information The mutual information (MI) is calculated as the information that is carried by individual neuron activity about the animals position¹⁵⁴:

$$(2.17) \quad MI = \sum_{i=1}^{n_b} \sum_{j=1}^{n_q} p_{i,j} \log_2 \left(\frac{p_{i,j}}{p_i p_j} \right),$$

where i runs over coarse grained position bins ($n_b = 20$), j runs over the quantized activity data ($n_q = 4$), $p_{i,j}$ is the joint probability of observing activity from quartile j while the animal is at position i , p_i is the probability of the animal being at position i ($\hat{=}$ normalized dwelltime) and p_j the probability of observing activity from bin j ($= \frac{1}{n_q}$ for all j).

Place field reliability Field location is determined as the position of the gaussian peak $\theta \pm 1\sigma_\theta$. Trials showing above baseline fluctuations ($A_0 \pm 2SD$, SD obtained from downward excursions) within this range are considered to be active trials of this place field. The reliability a of a place field is defined as the fraction of active trials within the total session.

2.7 Artificial activity data

I generate artificial activity data using model Eq. 2.14, describing the position-dependent activity $r(x)$ for any set of model parameters, Fig. 2.9. Natural behavior of the animal is included by using the position data $x(t)$ of one of the recorded sessions to generate time-dependent activity $r(x(t))$ from an inhomogeneous poisson point process $s(t)$ for each neuron.

Additional to the parameters of baseline activity A_0 and the field parameters (A, σ_θ, θ) introduced in Eq. 2.14, the model includes a randomly drawn trial-activation reliability a for each place fields. In trials with an inactive field, the respective gaussian modulation is not considered in the generation of the poisson process.

The poisson process is generated as follows: An initial set of events is generated from the maximum of the tuning curve ($r_{max} = A_0 + A$). Each of the events is then accepted with a probability of $\frac{r(t)}{r_{max}}$, depending on the instantaneous rate $r(t)$. Events are drawn from an inhomogeneous poisson distribution following the rate $r(x(t))$, with each event contributing an increment $s(t)$ at t , drawn from a gaussian distribution with mean A_s and standard deviation $\sigma_s = \frac{A_s}{4}$.

2.8 Discussion

This chapter has introduced the methods providing the data basis for any of the further work. The pipeline of neuron component detection (Sec. 2.4), cross-session registration (Sec. 2.5) and subsequent place field detection (Sec. 2.6) provides a rich dataset to analyze neuron parameters




Figure 2.9: Artificial neuron data is generated from place field model Eq.2.14; left: tuning curve for a set of randomly drawn parameters (A, A_0, σ, θ); bottom: underlying average firing rate at time t (black) generates an inhomogeneous poisson point process (red); top right: mouse location with superimposed spikes; place field is activated in some trials, only (green), drawn with probability a

and network dynamics in the further chapters. While the sections themselves were mostly restricted to presenting the methods themselves and the following chapter is dedicated to the discussion of the results from applying said methods, I want to dedicate a few lines to provide some further reasoning to the choice of methods and - where applicable - changes to earlier published methods.

Neuron activity The CNMF method provides two temporal components for each neuron: the denoised calcium traces \mathbf{C} , as well as their deconvolved counterparts describing neuron activation, \mathbf{S} , stripped off calcium indicator dynamics with their characteristic exponential shape. I chose to use \mathbf{S} in computations of firing rates and receptive fields, Sec. 2.6, instead of the commonly used dF/F , as it fulfills all the requirements of subtracted background and detrending, and - contrary to dF/F - does not contain calcium indicator dynamics. These dynamics are less important, when measuring constant stimuli, such as the orientation of gratings in the visual cortex, but are important when confronted with an ever changing set of stimuli, like locations in the virtual environment, where constant decay times of the calcium indicator interfere with the variable rate at which stimuli change, due to the movement velocity of the animal.

Neuron registration Neurons across different sessions are registered to one another, based on a model which I have adapted from an earlier work¹⁷. It provides a probabilistic substrate for any matching procedure by (1) acknowledging that statistics of spatial footprints of same neurons should be described by a different distribution than those of different neurons and (2) using information from both, footprint position and shape in a joint model instead of using information from a single variable, only. A major change towards the original model is the use of an optimally shifted footprint for the calculation of footprint correlations, in contrast to an unshifted version, even though it largely increases the uncertainty of matching based on this measure. The use of an unshifted footprint correlation and center-of-mass distance in a two dimensional model - as proposed in the original work - however ignores that these two values are highly correlated. The footprint correlation naturally decreases with increasing distance of footprints and therefore defies the idea of using neuron shape for estimating a matching probability. Sec. 2.5 shows how, even though anatomical neuron shapes are naturally correlated, footprint correlations from the same neurons clearly compose a distribution different from non-belonging neurons.

A challenge in matching neuron identities over a large number of sessions is, that a single mismatch can have a huge impact on the inferred neuron dynamics. One falsely assigned footprint will cause subsequent footprints to not match the identity of the original neuron, as well, introducing possible mis-interpretations, e.g. of place field turnover. To avoid this, I implemented a conservative update rule of the reference footprints, which maintains the neuron shape by carrying over information from previous sessions, and further minimizes the impact of matches with low probability to allow following the original neuron in subsequent sessions. The performance of the procedure will be further discussed in the following chapter.

Place cell detection Methods for detecting neurons with spatially modulated activity have used different characteristics of place fields: most commonly, place fields are identified from firing maps, by detecting locations of increased average activity^{14–16,18,85} with differing thresholds. Other works use an information theoretic approach, exploiting that neuron activity should carry significant information about its surroundings^{16,72,155}, but see Souza¹⁵⁴ for a critical discussion. Yet other approaches apply methods allowing to detect temporally changing fields, such as an adaptive point process estimation algorithm^{61,156}, recognizing that place fields can change over the course of a single measurement on the scale of seconds to minutes, shaped by experience^{61,63,64}.

Within a session, place fields were not only shown to change their location^{61,63,64,121}, but also to not be activated in a considerable number of exposures to given location^{18,157}, motivating the trial based place field detection applied here, similar to a previously published work¹²¹. This unreliability can originate from a number of sources, including an activity dependence not only on the spatial location, but also on the behavioral state²⁸, or previous experience^{158,159}. Other reasons could include the spatial code being dependent on a population code, rather than individual neuron activation, requiring reliability only on a network level.

Even without place fields, the stochastic nature of neuron activity¹⁶⁰ implies large fluctuations in neuron activity across different time bins, while the animal moves along the VR. Assuming that place cell activity can actually be described by an underlying one-dimensional function along location space, it is possible to compare the different models (place field vs non-place field) against one another. In this work, this is implemented via a Bayesian approach¹³⁶, inherently implying Occams razor¹⁶¹ and allows to quantify model evidences to provide a single quantity for place cell classification. It is implemented using nested sampling^{137,153}, an elegant method to provide both, the posterior parameter distribution and model evidence in an efficient way, while being able to identify multiple local maxima ($\hat{=}$ multiple place fields per neuron).

Chapter Three

Long-term recordings of hippocampal activity

3.1 Content

Previously introduced tools provide the means to extract neural components from raw data, collate sets of neurons from different sessions and identify and characterize functionally relevant activity. Applying these to the available data set requires some insight into the reliability and working of these, especially since the extent of the dataset with up to 100 sessions with over 1000 neurons each, surpasses most commonly presented data in recent publications.

It is the matter of this chapter to provide further insight into the working of presented methods: How do the methods perform on the given dataset? What are appropriate choices for threshold parameters? And what kind of data does the application of the tools leave me with to further investigate in the remaining chapters?

I start with the application of the introduced place cell detection algorithm to artificially generated neuron data to test the performance for a wide range of model parameters, Sec. 3.2, followed by some remarks on the detection of place fields in real neurons. Sec. 3.3 discusses the neuron detection, especially its combination with the matching procedure to test for the fate of undetected neurons. It comes with a discussion of thresholds for neuron evaluation. The results of the neuron registration procedure are discussed in Sec. 3.3.1 and I further elaborate on the distinct roles of cells in each session, Sec. 3.3.2. The last section present the data available to me after application of the methods above to the data of the six mice used for further analysis in Ch. 5 & 6 of this work.

3.2 Testing of the place cell detection algorithm

With this work, I introduced a novel way for place cell characterization (see Sec. 2.6), applying a trial-based approach, applicable to periodic, linear environments as present in the data available to me. It searches for gaussian-shaped firing-rate modulations, occurring in a subset of trials and




Figure 3.1: Example of place field detection on artificially generated activity data. Upper and lower left: as in Fig. 2.7, but showing artificially generated activity; trials with active place field are highlighted in green; in between: detected activated trials (marked in green); right: firing rate map of neuron data with original (black solid) vs inferred (black dashed) tuning curve

obtains according model parameters from nested sampling. To test the quality of this approach, I generated data of 100.000 neurons according to Sec. 2.7 with 10% cells without modulation (only baseline activity) and 90% cells with some modulation. Of the latter, 70% have a single and 30% have two modulations. Model parameters of each neuron and field were drawn uniformly from

$$A_0 \in [0, 20], \quad A \in [2, 40], \quad \sigma \in [2, 10], \quad \theta \in [0, 100], \quad a \in [0.1, 0.9]$$

Retrieval of model parameters. The activity rate parameters A_0 and A provided to the model are not actual firing rates, but activity events similar to the ones observed from the results of the CNMF method. Spike events are obtained from thresholding, following Sec. 2.4.4 with $SD = 1$, followed by the place field detection, Sec. 2.6 to obtain inferred model parameters \hat{A}_0 , \hat{A} , $\hat{\sigma}_\theta$ and $\hat{\theta}$. \hat{A}_0 and \hat{A} can only roughly be related to the underlying rates A_0 and A through a linear, though widely dispersed transformation, see Fig. 3.2(a,b). The effect of the thresholding procedure to obtain spike events becomes apparent, observing that neurons with higher median event amplitude (e.g. by regularly observing several events per time bin) in general have lower inferred firing rates due to higher thresholds (colorcode in Fig. 3.2(a-d)). Place field positions θ are reliably obtained with narrow deviations and very few mislocations (Fig. 3.2(c)) and place field widths $\hat{\sigma}_\theta$ show a linear relationship with the according input parameter σ_θ , though are spread more widely and slightly underestimate the original value, Fig 3.2(d). Lower peak amplitudes and larger place field widths usually come with a slight increase in the width of the confidence interval for place field position θ (upper part of Fig 3.2(a-d)). Finally, the activation rate \hat{a} is

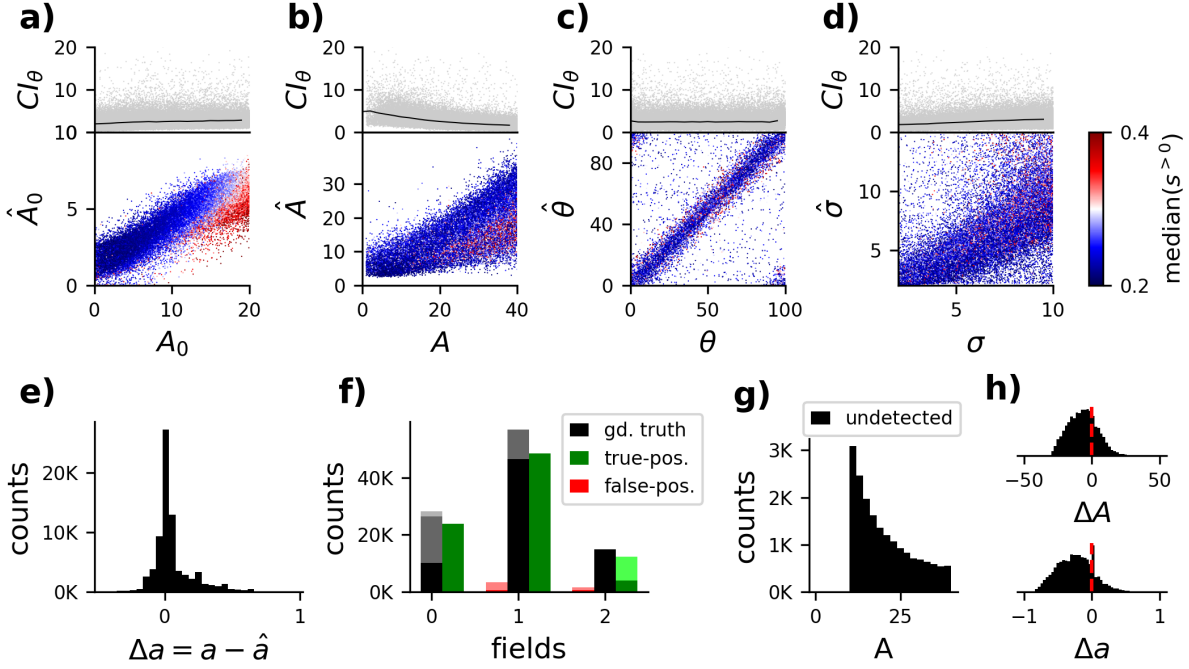


Figure 3.2: Place field model parameters from artificial data ($N = 100000$) are accurately retrieved by the detection method; **(a-d)** lower: inferred parameter \hat{x} against input model parameter x for all detected place fields (dots). Median amplitude of events (\hat{s} baseline for activity thresholding) colorcoded from single (blue) to two (red) events. Upper: width of 95% confidence interval of place field position, depending on input model parameter, average (black line) and all data (grey dots); **(e)** Histogram of activation rate difference Δa between input and inferred parameter; **(f)** ground truth and obtained place field numbers; ground truth: neurons with 0,1 or 2 generated significant fields (black), as well as one (dark grey) or two (light grey) insignificant fields. Number of neurons with true positives (green) for all (dark) or one (light, in case of multiple fields) significant field of a neuron. Number of neurons with false positive fields (red) due to non-existent ground-truth field (dark, barely visible) or an insignificant field which was correctly detected (light); **(g)** histogram of model parameter A (peak amplitude) of undetected place fields (other parameters not shown, but display homogeneous distribution); **(h)** difference of peak amplitudes A and activation a between place fields of neurons where only one of two fields was detected ($\Delta \hat{s} \hat{=}$ undetected minus detected); Insignificant place fields defined as $A < 10$ or $a < 0.1$; detected place fields defined as $\hat{A}_0 > 1\text{Hz}$, $\hat{A} > 5\text{Hz}$, $\hat{a} > 0.1$ and Bayes factor > 20 (see Fig. 3.4)

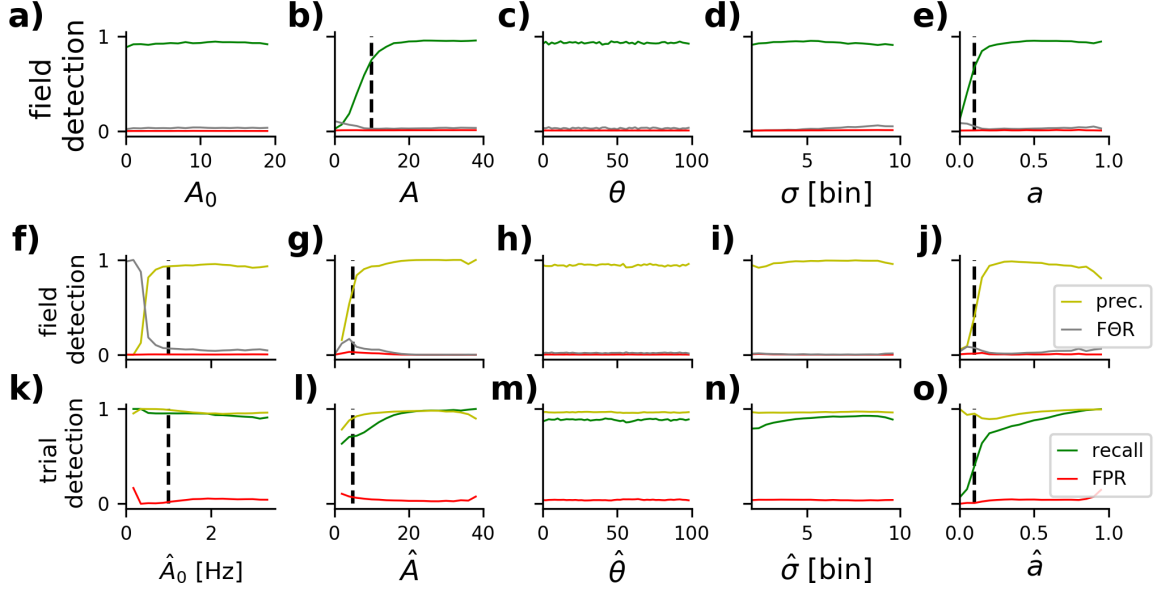


Figure 3.3: Detection performance of place fields and active trials is high, but place field parameter dependent; **(a-e)** recall, FPR and F0R of place field detection depending on input model parameters, thresholds used for classification as significant place field marked as dashed black line; **(f-j)** precision, FPR and F0R of place field detection depending on inferred parameters, thresholds used for place field detection marked as dashed black line; **(k-o)** recall, precision and FPR of trial detection, depending on inferred parameters; Rates computed based on significant fields and detection thresholds as specified in Fig. 3.2, unless according parameter is varied in respective plot

mostly well reproduced by the detection method, though in some cases underestimates the actual rate, Fig. 3.2(e). Overestimation of the actual activation rate usually happens by a few percent points, only, indicating one or two false positive trials.

Classification of detection quality. Within the set of artificially generated place cells, a number of cells have parameters which barely qualify as a proper place field. To avoid including insignificant fields in the assessment of detection quality, I require fields to have at least $A > 10$ peak amplitude and activation values $a > 0.1$, to be considered a place field. For classification of the analysis, I use the standard measures *recall* and *precision*, as well as the false-positive rate *FPR* and the rate of fields detected in a wrong position *F0R*:

$$(3.1) \quad recall = \frac{tp}{tp + fn + f\theta}, \quad precision = \frac{tp}{tp + fp + f\theta}, \quad FPR = \frac{fp}{fp + tn + f\theta}, \quad F\theta = \frac{f\theta}{tp + fn + f\theta},$$

with the number of true-positives tp (correctly detected place cells), false-negatives fn (neurons with undetected place fields), false-positives (neurons with incorrectly detected place fields), true-negatives tn (correctly identified non place cells) and misplaced fields $f\theta$ (original field center not within 5 bins or 95% confidence interval of detected field). Note, that a place field is only

classified as false positive, if the neuron has no other field, which is not yet assigned to another detected field, such that only neurons with more fields detected than actually present can count into this statistic. All other wrongly placed fields are assigned to the closest place field of the neuron which should have been detected, but was not and are counted towards $f\theta$.

While the false positive rate of place field detection remains very low for almost all parameters $< 1\%$, recall, precision and $F\theta R$ show some dependence on the parameters. In the following lines, the recall and precision value are discussed for the parameter space of the underlying model and the parameter space of the inferred model, respectively. Note that the parameters have different axes for the two rates, as there exists no exact match of model-parameter to inferred parameter - most strikingly caused by the nonlinearity in the applied event detection, Sec. 2.4.4.

Recall values remain high for most of the parameter space of generated place fields, showing values of $> 90\%$, Fig. 3.3(a-e) with a strong dependence on A and α , for which it rapidly drops to zero at low values.

Neurons with two fields take on a special role in the assessment of place field detection - specifically the recall. While for most of these neurons, at least one field is identified, the detection algorithm often misses out on the second one, Fig.3.2(f). Peak amplitude and activation of the missed out field on average are below the detected one, Fig.3.2(h), showing that usually the more prominent field is identified. This can be explained by how the detection algorithm works: it identifies correlated trials and searches for a single field in the average firing map of those. As trials are initially assigned to a unique cluster, they are usually counted towards the more prominent field, suggesting why two-field place cells are only partially identified.

Precision values are near perfect for most of the parameter-space of inferred parameters, dropping for low values of inferred baseline activity $\hat{A}_0 < 1$, as well as low activation rates $\hat{a} < 0.1$, Fig. 3.3(f-j). The former effect is due to highly sparse activity, causing a rugged firingmap in which localized fluctuations lead to pseudo place fields, wrongly detected by the algorithm, shown by the strong increase in $F\theta R$. The latter, on the other hand, is an effect of the detection algorithm, testing for a minimum of 20% correlated trial-firingmaps and, on the other hand the fact that it is easier to miss a single or few active trials due to unfavorable noise or poisson-process-realization, than it is for a larger number of trials.

As a last step, I analyzed the impact of the Bayes-factor (ratio of place cell vs non-place cell evidence) as a threshold for neuron detection, Fig. 3.4. Commonly,




Figure 3.4: Detection performance controlled by Bayes factor of field vs. non-field model; Recall, precision FPR and $F\theta R$ (colors as in Fig. 3.3) as a function of the Bayes factor. Applied threshold for later place field detection marked as black dashed line; detected place fields defined as $\hat{A}_0 > 1\text{ Hz}$, $\hat{A} > 3\text{ Hz}$, $\hat{a} > 0.1$

values of 1/2 or 2 are treated as sufficient evidence for a model¹⁶². In this case however, I find that both, precision and recall saturate at $BF \approx 20$, only. At the same time, the rate of falsely localized place field FθR falls off at $BF = 20$. Note that the ground truth is calculated from all neurons inferred to have a place field within this range of Bayes-factor values, therefore misses out on all neurons for which the threshold of sufficient correlated trials was not met and therefore no nested sampling was performed and thus slightly overestimates the actual recall value.

Summarized, the analysis of the place field detection method provides a clear estimate of thresholds for place field detection. Aiming for large precision and recall values, I require $\hat{A}_0 > 1Hz$, $\hat{A} > 5Hz$, $\hat{a} > 0.1$ and $BF > 20$.

Trial detection. The analysis does not only allow for a characterization of place fields, but also allows to identify the trials in which the place field is active. For correctly identified place fields, the precision of trial identification is very high for almost the complete parameter space, Fig. 3.2(k-o). The recall is negatively affected by low peak amplitudes A and low reliability α , but approaches almost perfect ratios for high values thereof, Fig. 3.2(l,o). Place field width barely impacts the recall value of trial detection, Fig. 3.2(m) and low inferred reliability α often seems to miss out on identifying a few activated trials, causing a low recall value, Fig. 3.2(o).

3.2.1 Place field detection in real data

Multiple modes I consistently find a fraction of up to 10% of neurons with multiple place fields, Fig. 3.5(a) - in the vast majority two. Field positions of multi-mode place cells are not significantly located around some preferred location and are similarly distributed as the location of single-mode place cell fields, Fig. 3.5(b).

The place cell detection algorithm rejects place fields closer than $\Delta\theta = 10$ bins within the same neuron, reducing repeated identification of the same place field. The distribution of place field distances within multi-modal place cells, however, shows an increased fraction of place fields with $\Delta\theta < 20$ bins. This could be the case, when place fields exhibit changes in shape or location between trials, as reported in an earlier study, quantifying inter-trial variability of place fields¹²¹, and originate from within-session experience as reported before^{61,63,64}. This is further supported by an increased trial coactivation coefficient c_a of nearby place fields, while more distant place fields are usually separately activated (low correlation coefficient c_a , Fig. 3.5(e)).

3.3 A complete set of neurons across all sessions

Spatial footprints and calcium activity of neurons are obtained from the available recordings of calcium activity, using the CaImAn toolbox¹³⁵, Sec. 2.4, which compiles a number of well-documented steps of the analysis procedure^{135,141,148,163}. It comprises a readily accessible interface, including

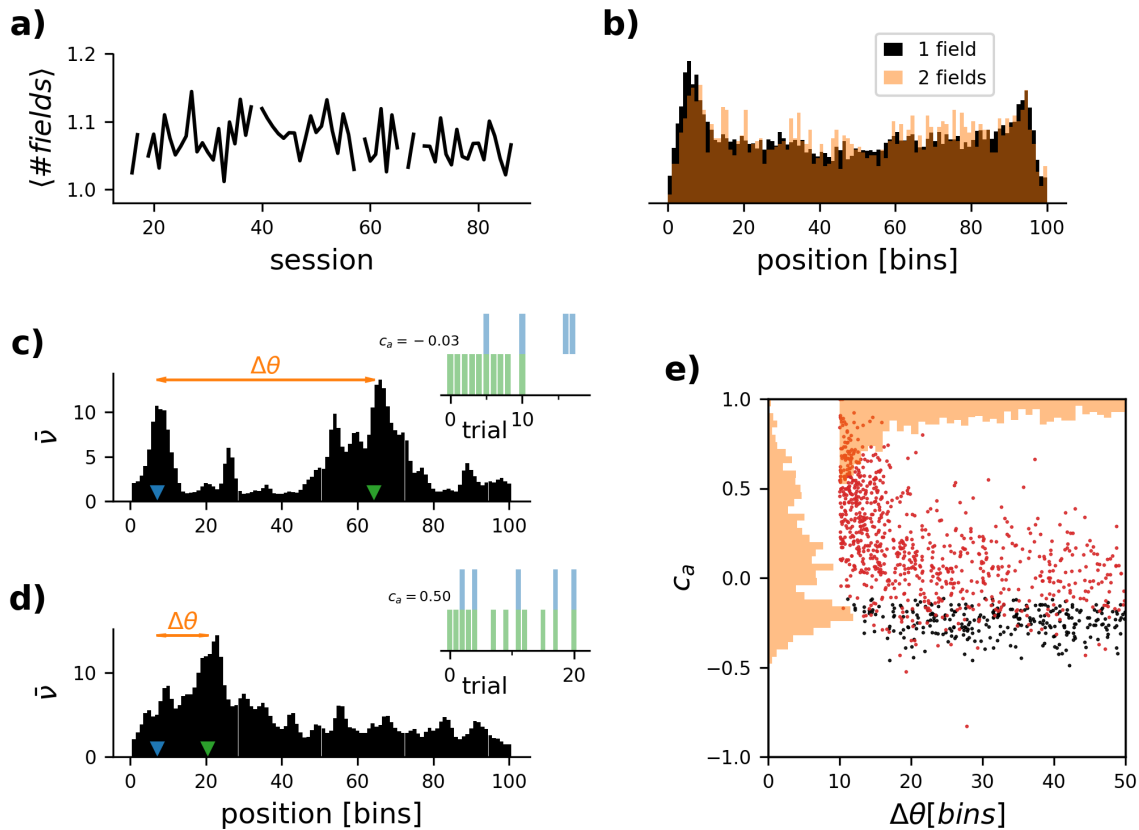


Figure 3.5: Some place cells display multiple fields; **(a)** session-wise average place field number per place cell; **(b)** distribution of place field locations of neurons with 1 (black) and 2 (orange) place fields; **(c)** firing map of a neuron with highly separated place fields (centers marked by green and blue triangle, place field distance in orange); inset: identified active trials of left (green) and right (blue) place field with Pearson correlation coefficient of trial activation c_a marked in plot; **(d)** same as (c) for a place cell with closely placed fields; **(e)** scatter plot of trial activation correlation (ordinate) and field distance (abscissa) for all neurons with two place fields; distinction between cases with no coactivated fields within trials (black) and cases with at least one coactivated field (red); histograms (orange) show distribution of according parameters

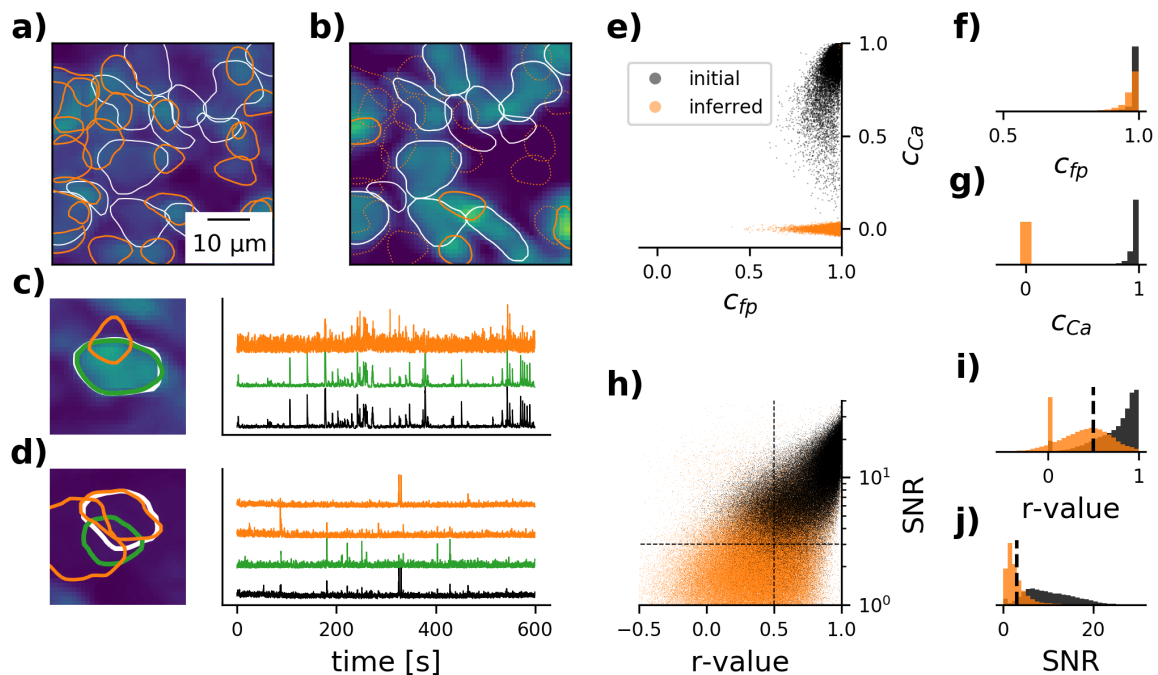


Figure 3.6: Initially undetected neurons are identified from other sessions’ footprints, but score bad on evaluation criteria, confirming inactivity in some cases, but provide activity in others; **(a)** initially detected set of footprints (white) and further footprints with shapes inferred from adjacent sessions (orange); **(b)** footprints from **(a)** after redetection by CNMF; dashed footprints do not pass thresholds of $\Theta_{SNR} = 3$, $\Theta_r = 0.5$; **(c)** left: initially detected footprint before (white) and after (green) CNMF, which maintains shape and temporal trace throughout the rerun of the detection. Footprint of overlapping neuron after CNMF (orange); right: Calcium traces of footprints displayed left; **(d)** similar to **(c)**, but for a neuron with low footprint- and calcium trace-correlation of components before and after redetection. Another component (orange) now represents the same neuron; **(e)** correlation of calcium trace c_{Ca} and footprint shape c_{fp} before and after redetection. Each dot represents one neuron, with initially detected in black and inferred ones in orange; **(f,g)** normalized distributions of c_{Ca} and c_{fp} **(h)** signal-to-noise-ratio and r-value for initially detected (black) and undetected (orange) neurons; **(i,j)** normalized distributions of r-value and SNR

a large number of customizable parameters to adapt to requirements of local machines, while the neuron detection performance remains largely independent of user-input, obtaining most parameters selfconsistently from the underlying data. The performance of neuron detection has been shown to be on par with human labeler performance in the paper presenting the toolbox¹³⁵, but an additional test of how exhaustively the algorithm detects the population of neurons will be subject of this section.

I decided to use the results from a single animal, *m762*, only, for displaying purposes, as this part deals with understanding the detection algorithm itself. All of the presented analyses, however, have been performed on all the animals analyzed in this work.

Redetection procedure preserves, refines and completes initial neuron detection After an initial run, yielding a first estimate of neuron components, footprints are matched to one another across different sessions, Sec. 2.5.2, allowing to assign neuron identities to footprints in each session and thus find, whether a neuron has been successfully identified in a given session, or not. Footprint shape and location of neurons in sessions where they were not detected are inferred from according footprints of the temporally most adjacent sessions, in which the given neuron appears, see Fig. 3.6(a-b). An additional run of the CNMF-method is initialized, using the initially detected neuron footprints, as well as the inferred locations and shapes of undetected neurons as input values.

Two examples in Fig. 3.6(c,d) display what kind of impact the redetection procedure can have on individual components. In most cases, the footprint and calcium trace of an initially detected neuron remains highly similar (white to green in Fig. 3.6(c)), expressed by high correlation values of footprints c_{fp} and calcium activity c_{Ca} (correlation between the in- and output arrays to the redetection), Fig. 3.6(e-g). c_{Ca} shows a wider distribution than c_{fp} , originating from stronger demixing of signals, as neuron density and overlap increases after adding the inferred neuron footprints. Calcium events can now be assigned to previously undetected neurons, see Fig. 3.6(d), right, causing the trend towards lower activity correlation values.

In some cases, however, the additional set of inferred neuron footprints is competing to represent initially detected neurons, in some cases resulting in a crippled footprint with noisy activity for an initially detected neuron (green footprint and trace in Fig. 3.6(d)). While in most cases the demixing successfully separates activity traces from one another, in such cases of two footprints representing a single neuron, the demixing tends to fail. I therefore added another health-check after data processing, removing footprints with lower SNR value in cases where footprint and activity correlation with a neighboring neuron surpass threshold values of $\tilde{c}_{fp} = 0.2$ and $\tilde{c}_{Ca} = 0.5$.

Footprint correlations of in- and output of inferred neurons show slightly lower correlation values c_{fp} than initially detected ones, Fig. 3.6(e,f), orange, as their shapes - inferred from other sessions - have to be fine-tuned to the exact location and form of the neuron at the current session, causing minor alterations. The correlation of in- and output calcium-trace c_{Ca} vanishes (Fig. 3.6(g)), since it - in lack of any prior knowledge on inferred neurons' activity - is initialized by gaussian noise.

Evaluation of detected neuron footprints An analysis of the evaluation parameters of signal-to-noise-ratio and r-value for all resulting neurons (Fig. 3.6(h-j)) reveals, why some neurons were not detected initially, but could only be found by guiding the algorithm to the position of neurons: Both, SNR- and r-value-distributions of inferred neurons center at low values, Fig. 3.6(h-j). Many of their temporal traces are noisy (low SNR) and do not add much improvement to approximating the raw data (low r-value). The overlap of distributions, however, suggests a




Figure 3.7: Evaluation thresholds impact neuron numbers, but not place field ratio; **(a,b)** number of detected (active) neurons (black) and place cells (blue, only for neurons considered in **(c)**). Dots represent values from single sessions, triangles the average over the animal; x-axis varies the threshold used for neuron detection, while other threshold is fixed at value marked by dashed black lines; **(c)** neuron numbers for different thresholds of matching probability with somata near FoV border removed (see text); blue line is the fraction of place cells within the population of active neurons (scale on the right)

non-negligible set of neurons which show significant activity and therefore should be added to the set of analyzed neurons.

Applying a threshold of SNR and r-value to the data allows to filter out noisy, which most likely represent noise, only and thus, a silent neuron in the inferred position. Characterizing components into active and silent this way, the number of active neurons per session decreases with increasing threshold, Fig. 3.7(a,b). I require neurons to have a signal-to-noise ratio allowing to properly identify significant spikes and to not contradict the actual raw data and thus choose thresholds of $\Theta_{SNR} = 3$ and $\Theta_r = 0.5$ for further analyses, Fig. 3.7(a,b), in line with the work introducing CaImAn¹³⁵. As a last step, I remove neurons which were detected too close to the border of the FoV, such that observed FoV shifts across different sessions would move parts of somata or complete ones out of the imaging section.

A complete set of detected neuron footprints is displayed in Fig. 3.8(a). Neuron sizes (diameter mean $19.1 \pm 1.8 \mu\text{m}$, Fig. 3.8(c)) are within an anatomically reasonable range ($\sim 20 \mu\text{m}$) and neurons are tightly packed, as expected from their arrangement within the stratum pyramidale³¹. Calcium activity correlations are very low on average, Fig. 3.8(d), lower center, but show an increase for low centroid distances, indicating pollution through overlapping or neighbouring pixels. However, the footprint correlation c_{fp} does not significantly impact activity correlation, Fig. 3.8(d), lower right, suggesting that the observed pollution rather originates in the dispersion of intensity across neighbouring pixels during recording, than from insufficient demixing.

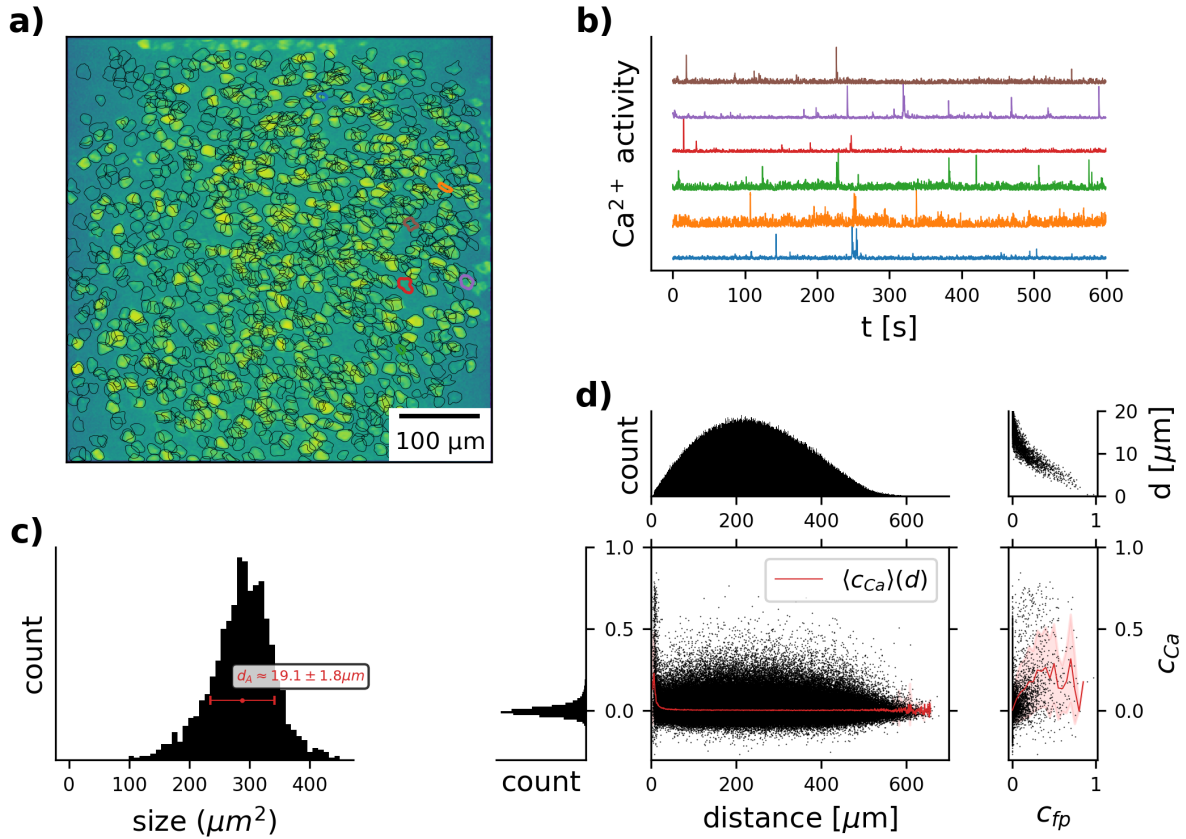


Figure 3.8: Neuron components as well as extracted calcium traces identified by CNMF uniformly match properties of hippocampal CA1 pyramidal neurons; **(a)** contours of identified neuron footprints (black) on top of a neighbourhood-correlation image; dotted contours are neuron components rejected due to low SNR or r-value; **(b)** extracted calcium activity of 6 neuron components highlighted in according colors in (a); **(c)** distribution of neuron size in pixels, counting all all entries in footprint matrix a_j above 10% maximum intensity; **(d)** distribution of Pearson correlation coefficient of calcium activity (lower left) and centroid distances (upper left) over all possible pairs; lower center: activity correlation plotted against centroid distance, distance-dependent average highlighted in red. Each dot is one pair; upper right: footprint correlation and centroid distance for closely neurons; lower right: activity correlation against footprint correlation with footprint-correlation dependent average (red, 95% confidence interval as shaded region)

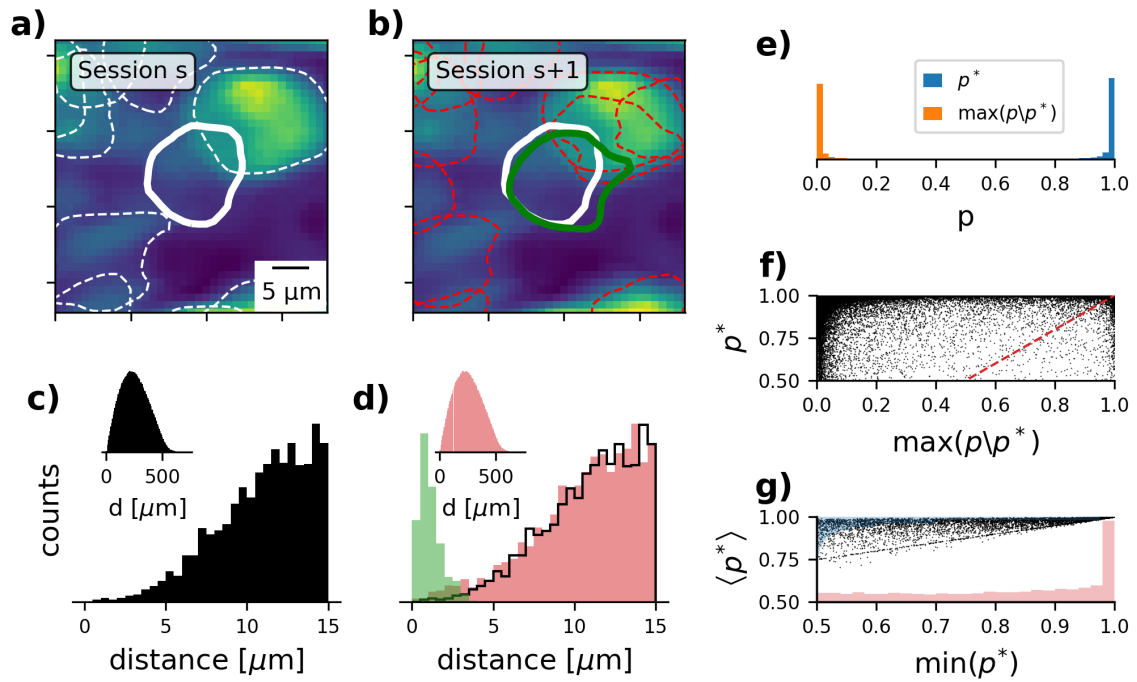


Figure 3.9: Registration procedure results in a large fraction of reliable matches; **(a)** contours (white) of several footprints in a section of the FoV, on top of a correlation image, center footprint is highlighted by solid line; **(b)** same section of the FoV with contours extracted from the next session: red - different footprints, green - matching footprint to the one from previous session (white); **(c)** histogram of pair-wise distance of footprint-centroids within one session (inset: overall distribution, large window: zoom in to low distances); **(d)** same as in (c), but distances between footprints of different sessions. Red - footprints belonging to different neurons, green - footprints belonging to the same neuron, black line - histogram from (c); **(e)** histogram of matching probabilities p^* (blue) of assigned matches and matching probabilities $p \setminus p^*$ of best non-matched candidate; data pooled over all sessions of displayed mouse; **(f)** p^* against $p \setminus p^*$ for each assigned pair. Case of equal matched and maximum non-matched probability highlighted by red dashed line; threshold used for further analysis ($p^* - \max(p \setminus p^*) > 0.5$, $p^* > 0.95$ highlighted in blue); **(g)** neuron registration scores: averaged score $\langle p^* \rangle$ against minimum matching score in neuron over all sessions

3.3.1 Matching footprints across sessions

The neuron registration procedure quantifies the probability of two neurons belonging together based on their footprint distances and shapes, Sec. 2.5.2. I can confirm that the distribution of footprint distances between consecutive sessions is composed of a set of footprint pairs which belong together (small distance, green in Fig. 2.5.2(d)) and a set of distinct neuron pairs, composing a distribution similar to the within-session footprint distribution (larger distance, red in Fig. 2.5.2(d)).

The distribution of the matching probability p^* of pairs assigned by the matching algorithm is sharply peaked at 1, indicating that neurons maintain similar, non-shifted footprints over consecutive sessions, leading to high certainty in matching for the vast majority of assigned pairs of footprints (Fig. 3.9(e), blue). For many neurons further candidates for matching with non-zero probability exist, which mostly remain well below the probability for the matched neuron, Fig. 3.9(e,f), orange. A few matches, however only have a small margin towards the second best matching candidate, or are even suboptimal, by not assigning the highest probability candidate. On closer inspection, however, I find that in those cases the better match usually has another neuron candidate that fits even better (tested on a few randomly picked cases).

Most neurons have a high average matching probability over the whole dataset (sharp peak at 1, Fig. 3.9(g), blue). However, a large number of neurons have at least one match with low probability in some session, Fig. 3.9(g), red. The matching algorithm allows for such outliers (matches with low probabilities), as they do not change the reference footprint, which matching candidates of following sessions will be compared against, Sec. 2.5.2. Single outliers can therefore be individually classified as proper matches or non-matches,




Figure 3.10: Total neuron numbers remain stable over all sessions with changing active population; top left: neuron count per session: initially detected (green), significant components detected after inference from other sessions (orange), thresholded with $\Theta_{SNR} = 3$ (black); bottom left: detected footprints at session given by x-axis for neuron ID given by y-axis, colors as before. vertical blank lines correspond to sessions removed from the dataset due to a large x-y-shift of the FoV; bottom right: number of sessions corresponding neuron on the left occurs in; top right: normalized, cumulative distribution of session-occurrence over all neurons displayed

based on their matching probability, and will only slightly change the matches of footprints from successive sessions.

To avoid obvious mismatches, I chose the minimum probability for accepting matches to be $\Theta_p^* = 0.5$. Even higher thresholds have almost no impact on the data, Fig. 3.7(c). A neuron is considered to be silent in a session, when a matching footprint could not be detected.

Tracking of individual neurons The set of matched neuron footprints not only allows to track the neurons functionality, but also to track movements of cell bodies over all sessions, examples displayed in Fig. 3.11(a). Detecting reasonable movement over the timescale of several sessions can serve as yet another way to provide certainty about the matching procedure. However, keep in mind that one of the two parameters used for matching - namely centroid distance - actually imposes a strong constraint, enforcing reasonable movement. Accordingly, I find a distribution of centroid movements between adjacent sessions strongly peaked around 0 in both, x- and y-direction, Fig. 3.11(b). I obtain overall movements of few μm over the period of over 100 sessions, Fig. 3.11(c), with maximum migration distances from the initial position of up to $10\mu\text{m}$, Fig. 3.11(d), indicating a very stable anatomical layout. Matching probabilities smaller than 0.9 within a neuron dataset result in wider shifts between sessions, Fig. 3.11(b), but not in an overall larger migration of the neuron, Fig. 3.11(d), red, indicating a very low probability of mismatches.

3.3.2 Cell types by functionality

Moving forward, for the rest of this work I will assume that the data is accurately described by the results of above methods and thresholding procedures.

A complete set of active and silent neurons Over all recording sessions of the given animal, the total number of active neurons remains stable around $\sim 900-1100$. The inference of footprints from other sessions brings an increase of $\sim 20-35\%$ in identified active footprint numbers with respect to the initial detection, Fig. 3.10, upper left. The total number of neurons identified within this animal amounts to over 3500, with most neurons active in few sessions, only.

Place coding cells In the displayed animal, I find an average fraction of $21.6 \pm 6.7\%$ of active neurons to be place coding per session, Fig. 3.12(c), in line with earlier reported measurements from CA1^{15,18}. There is no apparent anatomical correlate to the place field location, such that place cells occur in a salt-and-pepper manner, Fig. 3.12(a). Average signal-to-noise ratio of place cells and the ratio of in-field firing rate amplitude A to baseline rate A_0 remain constant across all sessions, Fig. 3.12(c).

