Doctoral theses at NTNU, 2023:45

Erik Hermansen

Finding topological structure in neural ensemble activity

Doctoral thesis

NTNU
Norwegian University of Science and Technology
Thesis for the Degree of
Philosophiae Doctor
Faculty of Information Technology and
Electrical Engineering
Department of Mathematical Sciences

NTNU
Norwegian University of
Science and Technology

Erik Hermansen

Finding topological structure in neural ensemble activity

Thesis for the Degree of Philosophiae Doctor

Trondheim, February 2023

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the Degree of Philosophiae Doctor

Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

ISBN 978-82-326-6835-9 (printed ver.)
ISBN 978-82-326-5627-1 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

Doctoral theses at NTNU, 2023:45

Printed by NTNU Grafisk senter

Preface

This thesis is submitted in partial fulfillment of the requirements for the degree of Philosophiae Doctor (Ph.D.) at the Norwegian University of Science and Technology (NTNU) in Trondheim.

During the years spent in working on this thesis, there are many who's support I have depended on. First, I want to express my gratitude to my advisor Benjamin Dunn for the guidance and advice he has given, both in science and in life in general. I would also like to thank my coadvisor Nils Baas for the ideas and discussions we have shared, and for introducing me to these topics and encouraging me to embark on this thesis. Moreover, I deeply appreciate all the people I have been fortunate to have had the pleasure to working with, including: Melvin Vaupel, Rich Gardner, Edvard Moser, May-Britt Moser, Paul Trygsland, the Neural Data Science Group at NTNU, David Klindt, the Topology group, Marius Thaule, the Department of Mathematical Sciences, the Kavli Institute for Systems Neuroscience, György Buzsaki and lab at NYU and the CINPLA group at UiO.

Furthermore, I would like to thank my friends and family for backing me and to my parents, Cécile and Finn Erik, for their loving support. Lastly, I hold Birgitte dear for keeping with me through ups and downs and all the moments we have shared the past few years.

Outline

In brief, the thesis is concerned with the application of topological data analysis (TDA) tools in studying firing rate activity of experimentally recorded neurons. We use and develop computational frameworks based on persistent cohomology for finding topological structure in the collective computations of neural ensembles. Specifically, we uncover the toroidal topology in grid cell population activity. The thesis opens with an introductory chapter describing TDA, data preprocessing and neural representations. This gives background and context for the main body of work, namely, the three articles which subsequently follow. These are briefly summarized below.

Article I

Uncovering 2-D toroidal representations in grid cell ensemble activity during 1-D behavior, Erik Hermansen, David A. Klindt and Benjamin A. Dunn. 2022. *bioRxiv* [1] We describe a novel framework combining a fuzzy topological representation of a point cloud (based on the uniformity assumption underlying UMAP) with persistent homology, to reveal topological structure in datasets. This allows uncovering the toroidal topology of grid cells in head-fixed mice during one-dimensional running tasks.

Article II

A topological perspective on the dual nature of the neural state space and the correlation structure, Melvin Vaupel, Erik Hermansen and Benjamin A. Dunn. In this paper, we discuss the duality of two approaches to finding topological structure of neural ensemble representations, via the population vector complex and the correlation complex. We propose a framework which combines the complementary information held in each approach to refine the topological inference of either. Furthermore, we suggest that a neural ensemble is identified when the two approaches coincide and provide examples of applications in different datasets.

Article III

Toroidal topology of population activity in grid cells, Rich Gardner*, Erik Hermansen*, Marius Pachitariu, Yoram Burak, Nils Baas, Benjamin A. Dunn, May-Britt Moser and Edvard Moser. *Equal contributions. 2022. *Nature*, 602(7895), pp.123-128 [2]. The toroidal topology of grid cell population activity in rats is demonstrated for the first time in experimental data. The structure is seen to be preserved both during different wake conditions, and sleep, suggesting the activity arises from intrinsically structured circuitry.

Contents

Introduction

0.1 Topology in data

It is common to think of data as a sample of some underlying space. Consider, for instance, the typical predator-prey cycle (Fig. 1). When the number of prey increases, more food is available for the predator. This leads to an increase in predators, taking the lives of more prey. In turn, the prey population decreases, leading to a decrease in number of predators until the balance again flips. Regular observations of the number of animals in each population results in a *circular* point cloud (Fig. 1a). However, as predators need many prey to survive, the predator population is usually smaller than that of the prey. Moreover, other factors, such as climate, may lead to seasonal variations, and the population observations may be made irregularly and unreliably. This will lead to a noisy, non-uniform, oval point cloud (Fig. 1b). Thus, while characterizing the relationship as *circular* leaves out details of the population balance (namely, its *geometry*), it gives an interpretable description of the relationship. The concept that the underlying shape is informative about the source from which the data originates is essentially the driving force of *topological data analysis* (TDA).

TDA utilizes mathematical tools from (algebraic) topology to describe data by computing topological invariants of algebraic structures associated to it. This allows describing the topological properties of the underlying space and distinguishing it from other spaces. The main algorithm used in this thesis is *persistent cohomology* [3–6]. Persistent cohomology defines a *barcode* of the data, a topolo-

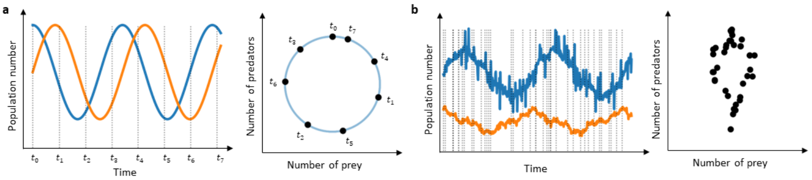

Fig. 1: The prey-predator feeding relationship is circular. **a.** Idealized prey-predator cycle. Left, the prey (blue) and predator (orange) population numbers follow each other in a cyclical manner. Dashed lines indicate observations. Right, plotting the number of prey vs. predators depicts the circular phenomenon. **b.** A hypothesized more realistic depiction of the cycle described in **a**.

gical signature colloquially describing the evolution of n-dimensional holes in the point cloud across increasing volumes of balls centered at each point in the data. This captures the notion of observing a circular relationship in the above example and extends it to higher-dimensional spaces and features. In this thesis, we will describe some of the shortcomings of this method and how to address these, but also see its formidable advantages. Prominently, we study the "most intriguing application of TDA" (Oct. 2022, *Scientific American* [7]), namely the shape of representations in the brain, and specifically, of the grid cell network (Fig. 2).

0.2 Contributions and outline

The main contributions in this work concern applications of TDA in neuroscience. We will describe the relevance of topology and the implications of our findings in neuroscience. Moreover, the analysis tools used and developed enabling these results are detailed.

In Section 1, we introduce the relevant concepts in TDA, describing the theoretical background of persistent (co)homology and coordinatization. In practice, the presence of noise in TDA may be detrimental. Hence, we discuss the influence of noise and the necessity to preprocess data in data analysis in Section 2. We also describe the theory and intuition behind the construction of a fuzzy topological representation of the data motivated by the dimensionality-reduction tool UMAP. This provides a well-founded denoising step for persistent cohomology and the combination is, in Article I, coined UMAPH, where we showcase its usefulness in neuroscience. Section 3 starts by describing neural representations, elucidating the importance of understanding shape in neural data. We discuss spatial representations and mention previous work in TDA analyzing the structure of these. Next, the two main approaches to infer the topology of neural representations are described. These are based on the dual constructs termed the *population vector* and *correlation* complexes in Article II. We describe how they may elicit different information and provide topological insights of the data, giving examples in both simulated and real neural data sets. In Section 4, we elaborate on the main finding of Article III. The characteristics of grid cell populations are first discussed, giving a background for the hypothesized shape of the network, before describing the results. Employing the pipeline detailed in the previous sections, we demonstrated in unprecedented recordings of the entorhinal cortex in rats, that the activity of populations of grid cells resides on a toroidal state space. Surprisingly, the analysis showed the toroidal description to be preserved across different environments and tasks and even during sleep. This portrays the significance of the topological approaches conveyed in this thesis. Finally, we summarize and give an outlook on TDA and neuroscience in Section 5.

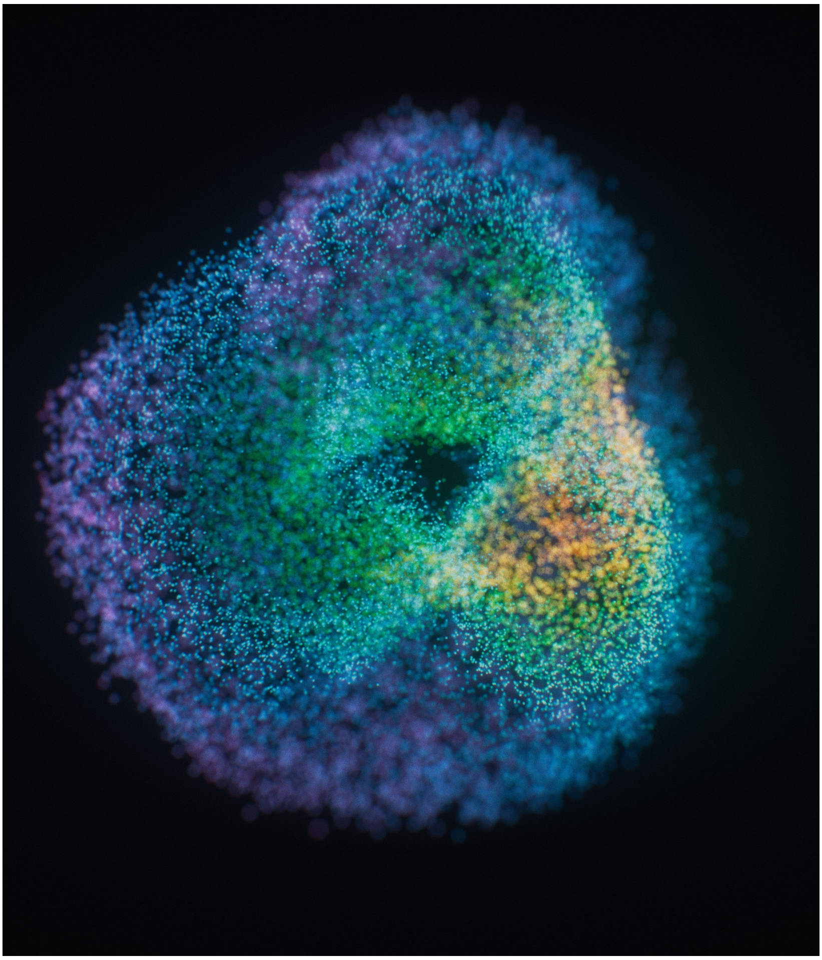

Fig. 2: A three-dimensional embedding of neural activity states of an ensemble of grid cells. Credits: Kavli Institute for Systems Neuroscience, NTNU and Helmet.

1 Persistent homology and cohomological coordinatization

Usually, data is given as a discrete set of points, X, with no inherent, interesting topology. However, the dataset often permits a relationship (such as pairwise distance or correlation), d, to be defined between the points. This allows estimating the underlying space of which the data is sampled from. We give a brief description of the main algorithm in TDA, persistent homology [3, 4, 8], assessing the shape of the underlying space through characterizing its n-dimensional (n-D) holes. The exposition draws from [9–12]. Next, we outline a circular coordinatization scheme based on persistent *co*homology introduced in [13], which is used in Sections 3 and 4 to *decode* neural representations.

1.1 Persistent homology

Simplicial complex. In TDA, the underlying space is commonly approximated as a *simplicial complex* [14]. A simplicial complex, K, is a set of vertices V_K and a set, S_K, of non-empty finite subsets of V_K called *simplices*, such that for every $\tau \in S_K$, all non-empty subsets $\sigma \subseteq \tau$ also belong to S_K and every vertex is a simplex (Fig. 3a,b). The dimension of $\sigma \in S_K$ is given as $|\sigma| - 1$. The dimension of K is given as the maximum of its simplices. A graph is a one-dimensional simplicial complex with vertices and edges (0- and 1-simplices) representing points and pairwise relations. A higher-dimensional complex additionally describes higher-order relations via its n-simplices. In fact, we may define a higher-dimensional complex (called the *clique complex*) from a graph by defining the sets of vertices in each clique as a simplex.

Rips and Čech complex. For all computations in this thesis we use the *Rips complex*, R_ϵ. Given a point cloud X embedded in some metric space (M, d), we choose a radius $\epsilon > 0$ and form balls around each point $x \in X$, $B_\epsilon(x) = \{x' \in M \mid d(x, x') \leq \epsilon\}$. Each set of points with pairwise intersecting balls defines a simplex of the Rips complex, i.e.:

$$S_{R_\epsilon} = \{\sigma \subseteq X; \quad B_\epsilon(x) \cap B_\epsilon(y) \neq \emptyset, \quad \forall x, y \in \sigma\}.$$

This may be seen as the clique complex of the graph formed by adding edges between points whose pairwise distance is less than ϵ. The Rips complex is said to approximate the *Čech complex*, $\check{C}_\epsilon$, in that any topological feature present across the Rips complexes of scales ϵ and $\sqrt{2}\epsilon$, is also a topological feature of $\check{C}_{\sqrt{2}\epsilon}$. The simplices of the *Čech* complex are the sets of points whose balls intersect, i.e.,

$$S_{\check{C}_\epsilon} = \{\sigma \subseteq X; \quad \cap_{x \in \sigma} B_\epsilon(x) \neq \emptyset\}.$$

If we assume X is nicely sampled from its underlying space, Y, and the radius ϵ is chosen such that $B := \cup_{x \in X} B_\epsilon(x) = Y$, we say that B forms a *good cover* of Y. The *Nerve lemma* then gives us that the *Čech* complex has the same *homotopy*

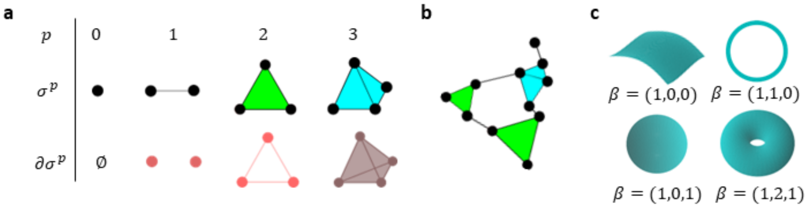

Fig. 3: a. Geometrical realization of simplices, σ^p, of dimension p, (top) and their respective boundaries, $\partial\sigma^p$ (bottom). **b.** Example of a simplicial complex. **c.** Different topological spaces (from left-right, top-bottom: a plane, a circle, a sphere and a torus) and the three first Betti numbers, β, associated to the spaces.

type as Y, i.e. the simplicial complex carries similar topological properties as the underlying space – in particular the same *homology*.

Homology. Homology is motivated by the observation that a topological space may be described by its n-dimensional holes. For instance, a 0-D hole describes a connected component, a circle additionally carries a 1-D hole and a sphere a 2-D hole, while a torus one 0-D, two 1-D and one 2-D hole (Fig. 3c). Notably, the homological signature of a torus and two circles attached to a sphere is the same. We will in the following describe *simplicial homology* with coefficients in some field $\mathbb{F}$. In all computations we used the coefficient field $\mathbb{Z}_{47} = \{0,1,2,\ldots,46\}$.

Chain complex. Given a simplicial complex K, a *simplicial p-chain* is defined as the weighted sum of p-dimensional simplices: $\sum_i a_i\sigma_i^p$, where $a_i \in \mathbb{F}$ and $\sigma_i^p \in S_K$. We denote $\sigma^p = (v_0, v_1, \ldots v_p)$, for $v_i \in V_K$, and note that the ordering matters - swapping v_i and v_j (for any choice of i, j) produces the opposite oriented simplex, $-\sigma^p$. The group of all p-chains with the p-simplices in K as basis is denoted $C_p(K)$. The *boundary operator* is defined on a p-simplex as $\partial_p(\sigma) = \sum_{i=0}^{p}(-1)^i(v_0,\ldots,\widehat{v_i},\ldots,v_p)$. This extends linearly: $\partial_p(\sum_i a_i\sigma_i^p) = \sum_{i=1} a_i\partial_p(\sigma_i^p)$. One may then see that $\partial_{p+1}\partial_p = 0$ for any $p \geq 0$. This induces a *chain complex*:

$$\begin{tikzcd}
\ldots \arrow[r] & C_{p+1}(K) \arrow[r, "\partial_{p+1}"] & C_p(K) \arrow[r, "\partial_p"] & C_{p-1}(K) \arrow[r] & \ldots ,
\end{tikzcd}$$

The kernel and image of ∂_p are subgroups of $C_p(K)$ called the *cycle* and *boundary* groups and we say the boundary operator takes a simplex to its *boundary*.

Betti numbers. Because $\partial^2 = 0$, all boundaries are also cycles. The cycles which are *not* boundaries represent to the p-dimensional holes. The homology groups of K are then defined as the quotient groups

$$H_p(K) = \frac{\mathrm{Ker}\partial_p}{\mathrm{Im}\partial_{p+1}} = \frac{\text{cycles}}{\text{boundaries}}.$$

An element of $H_p(K)$ is a class of *homologous* cycles in $C_p(K)$, i.e., $[z] = [x] \in H_p(K)$ if $z - x \in C_p(K)$ is a boundary. The dimension of $H_p(K)$ is called the p-

th *Betti number* of K. Homology has the *functorial* property that a map of simplicial complexes, $f : K \to J$, induces a map on homology [15]. I.e., we get a map of chain complexes $f_p : C_p(K) \to C_p(J)$ by extension, $f_p(\sum_{i=1} a_i \sigma_i^p) = \sum_{i=1} a_i f_p(\sigma_i^p)$. This commutes with the boundary map, $f_p \partial_{p+1} = \partial_p f_{p+1}$, giving rise to the map $f_p^* : H_p(K) \to H_p(J)$ by $f_p^*[z] = [f_p(z)]$.

Rips filtration. A *filtration* is a nested sequence of *subcomplexes* of K, e.g., $\emptyset = K_0 \subseteq K_1 \subseteq \cdots \subseteq K_n = K$. In the constructing of the Rips complex, increasing the radius around each point will only give rise to more simplices in the complex, hence, there is a natural chain of Rips complexes, called the Rips filtration:

$$R_{\epsilon_0} \subset R_{\epsilon_1} \subset \ldots \subset R_{\epsilon_n},$$

where $\epsilon_0 < \epsilon_1 < \ldots < \epsilon_n$ corresponds to values of ϵ in which new simplices are formed.

Persistent homology. The inclusion maps $i^l : K_l \to K_{l+1}$, induces maps on homology $(i_p^l)^* : H_p(K_l) \to H_p(K_{l+1})$ (denoted j_l in the following). Hence, we get a sequence of homology groups and maps called the p-th persistent homology of K:

$$H_p(K_0) \xrightarrow{j_0} H_p(K_1) \xrightarrow{j_1} \ldots \xrightarrow{j_{n-1}} H_p(K_n)$$

The induced maps track which classes are preserved in consecutive steps of the filtration. We then consider the longest composition of maps for which this holds true. E.g., if $[z] \in H_p(K_b)$ is not in the image of j_{b-1}, we say $[z]$ was *born* at b. If $j_b j_{b+1} \cdots j_d([z]) = 0$ is the shortest composition that does not preserve $[z]$, we say it *died* at d. The *lifetime* of the class is then $d - b$ [16]. The idea behind *persistent homology* is that classes which live the longest are considered more significant, while those that quickly disappear are considered noise (Fig. 4a).

It is possible to summarize this structure formally as a *persistence module*, $PH_p(K)$ [17]. Moreover, there exists a correspondence between the persistence modules over $\mathbb{F}$ and graded modules over a graded polynomial ring, $\mathbb{F}[t]$, and a classification theorem which allows us to decompose the persistence module into a sum of *persistence intervals* $I([b_i, d_i))$:

$$PH_p(K) \cong \bigoplus_k I([b_k, d_k)).$$

Here, k loops over all p-dimensional classes in $PH_p(K)$ and (b_k, d_k) denotes the filtration steps at which the classes are born and die. Thus, persistent homology may be given as bars starting at b_k and ending at d_k. The collection of bars for all dimensions is known as the *barcode*.

1.2 Cohomological coordinatization

In [18], De Silva et al. use persistent cohomology to extract circle-valued maps, giving a projection from the data to a (lower-dimensional) toroidal space. This will

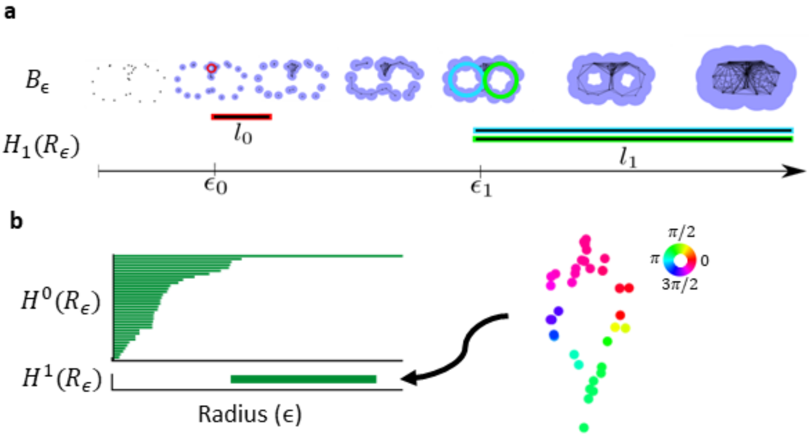

Fig. 4: Persistent homology describes the persistence of n-dimensional holes across increasing filtration scales. **a.** Example barcode of a figure-eight point cloud. Top, placing balls of increasing radii around each point and connecting points whose pairwise balls intersect gives rise to the Rips filtration of the point cloud (note, only the 1-skeleton is shown). Bottom, the lives of 1-dimensional persistent homology classes are described as bars starting at time of birth, ϵ_i, with lengths representing the lifetime, l_i. The noisy circular feature (red) in **a** is described by the short-lived bar, while the two long-lived bars (cyan and green) depict the prominent circles of the point cloud. **b.** The circular feature in the point cloud of Fig. 1b may be measured and parametrized through persistent cohomology and cohomological coordinatization. The barcode (left) shows a single long-lived bar, whose cocycle representative may be used to obtain circular coordinates of the point cloud (right).

provide well-suited coordinates of the neural representations studied in later sections. Here, we describe cohomology and the coordinatization algorithm, closely following the mentioned study.

Cohomology is a dual notion to homology, *reversing* all maps. I.e., a *cochain* complex:

$$\dots \longrightarrow C^{p-1}(K) \xrightarrow{\delta^{p-1}} C^p(K) \xrightarrow{\delta^p} C^{p+1}(K) \longrightarrow \dots ,$$

consists of groups of simplicial p-*co*chains, $C^p(K) = \{\alpha_{\mathbb{F}} : S_K^p \to \mathbb{F}\}$ and *co*boundary maps, δ^p. The function, $\alpha_{\mathbb{F}}$, assigns to each simplex, $\sigma^p \in S_K^p$, a coefficient value $a_i \in \mathbb{F}$, with coboundary defined by $\delta^p(\alpha_{\mathbb{F}}(\sigma^p)) = \alpha_{\mathbb{F}}(\partial_p(\sigma_p))$. The notions of *co*cycles, *co*boundaries and (persistent) *co*homology follow as above. Due to the universal coefficient theorem [19], the resulting barcodes are identical when working over a field.

The following result [19] motivates using cohomology to obtain circular coordinates:

$$H^1(X;\mathbb{Z}) \cong [X, S^1].$$

This says that the classes in the first cohomology group of a topological space X with coefficients in $\mathbb{Z}$, $[\alpha_{\mathbb{Z}}]$, are in bijection with classes of homotopy equivalent maps from X to the circle, $[\theta : K \to S^1]$. Intuitively, the desired circle-valued map is given by smoothing the coefficient edge values of the representative of the chosen cohomology class across the edges of the simplicial complex and defining the resulting values on the vertices as the coordinates of the point cloud. The algorithm is briefly described below [18]:

1. Given a cohomology class with coefficients $\mathbb{Z}_p$ (for some prime p), we *lift* the cocycle representative $\alpha_{\mathbb{Z}_p}$ to an integer cocycle by replacing the coefficient values with the corresponding integers in $\{-(p-1)/2, \dots, -1, 0, 1, \dots, (p-1)/2\}$. θ may then be defined through extension on the vertices and edges: $\theta(\sigma^0) = 0$ and $\theta(\sigma^1)$ goes around S^1 with winding number $\alpha_{\mathbb{Z}}(\sigma^1)$.
2. However, the vertices of the simplicial complex are the points in the data set, and as these are all mapped to zero, we rather replace the integer cocycle by $\alpha_{\mathbb{R}} = \alpha_{\mathbb{Z}} + \delta^0 b_{\mathbb{R}}$, for $b_{\mathbb{R}} \in C^0(K)$ with real coefficients. Then, θ sends each edge to an interval of length $\alpha_{\mathbb{R}}(\sigma^1)$ and vertices $\theta(\sigma^0) = b_{\mathbb{R}}(\sigma^0) \bmod \mathbb{Z}$.
3. To get a smooth map, we choose the cocycle which minimizes the variation over edges, $|\alpha_{\mathbb{R}}(\sigma^1)|^2$, solved by least-squares optimization.

Persistent cohomology gives an indication of the most prominent circular features in the dataset, and it is common to compute circular coordinates for only the features which are deemed the most relevant (Fig. 4b). For instance, if the barcode suggests a 2-D toroidal topology of the dataset, we may choose to coordinatize the two longest-lived classes found, thought to represent the two circles of the torus. However, we must first fix a simplicial complex for which the relevant features exist, and usually pick a scale close to the death of these, defining the coordinates on as large complex as possible.

2 Data preprocessing

While persistent (co)homology has been shown to be a robust tool in the sense that small perturbations of the data give small changes in the barcode [20], noise in data is often in the form of outliers or skewed sampling, which may distort the topological signature. For instance, a densely sampled circular point cloud with one noisy point appearing in the middle will cut the lifetime of the 1-D persistent homology class in half (Fig. 5). Another issue is the computational complexity, computing high-dimensional persistent homology has severe consequences on both the running time and memory usage of the algorithm [21]. Neural data is notoriously noisy, and with potentially hour-long recordings and *k*Hz-sampling covering multiple brain regions and different neural ensembles, it is necessary to process the data in order to extract the underlying topology using persistent (co)homology.

First, we discuss some of the common approaches of preprocessing in data analysis. We see how a different sets of tools were used prior to detecting the ring topology of the head direction cell network in mice during wakefulness and REM sleep [22, 23] with persistent (co)homology [24–26]. We emphasize the work in [24] which we are inspired by in the analyses in this thesis. Next, we sketch the theoretical justification of the construction of the 'UMAP complex', based on the theory of Spivak [27], McInnes and Healy [28], Barr [29] and Jardine [30–32]. This serves as a noise-robust approximation of the underlying space, and we discuss its practical benefits and show how it can be combined with persistent cohomology to detect the ring topology in the above-mentioned dataset also during SWS sleep.

2.1 Common techniques

Clustering and dimensionality reduction

One of the first decisions to be made in analyzing data is whether all dimensions of the data matrix are informative of the intended study. E.g., when predicting the weather in region 1 for the next day, including weather measurements in the far-away region 2 can lead the experimenter to infer conclusions about the other region or a mean of the two. Hence, we would like to separate the measurements of the regions into two *clusters* and analyze them independently. This may be done in various ways. Either the geographical locations of the measurements and the lack of information contained for the intended prediction may be known, and the irrelevant dimensions are naturally discarded. Alternatively, if we do not possess this information, but have historical information about the weather outcome, we may infer the information contained in each instrument by assessing its previous predictive power and keep only the dimensions which explain the most. In both these settings, the clustering is *supervised* by the experimenter's prior knowledge. However, in the shortage of such information, one may consider the correlations between the measurements. We expect all instruments predictive of the weather

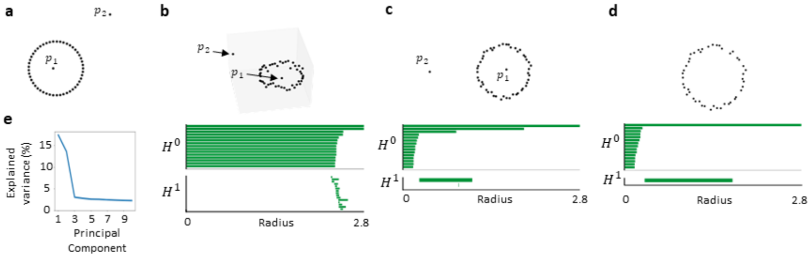

Fig. 5: Preprocessing of data is necessary to extract interesting shape information in the barcodes. **a.** Example circular point cloud with two outliers, p_1 and p_2. **b.** First three dimensions of an embedding in $\mathbb{R}^{202}$ of the point cloud in **a**, where 200 dimensions with small Gaussian noise have been appended (top). Applying persistent cohomology to the high-dimensional embedding (using Euclidean metric) gives an ambiguous barcode (bottom). **c.** Projecting the embedding to its two first principal components recovers the structure of the original point cloud (top). The corresponding barcode suggests a prominent circular feature in the barcode. **d.** Using 'fuzzy downsampling' to remove 5 points (top), leaves a 'denoised' point cloud. This denoising is seen in the barcode (bottom), where the additional H^1-bar and the two long-lived H^0-bars (caused by the two outliers, p_1 and p_2) seen in **c** are no longer present. **e.** The variance explained for the first 10 principal components after applying PCA to the point cloud in **b**.

in region 1 will have measurements which are more correlated with each other (implicitly through their correlation to the common weather) than they are with the measurements of region 2. This may guide the experimenter to find two groups of instruments with high inner group correlation.

In [26] and [24], both studies use cells from both the anterodorsal thalamic nucleus (ADn) and postsubiculum (PoS), while [25] only focuses on ADn, stating that PoS is a different brain area with a separate structure. Notably, in [24], the topological inference is done twice, where in the second iteration only neurons which are best explained by the decoded representation in the first iteration are kept, giving a much clearer predicted topology.

Nonetheless, even if the weather measurements are only from one region, some instruments may either be redundant (e.g., if they are located close to each other) and lead to overfitting, or be describing noisy or non-informative measurements (e.g., a broken instrument) leading to lower predictive performance (Fig. 5b). Additionally, working with high-dimensional data is inherently more difficult due to a statistical phenomenon known as the *curse of dimensionality* [33]. For instance, as the volume of a hypersphere in Euclidean space grows exponentially for increasing dimensions, the difference between the largest and smallest in a randomly generated point cloud will approach zero, and distinguishing noise from signal becomes impossible [34]. Hence, a common step is to try to project the data to a lower dimension whilst retaining the most relevant features (Fig. 5c). Principal component analysis (PCA) finds orthogonal linear combinations of the dimensions, called components, which maximizes the amount of variance captured in each component such that the first component explains most variance, the second the second most and so on. The user may then determines how many components (dimensions) to keep based on the variance explained (Fig. 5e).

In [24], PCA is used to project the n-dimensional data set (where n is the number of neurons) to its six first principal components. Chauhuri et al. use Isomap [35] to reduce the number of dimensions to 10 [25], whle in [26] two iterations of Laplacian Eigenmaps [36] is used to obtain a three-dimensional point cloud. Both are non-linear dimensionality reduction tools constructing a lower-dimensional embedding by optimizing the global and local distances, respectively, based on a neighborhood graph (connecting each point to its k nearest neighbors).

Downsampling

While dimensionality reduction is a way of denoising data, noisy or false measurements may still be mapped to outliers in the projection or may influence the performance of the dimensionality reduction tool (Fig. 5c,d). The dataset may also contain irrelevant or redundant measurements. Yet another problem is the computational cost of analyzing large amounts of data. Thus, one may try to reduce the number of points whilst preserving the critical information in the point cloud.

Rybakken et al. [24] simplified the point clouds to 100 points using both radial (distance) downsampling [37] and topological denoising [38]. Radial downsampling starts by picking a single point in the point cloud and removing the points within radius ϵ. The remaining closest point is then chosen, and the process repeated until no distances are less than ϵ. Topological denoising first chooses a random subset $X_0 \subset X$ and then moves the chosen points towards a topologically relevant region by recursively maximizing

$$D_n(x) = \frac{1}{|X|} \sum_{y \in X} \exp^{-\frac{||x-y||^2}{2\sigma^2}} - \frac{\lambda}{|X_n|} \sum_{y \in X_n} \exp^{-\frac{||x-y||^2}{2\sigma^2}},$$

where $n = 0, 1, \ldots, N$ (for some user-defined N) is the iteration step, λ and σ hyperparameters and X_n given by (for some learning parameter a)

$$X_{n+1} = \left\{ x + a \frac{\nabla D_n(x)}{\max\limits_{x' \in X_0} (|\nabla D_0(x')|)} \mid x \in X_n \right\}.$$

$D_n(x)$ may be understood as the difference between two kernel density estimators – the first term representing the whole data set X, moving the sampled points towards dense regions, and the second representing X_n, repelling the sampled points away from each other. In [25], the local density of each point was used to remove 5% of the 'sparsest' points, while Rubin et al. discarded all points for which the number of active neurons were less than 15 (of the 62 and 41 neurons in the mice analyzed) [26].

Note that other more or less well-motivated choices, such as time bin size, 1-D Gaussian smoothing, square rooting, binarizing etc. were also made to process the spike times in the different studies. Furthermore, the spike times analyzed were first sorted semi-automatically based on processed electrophysiological recordings, using the software KlustaKwik[39] followed by manual curation [23]. In summary, experimental (neural) data requires extensive preprocessing carefully designed to preserve the information contained in the data.

2.2 UMAP construction

Since its introduction in 2018, UMAP [28] has become a popular dimensionality-reduction method in neuroscience [40] and other fields [41]. The motivating idea behind UMAP is that the data are uniformly sampled from a low-dimensional manifold, but that this is not reflected in the distances given by the associated metric in the (high-dimensional) ambient space. UMAP tries to reconstruct the geodesic distances of the low-dimensional manifold by constructing local metrics for which the points in the dataset are uniform and then merge these to a global metric. However, the local metrics are usually incompatible. To solve this, UMAP converts the local metric spaces to *fuzzy simplicial sets*, for which there exists a natural way of merging the objects into a 'global' fuzzy simplicial set called the 'UMAP complex' [42]. While UMAP uses this construction to optimize the representation of a low-dimensional point cloud, we convert the UMAP complex back to an (extended pseudo-)metric space and apply persistent cohomology to the Rips filtration of the so-obtained geodesic distances. This may be seen as analyzing the UMAP complex directly and understanding its topological features through persistent cohomology. Thus, we circumvent the problem of initializing and optimizing a low-dimensional embedding [43, 44] while providing a useful preprocessing tool for persistent cohomology based on the uniformity assumption. In the following, we give a theoretical description of the construction of the UMAP complex, drawing largely from the work in [27] and [28].

Uniformity assumption. The main assumption behind UMAP is that a dataset is uniformly sampled on a (low-dimensional) manifold embedded in ambient space. This motivates the construction of a metric, d_g, for which this is true. Intuitively, this means the volume of a ball containing a fixed number of points, k, should be the same independent of where it is centered. Hence, the ambient distances are normalized,

$$d_g(x,y) = \frac{d(x,y)}{d(x,x_k)},$$

where d is the ambient space metric and x_k its k-th nearest neighbor.

Ep-metric space. However, this may give incompatible distances. Letting $d_x(x,y) = d_g(x,y)$ for $y \in N_x$, the neighborhood set of x, and infinite otherwise, (X, d_x) defines an *extended pseudo-metric space* for each point $x \in X$. Here, 'pseudo' means that if $d(x,y) = 0$, x and y are not necessarily the same, and 'extended' that we allow for infinite distances. A morphism between two ep-metric spaces (X, d_x) and

(Y, d_y) is a function $f : X \to Y$ such that $d_y(f(x), f(y)) \leq d_x(x, y)$. This defines a *category* of ep-metric spaces, **EPMet**.

Fuzzy set. Classically, a fuzzy set is defined in terms of a set X with a function $\mu : X \to I$, where $I = (0, 1]$, assigning to each $x \in X$ a 'probability' of the point being an element of X. The fuzzy sets form a category **Fuzz** with morphisms $(f, h) : (X, \mu) \to (Y, \rho)$ consisting of functions $f : X \to Y$ and a relation $h : \mu \leq \rho f$ such that $\mu(x) \leq \rho(f(x))$ for all $x \in X$. A *sheaf* S on I is a functor $S : I^{\mathrm{op}} \to \mathbf{Sets}$ with *restriction maps*, $r_{b,a} : S([0, b)) \to S([0, a))$ for $a \leq b \in I$, satisfying certain conditions [45]. Barr [29] shows that the category of fuzzy sets are in equivalence with the subcategory of sheaves on I where the restriction maps are injections. We may then think of a section $S([0, a))$, corresponding to (X, μ), as the set of elements $x \in X$ with probability greater than a, i.e., $\mu(x) \geq a$.

Simplicial set. The *simplex category* $\mathbf{\Delta}$ has objects $[n] = \{1, \ldots, n\}$ for $n \in \mathbb{N}$, and morphisms order-preserving functions. A *simplicial set* is then a functor $X^s : \mathbf{\Delta}^{op} \to \mathbf{Sets}$. This may be seen as a sequence of sets X_n with maps $d_i : X_n \to X_{n-1}$ and $s_j : X_n \to X_{n+1}$ satisfying certain relations [46]. The elements of X_n are called the n–simplices. The standard simplicial set Δ^n is the functor $\hom_{\Delta}(-, [n])$ with m-simplices $f : [m] \to [n] \in \mathbf{\Delta}$. By the Yoneda lemma, $x \in X_n$ is uniquely described as a natural transformation $x : \Delta^n \to X$.

Fuzzy simplicial set. A *fuzzy simplicial set* is a simplicial object in **Fuzz**, i.e., a functor $X^f : \mathbf{\Delta}^{\mathrm{op}} \to \mathbf{Fuzz}$. We may also define this as a sheaf $X^f : (\mathbf{\Delta} \times I)^{\mathrm{op}} \to \mathbf{Set}$. The set $X^f([n], [0, a))$ may then be seen as the n-simplices with probability $\geq a$. Denoting $\Delta^n_{<a}$ as the functor $\hom_{\Delta \times I}(-, ([n], [0, a)))$, we may write the fuzzy simplicial set X^f as the colimit

$$\operatorname*{colim}_{\Delta^n_{<a} \to X^f} \Delta^n_{<a} \overset{\cong}{\to} X^f.$$

X^f can thus be thought of as a simplicial set for which every simplex has probability less than or equal to its boundary simplices. The category of fuzzy simplicial sets and natural transformations is called **sFuzz**.

UMAP complex. The UMAP construction is defined as the (global) fuzzy simplicial set constructed from a family of ep-metric spaces $\{(X, d_x)\}_{x \in X}$. We will assume X is finite in the following.

First, we define a (realization) functor $\mathrm{Re} : \mathbf{sFuzz} \to \mathbf{EPMet}$ by setting

$$\mathrm{Re}(\Delta^n_{<a}) := (\{x_0, \ldots, x_n\}, d_a)$$

where

$$d_a(x_i, x_j) = \begin{cases} -\log(a) & \text{if } i \neq j \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

For $a \leq b$ we have $d_a \geq d_b$, and the map $F : \Delta^n_{<a} \to \Delta^m_{<b}$, induced by $f : [n] \to [m]$,

$$(\{x_0, \ldots, x_n\}, d_a) \mapsto (\{x_{f(0)}, x_{f(1)} \ldots, x_{f(n)}\}, d_b),$$

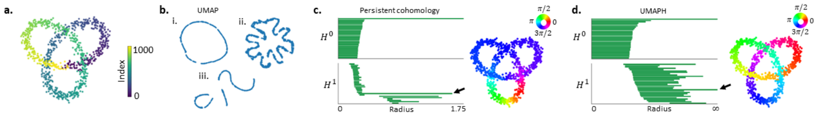

Fig. 6: Example application of UMAPH compared to UMAP and persistent homology. **a.** Two-dimensional projection of 1000 samples from a trefoil knot embedded in 3-D ambient space with added noise. Coloring depicts angular progression along the knot. **b.** Depending on different choices in the parameters, a 2-D UMAP projection unknots the underlying circular feature (i), though with possibly awkward geometry (ii). Using a random (instead of the default Laplacian Eigenmaps [36]) initialization of the low-dimensional embedding in UMAP breaks the knot into disjoint pieces (iii). **c.** Applying persistent cohomology gives a barcode showing three circular features (left). Each feature represents the three 'foils' of the knot, as seen in the circular coordinates of each, shown for one example feature (right). **d.** Constructing the Rips filtration based on the geodesic distances from the UMAP construction (using a small number of neighbors, k) and computing the persistent cohomology (i.e., applying UMAPH), gives a barcode showing an infinitely-lived circular feature (left). The corresponding angular coordinates describe the underlying circle of the knot (right).

is then a non-expansive map. This extends to any $X^f \in \mathbf{sFuzz}$,

$$\mathrm{Re}(X^f) := \operatorname*{colim}_{\Delta^n_{<a} \to X^f} \mathrm{Re}(\Delta^n_{<a}).$$

Conversely, there is a fuzzy singular set functor, $\mathrm{Si} : \mathbf{EPMet} \longrightarrow \mathbf{sFuzz}$, defined through extension on $\Delta \times I$:

$$\mathrm{Si}((Y, d_Y)) : ([n], [0, a)) \mapsto \hom_{\mathbf{EPMet}}(\mathrm{Re}(\Delta^n_{<a}), (Y, d_Y)).$$

We may thus obtain a global fuzzy representation – the UMAP complex of X – by converting $\{(X, d_x)\}_{x \in X}$ to fuzzy simplicial sets and taking the *fuzzy union*,

$$\bigcup_{x \in X} \mathrm{S}((X, d_x)).$$

Computations. In practice, each ep-metric space (X, d_x) is converted to a fuzzy simplicial set with probablity function $\mu_x(y) = \exp(-d_x(x, y))$. The UMAP complex may then be thought of as a simplicial complex with weighted simplices given by extending the symmetrized edge weights

$$\mu(x, y) = \mu_x(y) + \mu_y(x) - \mu_x(y)\mu_y(x).$$

To limit the effect of the choice of k in the approximation of the geodesic distances, d_x is defined as $d(x, y)/\sigma_x$ (for $y \in N_x$ and ∞ otherwise), where σ_x is determined by enforcing $\sum_{y \in N_x} \mu_x(y) = \log_2(k)$.

In UMAP, a dimension-reduced point cloud is initialized through Laplacian Eigenmaps (or randomly initialized, see Fig. 6b) and the *cross fuzzy entropy* between the UMAP complex of the low-dimensional point cloud and the high-dimensional one is minimized through stochastic gradient descent. This changes the layout of the low-dimensional embedding to approximate the topological representation of the original dataset [28].

2.3 Summary of Article I, part 1: Combining persistent homology and UMAP

Instead of projecting the data, X, to a lower-dimensional representation, we realize the UMAP complex: $\mathrm{Re}(\cup_x S((X, d_x))) = (X, d_g)$ with

$$d_g(x, y) = -\log(\mu(x, y)).$$

This is used in computing the Rips filtration and is equivalent to forming clique complexes of the 1-skeleton of the UMAP complex, filtering based on decreasing probability (i.e., the 'least likely' simplices are added towards end of the filtration). This assumes the probability of the higher-dimensional simplices is exactly the minimum of that of its boundary simplices. The persistent (co)homology (PH) of this filtration allows a visualization of the UMAPH complex beyond three dimensions and sidesteps the choice of initializing a low-dimensional embedding, as well as several hyperparameters used in the optimization, both of which may radically change the outcome of the method [43, 44]. At the same time, we show that it can reduce the computational cost of PH and provide an informative description of the underlying manifold even in the presence of noise. Extending the method with cohomological coordinatization additionally gives an informed way of dimension-reducing the data. We call this combination of tools *UMAPH* ('Uniform Manifold Approximation and Persistent Homology', Fig. 6).

These benefits are shown in the example of a noisy 8^3-sample of the three-torus $T^3 = S^1 \times S^1 \times S^1$ (see Extended Data Fig. 5c in the article), where the barcode describes the correct homology ($\beta_i = (1, 3, 3, 1)$) and cohomological parametrization matches the three circular features sampled. This showcases both how UMAPH can find and dimension-reduce the embedded (high-dimensional) space and how it can considerably reduce the computational cost of persistent cohomology by naturally simplifying the filtration through a limited number of neighbors (e.g., computing H^3 for the full filtration based on Euclidean distances was not permissible in Ripser with the available computing power due to memory requirements).

Furthermore, we introduce *fuzzy downsampling* based on the geodesic distance to reduce the point cloud (Fig. 7). Starting with a random initial point, we iteratively choose the $(i + 1)$-th point as

$$\max_{x \in X - X_i} \Big(\sum_{y \in X - X_i} \mu(x, y) \Big),$$

i.e., x_{i+1} is the point with the highest mean edge probabilities in the residual point

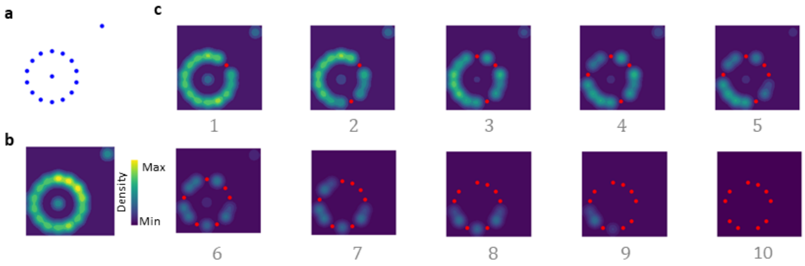

Fig. 7: Illustration of fuzzy downsampling technique. **a.** Example circular point cloud with two additional 'outliers'. **b.** Density plot of the objective function. Light colors indicate high density. **c.** The ten first iterations of the downsampling scheme. The sampled points (red) are chosen based on maximizing the density function in each step.

cloud $X - X_i$. This is equivalent to

$$\max_{x \in X}\Big(\sum_{y \in X} \mu(x,y) - \sum_{y \in X_i} \mu(x,y)\Big).$$

Inserting $\mu = \exp(-d_g(x,y))$ we see that this scheme iteratively maximizes two (unnormalized) kernel density estimators similar to D_n in the topological denoising algorithm.

By applying UMAPH to the downsampled (four first) principal components of firing rates of a cluster of the recorded head direction cell network previously described [22], we were able to detect and decode the internal head direction representation during wakefulness, REM *and* SWS sleep (see Extended Data Fig. 6 in the article).

3 Finding topological structure in neural data

External and internal input is somehow perceived and transformed by the neurons in our brains to elicit an action in the agent. A long-standing question is how this happens. To gain insight into the inner computations of the brain, neuroscientists have commonly recorded the electrochemical activity of single neurons and related it to some observable external variable such as visual or spatial stimuli. The neural responses to the variable may suggest the *function* of the recorded neurons - what piece of information each neuron produces or *encodes*. Neurons portraying an explicit response to a specific feature of the stimuli have often been understood through the representated variable. In 1959, Huber and Wiesel showed that single cells in the visual cortex of cats became selectively active to a particular orientation of visual bars presented to the cat [47]. We say these neurons are *tuned* to specific orientations. A population of such neurons are then thought to cover the range of orientations, thus collectively encoding this visual feature. Thus, we may define a *neural ensemble* as a group of neurons carrying a specific representation. Although the brain cannot directly compare the neural responses to the external feature, it can form relationships to the pattern of responses. For instance, downstream neurons are differentially influenced depending on which (and, possibly, to which degree the) neurons in the network are active. In this way, the information bestowed in the ensemble activity is conveyed or "read out".

3.1 Neural representations

"*Can one hear the shape of a drum?*" Similar to asking whether knowing the frequencies made by a drum suggests its shape [48], we ask whether we can infer the shape of the neural representation through recordings of the neural activity. Conceptually, both the activity *state space* (the space of possible activity patterns) and the *correlation structure* of an ensemble reflect the variable it holds (Fig. 8). The neural responses of the orientation network should thus correspond to the possible orientations and transitions between these. Explicitly, we think of the collective responses as constrained to a circular manifold and smoothly traversing this space in alignment with the visual stimulus. Similarly, seeing that the activity of each neuron is confined to a small set of angles, two correlated neurons should encode neighboring angles and share a strong 'functional connection'. This illustrates how knowing what covariate the network encodes gives us an indication of how the neurons should behave and what topology best describes this behavior. Conversely, knowing the topology of the neural activity should tell us something about what information the activity is storing [49, 50].

