




CZECH TECHNICAL UNIVERSITY IN PRAGUE
Faculty of Nuclear Sciences and Physical Engineering



Structure of neural activity in the hippocampus during learning

Struktura neurální aktivity v hipokampu během učení

Bachelor's Degree Project

Author: **Martin Mašek**
Supervisor: **Dr.rer.nat. Karolína Korvasová**
Consultant: **doc. Ing. Tomáš Hobza, Ph.D.**
Academic year: 2023/2024

I. OSOBNÍ A STUDIJNÍ ÚDAJE

Příjmení: **Mašek** Jméno: **Martin** Osobní číslo: **509415**
Fakulta/ústav: **Fakulta jaderná a fyzikálně inženýrská**
Zadávající katedra/ústav: **Katedra matematiky**
Studijní program: **Matematické inženýrství**
Specializace: **Matematická informatika**

II. ÚDAJE K BAKALÁŘSKÉ PRÁCI

Název bakalářské práce:

Struktura neurální aktivity v hipokampu během učení

Název bakalářské práce anglicky:

Structure of neural activity in the hippocampus during learning

Pokyny pro vypracování:

1. Seznamte se se základními poznatky o hipokampu a s topologickými metodami analýzy dat, zejména strukturálním indexem [1] a UMAP [2] a dále s analýzou hlavních komponent.
2. Proveďte rešerši literatury týkající se konkrétní topologické struktury neurální aktivity v hipokampu.
3. Popište základní statistiku aktivity neuronů z hipokampu snímané pomocí fluorescenční mikroskopie, např. variabilitu aktivity jednotlivých neuronů, stabilitu přes jednotlivé experimenty apod.
4. Zkoumejte strukturu dat pomocí analýzy hlavních komponent a výše zmíněných topologických metod.
5. Popište vliv učení na topologickou strukturu dat.

Seznam doporučené literatury:

- [1] Sebastian, E. R., Esparza, J. and de la Prida, L. M. Quantifying the distribution of feature values over data represented in arbitrary dimensional spaces. bioRxiv, 2022
- [2] McInnes, L., Healy, J., & Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction (Version 3). arXiv. <https://doi.org/10.48550/ARXIV.1802.03426>
- [3] Sizemore, A. E., Phillips-Cremins, J. E., Ghrist, R. and Bassett, D. S. The importance of the whole: Topological data analysis for the network neuroscientist. In Network Neuroscience 3, 2019, 656–673.
- [4] Centeno, E. G. Z., Moreni, G., Vriend, C., Douw, L. and Santos, F. A. N. A hands-on tutorial on network and topological neuroscience. In Brain Structure and Function 227, 2022, 741–762.
- [5] Dabaghian, Y., Mémoli, F., Frank, L. and Carlsson, G. A Topological Paradigm for Hippocampal Spatial Map Formation Using Persistent Homology. PLoS Comput Biol 8(8), 2012, e1002581.
- [6] Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M. and Tanila, H. The Hippocampus, Memory, and Place Cells. Neuron 23, 1999, 209–226.

Jméno a pracoviště vedoucí(ho) bakalářské práce:

Dr.rer.nat. Karolína Korvasová Katedra softwaru a výuky informatiky MFF UK


Jméno a pracoviště druhé(ho) vedoucí(ho) nebo konzultanta(ky) bakalářské práce:


doc. Ing. Tomáš Hobza, Ph.D. katedra matematiky FJFI

Datum zadání bakalářské práce: **31.10.2023**

Termín odevzdání bakalářské práce: **05.08.2024**

Platnost zadání bakalářské práce: **30.09.2025**


Dr.rer.nat. Karolína Korvasová
podpis vedoucí(ho) práce


prof. Ing. Zuzana Masáková, Ph.D.
podpis vedoucí(ho) ústavu/katedry

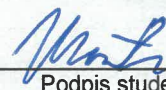

doc. Ing. Václav Čuba, Ph.D.
podpis děkana(ky)

III. PŘEVZETÍ ZADÁNÍ

Student bere na vědomí, že je povinen vypracovat bakalářskou práci samostatně, bez cizí pomoci, s výjimkou poskytnutých konzultací.
Seznam použité literatury, jiných pramenů a jmen konzultantů je třeba uvést v bakalářské práci.

27. 11. 2023

Datum převzetí zadání



Podpis studenta

Acknowledgment:

I would like to thank Dr.rer.nat. Karolína Korvasová for her expert guidance and express my gratitude to Dr. Sanja Bauer Mikulovic, Endre Levente Marosi and Pavol Bauer for their assistance in providing the data used in this thesis.

Author's declaration:

I declare that this Bachelor's Degree Project is entirely my own work and I have listed all the used sources in the bibliography.

Prague, August 5, 2024

Martin Mašek

Název práce:

Struktura neurální aktivity v hipokampu během učení

Autor: Martin Mašek

Studijní program: Matematické inženýrství

Specializace: Matematická informatika

Druh práce: Bakalářská práce

Vedoucí práce: Dr.rer.nat. Karolína Korvasová, Matematicko-fyzikální fakulta Univerzity Karlovy, Katedra softwaru a výuky informatiky

Konzultant: doc. Ing. Tomáš Hobza, Ph.D., Fakulta jaderná a fyzikálně inženýrská ČVUT, Katedra matematiky

Abstrakt: V této práci jsme zkoumali neuronální mechanismy kódování prostorových informací v hipocampu myši vystaveným apetitivnímu a averzivnímu podnětu. Naše zjištění ukazují, že aktivita place cells během tohoto úkolu vytváří stabilní, kruhovou reprezentaci prostoru, která však podléhá postupnému rozpadu v reakci na nový averzivní stimul. Ukazujeme, že tento rozpad je způsoben změnou referenčních bodů pro aktivitu place cells. Naše výsledky přispívají k pochopení neuronálních mechanismů reprezentace kognitivních map a ukazují, jak behaviorálně relevantní podněty mohou ovlivnit organizaci neuronálního kódování.

Klíčová slova: calcium-imaging, hipokampus, kognitivní mapy, prostorové kódování, prostorová reprezentace, topologická analýza dat

***Title:* Structure of neural activity in hippocampus during learning**

Author: Martin Mašek

Abstract: This thesis investigates the neural encoding of spatial information in the hippocampus of mice exposed to appetitive and aversive stimuli. Our findings confirm that activity of place cells during this task resides on a circular manifold, which remains stable in familiar environments but undergoes progressive breakdown when faced with novel aversive stimuli. We show that this breakdown results from a shift in place cell reference points. Our research extends our understanding of representations of cognitive maps, demonstrating how behaviorally relevant stimuli can induce a reorganization of neural encoding that prioritizes salient events over pure spatial mapping.

Key words: calcium-imaging, cognitive maps, hippocampus, spatial encoding, spatial representation, topological data analysis

Contents

Introduction	7
1 Literature review	9
1.1 The hippocampus	9
1.2 Imaging methods	10
1.2.1 Calcium imaging data preprocessing	11
1.3 Computational approaches	12
2 Methods	14
2.1 Experiment Description	14
2.2 Data Preprocessing	15
2.3 K-nearest neighbor	17
2.4 UMAP	17
2.4.1 Data	17
2.4.2 Calculation of UMAP	18
2.4.3 Initialization	20
2.4.4 Hyperparameters	21
2.4.5 Illustration	21
2.5 Structure Index	22
2.5.1 Calculation of the SI	23
2.5.2 Hyperparameters	23
2.5.3 Illustration	24
2.6 Rastermap	25
2.6.1 Calculation of the Rastermap	26
2.6.2 Hyperparameters	28
3 Results	29
4 Discussion	40
Conclusion	42

Introduction

Neuroscience has been instrumental in driving progress across various fields. Research on neural processes has produced important theoretical breakthroughs as well as practical applications in areas including artificial intelligence, medicine, and psychology.

For example, computational neuroscience has significantly advanced healthcare by providing insights into brain function and inspiring novel computational tools for medical applications (Wang et al. 2023). In machine learning, algorithms and architectures have been inspired by our brain's mechanisms since the field's inception (McCulloch and Pitts 1943) and continue to do so to this day. For instance, convolutional neural networks, which have revolutionized computer vision, were modeled after the mammalian visual cortex (Lindsay 2021).

Today, as AI systems become more complex, there is a growing need for innovative approaches, such as novel architectures or paradigms, to keep them sustainable and advance their capabilities (Ansar, Goswami, and Chakrabarti 2024). The brain's mechanisms for processing and encoding information could turn out to be a rich source of inspiration for further advances in machine learning and related fields (Schmidgall et al. 2023).

A major focus in computational neuroscience is dedicated to understanding how information is encoded in the brain. In this thesis, we will specifically deal with the problem of how spatial information is represented in the hippocampus and how it adapts to new environments.

One of the most articulated ideas in neuroscience today is that of a cognitive map (Tolman 1948). A cognitive map is a mental representation that allows an organism to acquire, code, store, recall, and decode information about the relative locations and attributes of phenomena in their environment (O'Keefe and Nadel 1978). One of the functions of cognitive maps is enabling animals, including humans, to navigate their environment by forming mental representations of spatial relationships. Continuous Attractor Networks (CANs) provide a computational framework to understand how such cognitive maps could be formed in the brain.

Specifically, CAN is a type of neural network in which connections are designed to keep the collective activity of the neurons within a limited, continuous range of patterns, such that the joint synaptic activity is constrained to a small set of possible co-activation patterns regardless of the various external stimuli (Hopfield 1982; Rolls 2010).

Recent papers have showed that this concept is particularly useful when studying neural circuits that handle spatial mapping, such as those in CA1 and CA3 locations in hippocampus and nearby areas (see literature review). These circuits are ideal for examining CAN dynamics because they encode information about space - a continuous and straightforward dimension to analyze. The data from these circuits are thus relative easy to interpret.

CAN models help to explain how these spatial variables are represented within the brain. These representations are visualized as existing on continuous manifolds such as rings (Chaudhuri et al. 2019) or tori (Gardner et al. 2022) in a low-dimensional spaces. The shape of this underlying structure varies depending on the type of information being encoded.

In this thesis, we collected data from mice running on a treadmill using calcium imaging and applied techniques from the field of topological data analysis on the neuronal time series to analyze its structure. We replicate previous studies (Sun et al. 2023) showing that the underlying structure of the space map of a mouse running on a track is encoded on a ring manifold. We also show that this structure remains stable in the absence of any external stimuli, confirming results from Wills et al. (2005) and Sun et al. (2023), but deteriorates when an external stimulus is introduced. At last, we put forward a possible explanations for the deterioration of the structure.

Chapter 1

Literature review

This literature review traces the evolution of our understanding of the hippocampus and its functions, particularly in the context of spatial cognition and the formation of cognitive maps. We will highlight key discoveries that have shaped our knowledge of hippocampal function and its role in spatial navigation. Additionally, we will describe the development of imaging methods used to record neuronal activity, with emphasis on those employed in our study. Lastly, we will discuss some recent applications of techniques from the field of topological data analysis in neuroscience.

1.1 The hippocampus

The human hippocampus is a sea-horse shaped structure that forms part of the limbic system. Anatomically, it lies above the subiculum and medial parahippocampal gyrus, forming a curved elevation approximately 5 cm long along the floor of the inferior horn of the lateral ventricle. The hippocampus proper can be divided into three distinct fields: CA1, CA2, and CA3. Each of these regions has unique cellular characteristics and connections. The hippocampal formation, a broader term, includes not only the hippocampus itself but also several closely related structures. The formation includes the dentate gyrus, hippocampus, subicular complex and entorhinal cortex. Phylogenetically, the hippocampus is a relatively old part of the brain, exhibiting functional similarities across all mammals (Mancall, Brock, and Gray 2011; Clark and Larry R. Squire 2013). The hippocampus is a critical brain region implicated in a wide range of cognitive and physiological processes. It plays a pivotal role in learning and memory (Bird and Neil Burgess 2008; L. R. Squire 1992), spatial navigation (Eichenbaum 2017), emotional regulation (Qasim et al. 2023), and even hypothalamic function (Bang et al. 2022). Understanding the hippocampus is essential as it is also implicated in several neurological disorders, such as Alzheimer's disease (Rao et al. 2022).

In 1971, O'Keefe along with his student Jonathan Dostrovsky, discovered new type of cells when they noticed that certain neurons in the hippocampus were activated only when a rat was in a specific position location in its environment. These neurons would fire consistently in that particular area, leading O'Keefe to the conclusion that these cells formed a map-like representation of the environment in the rat's brain. O'Keefe termed these neurons "place cells" because each seemed to correspond to a specific place in the environment. The discovery of place cells provided the first neural basis for a cognitive map in the brain (O'Keefe and Dostrovsky 1971).

Building on this, Muller and Kubie (1987) demonstrated that place cells maintain their relative spatial firing patterns when familiar environmental cues are rotated, indicating a strong link between external landmarks and internal representations. Moreover, they found that changing the arena's shape from circular to rectangular led to unpredictable alterations in firing patterns, suggesting distinct cognitive

maps for different environmental geometries. This discovery laid the groundwork for the concept of remapping. Expanding on these findings, Bostock, Muller, and Kubie (1991) explored the plasticity of place cell representations in response to subtle environmental alterations. Their work revealed that even minor changes could trigger significant shifts in place cell firing patterns, highlighting the hippocampus's mechanism for distinguishing between similar yet distinct spatial contexts. This flexibility enables the creation of unique neural signatures for each environment, underlining the hippocampus's critical role in forming and updating cognitive maps. This dynamic nature of these representations allows for efficient encoding of spatial information in different contexts.

Formalizing these observations, O'Keefe and N. Burgess (1996) introduced the concept of place cell remapping as a fundamental mechanism for encoding different environments. By systematically examining changes in place cell firing fields in response to environmental modifications, they solidified the idea of multiple, distinct cognitive maps within the hippocampus.

Expanding on the concept of place cells, subsequent research explored the temporal patterns of hippocampal activity. A critical discovery in this domain was the phenomenon of theta phase precession. O'Keefe and Recce (1993) among others (Skaggs et al. 1996) demonstrated that the phase of place cell spiking relative to theta rhythm, a neural oscillation in the hippocampus (Green and Arduini 1954), shifts to earlier phases of the theta phase as the rodent runs through a place field. This provided strong evidence for temporal coding of neural circuits by describing a possible mechanism behind dynamic updating of cognitive maps by integrating spatial and temporal aspects.

Nowadays, theta precession has also been shown to play a broader role in cognitive function. Recently, for instance, (Zheng et al. 2023) demonstrated that theta precession is a general coding mechanism supporting several aspects of memory, by observing changing strength of the theta precession in humans while watching short movie clips and found its correlation with memory retrieval.

Another breakthrough in understanding how cognitive maps are represented in the brain was made by nobel prize winning authors May-Britt and Edvard Moser in 2005 at the Norwegian University of Science and Technology with the discovery of grid cells. Grid cells generate a triangular grid of firing fields that cover the entire environment. Unlike place cells, which activate in specific locations, grid cells activate in multiple locations that form a hexagonal grid pattern across the environment (Hafting et al. 2005).