Cells in close proximity ($< 10-15\mu\text{m}$) tend to have an increased firing map correlation, Fig. 3.12(c), with a stronger effect in place cells. This could be either due to dispersion or

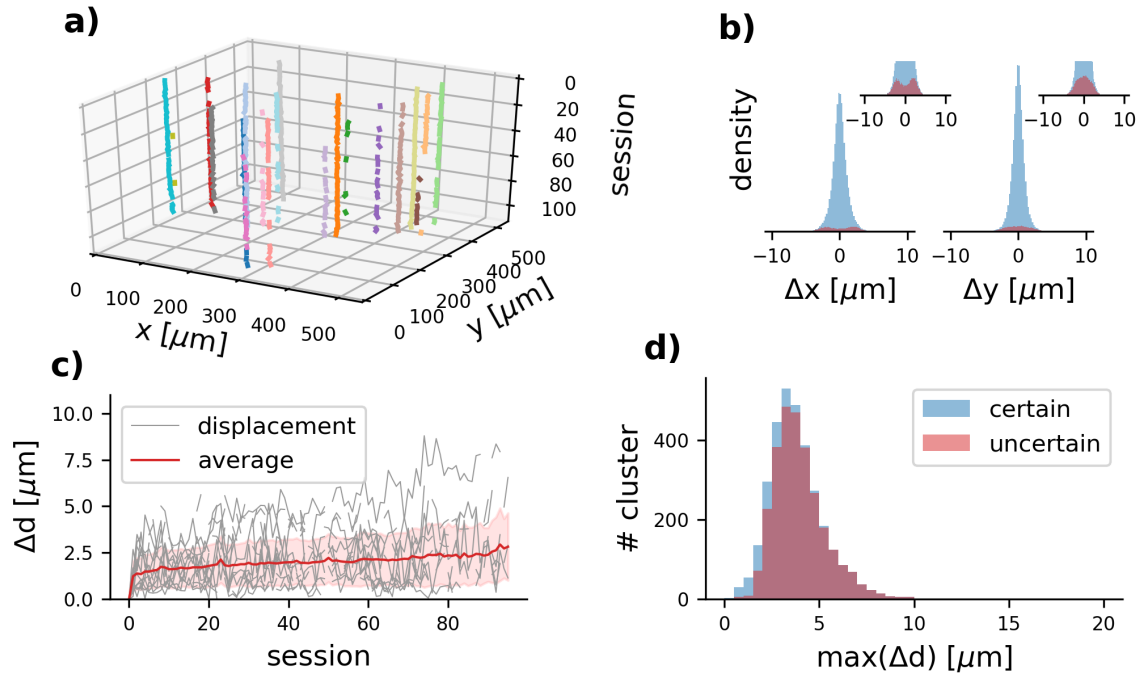


Figure 3.11: Neuron bodies remain stable but slowly migrate between sessions; **(a)** traces of a few randomly picked neuron centroids over all sessions. Each color corresponds to one neuron, displayed in sessions, where it occurs, only; **(b)** distribution of centroid shifts between adjacent sessions in x- and y-direction, certain matches (blue), uncertain matches ($p^* < 0.9$, red); inset: zoom in to highlight distribution of shifts of uncertain matches; **(c)** average migration of neuron centroids (red, SD displayed by shaded region), with single trajectories of neurons displayed in (a) as a function of the number of sessions passed since their first detection; **(d)** distribution of maximum shifts from original centroid position for all neurons, colors as in (b), with a neuron counting as uncertain, as soon as there is a single uncertain match in any of the sessions (see (b))

insufficient demixing of activity traces of overlapping regions, which has already been shown in Fig. 3.8(d), lower center, or could be the cause of correlated input to narrow regions of the CA1.

Finally, place fields from the population of place cells cover the whole virtual environment experienced by the animal, with changes to the distribution over the course of several sessions, which is further discussed in Sec. 5.4.

3.4 A vast data set

The data presented in Ch. 5 & 6 originates from 6 mice recorded in two batches, one from December 2013 to April 2014 (mice 762,231, 232, handled by Kotaro Mizuta) and one from December 2014 to May 2015 (mice 243, 245, 246, handled by Yukiko Sekine). After the data is processed by previously described steps, I obtain neuron and place field numbers widely varying across animals, Fig. 3.13(a): an average number of 700 to 1700 neurons are active per session,

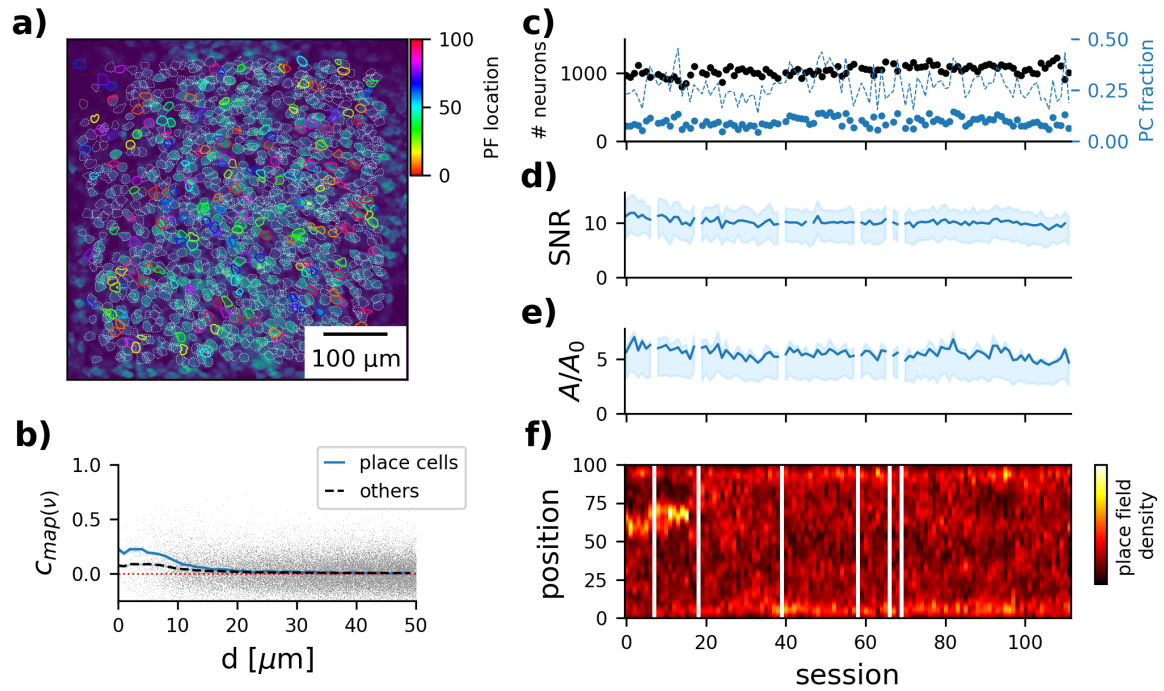


Figure 3.12: Place cell arrangement is not explained by a homotopic layout, but place fields reliably and completely cover the whole environment with only minor fluctuations across sessions; **(a)** identified footprints of neurons within one session with unspecific (dotted) and location specific (solid) activity; location of place fields color coded; **(b)** centroid distance-dependent pearson-correlation of firing maps; average of place cells (blue) and others (black) with confidence intervals pooled over all sessions, and single data points (grey dots) of place cells in a randomly picked session; vanishing correlation ($= 0$) marked by red dotted line; **(c)** number of active neurons (black) and PC (blue), with PC fraction (blue dashed line) on y-axis; **(d)** signal-to-noise ratio of place cells across all sessions with 2.5 to 97.5 percentile range (shaded areas); **(e)** ratio of in-field firing rate A to baseline firingrate A_0 of place fields with 2.5 to 97.5 percentile as in (d) **(e)** smoothed density of place fields (gaussian filter, $\sigma = 2$ bins) over the whole track with colorcode session-wise normalized to cover the whole color range; white data are sessions removed due to large FoV shifts or bad imaging conditions

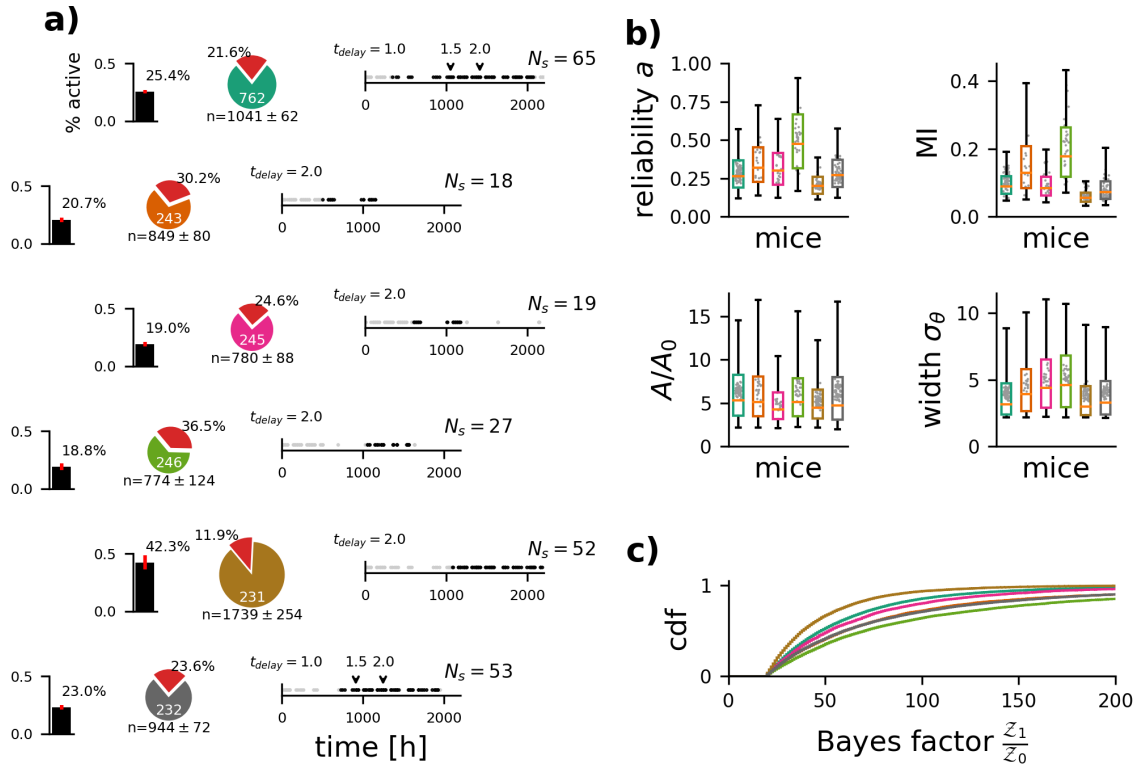


Figure 3.13: Neuron and place field parameters vary across different animals, **(a)** left: number of active neurons (mean + SD) beneath pie-chart and average fraction of place cells per session (red part with percentage); mice ID marked within pie-chart; right: time points of measurements after first exposure to environment; only sessions without major changes to the environment (black) are included in later analyses, Ch. 5, 6; waiting time for reward delivery after reward-zone entry t_{delay} indicated, with changes (in sec) indicated at first session with applied change; **(b)** neuron and place field parameters of place field trial-to-trial reliability a , mutual information MI , rate of within-field to out-of-field firing rate A/A_0 , width of place field (SD of gauss) σ_θ for different mice (colors as in pie-charts of a); box plots indicate median (line), 25-75 percentile (box) and 5-95 percentile (whiskers); grey dots show session average values; **(c)** cumulative distribution of Bayes factor values of detected place fields for different mice (colors as before)

corresponding to fractions of $\sim 20 - 40\%$ of the total population of detected neurons. This leaves a large fraction of $\sim 60 - 80\%$ silent neurons in each session, in line with experimental data^{52,164}. Averages of place cell fractions within an animal vary between 11.9% up to 36.5%, indicating individual differences in place coding across different mice.

Sessions for further analysis were taken from a larger set of recorded data, throughout which some animals experienced changes to the environment, such as shifts in the reward position. Only sessions after extensive familiarization with the environment (at least 15 experienced sessions), at least 3 sessions after the last change to the environment and without a change other than a slight increase in the delay period to receive a reward were included in the analysis, resulting in a number of $N_s = 18$ to $N_s = 65$ analyzed sessions, see Fig. 3.13(a).

Neuron and place field parameters vary widely between different mice: the mutual information ranges from 0.06 ± 0.02 bits to 0.21 ± 0.07 bits (mean \pm SD, Fig. 3.13(b), top right) - overall low values, reflecting the relatively poor environment of a linear, unidirectional track with frontal stimulation, only. Place field reliability takes on values between 0.21 ± 0.03 and 0.51 ± 0.10 (Fig. 3.13(b), top left). The ratio of field-to-baseline firing rates is relatively constant at a factor of 6.08 ± 0.61 across all mice and the average place field width takes on values of 4.57 ± 0.56 bins (mean \pm SD across animals, equal to 21.51 ± 2.68 cm when including ± 1.96 SD of the gaussian shaped field), Fig. 3.13(b), bottom left and right.

Bayes factor values used for the decision of a neuron being a place cell, or not, range from the very minimum value allowed to high values, Fig. 3.13(c).

3.5 Discussion

The CNMF algorithm¹³⁵ applied to detect neurons displays a natural bias towards more active neurons. A large population of silent neurons as has been reported before^{52,164} therefore remains undetected, but also active neurons with less pronounced activity appear to be missed out on. In this chapter I show how the results from the neuron registration procedure allows to infer the position of neurons in other sessions, to provide a better estimate to the initialization of the CNMF method, Sec. 3.3. I find an increase in detected active neurons of about 20 – 35% (depending on session and animal), allowing me to provide a more complete picture of neuron activity within sessions and more importantly, across different sessions.

Neuron registration The registration procedure presented in Sec. 2.5.2 results in a dataset, where several neurons can be tracked over a large number of, or even all recording sessions. The matching procedure is guided by a probabilistic model, allowing to specify a desired certainty of matching. A previous work, introducing the idea to the model has shown how the computed place field stability scales with the threshold for matching¹⁷, suggesting that low matching thresholds result in false matches.

Three factors allow to further control the matching procedure: (1) an adjustment to the model contributed within this work pronounces the shape of neurons as an important factor for registration in contrast to the original work. I suggest a computation of footprint correlations independent of neuron movement to obtain a more accurate model of matching probabilities. (2) a threshold of $\Theta_p^* = 0.5$ allows for a conservative estimate of matches. (3) the registration procedure with a variable update of the reference dataset allows to minimize the impact of false matches for further sessions, especially important in long datasets.

In lack of a ground truth dataset of neuron identities, I assess the matching performance from an indirect measure: the overall migration of neuron bodies. In adult brains, it is expected to remain low, which I can confirm from low average and maximum distance from the original neuron position over up to 100 sessions (Fig. 3.11(c,d)), even though the matching algorithm allows for deviations of up to a few μm per session without major penalty, Fig. 2.5(a).

A study providing a groundtruth dataset of either manually labeled neuron footprints, or using genetic markers to uniquely identify neuron bodies, or an artificially generated set of footprint data could provide a groundtruth for assessment of the matching procedure and support the field of Ca^{2+} -imaging to obtain certainty in the tracking of neurons over days, weeks and months.

Place cell detection The detection of place fields is at the heart of characterizing place field dynamics. Even though, earlier works do not show any assessment of place field detection quality, when analyzing the spatial representation within the hippocampus. I dedicated the first part of this chapter to assess the introduced approach of place field detection, using a surrogate dataset of artificially generated neuron activity.

I find that place fields are generally well characterized for a wide range of parameters, showing high recall and precision values, as well as a high accuracy in retaining the place field position and some accuracy in retaining other parameters such as baseline activity, peak amplitude and place field width. As could be expected, place fields with low amplitudes are more difficult to detect, while the place field position, width or baseline activity does not show an influence on the detection quality. I am further able to detect multiple place fields per neuron, even though a second, less pronounced field often remains undetected.

The introduced place field detection procedure uses a trial based approach, similar to what has been used in a recent study¹²¹, which makes the detection quality largely independent of the place field reliability, except for a sharp drop at very low activation probabilities. Active trials of correctly identified place cells are retrieved with high precision and recall, with similar parameter dependencies as the place fields themselves.

The assessment provides important insights into the quality of place field detection and suggests a very reliable process of identifying significant spatial modulations in the firing map of neurons. However, it has to be taken into account that the artificially generated data is generated from the same model that is used for place field detection. It therefore shares the underlying

assumptions that place fields originate from some underlying firing rate model with a fixed spatial modulation, and that fields are approximately gaussian. This can introduce a strong bias, causing favorable conditions for the detection algorithm.

Application of the place field detection to real data shows, how a fraction of $\sim 10\%$ of neurons displays multiple place fields, in line with studies reporting average place field number in CA1 neurons to increase along the proximo-distal axis³⁰. In many neurons displaying multiple place fields the fields are found in proximity with one another and often coactivate, suggesting that it is indeed a single place field. Causes for such a false classification could be the experience-dependent reshaping of place fields as reported before^{61,63,64}, but also a strong inter-trial variability of place field position¹²¹.

The exact size of place fields has been shown to depend on the size of the environment^{51,67}. Place field width are therefore more commonly reported as the covered fraction of the environment, found here to be $\sim 20 - 50\%$ (considering 95% mass of the gaussian field), in line with previous studies¹⁵.

In summary, I detect a large number of neurons for all animals presented and reliably track them over several dozens of sessions. The place field detection method provides a way to characterize spatial modulations, which has been tested successfully against surrogate data and shown to reliably retrieve underlying place field parameters.

Chapter Four

Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus

4.1 Content

A long list of studies has consistently found influences other than animal location to contribute to the shape of the hippocampal place map^{112,114–121}. Various studies have found visual cues (e.g. landmarks) or functional (e.g. rewards) saliency to result in an increased density of place cells, while the underlying mechanism of the formation of such a place map, as well as the role of visual or functional salience in the map dynamics so far has not been analyzed.

The following study has the goals of (1) analyzing the mechanisms underlying the formation of overrepresentation in the hippocampal place map and (2) to investigate the different role of landmarks and rewards in this. We investigate for the first time the conjoint effect of landmarks and rewards in place map formation and find that map formation is driven by local stabilization of salient regions, leading to the observed overrepresentation. Landmarks and rewards are found to have distinct contributions, identified by differing time scales at which respective overrepresentations appear. The ability to selectively remove landmark overrepresentation in a Shank2 knock-out mouse model suggests different underlying molecular mechanisms for reward and landmark coding.

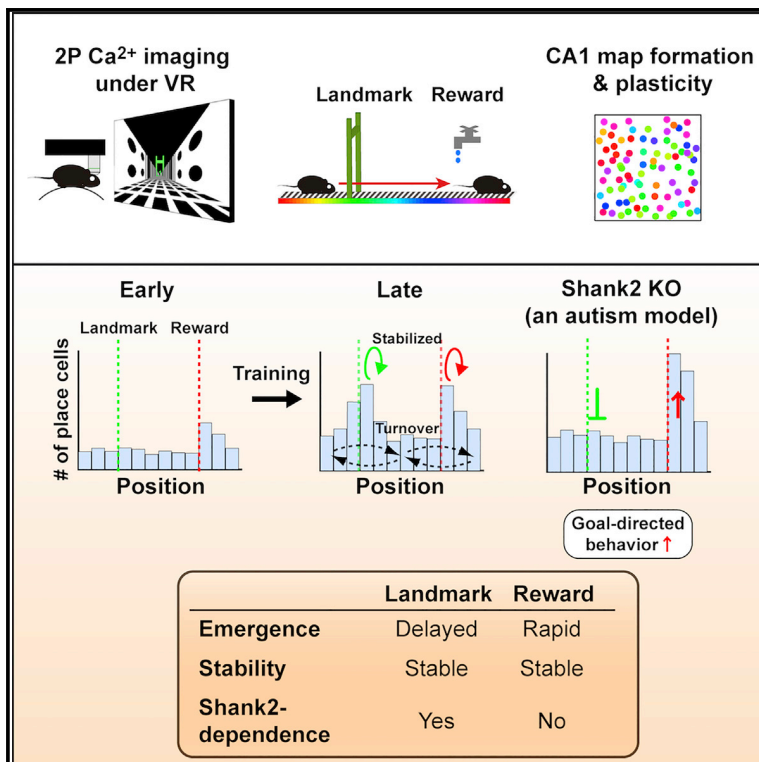
4.2 Citation and original contribution

Masaaki Sato, Kotaro Mizuta, Tanvir Islam, Masako Kawano, Yukiko Sekine, Takashi Takekawa, Daniel Gomez-Dominguez, **Alexander Schmidt**, Fred Wolf, Karam Kim, Hiroshi Yamakawa, Masamichi Ohkura, Min Goo Lee, Tomoki Fukai, Junichi Nakai and Yasunori Hayashi: *Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus*, Cell Reports 32.1 (2020), p. 107864.

I contributed to conceiving analysis strategies, contributed to the analysis of the data, performed independent additional checks of key conclusions and participated in preparing the manuscript. Experiments and data collection were performed by co-authors as indicated in the paper. This paper is a result of a long term collaboration between the Kyoto University and the Max Planck Society.

Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus

Graphical Abstract



Authors

Masaaki Sato, Kotaro Mizuta, Tanvir Islam, ..., Tomoki Fukai, Junichi Nakai, Yasunori Hayashi

Correspondence

masaaki.sato@riken.jp (M.S.),
yhayashi-ky@umin.ac.jp (Y.H.)

In Brief

Using longitudinal two-photon calcium imaging in mice during virtual navigation, Sato et al. demonstrate that persistent and separable neuronal subsets mediate the hippocampal over-representation of reward and landmark locations. Learning-induced over-representation of landmarks is absent while rapid over-representation of rewards is enhanced, in a mouse model of autism lacking *Shank2*.

Highlights

- CA1 over-representation of reward and landmark emerge with distinct time courses
- These cells form stable singularities during experience-dependent map consolidation
- The over-representation of landmark but not reward is dependent on *Shank2*



Article

Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus

Masaaki Sato,^{1,2,3,4,5,18,19,22,*} Kotaro Mizuta,^{5,6,18} Tanvir Islam,^{1,5} Masako Kawano,⁵ Yukiko Sekine,^{1,5} Takashi Takekawa,^{5,7} Daniel Gomez-Dominguez,^{5,8} Alexander Schmidt,^{6,9,10,11,12} Fred Wolf,^{9,10,11,12} Karam Kim,⁵ Hiroshi Yamakawa,^{5,13,14,15} Masamichi Ohkura,^{3,4,20} Min Goo Lee,¹⁶ Tomoki Fukai,¹⁷ Junichi Nakai,^{3,4,21} and Yasunori Hayashi^{4,5,6,*}

¹RIKEN Center for Brain Science, Wako, Saitama 351-0198, Japan

²PRESTO, Japan Science and Technology Agency, Kawaguchi, Saitama 332-0012, Japan

³Graduate School of Science and Engineering, Saitama University, Saitama 338-8570, Japan

⁴Brain and Body System Science Institute, Saitama University, Saitama 338-8570, Japan

⁵RIKEN Brain Science Institute, Wako, Saitama 351-0198, Japan

⁶Department of Pharmacology, Kyoto University Graduate School of Medicine, Kyoto 606-8501, Japan

⁷Faculty of Informatics, Kogakuin University, Tokyo 163-8677, Japan

⁸Instituto Cajal, CSIC, Madrid 28002, Spain

⁹Max Planck Institute for Dynamics and Self-Organization, Göttingen 37077, Germany

¹⁰Max Planck Institute for Experimental Medicine, Göttingen 37075, Germany

¹¹Campus Institute for Dynamics of Biological Networks, Göttingen 37075, Germany

¹²Center for Biostructural Imaging of Neurodegeneration, Göttingen 37075, Germany

¹³The University of Tokyo, Tokyo 113-8654, Japan

¹⁴Whole Brain Architecture Initiative, Tokyo 133-0057, Japan

¹⁵RIKEN Center for Biosystems Dynamics Research, Osaka 565-0874, Japan

¹⁶National Creative Research Initiative Center for Cell Membrane Transport, Yonsei University College of Medicine, Seoul 120-752, Republic of Korea

¹⁷Okinawa Institute of Science and Technology, Onna, Okinawa, 904-0495, Japan

¹⁸These authors contributed equally

¹⁹Present address: Department of Neuropharmacology, Hokkaido University Graduate School of Medicine, Sapporo, Hokkaido 060-8638, Japan

²⁰Present address: School of Pharmaceutical Sciences, Kyushu University of Health and Welfare, Nobeoka, Miyazaki 882-8508, Japan

²¹Present address: Division of Oral Physiology, Tohoku University Graduate School of Dentistry, Sendai, Miyagi 980-8575, Japan

²²Lead Contact

*Correspondence: masaaki.sato@riken.jp (M.S.), yhayashi-ky@umin.ac.jp (Y.H.)

<https://doi.org/10.1016/j.celrep.2020.107864>

SUMMARY

In the hippocampus, locations associated with salient features are represented by a disproportionately large number of neurons, but the cellular and molecular mechanisms underlying this over-representation remain elusive. Using longitudinal calcium imaging in mice learning to navigate in virtual reality, we find that the over-representation of reward and landmark locations are mediated by persistent and separable subsets of neurons, with distinct time courses of emergence and differing underlying molecular mechanisms. Strikingly, we find that in mice lacking *Shank2*, an autism spectrum disorder (ASD)-linked gene encoding an excitatory postsynaptic scaffold protein, the learning-induced over-representation of landmarks was absent whereas the over-representation of rewards was substantially increased, as was goal-directed behavior. These findings demonstrate that multiple hippocampal coding processes for unique types of salient features are distinguished by a *Shank2*-dependent mechanism and suggest that abnormally distorted hippocampal salience mapping may underlie cognitive and behavioral abnormalities in a subset of ASDs.

INTRODUCTION

Navigation and spatial memory are essential elements of animal behavior that allow animals to forage, return home, and avoid danger. The hippocampus plays a crucial role in these cognitive processes, as hippocampal neurons fire when an animal is located in a particular part of an environment, providing an allo-

centric cognitive map of space (O'Keefe and Nadel, 1978). Although whether these “place cells” (PCs) are indeed memory cells has been long debated, one line of evidence that favors this notion indicates that hippocampal place-specific firing exhibits dynamic changes according to context and experience on multiple timescales, ranging from a few minutes to days or weeks (Muller and Kubie, 1987; Bostock et al., 1991; Mehta



et al., 1997; Lever et al., 2002; Leutgeb et al., 2005). Furthermore, studies have reported that disproportionately large numbers of PCs are recorded in locations that are associated with reward, safety, or local cues (O'Keefe and Conway, 1978; Wiener et al., 1989; Hetherington and Shapiro, 1997; Hollup et al., 2001; Dupret et al., 2010; Danielson et al., 2016; Zaremba et al., 2017; Gauthier and Tank, 2018; Bourboulou et al., 2019), indicating that the environment surrounding an animal is not represented uniformly in the hippocampal cognitive map; representations are strongly influenced by the motivational and environmental salience of the locations.

These findings imply that the activity of an increased number of neurons encodes the presence of salience in the hippocampal map. This idea further proposes potential roles of such salience maps not only in spatial (Hollup et al., 2001; Dupret et al., 2010) or episodic-like memories (Komorowski et al., 2009; Eichenbaum and Cohen, 2014) but also in goal-directed and landmark-based navigation (Burgess and O'Keefe, 1996; Gothard et al., 1996). PCs are formed rapidly within minutes after initial exposure to a new environment (Hill, 1978; Wilson and McNaughton, 1993; Frank et al., 2004). However, how the over-representation of salient features is established and modified by experience remains to be fully understood. Several important questions regarding these salience maps remain to be addressed: for example, whether the over-representation of reward and that of other types of salient features are mediated by a single shared or multiple distinct mechanisms and how altered over-representation relates to behavioral and cognitive abnormalities in various brain disorders, including autism spectrum disorders (ASDs).

To elucidate the cellular and molecular mechanisms that govern the dynamics of hippocampal salience representation, we longitudinally imaged functional cellular maps of the deep sublayer of CA1 during training on a virtual linear track, in which two distinct locations were associated with reward or a visual landmark. We show that over-representations of motivationally and environmentally salient features are mediated by persistent and separable subsets of neurons with distinct time courses of emergence and differing molecular mechanisms. Remarkably, mice lacking *Shank2* (Won et al., 2012), a mouse model of ASD that lacks a glutamatergic postsynaptic scaffold protein, exhibit selective loss of learning-induced over-representation of landmark locations, while their rapid over-representation of reward locations and goal-directed behavior is further enhanced.

RESULTS

Mice and Behavioral Task

To reliably perform longitudinal imaging of large-scale functional hippocampal cellular maps, we generated a transgenic mouse line, herein termed Thy1-G-CaMP7, that coexpresses the fluorescent calcium indicator protein G-CaMP7 and the calcium-insensitive red fluorescent marker protein DsRed2 via 2A peptide-mediated bicistronic expression under the neuron-specific Thy1 promoter (Figure 1A; Ohkura et al., 2012; Sato et al., 2015; see also STAR Methods and Figure S1). In the dorsal CA1 of the hippocampus, the population of calbindin D-28K-negative pyramidal cells in the deep pyramidal cell sublayer was preferentially labeled with G-CaMP7 (Mizuseki et al., 2011;

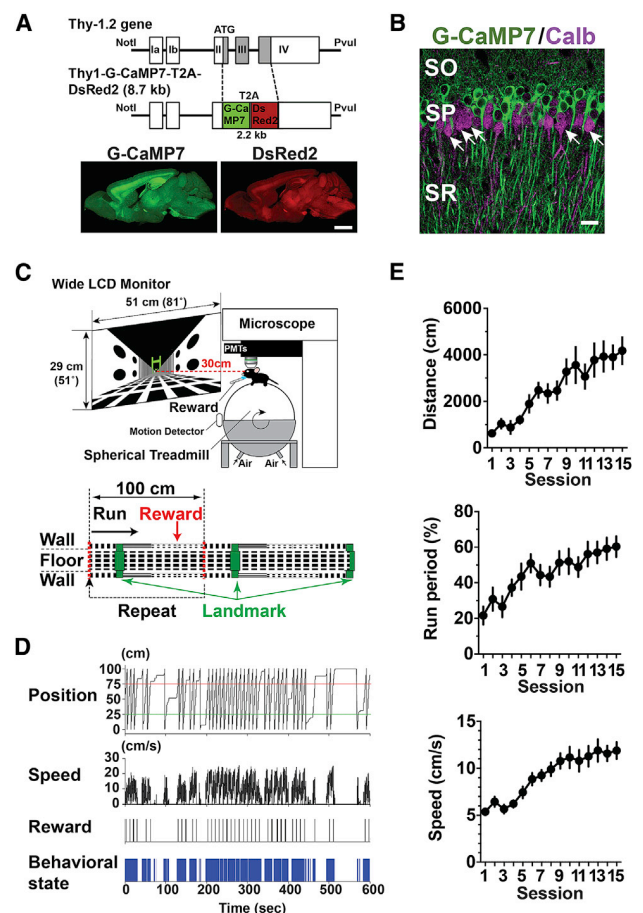


Figure 1. Transgenic Mice and Behavioral Task

(A) Transgene construct for Thy1-G-CaMP7 mice (top) and expression of G-CaMP7 (bottom left, green) and DsRed2 (bottom right, red) in a parasagittal section from a mouse at six months of age. Scale bar, 2 mm.

(B) G-CaMP7 expression (green) and calbindin immunofluorescence (calb, magenta) in the dorsal CA1 of the hippocampus of Thy1-G-CaMP7 transgenic mice. Arrows indicate examples of calbindin-positive G-CaMP7-negative cells. SO, stratum oriens; SP, stratum pyramidale; SR, stratum radiatum. Scale bar, 20 μ m.

(C) A schematic of the two-photon microscope and virtual reality setup (top) and virtual endless linear track task (bottom). The linear track segment contained a visual landmark (a green gate) and a reward delivery point at two distinct locations. When the mouse's virtual position reached the point indicated by the red dotted line in the middle, it returned to the origin, such that the same track segment was presented repeatedly.

(D) Example behavioral data from a single 10-min session. From top to bottom, the mouse's virtual position on the linear track, running speed, timing of reward delivery, and behavioral state are shown, where running is represented in blue.

(E) Behavioral changes induced by repeated training. Total distance traveled (distance, top), the fraction of time spent running (run period, middle), and running speed (speed, bottom) are shown. Data are expressed as mean \pm SEM.

Kohara et al., 2014; Lee et al., 2014; Valero et al., 2015; Danielson et al., 2016; Figures 1B and S1A). Immunofluorescence labeling of glutamic acid decarboxylase 65/67, parvalbumin, and somatostatin revealed that interneurons positive for these markers were devoid of G-CaMP7 expression (Figure S1B).

To allow imaging of hippocampal maps during repeated training of spatial behavior, we trained Thy1-G-CaMP7 mice in a virtual linear track task (Figures 1C and 1D; see STAR Methods for details). The mouse started running from the origin of the segment, passed under a green gate as a visual landmark, then received water at a reward point and returned to the origin instantaneously after reaching the other end. The visual landmark and reward delivery were associated with two distinct locations to examine the effects of two different types of salience separately. In training, behavioral performance as measured by time spent running, distance traveled, and running speed during 10-min sessions markedly increased as training proceeded (Figure 1E; distance, $p < 0.0001$, $F_{(14,154)} = 7.30$; run period, $p < 0.0001$, $F_{(14,154)} = 5.26$; speed, $p < 0.0001$, $F_{(14,154)} = 11.6$; $n = 12$ mice from three groups, one-way analysis of variance [ANOVA]). Furthermore, licking and slowing of running speed before the reward delivery point developed as training proceeded (Figure S2). The emergence of such anticipatory behavior indicates that the task involves goal-directed spatial learning.

Rapid and Delayed Emergence Distinguishes Over-Representations of the Reward and Landmark Locations

To examine whether and when the representations of the two salient locations become prominent in the map, we next visualized the emergence and establishment of hippocampal CA1 spatial maps during training on the virtual linear track task (Figure S3; see STAR Methods for details on imaging and analysis). The cells that exhibited virtual-location-specific activity (PCs) showed spatially more informative activity and higher event frequency than non-PCs (Figures S3H–S3K).

Consistent with previous studies in real and virtual environments, PCs were formed rapidly within the first session on the virtual linear track (Hill, 1978; Wilson and McNaughton, 1993; Frank et al., 2004; Chen et al., 2013) (Figures 2A and 2B). The fractions of PCs were initially low but then increased as the training proceeded. The fraction of PCs and that of time spent running showed a good overall correlation (Figure 2C). The slope of a regression line (termed the “PC formation factor”) significantly increased in the late phase of training compared with the early phase (Figures S4A and S4B). Furthermore, the sessions in the late phase contained larger fractions of PCs than the sessions in the early phase with comparable amounts of running time (Figures S4C and S4D). The average event frequency, mutual information content, and response amplitude of PCs did not notably change during training (Figures S4E–S4G).

We then examined whether the locations associated with salience were disproportionately represented in the hippocampal map. The histograms of PCs against positions typically exhibited two large peaks, which appeared more clearly in the late training phase; one peak corresponded to the location of the landmark, and the other corresponded to that of the reward (Figure 2D). While the first peak closely matched the landmark location, the second peak was slightly shifted in the direction the mouse was running, which likely reflects that the animals received the rewards in places that were slightly past the delivery point, as suggested by decreased running speed in this area (Figure S2A). Importantly, the over-representation of the reward

location was discernible even in the first session of training, whereas that of the landmark location gradually developed as the training proceeded (Figures 2D–2F). The fraction of PCs that encoded the location of the reward (herein termed “reward cells” [RW cells] for convenience) was not significantly different between the early and late phases of training, whereas that of PCs that encoded the location of the landmark (similarly termed “gate cells” [GT cells]) increased significantly, with a complementary decrease in the fraction of PCs that encoded other locations (termed “non-reward/gate cells” [non-RW/GT cells]; Figures 2E and 2F). The delayed emergence of increase in PCs that encode locations associated with salient visual cues is further supported by a more delayed and reduced increase in PCs that encode a location with less visual salience, such as a boundary of different wall patterns (termed “wall cells” [WL cells]; Figures S4H and S4I). Collectively, these results demonstrate that the over-representation of salient locations is formed and maintained at a population level, although the maps develop dynamically throughout the training period. The establishment and refinement of representations of salience depend on its nature; the representation of motivational salience is established rapidly, whereas that of environmental salience develops over the course of training.

RW Cells and GT Cells Form Stable Singularities during Experience-Dependent Map Consolidation

The PC maps imaged in each session appeared rather different from each other, even within the same animals, implying that hippocampal spatial representations are highly dynamic while being established (Figure 2A). To investigate whether representations of salient locations are more stable than those of nonsalient locations, we investigated training-induced changes in the maps at an individual cell level by comparing the place fields of the same cells across different sessions (Figure S5; see STAR Methods for details). In the early maps, only a small number of PCs were identified as common to both sessions (these cells are hereinafter called “common PCs”), but the fraction of common PCs increased significantly as additional PCs were imaged in the late phase of the training (Figures 3A–3C). Moreover, the fraction of PCs that had stable place fields in both sessions (“stable PCs”) also increased markedly as the training proceeded (Figures 3A, 3B, 3D, and 3E), indicating that the maps are consolidated in an experience-dependent manner. Image comparisons between adjacent sessions showed that the fractions of common cells were constant over time ($p = 0.57$, one-way ANOVA; Figure S5E). Furthermore, the fractions of common and stable PCs normalized to the number of PCs also increased significantly as the training proceeded, indicating that the training-induced increase in PC stability was not simply due to the increase in the number of PCs (Figure S6). We then calculated the fractions of stable PCs with respect to the number of common PCs as an index of stable representations at each location and found that this index was significantly higher for locations associated with the landmark or reward than for other locations (Figures 3F and 3G). In addition, the gain of PC stability between the early and late phases of training in individual animals exhibited a good linear correlation with their differences in time spent running between the two phases of the training ($r = 0.90$,

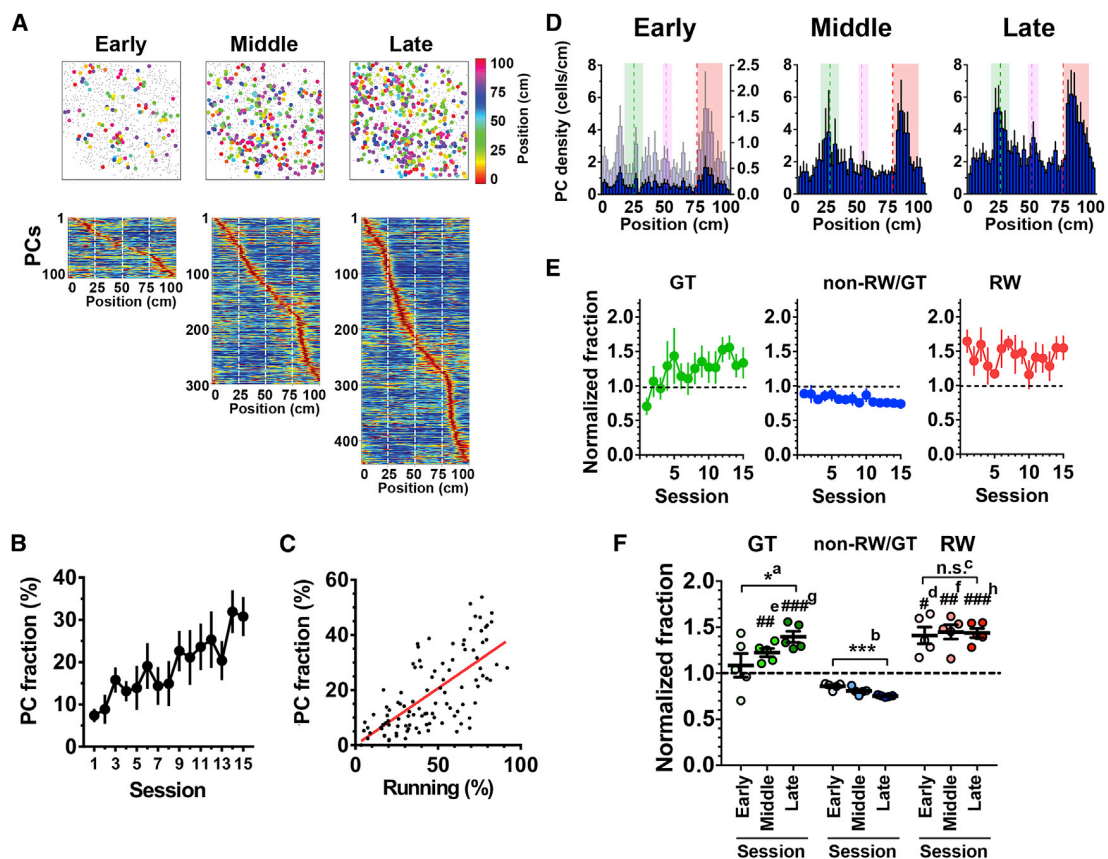


Figure 2. Over-Representations of the Reward and Landmark Locations Emerge with Different Time Courses

(A) Examples of place cell (PC) maps imaged in the same animal at the early (session 1), middle (session 9), and late (session 14) phases of training on the virtual endless linear track task (top). PCs and non-PCs are represented by filled circles of various colors and gray dots, respectively. The different colors of the filled circles represent different locations of the place fields. Heatmaps shown below are distributions of place fields of the corresponding sessions ordered by their positions (bottom).

(B) The fractions of PCs relative to the number of total identified cells imaged at each session. $n = 7$ mice.

(C) The fractions of PCs were plotted against the corresponding fractions of time spent running ($n = 105$ sessions from seven mice). The red line represents linear regression ($r = 0.59$).

(D) Histograms indicating the distribution of PCs with respect to track position in the early (session 1), middle (session 6), and late (session 12) phases of the training. The average data from seven mice are shown. For comparison, the histogram of the early phase was scaled to that of the late phase by their maximum values and is plotted in light blue on its right y axis. The green, red, and magenta dashed lines delineate the positions of the landmark, reward delivery, and boundary between different wall patterns, respectively. The areas shown in green, red, and magenta indicate those that define gate, reward, and wall cells, respectively.

(E) Hippocampal spatial representations as expressed by the fractions of gate cells (GT cells, green), non-reward/gate cells (non-RW/GT cells, blue), and reward cells (RW cells, red) relative to the number of total PCs identified in each session. Values were normalized to that obtained in the case of uniform distribution (i.e., 0.0125/bin), and values greater than 1 indicate that the locations are over-represented.

(F) Average normalized fractions of GT (green), non-RW/GT (blue), and RW (red) cells for the early, middle, and late phases of the training. *a , $p = 0.046$, $F_{(2,12)} = 3.26$; ***b , $p = 0.0003$, $F_{(2,12)} = 15.3$; n.s. c , $p = 0.96$, $F_{(2,12)} = 0.067$; one-way ANOVA, $n = 5$ sessions each; $\#^{d1}$, $p = 0.014$ versus non-RW/GT early, $F_{(1,059,4.235)} = 6.33$; $\#\#^e$, $p = 0.0037$ versus non-RW/GT middle; $\#\#^f$, $p = 0.0044$ versus non-RW/GT middle, $F_{(1,334,5.335)} = 27.8$; $\#\#\#^g$, $p = 0.0009$ versus non-RW/GT late; $\#\#\#^h$, $p = 0.0007$ versus non-RW/GT late, $F_{(1,142,4.566)} = 50.68$; one-way ANOVA; $n = 5$ sessions each.

(B and D–F) Data are expressed as mean \pm SEM.

Figure 3H). The results suggest that hippocampal place maps are more strongly stabilized if the animals learn to run the virtual linear track task more effectively.

Response Properties of RW Cells and GT Cells

In order to further characterize RW cells and GT cells, mice trained in the normal task were subjected to a task with 50%

reward delivery and a task with no gate ($n = 3$ mice). In the 50% reward task, reward was delivered randomly with 50% probability at the same delivery point, and the data were separately analyzed for rewarded trials and non-rewarded trials (Figure 4A). The fraction of time spent running increased slightly but not significantly in non-rewarded trials (rewarded trials $66.9\% \pm 4.2\%$ versus non-rewarded trials, $74.9\% \pm 3.9\%$, $p = 0.076$,

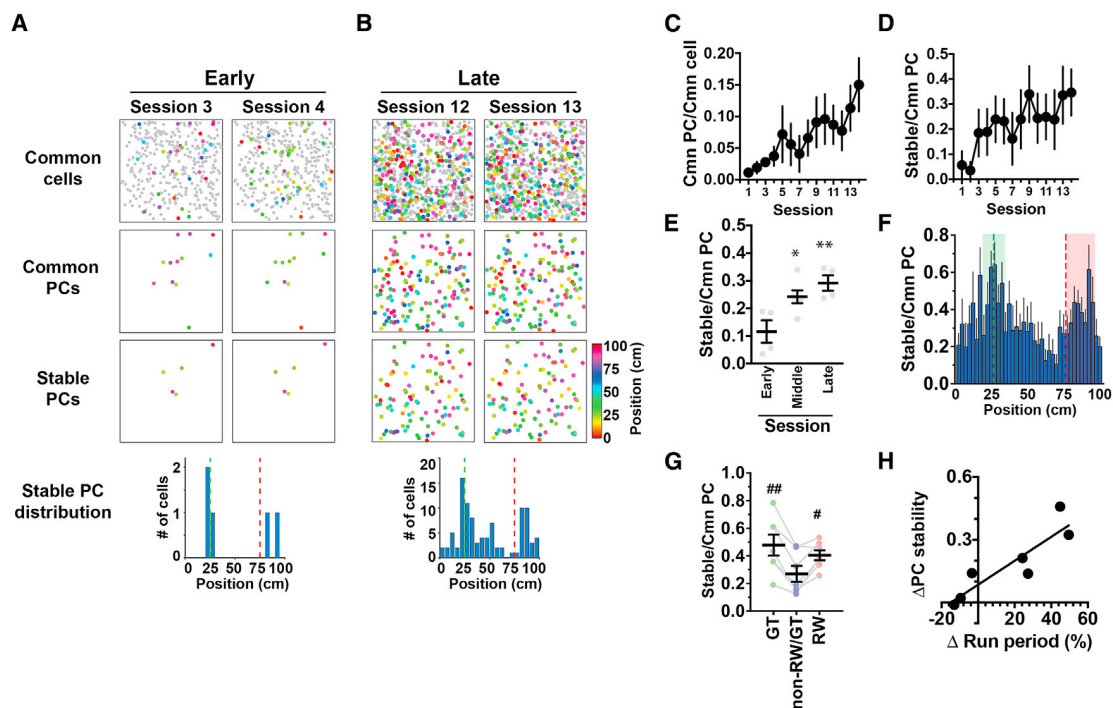


Figure 3. RW Cells and GT Cells Form Stable Singularities in Hippocampal Cognitive Maps

(A) Example hippocampal CA1 PC maps imaged in two consecutive sessions in the early phase of the training. Maps shown on top, middle, and bottom present cells identified in common to both sessions (common cells), cells identified as PCs in both sessions (common PCs), and cells identified as PCs with stable (< 10 cm difference) place fields in both sessions (stable PCs), respectively. The histogram shown at the bottom indicates the distributions of the stable PCs against track position.

(B) PC maps imaged in the late phase of training in the same animal as presented in (A).

(C) The fraction of common PCs relative to the number of common cells identified in the two consecutive sessions that were compared. The x axis indicates the earlier of the two sessions that were compared. $n = 7$ mice.

(D) PC stability calculated as the fraction of stable PCs relative to the number of common PCs identified in the two consecutive sessions that were compared.

(E) Average PC stability in the early (sessions 1–4, numbered according to the earlier of the two sessions that were compared), middle (sessions 5–10), and late (sessions 11–14) phases of training. * $p = 0.021$ versus early, ** $p = 0.0050$ versus early, $F_{(2,11)} = 8.04$, one-way ANOVA; $n = 4, 6,$ and 4 session pairs for the early, middle, and late phases, respectively.

(F) The average fractions of stable PCs relative to the number of common PCs plotted against track position. Values were calculated from data across all sessions and averaged for seven mice. The fractions were determined using the number of common PCs obtained separately for each spatial bin in the earlier of the two sessions compared. The green and red dashed lines delineate the positions of the landmark and reward delivery, respectively. The areas shown in green and red are those that define GT cells and RW cells, respectively.

(G) The average fractions for GT, non-RW/GT, and RW cells. # $p = 0.015$ versus non-RW/GT cells, ## $p = 0.0025$ versus non-RW/GT cells, $F_{(1,215,7,290)} = 11.6$, one-way ANOVA, $n = 7$ mice.

(H) The relationship between PC stability and learning to run along the track. The x axis presents the task performance of each mouse measured by the difference in the fraction of time spent running between the early (average of sessions 1–5) and late (average of sessions 11–15) phases of the training. The y axis presents the difference in PC stability between the early and late phases of training.

(C–G) Data are expressed as mean \pm SEM.

$t_{(5)} = 2.23$, paired two-tailed t test, $n = 6$ sessions from three mice). Licking within the reward zone decreased in non-rewarded trials with marginal statistical significance, whereas licking immediately before and far outside the reward zone was not altered between rewarded trials and non-rewarded trials (Figure 4B). Consistent with the notion that the presence of reward is encoded by the increased density of RW cells, the representation of the reward zone substantially decreased in non-rewarded trials (Figure 4C; rewarded trials, 1.19 ± 0.12 versus non-rewarded trials, 0.50 ± 0.09 -fold relative to a uniform distribution [see Figure 2E legend], $p = 0.0031$, $t_{(5)} = 5.35$, paired two-tailed t test, $n = 6$ sessions from three mice). Thus, RW cells

appear to be enriched in the reward location due to the presence of reward.

