

2023

Towards biologically plausible mechanisms of predictive learning

<https://hdl.handle.net/2144/48507>

Downloaded from DSpace Repository, DSpace Institution's institutional repository

BOSTON UNIVERSITY
GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

**TOWARDS BIOLOGICALLY PLAUSIBLE MECHANISMS OF
PREDICTIVE LEARNING**

by

G. WILLIAM CHAPMAN IV

B.S., Boston University, 2012
M.A., University of Colorado Boulder, 2018

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2023

© 2023 by
G. WILLIAM CHAPMAN IV
All rights reserved

Approved by

First Reader

Michael E. Hasselmo, D.Phil.
William Fairfield Warren Distinguished Professor
Professor of Psychological and Brain Sciences

Second Reader

Chandramouli Chandrasekaran, PhD
Assistant Professor of Anatomy and Neurobiology
Assistant Professor of Psychological and Brain Sciences

Third Reader

Mark Kramer, PhD
Professor of Mathematics and Statistics

EPIGRAPH

"So it goes."

- Kurt Vonnegut

TOWARDS BIOLOGICALLY PLAUSIBLE MECHANISMS OF PREDICTIVE LEARNING

G. WILLIAM CHAPMAN IV

Boston University, Graduate School of Arts and Sciences, 2023

Major Professor: Michael E. Hasselmo, D.Phil., William Fairfield Warren Distinguished Professor; Professor of Psychological and Brain Sciences

ABSTRACT

Animals perform a myriad of behaviors such as object tracking and spatial navigation, primarily in the absence of explicit target signals. In the absence of targets, neural circuits must implement a different target function. One primary theory for self-supervised learning is predictive learning, in which a system predicts feedforward signals over time, and in which internal representations emerge to provide longer-term structural information. While such theories are inspired by neural properties, they often lack direct links to low-level neural mechanisms.

In the first study, a model of the formation of internal representations is presented. I introduce the canonical microcircuit of cortical structures, including general connectivity and unique physiological properties of neural subpopulations. I then introduce a learning rule based on the contrast of feedforward potentials in pyramidal neurons with their feedback-controlled burst rates. Utilizing these two signals the learning rule instantiates a feedback-gated temporal error minimization. Combined with a set of feedforward-only units and organized hierarchically, the model learns to track the dynamics of external stimuli with high accuracy, and successive regions are shown to code temporal derivatives of their feedforward inputs.

The second study presents an electrophysiological experiment which showed a novel functional cell type in the retrosplenial cortex of behaving Long-Evans rats. Through rigorous statistical analysis we show that these neurons contain an egocentric representation of boundary locations. Combined with their location in the cortical hierarchy, this suggests that the retrosplenial neurons provide a mechanism for translating self-centered sensory information to the map-like representations present in subcortical structures.

In the final study I integrate the basic modular architecture of the first study with the specific afferent stimuli and macroscale connectivity patterns involved in spatial navigation. I simulate an agent in a simple virtual environment and compare the learned representations to tuning curves from experiments such as study two. I find the expected development of neural responses corresponding to egocentric sensory representations (retrosplenial cortex), self-oriented allocentric coding (postrhinal cortex) and allocentric spatial representations (hippocampus).

Together, these modeling results show how self-gated and guided learning in pyramidal ensembles can form useful and stable internal representations depending on the task at hand.

CONTENTS

Epigraph	iv
Abstract	v
List of Tables	xi
List of Figures	xii
List of Symbols and Abbreviations	xiv
1 Introduction	1
1.1 Learning Rules	3
1.2 Objective Functions	7
1.3 Dynamics & Architecture	8
1.4 Behavioral Tasks	10
1.4.1 Sequence Prediction	11
1.4.2 Spatial Navigation	12
1.5 Tying it all Together & Outline	14
2 Predictive Learning by a Burst-Dependent Learning Rule	17
2.1 Introduction	18
2.2 Methods	23
2.2.1 Task	23
2.2.2 Simulation Tools and Model Training	25
2.2.3 Baseline Model	27

2.2.4	Intermediate Models	28
2.2.5	Predictive Module	33
2.2.6	Spiking Model	38
2.3	Results & Discussion	46
2.3.1	Overall Model Comparisons	46
2.3.2	Baseline Model	48
2.3.3	Intermediate Models	50
2.3.4	Predictive Module	52
2.3.5	Spiking Model	56
2.4	Conclusions	57
2.5	Additional Methods	59
2.5.1	Task Parameters	59
2.5.2	Predictive Module	59
2.5.3	Spiking Model	59
2.5.4	FORCE	62
3	Egocentric boundary vector tuning of the retrosplenial cortex	65
3.1	Introduction	65
3.2	Results	68
3.2.1	RSC neurons exhibit stable spatial activity during free exploration	68
3.2.2	Egocentric boundary vector responsivity of RSC	71
3.2.3	Properties of RSC egocentric boundary vector receptive fields	72
3.2.4	EBC responses are localized within dysgranular RSC but lack topographic organization	74