The formation of spatial maps in the brain, particularly within regions like the dorsal hippocampus (dHPC), is significantly influenced by reward or other emotionally rich stimuli, which serve as critical motivational cues for spatial learning and memory (Barnstedt, Mocellin, and Remy 2024; Sosa, Plitt, and Giocomo 2024). The dHPC is well-known for its role in encoding spatial memories, such as navigating towards a learned reward location. When an individual receives a reward or encounters emotionally charged stimuli, these events can enhance the encoding of spatial information and the subsequent retrieval of these memories. This is because the reward-related signals from the dHPC are transmitted to areas like the nucleus accumbens (NAc), a region implicated in value-based action selection. This suggests that how spatial maps look is affected by motivational significance. These findings suggest that spatial representations in the brain are modulated by motivational significance. Neural spatial maps encode not just the physical layout of an environment but also its relevance based on emotional context. This integration of spatial and motivational information guides effective decision-making and action selection in a goal-directed manner.

1.2 Imaging methods

Neural recording techniques have improved substantially over the past 50 years. The amount of simultaneously recorded neurons approximately doubled every 7 years (Stevenson and Kording 2011).

According to Nicolelis (2008) (as cited in Stevenson and Kording (2011)), this has been made possible by advances in silicon processing techniques among other rapid computer hardware development such as data transfer speeds or storage capacity.

Traditional single-electrode extracellular recordings offered limited capacity, typically capturing the activity of only a few neurons. This approach has been surpassed by the emergence of multielectrode recording systems. These advanced technologies, as described by Kelly et al. (2007), allow for the simultaneous recording of activity from hundreds of neurons, significantly expanding the scope of neural data acquisition.

Alternative technique for recording activity from large neuronal populations is calcium imaging. This method offers great spatial and temporal resolution, allowing us to visualize individual neurons and capture their temporal dynamics with high precision (Giovannucci et al. 2019). Modern instruments allow us to monitor populations of up to thousands of neurons in various regions of the brain both *in vitro* and *in vivo* (Grienberger and Konnerth 2012; Nelson and Bonner 2021).

Calcium plays a critical role in neuronal signaling. When a neuron fires, calcium ions surge into the cell. By monitoring these changes in calcium levels using fluorescent calcium indicators, which emit light when they bind to calcium ions, we can indirectly observe neuronal activity (Grienberger and Konnerth 2012)).

The most common approach for calcium imaging is two-photon calcium imaging. Here, two lower-energy photons are absorbed by the fluorescent molecule, triggering it to emit light. This differs from methods like one-photon calcium imaging, which utilizes a single, higher-energy photon. In most cases, two-photon excitation offers superior depth penetration, leading to enhanced spatial resolution within the tissue (*ibid.*).

Calcium imaging data, while powerful, requires proper preprocessing to accurately extract neuronal firing rates. Several software libraries have been developed for this purpose, with CaImAn being a popular Python package offering advanced analysis tools. CaImAn (Giovannucci et al. 2019) employs a multi-step process to transform raw imaging data into time series representing neural activity. These steps encompass motion correction, source extraction and activity deconvolution.

1.2.1 Calcium imaging data preprocessing

As our research leverages calcium imaging data, a basic understanding of the methodology, including its preprocessing steps, is crucial for interpreting the results and drawing valid conclusions. This section will provide an overview of the main preprocessing steps typically used to transform raw fluorescence images into neural time series.

The initial step involves motion correction, addressing non-rigid motion artifacts. These artifacts can manifest as spatial shifts within the field of view (FOV) due to factors like breathing or movement. CaImAn utilizes the NoRMCorre algorithm (Pnevmatikakis and Giovannucci 2017) to tackle this issue. NoRMCorre estimates motion vectors with subpixel resolution from overlapping image patches within the FOV. These vectors are then used to infer a smooth motion field for each frame, crucial for subsequent data analysis.

Following motion correction, CaImAn performs source extraction step to identify individual neurons. This process leverages the constrained non-negative matrix factorization (CNMF) framework of (Pnevmatikakis, Soudry, et al. 2016). CNMF effectively separates individual cells from background noise, resulting in the identification of regions of interest (ROIs). These ROIs represent the specific areas within the image corresponding to individual neurons, allowing researchers to infer actual cellular activity.

Finally, CaImAn performs temporal deconvolution of the raw signal. This deconvolution process aims to infer the underlying spiking activity of the neurons, providing a more refined representation

of neuronal firing patterns. This is achieved using sparse non-negative deconvolution as described in (Vogelstein et al. 2010).

1.3 Computational approaches

The exponential growth of neuron recording capabilities had to be reflected in the way the data was being analysed (Stevenson and Kording 2011; Ann Sizemore Blevins and Bassett 2020). To interpret the complex information encoded in the activity of large neuronal ensembles, we need advanced decoding methods.

Ann S. Blevins et al. (2022) names several approaches utilized to deal with big amounts of data in neuroscience. Dimensionality reduction techniques, regression methods, cluster and non-linear analyses or topological approaches all allow analysis from different points of view. All these tools can ultimately be used to understand how the brain encodes information.

In this thesis, we will be concerned with methods based on ideas from topological data analysis (TDA). TDA is an emerging field that applies concepts from topology to extract meaningful insights from complex datasets. It focuses on understanding the shape and structure of data, rather than just its numerical properties. TDA uses techniques to identify and quantify topological features such as loops, holes, and clusters in high-dimensional data. These methods are particularly useful for analyzing large, noisy datasets where traditional statistical approaches might struggle. TDA offers several advantages, including robustness to noise, coordinate-free analysis, and the ability to capture global structure from local information. (ibid.)

TDA gained prominence in the early 2000s, with Mapper (Singh, Memoli, and Carlsson 2007) being one of the first widely recognized tools. Mapper’s successful application in identifying a new subset of breast cancer demonstrated the potential of TDA in real-world problems. This method creates a graph representation of high-dimensional data by clustering points in low-dimensional space into nodes and connecting them with edges based on their high-dimensional relationships. Following Mapper, several other TDA approaches were developed and popularized. Persistent homology (Edelsbrunner, Letscher, and Zomorodian 2002), which predates Mapper but gained wider recognition later, is a tool for quantifying topological features such as loops or holes in point cloud data across different scales. More recent additions to the TDA toolkit include Uniform Manifold Approximation and Projection (McInnes, Healy, and Melville 2020), a dimensionality reduction technique that preserves both local and global topological structure, and Structure Index (Sebastian, Esparza, and Prida 2024), both of which are utilized in this thesis.

The core concept lies in representing the system as a simplicial complex. This complex is built from nodes and simplices that capture how these nodes interact. Notably, unlike edges in a network that connect only two nodes, simplices can connect any number of them. (Ann S. Blevins et al. 2022)

Another fundamental concept here is that of homology. Homology provides a mathematical framework for analyzing the topological features of spaces or simplicial complexes. It uses vector spaces called homology groups (H_k) to represent and quantify K-dimensional so called holes - features of the topological space. The dimension of each homology group corresponds to the number of independent features of that dimension. For example, H_0 represents connected components, H_1 represents loops, and H_2 represents cavities. Higher-dimensional homology groups capture more abstract topological features. This algebraic approach allows for a systematic and quantitative analysis of shape and structure, making it particularly valuable in topological data analysis for uncovering hidden patterns and relationships in complex, high-dimensional datasets (Chazal and Michel 2021).

The dimensions of these homology groups are known as Betti numbers, which characterize different shapes and structures. For instance, a ring is characterized by the Betti number sequence (1,1,0), indi-

cating one connected component, one loop, and no cavities. A torus, on the other hand, is represented by $(1,2,1)$, signifying one connected component, two independent loops, and one cavity. Betti numbers thus uniquely identify different shapes based on their topological characteristics.

As Sizemore et al. (2019) states, topological methods will likely help elucidate data from systems characterized by:

- (a) complex, higher-order relationships among data points.
- (b) topologically meaningful cavities or loops.

With the discovery of place cells and grid cells grew the interest to further explore how these cells contribute to the brain's ability to the formation of cognitive maps. Authors Leutgeb et al. (2005) and Gothard, Skaggs, and McNaughton (1996) expanded on the previous findings about place field flexibility in changing 2D environments. Specifically, when the environment undergoes deformation, the place fields can alter in shape, size, and location due to remapping, yet they preserve crucial spatial relationships such as overlaps, adjacencies, and containments. Consequently, the firing sequences of the place cells during animal navigation remain mostly constant across a range of geometric transformations. This consistency supports the idea that the hippocampus provides a topological, rather than topographical, representation of space—similar to a subway map's depiction of routes rather than a detailed street map, as noted by Dabaghian (2021).

These insights motivated further use of topological techniques for analysis of cognitive maps and their representations in the brain. The manifolds discovered through these topological approaches could potentially represent the continuous attractors underlying the encoded information.

For instance, an important observations were made by Gardner et al. (2022) who employed simultaneous recordings from grid cells within the medial entorhinal cortex to reveal that the collective activity of grid cells organizes itself on a toroidal manifold. Using topological techniques (see UMAP section in methods), they demonstrated that different positions on the torus directly correspond to the animal's physical locations within its environment. Remarkably, the grid cells maintained their specific activity patterns consistently across different environments and between states of wakefulness and sleep.

Recently, Sun et al. (2023) utilized TDA to describe the formation of a cognitive map of a mouse running on a treadmill learning a novel task. First, they introduced new visual cue pairs to replace original indicators of reward on the treadmill, finding that mice learned these new cues significantly faster than the original set. Second, they occasionally presented "stretched trials" where the distance to reward zones was extended. During these trials, mice initially showed anticipatory licking at the usual reward locations, despite the altered environment. Neural recordings revealed that place cells adapted their firing patterns to these changes, with different responses observed for near and far trials. By applying UMAP, they showed that the cognitive representation gradually adapted to account for both short and stretched trials.

Chapter 2

Methods

2.1 Experiment Description

The experiment was conducted at the Leibniz Institute for Neurobiology in Magdeburg by Dr. Sanja Mikulovic and Endre Marosi. Data preprocessing was performed as described in the literature review chapter with additional assistance from Pavol Bauer.

The experimental apparatus consisted of 7 cm wide and 360 cm long textile belt including six differently textured zones: horizontal and vertical glue stripes, glue dots, Velcro dots, vertical tape stripes and upright nylon spikes. In this thesis, we employ a twilight color gradient to represent the position on the running track for improved clarity. The reward zone for imaging experiments was 30 cm long and was placed between the end of the horizontal glue stripes and the beginning of the vertical tape stripes (Figure 2.1)

Four mice, labeled mouse 1, 2, 3 and 4, were subjected to a two-week habituation period involving various handling procedures. Subsequently, they were trained to perform a task on the slowly moving treadmill apparatus. The task required mice to lick a metal spout located in a hidden reward zone on the treadmill belt to receive a liquid reward. After dispensing the reward, the mouse had to run another full lap for it to be dispensed again. This process repeated for 300 seconds per session.

Custom Python software managed reward delivery and task monitoring. Calcium activity and behavioral markers, including treadmill position, velocity, and licking behavior, were recorded following the habituation period. A detailed description of the experimental apparatus can be found in Barnstedt, Mocellin, and Remy (2024).

The experiment was conducted in two phases. During the initial three days of the recording, mice performed the previously learned task of running on a treadmill to reach a fixed reward location. In the subsequent three days, a novel obstacle was introduced in the form of a treadmill brake. This brake was activated over a specific treadmill distance (location 260cm-324cm), slowing the belt making it more difficult for the mice to run. Overall, we were working with 24 recordings.

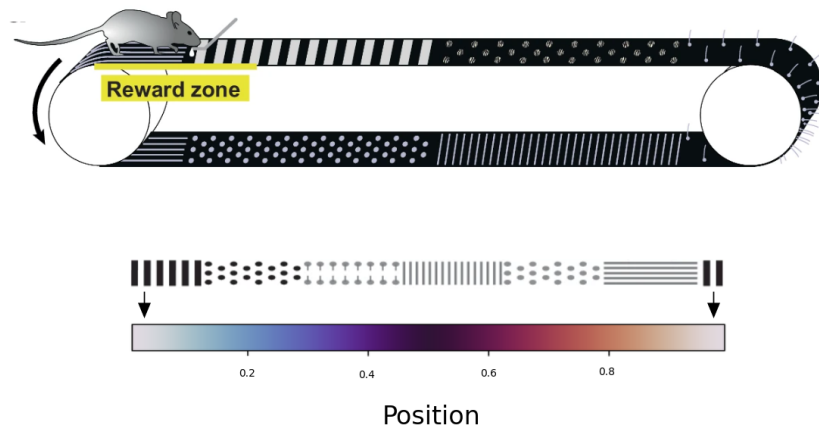


Figure 2.1: Running track layout. This color mapping provides a visual correspondence between locations on the track and their respective positions within the experimental setup. (Adapted from Barnstedt, Mocellin, and Remy (2024). Licensed under CC BY 4.0. <https://doi.org/10.1038/s41467-024-47361-x>. Modified from the original.)

2.2 Data Preprocessing

Neural data were originally preprocessed using the CaImAn pipeline as described in Literature review chapter. Following this, we have performed additional preprocessing steps. The CaImAn preprocessed dataset consisted of 300-second recordings sampled at 80 Hz, containing only deconvolved neural activity (spikes). The number of neurons recorded varied across sessions, with a maximum of 382, a minimum of 53, a mean of 221.5, and a median of 200.5 neurons. To standardize the data, all neural activity was z-scored.

Behavioral data was collected concurrently with the neural recordings. The treadmill position signal, originally ranging from 0 to 175, was normalized to values between 0 and 1, corresponding to a distance of 0 to 360 cm. We conducted a separate analysis of periods of active running behavior. To focus on these intervals, we excluded timepoints where the mouse's speed fell below 10 cm/s from our analysis.

To reduce data dimensionality and enhance computational efficiency, we implemented both temporal and spatial binning techniques. For temporal binning, we aggregated sets of 50 consecutive samples (0.625s) and calculated their mean values. For spatial binning, the treadmill positions were binned into 30 bins of 12 cm length. We then computed the mean value for each of these segments and concatenated the results across all laps.

All preprocessing was performed in Python.