We then tracked the changes in RW cells in non-rewarded trials. Twenty-one percent of RW cells (96 out of 467 cells) still had their place fields within the reward zone in non-rewarded trials (RW-RW cells; Figures 4D and 4E), while 37% (173 cells) became PCs that encoded other parts of the track and 42% (198 cells) lost their properties as PCs. The consistency of RW cells was markedly lower than the other categories of cells; 59% of GT cells (251 out of 423 cells) and 51% of non-RW/GT cells (552 out of 1,088 cells) in rewarded trials remained GT cells and non-RW/GT cells, respectively, in non-rewarded trials. However,

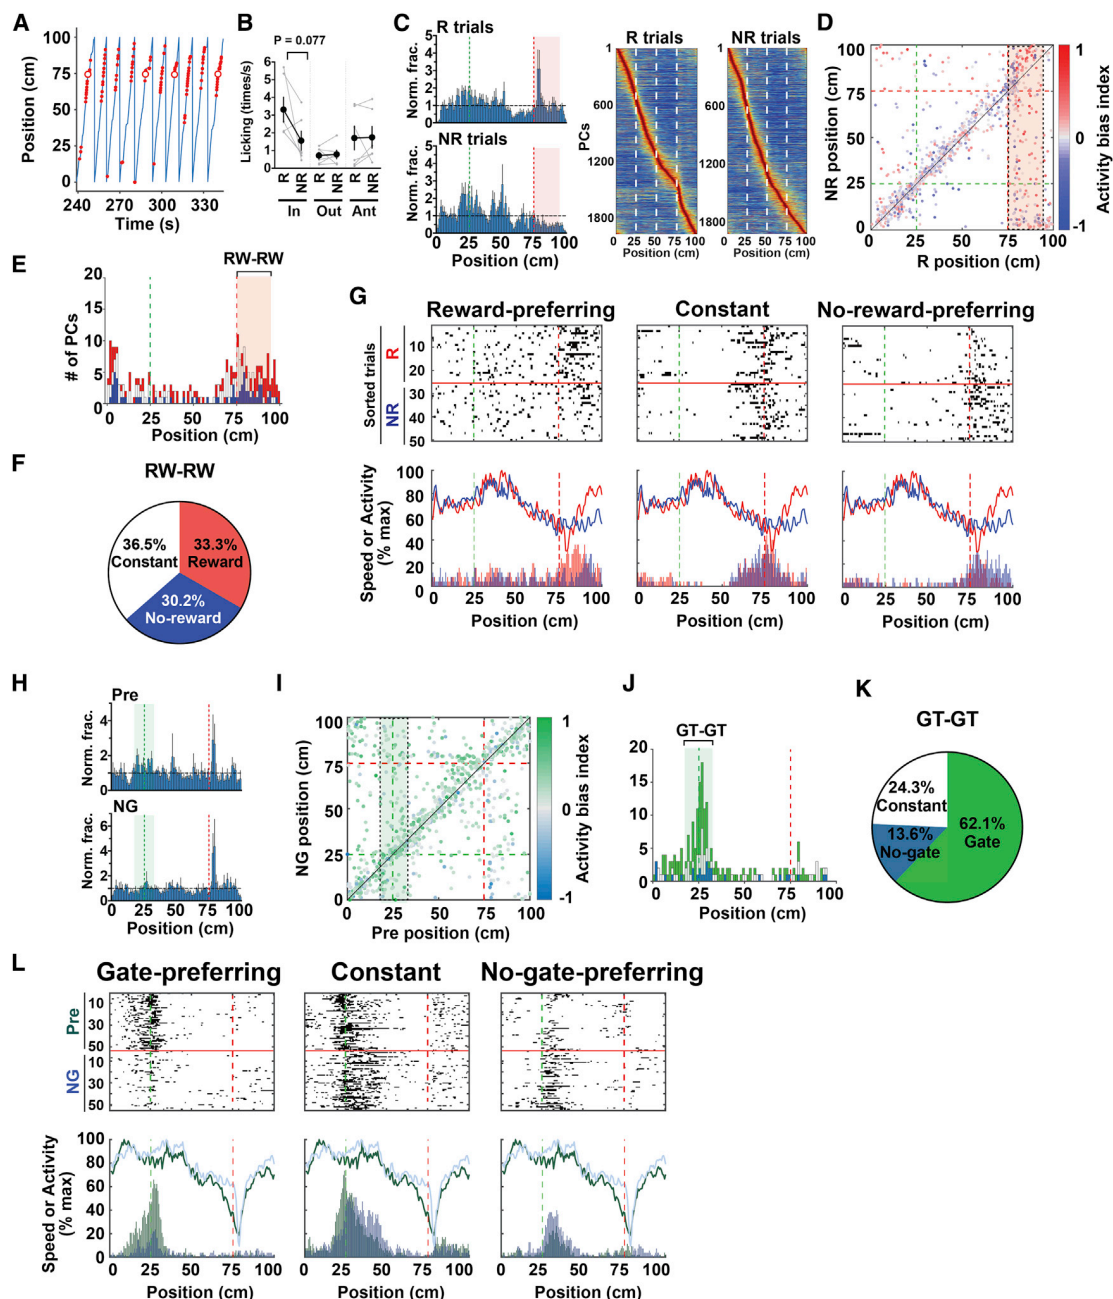


Figure 4. Response Properties of RW Cells and GT Cells

(A) An example segment of 50% reward delivery sessions. The blue line represents the animal's position. Red dots and red circles indicate licking and reward delivery, respectively.

(B) Licking within the reward zone (in), immediately before the reward zone (ant), and elsewhere (out) in rewarded (R) and non-rewarded (NR) trials (R in versus NR in, $p = 0.077$, $t_{(5)} = 2.22$, paired two-tailed t test, $n = 6$ sessions each). Data are expressed as mean \pm SEM.

(C) Histograms indicating the distribution of PCs with respect to track position for R and NR trials (left). The y axis represents PC fractions normalized relative to a uniform distribution. The average data of six sessions using three mice are shown. The green and red dashed lines delineate the positions of the landmark and reward delivery, respectively. The red area defines RW cells. Heatmaps shown on the right highlight the distributions of place fields in R and NR trials, ordered by their positions (right). $n = 1,978$ cells (R trials) and 1,915 cells (NR trials), respectively, from six sessions using three mice.

(D) A scatterplot showing the distribution of place-field positions in R and NR trials. Each dot represents the position of a cell's place field, and the color indicates the activity bias index of the cell. Cells from all relevant sessions are shown. The green and red dashed lines delineate the positions of the landmark and reward delivery, respectively. The red area enclosed by a black dashed rectangle defines RW cells in R trials that remained PCs in NR trials, and a histogram of the distribution of their place fields along the NR axis is shown in (E). The same convention applies hereinafter.

(legend continued on next page)

1.78-fold enrichment of RW–RW cells relative to a uniform distribution of PCs (96 out of 269 cells; [Figure 4E](#)) indicates that a part of RW cells still tend to be RW cells even in the absence of reward, which likely reflects a general tendency of PCs that their fields tend to remain in the same locations in non-rewarded trials ([Figure 4D](#)).

To further investigate how reward modulates RW cell activity, we compared the activity of each RW–RW cell between rewarded and non-rewarded trials and divided them into three groups according to the direction and extent of their activity biases ([Figures 4F and 4G](#)). Interestingly, we found that 30.2% of RW–RW cells (29 out of 96 cells) increased their activity in the absence of reward (no-reward-preferring RW cells), while approximately similar proportions of cells exhibited constant (termed “constant RW cells” after their constant activity regardless of the presence or absence of reward; this terminology does not mean they continuously encode the reward location, 36.5%) or decreased activity (reward-preferring RW cells, 33.3%; [Figures 4F and 4G](#)).

Next, we examined the effect of the visual landmark on GT cell activity in the no-gate task, in which the green gate was removed from its location throughout a test session that was conducted immediately after a preceding control session (pre). The fraction of time spent running did not change notably in no-gate sessions (pre, $76.4\% \pm 8.9\%$, versus no gate, $79.8\% \pm 9.5\%$, $n = 3$ sessions each from three mice). The representation of GT location decreased in no-gate sessions ([Figure 4H](#); pre, 1.37 ± 0.19 , versus no gate, 1.01 ± 0.06 -fold relative to a uniform distribution, $n = 3$ sessions each), supporting the idea that the increased density of GT cells encodes the presence of the landmark. In no-gate sessions, 36% of GT cells identified in pre (103 out of 287 cells, 3.3-fold relative to a uniform distribution) maintained their place fields within the area that defined GT cells (GT–GT cells; [Figures 4I and 4J](#)), while 36% (104 out of 287 cells) had significant place fields in other parts of the track and the remaining 28% (80 out of 287 cells) lost their properties as PCs. Thus, similar to RW cells, GT cells appear to accumulate at the landmark location when the landmark is present. Moreover, a part of GT cells tend to be GT

cells in the absence of landmark as PCs generally tend to have their fields at the same locations ([Figures 4I and 4J](#)). GT–GT cells were then categorized similarly into three groups by comparing their activity between the preceding control and no-gate sessions ([Figures 4K and 4L](#)). Notably, 62.1% (64 out of 103 cells) of GT–GT cells were gate-preferring GT cells, whereas only 13.6% of them (14 out of 103 cells) were no-gate-preferring GT cells ([Figures 4J and 4K](#)). This asymmetric modulation of GT cells by visual landmark indicates that the majority of GT cells decrease their activity when the visual landmark is not available.

Persistent Subsets of Cells Mediate Over-Representations of Reward and Landmark Locations

To elucidate the cellular mechanism of map establishment, we analyzed the transitions among RW, GT, non-RW/GT cells, and non-PCs between adjacent sessions ([Figure 5](#)). Although the conversion of PCs from non-PCs exhibited significant biases toward RW cells ([Figure 5A](#)), a subpopulation of PCs that became non-PCs also exhibited a similarly biased distribution ([Figure 5B](#)). This implies that the disproportionate formation of PCs is equilibrated by a similar bias in the disappearance of part of the PCs ([Figure 5C](#)). The distribution of PCs derived from former non-RW/GT cells appeared not to be significantly biased toward RW or GT cells ([Figure 5D](#)). In contrast, the distribution of PCs derived from former GT or RW cells was substantially biased toward the location by which each of the two PC subcategories was defined ([Figures 5E and 5F](#)). The persistence and relative independence of RW cells and GT cells demonstrate that selective stabilization of RW and GT cell subsets, rather than lateral recruitment of non-RW/GT cells, is a major mechanism that underlies the establishment of hippocampal salience maps. A further analysis of the cellular origins of RW cells and GT cells confirmed that most RW cells and GT cells in the middle and late sessions were cells of the same subcategories in preceding sessions, although these cells in the early sessions mainly originated from net formation from non-PCs ([Figures 5G and 5H](#)). Collectively, these findings indicate that hippocampal map formation consists of two steps: first, early populations of RW cells

(E) A histogram showing the distribution of place-field positions of PCs in NR trials that were RW cells in R trials. Red, white, and blue segments of the stacked bar represent the proportions of reward-preferring RW cells, constant RW cells, and no-reward-preferring RW cells, respectively, to the total number of cells in that bin. The red area defines RW cells in R trials that persisted as RW cells in NR trials (RW–RW), and the proportions of reward-preferring RW cells, constant RW cells, and no-reward-preferring RW cells in these cells are shown in (F).

(F) Proportions of reward-preferring RW cells, constant RW cells, and no-reward-preferring RW cells among RW–RW cells.

(G) Examples of reward-preferring RW cells (left), constant RW cells (middle), and no-reward-preferring RW cells (right). The three example cells shown were imaged during the same session. Top panels show raster plots against position, in which the plots for R and NR trials are divided by red horizontal lines and separately sorted. Bottom panels show activity histograms of each example cell against position for R (red) and NR (blue) trials and running speed averaged separately for R (red) and NR (blue) trials during the corresponding session.

(H) Histograms indicating the distribution of PCs with respect to track position for preceding control (pre) and no-gate (NG) sessions. The y axis represents PC fractions normalized relative to a uniform distribution. The average data from three sessions using three mice are shown. The green area defines GT cells.

(I) A scatterplot showing the distribution of place-field positions in pre and NG sessions. The green area enclosed by a black rectangle defines GT cells in pre sessions that remained PCs in NG sessions. A histogram of the distribution of their place fields along the NG axis is shown in (J).

(J) A histogram showing the distribution of place-field positions of PCs in NG sessions that were GT cells in pre sessions. Green, white, and turquoise segments of the stacked bar represent the proportions of gate-preferring GT cells, constant GT cells, and no-gate-preferring GT cells, respectively, to the total number of cells in that bin. The green area defines GT cells in pre that persisted as GT cells in the NG condition (GT–GT), and the proportions of gate-preferring GT cells, constant GT cells, and no-gate-preferring GT cells among these cells are shown in (K).

(K) Proportions of gate-preferring GT cells, constant GT cells, and no-gate-preferring GT cells among GT–GT cells.

(L) Examples of gate-preferring GT cells (left), constant GT cells (middle), and no-gate-preferring GT cells (right). Top panels show raster plots against position in which pre and NG sessions are shown above and below the red horizontal lines, respectively. Bottom panels show running speed averaged for pre (green) and NG (turquoise) sessions and activity histograms of each example cell against position for pre (green) and NG (turquoise) sessions.

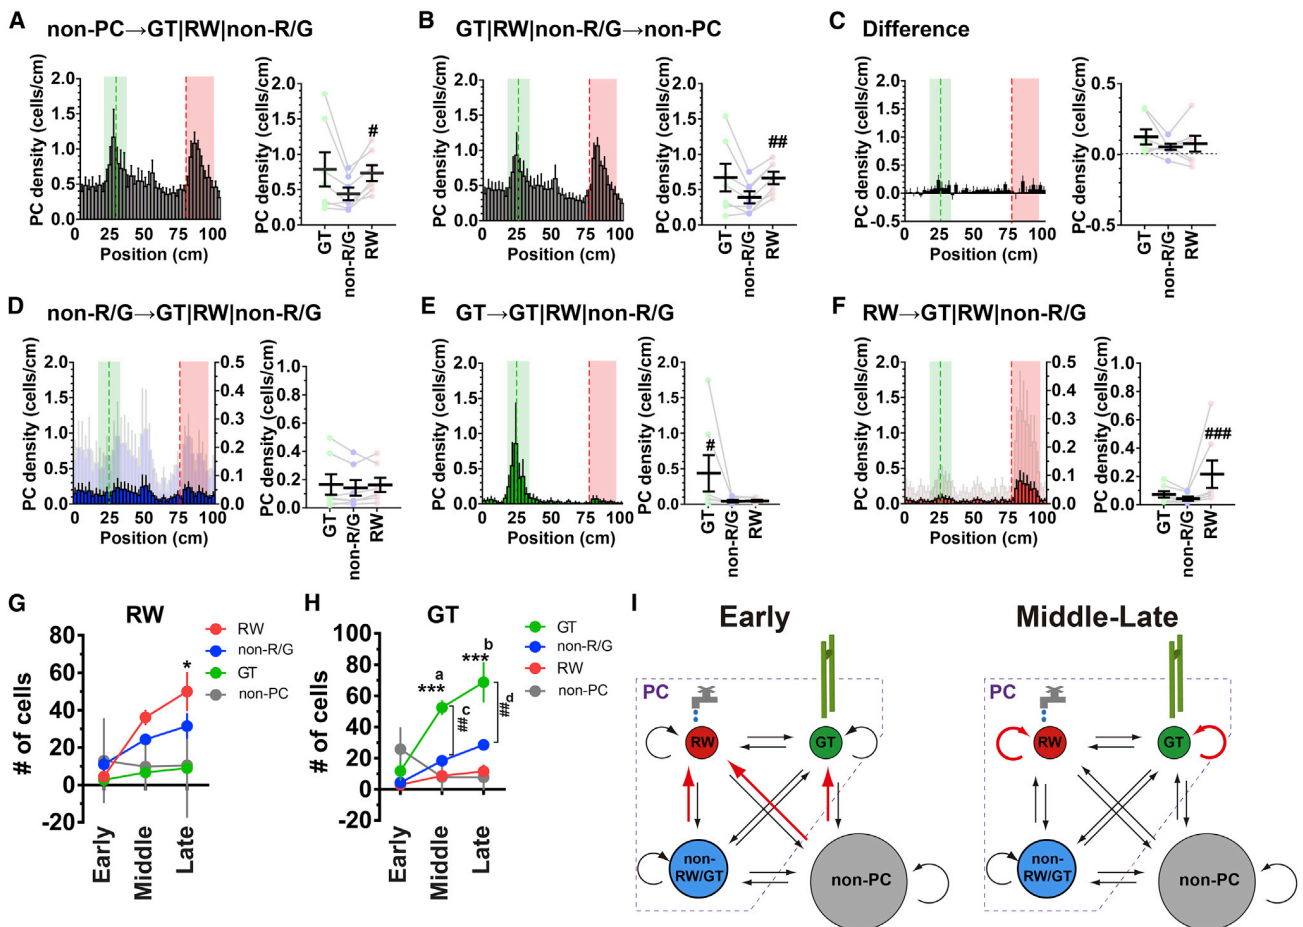


Figure 5. Persistent Neuronal Subsets Mediate Over-Representations of the Reward and Landmark Locations

(A) Formation of different PC categories from non-PCs. (Left) A histogram showing the distribution of PCs that were non-PCs in the previous sessions against position. The values were calculated from data across all sessions and averaged for seven mice. For comparison, the histograms shown in (A), (B), and (D–F) are plotted on the left y axes on the same scale. In addition, the histograms in (D) and (F) were scaled by their maximum values and are plotted in a light color on the right y axes. (Right) The average cell density of each PC subcategory formed from non-PCs. # $p = 0.023$ versus non-RW/GT, $\chi^2_{(2)} = 8.00$; Friedman test, $n = 7$ mice. (B) Elimination of different PC categories by conversion to non-PCs. (Left) A histogram showing the distribution of PCs that became non-PCs in subsequent sessions against position. (Right) The average cell density of each PC subcategory that became non-PCs in the subsequent sessions. ## $p = 0.0099$ versus non-RW/GT, $\chi^2_{(2)} = 8.86$; Friedman test, $n = 7$ mice. (C) Net PC formation. (Left) A histogram of the difference obtained by subtracting the histogram in (B) from that in (A). (Right) The average cell density of PCs of each subcategory.

(D) Transition and stability of non-RW/GT cells. (Left) A histogram showing the distribution of PCs that were non-RW/GT cells in the previous sessions against position. (Right) The average cell density of each PC subcategory that was derived from non-RW/GT cells.

(E) Transition and stability of GT cells. (Left) A histogram showing the distribution of PCs that were GT cells in the previous sessions against position. (Right) The average cell density of each PC subcategory that was derived from GT cells. # $p = 0.023$ versus non-RW/GT, $\chi^2_{(2)} = 7.14$, Friedman test, $n = 7$ mice.

(F) Transition and stability of RW cells. (Left) A histogram showing the distribution of PCs that were RW cells in the previous sessions against track position. (Right) The average cell density of each PC subcategory that was derived from RW cells. ### $p = 0.0005$ versus non-RW/GT, $\chi^2_{(2)} = 14.0$, Friedman test, $n = 7$ mice.

(G) The numbers of RW cells that were derived from RW cells, GT cells, non-RW/GT cells, and non-PCs during the early, middle, and late phases of training. The transitions from non-PCs represent net formation after subtracting the disappearance of the existing PCs. * $p = 0.021$ versus early RW, $F_{(3,33)} = 3.16$; two-way ANOVA, $n = 4$ –6 session pairs.

(H) The numbers of GT cells that were derived from RW cells, GT cells, non-RW/GT cells, and non-PCs during the early, middle, and late phases of training. **** $p = 0.0003$ versus early GT, *** $p < 0.0001$ versus early GT, $F_{(3,33)} = 17.6$; ## $p = 0.0011$ versus middle non-RW/GT, ### $p = 0.0018$ versus late non-RW/GT, $F_{(2,11)} = 4.96$; two-way ANOVA, $n = 4$ –6 session pairs.

(A–H) Data are expressed as mean \pm SEM.

(I) A model for the formation of the hippocampal salience map. The key processes in each phase are shown in red. During the early phase of map formation, *de novo* conversion from non-PCs primarily sets out a prototype map (left). Selective consolidation of GT cells and RW cells subsequently plays a dominant role in the establishment and maintenance of the salience map during the middle and late phases of training (right).

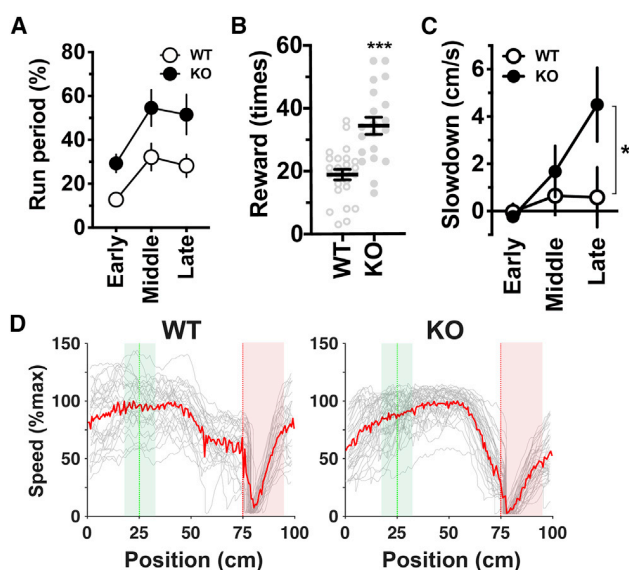


Figure 6. Enhanced Goal-Directed Behavior of Shank2-Deficient Mice

(A) The fraction of time spent running for the early (sessions 1–5), middle (sessions 6–10), and late (sessions 11–15) phases of training in wild-type (WT) and Shank2-deficient (KO) mice (WT versus KO, $p = 0.025$, $F_{(1,7)} = 8.14$, two-way ANOVA, $n = 5$ WT and 4 KO mice).
 (B) The number of rewards obtained in sessions in the late training phase. *** $p < 0.0001$ versus WT, $t_{(43)} = 5.07$; unpaired two-tailed t test, $n = 25$ and 20 sessions from 5 WT and 4 KO mice, respectively.
 (C) Slowdown before the reward delivery point * $p = 0.031$, $F_{(1,7)} = 2.199$, $t_{(21)} = 2.811$, two-way ANOVA with the Holm-Sidak test, $n = 5$ WT and 4 KO mice.
 (A–C) Data are expressed as mean \pm SEM.
 (D) Examples of trial-by-trial (gray) and average (red) running speed plotted against position for WT (left) and KO (right) mice in single sessions in the late training phase.

and GT cells are primarily derived from non-PCs until the number of cells is sufficient; second, thereafter, each subcategory becomes its own dominant source in the middle and late sessions (Figure 5I).

Since the analyses thus far considered active cells that exhibited detectable fluorescence changes, we further investigated the dynamics of inactive “silent” cells during the middle and late phases of map formation (for details, see STAR Methods). This independent analysis replicated the findings of selective stabilization of RW cells and GT cells and unbiased recruitment of non-RW/GT cells to these cell categories during map formation (recurrence of GT cells, $68.1\% \pm 7.0\%$ [4.5-fold relative to a uniform distribution]; recurrence of RW cells, $50.8\% \pm 7.5\%$ [2.5-fold relative to a uniform distribution]; recruitment of non-RW/GT cells to RW cells, $13.0\% \pm 3.2\%$ [0.9-fold relative to a uniform distribution]; recruitment of non-RW/GT cells to GT cells, $14.2\% \pm 2.2\%$ [0.7-fold relative to a uniform distribution]; $n = 10$ sessions of 893–1,219 cells from two mice). The analysis of silent cell dynamics revealed that the majority of silent cells ($65.7\% \pm 2.2\%$, mean \pm SD, $n = 10$ sessions of 1,576–1,682 cells from two mice) remained silent cells in the subsequent sessions. Among the silent cells that exited the pool of silent cells, $29.9\% \pm 4.6\%$ and $70.1\% \pm 4.6\%$ of them became PCs and

non-PCs, respectively, while $29.9\% \pm 4.8\%$ and $70.1\% \pm 4.8\%$ of active-cell-derived silent cells returned from PCs and non-PCs, respectively. Of those silent-cell-derived PCs, $24.1\% \pm 4.8\%$ (2.0-fold representation versus non-RW/GT cells) and $22.4\% \pm 4.5\%$ (1.4-fold versus non-RW/GT cells) became GT cells and RW cells, respectively, whereas $22.3\% \pm 3.2\%$ (1.8-fold versus non-RW/GT cells) and $22.6\% \pm 3.6\%$ (1.3-fold versus non-RW/GT cells) of PC-derived silent cells returned from GT cells and RW cells, respectively. These results demonstrate that salient-location-biased PC formation from silent cells is equilibrated by a similarly biased elimination, as in the case of non-PCs (Figures 5A–5C). Silent cells and non-PCs thus share a similar mode of dynamics when they convert to and revert from PCs.

Separable Reward and Landmark Coding Revealed by Shank2-Deficient Mice

To dissect the molecular mechanisms underlying the over-representation of salient locations, we examined mice deficient in *Shank2*, a gene that encodes a scaffold protein found at the excitatory postsynaptic density (PSD). These mice mimic the microdeletion of *Shank2* identified in a case of ASD with mild intellectual disability (Berkel et al., 2010), and exhibit ASD-like behaviors, including impaired social interaction and hyperactivity, reduced N-methyl-D-aspartate (NMDA) receptor function, and impaired long-term potentiation and long-term depression at hippocampal CA1 synapses (Won et al., 2012), providing a unique opportunity to test whether the integrity of postsynaptic signaling and synaptic plasticity is necessary for hippocampal map formation and plasticity.

Shank2-deficient mice displayed not only an increased fraction of time spent running (Figure 6A; knockout [KO] versus wild type [WT], $p = 0.025$, $F_{(1,7)} = 8.14$, two-way ANOVA, $n = 5$ WT and 4 KO mice) and acquisition of an increased number of rewards (Figure 6B) but also enhanced anticipatory slowdown before the reward zone (Figures 6C and 6D). This finding suggests that the hyperactivity of these mice is not merely a reckless run but rather reflects a more frequent repetition of highly stereotyped goal-directed behavior. However, the possibility that a higher level of attention and/or increased anxiety also contribute to this behavior cannot be excluded (Won et al., 2012). Consistent with the running-dependent increase in the fraction of PCs (Figure 2C), the overall PC fraction was elevated in Shank2-deficient mice compared to WT mice (Figure 7A; KO versus WT, $p = 0.025$, $F_{(1,7)} = 8.07$, two-way ANOVA, $n = 5$ WT and 4 KO mice). Remarkably, the hippocampal maps of Shank2-deficient mice exhibited substantially enhanced over-representation of the reward location, whereas their learning-induced over-representation of the landmark location was selectively and completely lacking (Figures 7B, 7H, and 7I). It is unlikely that the absence of over-representation of the landmark location in Shank2-deficient mice was due to their potential impairment in vision, as they exhibited anticipatory behavior before the landmark even when the position of the landmark that indicated the reward location was randomly shifted (Figure S7). Compared to the average non-RW/GT cell fractions, the average RW cell fractions of each mouse had an increased correlation with the average number of rewards obtained throughout the training (Figure 7C), which

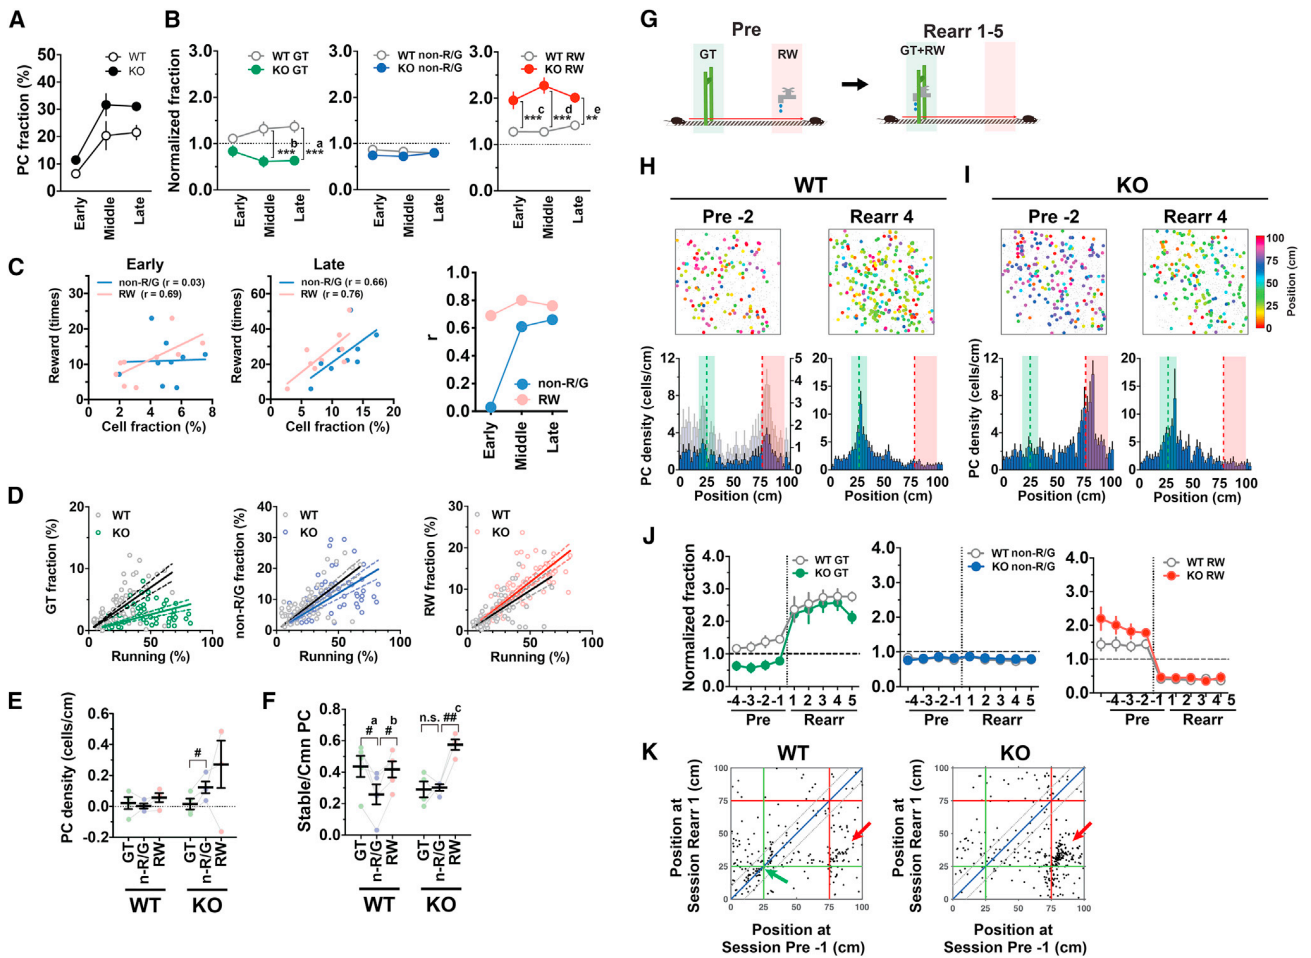


Figure 7. Loss of Over-Representation of the Landmark Location and Augmented Over-Representation of the Reward Location in Shank2-Deficient Mice

(A) PC fractions in the early, middle, and late phases of training ($n = 5$ WT and 4 KO mice).

(B) Normalized fractions of GT (left), non-RW/GT (middle), and RW (right) cells relative to the number of total PCs in each phase of training. **** a , $p = 0.0009$, **** b , $p < 0.0001$, $F_{(1,7)} = 22.39$; **** c , $p = 0.0014$, ** d , $p = 0.0003$, ** e , $p = 0.0003$, $F_{(1,7)} = 33.33$; two-way ANOVA, $n = 5$ WT and 4 KO mice.

(C) Correlations between the number of rewards obtained and the fractions of RW or non-RW/GT cells. The data for the early (average of sessions 1–5) and late (average of sessions 11–15) phases of training are shown in the left and middle panels, respectively. Each point represents the data from an individual animal, and the points and regression lines shown in red and blue represent those for RW and non-RW/GT cells, respectively. The graph shown right indicates the correlation coefficients of linear regressions for RW (red) and non-RW/GT (blue) cell fractions in the early, middle, and late phases of training. $n = 5$ WT and 4 KO mice.

(D) Scatterplots showing the relationship between the fraction of GT (left), non-RW/GT (middle), and RW (right) cells with respect to total identified cells and the fraction of time spent running. Each symbol represents an individual session. Solid and dotted lines indicate linear regression and 95% confidence intervals, respectively.

(E) Net formation of GT, non-RW/GT (n-R/G), and RW cells during the early training phase in WT and KO mice. # $p = 0.046$, $F_{(1,0.25,3,075)} = 2.35$, $t_{(3)} = 4.979$; one-way ANOVA with the Holm-Sidak test, $n = 4$ KO mice.

(F) The average PC stability for GT, non-RW/GT, and RW cells in WT and KO mice. # a , $p = 0.036$, # b , $p = 0.034$, $F_{(1,938,7,753)} = 11.36$, one-way ANOVA, $n = 5$ WT mice. ## c , $p = 0.0011$, $F_{(1,04,3,12)} = 16.5$, one-way ANOVA, $n = 4$ KO mice.

(G) Design of the reward-rearrangement task. Mice were first trained on the standard linear track, which included a visual landmark (GT) and reward delivery (RW) at separate locations, in preceding control sessions (pre). Once training was complete, the location of reward delivery was shifted to match the location of the landmark (GT+RW), and the mice were retrained in this new arrangement for the following five sessions (rearr 1–5).

(H and I) Example PC maps (top) and histograms indicating the average distribution of PCs with respect to position (bottom) before (pre -2) and after (rearr 4) reward rearrangement of WT (H) and KO mice (I).

(J) Normalized representations for GT (left), non-RW/GT (middle), and RW (right) cells in each session of reward-rearrangement experiments ($n = 5$ WT and 4 KO mice).

(K) Two-dimensional plots showing the positions of place fields of common PCs immediately before (pre -1) and after (rearr 1) the reward rearrangement in WT (left, $n = 259$ cells from five mice) and KO mice (right, $n = 339$ cells from four mice). Red arrows indicate clusters of cells that persistently encode the reward locations. The green arrow indicates a cluster of stable GT cells in WT mice.

(A), (B), (E), (F), and (H)–(J) Data are expressed as mean \pm SEM.

indicates that RW cell fractions are a better indicator of the amount of rewarded experience. The fraction of GT cells for a given amount of spatial behavior was substantially reduced in Shank2-deficient mice, while those of non-RW/GT and RW cells were slightly decreased and increased, respectively (Figure 7D). Consistent with these findings, the net formation of GT cells from non-PCs during the early phase was significantly less than that of non-RW/GT cells in these mice (Figure 7E). Moreover, the stability of GT cells was reduced to a level comparable to that of non-RW/GT cells (Figure 7F). These findings indicate that Shank2 plays an indispensable role in both initial formation and subsequent stabilization of GT cells. Finally, we tested whether Shank2 is required for the reorganization of preformed maps in reward-rearrangement experiments, where mice trained on the standard virtual linear track for 15 sessions were further trained for the following five sessions on the same track with the exception that the location of reward was shifted to match the landmark location (Figure 7G). In contrast to the pivotal role of Shank2 in synaptic plasticity, the peak of the density of RW cells rapidly shifted to the new reward location in Shank2-deficient mice as in WT mice (Figures 7H–7J); thus, Shank2 is dispensable for reward-relocation-induced plasticity of RW cells (Figure 7K). In summary, the results clearly demonstrate that the two forms of hippocampal over-representation are mechanistically separable. Shank2 is required for learning-induced over-representation of landmark locations but not for rapid emergence and plasticity of over-representation of reward locations. The absence of Shank2 elicits abnormally enhanced over-representation of reward locations, and this “super-representation” is associated with enhanced running and goal-anticipation behavior and increased acquisition of rewarded outcomes.

DISCUSSION

In this study, we found that hippocampal over-representations of locations associated with reward and a visual landmark are mediated by separable subsets of neurons with different time courses of emergence and different underlying molecular mechanisms. The disproportionate hippocampal maps were established as a result of *de novo* conversion of the cells that encode each salient location from nonspatially tuned cells in the early phase, followed by their selective stabilization in the late phase. The two types of hippocampal over-representations were distinguishable by the necessity of Shank2; the learning-induced over-representation of the landmark location is dependent on Shank2, whereas rapid emergence and plasticity of over-representation of the reward location is independent of Shank2. Moreover, the abnormally distorted representation toward the reward location in the Shank2-deficient mouse model of ASD was associated with enhanced running and goal-anticipation behavior and increased acquisition of the rewarded experience, which suggests that aberrant hippocampal mapping of salient features may underlie cognitive and behavioral abnormalities in a subset of ASD cases.

In contrast to the preceding studies that examined only a single type of salient feature (O’Keefe and Conway, 1978; Wiener et al., 1989; Hetherington and Shapiro, 1997; Hollup et al., 2001; Dupret et al., 2010; Danielson et al., 2016; Geiller et al., 2017; Zaremba et al., 2017; Gauthier and Tank, 2018, Bourbou-

lou et al., 2019), our study investigated the impacts of reward and a landmark separately side by side and provided strong evidence that multiple mechanistically distinct over-representations encode different types of salient features within a single hippocampal map. In our model, the number or fraction of a particular functional subset of neurons is determined by equilibrium between their stability and instability. During the early phase of map formation, new conversion of PCs from non-PCs creates a sparse prototype map in which “seed” GT cells and RW cells are laid out. The higher stability of each of these cell categories shifts the equilibrium upward to an elevated steady-state level and enables induction or maintenance of increased numbers of cells during the later phases. Representations of salient locations by these subsets of neurons constitute more stable singularities within hippocampal maps and may provide substrates by which different kinds of salient experience form lasting memory traces. Immediate emergence of over-representation for motivational salience is favorable for the rapid location of a source of positive reinforcement within the subject’s environment, whereas gradual mapping of environmental salience presumably reflects experience-dependent learning of the environment that could contribute to some form of landmark-based navigation (Sato et al., 2017; Le Merre et al., 2018). Thus, the multiplicity of functionally distinct salience coding revealed in our study expands the capacity of hippocampal maps to encode the presence and absence of a variety of nonspatial features onto their spatial representations.

We identified *Shank2* as a single gene that is essential for the establishment of hippocampal over-representation. Since the global KO of *Shank2* is thought to cause many changes throughout the brain, the relevant circuit mechanisms need to be elucidated in future studies. One possibility is that since the Shank2-deficient mice used in our study show regular synapse numbers yet impaired NMDA receptor function and synaptic plasticity in the hippocampal CA1 (Won et al., 2012), this suggests that normal NMDA receptor-mediated signaling fulfilled by Shank2 may play a role in the formation and stabilization of GT cells. Detailed circuit and molecular dissection of pathways that convey distinct salience signals for reward and landmark also remains to be conducted. Neurons in the lateral entorhinal cortex fire in the vicinity of objects when they are introduced into the environment (Deshmukh and Knierim, 2011; Tsao et al., 2013; Basu et al., 2016; Wang et al., 2018), which makes them a putative source of the salience signal for landmarks. The selective loss of landmark over-representation unveiled in the present study might be implicated in a specific type of navigation impairment in humans, landmark agnosia, which involves the temporal lobe and hippocampus as its major neural correlates (van der Ham et al., 2017), and could also shed light on spatial behavior in ASD (Smith, 2015). Salience signals may modulate PC stability via not only an NMDA-receptor-dependent mechanism but also a dopamine-dependent mechanism (Kentros et al., 1998, 2004). Loss of Shank2 did not abolish the rapid emergence and plasticity of RW cells. While a mouse model of 22q11.2 deletion syndrome that affects more than 20 genes shows deficits in hippocampal reward over-representation and goal-oriented learning (Zaremba et al., 2017), a specific deletion of the single *Shank2* gene in our study leads to substantially

augmented reward over-representation and enhanced running and goal-anticipation behavior, demonstrating that hippocampal reward representation and related behavior can be altered bidirectionally in mutually opposite directions in different mouse models of neurodevelopmental disorders. The appearance of reward over-representation preceding the emergence of anticipatory behavior during training raises the possibility that the former could play a role in the occurrence of the latter. Hypermotivated behavior and increased ventral striatal function have been reported in other Shank2-deficient rodent models (Pappas et al., 2017; Modi et al., 2018), suggesting a hypothesis that enhanced hippocampal reward mapping in the absence of Shank2 contributes to the development of augmented goal-anticipation behavior via increased interaction with striatal motivational circuits (Lansink et al., 2009).

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead Contact
 - Materials Availability
 - Data and Code Availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Generation of Thy1-G-CaMP7 transgenic mice
 - Analysis of transgene expression
 - Surgery
 - VR set-up
 - Behavior
 - Imaging
- QUANTIFICATION AND STATISTICAL ANALYSIS
 - Image analysis
 - Analysis of place fields
 - Alignment of cells across sessions
 - Statistics

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at <https://doi.org/10.1016/j.celrep.2020.107864>.

ACKNOWLEDGMENTS

We thank Charles Yokoyama, Thomas J. McHugh, Shigeyoshi Fujisawa, Liset Menendez de la Prida, and Hokto Kazama for helpful comments and discussion. This work was supported by Precursory Research for Embryonic Science and Technology (PRESTO) JPMJPR12A1 from the Japan Science and Technology Agency (JST), KAKENHI grants 24700403, 25116528, 26115530, 17H05985, 19H04942, and 20H03550 from the Ministry of Education, Culture, Sports, Science and Technology (MEXT)/Japan Society for the Promotion of Science (JSPS), and a grant from the Senshin Medical Research Foundation to M.S.; RIKEN, NIH grant R01DA17310, KAKENHI grants 22110006, 16H01292, and 18H05434, a Human Frontier Science Programme grant, and Fujitsu Laboratories to Y.H.; KAKENHI grants 25830023, 15H01571, 17H05695, and 19K16293 to K.M.; KAKENHI grant 26870577 to T.T.; KAKENHI grant 15H04265 to T.F.; KAKENHI grants 26115504, 15H05723, and 16H06536, and the program for Brain Mapping by Integrated Neurotechnol-

gies for Disease Studies (Brain/MINDS) from MEXT and the Japan Agency for Medical Research and Development (AMED) to J.N.; Regional Innovation Cluster Program grant (City Area Type, Central Saitama Area) from MEXT to J.N. and M.O.; and grants from the VW Foundation (ZN2632), BCCN (01GQ1005A and 01GQ1005B), DFG (CRC 1286, 889), the Ministry for Science and Culture of Lower Saxony, and the Max Planck Society to F.W. D.G.-D. is a recipient of the PhD fellowship BES-2013-064171 and a grant from the short-term visit program EEBB-I-15-09552.

AUTHOR CONTRIBUTIONS

M.S. and Y.H. conceived the study; M.S., K.M., and Y.H. designed the experiments; M.S., K.M., M.K., Y.S., D.G.-D., and K.K. performed the experiments; M.S., K.M., T.I., T.T., A.S., and F.W. analyzed the data; M.S. and T.I. built the virtual reality setup; M.O. and J.N. made the G-CaMP7-T2A-DsRed2 transgene; M.S. and M.K. generated Thy1-G-CaMP7 transgenic mice; M.G.L. generated Shank2-deficient mice; T.T., M.S., H.Y., and T.F. developed the image analysis software; and M.S. and Y.H. supervised the study and wrote the paper with input from other authors.

DECLARATION OF INTERESTS

Y.H. was supported in part by Takeda Pharmaceutical and DWANGO.

Received: December 3, 2019

Revised: April 6, 2020

Accepted: June 16, 2020

Published July 7, 2020

REFERENCES

- Basu, J., Zaremba, J.D., Cheung, S.K., Hitti, F.L., Zemelman, B.V., Losonczy, A., and Siegelbaum, S.A. (2016). Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science* *351*, aaa5694.
- Berkel, S., Marshall, C.R., Weiss, B., Howe, J., Roeth, R., Moog, U., Endris, V., Roberts, W., Szatmari, P., Pinto, D., et al. (2010). Mutations in the SHANK2 synaptic scaffolding gene in autism spectrum disorder and mental retardation. *Nat. Genet.* *42*, 489–491.
- Bostock, E., Muller, R.U., and Kubie, J.L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus* *1*, 193–205.
- Bourboulou, R., Marti, G., Michon, F.X., El Feghaly, E., Nouguiere, M., Robbe, D., Koenig, J., and Epsztein, J. (2019). Dynamic control of hippocampal spatial coding resolution by local visual cues. *eLife* *8*, e44487.
- Burgess, N., and O'Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus* *6*, 749–762.
- Chen, G., King, J.A., Burgess, N., and O'Keefe, J. (2013). How vision and movement combine in the hippocampal place code. *Proc. Natl. Acad. Sci. USA* *110*, 378–383.
- Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-specific coding dynamics during spatial navigation and learning in hippocampal area CA1. *Neuron* *91*, 652–665.
- Deshmukh, S.S., and Knierim, J.J. (2011). Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Front. Behav. Neurosci.* *5*, 69.
- Dombeck, D.A., Harvey, C.D., Tian, L., Looger, L.L., and Tank, D.W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat. Neurosci.* *13*, 1433–1440.
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* *13*, 995–1002.
- Eichenbaum, H., and Cohen, N.J. (2014). Can we reconcile the declarative memory and spatial navigation views on hippocampal function? *Neuron* *83*, 764–770.

- Feng, G., Mellor, R.H., Bernstein, M., Keller-Peck, C., Nguyen, Q.T., Wallace, M., Nerbonne, J.M., Lichtman, J.W., and Sanes, J.R. (2000). Imaging neuronal subsets in transgenic mice expressing multiple spectral variants of GFP. *Neuron* 28, 41–51.
- Frank, L.M., Stanley, G.B., and Brown, E.N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *J. Neurosci.* 24, 7681–7689.
- Gauthier, J.L., and Tank, D.W. (2018). A dedicated population for reward coding in the hippocampus. *Neuron* 99, 179–193.e7.
- Geiller, T., Fattahi, M., Choi, J.S., and Royer, S. (2017). Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* 8, 14531.
- Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Brown, B.L., Koay, S.A., Taxis, J., Najafi, F., Gauthier, J.L., Zhou, P., et al. (2019). CalmAn an open source tool for scalable calcium imaging data analysis. *eLife* 8, e38173.
- Gothard, K.M., Skaggs, W.E., Moore, K.M., and McNaughton, B.L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* 16, 823–835.
- Hetherington, P.A., and Shapiro, M.L. (1997). Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behav. Neurosci.* 111, 20–34.
- Hill, A.J. (1978). First occurrence of hippocampal spatial firing in a new environment. *Exp. Neurol.* 62, 282–297.
- Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., and Moser, E.I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci.* 21, 1635–1644.
- Kentros, C., Hargreaves, E., Hawkins, R.D., Kandel, E.R., Shapiro, M., and Muller, R.V. (1998). Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. *Science* 280, 2121–2126.
- Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., and Kandel, E.R. (2004). Increased attention to spatial context increases both place field stability and spatial memory. *Neuron* 42, 283–295.
- Kohara, K., Pignatelli, M., Rivest, A.J., Jung, H.Y., Kitamura, T., Suh, J., Frank, D., Kajikawa, K., Mise, N., Obata, Y., et al. (2014). Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits. *Nat. Neurosci.* 17, 269–279.
- Komorowski, R.W., Manns, J.R., and Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* 29, 9918–9929.
- Lansink, C.S., Goltstein, P.M., Lankelma, J.V., McNaughton, B.L., and Penartz, C.M.A. (2009). Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol.* 7, e1000173.
- Le Merre, P., Esmaeili, V., Charrière, E., Galan, K., Salin, P.A., Petersen, C.C.H., and Crochet, S. (2018). Reward-based learning drives rapid sensory signals in medial prefrontal cortex and dorsal hippocampus necessary for goal-directed behavior. *Neuron* 97, 83–91.e5.
- Lee, S.H., Marchionni, I., Bezaire, M., Varga, C., Danielson, N., Lovett-Barron, M., Losonczy, A., and Soltesz, I. (2014). Parvalbumin-positive basket cells differentiate among hippocampal pyramidal cells. *Neuron* 82, 1129–1144.
- Leutgeb, J.K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C.A., McNaughton, B.L., Moser, M.B., and Moser, E.I. (2005). Progressive transformation of hippocampal neuronal representations in “morphed” environments. *Neuron* 48, 345–358.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., and O’Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature* 416, 90–94.
- Markus, E.J., Barnes, C.A., McNaughton, B.L., Gladden, V.L., and Skaggs, W.E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* 4, 410–421.
- Mehta, M.R., Barnes, C.A., and McNaughton, B.L. (1997). Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc. Natl. Acad. Sci. USA* 94, 8918–8921.
- Mizuseki, K., Diba, K., Pastalkova, E., and Buzsáki, G. (2011). Hippocampal CA1 pyramidal cells form functionally distinct sublayers. *Nat. Neurosci.* 14, 1174–1181.
- Modi, M.E., Brooks, J.M., Guilmette, E.R., Beyna, M., Graf, R., Reim, D., Schmeisser, M.J., Boeckers, T.M., O’Donnell, P., and Buhl, D.L. (2018). Hyperactivity and hypermotivation associated with increased striatal mGluR1 signaling in a Shank2 rat model of autism. *Front. Mol. Neurosci.* 11, 107.
- Mower, A.F., Kwok, S., Yu, H., Majewska, A.K., Okamoto, K., Hayashi, Y., and Sur, M. (2011). Experience-dependent regulation of CaMKII activity within single visual cortex synapses in vivo. *Proc. Natl. Acad. Sci. USA* 108, 21241–21246.
- Mukamel, E.A., Nimmerjahn, A., and Schnitzer, M.J. (2009). Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron* 63, 747–760.
- Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* 7, 1951–1968.
- O’Keefe, J., and Conway, D.H. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp. Brain Res.* 31, 573–590.
- O’Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map* (Clarendon).
- Ohkura, M., Sasaki, T., Sadakari, J., Gengyo-Ando, K., Kagawa-Nagamura, Y., Kobayashi, C., Ikegaya, Y., and Nakai, J. (2012). Genetically encoded green fluorescent Ca²⁺ indicators with improved detectability for neuronal Ca²⁺ signals. *PLoS ONE* 7, e51286.
- Pappas, A.L., Bey, A.L., Wang, X., Rossi, M., Kim, Y.H., Yan, H., Porkka, F., Duffney, L.J., Phillips, S.M., Cao, X., et al. (2017). Deficiency of Shank2 causes mania-like behavior that responds to mood stabilizers. *JCI Insight* 2, e92052.
- Pneumatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* 89, 285–299.
- Sato, M., and Stryker, M.P. (2008). Distinctive features of adult ocular dominance plasticity. *J. Neurosci.* 28, 10278–10286.
- Sato, M., Kawano, M., Ohkura, M., Gengyo-Ando, K., Nakai, J., and Hayashi, Y. (2015). Generation and imaging of transgenic mice that express G-CaMP7 under a tetracycline response element. *PLoS ONE* 10, e0125354.
- Sato, M., Kawano, M., Yanagawa, Y., and Hayashi, Y. (2016). *In vivo* two-photon imaging of striatal neuronal circuits in mice. *Neurobiol. Learn. Mem.* 135, 146–151.
- Sato, M., Kawano, M., Mizuta, K., Islam, T., Lee, M.G., and Hayashi, Y. (2017). Hippocampus-dependent goal localization by head-fixed mice in virtual reality. *eNeuro* 4, e0369–16.2017.
- Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the same neurons across multiple days in Ca²⁺ imaging data. *Cell Rep.* 21, 1102–1115.
- Smith, A.D. (2015). Spatial navigation in autism spectrum disorders: a critical review. *Front. Psychol.* 6, 31.
- Takekawa, T., Asai, H., Ohkawa, N., Nomoto, M., Okubo-Suzuki, R., Ghandour, K., Sato, M., Hayashi, Y., Inokuchi, K., and Fukui, T. (2017). Automatic sorting system for large calcium imaging data. *bioRxiv*. <https://doi.org/10.1101/215145>.
- Tsao, A., Moser, M.B., and Moser, E.I. (2013). Traces of experience in the lateral entorhinal cortex. *Curr. Biol.* 23, 399–405.
- Valero, M., Cid, E., Averkin, R.G., Aguilar, J., Sanchez-Aguilera, A., Viney, T.J., Gomez-Dominguez, D., Bellistri, E., and de la Prida, L.M. (2015). Determinants of different deep and superficial CA1 pyramidal cell dynamics during sharp-wave ripples. *Nat. Neurosci.* 18, 1281–1290.
- van der Ham, I.J.M., Martens, M.A.G., Claessen, M.H.G., and van den Berg, E. (2017). Landmark agnosia: evaluating the definition of landmark-based navigation impairment. *Arch. Clin. Neuropsychol.* 32, 472–482.