Recent breakthroughs in electrophysiological recordings and calcium imaging have allowed for thousands of cells to be simultaneously recorded over consecutive days [51–53], giving the means to analyze neural ensembles during various brain states and tasks. For instance, Gallego et al. find the activity patterns of motor cortical recordings in monkeys to be confined to a low-dimensional "neural

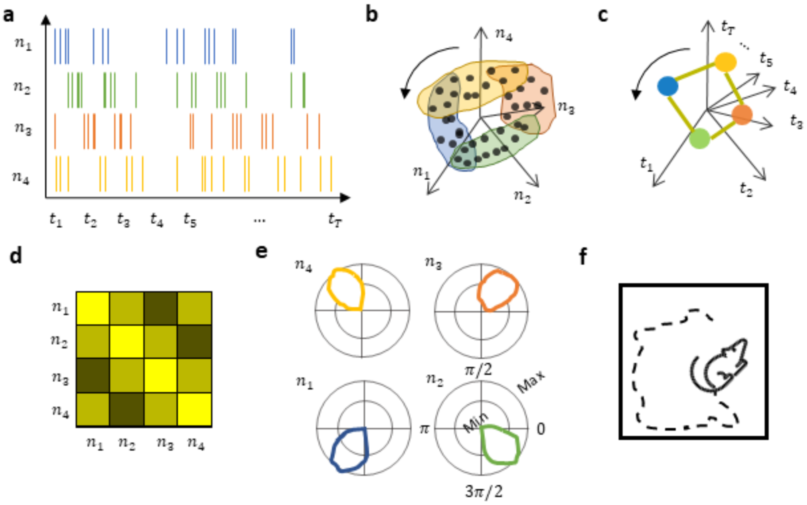

Fig. 8: The global shape of a neural representation is found in both the state space and the correlation structure of the neural ensemble. **a.** Example spike trains describing the spiking activity of four neurons n_i. The population activity described by the spike trains may be sampled at times t_i and arranged in an $N \times T$ data matrix, X, where $(N = 4)$ neurons. **b.** The columns of X, called *population vectors* (black), describe points in an N-dimensional space where each dimension describes the activity of neuron n_i. Time-dependent transitions (indicated by arrow) between the vectors are thought to describe (smooth) dynamics along the underlying manifold of the sampled *state space*. The colored groups indicate the vectors for which the neurons are (highly) active. **c.** The rows of X describe points in a T-dimensional space where each dimension gives the activity at time t_i. Transitions between neurons (black arrow) describe functional relations. Edges indicate high correlation between the neurons. Note how the resulting graph describes the overlap of the groups in **b**. **d.** Correlation diagram of pairwise correlations between activity traces of each neuron (i.e., between rows of X). **e.** The mean activity of each neuron as a function of an (internal or external) covariate shows peaked responses, i.e., tuning towards a given angle of the circular representation. **f.** The activity described in **a-e** may be seen as both the product and the cause of a disctinct behavior in the animal, for instance the head direction of a freely foraging rodent.

manifold" preserved across task [54, 55], and one of the first applications of topological data analysis in neuroscience showed that the topology of a population state space in primary visual cortex (V1) in monkeys is spherical [56]. In the following, we describe neural representations for navigation and localization and how the methods in the previous sections can and have been used to infer topological structure in neural data, particularly for studying spatial representations.

Spatial representations

Several spatially-tuned neurons, thought to partake in encoding representations of an animal's current whereabouts, have been discovered, such as place cells [57], grid cells [58], object vector cells [59], border [60] and boundary vector cells [61] and head direction cells [62] (Fig. 9). In experiments, these are usually found to describe the planar world of a rodent's experimental environment. For instance, the place cells are thought to encode the specific two-dimensional coordinate of the animal's location, with each place cell selectively active close to a preferred position in this coordinate system. The region of elevated activity is called the cell's *place field*. Curto et al. show that the union of place fields can form a good cover of the environment [9]. The nerve lemma then states that the simplicial complex formed by replacing n coactive place cells with n-simplices shares the topology of the physical environment. Moreover, by assuming a stereotypical shape of the place fields, they describe geometrical details of the animal's simulated trajectory. On the other hand, Dabaghian et al. suggest a purely *topological* neural representation [63] of the place cells, observing a preservation of the relative sequence of recorded place fields visited along a linear track across changes in the geometry, but not changes in its topology. In [16], Giusti et al. apply persistent homology to the clique complex of recordings of place cells in rat hippocampus. Using statistical summaries (*Betti curves*) of the barcodes they find a 'geometrical' organization in the complex. More puzzling, similar organization is also seen during wheel running and REM sleep, indicating that it is not the spatial stimulus that forms the network, but rather its intrinsic properties. How this relates to the re-configuration of map-like representation of place cells as the animal enters a novel environment [64] is still an open question. This exemplifies the need to understand the network structure without assuming specific covariate function [65].

We have already mentioned how the authors in [24–26] all expose the ring topology of the head direction cell network recorded from a population of cells in mouse thalamus [23]. These are cells specifically tuned to a single direction for which the animal's head is facing with respect to the environment [66]. Hence, the sequence of activity reflects the way the animal's head turns, confining the activity states to a ring manifold. While these studies look at the differences in activity patterns at different time frames to inferring the topology, the above-mentioned place cell studies look at differences in firing activity per neuron to understand the underlying representation. The comparison between these two ways of assessing

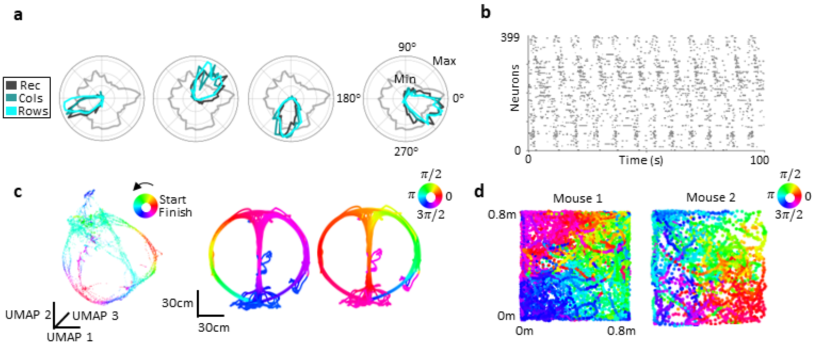

Fig. 9: Examples of neural representations for spatially-tuned neurons. **a.** Tuning to recorded head direction ('Rec') and decoded circular representation of correlation structure ('Rows') and state space ('Cols'), as in Fig. 8a of Article II, of four head direction cells recorded in the entorhinal cortex [52]. Gray line indicates occupancy of recorded head direction. **b.** Spike trains of hippocampal place cells [67] as a mouse runs on a virtual repeating track. The spike trains are sorted according to the decoded circular representation found in the correlation structure (Fig. 8b of Article II), implicitly showing the repeated progression along the track. **c.** Hippocampal place cell recording portray the environmental structure of the figure-eight task the mice are performing during recording (data from [68], unpublished analysis). Left, 3-D UMAP projection of the population vectors. Right, decoding the two longest-lived circular features in the barcodes of the UMAP projection reveals the coordinates of left and right circles of the environment. **d.** Decoding the longest-lived circular features of the correlation structure of entorhinal/parasubicular population recordings in two mice reveals border-vector representation (data from [69], unpublished analysis).

topological structure in neural data (Fig. 8) is the basis of Article II.

3.2 Summary of Article II: A topological perspective on the dual nature of the neural state space and the correlation structure

We try to relate two distinct ways of analyzing population recordings: 1) studying the correlation structure through the pairwise correlations of activity traces, and 2) understanding population coding by analyzing the neural state space as sampled by the population vectors. If we regard the population recordings as a data matrix, X, where the number of rows depict the number of neurons and the number of columns the time frames, the two approaches respectively correspond to analyzing the matrix row-wise or column-wise. We define the *population vector complex* by forming groups of neurons if they are active in the same population vector and connect groups if they overlap. Every group is then given by the population vector it stems from and the complex corresponds to connecting sets of k vectors with a k-simplex if they contain the same neuron(s). Dually, the correlation complex groups population vectors which share an active neuron and similarly connects overlapping groups (Fig. 8b,c). Correspondingly, every group represents a distinct neuron and the complex is given by letting l neurons form an l-simplex if they are all coactive in the same vector.

If we assume a binarized matrix, X, *Dowker duality* then says the corresponding correlation and population vector complex are homotopy equivalent [50]. However, these constructs are extremely sensitive to noise, e.g. a single random firing of a neuron may distort the topology. Hence, we assign to each simplex the corresponding number of overlapping neurons or shared vectors and set thresholds for the minimum number that elicits connecting the groups. We can assess all thresholds by observing that this gives rise to a filtration of simplicial complexes in decreasing thresholds – n groups that share λ neurons (vectors) also share $\lambda' < \lambda$ neurons (vectors), inducing an inclusion of the n-simplex. Computing the persistent homology of this filtration allows us to assess all thresholds and see which features are dominant. However, the resulting barcodes for the filtration of population vector complexes and correlation complexes are now not necessarily the same as there is no obvious link between the respective weights.

We study two cases in which the barcodes differ: 1) when we have multiple ensembles and 2) when there are neurons with disconnected (non-convex) receptive fields. In 1), the correlation complex may reveal the presence of each network, while the population vector complex will not, while in 2) the population vector complex may carry the shape of the representation, while the correlation complex gets distorted. Consequently, when both complexes indicate the same topology, this strongly indicates the true shape of the ensemble representation. We suggest a framework for how this complementary behavior may guide the analysis and give a simulated example of extracting the underlying structure from mixed population recordings without relation to an external covariate (Fig. 7 in the article). Next, we showcase a few examples of how this applies to experimental recordings (Fig. 2 and 8 in the article), namely: the entorhinal head direction cell network (Fig. 9a), orientation-tuned neurons in primary visual cortex, place cells during a recurring VR task (Fig. 9b) and a grid cell module.

4 The topology of the grid cell network

In this section, we discuss the spatial representation formed by grid cells and see how it gives rise to a toroidal topology in the ensemble activity. We will then see how, using the methods previously described, the proposed structure is found in the state space of recordings of grid cell ensembles.

4.1 Grid cells and the toroidal hypothesis

Grid cells are neurons found in the medial entorhinal cortex (MEC) specified by their firing pattern in space, each grid cell displaying multiple activity fields forming a 2-D hexagonal lattice in the environment [58, 70]. The pattern is characterized by its *scale* (the distance between fields and the size of a field) and orientation relative to the borders of the environment (Fig. 10a). These features are shared across ensembles (or *modules*) of grid cells, but each cell has its own spatial phase displaced from the other cells in the module [71]. There is a topographical organization to the scale of the distinct grid modules and the combination of modules may encode a vast amount of unique spatial codes [72]. They are therefore thought to provide the main spatial input to place cells [73, 74]. In contrast to the remapping of place fields in place cells [75], the functional relationships between the grid cells are preserved across different environments and in sleep [58, 76, 77]. There are two main theories for how the grid cell activity arises. In continuous attractor network (CAN) models, the network connectivity is a priori designed to follow the assumed geometry of the manifold of the network and the collective activity stabilizes to one or more localized bumps which may be translated smoothly along the state space. Grid cell CANs are shown to accurately integrate velocity and head direction inputs to generate the characteristic hexagonal responses of grid cells and drive a corresponding flow of the network pattern [78–80]. A second model hypothesizes a self-organization of the grid-cell pattern, arising due to interference between a global signal and cell-specific oscillations modulated by the rat's velocity [81].

A *2-torus* is the Cartesian product of two circles ($S^1 \times S^1$), allowing the periodic traversal along two dimensions. Hence, it has long been suggested that the population activity of grid modules should display a toroidal structure [80]. This may be seen by studying the fundamental rhomboidal tile of the grid cell pattern. We see that the tile repeats along the two spatial dimensions, and a (single) grid cell module cannot differentiate which of the tiles the animal is in, but only the location within each tile. Consequently, the representation of the module is identified by appropriately gluing the boundaries of the rhombus, giving rise to a torus. Each grid cell has a single receptive field on the torus (Fig. 11a) [49]. Assuming all phases are covered, the grid fields define a good cover of the torus, meaning the activity space should expose a toroidal topology.

However, navigation in natural environments is more challenging than in simple, familiar, obstacle-free enclosures, and the hexagonal pattern is seen to change de-

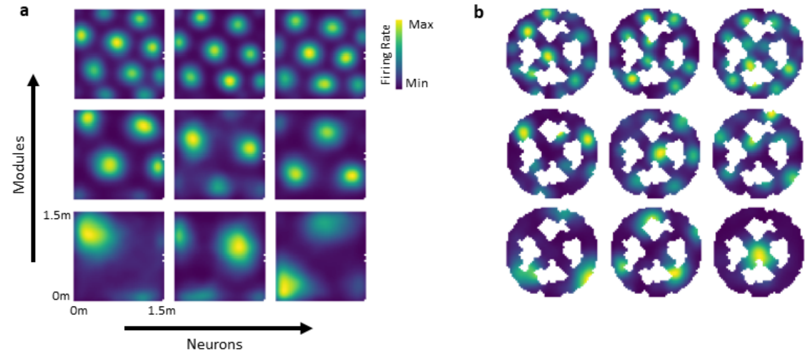

Fig. 10: The spatial specificity of grid cells display a hexagonal pattern in an open field and divide into ensembles or *modules* based on the pattern's grid scale. **a.** Nine example neurons of entorhinal grid cells from three modules during open field foraging (data from [2]). **b.** The mean firing pattern of the same neurons as in **a** during free foraging on a wagon-wheel track.

pending on the geometry of the environment (Fig. 10b). In novel environments, in environments of different shapes, merged arenas, near boundaries and objects and during goal-related tasks the hexagonal pattern of grid cells can appear distorted [82–86]. These distortions challenge the idea of the grid cell network procuring a precise metric of external space and begs the question whether the internal representation is similarly distorted – 1) is the neural representation toroidal and 2) is it so, regardless of the geometry of the environment? With a persistent internal representation, periodic in two dimensions, the proposed modular grid cell coding would still be permissible despite the hexagonal symmetry of the grid cell pattern broken in different environments. Moreover, it would support the continuous attractor model of grid cells where the pre-defined pairwise connections in the population may be preserved independent of environmental changes, while in self-organizing models it would require each neuron to learn a spatial modulation in such a way that the same pattern appears across the whole population for each new environment and brain state. In articles I (part 2) and III, we give an affirmative answer, showing that the topological structure of the population activity of a grid cell module is toroidal and that this structure is preserved across environments and brain states.

4.2 Summary of Article III: Toroidal topology of population activity in grid cells

With advances in electrophysiological recording technology [51], we were able to record a large number of grid cells simultaneously in the medial entorhinal cortex of three rats during different experimental conditions, with up to ~ 150 neurons in each of the six modules. Each session lasted up to several hours, giv-

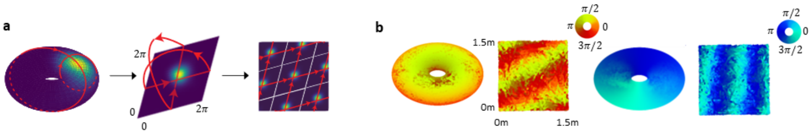

Fig. 11: The state space of a grid cell module is toroidal. **a.** The mean firing rate of a grid cell as a function of the toroidal ensemble representation. Left, the neuron's firing response reveals a preferred location on the torus. The red arrows indicate the two circles of the torus. Middle, unwrapping the toroidal surface may be seen as a rhombus with periodic boundaries. Right, stacking multiple rhomboidal tiles reveals the hexagonal firing rate pattern. **b.** Decoding the toroidal representation describes two circular coordinates (left and right) for each time frame. Coloring each spatial position visited by the rat with the simultaneous circular coordinate reveals a stripe-like pattern depicting the two spatial axes of the firing pattern of the module.

ing large and noisy datasets, so we used a similar pipeline as described in Section 2. First, excluding time frames of low activity, PCA was applied to each dataset, projecting the data to a six-dimensional embedding (see Extended Data Fig. 4 in the article). Next, we used the fuzzy downsampling technique and applied persistent cohomology to the UMAP complex of the reduced point cloud. Despite an apparent lack of hexagonality in the spatial selectivity of individual neurons in a wagon-wheel environment and a decoupling of neuronal activity to positional cues in sleep, we identified toroidal structure of the state space across open field and a wagon-wheel environments, as well as in sleep (both SWS and REM).

Using the cohomological coordinatization procedure described in Section 1.2, we obtained time-varying toroidal coordinates from the two most persistent circular features found in the barcode. Mapping these coordinates as a function of space displayed the two axes of the hexagonal pattern of the grid cell module (Fig. 11b). Computing single cell mean firing rate activity in binned toroidal coordinates we observed how each cell had single firing fields on the rhomboidal tile. This tuning was preserved across the different conditions, with the same preferred toroidal location, indicating a fixed network structure across each condition. We additionally tested how well the toroidal coordinates explained the data as compared to recorded physical position. Using both an information content score and the explained deviance of a GLM model fitted to each coordinate set, the toroidal coordinates were found to explain the data better.

However, during SWS, two modules lacked a clear toroidal signature of their state space. Further investigation of these revealed subpopulations of grid cells with different spiking statistics and with varying toroidal specificity. Confining the analysis to the subpopulation with bursty firing and weak theta-modulation again uncovered a torus for one of the modules.

4.3 Summary of Article I, part 2: Uncovering 2-D toroidal representations in grid cell ensemble activity during 1-D behavior

The UMAPH framework (described in Section 2) was first applied to calcium imaging data of the entorhinal cortex of freely moving mice [52, 69]. Without clustering, a toroidal structure was revealed in the whole population recording in two mice, showing that the locally confined Ca^{2+}-imaging permits recording a large number of cells from the same grid cell module. In three recording days of a third mouse, we first found a grid cell module by clustering the neurons based on their spatial pattern and showed that the toroidal structure in the open field recording was similarly carried in a head-fixed wheel-running task. That is, we found a preserved 2-D toroidal representation of the grid cell activity, despite the 1-D non-spatial nature of the task. Head-fixed animals in 1-D environments, often in combination with a virtual reality task, is a popular experimental setup with several benefits, allowing experiments with large equipment, easier preparations and better control of tracking [87]. However, the equipment usually does not allow 2-D tracking, making it difficult to assert the spatial tuning of the recorded neurons [88]. Having found the toroidal representation in a wheel-running task, motivated using UMAPH in such experiments despite a lack of grid cell identification. Thus, we studied the data recorded by Campbell et al. [89], where the head-fixed mice ran on a continuously repeating virtual linear track under different manipulations of the task. The neurons were clustered based on time-lag correlations of their firing rate activity. Using the largest clusters, we were able to obtain ensembles which exhibited toroidal representations. Each neuron had a single receptive field on the hexagonal torus, with preserved toroidal location across conditions. Furthermore, the decoded internal dynamics revealed smooth trajectories on the neural manifold evolving in line with the mouse's progression along the VR-track. In gain manipulation trials, the gain could be predicted by estimating the length of the internal trajectories. This establishes the advantages of using this approach to study the topological structure of neural ensemble activity and understand the dynamics and function of the ensemble computation without reference to external covariates.

5 Conclusion

5.1 Summary

In this thesis, we study topological structure found in neural data through the means of TDA. Persistent cohomology stands as a versatile descriptor of the underlying shape of data, and we demonstrate its efficacy in describing neural representations in population recordings in animals. However, naïvely applying the methods to a given dataset (as is) may not identify interpretable structure, and we have seen how careful preprocessing is often needed to meet the cacophony of experimental, high-dimensional data. To this end, we described how an intrinsic uniformity assumption may be helpful, creating a fuzzy topological representation of the data. Combining such a framework with cohomological coordinatization allows further insight into circular features detected in the given point cloud.

TDA is still an emerging field, also when it comes to practical applications requiring its adeptness. With its wealth of data and inherent complexity, neuroscience needs new ways of seeing and has for some time now been a field of active interest within the TDA community. We have seen how the topological structure of neural data may be connected to understanding the functions and computations of groups of neurons. More so, we observed how information may be found both in the rows and columns of the data matrix studied, with a neural ensemble exhibiting the same topology in either. Although this duality of coactivity matrices has been known and both approaches are used, we emphasize the benefit of their complementary information, and provide simulated and real examples in neural data.

We furthermore presented in this thesis the first demonstration of inferring the toroidal state space of grid cell modules. This serves as a clear example of the need for topological tools, where other approaches may fail to facilitate easy interpretation. Moreover, the discovery of a preserved structure during different behavioral tasks, such as head-fixed wheel running and in sleep, shows how the methods can reveal unprecedented information about neural ensembles. We end this thesis with some reflections on the drawbacks of the methods studied and outline some future directions.

5.2 Discussion and future directions

Pre- and postprocessing. We have already mentioned the sensitivity of persistent homology to the noise and size of experimental data, and while UMAPH may lessen the burden (by introducing the number of neighbors, k, as a parameter), it contains the same computational bottleneck. Iteratively computing PH on subsamples of a dataset and merging these into a global topological description of the data [90, 91] could be a way to improve on this. Additionally, the choice of k may radically change the outcome of UMAPH, and allow for barcodes with bars of infinite length, which often have a less straightforward interpretation. In general, it is difficult to know whether a performed analysis not showing the predicted results fails to do so because the hypothesis was wrong or because the parameters were set wrongly. Does this make the analysis a "*failed experiment ... producing data not proving anything*" [92]? We maintain that it is possible to make justified choices but believe a better theoretical reasoning behind parameter tuning is still needed [93]. Furthermore, the readout of barcodes is in any case difficult and the statistical properties of barcodes are actively studied [94]. Outlier detection [95], different statistical summaries [96, 97] and other heuristics [98, 99] have been used to distinguish signal from noise. In this thesis, we saw how the cohomological coordinatization was critical in interpreting the barcodes, and the intrinsic dynamics and well-behavedness of the decoding may be used to identify which features are of interest. Moreover, the coordinates may be actively used to uncover more information in the dataset. For instance, in Article II, we used the receptive fields on the inferred representation (similar to using information scores of the tuning as in [24]) to filter the neurons. Another option (as seen in [24, 100]), is using the coordinates in a GLM model to subtract the contribution of the (decoded) covariates from the neural data, to uncover topological information hidden in the residual activity. Developing the coordinatization scheme to other (higher-dimensional) spaces and homology groups will allow a more detailed account of the underlying structure to be revealed [101–103]. In this respect, incorporating statistical models in fitting the coordinates could give better robustness and interpretability of the inference.

Other tools. While (persistent) homology gives an important description of a topological space, it is not the only signature, nor does it fully describe the topology of the space. As previously mentioned, the homology of a torus and a sphere with two circles attached is equivalent, but these spaces are not (topologically) the same. To separate them, one may test different properties [104] and (persistent) cohomology can for instance be given a product structure which distinguish these spaces [105]. The homotopy type is a stronger characterization of the topology [19], but further theoretical and computational developments are needed [106, 107]. However, extensions of persistent homology and other (co)homology theories, such as sheaf cohomology, may refine and extend the topological inference [19, 108–110]. For instance, in [111] we study the flow of homology generators between height levels of a simplicial set. This is similar to levelset zigzag [112] and the MAPPER algorithm [56], and may give rise to a more refined version of the barcode [113, 114]. Furthermore, we note statistical approaches such as [115] and [116], where the specific latent space manifold(s) are assumed. Testing the fit to data for different manifolds allows determining the most likely topology of the underlying space. Incorporating tools from TDA to initialize the guessed manifold(s) could replace the assumptions made in these models.

Prospects for neuroscience and neural networks. The computational technique of grid cells seems to provide a useful representation for encoding spatial variables (Fig. 12). This may extend to higher-dimensional spaces [117] and non-spatial dimensions [118], and similar functioning is possibly replicated in other brain regions and for more conceptual spaces [119, 120]. The methods studied

Fig. 12: Fantastical illustration of the toroidal spatial representation of the grid cells. Credits: Kavli Institute for Systems Neuroscience, NTNU and Helmet.

and developed in this thesis may thus extend to other domains as well, such as in machine learning where the computational specifics of deep neural networks are difficult to interpret [121, 122]. For instance, in [123] the authors show that a recurrent neural network optimized for path integration suggests a toroidal population underlying such computation, but that these do not correspond to the grid cell population. Essentially, the functional role and the structure of a neural population may be hard to pin down, with potentially plastic or conjunctive selectivity [124–126]. To grasp deeper higher-order relationships in the brain, we believe it is necessary to start by separating neurons into ensembles (e.g., through a specific barcode signature as we propose in [127]) and understanding each ensemble individually. Topology may help describe the relations between ensembles [128], and allow studying, for example, the collective readout of the grid cell modules and how these are orchestrated in a joint network with other ensembles (see ’Peer Review File’ in [2] for combined grid module analysis). Understanding the brain can have great societal impact and we note that technological advances already allows large-scale human brain recordings with high resolution [129]. Early diagnostics of brain pathology [130, 131], leveraging algorithmic insights of the brain in computations [132] or enabling self-governed robotic limb movement for prostheses [133, 134] are all viable goals in the years to come.

Bibliography

[1] E. Hermansen, D. A. Klindt and B. A. Dunn, ‘Uncovering 2-d toroidal representations in grid cell ensemble activity during 1-d behavior,’ *bioRxiv*, 2022.

[2] R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser and E. I. Moser, ‘Toroidal topology of population activity in grid cells,’ *Nature*, vol. 602, no. 7895, pp. 123–128, 2022.

[3] H. Edelsbrunner, D. Letscher and A. Zomorodian, ‘Topological persistence and simplification,’ in *Proceedings 41st annual symposium on foundations of computer science*, IEEE, 2000, pp. 454–463.

[4] A. Zomorodian and G. Carlsson, ‘Computing persistent homology,’ in *Proceedings of the twentieth annual symposium on Computational geometry*, 2004, pp. 347–356.

[5] V. De Silva, D. Morozov and M. Vejdemo-Johansson, ‘Dualities in persistent (co) homology,’ *Inverse Problems*, vol. 27, no. 12, p. 124 003, 2011.

[6] U. Bauer, ‘Ripser: Efficient computation of vietoris–rips persistence barcodes,’ *Journal of Applied and Computational Topology*, vol. 5, no. 3, pp. 391–423, 2021.

[7] K. Houston-Edwards, *How squishy math is revealing doughnuts in the brain*, en. DOI: `10.1038/scientificamerican1022-36`. [Online]. Available: `https://www.scientificamerican.com/article/how-squishy-math-is-revealing-doughnuts-in-the-brain/`.

[8] R. Ghrist, ‘Barcodes: The persistent topology of data,’ en, *Bulletin of the American Mathematical Society*, vol. 45, no. 01, 61–76, 2007, ISSN: 0273-0979. DOI: `10.1090/S0273-0979-07-01191-3`.

[9] C. Curto and V. Itskov, ‘Cell groups reveal structure of stimulus space,’ *PLoS computational biology*, vol. 4, no. 10, e1000205, 2008.

[10] M. Vejdemo-Johansson, ‘Sketches of a platypus: persistent homology and its algebraic foundations,’ *ArXiv e-prints*, Dec. 2012. arXiv: `1212.5398 [math.AT]`.

[11] M. B. Botnan, 'Three approaches in computational geometry and topology: Persistent homology, discrete differential geometry and discrete morse theory,' Master's thesis, Norges Teknisk- og Naturvitenskapelige Universitet, 2011.

[12] E. Hermansen, 'Topological detection in spatially and directionally tuned neural network activity,' M.S. thesis, NTNU, 2017.

[13] V. De Silva, D. Morozov and M. Vejdemo-Johansson, 'Persistent cohomology and circular coordinates,' *Discrete & Computational Geometry*, vol. 45, no. 4, pp. 737–759, 2011.

[14] E. Munch, 'A user's guide to topological data analysis,' *Journal of Learning Analytics*, vol. 4, no. 2, pp. 47–61, 2017.

[15] N. Otter, M. A. Porter, U. Tillmann, P. Grindrod and H. A. Harrington, 'A roadmap for the computation of persistent homology,' *EPJ Data Science*, vol. 6, pp. 1–38, 2017.

[16] C. Giusti, E. Pastalkova, C. Curto and V. Itskov, 'Clique topology reveals intrinsic geometric structure in neural correlations,' *Proceedings of the National Academy of Sciences*, vol. 112, no. 44, pp. 13 455–13 460, 2015.

[17] G. Carlsson, 'Topology and data,' *Bulletin of the American Mathematical Society*, vol. 46, no. 2, pp. 255–308, 2009.

[18] V. De Silva and M. Vejdemo-Johansson, 'Persistent cohomology and circular coordinates,' in *Proceedings of the twenty-fifth annual symposium on Computational geometry*, 2009, pp. 227–236.

[19] A. Hatcher, *Algebraic Topology* (Algebraic Topology). Cambridge University Press, 2002. [Online]. Available: `https://pi.math.cornell.edu/~hatcher/AT/AT.pdf`.

[20] D. Cohen-Steiner, H. Edelsbrunner and J. Harer, 'Stability of persistence diagrams,' in *Proceedings of the twenty-first annual symposium on Computational geometry*, 2005, pp. 263–271.

[21] E. V. Somasundaram, S. E. Brown, A. Litzler, J. G. Scott and R. R. Wadhwa, 'Benchmarking r packages for calculation of persistent homology,' *The R journal*, vol. 13, no. 1, p. 184, 2021.

[22] G. B. A. Peyrache P. Petersen, *Peyrache 2015*, 2015. DOI: `10.6080/K0G15XS1`. [Online]. Available: `http://crcns.org/data-sets/thalamus/th-1`.

[23] A. Peyrache, M. M. Lacroix, P. C. Petersen and G. Buzsáki, 'Internally organized mechanisms of the head direction sense,' *Nature neuroscience*, vol. 18, no. 4, pp. 569–575, 2015.

[24] E. Rybakken, N. Baas and B. Dunn, 'Decoding of neural data using cohomological feature extraction,' *Neural computation*, vol. 31, no. 1, pp. 68–93, 2019.

[25] R. Chaudhuri, B. Gerçek, B. Pandey, A. Peyrache and I. Fiete, 'The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep,' *Nature neuroscience*, vol. 22, no. 9, pp. 1512–1520, 2019.

[26] A. Rubin, L. Sheintuch, N. Brande-Eilat, O. Pinchasof, Y. Rechavi, N. Geva and Y. Ziv, 'Revealing neural correlates of behavior without behavioral measurements,' *Nature communications*, vol. 10, no. 1, pp. 1–14, 2019.

[27] D. I. Spivak, 'Metric realization of fuzzy simplicial sets,' *Self published notes, available online*, 2012. [Online]. Available: `https://www.semanticscholar.org/paper/METRIC-REALIZATION-OF-FUZZY-SIMPLICIAL-SETS-Spivak/a73fb9d562a3850611d2615ac22c3a8687fa745e`.

[28] L. McInnes, J. Healy and J. Melville, 'Umap: Uniform manifold approximation and projection for dimension reduction,' *arXiv preprint arXiv:1802.03426*, 2018.

[29] M. Barr, 'Fuzzy set theory and topos theory,' *Canadian Mathematical Bulletin*, vol. 29, no. 4, pp. 501–508, 1986.

[30] J. F. Jardine, 'Fuzzy sets and presheaves,' *arXiv preprint arXiv:1904.10314*, 2019.

[31] J. Jardine, 'Metric spaces and homotopy types,' *arXiv preprint arXiv:2012.09026*, 2020.

[32] J. Jardine, 'Stability for umap,' *arXiv preprint arXiv:2011.13430*, 2020.

[33] R. Bellman, 'Dynamic programming,' *Science*, vol. 153, no. 3731, pp. 34–37, 1966.

[34] C. C. Aggarwal, A. Hinneburg and D. A. Keim, 'On the surprising behavior of distance metrics in high dimensional space,' in *International conference on database theory*, Springer, 2001, pp. 420–434.

[35] J. B. Tenenbaum, V. d. Silva and J. C. Langford, 'A global geometric framework for nonlinear dimensionality reduction,' *science*, vol. 290, no. 5500, pp. 2319–2323, 2000.

[36] M. Belkin and P. Niyogi, 'Laplacian eigenmaps for dimensionality reduction and data representation,' *Neural computation*, vol. 15, no. 6, pp. 1373–1396, 2003.

[37] E. De Koning, *Radial distance simplification*, 2011. [Online]. Available: `https://psimpl.sourceforge.net/radial-distance.html`.

[38] J. Kloke and G. Carlsson, 'Topological de-noising: Strengthening the topological signal,' *arXiv preprint arXiv:0910.5947*, 2009.

[39] K. D. Harris, D. A. Henze, J. Csicsvari, H. Hirase and G. Buzsaki, 'Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements,' *Journal of neurophysiology*, vol. 84, no. 1, pp. 401–414, 2000.

[40] E. Becht, L. McInnes, J. Healy, C.-A. Dutertre, I. W. Kwok, L. G. Ng, F. Ginhoux and E. W. Newell, ‘Dimensionality reduction for visualizing single-cell data using umap,’ *Nature biotechnology*, vol. 37, no. 1, pp. 38–44, 2019.

[41] A. Diaz-Papkovich, L. Anderson-Trocmé and S. Gravel, ‘A review of umap in population genetics,’ *Journal of Human Genetics*, vol. 66, no. 1, pp. 85–91, 2021.

[42] R. Jardine, ‘Umap for the working mathematician,’ 2021.

[43] T. Chari, J. Banerjee and L. Pachter, ‘The specious art of single-cell genomics,’ *bioRxiv*, 2021.

[44] D. Kobak and G. C. Linderman, ‘Initialization is critical for preserving global data structure in both t-sne and umap,’ *Nature Biotechnology*, vol. 39, no. 2, pp. 156–157, 2021.

[45] B. R. Tennison, *Sheaf theory*. Cambridge University Press, 1975, vol. 21.

[46] P. G. Goerss and J. F. Jardine, *Simplicial homotopy theory*. Springer Science & Business Media, 2009.

[47] D. H. Hubel and T. N. Wiesel, ‘Receptive fields of single neurones in the cat’s striate cortex,’ *The Journal of physiology*, vol. 148, no. 3, p. 574, 1959.

[48] M. Kac, ‘Can one hear the shape of a drum?’ *The american mathematical monthly*, vol. 73, no. 4P2, pp. 1–23, 1966.

[49] C. Curto, ‘What can topology tell us about the neural code?’ *Bulletin of the American Mathematical Society*, vol. 54, no. 1, pp. 63–78, 2017.

[50] C. Giusti, R. Ghrist and D. S. Bassett, ‘Two’s company, three (or more) is a simplex,’ *Journal of computational neuroscience*, vol. 41, no. 1, pp. 1–14, 2016.

[51] N. A. Steinmetz, C. Aydin, A. Lebedeva, M. Okun, M. Pachitariu, M. Bauza, M. Beau, J. Bhagat, C. Böhm, M. Broux *et al.*, ‘Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings,’ *Science*, vol. 372, no. 6539, eabf4588, 2021.

[52] W. Zong, H. A. Obenhaus, E. R. Skytøen, H. Eneqvist, N. L. de Jong, R. Vale, M. R. Jorge, M.-B. Moser and E. I. Moser, ‘Large-scale two-photon calcium imaging in freely moving mice,’ *Cell*, vol. 185, no. 7, pp. 1240–1256, 2022.

[53] C. Stringer, M. Pachitariu, N. Steinmetz, M. Carandini and K. D. Harris, ‘High-dimensional geometry of population responses in visual cortex,’ *Nature*, vol. 571, no. 7765, pp. 361–365, 2019.

[54] J. A. Gallego, M. G. Perich, L. E. Miller and S. A. Solla, ‘Neural manifolds for the control of movement,’ *Neuron*, vol. 94, no. 5, pp. 978–984, 2017.

[55] J. A. Gallego, M. G. Perich, S. N. Naufel, C. Ethier, S. A. Solla and L. E. Miller, 'Cortical population activity within a preserved neural manifold underlies multiple motor behaviors,' *Nature communications*, vol. 9, no. 1, pp. 1–13, 2018.

[56] G. Singh, F. Memoli, T. Ishkhanov, G. Sapiro, G. Carlsson and D. L. Ringach, 'Topological analysis of population activity in visual cortex,' *Journal of vision*, vol. 8, no. 8, pp. 11–11, 2008.

[57] J. O'Keefe and J. Dostrovsky, 'The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat.,' *Brain research*, 1971.

[58] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser and E. I. Moser, 'Microstructure of a spatial map in the entorhinal cortex,' *Nature*, vol. 436, no. 7052, pp. 801–806, 2005.

[59] Ø. A. Høydal, E. R. Skytøen, S. O. Andersson, M.-B. Moser and E. I. Moser, 'Object-vector coding in the medial entorhinal cortex,' *Nature*, vol. 568, no. 7752, pp. 400–404, 2019.

[60] T. Solstad, C. N. Boccara, E. Kropff, M.-B. Moser and E. I. Moser, 'Representation of geometric borders in the entorhinal cortex,' *Science*, vol. 322, no. 5909, pp. 1865–1868, 2008.

[61] C. Lever, S. Burton, A. Jeewajee, J. O'Keefe and N. Burgess, 'Boundary vector cells in the subiculum of the hippocampal formation,' *Journal of Neuroscience*, vol. 29, no. 31, pp. 9771–9777, 2009.

[62] J. Ranck Jr, 'Head direction cells in the deep layer of dorsal presubiculum in freely moving rats,' in *Society of neuroscience abstract*, vol. 10, 1984, p. 599.

[63] Y. Dabaghian, F. Mémoli, L. Frank and G. Carlsson, 'A topological paradigm for hippocampal spatial map formation using persistent homology,' *PLoS computational biology*, vol. 8, no. 8, e1002581, 2012.

[64] A. Samsonovich and B. L. McNaughton, 'Path integration and cognitive mapping in a continuous attractor neural network model,' *Journal of Neuroscience*, vol. 17, no. 15, pp. 5900–5920, 1997.

[65] C. Curto, E. Gross, J. Jeffries, K. Morrison, M. Omar, Z. Rosen, A. Shiu and N. Youngs, 'What makes a neural code convex?' *SIAM Journal on Applied Algebra and Geometry*, vol. 1, no. 1, pp. 222–238, 2017.

[66] J. S. Taube, R. U. Muller and J. B. Ranck, 'Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis,' *Journal of Neuroscience*, vol. 10, no. 2, pp. 420–435, 1990.

[67] N. L. Pettit, X. C. Yuan and C. D. Harvey, 'Hippocampal place codes are gated by behavioral engagement,' *Nature neuroscience*, vol. 25, no. 5, pp. 561–566, 2022.

[68] P. C. Petersen and G. Buzsáki, 'Cooling of medial septum reveals theta phase lag coordination of hippocampal cell assemblies,' *Neuron*, vol. 107, no. 4, pp. 731–744, 2020.

[69] H. A. Obenhaus, W. Zong, R. I. Jacobsen, T. Rose, F. Donato, L. Chen, H. Cheng, T. Bonhoeffer, M.-B. Moser and E. I. Moser, 'Functional network topography of the medial entorhinal cortex,' *Proceedings of the National Academy of Sciences*, vol. 119, no. 7, e2121655119, 2022.

[70] C. F. Doeller, C. Barry and N. Burgess, 'Evidence for grid cells in a human memory network,' *Nature*, vol. 463, no. 7281, pp. 657–661, 2010.

[71] H. Stensola, T. Stensola, T. Solstad, K. Frøland, M.-B. Moser and E. I. Moser, 'The entorhinal grid map is discretized,' *Nature*, vol. 492, no. 7427, pp. 72–78, 2012.

[72] I. R. Fiete, Y. Burak and T. Brookings, 'What grid cells convey about rat location,' *Journal of Neuroscience*, vol. 28, no. 27, pp. 6858–6871, 2008.

[73] E. I. Moser, E. Kropff and M.-B. Moser, 'Place cells, grid cells, and the brain's spatial representation system,' *Annu. Rev. Neurosci.*, vol. 31, pp. 69–89, 2008.

[74] T. Solstad, E. I. Moser and G. T. Einevoll, 'From grid cells to place cells: A mathematical model,' *Hippocampus*, vol. 16, no. 12, pp. 1026–1031, 2006.

[75] R. U. Muller and J. L. Kubie, 'The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells,' *Journal of Neuroscience*, vol. 7, no. 7, pp. 1951–1968, 1987.

[76] R. J. Gardner, L. Lu, T. Wernle, M.-B. Moser and E. I. Moser, 'Correlation structure of grid cells is preserved during sleep,' *Nature neuroscience*, vol. 22, no. 4, pp. 598–608, 2019.

[77] S. G. Trettel, J. B. Trimper, E. Hwaun, I. R. Fiete and L. L. Colgin, 'Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors,' *Nature neuroscience*, vol. 22, no. 4, pp. 609–617, 2019.

[78] Y. Burak and I. R. Fiete, 'Accurate path integration in continuous attractor network models of grid cells,' *PLoS Comput Biol*, vol. 5, no. 2, e1000291, 2009.

[79] J. J. Couey, A. Witoelar, S.-J. Zhang, K. Zheng, J. Ye, B. Dunn, R. Czajkowski, M.-B. Moser, E. I. Moser, Y. Roudi *et al.*, 'Recurrent inhibitory circuitry as a mechanism for grid formation,' *Nature neuroscience*, vol. 16, no. 3, pp. 318–324, 2013.

[80] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser and M.-B. Moser, 'Path integration and the neural basis of the'cognitive map',' *Nature Reviews Neuroscience*, vol. 7, no. 8, pp. 663–678, 2006.

[81] E. Kropff and A. Treves, 'The emergence of grid cells: Intelligent design or just adaptation?' *Hippocampus*, vol. 18, no. 12, pp. 1256–1269, 2008.

[82] C. Barry, L. L. Ginzberg, J. O'Keefe and N. Burgess, 'Grid cell firing patterns signal environmental novelty by expansion,' *Proceedings of the National Academy of Sciences*, vol. 109, no. 43, pp. 17 687–17 692, 2012.

[83] C. N. Boccara, M. Nardin, F. Stella, J. O'Neill and J. Csicsvari, 'The entorhinal cognitive map is attracted to goals,' *Science*, vol. 363, no. 6434, pp. 1443–1447, 2019.

[84] D. Derdikman, J. R. Whitlock, A. Tsao, M. Fyhn, T. Hafting, M.-B. Moser and E. I. Moser, 'Fragmentation of grid cell maps in a multicompartment environment,' *Nature neuroscience*, vol. 12, no. 10, pp. 1325–1332, 2009.

[85] J. Krupic, M. Bauza, S. Burton, C. Barry and J. O'Keefe, 'Grid cell symmetry is shaped by environmental geometry,' *Nature*, vol. 518, no. 7538, pp. 232–235, 2015.

[86] T. Stensola, H. Stensola, M.-B. Moser and E. I. Moser, 'Shearing-induced asymmetry in entorhinal grid cells,' *Nature*, vol. 518, no. 7538, pp. 207–212, 2015.

[87] Z. V. Guo, S. A. Hires, N. Li, D. H. O'Connor, T. Komiyama, E. Ophir, D. Huber, C. Bonardi, K. Morandell, D. Gutnisky *et al.*, 'Procedures for behavioral experiments in head-fixed mice,' *PloS one*, vol. 9, no. 2, e88678, 2014.

[88] G. Casali, S. Shipley, C. Dowell, R. Hayman and C. Barry, 'Entorhinal neurons exhibit cue locking in rodent vr,' *Frontiers in cellular neuroscience*, vol. 12, p. 512, 2019.

[89] M. G. Campbell, A. Attinger, S. A. Ocko, S. Ganguli and L. M. Giocomo, 'Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex,' *Cell reports*, vol. 36, no. 10, p. 109 669, 2021.

[90] Y. Solomon, A. Wagner and P. Bendich, 'A fast and robust method for global topological functional optimization,' in *International Conference on Artificial Intelligence and Statistics*, PMLR, 2021, pp. 109–117.

[91] A. Wagner, E. Solomon and P. Bendich, 'Improving metric dimensionality reduction with distributed topology,' *arXiv preprint arXiv:2106.07613*, 2021.

[92] R. Pirsig, 'On scientific method,' *Science and Its Ways of Knowing*, pp. 7–10, 1997.

[93] L. Wasserman, 'Topological data analysis,' *Annual Review of Statistics and Its Application*, vol. 5, pp. 501–532, 2018.

[94] C. Curto, J. Paik and I. Rivin, 'Betti curves of rank one symmetric matrices,' in *International Conference on Geometric Science of Information*, Springer, 2021, pp. 645–655.

[95] R. J. Adler and S. Agami, ‘Modelling persistence diagrams with planar point processes, and revealing topology with bagplots,’ *Journal of Applied and Computational Topology*, vol. 3, no. 3, pp. 139–183, 2019.

[96] P. Bubenik *et al.*, ‘Statistical topological data analysis using persistence landscapes.,’ *J. Mach. Learn. Res.*, vol. 16, no. 1, pp. 77–102, 2015.

[97] H. Adams, T. Emerson, M. Kirby, R. Neville, C. Peterson, P. Shipman, S. Chepushtanova, E. Hanson, F. Motta and L. Ziegelmeier, ‘Persistence images: A stable vector representation of persistent homology,’ *Journal of Machine Learning Research*, vol. 18, 2017.

[98] S. Vishwanath, B. K. Sriperumbudur, K. Fukumizu and S. Kuriki, ‘Robust topological inference in the presence of outliers,’ *arXiv preprint arXiv:2206.01795*, 2022.

[99] F. Jiang, T. Tsuji and T. Shirai, ‘Pore geometry characterization by persistent homology theory,’ *Water Resources Research*, vol. 54, no. 6, pp. 4150–4163, 2018.

[100] G. Spreemann, B. Dunn, M. B. Botnan and N. A. Baas, ‘Using persistent homology to reveal hidden covariates in systems governed by the kinetic ising model,’ *Physical Review E*, vol. 97, no. 3, p. 032 313, 2018.

[101] J. A. Perea, ‘Multiscale projective coordinates via persistent cohomology of sparse filtrations,’ *Discrete & Computational Geometry*, vol. 59, no. 1, pp. 175–225, 2018.

[102] L. Polanco and J. A. Perea, ‘Coordinatizing data with lens spaces and persistent cohomology,’ *arXiv preprint arXiv:1905.00350*, 2019.

[103] N. C. Schonsheck and S. C. Schonsheck, ‘Spherical coordinates from persistent cohomology,’ no. arXiv:2209.02791, 2022, arXiv:2209.02791 [cs, math]. [Online]. Available: `http://arxiv.org/abs/2209.02791`.

[104] S. Benas, X. Fernandez and E. Kropff, ‘Modeled grid cells aligned by a flexible attractor,’ *bioRxiv*, 2022.

[105] M. Contessoto, F. Mémoli, A. Stefanou and L. Zhou, ‘Persistent cup-length,’ *arXiv preprint arXiv:2107.01553*, 2021.

[106] A. J. Blumberg and M. A. Mandell, ‘Quantitative homotopy theory in topological data analysis,’ *Foundations of Computational Mathematics*, vol. 13, no. 6, pp. 885–911, 2013.

[107] J. F. Jardine, ‘Persistent homotopy theory,’ *arXiv preprint arXiv:2002.10013*, 2020.

[108] D. Cohen-Steiner, H. Edelsbrunner and J. Harer, ‘Extending persistence using poincaré and lefschetz duality,’ *Foundations of Computational Mathematics*, vol. 9, no. 1, pp. 79–103, 2009.

[109] F. Russold, ‘Persistent sheaf cohomology,’ *arXiv preprint arXiv:2204.13446*, 2022.

[110] J. M. Curry, *Sheaves, cosheaves and applications*. University of Pennsylvania, 2014.

[111] M. Vaupel, E. Hermansen and P. Trygsland, 'Reeb complexes and topological persistence,' *arXiv preprint arXiv:2201.12624*, 2022.

[112] G. Carlsson, V. De Silva and D. Morozov, 'Zigzag persistent homology and real-valued functions,' in *Proceedings of the twenty-fifth annual symposium on Computational geometry*, 2009, pp. 247–256.

[113] M. Vaupel, E. Hermansen and P. Trygsland, 'Section complexes of simplicial height functions,' *arXiv preprint arXiv:2201.12617*, 2022.

[114] P. Trygsland, 'Combinatorial models for topological reeb spaces,' *arXiv preprint arXiv:2109.05474*, 2021.

[115] K. Jensen, T.-C. Kao, M. Tripodi and G. Hennequin, 'Manifold gplvms for discovering non-euclidean latent structure in neural data,' *Advances in Neural Information Processing Systems*, vol. 33, pp. 22 580–22 592, 2020.