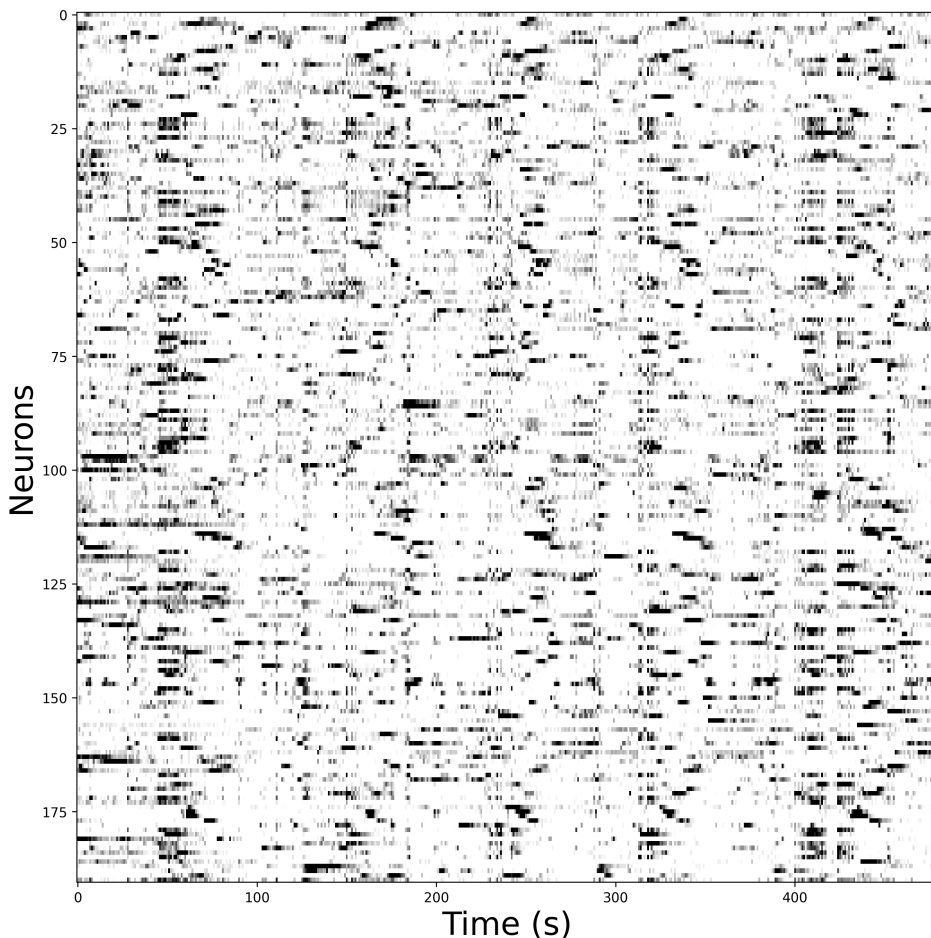


Figure 2.2: The heatmap visualizes the activity of recorded neurons. Binned time is displayed on the X-axis (480 bins), spanning the entire recording. The Y-axis represents individual neurons, with each row depicting a single neuron’s activity. This heatmap allows for a quick overview of how individual neurons’ firing rates change over time. Data displayed is from a single representative session (mouse 1, day 3, reward session)

Computational Techniques

In this section, we will describe the computational methods employed in this thesis. We begin with an explanation of the K-Nearest Neighbors (KNN) algorithm, as it is a key component in the computation of all topological methods used in this study. Following this, we will detail three main techniques. First, we’ll explore Uniform Manifold Approximation and Projection (McInnes, Healy, and Melville 2020), a manifold learning and visualization technique. Next, we’ll discuss the Structure Index (Sebastian, Esparza, and Prida 2024), a topological method for quantifying structure in point cloud data. Lastly, we’ll examine Rastermap (Stringer et al. 2023), a method for sorting and clustering neurons.

2.3 K-nearest neighbor

The KNN algorithm is a fundamental supervised learning method for both classification and regression tasks. It operates in a non-parametric way, meaning it makes no assumptions about the underlying data distribution. Here's a breakdown of how KNN works:

1. **Distance Calculation:** Given a new data point (x), KNN first calculates its distance to all points in the training dataset. This distance can be measured using various metrics like Euclidean distance, Manhattan distance, or cosine similarity.
2. **K -Nearest Neighbors Identification:** The algorithm then identifies the k closest data points (neighbors) to x based on the chosen distance metric. The value of k is a hyperparameter that needs to be tuned for optimal performance.
3. **Prediction:** KNN predicts the class label of (x) by performing a majority vote among the class labels of its (K) nearest neighbors. The class label that appears most frequently among the neighbors becomes the predicted class for (x).

2.4 UMAP

Uniform Manifold Approximation and Projection (UMAP) is a dimensionality reduction technique used to represent high-dimensional data in a lower-dimensional space. It is particularly useful for exploratory data analysis, clustering, and pattern recognition tasks.

UMAP is grounded in the ideas of category theory, a branch of mathematics that deals with abstract structures and relationships between them. Category theory provides a framework for understanding and formalizing the notion of similarity or equivalence between objects. UMAP uses these mathematical foundations to capture the intrinsic structure of data in a lower-dimensional space.

From a computational standpoint, UMAP is essentially a graph-based method. It constructs a high-dimensional graph representation of the data, where each data point connects to its nearest neighbors. This local connectivity captures the local structure of the data, which is crucial for preserving its inherent relationships. UMAP then defines an objective function that preserves desired characteristics of the constructed graph and finds a low-dimensional representation which optimizes this objective function.

2.4.1 Data

In summary, UMAP makes three key assumptions about the data:

1. **Manifold structure:** UMAP assumes that the data lies on a manifold, a lower-dimensional structure that locally resembles a Euclidean space near each point, embedded within the higher-dimensional space. While the data may not be inherently uniformly distributed across this manifold, UMAP assumes that a metric can be defined such that the data would be approximately uniformly distributed with respect to that metric. This means that any region of fixed volume within the data space should contain roughly the same number of points, ensuring a balanced representation across different areas.
2. **Local connectivity:** The data on the underlying manifold is assumed to be locally connected. This assumption allows UMAP to represent nearby or related points in high-dimensional space as neighbors in the low-dimensional space, capturing meaningful patterns and relationships within the data.

3. Topological preservation: UMAP's primary goal is to preserve the topological structure of the manifold. This means maintaining the overall shape and connectivity of the data, rather than focusing solely on preserving exact distances between points.

2.4.2 Calculation of UMAP

UMAP falls under the category of algorithms based on KNN graphs. This means understanding UMAP's inner workings can be done by looking at weighted graphs. Several dimensionality reduction techniques utilize KNN graphs. However, UMAP offers distinct advantages that make it particularly suitable for our analysis.

One key advantage is its ability to handle the "curse of dimensionality." In high-dimensional data, traditional distance metrics become unreliable. UMAP addresses this issue by focusing on preserving the relationships between neighboring points in the local area. This focus on local similarities makes UMAP effective even when dealing with data in very high dimensions.

Another advantage of UMAP is its ability to capture both local and global structures within the data. This is in contrast to some methods like t-SNE (Maaten and Hinton 2008), which primarily focus on local structure. For tasks where understanding the overall organization of the data is important, UMAP's ability to preserve both local and global information proves to be crucial. These advantages make UMAP a compelling choice for our analysis.

Constructing the low-dimensional representation involves two key stages: graph construction and graph layout.

Graph Construction

This stage establishes a weighted KNN graph that captures the relationships between data points. Let $X = x_1, \dots, x_n$ be the input data points, $d : X \times X \rightarrow \mathbb{R}_0^+$ a metric, and $k \in \mathbb{N}$. The process involves the following steps:

1. Identifying nearest neighbors: For each data point x_i , compute the set x_{i_1}, \dots, x_{i_k} of k nearest neighbors under the metric d_i .
2. Assigning edge weights while ensuring local connectivity: For each x_i , define

$$\rho_i = \min\{d_i(x_i, x_{i_j}) \mid 1 \leq j \leq k, d_i(x_i, x_{i_j}) > 0\},$$

and set σ_i such that

$$\sum_{j=1}^k \exp\left\{\frac{-\max(0, d_i(x_i, x_{i_j}) - \rho_i)}{\sigma_i}\right\} = \log_2(k).$$

This ensures that x_i connects to at least one other data point with an edge of weight 1, satisfying the local connectivity constraint. The max function ensures that the exponential term does not exceed 1.

3. Constructing the Directed Weighted Graph: Define a weighted directed graph $\overline{G} = (V, E, w)$ where:
 - Vertices V are the set X
 - Edges $E = \{(x_i, x_{i_j}) \mid 1 \leq j \leq k, 1 \leq i \leq N\}$

- Weight function w is defined as:

$$w((x_i, x_{ij})) = \exp \left\{ \frac{-\max(0, d(x_i, x_{ij}) - \rho_i)}{\sigma_i} \right\}$$

4. Symmetrizing the Graph: Unfortunately, the resulting graph is not necessarily symmetric. This asymmetry stems from UMAP's theoretical foundations, which assume a uniform distribution of the data on the manifold. To satisfy this assumption, the algorithm employs locally adaptive distance metrics, where the scale of distances is adjusted for each point based on its local neighborhood. As a result, the perceived distance from x_i to x_j may not equal the distance from x_j to x_i . This local adaptation of the metric leads to direction-dependent edge weights in the graph, reflecting the varying local geometries around each point.

Let A be the weighted adjacency matrix of \bar{G} . The symmetric matrix B is given by:

$$B = A + A^\top - A \circ A$$

where \circ is the pointwise product. Note that B_{ij} can be interpreted as the probability that at least one directed edge exists between x_i and x_j in either direction.

The final UMAP graph G is an undirected weighted graph with adjacency matrix B .

Graph Layout

The goal of this stage is to find a low-dimensional representation that preserves the relationships between data points captured in the KNN graph. To achieve this, UMAP utilizes a force-directed graph drawing algorithm, which iteratively applies a set of attractive and repulsive forces among vertices. This process amounts to a non-convex optimization problem, with convergence to a local minimum ensured by gradually diminishing both the attractive and repulsive forces.

Attractive forces pull connected points closer based on their weights in the KNN graph. This encourages points with stronger connections to be positioned closer in the low-dimensional space. In addition to the attractive forces, repulsive forces are used to push points apart, preventing overcrowding in the low-dimensional space. To manage computational complexity, these repulsive forces are approximated through sampling. In particular, when an attractive force is applied to an edge, one of its vertices is simultaneously repelled by a random sample of other vertices.

The forces are derived from gradients that minimize the edge-wise cross-entropy between the weighted KNN graph G and a low-dimensional equivalent graph H constructed from the embedded points y_i . This cross-entropy is defined as:

$$C((A, \mu), (A, \nu)) = \sum_{a \in A} \mu(a) \log \left(\frac{\mu(a)}{\nu(a)} \right) + (1 - \mu(a)) \log \left(\frac{1 - \mu(a)}{1 - \nu(a)} \right)$$

where μ and ν represent the edge weights in G and H respectively.

The optimization aims to position points y_i such that the induced weighted graph H closely approximates the topology captured by the source data in G .

The attractive and repulsive forces are derived from the cross-correlation, by calculating gradients of logarithms of the smooth approximation function of the edge strength $\Phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$, defined as:

$$\Phi(\mathbf{x}, \mathbf{y}) = \left(1 + a \left(\|\mathbf{x} - \mathbf{y}\|_2^2 \right)^b \right)^{-1}$$

The parameters a and b are derived by fitting non-linear least squares against the membership strength function $\Psi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$, where

$$\Psi(\mathbf{x}, \mathbf{y}) = \begin{cases} 1 & \text{if } \|\mathbf{x} - \mathbf{y}\|^2 \leq \text{min-dist} \\ \exp(-(\|\mathbf{x} - \mathbf{y}\|^2 - \text{min-dist})) & \text{otherwise} \end{cases}$$

$\Psi(\mathbf{x}, \mathbf{y})$ is used to define the desired properties of the low-dimensional embedding, which $\Phi(\mathbf{x}, \mathbf{y})$ approximates.

The resulting forces have the following forms:

1. Attractive Force:

$$F_A(y_i, y_j) = \frac{-2ab\|\mathbf{y}_i - \mathbf{y}_j\|_2^{2(b-1)}}{1 + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2} w((x_i, x_j))(\mathbf{y}_i - \mathbf{y}_j)$$

where $w((x_i, x_j))$ is the edge weight from the KNN graph.

2. Repulsive Force:

$$F_R(y_i, y_j) = \frac{2b}{(\epsilon + \|\mathbf{y}_i - \mathbf{y}_j\|_2^2)(1 + a\|\mathbf{y}_i - \mathbf{y}_j\|_2^{2b})} (1 - w((x_i, x_j)))(\mathbf{y}_i - \mathbf{y}_j)$$

where ϵ is a small number to prevent division by zero.

Finally, the new coordinates of the points are updated as:

$$y_{i+1} = y_i + \eta \left(\sum_{j \in \text{neighbors}(i)} F_A(y_i, y_j) - \sum_{k \in \text{random sample}} F_R(y_i, y_k) \right)$$

where η is the learning rate.

2.4.3 Initialization

:

While the algorithm can be initialized randomly, in practice we employ spectral layout for the initialization. The symmetric Laplacian of the graph G serves as a discrete approximation of the Laplace-Beltrami operator on the manifold represented by the graph. Using this spectral information to guide the initialization process enhances both the convergence speed and stability of the algorithm by capturing important structural information about the graph before the optimization begins.

To construct the Laplacian matrix of graph G , follow these steps:

1. Start with the Adjacency Matrix \mathbf{A} of graph G , where $A_{ij} = 1$ if nodes i and j are connected, and 0 otherwise.
2. Construct the Degree Matrix \mathbf{D} , a diagonal matrix where:

$$D_{ii} = \sum \text{ of row } i \text{ in } \mathbf{A}$$

$$D_{ij} = 0 \quad \text{for } i \neq j$$

3. Calculate the symmetric normalized Laplacian \mathbf{L} using the formula:

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$$

where \mathbf{I} is the identity matrix, and $\mathbf{D}^{-1/2}$ is the inverse square root of \mathbf{D} .

4. Compute k eigenvectors corresponding to the lowest (non-zero) eigenvalues of \mathbf{L} , where k is the target dimension of the UMAP embedding.

2.4.4 Hyperparameters

UMAP has four main hyperparameters:

1. **Number of nearest neighbors:** This represents a trade-off between preserving local (fewer neighbors) and global (more neighbors) structure in the data.
2. **Target dimension:** The dimension we want to embed our pointcloud in.
3. **Minimum Distance:** Minimum distance between points in the low-dimensional representation. It controls how tightly UMAP is allowed to pack points together with smaller values allowing for tighter clusters.
4. **Number of epochs:** Number of training epochs to use in the optimization.

UMAP analysis was performed using the official UMAP implementation with the following parameters: 15 nearest neighbors to balance local and global structure preservation, a 3-dimensional target space for visualization and interpretation, a minimum distance of 0 to allow for tight clustering, and the default 500 epochs for optimization.

2.4.5 Illustration

To illustrate UMAP's ability to reveal hidden structures in complex data, we constructed a synthetic dataset based on sine and cosine functions. Originally parameterizing perfect circles, these functions were masked by adding random noise and slight amplitude variations. We also add entirely random data to obfuscate the sine-cosine information even further. This results in a high-dimensional matrix where the circular patterns are effectively hidden. UMAP successfully recovered the underlying circular patterns from this high-dimensional, noisy data. By color-coding the projected points according to their original sine function values, we visually confirmed UMAP's ability to group data points corresponding to similar phases of the circular trajectory. (Figure 3)

The sine and cosine wave formulas with random amplitude $A \sim \mathcal{U}(0.5, 1.5)$, fixed frequency $f = 1.0$, and random noise $N \sim \mathcal{N}(0, \sigma^2)$ where $\sigma \sim \mathcal{U}(0.1, 0.3)$ can be represented as:

$$A \sin(2\pi ft) + N$$

$$A \cos(2\pi ft) + N$$




Figure 2.3: UMAP embedding of noisy sine/cosine data. Points are colored by their original sine function values, demonstrating the recovery of the underlying circular structure.