- Vogelstein, J.T., Packer, A.M., Machado, T.A., Sippy, T., Babadi, B., Yuste, R., and Paninski, L. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.* *104*, 3691–3704.
- Wang, C., Chen, X., Lee, H., Deshmukh, S.S., Yoganarasimha, D., Savelli, F., and Knierim, J.J. (2018). Egocentric coding of external items in the lateral entorhinal cortex. *Science* *362*, 945–949.
- Wiener, S.I., Paul, C.A., and Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *J. Neurosci.* *9*, 2737–2763.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* *261*, 1055–1058.
- Won, H., Lee, H.R., Gee, H.Y., Mah, W., Kim, J.I., Lee, J., Ha, S., Chung, C., Jung, E.S., Cho, Y.S., et al. (2012). Autistic-like social behaviour in Shank2-mutant mice improved by restoring NMDA receptor function. *Nature* *486*, 261–265.
- Youngstrom, I.A., and Stowbridge, B.W. (2012). Visual landmarks facilitate rodent spatial navigation in virtual reality environments. *Learn. Mem.* *19*, 84–90.
- Zaremba, J.D., Diamantopoulou, A., Danielson, N.B., Grosmark, A.D., Kaifosh, P.W., Bowler, J.C., Liao, Z., Sparks, F.T., Gogos, J.A., and Losonczy, A. (2017). Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat. Neurosci.* *20*, 1612–1623.
- Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., El Gomal, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* *16*, 264–266.

STAR★METHODS

KEY RESOURCES TABLE

| REAGENT or RESOURCE | SOURCE | IDENTIFIER |
|---|--|---|
| Antibodies | | |
| Rabbit anti-calbindin D-28K | Millipore | Cat#AB1778; RRID: AB_2068336 |
| Rabbit anti-GAD65/67 | Millipore | Cat#AB1511; RRID: AB_90715 |
| Mouse anti-parvalbumin (clone PARV-19) | Sigma | Cat#P3088; RRID: AB_477329 |
| Mouse anti-somatostatin | GeneTex | Cat#GTX71935; RRID: AB_383280 |
| Rabbit anti-glial fibrillary acidic protein | Dako | Cat#N1506; RRID: AB_10013482 |
| Rabbit anti-Iba1 | Wako | Cat#019-19741; RRID: AB_839504 |
| Alexa Fluor 647-labeled goat anti-rabbit IgG | Thermo Fisher | Cat#A-21245; RRID: AB_2535813 |
| Alexa Fluor 647-labeled goat anti-mouse IgG | Thermo Fisher | Cat#A-21236; RRID: AB_2535805 |
| Experimental Models: Organisms/Strains | | |
| Mouse: Thy1-G-CaMP: C57BL/6J-Tg(Thy1-G-CaMP7, -DsRed2)492Bsi | This paper | RIKEN BRC: RBRC06579 |
| Mouse: Shank2-deficient mice: B6N.129S4-Shank2 ^{tm1Mglc/CsbdJ} | Won et al., 2012 | JAX: 033667; RRID: IMSR_JAX033667 |
| Oligonucleotides | | |
| Genotyping primer for Thy1-G-CaMP7 mice: CTGCTGCCCGACAACCA | Sato et al., 2015 | N/A |
| Genotyping primer for Thy1-G-CaMP7 mice: GTCGTCCTTGAAGAAGATGG | Sato et al., 2015 | N/A |
| Genotyping primer for Shank2-deficient mice: Shank2 WT fwd: GCTAGCATGACGTGTGTTGTG | Won et al., 2012 | N/A |
| Genotyping primer for Shank2-deficient mice: Shank2 KO fwd: CCGACTGCATCTGCGTGTTTC | Won et al., 2012 | N/A |
| Genotyping primer for Shank2-deficient mice: Shank2 rev: ACCTGTGTGTGATTCTGAC | Won et al., 2012 | N/A |
| Recombinant DNA | | |
| Mouse Thy1.2 promoter | Feng et al., 2000 | RRID: Addgene_20736 |
| G-CaMP7-T2A-DsRed2 | Ohkura et al., 2012 Sato et al., 2015 | N/A |
| Software and Algorithms | | |
| ImageJ | NIH | RRID: SCR_003070 |
| TurboReg | Biomedical Imaging Group, Swiss Federal Institute of Technology Lausanne | RRID: SCR_014308 |
| MATLAB | Mathworks | RRID: SCR_001622 |
| LabVIEW | National Instruments | RRID: SCR_014325 |
| OmegaSpace | Solidray | http://www.solidray.co.jp/product/omega/index.html |
| Prism | GraphPad | RRID: SCR_002798 |
| Modified nonnegative matrix factorization | Takekawa et al., 2017 | https://www.biorxiv.org/content/10.1101/215145v1 |
| Constrained nonnegative matrix factorization | Pneumatikakis et al., 2016 Giovannucci et al., 2019 | https://github.com/flatironinstitute/CalMAn |
| Cell registration | Sheintuch et al., 2017 | https://github.com/zivlab/CellReg |

RESOURCE AVAILABILITY

Lead Contact

Further information and requests for resources and reagents should be directed to and will be fulfilled by the lead contact, Masaki Sato (masaaki.sato@riken.jp).

Materials Availability

The mouse line generated in this study is available from the RIKEN BioResource Center (<http://www.brc.riken.jp/lab/animal/en/>; stock number RBRC06579).

Data and Code Availability

The datasets and custom computer code associated with this study will be made available by the authors upon request.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All experiments were conducted in accordance with institutional guidelines and protocols approved by the RIKEN and Kyoto University Animal Experiments Committees. Adult male Thy1-G-CaMP7 transgenic mice and those lacking Shank2 were used for the experiments. Homozygous Shank2-deficient mice (Won et al., 2012) carrying the Thy1-G-CaMP7 transgene and age- and gender-matched wild-type control transgenic mice were obtained by crossing male homozygous Thy1-G-CaMP7 transgenic, heterozygous Shank2-deficient mice with female heterozygous Shank2-deficient mice lacking the Thy1-G-CaMP7 transgene. All mice used were at least 12 weeks old and weighed 28–35 g at the beginning of surgery. The mice were housed in groups of one to four per cage on a 12 h–12 h light–dark cycle (with lights on at 6 pm and off at 6 am the next day). Thy1-G-CaMP7 mice were genotyped by PCR using the primers 5'-CTGCTGCCCGACAACCA-3' and 5'-GTCGTCCTTGAAGAAGATGG-3', which provided a 465-bp product of the G-CaMP7 coding sequence from tail DNA samples of transgene-positive mice. The wild-type and Shank2-deficient alleles were detected by PCR using the primers WT fwd 5'-GCTAGCATGACGTGTGTTGTG-3' and rev 5'-ACCTGTGTGTGATTCTGAC-3', and the primers KO fwd 5'-CCGACTGCATCTGCGTGTTC-3' and rev 5'-ACCTGTGTGTGATTCTGAC-3', respectively (Won et al., 2012).

METHOD DETAILS

Generation of Thy1-G-CaMP7 transgenic mice

The cDNA encoding G-CaMP7 (Ohkura et al., 2012) ligated to the coding sequence of DsRed2 via a *Thosea asigna* virus-derived 2A peptide (T2A) sequence (Sato et al., 2015) was subcloned into the Xho I site of the modified mouse Thy1.2 promoter vector (Feng et al., 2000). The 8.7-kb DNA fragment was prepared by digestion with Not I and Pvu I restriction enzymes and subsequent gel purification and injected into the pronuclei of 466 fertilized eggs of C57BL/6J mice. From 32 offspring, 9 mice were identified as transgene positive, and 6 exhibited transgene expression in the brain. One founder mouse that expressed the transgene at a high level in the hippocampus was used for this study.

Analysis of transgene expression

Analysis of transgene expression in Thy1-G-CaMP7 mice was conducted essentially as described previously (Sato et al., 2015). The primary antibodies used in immunolabeling were rabbit anti-calbindin D-28K (1:500, AB1778, Millipore, Billerica, MA), rabbit anti-GAD65/67 (1:500, AB1511, Millipore), mouse anti-parvalbumin (1:1000, clone PARV-19, P3088, Sigma, St. Louis, MO), mouse anti-somatostatin (1:200, clone SOM-018, GTX71935, Gene Tex, Irvine, CA), rabbit anti-glial fibrillary acidic protein (GFAP) (1:1000, N1506, Dako, Denmark) and rabbit anti-Iba1 (1:1000, 019-19741, Wako Pure Chemical Industries, Ltd., Japan), and the secondary antibodies used were Alexa Fluor 647-labeled goat anti-rabbit and anti-mouse IgG antibodies (1:700–1000, A-21245 and A-21236, Thermo Fisher Scientific, Waltham, MA). Nuclear counterstaining was conducted in PBS containing 10 μ g/ml Hoechst 33258 (Calbiochem) and 0.1% Triton X-100 at room temperature for 5 min. The densities of CA1 pyramidal cells were determined in the pyramidal cell layer of the dorsal CA1 hippocampus in 4 fields per animal (field size, 212 \times 212 μ m) and expressed as cell number per 200 μ m length of the pyramidal cell layer. A quantitative analysis of G-CaMP7 and calbindin immunolabeling fluorescence signals was performed by averaging fluorescence intensities across cell bodies and normalizing them to those of the brightest cells within the field of view. High-calbindin cells were defined as those whose normalized calbindin signals were greater than or equal to 0.2, and the remaining cells were grouped as low-calbindin cells.

The histological analyses confirmed that the long-term transgenic expression of G-CaMP7 in Thy1-G-CaMP7 mice does not cause discernible toxicity. The pyramidal cell densities and overall distribution patterns of glial fibrillary acidic protein (GFAP)-positive astrocytes and Iba1-positive microglia were indistinguishable between Thy1-G-CaMP7 and wild-type (WT) C57BL/6 mice (Figure S1C; pyramidal cell density, 51.3 \pm 2.4 cells/200 μ m pyramidal cell layer in Thy1-G-CaMP7 mice versus 48.7 \pm 1.2 cells/200 μ m pyramidal cell layer in WT mice; $t_{(14)} = 0.975$, $p = 0.35$, $n = 8$ fields each from 2 Thy1-G-CaMP7 and 2 WT mice, unpaired two-tailed t test). It was also found that in addition to strong hippocampal expression (Figure S1D), Thy1-G-CaMP7 mice express G-CaMP7 in diverse brain areas, including the cerebral cortex, olfactory bulb, brainstem and cerebellum (Figures S1E–S1P).

Surgery

Mice were anesthetized with isoflurane in ambient air (3% induction, 1.5% maintenance) and placed in a custom-made stereotaxic frame. To reduce secretions and brain edema, we administered atropine (0.3 mg/kg, s.c.) and dexamethasone (2 mg/kg, s.c.) prior to anesthesia. A circular piece of the scalp was removed, and the underlying bone was cleaned and dried. Three small screws were then placed in the skull (two at the suture of the interparietal and occipital bones and one on the right frontal bone) to provide anchors for the head plate. A thin layer of cyanoacrylate was applied to provide a substrate to which the dental acrylic could adhere.

A stainless steel head plate (25 mm length, 4 mm width, 1 mm thickness) with a wide circular opening (7 mm inner diameter and 10 mm outer diameter, the center is 2.5 mm off relative to the middle of the long side of the plate) was affixed to the skull using dental cement. The center of the opening was targeted at 2 mm posterior to the bregma and 2 mm lateral to the midline in the left hemisphere. The cement was mixed with black ink to block light entry from the LCD monitor into the microscope and placed onto the skull such that it covered the entire skull, including the anchor screws, except for the area of skull inside the opening of the head plate.

Optical window preparation was performed as described previously with modifications (Dombeck et al., 2010; Sato et al., 2016). A few days after the head plate surgery, a 2.5-mm-diameter circular craniotomy was created on the skull overlying the dorsal hippocampus. The dura was removed with forceps, and the overlying cortex was aspirated in a small amount at a time using a blunted 25-gauge needle connected to a vacuum pump. This step was continued with occasional irrigation with cortex buffer (123 mM NaCl, 5 mM KCl, 10 mM glucose, 2 mM CaCl₂, 2 mM MgCl₂, 10 mM HEPES, pH 7.4) until the white matter, including the corpus callosum, was exposed. Then, the top-most layers of the white matter were gently peeled aside by holding with the vacuum-connected blunted needle such that its minimal thickness remained covering the dorsal surface of the hippocampus. To minimize bleeding, aspiration was initiated from a cortical area devoid of large vessels, and bleeding was treated immediately with a piece of gelatin sponge (Spongel, Astellas Pharma, Tokyo, Japan) wetted with cortex buffer. An imaging window was then inserted to mechanically support the cranial hole, its surrounding tissue and the hippocampal surface. The imaging window consisted of a stainless steel ring (2.5 mm outer diameter, 2.2 mm inner diameter and 1.0 mm height) with a round coverslip (2.5 mm diameter, 0.17 mm thickness, Matsunami Glass Ind., Osaka, Japan) attached to the bottom using a UV-curable adhesive (NOA81, Norland Products, Cranbury, NJ). To reduce brain movement during imaging, a small disk of medical grade clear silicone sheeting (0.13 mm thickness, 20-10685, Invotec International, Jacksonville, FL) was attached to the surface of the coverslip facing the hippocampal tissue (Mower et al., 2011). When the window was positioned, the bottom coverslip was approximately parallel relative to the head plate, and the hippocampal surface was clearly visible through the bottom coverslip without any trace of bleeding. The upper rim was then cemented to the skull with dental acrylic.

After surgery, a metal cover (0.3 mm thickness) was screwed onto the upper surface of the head plate to protect the imaging window from dust. The mice were placed in a warmed chamber until they fully recovered from anesthesia and were then returned to their home cages. They were housed for at least 4 weeks of postoperative recovery before the start of handling.

VR set-up

A VR system with an air-supported spherical treadmill for head-fixed mice was constructed as described previously (Sato et al., 2017). A 20-cm-diameter Styrofoam ball placed inside the bowl provided a freely rotating surface on which the mouse stood. The mouse was positioned near the top of the ball with its head fixed via the steel head plate that was screwed into a rigid cross bar and posts. A single wide-screen 23" LCD display (Dell U2312, Round Rock, TX) placed 30 cm in front of the mice presented VR scenes rendered by OmegaSpace 3.1 (Solidray Co. Ltd., Yokohama, Japan) running on a Windows 7 computer in 81° horizontal and 51° vertical fields of view. The LCD monitor was large enough to cover the major part of the mouse's binocular and monocular visual fields (Sato and Stryker, 2008). The use of a single LCD monitor for VR presentation effectively elicits visual cue-based virtual navigation behavior in head-fixed mice (Youngstrom and Stowbridge, 2012; Sato et al., 2017).

The movement of the ball was measured with a USB optical computer mouse (G400, Logitech, Newark, CA) via custom driver and LabVIEW software (National Instruments, Austin, TX). The optical mouse was positioned in front of the mouse and at the intersection of the mouse's sagittal plane and the equator of the ball. The signals along the horizontal axis (aligned parallel to the mouse's sagittal plane) generated by the running of the head-fixed mouse was used to compute rotational velocity in the forward-backward direction. This velocity signal was converted into analog control voltages (0–5 V) via a D/A converter and fed to a USB joystick controller (BU0836X, Leo Bodner, Northamptonshire, UK) connected to the OmegaSpace computer to move the mouse's position in VR.

Water rewards (5 μl/reward) were delivered by a microdispenser unit (O'Hara & Co., Ltd., Tokyo, Japan) attached to a water-feeding tube positioned directly in front of the mouse's mouth. The unit was triggered upon reward events in VR by 5 V TTL signals generated by an OmegaSpace script via a USB-connected D/A device (USB-6009, National Instruments). The behavioral parameters, such as the mouse's location in the virtual environment, the trigger signals for water rewards and the rotational velocity signals of the spherical treadmill, were recorded at 20-ms intervals using custom software in LabVIEW. The TTL signals for each frame sent by the microscope computer were recorded with the behavioral data to synchronize the imaging and behavioral data.

Behavior

At least 5 days before the start of imaging experiments, mice implanted with the head plate and the imaging window were acclimated to handling and the Styrofoam ball. During this acclimation, mice were handled by an experimenter for 5–10 min and then allowed to move freely on the top of the ball, which was rotated manually by the experimenter, for another 5–10 min. The procedure was

performed once a day and repeated for at least 3 days. The mice were then placed on a water restriction schedule 2–3 days before the start of the experiments. No other type of pre-training was conducted before the mice were exposed to the virtual environment in the first sessions. Body weight and general appearance were checked daily to ensure that the animals maintained at least ~85% of their preoperative body weight and exhibited no signs of abnormal behavior throughout the study. The experiments were performed during the dark phase of the cycle to enhance the locomotion of the mice.

The virtual endless linear track was created using an editor function in OmegaSpace. The mouse started at the origin of the virtual linear track segment and ran along the track unidirectionally with visual feedback rendered by OmegaSpace. The track segment was 100 cm long, measured as the number of rotations of the ball required to move from one end of the track to the other times the circumference of the ball. The mouse moved only one-dimensionally along the midline of the track, with its view angle fixed in the direction of movement. Different patterns were placed on the walls of each track subsegment as follows: vertical white and black stripes at 0–25 cm, horizontal white and black stripes at 25–50 cm, and black dots on a white background at 50–100 cm. The floor was patterned with a white grid on a black background. The space above the track was colored black. A green gate was placed 25 cm from the origin to serve as a salient landmark. Water rewards were delivered when the mouse reached a reward point located 75 cm from the origin. This reward point was located in the middle of a track zone with a certain wall pattern (i.e., black dots on a white background) and not denoted with any other salient visual cues. Upon reaching the other end of the segment, the mouse's virtual position was transferred back to the origin, and the same segment of the linear track was presented again. The approaching track segment following the current segment was always rendered on the monitor, such that the virtual linear track appeared infinitely long.

The mice underwent a total of 15 training sessions in the above task, with 1–2 sessions per day. Each session was 10 min long. When 2 sessions were performed in one day, the intervals between sessions were at least 4 h, and the mice were returned to their home cages between the sessions. The entire training period from the first to the last session was 225 ± 8 h (mean \pm SD, $n = 7$ mice). Each mouse was lightly anesthetized with isoflurane to detach the metal window cover screwed onto the head plate and clean the imaging window, after which the mouse was placed into the VR apparatus. The head was then fixed to the crossbar above the ball via the head plate, and the mouse was left on the ball in the dark for approximately 20 min until it had recovered fully from the anesthesia. During the behavioral session, the animal was allowed to behave freely in the head-fixed arrangement. G-CaMP7 fluorescence in hippocampal CA1 pyramidal neurons was simultaneously imaged as described below. Licking was detected using an infrared photo beam sensor (OPR-LKR, O'Hara & Co., Ltd.) in a subset of experiments. Slowdown of running speed before the reward delivery point was calculated as the difference between the maximum average running speed between the landmark location and the reward delivery point (32.5–65 cm from the origin) and the average running speed in the area immediately before the reward delivery point (65–75 cm from the origin).

For the reward-rearrangement task, mice first underwent 15 training sessions on the virtual linear track as described above. The mice were further trained for the following 5 sessions (Rearrangement 1–5) in the same virtual linear track except that the location of reward delivery (75 cm from the origin) was shifted to match the location of the landmark (25 cm from the origin). Data obtained from the last 4 sessions of the initial 15 training sessions before the shift (Sessions 12 through 15, also referred to as Pre –4 through –1) were analyzed as pre-rearrangement baseline sessions. The first rearrangement sessions were performed immediately after the last baseline sessions without releasing the mice from head fixation.

For the task with 50% reward delivery and the task with no gate, a separate cohort of mice ($n = 3$) were trained in 15 normal sessions and were then subjected to 2 sessions of the 50% reward task, in which the reward was delivered at the same delivery point with 50% probability, followed by 2 normal sessions and 1 no-gate session, where the green gate was removed from its location throughout the entire sessions. The first 50% reward sessions and the no-gate sessions were performed immediately after the preceding normal sessions without releasing the mice from head fixation.

For the random gate shift task, mice ($n = 3$ wild-type mice and 3 Shank2-deficient mice) were first trained on the virtual linear track in which the visual landmark indicated the reward location at 62.5 cm from the origin for 15 sessions. Then, each of the following 5 random shift sessions began with 12 baseline trials in which the position of the landmark plus reward was fixed at the pre-trained location, followed by random shift trials in which their position was shifted forward or backward by 12.5 cm or remained at the pre-trained (normal) position on a trial-by-trial basis in a random manner. Each of the three landmark plus reward position was tested four times for a total of 12 random shift trials. Slowdown and licking immediately (0–10 cm) before the landmark plus reward location were analyzed separately for baseline, forward, normal and backward trials in each session and averaged across the last three of the five sessions for each mouse.

Imaging

Imaging was performed using a Nikon A1MP (Nikon, Tokyo, Japan) equipped with a 16x, NA 0.8 water immersion objective lens. The microscope was controlled with Nikon NIS-elements software. G-CaMP7 and DsRed2 were excited using a Ti-sapphire laser (MaiTai DeepSee eHP, Spectra-Physics, Santa Clara, CA) at 910 nm. Typical laser power was approximately 40 mW at the objective lens. G-CaMP7 fluorescence was separated using a 560-nm dichroic mirror and collected with an external GaAsP photomultiplier tube (10770PB-40, Hamamatsu Photonics, Hamamatsu, Japan) mounted immediately above the objective lens. The calcium-insensitive DsRed2 fluorescence, which helped to identify G-CaMP7-labeled pyramidal neurons, was simultaneously imaged and recorded using another GaAsP photomultiplier tube. The DsRed2 images were checked by the experimenter for the on-site assessment of the quality of image acquisition but not used for offline quantitative image analysis, except for image alignment across sessions (Figure S5).

To image G-CaMP7-labeled CA1 pyramidal neurons, the microscope was focused at a depth of approximately 150 μm from the hippocampal surface. To prevent the entry of light from the LCD monitor into the microscope, a small sheet of aluminum foil was wrapped around the objective lens, so the foil completely covered the space between the objective and the skull. Images of 512 \times 512 pixels were acquired at a rate of 15 frames per second using a resonant-galvo scanner mounted on the microscope. Each imaging session was 10 min long. The size of the field of view was 532 \times 532 μm . In repeated chronic imaging, previously imaged cell populations usually re-appeared at similar depths in new sessions. We took reference images of DsRed2 fluorescence at the beginning of each session to confirm that the reference image of the current session was very similar to that of the previous session by ensuring that blood vessels and neurons arranged in unique patterns appeared in the same parts of the two images.

QUANTIFICATION AND STATISTICAL ANALYSIS

Image analysis

Each frame of a G-CaMP7 time-lapse movie was aligned to an average fluorescence image of the movie for motion correction using the TurboReg ImageJ plug-in. The registered movie was then denoised by a spatio-temporal median filter. This preprocessed movie $f(t, x)$ was reconstituted to the sum of fluorescence intensity of individual cells using a modified non-negative matrix factorization algorithm, as described in detail elsewhere (Vogelstein et al., 2010; Pnevmatikakis et al., 2016; Takekawa et al., 2017; Giovannucci et al., 2019). Briefly, this algorithm assumes that the fluorescence intensity of each cell can be deconvoluted to the spatial filter $a_c(x)$, which represents the position and shape of the cell, and the time variation $v_c(t)$ derived from spiking activities $u_c(t)$:

$$f(t, x) \sim N\left(a_0(x) + v_0(t) + \sum_c a_c(x)v_c(t), \sigma^2\right)$$

where a_0 , v_0 are baselines, and σ^2 is intensity of Gaussian noise. As is the case in cell identification using independent component analysis (Mukamel et al., 2009), this algorithm preferentially detects cells that change their fluorescence intensities over time ("active cells") because cells that barely do so are regarded as being near baseline. Each spike derives the transient elevation of fluorescence intensity with a double-exponential shape:

$$v_c(t) = \sum_{\tau=1}^t \left(\exp\left(-\frac{t-t'}{\tau_1}\right) - \exp\left(-\frac{t-t'}{\tau_2}\right) \right) u_c(t')$$

The exponential rise and decay time constants $\tau_1 = 0.09$ and $\tau_2 = 0.261$, respectively, were obtained by curve fitting of actual traces of cellular calcium transients in G-CaMP7-expressing CA1 pyramidal neurons in Thy1-G-CaMP7 mice *in vivo*. Spatial filters and spike timings were estimated by two iterative steps. In the first step, we prepared tentative spatial filters and estimated spike trains corresponding to respective filters by a least-squares approach with a non-negative restraint condition. Subsequently, spatial filters were estimated using the least-squares method on the condition that the estimated spike trains were feasible. In addition, we introduced L1 sparse regularization derived from priors that represented the typical cell size and spike frequency. To determine the mutual relationship between a and u , a regularized term was also introduced to the model. This condition guaranteed the uniqueness of the scale of a , u and v . As a consequence, a , u and v are presented in arbitrary units, while the product of a and v corresponds to the observed data.

In practice, 512 \times 512 pixel image data were divided into 4 \times 4 of 128 \times 128 pixel subareas with 32-pixel overlap regions. Each subarea was analyzed with the above algorithm, and the results were combined to cover the whole image area. After the initial calculation, the morphology of each spatial filter was defined as the region above 0.2 times its peak value, and the position of the filter was defined by its weighted centroid. We then removed the following filters as those that did not represent complete single pyramidal cell morphology: (1) filters whose areas were smaller than 25 pixels, (2) filters whose areas were larger than 400 pixels, (3) filters located on the edge of the image, (4) filters whose heights or widths were greater than 64 pixels because they often contained structures of multiple cells, and (5) smaller filters in filter pairs whose distances were closer than 10 pixels (10.4 μm) and whose temporal correlation coefficients of activities were greater than 0.3 because they were considered to be derived from the same cell.

After those non-cell filters were removed, we recalculated the activity time series for the new filter set. Visual inspection confirmed that nearly all active cells that were represented in a background-subtracted maximum-intensity projection image were identified with this procedure (Figures S3A and S3B). All images of the entire session, regardless of the mouse's behavioral state, were used for this image analysis. The average number of cells identified from a movie of a session was 900 \pm 246 (mean \pm SD, $n = 105$ sessions from 7 mice).

Analysis of place fields

Place fields were calculated using cellular activity during movement periods. We defined these periods as the time when the mouse moved at a speed of > 0.5 cm/s continuously for a duration of > 2 s to reject irrelevant movements, such as grooming and jittering on the ball. We divided the entire virtual linear track segment into 80 bins (bin size = 1.25 cm) and created a histogram of neuronal activity versus track position for each cell. The activity events were defined by binarizing the time series of inferred spike activity u at a threshold of 0.1, which was empirically determined to remove baseline noise. The counts on the histogram were then divided by

the mouse's occupancy time in each bin, and the resultant place fields were smoothed with a Gaussian function (Gaussian window size = 6.25 cm) and normalized to the maximum values. To test the significance of place-related activity, we calculated the mutual information content between neuronal activity and the mouse's virtual location for each cell (Markus et al., 1994; Ziv et al., 2013). We compared this value to a distribution of mutual information content calculated using 1000 randomly permuted data for the same cell. The permutation was conducted by rotating the activity event time series by a random amount relative to the time series of the mouse's virtual positions. Cells were considered place cells (PCs) if their overall activity rates within the session were no less than 0.1 events/s and their mutual information content in the real data was greater than the 95th percentile of the values obtained from the randomly permuted data. We defined the position of the place field of each PC by the position of the peak of the field. A PC was considered a "gate (GT) cell," "reward (RW) cell" or "wall (WL) cell" if the position of its place field was 17.5–32.5, 75–95 or 47.5–55 cm from the origin of the track segment. PCs with place fields outside the above zones were categorized as "non-reward, non-gate cells (non-RW/GT cells)." The PC formation factor was defined by the slope of a least-squares regression line fitted to a plot of the fraction of PCs against the fraction of time spent running, which contained data points from all tested animals in the session of interest. The linear regression model included no constant term under the assumption that no PCs were formed without running in any given session. When we calculated the fractions of GT, RW and non-RW/GT cells relative to the number of total PCs, we used data from sessions with at least 35 total PCs to avoid the effects of improperly large or small fractions caused by small numbers of cells.

Analysis of cellular activity in the 50% reward task was conducted by dividing the data from the full sessions into two subsets, namely, rewarded trials and non-rewarded trials, and analyzing them separately. The activity bias index between rewarded trials and non-rewarded trials was then calculated as $(R-N)/(R+N)$, where R and N represent the cell's in-field activity in rewarded trials and non-rewarded trials, respectively. The activity bias index ranges from -1 to 1 , where a positive value indicates a bias toward rewarded trials and a negative value indicates a bias toward non-rewarded trials. Cells that had activity bias indices < -0.1 (equivalent to an approximately 1.22-fold increase in non-rewarded trials), $-0.1 \leq$ activity bias indices < 0.1 (equivalent to an approximately 1.22-fold increase in rewarded trials), and activity bias indices ≥ 0.1 were subsequently categorized as no-reward-preferring RW cells, constant RW cells, and reward-preferring RW cells, respectively. In the analysis of cellular activity in the no-gate task, the same cells that were identified both in the preceding control (pre) and no-gate sessions were determined as described below ("Alignment of cells across sessions"), and the calculation of the activity bias index between pre and no-gate sessions and the categorization of gate-preferring GT cells, no-gate-preferring GT cells, and constant GT cells was conducted as in the 50% reward task.

Alignment of cells across sessions

To find a consistent population of cells in images that were acquired in two different sessions, we first estimated the extent of overall image displacement that existed between the two image datasets. We searched for a peak in the two-dimensional correlation coefficient calculated between the two DsRed2 reference images obtained at the beginning of each session within a range of 25×25 pixels ($26.0 \times 26.0 \mu\text{m}$) of displacement in the x and y dimensions (Figure S5A). In this analysis, we conservatively focused on comparisons between two consecutive sessions (i.e., sessions 1 and 2, 2 and 3, etc.) because the quality of the image alignment was gradually reduced as the number of sessions that separated the two images increased ($p < 0.0001$, $F_{(13,721)} = 10.94$, one-way ANOVA, Figure S5F). A preliminary assessment of the first cohort confirmed that every compared image pair displayed a peak within this range (average displacement in the x dimension, $5.9 \pm 4.5 \mu\text{m}$; average displacement in the y dimension, $4.7 \pm 4.2 \mu\text{m}$; average peak correlation coefficient, 0.77 ± 0.09 ; mean \pm SD; $n = 98$ image pairs). During the calculation of two-dimensional correlation coefficients, the image of one session (the "source" session) was systematically shifted relative to that of the other session (the "target" session). The map of the coordinates of all cell positions in the target session was then overlaid with that of the source session, shifted by the amount of the estimated displacement (Figure S5A). The cell closest to each cell in the target session was searched in the displaced source session map, and the cell that was found was regarded provisionally to be the same cell if they were separated by 5 pixels ($5.2 \mu\text{m}$) or less. Cells that were unable to find the closest cells within this range were rejected from the subsequent analysis. After finding the provisional counterparts in the displaced source session map, the same procedure was repeated for the cells in the displaced source session map to conversely find their closest partners in the target session. This step helped remove cell pairs that were redundantly assigned (e.g., two different cells in one session falsely assigned to the same single cell in the other session) and the resultant cell pairs that had mutually unique correspondence were considered to be the pairs that represented the same cells (termed hereafter "common cells"). When comparing PC maps, common PCs were defined as a subset of common cells that were identified as significant PCs in both consecutive sessions. Stable PCs were defined as a subset of common PCs with place-field positions in the consecutive sessions that were close to each other (i.e., place field distance < 10 cm). In the analysis of formation, recruitment and stabilization of PCs (Figure 5), we first identified a population of common cells that belonged to the cell category of interest in the reference session N and then tracked the position of the place field of each cell in the subsequent session $N+1$. The transition was quantified by determining the density of cells at each bin in session $N+1$.

Since the above-described pairwise cell matching can identify only the correspondence between active cells in adjacent sessions, a subset of data ($n = 2$ mice, 15 training sessions from each) was further analyzed including inactive ("silent") cells as follows. Spatial filters and spike timings of all sessions were estimated from G-CaMP7 time-lapse movies with the constrained nonnegative matrix factorization (CNMF) implementation developed by Pnevmatikakis et al. (2016) and Giovannucci et al. (2019). Common spatial filters were detected using a model that finds the most probable counterpart by jointly taking into account their centroid distances and correlations between spatial filters (Sheintuch et al., 2017). This procedure found common spatial filters in 77% of filters examined

($n = 32,816$) in two adjacent sessions. This search was then extensively conducted in every pair within a series of sessions, and clusters of common spatial filters detected in at least two different sessions were selected as candidate cells. These cells were considered “silent” in the sessions in which no corresponding spatial filters were identified. We confirmed that none of these missing regions contained individual cell-like activity by calculating the probability of obtaining the maximum amplitude deviation from baseline activity under given baseline fluctuations. Map comparisons including these silent cells were conducted as described above except that they additionally took into account the transitions from and to silent cells.

Statistics

Statistical details including sample sizes can be found in the figure legends and Results. When only two groups were compared, two-sided Student's *t* tests were used if the variances of the two groups were similar. Otherwise, two-tailed Mann-Whitney tests were used. When more than two groups were compared, analysis of variance (ANOVA) was used if variances of the groups compared were similar. Otherwise, a non-parametric version of ANOVA (Friedman test) was used. In both parametric and non-parametric ANOVA, *p* values were adjusted for post hoc multiple comparisons. Exact *p* values are shown unless $p < 0.0001$. Statistical tests were performed using GraphPad Prism versions 6 and 7 (GraphPad Software, Inc., La Jolla, CA).

Chapter Five

Turnover dynamics of the spatial map under baseline conditions

5.1 Content

Dynamics of the hippocampal place map have been subject to a long line of studies. A multitude of works has aimed to characterize the map's response to changes in the environment, such as room shapes^{57,66,165,166}, colours^{69,70}, odors⁶⁹ or the arrangement of visual cues^{70,71} but also to the internal behavioral state of the animal^{28,158,159,167,168}. Characterizing the response to changes in the environment depends on the assumption that the time in between recording sessions has no impact on the hippocampal map, supposedly based on anecdotal evidence of long-term stability of place fields^{14,85}. Experiments in the hippocampal region CA2, however, have shown that the region has strong temporal dynamics and even found that activity in two distinct environments is more correlated than in recordings within the same environment spaced only few hours apart⁸⁶. Even though CA1 has been found to display a higher degree of temporal stability than CA2⁸⁶, understanding the temporal dynamics is an integral step to understanding the mechanisms of turnover and responses to external changes.

Throughout the last decade, experiments using large scale imaging have allowed for recordings of hundreds to thousands of neurons from the hippocampus over several days to investigate baseline dynamics observed in non-changing environments^{15–18}. While they describe the observed turnover dynamics on a population level, the underlying mechanisms shaping the dynamics have not been characterized, motivating the content of this chapter and its guiding question: *What happens, when nothing happens?*

Within this chapter, I will present an in-depth analysis of hippocampal turnover dynamics. I will first discuss dynamics happening on a short-term scale within the 10 minutes of a single recording session, Sec. 5.2. Commonly observed turnover during exploration of a familiar environment is introduced in Sec. 5.3 and the set of data included in the analysis of such baseline behavior is further defined in Sec. 5.3.1. Overall transition probabilities between the different states that a neuron can take on within the hippocampal place coding population are described

in Sec. 5.4.1. I find, how the observed turnover can be further broken down into (1) a drift in the population of active neurons, (2) a drift in the population of neurons with identified place fields and (3) movements of place fields. These different levels of dynamics are dissected in Sec. 5.4. While previous sections describe dynamics on the basis of the number of experienced sessions, a number of other time scales along which hippocampal dynamics happen are feasible, which is further discussed in Sec. 5.5.

With the unique dataset available to me, I am able to characterize dynamics individually across different animals, allowing to identify individual variability and common features. Consistently across mice, I find that hippocampal turnover is dominated by a slow, gradual drift in the population of active cells, guided by a neuron-wise short-term increase of activation probability after an active session. On average, active neurons assume significant place fields randomly, but in the case of recurring place fields code for the same location with significantly increased probability over a timescale of days to weeks. This probability is increased by continuous coding, but interestingly also, by intermittent inactivity. A process similar to the gradual turnover across sessions can even be observed within single sessions, where the population of cells showing location-dependent activity becomes increasingly dissimilar over time.

5.2 Short-term dynamics of place representation

Place fields of neurons generally do not show perfect within-session reliability, Fig. 3.13(b)^{18,157}, such that individual neuron activity can be different in between trials, 5.1(a)). Using the trial-based identification of place field activity, I answer a few question about short-term dynamics of the place representation over the course of a single session.

I first asked, whether the place map of the environment obtained from single trial activity covers the whole space, similar to what has been reported for activity data from complete sessions. For this, I calculate the place map on a trial by trial basis, including only place fields active within a given trial for the corresponding firing map, Fig. 5.1(b). Generally, place maps of trials well reproduce the session-averaged place map, with minor deviations only, Fig. 5.1(c).

The number of place fields per session is defined by the total number of neurons which show a place field at some point within the session. Barely any place field is active in all trials, however, so I asked, whether there is some change in the place coding population over experienced trials. For this, I allowed to pool place fields from a range $\pm\Delta t$ around the currently examined trial, ranging from $\Delta t = 0$ (current trial only, purple) up to 3 (7 trials centered around the current one, red). Only $\approx 20 - 30\%$ of identified place fields are active in each trial, with a slight, but insignificantly decreasing trend towards the end of a session, Fig. 5.1(d-f) (slope of linear fit indiscernible from 0).

I then analyzed the overlap of the population of neurons with a place field in two distinct trials, separated by $\Delta trials$. In the case of comparing single trials, only ($\Delta t = 0$), two trials appear

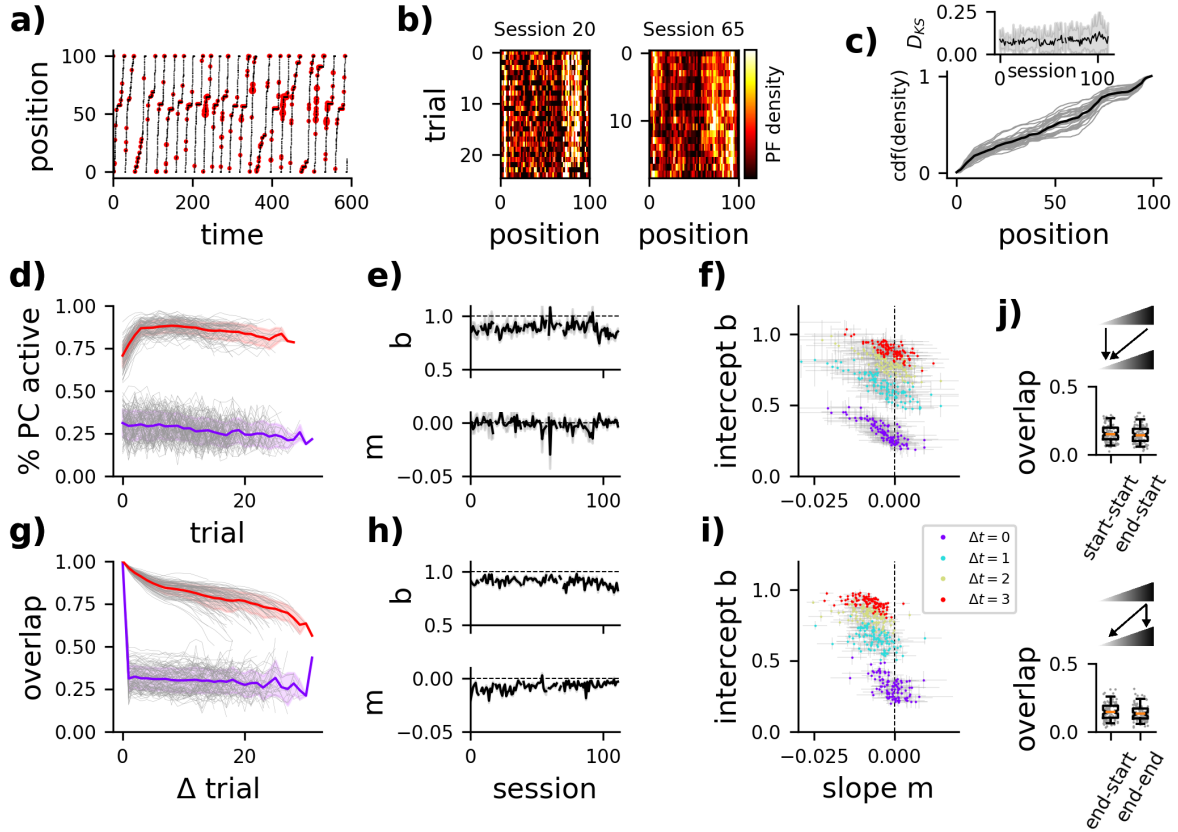


Figure 5.1: The population of place coding cells changes over the course of a single session; **(a)** example activity of a place cell coding for the central location in few trials, only; position of the animal marked by black dots, superimposed with red dots when the cell is active; size scales with Ca^{2+} -event amplitude; trials are defined as the time between entering at bin 0 and reaching bin 100 (followed by a teleport back); **(b)** two example sessions with single trial place field density along rows; **(c)** cumulative density of place fields of single trials (grey) and session-averaged density (black) from an example session; inset: maximum Kolmogorov-Smirnov-distance of single trials towards session-average for all sessions with 1SD-confidence interval over all trials (shaded region); **(d)** fraction of place fields active at adjacency-pooled trials for all sessions (grey lines) and averaged (coloured lines with shaded 1SD confidence interval) for different trial-pooling widths Δt ($=0$, single trial only, purple; $=3$, 7 trials centered around current indicated one, red); 100% are all place cells identified within this session; **(e)** parameters and confidence interval of linear least squares fit $\%PC = mx + b$ for $\Delta t = 3$ over all sessions; fit was performed excluding the first and last Δt trials, to avoid different trial pool sizes; **(f)** parameters from **(e)** for different Δt ; each datum represents one session; **(g)** overlapping fraction of place coding populations depending on trial distance $\Delta trial$ for all sessions (grey lines) and different trial pooling width $\Delta t = 0$ (purple), $\Delta t = 3$ (red); **(h)** parameters of linear least squares fit as in **(e)** for $\Delta t = 3$; **(i)** parameters from **(h)** for different Δt , as in **(f)**; **(j)** overlap in the place cell population between the start/end 5 trials of consecutive sessions, no significant difference detected (p-value 0.37 for start(s)/start(s+1), end(s)/start(s+1) comparison, p-value 0.12 for end(s)/start(s+1) vs end(s)/end(s+1) comparison)

to randomly recruit neurons from the overall population of place coding neurons available in this session (vanishing slope of linear fit in Fig. 5.1(g-i). Comparing active place cells pooled from several adjacent trials ($\Delta t > 0$), however, reveals a decreasing similarity of the place coding population with trial- (and thus time-) difference over all sessions (slope < 0 for almost all sessions).

Finally, I checked whether the turnover within a single session could have an impact on the following session. Is the population of active place fields at the end of a session more likely to be active in the next session than the population active at the start? I calculated the overlap of the coding population of the first and last 5 trials of each session with the 5 start and end trials of the following session, Fig. 5.1(j). I could not detect any significant difference (p-value < 0.05) between neither pair of start and end trials, suggesting that there is no impact on place field activity of the following session on when within a session a neuron is place coding.

A similar analysis of neuron activity turnover within a session could not find any dynamics, as detected neurons showed some kind of activity throughout the whole session. The results presented in Fig. 5.1 are for a single mouse, only, but are consistently detected in all mice.

5.3 Turnover of the spatial map in CA1

On linear tracks, 1D-projections of firing maps can be arranged according to the position of their respective place fields and stacked on top of each other to visualize this in a now common way of display, Fig. 5.2(a). In line with several studies^{15,16,18} I find that the virtual environment is represented by a unique set of neurons, which changes in between different sessions. This can be visualized by applying the neuron order of a session as in Fig. 5.2(a) to other sessions: The diagonal arrangement of place fields increasingly blurs out with session difference, indicating a change in the neural representation of the environment, Fig. 5.2(b).

Before further going into detail of describing observed changes, I will introduce a denotation for the different states a neuron can take on, see Fig. 5.2(d): the status of a neuron being active (some calcium activity detected in a session) is α^+ . If it is silent, it is α^- . Neurons with status α^+ are further differentiated into place coding (β^+) and non-coding β^- neurons, using the place field detection algorithm. Lastly, neurons with a place field can be stable $\gamma_{\Delta s, s}^+$, or unstable $\gamma_{\Delta s, s}^-$, where stability is with respect to session $s - \Delta s$, from which a place field is recovered (+) or not (-). The exact criteria for stability are discussed in a later section.

A status β^\pm always requires α^+ , same as $\gamma_{\Delta s}^\pm$ implies β^+ . The denomination α^+ therefore captures active neurons, whether they have a place field (which can be stable, or not), or not, while β^- denotes active neurons, which do not have a place field. The nomenclature of γ^\pm follows according rules, with $\gamma_{\Delta s, s}^-$ being a neuron with a place field which has not been present in session $s - \Delta s$.