[116] M. Bjerke, L. Schott, K. T. Jensen, C. Battistin, D. A. Klindt and B. A. Dunn, 'Understanding neural coding on latent manifolds by sharing features and dividing ensembles,' *arXiv preprint arXiv:2210.03155*, 2022.

[117] G. Ginosar, J. Aljadeff, Y. Burak, H. Sompolinsky, L. Las and N. Ulanovsky, 'Locally ordered representation of 3d space in the entorhinal cortex,' *Nature*, vol. 596, no. 7872, pp. 404–409, 2021.

[118] D. Aronov, R. Nevers and D. W. Tank, 'Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit,' *Nature*, vol. 543, no. 7647, pp. 719–722, 2017.

[119] N. J. Killian, M. J. Jutras and E. A. Buffalo, 'A map of visual space in the primate entorhinal cortex,' *Nature*, vol. 491, no. 7426, pp. 761–764, 2012.

[120] A. Mathis, M. B. Stemmler and A. V. Herz, 'Probable nature of higher-dimensional symmetries underlying mammalian grid-cell activity patterns,' *Elife*, vol. 4, 2015.

[121] A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil *et al.*, 'Vector-based navigation using grid-like representations in artificial agents,' *Nature*, vol. 557, no. 7705, pp. 429–433, 2018.

[122] G. Montavon, W. Samek and K.-R. Müller, 'Methods for interpreting and understanding deep neural networks,' *Digital signal processing*, vol. 73, pp. 1–15, 2018.

[123] V. S. Schøyen, M. B. Pettersen, K. Holzhausen, A. Malthe-Sørensen and M. E. Lepperød, 'Navigating multiple environments with emergent grid cell remapping,' *bioRxiv*, 2022.

[124] M. Shapiro, 'Plasticity, hippocampal place cells, and cognitive maps,' *Archives of neurology*, vol. 58, no. 6, pp. 874–881, 2001.

[125] F. Sargolini, M. Fyhn, T. Hafting, B. L. McNaughton, M. P. Witter, M.-B. Moser and E. I. Moser, 'Conjunctive representation of position, direction, and velocity in entorhinal cortex,' *Science*, vol. 312, no. 5774, pp. 758–762, 2006.

[126] D. Ledergerber, C. Battistin, J. S. Blackstad, R. J. Gardner, M. P. Witter, M.-B. Moser, Y. Roudi and E. I. Moser, 'Task-dependent mixed selectivity in the subiculum,' *Cell reports*, vol. 35, no. 8, p. 109 175, 2021.

[127] A. Myers, S. Utpala, S. Talbar, S. Sanborn, C. Shewmake, C. Donnat, J. Mathe, U. Lupo, R. Sonthalia, X. Cui *et al.*, 'Iclr 2022 challenge for computational geometry and topology: Design and results,' *arXiv preprint arXiv:2206.09048*, 2022.

[128] N. A. Baas, 'On the concept of space in neuroscience,' *Current Opinion in Systems Biology*, vol. 1, pp. 32–37, 2017.

[129] A. C. Paulk, Y. Kfir, A. R. Khanna, M. L. Mustroph, E. M. Trautmann, D. J. Soper, S. D. Stavisky, M. Welkenhuysen, B. Dutta, K. V. Shenoy *et al.*, 'Large-scale neural recordings with single neuron resolution using neuropixels probes in human cortex,' *Nature Neuroscience*, vol. 25, no. 2, pp. 252–263, 2022.

[130] J. Wang, R. Khosrowabadi, K. K. Ng, Z. Hong, J. S. X. Chong, Y. Wang, C.-Y. Chen, S. Hilal, N. Venketasubramanian, T. Y. Wong *et al.*, 'Alterations in brain network topology and structural-functional connectome coupling relate to cognitive impairment,' *Frontiers in Aging Neuroscience*, vol. 10, p. 404, 2018.

[131] T. Wadhera, 'Brain network topology unraveling epilepsy and asd association: Automated eeg-based diagnostic model,' *Expert Systems with Applications*, vol. 186, p. 115 762, 2021.

[132] N. K. Kasabov, *Time-space, spiking neural networks and brain-inspired artificial intelligence*. Springer, 2019.

[133] S. P. Levine, J. E. Huggins, S. L. BeMent, R. K. Kushwaha, L. A. Schuh, M. M. Rohde, E. A. Passaro, D. A. Ross, K. V. Elisevich and B. J. Smith, 'A direct brain interface based on event-related potentials,' *IEEE Transactions on Rehabilitation Engineering*, vol. 8, no. 2, pp. 180–185, 2000.

[134] A. Farshchian, J. A. Gallego, J. P. Cohen, Y. Bengio, L. E. Miller and S. A. Solla, 'Adversarial domain adaptation for stable brain-machine interfaces,' *arXiv preprint arXiv:1810.00045*, 2018.

Paper I

Uncovering 2-D toroidal representations in grid cell ensemble activity during 1-D behavior

Erik Hermansen[1*], David A. Klindt[1] and Benjamin A. Dunn[1*]

[1]Department of Mathematical Sciences, NTNU, Trondheim, Norway.

*Corresponding author(s). E-mail(s): erik.hermansen@ntnu.no; benjamin.dunn@ntnu.no;
Contributing authors: klindt.david@gmail.com;

Abstract

Neuroscience is pushing toward studying the brain during naturalistic behaviors with open-ended tasks. Grid cells are a classic example, where free behavior was key to observing their characteristic spatial representations in two-dimensional environments [1]. In contrast, it has been difficult to identify grid cells and study their computations in more restrictive experiments, such as head-fixed wheel running [2–6]. Here, we challenge this view by showing that shifting the focus from single neurons to the population level changes the minimal experimental complexity required to study grid cell representations. Specifically, we combine the manifold approximation in UMAP [7] with persistent homology [8] to study the topology of the population activity. With these methods, we show that the population activity of grid cells covers a similar two-dimensional toroidal state space during wheel running as in open field foraging [9, 10], with and without a virtual reality setup. Trajectories on the torus correspond to single trial runs in virtual reality and changes in experimental conditions are reflected in the internal representation, while the toroidal representation undergoes occasional shifts in its alignment to the environment. These findings show that our method can uncover latent topologies that go beyond the complexity of the task, allowing us to investigate internal dynamics in simple experimental settings in which the analysis of grid cells has so far remained elusive.

Keywords: grid cells, topological data analysis, neural population dynamics

Main

Introduction. Recent neuroscience research has pushed for the study of naturalistic open-ended behaviors [11], suggesting the experimental complexity must be sufficiently high to make novel insights into neural computations [12]. The question, however, is what the minimal complexity is in order to study the function and ethological relevance of specific neurons [13, 14]. Grid cells are a canonical example from the brain's navigational system [15] where the hexagonal grid patterns in their responses relative to space was first observed when the animals were allowed to freely move in large open field (OF) arenas [1, 16]. In contrast, the spatial responses of grid cells in one-dimensional (1-D) environments and head-fixed recordings are not yet well understood [17–19]: In 1-D linear tracks, these are speculated to be cross-sections of the 2-D patterns [6, 20], although recent results suggest that grid cells in 1-D tracks are tuned to integrated distance [21, 22]. Casali et al. even argued for 'the necessity of identifying grid cells from real open field environments' [2], thus implying that 2-D OF arenas are the minimal experimental complexity required to fully understand grid cell computations.

However, the brain performs computations not merely at the single-neuron level but also in populations of neurons (or *neural ensembles*) [reviewed in 23–25]. The internal representations encoded by neural ensembles often carry a distinct topology [26–29] which can be studied without reference to external covariates [30, 31]. The concept of understanding this topological structure is beginning to change the way we think about certain classes of neurons [32, 33], such as head direction cells where the 1-D ring structure of the network represents the encoding of an internal head direction system [34, 35]. In the case of grid cells, their population representation has been found to be *toroidal*, describing the two periodicities of the grid cell pattern in 2-D environments [10]. Therefore, in this study we asked whether focusing on the topology of grid cell representations would allow us to study their computations in experiments of lower complexity where it has so far been impossible to give a convincing classification of grid cell populations [2].

Population analysis. A common pipeline in population analysis starts by reducing the dimensionality of the neural activity, e.g., using uniform manifold approximation and projection (UMAP) [7], t-distributed stochastic neighbor embedding (t-SNE) [36] or (deep) latent variable models [37–39], with the goal of extracting a low-dimensional representation of the population activity at each time point. However, dimensionality-reduction requires specifying a range of parameters [40, 41], and in particular, choosing a dimensionality for the embedding space, often chosen as 2- or 3-D for interpretability [42, 43] which allows visualization but potentially discards crucial information [44]. Topological data analysis offers an alternative approach by characterizing high-dimensional point clouds through *persistent homology* (PH) [8, 45]. Specifically, we can compute a topological invariant called a barcode quantifying the presence of holes in representations of the data. This characterization helps to discern underlying topological structure, such as circular features [46, 47]. However, the method is sensitive to outliers (Extended Data Fig. 5b) and has a computational bottleneck, with polynomial growth of run-time and memory requirements in the number of points in the point cloud [48]. Therefore, preprocessing (noise removal and downsampling) of the data is necessary to uncover shape information from the barcodes [10, 34, 35]. Here, we propose uniform manifold approximation and persistent homology (UMAPH), which resolves these issues by applying persistent homology to the topological representation constructed in the first step of UMAP, thus eliminating the potentially problematic projection step of UMAP [44], while providing a topology-preserving [7] noise model for PH (Fig. 1a and Extended Data Fig. 5a, b).

The UMAPH algorithm. We first compute geodesic distances of the data based on the assumption that a dataset is uniformly sampled from an underlying manifold (e.g., the grid cell torus), but is presented as a non-uniform distribution embedded into an ambient space with observational noise (e.g., neural population activity). This motivates obtaining the intrinsic distances by enforcing the volume of a ball containing a fixed number of points to be constant independent of which point in the dataset it circumscribes [7]. However, this creates local irregularities – the distance between two neighboring points will depend on the local neighborhood of each – so these are averaged to be symmetric (Fig. 1a.ii). UMAP then proceeds with an approximation of the same distances on a lower-dimensional point cloud. In UMAPH, we instead, apply PH to the high-dimensional data using the geodesic distances to compute the barcode of the manifold. PH first constructs a combinatorial description of the proximity between points. By forming balls of a common radius around each point, it lists the *sets of balls* with pairwise intersections. Increasing the radius forms a sequence of such descriptions, called a *filtration* (Fig. 1a.iii). The n-dimensional homology captures how unions of balls form n-dimensional holes. Thus, applying homology to the filtration describes the evolution of holes for increasing volumes, and the collection of radial intervals (or *bars*) when different holes appear and disappear is called a *barcode* (Fig. 1a.v). Optionally, in a last step, we perform cohomological coordinatization [49], to derive a low-dimensional description of the data (*decoding* the time-varying internal state [34], Fig. 1a.vi). Based on the selected 1-D holes in the barcodes, we compute circle-valued maps, assigning to each population vector angular coordinates corresponding to the circular features represented by the bars.

In summary, the intuition behind UMAPH is that we apply homology to a filtration constructed by forming balls containing the same number of points, but of potentially different sizes in the ambient space (Extended Data Fig. 5). This is advantageous because the metric is then locally adaptive to data density. Comparing the results when applied to head direction data previously studied with PH [30, 34, 35], gives an example of the benefit of this construction (Extended Data Fig. 6). With UMAPH, we found the resulting barcodes to be more interpretable with clear head direction rings maintained across states, even during slow-wave sleep where this structure was previously not recovered.

Toroidal tuning of mouse entorhinal cells in open field calcium recordings. Zong et al. (2022, [50]) and Obenhaus et al. (2022, [51]) performed two-photon calcium imaging of the medial entorhinal cortex (MEC) in freely moving and head-fixed mice (Fig. 1b, c). In two mice (in 4/6 and 2/6 recording days), applying UMAPH to the activity of all recorded cells during OF firing resulted in barcodes with distinct toroidal signatures (one clear H^0-bar, two H^1-bars and one H^2-bar, all exceeding corresponding longest bars in 100 random perturbations of the data, Fig. 1d and Extended Data Fig. 1a). In a third mouse, we clustered the OF autocorrelograms to obtain a single grid cell ensemble (or *module* [9]), exhibiting similar toroidal expression in the barcode (in 4/7 OF recordings studied), also when an object was inserted (1 out of 3 object session recordings) (Fig. 1d and Extended Data Fig. 1a). The pair of circular coordinates attained from decoding of the two longest circular features in the barcode revealed stripe-like patterns in the OF corresponding to the two periodicities of the grid pattern (Fig. 1d and Extended Data Fig. 1a). Moreover, the internal dynamics displayed a clear relation to spatial movement (Video 1). We next studied each cell's tuning to the torus by fitting a generalized linear model (GLM) to the calcium activity, comparing the performance when using either the toroidal positions or the recorded positions as covariates. The activity was better modeled by the torus than physical space ($P \in [1.4 \cdot 10^{-34}, 0.0076]$, $W \in [7694, 168511]$, Wilcoxon signed-rank test), and we observed that cells with the best fit had a clear grid pattern and single tuning fields on the torus (Fig. 1e, f and Extended Data Fig. 1e and 7).

Linking toroidal grid cell tuning between open field and head-fixed treadmill recordings. We next performed the same analysis of grid cell module activity while the animal was running on a wheel. A 2-D toroidal topology as in the OF sessions was seen (in 3 out of 4 wheel recordings), despite the 1-D nature of the task (Fig. 1g and Extended Data Fig. 1a). Decoding the positions on the torus revealed mostly unidirectional trajectories in line with the spatial behavior (Fig. 1g and Video 2). This shows that internal position integrates over the animal's locomotion, even though the animal is not moving in space, supporting the idea of path integration. Each neuron had a single toroidal tuning field in the same location on the torus in OF foraging and wheel running (toroidal rate map correlations higher than shuffled comparisons, $n = 1000$ shuffles, $P < 0.001$, Fig. 1h, i and Extended Data Fig. 1c).

Stable toroidal tuning of mouse grid cells in virtual reality with Neuropixels recordings. Campbell et al. (2021, [22]) used Neuropixels probes to record extracellular spikes in the MEC of head-fixed mice engaged in a virtual-reality task on a running wheel under different conditions: baseline, dark, gain (wheel slowed down with respect to virtual reality (VR)) and contrast (reduced visual contrast) sessions (Fig. 2a) [22]. We first studied one exemplary recording day (mouse I1 day 0417) with a good experimental yield (having 75 'distance'/'putative grid' cells out of 225 cells recorded in the MEC using the classification given by Campbell et al.), before repeating our analysis for the remaining available data. Neuron ensembles were first clustered by their cross-correlation (maximized over < 0.9s time lags, Fig. 2b), thus avoiding any prior assumption about the spatial tuning of the studied neurons. We found two prominent clusters (henceforth termed 'M1' and 'M2', $n = 44$ and 41 neurons). Applying UMAPH to the firing rates of these clusters revealed barcodes with four long-lived bars (one H^0-, two H^1- and one H^2-bar) indicating toroidal structure ($P < 0.01$ in all conditions, except M2 baseline: $P = 0.06$, Fig. 2c). Decoding revealed primarily unidirectional trajectories on the toroidal manifold (Fig. 2d). The activity of most cells in both M1 and M2 was confined to a specific location on the toroidal surface (Fig. 2e) and each neuron's preferred location on the torus was preserved across the experimental conditions ($P < 0.001$ for toroidal rate map correlations compared to $n = 1000$ shuffles, Fig. 2e and Extended Data Fig. 8), consistent with grid cell populations in rats [10]. Lastly, to allow for unbiased comparison with the VR spatial track, we used the toroidal description found during the dark session to decode baseline trials and found the toroidal coordinates to have higher explained deviance than the VR positions during baseline trials (M1: $P = 3.8 \cdot 10^{-9}, W = 990$; M2: $P = 1.2 \cdot 10^{-8}, W = 861$, Wilcoxon signed-rank test; Fig. 2h). Taken together, our findings suggest that the neurons analyzed, a majority of which were classified as distance cells, were in fact, grid cells.

We then clustered the remaining (88) MEC recording days and classified ensembles as grid cell modules if the cells' toroidal tuning matched an idealized point source distribution on a hexagonal torus [52] (Extended Data Fig. 3b), as expected for grid cells [53] (Extended Data Fig. 3a)). We observed similar toroidal (grid cell) characteristics as for M1 and M2 in 6 more ensembles (Fig. 2h and Extended Data Fig. 2, 3b, c and 9 – 11).

Alignments of the toroidal representation in VR are unstable but follow (gain/contrast) manipulations. Usually, decoding of the internal state space dynamics is linked through the tuning to an external, known covariate (e.g., spatial position, [22, 54]). By contrast, our method performs unsupervised inference of state space dynamics derived directly from the population activity, allowing unbiased comparison with external covariates. Coloring the 1-D VR trajectories with the 2-D toroidal coordinates, we observed a clear, time dependent relation between spatial and toroidal coordinates (Fig. 2j, Extended Data Fig. 4 and Video 3). At times, the movement on the torus was aligned with the spatial movement (visible in the figures as segments with a stable horizontal pattern), but would occasionally shift or drift, supporting findings by Low et al. [19]. For ensembles M1 and M2, this seemed to happen in coordination, in line with recent work [54]. However, gain and contrast manipulations clearly elicited a shift in the alignment between the torus and virtual space. When the contrast changed from high to low, an aligned representation quickly disappeared until the contrast was reset, upon which the mapping returned to the former representation. During strong gain manipulation (gain < 0.8 and contrast either 10 or 100%), the toroidal path length was longer than baseline sessions ($P < 0.04, Z > 4$ for six ensembles; Fig. 2i). This conjunctive influence of self-motion and visual cues is in line with Campbell et al. (2021).

Discussion. By combining UMAP and PH, we obtained a powerful and noise-robust method that revealed toroidal topology of grid cell population activity in mice without reference to an external covariate, both in calcium imaging and with electrophysiology. Remarkably, this 2-D structure was found during head-fixation and 1-D wheel running, where it had so far not been observed... We observed a clear relation of the internal dynamics to movement, suggestive of path integration. While the toroidal structure was preserved across sessions, the alignment to the VR-track was affected by gain and contrast manipulations, and underwent shifts also during baseline trials.

Our results provide a proof-of-principle, demonstrating how a topological perspective on population coding allows us to study neural computations that go beyond the dimensionality of the task. This contributes to debates about whether simple, artificial stimuli are sufficient [e.g., in visual neuroscience 13, 14, 55] Thus, these results have the potential to shift the way we think about neural population coding, unlocking exciting possibilities for future studies of various neural systems using ideas from topology [32, 56]. As we have seen, the internal representation of grid cells can look distorted when mapped onto space, making spatial responses difficult to study. We propose grid cells and other cell types (e.g., head direction cells, see Extended Data Fig. 6) are best understood through their activity state space and not their external modulation.

By using an unsupervised state space discovery approach, we can study and compare population dynamics, not easily available at the single-cell level. This will be needed to understand recordings where the relevant covariates are of lower dimension than the neural representation, as in this study. Moreover, our approach is also useful when covariates are not known, for instance, in *cognitive spaces* [56], or where we have reason to assume higher-dimensional features in neural data (e.g., three- or higher-dimensional tori expected for conjunctive grid cells and combinations of modules [57, 58], Fig. 5c).

We look forward to future studies that embrace the toroidal (rather than 2-D spatial) nature of grid cells, which can be studied with a vast array of experimental methods in head-fixed 1-D and virtual environments [59, 60]. These simple task settings come with a number of benefits such as high-throughput experiments with large or movement-sensitive equipment, detailed animal tracking and permit trial-based, stereotyped analyses [11, 14]. We believe these findings open the path to new insights beyond what has been expected from such artificial settings [3, 4, 14], also in other brain regions and cognitive tasks [61, 62].

Acknowledgments. We thank Zong et al., Obenhaus et al., Campbell et al. and Peyrache et al. for sharing data publicly and Per Kristian Hove for setting up databases. We furthermore thank the Department of Mathematical

Sciences (NTNU) and the computing resources provided by IDUN. This work was supported by a grant by the Research Council of Norway (iMOD, NFR grant #325114).

Author contributions. E.H., D.A.K. and B.A.D. conceptualized and proposed analyses. E.H. developed and performed the analyses. E.H., D.A.K. and B.A.D. interpreted data and results. All authors wrote the paper. B.A.D supervised the project and obtained funding.

Supplementary Information. Not available for this paper.

Author Contact Information. Correspondence should be addressed to E.H. or B.A.D.

Competing interests statement. The authors declare they have no competing interests.

Methods

Preprocessing of entorhinal recordings

Calcium events found in the 'NAT.mat'-files of mouse 97045 days 20210305, 20210307, 20210308, 20210317 (renamed #1-4) and mouse 97046 days 20210308 (#1) and 20210312 (#2) were accessed from [63]. Duplicate cells listed in variable 'RepeatCell' of the corresponding 'NeuronInformation.mat'-file were removed. As the neurons were recorded in two separate planes with a temporal offset, the activity was interpolated at similar frames using the package *'scipy.stats.interp1d'*. Activity values less than 10^{-10} were subsequently set to 0 and neurons with average activity above 10 were excluded in the analysis. This resulted in 1-hour recordings of $n = 228, 350, 293$ and 428 (#1-4) neurons for mouse 97045 and 50- and 70-min. recordings of $n = 635$ and 637 (#1 and #2) neurons for mouse 97046. The data were speed-filtered at minimum 5 cm s^{-1}, so to remove moments of inactivity. Population vectors with no activity were excluded.

PCA whitening was next applied to the data, to denoise and standardize the n-dimensional data (n being the number of neurons), projecting the z-scored population vectors to its $d = 6$ first principal components and dividing by the square root of the eigenvalues (see [64] for a thorough discussion of PCA as a preprocessing tool).

Due to computational complexity of computing barcodes, the size of the point cloud was reduced. First, a *radial* downsampling scheme was performed (as in [34]). The point with maximum absolute summed value was chosen as initial landmark point, and the Euclidean distance to the rest of the point cloud computed. Points closer than $\epsilon = 0.5$ were discarded and the remaining closest point to all sampled landmark(s) (defined as the maximum distance to all landmarks) was then picked. This process was iterated until exhaustion, vastly reducing the size of the point cloud (leaving approx. $30 - 50\%$ of the points). While this method preserves the spread of the data, it is prone to keeping outliers. Thus, a second, density-based downsampling method was used (setting $\kappa = 1000$, see ‘Fuzzy Downsampling’), keeping $m = 1000$ points. Finally, UMAPH was applied to the reduced point cloud, using cosine metric, $k = 1000$ and $\mathbb{Z}_{47}$-coefficients (47 is chosen to not likely divide the torsion subgroup [49]).

The calcium events of Mouse 88529 were found in the ‘__filtered_spikes’ database table derived from the MySQL-dump file ‘dump.sql’ (available in [65]). The following sessions were used (and renamed): ‘419c1c6b319d0ddf’ (#1 W), ‘5b92b96313c3fc19’ (#1 OF), ‘d5a06b6a7630bb11’ (#2 W), ‘7e888f1d8eaab46b’ (#2 OF), ‘26fd0fbe1e205255’ (#3 W), ‘1f20835f09e28706’ (#3 OF), ‘59825ec5641c94b4’ (’#4 OF’), ‘c43d9bd004db772b’ (#4 Obj1),’9190f2fccd52497e’ (#4 Obj2). The ‘filtered_cells’-table was used to get cell IDs with a signal-to-noise ratio above 3.5 (see Supplementary Information in [51]). Grid modules were extracted by using agglomerative clustering, with average linkage, based on the ‘Manhattan’ distance between spatial autocorrelograms (in each recording day) found in the ‘__grid_score’-table (signal type ‘spikes’ and parameter set ‘B’). The largest cluster found, when setting a distance threshold of 2000 r, were analyzed, giving $n = 126, 160, 113$ and 162 neurons (recording days $1 - 4$) with sessions lasting 18-24 minutes. The events were temporally convolved with a Gaussian kernel of $\sigma = 0.27$ ms and square rooted, before speed-filtering the time frames using a threshold of 10 cm s^{-1}. The same pipeline as described above was subsequently applied, with $\epsilon = 0.7$ in recording days $1 - 3$ and $\epsilon = 0.55$ in #4.

Neuropixels recordings in the MEC were performed by Campbell et al. of 20 mice in a total of 90 sessions and retrieved from [66]. Each spike time was replaced by a delta function (valued 1 at the time of firing; 0 otherwise) and temporally convolved with a Gaussian kernel with $\sigma = 60$ ms, before summing over all spike times (for each neuron), giving a continuous firing rate function. Firing rates were sampled every 10 ms and square rooted, before speed-filtering at minimum 10 cm s^{-1}. Neurons not determined as ‘good’ (based on *“contamination, signal to noise ratio and firing rate”*, see Method details in [22]) in one recording or neurons with mean firing rates below 0.05 Hz or above 10 Hz were excluded. The remaining neurons were then clustered (see ‘Unsupervised clustering’), giving 65 clusters and a similar pipeline as described above was used, with $d = 7$, $\epsilon = 0.9$, $m = 800$ and $\kappa = 800$ for all ensembles.

UMAPH pipeline

Given a dataset $X \subseteq \mathbb{R}^N$, UMAPH first computes geodesic distances for a uniform manifold approximation of the data from which topological information is then studied using persistent homology. This approach was first used to find topological structure in grid cell population activity [10] and has similarities with concurrent work in [67], computing persistent homology based on geodesic distances.

Let X be a dataset and $d : X \times X \to \mathbb{R}$ a dissimilarity measure of the points inherited from the ambient space.

1. Choose a number k and define for each $x \in X$ its neighborhood set $N_x = \{x_1, x_2, \ldots, x_k\}$, such that for any $y \in X \setminus N_x$, $d(x,y) \geq d(x,x_k)$ and ordered by $d(x,x_1) \leq d(x,x_2) \leq \ldots \leq d(x,x_k)$. The extended-pseudo-metric, d_x, is defined as

$$d_x(y,z) := \begin{cases} d(y,z) & \text{if } y = x \text{ and } z \in N_x \sqcup \{x\} \text{ or } z = x \text{ and } y \in N_x \sqcup \{x\} \\ \infty & \text{otherwise.} \end{cases}$$

2. Next, for each x, define the (membership) function (or *fuzzy set*) $\mu_x : X \to I$, for $I = [0,1]$ by

$$\mu_x(y) = \exp\left(-\frac{d_x(x,y)}{\sigma_x}\right),$$

   where σ_x is found by letting $\sum_{y \in N_x} \mu_x(y) = \log_2(k)$. The value $\mu_x(y)$ may be seen as the probability of the point y being a member of the neigbourhood of x.
3. Construct a global membership function $\mu : X \times X \to I$:

$$\mu(x,y) = \mu_x(y) + \mu_y(x) - \mu_x(y)\mu_y(x)$$

   This operation is the *fuzzy* set union and defines the fuzzy topological representation of the data. This choice of union operation reflects the probabilistic sum, which describes the probability of the union of two independent events, and assigns probabilities to the edges of the neighbourhood graph of X. Other choices of fuzzy union operators are possible, and in [68] it is suggested that the fuzzy *intersection*, $\mu(x,y) = \mu_x(y)\mu_y(x)$, should be used for clustering. The standard fuzzy union is given by $\mu(x,y) = \max(\mu_x(y), \mu_y(x))$ [69] and is equivalent to taking the shortest path distance in merging the extended pseudo-metrics. This is used in [70] to describe the representation (or 'UMAP complex') as an iterated pushout of Vietoris-Rips systems.
4. Translate μ to a distance measure, $d_X : X \times X \to \mathbb{R}$,

$$d_X(x,y) = -\log(\mu(x,y)).$$

   This converts the fuzzy topological representation back to an extended-pseudo-metric space representation of the dataset and d_X reflects the negative log likelihood of the simplices of the UMAP complex.
5. Finally, use d_X to construct a Rips filtration (see 'Persistent Homology') and take its homology. This may be seen as applying persistent homology to the filtration of the UMAP complex given by decreasing probability, giving a visualizable description of the fuzzy topological representation of the dataset X (i.e., the barcode).
6. *(Optional)* If the barcode indicates circular features in the data (captured by the one-dimensional bars), circular coordinates representing

these may be computed through *co*homology. This requires defining the Rips complex for which the coordinates are computed, i.e., a choice of distance τ for which the circular feature exists, usually set to a value close to the end of the bar [71].

Fuzzy Downsampling

To reduce computational complexity and remove outliers in the dataset, the point cloud was downsampled before applying persistent homology. Given a point cloud X and neighborhood sets N_x, one for each $x \in X$ (containing the κ closest neighbors), the global membership function $\mu(x, y)$ was computed as described in 'UMAPH Pipeline'. Initializing X_0 as an empty set, a subsample $X_N \subset X$ was given recursively as follows.

For each iteration $n > 0$, define a function F_n as the summed membership strength for each point in the residual point cloud $x \in X'_n = X \setminus X_n$, given by

$$F_n(x) = \Sigma_{y \in X'_n} \mu(x, y) = \Sigma_{y \in X} \mu(x, y) - \Sigma_{y \in X_n} \mu(x, y). \tag{1}$$

The $(n+1)$-th point is then given by:

$$x_{n+1} = \max_x F_n(x).$$

In other words, the method iteratively keeps the point with the highest probability of being in the neighborhoods of all other points.

Defining $\mu(x, y) = \exp(-\frac{d(x,y)}{\sigma_y})$, there is a close relation to mean-shift clustering [72] and the objective function used in the topological denoising technique introduced by Kloke and Carlsson [73]. The latter uses this function to translate a subsample of points to topologically relevant positions, describing it as a weighted difference of two Gaussian kernel density estimators, one for the dataset X, serving to push the subsample, X_N, towards the densest regions of X and one for the subsample itself, repelling the points away from each other. Similarly, the fuzzy downsampling scheme picks points from the densest regions of X, but steers away from the regions already chosen.

Persistent Homology

The shape of the neural data was characterized using persistent *co*homology. Persistent cohomology results in the same barcodes as persistent homology (which is described below), but cohomology was necessary for decoding [49].

The homology of a topological space, T, is a sequence of vector spaces $H_n(T)$, for all natural numbers $n \in \mathbb{N}$ and the rank of $H_n(T)$ represents the number of n-dimensional holes [74]. A zero-dimensional hole describes a connected component, a 1-dimensional hole a circle, a 2-dimensional hole a void and so on for higher dimensions. The homology of a point cloud, X, only returns a count of the points. Thus, to elicit non-trivial homology reflecting the underlying space the dataset is sampled from, combinatorial spaces called

Rips complexes, $T_\tau(X)$, are associated to the data. The Rips complexes depend on a scale τ, commonly describing a dissimilarity relation between the points in the point cloud. Varying τ gives rise to an ordered sequence of complexes known as the *Rips filtration*:

$$T_{\tau_0} \subset T_{\tau_1} \subset \ldots \subset T_{\tau_n},$$

where $\tau_0 < \tau_1 < \ldots < \tau_n$. Applying homology to the Rips filtration gives a sequence of vector spaces and maps induced by inclusion in each dimension, n:

$$H_n(T_{\tau_0}(X)) \to H_n(T_{\tau_1}(X)) \to \ldots \to H_n(T_{\tau_n}(X)).$$

The totality of sequences is called persistent homology, $PH_*(X)$, and may be decomposed to a sum of *elementary persistence intervals*, $I([b_i, d_i))$:

$$PH_n(X) \cong \bigoplus_i I([b_i, d_i)).$$

Here, $b_i < d_i$ gives the scales for which an n-dimensional element hole in $PH_n(X)$ first appears and later disappears. Persistent homology may thus be represented as bars starting at b_i and ending at d_i. The collection of such bars (across all dimensions) is called the *barcode*.

To obtain a shuffled distribution of the barcodes, the activity of each neuron was independently rolled a random amount of time. The same pipeline was then applied to the shuffled population activity to give a barcode, and the process was repeated 100 times with different seeds.

We used `Ripser` [75, 76] for all computations of persistent cohomology.

Cohomological coordinatization

Circular coordinatization, as introduced by De Silva et al. [49], was performed to allow studying the internal dynamics of the population activity, This has previously been used to study head direction and grid cell activity [10, 34] and is motivated by a theoretical correspondence between 1-D cohomology and circle-valued maps of a topological space. By computing maps associated with the two longest-lived H^1-bars in the barcode and the Rips complexes at $\tau = b + 0.99 \cdot (d - b)$, where b and d correspond to the birth and death of the chosen bars, 2-D toroidal coordinates were computed for all vertices in the Rips complex. Note, for the head direction cell data only the single longest-lived bar was chosen.

The vertices correspond to the m points in the downsampled point cloud, thus, only m toroidal coordinates are obtained. To extrapolate these to the rest of the original point cloud or to a different session, the coordinates were first weighted by the values of the corresponding points, giving a distribution on the torus for each dimension. The toroidal coordinates were then computed, for each time step, by weighing the distribution by the corresponding value of each point in the full point cloud and finding the mass centers of the summed weighted distributions .

In visualizing the decoded toroidal positions as a function of VR-track (Fig. 2j and Extended Data Fig. 4), the coordinates were smoothed temporally with a Gaussian filter of $\sigma = 0.5$ s.

Rate maps and autocorrelograms

Spatial positions in the open field were binned into a 30^2 square grid, generating spatial rate maps. The mean neural activity in each bin was computed and spatially convolved with a Gaussian filter of width 2 bins (during which non-visited bins were assigned the mean value of the visited bins).

1D spatial autocorrelograms were computed for the linearized positions on the running wheel and the VR track, using 1 and 4 cm spatial bins, respectively. Autocorrelograms for each neuron were computed by finding the mean activity in each spatial bin and taking the dot product between this and a zero-padded copy of it, iteratively shifting the latter up to 300 bins.

Toroidal firing rate maps were calculated in the same way as the OF arena, first binning the toroidal surface into a square grid of $12^\circ \times 12^\circ$ bins and computing the average activity in each position bin. To spatially smooth the toroidal rate map, the 60° angle of the toroidal axes was addressed by first shifting the bins horizontally by a length equal to half the bins' vertical position. To address boundary conditions, nine copies of the shifted rate map were then tiled and spatially smoothed as the OF rate maps. The middle tile was then extracted and shifted back. For visualizations, the resulting rate maps were given 15° shear angles both horizontally and vertically.

Unsupervised Clustering

The electrophysiological data were clustered into groups of neurons through the time-lag cross-correlation values between pairs of neurons. This was computed, as in [31], for the entire population recording:

$$c_{ij}(t') = \int_0^T s_i(t)s_j(t+t')dt,$$

where $s_i(t)$ is the firing rate of neuron i at time t, converted from spike times as described in 'Preprocessing of entorhinal recordings', using a Gaussian kernel of $\sigma = 0.3$ s and sampled every 30 ms. T denotes the total duration of the recording. Setting $\tau_{\max} = 0.9$ s, the inverse, normalized cross-correlation was then given as:

$$C_{ij} = \frac{\min_\tau [c_{ij}(\tau), c_{ji}(\tau)]_0^{\tau_{\max}}}{\max_\tau [c_{ij}(\tau), c_{ji}(\tau)]_0^{\tau_{\max}}}.$$

To cluster the neurons, we performed agglomerative clustering with average linkage on the squared C_{ij}'s, averaged over all recordings of the same pair of neurons, using 0.6 r as distance threshold. Ensembles containing fewer than

20 neurons were disregarded, having too few neurons to confidently interpret the toroidal structure (see Extended Data Fig. 4e in [10]).

As described in the text, this was first performed for mouse I1, giving two ensembles, 'M1' and 'M2' (Fig. 6b). To see if these ensembles shared behavioral features, 1-D spatial autocorrelograms were computed. The mean autocorrelogram across the population showed repeating fields of activity along the track (Extended Data Fig. 3c). In both M1 and M2, the mean correlation between cells within the respective ensembles was higher than across the remaining MEC population in all three conditions (M1: 0.87 (B), 0.78 (D), 0.87 (G/C) v. 0.67, 0.58, 0.78; M2: 0.83, 0.82, 0.87 v. 0.71, 0.61, 0.77). Moreover, using the 'distance cells' classification in [22], 23 out of 44 (M1) and 27 out of 41 (M2) neurons were classified as 'distance cells'.

Hexagonal torus detection

To determine whether the decoded toroidal coordinates suggested a hexagonal torus in the VR-sessions, firing rates and rate maps were modeled for each neuron based on the analytical heat distribution on both a hexagonal and a square torus (Extended Data Fig. 3). The heat kernel on the hexagonal torus with point source at the origin is given as:

$$H_{hex}(x, y;\ t) = \frac{1}{t} \sum_{(k,l)\in\mathbb{Z}^2} \exp^{-\pi \frac{1}{t} \frac{2}{\sqrt{3}}((k+x)^2+(k+x)(l+y)+(l+y)^2)},$$

describing the temperature for (normalized) toroidal positions $(x, y) \in [0, 1]^2$ after time t, while

$$H_{sqr}(x, y;\ t) = \frac{1}{t} \sum_{(k,l)\in\mathbb{Z}^2} \exp^{-\pi \frac{1}{t}((k+x)^2+(l+y)^2)},$$

is the heat kernel on the square torus [52]. The fit of the original toroidal rate maps (generated by computing the mean activity in 12^2 bins of the m toroidal coordinates found prior to extrapolation, see 'Cohomological coordinatization') for the VR data to a toroidal point source heat distribution was tested as follows. First, m toroidal coordinates were sampled with even spacing. The origin of the sampled torus was shifted to each cell's peak activity on the torus (see 'Comparison of toroidal tuning'), and using $t = 0.1$, $k, l \in \{-1, 0, 1\}$ in the above equations, allowed computing temperature estimates for each toroidal position. The heat distribution on the sampled torus was defined as the mean temperature in 12^2 square bins. The linear correlation between the heat distribution and the firing rate maps asserted the hypothesized hexagonal toroidal tuning of each neuron. A hexagonal torus permits three periodic axes, and it was not known a priori pair of axes the decoded torus potentially described. However, this is reflected in the left vs right 45° angular shift of the distribution, so the reversed orientation was also tested by reversing one of the sampled coordinates. The maximum correlation between the modeled rate maps and the original one was used in the assessment and compared with a square torus rate map similarly modeled.

This analyses was first tested as a post hoc analysis of ensembles M1 and M2, showing that a majority of cells had rate maps more similar a distribution on the hexagonal torus than the square torus (M1: 38, 36, 36 out of 44 neurons in B, D, G/C and M2: 29, 35, 29 out of 41 neurons) and the median correlation was greater than by chance (M1 B, D, G/C: 0.81, 0.80, 0.75 r vs 0.08 – 0.16 r for 100 shuffles and M2: 0.71, 0.77, 0.68 r vs 0.02 – 0.11 r). The lowest median value for M1 and M2 was then used as a lower threshold to determine if the remaining ensembles had hexagonal toroidal structure, resulting in 16 such cases (6 mice and 8 ensembles, including M1 and M2, with $n = 25, 37, 40, 41, 43, 44, 53$ and 65 neurons for mouse and recording day: J4 0514, I5 0415, I1 0418, I1 0417 M2, J5 0506, I1 0417 M1, H3 0402, I3 0420. Recordings lasted between 9 and 159 minutes). Note, this heuristic eschews the computational problem of doing statistics on barcodes [77].

Replacing the sampled toroidal coordinates with the decoded toroidal positions in computing the heat kernel, H_{hex}, gave time-varying heat models for the firing rate of each neuron. This allowed computing spatial autocorrelograms for idealized hexagonal toroidal tuning, visually matching the original firing rate counterpart (Extended Data Fig. 3c).

Toroidal linearization and path length

The decoded coordinates were unwrapped onto a flat space to simplify the analysis of toroidal trajectories. This was done by iteratively assessing whether the next toroidal coordinate crossed either of the two circular origins (one for each dimension). All combinations (of 0 or 1 crossings per circle) were tested, and the next point was chosen to be the point closest to the previous one, as measured by the Euclidean metric.

To assess the influence of gain manipulation on the internal representation, the lengths of the linearized toroidal trajectories for each trial were estimated. First, the positions were fitted using linear regression for each axis and only trajectories with a fit of $r > 0.5$ in both axes were included. The length of each trial was then assessed as the Euclidean distance between the start and end point of the 2-D linear fit.

Toroidal alignment

As the decoded origins and orientations of the toroidal descriptions are arbitrary, the decoded toroidal coordinates were first pairwise reoriented (without reference to space), in order to compare these across modules and sessions for the VR recordings. Moreover, it was necessary to account for the hexagonal torus allowing for three axes (each axis being a linear combination of the two other), with the decoded pair of axes oriented at either 60° or 120° relative to each other.

For each comparison, one set of toroidal coordinates was held fixed, with the goal of obtaining the same orientation and axes for a second set (i.e., session). First, in assessing the hexagonality of the decoded rate maps (see 'Hexagonal Torus detection'), the angular orientation of the median rate map was determined for both sets of coordinates, limiting the possible axes-combinations (depending on the direction and if these were similar between the coordinate sets or not). To assess the remaining combinations, the coordinates of the fixed set were temporally smoothed with a 1-D Gaussian filter of width 200 ms and linearized. The time-varying directions of the toroidal trajectory was then computed as 'arctan2' of the derivatives of a cubic spline fitted to each unwrapped angle. The derivative of the sine and cosine of these directions were then compared with that of each combination of orientation and pair of axes for the second coordinate set. Note, in each assessment, the origins were aligned by subtracting the mean angular difference between the pair of coordinate sets. The combination that minimized the correlation between the two sets was recognized as the optimal alignment.

To align the toroidal coordinates of the ensembles with OF recordings, the combination of axes and orientation which looked most similar across different mice and sessions, when decoding the OF data and plotting the mean toroidal coordinate in 30^2 square bins of the OF coordinates (Fig. 1a), was chosen.

Generalized Linear Model

Bernoulli and Poisson generalized linear models were fitted to (binarized) calcium events and binned spike counts, respectively, using either the toroidal coordinates or the spatial positions as regressors, allowing a comparison of the explanatory power of the two descriptors. The setup was like Lederger et al. (2021) and Gardner et al. (2022), and included a smoothness prior to avoid overfitting.

Design matrices, D, were first constructed by binned the covariates and defining $D_i(t) = 1$ if the covariate value at time t was contained in the i-th bin and $D_i(t) = 0$ otherwise.

The probability of recording n spikes at time point t is given as:

$$P(n|\lambda_P(t,\gamma)) = \exp(-\lambda_P(t,\gamma))\frac{\lambda_P(t,\gamma)^n}{n!},$$

and the probability of recording a calcium event or not, $n \in \{0,1\}$, is given as:

$$P(n|\lambda_B(t,\gamma)) = \lambda_B(t,\gamma)^n(1-\lambda_B(t,\gamma))^{1-n},$$

Here, $\lambda_P(t,\gamma) = \exp(\sum_i \gamma_i D_i(t))$ and $\lambda_B(t,\gamma) = \exp(\sum_i \gamma_i D_i(t))/(1 + \exp(\sum_i \gamma_i D_i(t)))$ is the expected activity at time t and the parameters γ were optimized by minimizing the cost function

$$L(\gamma|\lambda(t,\gamma),\rho,k) = -\sum_t \ln(P(n(t)|\lambda_*(t,\gamma))) + \frac{1}{2}\rho \sum_{i,j\in N}(\gamma_i-\gamma_j)^2,$$

where N describes indices of neighboring bins. Note that the second term effectively minimizes large differences between neighboring parameters.

The minimization was run by first initializing each γ to zero and running two iterations of gradient descent on the above loss function before and after applying the 'l-bfgs-b'-algorithm (given in the 'scipy.optimize'-module with 'gtol'$=10^{-5}$).

The parameters ρ and number of bins were optimized for values $0, 1, \sqrt{10}$ and 10 and $8^2, 10^2, 12^2$ and 15^2, respectively for the OF arena of mouse 88592 day 2, and were found to be 1 and 10. These values were used for all recording days, and for mice 97045 and 97046. The 1-D VR-track positions were binned in 16 bins and $\rho = 1$ were used. To have the same number of bins for the 2-D toroidal positions, a 4×4 grid was used.

To avoid overfitting, the scheme was three-fold cross validated, i.e., the model was fitted to two-thirds of the data and its performance evaluated on the last third three times. The toroidal coordinates were decoded separately when used to fit each neuron's firing rate, omitting the activity of the neuron studied in the extrapolation (see 'Cohomological coordinatization'). Moreover, for the VR-sessions, baseline trials were decoded using the toroidal description found during dark session.

The deviance explained was computed, after similarly fitting a null model and a saturated model to the data, as:

$$1 - \frac{\sum L_s - \sum L_m}{\sum L_s - \sum L_0},$$

where L_s, L_m and L_0 are the log-likelihoods of the saturated, fitted and null model and the sums are over all time points in the given session.

Comparison of toroidal tuning

Peak toroidal positions were visually compared across session to assess the conservation of toroidal tuning. The preferred locations were computed as the mass center of each neuron's activity distribution on the torus (16^2 bins), given by:

$$T_{peak} = \text{arctan2}\left(\frac{\sum_i \sin\phi \cdot s_i}{\sum_i s_i}, \frac{\sum_i \cos\phi \cdot s_i}{\sum_i s_i}\right),$$

where s_i is the mean activity in the i-th toroidal bin ϕ_i.

A second comparison was made by flattening the 2-D activity distributions (as above) and computing the Pearson correlation coefficient (using the 'pearsonr'-function given in the 'scipy.stats'-library) between two distributions from the same neuron in different conditions. To generate a random distribution, the cell indices in one session were randomly shuffled before computing correlation with the second session, i.e., computing the correlation of the rate maps of two different neurons. This process was repeated 1000 times.

Continuous attractor network simulations

To study the toroidal structure of grid cell continuous attractor network (CAN) models, firing rate activity was simulated using three noiseless CAN models (Extended Data Fig. 3a).

First, a 56×44 grid cell network with purely inhibitory connectivity was simulated as proposed in [78]. Animal behavior was modeled by using the first 1000 s of the recorded trajectory of rat 'R' during OF foraging session day 1 found in [10]. The spatial positions were interpolated at 2 ms time steps and the speed $s(t)$ and head direction $\phi(t)$ were computed as a function of the difference in position between each time step. The firing rate for time step t_{n+1} was given as:

$$f_{t_{n+1}} = f_{t_n} + \frac{1}{\rho}(-f_{t_n} + [J + f_{t_n} \cdot W + \gamma s(t_n)\cos(\phi(t_n) - \overline{\phi})]_+),$$

where $[\ldots]_+$ is the Heaviside function, $\overline{\phi}$ is the preferred head direction and parameters were given as: $J = 1, \gamma = 0.15, l = 2, R = 20, W_0 = -0.01$ and $\rho = 10$. The activity pattern was initialized to random and subsequently stabilized by running 2000 iterations of the above equation with no movement. The firing rate was set to 0 if below 0.0001 and resampled at 10 ms time steps.

Next, the twisted torus model by Guanella et al [53] was used to simulate a 20×20 grid cell network for a simulated random walk in a square box. The parameter definitions and code can be found in the open-source implementation by Santos Pata (https://github.com/DiogoSantosPata/gridcells). Here, we used 5000 time frames and 'grid_gain=0.06'.

Finally, an untwisted version of the previous model created square grid cell patterns for a 10×10 network. This was simulated using a Python translation of the Matlab implementation of Zilli [79]. 20 ms time steps were used, and a total duration of 295 s was simulated, following an OF trajectory recorded by Hafting et al. [1], provided in the same code repository.

Head direction cell network

The UMAPH framework was applied to electrophysiological recordings from the mouse anterodorsal thalamic nucleus and subiculum of mouse 25 day 140130 ($n = 25$ neurons), mouse 17 day 130130 ($n = 37$ neurons) and mouse 28 day 140313 ($n = 62$ neurons) [80, 81]. The data contained 23- to 192-minute recordings during wakefulness, rapid-eye movement sleep (REM) and slow-wave sleep (SWS). Each mouse and brain state were processed and analyzed alike, similar to previously described in 'Preprocessing of neural data'.

First, delta functions (one per spike) were convolved with a Gaussian kernel of $\sigma = 250$ ms for wake and REM sleep recordings and $\sigma = 50$ ms for SWS sleep. Population vectors were then sampled at 250 ms (wake and REM) and 2 s (SWS) intervals. Neurons were clustered as described in 'Unsupervised clustering' (based on the mean cross-correlation of the firing rates across brain states) and the largest cluster in each session was analyzed ($n = 18, 25$ and 41 neurons). In contrast to the entorhinal recordings, radial downsampling was not performed (i.e., $\epsilon = 0$), and parameters were set to $d = 4, m = 500, \kappa = 300$, resulting in barcodes indicating ring topology (Extended Data Fig. 6a). For visualization, the 6 first (whitened) principal components of the firing rates were projected to 3-D using UMAP (default parameters, Extended Data Fig. 6b).