2.5 Structure Index

In neuroscience research, we frequently encounter the challenge of identifying how a specific feature is distributed across a point cloud. This arises in various contexts, such as investigating representations within neural manifolds (Barnstedt, Mocellin, and Remy 2024), analyzing neurophysiological signals (Rutkowski, Komendziński, and Otake-Matsuura 2024), and performing anatomical image segmentation (Momota et al. 2024).

The Structure Index (SI) addresses this challenge. It is a directed graph-based metric designed to quantify the distribution of feature values across arbitrary D -dimensional spaces within a point cloud. SI focuses on overlapping distributions of data points that share similar feature values within a defined local neighborhood. By analyzing these overlaps, SI allows us to identify structures, patterns, and relationships within the dataset. This approach enables us to explore how different features manifest and vary across the point cloud data. This is particularly relevant in neuroscience, where mapping features onto a point cloud can provide valuable insights into how the brain encodes and represents information.

While embedding a point cloud into a lower-dimensional space allows for visualization using dimensionality reduction techniques like UMAP, additional methods are necessary for effective comparison and accurate quantification of feature distributions and point cloud shapes across different datasets. SI offers a complementary approach specifically designed to analyze these distribution patterns within the original high-dimensional space.

2.5.1 Calculation of the SI

The aim of the SI is to quantify structure in the distribution of a feature over a point cloud. Consider a point cloud P embedded in a space of dimension D and a set of feature values f mapped onto P , $F = \{f_p | p \in P\}$.

The calculation of the SI consists of several steps:

1. Create a finite covering $B = \{B_i\}_{i \in \mathbb{N}}$ of a point cloud P by splitting it into bin-groups B_i according to their feature value such that $B_i = \{p \in P | t_i \leq f_p < t_{i+1}\}$, where $P = \{B_1, \dots, B_n\}$ and $B_i \cap B_j = \emptyset$ when $i \neq j$. Features can be both continuous or discrete.
2. Compute the overlap between each pair of bin-groups in terms of the K -nearest neighbors. Given two bin-groups, U and V , define an overlap score from U to V , $OS_{U \rightarrow V}$, as the ratio of k -nearest neighbors of all the points of U that belong to V in the point cloud space. That is,

$$OS_{U \rightarrow V}(k) = \frac{1}{|U|k} \sum_{u \in U} |\{N_u^j(U \cup V - \{u\}) | j = 1, \dots, k\} \cap V| \quad (2.1)$$

where $N_u^j(U \cup V)$ denotes the set of k -nearest neighbors of point u in the union of U and V .

3. Compute the overlap score for each pair of bin groups B_i and B_j , which yields an adjacency matrix $M_{n \times n}$, such that $M_{ij} = OS_{B_i \rightarrow B_j}$. Matrix M can be interpreted as a weighted directed graph where nodes are bin groups and edges represent overlap relationships. The diagonal elements of the matrix are set to zero, that is, $OS_{U \rightarrow U}(k) = 0$, indicating no self-overlap.
4. Finally, Structure Index SI is defined as 1 minus the mean weighted out-degree of the nodes after scaling:

$$SI(M) = 1 - \left(\frac{2}{n^2 - n} \sum_i^n \sum_j^n M_{i,j} \right) \quad (2.2)$$

To provide complementary quantification of the overlap patterns, we developed a method to summarize the overlap scores for statistical comparison across experimental days. Starting with each mouse's overlap matrix, we computed the sum of the matrix M with its transpose ($M + M^T$). From this symmetrized matrix, we extracted the upper triangular portion, excluding the main diagonal, and summed the elements of each diagonal separately. This process yielded an array of diagonal sums for each mouse, with each sum representing a different level of overlap. We then aggregated these arrays across all mice for each experimental day by element-wise summation. The resulting arrays, one for each day, were used to generate boxplots, providing a visual representation of the overlap score distributions.

2.5.2 Hyperparameters

Similar to UMAP, the SI utilizes hyper-parameters to influence its behavior:

1. **Number of neighbors:** Increasing the number of neighbors generally improves the algorithm's ability to capture global structures within the data. A larger neighborhood provides more context, enabling more accurate assessment of feature value distributions across the point cloud.
2. **Number of bins:** This parameter determines the granularity of the analysis. A higher number of bins allows for capturing finer details in the structure. However, it's essential to maintain a reasonable number of data points within each bin to prevent the SI from being unduly influenced by sparsely populated bins. Ideally, a relatively even distribution of data points across bins is preferred.

SI calculation was performed using the official library released by the authors of the SI paper. For our calculations, we used 10 nearest neighbors and 10 bins. These parameters balanced the capture of local and global structures while maintaining sufficiently populated bins for robust analysis.

2.5.3 Illustration

For illustration we generate a point cloud uniformly distributed on the surface of a sphere. By applying different labeling schemes, we achieve different values of SI. (Figure 4)

Using spherical coordinates (r, θ, ϕ) , where:

- r : constant radius of the sphere
- θ : azimuthal angle, $0 \leq \theta < 2\pi$
- ϕ : polar angle, $0 \leq \phi \leq \pi$

Points are generated by converting spherical to Cartesian coordinates:

$$x = r \sin \phi \cos \theta$$

$$y = r \sin \phi \sin \theta$$

$$z = r \cos \phi$$

Randomly sampling angles θ and ϕ ensures uniform point distribution on the sphere. For low SI labeling, points are labeled randomly with values from 0 to 5 (Figure 4A). For high SI labeling, the sphere is divided into 6 equal polar angle bands, and points are labeled accordingly (Figure 4B).

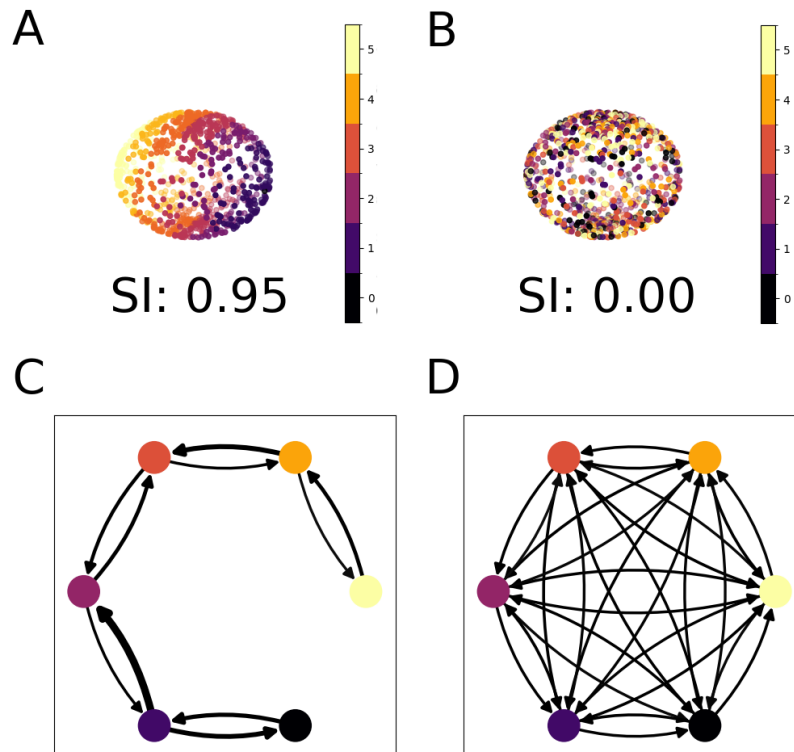


Figure 2.4: SI Illustration A) Structured labeling: A sphere divided into 6 equal-sized bins, each assigned a unique label. This structured arrangement results in a high SI of 0.95. B) Non-structured labeling: A sphere with randomly assigned labels, demonstrating a lack of spatial pattern and an SI of 0.00. C) A graphical representation of an SI calculation step of the orderly labeled sphere. Nodes symbolize labeled sections, and arrows depict overlap between sections. Thicker arrows signify greater overlap. Optimal SI is achieved when connections exist solely between adjacent sections, minimizing overlap with distant sections. D) Non-structured graph illustrating a step of SI calculation of the randomly labeled sphere. Arrows pointing to non-neighboring nodes indicate that points in the bin represented by the source node have neighbors from non-adjacent bins. This suggests less structure and thus lowers the SI score.

2.6 Rastermap

Rastermap (Stringer et al. 2023) is a visualization framework designed to find patterns in neural data. It sorts neural responses along a one-dimensional manifold, making it easier to identify and interpret activity patterns. The algorithm works by first normalizing and reducing the dimensionality of the data. Then, it clusters the data and computes an asymmetric similarity matrix between the clusters. Finally, it sorts the similarity matrix to optimize a matching score. This sorting helps to reveal patterns in the neural data that would be difficult to see otherwise.

The algorithm consists of five primary steps: dimensionality reduction, clustering, creating an asymmetric similarity matrix, sorting this matrix, and upsampling the cluster centers.

2.6.1 Calculation of the Rastermap

1. **Data Normalization and Dimensionality Reduction:** The data preparation process begins with normalization. Each neuron's activity is z-scored to achieve a mean of zero and a standard deviation of one. After normalization, the data undergoes dimensionality reduction. First, the normalized data is binned. Then, singular value decomposition (SVD) is performed on this binned data. Only a fraction of the left singular vectors representing the neurons are retained. These vectors are then scaled by their corresponding singular values.
2. **Clustering:** To identify groups of co-active neurons, a modified K-means algorithm called scaled K-means is used. This method assigns each neuron to a cluster iteratively re-estimating cluster means, similar to the regular K-means, with the difference of fitting an additional variable λ_i for each neuron i such that

$$\mathbf{x}_i = \lambda_i \mu_{\sigma_i} + \text{noise}$$

Here, \mathbf{x}_i is the activity vector for neuron i , σ_i is the cluster assigned to neuron i , and μ_j is the activity of cluster j . Similar to K-means, the model is optimized by iteratively assigning each neuron to the cluster that best explains its activity.

3. **Cluster Activity and Cross-Correlation:** For each cluster i , compute the mean cluster activity and z-score each cluster activity trace. Calculate the cross-correlation between all pairs of cluster activity traces using the following formula:

$$c_i \times c_j[\tau] = \frac{1}{T} \sum_t c_i(t - \tau) c_j(t)$$

for a specified number of time lags τ .

Next, construct an asymmetric similarity matrix S by taking the maximum value of the cross-correlation for each cluster pair over these τ values:

$$S_{i,j} = \max_{\tau} (c_i \times c_j[\tau])$$

4. **Sorting the Similarity Matrix:** Optimize the sorting of the similarity matrix of the cluster nodes to maximize the matching score. The matching score is defined as the dot product between the sorted similarity matrix and a prespecified matching matrix M :

$$\text{score} = \sum_{i,j} M_{i,j} S_{i,j}^{\text{sorted}}$$

The matching matrix M has two parts: global similarity part and local traveling salesman similarity part. The global similarity matrix is defined as:

$$M_{i,j}^{\text{global}} = -\log(|x_i - x_j| + 0.001)$$

where $x_i = \frac{i}{n}$ and the diagonal elements $M_{i,i}^{\text{global}}$ are set to zero.

This function generates a similarity measure between elements i and j , where the similarity decreases gradually as the difference between x_i and x_j increases. The logarithmic transformation ensures that even distant elements contribute to the overall similarity, creating a long-range effect. This is why it is termed "global." The small constant 0.001 prevents the logarithm from becoming undefined and controls the maximum similarity value.

In contrast, the local similarity matrix is close to 1 for the first off-diagonal and very small elsewhere:

$$M_{i,j}^{\text{local}} = \exp\left(-\frac{(x_i - x_j)^2}{2\sigma^2}\right)$$

where $\sigma = \frac{i}{2N}$ and the diagonal elements $M_{i,i}^{\text{local}}$ being zero.

This is a Gaussian function where similarity decreases rapidly as the difference between x_i and x_j grows. Consequently, M_{local} strongly emphasizes nearby elements and quickly approaches zero for distant ones. The parameter σ determines the width of this local influence.

Set the lower diagonal of each matrix to zero to enforce forward sequences of activity. Normalize the matrices by their sums across all entries. The final matching matrix is a weighted sum of the two matrices:

$$M = (1 - w)M^{\text{global}} + wM^{\text{local}}$$

where w is a parameter.

Intuitively, if $M_{i,j}$ is high, the score increases more by putting a high value of S at position (i, j) in the sorted matrix. On the other hand, if $M_{i,j}$ is low, putting a low value of S at (i, j) minimizes the reduction in score.

The sorting algorithm is initialized using the first singular vector weights for each cluster node. The process then iteratively optimizes the arrangement by computing the change in score for each potential movement of clusters within the matrix. Each potential movement practically means that each row and column, representing the clusters, is moved into every other row and column, resulting in $n * (n - 1)$ potential moves. Initially, all possible movements of individual nodes are tested, and the node movement that yields the highest score increase is implemented. If no single-node movement improves the score, the algorithm progresses to testing all possible moves of two consecutive nodes, and so on, gradually increasing the group size. In the implementation, this process of searching for beneficial moves of node groups is repeated for up to 400 iterations or until no movement of any group size results in a score increase. This approach ensures a thorough exploration of the solution space, balancing local and global optimizations to achieve the best possible sorting of the similarity matrix.

5. Upsampling: The granularity of the feature space representation in Rastermap is significantly enhanced through an upsampling technique. This process expands the initial set of N_{clusters} reference points to a denser grid of $10 \times N$ clusters points. These newly generated points serve as intermediate locations within the feature space, providing a more refined framework for subsequent neuron assignment.

To construct these intermediate points, a weighted, locally linear regression approach is employed. This method approximates a linear function mapping discrete cluster indices to features in the Principal Component Analysis (PCA) space, focusing on localized regions surrounding each cluster. In essence, this process interpolates between the original cluster centers to create a more detailed map of the feature space.

The extent and influence of these local neighborhoods are modulated by weights determined by the proximity of clusters in the Rastermap sorting. A Gaussian weighting function is utilized, based on Euclidean distance between cluster indices, with a standard deviation of $\sigma = 1/\sqrt{2}$. To maintain computational efficiency and focus on local relationships, only the 50 nearest neighbors are considered for each upsampled point.

This linear approximation is performed at every new upsampled position, resulting in a feature space map with tenfold increased resolution compared to the original data. Consequently, a significantly more detailed view of the feature space is obtained.

The ultimate goal of this upsampling process is to provide a more refined framework for neuron assignment. By creating a denser grid of reference points, neurons can be more precisely matched to their most similar location in feature space.

As a final step, the upsampled features are correlated with each neuron's activity profile. Neurons are then assigned to the position of their best-matching upsampled node. This assignment process ensures that each neuron is placed at the most appropriate location within the refined feature space, allowing for a more accurate representation of neuronal relationships and similarities.