Observed changes in the map representation can be broken down into three qualitatively

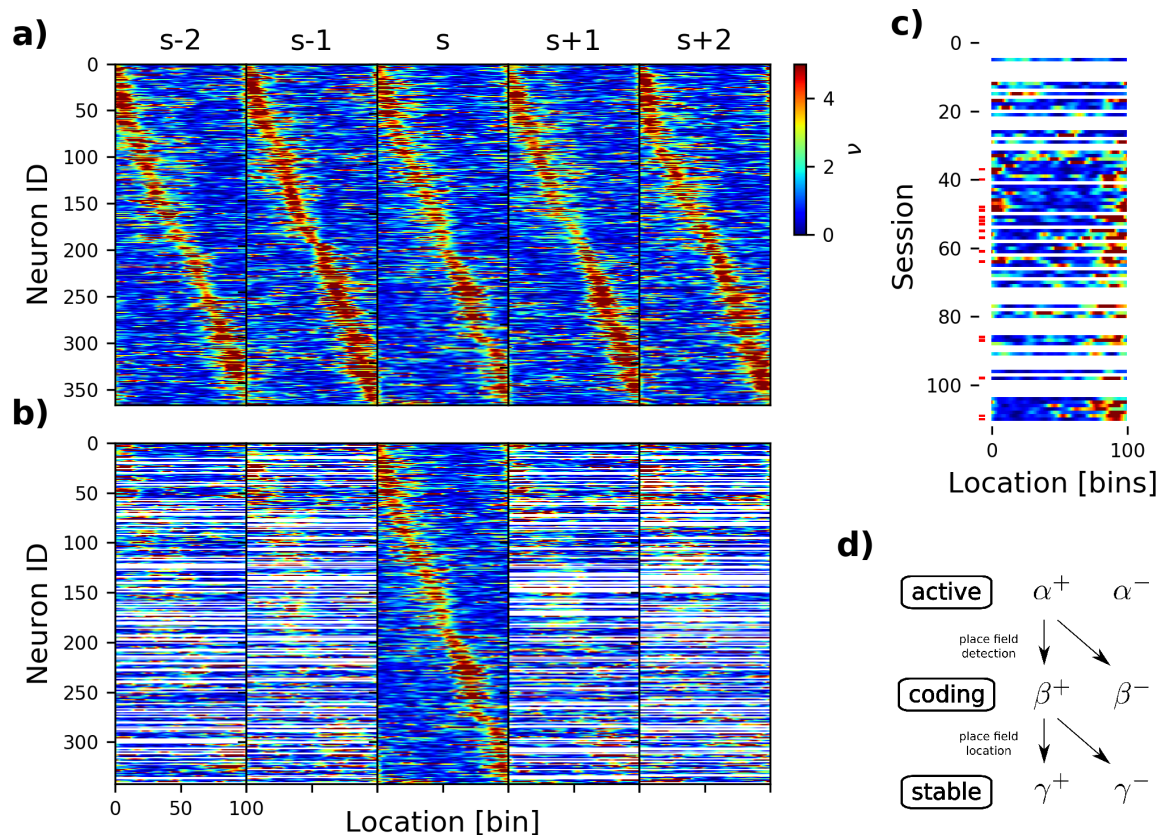


Figure 5.2: Hippocampal place representation changes over time; **(a)** firing maps (rows, smoothed with gaussian filter, $SD=2$ bins) of identified place cells arranged along the y-axis according to their place field location for five consecutive sessions; colorscale is firingrate in Hz (as defined in Sec. 2.6 with $SD=1$) above minimum firingrate (dark blue) of according neuron; **(b)** same sessions as in (a) with IDs and arrangement of neurons taken from the order of their place field locations in session s ; white rows are inactive neurons; **(c)** firing maps of a single neuron over all sessions, colorcode as before; red dots left to the map indicate that the cell was identified as a place cell in this session

different effects, which all can be observed by displaying the firing rate map of a single neuron over the course of all recorded sessions, Fig. 5.2(c). I will address these throughout the remainder of this chapter: Firstly, neurons active in one session (α^+) can become silent in the next (α^-), or have been silent in a previous session, resulting in a turnover of the active population, addressed in Sec. 5.4.2. Secondly, within the population of neurons remaining active, the spatial content in their activity can vary: place cells (β^+) can loose their place field between sessions (β^- , furtheron called non-coding, even though they might still be coding for some non-spatial stimulus), or previously non-coding neurons can obtain a place field, further discussed in Sec. 5.4.3. Finally, within the population of neuron with place fields in both sessions, place fields can change their location, ranging from small deviations from the original field γ^+ , up to long-distance relocation

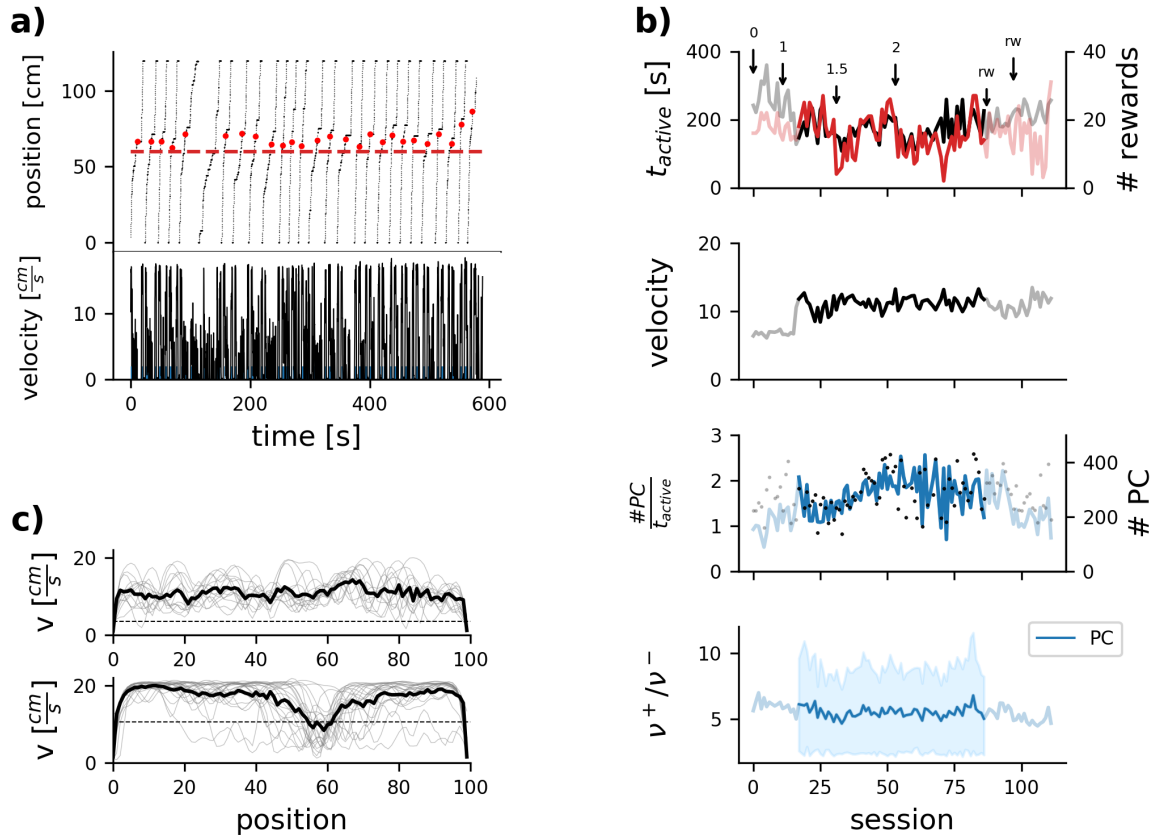


Figure 5.3: In non-changing environments, statistics of the hippocampal place map remain steady; **(a)** example of mouse behavior in a random session; top: animal position vs time, entry to reward zone marked by red dashed line, reward delivery highlighted in red; bottom: velocity of the animal; **(b)** session dependent behavioral and cell-activity statistics. Sessions identified as steady state for analysis of dynamics in solid colors, others transparent; top: active time per session (black, $\text{velocity} > 1 \frac{\text{cm}}{\text{s}}$) t_{active} and number of rewards received (red); sessions with changes to the environment highlighted by arrows: numbers indicate a change of waiting time for the reward after entering the reward area (in sec.), rw indicates a shift of the reward location (excluded from analyzed data); below: average velocity of the mouse; second from bottom: number of place cells (black dots) and ratio of place cells per t_{active} (blue line); bottom: ratio of in-field to out-of-field firing rate ratio $\frac{v^+}{v^-}$ of place cells with 1 SD confidence interval (displayed only for sessions within identified *baseline state*); **(c)** movement velocity of mouse during an early session (top) and a late session (bottom); grey lines: single trial data; black line: average over all trials within the session

γ^- . This behavior is further analyzed in Sec. 5.4.4.

These different impacts on place representation in the hippocampus form a type of hierarchical model, in that place field shifts can only be observed in neurons maintaining a place field, while the latter require a neuron to be active in both compared sessions.

5.3.1 Baseline conditions of the experiment

This chapter describes the dynamics of a network under baseline condition, after the animal has been familiarized with the environment. The first data entering the analysis is taken from at least 15 sessions after initial introduction to the environment and at least three sessions from the last introduced change to it (e.g. movement of the reward position). No further changes are introduced to the virtual environment, except for a slight increase of the waiting time for a reward at given timepoints in some mice (see Fig 5.3(b)).

To remove data from sessions introducing obvious inhomogeneities, the data is checked for transient behavior before further analysis. It is generally assumed for a mouse to assume stereotypic behavior in a familiar environment, resulting in repeated runs through the virtual environment to receive a reward at according positions, Fig. 5.3(a). The relative lack of explorative behavior results in an increased average velocity during active periods, found to stabilize in later sessions, Fig. 5.3(b), top, along with the time of active exploration per session (velocity $> 1 \frac{cm}{s}$), Fig. 5.3(b), second from top. Furthermore, mice adapt a characteristic slowing down behavior at the reward position after familiarization with the delayed reward delivery, Fig. 5.3(c).

Following the identification of place cells, their statistics provide useful information about the stabilization of the map, themselves. The number of place cells, as well as the number of cells coding for the reward is found to be slightly, but positively correlated with the time spent active t_{active} . Rescaling the number of place cells to the ratio of place cells per active time results in a slightly more stable reward- and place cell ratio over time, after an initial increase, Fig. 5.3(b), second from bottom. Mean values of in-field to out-of-field firing rate ratio $\frac{v^+}{v^-}$ remain quite stable for a wide range of sessions for place- and non-place cells alike, Fig. 5.3(b), bottom. Neither stability in firing rates, nor place cell number can be captured by means of statistical significance, reflecting the highly dynamic nature of hippocampal coding.

I performed similar analyses for all mice, identifying sessions for further analysis after initial effects in behavior and place field recruitment are not visible anymore.

5.4 Population dynamics across different levels

5.4.1 State transitions

Transitions between different states can be captured by the probabilities of a neuron occupying one state, given it has had a second state in the previous session. For example, the probability $p(\beta_{s+1}^+ | \beta_s^-)$ describes the probability of a neuron having a place field in a session after it has been active, but non-coding. Fig. 5.4 displays these neuron- and session-averaged probabilities for all possible combinations of states, allowing to discuss the impact of a neurons role in the network on its role in a subsequent session.

The graphs display the transition probabilities depending on previous states β^- (active, non-coding), γ^- (coding, non-stable) and γ^+ (stable) and thus capture the effects independent of

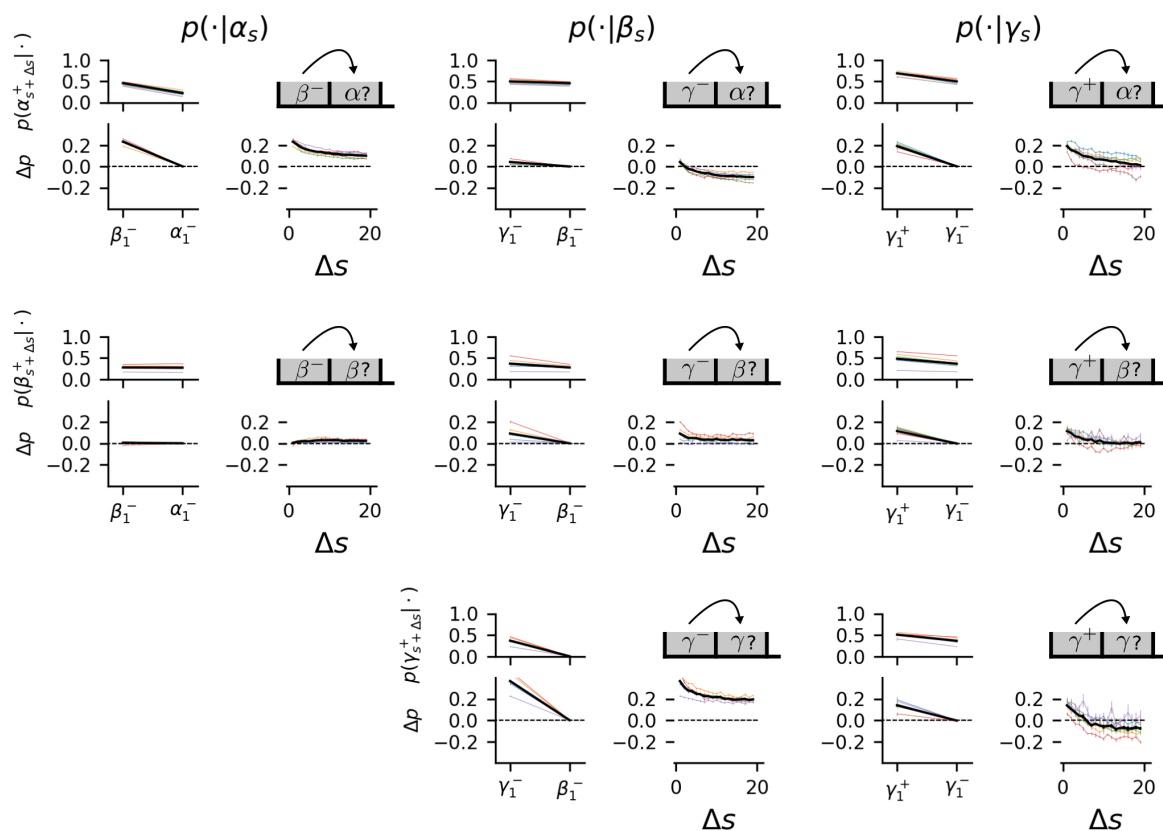


Figure 5.4: Contribution to network activity increases the chance of a neuron to contribute again; transition probabilities between different states from a previous (columns) to a following (rows) session. Probabilities are conditioned on lower states, each. Upper graph in each combination shows absolute probability values, lower left and right shows probability differences with respect to α^- , β^- , γ^- , respectively (see text for further description); left upper and lower graph show transition probabilities with $\Delta s = 1$, given state on y-axis in previous session; right graph shows long term impact on transition probability for up to 20 sessions difference; data from 6 mice, averaged over neurons and sessions

the other states (e.g. state α^+ instead of β^- would include states β^+ and γ^+). While the upper left graph for each combination of states shows the absolute probabilities for different mice, the lower left and right graphs, each, show the difference of probabilities towards the previous-state-dependent probabilities of lower hierarchy levels, such that $p(\gamma_{s+\Delta s}^+|\gamma_s^-)$ is subtracted by $p(\gamma_{s+\Delta s}^+|\beta_s^-)$, and so forth. Probabilities $p(x^+|\cdot)$ are further conditioned on the lower hierarchy, e.g. $p(\beta^+|\cdot)$ is the probability to encounter a place field, given that the neuron is at least in state α^+ in this session.

I find that, whenever a neuron plays a role in the network activity by being active, coding, or with a stable place field, it subsequently has an increased probability of being active in the following session, Fig. 5.4, top row. All states add some contribution, with the largest from β^-

and γ^+ (each adding $\sim 20\%$). Both states have a long-lasting effect on neuron activity, while the effect of just having a place field (γ^-) quickly decays, even below the level of $p(\alpha^+|\beta^-)$ after just 2-3 sessions.

Activity of a neuron itself does not have any impact on subsequent place coding, but having a place field (stable or instable) increases the probability of a neuron of having some place field again by $\sim 10\%$, each, Fig. 5.4, second row.

Transition probabilities to place field stability (γ^+) are a special case: A neuron is required to have a place field in session s to be able to retrieve it Δs sessions later, explaining the vanishing probability of $p(\gamma^+|\beta^-)$. The probability to encounter a stable place field decreases with time, but not to zero even after long waiting times (further details in Sec. 5.4.4), and encountering a stable place field further increases the probability for it to be stable again for a few sessions, Fig. 5.4, bottom row.

All effects described above are consistently found across all mice, though differ in exact values, with the largest spread encountered in the behavior after encountering a stable place field $p(\cdot|\gamma^+)$.

5.4.2 Turnover in the active population

Fig. 5.2(b) & Fig. 5.5(a) illustrate, how many neurons participating in the neural code at one given time turn silent in the next instance of the animal visiting the environment. In agreement with reports on a huge silent fraction of neurons in cortical areas¹⁶⁴ and with other works analyzing turnover dynamics in CA1^{15,16,18,52}, I find only a small fraction of neurons to be active during sessions ($\sim 20 - 40\%$, Fig. 3.10). The next lines are dedicated to understanding the underlying mechanism of neuron activation and how the population of neurons contributes to the overall dynamics.

Increased recurrence through heterogeneously distributed activation numbers Tracking individual neurons over all sessions, I find neurons, which are active in single, or few sessions only, up to neurons, which are active in all sessions, resulting in a broad distribution of active sessions N_{α^+} , Fig. 5.5(c). In agreement with earlier studies^{15,16}, the probability of a neuron to be active in two distinct sessions (also: recurrence probability, or active population overlap), $p(\alpha_{s+\Delta s}^+|\alpha_s^+)$, decreases with time passed in between them (here measured in session difference, Δs), Fig. 5.5(d).

Based on empirical data, I create two randomized datasets, to further investigate the origin of observed turnover. For the first one (R1), active sessions for all neurons are drawn randomly from the overall probability of encountering an active neuron in any session, $p_{\alpha^+} = \frac{\sum_{j=1}^{N_c} N_{\alpha^+}^j}{N_s N_c}$, resulting in a gaussian distribution of N_{α^+} , Fig. 5.5(c), grey, and a recurrence probability of $p_{\alpha^+}^{(0)}$ for all sessions, Fig. 5.5(d), black dashed line. For the second randomized dataset (R2), the active sessions of each neuron are randomly redistributed, such that it maintains N_{α^+} with respect to its empirical data, but destroys the order in which active sessions appear.

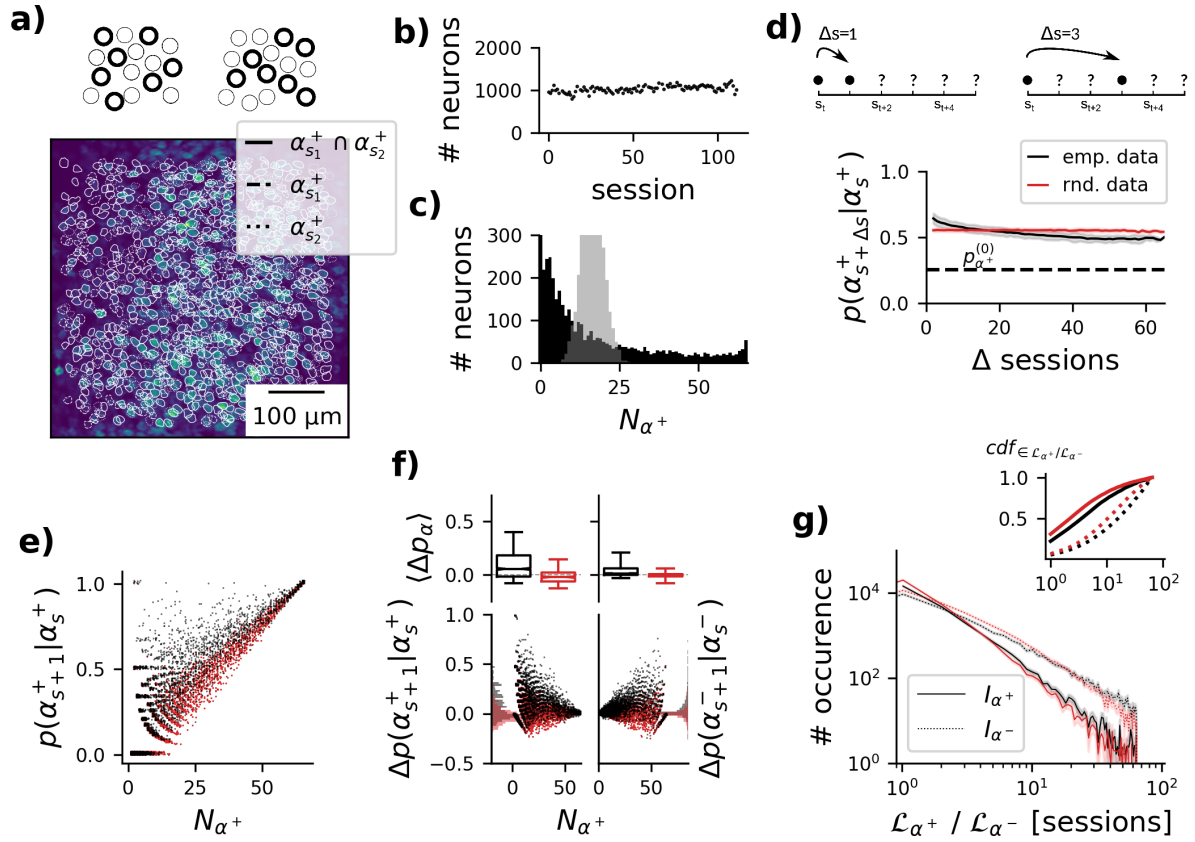


Figure 5.5: Turnover in the active population between sessions is dominated by heterogeneous levels of activity as well as an increase in activation probability after activity; **(a)** top: sketch showing how few neurons of a population (circles) are active (thick) in each given session (left and right), but the identity of active neurons changes; bottom: identified, active neurons in two consecutive sessions appearing in both (solid), the first only (dashed) or the second only (dotted); **(b)** total number of active neurons identified per session (dots) and total number of distinct neurons identified from all sessions (dashed line); **(c)** total number of active sessions per neuron N_{α^+} within identified steady state, as distributed over the whole, empirical dataset (black), and as distributed over randomized dataset R2 (see text); completely random distribution of active sessions (R1) displayed in grey; **(d)** top: sketch showing, how activity of neurons is compared between sessions with difference Δs , independent of the neurons state in other sessions; bottom: probability of a neuron to be active after Δs sessions, after being active with mean value (line) and 1-SD confidence interval (shaded area); randomly assigned activity with overlap fraction $p_{\alpha^+}^0$ marked in dashed black, randomly assigned activity maintaining neuron-wise activation levels N_{α} marked in red; colorcode for empirical (black) and random (red) data remains the same for (d-g); **(e)** probability of a neuron to be active in session $s+1$, after being active in s (data dithered by 1% for visibility); **(f)** bottom: difference of reactivation probability from (e) towards expected probability from random distribution of N_{α}^j (left) and the same for the probability of inactivity after an inactive session (right); top: boxplot of probability differences from plot below of empirical (black) vs randomized (red) data; line (average), box (25 – 75 prtile), whiskers (5 – 95 prtile); **(g)** distribution of periods of continuous activity of length \mathcal{L}_{α^+} (solid) and periods of continuous inactivity (dotted) of length \mathcal{L}_{α^-} ; inset: cumulative probability of an active session being part of a sequence of length \mathcal{L}_{α^+} (solid) and of an inactive session being part of a sequence of length \mathcal{L}_{α^-} (dotted); data from a single mouse

At all times, the empirically obtained recurrence probability remains above chance level obtained from R1, Fig. 5.5(d). The chance level of recurrence obtained by R2, however, is below the empirical data for small session differences Δs , only, and surpasses it for longer time intervals Δs , as it remains at a stable level.

Thus, the increased recurrence probability, with respect to R1 can partly be explained by the wide-spread distribution of activation numbers N_{α^+} . The increased number of neurons active in few, and in many sessions in the empirical dataset and R2, increases the probability of coactivation. Neurons active in few sessions have a low probability of being reactivated in another sessions, but contribute few counts to the calculation of the probability, only. Neurons with high N_{α^+} , on the other hand, contribute a high probability and many counts.

The observed inhomogeneity in the distribution of N_{α^+} already leads to a strong increase of the baseline recurrence, Fig. 5.5(d), red. However, it does not explain the time-dependency of the recurrence, observed in the empirical data.

Activation increases subsequent activation probability For many neurons, the probability of being active is increased right after they have been active in a previous session, Fig. 5.5(e,f). This can be expressed in the probability difference $\Delta p(\alpha_{s+1}^+|\alpha_s^+) = p(\alpha_{s+1}^+|\alpha_s^+) - \frac{N_{\alpha^+}}{N_s}$, with the first term being the observed probability of activation directly after an active session (Fig. 5.5(e)) and the second being the overall probability of this neuron being active. The distribution of Δp is significantly different from surrogate data R2 (two-sample KS-test, $I_{\alpha^+} : p = 6.2 \cdot 10^{-305}$) and reaches far into high probability values. A similar effect is found for inactive sessions (two-sample KS-test, $I_{\alpha^-} : p = 4.2 \cdot 10^{-277}$), which increase the probability of the next session being inactive.

Accordingly, the distribution of period lengths of continuous activity (I_{α^+} , of length \mathcal{L}_{α^+} sessions) as well as periods of continuous inactivity (inter-activity-intervals, I_{α^-} , of length \mathcal{L}_{α^-} sessions) show a decreased number of short periods and an increased number - or heavy tail - of long periods, with statistical significance compared to R2, Fig. 5.5(g). Active sessions thus are more likely to belong to long periods of continuous activity \mathcal{L}_{α^+} , with a similar effect in I_{α^-} , Fig. 5.5(g),inset.

5.4.3 Turnover in the population of place cells

Another factor contributing to the increasing dissimilarity of place maps with time is, that neurons remain active, but loose their place field. I find this to happen on a level very close to randomized data on a network level.

The fraction of neurons with identified place fields usually varies between $\approx 20 - 40\%$, depending on session (Fig. 5.6(b)) and mouse (Fig. 3.13(a)). N_{β^+} , the number of sessions with a field shows a skewed distribution with a large peak at 0 and a long tail, reaching up to 30 sessions for a few neurons and the probability of having a place field is largely uncorrelated with the overall

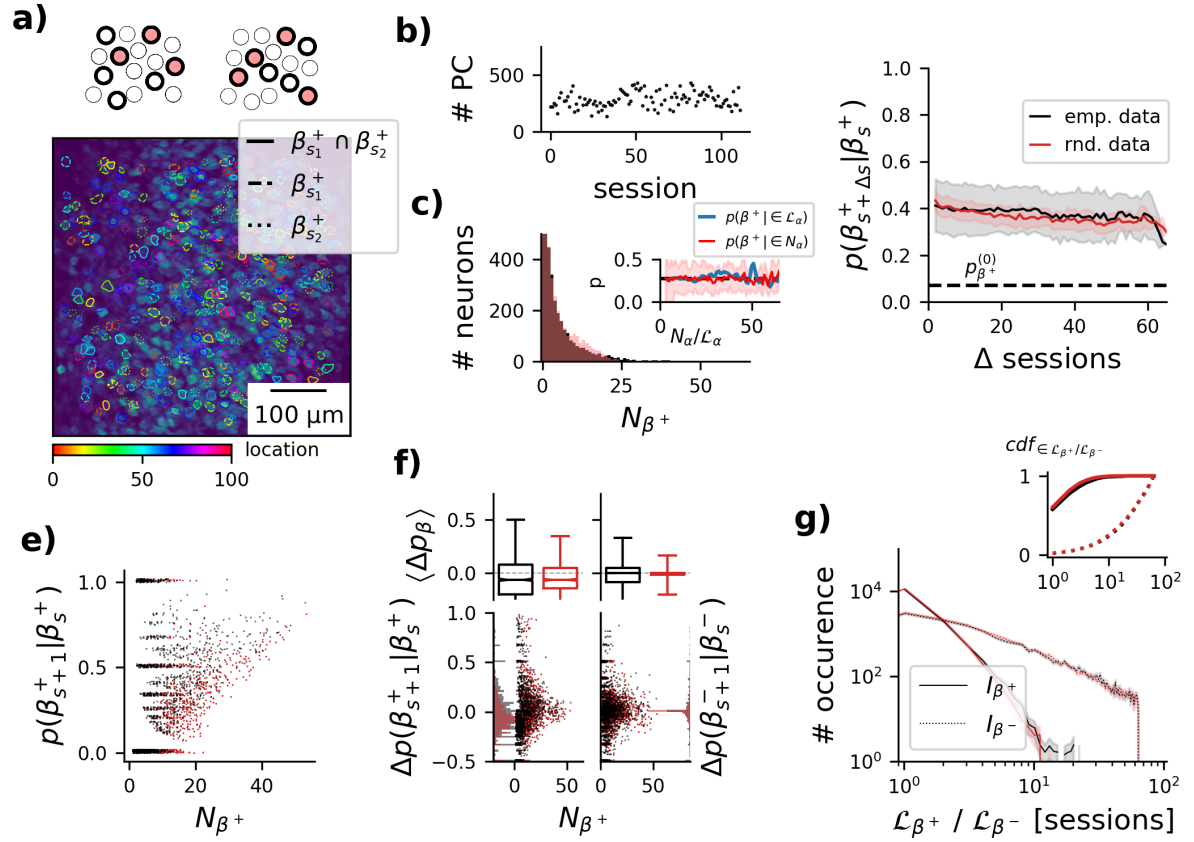


Figure 5.6: On a population level, neurons assuming the place coding property effectively follow a random process; **(a)** top: sketch of a neuron population with active neurons (thick) and a few place coding neurons (coloured), with identities changing in between two sessions (left and right); bottom: contours of place cells appearing in two consecutive sessions (solid), the first only (dashed) or the second only (dotted); colorcode describes location of identified place field center; **(b)** number of identified place cells per session; **(c)** total number of place coding sessions per neuron N_{β^+} ; inset: probability of place coding, given activation number N_{α^+} (red, shaded region: 1 SD confidence interval) or being part of a period of consecutive coding of length \mathcal{L}_{α^+} ; **(d)** recurrence of place coding property after Δs sessions, conditioned on the neuron being active for empirical (black) and surrogate (red) data, obtained from random shuffles of place coding sessions; **(e)** neuron-wise probability $p(\beta_{s+1}^+ | \beta_s^+)$ of having a place field, when being active in session $s+1$, after having a place field in session s (data dithered by 1% for visibility); **(f)** difference $\Delta p(\beta_{s+1}^+ | \beta_s^+)$ of data in (e) from average neuron-wise coding probability $\frac{N_{\beta^+}}{N_{\alpha^+}}$; distribution (bottom) and boxplot (top): line (average), box (25 – 75 percentile), whiskers (5 – 95 percentile); **(g)** distribution of periods of continuous place coding of length \mathcal{L}_{β^+} (solid) and periods of continuous non-coding (dotted) of length \mathcal{L}_{β^-} ; inset: cumulative probability of a place coding session being part of a sequence of length \mathcal{L}_{β^+} (solid) and of a non-coding session being part of a sequence of length \mathcal{L}_{β^-} (dotted); data from a single mouse

number of active sessions N_α , of this neuron or the duration of consecutive coding sessions \mathcal{L}_α , Fig. 5.6(c), inset.

Random activation of place coding Within the population of neurons active in two consecutive sessions, a fraction of $\approx 20 - 80\%$ of place cells are maintained, heavily dependent on individual mice, Fig. 5.6(c), Fig. 5.9(f). Similar to the active population, I find a decreasing recurrence of the place coding property in most mice, Fig. 5.6(d), black, however with an almost vanishing slope in most cases. As before, I generate a set of randomized data R3 of coding properties, to find the origin of the observed recurrence behavior. For this, I choose the empirically obtained data as a baseline, such that I redistribute sessions with coding properties randomly over the empirically observed active sessions. Interestingly, the observed temporal dependence of place cell recurrence can be almost perfectly reproduced by a randomized dataset R3, Fig. 5.6(d), red.

The average probability of coding after a coding session $p(\beta_{s+1}^+|\beta_s^+)$, which is further conditioned on the cell actually being active in session s is only slightly ($1 - 5\%$), but significantly increased (p-value < 0.01 for all mice, Mann-Whitney U test), Fig. 5.6(f), Fig. 5.9(g). The distribution of $\Delta p(\beta_{s+1}^+|\beta_s^+)$, is significantly wider spread in all mice than obtained from R3 (KS-test, p-value < 0.001), indicating that individual neurons should play a role in the network dynamics, obscured within network-averaged measures. Note that both, empirical and surrogate data have average probabilities $\Delta p(\beta_{s+1}^+|\beta_s^+) < 0$, caused by the need to remove a few sessions from within the dataset, due to bad imaging conditions, or large shifts of the FoV, disallowing for a few comparisons.

5.4.4 Place field dynamics

A last factor, contributing to the observed dynamics in the hippocampal place representation, Fig. 5.2 is, that place fields of active and coding neurons can change their position. The following lines will analyze this behavior within the population of neurons with place fields in compared sessions.

A probabilistic description of place field movement Place fields obtained according to Sec. 2.6 are not solely defined by a single point in the environment, but rather by the posterior distribution of their location, Fig. 5.7(a), left. In many cases (upper plot), it is sharply peaked, allowing the place field description by the single point of maximum likelihood, but in some cases (lower plot) there is a larger uncertainty - often non-gaussian - spanning several bins. Accordingly, instead of describing place field movements between two sessions by a single variable, I obtain a distribution of possible shifts between two place fields by sampling distances between reference- and remapped- place field. Start- and end- points of the remapping are drawn according to

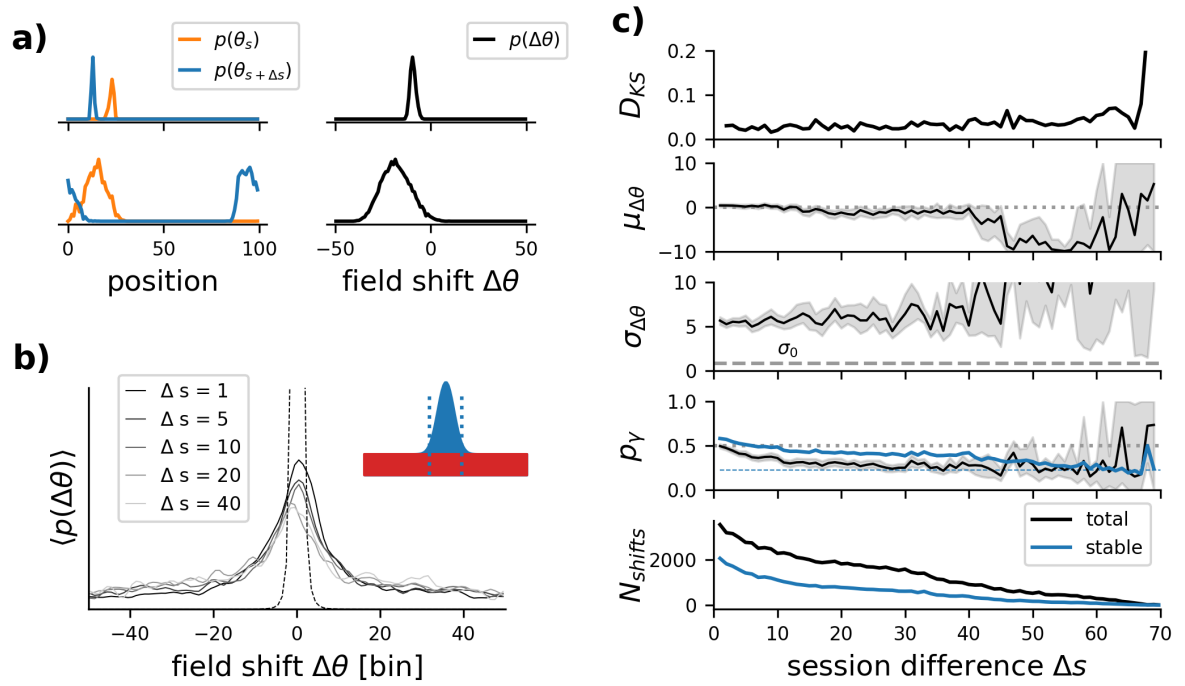


Figure 5.7: Place fields remain stable above chance level over long times; **(a)** two examples of place field shift $p(\Delta\theta)$ (black) obtained from the probability distribution $p(\theta)$ of reference and relocated field at session s (orange) and $s + \Delta s$ (blue), respectively; **(b)** distribution of place field shifts depending on time-interval Δs between sessions; uncertainty of place field location (dashed black) obtained from probability distribution averaged over all detected place fields; inset displays sketch of assumed behavior: a subset of place fields remains stable, tolerating minor fluctuations (blue, normal distributed with position μ , SD σ), while the others randomly remap (red, homogeneous distribution); **(c)** parameters obtained by fit of a Gaussian (stable) plus offset (relocation) to place field shift distributions as in (b) for all Δs : location $\mu_{\Delta\theta}$, SD $\sigma_{\Delta\theta}$, and probability mass accumulated in Gaussian, p_γ ; SD σ_0 of place field location accuracy displayed in $\sigma_{\Delta\theta}$ -plot (dashed grey); total fraction of stable place fields (probability mass of Gaussian plus shifts $|\Delta\theta| < 1.96\sigma_{\Delta\theta}$ not explained by Gaussian) added as blue line in plot for p_γ ; 95% confidence interval (grey) obtained from bootstrapping ($N_{bs} = 1000$) over all shifts; bottom: number of shifts available for calculation of shift distribution

respective reference- and remapped- posterior distributions of place field location and result in a distribution of the place field shift for this neuron, Fig. 5.7(a), right.

Place field shifts follow a characteristic pattern over long time The overall distribution of place field shifts for fixed Δs displays a characteristic shape, approximated by a gaussian plus constant offset, Fig. 5.7(b). It is obtained from the normalized sum over all observed field shifts with according session difference and is very similar to what has been reported in other studies^{15–17}. The central peak, reliably identified around $\Delta\theta = 0$ up to large Δs , Fig. 5.7(c), $\mu_{\Delta\theta}$, resembles a normal distribution, while field shifts above a certain distance seem to be almost homogeneously distributed. This behavior can be captured in the form of

$$(5.1) \quad p(\Delta\theta) = \underbrace{\frac{p_\gamma}{\sqrt{2\pi\sigma_{\Delta\theta}^2}} \exp\left(-\frac{(\Delta\theta - \mu_{\Delta\theta})^2}{2\sigma_{\Delta\theta}^2}\right)}_{\text{stability}} + \underbrace{\frac{1 - p_\gamma}{N_{bin}}}_{\text{random relocation}},$$

assuming a normalized distribution in which the probability mass p_γ is accumulated in the gaussian distribution and a fraction of $1 - p_\gamma$ forms a constant offset, well in agreement with the empirically observed distributions (low test statistic D_{KS} for all Δs , Fig. 5.7(c), top). Parameters are obtained from bootstrapping ($N = 1000$) over least squares fits to random samples of place field shift distributions, revealing a high certainty (narrow 95%-confidence interval) over a large range of session differences, ranging up to $\Delta s \approx 50$ sessions (40 – 50 days), where the number of observed shifts becomes too small, Fig. 5.7(c), bottom. Over the same range, both the center of the gaussian, $\mu_{\Delta\theta}$ and its standard deviation $\sigma_{\Delta\theta}$ remain almost constant, while the probability mass p_γ decays monotonously, Fig. 5.7(c).

Increased probability of place field stability for up to 2 months The probability mass of the gaussian p_γ which describes shifts being *anchored* (stable place field, contributing to observed gaussian peak) decays with time from $\approx 50\%$ to around 20% at $\Delta s \approx 50$. At the same time, the probability for random place field relocation $1 - p_\gamma$ rises accordingly, increasing the probability of long-range shifts, but also of a place field randomly reoccurring at the very same position. Considering all fields within a 1.96-SD interval (95% of the mass of the gaussian) to be stable, results in the Δs -dependent probability $p(\gamma_{\Delta s, s}^+ | \beta_s^+)$ of encountering a stable place field after Δs sessions. The fraction of stable place fields decreases from consecutive sessions ($\Delta s = 1$, $\approx 60 - 75\%$) to $\Delta s = 50$, where it approaches chance level ($= \frac{2\sigma_{\Delta\theta}}{n_{bin}}$), Fig. 5.7(f), dashed blue, Fig. 5.9(d), left. Unsurprisingly, increasing multiples of $\sigma_{\Delta\theta}$ (thus increasing the range of place field shifts to be considered *stable*) results in a monotonously increasing fraction of stable place fields with reduced growth after $SD = 1$, Fig. 5.9(d), right.

Overall, the turnover in active and coding population, as well as the described dynamics of place field relocation, leads to a mouse-dependent fraction of $\approx 10 - 25\%$ ($\Delta s = 1$) of neurons being active, and coding for the same location in two distinct sessions, decreasing with session

difference, Fig. 5.9(e), left. These results are also consistent with a commonly used alternative measure of population stability, a correlation of population vectors, in that they show a linear relationship, Fig. 5.9(e), right. Population vectors are calculated by stacking firing maps of place cells within one session on top of each other and correlating them binwise ($\hat{=}$ location wise) with the same stack of these cells in another session. For each pair of sessions, the mean value of these correlations can be computed and its average over sessions with the same Δs results in the value displayed.

Place fields either remain anchored or relocate randomly Above results suggest a two-fold behavior of place field stability: place fields either remain anchored around some reference position, allowing for some fluctuation with constant SD $\sigma_{\Delta\theta}$, or randomly relocate to any place on the linear track, with *any* also including the region in which a field is considered to be stable. To test if remaining and relocation are indeed two functionally different behaviors, I tested the probability of a field to reoccur around the same location within the next 10 sessions, depending on the shift distance towards the session before the tested interval, Fig. 5.8(a). I find, that indeed the probability of the place field to return to its original location drops rapidly to chance level once it shifts more than 2-SD (obtained from $\sigma_{\Delta\theta}(\Delta s = 1)$) from this location. Smaller shifts, however, result in high probabilities of coding for the same location again within the next few sessions. Neurons probability of subsequently being active, or coding show no significant dependency on the distance of place field shifts (data not shown).

5.4.5 Correlates of place field stability

Within the observed dynamics and neuron parameters, I find a few correlates of place field stability. Firstly, as discussed above, the probability of a field to reoccur in a given position, decreases with the distance of previously observed place field shifts, Fig. 5.8(a). Close vicinity to a previous field thus increases the chance of remaining anchored to a certain position.

Furthermore, the activity of a neuron in between two coding sessions has a major impact on its place fields stability: continuously coding neurons maintain a high probability of field stability ($\sim 60\%$, conditioned on the cell being active and having a place field in according sessions) over long periods of time, which even slightly increases, the longer the neuron maintains a place field. In contrast, place field stability of neurons which are non-coding (but active) in all intermittent sessions quickly decreases, Fig. 5.8(b). Interestingly, field stability is partly recovered when the neuron is found to be inactive in intermittent sessions, suggesting silence to play a role in the retention of functionality, for example through avoidance of hebbian mechanisms.

Finally, neuron parameters provide some information on how a place field will behave next, Fig. 5.8(c): stable place fields show a significantly increased field-to-baseline ratio of the firingrate A/A_0 , and slight, though non-significant increases in reliability and mutual information compared

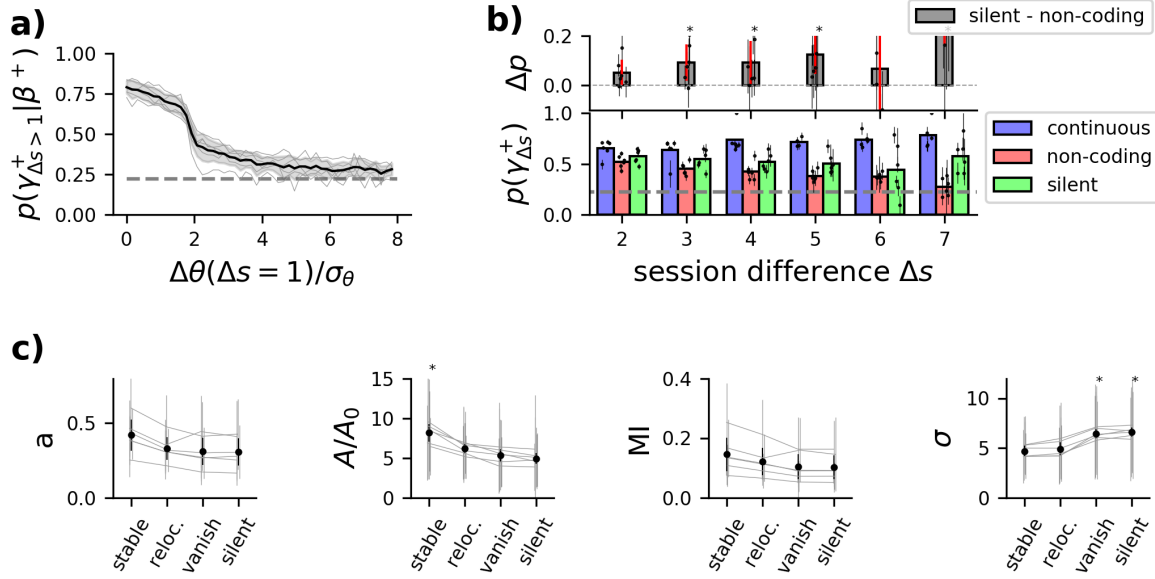


Figure 5.8: Place field stability is correlated with neuron behavior; **(a)** probability of returning to a reference place field ($\Delta\theta < \sigma_{\Delta\theta}$) within the next 10 sessions after an observed place field shift of magnitude indicated on the ordinate; **(b)** bottom: probability of a place field returning to its original location after Δs sessions, depending on the cells behavior in sessions lying in between: continuously having a place field (blue), active but non-coding (red) and inactive (green); top: difference in stability between inactive and active, but non-coding case; single mouse data displayed as dots; **(c)** single mouse (grey) and mice-averaged (black) data of neuron- and place field parameters in sessions before a place field remains stable, relocates, or vanishes; errorbars indicate 1SD; from left to right: place field reliability a , ratio of in-field to out-of-field firingrate A/A_0 , mutual information of activity with animal location MI , place field width σ . Significant deviations of a single case are marked by * (p-value < 0.05); data from $n = 6$ mice

to relocating or vanishing place fields, in line with what has been reported in an earlier work¹⁵. Place fields about to vanish show a significant increase in their width, while the activity.

5.4.6 Results are consistent, but quantitatively different across mice

Above described behavior for the turnover of the active and place coding population as well as for the stability of place fields has provided an overview of the different mechanisms in place. These can be found across all six mice used in this study, following the same qualitative behavior, but differing widely in some of the parameters.

For all mice, I find the distribution of activation numbers N_α to be widely different from a distribution of active sessions randomly drawn according to the network averaged probability of activity p_α^0 , Fig. 5.9(a). The distributions display a positive skew with many neurons active in few sessions only and a long tail with some neurons active in all sessions. For all mice, the probability of a neuron being active after an active session is significantly larger than

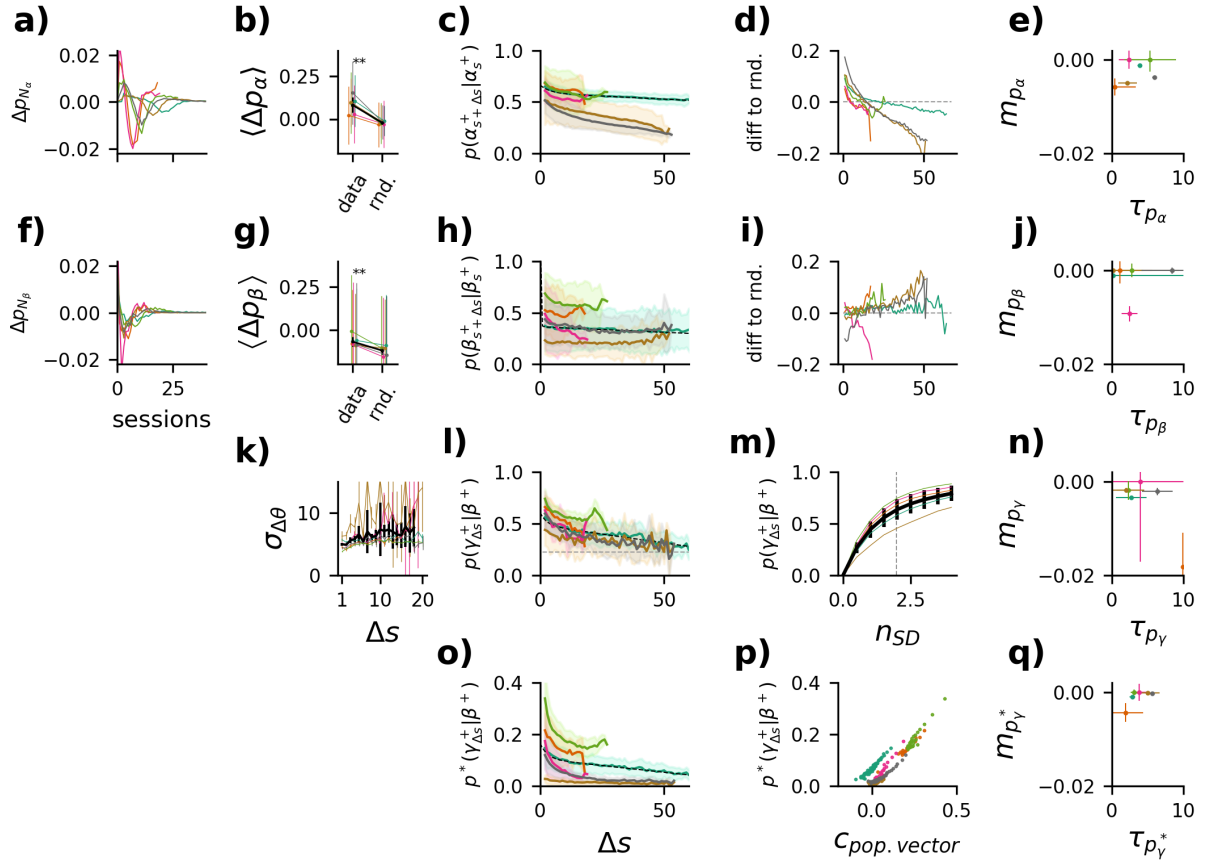


Figure 5.9: Network dynamics are qualitatively similar across mice, but show quantitative difference; all plots show data for six mice (different colors); **(a)** difference of empirically observed distribution of N_α towards randomized data (R1, see Sec. 5.4.2); **(b)** same as Fig. 5.5(f); **(c)** same as Fig. 5.5(d); fit of $a \exp\left(-\frac{\Delta s}{\tau_{p_\alpha}}\right) + m_{p_\alpha} \Delta s + b$ added to one dataset (dashed black line); **(d)** data of (c) subtracted by recurrence probability from random dataset R2; **(e)** parameters of fit from (c) for different mice; **(f)** difference of empirically observed distribution of N_β towards randomized data (randomly drawn place coding sessions within empirically observed active sessions with overall probability of finding a place coding session); **(g)** same as Fig. 5.6(f); **(h)** same as Fig. 5.6; **(i)** data of (h) subtracted by recurrence probability from random dataset R3; fit as in (c) added to one dataset (dashed black line); **(j)** parameters of fit from (h) for different mice; **(k)** same as Fig. 5.7(c), center; average across animals on top (black line); **(l)** same as Fig. 5.7(c), second from bottom; fit as in (c) added to one dataset (dashed black line); **(m)** dependence of stable fraction at $\Delta s = 1$ on definition of stability: multiples of gaussian width $\sigma_{\Delta\theta}$ beneath which shifts are defined to be stable; **(n)** parameters of fit from (l) for different mice; **(o)** overall probability of recurrence of a place field after Δs sessions, not conditioned on the neuron being active or coding in according session; fit as in (c) added to one dataset (dashed black line); **(p)** each data-point from (o) with according average population vector correlation (see text for description); **(q)** parameters of fit from (o) for different mice

expected from its activation number N_α (p-value $< 10^{-22}$ for all mice), Fig. 5.9(b) and the overlap of the active population decreases gradually with time difference, Fig. 5.9(c-e). I find that the overlap between consecutive sessions displays a range between $\sim 50 - 70\%$ across different mice, with a following decrease best described by an exponential plus a linear decay: $p(\alpha_{s+\Delta s}^+ | \alpha_s^+) = a \exp\left(-\frac{\Delta s}{\tau_{p_\alpha}}\right) + m_{p_\alpha} \Delta s + b$ and typically is less than 1% per session, Fig. 5.9(e).