Circular coordinatization of the longest-lived H^1-bar was used to decode both awake and sleep data (Extended Data Fig. 6c). The decoded angles were reoriented by testing clockwise and anti-clockwise orientation and origin fixed by minimizing the mean difference to the recorded head direction angles during wake sessions. The angular tuning curves were computed using 60 bins and max-normalized.

Three-torus

The applicability of the UMAPH framework in detecting higher-dimensional features was tested on a three-torus ($T^3 = S^1 \times S^1 \times S^1$). First, $8^3 = 512$ points (i.e., three-dimensional angular values) were evenly sampled from a three-torus (Extended Data Fig. 5c) and added small, independent noise sampled from a normal distribution of width 0.1. The point cloud was embedded in $\mathbb{R}^6$ by taking the sine and cosine of each angle. Applying UMAPH with Euclidean metric and $\kappa = 70$ revealed the expected homology and the three circular features were decoded based on the three longest-lived H^1-bars (see 'Cohomological coordinatization'). To visually compare with UMAP, the 6-D embedding was projected to 3-D using UMAP (with 150 as number of neighbors).

Data analysis and statistics

All data analyses were performed with custom-written scripts in Python 3.9. The following open-source Python packages were used: umap (version 0.5.2), ripser (0.6.1), numba (0.54.1), scipy (1.7.3), numpy (1.20.3), scikit-learn (0.24.2), matplotlib (3.4.2), h5py (3.6.0), gtda (0.5.1), cv2 (4.5.5), pandas (1.4.2) and datajoint (0.13.5).

The heaviest computational burdens were performed on resources provided by the NTNU IDUN/EPIC computing cluster [82] and that of the Department of Mathematical Sciences.

All statistical tests were one-sided.

Code Availability

Code used in this article will be made available upon publication at GitHub.

Data Availability

The data are publicly shared by the respective authors at: https://plus.figshare.com/articles/dataset/VR_Data_Neuropixel_supporting_Distance-tuned_neurons_drive_specialized_path_integration_calculations_in_medial_entorhinal_cortex_/15041316, https://archive.sigma2.no/pages/public/datasetDetail.jsf?id=10.11582/2022.00008, https://archive.sigma2.no/pages/public/datasetDetail.jsf?id=10.11582/2022.00005 and https://crcns.org/data-sets/thalamus/th-1.

Video 1. Comparison of open field foraging and toroidal population dynamics. Left, scatter plot of position in 2-D space of mouse 88592 day 2. Right, flattened torus, displaying toroidal population distribution at each time frame, given by first smoothing each neuron's toroidal rate map (in 50^2 bins) spatially with a Gaussian filter with $\sigma = 3$ bins. Next, every rate map was weighted by the firing rate of the corresponding neuron at each time point and the mean value for each bin across the entire population gave the time-dependent population distribution. Bright colors indicate high population activity.

Video 2. Comparison of wheel running and toroidal population dynamics for same neural ensemble as in Video 1. Left, wheel position (y axis) as a function of time (x axis). Right, as in Video 1 for internal toroidal representation decoded during wheel running.

Video 3. Comparison of VR running and toroidal population dynamics for gain sessions of mouse I3 recording day 0420. Left, wheel position as in Video 2 (top) and gain values (y axis) as a function of time (bottom). Right, inferred toroidal population dynamics, as in Video 1 and 2, found in the given recording.

References

[1] Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. Nature **436**(7052), 801–806 (2005)

[2] Casali, G., Shipley, S., Dowell, C., Hayman, R., Barry, C.: Entorhinal neurons exhibit cue locking in rodent vr. Frontiers in cellular neuroscience **12**, 512 (2019)

[3] Gao, P., Ganguli, S.: On simplicity and complexity in the brave new world of large-scale neuroscience. Current opinion in neurobiology **32**, 148–155 (2015)

[4] Musall, S., Urai, A.E., Sussillo, D., Churchland, A.K.: Harnessing behavioral diversity to understand neural computations for cognition. Current opinion in neurobiology **58**, 229–238 (2019)

[5] Gao, P., Trautmann, E., Yu, B., Santhanam, G., Ryu, S., Shenoy, K.,

Ganguli, S.: A theory of multineuronal dimensionality, dynamics and measurement. BioRxiv, 214262 (2017)

[6] Yoon, K., Lewallen, S., Kinkhabwala, A.A., Tank, D.W., Fiete, I.R.: Grid cell responses in 1d environments assessed as slices through a 2d lattice. Neuron **89**(5), 1086–1099 (2016)

[7] McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426 (2018)

[8] Carlsson, G.: Topology and data. Bulletin of the American Mathematical Society **46**(2), 255–308 (2009)

[9] Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., Moser, E.I.: The entorhinal grid map is discretized. Nature **492**(7427), 72–78 (2012)

[10] Gardner, R.J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N.A., Dunn, B.J., Moser, M.-B., Moser, E.I.: Toroidal topology of population activity in grid cells. bioRxiv (2021)

[11] Datta, S.R., Anderson, D.J., Branson, K., Perona, P., Leifer, A.: Computational neuroethology: a call to action. Neuron **104**(1), 11–24 (2019)

[12] Paninski, L., Cunningham, J.P.: Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. Current opinion in neurobiology **50**, 232–241 (2018)

[13] Olshausen, B.A., Field, D.J.: How close are we to understanding v1? Neural computation **17**(8), 1665–1699 (2005)

[14] Rust, N.C., Movshon, J.A.: In praise of artifice. Nature neuroscience **8**(12), 1647–1650 (2005)

[15] Moser, E.I., Kropff, E., Moser, M.-B., *et al.*: Place cells, grid cells, and the brain's spatial representation system. Annual review of neuroscience **31**(1), 69–89 (2008)

[16] Rowland, D.C., Roudi, Y., Moser, M.-B., Moser, E.I., *et al.*: Ten years of grid cells. Annu Rev Neurosci **39**, 19–40 (2016)

[17] Campbell, M.G., Ocko, S.A., Mallory, C.S., Low, I.I., Ganguli, S., Giocomo, L.M.: Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. Nature neuroscience **21**(8), 1096–1106 (2018)

[18] Kinkhabwala, A.A., Gu, Y., Aronov, D., Tank, D.W.: Visual cue-related

activity of cells in the medial entorhinal cortex during navigation in virtual reality. Elife **9**, 43140 (2020)

[19] Low, I.I., Williams, A.H., Campbell, M.G., Linderman, S.W., Giocomo, L.M.: Dynamic and reversible remapping of network representations in an unchanging environment. Neuron **109**(18), 2967–2980 (2021)

[20] Pröll, M., Häusler, S., Herz, A.V.: Grid-cell activity on linear tracks indicates purely translational remapping of 2d firing patterns at movement turning points. Journal of Neuroscience **38**(31), 7004–7011 (2018)

[21] Jacob, P.-Y., Capitano, F., Poucet, B., Save, E., Sargolini, F.: Path integration maintains spatial periodicity of grid cell firing in a 1d circular track. Nature Communications **10**(1), 1–13 (2019)

[22] Campbell, M.G., Attinger, A., Ocko, S.A., Ganguli, S., Giocomo, L.M.: Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex. Cell reports **36**(10), 109669 (2021)

[23] Kriegeskorte, N., Wei, X.-X.: Neural tuning and representational geometry. Nature Reviews Neuroscience **22**(11), 703–718 (2021)

[24] Rust, N.C., Cohen, M.R.: Priority coding in the visual system. Nature Reviews Neuroscience **23**(6), 376–388 (2022)

[25] Panzeri, S., Moroni, M., Safaai, H., Harvey, C.D.: The structures and functions of correlations in neural population codes. Nature Reviews Neuroscience, 1–17 (2022)

[26] Fuhs, M.C., Touretzky, D.S.: A spin glass model of path integration in rat medial entorhinal cortex. Journal of Neuroscience **26**(16), 4266–4276 (2006)

[27] McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., Moser, M.-B.: Path integration and the neural basis of the'cognitive map'. Nature Reviews Neuroscience **7**(8), 663–678 (2006)

[28] Lehky, S.R., Tanaka, K.: Neural representation for object recognition in inferotemporal cortex. Current opinion in neurobiology **37**, 23–35 (2016)

[29] Chung, S., Lee, D.D., Sompolinsky, H.: Linear readout of object manifolds. Physical Review E **93**(6), 060301 (2016)

[30] Rubin, A., Sheintuch, L., Brande-Eilat, N., Pinchasof, O., Rechavi, Y., Geva, N., Ziv, Y.: Revealing neural correlates of behavior without behavioral measurements. Nature communications **10**(1), 1–14 (2019)

[31] Giusti, C., Pastalkova, E., Curto, C., Itskov, V.: Clique topology reveals

intrinsic geometric structure in neural correlations. Proceedings of the National Academy of Sciences **112**(44), 13455–13460 (2015)

[32] Curto, C.: What can topology tell us about the neural code? Bulletin of the American Mathematical Society **54**(1), 63–78 (2017)

[33] Giusti, C., Ghrist, R., Bassett, D.S.: Two's company, three (or more) is a simplex. Journal of computational neuroscience **41**(1), 1–14 (2016)

[34] Rybakken, E., Baas, N., Dunn, B.: Decoding of neural data using cohomological feature extraction. Neural computation **31**(1), 68–93 (2019)

[35] Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., Fiete, I.: The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. Nature neuroscience **22**(9), 1512–1520 (2019)

[36] Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008)

[37] Hinton, G.E., Salakhutdinov, R.R.: Reducing the dimensionality of data with neural networks. science **313**(5786), 504–507 (2006)

[38] Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)

[39] Pei, F., Ye, J., Zoltowski, D., Wu, A., Chowdhury, R.H., Sohn, H., O'Doherty, J.E., Shenoy, K.V., Kaufman, M.T., Churchland, M., et al.: Neural latents benchmark'21: Evaluating latent variable models of neural population activity. arXiv preprint arXiv:2109.04463 (2021)

[40] Kobak, D., Linderman, G.C.: Initialization is critical for preserving global data structure in both t-sne and umap. Nature Biotechnology **39**(2), 156–157 (2021)

[41] Kobak, D., Berens, P.: The art of using t-sne for single-cell transcriptomics. Nature communications **10**(1), 1–14 (2019)

[42] Becht, E., McInnes, L., Healy, J., Dutertre, C.-A., Kwok, I.W., Ng, L.G., Ginhoux, F., Newell, E.W.: Dimensionality reduction for visualizing single-cell data using umap. Nature biotechnology **37**(1), 38–44 (2019)

[43] Dorrity, M.W., Saunders, L.M., Queitsch, C., Fields, S., Trapnell, C.: Dimensionality reduction by umap to visualize physical and genetic interactions. Nature communications **11**(1), 1–6 (2020)

[44] Chari, T., Banerjee, J., Pachter, L.: The specious art of single-cell genomics. bioRxiv (2021)

[45] Edelsbrunner, H., Letscher, D., Zomorodian, A.: Topological persistence and simplification. In: Proceedings 41st Annual Symposium on Foundations of Computer Science, pp. 454–463 (2000). IEEE

[46] Georgopoulos, A.P., *et al.*: On reaching. Annual review of neuroscience **9**(1), 147–170 (1986)

[47] Taube, J.S., Muller, R.U., Ranck, J.B.: Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. Journal of Neuroscience **10**(2), 420–435 (1990)

[48] Somasundaram, E.V., Brown, S.E., Litzler, A., Scott, J.G., Wadhwa, R.R.: Benchmarking r packages for calculation of persistent homology. The R journal **13**(1), 184 (2021)

[49] De Silva, V., Morozov, D., Vejdemo-Johansson, M.: Persistent cohomology and circular coordinates. Discrete & Computational Geometry **45**(4), 737–759 (2011)

[50] Zong, W., Obenhaus, H.A., Skytøen, E.R., Eneqvist, H., de Jong, N.L., Vale, R., Jorge, M.R., Moser, M.-B., Moser, E.I.: Large-scale two-photon calcium imaging in freely moving mice. Cell **185**(7), 1240–1256 (2022)

[51] Obenhaus, H.A., Zong, W., Jacobsen, R.I., Rose, T., Donato, F., Chen, L., Cheng, H., Bonhoeffer, T., Moser, M.-B., Moser, E.I.: Functional network topography of the medial entorhinal cortex. Proceedings of the National Academy of Sciences **119**(7), 2121655119 (2022)

[52] Faulhuber, M.: An application of hypergeometric functions to heat kernels on rectangular and hexagonal tori and a "weltkonstante"-or-how ramanujan split temperatures. The Ramanujan Journal **54**(1), 1–27 (2021)

[53] Guanella, A., Kiper, D., Verschure, P.: A model of grid cells based on a twisted torus topology. International journal of neural systems **17**(04), 231–240 (2007)

[54] Waaga, T., Agmon, H., Normand, V.A., Nagelhus, A., Gardner, R.J., Moser, M.-B., Moser, E.I., Burak, Y.: Grid-cell modules remain coordinated when neural activity is dissociated from external sensory cues. Neuron **110**(11), 1843–1856 (2022)

[55] Baden, T., Euler, T., Berens, P.: Understanding the retinal basis of vision across species. Nature Reviews Neuroscience **21**(1), 5–20 (2020)

[56] Constantinescu, A.O., O'Reilly, J.X., Behrens, T.E.: Organizing conceptual knowledge in humans with a gridlike code. Science **352**(6292),

1464–1468 (2016)

[57] Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.-B., Moser, E.I.: Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science **312**(5774), 758–762 (2006)

[58] Kang, L., Xu, B., Morozov, D.: Evaluating state space discovery by persistent cohomology in the spatial representation system. Frontiers in computational neuroscience **15**, 28 (2021)

[59] Cogno, S.G., Obenhaus, H.A., Jacobsen, R.I., Donato, F., Moser, M.-B., Moser, E.I.: Minute-scale oscillatory sequences in medial entorhinal cortex. bioRxiv (2022)

[60] Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D., *et al.*: Procedures for behavioral experiments in head-fixed mice. PloS one **9**(2), 88678 (2014)

[61] Killian, N.J., Jutras, M.J., Buffalo, E.A.: A map of visual space in the primate entorhinal cortex. Nature **491**(7426), 761–764 (2012)

[62] Mathis, A., Stemmler, M.B., Herz, A.V.: Probable nature of higher-dimensional symmetries underlying mammalian grid-cell activity patterns. Elife **4** (2015)

[63] Zong, W., Obenhaus, H., Skytøen, E., Eneqvist, H., De Jong, N., Jorge, M., Moser, M.-B., Moser, E.: Zong 2022. Archive2014 (2022). https://doi.org/10.11582/2022.00008

[64] Whiteley, N., Gray, A., Rubin-Delanchy, P.: Discovering latent topology and geometry in data: a law of large dimension. arXiv preprint arXiv:2208.11665 (2022)

[65] Obenhaus, H., Zong, W., Jacobsen, R.I., Rose, T., Donato, F., Chen, L., Cheng, H., Bonhoeffer, T., Moser, M.-B., Moser, E.I.: Obenhaus 2022. Archive2014 (2022). https://doi.org/10.11582/2022.00005

[66] Giocomo, L., Campbell, M.G., Attinger, A.: VR Data Neuropixel supporting "Distance-tuned neurons drive specialized path integration calculations in medial entorhinal cortex" (2021). https://doi.org/10.25452/figshare.plus.15041316

[67] Beshkov, K., Tiesinga, P.: Geodesic-based distance reveals nonlinear topological features in neural activity from mouse visual cortex. Biological Cybernetics **116**(1), 53–68 (2022)

[68] McInnes, L.: Topological methods for unsupervised learning. In: International Conference on Geometric Science of Information, pp. 343–350 (2019). Springer

[69] Zadeh, L.A.: Fuzzy sets. Information and control **8**(3), 338–353 (1965)

[70] Jardine, J.: Stability for umap. arXiv preprint arXiv:2011.13430 (2020)

[71] Perea, J.A.: Sparse circular coordinates via principal z-bundles. In: Baas, N.A., Carlsson, G.E., Quick, G., Szymik, M., Thaule, M. (eds.) Topological Data Analysis, pp. 435–458. Springer, Cham (2020)

[72] Fukunaga, K., Hostetler, L.: The estimation of the gradient of a density function, with applications in pattern recognition. IEEE Transactions on information theory **21**(1), 32–40 (1975)

[73] Kloke, J., Carlsson, G.: Topological de-noising: Strengthening the topological signal. arXiv preprint arXiv:0910.5947 (2009)

[74] Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2002)

[75] Bauer, U.: Ripser: efficient computation of vietoris–rips persistence barcodes. Journal of Applied and Computational Topology **5**(3), 391–423 (2021)

[76] Tralie, C., Saul, N., Bar-On, R.: Ripser. py: A lean persistent homology library for python. Journal of Open Source Software **3**(29), 925 (2018)

[77] Robins, V., Turner, K.: Principal component analysis of persistent homology rank functions with case studies of spatial point patterns, sphere packing and colloids. Physica D: Nonlinear Phenomena **334**, 99–117 (2016)

[78] Couey, J.J., Witoelar, A., Zhang, S.-J., Zheng, K., Ye, J., Dunn, B., Czajkowski, R., Moser, M.-B., Moser, E.I., Roudi, Y., *et al.*: Recurrent inhibitory circuitry as a mechanism for grid formation. Nature neuroscience **16**(3), 318–324 (2013)

[79] Zilli, E.A.: Models of grid cell spatial firing published 2005–2011. Frontiers in neural circuits **6**, 16 (2012)

[80] Peyrache, A., Lacroix, M.M., Petersen, P.C., Buzsáki, G.: Internally organized mechanisms of the head direction sense. Nature neuroscience **18**(4), 569–575 (2015)

[81] A. Peyrache, G.B. P. Petersen: Peyrache 2015. CRCNS.org (2015). https://doi.org/10.6080/K0G15XS1

[82] Själander, M., Jahre, M., Tufte, G., Reissmann, N.: Epic: An energy-efficient, high-performance gpgpu computing research infrastructure. arXiv preprint arXiv:1912.05848 (2019)

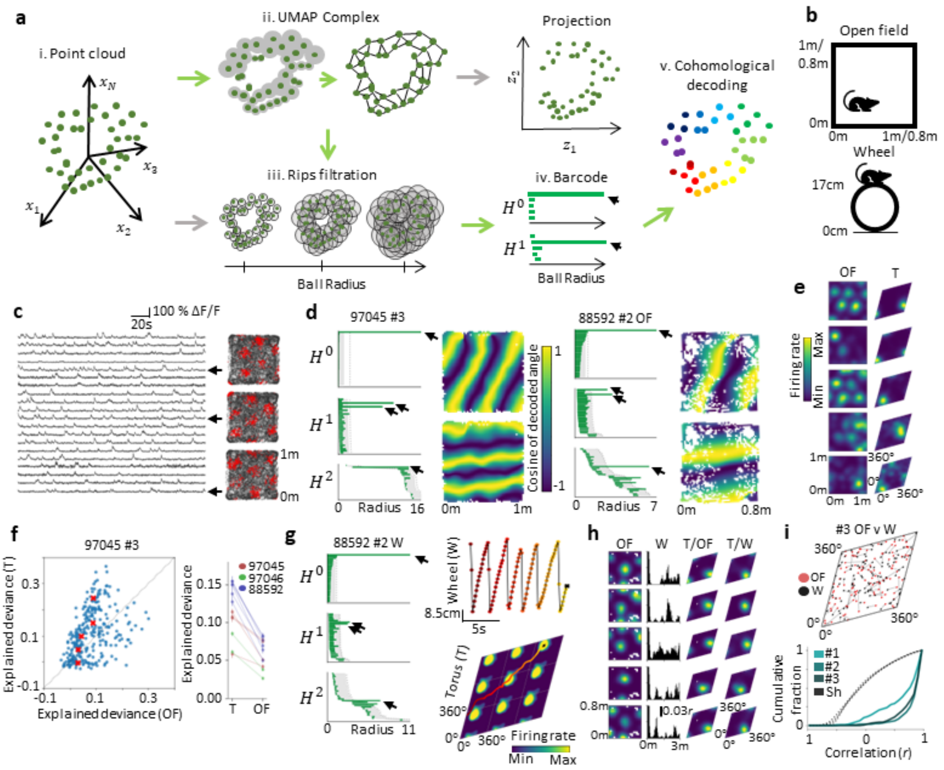

Fig. 1 Toroidal representations in calcium imaging of mouse entorhinal cortex in 1- and 2-D environments. **a.** UMAPH (green arrows) combines aspects of UMAP (top arrows) and persistent homology (bottom arrows). Given a point cloud (green dots) in ambient space (i), UMAP computes a topological representation of the data (ii) based on local neighborhoods and approximates this in a low-dimensional projection of the data. Persistent homology constructs a sequence of combinatorial spaces from pairwise distances (iii) and tracks the evolution of n-dimensional holes (H^n) in these spaces for increasing distances, displayed as bars in a barcode (iv). The long bars depict the most prominent features in the data (black arrows). 1-D *co*homology features allow for circular coordinatization (v). **b.** The mice ran in an OF arena (top) and on a Styrofoam wheel (bottom). **c.** 20 example calcium traces from the MEC (left). Coloring the spatial positions by heightened activity (red) of three example neurons reveals hexagonal grid patterns (right). **d.** Left in each block, barcodes from OF sessions in mice 97045 (day 3, $n = 293$ neurons) and 88592 (day 2, $n = 160$), indicating signature homology of a torus (one H^0-bar, two H^1-bars and one H^2-bar). Only the 30 longest bars in each dimension are shown. Shaded gray shows 95th percentile and dotted line the maximum of the longest lifetimes of 100 shuffles. Right in each block, OF positions color-coded by mean decoded circular coordinates. **e.** Tuning of five example neurons (mouse 97045 day 3, see Fig. 7 for all neurons) to position in OF arena and on torus (T), sorted by explained deviance scores. **f.** Left, explained deviance scores for all recorded neurons in mouse 97045 day 3 describing goodness-of-fit for GLM models fitted to calcium events for T vs. OF positions as covariates. Red crosses mark the five cells shown in **e**. Right, mean value (± s.e.m.) for each OF (and object) session where toroidal topology was detected (Extended Data Fig. 1e, $n \in [113, 637]$ neurons). **g.** Barcode from grid cell activity during wheel running of mouse 88592 day 2 indicates toroidal topology (left). Right, internal (bottom) and external (top) trajectory, color-coded by path progression, on the unwrapped torus during wheel running. Shading on the torus indicates mean population activity (indicated by color bar) at end of the path (black cross). **h.** Rate maps of top five neurons (mouse 88592 day 2), ranked according to toroidal explained deviance. From left: OF tuning; linearized wheel position autocorrelogram and tuning to torus inferred from corresponding sessions. Color-coding as in **e**. **i.** Top, toroidal tuning peaks of grid cell population in OF (red) and wheel (black) recordings (day 3). Lines connect peaks of the same neurons. Bottom, cumulative line plots of the correlation between rate maps from OF and wheel sessions for mouse 88592 days 1 – 3 ($n = 126, 160$ and 113 neurons). Dotted lines show shuffled distributions ($n = 1000$ shuffles).

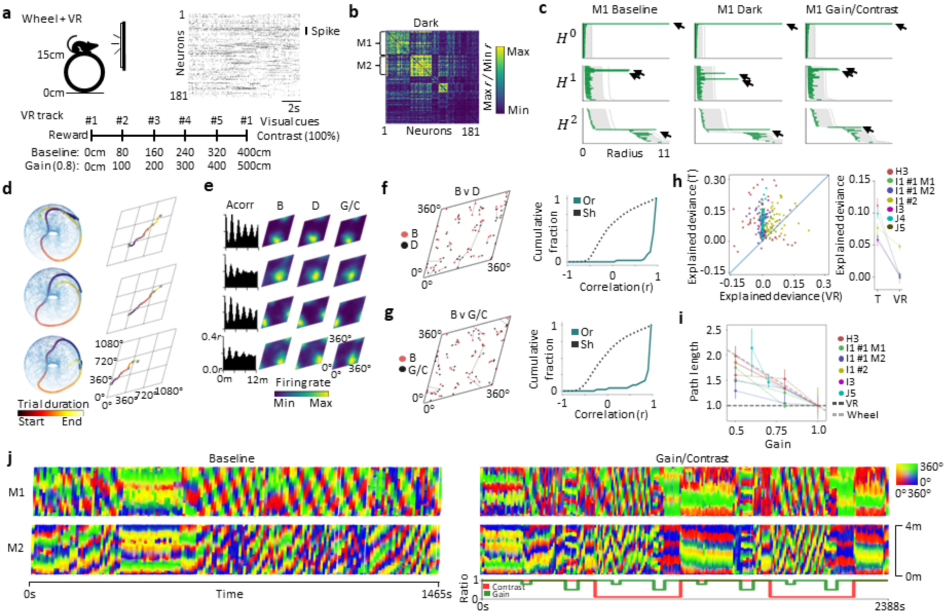

Fig. 2 UMAPH detects tori in spiking activity of head-fixed mice during a VR task. **a.** Electrophysiological recordings were made while head-fixed mice engaged in a virtual reality task showing (for baseline trials) a continuously repeating linear VR track of 400 cm with five evenly spaced visual cues with a water reward at the end. During gain and contrast trials, the gain of the visual flow to the running wheel and the contrast of the visual cues were reduced. **b.** Cross-correlation matrix during dark session (monitors turned off and no rewards given) of mouse I1. Values indicated by color bar, minimum and maximum defined as 5th and 95th percentile. The two largest clusters, 'M1' and 'M2' ($n = 44$ and 41 neurons) are indicated. **c.** Barcodes (as in Fig. 1d, g) indicate a toroidal state space in every condition for M1. **d.** Example trajectories on the decoded neural torus for three different trials, embedded in 3-D (left) and linearized to 2-D (right). **e.** Virtual track autocorrelogram and toroidal firing rate maps (as in 1e and h) in baseline (B), dark (D) and gain/contrast (G/C) sessions for four example neurons of M1 (see Extended Data Fig. 8 – 11 for all neurons). **f, g.** Peak activity locations (left) for all neurons in M1 in B (red) vs. D and G/C sessions (black), and cumulative distribution of Pearson correlation between rate maps of M1 of corresponding conditions (blue) (as in Fig. 1i). **h.** Explained deviance scores (as in Fig. 1f) for T vs. VR positions, for GLM models fitted to spike train data during baseline trials. Right, mean explained deviance (± s.e.m.) per ensemble ($n = 25-65$). The toroidal description explained the data better than virtual track in each case ($P \in [3.8 \cdot 10^{-9}, 7.5 \cdot 10^{-6}], W \in [307, 1747]$, Wilcoxon signed-rank test). **i.** Median and interquartile path length of the internal toroidal trajectory per trial relative to baseline trials of similar contrast (10 or 100%) for different gain values (x axis) ($3-8$ gain trials for each gain and contrast combination, and $44-178$ baseline trials). Dashed lines indicate expected trial length if strictly aligned to distance on wheel (gray) or to VR track (black). **j.** Each VR-position (y axis) during B (left) and G/C (right) sessions for modules M1 (top) and M2 (bottom), color-coded by the instantaneous toroidal coordinates. Bottom right, time-dependent gain (green) and contrast (red) values, given as ratio of baseline.

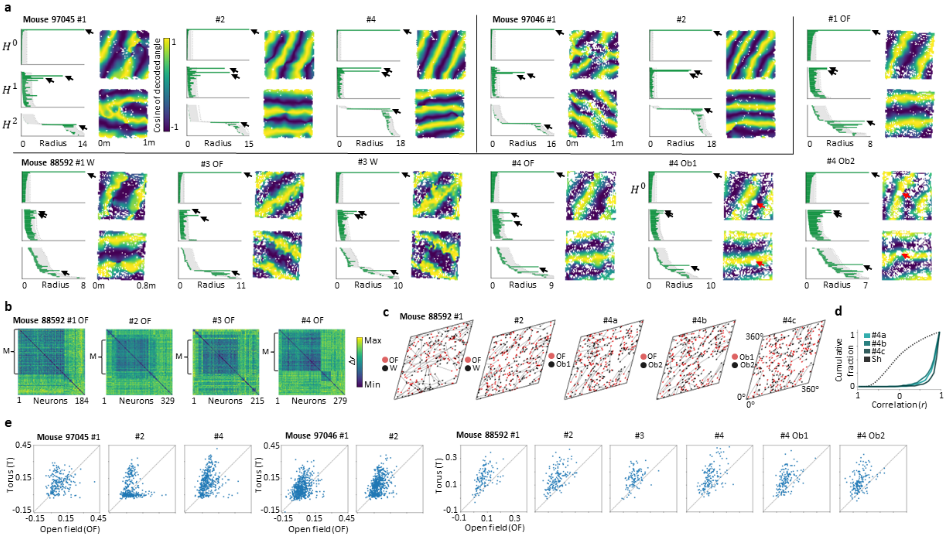

Extended Data Fig. 1 Barcodes, spatial mappings, distance matrices and toroidal tuning statistics for calcium recordings not shown in Fig. 1. **a.** Barcodes and mean toroidal coordinates as a function of OF locations (as in Fig. 1d) showing the results of UMAPH analysis for days 1 − 4 ($n = 113 - 637$ neurons). For wheel sessions, OF toroidal coordinates are found by using the longest-lived H^1-bars to decode the activity during the OF session. The toroidal coordinates are reoriented to match across sessions. In two sessions (day 4 Obj1 and 2), an object was inserted into the environment (red arrows). **b.** Distance matrices of spatial autocorrelograms during OF foraging session for mouse 88592 across four different recording days. Values indicated by color bar. The clusters analyzed are marked as 'M' ($n = 113 - 162$ neurons). **c, d.** Distributions of receptive field centers and cumulative distributions of Pearson correlation of pairs of toroidal rate maps (as in Fig. 1i) for all neurons in each recording and comparison across environments. **e.** Comparison of explained deviance scores between OF locations and toroidal coordinates (as in Fig. 1f) for each neuron included in the UMAPH analyses in **a**.

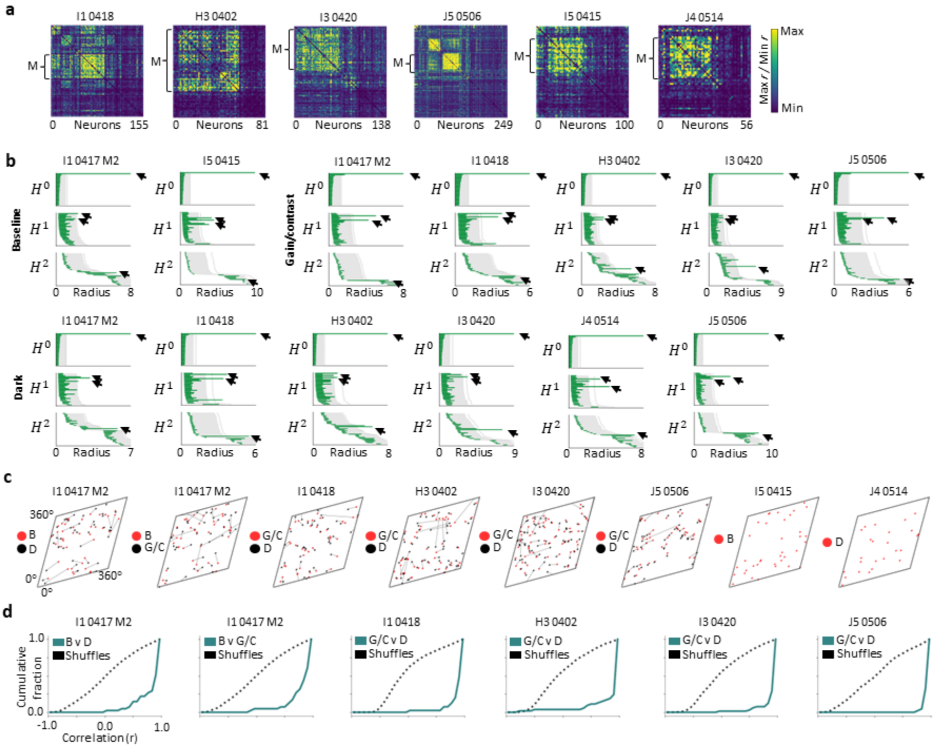

Extended Data Fig. 2 Correlation matrices, barcodes and toroidal tuning statistics for Neuropixels recordings classified as toroidal, not shown in Fig. 2. **a.** Correlation matrices (maximum divided by minimum) during dark session (as in Fig. 2b) sorted by cluster indices. Minimum and maximum range from 1.02 – 1.05 r/r and 1.59 – 2.86 r/r respectively. The clusters analyzed are marked as 'M' ($n = 25 - 65$ neurons). **b.** Barcode diagrams (as in Fig. 2c) for each session. **c, d.** Distributions of receptive field centers on inferred torus and cumulative distributions showing Pearson correlation of pairs of toroidal rate maps (as in Fig. 2f, g) for all neurons in each ensemble, compared across sessions.

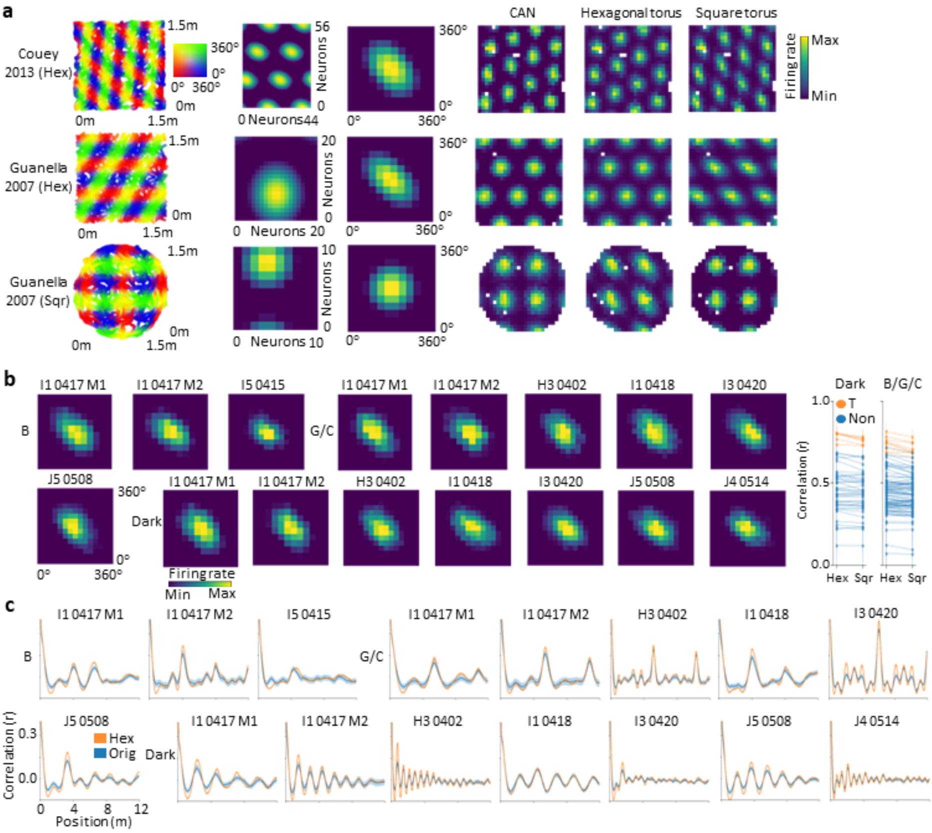

Extended Data Fig. 3 Decoding of toroidal coordinates in simulated grid cell and entorhinal ensemble activity allows comparing square and hexagonal toroidal structure. **a.** Circular coordinates from UMAPH analyses on data generated by three CAN models (row-wise, from top): an inhibitory, multi-bump network model ($n = 2464$ neurons), a twisted ($n = 400$) and an un-twisted (square, $n = 100$) torus network model. Each row displays (from left): a 2-D spatial trajectory colored according to toroidal position (2-D color map); firing rates of each neuron at a chosen time frame, ordered according to network connectivity (firing rate given by color bar); stacked, centered toroidal rate maps of all neurons in the network; spatial rate maps of activity from a CAN grid cell and fitted data using hexagonal and square torus point source models. **b.** Stacked, centered toroidal rate maps of all neurons in each ensemble ($n = 25-65$ neurons). Note angle of the firing fields. Right, median (and interquartile) correlation, per experimental condition, between each neuron's rate map and the heat distribution on a hexagonal (Hex) and a square (Sqr) torus for all ensembles. Ensembles determined as toroidal are colored in orange. **c.** Mean spatial autocorrelograms (+- s.e.m) for each toroidal ensemble for recorded data (blue) and data generated by the hexagonal torus model (orange).

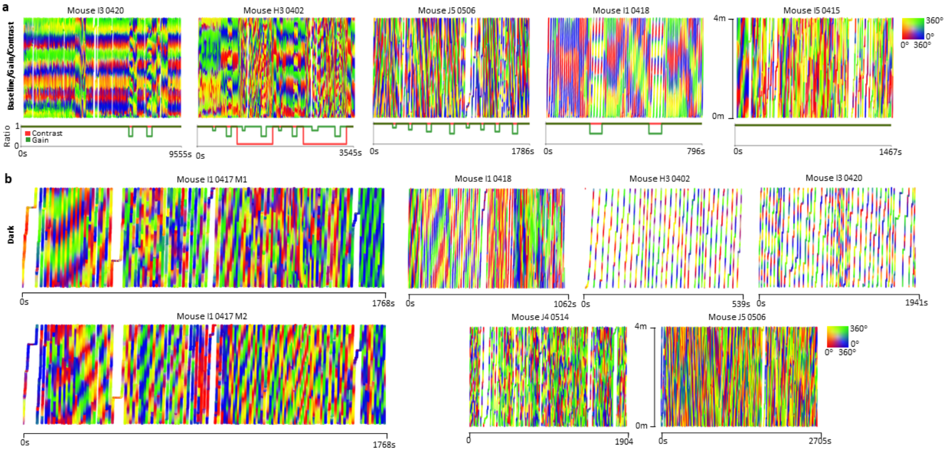

Extended Data Fig. 4 Internal toroidal dynamics as a function of the VR track for sessions not shown in Fig. 1j. **a, b.** Spatio-temporal positions colored by the 2-D toroidal positions across the entire session for baseline, gain, contrast and dark sessions. Gain and contrast manipulations are indicated in green and red line plots below relevant sessions. Note, in dark, the mice cannot see a progression in VR.

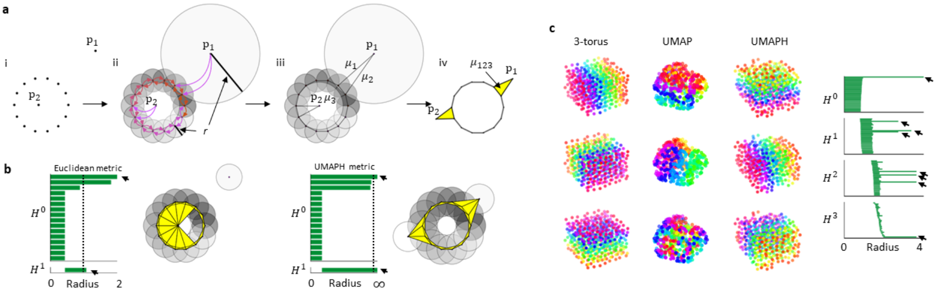

Extended Data Fig. 5 UMAPH visual description and example application for high-dimensional manifold. **a.** The UMAP complex construction for a 'noisy' circular point cloud. i. Sample point cloud with two outliers, p_1 and p_2. ii. The geodesic distance of the underlying manifold is approximated by first computing local metric spaces (one for each point) through scaling the ambient space metric with respect to each point's $\kappa = 2$ nearest neighbors. iii. The local metric spaces are merged to a global representation by translating each to a fuzzy set and then taking the fuzzy union of these. This creates a graph with weights μ_i, describing edge (i) probabilities based on the scaled distance. iv. The UMAP complex is the weighted clique complex of this graph. I.e., the n-cliques form $(n-1)$-simplices with weights equal to the minimum of that of its edges. Intuitively, p_1 and p_2 will lie similarly distanced from the circle. **b.** Left, applying PH to the point cloud in **a**, using the Euclidean metric, reveals a short-lived circle (H^1-bar) and three long-lived clusters (H^0-bars). Black arrows indicate the homology of a circle. The two additional clusters represent p_1 and p_2. The dashed line indicates the radius of the Rips complex shown to the right of the barcode. Right, the resulting barcode from applying PH to the UMAP complex (i.e., using $-\log \mu_i$ as geodesic distances). The circle never closes, as seen in the corresponding Rips complex. **c.** Left, each column shows (from left): a noisy sample from a three-torus, $T^3 = S^1 \times S^1 \times S^1$; 3-D UMAP projection of the 6-D embedding and 3-D embedding of the circular coordinates obtained from UMAPH. Each row is colored by the three sampled angles in the two first columns and by the decoded angles in the third. Right, barcode from applying UMAPH to the sampled 3-torus, with long bars implying the homology of a 3-torus (one H^0-bar, three H^1-bars, three H^2-bars and one H^3-bar, as indicated by arrows).

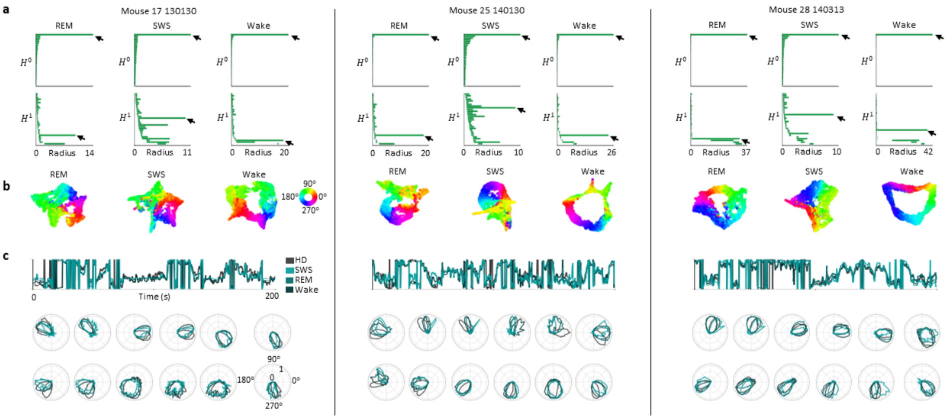

Extended Data Fig. 6 Ring topology of electrophysiological data in ADn and Subiculum representing internal head direction during OF foraging, REM sleep and SWS. **a.** Barcode diagrams from applying UMAPH in recordings of three mice ($n = 18, 25$ and 41 neurons). Arrows point to bars indicating ring topology. Each block (separated by black lines) contains the results from a single mouse. **b.** UMAP projection of firing rate activity colored by decoded circular coordinates (describing the circular features captured by the longest bar in the respective barcodes). **c.** Top, 200 s snippet of decoded circular coordinates and the recorded head direction. Note, the longest-lived circular feature found in each brain state was used to decode awake data. Bottom, single-cell tuning to the decoded coordinates and awake head direction recording for twelve example cells per animal (now, decoding each brain state separately).

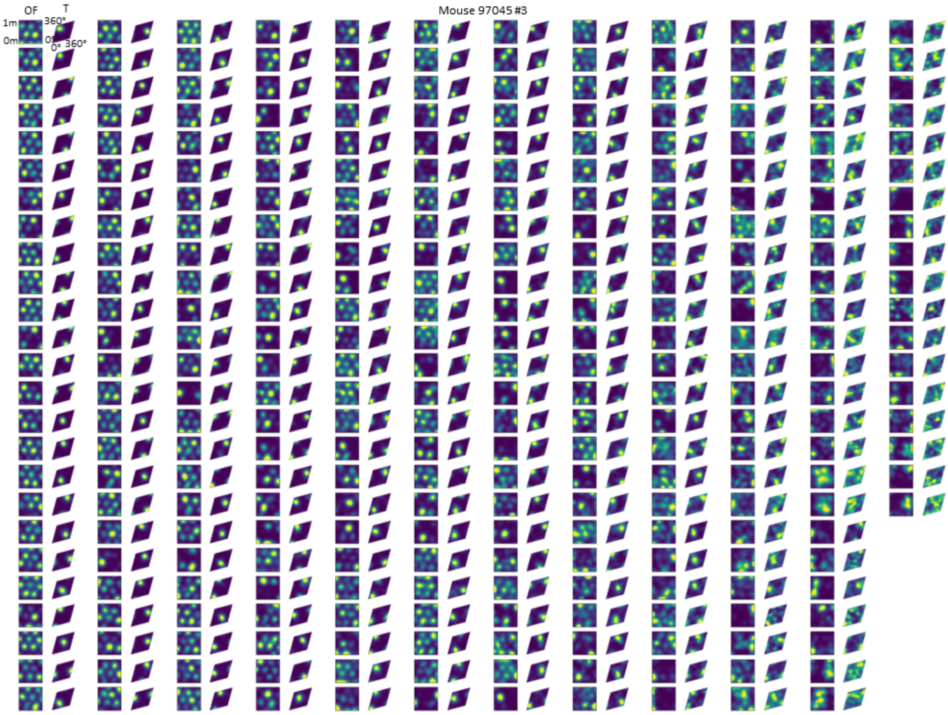

Extended Data Fig. 7 Tuning to coordinates in space and on the inferred torus for all neurons of mouse 97045 (as in Fig. 1e), recording day 2 (293 neurons) sorted top-bottom and left-right according to toroidal explained deviance (Fig. 1f).

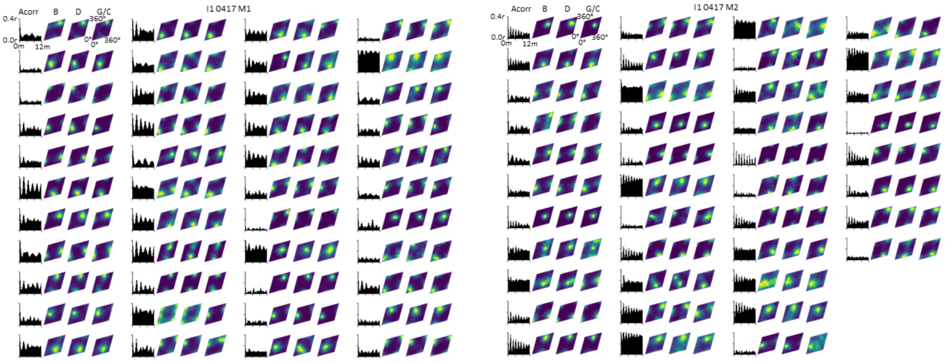

Extended Data Fig. 8 Tuning to coordinates in space and on the inferred torus (as in Fig. 2e) for all neurons of mouse and recording day (and module): I1 0417 (M1, 44 neurons) and (M2, 41 neurons). Each row of four shows 1-D VR-track autocorrelogram during dark session (left) and toroidal firing rate maps for each experimental condition (right).

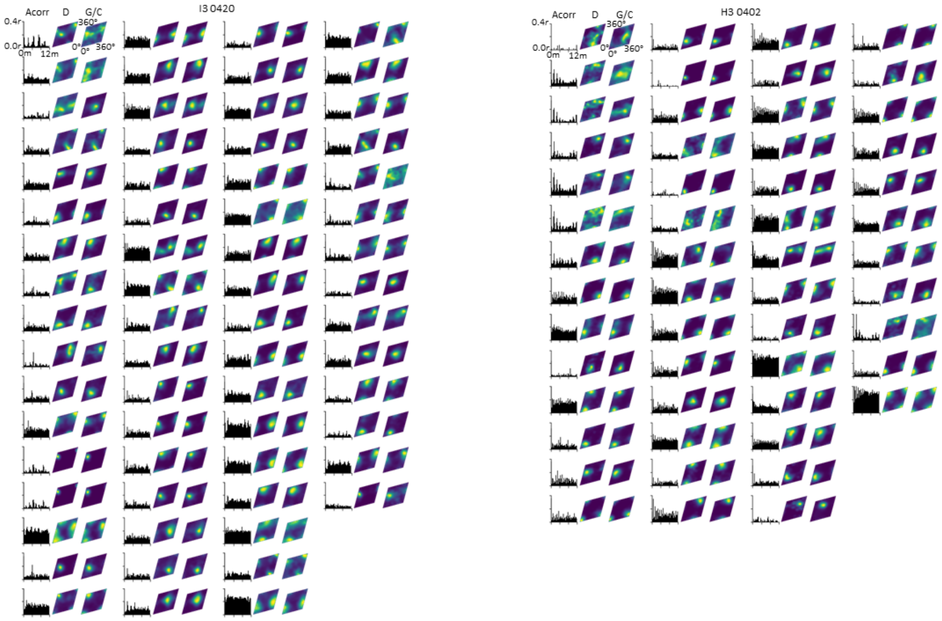

Extended Data Fig. 9 Tuning to coordinates in space and on the inferred torus for all neurons of mouse and experimental day: I3 0420 (65 neurons) and H3 0402 (53 neurons). Plots from left to right: 1-D VR-track autocorrelogram during dark session, toroidal firing rate map for dark (D) and gain/contrast (G/C) sessions.

