

INTERACTIONS BETWEEN CONTINUOUS ATTRACTORS ON MULTIPLE SCALES IN A GRID-CELL NETWORK MODEL

Avleen Sahni



Graduate School of
Systemic Neurosciences

LMU Munich



Dissertation at the
Graduate School of Systemic Neurosciences
Ludwig-Maximilians-Universität München

July, 2024

Supervisor

Prof. Dr. Andreas V.M. Herz

Computational Neuroscience – Department Biology II

Ludwig-Maximilians-Universität München

First Reviewer: Prof. Dr. Andreas V. M. Herz

Second Reviewer: Prof. Dr. Anton Sirota

External Reviewer: Prof. Dr. Christian Leibold

Date of Submission: 17.07.2024

Date of Defense: 22.10.2024

Abstract

Grid cells in the mammalian entorhinal cortex have spatial firing fields that repeat periodically, forming a hexagonal lattice. Grid cells are organized into modules, with cells in each module sharing the same spacing between fields, called the grid scale, and the same lattice orientation. Across modules, spatial scales follow an approximately geometric progression with a ratio of about 3 : 2.

This thesis addresses the possible role of intermodular connections for maintaining a consistent spatial representation across the hierarchy of modules, using continuous attractor network (CAN) models. In the absence of landmarks or other spatial cues to calibrate the firing rate maps, I show how the spatial representation in such models will invariably drift, even in deterministic CAN models with no noise. When combining information across different spatial scales, independent drift in different modules can lead to catastrophic errors. For this reason, in this study I explore the role of biologically inspired, yet weak inter-modular connections could play in canceling or reducing the relative drift. I show that when the inter-modular connections reinforce the 3:2 ratio of the spatial scales across modules the relative drift is reduced by an order of magnitude. Moreover, I show that this reduction does not occur when the inter-modular connections are randomized. Targeted inter-modular connections, therefore, are a simple yet powerful means to mitigate drift and prevent catastrophic errors in multi-scale grid codes.

Abstract

Gitterzellen im entorhinalen Kortex von Säugetieren haben räumliche Feuerfelder, die sich periodisch wiederholen und ein hexagonales Gitter bilden. Diese Zellen sind in Modulen organisiert, wobei die Zellen in jedem einzelnen Modul die gleichen Abstände zwischen den Feuerfeldern und die gleichen Gitterachsen haben. Die räumlichen Skalen zwischen den verschiedenen Modulen folgen näherungsweise einer geometrischen Progression mit einem Skalenverhältnis von 3 : 2.

Diese Arbeit befasst sich mit der möglichen Rolle intermodularer Verbindungen für die Aufrechterhaltung einer konsistenten räumlichen Repräsentation über die Hierarchie der Module hinweg, wobei Modelle eines kontinuierlichen Attraktorennetzwerks (CAN) verwendet werden. In Ermangelung von Landmarken oder anderen räumlichen Anhaltspunkten zur Kalibrierung der Feuerratenkarten zeige ich, wie die räumliche Repräsentation in solchen Modellen unweigerlich driftet, selbst in deterministischen CAN-Modellen ohne jegliches Rauschen. Wenn Information über verschiedene räumliche Skalen hinweg kombiniert wird, kann eine unabhängige Drift in verschiedenen Modulen zu katastrophalen Fehlern bei der Ortsbestimmung führen. Aus diesem Grund untersuche ich in dieser Studie, welche Rolle biologisch inspirierte, schwache intermodulare Verbindungen bei der Reduzierung der relativen Drift spielen können.

Ich zeige, dass die relative Drift um eine Größenordnung reduziert wird, wenn die intermodularen Verbindungen das $3/2$ -Verhältnis der räumlichen Skalen zwischen den Modulen verstärken. Außerdem zeige ich, dass diese Verringerung nicht auftritt, wenn die intermodularen Verbindungen randomisiert sind. Gezielte intermodulare Verbindungen sind daher ein einfaches, aber wirkungsvolles Mittel, um die Drift zu mindern und katastrophale Fehler in mehrskaligen Gitterzellen-Darstellungen zu verhindern.

Contents

1	INTRODUCTION	6
1.1	Aim and scope of the thesis	10
1.2	Outline of thesis	12
2	LITERATURE REVIEW	15
2.1	Grid cells in the entorhinal cortex	16
2.2	Relationship of grid cells to other hippocampal spatial cells	21
2.3	Environmental and behavioral modulation	23
2.4	Continuous attractor models	28
2.5	Oscillatory interference models	33
2.6	Hybrid interference and attractor models	35
2.7	Single cell plasticity models	36
2.8	Example CAN model: Burak and Fiete (2009)	36
2.9	Open Questions	38
3	MODIFIED BURAK AND FIETE MODEL	41
3.1	Problems with the original model of Burak and Fiete	42
3.2	The solution: a four-sheet model	49
3.3	Decoding	52
3.4	Discrete attractor states	56
3.5	Response to constant input velocities	66
4	MODEL RESPONSE TO RANDOM TRAJECTORIES	69
4.1	Inputs and decoding	70
4.2	Model Dynamics	72
5	INTERMODULAR CONNECTIONS	81
5.1	A geometrical hypothesis	82
5.2	Correlations between networks	83

5.3	Connectivity scheme	84
5.4	Fixed Points in a two-module network	86
5.5	Multi-module model and drift	88
5.6	Effect of coupling	91
6	DISCUSSION	97
6.1	Conclusions	98
6.2	Relation to existing models	99
6.3	Biological significance	102
6.4	Limitations of our approach	105
6.5	Further work	105
	REFERENCES	107

1

Introduction

We initially feel disoriented when we visit a new environment, such as a foreign city. To navigate to a restaurant or museum, we rely on printed maps from tourist kiosks or electronic maps on our smartphones. However, as we become more familiar with the environment, we start to recognize previously encountered routes and landmarks, eventually navigating back to our hotel using shortcuts or alternative routes. But how is physical space

represented in the brain and used for navigation?

A well-known theory in experimental psychology suggests that humans and other animals internalize spatial experiences in the form of a ‘cognitive map’—a mental representation of space that incorporates known places and their relationships into a common reference frame (Tolman, 1948). For instance, when trying to remember where we parked our car in a large parking lot, we can mentally visualize the lot’s layout, the position of landmarks, and the relative location of our car. This visualization is an example of a cognitive map.

While a cognitive map relies on stored landmarks, animals can also keep track of location on-the-fly by path integration. This process involves continuously updating the current position vector by summing every movement’s distance and direction relative to its starting point. Thus, it enables an animal to return directly to its nest or food source even in complete darkness (Mittelstaedt and Mittelstaedt, 1982; Etienne and Jeffery, 2004). Path integration is often used when there are no visual landmarks, or when following an odor trail would be inefficient. It relies on internal (idiothetic) cues, such as proprioceptive feedback (body position and movement), vestibular information (balance and spatial orientation), and motor efference copies (parallel to signals sent from the brain to muscles).

Tolman’s cognitive map concept remained a theoretical construct until O’Keefe and Dostrovsky (1971) discovered hippocampal place cells. By recording from freely-moving rats, O’Keefe and Dostrovsky (1971) found neurons in the dorsal hippocampus that responded specifically when the rat was in a particular location facing a specific direction. These place cells were proposed as neural correlates of Tolman’s cognitive map, providing evidence for the idea that spatial representations exist in the brain.

Place cells contribute to cognitive mapping and have been implicated to play a role in path integration because they encode the animal's current location. They fire when an animal is in a specific location, forming a spatial map of the environment. While place cell activity is modulated by self-motion cues (such as proprioception, vestibular input, and motor efference) (McNaughton et al., 2006; Gothard et al., 1996), it remains debated whether they themselves perform path integration or reflect computations occurring elsewhere in the navigation system (Haas et al., 2019).

The discovery of hippocampal place cells generated significant interest in systems neuroscience. It provided a way to study high-level cognitive concepts like spatial perception at a mechanistic level.

Some thirty years later, in their investigation of the inputs to hippocampal place cells, Hafting et al. (2005) discovered neurons in the medial entorhinal cortex (MEC) that each fired at multiple spatial locations with a regular triangular pattern and specific length scale and orientation. These neurons, termed grid cells, represent physical space with periodic tuning curves at multiple spatial scales. Grid cell activity is stable, persisting without external landmarks and in darkness (Hafting et al., 2005). Grid-like firing is expressed immediately in new environments Hafting et al. (2005) and Barry et al. (2007) have shown that the grid spacing can adapt, when the environment undergoes modest rescaling. Because grid cells can maintain spatially regular firing in darkness and respond to velocity input, they have been proposed to contribute to path integration.

Several mechanisms have been proposed. Continuous-attractor network (CAN) models explain grid translation as a velocity-driven shift of an activity bump across a recurrent sheet (Burak and Fiete, 2009). Alternative frameworks include oscillatory-interference

models, in which phase differences between theta-modulated oscillators generate the grid periodicity (Burgess et al., 2007), and single cell plasticity models 2.7. Using Fisher information, Mathis et al. (2012) showed that a multi-scale grid code is more efficient than a place-cell code of comparable size.

Animals possess many grid cells that span a wide range of scales with scales increasing from dorsal to ventral along the longitudinal axis of MEC (Hafting et al., 2005). These scales are not smoothly distributed but cluster around specific values (Stensola et al., 2012; Burgess et al., 2007); each of these clusters constitutes a module. Consecutive pairs of modules have consistent scale ratios, suggesting a universal scale ratio in the range of 1.4 – 1.7 (Stensola et al., 2012; Krupic et al., 2015; Burgess et al., 2007). Within a module, the grid lattices have the same spacing between nodes and nearly identical axes (Hafting et al., 2005). Across modules, grid lattices have clustered orientations (Stensola et al., 2012; Barry et al., 2007; Gardner et al., 2022).

Previous work from our lab (Mathis et al., 2012; Stemmler et al., 2015) and others (Fiete et al., 2008; Burak, 2014; Sreenivasan and Fiete, 2011) have shown that this modular structure of the representation is advantageous because it gives the grid code a large capacity and reduces ambiguity in the decoded position. A geometric series of the grid scales is the optimal code (Mathis et al., 2012).

In figure 1.1, one can see the one dimensional (1D) analog of 3 modules with different scales and a periodic code representing position. The inferred position is given by finding the location where neurons in all modules overlap. In the absence of sensory cues, the phase represented in each module will probably drift gradually (Hardcastle et al., 2015), relative to the value that matches the true position of the animal. These drifts may occur

independently in different modules, leading to false position estimates (Burak, 2014; Fiete et al., 2008; Mosheiff and Burak, 2019).




Figure 1.1: In a one dimensional (1D) analog of the grid cell code, all neurons in a single module represent position as a 1D phase, relative to the grid spacing. The thin black bars represent the encoded phase in each module, and the red, blue and green bars represent the possible locations compatible with the activity in each module. The width of the colored bars schematically represents the accuracy of readout. Combining the activity from all neurons in the three modules can greatly reduce the ambiguity. Here, the joint activity from the three modules is compatible with a single location (yellow bar) within the range shown. The full range over which positions can be unambiguously inferred is the capacity of the code, whereas the width of the yellow bar roughly represents the resolution of the code. Figure taken from Burak (2014) with permission.

The emergence of the grid pattern and the implementation of path integration has been modeled for single modules (Burak and Fiete, 2009; Guanella et al., 2007). However, how this geometric ratio of scales emerges is not well understood and no current experimental data about the connectivity between different modules exist, other than weak intermodular correlations reported by Gardner et al. (2019). In addition, it is not known how the grid cells remain stable in the dark in the absence of any sensory cues.

1.1 AIM AND SCOPE OF THE THESIS

We propose that geometrically derived connections between specific neurons in different modules can reduce relative error between them and make the drift coherent. For this, we will model each grid cell module as a continuous attractor network and determine the ge-




Figure 1.2: Grid cell properties. A: Left: schematic of single unit recording. Middle left: raw data from a sample grid cell. The animal's path is in black, and the positions at which grid cells fire are superimposed in blue. Middle right: firing rate map for the same grid cell. Right: the regular firing pattern can be characterized by its orientation, scale, and offset or spatial phase. B: Two grid cells co-recorded on a single tetrode in different environments exhibit the same grid scale and orientation but differ in their offset or relative spatial phase. Top row: firing rate maps for a pair of grid cells recorded in a familiar (left) and novel (right) environment. Bottom row: spatial cross-correlation of the grid cell firing rate maps in each environment. Black dashed lines indicate the central six peaks of the cross-correlation; colored line shows the distance and direction from the central peak to the origin of the spatial cross-correlation. C: Grid cells appear to be organized into discrete functional modules whose scale increases in discrete steps along the dorso-ventral axis of MEC. D: Grid field orientation of grid cells recorded in three different rats. The orientations of grid firing patterns are significantly clustered within and between modules. Figure from: [Bush et al. \(2015\)](#) (License: CC BY 4.0).

ometric coupling between the modules which could be set up by some plausible plasticity mechanism based on correlations. We will compare this response with other possible couplings (such as randomly coupled and uncoupled).

Our proposed model will play a crucial role in making experimental predictions about the type of couplings between grid cell modules and the role the combined network plays in path integration. We will investigate how the coupling found using the above approach may be implemented by recurrent synaptic connectivity within the MEC with a simple biologically plausible structure.

Thus, we aim to:

- Estimate the response of grid cells in a continuous attractor network model and the accumulation of error in a deterministic network
- Identify and compare possible couplings between grid modules that allow for coherent drift across modules and make experimental predictions for existing couplings.

1.2 OUTLINE OF THESIS

The thesis is structured as follows. We will begin by reviewing the current literature, particularly a model of interest (Burak and Fiete, 2009). Then, we model each grid cell module as a 2D sheet of recurrently connected neurons such that their activity pattern forms a continuous attractor network. Continuous attractor network models have been shown to explain the emergence and update of grid cells, but the models are prone to diffusive drift over time (Burak and Fiete, 2009). The network neurons will represent a population of nearby grid cells of the MEC, whose grids share the same orientation and spacing but have different

phases. The different grid scales will emerge when using different gains (with the correct gain ratio) to modulate the external velocity input. We will conduct extensive tests on our model to quantify the accumulation of error with time, distance traveled, and grid module parameters of a single module.

Even though a single module can be modeled as an attractor with N neurons, the network activity converges to a state in a 2 dimensional manifold. With a network consisting of M uncoupled modules, there is a $2M$ dimensional space of phases accessible to the network. Using the above M uncoupled modules, the drifting neuronal representations due to independent errors can lead to readout errors in the animal's location.

Based on earlier experimental [Stensola et al. \(2012\)](#); [Barry et al. \(2007\)](#); [Krupic et al. \(2015\)](#) and theoretical ([Stemmler et al., 2015](#); [Mosheiff et al., 2017](#); [Wei et al., 2015b](#)) work, we argue that the scale ratio of 1.5 or $3/2$ gives a specific geometric relationship between the modules and thus makes the set of accessible states much smaller than $2M$. This helps to reduce the drift and allows for error correction.

In 1D, the coupling between the different modules based on the specific scale ratio can be determined easily, which may provide intuitive insight into the geometrical relationship between them.

The coupling has to be strong enough to allow for mutual error correction but not so strong that it interferes with the stability of the grid pattern of a given module. In addition, the network needs to update network activity based on self-motion signal (velocity) coherently to allow for the correct decoding of position from multiple modules. We will come up with a weak correlation-based coupling between modules.

We will compare the stability of the grid pattern and the error accumulation in the po-

sition estimate of our derived coupling with uncoupled modules and modules with other connection schemes (unidirectional all-to-all, random). Thus, we will determine whether such a coupling allows for coherent drift.

Based on the derived coupling, we can design an in-silico mock-lesion experiment to make specific experimental predictions about the nature of the coupling that may exist in the MEC between grid cell modules. This will allow experimentalists to test for connections between modules indirectly.

2

Literature Review

This chapter reviews current empirical knowledge on grid cells, explicitly delving into their firing patterns, the modulation of their activity by sensory inputs, and their relation to other spatial cells within the hippocampal system. Next, it explores various classes of computational models of grid cells, including continuous attractor network models highlighting the significance of recurrent connectivity, oscillatory-interference models emphasizing

the role of rhythmic oscillations, and single cell plasticity models emphasizing the impact of external sensory cues and learning.

2.1 GRID CELLS IN THE ENTORHINAL CORTEX

The discovery of grid cells stemmed from the quest to understand the origins of place cell activity in the hippocampus. A key finding was that hippocampal place fields in CA1 were partially maintained even after the CA3–CA1 connections were severed, suggesting that the entorhinal cortex (EC), the other major input to CA1, might provide spatial signals to the hippocampus (Brun et al., 2002). Recordings from the dorsal MEC by the Mosers’ group revealed sharp spatial firing, with single cells exhibiting multiple firing fields arranged in a regular triangular pattern (Fyhn et al., 2004; Hafting et al., 2005). These were termed grid cells.

The spatial firing pattern of a grid cell has three fundamental properties:

- Grid scale: the distance between neighboring firing fields.
- Grid orientation: the angle between one of the grid axes and a reference direction.
- Grid spatial phase: the two-dimensional spatial offset between the firing fields and a reference point.

Nearby grid cells recorded from the same tetrode show similar scale and orientation but have scattered phases, indicating that while grid scale and orientation are topographically organized in the cortex, grid phase is not (Hafting et al., 2005).

The remainder of the chapter reviews (i) the anatomical and modular organization of grid cells, (ii) local versus long-range circuitry and its relation to theta oscillations, (iii) in-

teractions with other spatial cell types, (iv) environmental and behavioral influences, and (v) empirical limits on path-integration accuracy—providing the context for the modeling work in Chapters 3–5.

2.1.1 ANATOMICAL ORGANIZATION OF THE MEC

The medial entorhinal cortex is a six-layered cortical area. Superficial layers convey neocortical output to the hippocampus, whereas deep layers return hippocampal output to cortex (Witter et al., 2017).