2.6.2 Hyperparameters

1. **Number of clusters:** This parameter determines the number of clusters to be computed using scaled K-means clustering. It represents the initial number of reference points before upsampling.
2. **Number of principle components:** This parameter specifies the number of top left singular vectors to keep after performing singular value decomposition on the normalized activity matrix.
3. **Locality:** This parameter determines the weight given to the local "traveling-salesman" similarity matrix when creating the final matching matrix. It balances between global and local similarity in the optimization process.
4. **Time lag window:** The range of time lags (τ) used when computing cross-correlations between cluster activity traces.
5. **Temporal bin:** Determines whether and how to bin the data in time before computing singular vectors.

We employed the Python library provided by the Rastermap authors for our analysis. Our calculations used the following parameter values: 30 clusters, 64 principal components, a locality of 0.15, and a time lag window of 15, temporal bin was left 1, as our data was already binned.

Chapter 3

Results

To understand how the hippocampus encodes spatial information, we employed UMAP for dimensionality reduction and visualization of neural signals recorded from mice running on a treadmill. Prior to applying UMAP, we preprocessed the data using two binning approaches: temporal binning, where signals were mean-binned in time using 50 samples per bin (each bin representing 0.625 seconds), and spatial binning, where data was divided into 30 bins (each bin representing 12 cm) for each lap of the treadmill. The time-binned and place-binned signals were then separately analyzed.

Analyzing data of mice running on the treadmill with familiar reward stimulus at the track's end, UMAP successfully captured the underlying structure of the neural activity. The resulting three-dimensional embedding revealed a circular pattern, suggesting that the neural representation of the mouse's position during laps on the treadmill followed a circular manifold (Figure 3.1).

To further investigate the spatial representation within the UMAP embedding, we mapped the mouse's position on the treadmill onto the embedding using a color gradient. This gradient transitioned smoothly from white at the starting position to white again at the final position, with a blue-orange color spectrum representing intermediate locations. Remarkably, this color mapping aligned perfectly with the circular UMAP projection. Each loop within the circle corresponded to a single lap completed by the mouse, and the color gradient remained consistent across all laps.

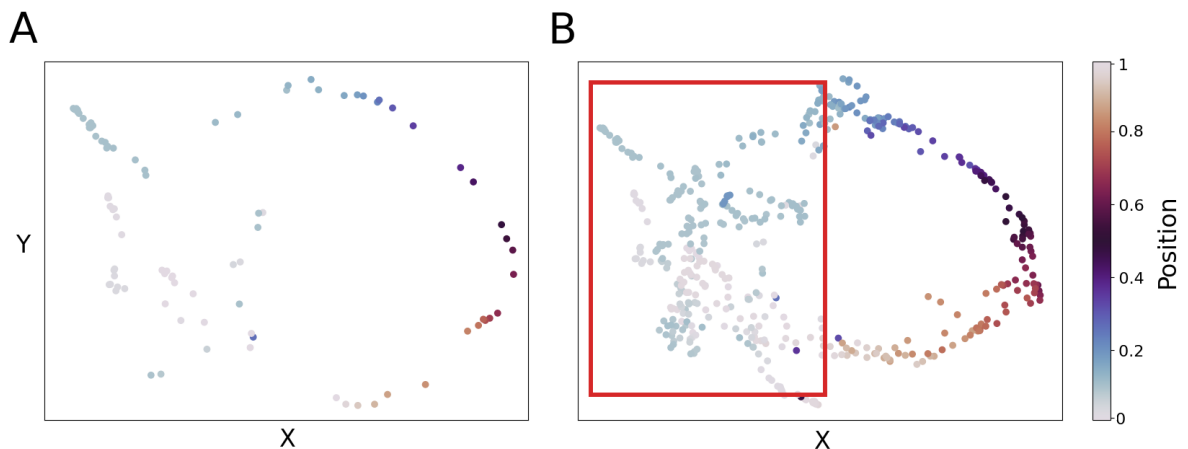


Figure 3.1: UMAP embeddings of neural activity during treadmill running of mouse 1, day 3, reward run. A) UMAP projection of a single (third) lap on the treadmill. Individual points represent temporal bins, colored with a twilight gradient indicating position on the treadmill (refer to Figure 1 for color scale mapping). This visualization demonstrates how neural activity patterns change with respect to the animal’s position during a single lap. This embedding consists of 67 points. B) Full UMAP embedding of the entire session, comprising 480 points. A distinct cluster of points (highlighted by a red rectangle) represents times when the mouse was consuming the reward. These points are concentrated at the part of the embedding representing the beginning of the treadmill, indicating that the mouse slowed down or stopped to receive and consume the reward, thus spending more time in this position. Principal Component Analysis (PCA) was applied to align the first two components with the X and Y axes, resulting in this 2D visualization.

To investigate how different binning methods affect UMAP embeddings, we evaluated two distinct approaches:

1. Temporal binning: In this method, neural activity data points within specific time intervals are aggregated and averaged to form a single data point.
2. Spatial binning: This approach groups and averages neural activity data from adjacent locations on the treadmill to create a single data point.

A key distinction between these methods lies in how they handle periods of mouse inactivity, such as during reward consumption. In spatial binning, these stationary periods are typically consolidated into a single data point. In contrast, temporal binning may generate multiple data points during these inactive intervals.

When using temporal binning, a distinct cluster emerged within the circular UMAP projection. This cluster corresponded to time points when the mouse was stationary or slow, likely due to eating behavior (Figure 3.1B). To eliminate this cluster we focused on the bins when the mouse was running (velocity $> 10m/s$). Embedding the running-only signal disrupted the circular structure by splitting it at the point of removal (Figure 3.2B).

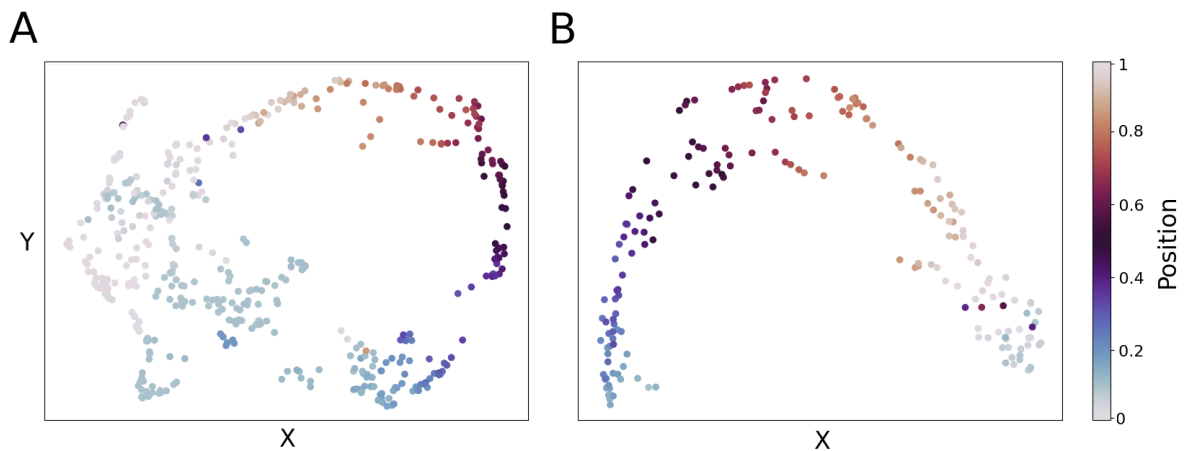


Figure 3.2: Comparison of full run to running-only UMAP embeddings. A) UMAP embedding of the entire session (480 points), including both running and stationary periods. The embedding shows a complete circular structure. B) UMAP embedding after removing stationary periods (241 points). The circular structure is disrupted, revealing a split at the point where stationary temporal bins were excluded. This comparison illustrates the impact of isolating running behavior on the overall neural state space representation. PCA was employed to align the first two principal components of each embedding with the X and Y axes.

Spatial binning of neural activity did not yield a distinct cluster for stationary periods, as expected, since all immobile time points corresponded to a single location (Figure 3.3A). Interestingly, the effect of bin size on the UMAP embedding differed between temporal and spatial binning strategies. While finer temporal binning generally produced a more cohesive circular structure, spatial binning exhibited a more complex relationship. For some mice, coarser spatial bins were necessary to maintain a connected circular embedding. Increasing the spatial resolution beyond a certain point led to fragmentation of the embedding structure, contrary to what was observed with temporal binning (Figure 3.3B).

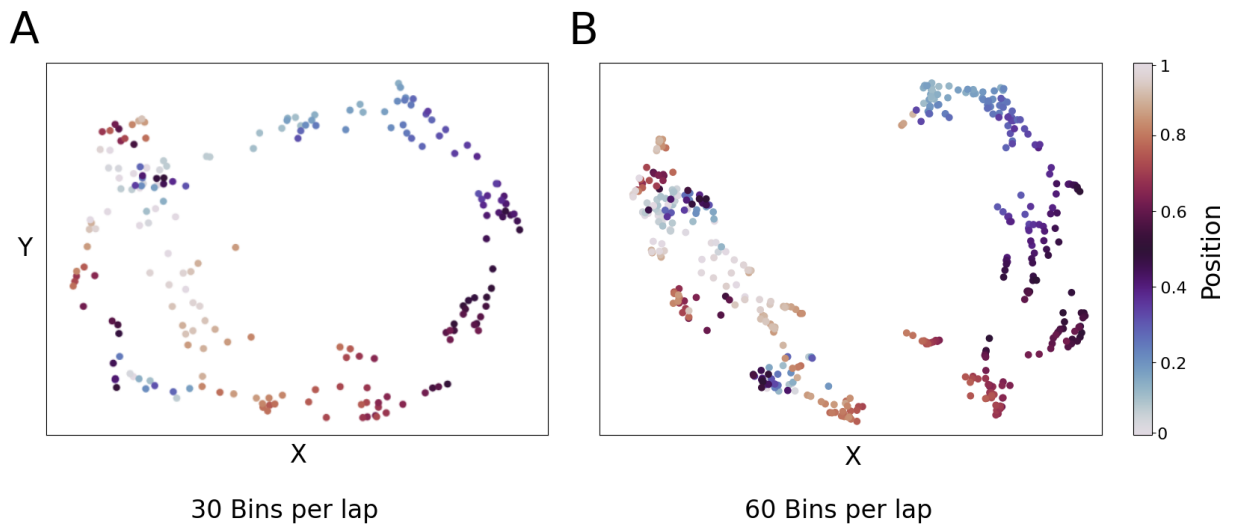


Figure 3.3: Impact of spatial binning resolution on UMAP embeddings of neural activity. A) 30 bins per lap. This coarser spatial resolution maintains a more continuous circular structure in the embedding. B) 60 bins per lap. The finer spatial resolution leads to a more fragmented embedding, with breaks in the circular pattern.

To investigate how the introduction of a novel stimuli affects the UMAP representation, we analyzed neural signals recorded while mice ran on a treadmill with a braking mechanism that made it harder for the mouse to move the belt at a specific location. This intervention was applied for three consecutive days.

On the initial day, after introducing the novel braking stimulus, mouse 1 displayed a UMAP embedding with a well-defined structure, closely resembling the patterns observed during standard reward runs (Figure 3.4A). The position gradient mapped onto this structure also mirrored that of reward runs, suggesting minimal initial disruption to the neural representation of spatial information. In contrast, mice 2, 3, and 4 exhibited less coherent structures on this first day (Figure 3.4D), indicating a more immediate impact of the braking stimulus on their neural representations. The second day brought further changes (Figure 3.4B,D): Mouse 1's previously stable structure deteriorated, while mouse 2 showed improved structural coherence. Mice 3 and 4 maintained their low structural integrity from the previous day. By the third day, a marked deterioration in UMAP structure was evident across all mice compared to their baseline reward runs. The characteristic circular pattern and associated position gradient, previously observed, were no longer discernible (Figure 3.4C,F).

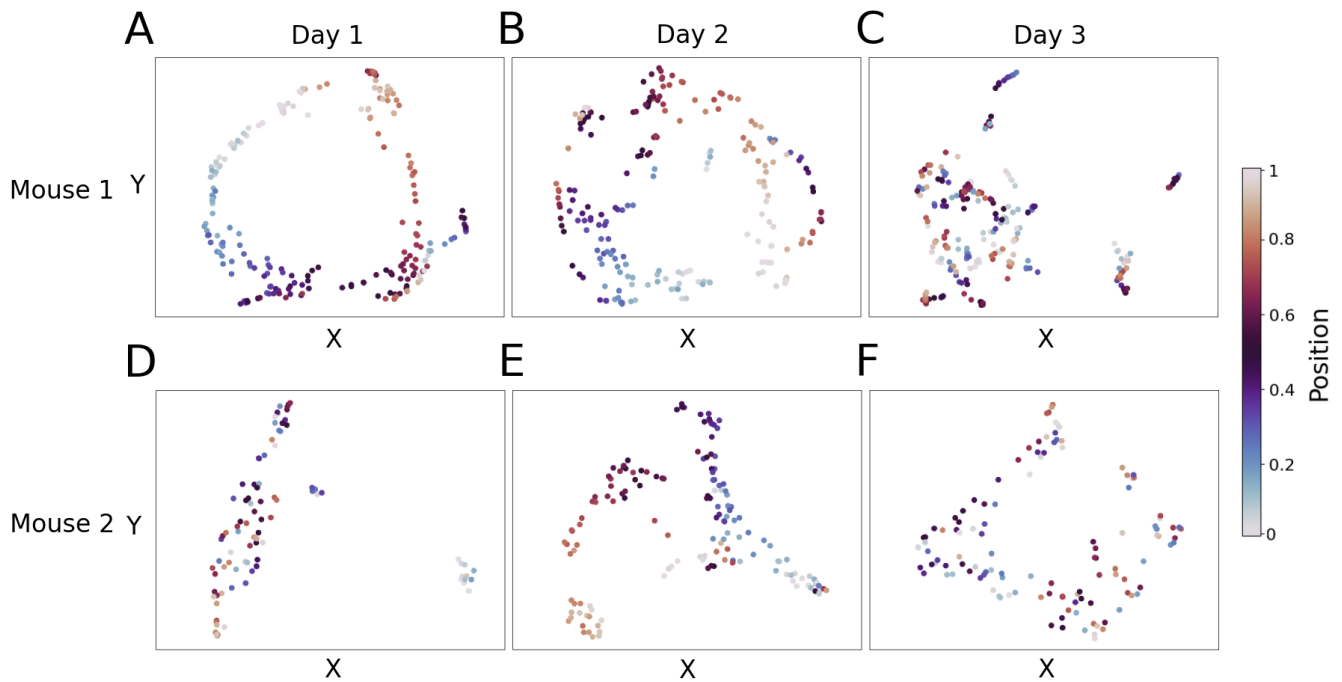


Figure 3.4: Progressive impact of the braking stimulus on UMAP representations in two mice over three consecutive days. Mouse 1: A) Day 1 (243 points): The UMAP embedding shows a well-defined circular structure with a clear position gradient, closely resembling the pattern observed in standard reward runs. This suggests minimal initial disruption to spatial information encoding. B) Day 2 (243 points): The previously stable structure begins to show signs of deterioration. The circular pattern is less pronounced, and the position gradient appears more diffuse. C) Day 3 (210 points): A marked deterioration is evident. The characteristic circular structure is largely lost, and the position gradient is no longer discernible, indicating significant disruption to spatial encoding. Mouse 2: D) Day 1 (95 points): The UMAP embedding displays a less coherent structure compared to mouse 1. The embedding consists of fewer points, because mouse 2 completed fewer laps on the treadmill. This reduced activity is likely due to the surprise effect of the novel braking stimulus, which may have caused the mouse to hesitate or pause more frequently. The lack of visible structure in this embedding is primarily attributed to this reduced running behavior rather than solely to disrupted neural representations. E) Day 2 (189 points): The embedding shows improved structural coherence compared to Day 1. A more defined circular pattern emerges, though not as pronounced as mouse 1’s initial state. This improvement suggests that mouse 2 had become more comfortable with the novel braking stimulus. F) Day 3 (137 points): Similar to mouse 1, the embedding structure has significantly deteriorated. The circular pattern and position gradient are no longer visible, indicating a breakdown in the organized encoding of spatial information. PCA was applied to align the first two principle components of each embedding with the X and Y axes, resulting in this 2D visualization.