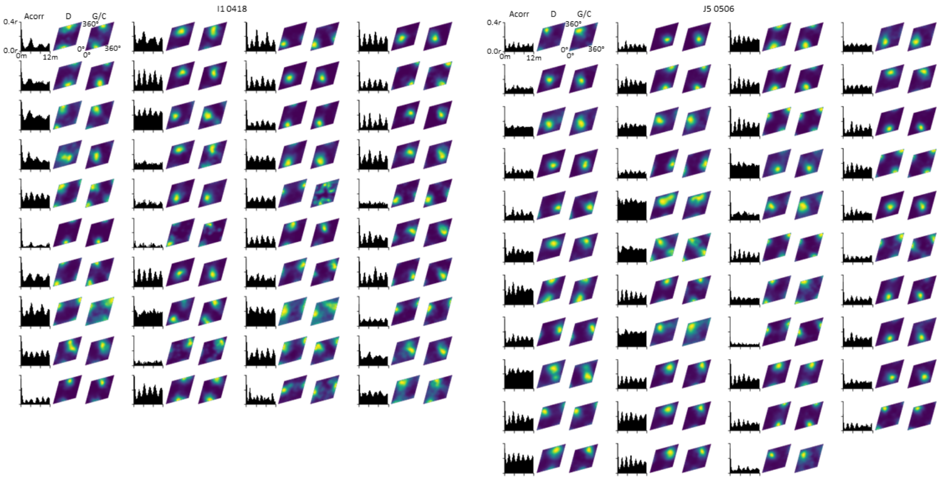

Extended Data Fig. 10 Tuning to coordinates in space and on the inferred torus for all neurons of mouse and experimental day: I1 0418 (40 neurons) and J5 0506 (47 neurons). Plots from left to right: 1-D VR-track autocorrelogram during dark session, toroidal firing rate map for dark (D) and gain/contrast (G/C) sessions.

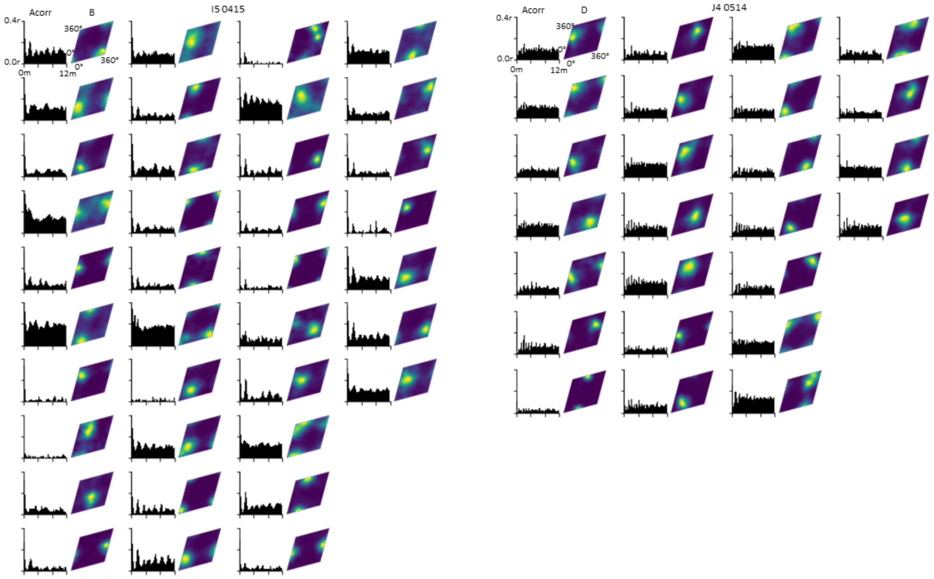

Extended Data Fig. 11 Tuning to coordinates in space and on the inferred torus for all neurons of mouse and experimental day: I5 0415 (37 neurons) and J4 0514 (25 neurons). Plots from left to right: 1-D VR-track autocorrelogram during dark session, toroidal firing rate map for baseline (B) or dark (D).

Paper II

Paper III

R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser and E. I. Moser, ‘Toroidal topology of population activity in grid cells,’ *Nature*, vol. 602, no. 7895, pp. 123–128, 2022

Article

Toroidal topology of population activity in grid cells

https://doi.org/10.1038/s41586-021-04268-7

Received: 24 February 2021

Accepted: 19 November 2021

Richard J. Gardner[1,6]✉, Erik Hermansen[2,6], Marius Pachitariu[3], Yoram Burak[4,5], Nils A. Baas[2]✉, Benjamin A. Dunn[1,2]✉, May-Britt Moser[1] & Edvard I. Moser[1]✉

The medial entorhinal cortex is part of a neural system for mapping the position of an individual within a physical environment[1]. Grid cells, a key component of this system, fire in a characteristic hexagonal pattern of locations[2], and are organized in modules[3] that collectively form a population code for the animal's allocentric position[1]. The invariance of the correlation structure of this population code across environments[4,5] and behavioural states[6,7], independent of specific sensory inputs, has pointed to intrinsic, recurrently connected continuous attractor networks (CANs) as a possible substrate of the grid pattern[1,8–11]. However, whether grid cell networks show continuous attractor dynamics, and how they interface with inputs from the environment, has remained unclear owing to the small samples of cells obtained so far. Here, using simultaneous recordings from many hundreds of grid cells and subsequent topological data analysis, we show that the joint activity of grid cells from an individual module resides on a toroidal manifold, as expected in a two-dimensional CAN. Positions on the torus correspond to positions of the moving animal in the environment. Individual cells are preferentially active at singular positions on the torus. Their positions are maintained between environments and from wakefulness to sleep, as predicted by CAN models for grid cells but not by alternative feedforward models[12]. This demonstration of network dynamics on a toroidal manifold provides a population-level visualization of CAN dynamics in grid cells.

The idea of a CAN has become one of the most influential concepts in theoretical systems neuroscience[13–15]. A CAN is a network in which recurrent synaptic connectivity constrains the joint activity of cells to a continuous low-dimensional repertoire of possible coactivation patterns in the presence of a wide range of external inputs. Few systems are more suitable for analysis of CAN dynamics than the spatial mapping circuits of the rodent brain, owing to the continuous, low-dimensional nature of space, and the availability and interpretability of data from these circuits[1–6]. In medial entorhinal cortex (MEC) and surrounding areas, head direction cells[16] encode orientation whereas grid cells[2] encode position. CAN models conceptualize the neural representations of these variables as spanning periodic one- or two-dimensional (1D or 2D) continua on a ring[17–19] or a torus[1,8–11], respectively. In this scheme, activity within the neural network stabilizes as a localized bump when cells are ordered according to their preferred firing directions or locations in physical space. The activity bump may be smoothly translated along the network continuum by speed and direction inputs, or by external sensory cues.

In agreement with CAN models[1,8–11], head direction cells[16,20,21] and modules of grid cells[4–7] maintain fixed correlation structures. In head direction cells, cell samples of a few dozen have been sufficient to demonstrate that the network activity traverses a ring[22–24], but for grid cells, the number of possible locations in the two-dimensional state space has been too large for the topology of the manifold to be uncovered. Here we take advantage of recently developed high-site-count Neuropixels silicon probes[25,26] to determine in many hundreds of simultaneously recorded grid cells whether, as predicted by two-dimensional CAN models[8–11], the population activity in an individual grid-cell module resides on a toroidal manifold, independently of behavioural tasks and states and decoupled from the position of the animal in physical space. We focused on individual modules because (i) these are the unit networks of CAN models[1,8–10]; and (ii) topological analysis of multi-module representations would require even larger numbers of cells[27].

Visualization of toroidal manifold

We recorded extracellular spikes of a total of 7,671 single units in layers II and III of the MEC–parasubiculum region in freely moving rats with unilateral or bilateral implants (total of 4 recordings, in 2 rats with bilateral single-shank probes and 1 rat with a unilateral 4-shank probe; from 546 to 2,571 cells per recording; Extended Data Fig. 1). During recordings, the rats were engaged in foraging behaviour in a square open-field (OF) enclosure or on an elevated track, or they slept in a small resting box. Using a clustering-based approach, we identified six grid modules across all rats (4 recording sessions, from 140 to 544

---

[1]Kavli Institute for Systems Neuroscience and Centre for Neural Computation, Norwegian University of Science and Technology, Trondheim, Norway. [2]Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway. [3]HHMI Janelia Research Campus, Ashburn, VA, USA. [4]Edmond and Lily Safra Center for Brain Sciences, The Hebrew University of Jerusalem, Jerusalem, Israel. [5]Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem, Israel. [6]These authors contributed equally: Richard J. Gardner, Erik Hermansen. ✉e-mail: richard.gardner@ntnu.no; nils.baas@ntnu.no; benjamin.dunn@ntnu.no; edvard.moser@ntnu.no

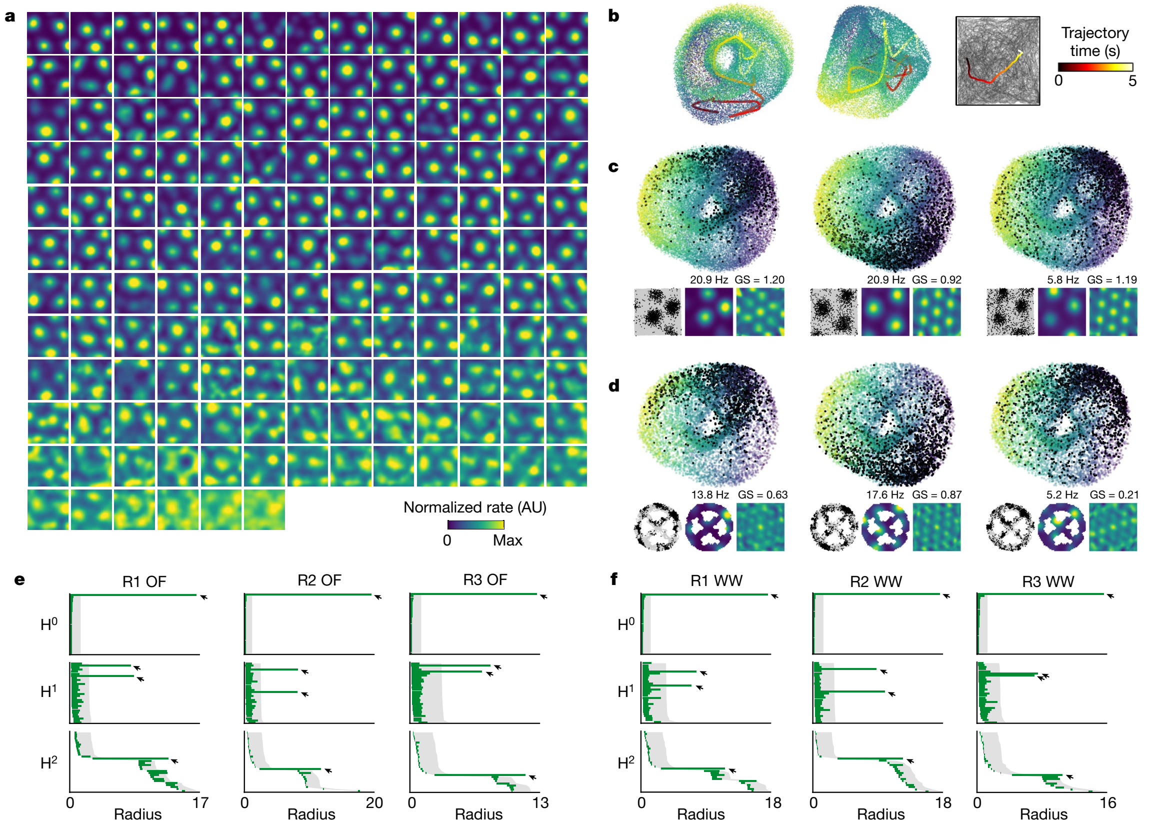

Fig. 1 | Signatures of toroidal structure in the activity of a module of grid cells. a, Firing rates of 149 grid cells co-recorded from the same module and shown, in order of spatial information content, as a function of rat position in OF arena (rates colour-coded, max 0.2–35.0 Hz; rat 'R' day 1, module 2; Extended Data Fig. 2b). **b**, Nonlinear dimensionality reduction reveals torus-like structure in the population activity of a single grid module (same 149 cells; 3 different views of same point cloud). Each dot represents the population state at one time point (dots coloured by first principal component). Bold line shows a 5-s trajectory, demonstrating smooth movement over the toroidal manifold. Right, corresponding trajectory in OF. **c**, Toroidal positions of spikes from three grid cells from the module in **a**. Each panel shows the same 3D point cloud of population states as in **b**, with black dots indicating when the cell fired. Insets show: left: the cell's 2D firing locations in OF (black dots on grey trajectory); middle: colour-coded firing rate map in OF (range 0 to max); right: colour-coded autocorrelogram of the rate map (range −1 to +1). Maximum rate and grid score (GS) are indicated. **d**, Same as in **c** (same cells) but with the rat running on an elevated, wheel-shaped track ('wagon-wheel track'; WW). Note preserved toroidal field locations. **e**, **f**, Barcodes indicate toroidal topology of grid-cell population activity. Results of persistent cohomology analyses (30 longest bars in the first three dimensions: H^0, H^1 and H^2) are shown for three grid modules from one rat (R1–R3 day 1, $n = 93$, 149 and 145 cells, respectively), in OF (**e**) and WW (**f**). Grey shading indicates longest lifetimes among 1,000 iterations in shuffled data (aligned to lower values of original bars). Arrows show four most prominent bars across all dimensions (all longer than in shuffled data). One prominent bar in dimension 0, two in dimension 1 and one in dimension 2 indicates cohomology equal to that of a torus.

grid cells per session; 7.8% to 25.6% of total number of cells; Extended Data Fig. 2a–d, g, h). Each grid module cluster contained a mixture of nondirectional ('pure') grid cells and conjunctive grid × direction cells[28], from 66 to 189 grid cells per module (total pure and conjunctive grid cells; Extended Data Fig. 2g). We initially limited our analyses to the subset of pure grid cells because (i) the expected toroidal topology might be distorted by additional directional modulation; and (ii) detection of topology in conjunctive cells may require a larger number of cells than recorded here[27].

To visually inspect the structure of the population activity of grid cells for signatures of toroidal topology, we constructed a three dimensional (3D) embedding of the n-dimensional population activity of a module of $n = 149$ pure grid cells (Fig. 1a). For this, we applied a two-stage dimensionality reduction procedure on the matrix of firing rates. First, to improve robustness to noise, we conducted a principal component analysis (PCA). We retained the first six principal components, which explained a particularly large fraction of the variance for all grid modules in the OF condition (with a similar tendency seen during sleep; Extended Data Fig. 4a). Next, we applied uniform manifold approximation and projection (UMAP) to reduce the six principal components into a 3D visualization. This visualization revealed a torus-like structure (Fig. 1b, Supplementary Video 1). Movement of the rat in the OF was accompanied by similarly continuous movement of the population activity across the toroidal manifold (Fig. 1b). When the activity of individual cells was plotted with reference to the 3D population representation, spikes for each cell were localized within a single patch of the population state space (Fig. 1c). The offsets between the firing locations of individual cells in the arena corresponded with the relative firing locations of the cells in the toroidal state space.

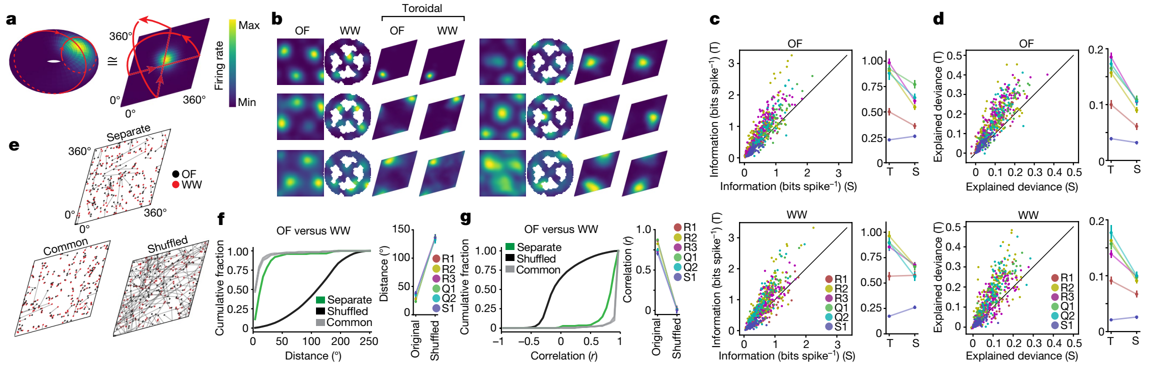

Fig. 2 | Cohomological decoding of position on an inferred state space torus. a, **b**, Individual grid cells have distinct firing fields on the inferred torus (Extended Data Fig. 5). Toroidal coordinates for population activity vectors were decoded from the two significant 1D holes (red circles in **a**) in the barcodes in Fig 1e, f. **a**, Left, 3D embedding of the toroidal state space displaying colour-coded mean firing rate of one grid cell as a function of toroidal position. Right, a 2D torus may be formed by gluing opposite sides of a rhombus. **b**, Representative grid cells from module R2 day 1 showing tuning to toroidal coordinates (all R2 cells: Supplementary Fig. 1). Each row of four plots corresponds to one cell. Left to right, colour-coded maps of cells' firing rates across the environment (OF or WW) and on the inferred torus (toroidal OF, toroidal WW, aligned to common axes). **c**, **d**, Toroidal information content (**c**) and explained deviance (**d**) for toroidal position (T) versus spatial position (S) in OF (top) and WW (bottom). Explained deviance is an R^2-statistic (range 0–1) expressing goodness-of-fit of GLM models for S or T. Left, scatterplots with dots showing individual cells; colour indicates module (inset). Right, mean ± s.e.m. for each module. n = 93 (R1), 149 (R2), 145 (R3), 94 (Q1), 65 (Q2) and 73 (S1) cells. **e**, **f**, Distances between toroidal firing field locations. **e**, Field locations of all R2 cells in OF and WW. Lines connect fields of the same cell. Toroidal OF and WW axes were aligned either separately ('separate') or commonly to OF ('common'). **f**, Left, cumulative frequency distribution of field distances (all R2 cells; green curve, separate alignment; grey lines, common alignment (to either OF or WW); black curve, shuffled data, n = 1,000 shuffles). Right, mean distance between field centres (±s.e.m.) for all modules. n cells as in **c**, **d**. **g**, Same as **f**, but showing Pearson correlations between pairs of toroidal rate maps.

Quantification of toroidal topology

Although the UMAP projection allowed a toroidal point cloud to be visualized, the method does not lend itself to straightforward quantification of the topology of the state space or comparison of representations across experiments. We therefore turned to the framework of persistent cohomology, a toolset from topological data analysis in which the structure of neural data can be classified by identifying holes of varying dimensionality in topological spaces assigned to point clouds of the cells' firing rates[22,23]. In applying this toolset, we replace each point of the point cloud by a ball of common radius. The union of balls results in a topological space in which the number of holes of different dimensions can be counted. By increasing this radius from zero until all the balls intersect, we observe the lifetime of each hole—the range of radii from when the hole first appears until it disappears (see Extended Data Fig. 3C). The lifetimes of the holes are depicted as bars and the totality of bars referred to as the barcode. For a torus, the barcode must display four bars of substantial length: a 0D hole (a single component connecting all points); two 1D holes (describing circular features); and a 2D hole (a cavity; Extended Data Fig. 3B).

Persistent cohomology analyses allowed us to classify the shape of the six-dimensional representation that serves as an intermediate step in UMAP (Extended Data Fig. 3A). We constructed barcodes for each of the six individual modules of grid cells recorded in the OF arena (three modules from rat 'R', 2 from rat 'Q' and 1 from rat 'S', henceforth named R1, R2, R3, Q1, Q2 and S1). The barcodes showed clear indications of toroidal characteristics. For all six modules, we detected four long-lived bars representing a single 0D hole, two 1D holes and a 2D hole. Their lifetimes were significantly longer than the lifetime of any bar obtained in 1,000 shuffles of the data in which spike times were randomly rotated (Fig. 1e, f, Extended Data Fig. 6Aa; $P < 0.001$). The findings suggest that network dynamics during OF foraging resides on a low-dimensional manifold with the same barcode as a torus. We noted the appearance of additional short bars in the barcodes for all modules, but these are expected for toroidal point clouds[27], as we confirmed with simulated data from several CAN models[10,11] and point clouds sampled from idealized tori, which in each case exhibited similar features (see Extended Data Fig. 7).

Tori persist despite grid distortions

The appearance of a torus in the point cloud, and the mapping of the activity of individual grid cells onto the torus (Fig. 1c), are consistent with a relationship between position in 2D physical space and position in the dimensionality-reduced neural state space. However, in many environments, this relationship may not be isometric, as the grid pattern is distorted by geometrical features of the environment, such as walls and corners[29–31] or discrete landmarks and reward locations[32,33]. We thus asked whether such geometric features could similarly distort the toroidal organization of network activity in the point cloud. We tested rats on an elevated running track shaped like a wagon wheel with four radial spokes ('wagon-wheel track' (WW); Fig. 1d, f). Spatial autocorrelation analyses confirmed that the strict periodicity of the grid pattern was compromised in this task (Extended Data Fig. 2e, f). Despite these distortions of the grid pattern in individual cells, toroidal tuning was maintained in the transformed population activity (Fig. 1d). The persistent cohomology analysis continued to identify one 0D hole, two 1D holes and one 2D hole with lifetimes that substantially exceeded those of shuffled data (Fig. 1f, Extended Data Fig. 6Ab). We also determined how the neural population activity mapped onto the torus by calculating angular coordinates from each of the two 1D holes identified by the barcode ('cohomological decoding'; Extended Data Fig. 5). The two angular coordinates defined directions intersecting at 60°, identifiable as a twisted torus (Fig. 2a). Consistent with CAN models, the vast majority of grid cells were tuned to a single location on the torus in each module and across environments, independent of geometry and local landmarks (Fig. 2b, Extended Data Fig. 4f, Supplementary Information).

To test how faithfully location in the environment is mapped onto the toroidal representation, we next asked whether grid-cell activity is predicted better by the cells' tuning to the inferred torus than by their

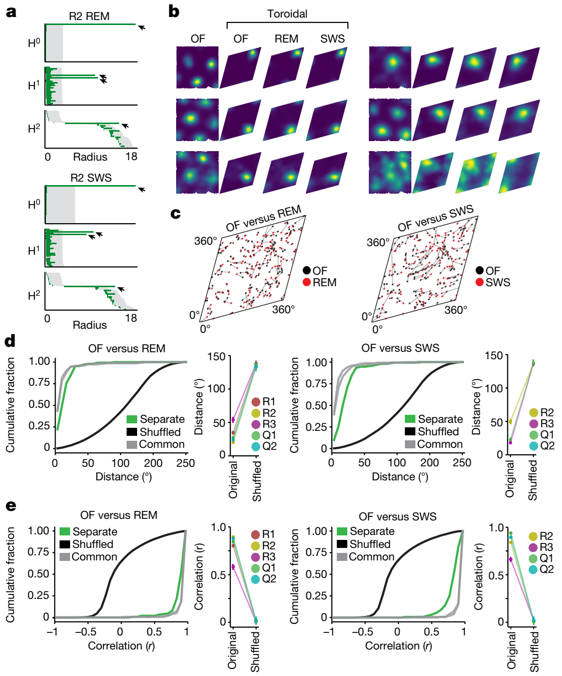

Fig. 3 | Preservation of toroidal structure during sleep. **a**, Barcodes indicating toroidal topology for grid-cell module R2 day 2 ($n = 152$ cells) during REM sleep and SWS (as in Fig. 1e, f). **b**, Toroidal rate maps showing preserved toroidal tuning for individual cells across environments and brain states (as in Fig. 2b; all cells shown in Extended Data Fig. 10). From left: rate map for OF in physical coordinates; and rate maps for OF, REM sleep and SWS in toroidal coordinates. **c**, Distribution of toroidal field centres (as in Fig. 2e) in OF and sleep (n as in a). **d**, **e**, Left, cumulative distributions of distances between toroidal field centres (**d**) and Pearson correlation r values (**e**) of rate maps for all R2 grid cells, as in Fig. 2f, g, but comparing OF with REM or SWS. Right, mean value ± s.e.m. for all modules. $n = 111$ (R1), 152 (R2), 165 (R3), 94 (Q1), 65 (Q2) and 72 (S1) cells. $n = 1{,}000$ shuffles.

tuning to physical space. For five out of six grid modules in OF and four out of six in WW, the information content conveyed about position, in bits per spike, was higher for position on the torus than for position in physical space (Fig. 2c; R2, R3, Q1, Q2: all $P < 0.001$, $W > 1{,}932$ in OF and WW; R1: $P < 0.001$, $W = 4{,}010$ in OF, $P = 0.586$, $W = 2{,}129$ in WW; S1: $P = 1.000$ in OF and WW, $W = 620$ in OF, $W = 129$ in WW; Wilcoxon signed-rank test). We verified this difference by comparing the cross-validated prediction of two Poisson generalized linear model (GLM)-based encoding models of each cell's activity that included toroidal position (decoded as above) and 2D spatial position. For both environments (OF and WW), the toroidal covariate was closer to a perfectly fitted model of the data than was the physical position covariate in five out of six grid-cell modules (Fig. 2d; R1, R2, R3, Q1, Q2: $P < 0.001$, $W > 2{,}045$ in OF and WW; S1: $P < 0.001$, $W = 1{,}941$ in OF, $P = 1.000$, $W = 727$ in WW; Wilcoxon signed-rank test). Together, these differences point to toroidal structure as the primary feature of the population activity of grid cells, superior to that of the 2D coordinates of the animal's current position in the physical environment.

If grid cells operate on a toroidal manifold determined by intrinsic network features, this manifold may be expressed universally across environments, independently of sensory inputs. We tested this proposition by assessing, on the inferred tori, whether the locations of firing fields of different grid cells were maintained between OF and WW (Fig. 2b, Supplementary Information). To compare the toroidal parametrizations, we aligned the axes of the toroidal coordinates (Extended Data Fig. 5b). First, we compared, for each cell, the distance between the centres of mass of the toroidal rate maps in OF and WW (Fig. 2e, f, Extended Data Fig. 6Ba). This distance was substantially shorter (mean ± s.e.m. of mean distances for all modules: 31.5 ± 6.3 degrees) than that of control data in which the order of the rate maps in one environment was shuffled (135.8 ± 1.7 degrees; maximum possible distance $\sqrt{2}\cdot 180 \approx 254.6$ degrees; data versus shuffled: $P < 0.001$ in all modules). Second, we calculated the pairwise Pearson correlations of binned toroidal rate maps across the two environments (Fig. 2g, Extended Data Fig. 6Ba). Consistent with the centre-of-mass comparison, the correlations between OF and WW were higher in observed data (mean ± s.e.m. of mean r values for all modules: 0.79 ± 0.07) than in shuffled data ($r = 0.01 \pm 0.01$; $P < 0.001$ for all modules). Very similar results were obtained when applying the toroidal parametrization from the same environment (either OF or WW) to activity from both environments (Fig. 2f, g, 16.0 ± 3.4 degrees; $r = 0.95 \pm 0.02$; $P < 0.001$ for all modules and both mappings). Together, these findings suggest that physical space is mapped onto the same internal low-dimensional manifold irrespective of the specific environment.

Toroidal topology persists during sleep

If population activity is mapped onto the same toroidal manifold independently of sensory inputs, the toroidal topology should also be maintained during sleep. To test this idea, the rats rested in a high-walled, opaque box placed in the centre of the OF or WW track. Periods of rapid-eye-movement (REM) sleep and slow-wave sleep (SWS) were identified on the basis of the low-frequency rhythmic content of the aggregated multi-unit activity in combination with prolonged behavioural immobility (Extended Data Fig. 9).

Persistent cohomology analysis of the sleep population activity suggested toroidal topology in five of the six grid modules during REM and four out of six modules during SWS (modules R2, R3, Q1 and Q2 for both sleep stages and module R1 only in REM; Fig. 3a, Extended Data Fig. 6Ac, d). In the remaining module (S1), there were no long-lived bars in dimensions 1 or 2 (Extended Data Fig. 6Ac, d), indicating an absence of toroidal structure during sleep, perhaps because of an insufficient number of cells in this module (72 cells; Extended Data Fig. 4e). The barcode results were supported by the toroidal mapping, which revealed sharply tuned firing fields on the REM and SWS tori (99.3 ± 1.6% and 99.1 ± 1.8%, respectively, of the grid cells in each module had higher information content than shuffled data, and in 95.3 ± 7.2% and 98.6 ± 2.4% of cells the toroidal tuning explained the activity better than a null model that assumes a constant firing rate; Fig. 3b, Extended Data Figs. 6C, 10, Supplementary Information). In addition, the spatial arrangements of toroidal firing locations of different cells were maintained between wake, REM and SWS states (Fig. 3c, Extended Data Fig. 6Bb, c). For between-condition pairs of rate maps, the mean distance (±s.e.m.) between the peak firing locations (OF versus REM 31.5 ± 15.4 degrees, OF versus SWS 29.8 ± 14.3 degrees) was well below the distribution of shuffled distances (Fig. 3d, Extended Data Fig. 6Bb, c; 135.8 ± 2.3 degrees in both REM and SWS, $P < 0.001$ for all 5 and 4 modules, respectively). Similarly, the mean correlations of pairs of toroidal rate maps (REM versus OF $r = 0.80 \pm 0.15$, SWS versus OF $r = 0.83 \pm 0.12$) were substantially larger than in shuffled versions of the data (Fig. 3e, Extended Data Fig. 6Bb, c; $r = 0.01 \pm 0.01$ in both REM and SWS, $P < 0.001$ for all 5 and 4 modules, respectively). Thus, the toroidal structure is maintained in both sleep conditions, despite the lack of external spatial inputs.

Classes of grid cells

We next investigated why toroidal structure was not visible during REM in module S1 and during SWS in modules R1 and S1 (Fig. 4a, Extended

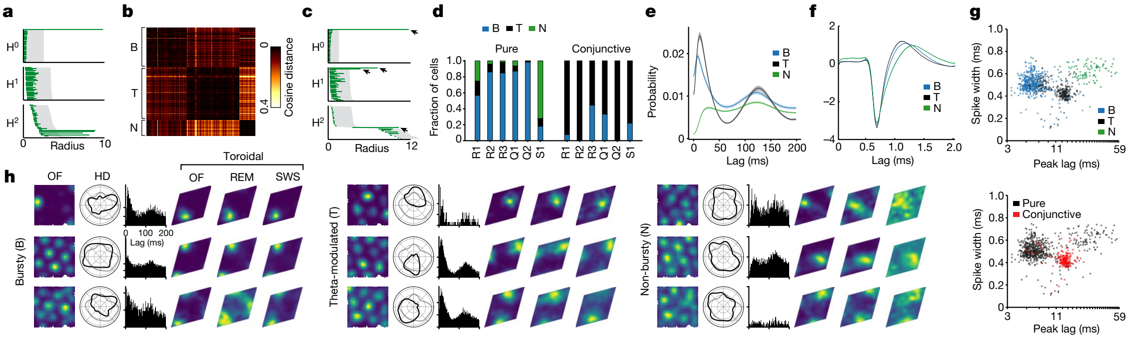

Fig. 4 | Differential toroidal tuning of grid-cell subpopulations. a, Barcode of all pure R1 grid cells (day 2, $n = 111$ cells) does not indicate toroidal structure during SWS. **b**, Matrix of cosine distances between pairs of spike-train autocorrelograms of grid cells in module R1. Rows and columns show 189 grid cells (pure and conjunctive) sorted by cluster identity. Three clusters were identified, appearing as dark (that is, similar) squares along the matrix diagonal. On the basis of temporal firing patterns (**e**), they were named 'bursty' (B), 'theta-modulated' (T) and 'non-bursty' (N). **c**, Barcode of the 'bursty' class of R1 ($n = 69$ cells) indicates toroidal structure. Symbols as in **a**. Arrows point to the four most persistent features. **d**, Fractions of grid cells in each class, shown for each grid module. Left, pure grid cells only, right, conjunctive grid × head-direction cells only. For n see Extended Data Fig. 2g. **e**, Average temporal autocorrelogram for cells in each class. Shaded area shows mean ± s.e.m. (bursty $n = 523$, theta-modulated $n = 229$, non-bursty $n = 95$ cells). For each class, note short-latency peak (burst-firing) and long-latency peak (theta-modulation). **f**, Average spike template waveforms of cells from each class (n as in **e**). Shaded area indicates mean ± s.e.m. **g**, Cell classes have different burst-firing characteristics, as expressed by latency of first autocorrelogram peak (x axis) and peak-to-peak spike width (y axis). Cells (dots) are colour-coded by class (n as in **e**) or by identity (pure or conjunctive, $n = 659$ or 188 cells, respectively). **h**, Example cells from each class (one row of plots per cell). Plots from left to right: OF firing rate map; head-direction (HD) tuning curve (black) compared to occupancy of head directions (light grey); temporal autocorrelogram; toroidal firing rate maps for OF, REM and SWS.

Data Fig. 6Ad). Previous studies of medial entorhinal spiking activity have described cell populations with distinct burst-firing and theta-modulation characteristics[34–36]; therefore, we asked whether a lack of toroidal structure was due to heterogeneity in the composition of the module. We quantified each cell's temporal modulation characteristics using the spike train temporal autocorrelogram from the OF session, and by applying clustering to the matrix of autocorrelograms we obtained three cell classes (Fig. 4b). Each class was distributed across multiple modules (Fig. 4d). Within each module, cells from the three classes showed overlapping grid spacing and orientation properties (Extended Data Fig. 8a). We named the classes 'bursty' (B), 'non-bursty' (N) and 'theta-modulated' (T), following the most prominent autocorrelogram feature of each class (Fig. 4e). We also examined the spike waveforms of the cells, and found that each class showed a characteristic spike width (Fig. 4f, g), suggesting that they differ in morphology or biophysical properties.

The firing rates of the cells during SWS exhibited marked correlation structure within—but not between—classes (Extended Data Fig. 8b). Even though our classification strategy was not influenced by the cells' directional tuning, class T contained 80% of all conjunctive grid cells and only 11% of all pure grid cells, supporting the idea that conjunctive grid cells are a distinct population. Accordingly, in modules R1 and S1, which contained the largest numbers of T cells, pairwise correlations of T cells' spike trains were more strongly related to head-direction tuning than to toroidal tuning (Extended Data Fig. 8c). When we subdivided module R1 into the three classes (Fig. 4b), we found that during SWS toroidal topology was detectable only in B cells (Fig. 4c). By decoding toroidal position from B cells, we were able to recover the selectivity of each cell with respect to toroidal position in module R1 (Fig. 4h). The toroidal tuning locations were preserved between OF and SWS in each cell class in R1 (Extended Data Fig. 8d, B: distance of 26.4 ± 6.1 degrees and correlation of $r = 0.85 \pm 0.02$, T: 43.6 ± 3.9 degrees and $r = 0.74 \pm 0.02$, N: 29.9 ± 3.5 degrees and $r = 0.80 \pm 0.02$; mean values of shuffled versions of each class were between 135.4 ± 5.2 and 136.4 ± 6.2 degrees, and between $r = 0.00 \pm 0.07$ and $r = 0.02 \pm 0.03$; comparison between observed and shuffled $P < 0.001$ for all 3 classes and both measures). However, in R1 as well as all other modules, toroidal spatial information and explained deviance were highest for B cells and lower for N and T cells in OF, REM and SWS (Extended Data Fig. 8e) (information content: $P < 10^{-56}$, $H > 255$; Kruskal–Wallis test; $P < 10^{-9}$, $Z > 6.4$; Dunn test with Bonferroni correction; explained deviance: $P < 10^{-20}$, $H > 96$; Kruskal–Wallis test; $P < 10^{-12}$, $Z > 7.4$; Dunn test with Bonferroni correction, for OF, REM and SWS). Collectively, these results show that the B cell population (containing the majority of our grid cells) represents the torus most robustly across behavioural conditions. The weaker toroidal representation in T cells may partly be an effect of the higher dimensionality of the code carried by conjunctive grid × direction cells. Indeed, running cohomology analysis on T cells from modules S1 and R1 (which contained the most T cells) revealed a circular feature that corresponded to the animal's head direction (Extended Data Fig. 8f, g).

Discussion

Our findings, from many hundreds of simultaneously recorded grid cells, show that population activity in grid cells invariably spans a manifold with toroidal topology, with movement on the torus matching the animal's trajectory in the environment. The toroidal representation was most stably encoded by the bursting subclass of grid cells. Toroidal topology was not simply inherited from the encoded variable, as 2D space is not characterized by toroidal topology, as opposed to pitch and azimuth of head orientation, which in bats together span a torus and thus naturally map onto a toroidal neural code[37]. Using cohomological decoding, we were able to demonstrate, in each environment and in both sleep and awake states, that the toroidal coordinates of individual grid cells in individual grid modules were maintained, independently of external sensory inputs or environment-induced deformations of hexagonal symmetry in the rate maps[29–33]. The uniform and consistent toroidal structure of the manifold suggests that distortions in grid patterns occur in the mapping between physical space and the toroidal grid code rather than in the grid code itself.

The invariance of the toroidal manifold across environments and brain states is informative about the mechanisms that underlie grid-cell activity. Although toroidal topology can be generated by both CAN[1,8–10] and feedforward[12] mechanisms, the persistence of an invariant toroidal manifold under conditions that give rise to changes in the correlation structure of place-cell activity in the hippocampus[6,7] is predicted only by CAN models. While the findings do not exclude co-existing feedforward mechanisms[12,38], they point to intrinsic network connectivity as the mechanism that underlies the rigid toroidal dynamics of the grid-cell system. What kind of network architecture keeps the activity on a toroidal manifold—whether it is geometrically organized[1,8–10] or acquired from random connectivity by synaptic weight adjustments through learning[39–41]—remains to be determined, as does the mode of connectivity with other CANs in the entorhinal–hippocampal system[22,23].

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41586-021-04268-7.

1. McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I. & Moser, M.-B. Path integration and the neural basis of the 'cognitive map'. *Nat. Rev. Neurosci.* **7**, 663–678 (2006).
2. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).
3. Stensola, H. et al. The entorhinal grid map is discretized. *Nature* **492**, 72–78 (2012).
4. Fyhn, M., Hafting, T., Treves, A., Moser, M.-B. & Moser, E. I. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* **446**, 190–194 (2007).
5. Yoon, K. et al. Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat. Neurosci.* **16**, 1077–1084 (2013).
6. Gardner, R. J., Lu, L., Wernle, T., Moser, M.-B. & Moser, E. I. Correlation structure of grid cells is preserved during sleep. *Nat. Neurosci.* **22**, 598–608 (2019).
7. Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R. & Colgin, L. L. Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nat. Neurosci.* **22**, 609–617 (2019).
8. Fuhs, M. C. & Touretzky, D. S. A spin glass model of path integration in rat medial entorhinal cortex. *J. Neurosci.* **26**, 4266–4276 (2006).
9. Burak, Y. & Fiete, I. R. Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput. Biol.* **5**, e1000291 (2009).
10. Guanella, A., Kiper, D. & Verschure, P. A model of grid cells based on a twisted torus topology. *Int. J. Neural Syst.* **17**, 231–240 (2007).
11. Couey, J. J. et al. Recurrent inhibitory circuitry as a mechanism for grid formation. *Nat. Neurosci.* **16**, 318–324 (2013).
12. Kropff, E. & Treves, A. The emergence of grid cells: intelligent design or just adaptation? *Hippocampus* **18**, 1256–1269 (2008).
13. Amari, S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**, 77–87 (1977).
14. Ben-Yishai, R., Bar-Or, R. L. & Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proc. Natl Acad. Sci. USA* **92**, 3844–3848 (1995).
15. Seung, H. S. How the brain keeps the eyes still. *Proc. Natl Acad. Sci. USA* **93**, 13339–13344 (1996).
16. Taube, J. S., Muller, R. U. & Ranck, J. B. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J. Neurosci.* **10**, 420–435 (1990).
17. Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S. & McNaughton, B. L. A model of the neural basis of the rat's sense of direction. *Adv. Neural Inf. Process. Syst.* **7**, 173–180 (1995).
18. Redish, A. D., Elga, A. N. & Touretzky, D. S. A coupled attractor model of the rodent head direction system. *Network* **7**, 671–685 (1996).
19. Zhang, K. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.* **16**, 2112–2126 (1996).
20. Yoganarasimha, D., Yu, X. & Knierim, J. J. Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. *J. Neurosci.* **26**, 622–631 (2006).
21. Peyrache, A., Lacroix, M. M., Petersen, P. C. & Buzsáki, G. Internally organized mechanisms of the head direction sense. *Nat. Neurosci.* **18**, 569–575 (2015).
22. Rybakken, E., Baas, N. & Dunn, B. Decoding of neural data using cohomological feature extraction. *Neural Comput.* **31**, 68–93 (2019).
23. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. & Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **22**, 1512–1520 (2019).
24. Rubin, A. et al. Revealing neural correlates of behavior without behavioral measurements. *Nat. Commun.* **10**, 4745 (2019).
25. Jun, J. J. et al. Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
26. Steinmetz, N. A. et al. Neuropixels 2.0: a miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**, eabf4588 (2021).
27. Kang, L., Xu, B. & Morozov, D. Evaluating state space discovery by persistent cohomology in the spatial representation system. *Front. Comput. Neurosci.* **15**, 616748 (2021).
28. Sargolini, F. et al. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* **312**, 758–762 (2006).
29. Barry, C., Hayman, R., Burgess, N. & Jeffery, K. J. Experience-dependent rescaling of entorhinal grids. *Nat. Neurosci.* **10**, 682–684 (2007).
30. Stensola, T., Stensola, H., Moser, M.-B. & Moser, E. I. Shearing-induced asymmetry in entorhinal grid cells. *Nature* **518**, 207–212 (2015).
31. Krupic, J., Bauza, M., Burton, S., Barry, C. & O'Keefe, J. Grid cell symmetry is shaped by environmental geometry. *Nature* **518**, 232–235 (2015).
32. Butler, W. N., Hardcastle, K. & Giocomo, L. M. Remembered reward locations restructure entorhinal spatial maps. *Science* **363**, 1447–1452 (2019).
33. Boccara, C. N., Nardin, M., Stella, F., O'Neill, J. & Csicsvari, J. The entorhinal cognitive map is attracted to goals. *Science* **363**, 1443–1447 (2019).
34. Latuske, P., Toader, O. & Allen, K. Interspike intervals reveal functionally distinct cell populations in the medial entorhinal cortex. *J. Neurosci.* **35**, 10963–10976 (2015).
35. Newman, E. L. & Hasselmo, M. E. Grid cell firing properties vary as a function of theta phase locking preferences in the rat medial entorhinal cortex. *Front. Syst. Neurosci.* **8**, 193 (2014).
36. Csordás, D. É., Fischer, C., Nagele, J., Stemmler, M. & Herz, A. V. M. Spike afterpotentials shape the in vivo burst activity of principal cells in medial entorhinal cortex. *J. Neurosci.* **40**, 4512–4524 (2020).
37. Finkelstein, A. et al. Three-dimensional head-direction coding in the bat brain. *Nature* **517**, 159–164 (2015).
38. Ginosar, G. et al. Locally ordered representation of 3D space in the entorhinal cortex. *Nature* **596**, 404–409 (2021).
39. Sussillo, D. & Barak, O. Opening the black box: low-dimensional dynamics in high-dimensional recurrent neural networks. *Neural Comput.* **25**, 626–649 (2013).
40. Sorscher, B., Mel, G. C., Ocko, S. A., Giocomo, L. & Ganguli, S. A unified theory for the computational and mechanistic origins of grid cells. Preprint at https://doi.org/10.1101/2020.12.29.424583 (2020).
41. Darshan, R. & Rivkind, A. Learning to represent continuous variables in heterogeneous neural networks. Preprint at https://doi.org/10.1101/2021.06.01.446635 (2021).

Methods

Rats

The data were collected from three experimentally naive male Long Evans rats (Rats Q, R and S, 300–500 g at time of implantation). The rats were group-housed with three to eight of their male littermates before surgery and were singly housed in large Plexiglas cages (45 × 44 × 30 cm) thereafter. They were kept on a 12-h light–12-h dark schedule, with strict control of humidity and temperature. All procedures were performed in accordance with the Norwegian Animal Welfare Act and the European Convention for the Protection of Vertebrate Animals used for Experimental and Other Scientific Purposes. Protocols were approved by the Norwegian Food Safety Authority (FOTS ID 18011 and 18013).

Electrode implantation and surgery

The rats were implanted with Neuropixels silicon probes[25,26] targeting the MEC–parasubiculum (PaS) region. Two rats were implanted bilaterally with prototype Neuropixels 'phase 3A' single-shank probes and with one probe targeting MEC–PaS in each hemisphere; the third rat was implanted with a prototype Neuropixels 2.0 multi-shank probe in the left hemisphere. Probes were inserted at an angle of 25° from posterior to anterior in the sagittal plane. Implantation coordinates were AP 0.05–0.3 mm anterior to the sinus and 4.2–4.7 mm lateral to the midline. The probes were inserted to a depth of 4,200–6,000 µm. The implant was secured with dental cement. The detailed implantation procedure has been described elsewhere[6,26]. After surgery, the rats were left to recover for approximately 3 h before beginning recording. Postoperative analgesia (meloxicam and buprenorphine) was administered during the surgical recovery period.

Recording procedures

The details of the Neuropixels hardware system and the procedures for freely moving recordings have been described previously. In brief, electrophysiological signals were amplified with a gain of 500 (for phase 3A probes) or 80 (for 2.0 probes), low-pass-filtered at 300 Hz (phase 3A) or 0.5 Hz (2.0), high-pass-filtered at 10 kHz, and then digitized at 30 kHz (all steps performed by the probe's on-board circuitry). The digitized signals were multiplexed by an implant-mounted 'headstage' circuit board and were transmitted along a lightweight 5-m tether cable, made using either micro-coaxial (phase 3A) or twisted pair (2.0) wiring.

Three-dimensional motion capture (OptiTrack Flex 13 cameras and Motive recording software) was used to track the rat's head position and orientation, by attaching a set of five retroreflective markers to implant during recordings. The 3D marker positions were projected onto the horizontal plane to yield the rat's 2D position and head direction. An Arduino microcontroller was used to generate digital pulses, which were sent to the Neuropixels acquisition system (via direct TTL input) and the OptiTrack system (via infra-red LEDs), to permit precise temporal alignment of the recorded data streams.

Behavioural procedures

Data were obtained from four recording sessions performed within the first 72 h after recovery from surgery. The recordings were performed while the rats engaged in three behavioural paradigms, each in a different arena within the same room. Abundant distal visual and sonic cues were available to the rat. On each day of recording, the rat remained continuously connected to the recording apparatus across the various behavioural sessions that were performed. Occasionally it was necessary to remove twists that had accumulated in the Neuropixels tether cable. In such cases, the ongoing behavioural task was paused while the experimenter gently turned the rat to remove the twists. During pre-surgical training, the rats were food-restricted, maintaining their weight at a minimum of 90% of their free-feeding body weight. Food was generally removed 12–18 h before each training session. Food restriction was not used at the time of recording.

Open-field foraging task

Rats foraged for randomly scattered food crumbs (corn puffs) in a square open-field (OF) arena of size 1.5 × 1.5 m, with black flooring and enclosed by walls of height 50 cm. A large white cue card was affixed to one of the arena walls (height same as the wall; width 41 cm; horizontal placement at the middle of the wall). At the time of the surgery, each rat was highly familiar with the environment and task (10–20 training sessions lasting at least 20 min each).