Layer II contains two excitatory populations: reelin-positive stellate cells and calbindin-positive pyramidal cells, both capable of grid firing (Fuchs et al., 2016). Layer II exhibits a striking anatomical modularity where these segregate into cell clusters referred to as ‘islands’ embedded in an ‘ocean’ of cells. Island cells are the calbindin-positive pyramidal neurons located in clusters (‘islands’), while ocean cells are the surrounding reelin-positive stellate neurons (Fuchs et al., 2016; Kitamura et al., 2014). These populations differ in connectivity and projection targets: ocean cells largely project to the dentate gyrus and CA3, while island cells project to distal CA1 where they provide feed-forward inhibition through local interneurons (Kitamura et al., 2014; Sun et al., 2015). This anatomical modularity creates parallel output channels from MEC layer II and may underlie functional differences in spatial information processing (Domnisoru et al., 2013; Kropff et al., 2015).

Layer III is composed primarily of pyramidal neurons that target CA1 and the subiculum Witter et al. (2017). Deep layers V–VI receive hippocampal feedback and project widely to MEC and neocortical/subcortical targets, completing the cortico-hippocampal loop (Witter et al., 2017).

2.1.2 MODULAR ORGANIZATION OF GRID CELL ACTIVITY

Two years after grid cells were first reported, [Barry et al. \(2007\)](#) found that grid scale increases along the dorso-ventral axis of the MEC. This increase was not continuous, but in discrete steps ([Barry et al., 2007](#)). The ratio between subsequent grid scales in an animal was a constant, about 1.7, matching theoretical predictions for optimal spatial coding with periodic codes ([Stemmler et al., 2015](#); [Wei et al., 2015a](#)). Independent studies reported similar ratios (1.4–1.7) ([Stensola et al., 2012](#); [Krupic et al., 2015](#)).

[Stensola et al. \(2012\)](#) found that grid cell activity is organized into a small number of discrete functional modules. Within a module, grid patterns share similar scale, orientation, and elliptic distortion, but these properties vary across modules. They identified four to five modules per animal, suggesting a limited number of grid modules. Grid fields from different cells within a module drift together from trial to trial, regardless of whether the environment is real or virtual or whether the animal moves in light or darkness ([Nagele et al., 2020](#)).

2.1.3 LOCAL AND LONG-RANGE CONNECTIVITY OF MEC GRID MODULES

Paired recordings show that excitatory connections between neighboring stellate cells in layers II and III are sparse, whereas fast-spiking interneurons provide dense recurrent inhibition ([Couey et al., 2013](#)). Principal cells therefore influence one another mainly via di-synaptic inhibitory loops rather than direct excitation. Within a grid module, excitatory and inhibitory neurons form a locally connected sub-network whose spatial organization mirrors the module’s hexagonal firing geometry ([Dunn et al., 2015](#)). This architecture is consistent with continuous-attractor models in which short-range inhibition stabilizes a

locally coherent grid pattern.

Recordings that track multiple modules simultaneously show that modules are internally coherent yet only loosely tethered to one another. When a boundary of an enclosure was shifted, cells within a module rescaled coherently, but distinct modules within the same animal often rescaled by very different amounts (Stensola et al., 2012). Across contexts, cells in a module preserve fixed phase relationships, whereas cross-module relationships drift (Yoon et al., 2013). Pairwise functional-connectivity analyses reveal strong interactions within modules but near-zero coupling across modules (Dunn et al., 2015). During cue-free navigation, grid phases in different modules drift almost in register, remaining within a few cm of each other, even as the entire map slips against the world, implying a weak but non-zero inter-module tether (Waaga et al., 2021; Hardcastle et al., 2015). A parsimonious interpretation is that modules are stabilized primarily by their own local attractor dynamics and only weakly tethered to one another—enough to limit catastrophic phase divergence, but insufficient to force identical responses to environmental distortions. Together, these findings are consistent with a “weak-coupling” regime: modules behave as autonomous attractors whose small relative drifts are corrected by modest shared inputs or sparse long-range interactions. This “weak-coupling” regime is the working assumption of the models explored in this thesis.

2.1.4 THETA OSCILLATIONS AND SEPTAL INPUT

During active exploration, local field potentials in MEC exhibit a sustained 6–10 Hz theta rhythm that is phase-locked with hippocampal theta. Grid-cell spikes are modulated by this rhythm and often show theta phase precession within their firing fields, analogous to

hippocampal place cells (Hafting et al., 2008).

SEPTAL PACEMAKER

The medial septum–diagonal band (MS-DB) provides convergent cholinergic and GABAergic input to MEC. Septal GABAergic neurons rhythmically inhibit entorhinal interneurons, while septal cholinergic tone depolarizes principal cells, together pacing the theta cycle (Robinson et al., 2024). When the MS-DB is pharmacologically or optogenetically inactivated, the network theta is abolished and grid periodicity collapses, even though head-direction tuning in the same recordings remains intact (Brandon et al., 2011; Koenig et al., 2011; Robinson et al., 2024). These experiments demonstrate that theta-paced inhibition is essential for maintaining the grid attractor, whereas head-direction signaling can persist without it.

INTRINSIC RESONANCE GRADIENT

Layer-II stellate cells exhibit intrinsic theta-range oscillations; their resonant frequency decreases from dorsal to ventral MEC in parallel with the increase in grid scale (Giocomo et al., 2007), suggesting that intrinsic resonance contributes to module-specific spatial periods.

NETWORK COORDINATION

Theta provides a common temporal reference that aligns spikes across modules and between MEC and hippocampus, supporting coherent phase updates during path integration. Removing the septal drive eliminates this shared clock, disorganizing grid patterns

while leaving local circuitry intact (Koenig et al., 2011). Thus, weakly coupled modules depend on a common theta oscillation to limit relative drift and maintain a unified metric. Consistent with the weak-coupling framework, modules rely on a global theta drive to maintain synchrony; when that drive is removed, inter-module coordination weakens and grid patterns degrade.

2.2 RELATIONSHIP OF GRID CELLS TO OTHER HIPPOCAMPAL SPATIAL CELLS

Grid cells are only one type of many spatially-modulated cells in the hippocampal system. Other types include place cells (O'Keefe and Dostrovsky, 1971), head-direction cells Taube et al. (1990), border cells (Solstad et al., 2008), speed cells (Kropff et al., 2015) and object-vector cells (Høydal et al., 2019). This section summarizes the key findings that link grid cell activity to the cells most relevant for the present work, specifically place, head-direction, and border cells.

2.2.1 GRID AND PLACE CELLS

The hippocampus receives direct input from the superficial layers of the MEC, where grid cells are abundant (Hafting et al., 2005; Sargolini et al., 2006). This led to the hypothesis that entorhinal grid fields generate hippocampal place fields (Solstad et al., 2006; Cheng and Frank, 2011). This is supported by evidence of monosynaptic connections from grid cells to hippocampal neurons (Zhang et al., 2013) and the spatial scale of both grid and place fields being affected by genetic manipulations in the entorhinal cortex (Mallory et al., 2018).

However, evidence also suggests that grids are not necessary to form place fields. For ex-

ample, place cells develop before grid cells (Wills et al., 2010) and place fields remain largely intact after entorhinal lesions (Van Cauter et al., 2008) or disruption of grid cell activity via septal inactivation (Koenig et al., 2011). However, place fields might still form from degraded grids (Azizi et al., 2014).

Conversely, grid fields could be influenced or anchored by hippocampal place fields. Hippocampal inactivation (Bonnievie et al., 2013) or lesions (Fyhn et al., 2004) disrupt grid cell firing. Together, these findings indicate an overall bidirectional influence between these cell types.

2.2.2 GRID AND HEAD DIRECTION CELLS

Head direction (HD) cells fire based on the orientation of the animal's head, independent of its location (Taube et al., 1990). First discovered in the dorsal presubiculum (Taube et al., 1990), HD cells have also been found in the anterior thalamic nucleus (ATN) (Taube, 1995), parasubiculum (Taube, 1995b), and entorhinal cortex (Sargolini et al., 2006). HD cells are prevalent in MEC layers III, V, and VI, where grid and conjunctive grid-HD cells are also located (Sargolini et al., 2006). Conjunctive cells are used by many models (Burak and Fiete, 2009), since HD cells likely provide grid cells with self-motion information, which is crucial for computational models of grid cell activity.

Lesions or inactivation of the ATN degrades both HD fields and grid firing patterns in the MEC (Winter et al., 2015). Conversely, septal inactivation abolishes theta and collapses grid firing while leaving head-direction coding largely intact (see § 2.1.4 for details).

2.2.3 GRID AND BORDER CELLS

Border cells are neurons that ‘selectively fire as an animal approaches the boundaries of an enclosure’ (Solstad et al., 2008; D’Albis, 2018). Found in the MEC (Solstad et al., 2008; Savelli et al., 2008) and parasubiculum (Boccara et al., 2010), border cells are functionally related to the boundary-vector cells in the subiculum, which fire whenever the animal is at a certain distance to a boundary wall (Barry et al., 2006).

Since boundaries influence both the symmetry and alignment of grid fields, border cells are believed to provide input to grid cells. This is why border cells have been proposed to correct path-integration errors in grid cell models (Hardcastle et al., 2015).

2.3 ENVIRONMENTAL AND BEHAVIORAL MODULATION

2.3.1 SLEEP AND AWAKE CORRELATIONS

Grid cells maintain similar correlation patterns both during active navigation and rest states, such as sleep (Trettel et al., 2019; Gardner et al., 2019). This consistency supports the idea that grid cell networks rely on intrinsic connectivity, possibly organized into a continuous attractor network (CAN) structure, to maintain spatial coherence. These models predict that grid cells will exhibit similar co-activity patterns during different modes of behavior. Intramodular grid cell pairs (cells within the same module) have higher correlation amplitudes than trans-modular pairs (cells from different modules) across active and sleep states (Trettel et al., 2019; Gardner et al., 2019). This finding indicates that recurrent connectivity is stronger within grid cell modules than between them. Also, Gardner et al. (2022) showed that grid cell population activity can span a toroidal manifold, which is in-

variant across different environments and brain states.

2.3.2 BOUNDARY AND LOCAL GEOMETRY INFLUENCES

Proximal boundaries and their geometric configuration are essential for anchoring (Savelli et al., 2008; Solstad et al., 2008; Stensola et al., 2015; Krupic et al., 2015), shaping boundaries (Barry et al., 2007; Krupic et al., 2015) and correction (Hardcastle et al., 2015) of the grid pattern. Theoretical work suggests that grid cells primarily represent geometric information about the local enclosure and that boundaries help stabilize grid maps (Krupic et al., 2016; Stensola and Moser, 2016).

Grid cell activity is influenced by external sensory cues, as shown by studies manipulating the local geometry of the environment. Grid cells rescale along with environmental compression (Barry et al., 2007; Solstad et al., 2008). Similarly, grid fields locally shift when animals move between previously separate enclosures (Wernle et al., 2018).

Krupic et al. (2015) showed that in square environments, grid orientation is anchored to the walls: rotating the enclosure by 45° rotated the grids by the same amount, despite stationary distal cues. In these square arenas, grid axes were typically offset by about 9° from the walls, and simultaneously recorded modules exhibited preferred relative orientations clustering around 0° and 30° , with 0° being most common. Importantly, these relative orientation relationships persisted when animals were tested in circular or hexagonal enclosures, even though absolute grid orientations were less tightly clustered in circular enclosures.

In more complex arenas, such as hairpin mazes (Derdikman et al., 2009) or in trapezoidal arenas (Krupic et al., 2015), the characteristic triangular symmetry of grid cell firing

is largely lost, further supporting the significant influence of environmental factors on grid cell activity.

2.3.3 ERROR CORRECTION MECHANISMS

While the modular structure of grid cells enhances coding capacity and reduces ambiguity, it also poses challenges for error correction. Small drifts in the decoded positions from different modules can accumulate over time, potentially leading to errors in spatial representation (Burak, 2014; Sreenivasan and Fiete, 2011; Welinder et al., 2008). Mechanisms for correcting these errors, particularly in the absence of sensory cues, remain an area of active research.

2.3.4 SELF-MOTION INPUTS VERSUS EXTERNAL SENSORY CUES

When grid cells were first discovered, Hafting et al. (2005) found that grid fields remained stable in darkness, suggesting they might arise from integrating self-motion cues such as vestibular signals or proprioceptive feedback. However, the same study showed that grid patterns rotated with polarizing visual cues. These findings led to the hypothesis that grid patterns emerge from path integration and anchor to external landmarks with experience (McNaughton et al., 2006).

Despite this, the initial emphasis on self-motion inputs might have been overestimated, as grid patterns in darkness could also be influenced by olfactory or somatosensory cues from arena boundaries (Hafting et al., 2005; Barry et al., 2012). This is supported by grid patterns quickly disrupted in darkness in elevated arenas and minimal olfactory conditions in mice but not in rats (Chen et al., 2016; Pérez-Escobar et al., 2016). This suggests that the

balance between self-motion and external cues might differ across species.

2.3.5 GRID CELL PATTERNS ACROSS ENVIRONMENTS

Grid cell patterns change across environments in a process that can be compared to the ‘remapping’ observed in hippocampal place cells. Remapping in place cells occurs in two ways depending on the manipulation of the spatial context: global remapping and rate remapping.

When a rat encounters two different boxes in a room or similar boxes in different rooms, hippocampal activity undergoes global remapping, where place cells active in one condition are either silent or have fields in different locations in the other (Wills et al., 2005; Leutgeb et al., 2005). In contrast, when the same room is used but salient features of the arena such as wall colors, arena shape, or odors are changed, place cells undergo rate remapping, where the firing fields remain stable but firing rates vary substantially (Hayman et al., 2003; Leutgeb et al., 2005; Latuske et al., 2018).

Fyhn et al. (2007) recorded entorhinal grid cells in conditions where hippocampal firing undergoes either global or rate remapping. They found that grid cell activity was significantly more stable compared to hippocampal activity, maintaining the cells’ spatially periodic firing across contexts even when place fields were globally remapped. In the same room with different boxes, grid patterns shifted coherently within modules, meaning cells recorded from the same electrode had similar changes in phase offsets. When identical boxes were in different rooms, grids rotated and translated but maintained constant relative-phase relationships within modules, termed ‘coherent remapping’ (Fyhn et al., 2007; Yoon et al., 2013).

In scenarios where place fields underwent rate remapping, grid cells either remained stable or shifted coherently, though peak firing rates within fields varied (Ismakov et al., 2017; Diehl et al., 2017). Grid cell remapping is typically studied in familiar environments, where grid fields maintain similar size and spacing (Fyhn et al., 2007; Marozzi et al., 2015; Ismakov et al., 2017; Diehl et al., 2017). However, grid patterns expand in novel environments and become less regular (Barry et al., 2012).

2.3.6 BEHAVIORAL ACCURACY OF PATH INTEGRATION IN RODENTS

Rodents can navigate back to a home refuge after complex outbound paths even in complete darkness. Classic homing-table studies showed that rats keep angular error within $10\text{--}15^\circ$ and distance error below $\sim 10\%$ of path length over returns of 1–3 m (Etienne and Jeffery, 2004). Error grows roughly linearly with traveled distance until the animal contacts a boundary or landmark that re-anchors its internal map.

Hippocampal integrity is essential for path integration: rats with hippocampal lesions show impaired homing when only self-motion cues are available (Maaswinkel et al., 1999). Medial entorhinal cortex lesions likewise disrupt homing, whereas lateral entorhinal lesions do not (Van Cauter et al., 2013). Genetic disruption of grid cell firing further impairs path integration, as shown in mice lacking GluA1 AMPA receptors or NMDA receptor function in MEC circuits (Gil et al., 2018; Allen et al., 2014). Taken together, these behavioral studies demonstrate that both hippocampus and entorhinal grid circuits are required for accurate path integration.

Physiological recordings for rats during foraging in a circular arena show that in darkness, the decoded position deviated from the true position, with a mean absolute error

(MAE) substantially larger than in light (Waaga et al., 2021), with errors reached ~ 30 cm in high-MAE segments. Importantly, this drift occurred coherently across modules: decoded positions from different grid scales remained aligned with each other even as they collectively deviated from the animal's actual path (Waaga et al., 2021). Thus, while the absolute decoded position drifted from the animal's true trajectory, relative alignment across modules was preserved (Waaga et al., 2021).

Peng et al. (2023) extends these findings: grid cells in mice can represent self-motion in both an arena-centred and an object- (lever-) centred reference frame, and firing patterns switch between these frames depending on task demands. Thus, grid activity supports path integration across changing reference frames while preserving coherent population coding.

2.4 CONTINUOUS ATTRACTOR MODELS

Continuous attractor networks (CANs) underlie one of the most prominent theories for the origin of grid cell patterns. The core idea behind such models is that a properly wired recurrent neural network has a continuum of neutrally stable states in the activity of its neurons (Amari, 1977). Because all states are neutrally stable, small external perturbations can push the network to a nearby stable configuration. Thus, the network activity at any given time reflects the cumulative effects of the external inputs onto the system (Wu et al., 2008).

As CANs can potentially integrate well, they have been employed in models of various brain networks that require tracking continuous variables. For instance, CAN models have been used in the oculomotor system to monitor eye position (Cannon et al., 1983; Seung, 1998), in the thalamus to monitor head direction (Skaggs et al., 1994; Zhang, 1996; Blair

and Sharp, 1996), and in the hippocampal system to monitor position (McNaughton et al., 1996, 2006; Fuhs and Touretzky, 2006).

Because CANs share the same basic architecture, we start with the network setup for the simplest scenario, the head direction system.

2.4.1 CONTINUOUS ATTRACTOR MODELS OF HD CELLS

The population activity of HD cells forms a distributed representation of the animal’s current heading using population vector encoding. The current heading estimate, θ , is the weighted circular mean of unit vectors φ_i and normalized firing rate r_i :

$$\theta(t) = \sum_j r_j(t) e^{i\varphi_j} \quad (2.1)$$

In rats, the HD system remains active even in darkness, indicating that non-visual inputs are sufficient for updating the heading estimate. In the model, neurons that acquire head direction tuning are arranged on a ring, with nearby neurons having strong excitatory connections and distant neurons having inhibitory connections. This connectivity pattern generates a localized activity bump on the ring, which moves with the animal’s head rotation due to two groups of ‘rotation neurons’: clockwise-rotation and anticlockwise-rotation neurons.