To quantify the changes in spatial encoding observed in the UMAP embeddings, we employed a graph-based method called the Structure index (Figure 3.5). This method analyzes the topological structure of the data and provides a numerical value to represent the level of organization.

As part of the SI calculation, we first computed overlap scores between different positions in the

environment. To visualize the distribution of these overlap scores across days and conditions, we developed a method to summarize them for statistical comparison. This involved symmetrizing each mouse's overlap matrix, extracting the upper triangular portion, and summing diagonal elements to create an array of overlap scores. These arrays were aggregated across mice for each day and visualized as boxplots (Figure 3.5A,C).

For a more detailed view of individual mouse performance, we also calculated the SI for each mouse separately across the three days (Figure 3.5B,D).

We observed a consistently high SI value across all three days for the reward runs in all mice except mouse 3, which had low structure on the first day and continued to learn throughout the runs (Figure 3.5A,B). This indicates that the neural representation of spatial information remained stable during these runs with only reward delivery.

Consistent with the deterioration observed in the UMAP structure, the SI value associated with the braking stimulus condition exhibited a gradual decline over the three days in all mice except for mouse 4, which SI sharply dropped immediately in day 1 (Figure 3.5C,D), ultimately reaching a point where no significant structure was detectable. This finding suggests that repeated exposure to the novel braking stimulus disrupted the organized encoding of spatial information in the hippocampus progressively in some mice and more rapidly in others.

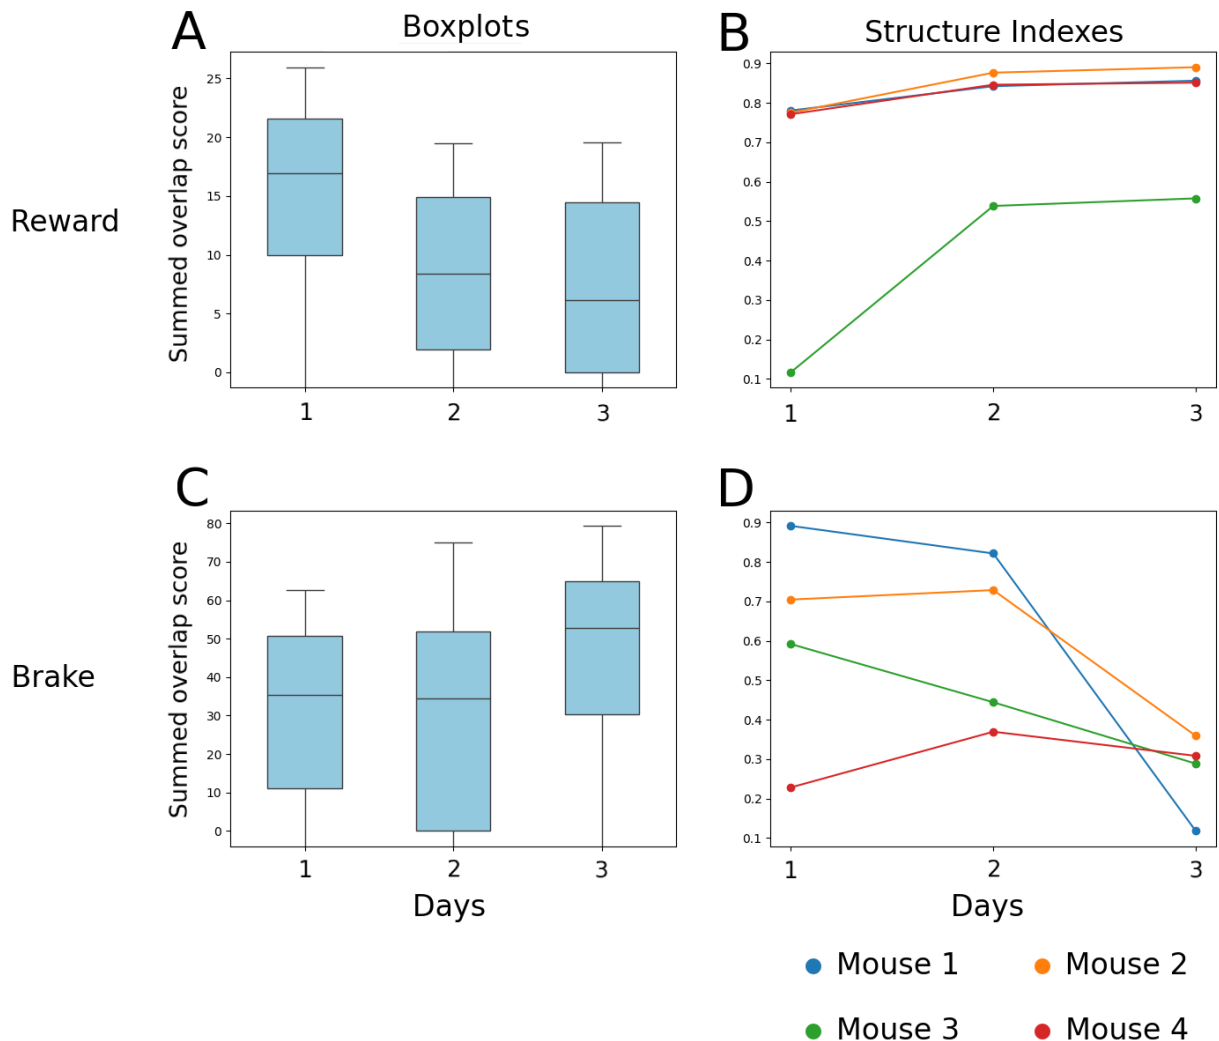


Figure 3.5: Comparison of Structure Index between reward and brake conditions over three days. This figure illustrates the changes in neural representation structure across reward and brake conditions: Reward Condition: A) Boxplots show decreasing overlap score across all three days, indicating improving structure. The improvement is the result of the improvement of structure of mouse 3 (Green) in days 2 and 3. B) Individual mouse trajectories demonstrate minimal fluctuation (mouse 1, 2, 4) or improving trend (mouse 3) in SI scores, further supporting the stability of neural structure in the reward condition. Brake Condition: C) Boxplots reveal a declining trend in SI scores over the three days, suggesting a progressive loss of structure in neural representations. D) Individual mouse trajectories show a general downward trend, particularly pronounced by day 3, indicating a consistent degradation of structural organization across subjects.

To gain further insights into the changes in UMAP structure, we employed Rastermap clustering (Figure 3.6). This method groups neurons with similar activity patterns. During reward runs, place cell sequences exhibited a consistent firing pattern, activating once per lap (Figure 3.6A,C). This regularity

aligns with the stable UMAP structure and high SI values observed across all three days for reward runs.

By the third day of brake runs, we observed notable alterations in place cell activation patterns. In mouse 1, we detected skips in the place cell sequence, consistent with the progressive deterioration of the UMAP structure and decreasing SI. Mouse 2 exhibited an increased frequency of place cell iterations, with six complete sequences occurring within just three full laps, aligning with the overall lower structure scores during brake runs compared to reward runs (Figure 3.6D). The sequence of place cell activations was not entirely disrupted by the novel stimuli, but rather the mapping of this sequence onto the spatial representation was distorted. This distortion appeared as a shift and disrupted frequency of firing patterns.

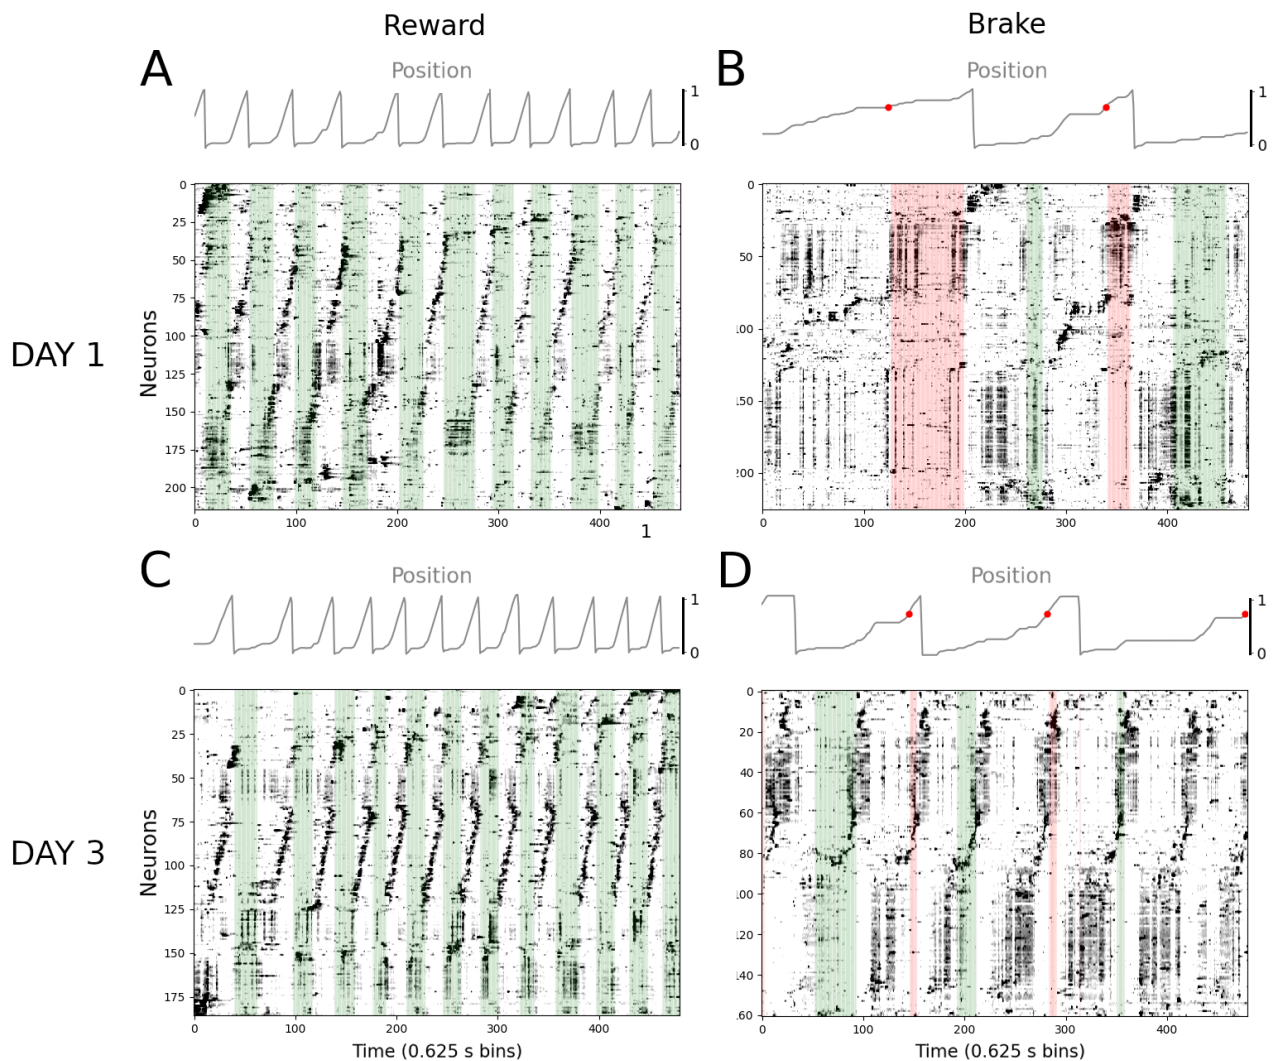


Figure 3.6: Rastermap plot displaying neuronal activity patterns across reward and brake runs of mouse 2 on days 1 and 3. A) Day 1, Reward condition: Stable place cell activation patterns are evident, with each sequence corresponding to a single lap. Green highlights indicate reward dispensing periods. B) Day 1, Brake condition: Place cell activations are present but less pronounced due to fewer laps run. Red highlights show periods when the treadmill brake was on. C) Day 3, Reward condition: Similar to day 1, stable place cell activation patterns are observed, with each sequence matching a single lap. D) Day 3, Brake condition: Sequential place cell activations outnumber the lap count, suggesting increased frequency of place cell iterations. Two distinct bands of cell activities are visible during brake runs: one underlying place cell activity, the other corresponding to periods of no movement. In all panels, the position signal ranges from 0 (beginning of the treadmill) to 1 (end of the treadmill). Red dots on the position plots mark the start of brake signals. The heatmap below each position plot represents neuronal time series data. Green highlights in the heatmap indicate temporal bins when reward was being dispensed, while red highlights show temporal bins with the treadmill brake on.

Theoretically, such a distorted place cell mapping should disrupt the interpretability of the color

gradient in the UMAP embedding, while keeping the underlying UMAP structure itself intact. However, our observations did not fully align with this expectation (Figure 3.4C).

To further investigate the UMAP distortion, we performed a detailed analysis of the rastermap clusters. In this analysis, we identified three distinct firing patterns within the recorded neural signals. The first firing pattern closely resembled that of place cells, exhibiting repeated activation at specific locations on the treadmill. This suggests that these neurons were primarily responsible for encoding spatial information. The second pattern consisted of two disjunct bands of activity. One of these bands exhibited neuronal activity predominantly during periods when the mouse was in motion, ostensibly supporting the activation observed within the place cells. In contrast, the second non-place-cell cluster demonstrated activity primarily when the mouse was either stationary or engaged in feeding behaviors, as visible in Figure 3.6D, possibly encoding information related to rest or some other behavioral state.

Based on these findings, we proceeded to manually extract place cells whenever they were identifiable within the Rastermap clusters. This allowed us to apply UMAP separately to two distinct neural populations: the extracted place cells and the cells exhibiting the disjunct firing patterns (Figure 3.7).