Wagon-wheel track foraging task

The wagon-wheel (WW) track task was designed to function as a 1D version of the 2D OF foraging task. The track's geometry comprised an elevated circular track with two perpendicular cross-linking arms spanning the circle's diameter. The track was 10 cm wide and was bounded on both sides by a 1-cm-high lip. Each section of the track was fitted with a reward point, placed halfway between the two nearest junctions, in the centre of the section. Each reward point consisted of an elevated well that could be remotely filled with chocolate milk via attached tubing. To encourage foraging behaviour, a pseudorandom subset of the wells (between one and four of the eight wells) was filled at a given time, and the rat was allowed to explore the full maze freely and continuously. Wells were refilled as necessary when the rat consumed rewards. Each rat was trained to high performance on the foraging task before the surgery (collecting at least 30 rewards within a 30-minute session). Training to this level of performance took 5–10 half-hour sessions.

Natural sleep

For sleep sessions, the rat was placed in a black acrylic 'sleep box' with a 40 × 40-cm square base and 80-cm-high walls. The black coating of walls was transparent to infrared, which allowed the 3D motion capture to track the rat through the walls. The bottom of the sleep box was lined with towels, and the rat had free access to water. During recording sessions in the sleep box, the main room light was switched on and pink noise was played through the computer speakers to attenuate disturbing background sounds. Sleep sessions typically lasted 2–3 h, but were aborted prematurely if the rat seemed highly alert and unlikely to sleep.

Spike sorting and single-unit selection

Spike sorting was performed with KiloSort 2.5[26]. In brief, the algorithm consists of three principal stages: (1) a raw-data alignment procedure that detects and corrects for shifts in the vertical position of the Neuropixels probe shank relative to the surrounding tissue; (2) an iterative template-matching procedure that uses low-rank, variable-amplitude waveform templates to extract and classify single-unit spikes; and (3) a curation procedure which detects appropriate template merging and splitting operations based on spike train auto- and cross-correlograms. Some customizations were made to the standard KiloSort 2.5 method to improve its performance on recordings from the MEC–PaS region, where there is a particularly high spatiotemporal overlap of spike waveforms owing to the high density of cells. Therefore, the maximum number of spikes extracted per batch in step 1 above was increased, as was the number of template-matching iterations in step 2. To improve the separation between cells with very similar-looking waveforms, the upper limit on template similarity was raised from 0.9 to 0.975 in step 2 and to 1.0 on step 3, while supervising manually all merge and split operations from step 3, using a custom-made GUI running in MATLAB. The manual supervision ensured that Kilosort 2.5 did not automatically merge pairs of units with a dip in the cross-correlogram, which in our data was often due to out-of-phase spatial tuning. The merge and split operations were repeated several times to ensure the best separation between single units.

Single units were discarded if more than 1% of their interspike interval distribution consisted of intervals less than 2 ms. In additions, units were excluded if they had a mean spike rate of less than 0.05 Hz or greater than 10 Hz (calculated across the full recording duration).

Single-unit spike waveforms

During spike sorting, Kilosort assigned each unit with a 2 ms spike waveform template on each recording channel. To calculate a representative single waveform for each unit, the peak-to-peak amplitude of the template was calculated on every channel, and the templates from the three highest-amplitude channels were averaged to generate the representative spike waveform. To calculate spike width, a unit's representative waveform was finely interpolated (from 61 to 1,000 points) using a cubic spline function. Spike width was defined as the time difference between the waveform's negative peak (to which the waveform was aligned by Kilosort), and the following positive peak.

Spatial position and direction tuning

During awake foraging sessions in the OF arena or wagon-wheel track, only time epochs in which the rat was moving at a speed above 2.5 cm s^{-1} were used for spatial or toroidal analyses. To generate 2D rate maps for the OF arena, position estimates were binned into a square grid of 3 × 3-cm bins. The spike rate in each position bin was calculated as the number of spikes recorded in the bin, divided by the time the rat spent in the bin. To interpolate the values of unvisited bins, two auxiliary matrices were used, M_1 and M_2, setting visited bins equal to the value of the original rate map in M_1 and to 1 in M_2, and setting unvisited bins to zero in both. One iteration of the image-processing 'closing' operation was then performed (binary dilation followed by erosion, filling out a subset of the non-visited bins) on $M_{2,}$ using a disk-shaped structuring element, first padding the matrix border by one bin. Both matrices were then spatially smoothed with a Gaussian kernel of smoothing width 2.75 bins. Finally, the rate map was obtained by dividing M_1 by M_2. Rate-map spatial autocorrelograms and grid scores were calculated as described previously[28]. The selectivity of each cell's position tuning was quantified by computing its spatial information content[42], measured in bits per spike (see 'Information content').

Head-direction tuning curves were calculated by binning the head-direction estimates into 6° bins. The spike rate in each angular bin was calculated as the number of spikes recorded in the bin divided by the time that the rat spent in the bin. The resultant tuning curve was smoothed with a Gaussian kernel with $\sigma = 2$ bins, with the ends of the tuning curve wrapped together. The selectivity of head-direction tuning was quantified using the mean vector length (MVL) of the tuning curve. This was calculated according to:

$$\mathrm{MVL} = \frac{|\sum_{j=1}^{M} \mathbf{f}_j \exp(i\boldsymbol{\alpha}_j)|}{\sum_{j=1}^{M} \mathbf{f}_j},$$

where vector $\mathbf{f}$ represents the tuning curve values (firing rates), vector $\boldsymbol{\alpha}$ represents the corresponding angles, M is the number of tuning curve values, and $|\cdot|$ represents the absolute value of the enclosed term.

Grid module classification

A novel method was implemented to detect populations of cells corresponding to grid modules by finding clusters of cells that expressed similar spatially periodic activity in the open field (Extended Data Fig. 2). Contrary to previous clustering-based methods for grid modules[3], this approach makes no assumptions about the specific geometry of the grid pattern, thus making it less susceptible to the detrimental effects of geometric distortions such as ellipticity[3,30].

For each MEC–PaS cell in a given recording, a coarse-resolution rate map of the OF session was constructed, using a grid of 10 × 10-cm bins, with no smoothing across bins. The 2D autocorrelogram of this rate map was calculated, and the central peak was removed by excluding all bins located less than 30 cm from the autocorrelogram centre. Bins located more than 100 cm from the autocorrelogram centre were also excluded. The autocorrelograms for all cells were subsequently converted into column vectors, z-standardized, then concatenated to form a matrix with spatial bins as rows and cells as columns. The nonlinear dimensionality reduction algorithm UMAP[43,44] was then applied to this matrix, yielding a two-dimensional point cloud in which each data point represented the autocorrelogram of one cell (Extended Data Fig. 2a–d; UMAP hyperparameters: 'metric'='manhattan', 'n_neighbors'=5, 'min_dist'=0.05, 'init'='spectral'). In the resultant 2D point cloud, cells with small absolute differences between their autocorrelogram values were located near to one another. The point cloud was partitioned into clusters using the DBSCAN clustering algorithm (MATLAB function 'dbscan', minimum 30 points per cluster, eta = 0.6–1.0). In every recording, the largest cluster was mainly composed of cells that either lacked strong spatial selectivity or were spatially selective but without clear periodicity. All remaining clusters contained cells with high grid scores, and with similar grid spacing and orientation (Extended Data Fig. 2a–d); cluster membership was therefore used as the basis for grid module classification. In one recording (rat 'R' day 1), two clusters were identified that had similar average grid spacing and orientation (labelled as 'R1a' and 'R1b' in Extended Data Fig. 2a–d), suggesting that they represented the same grid module. R1b appeared to comprise cells with higher variability in the within-field firing rates of the spatial rate maps, accompanied by more irregularities in the autocorrelograms. These two clusters were merged together in subsequent analysis (in which the resultant cluster is called 'R1').

A subset of the cells that were assigned to grid module clusters by the above procedure were tuned to both location and head direction (conjunctive grid × direction cells). These cells, which were defined as having a head-direction tuning curve with mean vector length above 0.3, were discarded from further analysis.

Classification of sleep states

SWS and REM periods were identified on the basis of a combination of behavioural and neural activity, following previously described approaches[6,45,46]. First, sleep periods were defined as periods of sustained immobility (longer than 120 s with a locomotion speed of less than 1 cm s^{-1} and head angular speed of less than 6° s^{-1}). Qualifying periods were then subclassified into SWS and REM on the basis of the amplitude of delta- and theta-band rhythmic activity in the recorded MEC–PaS cells. Spike times for each cell were binned at a resolution of 10 ms and the resultant spike counts were binarized, such that '0' indicated the absence of spikes and '1' indicated one or more spikes. The binarized spike counts were then summed across all cells (Extended Data Fig. 9A). The rhythmicity of this aggregated firing rate with respect to delta (1–4 Hz) and theta (5–10 Hz) frequency bands was quantified by applying a zero-phase, fourth-order Butterworth band-pass filter, then calculating the amplitude from the absolute value of the Hilbert transform of the filtered signal, which was smoothed using a Gaussian kernel with $\sigma = 5$ s and then standardized ('z-scored'). The ratio of the amplitudes of theta and delta activity was hence calculated (theta/delta ratio, 'TDR'). Periods in which TDR remained above 5.0 for at least 20 s were classified as REM; periods in which TDR remained below 2.0 for at least 20 s were classified as SWS (Extended Data Fig. 9B).

Spectral analysis was performed on 10-ms-binned multi-unit activity using the multi-tapered Fourier transform, implemented by the Chronux toolbox (http://chronux.org/, function 'mtspectrumsegc'). Non-overlapping 5-second windows were used, with a frequency bandwidth of 0.5 Hz and the maximum number of tapers.

Visualization of toroidal manifold

For each module of grid cells, spike times of co-recorded cells in the OF were binned for each cell at a resolution of 10 ms, and the binned spike counts were convolved with a Gaussian filter with $\sigma = 50$ ms. Time bins in which the rat's speed was below 2.5 cm s^{-1} were then discarded. To account for variability of average firing rates across cells, the smoothed firing rate of each cell was z-scored. For computational reasons, the time bins were downsampled, taking every 25th time bin (equating to 250-ms intervals between selected samples). Collectively, the downsampled firing rates of the full population of cells formed a matrix with time bins in rows and cells in columns. PCA was applied to this matrix (treating time bins as observations and cells as variables), and the first six principal components were retained (Extended Data Figs. 3Aa–c, 4a–d). UMAP[43,47] was then run on these six principal components (with time bins as observations and principal components as variables). The hyperparameters for UMAP were: ‘n_dims’=3, ‘metric’=‘cosine’, ‘n_neighbours’=5000, ‘min_dist’=0.8 and ‘init’=‘spectral’.

For visualizing the toroidal manifold during WW, smoothed firing rates were first calculated by the same procedure described above for OF. Subsequently, to allow comparison of the toroidal manifold between OF and WW, the same PCA and UMAP transformations calculated for the OF data were re-applied to the WW data, by supplying the fitted OF UMAP transformation as the argument ‘template_file’ to the ‘run_umap’ function in the MATLAB implementation[47].

Preprocessing of population activity

Each topological analysis was based on the activity of a single module of grid cells, during a single experimental condition in one recording session. Topological analysis of multi-module and conjunctive grid × direction cell activity was not considered as we expect such data to exhibit higher-dimensional topological structure requiring a higher number of cells[27]. The experimental conditions were: open-field foraging (OF), wagon-wheel track foraging (WW), slow-wave sleep (SWS), and rapid eye-movement sleep (REM). Sleep epochs of the same type were collected from across the recording and concatenated for analysis purposes. Similarly, in one case (rat 'S'), two WW task sessions were concatenated to increase the sample size.

In total there were 27 combinations of module (Q1, Q2, R1, R2, R3, S1) and experimental condition (OF day 1, OF day 2, WW, REM, SWS).

Preprocessing of spike trains began by computing delta functions centred on the spike times (valued 1 at time of firing; 0 otherwise), and convolving these temporally with a Gaussian kernel with $\sigma = 50$ ms (OF, WW and REM) or 25 ms (SWS). Samples of the smoothed firing rates of all cells (‘population activity vectors’) were then computed at 50-ms intervals. The awake states were further refined by excluding vectors which originated from time periods when the rat’s speed was below 2.5 cm s^{-1}.

Computing the persistent cohomology of a point cloud is computationally expensive and may be sensitive to outliers (for example, spurious points breaking the topology of the majority of points in the point cloud). For this reason, it is common to preprocess the data by downsampling and dimension-reducing the point cloud. The same preprocessing procedure was used for all datasets in the present study.

First, the data points were downsampled by keeping the 15,000 most active population activity vectors (as measured by the mean population firing rate). During SWS, this selection criterion had the consequence of automatically discarding population activity vectors during down-states, when neural activity is near-silent. As noise is inherently more prevalent and cosine distances less reliable in high-dimensional spaces (“the curse of dimensionality”)[48], dimensionality-reduction and a normalization of distances were subsequently performed. The reduced point cloud was z-scored and projected to its six first principal components, thus reducing noise while keeping much of the variance (see Extended Data Fig. 4a). This was supported by the lack of grid structure and the clear drop in explained deviance after six components (see Extended Data Fig. 4b, c). The explained deviance was computed by fitting a GLM model to each component individually, using the spatial coordinates as covariate, suggesting that the higher components are less spatially modulated and possibly better described by other (unknown) covariates. Consistent with this, the toroidal structure was most clearly detected in the barcodes when comparing the ratio of the lifetimes of the two most persistent H^1 bars versus the third longest-lived H^1 bar for the barcodes obtained when using different numbers of components in the analysis (see Extended Data Fig. 4d). These analyses both indicated that dimensionality reduction was required to firmly demonstrate the toroidal topology in the grid cells. The empirical findings are supported theoretically; see ‘Theoretical explanation of the six-dimensionality proposed by PCA’ in Supplementary Methods.

To further simplify the low-dimensional point cloud, a different downsampling technique was introduced, based on a point-cloud density strategy motivated by a topological denoising technique introduced previously[49] and a fuzzy topological representation used in UMAP[43,50]. Parts of the open-source implementation of the latter were copied in this computation. This approach consisted of assigning, for each point, a neighbourhood strength to its k nearest neighbours, and subsequently sampling points that represent the most tight-knit neighbourhoods of the point cloud in an iterative manner. First, we defined $m'_{i,i_j} = \exp\left(-\frac{d_{i,i_j}}{\sigma_i}\right)$, where d_{i,i_j} is the cosine distance between point x_i and its j-th nearest neighbour and σ_i is chosen to make $\sum_{j=1}^{k} m'_{i,i_j} = \log_2 k$, using $k = 1{,}500$. The neighbourhood strength was then obtained by symmetrizing: $m_{i,i_j} = m'_{i,i_j} + m'_{j,j_i} - m'_{i,i_j} \cdot m'_{j,j_i}$. Finally, the point cloud was reduced to 1,200 points by iteratively drawing the i-th point as: $\max_{x_i} \sum_{j \in \tilde{I}} m_{j,j_i}$, where $\tilde{I}$ denotes the indices of the points not already sampled. In other words, for each iteration, the sampled point is the one with the strongest average membership of the neighbourhoods of the remaining points.

To compute the persistent cohomology of the downsampled point cloud, the neighbourhood strengths were first computed for the reduced point cloud (using $k = 800$) and its negative logarithm was taken, obtaining a distance matrix. This matrix was then given as input to the Ripser implementation[51,52] of persistent cohomology, returning a barcode. In short, the barcode gave an estimate of the topology of the fuzzy topological representation of the six principal components of the grid-cell population activity. Thus, in essence, the first step of UMAP was applied before describing the resulting representation with persistent cohomology, instead of using it to project each point of the point cloud to a representation of user-specified dimensionality for visualization (Extended Data Fig. 3Ad, e). This gives a more direct and stable quantification of the global data structure, without having to choose an initialization[53] or optimize a lower-dimensional representation.

Persistent cohomology

Persistent cohomology, a tool in topological data analysis, was used to characterize the manifold assumed to underlie the data. This has clear ties with persistent homology and the main result (the barcode) is identical, thus the two terms are often used interchangeably. Persistent cohomology was chosen because the computation is (to our knowledge) faster and is required to obtain cocycle representatives, which are necessary to perform decoding (see ‘Cohomological decoding’). Persistent (co-)homology has previously been successful in analysing neural data, describing the ring topology of head direction cell activity[22–24], the spherical representation of population activity in primary visual cortex[54], and the activity of place cells[55–58].

The general outline of the algorithm is as follows. Each point in the cloud is replaced by a ball of infinitesimal radius, and the balls are gradually expanded in unison. Taking the union of balls at a given radius results in a space with holes of different dimensions. The range of radii for which each hole is detected is tracked; this is referred to as the ‘lifetime’ of the hole and is represented by the length of a bar. The totality of bars is referred to as the barcode.

The software package Ripser[51,52] was used for all computations of persistent cohomology. Ripser computes the persistent cohomology of ‘Vietoris-Rips complexes’ (which approximate the union of balls for different radii), constructed based on the input distance matrix and a choice of coefficients (in our case, $\mathbb{Z}_{47}$-coefficients), and outputs the barcode and cocycle representatives for all bars. The prime 47 was chosen as homology and cohomology coincide in this case and as it is unlikely that this divides the torsion subgroup of the homology of the space. Torsion may indicate, for example, orientability of a manifold and in choosing 47 as our prime, we disregard all but 47-torsion. Testing with other primes (for example, 43) gave similar results (data not shown) and the Betti numbers stayed the same regardless of choice of prime.

To verify that the lifetimes of prominent bars in the barcodes were beyond chance, shuffled distributions were generated for the persistence lifetimes in each dimension. In each shuffling, the spike train of each cell was shifted independently in time by rolling the firing rate arrays a random length between 0 and the length of the session. The same preprocessing and persistence analysis were then performed on the shifted spike trains as for the unshuffled data. This was performed 1,000 times, and each time a barcode was obtained. The barcodes were concatenated for all shuffles and the maximum lifetime was found for each dimension. This lifetime served as a significance criterion for the bar lifetimes. It is noted, however, that this is a heuristic and that statistics of barcodes are still not well established.

Cohomological decoding

As there are other spaces with similar barcodes as for a torus, the results identified by the barcode were further investigated, using the 'cohomological decoding' procedure introduced previously[59] to calculate a toroidal parametrization of the point clouds of population activity. This assigns to each point corresponding positions on each of the two circular features identified by the 1D bars with the longest lifetime, resulting in coordinates that further characterize the underlying shape of the data.

Cohomological decoding is motivated by the observation that the 1D cohomology (with integer coefficients) of a topological space X is equivalent to the set of homotopy-equivalent classes of continuous maps from X to the circle (S^1)[60]; that is:

$$H^1(X;\mathbb{Z}) \cong [X, S^1].$$

This subsequently means that for each 1D bar existing at a given radius, there exists a corresponding continuous map from the Vietoris-Rips complex of that radius to the circle. Thus, we may first use persistent cohomology to detect which elements represent meaningful (long-lived) features of the data and choose a radius for which these features exist. As the vertices of the Vietoris-Rips complex are points in the point cloud, the circular values of the corresponding maps at the vertices describe circular coordinates of the data.

In the present case, persistent cohomology was first applied to the grid-cell population activity and X was identified as the Vietoris-Rips complex for which the two longest-lived one-dimensional bars in the barcode (representing each of the two circles of the torus) existed. To define the desired toroidal coordinates on a domain that was as large as possible, we chose the complex given at the scale of the birth plus 0.99 times the lifetime of the second longest-lived one-dimensional bar in the barcode[22,59,61]. Next, the cocycle representatives (given by the persistent cohomology implementation of Ripser[51,52]) of each of the chosen 1D bars defined $\mathbb{Z}_{47}$-values for each of the edges in the complex. These edge values were then lifted to integer coefficients and subsequently smoothed by minimizing the sum over all edges (using the scipy implementation 'lsmr'). The values on the vertices (points) of each edge followed from the edge values and gave the circular parametrizations of the point cloud. The product of the two parametrizations thus provided a mapping from the neural activity to the two-dimensional torus—that is, giving a toroidal coordinatization (decoding) of the data.

As persistent cohomology was computed for a reduced dataset of 1,200 points and therefore circular parametrizations were obtained only for this point cloud, each parametrization was interpolated to the population activity from the rest of the session(s). First, the 1,200 toroidal coordinates were weighted by the normalized ('z-scored') firing rates of the cells at those time points, obtaining a distribution of the coordinates for each grid cell. The decoded toroidal coordinates were then computed by finding the mass centre of the summed distributions, weighted by the population activity vector to be decoded. These activity vectors were calculated by first applying a Gaussian smoothing kernel of 15-ms standard deviation to delta functions centred on spike times, sampling at 10-ms intervals and then z-scoring the activity of each cell independently. Time intervals that contained no spikes from any cell were subsequently excluded. When decoding was used to assess or compare the tuning properties of single cells (for example, comparison of toroidal versus spatial description), the coordinates were computed using the weighted sum of the distributions of the other cells; that is, the contribution of the cell to be assessed or compared was removed. When comparing preservation of toroidal tuning across two sessions, coordinates were interpolated either using the toroidal parametrization in each session independently ('Separate') or using the same toroidal parametrization in both sessions ('Common').

Toroidal rate map visualization

For visualization, toroidal firing rate maps were calculated in the same way as the physical space covariate (see 'Spatial position and direction tuning'), first binning the toroidal surface into a square grid of 7.2° × 7.2° bins and computing the average spike rate in each position bin. However, for toroidal maps, it was necessary to address the 60° angle between the toroidal axes before smoothing. After binning the toroidal coordinates, the rate map was 'straightened' by shifting the bins along the x axis ('horizontally') the length of (y mod 2)/2 bins, where y is the vertical enumeration of the given bin. Copies of the rate map were then tiled in a three-by-three square (similar to Extended Data Fig. 5d), before applying the closing and smoothing operations as for the spatial firing rate map. The single toroidal rate map was finally recovered by cutting out the centre tile, rotating it 90° and defining 15° shear angles along both the x and the y axis to correct for the 60° offset between them.

Comparison of spatial periodicity

Differences in grid periodicity between OF and WW environments were quantified for a given cell by comparing the grid scores in the two behavioural conditions. Two alternative methods were used to generate the spatial autocorrelograms for this comparison: (1) comparing the autocorrelograms for OF and WW directly; and (2) comparing autocorrelograms for OF and WW after first equalizing the spatial coverage between the two conditions.

For method (1), rate maps were calculated as specified in the above section 'Spatial position and direction tuning', using the same grid of 3 × 3-cm bins for both environments. This set of bins spanned the entirety of the OF arena and covered most of the WW track apart from some small regions at the outer extrema, which were discarded for the purpose of this analysis. For each of the two rate maps, the autocorrelogram was computed and the grid score was calculated.

Method (2) was similar to method (1), except that the cell's OF rate map was converted into a 'masked OF' rate map, by removing all bins that were unvisited by the rat in the WW session. This effectively equalized the position coverage between the two conditions, and thus allowed for a more valid comparison.

Toroidal versus spatial description

The explanatory significance of the toroidal description was evaluated by comparing statistical measures of how well the toroidal coordinates explained neural activity on the torus and in physical space. For a fair comparison, it was important to avoid overfitting, which might occur if a toroidal parametrization of a point cloud is used to describe that same set of data points. Two precautions were taken to avoid such overfitting: first, the data were decoded using the toroidal parametrization from a different condition (an OF session for a WW recording and a WW session for an OF recording), and second, the cell for which the statistical measurement was made was omitted from the decoding.

The comparison of toroidal and environmental representations also accounted for tracking error in the physical position estimate, which mainly resulted from the approximately 4 cm vertical offset of the tracking device above the rat's head. This causes a discrepancy when the angle α between the animal's zenith and the axis of gravitation is different from 0°, measured as 4 $\tan(\alpha)$ cm. The mean discrepancy in the recorded position data was measured to 1.5 cm. To account for this error of the position estimate, proportional Gaussian noise was added to the toroidal coordinates, using a standard deviation of 1.5 cm/Ω, where Ω denotes the grid spacing of the particular grid-cell module, estimated from the mean period of the fitted cosine waves of the toroidal coordinates in the open field (see 'Toroidal alignment').

Information content

The information content (I) was calculated as previously described[42], to quantify and compare the amount of information carried by single-cell activity about the location on the torus and physical space per spike. Both covariates were binned in a $M = 15 \times 15$ grid of square bins. For each bin j, the average firing rate $\mathbf{f}_j$ (given in spikes per second), and the occupancy ratio, $\mathbf{p}_j$, were computed. The information content for each grid cell was then given as:

$$I = \frac{1}{\bar{\mathbf{f}}} \sum_{j=1}^{M} \mathbf{f}_j \log_2 \frac{\mathbf{f}_j}{\bar{\mathbf{f}}} \mathbf{p}_j,$$

where $\bar{\mathbf{f}}$ is the mean firing rate of the cell across the entire session.

Note that although the rate maps for physical space have multiple firing fields, whereas the toroidal rate maps have single firing fields, we expect the spatial information to be comparable, as the measure primarily depends on the ratio of bins with high firing activity. This number should be comparable as the firing field size (in bins) will be inversely related to the number of fields in the rate map, assuming that the discretization of the map captures the relevant firing rate variations. For example, given a similar binning of space, a larger OF environment will include more fields, but the number of bins per field will decrease correspondingly. The binning used should be sufficient to resolve the smallest fields, as the same discretization was used in classifying the grid cells in the recorded population.

Deviance explained

Deviance explained was computed to measure how well a Poisson GLM model fitted to the spike count was at representing the data, using either the toroidal coordinates or the tracked position as regressors. A similar set-up was used to that of a previous study[62], with a smoothness prior for the GLM to avoid overfitting.

Both the toroidal and spatial coordinates were binned into a 15 × 15 grid of bins, and GLM design matrices were built with entries $X_i(t) = 1$ if the covariate at time t fell in the i-th bin and $X_i(t) = 0$ otherwise.

The Poisson probability of recording k spikes in time bin t is:

$$P(k|\mu(t), \beta) = \exp(-\mu(t)) \frac{\mu(t)^k}{k!},$$

where $\mu(t) = \exp\left(\sum_i \beta_i X_i(t)\right)$ is the expected firing rate in time bin t. The parameters β of the Poisson GLM were optimized for each covariate by minimizing the cost function:

$$L(\beta|\mu(t), \gamma, k) = -\sum_t \ln(P(k(t)|\mu(t), \beta)) + \frac{1}{2}\gamma \sum_{i,j \in N} (\beta_i - \beta_j)^2,$$

where N is the set of neighbour pairs. The first term is the negative log-likelihood of the spike count in the given time bin, whereas the second term puts a penalty on large differences in neighbouring parameters, enforcing smoothness in the covariate response of the predicted spike count.

The parameters, β, were initialized to zero and then modified to minimize the loss function by first running two iterations of gradient descent, before optimizing using the 'l-bfgs-b'-algorithm (as implemented in the 'scipy.optimize'-module) with 'gtol'=1e-5 as the cut-off threshold, and finally running two more iterations of gradient descent. A three-fold cross validation procedure was used, repeatedly fitting the model to two-thirds of the data and testing on the held-out last third.

The smoothness hyperparameter γ was optimized a priori on each grid-cell module based on the summed likelihood, testing $\gamma \in (1, \sqrt{10}, 10, \sqrt{1{,}000})$, and found to be either 1 or $\sqrt{10}$ in all cases.

Similarly, after fitting a null model (using only the intercept term) and the saturated model (perfectly fitting each spike count), the deviance explained could be computed as:

$$1 - \frac{\mathrm{ll_s} - \mathrm{ll_p}}{\mathrm{ll_s} - \mathrm{ll_0}},$$

where $\mathrm{ll_p}$, $\mathrm{ll_0}$ and $\mathrm{ll_s}$ denote the cross-validated log likelihood of the fitted model, the null model and the saturated model, respectively. This provides a normalized comparison describing the difference between the fitted model and the idealized model.

Toroidal alignment

To infer a geometric interpretation of the tori, as characterized via the cohomological decoding, and compare the toroidal parametrizations across modules and conditions, two cosine waves of the form $\cos(\omega t + k)$ were fitted to the OF mappings of the decoded circular coordinates (Extended Data Fig. 5a), where t is the centre 100^2-bins of a 540° × 540°-valued 150^2-bin grid rotated θ degrees. The parameters (ω, k, θ) were optimized by minimizing the square difference between the cosine waves and the cosine of the mean of the circular coordinates in 100^2 bins of the physical environment (smoothed using a Gaussian kernel with 1-bin standard deviation). Estimates were first obtained by finding the minimum when testing all combinations in the following intervals, each discretized in 10 steps: $\omega \in [1,6]$, $\phi \in [0, 360)$ and $\theta \in [0,180)$. The parameters of the cosine waves were further optimized using the 'slsqp'-minimization algorithm (as implemented in the 'scipy.optimize'-module using default hyperparameters). The period of each cosine wave was computed as 1.5 m/ω, giving a spatial scale estimate of the grid-cell modules.

As circular coordinates have arbitrary origin and orientation (that is, clockwise or counterclockwise evolution) we needed to realign the directions of the circular coordinates to compare these across modules and sessions (see Extended Data Fig. 4b). The clockwise orientation of each circular coordinate was first determined by noting whether $(\omega t + k)$ or $360° - (\omega t + k)$ best fit the spatial mapping of the circular means of the toroidal coordinates, and subsequently reoriented to obtain the same orientation for both coordinates. The coordinate for which $\cos(\theta)$ was largest (intuitively, the 'x axis') was then defined as the first coordinate (denoted ϕ_1, with parameters $(\omega_1, k_1, \theta_1)$) and the other as the second coordinate (ϕ_2). Although (ϕ_1, ϕ_2) fully describe the toroidal location, the hexagonal torus allows for three axes, and the two axes obtained are thus oriented at either 60° or 120° relative to each other (see Extended Data Fig. 5b). The difference in directions was given by $\theta_1 - \theta_2$ and if this difference was greater than 90°, ϕ_2 was replaced with $\phi_2 + 60° \cdot \phi_1$. Finally, the origin of the coordinates was aligned to a fixed reference, by subtracting the mean angular difference between the decoded coordinates and the corresponding coordinates obtained when using the toroidal parametrization of the reference OF session.

For visualization (Extended Data Fig. 5), it was furthermore necessary, in some cases, to rotate both vectors of the rhombi 30 degrees depending on whether one of the axes was directed outside of the box.

Preservation of toroidal tuning

Centre-to-centre distance and Pearson correlation were computed between toroidal tuning maps of different sessions to measure the degree of preservation between the toroidal descriptions.

First, the preferred toroidal firing location for each cell was computed as the centre of mass of the toroidal firing distribution:

$$T_{\mathrm{c}} = \arctan 2\left(\frac{\sum_i \sin\theta_i \cdot \mathbf{y}_i}{\sum_i \mathbf{y}_i}, \frac{\sum_i \cos\theta_i \cdot \mathbf{y}_i}{\sum_i \mathbf{y}_i}\right),$$

where $\mathbf{y}_i$ denotes the mean spike count of the given cell in the i-th bin whose binned toroidal coordinates are given by θ_i. The distance between mass centres found in two sessions (‘S_1’ and “S_2”) was then defined as:

$$d = \|\arctan 2(\sin(T_{\mathrm{c}}^{S_2} - T_{\mathrm{c}}^{S_1}),\ \cos(T_{\mathrm{c}}^{S_2} - T_{\mathrm{c}}^{S_1}))_2\|$$

where $\| \cdot \|_2$ refers to the L_2-norm.

Pearson correlation between two tuning maps was computed by flattening the smoothed 2D rate maps to 1D arrays and calculating the correlation coefficient, r, using the ‘pearsonr’-function given in the ‘scipy.stats’-library.

To determine how much the preservation of the toroidal representations across two sessions (measured with Pearson correlation and peak distance) differed from a random distribution, the indices of the cells in one of the sessions were randomly re-ordered before computing correlation and distance for the pair of conditions. This process was repeated 1,000 times, and the P value was calculated from the rank of the original r value or distance with respect to the shuffled distribution.

Classification of grid cells

Temporal autocorrelograms were computed, for each cell, by calculating a histogram of the temporal lags between every spike and all surrounding spikes within a 200 ms window, using 1 ms bins. The histogram was then divided by the value of the zero-lag bin, which was subsequently set to zero. The autocorrelogram was smoothed using a gaussian kernel with smoothing window 4 ms. Considering the autocorrelograms of all modules during OF foraging (day 2 for R1–3) as a point cloud, the cosine distances between all points were calculated, and hence each point’s 80 nearest neighbours were found. This defined a graph in which each point described a vertex and the neighbour pairs gave rise to edges. A density estimate was then calculated as the exponential of the negative distances summed over each neighbour for each point. The graph and the density estimate were given as the input to the Gudhi implementation[63] of ToMATo[64]. ToMATo uses a hill-climbing procedure to find modes of the density function and uses persistence to determine stable clusters. In the present case, the algorithm finds three long-lived clusters.

Minimum number of cells for torus detection

To address the question of how many cells are minimally needed to expect to see toroidal structure, random samples of $n = 10, 20, \ldots, 140$ cells were taken from R2 ($n = 149$ cells) during OF foraging, and the same topological analysis was repeated as for the whole population. The cells were resampled 1,000 times for each number of cells in the subsample. To determine whether toroidal structure was detected, a heuristic was introduced based on the circular parameterization given by the two most persistent 1D bars in the barcode mapped onto physical space. An estimate of the resulting planar representation of the torus was obtained by fitting planar cosine waves to each mapping (see ‘Toroidal alignment’). For the analysis to be determined ‘successful’ in detecting toroidal structure, we required: (i) the mean value of the least-squares fitting (across bins of the mapping) to be less than 0.25; (ii) the angle of the rhombus to be close to 60° (between 50° and 70°); and (iii) the side lengths to be within 25% of each other.

Toroidal peak detection

The number of peaks per toroidal rate map was detected to assert the number of grid cells whose toroidal rate map portrayed single fields. First, 1,000 points were sampled from the toroidal distribution given by the mean activity of each cell in 150 × 150 bins of the stacked toroidal surface (that is, as described in ‘Toroidal rate map visualization’, each 50 × 50-binned toroidal rate map is first ‘straightened’ and subsequently stacked in 3 × 3 to address the toroidal boundaries) and then spatially smoothed using a Gaussian kernel with smoothing widths 0, 1, 2, …, 10 bins with mode set to ‘constant’ in the ‘scipy.gaussian_filter’ function. Next, the points were clustered by computing a density estimate, using the Euclidean distance, and defining neighbours as points closer than 5 bins. Cluster labels were iteratively assigned to each point and all its neighbours in a downhill manner, instantiating a new cluster identity if the point was not already labelled. Finally, the centroids for each cluster were computed and counted as a peak depending on whether its position fell within the centre 50 × 50 bins of the stacked rate maps.

Simulated CAN models

To confirm the expected outcomes of topological analyses of grid cell CAN models, grid cells were simulated using two different, noiseless CAN models (Extended Data Fig. 7).

First, a 56 × 44 grid cell network was simulated based on the CAN model proposed previously[9], but using solely lateral inhibition (for details see ref. [11]) in the connectivity matrix, W. The animal movement was given as the first 1,000 s of the recorded trajectory of rat ‘R’ during OF session, originally sampled at 10 ms, and interpolated to 2-ms time steps. The speed, $v(t)$, and head direction $\theta(t)$ of the animal was calculated as the (unsmoothed) displacement in position for every time step. The activity, $\mathbf{s}$, was updated as:

$$\mathbf{s}_{i+1} = \mathbf{s}_i + \frac{1}{\tau}(-\mathbf{s}_i + (I + \mathbf{s}_i \cdot W + \alpha v(t)\cos(\theta(t) - \tilde{\theta}))_+),$$

where $(\ldots)_+$ is the Heaviside function and $\tilde{\theta}$ is the population vector of preferred head directions. The following parameters were used: $I = 1$, $\alpha = 0.15$, $l = 2$, $W_0 = -0.01$, $R = 20$ and $\tau = 10$, and let the activity pattern stabilize by first initializing to random and performing 2,000 updates, disregarding animal movement. For computational reasons, the activity was set to 0 if $\mathbf{s}_i < 0.0001$. The simulation was subsequently downsampled keeping only every 5th time frame.

Next, a 20 × 20 grid-cell network was simulated, for a synthetically generated OF trajectory (‘random walk’), based on the twisted torus model formulated in a previous study[10]. The parameter values and the code for computing both the grid cell network (choosing a single grid scale by defining the parameter ‘grid_gain’ = 0.04) and the random navigation (using 5,000 time steps) were given by the implementation by Santos Pata[65].

Idealized torus models

To compare the results of both the original and simulated grid cell networks with point clouds where the topology is known, a priori, to be toroidal, points were sampled from a square and a hexagonal torus. First, a 50 × 50 (angle) mesh grid (θ_1, θ_2) was created in the square $[0,2\pi)\times[0,2\pi)$ and slight Gaussian noise ($\epsilon = 0.1 \cdot N(0,1)$) was added to each angle. The square torus was then constructed via the 4D Clifford torus parametrization: $(\cos(\theta_1), \sin(\theta_1), \cos(\theta_2), \sin(\theta_2))$. The hexagonal torus was constructed using the 6D embedding: $(\cos(\theta_1), \sin(\theta_1), \cos(a_1\theta_1 + \theta_2), \sin(a_1\theta_1 + \theta_2), \cos(a_2\theta_1 + \theta_2), \sin(a_2\theta_1 + \theta_2))$, where $a_1 = 1/\sqrt{3}$ and $a_2 = -1/\sqrt{3}$.

Histology and recording locations

Rats were given an overdose of sodium pentobarbital and were perfused intracardially with saline followed by 4% formaldehyde.

The extracted brains were stored in formaldehyde and a cryostat was used to cut 30-μm sagittal sections, which were then Nissl-stained with cresyl violet. The probe shank traces were identified in photomicrographs, and a map of the probe shank was aligned to the histology by using two reference points that had known locations in both reference frames: (1) the tip of the probe shank; and (2) the intersection of the shank with the brain surface. In all cases, the shank traces were near-parallel to the cutting plane, therefore it was deemed sufficient to perform a flat 2D alignment in a single section where most of the shank trace was visible. The aligned shank map was then used to calculate the anatomical locations of individual electrodes (Extended Data Fig. 1).

Data analysis and statistics

Data analyses were performed with custom-written scripts in Python and MATLAB. Open-source Python packages used were: umap (version 0.3.10), ripser (0.4.1), numba (0.48.0), scipy (1.4.1), numpy (1.18.1), scikit-learn (0.22.1), matplotlib (3.1.3), h5py (2.10.0) and gudhi (3.4.1.post1). Samples included all available cells that matched the classification criteria for the relevant cell type. Power analysis was not used to determine sample sizes. The study did not involve any experimental subject groups; therefore, random allocation and experimenter blinding did not apply and were not performed. All statistical tests were one-sided.

The most intensive computations were performed on resources provided by the NTNU IDUN/EPIC computing cluster[66].

Additional discussion

The demonstration that populations of grid cells operate on a toroidal manifold, which is preserved across environments and behavioural states, confirms a central prediction of CAN models. The present observations provide the first—to our knowledge—population-level visualization of a two-dimensional CAN manifold, though there is accumulating evidence for one-dimensional CANs in a number of neural systems. The most powerful support for the latter has been obtained in fruit flies, in which CAN-like dynamics can be visualized in a ring of serially connected orientation-tuned cells of the central complex[67–69]. In mammals, analysis of data from dozens of simultaneously recorded head direction cells has shown that population activity in these cells faithfully traverses a conceptual ring[22–24], in accordance with ring-attractor models[17–19]. Dynamics along low-dimensional manifolds with line, ring, or sheet topologies is also thought to underlie a wide range of other mammalian brain functions that operate on continuous scales, spanning from visual orientation tuning[14] to neural operations underlying place-cell formation[70–72], as well as motor control[73], decision making and action selection[74–76], and certain forms of memory[39,77–80]. The present analyses provide a visualization of 2D CAN dynamics in pure grid cells within a module and, together with the previous work, point to a widespread implementation of CAN dynamics in the brain. The existence of CAN structure to constrain activity to low-dimensional manifolds does not preclude additional mechanisms for pattern formation, however. Grid cell patterns may emerge also by feedforward mechanisms[12,38,81–86]. Such mechanisms may operate in parallel with recurrent networks[87] and may even be the primary mechanism for grid-like firing at early stages of development, before the full maturation of recurrent connectivity[11,88–90].

Reporting summary

Further information on research design is available in the Nature Research Reporting Summary linked to this paper.

Data availability

The datasets generated during the current study are available at https://figshare.com/articles/dataset/Toroidal_topology_of_population_activity_in_grid_cells/16764508. Source data are provided with this paper.

Code availability

Code for reproducing the analyses in this article is available at https://figshare.com/articles/dataset/Toroidal_topology_of_population_activity_in_grid_cells/16764508.

42. Skaggs, W., Mcnaughton, B. & Gothard, K. An information-theoretic approach to deciphering the hippocampal code. *Adv. Neural Inf. Process. Syst.* **5**, 1030–1037 (1992).
43. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform manifold approximation and projection for dimension reduction. Preprint at https://arxiv.org/abs/1802.03426 (2018).
44. McInnes, L. UMAP https://github.com/lmcinnes/umap.
45. Lee, A. & Wilson, M. A. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183–1194 (2002).
46. Montgomery, S. M., Sirota, A. & Buzsáki, G. Theta and gamma coordination of hippocampal networks during waking and rapid eye movement sleep. *J. Neurosci.* **28**, 6731–6741 (2008).
47. Meehan, C., Ebrahimian, J., Moore, W. & Meehan, S. *Uniform Manifold Approximation and Projection (UMAP)* https://www.mathworks.com/matlabcentral/fileexchange/71902 (MATLAB, 2021).
48. Bellman, R. *Dynamic Programming* (Princeton Univ. Press, 1957).
49. Kloke, J. & Carlsson, G. Topological de-noising: strengthening the topological signal. Preprint at https://arxiv.org/abs/0910.5947 (2009).
50. Spivak, D. I. Metric realization of fuzzy simplicial sets https://www.semanticscholar.org/paper/metric-realization-of-fuzzy-simplicial-sets-spivak/a73fb9d562a3850611d2615ac22c3a8687fa745e (Semantic Scholar, 2009).
51. Tralie, C., Saul, N. & Bar-On, R. Ripser. py: a lean persistent homology library for python. *J. Open Source Softw.* **3**, 925 (2018).
52. Bauer, U. Ripser: efficient computation of Vietoris-Rips persistence barcodes. *J Appl. Comput. Topol.* **5**, 391–423 (2021).
53. Kobak, D. & Linderman, G. C. Initialization is critical for preserving global data structure in both t-SNE and UMAP. *Nat. Biotechnol.* **39**, 156–157 (2021).
54. Singh, G. et al. Topological analysis of population activity in visual cortex. *J. Vis.* **8**, 1–18 (2008).
55. Giusti, C., Pastalkova, E., Curto, C. & Itskov, V. Clique topology reveals intrinsic geometric structure in neural correlations. *Proc. Natl Acad. Sci. USA* **112**, 13455–13460 (2015).
56. Dabaghian, Y., Mémoli, F., Frank, L. & Carlsson, G. A topological paradigm for hippocampal spatial map formation using persistent homology. *PLoS Comput. Biol.* **8**, e1002581 (2012).
57. Spreemann, G., Dunn, B., Botnan, M. B. & Baas, N. A. Using persistent homology to reveal hidden covariates in systems governed by the kinetic Ising model. *Phys. Rev. E* **97**, 032313 (2018).
58. Baas, N. A. On the concept of space in neuroscience. *Curr. Opin. Syst. Biol.* **1**, 32–37 (2017).
59. De Silva, V., Morozov, D. & Vejdemo-Johansson, M. Persistent cohomology and circular coordinates. *Discrete Comput. Geom.* **45**, 737–759 (2011).
60. Hatcher, A. *Algebraic Topology* (Cambridge University Press, 2002).
61. Perea, J. A. in *Topological Data Analysis: The Abel Symposium 2018* (eds Baas, N. A. et al.) 435–458 (Springer, 2020).
62. Ledergerber, D. et al. Task-dependent mixed selectivity in the subiculum. *Cell Rep.* **35**, 109175 (2021).
63. The GUDHI Project. *GUDHI User and Reference Manual* (GUDHI Editorial Board, 2021).
64. Chazal, F., Guibas, L. J., Oudot, S. Y. & Skraba, P. Persistence-based clustering in Riemannian manifolds. *J. ACM* **60**, 41 (2013).
65. Santos Pata, D. *Grid Cells* https://github.com/DiogoSantosPata/gridcells (2020).
66. Själander, M., Jahre, M., Tufte, G. & Reissmann, N. EPIC: an energy-efficient, high-performance GPGPU computing research infrastructure. Preprint at https://arxiv.org/abs/1912.05848 (2020).
67. Seelig, J. D. & Jayaraman, V. Neural dynamics for landmark orientation and angular path integration. *Nature* **521**, 186–191 (2015).
68. Kim, S. S., Rouault, H., Druckmann, S. & Jayaraman, V. Ring attractor dynamics in the *Drosophila* central brain. *Science* **356**, 849–853 (2017).
69. Green, J. et al. A neural circuit architecture for angular integration in *Drosophila*. *Nature* **546**, 101–106 (2017).
70. McNaughton, B. L. et al. Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J. Exp. Biol.* **199**, 173–185 (1996).
71. Samsonovich, A. & McNaughton, B. L. Path integration and cognitive mapping in a continuous attractor neural network model. *J. Neurosci.* **17**, 5900–5920 (1997).
72. Tsodyks, M. & Sejnowski, T. J. Associative memory and hippocampal place cells. *Int. J. Neural Syst.* **6**, 81–86 (1995).
73. Aksay, E. et al. Functional dissection of circuitry in a neural integrator. *Nat. Neurosci.* **10**, 494–504 (2007).
74. Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
75. Machens, C. K., Romo, R. & Brody, C. D. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* **307**, 1121–1124 (2005).
76. Lukashin, A. V. & Georgopoulos, A. P. A dynamical neural network model for motor cortical activity during movement: population coding of movement trajectories. *Biol. Cybern.* **69**, 517–524 (1993).
77. Romani, S. & Tsodyks, M. Continuous attractors with morphed/correlated maps. *PLoS Comput. Biol.* **6**, e1000869 (2010).
78. Compte, A., Brunel, N., Goldman-Rakic, P. S. & Wang, X. J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* **10**, 910–923 (2000).
79. Wimmer, K., Nykamp, D. Q., Constantinidis, C. & Compte, A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.* **17**, 431–439 (2014).

80. Koyluoglu, O. O., Pertzov, Y., Manohar, S., Husain, M. & Fiete, I. R. Fundamental bound on the persistence and capacity of short-term memory stored as graded persistent activity. *eLife* **6**, e22225 (2017).
81. Stepanyuk, A. Self-organization of grid fields under supervision of place cells in a neuron model with associative plasticity. *Biol. Inspired Cogn. Archit.* **13**, 48–62 (2015).
82. Dordek, Y., Soudry, D., Meir, R. & Derdikman, D. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife* **5**, e10094 (2016).
83. Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
84. D'Albis, T. & Kempter, R. A single-cell spiking model for the origin of grid-cell patterns. *PLoS Comput. Biol.* **13**, e1005782 (2017).
85. Monsalve-Mercado, M. M. & Leibold, C. Hippocampal spike-timing correlations lead to hexagonal grid fields. *Phys. Rev. Lett.* **119**, 038101 (2017).
86. Weber, S. N. & Sprekeler, H. Learning place cells, grid cells and invariances with excitatory and inhibitory plasticity. *eLife* **7**, e34560 (2018).
87. Si, B., Kropff, E. & Treves, A. Grid alignment in entorhinal cortex. *Biol. Cybern.* **106**, 483–506 (2012).
88. Langston, R. F. et al. Development of the spatial representation system in the rat. *Science* **328**, 1576–1580 (2010).
89. Wills, T. J., Cacucci, F., Burgess, N. & O'Keefe, J. Development of the hippocampal cognitive map in preweanling rats. *Science* **328**, 1573–1576 (2010).
90. Donato, F., Jacobsen, R. I., Moser, M.-B. & Moser, E. I. Stellate cells drive maturation of the entorhinal–hippocampal circuit. *Science* **355**, eaai8178 (2017).
91. Gray, C. M., Maldonado, P. E., Wilson, M. & McNaughton, B. Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J. Neurosci. Methods* **63**, 43–54 (1995).
92. Hardcastle, K., Maheswaranathan, N., Ganguli, S. & Giocomo, L. M. A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. *Neuron* **94**, 375–387 (2017).