HD cells provide the current heading, while vestibular angular velocity input supplies head-turning velocity. When the animal turns its head, rotation neurons are activated proportional to the turning speed. These neurons project back to the HD ring asymmetrically, effectively generating an HD signal by integrating angular velocity from the vestibular system.

At any given time, the location of the activity bump on the ring can be used to determine head direction, where each cell has a bell-shaped directional tuning curve. The network's topology matches the encoded signal's nature: cells are connected on a ring because HD is a one-dimensional periodic variable. This ring-like connectivity does not imply a topographical arrangement of preferred head orientations on neural tissue, as observed in mammals (Taube et al., 1990; Taube, 1995). First theorized in the 1990s, CAN models of HD cells have more recently gained experimental support from imaging studies in the brains of flies (Seelig and Jayaraman, 2015; Turner-Evans et al., 2017) and zebrafish (Petrucco et al., 2023).

2.4.2 CONTINUOUS-ATTRACTOR MODELS OF GRID CELLS

Since the discovery of the periodic grid cell tuning in the MEC (Hafting et al., 2005), CAN models of place cells have been adapted to explain grid cell activity (McNaughton et al., 2006; Fuhs and Touretzky, 2006).

Two classes of models have been studied: models where a single activity bump moves on a network with twisted-torus topology (Guanella et al., 2007) and models with periodic boundary conditions that exhibit multiple activity bumps (Fuhs and Touretzky, 2006; Burak and Fiete, 2009; Couey et al., 2013; Widloski and Fiete, 2014).

Multi-bump attractor networks can form triangular grids at the single cell level without periodic boundary conditions. However, they are poor velocity integrators and unlikely to generate stable grids at the single cell level due to edge effects (Burak and Fiete, 2009).

The grid scale can be controlled by different scales of connectivity (Kang and Balasubramanian, 2019) or by the speed modulation gain of conjunctive cells, which translate the

activity pattern (bumps) on the neuronal sheet with the animal's movements. High speed gains result in faster bumps and small grids, while lower gains produce slower bumps and larger grids.

To address the issue of topographic organization, McNaughton et al. (2006) proposed a learning process for forming CANs during development. In the first step, a transient population activity pattern initially emerges in a 'teaching layer' of neurons via Mexican-hat (characterized by an excitatory center and inhibitory surround) connectivity. Then, synaptic noise translates this pattern randomly across the teaching layer, enabling downstream neurons to form locally periodic recurrent connections. While intriguing, this model has faced mathematical challenges (Burak and Fiete, 2009) and limited experimental support.

Widloski and Fiete (2014) suggested an alternative method for learning grid cell CANs, requiring spatial exploration during development. Place-selective and self-motion inputs activate an initial random network as the animal explores. Temporally asymmetric synaptic plasticity then generates a network connectivity akin to hard-wired CANs. After learning, synaptic plasticity is turned off, and the mature network's dynamics are driven by self-motion inputs alone, sustaining multiple coherent activity bumps.

However, this model has several issues. First, the population-activity pattern in the mature network does not consistently align with the spatial inputs active during learning (Widloski and Fiete, 2014). While this is presented as a feature, it seems conceptually odd for feed-forward and recurrent inputs to contradict each other in a familiar, cue-rich environment. Second, the aperiodic learned connectivity makes the model susceptible to noise Burak and Fiete (2009).

SINGLE CELL GRID FORMATION IN CAN MODELS

CAN models require precise integration of the animal's velocity and heading to form single cell grids from population patterns. However, any integrating system is prone to errors without external calibrating signals. For CANs, an integration error equivalent to the grid spacing (tens of centimeters) can completely disrupt single cell firing patterns. [Burak and Fiete \(2009\)](#) estimated that a noisy periodic network with approximately 10^4 neurons accumulates errors at a rate of about $\approx 1\text{cm/min}$, indicating that velocity integration must be corrected or reset every few minutes of exploration. In biological systems, errors may accumulate faster due to smaller network sizes, inhomogeneous connectivity patterns, or spatially irregular feed-forward inputs. Also, coherent grid firing across exploration sessions requires anchoring network activity to stable environmental landmarks.

Landmark based error correction and anchoring have been discussed in grid cell literature (e.g., [Welinder et al. \(2008\)](#)), with few explicit models ([Guanella et al., 2007](#); [Hardcastle et al., 2015](#); [Pastoll et al., 2013](#)). One possibility is that hippocampal place cells constantly calibrate grid cell CANs ([Guanella et al., 2007](#); [Pastoll et al., 2013](#)). Hebbian synaptic plasticity could rapidly form biologically plausible associations between simultaneously active grid and place cells as the animal explores a new environment. Place cell firing could, in this case, correct CAN integration errors within sessions and anchor grid fields across sessions ([Guanella et al., 2007](#); [Pastoll et al., 2013](#)). However, since place cell activity remaps across environments, this solution requires 'one-shot learning' of place to grid associations in each environment.

A more straightforward solution involves using border cell inputs for calibration ([Hardcastle et al., 2015](#)). These authors found that grid cell spike dispersion aligns with the in-

tegration drifts predicted by CAN models, which reduce after the animal encounters an arena boundary. Furthermore, grid cell firing is more precise in the direction perpendicular to the last-encountered boundary, indicating that border cells provide error-correction signals to grid cells. Since border cells remain stable across different spatial contexts (Solstad et al., 2008), this border-based calibration does not require rapid learning in new environments. However, border inputs only provide spatial information perpendicular to the boundary, which is insufficient for anchoring single cell grids across sessions.

EVIDENCE

As summarized in § 2.3.4, behavioral and lesion work show that self-motion cues can sustain grids in darkness and that disrupting upstream HD signals or septal theta degrades them, consistent with CAN assumptions. The attractor theory also aligns with the modular organization of grid cell activity and its coherent remapping across environments (Stensola et al., 2012; Yoon et al., 2013) and behavioral states (Trettel et al., 2019; Gardner et al., 2019). Such models predict a topographic organization of grid spatial phases in the cortex, which is observed experimentally (Gu et al., 2018) and toroidal topology (Gardner et al., 2022). However, attractor models require a sophisticated neural architecture with an unclear developmental origin. They need anchoring to physical space and fail to explain how external sensory cues and environmental geometry distort grid fields.

2.5 OSCILLATORY INTERFERENCE MODELS

Oscillatory interference (OI) models propose that grid patterns arise from temporal interference between multiple velocity-controlled oscillations and a baseline theta rhythm, pre-

dicting that amplitude modulation of intracellular theta drives grid cell firing. This concept builds on the theta modulation of grid-cell spiking described in Section 2.1.4, where septal input entrains a 610 Hz rhythm across MEC and hippocampus. In OI models, these oscillations occur within individual cells and their phase relationships encode spatial location.

The model was extended to 2D environments (Burgess et al., 2007), where a theta-range oscillation of the somatic membrane potential interferes with multiple velocity-controlled oscillations (VCOs) of the dendritic membrane potential. However, phase-locking issues can be overcome by modeling VCOs as separate units or ‘theta cells’ (Burgess et al., 2007).

VCOs must be tuned to movement directions differing by 60 degrees to generate a triangular grid. While two VCOs can create grid-like patterns, we need six for omnidirectional phase precession (Climer et al., 2015; Jeewajee et al., 2014; Reifenstein et al., 2014). However, these models cannot account for path-dependent phase precession features (Reifenstein et al., 2014).

A fundamental weakness of OI models is the manual insertion of the 60-degree periodicity and the requirement for precise phase matching of multiple VCOs. While self-organizing processes have been proposed (Burgess et al., 2007), they have yet to be explicitly simulated. Additionally, grid-like periodicity is susceptible to disruption by noise (Welinder et al., 2008; Zilli et al., 2009), and it is unclear if network synchronization can mitigate this in realistic settings (Zilli and Hasselmo, 2010).

EVIDENCE

OI models are supported by evidence of sub-threshold oscillations in entorhinal neurons (Alonso and Llinás, 1989; Alonso and Klink, 1993), theta-modulated and phase-precessing

grid cell activity (Hafting et al., 2005; Schmidt-Hieber and Häusser, 2013; Domnisoru et al., 2013), and experimentally identified velocity-modulated theta cells (Welday et al., 2011). Also, the dorso-ventral frequency gradient in stellate cells corresponds to increasing grid spacing (Giocomo et al., 2007), and abolishing theta modulation disrupts grid cell firing (Brandon et al., 2011).

However, there is solid evidence against OI models: in Reifenstein et al. (2012), it is shown that the observed phase precession in grid cells is incompatible with global oscillator interference.

2.6 HYBRID INTERFERENCE AND ATTRACTOR MODELS

Schmidt-Hieber and Häusser (2013) and Domnisoru et al. (2013) recorded membrane voltages of grid cells during virtual exploration, finding depolarizing ramps within firing fields. Both membrane voltage and output spikes were theta-modulated, with increased intracellular power within firing fields. Further analysis indicated that depolarizing ramps, rather than theta power, predicted output spiking, favoring CAN models. The observed intracellular theta modulation led to hybrid interference-attractor models (Schmidt-Hieber and Häusser, 2013; Bush and Burgess, 2014).

Hybrid models connect grid cells driven by temporal interference, with strong connections between cells with similar spatial phases. Unlike CAN models, VCOs control activity bump shifts, generating spatial periodicity and intracellular theta modulation, while recurrent connections produce depolarization ramps.

Hybrid models also generate phase precession (Schmidt-Hieber and Häusser, 2013; Bush and Burgess, 2014). Navratilova et al. (2012) suggested that CAN models alone could

generate grid fields and phase precession using theta-modulated conjunctive cells and intrinsic synaptic currents. However, realistic grid scales required long synaptic delays and running-speed-modulated intrinsic currents.

2.7 SINGLE CELL PLASTICITY MODELS

A final class of models proposes that grid patterns emerge at the single cell level through spatially-tuned feed-forward inputs and Hebbian synaptic plasticity (D’Albis and Kempter, 2017; Monsalve-Mercado and Leibold, 2017). Unlike previous models, these theories suggest that periodic spatial patterns emerge from integrating external sensory inputs rather than self-motion signals.

2.8 EXAMPLE CAN MODEL: BURAK AND FIETE (2009)

2.8.1 NETWORK ARCHITECTURE

The model consists of $(n_x \times n_y)$ neurons, arranged on a 2D sheet with coordinates $\vec{r} = (x, y)$, where $x, y \in N, 1 \leq x \leq n_x, 1 \leq y \leq n_y$.

Neuron i is assigned a directional preference along the angle $\theta_i = \text{N, S, E or W}$ depending on its position on the neuronal sheet. Neurons with the four directional preferences are arranged in repeating blocks of 2×2 on the sheet, and the unit vector along the direction preference of a particular neuron i is denoted by $\hat{e}_{\theta_i}$.

2.8.2 NEURONAL DYNAMICS

The activation of the i th neuron is given by s_i . The dynamics is governed by the following equation:

$$\tau \frac{ds_i}{dt} + s_i = f \left[\sum_j W_{ij} s_j + B_i \right] \quad (2.2)$$

Here, τ is the neuronal time constant and the activation function $f(x) = \max(0, x)$.

W_{ij} denotes the weight from neuron j to neuron i and B_i defines the feed-forward input to the i th neuron. The synaptic weight from neuron j to neuron i within a module is given by

$$W_{ij} = W_0(\vec{r}_i - \vec{r}_j - l\hat{e}_{\theta_j}) \quad (2.3)$$

where $W_0(\vec{r}) = e^{-\gamma|\vec{r}|^2} - e^{-\beta|\vec{r}|^2}$. The symmetric term W_0 is the difference of two Gaussians with $\gamma > \beta$, which allows for local inhibition with a Mexican Hat shape. $\vec{r}_i, \vec{r}_j$ are the positions of the neurons i, j respectively in the sheets for both modules as described above and $|\vec{r}|$ is the Euclidean distance between the neurons.

Thus, the outgoing weights of neuron j (defined by W_{ij}) have a center-surround structure, with an asymmetry introduced through the shifting of the center by l units in the direction along the preferred direction of neuron j .

If the system received only this input, it would relax to a state where all neurons are inactive, since the connections are purely inhibitory. To offset the inhibition, a constant feed-forward excitation is supplied to all neurons, and this leads to static bumps being formed.

The total feed-forward input to the i th neuron is:

$$B_i = (1 + \alpha \hat{e}_{\theta_i} \cdot \vec{v}_{in}) \quad (2.4)$$

where, $\vec{v}_{in}$ is the input velocity of the animal and α is the gain factor.

The movement of the patterns is controlled by the second term in B_i and the pattern flow does not destroy the stability of the lattice if $\alpha |\vec{v}_{in}| \ll 1$. Small velocity gain factors α make the integrator respond slowly for the same input velocity. As a result, for smaller α , the pattern on the neural sheet would need to cover more distance in physical space in order to return to its starting position, thus forming larger grids in physical space. If the gain factors are chosen to be in the ratio of 3 : 2, then the grid cells in two modules have their spatial scales in the ratio ≈ 1.5 , lying in the experimentally observed range of scale ratios (1.4 – 1.7) (Stensola et al., 2012; Barry et al., 2007; Krupic et al., 2015).

2.9 OPEN QUESTIONS

2.9.1 PROBLEM OF CATASTROPHIC READOUT

In general, the neural representation in each module is prone to gradual diffusive drift. Such errors are bound to accumulate and corrupt the estimate of position. This problem is of special significance in the absence of salient sensory cues, since the phases represented in different modules might drift independently of each other under these conditions. A mechanism is needed to prevent the position estimate from rapidly deviating from the actual location of the animal.

The key questions for our study are:




Figure 2.1: Catastrophic read-out when two grid cell modules drift independently. Top row: At t_0 , activity in Module 1 (cyan, larger period) and Module 2 (orange, smaller period) both encode the animal's true position (black dot). Their likelihood maps overlap at the same location (black dot). Bottom row: At t_1 , small independent phase drifts shift Module 1 leftward and Module 2 upward. The joint likelihood could lie at a distant yellow dot which is far from the animal, illustrating a catastrophic read-out error. Cropped and recolored from Fig. 1a of [Mosheiff and Burak \(2019\)](#) under CC BY 4.0 (creative commons) license.

- Do different modules in biological grid-cell systems function independently?
- What happens when different modules are drifting independently?
- Can intermodular coupling between modules contribute to error correcting mechanisms in the absence of external sensory cues?

Now, if the modules are independent, the joint estimate gets worse and worse with time and the incoherent drift between the 2 modules would accumulate and make the decoded position be very far away from the actual position. This was referred to as catastrophic read-out by [Burak \(2014\)](#) and is illustrated in figure 2.1.

2.9.2 NON-STATIONARY VELOCITY INPUT

In natural foraging, a rodent’s velocity vector varies continuously: speed accelerates and decelerates, brief pauses occur, and heading changes are frequent (Sargolini et al., 2006; Kropff et al., 2015). Such non-stationary input contrasts with the constant-velocity assumption used in many early continuous-attractor (CAN) models of grid cells. Empirically, running speed modulates both grid-cell firing rate and theta phase precession (Jeejee et al., 2014), implying that integration mechanisms must cope with variable speed as well as direction.

Classical CAN simulations show that very low speeds can fail to translate the activity bump or produce non-linear updates, leading to phase errors (Burak and Fiete, 2009). Subsequent models therefore introduced velocity-gated dynamics or trained agents on naturalistic rodent trajectories to maintain stable grids under realistic movement statistics (Fuhs and Touretzky, 2006; Raudies and Hasselmo, 2012). These studies highlight a practical requirement for grid-module coupling: if each module integrates the same noisy, non-stationary velocity signal, weak inter-module interactions can help prevent the independent phase drifts that would otherwise accumulate during slow or irregular motion.

3

Modified Burak and Fiete model

In this chapter, we identify and analyze sources of drift in the deterministic version of the [Burak and Fiete \(2009\)](#) model and propose a modification to reduce it. Additionally, we examine the model's response to constant velocity inputs and identify its attractor states.

3.1 PROBLEMS WITH THE ORIGINAL MODEL OF BURAK AND FIETE

Since one aim of this thesis is to study drift as a limiting factor for path integration in the grid-cell system, we aim to quantify and understand its sources in the deterministic [Burak and Fiete \(2009\)](#) model. This section explains our parameter choices and model adjustments as we identify or quantify different sources of drift.

3.1.1 SQUEEZING/ STRETCHING OF EMERGENT ACTIVITY PATTERNS

In a translation-invariant continuous attractor, such as an infinite neuronal sheet with zero velocity input, the homogeneous firing state is stable if small perturbations of any spatial wavelength decay. Linear stability analysis tests this by decomposing perturbations into spatial Fourier modes, each with a growth rate $\lambda(k)$ given by the connectivity spectrum $\hat{W}(k)$ minus a uniform decay term ([Amari, 1977](#); [Zhang, 1996](#)). If $\lambda(k) > 0$ for some k , modes of that wavelength grow. The wavenumber k_0 at which $\lambda(k)$ is maximal is the first unstable mode: it grows fastest after instability onset and sets the natural wavelength $\lambda_{nat} = 2\pi/k_0$ of the emergent bump lattice. This wavelength reflects the network's preferred spacing in the absence of boundary constraints, and serves as the benchmark against which we can detect global stretching or squeezing in finite networks.

The recurrent coupling function $W(x)$ and its Fourier transform, as well as its derivatives, help us identify the magnitude of the wavevector k_0 of the first unstable mode. Given the coupling function:

$$W(|\vec{x}|) = \exp(-\gamma|\vec{x}|^2) - \exp(-\beta|\vec{x}|^2)$$

with $\gamma = c\beta$. Since the coupling is radial, the Fourier transform only depends on $k = |\vec{k}|$ and is given by:

$$\hat{W}(k) = \frac{\pi}{c\beta} \exp\left(-\frac{k^2}{4c\beta}\right) - \frac{\pi}{\beta} \exp\left(-\frac{k^2}{4\beta}\right)$$


The first unstable mode (minimum k) will occur for $d\hat{W}(k)/dk = 0$. The magnitude of the wavevector of the first unstable mode is:

$$k_0 = \sqrt{\frac{8\beta c \ln(c)}{c-1}} \quad (3.1)$$

For the parameters of [Burak and Fiete \(2009\)](#), $\beta = 3/13^2$ and $c = 1.05$, linear stability analysis gives $2\pi/k \approx 16.47$ which deviates from the spacing of the actual pattern (≈ 21.33) (fig. 3.1(a)). In other words, the periodic boundary conditions force the network pattern to be globally stretched in the x- and y- directions by roughly 23%. To use the nomenclature of physics, the states formed are “frustrated” (forced into a spacing or orientation that differs from the intrinsic one) and this mismatch has the potential to derail the system’s path-integration capabilities, and one should try to choose network parameters that allow for a “relaxed” (at its intrinsic spacing without external constraints) activity pattern.