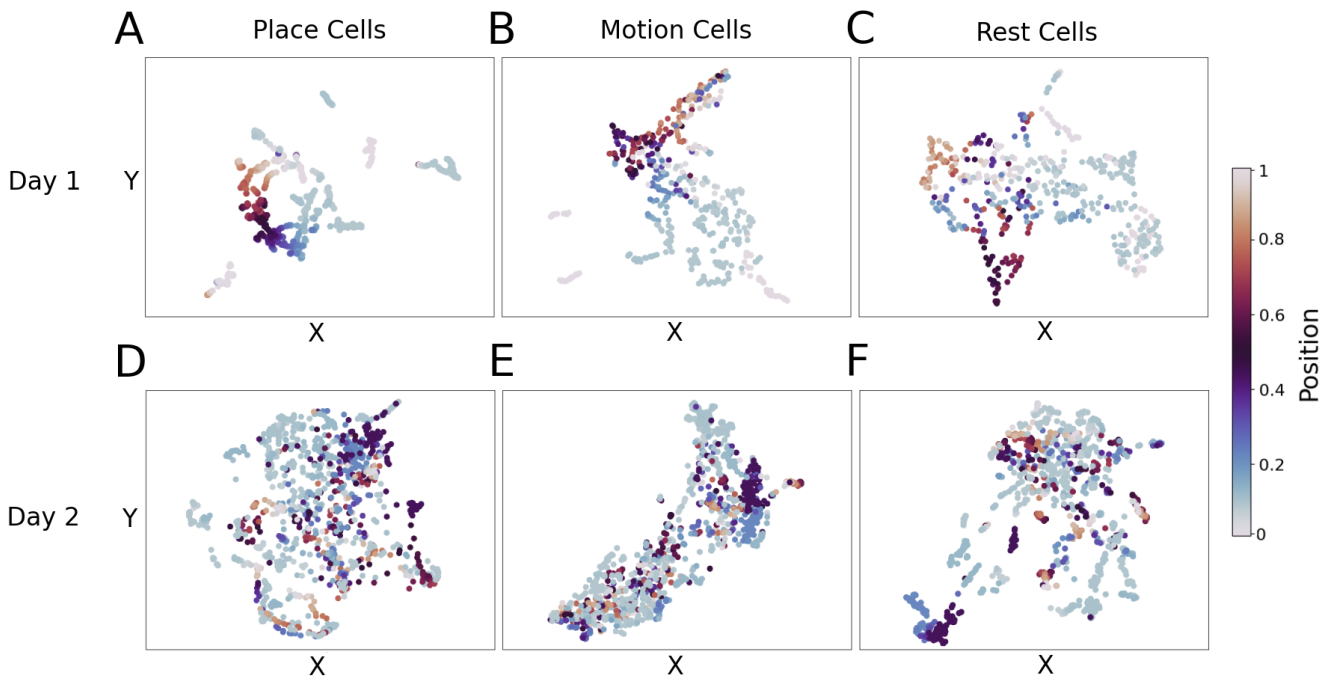


Figure 3.7: UMAP projections of different neuronal populations during reward runs. Each embedding has 480 points. (A) UMAP embedding of extracted place cells, showing a circular structure. (B,E) UMAP projections of motion cells, which are active alongside place cells during movement. (C,F) UMAP projections of rest cells, active when the motion cells are inactive. (D) UMAP embedding of all neurons in a case where no circular structure was originally observed, demonstrating that extraction of place cells did not restore circular structure in this instance. Note that non-place-cell clusters (B,C,E,F) display no visually meaningful embeddings.

Interestingly, a circular embedding was only observed for the UMAP projection of the place cells (Figure 3.7A). This finding aligns with our initial observations, where a circular structure was present

in the UMAP embedding during reward runs with well-defined spatial information. In contrast, the application of UMAP to the non-place-cell clusters did not yield any visually meaningful embeddings (Figure 3.7B,C,E,F). Notably, extracting the place cells did not restore the circular structure in cases where it was previously absent in the UMAP embedding of all neurons (Figure 3.7D). No structure emerged, even with the removal of seemingly redundant neurons.

Chapter 4

Discussion

In our study, we performed an in-depth examination of neuronal signal data using three advanced analytical techniques: Uniform Manifold Approximation and Projection, RasterMap, and Structure Index. This analysis has provided significant insights into how spatial information is encoded within the neurons of mice.

We began by replicating previous findings (Sun et al. 2023) through the application of UMAP to visualize spatial information encoded in hippocampal neuronal activity. Our results demonstrated that the spatial position of mice on a treadmill is represented by a circular embedding. As expected from previous works (Wills et al. 2005), we observed that this encoded structure remains stable in the absence of novel stimuli.

Our investigation into binning strategies has yielded insights into the neural representation of space in the hippocampus. Temporal binning revealed a distinct cluster in the signal embedding corresponding to stationary periods, highlighting the importance of considering behavioral states when analyzing neural data. Analysis of running-only signals disrupted the circular structure observed in the UMAP projection, emphasizing the continuous nature of spatial encoding during movement.

These findings suggest that in addition to visualizing spatial information encoded in hippocampus, topological techniques enable us to also infer behavioral and environmental states associated with them. Previous studies have demonstrated that spatial continuity is achieved through partial coverage of spatial fields modulated by phase precession (O'Keefe and Recce 1993; Skaggs et al. 1996). Our findings suggest UMAP might be sensitive to different spatial bin size selections. We hypothesize that this sensitivity arises because optimal representation may require each spatial bin to cover a single place field. This could be analyzed in future using additional techniques for precise identification of place fields. Our observation emphasizes the importance of careful methodological considerations in neural data analysis and highlights the complex relationship between analytical spatial resolution and the underlying neural code.

Our study also demonstrated that the introduction of a novel stimulus leads to a deterioration of the structure in low dimensional spatial representations. This was evidenced by changes in the UMAP visualizations over time and quantified using SI. This contrasts with the results of Sun et al. (2023), who observed a gradual adaptation of representations to a novel task. This discrepancy may be attributed to differences in the nature of the novel stimuli. While their novel stimulus consisted of the introduction of unfamiliar visual cues within a familiar task context, our study involved a novel stimulus that imposed increased behavioral demands on the animals. The aversive nature of our task may have induced a more pronounced disruption to the underlying neural code.

Interestingly, we observed variability in the response to the novel stimulus across individual mice. Mouse 1 exhibited initial resilience, maintaining its UMAP structure initially before showing deteri-

oration. This contrasted with other mice, which displayed more immediate disruption to their spatial representations.

Despite individual differences, a consistent pattern emerged across all mice over the three-day period following the introduction of the novel stimulus. We observed a progressive breakdown of the circular UMAP structure, indicating a fundamental reorganization of spatial representations. This reorganization appears to be in response to the new temporal dynamics introduced by the braking mechanism, which served as our novel stimulus.

To quantify these observations, we employed the Structure Index to assess the topological structure of the data. This analysis supported our qualitative observations from the UMAP visualizations. Embeddings that were well-visualized and maintained clear structure exhibited higher Structure Index values, while more chaotic embeddings showed lower values.

Our study builds upon and complements previous findings, such as those reported by Sosa, Plitt, and Giocomo (2024), which demonstrated that hippocampal place cells encode not only spatial locations but also salient events within an environment. While Sosa et al. observed that changing reward locations induced remapping of place cell activity, our research extends this concept by examining the effects of introducing novel stimuli. It is known (Wills et al. 2005), that place cells are able to rapidly switch between different representations of different environments. Our research offers new insights into the explanation of what happens to a stable representation during learning of another.

Our analysis, utilizing RasterMap clustering techniques, has revealed a possible mechanism underlying the observed deterioration of spatial representations following the introduction of a novel stimulus. We found that this deterioration is primarily associated with a fundamental change in place cell encoding patterns.

In familiar environments, place cells predominantly encoded reward-to-reward intervals at the beginning of each lap. However, upon introduction of a novel stimulus, this encoding pattern shifted dramatically. Place cells began to represent reward-to-stimulus and stimulus-to-reward intervals instead. The stimulus being the beginning of the brake section of the treadmill or its anticipation. This alteration in the reference points for spatial encoding may be the key factor disrupting the stable attractor dynamics observed in familiar settings.

Anticipation of mouse is well documented phenomena (Barnstedt, Mocellin, and Remy 2024). Given that the braking event represents an aversive stimulus, it is plausible that the location where the mouse anticipates this event could serve as a salient point in the animal's cognitive map. This anticipation of a surprising or aversive event might trigger changes in neuronal firing patterns, potentially leading to a remapping of spatial representations in the hippocampus (Blair et al. 2023).

This temporal reorganization of place cell activity provides a possible explanation for the breakdown of the circular UMAP structure observed over the three-day period. It suggests that the hippocampal representation is dynamically adjusting to incorporate the new temporal dynamics introduced by the braking mechanism, potentially at the expense of maintaining a stable spatial map. These findings extend our understanding of hippocampal plasticity, demonstrating how the introduction of a novel, behaviorally relevant stimulus can lead to a reorganization of neural encoding that prioritizes salient events over pure spatial mapping.

Lastly, we found that removing non-place cells did not help recover structure in data where it was not previously present. This observation suggests that the lack of structure in certain datasets was not due to the presence of non-place cells masking an underlying spatial representation. Rather, it indicates that the absence of structure reflects a genuine reorganization or disruption of the spatial encoding.

Conversely, when we removed non-place cells from structured data, we observed that the place cell encoding remained intact. This preservation of structure suggests that the spatial information is primarily encoded by the place cells, as expected.

Conclusion

This study aimed to investigate the neural encoding of spatial maps with a particular focus on the stability and adaptability of these representations in the hippocampus.

Our UMAP analysis revealed a ring-like structure in the spatial representations of CA1 regions, confirming previous studies that suggested the existence of a ring attractor in hippocampal spatial encoding.

Under familiar environmental conditions, we observed that these ring attractors remained highly stable, aligning with theoretical predictions of CAN behavior as well as previous studies. This stability is crucial for maintaining consistent spatial representations and supports the role of CANs in reliable spatial navigation and memory.

One of the contributions of this study is the observation of the degradation of these attractor states when animals were exposed to novel stimuli. The Structure Index quantitatively demonstrated this deterioration, while UMAP visualizations clearly illustrated the breakdown of the ring structure. This degradation of the attractor state in novel environments supports the idea that different environments are represented by distinct CAN representations in the hippocampus. It suggests that there are limits to what a single CAN can represent, and that the brain may need to form new or modified CANs to accurately encode novel spatial contexts.

A key contribution of our research is the identification of a change in place cell encoding as the primary mechanism underlying this attractor degradation. Through RasterMap clustering techniques, we discovered that in familiar environments, place cells primarily encoded the reward-to-reward interval at the beginning of each lap. However, when exposed to novel stimuli, this encoding pattern changed dramatically. Place cells tended to shift to represent reward-to-stimulus and stimulus-to-reward intervals. We hypothesize that the observed disruption of stable attractor dynamics in familiar environments may primarily result from this shift in spatial encoding reference points.

These observations challenge the view of CANs as rigid structures and instead suggest a more flexible model where attractor states can be temporarily destabilized to accommodate new information. This plasticity may be crucial for adaptive behavior, allowing the hippocampus to balance stable representations with the ability to encode novel experiences.

A potential limitation of our study is the specific experimental paradigm used, particularly the focus on a single type of salient stimulus (the brake) and the sample size. To address these limitations and further validate our findings, future research should aim to replicate these results with a larger number of mice and explore the effects of diverse salient stimuli beyond the brake used in this study. This could include various sensory modalities such as visual, auditory, or olfactory cues, as well as other types of behaviorally relevant events. Such investigations would help confirm the generalizability of our observations and provide a more comprehensive understanding of how different novel stimuli impact hippocampal spatial representations and attractor dynamics.

Bibliography

- Ansar, Wazib, Saptarsi Goswami, and Amlan Chakrabarti (May 2024). *A Survey on Transformers in NLP with Focus on Efficiency*. arXiv:2406.16893 [cs]. doi: 10.48550/arXiv.2406.16893. URL: <http://arxiv.org/abs/2406.16893> (visited on 07/31/2024).
- Bang, Jee Yoon et al. (June 2022). “Hippocampus-Anterior Hypothalamic Circuit Modulates Stress-Induced Endocrine and Behavioral Response.” In: *Frontiers in Neural Circuits* 16, p. 894722. ISSN: 1662-5110. doi: 10.3389/fncir.2022.894722. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9251012/> (visited on 07/27/2024).
- Barnstedt, Oliver, Petra Mocellin, and Stefan Remy (Apr. 2024). “A hippocampus-accumbens code guides goal-directed appetitive behavior.” en. In: *Nature Communications* 15.1. Publisher: Nature Publishing Group, p. 3196. ISSN: 2041-1723. doi: 10.1038/s41467-024-47361-x. URL: <https://www.nature.com/articles/s41467-024-47361-x> (visited on 07/28/2024).
- Bird, Chris M. and Neil Burgess (Mar. 2008). “The hippocampus and memory: insights from spatial processing.” eng. In: *Nature Reviews. Neuroscience* 9.3, pp. 182–194. ISSN: 1471-0048. doi: 10.1038/nrn2335.
- Blair, Garrett J. et al. (July 2023). *Hippocampal place cell remapping occurs with memory storage of aversive experiences*. en. Publisher: eLife Sciences Publications Limited. doi: 10.7554/eLife.80661. URL: <https://elifesciences.org/articles/80661> (visited on 08/02/2024).
- Blevins, Ann S. et al. (Oct. 2022). “From calcium imaging to graph topology.” In: *Network Neuroscience* 6.4, pp. 1125–1147. ISSN: 2472-1751. doi: 10.1162/netn_a_00262. URL: https://doi.org/10.1162/netn_a_00262 (visited on 07/20/2024).
- Blevins, Ann Sizemore and Danielle S. Bassett (2020). “Topology in Biology.” en. In: *Handbook of the Mathematics of the Arts and Sciences*. Ed. by Bharath Sriraman. Cham: Springer International Publishing, pp. 1–23. ISBN: 978-3-319-70658-0. doi: 10.1007/978-3-319-70658-0_87-1. URL: https://doi.org/10.1007/978-3-319-70658-0_87-1 (visited on 07/28/2024).
- Bostock, E., R. U. Muller, and J. L. Kubie (Apr. 1991). “Experience-dependent modifications of hippocampal place cell firing.” eng. In: *Hippocampus* 1.2, pp. 193–205. ISSN: 1050-9631. doi: 10.1002/hipo.450010207.
- Chaudhuri, Rishidev et al. (Sept. 2019). “The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep.” en. In: *Nature Neuroscience* 22.9. Publisher: Nature Publishing Group, pp. 1512–1520. ISSN: 1546-1726. doi: 10.1038/s41593-019-0460-x. URL: <https://www.nature.com/articles/s41593-019-0460-x> (visited on 08/02/2024).
- Chazal, Frédéric and Bertrand Michel (Sept. 2021). “An Introduction to Topological Data Analysis: Fundamental and Practical Aspects for Data Scientists.” English. In: *Frontiers in Artificial Intelligence* 4. Publisher: Frontiers. ISSN: 2624-8212. doi: 10.3389/frai.2021.667963. URL: <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2021.667963/full> (visited on 07/20/2024).