Acknowledgements We thank M. P. Witter for help with evaluation of recording locations, and A. M. Amundsgård, K. Haugen, K. Jenssen, E. Kråkvik, I. Ulsaker-Janke and H. Waade for technical assistance. The work was supported by a Synergy Grant to E.I.M. and Y.B. from the European Research Council ('KILONEURONS', grant agreement no. 951319); an RCN FRIPRO grant to E.I.M. (grant no. 286225); a Centre of Excellence grant to M.-B.M. and E.I.M. and a National Infrastructure grant to E.I.M. and M.-B.M. from the Research Council of Norway (Centre of Neural Computation, grant number 223262; NORBRAIN, grant number 295721); the Kavli Foundation (M.-B.M. and E.I.M.); the Department of Mathematical Sciences at the Norwegian University of Science and Technology (B.A.D., E.H. and N.A.B.); a direct contribution to M.-B.M. and E.I.M. from the Ministry of Education and Research of Norway; and grants to Y.B. from the Israel Science Foundation (grant no. 1745/18) and the Gatsby Charitable Foundation. Some of the computations were performed on resources provided by the NTNU IDUN/EPIC computing cluster.

Author contributions R.J.G., M.-B.M. and E.I.M. designed experiments. R.J.G. performed experiments. N.A.B., E.H., B.A.D., R.J.G., Y.B. and E.I.M. conceptualized and proposed analyses. E.H. and R.J.G. developed and performed the analyses. M.P. shared unpublished Kilosort software. R.J.G., E.H., B.A.D., Y.B., M.-B.M. and E.I.M. interpreted data. E.H. and R.J.G. visualized data. R.J.G., E.H., B.A.D., Y.B. and E.I.M. wrote the paper, with periodic input from all authors. E.I.M., M.-B.M., B.A.D. and N.A.B. supervised the project. E.I.M., M.-B.M. and Y.B. obtained funding.

Competing interests The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41586-021-04268-7.

Correspondence and requests for materials should be addressed to Richard J. Gardner, Nils A. Baas, Benjamin A. Dunn or Edvard I. Moser.

Peer review information *Nature* thanks Carina Curto, Arseny Finkelstein and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

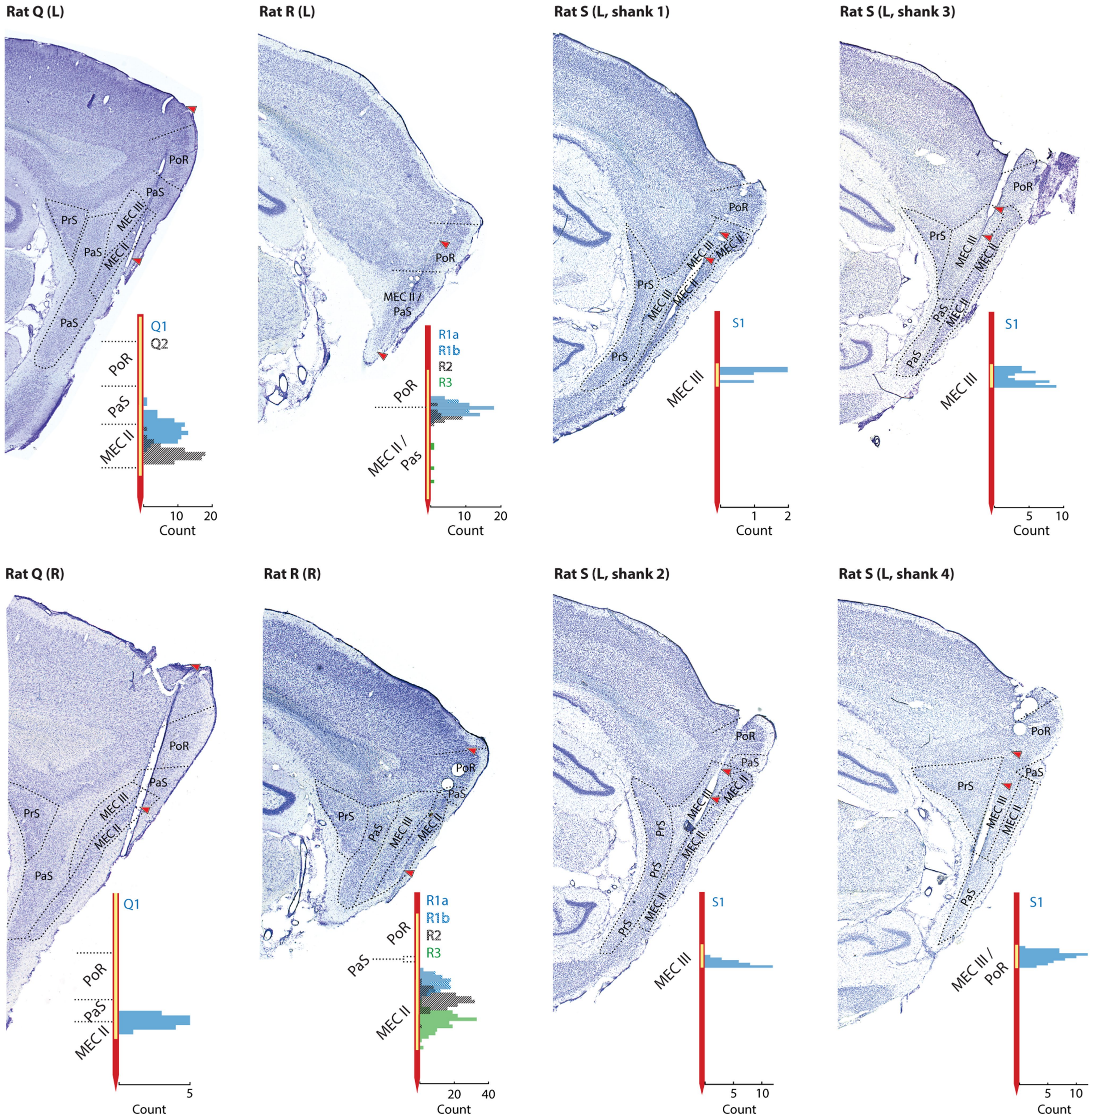

Extended Data Fig. 1 | Nissl-stained sagittal brain sections showing recording locations for rats Q, R and S. Red arrows indicate the dorsoventral range of the probe's active recording sites (corresponding to the yellow stripe in the inset). Stippled lines indicate borders between brain regions (MEC, medial entorhinal cortex; PaS, parasubiculum, PrS, presubiculum; PoR, postrhinal cortex). Layers are indicated for MEC (MECII, MECIII). Animal name, hemisphere (L, left; R, right) and shank number (for Rat 'S') are indicated in text above each section. Insets show, for each section, the number of grid cells recorded at each depth on the probe shank (histogram bin sizes 100 µm for Rats 'Q' and 'R', 75 µm for Rat 'S'; total numbers of cells are given in Extended Data Fig. 2g). Only the implanted portion of the probe shank is shown. Counts are colour-coded according to module identity. Module R1 is subdivided into the two UMAP clusters R1a and R1b (as shown in Extended Data Fig. 2), shown here as two stacked histograms. The yellow stripe on the probe shank indicates the range of active recording sites. The indicated locations of units are subject to measurement error, because the anatomical registration of probe shanks can only be approximately estimated, and furthermore because units may be detected on electrodes up to 50 µm away[91]. Note that several modules spanned across hemispheres (see Extended Data Fig. 2g). The cell counts shown for Rat 'R' are from Recording Day 1. The same set of recording sites was used for both recording sessions, and therefore the anatomical distributions of recorded cells were similar between the two sessions.

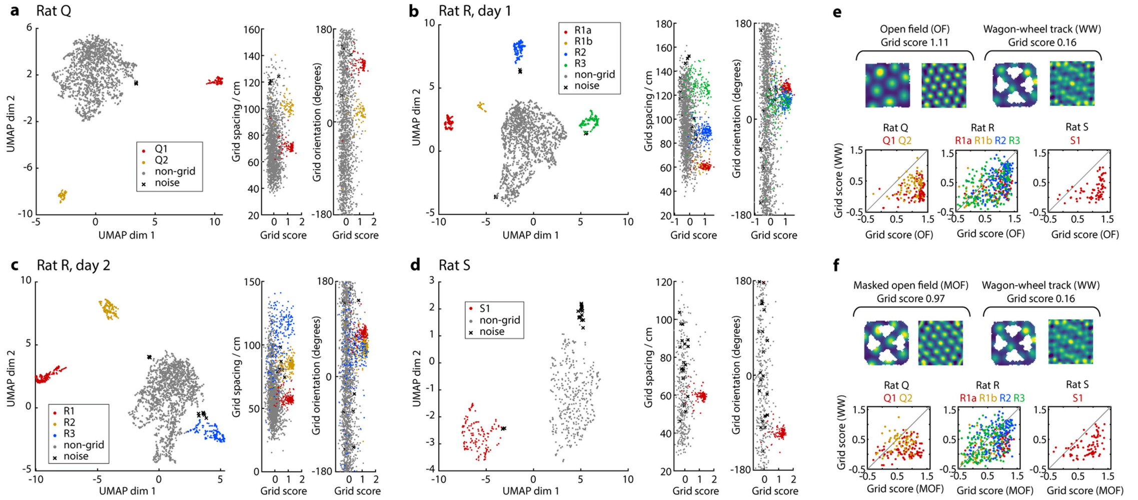

g

Cluster	Grid	Pure grid	Grid-HD	Class B	Class N	Class T	Hemisphere L	Hemisphere R
Q1	97	94	3	83	6	8	79	18
Q2	66	65	1	64	0	2	66	0
Q non-grid	1926	n/a	n/a	n/a	n/a	n/a	1232	1097
R1a, day 1	134	62	72	n/a	n/a	n/a	66	68
R1b, day 1	32	31	1	n/a	n/a	n/a	4	28
R2, day 1	168	149	19	n/a	n/a	n/a	21	147
R3, day 1	149	145	4	n/a	n/a	n/a	4	145
R non-grid, day 1	1956	n/a	n/a	n/a	n/a	n/a	1310	1111
R1, day 2	189	111	78	69	28	92	75	114
R2, day 2	172	152	20	131	6	35	23	149
R3, day 2	183	165	18	148	3	32	13	170
R non-grid, day 2	2011	n/a	n/a	n/a	n/a	n/a	1533	1190
S1	140	72	68	28	52	60	140	n/a
S non-grid	384	n/a	n/a	n/a	n/a	n/a	499	n/a
S1 shank 1	6	6	0	2	3	1	6	n/a
S1 shank 2	30	30	0	2	28	0	30	n/a
S1 shank 3	45	24	21	10	21	14	45	n/a
S1 shank 4	59	12	47	14	0	45	59	n/a
S non-grid shank 1	142	n/a	n/a	n/a	n/a	n/a	142	n/a
S non-grid shank 2	86	n/a	n/a	n/a	n/a	n/a	86	n/a
S non-grid shank 3	66	n/a	n/a	n/a	n/a	n/a	66	n/a
S non-grid shank 4	90	n/a	n/a	n/a	n/a	n/a	90	n/a

h

Rat	Cells	Grid	Pure grid	Grid-HD	Min. OF	Min. WW	Min. SWS	Min. REM	Min. total
Rat Q	2094	163	159	4	31.4	42.4	105.0	17.5	360.8
Rat R day1	2460	483	387	96	127.7	44.9	0	0	236.4
Rat R day2	2571	544	428	116	21.1	0	137.0	30.9	259.6
Rat S	546	140	72	68	30.64	25.0	87.7	9.8	250.0

Extended Data Fig. 2 | Grid module identification and properties. **a**–**d**, Clustering of grid modules (**a**, Rat 'Q'; **b**, Rat 'R', day 1; **c**, Rat 'R', day 2; **d**, Rat 'S'). For all experiments, coarse spatial autocorrelograms were first calculated from all cells' OF firing rate maps (*n* cells as shown in **g**). UMAP was then used to reduce the *M*-dimensional autocorrelograms (where $M = 668$ spatial bins) to a two-dimensional point cloud, where each point represented the autocorrelogram of a single cell, and distances between points represented the similarity between autocorrelograms. Left scatterplot in **a**–**d**: 2D point cloud, with points colour-coded according to cluster ID. Clusters were identified by applying the density-based clustering algorithm DBSCAN to the 2D point cloud. In every recording, the largest cluster (in grey, labelled "main") comprised mainly non-grid cells, and the remaining smaller clusters (coloured) represented different modules of grid cells. The black crosses ("noise") are identified as outlier data points. The well-isolated clusters formed by grid cells support the notion that these cells are a distinct functional class, in contrast to the claim that grid-like characteristics are expressed by MEC cells to different extents[92]. Right pair of scatterplots in **a**–**d:** Combinations of three grid parameters (grid score, grid spacing and grid orientation) for co-recorded cells from each recording. Each dot corresponds to one autocorrelogram (one cell). Dots are coloured by cluster ID as in **a**. **e**, Comparison of grid-cell spatial periodicity in the open-field arena (OF) and on the wagon-wheel track (WW). Top: firing rate map and corresponding autocorrelogram for an example grid cell in OF (left) and WW (right). For the purposes of this comparison, the same position bins were applied to both environments, resulting in cropping of the outermost parts of WW. Colour coding as indicated by scale bar; peak rates 16.1 Hz (OF) and 15.8 Hz (WW); range of autocorrelation values: −0.56 to 0.83 and −0.58 to 0.71, respectively. Note the more irregular appearance of the autocorrelogram for WW. Bottom: scatter plots showing grid scores of all grid cells in OF (*x* axis) and WW (*y* axis). Colours refer to the module assignment in **a**. Note the bias for points to lie in the lower-right quadrant, reflecting generally higher grid scores in OF than in WW. **f**, As for **e**, but controlling for differences in behavioural coverage of OF and WW environments. It is possible that the lower WW grid scores in **e** were a product of sparser behavioural coverage of the WW environment (animals visited only positions on the track). To control for this possibility, we created "masked OF" (MOF) rate maps by removing spatial bins from the original OF rate map which were not visited by the animal in WW. In all modules, grid scores in the "masked" OF condition were higher than in WW (grid score mean ± S.E.M. across all cells: OF: 0.677 ± 0.017, WW: 0.360 ± 0.017, $N = 618$ cells, *P* values for the 6 modules ranged from 1.26×10^{-14} to 0.03, *Z*-values ranged from 2.12 to 7.71, Wilcoxon signed-rank test). Top row shows the same example cell as in **e** after leaving the same subset of position bins in OF as in WW. Bottom row shows comparison of grid scores for MOF and WW. As in **e**, grid scores are lower for WW, indicating that grid periodicity is reduced in WW even when differences in spatial coverage are accounted for. **g**, Table showing total number of cells and number of pure grid cells and conjunctive grid × direction cells. **h**, Number of cells (as in **g**) broken down on recording sessions, with session lengths in minutes indicated for open field (OF), wagon wheel (WW), slow-wave sleep (SWS) and REM sleep.

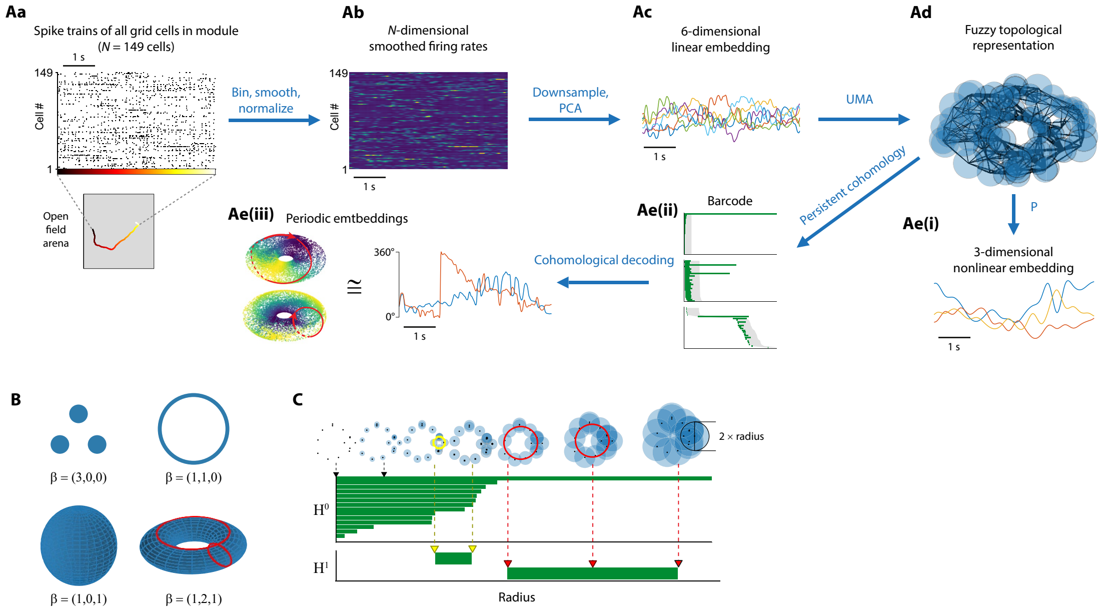

Extended Data Fig. 3 | Preprocessing steps for visualization and detection of toroidal topology. A, Flow diagram showing method for extracting low-dimensional embeddings of neural activity. The animal foraged in an OF arena while spikes from 149 grid cells shown in Fig. 1a were recorded (**Aa**; cells are ordered arbitrarily). A 5-second example behavioural trajectory is highlighted, with colour indicating elapsed time. The spike trains were binned in time (*N* bins) and then smoothed and normalized, yielding a matrix of *N*-dimensional population activity vectors (**Ab**). After temporally downsampling and *z*-scoring the neural activity, PCA was applied to the *N*-dimensional neural activity, yielding a six-dimensional linear embedding (**Ac**). This preserved the grid structure in the activity (Extended Data Fig. 4b, c), while mitigating drawbacks associated with high-dimensional spaces (the "curse of dimensionality")[48]. The six principal components were then passed through a second, nonlinear, dimensionality reduction step by UMAP, which generated a three-dimensional nonlinear embedding (**Ae(i)**) allowing the toroidal structure to be visualized. UMAP consists of two steps: first, a fuzzy topological graph representation is constructed (i.e. a "Uniform Manifold Approximation" - UMA) using a distance metric in the high-dimensional space (**Ad**); second, to obtain the lower-dimensional projection (P), the coordinates of corresponding points in fewer dimensions are optimized to have a similar fuzzy topological representation. In the persistence analysis, we applied persistent cohomology to the fuzzy topological representation of the high-dimensional point cloud (**Ae(ii)**) and subsequently used cohomological decoding to obtain a two-dimensional projection of the original *N*-dimensional point cloud (**Ae(iii)**; right, showing a 5-second snippet; left, embedded in 3D, points are coloured by each angular coordinate, whose direction is indicated by a red arrow). **B**, Cohomology can help differentiate topological spaces such as the union of three discs (upper left), a circle (upper right), a sphere (lower left) and a torus (lower right) by counting the number of topological holes (β) in different dimensions. A disc has a 0D hole (a connected component); a circle additionally has a 1D hole; a (hollow) sphere is a connected component and has a 2D hole (a cavity); a torus is a connected component with two 1D holes (illustrated with red circles) and one 2D hole (a cavity in the interior of the torus). **C**, Persistent cohomology tracks the lifetime of topological holes in spaces associated with point clouds. Top: The radius of balls centred at each data point in the point cloud is continuously increased (left to right). The union of the balls forms a space with possible holes. The lifetime of a hole during expansion of the radius is defined as the radial interval from when the hole first appears until it is filled in. Note the short lifetime of the hole marked with a red circle and the long lifetime of the hole indicated with a yellow circle. Second and third row: The lifetime of each hole of dimension zero (H^0) and one (H^1) in the example in the top row is indicated by the length of a bar (in green) in the barcode diagram. Two 1D holes are detected: the first bar, corresponding to the red hole in the top row, is short and regarded as noise, and the second, corresponding to the yellow hole, is substantially longer and captures the prominent topology of the point cloud.

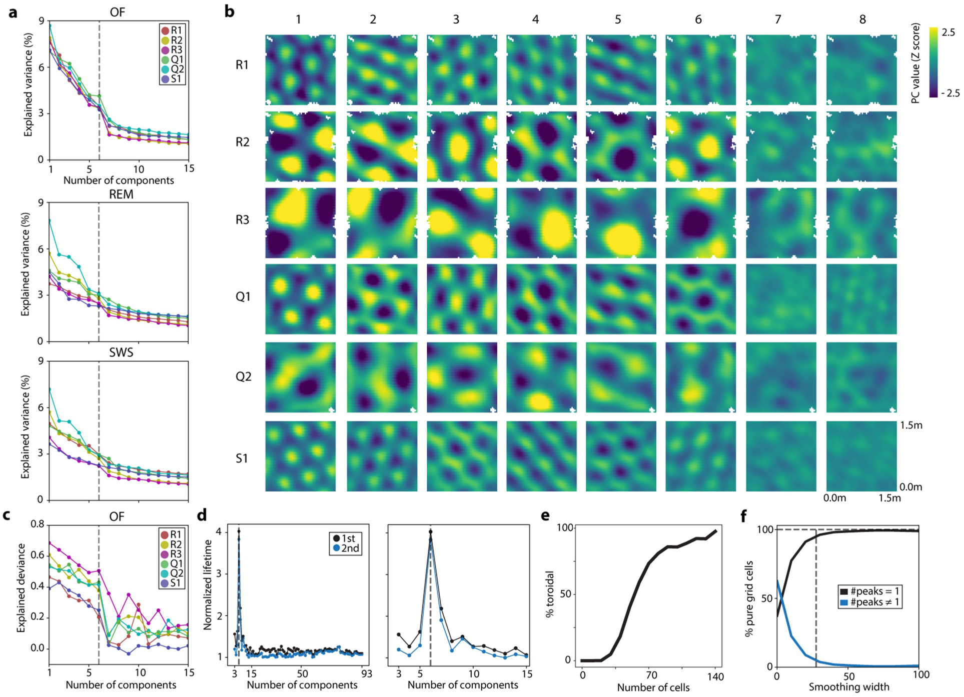

Extended Data Fig. 4 | Analysis of principal components, number of cells and number of toroidal peaks. **a**, Variance explained by the first 15 principal components (PCs) after applying PCA to the *n*-dimensional neural activity, shown for each module. Note that during OF, a particularly large amount of variance is explained by the first 6 PCs, followed by a sharp drop in the 7th PC, in all modules. A drop in variance explained is also seen after the 6th PC in REM and SWS. **b**, The first six PCs contain a grid-like representation at the population level. Each panel shows the mean value of one PC as a function of the animal's position in the OF. PC value is colour-coded as indicated by the scale bar. The 8 first PCs are arranged in descending order of explained variance (columns, from left to right), and are shown for each module (in rows). Note the presence of grid-like structure, which is particularly strong in the first six PCs, irrespective of the grid spacing. These six grid-like PCs correspond to the set with the highest explained variance in **a**. *z*-scored PC values are indicated by the scale bar (see Supplementary Methods for theoretical explanation of the six-dimensionality). **c**, Line plots showing the goodness-of-fit of a Gaussian GLM model based on the position in the spatial environment (OF) fitted to each principal component (components as in **a**). This is measured (as in Fig. 2d) as the explained deviance of the model showing that the six first components are better explained by space than the subsequent components for each module. **d**, Line plots showing the lifetime of the two longest-lived H^1-bars (longest-lived – "1st", black; second longest-lived – "2nd", blue) divided by the lifetime of the third longest-lived H^1-bar as a function of number of principal components kept in the persistence analysis of R1 day 1 OF ($n = 93$ cells). This heuristic measures how clearly the two longest-lived H^1-bars (expected to be long for a torus) separates from the third (expected to be short), thus indicating how clearly the barcode displays toroidal topology. This is clearly the case when using 6 principal components in this dataset. **e**, The percentage of subsamples of R2 (resampled randomly 1,000 times per number of cells; total $n = 149$ cells) for which toroidal structure was detected in the parameterization given by the two most persistent 1D bars in the barcode (as in Extended Data Fig. 5). Note that approximately 60 cells were needed for the probability of detecting toroidal structure exceed 50%. **f**, Effect of varying spatial smoothing on the number of peaks in toroidal rate maps. The *y* axis displays the percentage of single-peaked (black) and multi-peaked (blue) toroidal rate maps of all grid cells ($n = 2{,}727$ cells) pooled across modules and behaviour conditions. The vertical dashed line marks the smoothing width used in Extended Data Fig. 10, and the horizontal dashed line marks 100%. Note that cells with single peaks quickly describe the majority of the pooled cells.

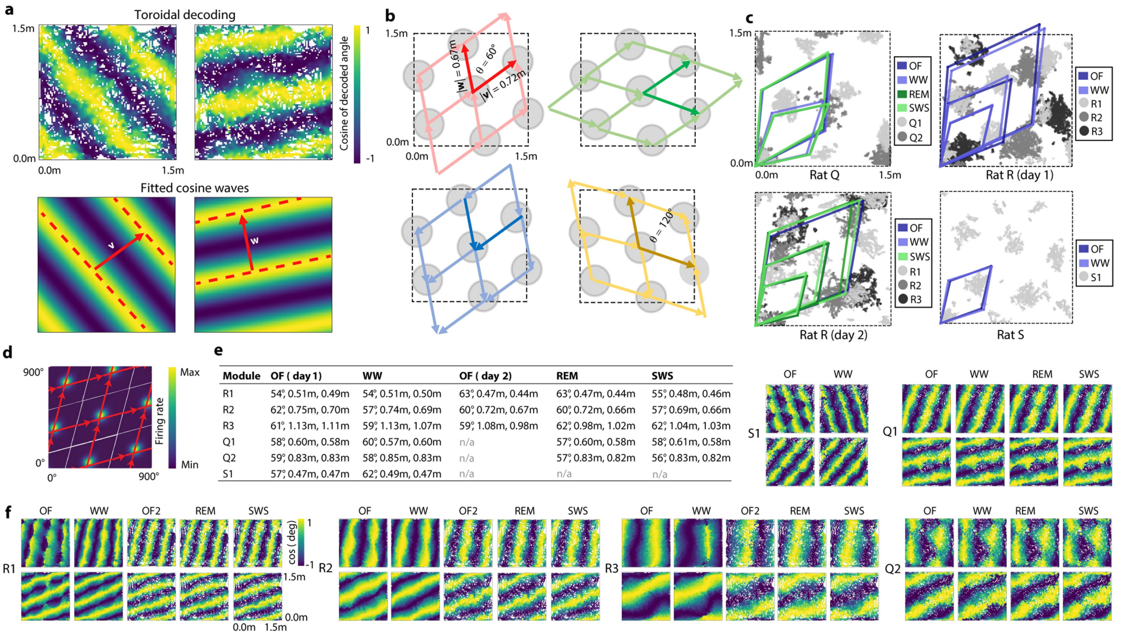

Module	OF (day 1)	WW	OF (day 2)	REM	SWS
R1	54°, 0.51m, 0.49m	54°, 0.51m, 0.50m	63°, 0.47m, 0.44m	63°, 0.47m, 0.44m	55°, 0.48m, 0.46m
R2	62°, 0.75m, 0.70m	57°, 0.74m, 0.69m	60°, 0.72m, 0.67m	60°, 0.72m, 0.66m	57°, 0.69m, 0.66m
R3	61°, 1.13m, 1.11m	59°, 1.13m, 1.07m	59°, 1.08m, 0.98m	62°, 0.98m, 1.02m	62°, 1.04m, 1.03m
Q1	58°, 0.60m, 0.58m	60°, 0.57m, 0.60m	n/a	57°, 0.60m, 0.58m	58°, 0.61m, 0.58m
Q2	59°, 0.83m, 0.83m	58°, 0.85m, 0.83m	n/a	57°, 0.83m, 0.82m	56°, 0.83m, 0.82m
S1	57°, 0.47m, 0.47m	62°, 0.49m, 0.47m	n/a	n/a	n/a

Extended Data Fig. 5 | Mapping of decoded circular coordinates onto the open field allows geometrical interpretation of toroidal structure. a, Top row: Toroidal coordinates given by cohomological decoding from activity of grid module R2 during OF foraging, mapped onto the recording box. In each plot, colour indicates the mean value of the cosine of each of the two circular coordinates. The mappings of both coordinates show 2D striped patterns, with similar periods but distinct angles. Bottom row: A cosine wave is fitted to each coordinate to obtain the direction of the toroidal axes. The period and angle of the cosine wave in the plane may be represented by spatial vectors, **v** and **w**, with corresponding length and orientation. Note the clear transversality of the two circles, expressed in the directions of the two vectors, further confirming the toroidal identification of the data. **b**, The periods and angles of the cosine waves in **a** reflect the scale and orientation of the grid module. Taking the origin of the vectors in **a** to be alike, we see that the vectors span a parallelogram with approximately equal side lengths (0.67m and 0.72m) and an angle of 60 degrees, suggesting a rhomboidal tile representing the toroidal structure (top left). When repeated across the environment, the tile depicts the hexagonal grid pattern of the grid-cell module, confirming that the product of the two decoded circles defines a hexagonal ("twisted") torus. As the orientation of the circular coordinates is arbitrary, the directions of the axes may be any of the following: reversely oriented (blue arrows), a different 60-degree pair of axes (green), or have a relative angle of 120 degrees (yellow). **c**, Rhombi of each module for each OF session (*n* cells as in Extended Data Fig. 2g), given by the cosine wave fitted to the toroidal coordinates (as in **b**). The toroidal parametrizations were obtained independently in different behavioural conditions (colour-coded), then used to decode the module's activity during OF foraging, and subsequently mapped as a function of the rat's position in the environment (see **f**). Positions of downsampled spikes from example cells of each module are shown in greyscale to illustrate grid scale and orientation. The consistent angle and side lengths suggest the geometry of the rhombus is retained across brain states and environments, with a constant scale relationship between modules. **d**, Mean value of a single neuron in rhomboidal coordinates displays a single bump (as in Fig. 2a), which, when repeated and arranged to tesselate a 2D surface, reveals a grid-like pattern in the activity of the grid cell, akin to its spatial firing. **e**, Table of side lengths and angles of the cosine waves that form the rhombi in **c**, shown for each grid module and each condition (*n* cells as in Extended Data Fig. 2g). **f**, Visualization of the cohomological decoding of toroidal coordinates as a function of physical space (one visualization for each grid module during each condition, with the toroidal parametrizations aligned to the same axes before creating the rate maps; *n* cells as in Extended Data Fig. 2g). All barcodes which indicated toroidal structure exhibited periodic stripes in the OF, with phase and orientation corresponding to the two-dimensional periodicity of the grid pattern of the respective module. SWS* refers to the decoding when considering only "bursty" (B) cells of R1 as given by the correlation clustering method described in Fig 4b.

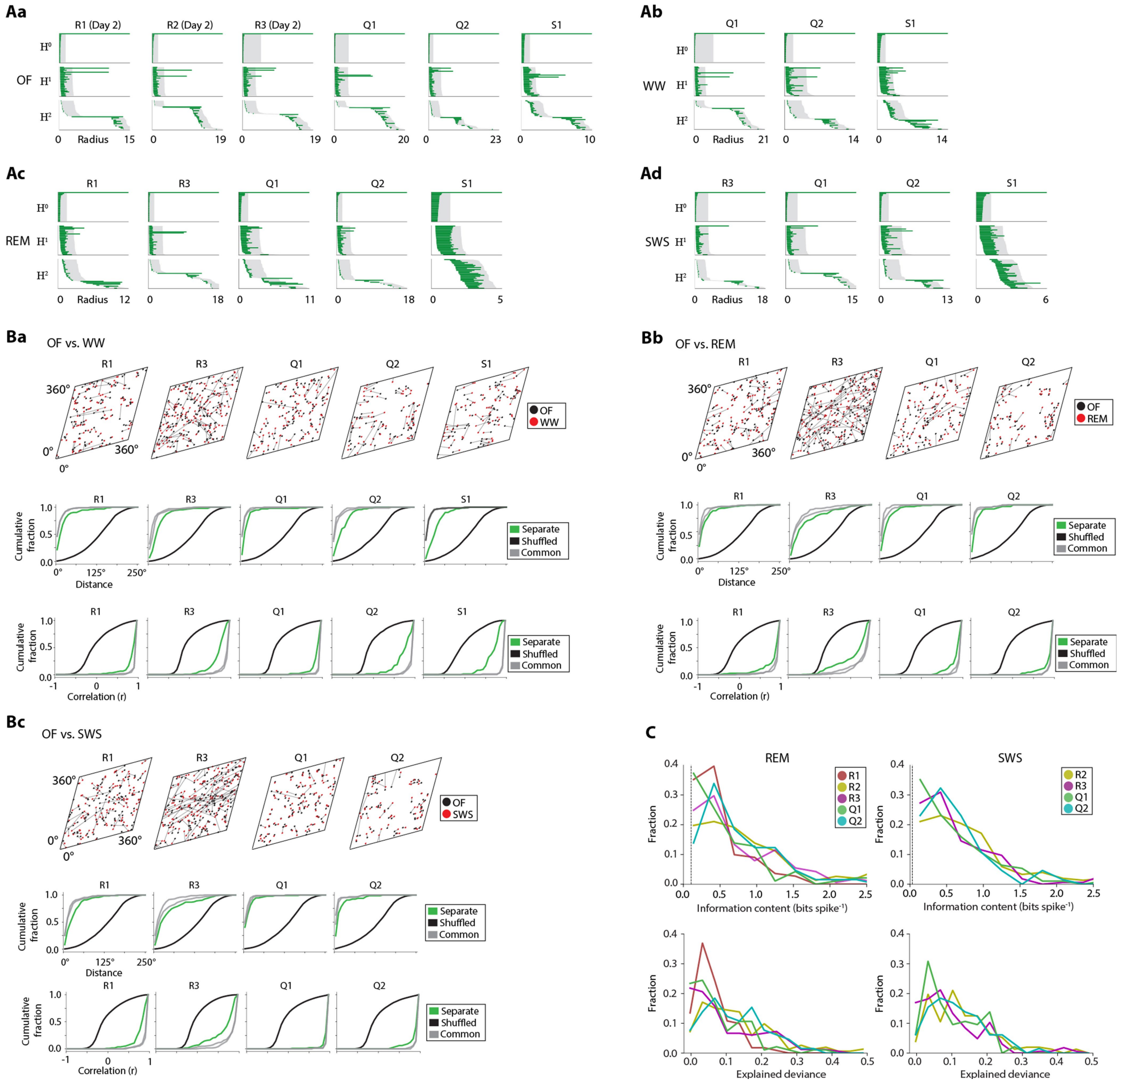

Extended Data Fig. 6 | Barcodes and toroidal tuning statistics for grid modules or recording sessions not included in Figs. 2–4. Data are shown for six grid-cell modules: R1, R3, Q1, Q2, S1 and R2 (*n* cells as in Extended Data Fig. 2g). Toroidal structure is clearly present across environments and behavioural states. **Aa**–**Ad**, Barcode diagrams (as in Fig 1e, f) showing the results of the persistent cohomology analysis on open-field (OF), wagon-wheel track (WW) or sleep (REM or SWS) data. **Ba**–**Bc**, Preservation of toroidal field centres between conditions: OF vs WW (1), OF vs REM (2) OF vs SWS (3). Top row in each panel: Distribution of grid cells' receptive field centres on the inferred torus for OF and WW as well as sleep states, similar to Fig 2e. Each dot signifies the field centre of an individual grid cell. Grey lines connect field centres of the same cell across conditions. Note the proximity of red-black pairs (after separate alignment for the two recording sessions of each panel). Middle and bottom rows: Cumulative distributions showing stability of grid cells' toroidal tuning between brain states, as in Fig. 2f, g. Distributions show peak field distance (middle) and Pearson correlation of pairs of toroidal rate maps (bottom). Labelling as in Fig. 2e–g. **C**, Top: Histograms of the information content carried by individual cells' activity about position on the inferred torus during REM (left) and SWS (right). Counts (fractions of the cell sample) are shown as a function of information content (in bins of 0.28 bits/spike) for all grid modules (colour-coded). The vertical dashed line (close to zero) shows mean information content for shuffled distributions (*n* = 1,000 shuffles). The majority of cells have a higher information content. Bottom: Explained deviance of a GLM model fitted to the spike count with toroidal coordinates during REM (left) and SWS (right) as regressor. Distributions show counts (fractions of the cell sample) as a function of explained deviance, in bins of 0.035, for all grid modules. Values larger than 0 indicate that the fitted model explains the data better than a null model that assumes a constant firing rate.

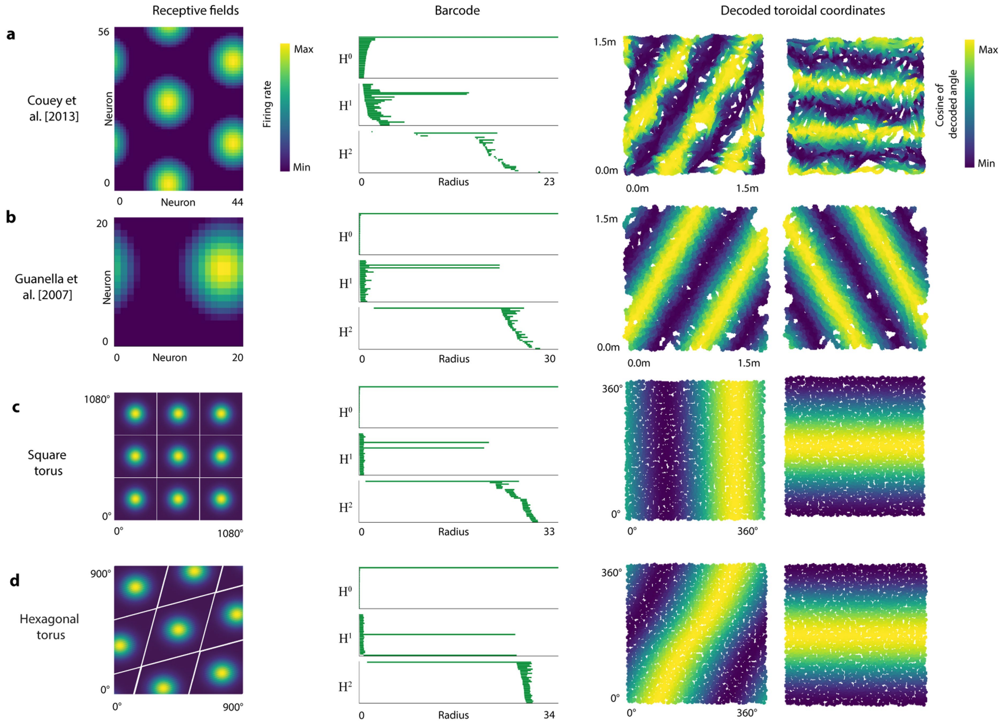

Extended Data Fig. 7 | Barcodes and decoding of simulated firing activity for two grid-cell CAN models (with no noise), and for two point clouds randomly sampled on a hexagonal and a square torus. a, Persistent cohomology analysis of a simulated grid-cell network based on the CAN model from Couey et al (2013)[11] during OF foraging. Left: Colour-coded firing rates for a single time frame of the 56 × 44 grid cells, shown at their respective positions on the neural sheet. Middle: Barcode of the simulated data. Arrows point to one 0D, two 1D and one 2D bar with long lifetimes, indicating toroidal structure. Right: Each coordinate of the toroidal parametrization of the two longest lived 1D features is mapped onto the spatial trajectory, colour-coded by its cosine value (as in Extended Data Fig. 5a, f). The resulting striped patterns of the two maps are oriented approximately 60 degrees relative to each other, as expected from a hexagonal torus network structure (see **d**). **b**, Analysis of a random sample of 100 grid cells (of a total of 400 cells) of a simulated grid cell network, using the twisted torus CAN model formulated by Guanella et al (2007)[10]. Left: Firing rates of the cells in the network at a single time frame. The model generates a single bump of activity based on both inhibitory and excitatory, asymmetric connections representing a twisted torus. Barcode (middle) and cohomological decoding of toroidal position (right) are shown as in **a**. The barcode shows four prominent bars: one 0D bar, two 1D bars and one 2D bar, similar to that of a torus. Note that the pair of stripes in toroidal coordinates are oriented 60 degrees relative to each other. **c**, **d**, To verify the expected barcodes and decoding of a torus and compare with both real and synthetic grid cell data, we performed the same topological analysis on point clouds sampled from two idealized toroidal parametrizations ($n = 2{,}500$ points): a 4D description of a square torus (**c**) and a 6D embedding of a hexagonal torus (**d**). Left: Representing the firing of a cell as a Gaussian function centred at a single toroidal coordinate on the toroidal sheet results in a square (**c**) and hexagonal (**d**) firing pattern, when arranged to tesselate a 2D surface. Middle: The expected barcode of a torus (one 0D, two 1D, and one 2D bar clearly longer than the other bars) is seen in both cases. Right: each sampled angle is coloured according to the decoded toroidal coordinates. Note the difference in the relative angle of the pair of stripes between the square and the hexagonal torus.

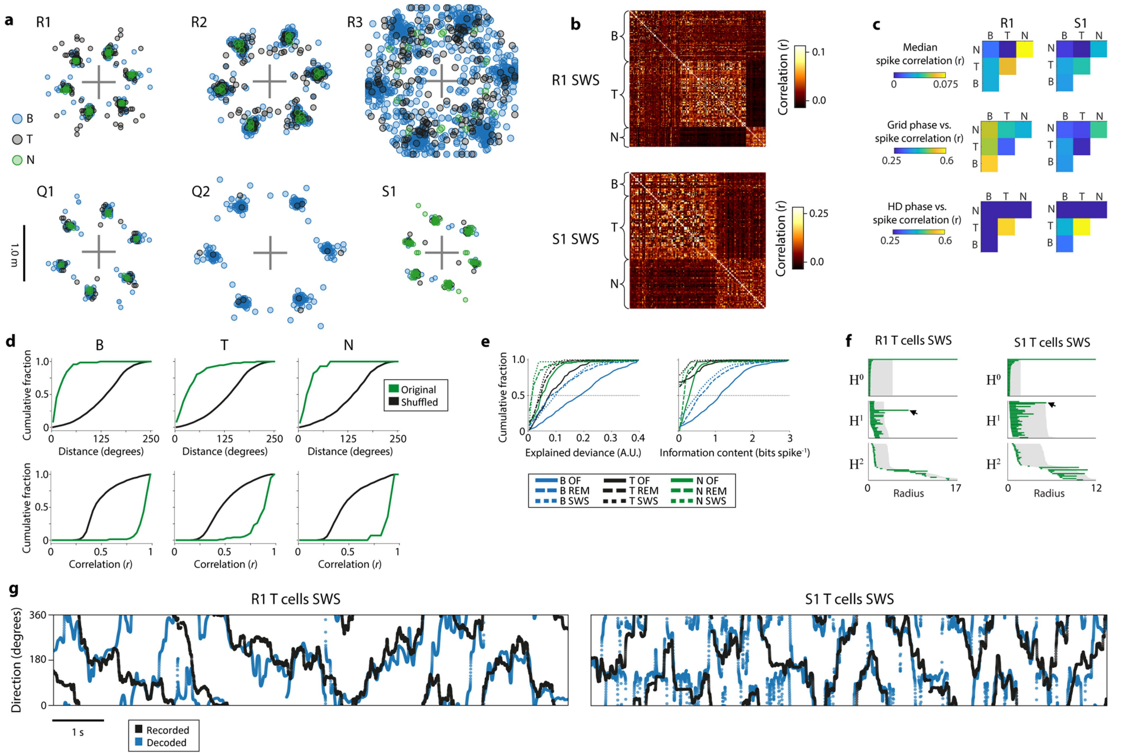

Extended Data Fig. 8 | Subpopulations of grid cells with different temporal spiking statistics have different degrees of toroidal selectivity. **a**, Geometry of grid-cell pattern of all six modules with classes of grid cells (B, bursty; T, theta-modulated; N, non-bursty; as defined in Fig. 4). Each plot shows the locations of the innermost six peaks of the spatial autocorrelogram for every grid cell in one module. Each dot indicates the position of one peak from one cell (total of 6 dots per cell); dots are coloured by the cell's class. The grey crosshair indicates the centre of the autocorrelogram. **b**, Correlation matrix showing pairwise correlation of firing rates for all grid cells belonging to S1 (left; *n* = 73 cells) and R1 (right – same data as for autocorrelogram distance matrix in Fig. 4b; *n* = 111 cells). Correlation is colour-coded according to the scale bar, with minimum and maximum defined as the 1st and 99th percentile, respectively, of the pairwise correlation distribution for each module. Rows and column (cells) are ordered according to class, as assigned by the clustering analysis shown in Fig. 4. Each cluster displays strong inner correlation structure for both modules during SWS. Cluster boundaries are indicated on the *x* axis of the correlation matrix. **c**, Summary of pairwise correlations of SWS activity for grid cells in modules R1 and S1, shown according to cell class. In each matrix plot, rows and columns indicate cell classes, and each element represents all pairs of grid cells from the classes corresponding to the row and column. Matrix elements are colour-coded to represent (top) the median of the spike train Pearson correlation *r* value across all cell pairs, (middle) Spearman rank correlation between cell pairs' grid (toroidal) phase offsets and their spike train Pearson correlation *r* values, (bottom) same as middle, but for head-direction phase instead of grid phase. Number of cell pairs were as follows: module R1, B-B 2346, B-T 6348, B-N 1932, T-T 4186, T-N 2576, N-N 378; module S1 B-B 378, B-T 1680, B-N 1456, T-T 1770, T-N 3120, N-N 1326. Note that, in agreement with the topological analyses, the correlation between cell pairs' grid phases and their spike-time correlations are weaker for theta-modulated cells than non-bursty and particularly bursty cells. This drop is explained by an increase in the correlation with head direction, suggesting, as expected in conjunctive cells, that head direction accounts for much of the variation in these cells, unlike the other classes. Furthermore, the median spike correlation for pairs of theta-modulated and non-bursty cells is higher than for bursty cells, indicating a stronger positive correlation bias, consistent with more global fluctuations of activity in these populations. **d**, Cumulative distributions showing distance between toroidal field centres (upper) and Pearson correlation *r* values (lower) for toroidal rate maps of grid cells in each class as in Fig. 2f, g, but here comparing awake behaviour in OF with SWS, *n* cells = 523(B), 229(T) and 95(N) cells for OF and 495(B), 169(T), 43(N) cells for REM and SWS. *n* = 1,000 shuffles. **e**, Cumulative distributions showing toroidal explained deviance (left) and information content (right) for all grid cells in each class – bursty (B), theta-modulated (T) and non-bursty (N) – and for each of three conditions – OF, REM and SWS. Cells are from all modules. *n* cells as in **d**. **f**, Barcode of T-class grid cells from modules R1 (left; *n* = 92 cells) and S1 (right; *n* = 60 cells) during SWS reveals a single prominent long-lived H^1 bar (indicated by black arrow). **g**, Cohomological decoding of the longest-lived H^1 bar in each barcode in **f** reveals strong correlation with recorded head direction. Recorded head direction (black) and decoded direction (blue) are shown as a function of time (total snippet length 10 s).