To address this issue, the network size was adapted instead, as detailed in the next section. The period for the 2D case was determined numerically for a large sheet size to avoid edge effects, resulting in good agreement with theoretical values (Figure 3.2), even though we used linear stability analysis as an approximation for our non-linear system. As a consequence, the pattern that forms no longer has to be squeezed or stretched to fit into the network. This minimizes the “quantization” effect, and the result is in good agreement

(a) Burak and Fiete network



(b) Bush and Schmidt-Hieber network




Figure 3.1: (a-b) 128×128 neurons grid cell networks from [Burak and Fiete \(2009\)](#) and [Bush and Schmidt-Hieber \(2018\)](#) are not composed of regular hexagons with angle 60° . Note: There is a slight shearing in (a) (which is absent in (b)) - in one dimension, the pattern is aligned with the coordinate axis but not in the other, as indicated by the blue lines. The extra line in (a) is parallel to the y-direction and the two bumps not aligned helps to see the deviation from perfect hexagonality clearly. Figures modified from [Burak and Fiete \(2009\)](#) and [Bush and Schmidt-Hieber \(2018\)](#).

with the first unstable mode.

3.1.2 SIZE OF NEURAL SHEET

When the network has periodic boundary conditions, the formation of the pattern depends on the size of the sheet considered. If the pattern period is incommensurate with the size of the neuronal sheet, the pattern would be deformed and deviate from an exact multi-bump hexagonal pattern. Such distortion is not merely cosmetic: in a continuous attractor, exact translational symmetry ensures unbiased velocity integration, whereas stretching or shearing breaks this symmetry. The resulting anisotropy in the attractor landscape induces direction-dependent restoring forces that bias bump motion, leading to potential decoding errors ([Fuhs and Touretzky, 2006](#)). This happens for ([Burak and Fiete, 2009](#)) where




Figure 3.2: Zero input equilibrium activity pattern for a network with 120×104 neurons (left) for which the central peaks (red dots) of the autocorrelogram (right) are used to calculate the spatial period of the grid pattern. This pattern has an average period of 15.009 for a system with $c = 1.1$ and theoretical spacing = 15 neurons. We consider the large network to reduce spurious influences of the boundary conditions.

the instantaneous population activity on the neural sheet, although grid-like, is not truly hexagonal.

One solution is using the size of the sheet such that it is a multiple of a unit cell. The pattern in Figure 3.1(a) has a unit cell of base of size $128/6 = 21.33$ and height of $128/7 = 18.28$ neurons. This leads to a height-to-base ratio of 0.8571, and the angle would be $\arcsin(0.8571) = 58.992^\circ$, a deviation of 1.68% from the desired value of 60° . An alternate implementation by [Bush and Schmidt-Hieber \(2018\)](#) replaces the 7 rows with 6, and leads to a pattern with height-to-base ratio 0.833 without the shearing distortion but has the stretching discussed in the previous section. [Couey et al. \(2013\)](#) also use a sheet with 128×128 neurons but have a height to base ratio of 0.93.

Since the neurons are located at integer points on a rectangular 2D sheet, our unit cell would be of size $(\lambda, \lambda \sin(60^\circ))$. To minimize distortion, the neural sheet should ideally have integer unit cells, leading to a more accurate and stable pattern. For an ideal sheet,

$$h/b = \sin(60^\circ) = 0.8660$$

We considered which integral combinations would allow us to obtain a ratio as close as possible to the above. The continued fraction representation of $\sqrt{3}/2$ yields $1, 6/7, 13/15, 84/97, 181/209$ and so on. And we know, from Dirichlet's theorem, that for any of the p/q rational numbers above, the difference to $\sqrt{3}/2$ is less than $1/q^2$ in magnitude.

Choosing $b=15$ and $h=13$ results in a h/b ratio of 0.8666 , corresponding to an angle of 60.07° , i.e., a relative distortion of only about 0.1% . Thus, the sheet is set to have a unit cell (15×13) or (30×26) .

Side note: this choice comes with the complication that the bumps on a sheet would not be identical - since if one maximum lies at (x, y) , then another would have to be at $(x + \lambda/2, y+13)$ and for an odd λ , one maximum would be at a single neuron but in other cases, it would be shared between 2. This is because the rational h/b , in terms of their reduced form, have both b and h odd. Therefore, using unit cells of 30×26 neurons is advantageous but costlier simulation-wise.

Our findings open up new possibilities for future research, such as the use of hexagonal tiling in the neuronal sheet itself. While this approach is beyond the scope of our current work, it could be a promising direction for further exploration in such modeling.

3.1.3 DISCRETE STATES IN FINITE "CONTINUOUS" ATTRACTOR NETWORKS

For the activity patterns of the form shown in fig. 3.1, when an external input is applied, the whole pattern moves instead of distorting. The network activity dynamics of a continuous attractor network as described in [Burak and Fiete \(2009\)](#) is thus conceptualized as minimizing an energy functional, and the attractor states correspond to local minima of

this functional. *

In a continuous-attractor network, any rigid translation of the steady-state activity pattern is another steady state with the same Lyapunov energy (Amari, 1977; Seung, 1998). Appropriately tuned velocity input therefore moves the bump along a neutrally-stable, low-dimensional manifold, enabling path integration (Burak and Fiete, 2009). By contrast, displacements orthogonal to that manifold raise the energy, so the landscape is flat only within the translational subspace; in finite discretized networks the resulting pinning barriers may become appreciable (Noorman et al., 2022).

In general, CAN models require many neurons to approximate a smooth continuous manifold and achieve accurate persistence and integration (Zhang, 1996). Strikingly, biological circuits can implement ring-attractor-like dynamics with far fewer neurons: Seelig and Jayaraman (2015) showed that a bump of activity in the fly ellipsoid body tracks heading even in darkness, consistent with a continuous attractor implemented by only a few dozen neurons. Noorman et al. (2022) also demonstrated that, with exact tuning of threshold-linear dynamics, such bumps can be eliminated, yielding mathematically perfect integration even with few neurons. However, these solutions are highly sensitive to parameter mistuning, suggesting that biological systems may rely on large populations or corrective sensory inputs for robust function.

The limitations of such networks have been formalized: Noorman et al. (2022) showed that a finite one-dimensional network has a “bumpy” energy landscape in general, leading to a discrete set of stable states. This discreteness produces velocity thresholds, saturation at

*Note that the terms “energy” and “energy landscape” will be used as metaphors; to date, there is no proof that the dynamics of the Burak and Fiete (2009) system is actually governed by a Lyapunov (or “energy”) function for non-zero velocity input.

low inputs, and inconsistent responses to constant drive, in contrast to the unbiased continuous integration expected from an ideal attractor.

Specifically, the bumpy energy landscape of a discrete 1D network leads to the following properties for the network dynamics (Noorman et al., 2022):

- In the absence of input, the network will be attracted to local energy minima, resulting in a discrete number of stable states.
- In the presence of low input velocity, the system cannot follow the input, and the output saturates at the nearest stable point. This can create a minimum velocity threshold that the system needs to exceed in order to start moving. Therefore, if the input is sufficiently small, the system may fail to integrate, leaving the bump pinned at a fixed state until input exceeds a threshold. This can be likened to a car being pushed uphill, experiencing decreasing speed for the same force, and eventually reaching a point where it can no longer move unless the force (input) is increased.
- The network processes higher input velocities so that, depending on the duration of the input, the system transitions to one of multiple distinct attractor states. As a result, the response to a constant input is not consistent, but rather displays peaks and troughs that correspond to the discrete nature of the energy landscape.

These effects are explored in detail and characterized for the 2D case in later sections in this chapter.




Figure 3.3: Example to illustrate problems in path integration caused by lack of symmetry. The difference in the y-positions of the center of mass of an activity bump on the neuronal sheet from two simulations where the network receives a constant input of 1m/s along 5° and -5° relative to the x-axis for 50s. During that time period, the patterns move by a total distance of ≈ 1750 neurons.

3.1.4 ASYMMETRY OF GRID CELL LOCATIONS

In the model of [Burak and Fiete \(2009\)](#), neurons on the sheet are arranged in tiles of 2×2 , with neurons of each of the 4 preferred directions arranged in tiles of $\begin{pmatrix} N & S \\ E & W \end{pmatrix}$.

This leads the system to behave differently depending on the angle of the input. The elementary NSEW-quadruple is not mirror symmetric. For example, when the network receives a constant velocity input along 5° and along -5° , the output (in this case measured as the position of the center of mass of an activity bump as a function of time) is the same in the x- but varies in the y-direction. This difference is shown in figure 3.3.

3.2 THE SOLUTION: A FOUR-SHEET MODEL

Following [Widloski and Fiete \(2014\)](#); [Kang and DeWeese \(2019\)](#) and others, we propose a modification to the model of [Burak and Fiete \(2009\)](#), where a grid module consists of




Figure 3.4: The (a) architecture and (b) activity of our modified model with 4 sub-populations corresponding to the directional preferences θ of the neurons. Thus, at every x, y location on the neuronal sheet, there are 4 neurons. The contribution each neuron receives from neurons of all directional preferences is equal; the only difference comes from the external input. So, the activity pattern looks the same with zero input and can look slightly different when v_{in} is non-zero, depending on the direction. In (b) the velocity is towards only along the -y (south direction) and no input in the x-direction.

a 2D sheet, and at each location on this neuronal sheet, we have 4 neurons, each with a different directional preference. Thus, the position of neuron i , θ on the neuronal sheet is characterized as $(x_i, y_i, \theta) = (\vec{r}_i, \theta)$ and where $\theta \in \{E, W, N, S\}$.

The activity of each neuron is a function of time and its position on the sheet. Analogous to equation 2.3, the synaptic weights are given by:

$$W_{ij\theta} = W_0(\vec{r}_i - \vec{r}_j - l\hat{e}_\theta) \quad (3.2)$$

where W_0 is a function of the relative distance and direction between neurons i and j , and $\hat{e}_\theta$ represents the unit vector in the direction θ of the sending neuron.

For the network depicted in Figure 3.4, the weights W_{ij} are defined similarly to equation 2.3, with one key modification: neurons with all four directional preferences lie at the same position $\vec{r}_i$ on the neural sheet. This effectively removes any asymmetry arising from the positions of the sub-populations on the 2D sheet. We modify the weight matrix to maintain consistency with the original system by setting the shift l to 1 instead of 2 neurons.

Thus, the network equation for one module is:

$$\tau \frac{ds_{i\theta}}{dt} + s_{i\theta} = f \left[\sum_j \sum_{\theta'} W_{ij\theta'} s_{j\theta'} + B_{i\theta} \right] \quad (3.3)$$

where $s_{i,\theta}$ is the activity of neuron i with directional preference θ , and θ' represents summation over all directional preferences.

The external input $B_{i\theta}$ is defined as:

$$B_{i\theta} = 1 + \alpha(\hat{e}_\theta \cdot \vec{v}_{\text{in}}) \quad (3.4)$$

where α is a scaling factor, $\hat{e}_\theta$ is the unit vector in the direction θ , and $\vec{v}_{\text{in}}$ is the input velocity vector.

3.2.1 INVESTIGATED MODEL PARAMETERS

Final parameters used with the basis of [Burak and Fiete \(2009\)](#) and the considerations from the previous section are listed in table 3.1. The coordinate Y' corresponds to an axis inclined at 60° with respect to the x-direction.

Note: A pinning input is applied during the pattern stabilization period to control the

Table 3.1: Default simulation parameters for a single module model.

Number of neurons along x	$n_x = 30^*$
Number of neurons along y	$n_y = 26^*$
Velocity gains	$\alpha^{(1)} = 0.2^*$ $\alpha^{(2)} = 0.3^*$
Time Step (ms)	$dt = 1$
Total simulation time (ms)	$simdur = 50000$
Time without velocity input (ms)	$t_{stab} = 2000$
Neuronal time constant (ms)	$\tau = 10$
Weight matrix within the module	$\gamma/\beta = 1.1$ $\beta = 3/\Lambda^2, \Lambda = 15^*$
Number of spatial bins along X and Y for rate map	80 or 150
Arena Radius (cm)	90

orientation of the output pattern. This is a two-point input along the x-direction with a period of Λ .

Our second configuration uses the same parameters as in Table 3.1 except for the following modifications: $\Lambda = 30$, $n_x = 60$, and $n_y = 52$ to halve the discreteness $\delta n/2$, and $\alpha^{(1)} = 0.4$, $\alpha^{(2)} = 0.6$ to ensure the distance moved on the neuronal sheet by a given input is identical in both cases.

3.3 DECODING

3.3.1 NETWORK PHASES

In order to quantify the movement of the activity pattern, one can either track the movement of a single bump or of the entire sheet. A single bump can be tracked by looking at the center of mass of a bump and how it changes as a function of time. However, this approach relies on the shape of the bump, the tracking window, and leads to discretization

errors.

Thus, we tracked the ‘phase’ of the entire pattern using the fundamental idea that any translations of the pattern only change its Fourier phase spectrum, and not the amplitude spectrum. In our simulations the amplitude remained constant over analysis time window.

Let $f(x, y)$ be a hexagonal pattern on the neuronal sheet. When the pattern moves by $\vec{c} = [c_x, c_y]$, the new pattern can be described by $g(x, y) = f(x + c_x, y + c_y)$. In the Fourier space, one can determine the movement vector $\vec{c}$ using the difference in the phases of the Fourier transforms of f and g :

$$k_x c_x + k_y c_y = \arg(\tilde{f}(k_x, k_y)) - \arg(\tilde{g}(k_x, k_y)) \quad (3.5)$$

where $\vec{k} = [k_x, k_y]$ is the wavevector of the hexagonal pattern and $\tilde{f}, \tilde{g}$ are the Fourier transforms of f and g respectively. Determining this at each time step allows one to get a comprehensive description of the pattern movement over time. Fig 3.6 shows the behavior of $\vec{c}(t)$ as a function of the initial phase of the pattern. Note that phases will generally be labeled in terms of inter-neuron distances, i.e., a phase difference of one neuron corresponds to 2π radians.

3.3.2 NEURONAL PHASES

The system is now simulated with all N_{traj} input trajectories, to get each neuron’s firing rate map. These maps are determined within the unit cell in physical space. The unit cell is a minimal version of the pattern, which reconstructs the entire periodic pattern when repeated. In this case, we are considering a rhombic unit cell with one axis along 0° and the y' -axis 60° inclined (fig 3.5) with side corresponding to the length scale of the module. We




Figure 3.5: Example rate maps of a neuron from a particular module, within the unit cell binned (150×150). One can see that there is only one firing field and since module 2 has a smaller scale, the average activity per bin is higher.

bin the unit cell with $nBins_x \times nBins_y$ spatial bins to get the firing rate map. The center of mass of the firing field in the rate map within a unit cell gives the 2D neuronal phase $\varphi = (\varphi_x, \varphi_{y'})$.

This is calculated using the circular mean - with each bin assigned a phase between 0 and 2π and weighted by the activity of the neuron in the respective bin. The x-phase of the neuron n with directional preference θ in module k is determined using:

$$\varphi_x^{(k)}(n, \theta) = \arg \left(\sum_{j_x=1}^{nBins_x} \exp \left(\frac{2\pi i \cdot j_x}{nBins_x} \right) r_{n\theta}^{(k)}(j_x) \right) \quad (3.6)$$

where, $r_{n\theta}^{(k)}$ is the spatial firing rate normalized by occupancy as shown in the rate map for neuron n in module k . Similarly, one can calculate the phase along the y' axis.

The distribution of phases has redundancy corresponding to the multi-bump structure. Thus, for each module, the phase distribution consists of 390 unique points, corresponding to 30 neurons along the x direction and $15 \sin(\pi/3) \approx 13$ neurons along the

y-direction. Each phase is shared by 8 other neurons (2 bumps in the same row $\times$ 4 directional sheets), and thus the phase distribution accounts for all the neurons of the module. As discussed in Section 3.1, only the bumps with same x -coordinates share the phases due to the period being odd ($\Lambda = 15$).

These phase distributions can then be used to decode the trajectories directly using the activities of the corresponding neurons at each time step.

3.3.3 DECODING POSITION

Using the phases of the neurons, one can get the position encoded by each neuron in a fundamental domain. The decoded position (for novel and already seen trajectories) at time t , $X^{(k)}(t)$ is the circular mean of the phases weighted by the activity of the respective neurons. This is an approximation of the formula derived in [Stemmler et al. \(2015\)](#) with the spike number being replaced by neural activation.

$$X^{(k)}(t) = \frac{\lambda^{(k)}}{2\pi} \arg \left(\sum_{\theta \in \{E, W, N, S\}} \sum_{j=1}^{n_x n_y} s_{j\theta}^{(k)}(t) \exp(i\phi_{j\theta}^{(k)}) \right) \quad (3.7)$$

Here, the spatial phase of neuron j with directional preference θ in module k is $\phi_{j\theta}^{(k)}$ and $\exp(i\phi_{j\theta}^{(k)}) = \cos(\phi_{j\theta}^{(k)}) + i \sin(\phi_{j\theta}^{(k)})$ whose argument gives the circular mean.