The number of place coding sessions per neuron N_β also shows a distribution different from randomly chosen place coding sessions according to the network averaged probability p_β^0 , even when the sessions are chosen from empirically obtained active sessions, only (random set R3, see Sec. 5.4.3), Fig. 5.9(f). The probability of encountering another place coding session after a neuron has expressed a place field is only slightly, but still significantly enhanced (p-value < 0.01 for all mice). The overlap of the place coding population within the population of neurons active in two distinct sessions remains remarkably constant over arbitrary session differences Δs for most mice, Fig. 5.9(h,j). I find that with increasing time difference, random data (R3) increasingly underestimates the overlap, Fig. 5.9(i), suggesting a long-term effect of place coding recurrence, despite a vanishing short-term increase over random data.

Place field recurrence amongst place coding neurons appears to be most consistent across different animals: the standard deviation of place field shifts considered to be stable ($\sigma_{\Delta\theta}$) displays similar values for consecutive sessions (4.9 ± 0.82 bins, mean \pm SD over all mice) and shows a remarkably common temporal evolution for 4 out of 6 mice over several dozen session, Fig. 5.9(k). The probability of place field stability varies between 40 – 75% for consecutive sessions (stable place fields are defined as field shifts $< n_{SD} \sigma_{\Delta\theta}$; mice show similar dependence on the factor n_{SD} , Fig. 5.9(m)) and slowly decreases over the whole range of the data, Fig. 5.9(l,n).

The previous paragraphs have described the observed turnover in the active and place coding population, as well as the place field locations, conditioned on a subpopulation, each. Place coding recurrence is conditioned on respective neurons being active in the respective sessions and place field stability depends on the neurons expressing a place field in the sessions, at all. The probability of the recurrence of a place field independent of these conditions therefore is composed of all of the above effects. The probability of a neuron expressing a place field in the same location Δs sessions after it has expressed the place field shows a wide spread between different mice. While one mouse barely recovers any place field in the next session ($\sim 2\%$ recurrence), another one displays a strong consistency in the spatial code with $\sim 35\%$ of all place fields recovered in the next session, Fig. 5.9(o). Probabilities of recurrence for most mice drop rapidly with increasing time, but show a trend of gradual, slow decay after an initial exponential drop with a characteristic time of 3-5 sessions, Fig. 5.9(q).

In a last test, I compare the above described overall place field recurrence with a commonly used measure of spatial map coherence: the population vector correlation. Population vectors are composed from the firing maps of all neurons within the dataset. For this, all firing maps (gaussian filtered with $SD=2$) are stacked on top of each other and population vectors are

extracted as all the entries corresponding to one bin of the environment (vectors therefore have the length of the overall number of neurons in the dataset). Silent neurons contribute zeros in the according positions. Population vectors can then be correlated between arbitrary sessions and yield an alternative description of the consistency of the spatial map. I find that the probability of place field recurrence $p^*(\gamma_{\Delta s}^+ | \beta^+)$ is linearly related to the population vector correlation, showing that the analysis provides results consistent with alternative measures.

5.5 Time in the hippocampus

Throughout this chapter, I have provided an analysis of network dynamics, based on changes between sessions. The assumed time axis of experienced sessions, however, does not necessarily represent the scale along which the dynamics in the hippocampus happen. In the case of the underlying data, measurements were obtained twice per weekday, with about 4 hour difference between same-day sessions and accordingly 20h difference over night and 68h over a weekend. This introduces a huge inhomogeneity in the time axis, motivating a look at alternative time representations, only to find non-conclusive evidence.

I compared the turnover of active and coding population, as well as place field stability over time for four different candidates of the temporal axis: session difference Δs , nights spent in between sessions $\Delta nights$, number of trials experienced between sessions $\Delta trials$ (coarse scaled to multiples of 10 to allow calculating dynamics) and the absolute time in hours between sessions Δt , Fig. 5.10(a,c,f,i). For each given session difference Δs , there is a number of different corresponding points on other time scales (e.g. $\Delta s = 1$ can be realized by $\Delta t \in [4, 20, 68]$), allowing to test, whether different points on one time scale have an impact on the dynamics corresponding to a single point on the other.

To remain within the given example, I pooled the values obtained for the three different dynamics from all sessions with difference $\Delta s = 1$ and $\Delta t = 4h$ into one set, all with $\Delta s = 1$ and $\Delta t = 20h$ into another, and so forth. A one-way Kruskal-Wallis test then reveals, whether these sets can be described as significantly different from one another, or not, revealing whether the absolute time passed has an impact, when the number of sessions Δs is fixed.

A test for all three dynamics (Fig. 5.10(b,d,g,j)) and all combinations of the time axes reveals, that for none of the time scales the hypothesis of having no impact on the dynamics can be rejected for all Δ , in neither direction, Fig. 5.10(e,h,k). Interestingly, however, it is mostly only one or two types of dynamics, passing the test with a p-value < 0.01 for one combination of compared axes, e.g. activity turnover in $\Delta nights$ and Δt vs. Δs and place cell turnover/place field stability in $\Delta trials$ vs Δs (and vice versa), suggesting the possibility of the observed dynamics happening on separate timescales.

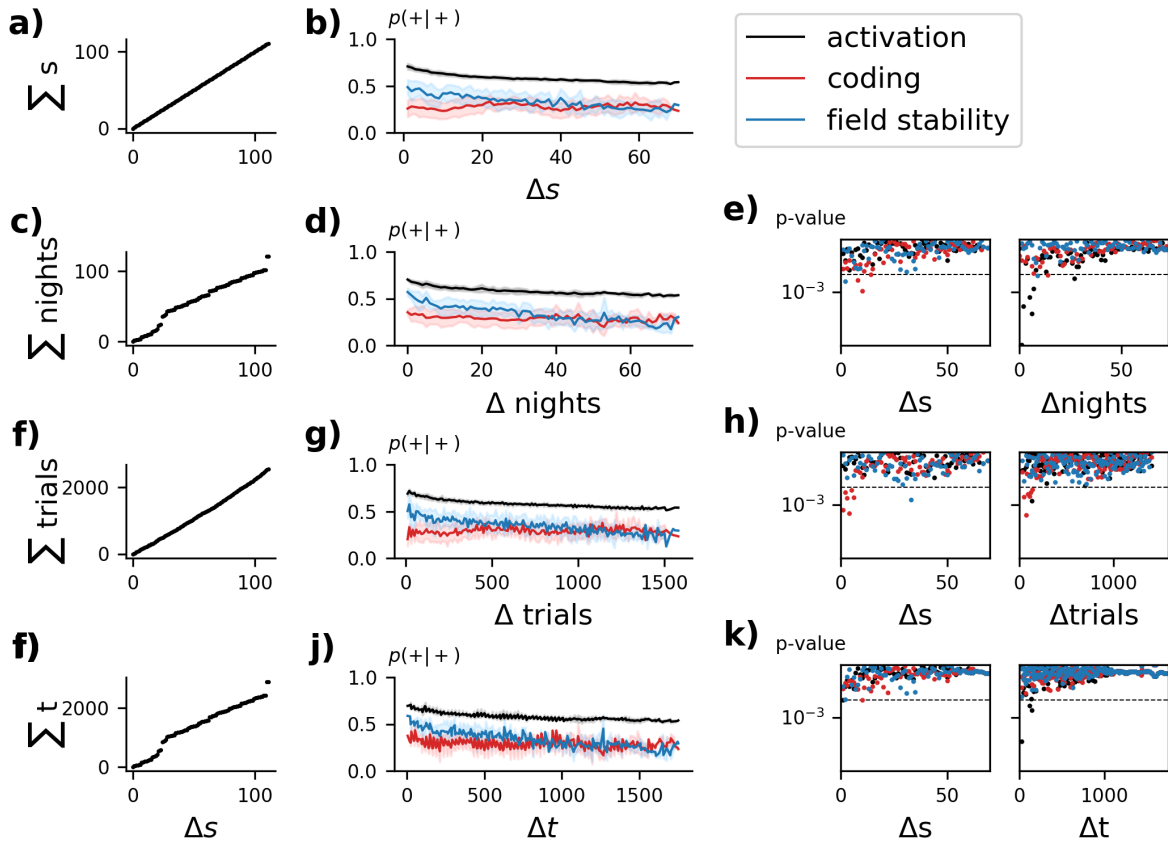


Figure 5.10: Neither time scale of session difference, absolute time, nights spent, or experienced trials can be identified as the unique dependency for dynamics; **(a)** the number of passed sessions is the parameter along which previous dynamics have been explored. Naturally, it is linear with itself; **(b)** dynamics of turnover in the active (black, see Fig. 5.5(d)) and coding (red, see Fig. 5.6(d)) population, as well as from place field recurrence (blue, see Fig. 5.7(c), second from bottom), as described before; **(c)** passed nights as a function of experienced sessions; **(d)** dynamics as in (b), calculated as a function of nights spent between them; **(e)** left: p-values of hypothesis tests (Kruskal-Wallis) examining whether there is significant impact of the number of nights when the number of sessions is held constant for the different dynamics as described before (colorcode as before); rejection ($p_{val} < 0.01$) signals significant impact of nights; right: p-values of hypothesis tests examining whether there is significant impact of session numbers when the number of nights is held constant; **(f)** number of trials as a function of sessions passed; **(g)** same as (d) with the number of experienced trials as abscissa (coarse grained in steps of 10 trials); **(h)** same as (e) with hypothesis tests examining the mutual impact of session numbers and experienced trials; **(i)** passed time in hours as a function of experienced sessions; **(j)** same as (d) with absolute time t as abscissa; **(k)** same as (e) with hypothesis tests examining the mutual impact of session numbers and absolute time passed

5.6 Discussion

In this chapter, I described how the turnover observed in the spatial map across different recording sessions can be dissected into three different levels: a slow, gradual change of the population of active neurons, a close to random process of place field activation and a long-term stability of place field location, when a neuron exhibits a place field again after some time. I analyzed transitions between different states of a neuron between consecutive sessions to find, how activity of neurons is commonly clustered in adjacent sessions and how neurons express place fields almost randomly in time. I find, how turnover does not only appear in between sessions, but how the spatial representation also shows turnover on a short time scale within a single session. The later Ch. 7 elaborates on the different effects in detail, and how they can be interpreted in the light of current state of research.

Correlates of place field stability I find individual transitions to be correlated with certain neuron parameters, indicating stability: a neuron with a place field about to remain stable usually carries more information content, is more reliable and has a larger peak to baseline firing rate ratio than neurons whose place field vanish in the next session. The ability to identify such predictive parameter values is conflicting with the notion of place coding turnover being due to upstream dynamics happening on a time scale of hours in between distinct sessions. Instead cells might be driven by transient synapses, which are more prone to be subject to spine turnover¹³.

The role of silence Place field stability over long times is strongly influenced by the behavior of the neuron in between two distinct sessions: neurons consistently maintaining a place field have a high probability to remain stable and anchored to a reference position. In contrast, the probability to remain stable decreases gradually in non-place coding neurons, with the effect being slightly stronger in active, but non-coding neurons, than in neurons remaining completely silent. This suggests a role of hebbian learning in the role of place field stability, such that a learned place field is weakened, or abolished by random input during an active, non-coding session¹⁰, while synaptic weights are preserved when a neuron instead remains silent. A similar effect has been identified in the vocal motor cortex of songbirds, where learned sequences are preserved by targeted silencing of single neurons¹⁶⁹.

The existence of a silent population within cortical regions is quite well established, the function however remains unclear. Maintenance of neurons even without activity comes at a huge metabolic cost - one of the primary limiting factors of brain size in mammals¹⁷⁰⁻¹⁷³ and therefore suggests some underlying evolutionary advantage. An early study, attempting to activate silent neurons in the hippocampus by antidromic stimulation, failed to do so, suggesting an active effort of the network to keep them silent.

Results of this, and previous works, suggest that silence might serve two functions: (1) it allows sparse representations of environments to undergo gradual change by a change in

the active subpopulation, allowing the network to time-stamp events^{16,27,174}, realizing episodic memory. (2) silent neurons might contain connections, which should be maintained and not be altered by ongoing network activity. It would be instructive to obtain a further understanding of the mechanism causing large populations of the hippocampus to remain silent under most conditions.

Time in the hippocampus This work has described network dynamics on a time axis defined by the number of exposures to the virtual environment, the session number. Hippocampal functions, however, occur on a number of different timescales: Synaptic turnover is usually reported in the absolute time scale of hours or days^{12,13}, assuming biological functions to be an ever ongoing process, with experiments of place field formation suggesting a strong dependence on time spent within the environment⁶¹. Stabilization of representations, however, are reported to be rather influenced by exposures to the environment (session numbers, trials completed)¹⁰, but also affected by SWR sleep (nights), during which experienced activity patterns are reactivated and strengthened¹⁰.

An analysis of above discussed dynamics on these different time axes did not yield concluding evidence of a preferred time axis, being unable to reject significant impact of any of these on all three levels of turnover. It is possible, however, that the dynamics of turnover on the different levels happen along different time axes: turnover dynamics of the active population are accurately described by session difference, with no significant impact of nights, trial numbers, or absolute time, while the turnover in the population of place coding cells is more accurately described by absolute time differences.

Data used in this study only offers a small range of different combinations of time and experience, while one of them is fixed. The possibly distinct impact of these two factors on different levels of hippocampal dynamics, however, proposes an interesting subject for further studies, where these differences could be studied systematically.

The impact of time on the hippocampal dynamics might further be influenced by the conditions under which animals are held. Enriched cages are reported to have a strong impact on animal performance and neuronal dynamics^{175,176}, while standard cages, combined with single breeding of mice provides little to no stimuli to the animal in between sessions. It is feasible, that hippocampal dynamics are influenced by these conditions, even though spatial representations in general are reported to be independent of one another and stable towards learning experiences in other environments^{20,54}.

Chapter Six

Influences on network dynamics

6.1 Content

Network level dynamics previously reported^{15,18,20,86} and further described and analyzed in Ch. 5 have identified a turnover on different levels of neuron activity. The description has been based on network averaged statistics of the population of recorded neurons, while omitting effects potentially introduced by heterogeneities within the population, or over time. A number of factors, however, suggest there to exist specific roles of individual neurons: strengthening of synapses by experience and subsequent consolidation can pronounce the role of individual neurons by undergoing significant long term potentiation, as has been reported for connections between CA3 and CA1¹⁰. Reliable replay of sequences, as observed in a wide range of studies⁴⁷ obviously requires specific roles of neurons. Furthermore, studies have reported the existence of subpopulations of neurons specialized for specific tasks, such as *reward* cells in CA1, uniquely responsive for reward coding across different environments¹²⁰. In contrast to these findings, studies analyzing synaptic spine turnover have identified timescales of few days, over which the whole population of spines within CA1 is undergoing complete turnover^{12,13,177}, questioning how neurons can take on prominent roles over long time scales.

The contribution of such specific neurons, supposedly carrying the functions of spatial representation and memory processes, has not been analyzed over longer time scales in a large scale dataset as available to me within the scope of this work. In this chapter, I will depart from a population level description of dynamics to elucidate the role of individual neurons in observed dynamics.

What are the effects of salient regions on observed long-term dynamics? Is there a constant subpopulation of neurons, contributing to representational stability over long time scales, or is it a fluctuating set of neurons contributing to a reliable spatial representation? What is the contribution of individual neurons to the overall dynamics and are there characteristic timescales?

To answer these questions, Sec. 6.2 first analyzes the impact of specific regions of the experimental setup on the dynamics of neurons. I then define individual stability parameter $r_\alpha, r_\beta, r_\gamma$ to find how individual neurons contribute to place coding overall and at different times within the

experiment (Sec. 6.3)

In contrast to the study in Ch. 4⁶² which has analyzed the formation of a place map, I find that other than edge locations of the linear track, there is no special role of some location in the place representation after extensive familiarization with an environment. Furthermore, only a small, swiftly changing population of neurons at each given time participate strongly to a stable representation of space, with reliable contributions to the overall network activity lasting over several days, to weeks, and reliable contributions to place coding lasting over few sessions or days, only.

6.2 The impact of different locations

The representation of space in the hippocampus is obviously shaped by features of the environment. A range of studies have shown how strong visual cues lead to an overrepresentation of according regions in the cognitive map^{3,122,124,125}. Also locations of behavioral relevance, such as the possibility to receive a reward have been reported to result in overrepresentations^{62,112,116–118,120,121}, however less reliably and have been suggested to depend on (1) no additional visual cue at rewarded locations and (2) repeated relocation of the reward¹²⁰.

Only most behaviorally relevant regions maintain a specific role In the mice data of this study, the animals are subject to an unchanging environment, with a single, constant reward location (see sketch, Fig. 6.1(a)), the first condition of no additional cue is fulfilled, while the second is not at all. Accordingly, only a single one of six mice shows a significant overrepresentation of the reward location after removing initial training sessions, Fig. 6.1(d) right. Other mice either lose their overrepresentation after a few initial sessions (see example in Fig. 6.1(c)), or don't exhibit an overrepresentation at all (not shown). They even show a slightly decreased density of place fields at the reward location, Fig. 6.1(d), right, possibly caused by the slowing down while approaching the region. In fact, the place field density adopts a stereotypical shape, which it maintains over wide range of sessions (Fig. 6.1(d)) with almost indistinguishable non-rewarded and rewarded locations. The boundaries of the linear track at the beginning and the end, however, maintain an increased density over all sessions in agreement with previous studies¹⁵, Fig. 6.1(d), supposedly caused by the considerable behavioral manipulation of teleport and three seconds disabled advancement of the VR (see Sec. 3.4). Neuron parameters are not found to be significantly different in any of the regions, Fig. 6.1(e), except for the place field width, which is strongly reduced at the edges of the linear track and slightly reduced at rewarded positions, again possibly caused by the slowed down movement.

Place field turnover becomes largely independent of location The status of a neuron preceding and succeeding a session in which the neuron has a place field can display a number of different cases: the cell can be silent, active but non-coding or have a place field in the same, or

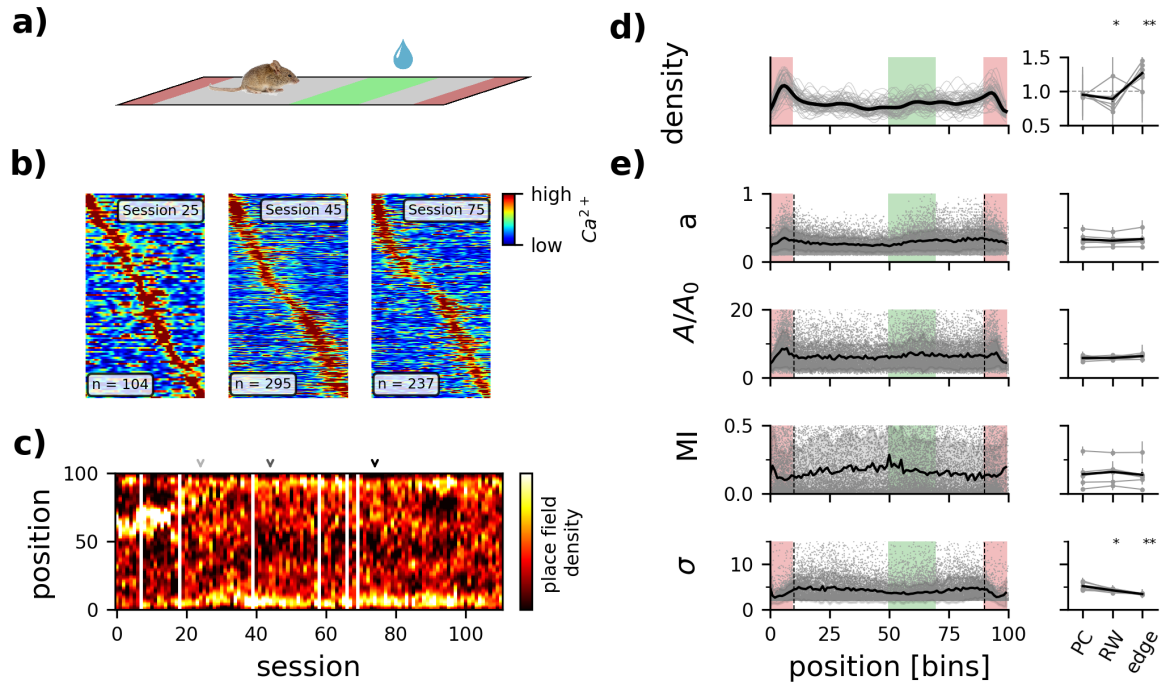


Figure 6.1: Hippocampal place map displays local inhomogeneities, while place field parameters are mostly homogeneously distributed at all locations; **(a)** sketch of linear virtual environment: mouse receives reward at green zone when it stops long enough; when the mouse reaches the end of the track, it is teleported back to the beginning after some delay (red); **(b)** place field map in 3 example sessions; **(c)** density of place fields in all sessions, example sessions highlighted by arrows; **(d)** left: location-dependent place field density for all sessions (grey) and average (black) for one animal; reward zone (RW) and edge zone highlighted by colors according to (a); right: density averaged over general place field locations (PC: non-RW, non-edge locations) for different mice (grey) and average (black); **(e)** place field parameters and mutual information across different place field locations: a : chance of place field activation, A/A_0 : ratio of in-field to out-of-field firing rate, MI : mutual information, σ : place field width; colors, averaging as in (d); data from six mice, significant differences to PC location marked by * (p-value < 0.05), ** (p-value < 0.01)

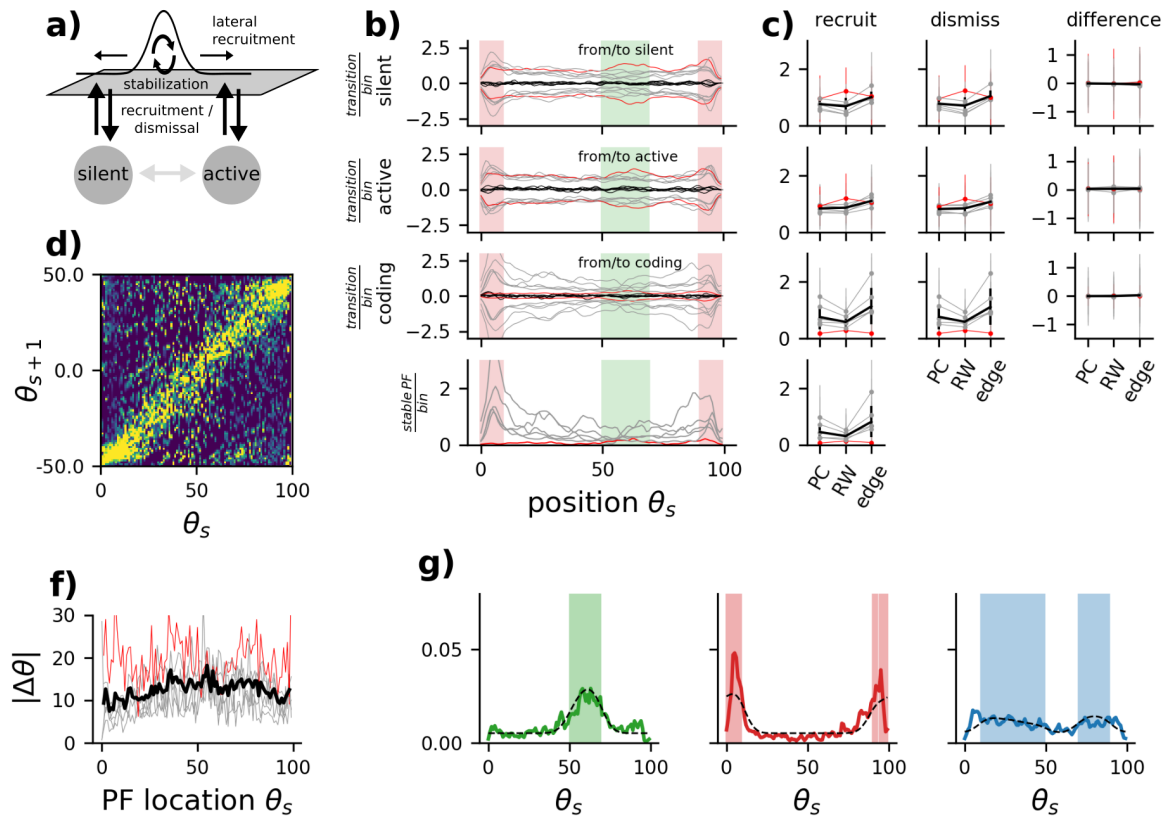


Figure 6.2: Dynamics under baseline conditions are location independent; **(a)** sketch of transitions: silent or active neurons can become place coding at some point (indicated by gaussian place field) and place fields can remain in place (stabilization), move to other positions (lateral recruitment) or turn silent or active, non-coding (dismissal); **(b)** location-wise transition rates of recruitment and dismissal for populations of silent, active (but non-coding) and coding neurons: recruitment (positive axis) meaning the neuron had the according state preceding a session in which it has a place field at location θ_s , dismissal (negative axis) meaning the neuron has an according state succeeding a session in which it had a place field at location θ_s ; location-wise difference displayed as red lines; each line is one mouse; bottom: stability of place fields depending on position, (place field occurs within $\pm 1.96\sigma_{\Delta\theta}$ in next session); data from mouse with divergent behavior (overrepresentation of RW location) highlighted in red; **(c)** regional averages of transition rates for place fields at reward position (RW), border position (edge) and other positions (PC); **(d)** probability of relocation from position θ_s (abscissa) to θ_{s+1} (ordinate); **(e)** average absolute place field shift distance per location; **(f)** distribution of place field relocation from RW (green), edge (red) and other (blue) locations, averaged over all sessions; **(g)** inferred distribution of place field relocation, based on empirical density at according locations, convolved with model Eq. 5.1

a different location, Fig. 6.2(a). I speak of recruitment, when referring to the previous state of a neuron before it has a place field, and of dismissal when talking about the session after, such that e.g. a neuron which has been silent and then starts to be active and have a place field is *recruited* from the silent population.

The different possible transitions can be calculated depending on the position of the place field and I find a shape similar to the place field density for all transitions. Furthermore, I can report huge similarities of the exact shape and rate across different mice for transitions from silent and active, but non-coding neurons, which are very homogeneously distributed over the whole range of non-border positions, Fig. 6.2(b,c), except for one mouse. The difference between recruitment and dismissal for all transitions is almost perfectly vanishing for all transitions, suggesting a balanced, independent mechanism for all three cases of recruitment/dismissal of silent, active non-coding and coding neurons, Fig. 6.2(b,c). Position-dependent stability varies most strongly across different animals and across locations, supporting findings of Ch. 4 that suggest the built-up of specific place-map shapes is driven by localized differences in stability and suggests that a similar mechanism might be in place during baseline conditions.

A previous work has identified place field shifts to be significantly reduced at rewarded positions¹¹⁸, contributing to the formation of overrepresentation. Consistent with the missing reward overrepresentation, I do not find a significant decrease of the shift distance $|\Delta\theta|$, Fig. 6.2(e). Indeed, all locations consistently show a preference to maintain their field within vicinity of the previous one, Fig. 6.2(d), suggesting that the location-dependent remapping is an effect present during map formation.

The model of place field shifts introduced in Ch. 5, Eq. 5.1 allows to predict the distribution of place field shifts from a sessions current population of place coding cells: it results from a convolution of the place map with the functional shape of Eq. 5.1. Again, separating the track into place cells coding for the borders (*edge*), for rewarded positions (*RW*) and others (*PC*), I can calculate the expected distribution of place fields from the according regions in the next session. Results show that all regions place field shifts are well reproduced by the prediction from the model with only minor underestimations of the observed data for stability at the edges, Fig. 6.2(f), suggesting no significant difference in the dynamics between any of the locations.

6.3 Individual contributions to observed population dynamics

The description of network level dynamics of Ch.5 has utilized averages over both temporal and individual behavior to obtain an overall understanding of what is going on in the network. Temporal, as well as individual behavior, however, can strongly deviate from the average description, already suggested by the wide spread of activation numbers N_α and N_β .

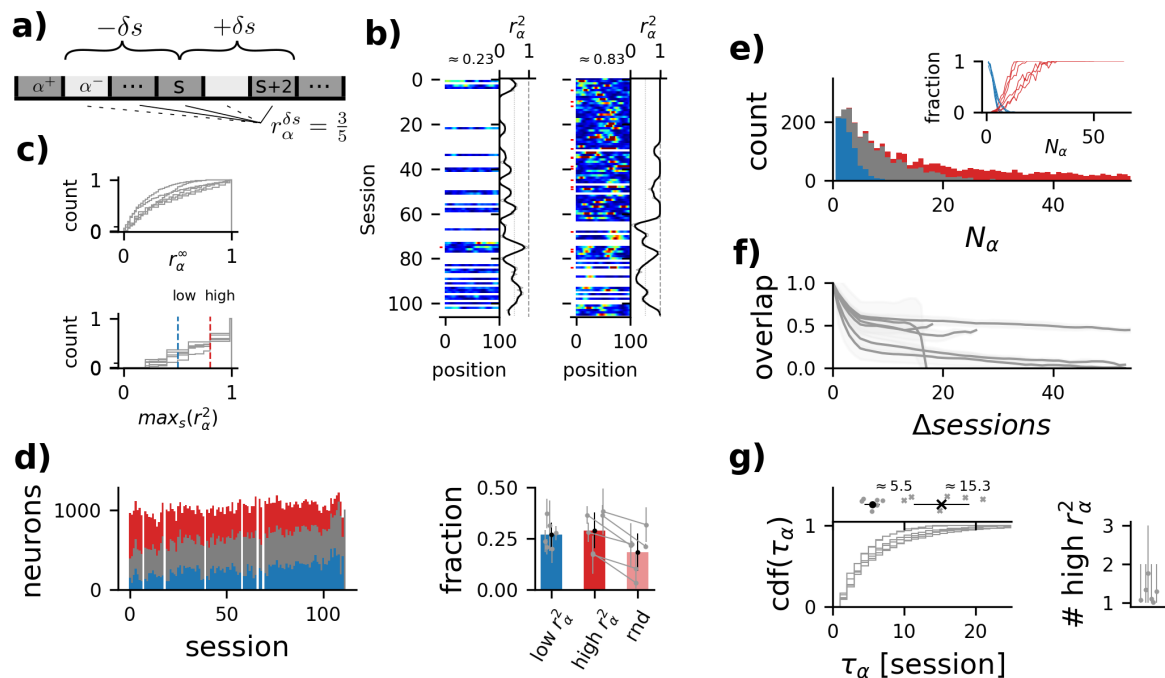


Figure 6.3: Consistency in network activity is carried by an ever changing population of neurons with temporary high reliability in activation; **(a)** sketch of the construction of $r_\alpha^{\delta s}$; **(b)** examples of neuron firingmaps over all sessions of a neuron with at most intermediate activation reliability $r_\alpha^{\delta s}$ (left) and a neuron with high contribution over a long time (right); local measure $r_\alpha^{\delta s}$ displayed next to firingmaps; raw value (grey, thin line), smoothed value (black) used for further computations; global measure r_α^∞ marked above firingmaps; **(c)** log-linear plot of distribution of r_α^∞ (top) and $\max_s(r_\alpha^{\delta s})$ (bottom) over all neurons, thresholds for determining low/high (blue/red) reliability marked by dashed lines; **(d)** left: number of active neurons per session, differentiated between highly reliable neurons ($r_\alpha^{\delta s}(s) > 0.8$, red), neurons with low reliability ($r_\alpha^{\delta s} < 0.5$, blue) and intermediate values (grey); right: average fraction of neurons with high/low reliability per session for all mice; compared with high reliability observed in randomized data (light red, R2, see Sec. 5.4.2), significant difference in 5 out of 6 mice (p-value < 0.001); **(e)** number of neurons with low/high/intermediate (blue/red/grey) maximum reliability, depending on overall activation N_α ; inset: fraction of neurons with low/high (blue/red) maximum reliability for different mice; **(f)** overlap of population of neurons with high reliability after Δs sessions; **(g)** lengths of periods of high reliability (for neurons with $\max_s(r_\alpha^{\delta s}(s)) > 0.8$); bottom: cumulative distribution of period lengths τ_α ; top: animal-average (dot) and 95 percentile of duration of consecutive periods with high reliability, grey: single animal data, black: animal average with SD; right: number of periods with high reliability per neuron (mean \pm SD); data for six mice

6.3.1 Dynamics of neuron activity on an individual level

Sec. 5.4.2 has shown, how the population of active neurons is subject to a gradual turnover and how it has an increased probability of a neuron to be active after an active session. To further investigate the role of individual neurons in this process, I define a stability parameter for neuron activity:

$$(6.1) \quad r_{\alpha}^{\delta s}(s) = \frac{1}{N_s^{\delta s}} \sum_{s' \in \delta S_{\cup}} \alpha_{s'}^+$$

with $N_s^{\delta s}$ the number of sessions within the set of sessions δS_{\cup} and δS_{\cup} the set of sessions within range $\pm \delta s$ of session s . The sum increments by 1, whenever session s' is active ($\alpha_{s'}^+$). For each session it computes the fraction of active sessions within range $\pm \delta s$ (see Fig. 6.3(a-b)) and is thus a measure for how reliably the neuron contributes to the network activity in a certain interval. For $\delta s \rightarrow \infty$ it is a global, session-independent measure r_{α}^{∞} and represents the normalized N_{α} .

The distribution of r_{α}^{∞} , as previously shown (Fig. 5.5(c)) monotonously decays towards larger values. However, the distribution of the local measure $r_{\alpha}^{\delta s}$ (here exemplary for $\delta s = 2$) shows how a large fraction of neurons has some time interval, in which it reliably contributes to neuron activity ($\max(r_{\alpha}^2) = 1$), while only few neurons follow apparently random activation over the whole dataset ($\max(r_{\alpha}^{\delta s}) < 0.5$), Fig. 6.3(c).

A fraction of $27 \pm 6\%$ (mean + SD) neurons display low reliability ($r_{\alpha}^{\delta s} < 0.5$) in surrounding sessions, see Fig. 6.3(d) while $29 \pm 9\%$ (mean + SD) of active neurons per session reliably contribute to the neural code ($r_{\alpha}^{\delta s} \leq 0.8$). The latter fraction is larger than expected from a randomized data set with randomly redistributed active sessions in each neuron (data set R2, see Sec. 5.4.2), with statistical significance (p-value < 0.001) in 5 out of 6 mice, Fig. 6.3(d), right.

Across different animals, I can not find a consistent trend of changes to these contributions over the course of sessions. Unsurprisingly, most neurons contributing reliably to the network activity at some point are those with a high number of active session N_{α} , Fig. 6.3(e). Just like the overall population of active neurons is subject to a gradual turnover, the population of neurons contributing with high reliability in some time interval is subject to a similar turnover, such that the population becomes increasingly dissimilar with increasing time, Fig. 6.3(f). There is a sharp drop at low Δs , somewhat masked by the overlap of intervals for calculating $r_{\alpha}^{\delta s}$ for adjacent sessions.

I find neurons to have a characteristic timescale over which they maintain a high contribution ($r_{\alpha}^{\delta s} \leq 0.8$). Intervals on average span 5.5 ± 1.1 sessions (mean \pm SD across mice), which corresponds to a range of 9.5 ± 1.1 sessions over which the neuron is active (~ 1 week), as high contributions require sessions in an interval ± 2 sessions to also be active. The distribution displays a wide tail into longer intervals (95-percentile: 19.1 ± 3.9 sessions, value padded with 4 sessions), Fig. 6.3(g), left. Neurons that contribute reliably to the network activity at some point mostly do so once, only (Fig. 6.3(g), right), with the exception of one mouse, where the average number of reliable intervals approaches 2 (only neurons with at least one session with high $r_{\alpha}^{\delta s}$ are counted). This

mouse is the one with the longest measuring period, suggesting the cause to be either a higher probability of randomly encountering a constellation of highly reliable sessions, or that neurons tend to re-engage in the network activity again, after some time.

6.3.2 Turnover of place field expression in individual neurons

Similar to the activation of neurons, I define a measure for how reliably neurons display a place field globally, or locally within some time interval:

$$(6.2) \quad r_{\beta}^{\delta s}(s) = \frac{1}{N_{\alpha}^{\delta s}} \sum_{s' \in \delta S_{\cup}} \beta_{s'}^{+}$$

with $N_{\alpha}^{\delta s}$ the number of active sessions within the set of sessions δS_{\cup} , δS_{\cup} as above and the sum increments by 1, whenever session s' has a place field ($\beta_{s'}^{+}$). The measure is very similar to $r_{\alpha}^{\delta s}(s)$, Eq.6.1, but counts the fraction of place coding sessions within the active sessions of a neuron in session range $\pm \delta s$ around the current session. A value of 1 represents the case, where a neuron - whenever it is active within the session interval δS_{\cup} - has a place field, including also cases, where a neuron is only active in one or few sessions, see Fig. 6.4(a,b) for a sketch and examples. Accordingly, I find many neurons to have a maximum value of $\max(r_{\beta}^{\delta s}) = 1$ (Fig. 6.4(c), while these values often occur in sessions, where the neuron has low activation reliability $r_{\alpha} < 0.5$, Fig. 6.4(e), inset. To make sure the neuron actually shows place fields in several sessions when reporting high values of r_{β} , I condition the analysis additionally on the according sessions to have a value of $r_{\alpha} > 0.5$ when identifying high values $r_{\beta}^{\delta s}$. Contributions of high reliability then originate from neurons almost independent of their total activation number N_{α} , Fig. 6.4(e).

I find only few neurons per session to contribute a high place field reliability (high: $11.7 \pm 8.2\%$, low: $28.6 \pm 8.8\%$ of place cells, mean \pm SD) to the spatial code, see Fig. 6.4(d). These contributions are above what is expected from a random redistribution of place coding sessions (random data set R3, see Sec. 5.4.3) for 4 out of 6 mice (p-value < 0.01 , Mann-Whitney U test), Fig. 6.4(d), right.

The overlap of the neuron populations with high coding reliability between different sessions that are Δs apart drops with the length of the interval $\pm \delta s$, but remains almost stable thereafter over a large number of sessions, Fig. 6.4(f). This indicates a small population of neurons that are highly reliable in their place coding property. The few cases of high contributions on average last 2.5 ± 0.5 sessions, up to the 95 percentile 7.5 ± 2.1 sessions (mean \pm SD, each). Note, that these intervals do not include the padding of δs , as included in $r_{\alpha}^{\delta s}$, as adjacent session do not necessarily contribute a place field. Finally, barely any neuron shows a high contribution more than once in the data set (number of periods with high $r_{\beta}^{\delta s} \approx 1$), Fig. 6.4(g).

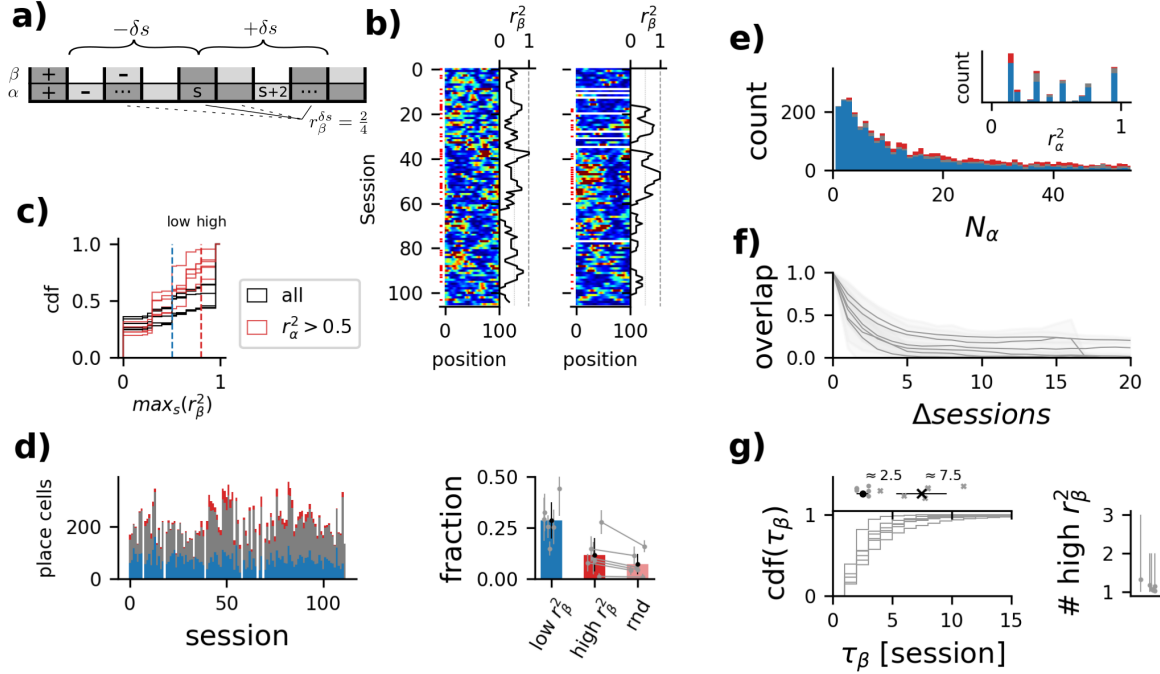


Figure 6.4: Only few neurons show high recurrence of coding property; **(a)** sketch showing how $r_\beta^{\delta s}$ is calculated: the total number of encountered place coding sessions β^+ (independent of place field position) divided by the total number of active sessions α^+ in the interval $\pm\delta s$; **(b)** two example neurons: firing maps of all sessions (left), values of $r_\beta^{\delta s}(s)$ (right, here and for remainder of figure $\delta s = 2$); **(c)** cumulative density function (cdf) of maximum values $r_\beta^{\delta s}(s)$ over all sessions (black); cdf for maximum values $r_\beta^{\delta s}(s)$ for sessions where $r_\beta^{\delta s}(s) > 0.5$ (red); data for all six mice shown; thresholds for low (blue) and high (red) $r_\beta^{\delta s}(s)$ highlighted by dashed lines; **(d)** left: number of place cells per session with low (blue), high (red) and intermediate (grey) $r_\beta^{\delta s}(s)$; right: average fraction of neurons with low/high $r_\beta^{\delta s}(s)$ of all place cells per session; data of different mice (grey dots) and average over mice (black) with errorbars indicating 1SD; **(e)** distribution of activation numbers N_α for neurons with low, intermediate and high maximum $r_\beta^{\delta s}(s)$ (colors as before); inset: distribution of activation reliability measures $r_\alpha^{\delta s}$ with fraction of low, intermediate and high $r_\beta^{\delta s}(s)$ highlighted (colors as before); **(f)** overlap of population with high coding reliability as function of session difference Δs for different mice; **(g)** lengths of periods of high coding reliability (for neurons with $\max_s(r_\beta^{\delta s}(s)) > 0.8$); bottom: cumulative distribution of period lengths τ_β ; top: animal-average (dot) and 95 percentile (cross) of duration of consecutive periods with high coding reliability, grey: single animal data, black: animal average with 1SD; right: number of periods with high coding reliability per neuron (mean \pm SD); data for six mice

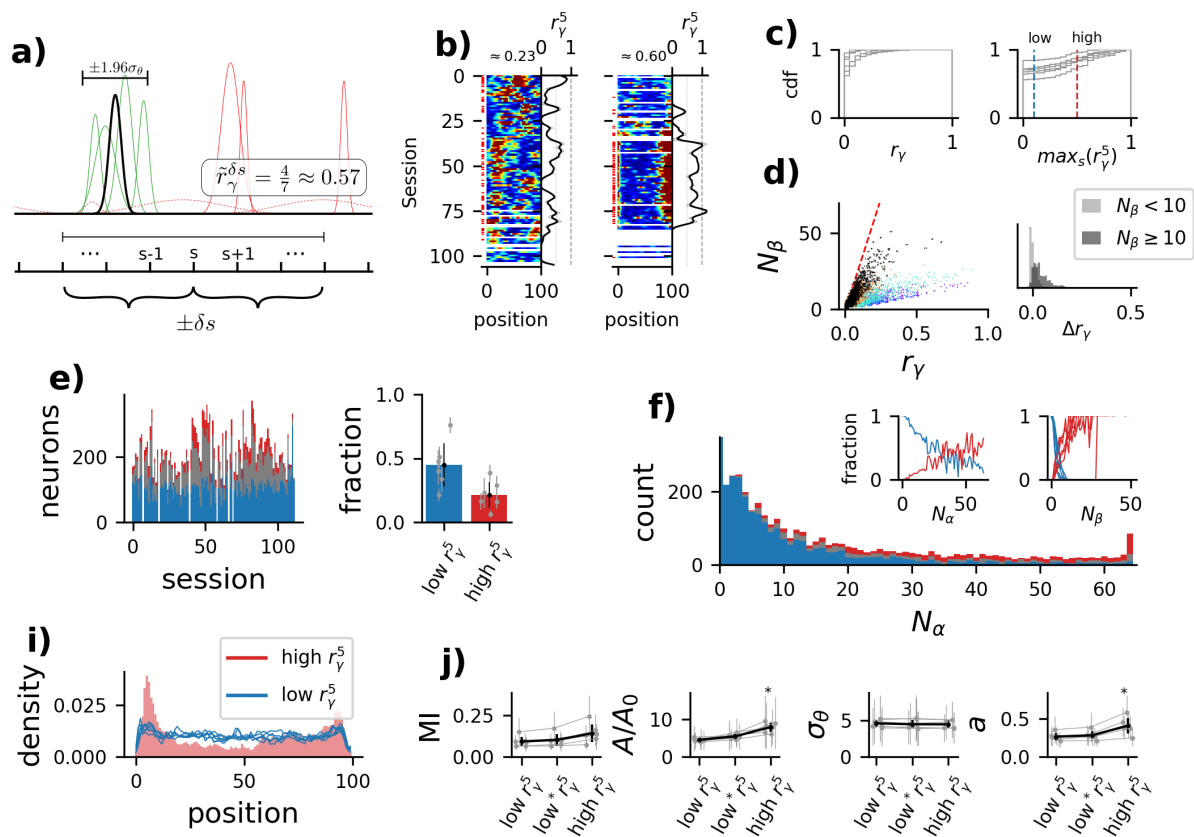


Figure 6.5: Reliable spatial coding is provided by a population with continuously distributed place field reliability; **(a)** sketch of parameter calculation of $r_\gamma^{\delta s}$: all fields detected in sessions $s \pm \delta s$ are registered (gaussian shapes, dashed lines indicate no place field) and the fraction of fields occurring within range $\pm 1.96\sigma_{\Delta\theta}$ (95% range, green, fields outside the range highlighted as red) of the field in session s (black) is calculated; if a neuron has multiple fields, all fields within the range of any of these are counted; **(b)** examples of two neurons with intermediate (left) and high (right) local (r_γ^5 , plotted next to firingmaps, $\delta s = 5$ used for rest of this figure) and global (r_γ indicated above according firingmaps) stability; **(c)** left: cumulative density function of r_γ for different mice; right: cumulative density function of $\max(r_\gamma^5)$ for different mice; **(d)** left: scatter plot of neuron-wise r_γ against number of place coding sessions N_β . Value of r_γ expected from random distribution of place fields indicated by red dashed line for one mouse (with black scatter points); right: histogram of differences between empirically observed r_γ and r_γ from randomly distributed place fields for black mouse data in left plot; distinguished between neurons with few place coding sessions ($N_\beta < 10$) and many place coding sessions ($N_\beta \geq 10$); **(e)** fraction of neurons with low/high r_γ^5 (blue/red) per session, conditioned on $r_\alpha^2 > 0.5$, as in Fig. 6.4; **(f)** distribution of activation numbers with fraction of neurons showing low/high $\max(r_\gamma^5)$ colored in blue/red; left inset: fraction of neurons with low/high $\max(r_\gamma^5)$ depending on activation number N_α ; right inset: fraction of neurons with low/high $\max(r_\gamma^5)$ depending on number of place coding sessions N_β ; **(i)** distribution of place field positions when neurons have high r_γ^5 (red, example of one mouse) and when they have low r_γ^5 (blue, all mice); **(j)** neuron and place field parameters when neurons have low r_γ^5 , when neurons with high $\max(r_\gamma^5)$ have a low r_γ^5 , and when neurons have a high r_γ^5 : mutual information (MI), in-field to out-of-field firing rate ratio A/A_0 , place field width σ_θ , place field reliability a for all mice (grey), with average (black). significant differences to other data marked by * (p-value < 0.05)

6.3.3 Place field stability in individual neurons

In a last step, I define a stability parameter capturing the stability of place fields of a neuron. It is defined as

$$(6.3) \quad r_\gamma = \lim_{\delta s \rightarrow \infty} \frac{1}{N_s} \sum_{s=1}^{N_s} \tilde{r}_\gamma^{\delta s}$$

$$(6.4) \quad \tilde{r}_\gamma^{\delta s} = \frac{1}{N_\theta^{\delta s}} \sum_{s' \in \delta S_\cup} \gamma_{s'}^+$$

with N_θ the number of place fields identified in sessions δS_\cup (multiple fields per session are possible) and δS_\cup as above. The sum in Eq.6.4 increments by one for each session in δS_\cup , in which a place field from session s is detected (place field center within $\pm 1.96\sigma_{\Delta\theta}$ of the place field in session s), see Fig. 6.5(a). When used as a local measure, $\tilde{r}_\gamma^{\delta s}$ is gaussian filtered (SD = 1 session) to obtain $r_\gamma^{\delta s}$.