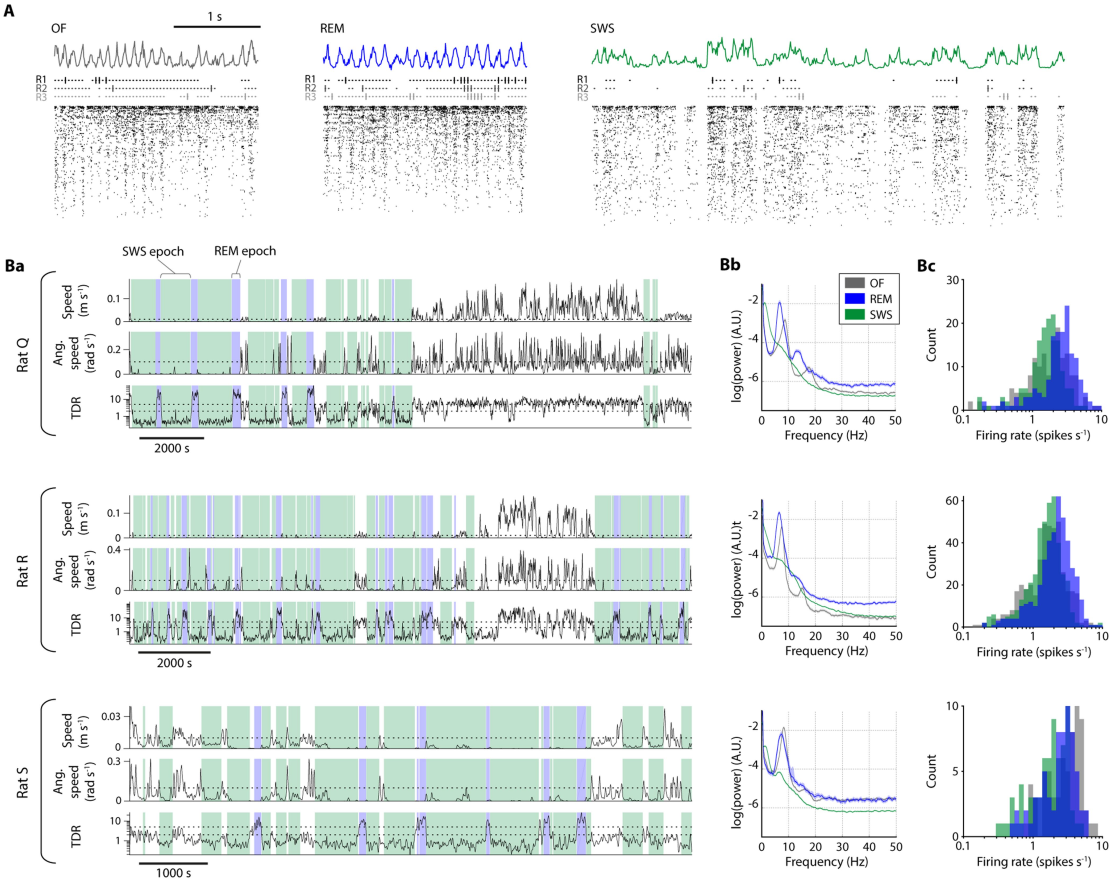

C

	Rat Q	Rat R	Rat S
Total time SWS (s)	6300	8218	5259
Total time REM (s)	1052	1857	588
Median bout length SWS (s)	332	178	190
Median bout length REM (s)	175	79	96

Extended Data Fig. 9 | Classification of sleep and wake states based on behavioural and neural activity during rest sessions. A, Example traces of MEC multi-unit activity (upper; coloured lines), and rasters of spike times of 444 grid cells (lower; black dots) recorded from rat 'R' during OF foraging, REM sleep and slow-wave sleep (SWS). Cells are ranked from top to bottom by the number of spikes fired during the example time window. Note the presence of regular theta waves (5–10 Hz) during OF and REM, and presence of slower, more irregular fluctuations between active "up-states" and silent "down-states" during SWS. Middle: times of population activity vectors (calculated in 10 ms time bins) which were selected for persistent cohomology analysis, for each module (R1-R3). Each dot indicates a vector which was included in the initial downsampled set of 15,000 vectors with the highest mean firing rate across cells in the module. Vertical ticks indicate the subset of these vectors which were retained after using a density-based method to reduce the data to a representative point cloud. Note that during SWS, all of the selected population activity vectors occurred during up-states. **B**, Classification of sleep/wake states based on behavioural and neural activity during rest sessions. Each of the three horizontal blocks shows a recording from one animal. Rat 'R' day 1 did not contain a rest session and is not shown on this figure. **Ba**, Detection of REM and SWS sleep epochs in the rest session. The plots show the time courses of the three variables used for detecting REM and SWS epochs. Top panel of each block: animal locomotion speed; middle panel: the animal's head angular speed; bottom panel: the ratio of the amplitude of theta (5–10 Hz) and delta (1–4 Hz) frequency bands in the multi-unit spiking activity (theta/delta ratio, TDR). **Bb**, Log-power spectra of MEC multi-unit activity during each sleep/wake state. The line and shaded area indicate the mean and 95% bootstrap confidence intervals, calculated across time windows (confidence intervals are narrow). Note the pronounced peak corresponding to the theta band (5–10 Hz) during OF and REM, and the higher power in the delta band (1–4 Hz) during SWS. **Bc**, Histograms showing distributions of firing rates for all grid cells during each sleep/wake state (number of grid cells: rat 'Q' 159, rat 'R' 428, rat 'S' 72). **C**, Table showing total time and median bout length of recorded sleep for each animal.

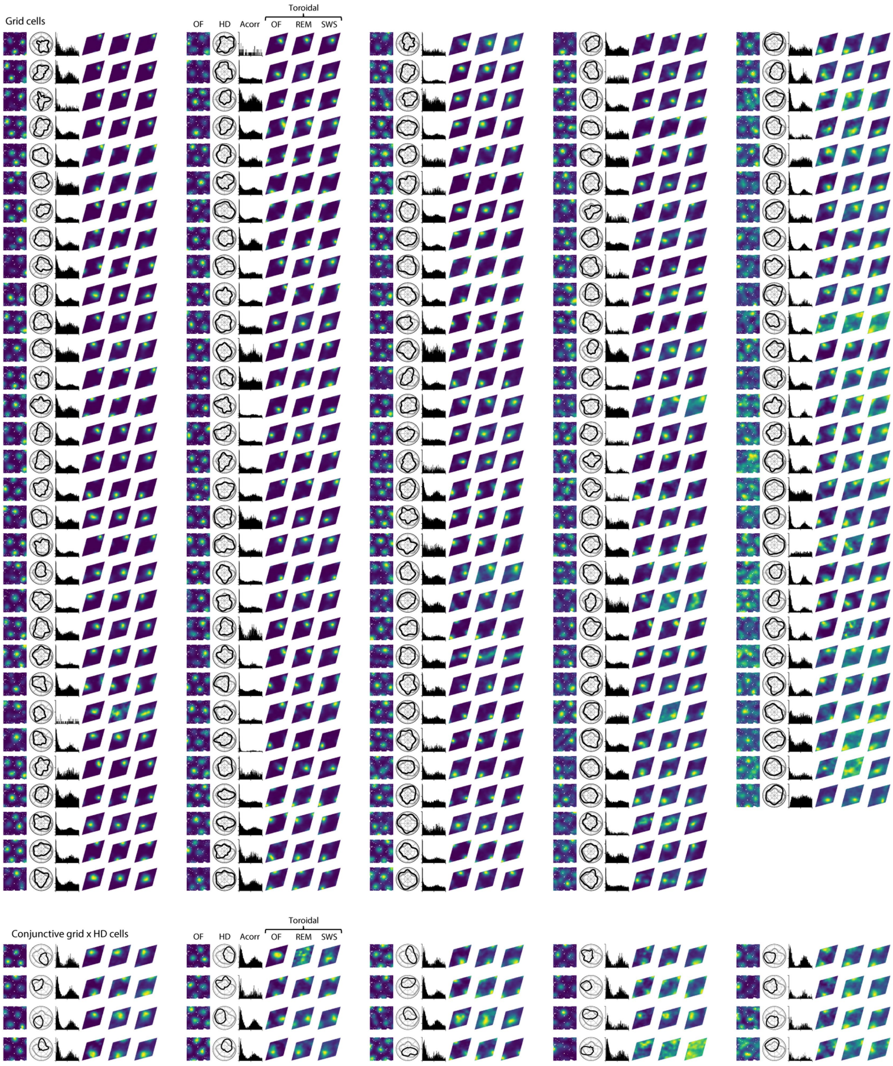

Extended Data Fig. 10 | Tuning to coordinates in space and on the inferred torus for all grid cells of module R2 (separated into pure and conjunctive categories) on recording day 2. Plots show all 152 cells in module R2, a subset of which is shown in Fig. 3b. Plots from left to right: OF firing rate map, head-direction tuning curve (black) compared to occupancy of head directions (light grey), temporal autocorrelogram, toroidal firing rate maps for OF, REM and SWS. The full set of plots, for all remaining grid cells of all recordings, is shown in Supplementary Information.

nature research

Corresponding author(s): Richard J. Gardner
Nils A. Baas
Benjamin A. Dunn
Edvard I. Moser

Last updated by author(s): Oct 07 2021

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a | Confirmed

☐ ☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement

☐ ☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly

☐ ☒ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.

☐ ☒ A description of all covariates tested

☐ ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons

☐ ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)

☐ ☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted
Give P values as exact values whenever suitable.

☒ ☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings

☒ ☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes

☐ ☒ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Commercial software: Motive (OptiTrack) version 2.2.0; MATLAB (MathWorks) version r2019b Open-source software: SpikeGLX (https://billkarsh.github.io/SpikeGLX) versions 20190724 and 20190919
Data analysis	Commercial software: MATLAB (MathWorks) version r2020b, Python version 3.7 Open-source code (for MATLAB): Kilosort version 2.5 (https://github.com/MouseLand/Kilosort) UMAP version 1.4.1 (https://www.mathworks.com/matlabcentral/fileexchange/71902-uniform-manifold-approximation-and-projection-umap) Chronux version 2.10 (http://chronux.org/) Open-source Python packages: umap 0.3.10 ripser 0.4.1 numba 0.48.0 scipy 1.4.1 numpy 1.18.1 scikit-learn 0.22.1 matplotlib 3.1.3 h5py 2.10.0 gudhi 3.4.1.post1

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The datasets generated during the current study will be available after publication at https://figshare.com/articles/dataset/Toroidal_topology_of_population_activity_in_grid_cells/16764508

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Samples included all available cells that matched the classification criteria for the relevant cell type.
Data exclusions	Cells with very low firing rates (below 0.05 Hz) were excluded because of their unsuitability for spike-train analysis. All non-grid cells were excluded, because they were irrelevant for analyses of toroidal structure.
Replication	For the six grid modules included in the study, in the results text we indicate for each result the number of modules in which the effect was found. For each statistical test we state the sample size (n) in the manuscript.
Randomization	The study did not involve any experimental subject groups; therefore, random allocation did not apply and was not performed.
Blinding	The study did not involve any experimental subject groups; therefore, experimenter blinding did not apply and was not performed.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study	
☒	☐	Antibodies
☒	☐	Eukaryotic cell lines
☒	☐	Palaeontology and archaeology
☐	☒	Animals and other organisms
☒	☐	Human research participants
☒	☐	Clinical data
☒	☐	Dual use research of concern

Methods

n/a	Involved in the study	
☒	☐	ChIP-seq
☒	☐	Flow cytometry
☒	☐	MRI-based neuroimaging

Animals and other organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research

Laboratory animals	Long Evans rats, male, age 3-4 months (300-500 g)
Wild animals	None
Field-collected samples	None
Ethics oversight	Protocols approved by the Norwegian Food Safety Authority (FOTS ID 18011 and 18013) .

Note that full information on the approval of the study protocol must also be provided in the manuscript.

nature portfolio

https://doi.org/10.1038/s41586-021-04268-7

Supplementary information

Toroidal topology of population activity in grid cells

nature portfolio

https://doi.org/10.1038/s41586-021-04268-7

Supplementary information

Toroidal topology of population activity in grid cells

Supplementary Data

Tuning to coordinates in space and on the inferred torus for all grid cells of all modules (separated into ‘pure’ and ‘conjunctive’ categories), except R2 day 2 (Extended Data Fig. 10). Subsets of these plots are shown in Fig. 2b and 4h. Plots from left to right: open-field firing rate map, head-direction tuning curve (black) compared to occupancy of head directions (light gray), temporal autocorrelogram, wagon-wheel firing rate map (depending on recording session), toroidal firing rate maps for OF, WW, REM and SWS (where available)


R1 day 1
OF
HD
Acorr
WW
Toroidal
OF
WW

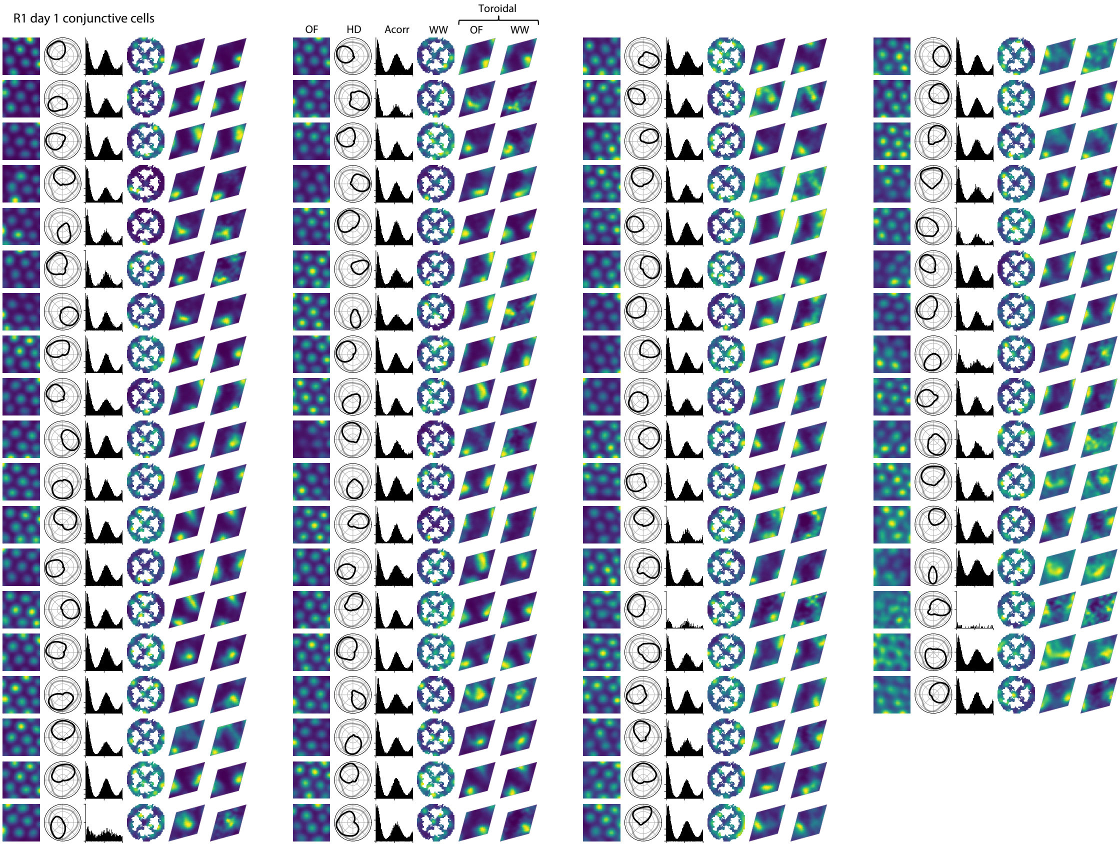
R1 day 1 conjunctive cells


R1 day 2
Toroidal
OF
HD
Acorr
OF
REM
SWS

R1 day 2 conjunctive cells

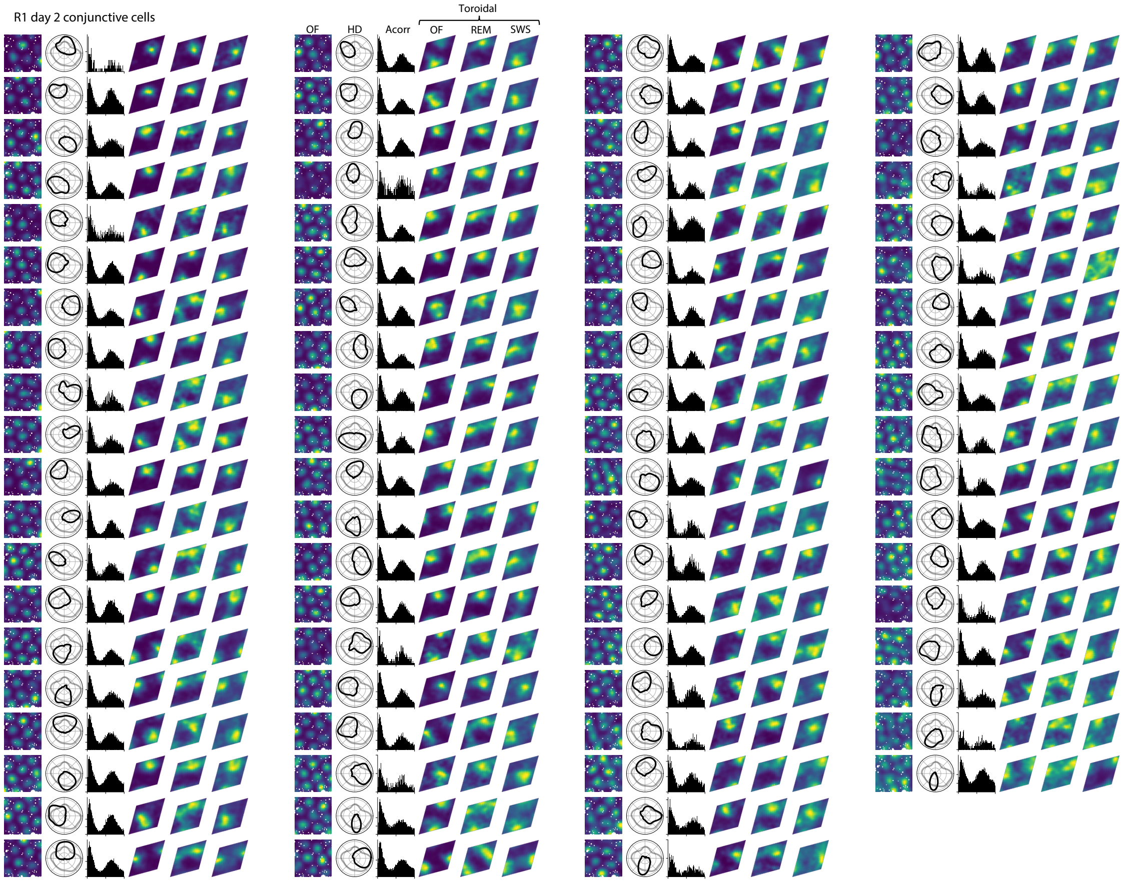

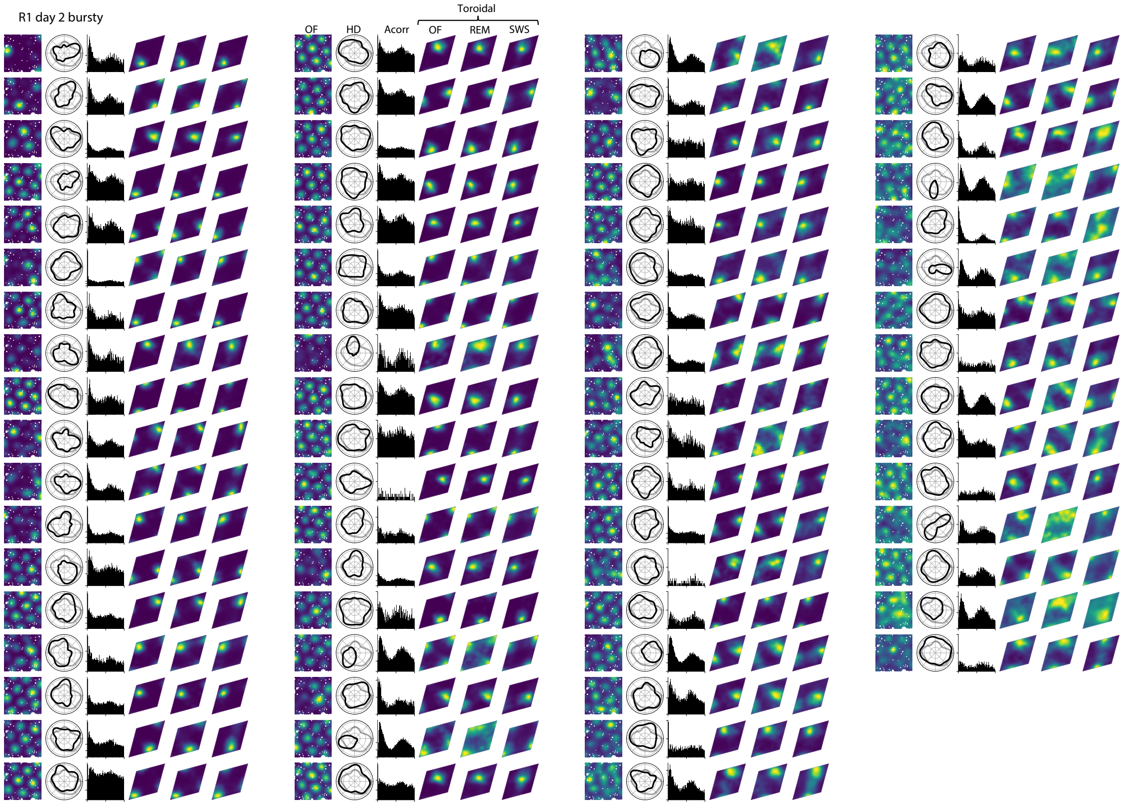
R1 day 2 bursty
Toroidal
OF
HD
Acorr
OF
REM
SWS


R1 day 2 theta-modulated
Toroidal
OF
HD
Acorr
OF
REM
SWS

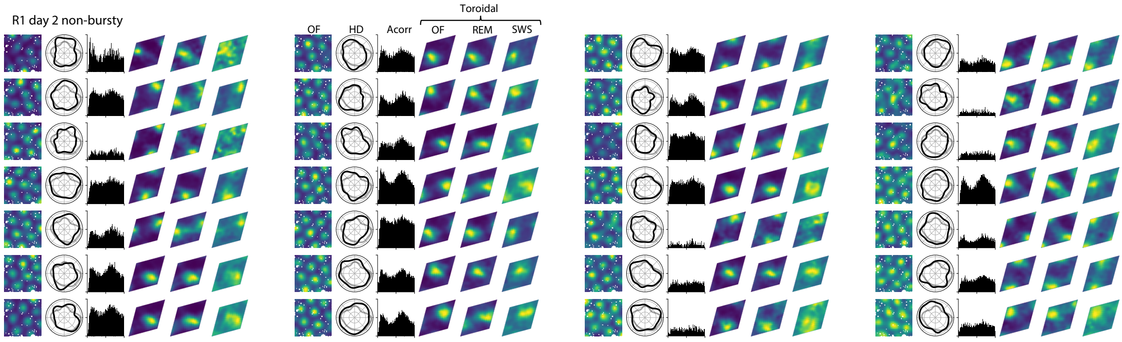
R1 day 2 non-bursty
Toroidal
OF
HD
Acorr
OF
REM
SWS


R2 day 1
Toroidal
OF
HD
Acorr
WW
OF
WW

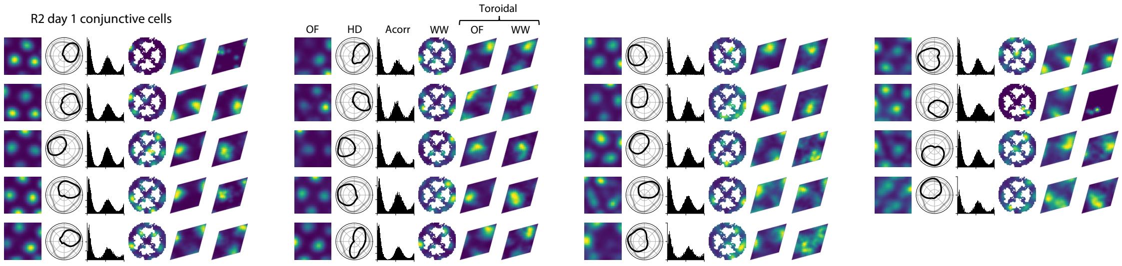
R2 day 1 conjunctive cells
Toroidal
OF
HD
Acorr
WW
OF
WW

R3 day 1

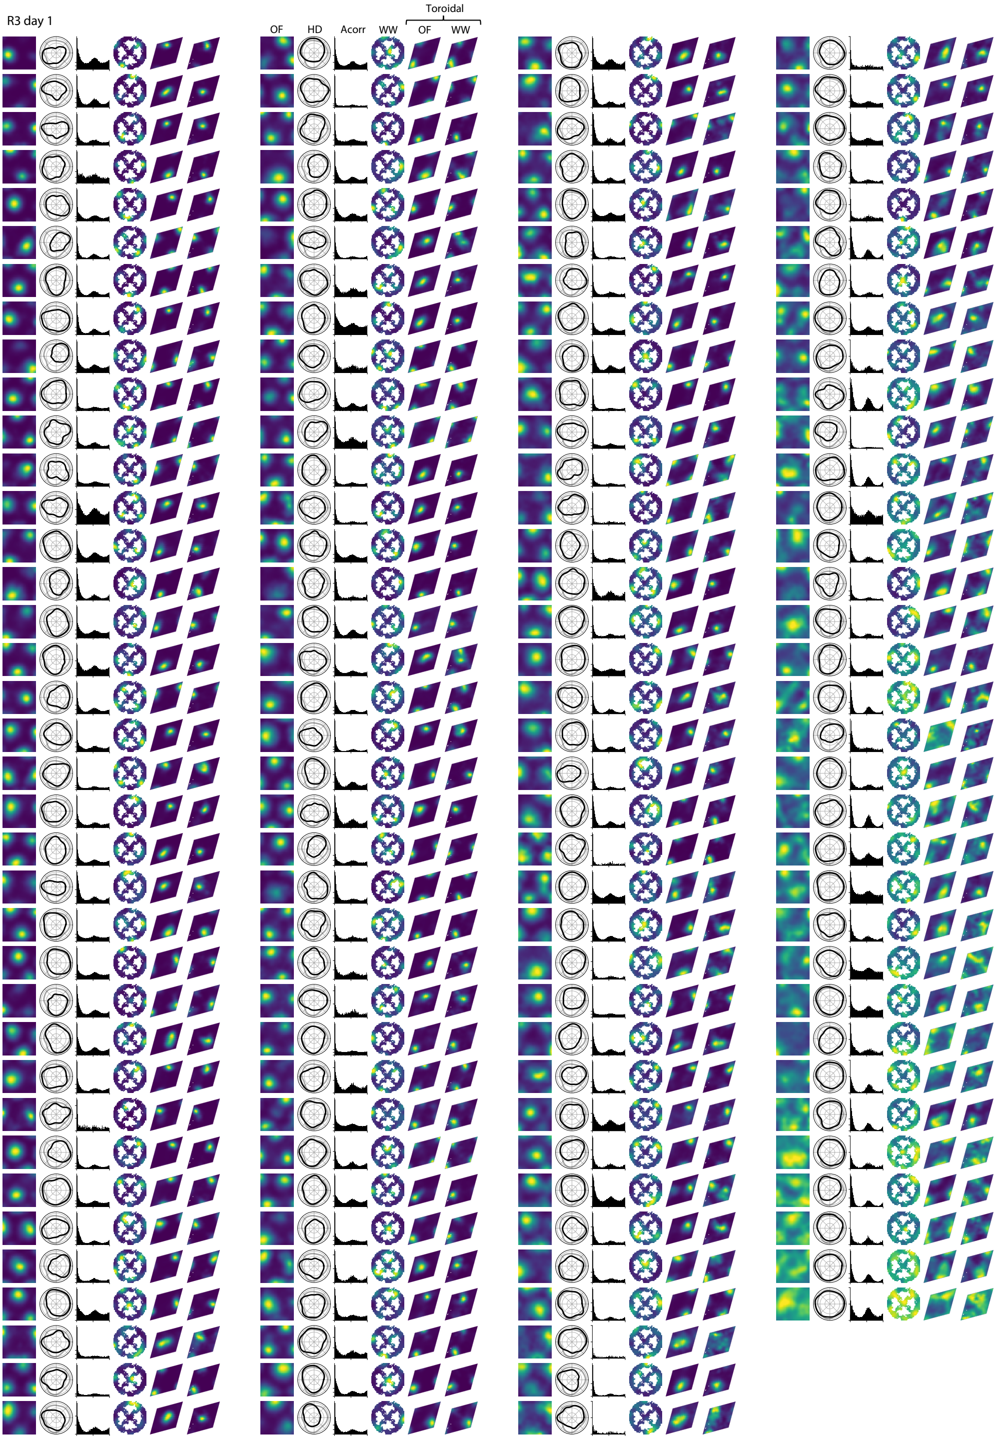

R3 day 1 conjunctive cells

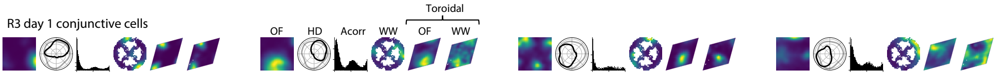


R3 day 2
Toroidal
OF
HD
Acorr
OF
REM
SWS

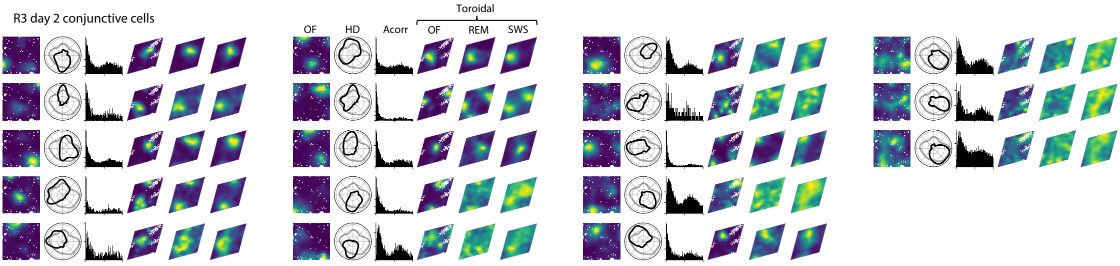
R3 day 2 conjunctive cells
Toroidal
OF
HD
Acorr
OF
REM
SWS

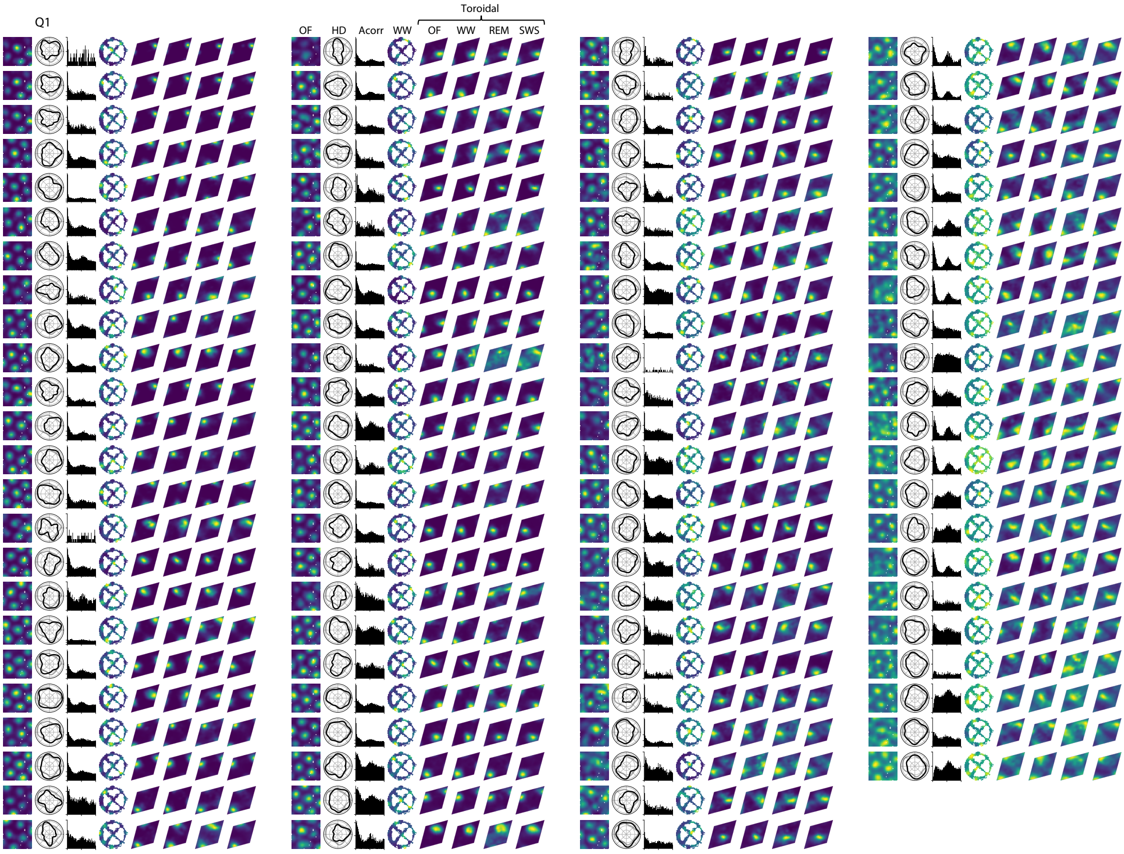
Q1
OF
HD
Acorr
WW
Toroidal
OF
WW
REM
SWS

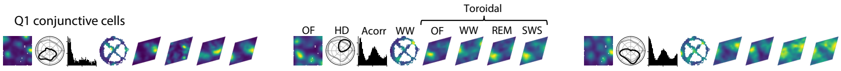
Q1 conjunctive cells
OF
HD
Acorr
WW
Toroidal
OF
WW
REM
SWS

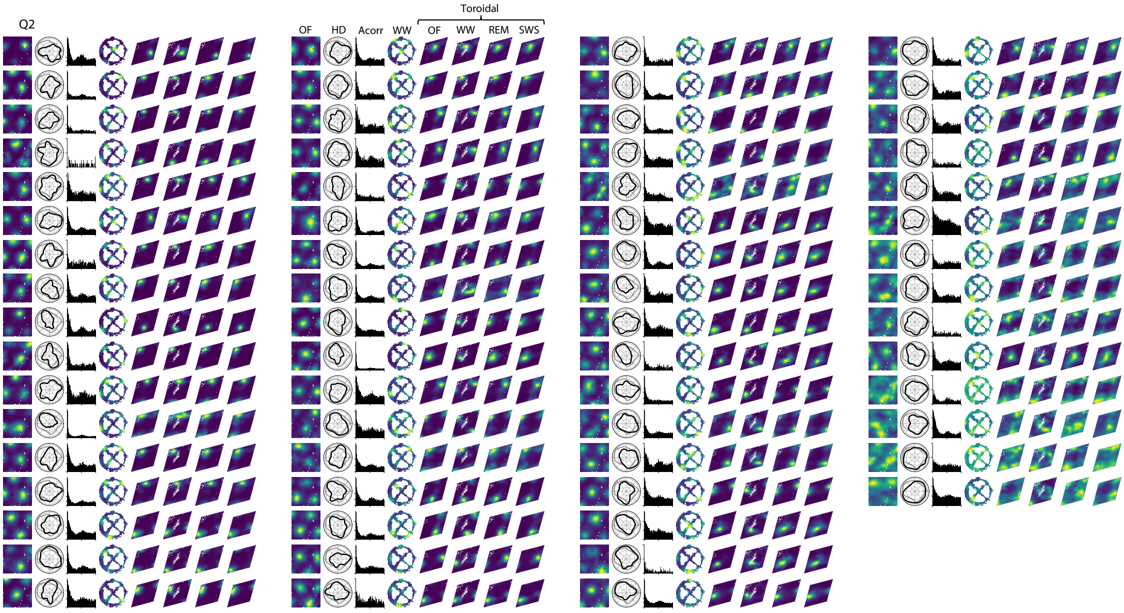
Q2
Toroidal
OF
HD
Acorr
WW
OF
WW
REM
SWS

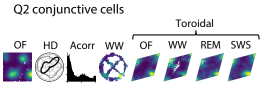
Q2 conjunctive cells
Toroidal
OF
HD
Acorr
WW
OF
WW
REM
SWS

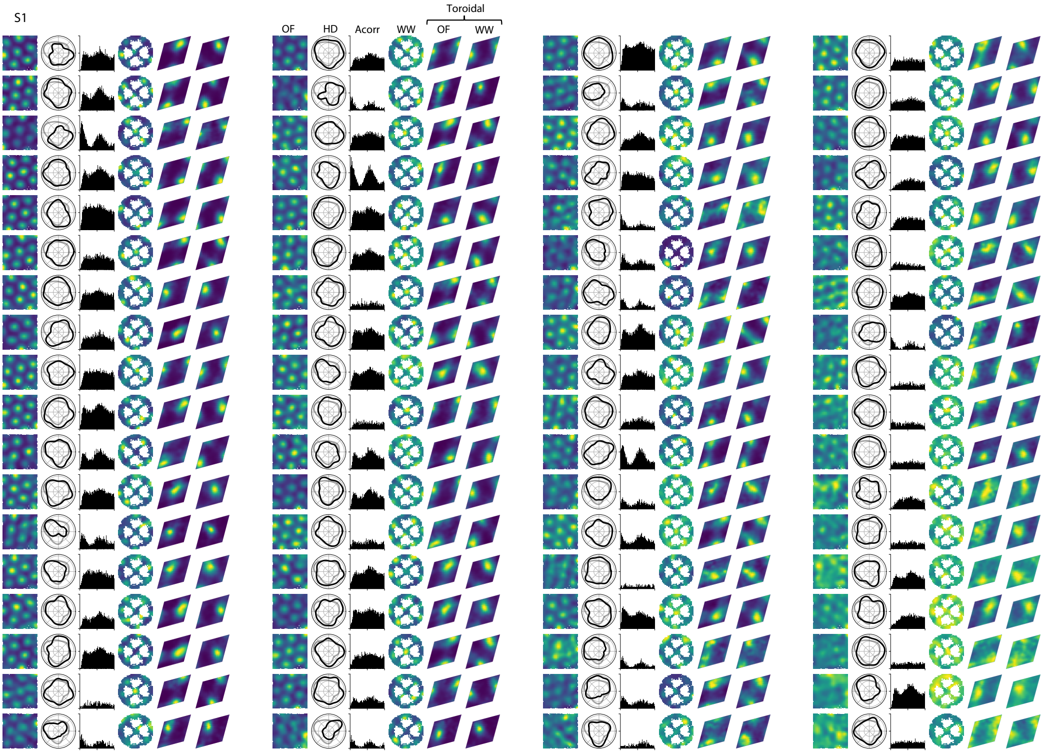
S1

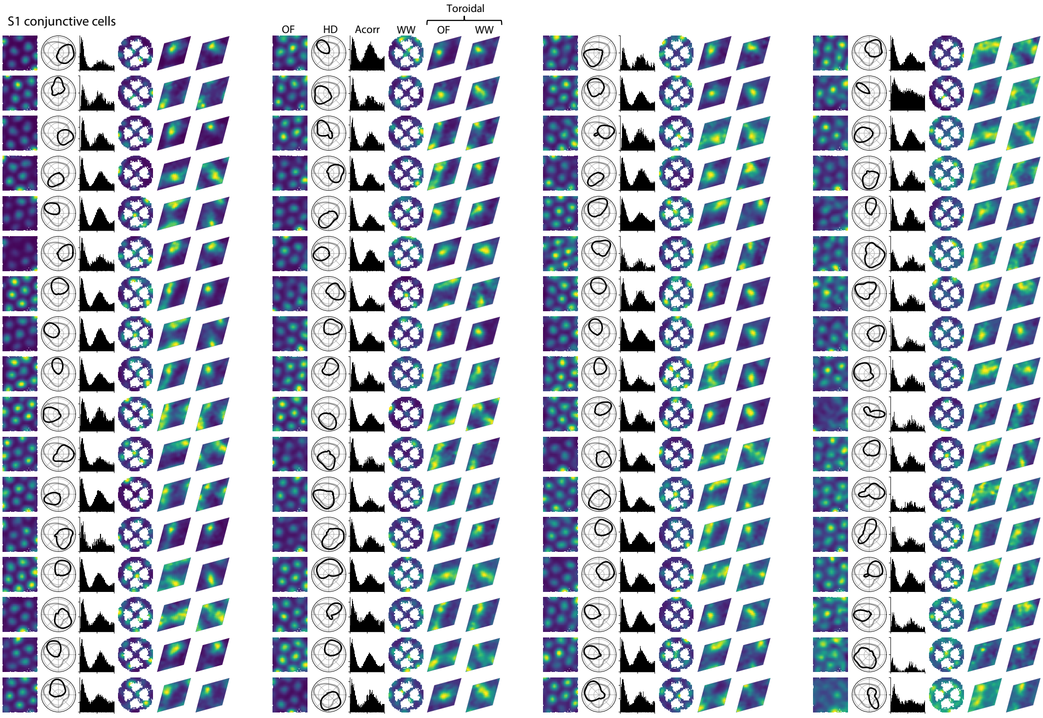
S1 conjunctive cells

Supplementary Methods

Theoretical explanation of the six-dimensionality proposed by PCA

To understand why a minimum of six PCA components is necessary in order to account for a large fraction of the variance in the population patterns of grid cells, we shall consider an idealized model of grid cell firing, with the following three assumptions:

First, grid cell population activity patterns lie on a two-dimensional manifold with toroidal topology. This 2-D surface can be mapped to a rhombus with an angle of 60 degrees and periodic boundaries. Specifically, we can parametrize positions on the rhombus in the form $\alpha_1 \vec{u}_1 + \alpha_2 \vec{u}_2$, where $\vec{u}_1$, and $\vec{u}_2$ are unit vectors whose orientations differ by 60 degrees, and $\alpha_{1,2} \in [0,1]$.

Second, the tuning of an individual grid cell to the toroidal coordinates (the position on the rhombus) is identical in all cells, up to a translation:

$$r_i = \phi\left(\left[\vec{\theta} - \vec{\theta}_i\right]_p\right)$$

where $\vec{\theta}$ is the 2-D position on the rhombus, $\vec{\theta}_i$ is the center of the receptive field of cell i, $\left[\vec{\theta} - \vec{\theta}_i\right]_p$represents a shift of $\vec{\theta}$ by $\vec{\theta}_i$ with periodic boundary conditions on the rhombus and ϕ represents the structure of the tuning function.

Third, activity patterns of the neural population uniformly sample the states that correspond to different positions on the rhombus.

Under these conditions, the covariance matrix C of neural activity has the structure

$$C_{ij} = c\left(\left[\vec{\theta}_i - \vec{\theta}_j\right]_p\right).$$

If we further assume that receptive field centers $\vec{\theta}_i$ uniformly and regularly sample the rhombus, the covariance matrix commutes with periodic rigid translation operators on the rhombus. A full basis of eigenvectors of C can then be obtained such that the eigenvectors are also eigenvectors of the translation operators. Thus, these eigenvectors are Fourier modes of the form

$$\psi_i[\vec{k}] = \mathrm{A}\exp\left(i\,\vec{k} \cdot \vec{\theta}_i\right),$$

where the wavevector $\vec{k}$ must be selected such that ψ is periodic on the rhombus. To achieve this requirement, $\vec{k}$ must be a vector lying on a vertex of a triangular lattice: the reciprocal of the lattice with basis vectors $\vec{u}_1, \vec{u}_2$.

The eigenvalues of C are the corresponding Fourier transform components of c. Typically, for a unimodal tuning curve, these eigenvalues will be a monotonically decreasing function of $|\vec{k}|$ for $|\vec{k}| > 0$. Furthermore, if the tuning of individual cells is isotropic, the eigenvalues depend only on $|\vec{k}|$. Therefore, the six PCA modes that correspond to the smallest value of $|\vec{k}|$ must contribute equally to the variance. If the tuning curve of individual cells is sufficiently wide, the eigenvalues are expected to decay rapidly with $|\vec{k}|$, and the modes that correspond to the first six PCA components are expected to capture a large fraction of the variance.

Theoretical background to persistent (co)homology

As persistent cohomology is a cornerstone of the current analyses, we wish to elaborate on its theoretical background to elucidate how and why it works (broader perspectives on topological data analysis and its application in biology are available elsewhere[57,58,93]). We start by introducing (co)homology and the Vietoris-Rips complex before we turn to persistent homology.

Given a topological space X, we can assign, for all natural numbers n, a vector space $H_i(X)$, called the i-th homology group of X, such that if $f: X \to Y$ is a continuous map between topological spaces X and Y, then $f_*: H_i(X) \to H_i(Y)$ is a linear map between vector spaces. Similarly, we may define cohomology groups, $H^i(X)$ by reversing arrows. These are dual notions and give the same results in our case. We will thus only continue describing the former. The dimension of $H_i(X)$ is called the i-th Betti number, representing the number of i-dimensional holes (for further details see Hatcher, 2002[60]). However, this may vary depending on the choice of coefficients of the vector space (note that it is choosing algebraic fields as coefficients that makes the homology groups vector spaces). For example, using $\mathbb{Z}_2$-coefficients, the Klein bottle will have the same homology as a torus. To separate these, we use $\mathbb{Z}_{47}$-coefficients in our computations. The choice of field coefficients simplifies computations at the risk of losing topological information known as *torsion*, measuring the orientability of a space. However, the same number of holes (Betti numbers) is computed, which is what is here used to distinguish spaces.

As a point cloud is finite and discrete, its homology only returns the number of points in the point cloud (its 0-th Betti number). Thus, we associate combinatorial spaces known as simplicial complexes to the point cloud which may have non-trivial topology reflecting interesting structure and information of the data set and whose homology is easy to compute. A simplicial complex is a set V of vertices and a set S of finite non-empty subsets of V called simplices such that any vertex is a simplex and any non-empty subset of a simplex is a simplex. A simplex of cardinality $p + 1$ is referred to as a p-simplex (of simplicial dimension p) and geometrically, we may refer to a 0-simplex as a point, a 1-simplex as an edge, a 2-simplex a triangle, a 3-simplex a tetrahedron and so on in higher dimensions.

There are different choices in constructing simplicial complexes associated with the data. We used what is known as the Vietoris-Rips complex, here denoted R_r. The vertices of the Vietoris-Rips complex are the points in the point cloud and the simplices are the sets of points whose pairwise distance is less than the scale value, r. This is equivalent to replacing each point by a ball of common radius r and connecting two points with an edge if their balls intersect. p-simplices are then formed if each point of a subset of $p + 1$ points have edges to all other points in the subset (i.e. a $p + 1$-clique). Although this simplicial complex is not homotopy equivalent to taking the union of balls (as e.g. the Čech complex is) and thus does not necessarily have the same homology, the basic topological information is preserved under this correspondence.

One way to construct the Vietoris-Rips complex in detecting the topology of neural data is to regard individual cells as points and their pairwise dissimilarity (e.g. correlation) as scale. In our case, we rather considered the population activity vectors as the points. There is a subtle correspondence between these constructions, where the resulting barcodes are the same when applying persistent cohomology. This relationship is only valid when the tuning of the cells is such that the response is convex, seemingly invalidated by the firing patterns of grid cells with respect to its physical position in the environment. However, when considering the tuning to be a function of the toroidal state space, we find it indeed to be convex (indicated by the single bumps in the toroidal rate maps for each grid cell – Fig 3a, Extended Data Fig. 9)[94], suggesting both constructions should give rise to the same barcodes.

We consider the nested chain of Vietoris-Rips complexes, R, constructed for all increasing values of r in which new simplices are formed:

$$R_{r_0} \subset R_{r_1} \subset \cdots \subset R_{r_n},$$

where $r_0 = 0$ and r_n is the largest pairwise distance in the point cloud, and apply homology to get a sequence of vector spaces and maps, for all dimensions i:

$$H_i(R_{r_0}) \to H_i(R_{r_1}) \to \cdots \to H_i(R_{r_n}),$$

where the maps are induced by the inclusion maps (note that we have omitted compositions and identity maps), called the i-th persistent homology. This may again be decomposed into a sum of elementary *persistence modules*[95]:

$$H_i(R) \cong \bigoplus_k I([b_k, d_k)),$$

where $b_k < d_k$ give the scales in which a class in $H_i(R)$ first appears and later disappears. Thus, we may represent the persistent (co)homology by displaying the intervals as bars starting at b_k and ending at d_k. The collection of bars for all dimensions results in what is known as the *barcode*.

The barcodes are shown to be stable under influence of noise (given certain assumptions on the construction)[96]. This means that small perturbations to the point cloud lead to small changes in the barcode. Thus, we note that in dimensionality reduction, quantifying the dissimilarity between the barcode of the high-dimensional representation and its embedding may address the challenge of measuring how faithful an embedding is[53].

Supplementary References

93. Rabadan, R. & Blumberg, A. J. *Topological Data Analysis for Genomics and Evolution: Topology in Biology*. (Cambridge University Press, 2019). doi:10.1017/9781316671665.

94. Curto, C. What can topology tell us about the neural code? *Bull. Am. Math. Soc.* **54**, 63–78 (2017).

95. Chazal, F., De Silva, V., Glisse, M. & Oudot, S. *The structure and stability of persistence modules*. (Springer, 2016).

96. Cohen-Steiner, D., Edelsbrunner, H. & Harer, J. Stability of persistence diagrams. *Discrete Comput. Geom.* **37**, 103–120 (2007).

ISBN 978-82-326-6835-9 (printed ver.)
ISBN 978-82-326-5627-1 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (online ver.)

NTNU

Norwegian University of
Science and Technology