3.3.4 DECODING ERROR

For trial i at time t , define the decoding error where $\vec{p}_i(t)$ is the current position and $\hat{\vec{p}}_i(t)$ is the decoded position:

$$d_i(t) = \|\hat{\vec{p}}_i(t) - \vec{p}_i(t)\|_2.$$

We use mean decoding error across N trials:

$$\bar{d}(t) = \frac{1}{N} \sum_{i=1}^N d_i(t).$$

3.3.5 THEORETICAL LIMIT ON DECODING ERROR

If we assume that the actual and the decoded positions can be approximated as two random walk processes on a circular disk of radius R , then the expectation value $\langle e \rangle$ of the distance between two points drawn from a uniform distribution on the disk would give us an upper bound for the average asymptotic decoding error. This is the solution of an integral of the distance between two points at positions $(r_1, 0)$ and (r_2, θ_2) :

$$\langle e \rangle = \frac{4}{\pi R^2} \int_0^{2\pi} \int_0^R \int_0^R \sqrt{r_1^2 + r_2^2 - 2r_1r_2 \cos(\theta_2)} r_1 r_2 d\theta_2 dr_1 dr_2 \quad (3.8)$$

This is given by [Dunbar \(1997\)](#) to be: $\frac{128R}{45\pi}$, which for our system is ≈ 81 cm.

3.4 DISCRETE ATTRACTOR STATES

We investigated two different 2D neuronal sheets, both in the absence and presence of external inputs, and observed their asymptotic behavior to understand attractor states. We examined the final states and differences between the principal directions of the network.

3.4.1 ATTRACTORS AND THEIR BASINS

For our first in-silico experiment, we gave the system no input but varied initial conditions. To examine the nature of our system’s energy landscape, we used rigidly translated versions of the stationary activity pattern as initial conditions. The system was initialized using 20 translations in the x -direction and 10 in the y -direction.

Without any external input, the pattern on the neuronal sheet (30×26 neurons) was then allowed to settle, and the dynamics of these 20×10 patterns were observed. Due to the odd number of neurons corresponding to a period (15 neurons) in the x -direction, bump centers of mass would have $(15, 0)$ or $(7.5, 13)$ as their phase differences. Thus, stable states corresponding to a movement by 0.5 neurons in x are equivalent (because of the unit cell 15×13 , the x -maxima are at 0, 7.5, 15 etc), whereas only a movement by integer number of neurons in y would be equivalent to each other. This produces an inherent asymmetry in the x - and y -directions.

Each attractor’s basin of attraction (the set of initial conditions whose long-term behavior approaches the given attractor state) is shown in figure 3.6. Note that the system consists of a finite number of point attractor states and is thus a discrete attractor (Noorman et al., 2022).

The presence of the stable states at $x = 0.25$ neurons and $x = 0.75$ neurons is explained by A simple explanation is the discretization imposed by the 15×13 unit cell, possibly via a second-harmonic component of the pattern. We did not quantify this component.

For our second in-silico experiment, the network received a step input of 0.5m/s at $t = 2$ s which was withdrawn at different times ranging from $t = 4$ s to $t = 6$ s in steps of 0.1s. Since the network’s response to a constant input is not constant, we expected the system to




Figure 3.6: Trajectories for the 30×26 network for different initial conditions, color-coded with the normalized speed along the transients. The red dots represent the equilibrium states of the network. Four points — $(0, 0)$, $(0.25, 0)$, $(0.50, 0)$ and $(0.75, 0)$ — are attractors. The points at $(0.25, 0.5)$ and $(0.75, 0.5)$ are *saddles*: stable for perturbations along x but unstable along y . The remaining two red dots, $(0, 0.5)$ and $(0.50, 0.5)$, are unstable. Thus, initial states $(0, 0.499)$ and $(0, 0.501)$ converge to $(0, 0)$ and $(1, 0)$ respectively. The x -axis is aligned with one basis vector of the hexagonal lattice, whereas the y -axis is orthogonal to x and not parallel to the second lattice vector.




Figure 3.7: For a step velocity input of different durations (inset), the response of the system (30×26 , main figure) has the same attractor states as that of the one with shifted initial states as would be expected. This can be seen as the attractor states match the red dots for zero input above. Here, the initial y-phase was set to zero, and velocity input was applied only in the x-direction for the sake of simplicity.

end up in different basins of attraction depending on when the input was withdrawn (see fig. 3.7). This is indeed the case and the asymptotic states in this case are identical to the red dots in fig. 3.6 with $y\text{-phase} = 0$. This is because the initial $y\text{-phase}$ was set to zero, and velocity input was applied only in the $x\text{-direction}$ for the sake of simplicity.

Keeping in mind the discreteness described by Noorman et al. (2022) for head direction systems, we looked at the response of the network to different input velocities (see figures 3.8 and 3.9). In our system, the 2D analogues are:

- If the system does not start at a local energy minimum, it will move towards it, even in the absence of external input.
- For small inputs, the system cannot climb over the neighboring maximum (and the phase comes to a stand still even when the non-zero input persists).
- For larger inputs, the system has a jagged movement over different stable states and moves to the nearest stable state once the input is removed.

All these discretization effects in this deterministic system lead to a drift when time-varying inputs are given. Other than the size of the input, response differences between the different input directions can also result from the discretization.

3.4.2 LARGER NETWORK

To quantify the effect of discretization, we replicated the simulation with neurons placed at half the original distance, yielding a 60×52 sheet whose unit-cell length is $\Lambda \approx 30$ neurons (Fig. 3.10). This denser lattice alters the basins of attraction, the number of stationary solutions, and their stability.




Figure 3.8: For a fixed shift of the equilibrium bump-pattern along the y-axis (0.3 neurons), we checked the response of the 30×26 system to a constant input v , introduced at $t = 4$ s. (a) For very low inputs, the initial states approach nearest stable state. (b) For low inputs, the system cannot overcome the energy barrier and responds with a constant phase. (c) A transition state where one of the energy maxima is overcome, but the velocity is not large to overcome the next one. (d) The network integrates larger inputs non-linearly in this jagged manner. We only need to show $\bar{X}$ phases with initial states in $(0, 0.5)$ due to the symmetry of the system.




Figure 3.9: With zero input in x and a constant input in y introduced at $t = 4s$, (a) the x -phases from different initial conditions converge during the transient and then remain at fixed values. (b) the y -phase first relaxes to 0 in the absence of drive and, shifts to and remains at a nearby constant value near 0, after the input is introduced. The traces illustrate three effects of a discretized attractor: pinning (the bump remains locked once it reaches a local basin), thresholds (weak velocities are insufficient to move the bump), and shifted equilibrium (the fixed point is displaced without producing continuous drift).

Firstly, doubling the neuron density sharpens the spatial representation: bumps are now identical copies of one another, whereas in the 30×26 sheet odd spacing forces slight shape differences.

Secondly, increasing the sheet to 60×52 neurons changes both the number and the spacing of equilibrium points. In the 30×26 sheet four fully-stable minima are separated by $\Delta x = 0.25$ neurons (0.017λ), whereas in the 60×52 sheet, the nearest-neighbour distance between successive minima is $\Delta x = 0.17$ neurons (the smaller of the two arcs $0.330.50$ and $0.500.66$), producing a denser lattice of pinning wells (Fig. 3.6 vs 3.10). This trend matches the finite-size analysis of pinning barriers in continuous attractors by [Noorman et al. \(2022\)](#).

Lastly, some equilibria are saddle points: they are restoring in one phase direction (e.g. y) but repelling in the other. The relative proportions as well as the locations of minima and saddles differ between the two networks.




Figure 3.10: Trajectories of the 60×52 network launched from a grid of initial translations (colour indicates normalized speed along the transient). Red dots mark the equilibrium positions reached with zero external input. Three fully-stable attractors lie at phase coordinates $(x, y) = (0.33, 0)$, $(0.50, 0)$, $(0.66, 0)$ and so on. Saddle points, at $(0, 0.5)$, and $(0.5, 0.5)$, are stable along x but unstable along y and vice-versa for the points at $(0, 0)$.




Figure 3.11: 60×52 network response with different αv values for inputs along x-direction with different starting x-phases and initial y-phase 0.3 neurons. At $t = 4$ s, the network receives a constant input speed v along the x-direction. (a-b) For very low inputs, the initial states approach nearest stable state in the x- and y-directions. (c-d) For low inputs, there are two types of response: in one case, the system cannot overcome the energy barrier and responds with a constant phase (see top two traces in (c)). In the other, a transition state where one of the energy maxima is overcome, but the velocity is not large to overcome the next one (in this case because it is an unstable state in the x-direction at $x = 0$). Note that since $y = 0$ is stable in the y-direction, this behavior is similar to (b). (e-f) The network integrates larger inputs non-linearly in this jagged manner in the x-direction and shows qualitatively the same response as (b,d) in the y-direction.

In contrast, all the three properties corresponding to the three levels of input mentioned in the previous section can be seen in figure 3.11. Note that, since $x = 0$ is unstable in the x -direction, small inputs move the pattern to the nearest stable fixed point.

The threshold values of the input required to observe different dynamical regimes vary slightly between the two networks. For example, in the larger system, the lower threshold for low inputs and the threshold for a high-velocity response is higher. While one possible explanation is that finer discretization in the larger network produces shallower basins for low inputs and higher barriers at higher inputs, we have not directly quantified the valley depths or barrier heights to confirm this.

3.4.3 IMPLICATIONS FOR THE GRID CELL SYSTEM

When the velocity input is set to zero, the bump simply relaxes to a handful of fixed points (Fig.3.6). These equilibria are not interpreted as dedicated “memory” states; they emerge solely because a finite lattice samples a continuous attractor imperfectly, producing shallow pinning wells at a few phase offsets. Their direct biological relevance is therefore limited. Nonetheless, the wells have clear consequences once motion resumes: (i) local minima cause drift in the absence of input, (ii) inputs below the barrier height leave the decoded position locked, and (iii) very strong inputs bypass the wells but still introduce phase-dependent gain, so drift accumulates over closed paths.

In a strictly deterministic network the response to velocity is therefore non-linear and higher resolution generally improves precision.




Figure 3.12: The slopes of the linear fits between input and output differ across angles. This shows that as a function of the movement direction, velocity inputs vary in their effect on the displacement of the grid pattern. As a consequence, even in its deterministic form, the investigated CAN model cannot support path integration.

3.5 RESPONSE TO CONSTANT INPUT VELOCITIES

The system is given constant input speeds along different directions for 50s and the average flow rate of the pattern on the neuronal sheet is measured to get the response. This flow rate is fit as a linear function of the input speed for a fixed angle in the range of velocities considered ($[0, 1]$ m/s) as seen in fig. 3.12 for different input speeds every 0.1 m/s. The average of different slope curves along different input directions (sampled every 9°) is how we get the scaling factor between neural and physical space. The x -intercepts are used to determine the minimum thresholds along different directions. The summary is given in Fig. 4.2.




Figure 3.13: Direction-dependent decoding bias in Module 1. (a) Angular decoding error (decoded direction – true direction) as a function of input direction. (b) Flow rate expressed as a fraction of the flow rate at 0°. Both panels show curves for input speeds 0.1–1.0 m/s in 0.1 m/s steps; color indicates input speed.

3.5.1 DEPENDENCE ON THE ANGLE OF THE INPUT TRAJECTORY

The direction of the input velocity does not have to correspond to the direction of the pattern flow. We want to quantify the error in flow direction as a function of the input speed and input angle. Across speeds from 0.1 m/s to 1 m/s the angular decoding error follows an approximate $\sin 2\theta$ pattern (Fig. 3.13(a)).

For each curve, we divide the flow rate at a particular input direction by the flow rate at the same input speed along 0° to determine shape differences for all the speeds. The decoded speed shows a systematic angle dependence: it is highest near 90°/270° and lowest around 0°/180° axes, with the relative gain profile similar for all tested speeds (Fig. 3.13(b)).

3.5.2 SHEET-DEPENDENT RESPONSE DIFFERENCES

Depending on the input magnitude and direction, different sheets show a change in their shape. This effect is most pronounced in the case of an input along the preferred directions.




Figure 3.14: The activity pattern formed differs across the four neuronal sheets corresponding to preferred directions (East, West, North, South). The input is applied along 0° (East), and its speed is varied. The resulting bump size, measured as the full width at half maximum (FWHM), is plotted for each sheet. (a) FWHM values for sheets with preferred directions perpendicular to the input (North and South) show minimal variation; their curves overlap in the central region. (b) In the sheet aligned with the input (East), bump width increases with input magnitude. (c) In the sheet opposite to the input (West), bump width decreases with increasing input. The color coding is as follows: green represents the East sheet (aligned with input), red represents the West sheet (opposite to input), and the two overlaid curves in grey represent North and South (perpendicular to input). The three line styles in each plot correspond to FWHM estimates along the x- and y-directions, and their mean.

Figure 3.14 shows the full-width at half-maximum as a function of the input speed in the x -direction.

4

Model response to random trajectories

In this chapter, we focus on the response of the model to real and simulated trajectory inputs. We describe our decoding methodology and discuss the problem of catastrophic read-out.

4.1 INPUTS AND DECODING

4.1.1 INPUT TRAJECTORIES

A large number $N_{traj} \approx 100$ of random walk trajectories are used as inputs to ensure that the two modules' unit cells are completely covered (occupancy of bins is non-zero everywhere in the arena). Our trajectories are generated using the algorithm specified in [Raudies and Hasselmo \(2012\)](#) for a circular arena of radius 90 cm.

The system has some inertia and cannot follow instantaneous changes in the input (see Section 3.1.3 for details on pattern inertia). This allows us to reduce the frequency of speed updates (we chose every 20th time step), to reduce the computational costs, and linearly interpolate in between.

Generating the trajectories involves a combination of deterministic algorithms and random components. The process starts by defining the initial position to be at the center of the arena and a randomly chosen heading direction of the simulated point-sized agent. The running speeds are drawn from a Rayleigh distribution, and rotational velocities are sampled from a normal distribution with mean $\mu = -2.5^\circ/\text{s}$ and standard deviation $\sigma = 350^\circ/\text{s}$ from [Raudies and Hasselmo \(2012\)](#). The probability density of the Rayleigh distribution is:

$$p(v_{in}; \sigma) = \frac{v_{in}}{\sigma^2} \left(e^{\frac{-v_{in}^2}{2\sigma^2}} \right) \quad (4.1)$$

Here σ is the scale parameter of the distribution, such that the mean of the input distribution is 0.17m/s. We chose this based on the running speeds of animals as in [Raudies and Hasselmo \(2012\)](#).

At each 20-ms step, we check whether the agent is within 2-cm of the nearest wall. If




Figure 4.1: Trajectories are generated using an algorithm based on [Raudies and Hasselmo \(2012\)](#) with an arena of radius 90 cm resembling that of [Hafting et al. \(2005\)](#). In (a) we see the histogram corresponding to random draws from a Rayleigh distribution for the input speeds. (b) shows the histogram for the angular speed which is approximately a Gaussian with zero mean. (c) shows a representative 5-minute trajectory showing the animal's position in 2D space, with the center of the arena (0, 0) as the animal's starting point. (x and y in meters). (d) shows 30 different simulated trajectories overlaid to show that together they sample the entire arena.

it is and its heading points toward the wall (angle to the outward wall normal $< 90^\circ$), we apply a wall-avoidance update: (i) rotate the heading by the shortest amount needed to be parallel to the wall, and, (ii) compress the forward speed toward a floor $v_{min} = 0.05$ m/s via $v \leftarrow v - 0.5(v - v_{min})$. When the agent is outside the border band or not headed toward the wall, it follows the baseline random-walk dynamics. The resulting speed and turn-angle distributions and a sample trajectory are shown in Fig. 4.1.

4.2 MODEL DYNAMICS

4.2.1 SUMMARY OF RESPONSE

As in [Burak and Fiete \(2009\)](#), the response to different input speeds is linear in the intermediate range of input speeds considered, but the response slope depends on movement direction. As explained in section 3.6, we have a threshold effect at low input speeds, and thus a non-zero x-intercept in the input vs. output speed curves.

In order to determine the input-output relationship of the system and a scaling factor between neuronal and physical space, we gave the system constant inputs for 52000 time steps (1 time step = 1ms) in steps of 0.1 m/s, in the range of 0 to 1 m/s. The first 2000 time steps were ignored to avoid any transients resulting from the change in input speed. We then calculated the average flow rate of the network (in neurons/s) over the remaining 50,000 time steps and plotted the input-output curve. The average over all the steps allows us to not have to consider short-term effects resulting from the inertia of the pattern movement.

The above is done for inputs along angles from $(0^\circ, 360^\circ)$ in steps of 4.5° . We then get the slopes for the input-output relation along these angles and take an average. The means




Figure 4.2: (a) The average slopes of a linear fit to input-output curves for each input direction and (b) the x-intercept from quantifying the thresholding for low input speeds.

were identical for corresponding bumps from all four sub-networks - and thus do not depend on the directional preference of a particular neuron. This slope is the one considered for determining the size of the unit cell and the calculation of phases.

The decoding error for module 1 is shown in figure 4.3(a). This is simply the distance between the actual and decoded positions.

4.2.2 MODEL OF MODEL

Given that we know the input-output characteristics of the system as a function of the input speed and angle as determined previously (see figures 4.2 and 3.13), we can model the decoded position and hence the decoding error directly. This allows us to understand the form of figure 4.3(a) without further simulations and lets us separate the contributions of speed mis-scaling and direction bias to the decoding error.

- Speed-only model: We kept the true heading for every time step and applied the above fitted gain/offset to the speed. The resulting error 4.3(b) stays well below the

curve in panel (a), indicating that speed mis-scaling alone adds only a modest contribution.


- Angle-only model: Here the true speed was preserved while each heading was shifted by the direction-dependent bias $\Delta\theta(\theta)$. Angular mismatch clearly dominates: panel (c) reproduces the bulk of the error seen in panel (a),
- Full model: We combine both terms: speed mismatch and directional bias.

Most of the error is explained by angular mismatch as opposed to speed mismatch (see fig4.3). Since our drawn speeds have a low probability of lying in the non-linear higher speed range (probability of speed $> 0.8 \text{ m/s} = 1.6 \times 10^{-5}$), the small contribution of the speed mismatch makes sense.


4.2.3 EFFECT OF BOUNDARIES

The errors in Figure 4.3 all show a noticeable kink in the mean error at approximately 250 time steps, followed by a subsequent slowdown. This kink, a sudden change in the error behavior, is a significant observation as it suggests a potential point of interest in the trajectory. We investigated the relationship between this kink and the arena's size to understand why this occurs. We found that the peak error time strongly correlates with the arena's size, suggesting that the kink may be linked to boundary encounters. In what follows, we investigate this hypothesis in detail.