Long-term contribution to place field stability is low for all neurons While $r_\gamma^{\delta s}$ is a localized measure, describing place field stability only within an interval $\pm\delta s$ around session s , r_γ is a global measure, which uses the limit of $\delta s \rightarrow \infty$ to characterize for each place field observed within the whole dataset, whether it is recovered in some session, or not. r_γ describes the average number of sessions, in which each of the place fields of the neuron occurs, with high values when most place fields are in the same position, and low values when they are scattered randomly across the virtual environment. As a session-averaged measure, r_γ depends on the number of sessions in which the neuron has a place field N_β , Fig. 6.5(a). An expectation value in the case of randomly distributed fields can be obtained, by assuming a uniform distribution of place fields over the whole track and at most one place field per session, thus $N_\theta^\infty = N_\beta - 1$:

$$(6.5) \quad r_\gamma^{rnd.} = \frac{1}{N_s} \sum_{s=1}^{N_s} \frac{1}{N_\beta - 1} N_\beta \frac{2 \cdot 1.96\sigma_{\Delta\theta}}{L} = \frac{1}{N_s} \frac{N_\beta^2}{N_\beta - 1} \frac{2 \cdot 1.96\sigma_{\Delta\theta}}{L},$$

with L the length of the track. The result for one mouse is displayed in Fig. 6.5(d), red line and shows how neurons with place fields in few sessions, only (low N_β , have their fields distributed mostly randomly (Fig. 6.5(d), inset). Neurons with higher N_β show a wider distribution of differences from random place field locations (Δr_γ), with a good portion of neurons above average chance of place field recurrence in the same position. However, only few neurons display a consistent contribution to the place code, indicated by the overall low value of r_γ for all neurons, Fig. 6.5(c), left.

Reliable contributions to stability happen on short time scale Previous paragraphs have used a small interval of sessions ($\delta s = 2$) to describe reliability in activity and expression of a place field. The stability of place fields described here, however, naturally depends on the neuron to express a place field, which has already been shown above to be quite unreliable for small

intervals. As discussed in Sec. 5.4.3, place fields on average reoccur with increased probability over longer periods of time. To investigate this behavior on the single neuron level and over intermediate time intervals, I chose $\delta s = 5$ for the following lines.

A fraction of $21.4 \pm 10.4\%$ (mean \pm SD) of place cells, strongly varying between different sessions, contributes high values of r_γ^5 to the network activity, while the majority of place cells ($44.9 \pm 17.0\%$, mean \pm SD) contributes low values, indicating that most place fields present in one session are not recovered within the interval $\delta s = 5$, Fig. 6.5(e). Not only neurons with high, but also some with low and intermediate activation numbers N_α at some point contribute a high reliability of place coding to the network dynamics, Fig. 6.5(f). The probability to obtain a high value of r_γ is large for neurons with large N_α and N_β (Fig. 6.5(f), insets), indicating that, when a neuron often contributes to the network activity, it is also highly probable to do so reliably at least over some period of time.

Time courses and characteristic time scales are highly dependent on the previously described dynamics of neuron activation and place field expression and are therefore neglected for the parameter $r_\gamma^{\delta s}$.

In agreement with previously discussed location-dependent transition probabilities, contributions of high r_γ^5 are usually provided by neurons with place fields at specifically stable positions, such as the edges of the VR, while place cells offering low contributions of $r_\gamma^{\delta s}$ have place fields homogeneously distributed over the complete environment, Fig. 6.5(i). High contributions are further contributed by neurons with an increased ratio of in-field to baseline firing rate, as well as increased reliability, Fig. 6.5(j), while place field width and mutual information σ shows no significant deviations. Neurons with high place field reliability at some point do not show any different behavior from other neurons in sessions where their reliability r_γ^5 is low.

6.4 Discussion

While the population- and session averaged dynamics described in Ch. 5 provide a description of how different levels of turnover are present in the spatial map of the hippocampus, this chapter explores the contributions of individual neurons in these dynamics.

Missing overrepresentation of salience The special role of neurons coding for rewarded locations or locations with landmarks described in Ch. 4 can not be detected in the analyzed dataset. This apparent contradiction is discussed in Ch.7, where the results of both chapters are shown to fit into currently ongoing research into effects of overrepresentations. In short, the different setup of the experiments with a visually marked reward location and an exhaustive time of familiarization with the environment is reported to lead to an abolishment of the overrepresentation at the reward location^{112,118,120}.

Interestingly, one mouse still displays transition rates modulated by the reward position and also shows significantly lower rates of lateral recruitment (relocating place fields) and place field

stability (data highlighted in red in Fig. 6.2(b)) than the other mice. This might indicate that the mouse is in a different behavioral state from the others, continuously recreating a unique spatial representation upon exposure to the environment by recruiting previously non-coding neurons. For some reason, in contrast to the other animals, this mouse might experience the environment as a novel experience in each session. Previous studies have suggested that rewards are preferentially expressed upon learning experiences¹²⁰, providing a possible explanation to the special behavior of this mouse, with respect to the other five individuals.

Individual contributions in line with network averaged dynamics The introduction of three parameters for the characterization of reliability in individual activation, place cell expression and place cell location has allowed me to describe how individual neurons contribute to the dynamics in general, but more importantly, how their roles change throughout the time of the experiment. Consistently across all mice I find that a large number of neurons takes on an important role to the dynamics at least once in their lifetime. About half of the neurons have a period in which they are reliably activated across several consecutive sessions, with an average period length of ~ 9 sessions, corresponding to ~ 1 week. In contrast, only few neurons provide a reliable contribution by expression place fields over several sessions ($\sim 11.7 \pm 8.2\%$ per session, corresponding to 10-50 neurons per session). The few neurons that do express high reliability, often do so over short periods of time, only, with an average of ~ 2.5 sessions. This further supports the hypothesis that overall dynamics are lead by a slow turnover in the active population, with a very volatile population of place coding neurons. Ch. 7 further elaborates on how this fits into current findings on hippocampal functional and synaptic turnover.

Chapter Seven

Summary and discussion

In this thesis, I applied and extended state of the art methods for the analysis of neuron activity from imaging data to elucidate mechanisms underlying the generation and sustaining of the hippocampal spatial representation in the deep sublayer of CA1. An extensive pipeline for analysis presented in Ch. 2 allows reliably obtaining data from and track identities of thousands of neurons over arbitrarily long sequences of recording sessions and characterize their spatial coding properties, as presented in Ch. 3. These tools of analysis were successfully applied in a publication on spatial map formation and the development of overrepresentation of visually and functionally salient regions therein, see Ch. 4⁶². Furthermore they provide the basis for the analysis of baseline dynamics in the hippocampal representation of space, both on a network (Ch. 5) and an individual neuron level (Ch.6).

7.1 Data / Methods

The underlying data used in Ch. 5 & 6 has been obtained from six mice in several dozen recording sessions per animal, up to over 100 in 3 individuals. It is worth noting that all of the data has been obtained long before the start of this work and the experiments were not devised with the presented study in mind. Still, the data provides a rich set of recordings, suitable for the needs of this study.

The processed data was obtained from raw intensity images by a series of steps: neuron extraction, demixing and deconvolution has been performed by the established tool CaImAn¹³⁵ using a CNMF approach, see Sec. 2.4 and App. A. The registration of neuron identities has been adapted from previous works^{17,135} to allow applying it to the underlying large-scale data set. I introduced a novel method to test and improve completeness of the detected neuron data in each session. Its advantages as well as ways for further development and assessment are outlined in Sec. 3.5.

The registration procedure provides a set of unique neurons, which are active in some and silent (and therefore undetected) in other sessions. The fate of undetected neurons was further evaluated by applying the neuron detection procedure on assumed anatomical locations of

neurons, resulting in an increase of $\sim 20 - 30\%$ in active neuron numbers per session, while confirming silence for other neurons. Resulting ratios of $\sim 60 - 80\%$ of silent neurons per session are well in agreement with earlier reported works^{52,164}. The large population of additionally detected neurons suggests that common imaging analyses might miss out on a significant fraction of neurons, and could profit from similar steps as proposed in this work for the inference of according neuron components. Especially when tracking the behavior of individual neurons a complete detection algorithm is necessary to not misinterpret apparent silent periods, while the neuron activity in fact was just missed out on a first detection.

I developed an algorithm for place field detection based on a previous work on tuning curve identification using a Hierarchical Bayesian model¹³⁶ solved by nested sampling¹³⁷, an approach commonly applied in astrophysics, exceeding most commonly applied place field detection algorithms in complexity, but allowing for a well-controlled detection characterization of place field parameters, further elaborated in Sec. 2.8. I evaluated the new approach in Sec. 3.2 using a large set of artificially generated neuron activity and found that it accurately recovers place field parameters such as position and place field width, while showing a low rate of false positive discoveries. Resulting fractions of $\sim 20 - 40\%$ place cells per session are in line with previously reported fractions^{15,18}. The presented method is not only able to accurately obtain place field parameters, but also to identify multiple place fields per neuron (which were found to increasingly occur along the proximo-distal axis of CA1⁴⁴) and characterize whether a field is active or inactive on a trial-by-trial base.

In summary, the data is processed by a number of well assessed tools that provide results in line with earlier published reports on CA1 place cells. The methods result in a set of several hundreds, up to thousands of neurons per mouse and session which can be tracked over an arbitrary number of session.

7.2 Dynamics of the spatial representations in CA1

This work elucidates different processes of observed turnover, both on a network level and on an individual neuron level. Unlike many other studies that describe the immediate response of CA1 to new learning experiences^{57,88,107}, the later chapters 5, Ch. 6 set a focus on animals after exhaustive familiarization with an environment and thus capture baseline dynamics. These dynamics are important for the interpretation of experiments characterizing dynamics in response to changes in the environment, e.g. when turnover is interpreted as a reaction of the network to some stimulus, while in fact it is part of the ever ongoing dynamics of synaptic turnover^{12,13}.

7.2.1 Place map formation and the role of salience

Visual salience supports animals in the navigation through an environment, which is often motivated by some reward. It comes as no surprise that the hippocampal spatial map is reported

to be shaped by such salient cues³, often expressed by an overrepresentation in place field density of the according location^{62,112,116–118,120,121}.

Ch. 4 presents the first study in which the underlying mechanism to the formation of such overrepresentation has been investigated with a focus on session-to-session transitions of neuron states, reporting rates of recruitment, dismissal and stability of place fields depending on the place field location⁶². Overrepresentations of landmarks and rewards are found to be formed by a localized increase in place field stabilization and form over distinct timescales. An experiment using Shank2-KnockOut mice shows how the landmark overrepresentation is completely abolished in these animals, while the reward location is still encoded by an increased place field density, suggesting distinct mechanisms for reward and landmark overrepresentation.

In a variant of the experiment where reward delivery is determined by chance, we report how responses to the reward location can be partitioned into subpopulations of neurons coding for the expectation of a reward, the reception thereof and lastly, the location itself. An overrepresentation is thus not necessarily a more dense coverage of the environment in the spatial map, but rather a superposition of various stimuli coded for in the hippocampus.

While the overrepresentation and its formation is well described in Ch. 4, the data analyzed for the baseline dynamics of the spatial map were found to display no such overrepresentation. Such inconsistent reports of reward overrepresentation have been identified and discussed before¹²⁰, suggesting two preconditions for an increased place field density at rewarded locations: (1) rewards are not additionally marked by visual salience¹¹⁷ and (2) the reward location is frequently shifted to other parts of the environment¹¹⁸. Neither of these conditions is fulfilled in the experiments underlying the data in Ch. 5 and 6, where the reward location remains within the same location, marked by green wall patterns. Additionally, the cognitive demands of navigating this highly familiar environment are supposedly low, which has further been shown to negatively influence place coding in general⁷³, and also coding for reward locations¹¹⁷. Accordingly, only one of 6 mice shows a slightly increased density at the reward location. In contrast, edges of the linear track, representing the start and end of a trial, are significantly and consistently overrepresented and show differences in place field parameters, transition rates and stability.

It would be interesting to see if a study similar to this one but with different reward conditions could show a lasting overrepresentation of the according position over an extended period of time, and whether reward regions and visually salient regions (e.g. landmarks) show qualitative differences under baseline conditions.

7.2.2 Turnover dynamics on different levels suggests influence from different regions

This work introduces a new, detailed characterization of commonly observed turnover dynamics of the spatial map in CA1^{15,16} and describes how it can be dissected into three different processes

with unique dynamics, each: (1) slow changes in the population of active neurons, (2) activation of place coding properties that on average occur randomly, and (3) place field movement characterized by fluctuations around a reference point, interrupted by random relocation.

I describe how the overall network dynamics are dominated by a slow, gradual change of the active population, ranging over the whole dataset of each mice with up to 70 sessions under baseline condition and taken over a period of more than 2 months. In contrast, the population of place cells appears to follow an almost random process, such that recurrence is only mildly time-dependent and mostly explained by the turnover of the active population. Finally, I find place field movement to be explained by a binary model in which a field either remains stable (remaining within the same position while allowing for some minor fluctuations) or randomly relocates to any position on the linear track. The probability of a place field to reoccur in the same position is significantly increased over a period of up to several weeks.

Contributions to observed dynamics The two very different dynamics observed in the population of active and the population of place coding neurons can originate from two different mechanisms: the frequency at which a neuron adapts the according status (active, coding) over all sessions and the arrangement at which it adapts these. I find the distribution of activation numbers across neurons to be highly heterogeneous, with some neurons active across the whole data set and some active in few sessions only. Place coding sessions do not show such a broad distribution, but are still significantly different from random activation of place fields.

Individual contributions to the network activity are clustered in time, such that neurons having some role in the network in one session are more probable to be active in the next session, causing the gradual, time-dependent turnover. A large number of individual neurons displays at least one such period of highly reliable contribution to the network, with an average duration of 9-10 sessions (~ 1 week) and some ranging up to several dozens of sessions. Works have suggested an important functional implication of such structured turnover: an ever changing population of active neurons allows for a unique identification of experienced events with times, effectively providing memories with time stamps^{16,174}. Studies reporting global remapping with statistically independent active samples in different environments^{53,54,89,178} and gradual turnover in unchanging environments^{15,16} suggest that a structured activation of a subpopulation of neurons plays a major role in the encoding of environments. It might be guided by Hebbian effects, both during experience and during sequence-reactivation during short-wave-ripple (SWR) sleep, which is reported to strengthen connections from CA3 to CA1¹⁰. Other studies suggest that silent neurons within the hippocampal spatial code are actively inhibited, supported by the inability to make silent neurons elicit action potentials, e.g. using antidromic stimulation⁵².

In contrast to the highly structured neuron activity, the appearance of significant place fields within the population of active neurons can almost be reproduced by a random redistribution of place coding sessions in each neuron. The probability to encounter a place field in a session after a neuron has expressed a place field is only marginally increased against the expectation

from a random arrangement of place coding sessions. Furthermore, the number of neurons with highly reliable place field recurrence per session is low, though significantly higher than expected from randomized data in most mice. A few neurons display temporally highly reliable place field recurrence, with average period lengths of ~ 2 -4 sessions (1-3 days). It is possible that the relatively strict identification of place fields could have some impact on these statistics by not detecting a few place fields. However, I have shown that in general only less reliable or less pronounced place fields remain undetected.

7.2.3 Origins of stability and turnover

These findings highlight how two timescales of distinct mechanisms coexist in CA1: On the one hand I can observe a slow, gradual turnover in the active population as well as maintenance of place field locations over up to several weeks. On the other hand, there is almost random activation of place field properties, with low reliability on short timescales. While this study does not allow for any further conclusions on the origin of such distinct mechanisms, there are two probable sources, external or internal to CA1:

Traces from dynamics of upstream regions Firstly, CA1 is a region integrating input from different upstream regions within the hippocampal loop. Footprints of the dynamics of upstream regions might be inherited by CA1, while further processing of information within the region should be minimal, suggested by vanishing or no recurrent excitatory feedback within CA1³¹. Its pyramidal neurons are targeted by synapses from upstream regions of CA3 and MEC, which are reported to be responsible for temporally stable, spatially modulated activity^{20,56,86}. Synaptic connections from CA3 to CA1 have been shown to stabilize with experience¹⁰, suggesting how the stable code of CA3, supposedly promoted by autoassociative properties¹⁰²⁻¹⁰⁴, results in the observed long-term stability of CA1.

This stable code might be modulated by temporally highly dynamic input originating from CA2 and the lateral entorhinal cortex (LEC)^{27,86}, causing the high volatility in maintaining a place field. Furthermore, I observed short-term turnover dynamics with gradual changes to the spatial code within a single session, in line with earlier reports of CA1 turnover between tightly spaced sessions⁸⁶. These might be a footprint of LEC input, where changes in neuron activity have been reported on a seconds to minutes timescale²⁷.

The impact of temporally modulated code from CA2 on the functionality of CA1 is further supported by a study using vasopressin 1b receptor knock-out mice¹¹¹: The receptor is expressed at significantly higher levels in CA2 than in other hippocampal regions and was shown to modulate social behavior, which is believed to be one of the major functions of CA2¹⁰⁸. Knock-out mice could still perform tasks depending on object locations (*where*) and object identities (*what*), but were unable to perform tasks requiring episodic memory (*when*), in contrast to wild type mice¹¹¹. Suppression of CA2 functionality therefore appears to reduce the ability for remembering

the order of events. This is in line with the notion that CA2 causes increased dynamics in CA1, which were previously suggested to provide a mechanism for time stamping events¹⁶.

Further studies researching the impact of lesions or inhibition of upstream regions (MEC/LEC, CA3 and CA2) on CA1 dynamics could provide insight into respective influences on CA1 turnover.

Stable neuronal function despite synaptic turnover Secondly, observed dynamics of long-term stability and short-term turnover might arise from synaptic rewiring within CA1. The maintenance of network function underlying the turnover of synapses has received increased attention throughout the last years. Studies have reported how large fractions of synaptic spines in CA1 turn over within few days^{12,13} (~ 9 days¹² and 40% turnover over 4 days¹³), which is in line with the here reported short periods of high reliability in place coding. These studies suggest that - unlike in sensory cortical areas - CA1 does not have a fraction of long-term stable spines¹², even though certain parameters like spine size influence the stability¹³, in agreement with reports from other cortical areas¹⁷⁷. Hebbian effects of synaptic plasticity that increase synaptic strengths as CA1 neurons co-activate with their input from CA3¹⁰ might therefore contribute to the stabilization of spines. Such effects of long-term stabilization are necessary for the here observed stability of neuron activation and place field position over up to several weeks. While the deterministic nature of Hebbian plasticity would result in a subset of consistently stable neurons, it was previously reported that spine parameters such as size and shape only provide moderately accurate predictions of spine stability¹⁷⁹. These reports suggest a stochastic contribution to spine turnover and thus provide a possible explanation to the short- and long-term timescales identified in this work.

Further support for the interplay of deterministic and stochastic spine turnover is provided by the observation made in Ch. 5: intermittent sessions of non-coding (no place field identified) decrease the probability of a place field recurring in the same position, indicating how the consistent expression of a place field is required to reliably maintain it. The effect is slightly reduced in silent neurons, suggesting that non-specific activity of active, non-coding neurons destroys previously learned connectivity using spike time dependent plasticity (STDP) mechanisms. Such effects are not triggered in silent neurons, where post-synaptic spikes are missing. The effect of silence preserving neuronal function has been observed in a work where previously learned sequences in the high vocal center of zebra finches were preserved by targeted inhibition¹⁶⁹.

A number of theoretical studies have researched how network function or memories encoded in learned sequences can be maintained under synaptic turnover^{21,180–182}. One study has found how a backbone of strong, stable synapses can reproduce cortical data such that the network remains relatively robust to large subpopulations of volatile spines²¹. Even though CA1 is assumed not to have such a stable population¹², it is possible that subsets of temporarily stable connections could take on such a role over days or few weeks. A simulation of a balanced state network with an inhibitory and excitatory population has shown how network function can be sustained despite large fluctuations in excitatory-excitatory connectivity, as long as inhibitory

connectivity remains stable¹⁸¹. Long-term stability of place fields might therefore be encoded in stable inhibitory connections, whose turnover dynamics in CA1 have not yet been reported. Finally, a model of synaptic plasticity in which the number of synaptic connections per neuron pair increases spine stability finds that connections between neurons can be upheld reliably despite high rates of synaptic turnover¹⁸⁰. Summarized, these works highlight how network function can be preserved despite strong synaptic turnover. The detailed descriptions of individual and network level dynamics provided by this thesis could help identify whether these models agree with the observed dynamics on these two different levels.

7.2.4 Uniform population vs distinct populations

Within this work I have characterized both network and individual dynamics of neurons. While I am able to identify major heterogeneities between neurons, the data does not point towards distinct subpopulations of neurons, displaying different parameters or characteristics of stability throughout baseline dynamics. Parameters of long-term stability display a wide distribution across neurons and - depending on the animal - different numbers of neurons with high long-term reliability, which in principle can suffice to drive animal behavior¹⁸³. Most neurons, however, show high temporary stability only once throughout the period of recording with characteristic timescales as discussed above.

In contrast to a population described as a continuum, localized stabilization of place fields at salient locations during place map formation shown in Ch. 4 suggest specific roles of neurons, in line with a previous study reporting a population of reward coding neurons¹²⁰. It is a small population of neurons in CA1 which uniquely, but consistently responds to reward locations across different environments, which we could also identify in our study⁶². During baseline condition, only neurons coding for locations next to the beginning and end of the linear track could be identified to display specific parameters and increased stability over other locations, while differences in place field parameters remain largely low and insignificant between different locations.

7.3 Overall effect / Future works

This work has highlighted how hippocampal dynamics in the representation of space can be described by different levels of turnover. It provides a detailed description of how stability over long timescales in the population of neurons contributes to the spatial code. Robust locations of place fields are complemented by short term fluctuations that almost randomly activate and inactivate a neuron's place field. These effects could originate from two distinct sources: according dynamics in upstream regions, which are combined in CA1 to express the observed activity, or synaptic dynamics which are known to show a high degree of turnover within CA1. Most probably, the dynamics arise from an interplay of these two effects, which is supported by the inability to

identify a unique time axis along which the network dynamics appear, Sec. 5.5. It is possible that, while mechanisms of synaptic turnover occur with time measured in minutes and hours, the connectivity to other regions, as well as their dynamics, rather occur along a time axis measured in the amount of experience or nights slept. I have outlined in Sec. 5.6, how future works can provide useful insight to disentangling the effects of time on the hippocampal dynamics.

Future lesion or inhibition studies of upstream regions can test the impact on animal behavior and performance in the navigation of environments over long times. In behavioral studies it has been reported that CA2 inhibition has no impact on spatial learning¹⁰⁸, while it shows a striking impact on temporal coding¹¹¹. According imaging studies of neuronal activity can provide insight into how the suppressed input from CA2 influences CA1 dynamics over extended timescales and similar studies could be performed with input from CA3 as well as the input regions of the entorhinal cortex (layer 3 of LEC, MEC). If the hypothesis of footprints from upstream regions in CA1 is correct, suppression of the more dynamic, non-spatial regions CA2 and LEC should result in less temporal variability of place coding and more individual stability of neurons expressing place fields. Moreover, a suppression of CA3 or MEC should result in decreased long-term stability, increased individual place field movement and erased long-term effects in the gradual turnover of the active population and the maintenance of place field locations.

Studies on synaptic turnover have provided valuable insight into timescales and the volatility of the underlying medium for neuronal dynamics. Further studies, drawing a more detailed picture of synaptic connectivity and their long-term stability, could provide further insight into the origin of the different dynamics: Is input from temporally more reliable regions (CA3, MEC) also transmitted via more stable spines than input from CA2 or LEC? Do synapses from inhibitory interneurons display increased stability, possibly stabilizing the overall network dynamics¹⁸¹? Such studies can exploit the different locations at which synapses from different regions reach the dendritic tree of CA1 neurons³¹ to identify according spines. Furthermore, differences in connectivity along the CA1 proximodistal axis have been reported, with the more proximal region preferentially receiving input from MEC and the more distal region receiving input from LEC, resulting in differences in the reliability of the spatial code⁴⁴. It is possible that differences in place field stability of individual neurons might arise from similar differences in individual connectivity, which requires further testing.

Finally, the detailed description of network and individual dynamics provided in this work allows for the development of a model of network dynamics in CA1, similar to the kinetic model previously used to describe spine turnover¹². Such a model could provide invaluable insight into the behavior of a network under manipulation of certain parameters, possibly explaining quantitative animal differences in the observed dynamics.

Appendix One

Constrained non-negative matrix factorization (CNMF)

Neuron detection algorithm (CaImAn)

Calcium imaging experiments provide a visual readout of neuron activity, which is mediated by calcium indicators and captured by fluctuating pixel intensities. Obtained recordings - sequences of frames, each with an image dimension of $X \times Y$ pixels can be used to obtain information on single neurons. It comes with the challenge of identifying the locations and shapes of neurons (spatial footprints or regions of interest, ROI), which is a topic of ongoing research.

Calcium imaging recordings provide information in the form of a high-dimensional space ($x \times y = d$ pixels) of intensity values fluctuating over the measurement time T . In general, neuron detection algorithms aim to approximate such observations by a lower-dimensional, biologically meaningful representation. On activation, neurons show a relatively homogeneous intensity response and can be observed on top of some background, representing both, anatomical landmarks or neuropil and effects of the microscope. CaImAn¹³⁵ decomposes the recorded image $\mathbf{Y}(t) \in \mathcal{R}^{d \times T}$ into such two types of objects, with K neurons represented by their locally constrained spatial footprints $\mathbf{A} \in \mathcal{R}^{d \times K}$ and temporal dynamics $\mathbf{C} \in \mathcal{R}^{K \times T}$ and K_b background components, represented by the corresponding footprints $\mathbf{b} \in \mathcal{R}^{d \times K_b}$ and temporal dynamics $\mathbf{f} \in \mathcal{R}^{K_b \times T}$. Each objects' footprint describes how much the pixels of the FoV are manipulated by its temporal dynamics. Photon shot noise and further stochastic fluctuations e.g. of ion density, introduce noise which can be described by some random, pixelwise independent process $\varepsilon(t)$ and provides the last building block of a matrix factorization framework of the recorded image Y :

$$(A.1) \quad \mathbf{Y}(t) = \mathbf{A} \cdot \mathbf{C} + \mathbf{b} \cdot \mathbf{f} + \varepsilon(t)$$

This framework has been introduced to calcium imaging analysis with the requirement of positive matrix entries, only, restricting the reconstructed intensities to positive values (non-negative matrix factorization, NMF¹⁴⁰). Additionally, the constrained non-negative matrix factorization (CNMF)-method applied by CaImAn incorporates a number of further constraints:

- Residuals of fitting the model to the original recordings have to obey the statistics of gaussian white noise with pixelwise standard deviation σ estimated from the data
- Neuron shapes should be sparse but connected, accounting for locally constrained somata
- Calcium activity $c(t)$ should obey calcium indicator dynamics with underlying neuronal activity $s(t)$ and be sparse, allowing to effectively demix the activity of overlapping neurons by avoiding overfitting

After an initialization process in which first guesses of neuron footprints are obtained, the algorithm iteratively refines temporal and spatial components under the given constraints to give a final estimation of the set of neuron candidates¹⁴¹, ready to be further processed by deconvolution and evaluation. The individual steps are explained in brief detail below.

Measurement noise statistics can be obtained directly from power spectral density (PSD) of the recorded image. It contains contributions from the analyzed, biological processes, and calcium indicator dynamics, but also from the noise present in the data. While the former two processes enter Eq. A.1 via underlying models to neuron (**A,C**) and background (**b,f**) components and are assumed to have vanishing high-frequency components, the latter part is captured by the gaussian random process ε , assumed to be white, thus with a flat PSD over all frequencies. The high frequency tail - indeed found to be flat - therefore offers an estimate of the noise PSD by interpolating the flat power spectrum towards lower frequencies, Fig. A.1(b). A complete characterization of white noise is provided by its variance, which - according to Parseval's theorem - can be obtained from integrating the PSD over all frequencies:

$$(A.2) \quad \sigma^2 = \frac{1}{\pi} \int_0^\infty S_x(\omega) d\omega \quad , \quad S_{xx}(\omega) = \lim_{T \rightarrow \infty} \mathbb{E} [|\hat{x}(\omega)|^2] \quad ,$$

with \hat{x} the Fourier transformed signal, ω the frequency, $\mathbb{E}[\dots]$ the expectation value and S_{xx} the PSD.

Calcium indicator dynamics introduce a characteristic decay pattern with a fixed timescale into the temporal components of neurons. In CaImAn calcium traces $c(t)$ and resulting intensity recordings $y(t)$ are captured by an autoregressive process (AR) of low order p (here = 1),^{141,148}:

$$(A.3) \quad c(t) = \sum_{k=1}^p \gamma_k c(t-k) + s(t)$$

$$(A.4) \quad y(t) = \alpha(c(t) + b) + \varepsilon(t),$$

with underlying neuronal activity $s(t)$, α a non-negative scalar, b the baseline concentration and ε as above. The calcium indicator time constants γ_k define rise- and decay- time and thereby the characteristic intensity response to neuronal activity, allowing to obtain $s(t)$ by deconvolution of $c(t)$. Usually, $s(t)$ is reported as the spike count at time t , containing integer values and translating

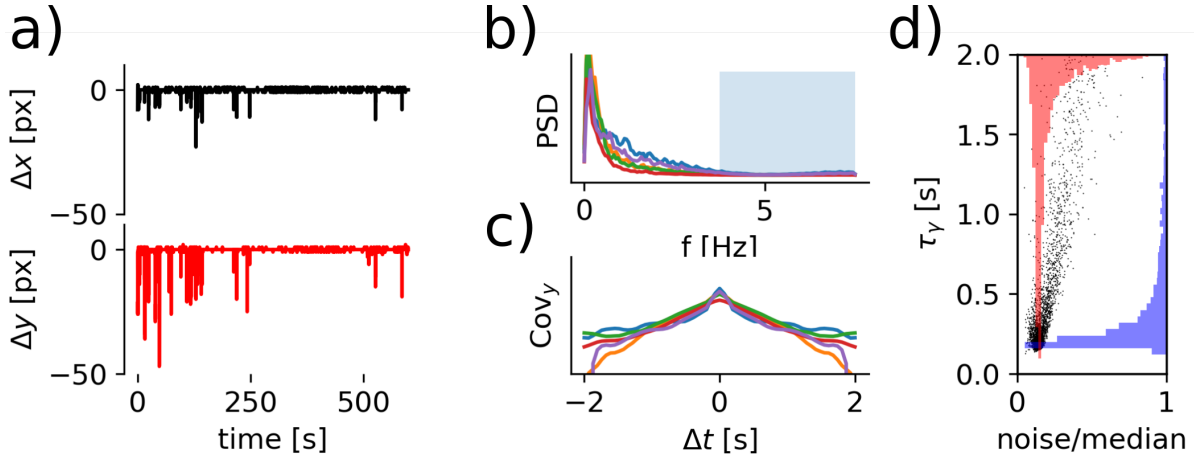


Figure A.1: Preparation and parameters for the CNMF method; **(a)** detected x - and y -motion over several sessions; **(b)** power spectral densities of five neurons, high frequency region (blue area) used for noise level estimation; **(c)** autocovariance of five neurons (logscale in y -axis) used for estimation of time constant γ ; **(d)** estimated time constants $\tau_\gamma = -(f \cdot \log(\gamma))^{-1}$ against noise level of Ca^{2+} -trace (fraction of median Ca^{2+} -level) for each neuron (black dots). Histograms show distributions of respective values

to calcium values via the nonnegative scalar value α . Following¹⁸⁴, the integer requirement is dropped and $s(t)$ can take on any nonnegative value.

The time constants can be estimated from the autocovariance through

$$(A.5) \quad \text{Cov}_c(\tau) = \sum_{k=1}^p \gamma_k \text{Cov}_c(\tau - k)$$

$$(A.6) \quad \text{Cov}_y(\tau) = \sum_{k=1}^p \gamma_k \text{Cov}_y(\tau - k) - \begin{cases} \sigma^2 \gamma_k & \text{if } 1 \leq \tau \leq p \\ 0 & \text{if } \tau \geq p \end{cases},$$

where Eq. A.6 is obtained from inserting the autocovariance of Eq. A.4, $\text{Cov}_y(\tau) = \alpha^2 \text{Cov}_c(\tau) + \sigma^2 \delta(\tau)$ into Eq. A.5. A fit to empirically obtained autocovariances, Fig. A.1(c), yields the time constant(s) γ_k .

Initialization of the neuron detection algorithm can be achieved by a number of different methods to yield a first guess of neuron positions and their temporal traces. I used an approach termed greedyROI¹⁴¹, which exploits that neuron cell bodies can be approximated by a 2D-gaussian distribution and show correlated activation of pixels belonging to said neuron.

The recorded data is filtered by a gaussian kernel with the expected size of a cell body (provided by the user), highlighting regions of locally correlated activity. An iteration of the NMF¹⁴⁷ is applied to the neighborhood around the global maximum of the filtered data and provides the initial guess of spatial and temporal component of this neuron. The contributions from this neuron are removed from the data matrix and the procedure is repeated on the

residual matrix, continuing until a predetermined number of components is detected. Finally, the background components b and f are initialized from the residual. It is also possible to provide estimations of spatial footprints directly.

Spatiotemporal updates including deconvolution refine the initial estimates. Both, spatial- and temporal- updates use Eq. A.1 to solve a minimization problem with the constraints

$$(A.7) \quad \|\mathbf{Y}(i, :) - \mathbf{A}(i, :)\mathbf{C} - \mathbf{b}(i, :)\mathbf{f}\| \leq \sigma\sqrt{T}, \quad i = 1 \dots d.$$

During the spatial update, \mathbf{C} and \mathbf{f} of the previous iteration are held fixed, while solving the convex problem of finding sparse spatial footprints via

$$(A.8) \quad \min_{a_j, b} \|a_j\|_1, \quad \text{subject to: } a_j, \mathbf{b} \geq \mathbf{0}, \quad j = 1 \dots K$$

and subject to Eq. A.7. It uses a_j as the j th spatial footprint, thus the j th column of \mathbf{A} . Throughout the update, spatial footprints a_j are confined to take on non-zero values in a neighborhood of their previous estimates, only.

The temporal update holds \mathbf{A} and \mathbf{b} of the previous iteration fixed, and actively demixes calcium signals of overlapping footprints by solving another convex problem of finding

$$(A.9) \quad \min_{c_1, c_2, \dots, c_K, \mathbf{f}} \sum_{j=1}^K \mathbf{1}^T \mathbf{G} c_j, \quad \text{subject to: } \mathbf{G} c_j \geq \mathbf{0}, \quad j = 1 \dots K$$

and subject to Eq. A.7. It uses c_j as the temporal trace of the j th component, thus the j th row of \mathbf{C} and introduces the matrix \mathbf{G} , which couples the calcium activity to the underlying neuronal activity via Eq. A.3:

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ -\gamma_1 & 1 & 0 & \dots & 0 \\ -\gamma_2 & -\gamma_1 & 1 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \dots & -\gamma_2 & -\gamma_1 & 1 \end{pmatrix}$$

The deconvolved activity s_j is obtained with the constraint of sparse activity:

$$\min_{c_j, \mathbf{b}, s_j} \sum_{t_1}^T s_j(t), \quad \text{subject to: } \begin{cases} s_j \geq \mathbf{0}, s_j = \mathbf{G} c_j \\ \|y - c_j - b\| \leq \sigma\sqrt{T} \end{cases}, \quad j = 1 \dots K$$

It is implemented via the Online Active Set method to Infer Spikes (OASIS¹⁴⁸), which uses a variant of the pool adjacent violators algorithm (PAVA) adapted to monotonously decaying patterns as observed after spikes for simultaneous denoising and deconvolution of calcium traces.

The algorithm alternates between temporal and spatial updates and includes several checks: If any component has a vanishing contribution to the overall fit, it is eliminated by absorbing it into the background. If two components show significant spatial and temporal correlation, they are merged into a single component.

Online processing

CaImAn provides the option to perform these steps in an offline, batch-framework, or in an online framework, OnACID, in which the algorithm iterates through the data along the time axis. In this work I have made use of both frameworks and, while the basic methods are similar for both and were explained for the batch-processing in the previous pages, the following lines will outline the adaptation made in the online-algorithm.

While the batch processing of CaImAn relies heavily on a first estimate, as it does not allow to add neurons in a later process, the Online Analysis of Calcium Imaging Data-method (OnACID¹⁶³) provides a way to dynamically add components and is performant enough, to process incoming data in real time.

Methods for temporal and spatial updates are similar to the ones described above for batch processing, but with few changes necessary, as it iterates through data frames one at a time. While denoising and deconvolution are handled by the OASIS algorithm, which is iterating along the time axis, anyway, procedures for updates of spatial and temporal components have to be adapted to allow for online processing. This is achieved through an online NMF algorithm¹⁸⁵, updating the temporal components using Eq. A.7 & A.9 for frame $Y(:, t)$ corresponding to time t , one at a time. For the spatial update, the algorithm keeps the sufficient statistics of recent time steps in memory and performs the spatial update every few iteration steps as in Eq. A.8.

Registration of new candidates during runtime is possible by keeping the residual of a few timesteps in memory and applying to it a method similar to the initialization of the CNMF described above. Once the gaussian kernel-filtered residual displays a significant maximum, the component candidate is evaluated by the tools described above and if it passes the provided thresholds, added to the set of neurons. While this process allows OnACID to find neurons once they become active, it is usually started with a first estimate, provided for example by a CNMF-batch processing on the first few timesteps. It is worth noting, that calcium traces obtained from the online algorithm start at the timepoint, where the neuron was first detected and do not contain any information about its activity before that.

OnACID¹³⁵ was compared to the CNMF batch processing and the output of a number of human labelers. It performed on par with the batch processing (while being significantly faster) and slightly below human performance, however keeping in mind that the consensus of human labelling was taken as a ground truth, biasing results in favor of human performance. Both, spatial footprints and calcium transients could be reliably reproduced by either method.

Appendix Two

Nested sampling

Bayesian model evaluation

In general, the probability of observing two variables x and y can be expressed using joint ($p(\cdot, \cdot)$) and conditional ($p(\cdot|\cdot)$) probabilities

$$(B.1) \quad \begin{aligned} p(x, y) &= p(x|y)p(y) = p(y|x)p(x) \\ p(x|y) &= \frac{p(y|x)p(x)}{p(y)} \end{aligned}$$

with the latter commonly referred to as Bayes' Theorem. In the case of experimental data D , being fit to a model M with parameters Ω , it reads as

$$\underbrace{p(\Omega|D, M)}_{\text{posterior}} = \frac{\overbrace{p(D|\Omega, M)}^{\text{likelihood}} \overbrace{p(\Omega|M)}^{\text{prior}}}{\underbrace{p(D|M)}_{\text{evidence}}},$$

equivalent to Eq. B.1, with all probabilities additionally conditioned on the model M . The *posterior* $p(\Omega|D, M)$ describes the distribution of the parameters Ω of model M , after the *likelihood* $\mathcal{L}(\Omega) = p(D|\Omega, M)$ is evaluated and weighted by the *prior* $\pi(\Omega) = p(\Omega|M)$ - the initial belief of how parameters Ω should be distributed - and subsequently normalized by the *evidence* - or marginal likelihood - $\mathcal{Z} = p(D|M)$ - the probability to observe data like D at all, given an underlying model M .

In cases when the posterior distribution of parameters for a single model, only, should be obtained, the evidence can be omitted because it enters as a normalizing factor, only. The application in this work, however, aims to not only estimate parameters of a model, but also to compare the two different models of non-coding- and place-cells with one another. In this case, the evidence \mathcal{Z} can be obtained as the marginal likelihood by integrating over the d-dimensional parameter space

$$(B.2) \quad \mathcal{Z} = \int \mathcal{L}(\Omega)\pi(\Omega)d^d\Omega.$$

To obtain a criterion for model selection, the probability of the data resulting from either underlying model again is described by Bayes' Theorem as

$$p(M|D) = \frac{p(D|M)p(M)}{p(D)}$$

and the ratio of the two models provides the desired selection criterion via the Bayes factor and ratio of prior beliefs about the models as

$$(B.3) \quad \frac{p(M_1|D)}{p(M_0|D)} = \frac{p(D|M_1)p(M_1)}{p(D|M_0)p(M_0)} = \underbrace{\frac{\mathcal{Z}_1}{\mathcal{Z}_0}}_{\text{BayesFactor}} \frac{p(M_1)}{p(M_0)}.$$

Bayesian model selection inherently contains Occams Razor¹⁶¹, penalizing more complex - potentially overfitting - models, as those will always give a more accurate description of the data, at the cost of general applicability. In this, it is similar to other commonly used selection criteria such as Akaike information criterion or Bayesian information criterion, which use the estimated maximum likelihood, only, while the Bayes factor provides an estimate over the whole parameter space. The evidence is obtained by integrating the likelihood over the whole parameter space, which has the dimensionality of the number of free parameters, thus grows rapidly with each added model-parameter. While the detection of a fitting solution in a large parameter space becomes more and more probable the larger it grows, the volume-fraction of parameter space covered by this solution becomes smaller and smaller. Normalization of the evidence over the whole parameter space thereby implements the penalty towards more complex models.

Nested sampling

While the estimation of the posterior distribution and the selection criterion can be written down as direct results of applying Bayes' Theorem, the evaluation often turns out to be not as straight forward. It is often realized by sampling the parameter space through Markov chain Monte-Carlo (MCMC) methods, as has been done in¹³⁶, providing an estimate of the posterior distribution. However, the nature of MCMC methods provides an exhaustive sampling of the parameter space at high likelihoods, only, while large fractions of the parameter space remain unexplored. The evidence - as an integration of the likelihood over the parameter space, therefore can only be roughly approximated and comes at additional computational cost. Furthermore, MCMC methods often fail to capture multimodal shapes of the likelihood landscape.

In this work, I am using nested sampling¹³⁷, an approach commonly applied in astrophysics. It provides an elegant way of sampling the complete parameter space to obtain the evidence including an estimate of numerical uncertainty. The posterior distribution comes as a by-product and can be computed at almost no additional cost. Nested sampling is able to sample efficiently from high-dimensional parameter-spaces and has a natural way to deal with multimodal landscapes, such that it is able to detect multiple local maxima^{186,187}. Even though the method has

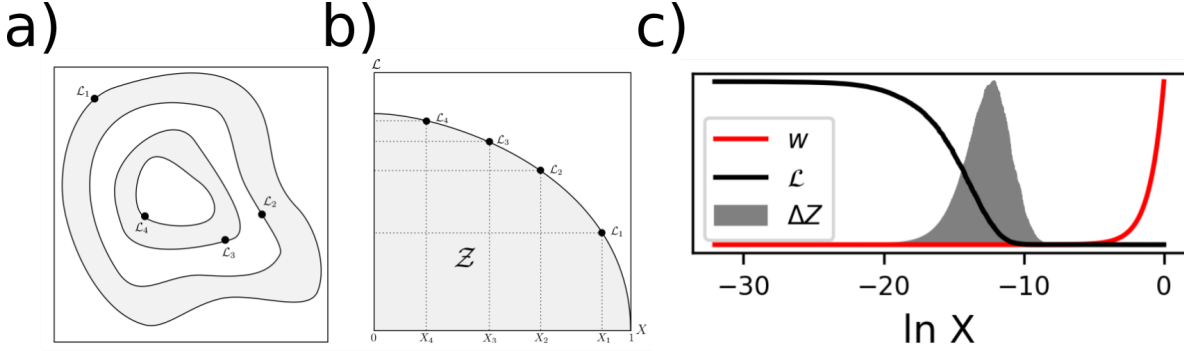


Figure B.1: Nested sampling explores the prior phase space by sampling from contours of iteratively increasing likelihood; **(a)** sketch of simple 2-dimensional phase space, with iso-likelihood contours at \mathcal{L}_{1-4} ; **(b)** increments of the evidence are estimated from the factor of the theoretically obtained prior volume shell w_i between sampling points enclosing prior volumes X_{i-1} and X_i and sampled likelihood \mathcal{L}_i , Eq. B.5; **(c)** volume shell w_i (red) decreases with sampling, while the likelihood (black) monotonously grows; accordingly, the bulk of the evidence - as a factor of the two values - is accumulated at intermediate values and saturates; figures (a,b) taken from Feroz et al, 2008¹⁸⁶

been subject to a number of refinements from further works, the general ideas which I am going to briefly go over throughout the next lines still remain the same.

The main idea of nested sampling is, to map the potentially multi-dimensional parameter landscape to a single variable - the volume of parameter space surrounded by iso-likelihood contours, see Fig. B.1(a), transforming Eq. B.2 into:

$$(B.4) \quad \mathcal{Z} = \int_0^1 \mathcal{L}(X) dX \quad , \quad X(\lambda) = \int_{\mathcal{L}(\Omega) > \lambda} \pi(\Omega) d^d \Omega,$$

$$(B.5) \quad \approx \sum_{i=1}^M \mathcal{L}_i w_i$$

with λ defining the iso-likelihood levels and X the normalized phase space volume. The integral is approximated by a summation over the product of likelihood \mathcal{L}_i at sample i and the weight w_i as the prior volume enclosed between iso-likelihood contours \mathcal{L}_{i-1} and \mathcal{L}_i . Samples are obtained by requiring increasing minimum likelihood values, starting with no constraint (phase space volume $X = 1$ at $i = 0$) and potentially ending with the maximum likelihood (or a point close to it), for which virtually $X \rightarrow 0$ at $i = M$.

In practice, the sampling starts out with a large number N of unconstrained samples which are iteratively updated. At each step, the sample with lowest likelihood is used to update the sum, Eq. B.5, Fig. B.1(b), and subsequently uniformly resampled from a reduced parameter space, constrained by the requirement to surpass previously evaluated likelihood-values. It thereby segments the phase space into nested shells around contours of constant likelihood \mathcal{L}_i (see Fig. B.1(a)), whose average occupation of phase space $w_i = X_{i-1} - X_i$ can be estimated:

The largest value of N uniform samples from a linear space of length X_{i-1} (with $X_0 = 1$) is described by the random variable t_i :

$$X_i = t_i X_{i-1} \quad , \quad p(t_i) = N t_i^{N-1}, t_i \in (0, 1)$$

with average $E(\ln t_i) = -\frac{1}{N}$ and standard deviation $\sigma(\ln t_i) = \frac{1}{N}$, thus describing the stochastic shrinking of phase space volume with increasing likelihood constraints. Consequently, the encircled volumes X_i can be described as

$$X_i = t_i X_{i-1} = t_i \cdots t_1 \cdot t_0 X_0 \quad , \quad t_j \in (0, 1).$$

A single run of the nested sampling procedure consists of a set of samples, for each of which the likelihood is evaluated and the weights are estimated by their expectation value. The product $\mathcal{L}_i w_i$ yields the contribution to the overall evidence.