- Clark, Robert E. and Larry R. Squire (June 2013). “Similarity in form and function of the hippocampus in rodents, monkeys, and humans.” In: *Proceedings of the National Academy of Sciences of the United States of America* 110.Suppl 2, pp. 10365–10370. ISSN: 0027-8424. DOI: 10.1073/pnas.1301225110. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3690603/> (visited on 08/03/2024).
- Dabaghian, Yuri (Jan. 2021). “From Topological Analyses to Functional Modeling: The Case of Hippocampus.” en. In: *Frontiers in Computational Neuroscience* 14, p. 593166. ISSN: 1662-5188. DOI: 10.3389/fncom.2020.593166. URL: <https://www.frontiersin.org/articles/10.3389/fncom.2020.593166/full> (visited on 07/19/2024).
- Edelsbrunner, Letscher, and Zomorodian (Nov. 2002). “Topological Persistence and Simplification.” en. In: *Discrete & Computational Geometry* 28.4, pp. 511–533. ISSN: 0179-5376, 1432-0444. DOI: 10.1007/s00454-002-2885-2. URL: <http://link.springer.com/10.1007/s00454-002-2885-2> (visited on 08/03/2024).
- Eichenbaum, Howard (Apr. 2017). “The role of the hippocampus in navigation is memory.” eng. In: *Journal of Neurophysiology* 117.4, pp. 1785–1796. ISSN: 1522-1598. DOI: 10.1152/jn.00005.2017.
- Gardner, Richard J. et al. (Feb. 2022). “Toroidal topology of population activity in grid cells.” en. In: *Nature* 602.7895. Publisher: Nature Publishing Group, pp. 123–128. ISSN: 1476-4687. DOI: 10.1038/s41586-021-04268-7. URL: <https://www.nature.com/articles/s41586-021-04268-7> (visited on 07/31/2024).
- Giovannucci, Andrea et al. (Jan. 2019). “CaImAn an open source tool for scalable calcium imaging data analysis.” In: *eLife* 8. Ed. by David Kleinfeld and Andrew J King. Publisher: eLife Sciences Publications, Ltd, e38173. ISSN: 2050-084X. DOI: 10.7554/eLife.38173. URL: <https://doi.org/10.7554/eLife.38173> (visited on 07/20/2024).
- Gothard, K. M., W. E. Skaggs, and B. L. McNaughton (Dec. 1996). “Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues.” eng. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 16.24, pp. 8027–8040. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.16-24-08027.1996.
- Green, J. D. and A. A. Arduini (Nov. 1954). “Hippocampal electrical activity in arousal.” eng. In: *Journal of Neurophysiology* 17.6, pp. 533–557. ISSN: 0022-3077. DOI: 10.1152/jn.1954.17.6.533.
- Grienberger, Christine and Arthur Konnerth (Mar. 2012). “Imaging Calcium in Neurons.” en. In: *Neuron* 73.5, pp. 862–885. ISSN: 08966273. DOI: 10.1016/j.neuron.2012.02.011. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0896627312001729> (visited on 07/20/2024).
- Hafting, Torkel et al. (Aug. 2005). “Microstructure of a spatial map in the entorhinal cortex.” en. In: *Nature* 436.7052. Publisher: Nature Publishing Group, pp. 801–806. ISSN: 1476-4687. DOI: 10.1038/nature03721. URL: <https://www.nature.com/articles/nature03721> (visited on 08/03/2024).
- Hopfield, J J (Apr. 1982). “Neural networks and physical systems with emergent collective computational abilities.” In: *Proceedings of the National Academy of Sciences of the United States of America* 79.8, pp. 2554–2558. ISSN: 0027-8424. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC346238/> (visited on 07/31/2024).
- Kelly, Ryan C. et al. (Jan. 2007). “Comparison of Recordings from Microelectrode Arrays and Single Electrodes in the Visual Cortex.” en. In: *The Journal of Neuroscience* 27.2, pp. 261–264. ISSN: 0270-6474, 1529-2401. DOI: 10.1523/JNEUROSCI.4906-06.2007. URL: <https://www.jneurosci.org/lookup/doi/10.1523/JNEUROSCI.4906-06.2007> (visited on 07/20/2024).
- Leutgeb, Stefan et al. (July 2005). “Independent codes for spatial and episodic memory in hippocampal neuronal ensembles.” eng. In: *Science (New York, N.Y.)* 309.5734, pp. 619–623. ISSN: 1095-9203. DOI: 10.1126/science.1114037.

- Lindsay, Grace W. (Sept. 2021). “Convolutional Neural Networks as a Model of the Visual System: Past, Present, and Future.” In: *Journal of Cognitive Neuroscience* 33.10, pp. 2017–2031. issn: 0898-929X. doi: 10.1162/jocn_a_01544. url: https://doi.org/10.1162/jocn_a_01544 (visited on 07/31/2024).
- Maaten, Laurens van der and Geoffrey Hinton (2008). “Visualizing Data using t-SNE.” In: *Journal of Machine Learning Research* 9.86, pp. 2579–2605. issn: 1533-7928. url: <http://jmlr.org/papers/v9/vandermaaten08a.html> (visited on 08/04/2024).
- Mancall, Elliott L., David G. Brock, and Henry Gray, eds. (2011). *Grays Clinical Neuroanatomy: The Anatomic Basis for Clinical Neuroscience*. en. Philadelphia, Pa: Elsevier Saunders. isbn: 978-1-4160-4705-6.
- McCulloch, Warren S. and Walter Pitts (Dec. 1943). “A logical calculus of the ideas immanent in nervous activity.” en. In: *The bulletin of mathematical biophysics* 5.4, pp. 115–133. issn: 1522-9602. doi: 10.1007/BF02478259. url: <https://doi.org/10.1007/BF02478259> (visited on 07/31/2024).
- McInnes, Leland, John Healy, and James Melville (Sept. 2020). *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. arXiv:1802.03426 [cs, stat]. doi: 10.48550/arXiv.1802.03426. url: <http://arxiv.org/abs/1802.03426> (visited on 07/28/2024).
- Momota, Yuki et al. (Apr. 2024). “Amyloid- prediction machine learning model using source-based morphometry across neurocognitive disorders.” en. In: *Scientific Reports* 14.1. Publisher: Nature Publishing Group, p. 7633. issn: 2045-2322. doi: 10.1038/s41598-024-58223-3. url: <https://www.nature.com/articles/s41598-024-58223-3> (visited on 07/29/2024).
- Muller, R. U. and J. L. Kubie (July 1987). “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells.” eng. In: *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 7.7, pp. 1951–1968. issn: 0270-6474. doi: 10.1523/JNEUROSCI.07-07-01951.1987.
- Nelson, Carl J. and Stephen Bonner (June 2021). “Neuronal Graphs: A Graph Theory Primer for Microscopic, Functional Networks of Neurons Recorded by Calcium Imaging.” English. In: *Frontiers in Neural Circuits* 15. Publisher: Frontiers. issn: 1662-5110. doi: 10.3389/fncir.2021.662882. url: <https://www.frontiersin.org/journals/neural-circuits/articles/10.3389/fncir.2021.662882/full> (visited on 07/20/2024).
- Nicolelis, Miguel AL, ed. (2008). *Methods for Neural Ensemble Recordings*. eng. 2nd. Frontiers in Neuroscience. Boca Raton (FL): CRC Press/Taylor & Francis. isbn: 978-0-8493-7046-5. url: <http://www.ncbi.nlm.nih.gov/books/NBK1985/> (visited on 08/03/2024).
- O’Keefe, J. and N. Burgess (May 1996). “Geometric determinants of the place fields of hippocampal neurons.” eng. In: *Nature* 381.6581, pp. 425–428. issn: 0028-0836. doi: 10.1038/381425a0.
- O’Keefe, J. and J. Dostrovsky (Nov. 1971). “The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat.” eng. In: *Brain Research* 34.1, pp. 171–175. issn: 0006-8993. doi: 10.1016/0006-8993(71)90358-1.
- O’Keefe, J. and L. Nadel (1978). *The Hippocampus as a Cognitive Map*. eng. Publication Title: Oxford University Press: Oxford, UK. (1978). Oxford, UK: Oxford University Press. url: <http://www.cognitivemap.net/> (visited on 08/02/2024).
- O’Keefe, J. and M. L. Recce (July 1993). “Phase relationship between hippocampal place units and the EEG theta rhythm.” eng. In: *Hippocampus* 3.3, pp. 317–330. issn: 1050-9631. doi: 10.1002/hipo.450030307.
- Pnevmatikakis, Eftychios A. and Andrea Giovannucci (Nov. 2017). “NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data.” eng. In: *Journal of Neuroscience Methods* 291, pp. 83–94. issn: 1872-678X. doi: 10.1016/j.jneumeth.2017.07.031.

- Pnevmatikakis, Eftychios A., Daniel Soudry, et al. (Jan. 2016). “Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data.” eng. In: *Neuron* 89.2, pp. 285–299. ISSN: 1097-4199. DOI: 10.1016/j.neuron.2015.11.037.
- Qasim, Salman E. et al. (May 2023). “Neuronal activity in the human amygdala and hippocampus enhances emotional memory encoding.” en. In: *Nature Human Behaviour* 7.5. Publisher: Nature Publishing Group, pp. 754–764. ISSN: 2397-3374. DOI: 10.1038/s41562-022-01502-8. URL: <https://www.nature.com/articles/s41562-022-01502-8> (visited on 07/27/2024).
- Rao, Y. Lakshmisha et al. (Feb. 2022). “Hippocampus and its involvement in Alzheimer’s disease: a review.” In: *3 Biotech* 12.2, p. 55. ISSN: 2190-572X. DOI: 10.1007/s13205-022-03123-4. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8807768/> (visited on 07/27/2024).
- Rolls, Edmund T. (Jan. 2010). “Attractor networks.” en. In: *WIREs Cognitive Science* 1.1, pp. 119–134. ISSN: 1939-5078, 1939-5086. DOI: 10.1002/wcs.1. URL: <https://wires.onlinelibrary.wiley.com/doi/10.1002/wcs.1> (visited on 07/31/2024).
- Rutkowski, Tomasz M., Tomasz Komendziński, and Mihoko Otake-Matsuura (Jan. 2024). “Mild cognitive impairment prediction and cognitive score regression in the elderly using EEG topological data analysis and machine learning with awareness assessed in affective reminiscent paradigm.” English. In: *Frontiers in Aging Neuroscience* 15. Publisher: Frontiers. ISSN: 1663-4365. DOI: 10.3389/fnagi.2023.1294139. URL: <https://www.frontiersin.org/journals/aging-neuroscience/articles/10.3389/fnagi.2023.1294139/full> (visited on 07/29/2024).
- Schmidgall, Samuel et al. (May 2023). *Brain-inspired learning in artificial neural networks: a review*. arXiv:2305.11252 [cs, q-bio]. DOI: 10.48550/arXiv.2305.11252. URL: <http://arxiv.org/abs/2305.11252> (visited on 07/31/2024).
- Sebastian, Enrique R., Julio Esparza, and Liset M. de la Prida (Jan. 2024). “Quantifying the distribution of feature values over data represented in arbitrary dimensional spaces.” en. In: *PLOS Computational Biology* 20.1. Publisher: Public Library of Science, e1011768. ISSN: 1553-7358. DOI: 10.1371/journal.pcbi.1011768. URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1011768> (visited on 07/28/2024).
- Singh, Gurjeet, Facundo Memoli, and Gunnar Carlsson (2007). *Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition*. en. ISSN: 1811-7813. The Eurographics Association. ISBN: 978-3-905673-51-7. URL: <https://doi.org/10.2312/SPBG/SPBG07/091-100> (visited on 08/03/2024).
- Sizemore, Ann E. et al. (July 2019). “The importance of the whole: Topological data analysis for the network neuroscientist.” In: *Network Neuroscience* 3.3, pp. 656–673. ISSN: 2472-1751. DOI: 10.1162/netn_a_00073. URL: https://doi.org/10.1162/netn_a_00073 (visited on 07/20/2024).
- Skaggs, W. E. et al. (1996). “Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences.” eng. In: *Hippocampus* 6.2, pp. 149–172. ISSN: 1050-9631. DOI: 10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K.
- Sosa, Marielena, Mark H. Plitt, and Lisa M. Giocomo (Feb. 2024). *Hippocampal sequences span experience relative to rewards*. en. Pages: 2023.12.27.573490 Section: New Results. DOI: 10.1101/2023.12.27.573490. URL: <https://www.biorxiv.org/content/10.1101/2023.12.27.573490v2> (visited on 08/02/2024).
- Squire, L. R. (Apr. 1992). “Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans.” eng. In: *Psychological Review* 99.2, pp. 195–231. ISSN: 0033-295X. DOI: 10.1037/0033-295x.99.2.195.
- Stevenson, Ian H and Konrad P Kording (Feb. 2011). “How advances in neural recording affect data analysis.” en. In: *Nature Neuroscience* 14.2, pp. 139–142. ISSN: 1097-6256, 1546-1726. DOI: 10.1038/nn.2731. URL: <https://www.nature.com/articles/nn.2731> (visited on 07/20/2024).

- Stringer, Carsen et al. (July 2023). *Rastermap: a discovery method for neural population recordings*. en. Pages: 2023.07.25.550571 Section: New Results. doi: 10.1101/2023.07.25.550571. URL: <https://www.biorxiv.org/content/10.1101/2023.07.25.550571v1> (visited on 08/03/2024).
- Sun, Weinan et al. (Aug. 2023). *Learning produces a hippocampal cognitive map in the form of an orthogonalized state machine*. en. Pages: 2023.08.03.551900 Section: New Results. doi: 10.1101/2023.08.03.551900. URL: <https://www.biorxiv.org/content/10.1101/2023.08.03.551900v1> (visited on 07/31/2024).
- Tolman, Edward C. (1948). "Cognitive maps in rats and men." In: *Psychological Review* 55.4. Place: US Publisher: American Psychological Association, pp. 189–208. issn: 1939-1471. doi: 10.1037/h0061626.
- Vogelstein, Joshua T. et al. (Dec. 2010). "Fast Nonnegative Deconvolution for Spike Train Inference From Population Calcium Imaging." In: *Journal of Neurophysiology* 104.6. Publisher: American Physiological Society, pp. 3691–3704. issn: 0022-3077. doi: 10.1152/jn.01073.2009. URL: <https://journals.physiology.org/doi/full/10.1152/jn.01073.2009> (visited on 08/03/2024).
- Wang, Zhijin et al. (Nov. 2023). "Editorial: Crosstalk between computational medicine and neuroscience in healthcare." English. In: *Frontiers in Neuroscience* 17. Publisher: Frontiers. issn: 1662-453X. doi: 10.3389/fnins.2023.1333227. URL: <https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2023.1333227/full> (visited on 07/31/2024).
- Wills, Tom J. et al. (May 2005). "Attractor Dynamics in the Hippocampal Representation of the Local Environment." In: *Science* 308.5723. Publisher: American Association for the Advancement of Science, pp. 873–876. doi: 10.1126/science.1108905. URL: <https://www.science.org/doi/10.1126/science.1108905> (visited on 08/03/2024).
- Zheng, Jie et al. (June 2023). *Hippocampal Theta Phase Precession Supports Memory Formation and Retrieval of Naturalistic Experience in Humans*. en. Pages: 2023.06.05.543539 Section: New Results. doi: 10.1101/2023.06.05.543539. URL: <https://www.biorxiv.org/content/10.1101/2023.06.05.543539v1> (visited on 07/27/2024).