(a) Result from Simulations



(b) Only Model Speed



(c) Only Angle



(d) Both Speed and Angle




Figure 4.3: Decoding error (euclidean distance between actual and decoded positions) for module 1 using (a) Simulations, (b) Modeling the variation in speed only and direction fixed to true heading, (c) Modeling variation in angle only and retaining true speed, (d) Using the complete model including speed miscaling of (b) and direction bias of (c). Gray traces: RMS position error of individual trajectories (20 ms bins). Colored traces: mean across trajectories in each panel: panel (a) solid blue, panel (b) dashed red, panel (c) dashed green, panel (d) dashed dark blue. Note, that the y-axes have different ranges in the 4 panels.




Figure 4.4: The mean and median times of the first boundary encounter for 3 radii of the arena, over a 100 trajectories each. Note that the right most panel has reflective boundary conditions and the y-axes have different scales. Since the distributions are long tailed, the median provides a better central tendency measure than the mean.

4.2.4 DIFFERENT ARENA SIZES

We asked whether the timing of the first boundary encounter tracks the timing of the first peak in the (trial-averaged) decoding error across arena sizes. For clarity, by interruption we mean a boundary-locked change in the mean error trajectory (the kink) occurring within a short window of the first wall contact.

- Circular arena, radius $R = 90$ cm: median first boundary time = 5.10 s; first peak of mean decoding error at 5.06 s ($\Delta = -0.04$ s).
- Smaller arena, $R = 45$ cm: median first boundary time = 2.56 s; first peak at 2.74 s ($\Delta = 0.18$ s).
- Larger arena, $R = 180$ cm: median first boundary time = 12.08 s; first peak at 13.14 s ($\Delta = 1.06$ s).
- Reflective boundary, $R = 90$ cm: median first boundary time = 5.30 s; first peak at 5.46 s ($\Delta = 0.16$ s).




Figure 4.5: By aligning the decoding errors from 50 trajectories at their first boundary encounter (grey line), we analyze the error behavior before and after the encounter. The error averaged across trajectories (blue line) falls after the boundary encounter.

Across arena sizes and boundary conditions, the time of the first peak in the mean decoding error occurs close to the first boundary encounter and co-scales with arena size, consistent with a boundary-locked interruption of otherwise diffusive error growth. We did not, however, perform single-trial correspondences between first-hit times and peak times; establishing such within-trial linkage is outside the scope of the present analysis.

4.2.5 ERROR REDUCTION AT BOUNDARY

We aligned multiple trajectories at their first boundary encounter to analyze the error behavior before and after the encounter. Figure 4.5 shows that the average error decreases after a boundary encounter.

To further investigate this phenomenon, we looked at the input statistics near boundaries. One property of our simulated trajectories is that, by design, they slow down in the boundary region. We wanted to check if this effect is also observed in running trajectories of rodents in experiments to see if this effect could be generalizable.




Figure 4.6: Panels (a-b) use data from a single rat trajectory reported by (Stensola et al., 2012) in a square arena of size 150×150 cm. (a) Positions < 10 cm from the enclosure boundary, color-coded by speed at those points. Grey dots are speeds > 100 cm/s. (b) Kernel-density estimates of instantaneous speed for samples near the boundary (distance 10 cm to the nearest wall; red) versus the interior (> 10 cm; blue). Inset: Within the boundary subset, speeds are split by direction relative to the nearest wall's inward normal. Along Wall (dashed) = the middle 50% of angles between the velocity vector and the wall normal (i.e., roughly tangential); Towards/Away (solid) = angles outside this interval.

4.2.6 EXPERIMENTAL TRAJECTORY

To examine the input statistics of an actual trajectory, we used data from Stensola et al. (2012) where a rat runs in a $3\text{m} \times 3\text{m}$ square enclosure. The boundary region was defined as positions within 10 cm of the nearest wall, as shown in Figure 4.6(a). This choice was motivated by the larger arena and that unlike the simulated trajectories, the experiments were not carried out with point-like agents but with (spatially extended) rats. Thus, the size of the boundary region is considered based on the choice by Hardcastle et al. (2015). As expected, the probability of lower speeds is higher near the boundary (Figure 4.6(b)).

Thus, both experimental data from Stensola et al. (2012) and simulated trajectories show slower speeds near the boundary. This observation raised the question: Could the slowing down near boundaries explain the observed effect?




Figure 4.7: For three different τ , the network received different velocity inputs for 3 s and then subsequently the agent came to an abrupt stop. We did this to see the effect of slowing down without encountering boundaries. We observe a transient error reduction of (a) 0.105 cm, (b) 0.427 cm, and (c) 1.03 cm for the three time constants. The time behavior after the stopping is reminiscent of Fig. 3.8, corresponding to the network approaching different attractor states after input is withdrawn differently.

4.2.7 ILLUSTRATION

An analogy can help understand this phenomenon: When a swimmer moves quickly through a pool, the wave (representing the decoded position) trails behind. The waves catch up as the swimmer stops at the pool's walls. To quantify the effect on the decoding error, we used a stopping protocol, where the network received different velocity inputs for 3 s and then subsequently the agent came to an abrupt stop.

Using the leaky integrator x as a rough approximation, we get:

$$\tau \frac{dx}{dt} = -x(t) + v(t) \quad (4.2)$$

The effect up to the first order will be proportional to $v\tau$. Since for our system $\tau = 10$ ms, the magnitude is small and timescale is on the order of τ . Thus, the leading order effects are small and fast, and thus insufficient to explain the observed kink and subsequent slowdown in the error. Figure 4.7 demonstrates that varying τ shows some error reduction but does not fully account for the observed effect.

Given the bound above, simply slowing down near boundaries does not entirely account for the error reduction. A plausible additional mechanism is attractor relaxation at low drive: When $v(t) \approx 0$, recurrent dynamics of the continuous attractor relax the phase toward a nearby stable state with a network time constant $\tau_{\text{net}} > \tau$. Also, motion within a thin region around the boundary is directionally biased (towards/away constrained, along allowed), so $\sqrt{\langle v_{\parallel}^2 \rangle + \langle v_{\perp}^2 \rangle}$ is reduced when $\langle v_{\perp}^2 \rangle$ is small (with $v_{\parallel} = \mathbf{v} \cdot \hat{\theta}$ and $v_{\perp} = \mathbf{v} \cdot \hat{\mathbf{r}}$). However, exploring these other possibilities falls beyond the scope of this study.

5

intermodular connections

This chapter describes the intermodular connection scheme we propose. Then, we discuss the results for the uncoupled systems and the effect of different types of coupling.




Figure 5.1: Conceptual sketch of geometric coupling in one dimension. The agent moves along a linear track in physical space (black line, bottom), with positions at t_1 , t_2 , and t_3 marked. Each row of circles represents neurons from two modules: module 1 (blue) and module 2 (orange), whose spatial scales form a 3:2 ratio. Neurons are arranged periodically: after each module's spatial period (blue and orange boxes), the numbering restarts, reflecting the repeating structure of grid fields. Filled circles indicate neurons active at the agent's current position. Because of the integer scale ratio, module 2 neuron 4 is coactive with module 1 neurons 4, 8, and 12 at successive agent positions, motivating excitatory couplings between them (arrows). This schematic illustrates the coupling rule; in simulations both modules contain the same number of neurons, unlike in this simplified cartoon. Figure code generated with the assistance of ChatGPT.

5.1 A GEOMETRICAL HYPOTHESIS

We hypothesize that a coupling between different modules that ‘correctly’ wires the different activities can potentially serve as a mechanism to prevent the position estimate from rapidly deviating from the actual location of the animal. What happens when there are drifts and/or errors, and is it possible that having coupled modules helps with error correction?

Here we implement a model where we couple 2 networks of neurons, both exhibiting different spatial scales, and then we ask - can the coupling that exploits the geometrical relationship between modules correct for the drifts?

Figure 5.1 shows a simplified 1D sketch of activity in two grid-cell modules as an agent moves along a linear track. The horizontal black line at the bottom represents the phys-

ical track. Black dots mark the agent's positions at times t_1 , t_2 , and t_3 ; for example, at t_1 the agent is at position 4. Above the track, two rows of circles represent the two modules. Their periodic spatial scales are indicated by the blue (module 1) and orange (module 2) boxes. Numbers label the neurons within one spatial period of module 1 (1–12) and module 2 (1 – 8). Filled circles denote active neurons, open circles inactive ones.

At t_1 , neurons 4 in both modules are coactive. At later positions t_2 and t_3 , module 2 neuron 4 remains active while neurons 12 and 8 in module 1 are coactive with it. Thus, module 2 neuron 4 aligns with module 1 neurons 4, 8, and 12 (highlighted in blue).

Because the module scales form a ratio of $3/2$, our hypothesis is that excitatory connections link coactive neurons. In this example, module 2 neuron 4 couples to module 1 neurons 4, 8, and 12 as indicated by the arrows. All remaining excitatory projections from neurons in module 2 to module 1 can be determined in a similar manner.

In simulations both modules contain the same number of neurons, unlike in this simplified cartoon. Additionally, in order to not disturb the formation of the pattern, we introduce weak inhibitory connections between neuron 4 in module 2 and all other neurons in module 1 to balance excitation and inhibition. In one dimension, this yields 3 couplings; in two dimensions, the rule extends to $3 \times 3 = 9$ couplings.

5.2 CORRELATIONS BETWEEN NETWORKS

A good starting point to get the connections between the modules would be to determine the correlations between the activities of the neurons between the modules. Since the ratios of the scales is chosen to be exactly 1.5, we would expect a regular pattern similar to fig 5.2. This is the consequence of uniform and exact translation of the activity pattern across the




Figure 5.2: Panel (a) shows the correlations of module 2 neuron $(0, 0, E)$ with neurons of the E subpopulation in module 1. In this case the number of firing fields in physical space is four times the number of activity bumps on the neuronal sheet, which gives 16 correlation peaks. Panel (b) shows the correlations of module 1 neuron $(0, 0, E)$ with neurons of the E subpopulation in module 2. Here the number of firing fields in physical space is nine times the number of activity bumps on the neuronal sheet, which gives 36 correlation peaks. These factors arise because the arena is set equal to the multi-module fundamental domain, which consists of four unit cells of module 1 and nine unit cells of module 2. The same correlation patterns are observed for the other three subpopulations, which are not shown. The results are illustrated for a 60×52 network with a period of 30 neurons.

two modules and the fact that the scale ratio is exactly $3/2$. This is then be used to get the connections between the modules (motivating a Hebbian-like correlation based connection scheme).

5.3 CONNECTIVITY SCHEME

Figure 5.3 illustrates the two-dimensional connectivity scheme. We show the connections from a single reference neuron in module 2, neuron $(0, 0, E)$, to neurons in module 1. Excitatory targets are shown as black dots and inhibitory targets as blue dots. The excitatory targets form 9 clusters, as expected from the 3×3 combinations of coactive neurons along the two axes. This generalizes the one-dimensional case (Figure 5.1), where one neuron in module 2 aligned with 3 neurons in module 1; in two dimensions, the same principle applies independently along x and y' , giving $3 \times 3 = 9$ clusters.




Figure 5.3: This figure shows the connectivity scheme in two dimensions. Connections from a single reference neuron in module 2 at position $(0, 0, E)$ to neurons in module 1 are plotted in phase space $(0, 2\pi)$ along the x and y' axes as defined in Eq. 3.6, where y' is oriented 60° relative to x . Black dots indicate excitatory targets, and blue dots indicate inhibitory targets. The excitatory targets form 9 distinct clusters, which correspond to the 3×3 combinations of coactive neurons along x and y' (compare to the one-dimensional case in Figure 5.1). Each cluster contains multiple neurons, because with an odd value of λ the activity bumps can produce either one or two maxima depending on their shape. All four module 1 populations (east, west, north, and south) exhibit the same clustered distribution. Phases are determined using the actual data and are not assigned automatically and hence have some undulations reflecting decoding errors.

To display these connections we use the phase coordinates defined in Eq. 3.6. In this representation, each neuron is assigned a pair of phases $(\varphi_x, \varphi_{y'})$ modulo the lattice periods, so that all repeating copies of the same grid field collapse into a single unit cell. This provides a compact view in which the repeating structure of the grid pattern is represented once within the unit cell $(0, 2\pi) \times (0, 2\pi)$. Phase space is therefore the natural coordinate system for identifying clusters of coactive neurons, which would be spread across many locations if plotted directly in physical space.

Each cluster contains more than one dot because neurons that are equivalent in phase can differ depending on the shape of the underlying activity bump. With an odd value of λ , some bumps on the neuronal sheet have a single central maximum, whereas others exhibit two local maxima. These additional maxima show up as more than one co-active neuron. As a result, multiple dots appear within each cluster in the phase-space plot.

There are exactly 13 rows because the unit cell repeats after 13 steps in the y' -direction. Along the x -direction the discretization step is $1/30$ instead of $1/15$, because the presence of two local maxima for some bumps effectively halves the spacing when λ is odd. This asymmetry, causes the clusters to contain different numbers of neurons across rows. The overall pattern of 9 excitatory clusters is preserved and is the same across all four module 1 populations (E, W, N, S) .

5.4 FIXED POINTS IN A TWO-MODULE NETWORK

How are the stable states affected by weak intermodular coupling? How do they change based on the coupling strength?

For a two-module system, we introduced weak coupling between the modules with dif-




Figure 5.4: Attractor states with intermodular coupling from module 2 to module 1 are shown for the 30×26 system. (a) Module 2 states are uniformly distributed as expected. (b) Very weak coupling, module 1 states are unaffected. (c) Weak coupling, resulting in a superposition of the attractor states from both modules. (d) If the coupling is strong enough, then the states of module 1 resemble those of module 2. Note: We still need to keep the connections weak enough that they do not destroy the grid pattern of the network. See Fig. 5.7 for activity patterns at different coupling strengths.

ferent strengths (η) from module 2 (smaller scale) to module 1 (larger scale) once both systems settled at $t = 2s$. The stable states are shown in fig. 5.4 for an initial set of states with different shifts in the x-direction and no external input. Since the module scales have a ratio of 1.5, a shift in neuronal space of 2 neurons in module 1 corresponds to a shift of 3 neurons in module 2.

For very weak coupling ($\eta = 10^{-4}$), the stable states of module 1 are not affected by the input from module 2. For a coupling strength of $\eta = 10^{-3}$, module 1 stable states depend on those of module 2, especially when they belong to different basins of attraction, to begin with. For an even higher coupling strength, module 2 dominates module 1, making the final stable states of module 1 progressively resemble those of module 2.

A higher density of stable states allows the network to represent a larger variety of patterns or configurations. This can improve the network's ability to encode and distinguish between different inputs, enhancing its spatial resolution. In principle, the intermediate stable states could allow module 1 to drift less in cases with coupling as opposed to without coupling.

5.5 MULTI-MODULE MODEL AND DRIFT

5.5.1 MULTI-MODULE MODEL

In order to describe a multi-module model, we have two terms, one for within module (see eq 3.3) and the intermodule term with a connection strength η and an intermodular weight matrix $\tilde{W}$.

$$\tau \frac{ds_{i\theta}^{(k)}}{dt} + s_{i\theta}^{(k)} = f \left[\sum_{j\theta'} W_{ij\theta'} s_{j\theta'}^{(k)} + B_{i\theta}^{(k)} + \sum_{n \neq k} \eta_{kn} \sum_{j'\theta''} \tilde{W}_{ij'\theta''} s_{j'\theta''}^{(n)} \right] \quad (5.1)$$

where: $s_{i\theta}^{(k)}$ is the activity of neuron i with directional preference θ in module k .

The input to neuron (i, θ) in module k contains three terms:

- $\sum_{j,\theta'} W_{ij\theta'} s_{j\theta'}^{(k)}$: Sum of contributions from all neurons with all directional preferences in the same module k .
- $B_{i\theta}^{(k)} = 1 + \alpha(e_{\theta}^{(k)} \cdot \vec{v}_{\text{in}})$: External input to neuron i with directional preference θ in module k . By changing $\alpha^{(k)}$, we can control the scale of the module.
- Coupling term representing input from other modules $n \neq k$.
 - η_{kn} : Connection strength for coupling from module n to module k .
 - Sum of contributions from all neurons with all directional preferences in module n with the intermodular weight matrix $\tilde{W}$.

Since the modules have the same connectivity, the superindex k has been dropped above. Unless stated otherwise, we consider a 2 module system, with connections from module 2 to module 1. The form of the intermodular weight matrix is discussed in detail in section 5.3.




Figure 5.5: Mean squared error increases as a function of time. Each panel shows the fit to a piece-wise linear function with slopes a and b and x and y intercepts for the second fit t_0 and c_2 respectively. For different connection strengths, the rate of the error increase differs for module 1. By itself, module 2 provides the highest accuracy, as indicated by the small size of slope b shown in (b). But even very weak coupling to module 1 strongly improve the accuracy of that module, unless the coupling is too strong so that it destroys the hexagonal grid-field structure. All values were determined using 100 simulated trajectories.




Figure 5.6: The slope (b) for different coupling strengths summarizing the results in 5.5. When coupling is weak, the rate of error accumulation in module 1 goes down.

5.6 EFFECT OF COUPLING

5.6.1 SMALLER TO LARGER MODULE

Since the module with the smaller scale has a better resolution, we examined uni-directional projections from module 2 to module 1. This would result in intermediate steady states in module 1 (Fig. 5.4). The results for different coupling strengths are shown in Fig. 5.5. Since module 2 receives no additional input, its rate of error accumulation does not depend on η . But the slope of error accumulation in module 1 varies - with even a very weak coupling resulting in a sharp fall in the rate of error accumulation. In order to understand how the coupling affects module 1 activity patterns, we can look at the activity patterns in Fig. 5.7 at a time step t for module 1 for different coupling strengths. While very weak coupling has little to no effect (panels c-d), stronger coupling can change the bump shapes and even destroy the hexagonal pattern entirely. The cumulative effect of the coupling can be seen through the rate maps in Fig. 5.11 for single neurons in module 1 for different connection strengths. To be able to compare, we also depict the above for module 2. We see that for

module 1, the size of the fields (Fig. 5.8) depends on the coupling - with larger coupling strengths they start to get sharper and ultimately the pattern breaks down for a coupling strength 10^{-1} .