Even though large sample numbers N allow to approximate the shrinkage well, the actual set of shrinkage parameters t_i is a random process, of which the assumed $t_i = 1/N$ for all i is only one (very unlikely) realization. Vanishing likelihood values for low i (covering large parts of the prior space) and vanishing prior weights for later i (at high likelihood values) result in the bulk of the evidence concentrated in intermediate values. The uncertainty in the shrinkage of prior volume in this region of large contributions, but mostly the uncertainty in the number of steps to reach this region, resulting from the underlying poisson process, is the main contribution to the overall uncertainty and can be estimated from evaluating Z from subsamples of the obtained total distribution,¹³⁷

The major challenge in nested sampling is the efficient and unbiased sampling of new points from the constrained prior. A number of methods are available and have been evaluated against analytically tractable likelihood landscapes^{152,188}. Other changes, such as dynamic nested sampling, in which the number of samples is dynamically adapted during runtime¹⁸⁹, have been proposed and are mostly increasing sampling efficiency to obtain better accuracy at less cost. I chose to use the readily available algorithm `UltraNest`¹⁵³, using a sampling procedure termed `MLFriends`, in which new samples are drawn from ellipsoids whose shape and sizes are determined dynamically from the current sample distribution.

REFERENCES

- [1] Larry R. Squire. “Memory and the Hippocampus: A Synthesis From Findings With Rats, Monkeys, and Humans”. In: *Psychological Review* 99.2 (1992), pp. 195–231.
- [2] Yadin Dudai, Avi Karni, and Jan Born. “The Consolidation and Transformation of Memory”. In: *Neuron* 88.1 (2015), pp. 20–32.
- [3] John O’Keefe and Lynn Nadel. *The Hippocampus as a Cognitive Map*. 1978.
- [4] Edward C. Tolman. “Cognitive maps in rats and men.” In: *Psychological Review* 55.4 (1948), pp. 189–208.
- [5] W. B. SCOVILLE and B. MILNER. “Loss of recent memory after bilateral hippocampal lesions.” In: *Journal of neurology, neurosurgery, and psychiatry* 20.1 (1957), pp. 11–21.
- [6] Larry R. Squire, Lisa Genzel, John T. Wixted, and Richard G. Morris. “Memory consolidation”. In: *Cold Spring Harbor Perspectives in Biology* 7.8 (2015).
- [7] Eric R. Kandel, Yadin Dudai, and Mark R. Mayford. “The molecular and systems biology of memory”. In: *Cell* 157.1 (2014), pp. 163–186.
- [8] Shanping Chen, Diancai Cai, Kaycey Pearce, Philip Y.W. Sun, Adam C. Roberts, and David L. Glanzman. “Reinstatement of long-term memory following erasure of its behavioral and synaptic expression in Aplysia”. In: *eLife* 3.November (2014), pp. 1–21.
- [9] David J. Foster and Matthew A. Wilson. “Reverse replay of behavioural sequences in hippocampal place cells during the awake state”. In: *Nature* 440.7084 (2006), pp. 680–683.
- [10] Joseph O’Neill, Timothy J. Senior, Kevin Allen, John R. Huxter, and Jozsef Csicsvari. “Reactivation of experience-dependent cell assembly patterns in the hippocampus”. In: *Nature Neuroscience* 11.2 (2008), pp. 209–215.
- [11] D. O. Hebb. “The Organization of Behavior; A Neuropsychological Theory”. In: *Psychology Press edition* (Oct. 1949).
- [12] Alessio Attardo, James E. Fitzgerald, and Mark J. Schnitzer. “Impermanence of dendritic spines in live adult CA1 hippocampus”. In: *Nature* 523.7562 (2015), pp. 592–596.
- [13] Thomas Pfeiffer, Stefanie Poll, Stephane Bancelin, Julie Angibaud, V. V.G.Krishna Inavalli, Kevin Keppler, Manuel Mittag, Martin Fuhrmann, and U. Valentin Nägerl. “Chronic 2P-STED imaging reveals high turnover of dendritic spines in the hippocampus in vivo”. In: *eLife* 7 (2018), pp. 1–17.

REFERENCES

- [14] L T Thompson and P J Best. “Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats”. In: *Brain Research* 509.2 (1990), pp. 299–308.
- [15] Yaniv Ziv, Laurie D Burns, Eric D Cocker, Elizabeth O Hamel, Kunal K Ghosh, Lacey J Kitch, Abbas El Gamal, and Mark J Schnitzer. “Long-term dynamics of CA1 hippocampal place codes.” In: *Nature neuroscience* 16.3 (2013), pp. 264–266.
- [16] Alon Rubin, Nitzan Geva, Liron Sheintuch, and Yaniv Ziv. “Hippocampal ensemble dynamics timestamp events in long-term memory”. In: *eLife* 4.DECEMBER2015 (2015), pp. 1–16.
- [17] Liron Sheintuch, Alon Rubin, Noa Brande-Eilat, Nitzan Geva, Noa Sadeh, Or Pinchasof, and Yaniv Ziv. “Tracking the Same Neurons across Multiple Days in Ca²⁺ Imaging Data”. In: *Cell Reports* 21.4 (2017), pp. 1102–1115.
- [18] Thomas Hainmueller and Marlene Bartos. “Parallel emergence of stable and dynamic memory engrams in the hippocampus”. In: *Nature* 558.7709 (2018), pp. 292–296.
- [19] J. O’Keefe and J. Dostrovsky. “The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat”. In: *Brain Research* 34.1 (1971), pp. 171–175.
- [20] Emily A. Mankin, Fraser T. Sparks, Begum Slayyeh, Robert J. Sutherland, Stefan Leutgeb, and Jill K. Leutgeb. “Neuronal code for extended time in the hippocampus”. In: *Proceedings of the National Academy of Sciences of the United States of America* 109.47 (2012), pp. 19462–19467.
- [21] Claudia Clopath, Tobias Bonhoeffer, Mark Hübener, and Tobias Rose. “Variance and invariance of neuronal long-term representations”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 372.1715 (2017).
- [22] William A Roberts, Miranda C. Feeney, Krista MacPherson, Mark Petter, Neil McMillan, and Evanya Musolino. “Episodic-like memory in rats: Is it based on when or how long ago?” In: *Science* 320.5872 (2008), pp. 113–115.
- [23] D H Hubel and T N Wiesel. “Receptive fields of single neurones in the cat’s striate cortex.” In: *The Journal of Physiology* 148.3 (Oct. 1959), pp. 574–591.
- [24] Manuel Schottdorf, Wolfgang Keil, David Coppola, Leonard E. White, and Fred Wolf. “Random Wiring, Ganglion Cell Mosaics, and the Functional Architecture of the Visual Cortex”. In: *PLoS Computational Biology* 11.11 (2015), pp. 1–40.
- [25] Dmitriy Aronov, Rhino Nevers, and David W. Tank. “Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit”. In: *Nature* 543.7647 (2017), pp. 719–722. arXiv: nature21692 [10.1038].
- [26] Christopher J. MacDonald, Kyle Q. Lepage, Uri T. Eden, and Howard Eichenbaum. “Hippocampal “time cells” bridge the gap in memory for discontinuous events”. In: *Neuron* 71.4 (2011), pp. 737–749. arXiv: arXiv:1011.1669v3.

- [27] Albert Tsao, Jørgen Sugar, Li Lu, Cheng Wang, James J. Knierim, May Britt Moser, and Edvard I. Moser. “Integrating time from experience in the lateral entorhinal cortex”. In: *Nature* 561.7721 (2018), pp. 57–62.
- [28] Alejandro Terrazas, Michael Krause, Peter Lipa, Katalin M. Gothard, Carol A. Barnes, and Bruce L. McNaughton. “Self-motion and the hippocampal spatial metric”. In: *Journal of Neuroscience* 25.35 (2005), pp. 8085–8096.
- [29] R. (1934) Lorente De Nó. “Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system.” In: *Journal für Psychologie und Neurologie* 46 (1934), pp. 113–177.
- [30] Kei M. Igarashi, Hiroshi T. Ito, Edvard I. Moser, and May Britt Moser. “Functional diversity along the transverse axis of hippocampal area CA1”. In: *FEBS Letters* 588.15 (2014), pp. 2470–2476.
- [31] Menno P. Witter and David G. Amaral. “Hippocampal Formation”. In: *The Rat Nervous System*. 2004, pp. 635–704.
- [32] Vivien Chevaleyre and Steven A. Siegelbaum. “Strong CA2 pyramidal neuron synapses define a powerful disinaptic cortico-hippocampal loop”. In: *Neuron* 66.4 (2010), pp. 560–572.
- [33] Keigo Kohara, Michele Pignatelli, Alexander J. Rivest, Hae Yoon Jung, Takashi Kitamura, Junghyup Suh, Dominic Frank, Koichiro Kajikawa, Nathan Mise, Yuichi Obata, Ian R. Wickersham, and Susumu Tonegawa. “Cell type-specific genetic and optogenetic tools reveal hippocampal CA2 circuits”. In: *Nature Neuroscience* 17.2 (2014), pp. 269–279.
- [34] Marianne Fyhn, Sturla Molden, Menno P. Witter, Edvard I. Moser, and May Britt Moser. “Spatial representation in the entorhinal cortex”. In: *Science* 305.5688 (2004), pp. 1258–1264.
- [35] Hong Wei Dong, Larry W. Swanson, Lin Chen, Michael S. Fanselow, and Arthur W. Toga. “Genomic-anatomic evidence for distinct functional domains in hippocampal field CA1”. In: *Proceedings of the National Academy of Sciences of the United States of America* 106.28 (2009), pp. 11794–11799.
- [36] Min W. Jung, Sidney I. Wiener, and Bruce L. McNaughton. “Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat”. In: *Journal of Neuroscience* 14.12 (1994), pp. 7347–7356.
- [37] Menno P. Witter and Edvard I. Moser. “Spatial representation and the architecture of the entorhinal cortex”. In: *Trends in Neurosciences* 29.12 (2006), pp. 671–678.
- [38] Norio Ishizuka, W Maxwell Cowan, and David G. Amaral. “A quantitative analysis of the dendritic organization of pyramidal cells in the rat hippocampus”. In: *Journal of Comparative Neurology* 362.1 (1995), pp. 17–45.
- [39] Norio Ishizuka, Janet Weber, and David G Amaral. “Organization of intrahippocampal projections originating from CA3 pyramidal cells in the rat”. In: *Journal of Comparative Neurology* 295.4 (1990), pp. 580–623.

REFERENCES

- [40] Kenji Mizuseki, Kamran Diba, Eva Pastalkova, and György Buzsáki. “Hippocampal CA1 pyramidal cells form functionally distinct sublayers”. In: *Nature Neuroscience* 14.9 (2011), pp. 1174–1183.
- [41] Nobuaki Tamamaki and Yoshiaki Nojyo. “Preservation of topography in the connections between the subiculum, field CA1, and the entorhinal cortex in rats”. In: *Journal of Comparative Neurology* 353.3 (1995), pp. 379–390.
- [42] Menno P. Witter, Floris G. Wouterlood, Pieterke A. Naber, and Theo Van Haften. “Anatomical organization of the parahippocampal-hippocampal network”. In: *Annals of the New York Academy of Sciences* 911 (2000), pp. 1–24.
- [43] Pieterke A Naber, Fernando H. Lopes Da Silva, and Menno P Witter. “Reciprocal connections between the entorhinal cortex and hippocampal fields CA1 and the subiculum are in register with the projections from CA1 to the subiculum”. In: *Hippocampus* 11.2 (2001), pp. 99–104.
- [44] Espen J. Henriksen, Laura L. Colgin, Carol A. Barnes, Menno P. Witter, May Britt Moser, and Edvard I. Moser. “Spatial representation along the proximodistal axis of CA1”. In: *Neuron* 68.1 (2010), pp. 127–137.
- [45] Edvard I. Moser, Emilio Kropff, and May-Britt Moser. “Place Cells, Grid Cells, and the Brain’s Spatial Representation System”. In: *Annual Review of Neuroscience* 31.1 (2008), pp. 69–89.
- [46] Joseph R. Manns and Howard Eichenbaum. “A cognitive map for object memory in the hippocampus”. In: *Learning and Memory* 16.10 (2009), pp. 616–624.
- [47] David J. Foster. “Replay Comes of Age”. In: *Annual Review of Neuroscience* 40.1 (2017), pp. 581–602.
- [48] M. W. Jung and B. L. McNaughton. “Spatial selectivity of unit activity in the hippocampal granular layer”. In: *Hippocampus* 3.2 (1993), pp. 165–182.
- [49] Eric L. Hargreaves, Geeta Rao, Inah Lee, and James J. Knierim. “Major dissociation between medial and lateral entorhinal input to dorsal hippocampus”. In: *Science* 308.5729 (2005), pp. 1792–1794.
- [50] Jill K. Leutgeb, Stefan Leutgeb, May Britt Moser, and Edvard I. Moser. “Pattern separation in the dentate gyrus and CA3 of the hippocampus”. In: *Science* 315.5814 (2007), pp. 961–966. arXiv: 1002.1037.
- [51] Kirsten Brun Kjelstrup, Trygve Solstad, Vegard Heimly Brun, Torkel Hafting, Stefan Leutgeb, Menno P. Witter, Edvard I. Moser, and May Britt Moser. “Finite scale of spatial representation in the hippocampus”. In: *Science* 321.5885 (2008), pp. 140–143.
- [52] L T Thompson and P J Best. “Place cells and silent cells in the hippocampus of freely-behaving rats.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 9.7 (1989), pp. 2382–2390.

- [53] A. Vazdarjanova. “Differences in Hippocampal Neuronal Population Responses to Modifications of an Environmental Context: Evidence for Distinct, Yet Complementary, Functions of CA3 and CA1 Ensembles”. In: *Journal of Neuroscience* 24.29 (2004), pp. 6489–6496.
- [54] Charlotte B. Alme, Chenglin Miao, Karel Jezek, Alessandro Treves, Edvard I. Moser, and May Britt Moser. “Place cells in the hippocampus: Eleven maps for eleven rooms”. In: *Proceedings of the National Academy of Sciences of the United States of America* 111.52 (2014), pp. 18428–18435.
- [55] Francesca Sargolini, Marianne Fyhn, Torkel Hafting, Bruce L. McNaughton, Menno P. Witter, May Britt Moser, and Edvard I. Moser. “Conjunctive representation of position, direction, and velocity in entorhinal cortex”. In: *Science* 312.5774 (2006), pp. 758–762.
- [56] Marianne Fyhn, Torkel Hafting, Menno P. Witter, Edvard I. Moser, and May Britt Moser. “Grid cells in mice”. In: *Hippocampus* 18.12 (2008), pp. 1230–1238.
- [57] Colin Lever, Tom Wills, Francesca Cacucci, Neil Burgess, and John O’Keefe. “Long-term plasticity in hippocampal place-cell representation of environmental geometry”. In: *Nature* 416.6876 (2002), pp. 90–94.
- [58] S. Leutgeb. “Distinct Ensemble Codes in Hippocampal Areas CA3 and CA1”. In: *Science* 305.5688 (2004), pp. 1295–1298.
- [59] A. J. Hill. “First occurrence of hippocampal spatial firing in a new environment”. In: *Experimental Neurology* 62.2 (1978), pp. 282–297.
- [60] M. Wilson and B. McNaughton. “Dynamics of the hippocampal ensemble code for space”. In: *Science* 261.5124 (1993), pp. 1055–1058.
- [61] L M Frank. “Hippocampal Plasticity across Multiple Days of Exposure to Novel Environments”. In: *Journal of Neuroscience* 24.35 (2004), pp. 7681–7689.
- [62] Masaaki Sato, Kotaro Mizuta, Tanvir Islam, Masako Kawano, Yukiko Sekine, Takashi Takekawa, Daniel Gomez-Dominguez, Alexander Schmidt, Fred Wolf, Karam Kim, Hiroshi Yamakawa, Masamichi Ohkura, Min Goo Lee, Tomoki Fukai, Junichi Nakai, and Yasunori Hayashi. “Distinct Mechanisms of Over-Representation of Landmarks and Rewards in the Hippocampus”. In: *Cell Reports* 32.1 (2020), p. 107864.
- [63] M. R. Mehta, C. A. Barnes, and B. L. McNaughton. “Experience-dependent, asymmetric expansion of hippocampal place fields”. In: *Proceedings of the National Academy of Sciences* 94.16 (1997), pp. 8918–8921.
- [64] Mayank R. Mehta, Michael C. Quirk, and Matthew A. Wilson. “Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields”. In: *Neuron* 25.3 (2000), pp. 707–715.
- [65] Kimberly L Stachenfeld, Matthew M Botvinick, and Samuel J Gershman. “The hippocampus as a predictive map”. In: *IEEE Transactions on Geoscience and Remote Sensing* 42.11 (2017), pp. 2555–2568. arXiv: arXiv:1011.1669v3.
- [66] R. U. Muller and J. L. Kubie. “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells”. In: *Journal of Neuroscience* 7.7 (1987), pp. 1951–1968.

REFERENCES

- [67] J. O'Keefe and N. Burgess. "Geometric Determinants Of Place Cell Firing". In: *Nature* 381.May (1996), pp. 435–428.
- [68] Benjamin J. Kraus, Robert J. Robinson, John A. White, Howard Eichenbaum, and Michael E. Hasselmo. "Hippocampal "Time Cells": Time versus Path Integration". In: *Neuron* 78.6 (2013), pp. 1090–1101.
- [69] Michael I Anderson and Kathryn J Jeffery. "Heterogeneous modulation of place cell firing by changes in context." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 23.26 (2003), pp. 8827–8835.
- [70] Elizabeth Bostock, Robert U. Muller, and John L. Kubie. "Experience dependent modifications of hippocampal place cell firing". In: *Hippocampus* 1.2 (1991), pp. 193–205.
- [71] Nathaniel R. Kinsky, David W. Sullivan, William Mau, Michael E. Hasselmo, and Howard B. Eichenbaum. "Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales". In: *Current Biology* 28.22 (2018), 3578–3588.e6.
- [72] G. Chen, J. A. King, N. Burgess, and J. O'Keefe. "How vision and movement combine in the hippocampal place code". In: *Proceedings of the National Academy of Sciences* 110.1 (2013), pp. 378–383.
- [73] Clifford G Kentros, Naveen T Agnihotri, Samantha Streater, Robert D Hawkins, and Eric R Kandel. "Increased attention to spatial context increases both place field stability and spatial memory". In: *Neuron* 42.2 (2004), pp. 283–295.
- [74] Tobias Rose, Juliane Jaepel, Mark Hübener, and Tobias Bonhoeffer. "Cell-specific restoration of stimulus preference after monocular deprivation in the visual cortex." In: *Science (New York, N.Y.)* 352.6291 (2016), pp. 1319–1322.
- [75] Johannes M. Mayrhofer, Florent Haiss, Fritjof Helmchen, and Bruno Weber. "Sparse, reliable, and long-term stable representation of periodic whisker deflections in the mouse barrel cortex". In: *NeuroImage* 115 (2015), pp. 52–63.
- [76] Ashesh K Dhawale, Rajesh Poddar, Steffen B.E. Wolff, Valentin A Normand, Evi Kopelowitz, and Bence P. Ölveczky. "Automated long-Term recording and analysis of neural activity in behaving animals". In: *eLife* 6 (2017), pp. 1–40.
- [77] Simon P. P. Peron, Jeremy Freeman, Vijay Iyer, Caiying Guo, and Karel Svoboda. "A Cellular Resolution Map of Barrel Cortex Activity during Tactile Behavior". In: *Neuron* 86.3 (2015), pp. 783–799.
- [78] Jasper Poort, Adil G. Khan, Marius Pachitariu, Abdellatif Nemri, Ivana Orsolich, Julija Krupic, Marius Bauza, Maneesh Sahani, Georg B. Keller, Thomas D. Mrsic-Flogel, and Sonja B. Hofer. "Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex". In: *Neuron* 86.6 (2015), pp. 1478–1490.
- [79] Jerry L. Chen, David J. Margolis, Atanas Stankov, Lazar T. Sumanovski, Bernard L. Schneider, and Fritjof Helmchen. "Pathway-specific reorganization of projection neurons in somatosensory cortex during learning". In: *Nature Neuroscience* 18.8 (2015), pp. 1101–1108.

-
- [80] Sonja B. Hofer, Thomas D. Mrsic-Flogel, Tobias Bonhoeffer, and Mark Hübener. “Experience leaves a lasting structural trace in cortical circuits”. In: *Nature* 457.7227 (2009), pp. 313–317.
- [81] Anthony J.G.D. Holtmaat, Joshua T. Trachtenberg, Linda Wilbrecht, Gordon M. Shepherd, Xiaoqun Zhang, Graham W. Knott, and Karel Svoboda. “Transient and persistent dendritic spines in the neocortex in vivo”. In: *Neuron* 45.2 (2005), pp. 279–291.
- [82] Tonghui Xu, Xinzhu Yu, Andrew J. Perlik, Willie F. Tobin, Jonathan A. Zweig, Kelly Tennant, Theresa Jones, and Yi Zuo. “Rapid formation and selective stabilization of synapses for enduring motor memories”. In: *Nature* 462.7275 (2009), pp. 915–919.
- [83] Guang Yang, Feng Pan, and Wen Biao Gan. “Stably maintained dendritic spines are associated with lifelong memories”. In: *Nature* 462.7275 (2009), pp. 920–924.
- [84] Yi Zuo, Guang Yang, Elaine Kwon, and Wen Biao Gan. “Long-term sensory deprivation prevents dendritic spine loss in primary somatosensory cortex”. In: *Nature* 436.7048 (2005), pp. 261–265.
- [85] R U Muller and John L Kubie. “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells”. In: *Journal of Neuroscience* 7.7 (1987), pp. 1951–1968.
- [86] Emily A Mankin, Geoffrey W Diehl, Fraser T Sparks, Stefan Leutgeb, and Jill K Leutgeb. “Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts”. In: *Neuron* 85.1 (2015), pp. 190–201.
- [87] Li Lu, Kei M. Igarashi, Menno P. Witter, Edvard I. Moser, and May Britt Moser. “Topography of Place Maps along the CA3-to-CA2 Axis of the Hippocampus”. In: *Neuron* 87.5 (2015), pp. 1078–1092.
- [88] Patrick Latuske, Olga Kornienko, Laura Kohler, and Kevin Allen. “Hippocampal Remapping and Its Entorhinal Origin”. In: *Frontiers in Behavioral Neuroscience* 11.January (2018), pp. 1–13.
- [89] Stefan Leutgeb, Jill K Leutgeb, Carol A Barnes, Edvard I Moser, Bruce L McNaughton, and May-Britt Moser. “Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles”. In: *Science* 309.July (2005), pp. 619–623.
- [90] Tom J. Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O’Keefe. “Attractor dynamics in the hippocampal representation of the local environment”. In: *Science* 308.5723 (2005), pp. 873–876.
- [91] Emma R. Wood, Paul A. Dudchenko, and Howard Eichenbaum. “The global record of memory in hippocampal neuronal activity”. In: *Nature* 397.6720 (1999), pp. 613–616.
- [92] R. W. Komorowski, J. R. Manns, and H. Eichenbaum. “Robust Conjunctive Item-Place Coding by Hippocampal Neurons Parallels Learning What Happens Where”. In: *Journal of Neuroscience* 29.31 (2009), pp. 9918–9929.

REFERENCES

- [93] K. Allen, J. N. P. Rawlins, D. M. Bannerman, and J. Csicsvari. “Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping”. In: *Journal of Neuroscience* 32.42 (2012), pp. 14752–14766.
- [94] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. “Microstructure of a spatial map in the entorhinal cortex”. In: *Nature* 436.7052 (2005), pp. 801–806. arXiv: /dx.doi.org/10.1038/nature01964 [http:].
- [95] Marianne Fyhn, Torkel Hafting, Alessandro Treves, May Britt Moser, and Edvard I Moser. “Hippocampal remapping and grid realignment in entorhinal cortex”. In: *Nature* 446.7132 (2007), pp. 190–194.
- [96] Sachin S. Deshmukh and James J. Knierim. “Representation of Non-Spatial and Spatial Information in the Lateral Entorhinal Cortex”. In: *Frontiers in Behavioral Neuroscience* 5 (2011).
- [97] James J. Knierim, Joshua P. Neunuebel, and Sachin S. Deshmukh. *Functional correlates of the lateral and medial entorhinal cortex: Objects, path integration and local - Global reference frames*. 2014.
- [98] Joshua P. Neunuebel, D. Yoganarasimha, Geeta Rao, and James J. Knierim. “Conflicts between local and global spatial frameworks dissociate neural representations of the lateral and medial entorhinal cortex”. In: *Journal of Neuroscience* 33.22 (2013), pp. 9246–9258.
- [99] Albert Tsao, May Britt Moser, and Edvard I. Moser. “Traces of experience in the lateral entorhinal cortex”. In: *Current Biology* 23.5 (2013), pp. 399–405.
- [100] Trygve Solstad, Trygve Solstad, Charlotte N Boccara, Charlotte N Boccara, Emilio Kropff, Emilio Kropff, May-Britt Moser, May-Britt Moser, Edvard I Moser, and Edvard I Moser. “Representation of geometric borders in the entorhinal cortex”. In: *Science* 322.December (2008), pp. 1865–1868.
- [101] Francesco Savelli, D. Yoganarasimha, and James J. Knierim. “Influence of boundary removal on the spatial representations of the medial entorhinal cortex”. In: *Hippocampus* 18.12 (2008), pp. 1270–1282.
- [102] S. J.Y. Mizumori, B. L. McNaughton, C. A. Barnes, and K. B. Fox. “Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3c output: Evidence for pattern completion in hippocampus”. In: *Journal of Neuroscience* 9.11 (1989), pp. 3915–3928.
- [103] David Marr. “Simple Memory: A Theory For Archicortex”. In: *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 262.841 (1971), pp. 23–81.
- [104] Alessandro Treves and Edmund T. Rolls. “Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network”. In: *Hippocampus* 2.2 (1992), pp. 189–199.

-
- [105] Heekyung Lee, Cheng Wang, Sachin S. Deshmukh, and James J. Knierim. “Neural Population Evidence of Functional Heterogeneity along the CA3 Transverse Axis: Pattern Completion versus Pattern Separation”. In: *Neuron* 87.5 (2015), pp. 1093–1105.
- [106] Inah Lee, D. Yoganarasimha, Geeta Rao, and James J. Knierim. “Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3”. In: *Nature* 430.6998 (2004), pp. 456–459.
- [107] Stefan Leutgeb. “Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles”. In: July (2005), pp. 619–624.
- [108] Frederick L. Hitti and Steven A. Siegelbaum. “The hippocampal CA2 region is essential for social memory”. In: *Nature* 508.1 (2014), pp. 88–92.
- [109] Erica L. Stevenson and Heather K. Caldwell. “Lesions to the CA2 region of the hippocampus impair social memory in mice”. In: *European Journal of Neuroscience* 40.9 (2014), pp. 3294–3301.
- [110] Matthew W. Jones and Thomas J. McHugh. “Updating hippocampal representations: CA2 joins the circuit”. In: *Trends in Neurosciences* 34.10 (2011), pp. 526–535.
- [111] Loren M. DeVito, Rachael Konigsberg, Christine Lykken, Magdalena Sauvage, W. Scott Young, and Howard Eichenbaum. “Vasopressin 1b receptor knock-out impairs memory for temporal order”. In: *Journal of Neuroscience* 29.9 (2009), pp. 2676–2683.
- [112] Nathan B Danielson, Jeffrey D Zaremba, Patrick Kaifosh, John Bowler, Max Ladow, and Attila Losonczy. “Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1”. In: *Neuron* 91.3 (2016), pp. 652–665.
- [113] J. O’Keefe and D. H. Conway. “Hippocampal place units in the freely moving rat: Why they fire where they fire”. In: *Experimental Brain Research* 31.4 (1978), pp. 573–590.
- [114] S I Wiener, C A Paul, and H Eichenbaum. “Spatial and behavioral correlates of hippocampal neuronal activity”. In: *J Neurosci* 9.8 (1989), pp. 2737–2763.
- [115] Phil A. Hetherington and Matthew L. Shapiro. “Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner.” In: *Behavioral Neuroscience* 111.1 (1997), pp. 20–34.
- [116] S a Hollup, S Molden, J G Donnett, M B Moser, and E I Moser. “Accumulation of hippocampal place fields at the goal location in an annular watermaze task.” In: *The Journal of neuroscience* 21.5 (2001), pp. 1635–1644.
- [117] David Dupret, Joseph O’Neill, Barty Pleydell-Bouverie, and Jozsef Csicsvari. “The reorganization and reactivation of hippocampal maps predict spatial memory performance”. In: *Nature Neuroscience* 13.8 (2010), pp. 995–1002.
- [118] Jeffrey D. Zaremba, Anastasia Diamantopoulou, Nathan B. Danielson, Andres D. Gromark, Patrick W. Kaifosh, John C. Bowler, Zhenrui Liao, Fraser T. Sparks, Joseph A. Gogos, and Attila Losonczy. “Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion”. In: *Nature Neuroscience* 20.11 (2017), pp. 1612–1623.

REFERENCES

- [119] Tristan Geiller, Mohammad Fattahi, June-Seek Choi, and Sébastien Royer. “Place cells are more strongly tied to landmarks in deep than in superficial CA1”. In: *Nature Communications* 8 (2017), p. 14531.
- [120] Jeffrey L. Gauthier and David W. Tank. “A Dedicated Population for Reward Coding in the Hippocampus”. In: *Neuron* 99.1 (2018), 179–193.e7.
- [121] Romain Bourboulou, Geoffrey Marti, François Xavier Michon, Elissa El Feghaly, Morgane Nouguié, David Robbe, Julie Koenig, and Jerome Epszstein. “Dynamic control of hippocampal spatial coding resolution by local visual cues”. In: *eLife* 8 (2019), pp. 1–30.
- [122] Sachin S Deshmukh and James J Knierim. “Influence of local objects on hippocampal representations: Landmark vectors and memory”. In: *Hippocampus* 23.4 (2013), pp. 253–267.
- [123] Vincent Hok, Pierre Pascal Lenck-Santini, Sébastien Roux, Etienne Save, Robert U. Muller, and Bruno Poucet. “Goal-related activity in hippocampal place cells”. In: *Journal of Neuroscience* 27.3 (2007), pp. 472–482.
- [124] James J Knierim and Geeta Rao. *Distal landmarks and hippocampal place cells: Effects of relative translation versus rotation*. 2003.
- [125] Sara N. Burke, Andrew P. Maurer, Saman Nematollahi, Ajay R. Uprety, Jenelle L. Wallace, and Carol A. Barnes. “The influence of objects on place field expression and size in distal hippocampal CA1”. In: *Hippocampus* 21.7 (2011), pp. 783–801.
- [126] Marianne Fyhn, Sturla Molden, Stig Hollup, May Britt Moser, and Edvard I Moser. “Hippocampal neurons responding to first-time dislocation of a target object”. In: *Neuron* 35.3 (2002), pp. 555–566.
- [127] Matthijs A.A. van der Meer, Adam Johnson, Neil C. Schmitzer-Torbert, and A David Redish. “Triple dissociation of information processing in dorsal striatum, ventral striatum, and hippocampus on a learned spatial decision task”. In: *Neuron* 67.1 (2010), pp. 25–32.
- [128] Brad E. Pfeiffer and David J. Foster. “Hippocampal place-cell sequences depict future paths to remembered goals”. In: *Nature* 497.7447 (2013), pp. 74–79.
- [129] C. Hölscher, A. Schnee, H. Dahmen, L. Setia, and H. A. Mallot. “Rats are able to navigate in virtual environments”. In: *Journal of Experimental Biology* 208.3 (2005), pp. 561–569.
- [130] Pascal Ravassard, Ashley Kees, Bernard Willers, David Ho, Daniel Aharoni, Jesse Cushman, Zahra M Aghajan, and Mayank R Mehta. “Multisensory Control of Hippocampal Spatiotemporal Selectivity”. In: *June* (2013), pp. 1342–1347.
- [131] Dmitriy Aronov and David W. Tank. “Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System”. In: *Neuron* 84.2 (2014), pp. 442–456.
- [132] Kay Thurley and Asli Ayaz. “Virtual reality systems for rodents”. In: *Current Zoology* 63.1 (2017), pp. 109–119.

- [133] Isaac A. Youngstrom and Ben W. Strowbridge. “Visual landmarks facilitate rodent spatial navigation in virtual reality environments”. In: *Learning and Memory* 19.3 (2012), pp. 84–90.
- [134] Masaaki Sato, Masako Kawano, Kotaro Mizuta, Tanvir Islam, Min Goo Lee, and Yasunori Hayashi. “Hippocampus-Dependent Goal Localization by Head-Fixed Mice in Virtual Reality”. In: *eneuro* 4.3 (2017), ENEURO.0369–16.2017.
- [135] Andrea Giovannucci, Johannes Friedrich, Pat Gunn, Jérémie Kalfon, Brandon L Brown, Sue Ann Koay, Jiannis Taxisidis, Farzaneh Najafi, Jeffrey L Gauthier, Pengcheng Zhou, Baljit S Khakh, David W Tank, Dmitri B Chklovskii, and Eftychios A Pnevmatikakis. “CaImAn an open source tool for scalable calcium imaging data analysis”. In: *eLife* 8 (2019), pp. 1–45.
- [136] Beau Cronin, Ian H. Stevenson, Mriganka Sur, and Konrad P. Körding. “Hierarchical bayesian modeling and Markov chain Monte Carlo sampling for tuning-curve analysis”. In: *Journal of Neurophysiology* 103.1 (2010), pp. 591–602.
- [137] John Skilling. “Nested sampling for general Bayesian computation”. In: *Bayesian Analysis* 1.4 (2006), pp. 833–860.
- [138] Marius Pachitariu, Carsen Stringer, Sylvia Schröder, Mario Dipoppa, L. Federico Rossi, Matteo Carandini, and Kenneth D Harris. “Suite2p: beyond 10,000 neurons with standard two-photon microscopy”. In: *bioRxiv* (2016), p. 61507.
- [139] Sebastián A. Romano, Verónica Pérez-Schuster, Adrien Jouary, Jonathan Boulanger-Weill, Alessia Candeo, Thomas Pietri, and Germán Sumbre. “An integrated calcium imaging processing toolbox for the analysis of neuronal population dynamics”. In: *PLoS Computational Biology* 13.6 (2017).
- [140] Ryuichi Maruyama, Kazuma Maeda, Hajime Moroda, Ichiro Kato, Masashi Inoue, Hiroyoshi Miyakawa, and Toru Aonishi. “Detecting cells using non-negative matrix factorization on calcium imaging data”. In: *Neural Networks* 55 (2014), pp. 11–19.
- [141] Eftychios A. Pnevmatikakis, Daniel Soudry, Yuanjun Gao, Timothy A. Machado, Josh Merel, David Pfau, Thomas Reardon, Yu Mu, Clay Lacefield, Weijian Yang, Misha Ahrens, Randy Bruno, Thomas M. Jessell, Darcy S. Peterka, Rafael Yuste, and Liam Paninski. “Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data”. In: *Neuron* 89.2 (2016), p. 299. arXiv: arXiv:1409.2903v1.
- [142] Masamichi Ohkura, Takuya Sasaki, Junko Sadakari, Keiko Gengyo-Ando, Yuko Kagawa-Nagamura, Chiaki Kobayashi, Yuji Ikegaya, and Junichi Nakai. “Genetically Encoded Green Fluorescent Ca²⁺ Indicators with Improved Detectability for Neuronal Ca²⁺ Signals”. In: *PLoS ONE* 7.12 (2012), pp. 1–10.
- [143] Masaaki Sato, Masako Kawano, Masamichi Ohkura, Keiko Gengyo-Ando, Junichi Nakai, and Yasunori Hayashi. “Generation and imaging of transgenic mice that express G-CaMP7 under a tetracycline response element”. In: *PLoS ONE* 10.5 (2015), pp. 1–13.
- [144] Masaaki Sato and Michael P. Stryker. “Distinctive features of adult ocular dominance plasticity”. In: *Journal of Neuroscience* 28.41 (2008), pp. 10278–10286.

REFERENCES

- [145] Bruce D. Lucas and Takeo Kanade. “ITERATIVE IMAGE REGISTRATION TECHNIQUE WITH AN APPLICATION TO STEREO VISION.” In: vol. 2. 1981, pp. 674–679.
- [146] Patrick Kaifosh, Jeffrey D. Zaremba, Nathan B. Danielson, and Attila Losonczy. “SIMA: Python software for analysis of dynamic fluorescence imaging data”. In: *Frontiers in Neuroinformatics* 8.September (2014), pp. 1–10.
- [147] Eran A.; Mukamel, Axel; Nimmerjahn, and Mark J. Schnitzer. “Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data”. In: *Neuron* 63.6 (2009), pp. 747–760.
- [148] Johannes Friedrich, Pengcheng Zhou, and Liam Paninski. “Fast online deconvolution of calcium imaging data”. In: *PLOS Computational Biology* (2017).
- [149] Borbala Podor, Yi-ling Hu, Roger Croll, Masamichi Ohkura, Alan Fine, and Junichi Nakai. “Comparison of genetically encoded calcium indicators for monitoring action potentials in mammalian brain by two-photon excitation fluorescence microscopy”. In: *Neurophotonics* 2.2 (2015), p. 021014.
- [150] Gunnar Farneback. “Two-frame motion estimation based on polynomial expansion”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2749.1 (2003), pp. 363–370.
- [151] H. W. Kuhn. “The Hungarian method for the assignment problem”. In: *Naval Research Logistics Quarterly* 2.1-2 (1955), pp. 83–97.
- [152] Johannes Buchner. “A statistical test for Nested Sampling algorithms”. In: *Statistics and Computing* 26.1-2 (2016), pp. 383–392. arXiv: 1407.5459.
- [153] Johannes Buchner. “Collaborative nested sampling: Big data versus complex physical models”. In: *Publications of the Astronomical Society of the Pacific* 131.1004 (2019). arXiv: 1707.04476.
- [154] Bryan C Souza, Rodrigo Pavão, Hindiael Belchior, and Adriano B.L. Tort. “On Information Metrics for Spatial Coding”. In: *Neuroscience* 375.Ii (2018), pp. 62–73.
- [155] William E. Skaggs and Bruce L. McNaughton and Katalin M. Gothard and Etan J. Markus. “An Information-Theoretic Approach to Deciphering the Hippocampal Code”. In: *Proceedings of the IEEE* 1990 (1993), pp. 1030–1037.
- [156] Emery N. Brown, David P. Nguyen, Loren M. Frank, Matthew A. Wilson, and Victor Solo. “An analysis of neural receptive field plasticity by point process adaptive filtering”. In: *Proceedings of the National Academy of Sciences of the United States of America* 98.21 (2001), pp. 12261–12266.
- [157] F. Cacucci, T. J. Wills, C. Lever, K. P. Giese, and J. O’Keefe. “Experience-Dependent Increase in CA1 Place Cell Spatial Information, But Not Spatial Reproducibility, Is Dependent on the Autophosphorylation of the -Isoform of the Calcium/Calmodulin-Dependent Protein Kinase II”. In: *Journal of Neuroscience* 27.29 (2007), pp. 7854–7859.
- [158] Loren M. Frank, Emery N. Brown, and Matthew Wilson. “Trajectory encoding in the hippocampus and entorhinal cortex”. In: *Neuron* 27.1 (2000), pp. 169–178.

- [159] Eva Pastalkova, Vladimir Itskov, Asohan Amarasingham, and György Buzsáki. “Internally generated cell assembly sequences in the rat hippocampus”. In: *Science* 321.5894 (2008), pp. 1322–1327.
- [160] W R Softky and C Koch. “The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 13.1 (1993), pp. 334–350.
- [161] David J C MacKay. *Information Theory, Inference, and Learning Algorithms*. Vol. 100. 2005, pp. 1–640.
- [162] Harold Jeffreys. *The Theory of Probability*. Oxford University Press, 1961.
- [163] Andrea Giovannucci, Johannes Friedrich, Matt Kaufman, Anne Churchland, Dmitri Chklovskii, Liam Paninski, and Eftychios A Pnevmatikakis. “Onacid: Online analysis of calcium imaging data in real time”. In: *Advances in Neural Information Processing Systems Nips* (2017), pp. 2381–2391.
- [164] Shy Shoham, Daniel H O’Connor, and Ronen Segev. “How silent is the brain: Is there a “dark matter” problem in neuroscience?” In: *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 192.8 (2006), pp. 777–784.
- [165] Jill K. Leutgeb, Stefan Leutgeb, Alessandro Treves, Retsina Meyer, Carol A. Barnes, Bruce L. McNaughton, May Britt Moser, and Edvard I. Moser. “Progressive transformation of hippocampal neuronal representations in “morphed” environments”. In: *Neuron* 48.2 (2005), pp. 345–348.
- [166] Laura Lee Colgin, Edvard I. Moser, and May Britt Moser. “Understanding memory through hippocampal remapping”. In: *Trends in Neurosciences* 31.9 (2008), pp. 469–477.
- [167] Tsuneyuki Kobayashi, Hisao Nishijo, Masaji Fukuda, Jan Bures, and Taketoshi Ono. “Task-dependent representations in rat hippocampal place neurons”. In: *Journal of Neurophysiology* 78.2 (1997), pp. 597–613.
- [168] Emma R. Wood, Paul A. Dudchenko, R. Jonathan Robitsek, and Howard Eichenbaum. “Hippocampal neurons encode information about different types of memory episodes occurring in the same location”. In: *Neuron* 27.3 (2000), pp. 623–633.
- [169] Daniela Vallentin, Georg Kosche, Dina Lipkind, and Michael A. Long. “Inhibition protects acquired song segments during vocal learning in zebra finches”. In: *Science* 351.6270 (2016), pp. 267–271.
- [170] S B Laughlin, R R de Ruyter van Steveninck, and J C Anderson. “The metabolic cost of neural information.” In: *Nature neuroscience* 1.1 (1998), pp. 36–41.
- [171] David Attwell and Simon B. Laughlin. “An energy budget for signaling in the grey matter of the brain.” In: *Journal of cerebral blood flow and metabolism : official journal of the International Society of Cerebral Blood Flow and Metabolism* 21.10 (Oct. 2001), pp. 1133–1145.
- [172] Peter Lennie. “The cost of cortical computation.” In: *Curr. Biol.* 13.6 (Mar. 2003), pp. 493–497.

REFERENCES

- [173] Clare Howarth, Claire Peppiatt-Wildman, and David Attwell. “The Energy Use Associated with Neural Computation in the Cerebellum”. In: *J Cereb Blood Flow Metab* 30.2 (2010), pp. 403–414.
- [174] Joseph R Manns, Marc W Howard, and Howard Eichenbaum. “Gradual Changes in Hippocampal Activity Support Remembering the Order of Events”. In: *Neuron* 56.3 (2007), pp. 530–540.
- [175] Jess Nithianantharajah and Anthony J. Hannan. “Enriched environments, experience-dependent plasticity and disorders of the nervous system”. In: *Nature Reviews Neuroscience* 7.9 (2006), pp. 697–709.
- [176] Henriette van Praag, Gerd Kempermann, and Fred H. Gage. “Neural consequences of environmental enrichment”. In: *Nature Reviews Neuroscience* 1.3 (2000), pp. 191–198.
- [177] Heinz Steffens, Alexander C Mott, Siyuan Li, Waja Wegner, Š Pavel, and W Y Vanessa. “Chronic in vivo STED nanoscopy uncovers multiple drivers of shape volatility in stable cortical spines”. In: (2020).
- [178] Marie E. Wintzer, Roman Boehringer, Denis Polygalov, and Thomas J. McHugh. “The hippocampal CA2 ensemble is sensitive to contextual change”. In: *Journal of Neuroscience* 34.8 (2014), pp. 3056–3066.
- [179] Yonatan Loewenstein, Uri Yanover, and Simon Rumpel. “Predicting the dynamics of network connectivity in the neocortex”. In: *Journal of Neuroscience* 35.36 (2015), pp. 12535–12544.
- [180] Michael Fauth, Florentin Wörgötter, and Christian Tetzlaff. “Formation and Maintenance of Robust Long-Term Information Storage in the Presence of Synaptic Turnover”. In: *PLoS Computational Biology* 11.12 (2015), pp. 1–22.
- [181] Gianluigi Mongillo, Simon Rumpel, and Yonatan Loewenstein. “Inhibitory connectivity defines the realm of excitatory plasticity”. In: *Nature Neuroscience* 21.10 (2018), pp. 1463–1470.
- [182] Lee Susman, Naama Brenner, and Omri Barak. “Stable memory with unstable synapses”. In: *Nature Communications* 10.1 (2019). arXiv: 1808.00756.
- [183] Michael London, Arnd Roth, Lisa Beeren, Michael Häusser, and Peter E. Latham. “Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex”. In: *Nature* 466.7302 (2010), pp. 123–127.
- [184] Joshua T Vogelstein, Adam M Packer, Timothy A Machado, Tanya Sippy, Baktash Babadi, Rafael Yuste, and Liam Paninski. “Vogelstein et al 2010-Fast Nonnegative Deconvolution for Spike Train Inference From Population Calcium Imaging”. In: (2010), pp. 3691–3704.
- [185] Julien Mairal, Francis Bach FRANCISBACH, Jean Ponce JEANPONCE, and Guillermo Sapiro. “Online Learning for Matrix Factorization and Sparse Coding”. In: *Journal of Machine Learning Research* 11 (2010), pp. 19–60.

- [186] F. Feroz and M. P. Hobson. “Multimodal nested sampling: An efficient and robust alternative to Markov Chain Monte Carlo methods for astronomical data analyses”. In: *Monthly Notices of the Royal Astronomical Society* 384.2 (2008), pp. 449–463. arXiv: 0704.3704.
- [187] Farhan Feroz and John Skilling. “Exploring multi-modal distributions with nested sampling”. In: *AIP Conference Proceedings* 1553.2 (2013), pp. 106–113. arXiv: 1312.5638.
- [188] Edward Higson, Will Handley, Michael Hobson, and Anthony Lasenby. “NESTCHECK: Diagnostic tests for nested sampling calculations”. In: *Monthly Notices of the Royal Astronomical Society* 483.2 (2019), pp. 2044–2056. arXiv: arXiv:1804.06406v3.
- [189] Edward Higson, Will Handley, Michael Hobson, and Anthony Lasenby. “Dynamic nested sampling: an improved algorithm for parameter estimation and evidence calculation”. In: *Statistics and Computing* 29.5 (2019), pp. 891–913. arXiv: 1704.03459.

ACKNOWLEDGEMENTS

First of all, I want to thank my supervisor and mentor Fred Wolf, who has not only sparked my interest into the the research of the brain, but has also kept me interested and encouraged over long years to pursue a variety of different topics within the field of theoretical and computational neurosciences. I have greatly profited from every exchange with him on so many levels and have been left nothing short of inspired from our discussions.

Within the scope of this work I have been granted a unique opportunity for a collaboration with the Hayashi Lab at the Department of Pharmacology, Kyoto University Graduate School of Medicine. My special thanks go out to Yasunori Hayashi, who has given me a warm welcome at his working group in Kyoto for over a year and supported me not only in numerous scientific discussions, but also with the challenges arising from a year in a foreign country. It is thanks to this collaboration that I was able to develop the methods and analysis presented in this work, that I could contribute to the publication presented within the scope of this thesis and that I could use the data recorded by members of the group, which provided the basis to the presented analyses.

Not only during my year in Kyoto, but also since then, my supervisor Kotaro Mizuta has played an important role in the development of this thesis. Countless discussions, lunch breaks and long nights at the institute allowed me to profit from his experience and knowledge in the field of experimental neurosciences. He has also been the one that has recorded or has supervised the recording of the mouse data used in this work in endless sessions. Thank you for your support!

I further want to thank Masaaki Sato, the main author of the paper presented in this work. In many discussions we have not only shaped the final outcome of this paper, but could also develop a bunch of ideas feeding into this work.

The process and story behind this work has also significantly been shaped by my colleagues, which have accompanied me in the two working groups on opposite sides of the world that I have grown to call home: Thank you to Risa Takamura, Hiroto Kojima, Ayaka Bota, Morgane Rosendale and the rest of the Hayashi lab for making my stay worthwhile and countless fun experiences. Thank you Julian Vogel, Nicolas Lenner, Matthias Häring, Charlotte Burmeister, Jonas Franz, Michael Sternbach and the rest of the Wolf Group in Göttingen for all the fun times and memorable Winter Seminars and PhD excursions.

A special thanks goes out to Chepe - Juan Daniel Flórez Weidinger - who got me excited about the hippocampus in the first place and got me started with countless discussions, programme

REFERENCES

code and methods at the very beginning of this journey.

It is impossible to work in the Wolf Group without extending a huge and heartfelt thanks to Christina von Behren, who is working tirelessly to allow everyone to fully concentrate on their scientific work. Thank you for your dedication and for everything that is being taken care of in the background!

Finally, I want to thank Maike Thomas, who has accompanied me on every step of the way to this thesis. It has been a long and sometimes tough ride, but you bore with me and supported me wherever possible and often when I didn't even know that I needed help. This work would not have been possible without you - Thank you so much for everything!