5.6.2 RANDOM CONNECTIONS

We compared our connection scheme with a control condition with the same number and strengths of connections as the previous case, but with randomized locations. The mean squared errors for two different connection strengths are shown in figure 5.9 for two connection strengths. They are clearly worse than correlation-based connections.


5.6.3 ONE-TO-ONE CONNECTIONS

Alternatively, we tried one-to-one connections for neurons at the same positions $(\vec{r}, \theta)$ on the neuronal sheet (fig. 5.10). Trivially, this is a subset of connections considered in the case of the correlations if the 2 sheets are initialized to be exactly in sync with each other before the input is given.


(a) Module 1, $\eta = 0$




(b) Module 2




(c) $\eta = 10^{-6}$




(d) $\eta = 10^{-5}$




(e) $\eta = 10^{-4}$



(f) $\eta = 10^{-3}$



(g) $\eta = 10^{-2}$



(h) $\eta = 10^{-1}$




Figure 5.7: The activity patterns for Module 1 subpopulation E for different connection strengths. While for very weak coupling strengths η (c-d), we do not see a big change in the activity pattern as compared to the uncoupled one (a). However, the peak activity increase with $\eta = 10^{-2}$ (g) and the pattern is completely destroyed for $\eta = 10^{-1}$ (h)





Figure 5.8: The average size of the fields in the rate maps characterized by the full-width at half-maximum of a fitted Gaussian. The size and the variability of the fields falls for larger coupling strengths, as can be seen in the figure 5.11. A part of the wiggles is due to discretized bins.

(a) $\eta = 10^{-3}$



(b) $\eta = 10^{-2}$





Figure 5.9: Mean squared errors for module 1 increase much more rapidly when we use random couplings from module 2 to module 1. In that case the performance of module 1 is worse than without any inter-module coupling.

(a) $\eta = 10^{-3}$



(b) $\eta = 10^{-2}$






Figure 5.10: Mean squared errors for module 1 increase much more rapidly when we use couplings from module 2 to module 1 that connect neurons with the same position $(\vec{r}, \theta)$ on the neuronal sheet. These are much worse than the uncoupled network.


(a) $\eta = 0$




(b) Module 2




(c) $\eta = 10^{-6}$




(d) $\eta = 10^{-5}$




(e) $\eta = 10^{-4}$



(f) $\eta = 10^{-3}$



(g) $\eta = 10^{-2}$



(h) $\eta = 10^{-1}$




Figure 5.11: The rate maps for different connection strengths using 10 simulated trajectories over 5 minutes. The white regions indicate unvisited bins in the arena. For the uncoupled and very weakly coupled systems, we can see that the fields are diffusing, whereas (b) and (c) are much sharper. (h) shows a combination of the fields from module 1 and module 2.

6

Discussion

We have shown that weak correlation-based coupling between modules improves the accuracy of decoding position for a larger module compared to an uncoupled network. Specifically, our model with a configuration of $26 \times 30 \times 4$ neurons shows a mean error accumulation slope for different coupling strengths is shown in Fig. 5.6. For context, the [Burak and Fiete \(2009\)](#) model with 40×40 neurons accumulates error at a rate of 0.03 cm/s, and

their larger 128×128 neuron model has an error rate of 0.01 cm/s.

We note that while their reported errors are based on a single trajectory, our values are averaged across multiple trajectories.

Furthermore, we characterized the discrete nature of the attractor network’s energy landscape, building upon the work of [Noorman et al. \(2022\)](#), which examined this landscape in 1D and extended it to 2D. Our model serves as a proof of concept for a simple yet novel mechanism for coordination between modules that uses their geometrical relationships.

6.1 CONCLUSIONS

6.1.1 ERROR REDUCTION THROUGH GEOMETRIC COUPLING

In our deterministic two-module model (3 : 2 scales), weak geometric input from the smaller-scale module reduced drift in the larger-scale module and yielded sharper fields. The effect reflects improved of the larger-scale readout; it does not increase resolution beyond the finer module.

6.1.2 BOUNDARY ERROR REDUCTION WITHOUT BOUNDARY CELLS

We observed a notable reduction in errors at the boundaries without including boundary cells as in [Hardcastle et al. \(2015\)](#). The deceleration near boundaries observed in actual trajectories is a partial explanation for this reduction in error. This suggests that the dynamics and input statistics improve boundary accuracy, providing an avenue for further research.

6.2 RELATION TO EXISTING MODELS

6.2.1 MOSHIEFF AND BURAK

Mosheiff and Burak (2019) proposed a model where the position estimate of one grid cell module is used as an additional input to other modules, effectively reducing error through dynamic feedback mechanisms. This synaptic connectivity between modules helps suppress incompatible drifts driven by noise while preserving the capacity of the grid cell code. This model dynamically adjusts to maintain stability, ensuring the representation remains robust against noise.

Mosheiff and Burak (2019) refer to fixed-phase relation as a condition where the phases (or positions) of grid cell modules have a constant relationship, akin to two gears meshing together such that their rotation angles are always identical. When modules are phase-coupled this way, there is effectively only one encoded angle, limiting the dynamic range of the representation. This means that every neuron's phase in one module directly corresponds to a fixed phase in another, reducing the system's capacity to represent a wide range of positions.

In contrast, our fixed geometrical coupling model simplifies the design and implementation by establishing a correlation-based relationship across modules, avoiding the complexity of continuous exchange of the position estimate and dynamic feedback. The primary advantage is its simplicity and ease of implementation using a Hebbian-like mechanism, making it biologically plausible.

Mosheiff and Burak (2019) suggest that synaptic connectivity between grid cell modules within the MEC can enforce the desired coupling, relying on broad, relatively unstructured

connections based on heading preference; these connections ignore absolute phase and allow any phase offset as long as the modules move at the same speed. By contrast, our geometric coupling links each coarse-scale neuron weakly to the subset of fine-scale neurons whose phases are compatible with the 3:2 scale ratio. This many-to-many pattern favours certain relative phases but does not impose a unique, fixed offset. Consequently it constrains inter-module drift without the strong phase locking, or the broad, phase-agnostic velocity coupling, used in Mosheiff and Burak (2019)’s formulation.

Our approach lacks the robustness of Mosheiff and Burak (2019)’s model. It cannot dynamically adjust to changes in neural activity patterns due to external stimuli or internal fluctuations, which could lead to a loss of accuracy and coherence in the system’s performance.

Mosheiff and Burak (2019)’s model significantly reduces error accumulation due to coupling. Without coupling, position representations by different modules can diverge due to noise, leading to catastrophic readout errors (Burak, 2014; Mosheiff and Burak, 2019). The mean squared error accumulation rate is reduced by a factor equal to the number of modules in the coupled system (Mosheiff and Burak, 2019), effectively eliminating catastrophic errors over time.

Future work should explore how weak couplings in our model affect the capacity of the grid cell code and whether they adapted in response to neural activity changes. Comparing both models’ long-term stability and error rates under various conditions could provide deeper insights into their respective advantages and limitations. Another approach could be combining fixed geometrical coupling for baseline stability with dynamic feedback for adaptability and noise suppression to achieve a robust and flexible system for grid cell repre-

sensation.

6.2.2 KANG AND BALASUBRAMANIAN

Kang and Balasubramanian (2019) propose a dynamical self-organization mechanism that may explain how grid modules with characteristic scale ratios naturally emerge in the MEC. Each module resembles a Burak and Fiete (2009)-like network, and different spatial scales are generated by varying inhibition distances. Their approach primarily focuses on achieving modular structures with fixed scale ratios through competition between intermodular excitation and intra-modular inhibition by incorporating local excitatory coupling between modules (Kang and Balasubramanian, 2019).

In contrast, our model emphasizes geometric coupling to reduce decoding errors and enhance robustness. While Kang and Balasubramanian (2019)’s model achieves modular self-organization, it does not explicitly address error correction mechanisms. Their reliance on geometric relationships between triangular grids to maintain constant scale ratios shows limitations, particularly when producing discrete scale ratios using a velocity gain approach. Our model starts with predefined scale ratios to ensure coherent activity across modules, simplifying the system while maintaining robust spatial representation and error correction.

Kang and Balasubramanian (2019) argue that using a velocity gain approach can result in slips or independent drifts in the relative phases of the modules, causing occasional catastrophic errors. They exclusively use excitatory coupling, which may contribute to slips in the relative phases of the modules when using velocity-based scaling. In contrast, our model employs a dual coupling strategy focusing on excitatory and inhibitory interactions be-

tween modules, where inhibitory couplings between out-of-phase neurons may result in a stabilizing effect and better error correction.

6.2.3 WAAGA ET AL. 2021

Waaga et al. (2021) analyzed grid-cell activity during foraging in light and darkness, quantifying both the absolute error of the decoded trajectory (MAE) and the relative alignment between modules (δ_i, Δ_{ij}). They found that while MAE increased substantially in darkness, inter-module distances remained unchanged, indicating preserved relative alignment even during high-error epochs. In contrast, our model tracks the absolute error of each module's estimate relative to ground truth, but does not quantify inter-module distances. Thus, our results can be directly compared to their MAE measures, but not to their (δ_i, Δ_{ij}) metrics. Future extensions of our model could incorporate inter-module distance analyses to enable a more direct comparison to their findings on preserved coordination across modules.

6.3 BIOLOGICAL SIGNIFICANCE

6.3.1 PREDICTED CORRELATIONS

The experimentally measured correlations between neurons from the same module are quite strong (Gardner et al., 2019; Trettel et al., 2019) and change as a function of the phase difference between the neurons (Barry et al., 2007; Stensola et al., 2012). However, intermodular correlations are more complex and variable (Trettel et al., 2019; Gardner et al., 2019). For example, environmental context, behavioral states, and cognitive demands can influence intermodular correlations. This flexibility allows the grid cell network to adapt to various spatial tasks and environments.

The circuit mechanisms that generate intermodular correlations remain unclear. One possibility is indirect coordination via the hippocampus and connected cortical areas; hippocampo-entorhinal feedback can disrupt or realign grid firing (Zhang et al., 2013). Another proposal is that weak, spatially patterned recurrent projections within the medial entorhinal cortex couple modules directly (Couey et al., 2013; Kang and Balasubramanian, 2019). In either case, synaptic plasticity and theta-phase-locked excitation are likely contributors to the residual synchrony observed across modules (Bush and Burgess, 2014; Pastoll et al., 2013). Our correlation-based coupling scheme is consistent with the empirical finding that inter-module grid-cell pairs exhibit only weak synchrony: the model employs very small inter-module weights, sufficient to reduce relative drift without inducing the strong, phase-locked correlations characteristic of cells within the same module.

6.3.2 BOUNDARY EFFECTS

Other grid cell models, including Burak and Fiete (2009) also observe a slowing down of the drift after a rapid initial increase (Figure 1a). Here, we identify and quantify the factors contributing to both the rapid increase (the geometry and size of the arena) and the subsequent slowing down (boundary encounters) of the drift.

We also observe a decrease in error immediately after boundary encounters, similar to the findings of Hardcastle et al. (2015) and consistent with grid-pattern shear reported in square arenas (Stensola et al., 2015). Unlike Hardcastle et al. (2015), we did not incorporate border cells into our model, demonstrating that the observed effects are at least partially intrinsic to the network’s dynamics and input statistics. Hardcastle et al. (2015) focused on quantifying errors due to intrinsic neuronal noise and spiking, which are factors that can

also be assessed in an extension to our model. Error reduction near boundaries may also be attributed to a combination of the two factors.

Since rodents tend to spend a significant amount of time near the walls (Chen et al., 2016; Pérez-Escobar et al., 2016), further exploration of the changes in statistical properties of the input near boundaries or the drift could be promising areas for future research. This includes examining how boundary conditions affect network stability and error rates and how these factors interact with input noise and spiking activity. By extending our analysis to include these elements, we can better understand the robustness and adaptability of grid cell networks in various spatial environments.

Finally, integrating our findings with experimental data from real-world navigational tasks is crucial. This integration can help validate and refine our models and more importantly, it can contribute to the development of biologically plausible models of spatial navigation, with potential real-world applications.

We measure and formalize the drift in deterministic variants of the Burak and Fiete (2009) model, providing important insights for other modelers. Our findings show that errors depend more significantly on the direction of movement relative to the network-intrinsic flow than on the flow rate itself. Additionally, we elucidate the inertia of the pattern flow, highlighting the limitations of this type of model. While parameter regimes where the network operates in a continuous state may exist, exploring such regimes would be important for future work. The relationship between the number, locations, and type of fixed points and the rate of error accumulation needs to be more carefully explored.

6.4 LIMITATIONS OF OUR APPROACH

Our model relies on the assumption that the ratio of the scales of the modules is exactly 1.5. Because the scale ratio is a ratio of two small integers (Stemmler et al., 2015; Wei et al., 2015a), we have a unique correlation structure with peaks that repeat periodically. If the ratio is not exactly 1.5, this correlation structure would become much less stable (Fiete et al., 2008; Stemmler et al., 2015; Wei et al., 2015a), with correlations that are not consistent but depend on the size of the arena considered.

We did not include input or spiking noise. We did not compute a joint multi-module decoder, measure inter-module alignment, or quantify changes in spatial period.

6.5 FURTHER WORK

Possible directions to explore include further examining the error at the boundaries, the response of the couplings to input noise, and exploring a spiking network to study the interplay between inter-module coupling and intrinsic neuronal noise. Finally, it would be illuminating to see if bidirectional correlation-based intermodular couplings might reduce the absolute drift.

Acknowledgments

Completing this PhD has been an unexpected and incredible journey, especially considering my family's entrepreneurial background. I am profoundly grateful to everyone who has played a pivotal role in this achievement.

First, I sincerely thank my parents, grandmothers, and brother for their unwavering support, patience, and positivity. Your belief in me has given me the courage and wings to explore new horizons.

I want to express my heartfelt thanks to my supervisor, Andreas Herz, for his relentless guidance, belief in my potential, and continuous push to excel scientifically.

To my TAC members, Caroline Geisler, Anton Sirota, and Christian Leibold thank you for your invaluable feedback and guidance. Martin Stemmler, your enthusiastic feedback has been incredibly motivating and always thought-provoking. Special appreciation goes to Florian Eberhardt for being a steadfast friend and confidant, Dora Csordas for her superhero-like support, Michaela Poth for her kind and wise words, Florian Mueller for his indispensable software and IT support and Aditya Chowdhury, for believing in me and helping expedite my coding processes.

To Sharmistha, Neha, Sanskriti, and Hasneet, your support through the highs and lows of this challenging journey has been invaluable. Thank you to all my teachers, from school and university, for their guidance and encouragement throughout my academic journey.

I also thank the GSN team: Lena, Verena, Stefanie, Katrin, Raluca, and the administrative staff: Silke, Martina, and Monika. Your assistance and support have been crucial. Thanks also to the IT department, Mr. Schimmel and Mr. Molter.

Last but certainly not least, I want to acknowledge the incredible women who have supported and inspired me throughout this journey. Your laughter, food, and unwavering strength as role models have been a constant source of inspiration: Elena Kalantzis, Reema Gupta, Naiiri Ghazari, Preshika Wright, Erika Lupieri, Anna Umurzakova, Mary Antonova, Poornima Ramesh, Rangoli Saxena, Sai Ayachit, Julia Mayer, Miao Wang, Mariam Museridze, Katharina Bracher, Ann Kotkat and Magdalena Kautzky.

Thank you all for being a part of this journey. Your support and encouragement have made this achievement possible.

References

- Allen, K., Gil, M., Resnik, E., Toader, O., Seeburg, P., and Monyer, H. (2014). Impaired path integration and grid cell spatial periodicity in mice lacking glua1-containing ampa receptors. *Journal of Neuroscience*, 34(18):6245–6259.
- Alonso, A. and Klink, R. (1993). Differential electroresponsiveness of stellate and pyramidal-like cells of medial entorhinal cortex layer ii. *Journal of Neurophysiology*, 70(1):128–143.
- Alonso, A. and Llinás, R. R. (1989). Subthreshold na⁺-dependent theta-like rhythmicity in stellate cells of entorhinal cortex layer ii. *Nature*, 342(6246):175–177.
- Amari, S.-i. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, 27(2):77–87.
- Barry, C., Ginzberg, L. L., O’Keefe, J., and Burgess, N. (2012). Grid cell firing patterns signal environmental novelty by expansion. *Proceedings of the National Academy of Sciences*, 109(43):17687–17692.
- Barry, C., Hayman, R., Burgess, N., and Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature neuroscience*, 10(6):682–684.
- Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O’Keefe, J., Jeffery, K., and Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, 17(1-2):71–98.
- Blair, H. T. and Sharp, P. E. (1996). Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat. *Behavioral neuroscience*, 110(4):643.
- Boccaro, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., and Moser, M.-B. (2010). Grid cells in pre-and parasubiculum. *Nature neuroscience*, 13(8):987–994.

- Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J. L., Roudi, Y., Moser, E. I., and Moser, M.-B. (2013). Grid cells require excitatory drive from the hippocampus. *Nature neuroscience*, 16(3):309–317.
- Brandon, M. P., Bogaard, A. R., Libby, C. P., Connerney, M. A., Gupta, K., and Hasselmo, M. E. (2011). Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. *Science*, 332(6029):595–599.
- Brun, V. H., Otnæss, M. K., Molden, S., Steffenach, H.-A., Witter, M. P., Moser, M.-B., and Moser, E. I. (2002). Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. *Science*, 296(5576):2243–2246.
- Burak, Y. (2014). Spatial coding and attractor dynamics of grid cells in the entorhinal cortex. *Current opinion in neurobiology*, 25:169–175.
- Burak, Y. and Fiete, I. R. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol*, 5(2):e1000291.
- Burgess, N., Barry, C., and O’keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9):801–812.
- Bush, D., Barry, C., Manson, D., and Burgess, N. (2015). Using grid cells for navigation. *Neuron*, 87(3):507–520. This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.
- Bush, D. and Burgess, N. (2014). A hybrid oscillatory interference/continuous attractor network model of grid cell firing. *Journal of Neuroscience*, 34(14):5065–5079.
- Bush, D. and Schmidt-Hieber, C. (2018). Computational models of grid cell firing. *Hippocampal Microcircuits: A Computational Modeler’s Resource Book*, pages 585–613.
- Cannon, S. C., Robinson, D. A., and Shamma, S. (1983). A proposed neural network for the integrator of the oculomotor system. *Biological cybernetics*, 49(2):127–136.
- Chen, G., Manson, D., Cacucci, F., and Wills, T. J. (2016). Absence of visual input results in the disruption of grid cell firing in the mouse. *Current Biology*, 26(17):2335–2342.
- Cheng, S. and Frank, L. M. (2011). The structure of networks that produce the transformation from grid cells to place cells. *Neuroscience*, 197:293–306.

- Climer, J. R., DiTullio, R., Newman, E. L., Hasselmo, M. E., and Eden, U. T. (2015). Examination of rhythmicity of extracellularly recorded neurons in the entorhinal cortex. *Hippocampus*, 25(4):460–473.
- 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. (2013). Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature neuroscience*, 16(3):318–324.
- D’Albis, T. (2018). *Models of spatial representation in the medial entorhinal cortex*. PhD thesis, Humboldt-Universität zu Berlin, Lebenswissenschaftliche Fakultät.
- D’Albis, T. and Kempter, R. (2017). A single-cell spiking model for the origin of grid-cell patterns. *PLoS Computational Biology*, 13(10):e1005782.
- Derdikman, D., Whitlock, J. R., Tsao, A., Fyhn, M., Hafting, T., Moser, M.-B., and Moser, E. I. (2009). Fragmentation of grid cell maps in a multicompartment environment. *Nature neuroscience*, 12(10):1325–1332.
- Diehl, G. W., Hon, O. J., Leutgeb, S., and Leutgeb, J. K. (2017). Grid and nongrid cells in medial entorhinal cortex represent spatial location and environmental features with complementary coding schemes. *Neuron*, 94(1):83–92.
- Domnisoru, C., Kinkhabwala, A. A., and Tank, D. W. (2013). Membrane potential dynamics of grid cells. *Nature*, 495(7440):199–204.
- Dunbar, S. R. (1997). The average distance between points in geometric figures. *The College Mathematics Journal*, 28(3):187–197.
- Dunn, B., Mørreaunet, M., and Roudi, Y. (2015). Correlations and functional connections in a population of grid cells. *PLoS computational biology*, 11(2):e1004052.
- Etienne, A. S. and Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus*, 14(2):180–192.
- Fiete, I. R., Burak, Y., and Brookings, T. (2008). What grid cells convey about rat location. *Journal of Neuroscience*, 28(27):6858–6871.
- Fuchs, E. C., Neitz, A., Pinna, R., Melzer, S., Caputi, A., and Monyer, H. (2016). Local and distant input controlling excitation in layer ii of the medial entorhinal cortex. *Neuron*, 89(1):194–208.
- Fuhs, M. C. and Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, 26(16):4266–4276.

- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., and Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194.
- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264.
- Gardner, R. J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N. A., Dunn, B. A., Moser, M.-B., and Moser, E. I. (2022). Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128.
- Gardner, R. J., Lu, L., Wernle, T., Moser, M.-B., and Moser, E. I. (2019). Correlation structure of grid cells is preserved during sleep. *Nature Neuroscience*, 22:25–37.
- Gil, M., Ancau, M., Schlesiger, M. I., Neitz, A., Allen, K., De Marco, R. J., and Monyer, H. (2018). Impaired path integration in mice with disrupted grid cell firing. *Nature neuroscience*, 21(1):81–91.
- Giocomo, L. M., Zilli, E. A., Fransén, E., and Hasselmo, M. E. (2007). Temporal frequency of subthreshold oscillations scales with entorhinal grid cell field spacing. *Science*, 315(5819):1719–1722.
- Gothard, K., Skaggs, W., and McNaughton, B. (1996). Dynamics of hippocampal ensemble coding during space coding tasks. *Journal of Neuroscience*, 16(24):8027–8040.
- Gu, Y., Lewallen, S., Kinkhabwala, A. A., Domnisoru, C., Yoon, K., Gauthier, J. L., Fiete, I. R., and Tank, D. W. (2018). A map-like micro-organization of grid cells in the medial entorhinal cortex. *Cell*, 175(3):736–750.
- Guanella, A., Kiper, D., and Verschure, P. (2007). A model of grid cells based on a twisted torus topology. *International journal of neural systems*, 17(04):231–240.
- Haas, B., Givon-Mayo, R., Nachmani, E., and Ulanovsky, N. (2019). Can hippocampal place cells support path integration? *Hippocampus*, 29(7):599–607.
- Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.-B., and Moser, E. I. (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature*, 453(7199):1248–1252.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806.
- Hardcastle, K., Ganguli, S., and Giocomo, L. M. (2015). Environmental boundaries as an error correction mechanism for grid cells. *Neuron*, 86(3):827–839.

- Hayman, R. M., Chakraborty, S., Anderson, M. I., and Jeffery, K. J. (2003). Context-specific acquisition of location discrimination by hippocampal place cells. *European Journal of Neuroscience*, 18(10):2825–2834.
- Høydal, Ø. A., Skytøen, E. R., Andersson, S. O., Moser, M.-B., and Moser, E. I. (2019). Object-vector coding in the medial entorhinal cortex. *Nature*, 568(7752):400–404.
- Ismakov, R., Barak, O., Jeffery, K., and Derdikman, D. (2017). Grid cells encode local positional information. *Current Biology*, 27(15):2337–2343.
- Jeewajee, A., Barry, C., Douchamps, V., Manson, D., Lever, C., and Burgess, N. (2014). Theta phase precession of grid and place cell firing in open environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635):20120532.
- Kang, L. and Balasubramanian, V. (2019). A geometric attractor mechanism for self-organization of entorhinal grid modules. *Elife*, 8:e46687.
- Kang, L. and DeWeese, M. R. (2019). Replay as wavefronts and theta sequences as bump oscillations in a grid cell attractor network. *Elife*, 8:e46351.
- Kitamura, T., Pignatelli, M., Suh, J., Kohara, K., Yoshiki, A., Abe, K., and Tonegawa, S. (2014). Island cells control temporal association memory. *Science*, 343(6173):896–901.
- Koenig, J., Linder, A. N., Leutgeb, J. K., and Leutgeb, S. (2011). The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science*, 332(6029):592–595.
- Kropff, E., Carmichael, J. E., Moser, M.-B., and Moser, E. I. (2015). Speed cells in the medial entorhinal cortex. *Nature*, 523(7561):419–424.
- Krupic, J., Bauza, M., Burton, S., Barry, C., and O’Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538):232–235.
- Krupic, J., Bauza, M., Burton, S., and O’Keefe, J. (2016). Framing the grid: effect of boundaries on grid cells and navigation. *The Journal of physiology*, 594(22):6489–6499.
- Latuske, P., Kornienko, O., Kohler, L., and Allen, K. (2018). Hippocampal remapping and its entorhinal origin. *Frontiers in behavioral neuroscience*, 11:253.
- Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., and Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, 309(5734):619–623.

Maaswinkel, H., Jarrard, L. E., and Whishaw, I. Q. (1999). Hippocampectomized rats are impaired in homing by path integration. *Hippocampus*, 9(5):553–561.

Mallory, C. S., Hardcastle, K., Bant, J. S., and Giocomo, L. M. (2018). Grid scale drives the scale and long-term stability of place maps. *Nature neuroscience*, 21(2):270–282.

Marozzi, E., Ginzberg, L. L., Alenda, A., and Jeffery, K. J. (2015). Purely translational realignment in grid cell firing patterns following nonmetric context change. *Cerebral Cortex*, 25(11):4619–4627.

Mathis, A., Herz, A. V., and Stemmler, M. (2012). Optimal population codes for space: grid cells outperform place cells. *Neural computation*, 24(9):2280–2317.

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., Kudrimoti, H., Qin, Y., Skaggs, W., Suster, M., et al. (1996). Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *Journal of Experimental Biology*, 199(1):173–185.

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

Mittelstaedt, H. and Mittelstaedt, M.-L. (1982). Homing by path integration. In *Avian Navigation: International Symposium on Avian Navigation (ISAN) held at Tirrenia (Pisa), September 11–14, 1981*, pages 290–297. Springer.

Monsalve-Mercado, M. M. and Leibold, C. (2017). Hippocampal spike-timing correlations lead to hexagonal grid fields. *Physical Review Letters*, 119(3):038101.

Mosheiff, N., Agmon, H., Moriel, A., and Burak, Y. (2017). An efficient coding theory for a dynamic trajectory predicts non-uniform allocation of entorhinal grid cells to modules. *PLoS computational biology*, 13(6):e1005597.

Mosheiff, N. and Burak, Y. (2019). Velocity coupling of grid cell modules enables stable embedding of a low dimensional variable in a high dimensional neural attractor. *Elife*, 8:e48494.

Nagele, J., Herz, A. V., and Stemmler, M. B. (2020). Untethered firing fields and intermittent silences: Why grid-cell discharge is so variable. *Hippocampus*, 30(4):367–383.

- Navratilova, Z., Giocomo, L. M., Fellous, J.-M., Hasselmo, M. E., and McNaughton, B. L. (2012). Phase precession and variable spatial scaling in a periodic attractor map model of medial entorhinal grid cells with realistic after-spike dynamics. *Hippocampus*, 22(4):772–789.
- Noorman, M., Hulse, B. K., Jayaraman, V., Romani, S., and Hermundstad, A. M. (2022). Accurate angular integration with only a handful of neurons. *bioRxiv*, pages 2022–05.
- O’Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*.
- Pastoll, H., Solanka, L., van Rossum, M. C., and Nolan, M. F. (2013). Feedback inhibition enables theta-nested gamma oscillations and grid firing fields. *Neuron*, 77(1):141–154.
- Peng, J.-J., Throm, B., Jazi, M. N., Yen, T.-Y., Monyer, H., and Allen, K. (2023). Grid cells perform path integration in multiple reference frames during self-motion-based navigation. *bioRxiv*, pages 2023–12.
- Pérez-Escobar, J. A., Kornienko, O., Latuske, P., Kohler, L., and Allen, K. (2016). Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. *Elife*, 5:e16937.
- Petrucchio, L., Lavian, H., Wu, Y.-K., et al. (2023). Neural dynamics and architecture of the heading direction circuit in zebrafish. *Nature Neuroscience*, 26(5):765–773.
- Raudies, F. and Hasselmo, M. E. (2012). Modeling boundary vector cell firing given optic flow as a cue. *PLoS computational biology*, 8(6):e1002553.
- Reifenstein, E., Stemmler, M., Herz, A. V., Kempter, R., and Schreiber, S. (2014). Movement dependence and layer specificity of entorhinal phase precession in two-dimensional environments. *PLoS ONE*, 9(6):e100638.
- Reifenstein, E. T., Kempter, R., Schreiber, S., Stemmler, M. B., and Herz, A. V. M. (2012). Grid cells in rat entorhinal cortex encode physical space with independent firing fields and phase precession at the single-trial level. *Proceedings of the National Academy of Sciences*, 109(16):6301–6306.
- Robinson, J. C., Ying, J., Hasselmo, M. E., and Brandon, M. P. (2024). Optogenetic silencing of medial septal gabaergic neurons disrupts grid cell spatial and temporal coding in the medial entorhinal cortex. *Cell Reports*, 43(8).

- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., and Moser, E. I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762.
- Savelli, F., Yoganarasimha, D., and Knierim, J. J. (2008). Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus*, 18(12):1270–1282.
- Schmidt-Hieber, C. and Häusser, M. (2013). Cellular mechanisms of spatial navigation in the medial entorhinal cortex. *Nature Neuroscience*, 16(3):325–331.
- Seelig, J. D. and Jayaraman, V. (2015). Neural dynamics for landmark orientation and angular path integration. *Nature*, 521(7551):186–191.
- Seung, H. S. (1998). Continuous attractors and oculomotor control. *Neural Networks*, 11(7-8):1253–1258.
- Skaggs, W., Knierim, J., Kudrimoti, H., and McNaughton, B. (1994). A model of the neural basis of the rat’s sense of direction. *Advances in neural information processing systems*, 7.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., and Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909):1865–1868.
- Solstad, T., Moser, E. I., and Einevoll, G. T. (2006). From grid cells to place cells: a mathematical model. *Hippocampus*, 16(12):1026–1031.
- Sreenivasan, S. and Fiete, I. (2011). Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature neuroscience*, 14(10):1330.
- Stemmler, M., Mathis, A., and Herz, A. V. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. *Science Advances*, 1(11):e1500816.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., and Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492(7427):72–78.
- Stensola, T. and Moser, E. I. (2016). Grid cells and spatial maps in entorhinal cortex and hippocampus. *Micro-, meso- and macro-dynamics of the brain*, pages 59–80.
- Stensola, T., Stensola, H., Moser, M.-B., and Moser, E. I. (2015). Shearing-induced asymmetry in entorhinal grid cells. *Nature*, 518(7538):207–212.

- Sun, C., Kitamura, T., Yamamoto, J., Martin, J., Pignatelli, M., Kitch, L. J., Schnitzer, M. J., and Tonegawa, S. (2015). Distinct speed dependence of entorhinal island and ocean cells, including respective grid cells. *Proceedings of the National Academy of Sciences*, 112(30):9466–9471.
- Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15(1):70–86.
- Taube, J. S., Muller, R. U., and Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. *Journal of Neuroscience*, 10(2):420–435.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, 55(4):189.
- Trettel, S. G., Trimper, J. B., Hwaun, E., Fiete, I. R., and Colgin, L. L. (2019). Grid cell co-activity patterns during sleep reflect spatial overlap of grid fields during active behaviors. *Nature Neuroscience*, 22:25–37.
- Turner-Evans, D., Wegener, S., Rouault, H., Franconville, R., Wolff, T., Seelig, J. D., and Jayaraman, V. (2017). Angular velocity integration in a fly heading circuit. *Elife*, 6:e23496.
- Van Cauter, T., Camon, J., Alvernhe, A., Elduayen, C., Sargolini, F., and Save, E. (2013). Distinct roles of medial and lateral entorhinal cortex in spatial cognition. *Cerebral Cortex*, 23(2):451–459.
- Van Cauter, T., Poucet, B., and Save, E. (2008). Unstable ca1 place cell representation in rats with entorhinal cortex lesions. *European journal of Neuroscience*, 27(8):1933–1946.
- Waaga, T., Agmon, H., Normand, V. A., Nagelhus, A., Gardner, R. J., Moser, M.-B., Moser, E. I., and Burak, Y. (2021). Grid-cell modules remain coordinated when neural activity is dissociated from external sensory cues. *bioRxiv*.
- Wei, X.-X., Prentice, J., and Balasubramanian, V. (2015a). An optimal grid cell code for space. *Neuron*, 84(6):1183–1190.
- Wei, X.-X., Prentice, J., and Balasubramanian, V. (2015b). A principle of economy predicts the functional architecture of grid cells. *Elife*, 4:e08362.
- Welday, A. C., Shlifer, I. G., Bloom, M. L., Zhang, K., and Blair, H. T. (2011). Cosine directional tuning of theta cell burst frequencies: evidence for spatial coding by oscillatory interference. *Journal of Neuroscience*, 31(45):16157–16176.

- Welinder, P. E., Burak, Y., and Fiete, I. R. (2008). Grid cells: the position code, neural network models of activity, and the problem of learning. *Hippocampus*, 18(12):1283–1300.
- Wernle, T., Waaga, T., Mørreaunet, M., Treves, A., Moser, M.-B., and Moser, E. I. (2018). Integration of grid maps in merged environments. *Nature neuroscience*, 21(1):92–101.
- Widloski, J. and Fiete, I. R. (2014). A model of grid cell development through spatial exploration and spike time-dependent plasticity. *Neuron*, 83(2):481–495.
- Wills, T. J., Cacucci, F., Burgess, N., and O’Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. *science*, 328(5985):1573–1576.
- Wills, T. J., Lever, C., Cacucci, F., Burgess, N., and O’Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876.
- Winter, S. S., Clark, B. J., and Taube, J. S. (2015). Disruption of the head direction cell network impairs the parahippocampal grid cell signal. *Science*, 347(6224):870–874.
- Witter, M. P., Doan, T. P., Jacobsen, B., Nilssen, E. S., and Ohara, S. (2017). Architecture of the entorhinal cortex a review of entorhinal anatomy in rodents with some comparative notes. *Frontiers in systems neuroscience*, 11:46.
- Wu, S., Hamaguchi, K., and Amari, S.-i. (2008). Dynamics and computation of continuous attractors. *Neural computation*, 20(4):994–1025.
- Yoon, K., Buice, M. A., Barry, C., Hayman, R., Burgess, N., and Fiete, I. R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature neuroscience*, 16(8):1077–1084.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *Journal of Neuroscience*, 16(6):2112–2126.
- Zhang, S.-J., Ye, J., Miao, C., Tsao, A., Cerniauskas, I., Ledergerber, D., Moser, M.-B., and Moser, E. I. (2013). Optogenetic dissection of entorhinal-hippocampal functional connectivity. *Science*, 340(6128):1232627.
- Zilli, E. A. and Hasselmo, M. E. (2010). Coupled noisy spiking neurons as velocity-controlled oscillators in a model of grid cell spatial firing. *Journal of Neuroscience*, 30(41):13850–13860.

Zilli, E. A., Yoshida, M., Tahvildari, B., Giocomo, L. M., and Hasselmo, M. E. (2009). Evaluation of the oscillatory interference model of grid cell firing through analysis and measured period variance of some biological oscillators. *PLoS Computational Biology*, 5(11):e1000573.

Thesis: Interactions between continuous attractors on multiple scales in a grid cell network model

Avleen Sahni, Martin Stemmler and Andreas V.M. Herz

A.V.M.H. conceptually designed the project; A.S., developed and implemented the model; A.S. performed research and simulated and analyzed data under supervision from A.V.M.H. and M.S.; A.S. wrote the thesis, A.V.M.H. and M.S. provided corrections.

During the preparation of this work A.S. used ChatGPT-3.5 and Grammarly in order to check grammar and spelling. ChatGPT-3.5 was also used for generating the code for some figures as indicated in the main text. After using this tool, A.S. reviewed and edited the content as needed and takes full responsibility for the content of the publication.