3.2.5	Egocentric boundary vector tuning in secondary motor cortex and posterior parietal cortex but not medial entorhinal cortex	75
3.2.6	EBC responsivity is not explained by self-motion correlates .	76
3.2.7	Generalized linear models demonstrate robust egocentric vector tuning of RSC EBCs	79
3.2.8	GLM confirms vectorial representation	83
3.2.9	EBCs respond to local, not distal, environmental features . . .	84
3.2.10	EBC responsivity is anchored to boundaries, not the center of the environment	87
3.2.11	EBC responsivity is stable in novel environments	88
3.2.12	Stability of EBC subpopulations requires physical boundaries	89
3.2.13	RSC EBCs are insensitive to environmental geometry, which yields a directional representation of environment shape . . .	91
3.2.14	A subpopulation of RSC EBCs are theta-modulated	94
3.3	Discussion	96
3.3.1	RSC spatial representations facilitate reference frame transformations	96
3.3.2	The RSC egocentric boundary vector code is context independent, which generates a directional code that reflects environment geometry	97
3.3.3	EBCs are primarily restricted to the dysgranular RSC	99
3.3.4	A network of vector-based egocentric spatial representation .	100
3.4	Materials and Methods	102
3.4.1	Subjects	102

3.4.2	Data Analysis	102
4	Allocentric Representation and Transformation Learning from Egocentric	
	Sensory Prediction	112
4.1	Abstract	112
4.2	Introduction	113
4.3	Methods	117
4.3.1	General Approach	117
4.3.2	Neural Dynamics	120
4.3.3	Learning Rule	122
4.3.4	Tuning Curves	124
4.4	Results	127
4.4.1	Learned Responses	127
4.4.2	Response to Novel Environment	131
4.5	Discussion	135
5	Discussion	138
5.1	Summary of Findings	138
5.2	Current Work	140
5.3	Future Directions	140
5.4	Conclusion & Broader Implications	141
	Bibliography	143
	Curriculum Vitae	160

LIST OF TABLES

2.1	Organization of connections in the laminar intermediate model . . .	32
2.2	Initialization of weights in the predictive module	36
2.3	Spiking Connections	41
2.4	All Model Performances	47
2.5	Task Parameters	59
2.6	Stellate and interneuron parameters	60
2.7	Pyramidal Neuron Parameters	60
2.8	Parameter sets used for the short-term dependent plasticity.	60
2.9	Noise Parameters	61
2.10	Force Parameters	64
4.1	Behavioral Variables	118
4.2	Tuning Curve Parameters	125

LIST OF FIGURES

2.1	Overall setup for training of models	25
2.2	Intermediate model architectures	28
2.3	Architecture of the predictive module	33
2.4	Spiking Model Architecture	38
2.5	RNN Performance	48
2.6	Intermediate Performance	50
2.7	Predictive Module Performance	52
2.8	Lorenz Performance	54
2.9	Spiking Decodability	56
2.10	FORCE Training	63
3.1	Egocentric boundary vector representations of RSC neurons.	69
3.2	RSC egocentric boundary vector representations cannot be explained purely by self-motion correlates.	77
3.3	Egocentric vector tuning is more robust than allocentric or self-motion correlates using a generalized linear modeling framework.	81
3.4	EBCs are anchored to local boundaries, respond in novel environ- ments, and lose sensitivity in arenas without explicit borders	86
3.5	RSC EBCs are insensitive to environmental geometry, which gener- ates a directional representation of environment shape	92
3.6	A subset of RSC EBCs is theta-modulated	95
4.1	Simulation Setup	119

4.2	Learned Representations	129
4.2	Representation Learning Timescale	131
4.3	Novel Environment Tuning	132
4.4	Behavioral Decoding	134

LIST OF SYMBOLS AND ABBREVIATIONS

ANN	Artificial Neural Network
BAP	Backpropagating Action Potential
BDP	Burst-Dependent Plasticity
BPTT	Backpropagation Through Time
EC .	Entorhinal Cortex
HPC	Hippocampus
L2/3	Layer 2 / 3 Cortex, Superficial
L4 .	Layer 4 Cortex, Granular
L5/6	Layer 5 & 6 Cortex, Infragranular
LTP	Long-term Potentiation
M2 .	Motor Cortex
POV	Parvalbumin Interneurons
PR .	Perirhinal Cortex
RSC	Retrosplenial Cortex
SOM	Somatostatin Interneurons
STD	Short-term Depression
STDP	Spike Timing Dependent Plasticity
STF	Short-term Facilitation
STP	Short-term Plasticity

CHAPTER 1

Introduction

Animals perform a wide variety of behaviors with currently unmatched level of fidelity. However, for the majority of daily tasks, there is rarely an explicit target signal such as a motor sequence to a food source, and even a reinforcing signal such as food rewards. Although such explicit targets do exist, particularly in the form of reinforcement, using such sparse signals to create large-scale distributed representations of sensory and motor information is a highly difficult and inefficient process. Instead, it is theorized that a large portion of learning occurs in a ‘self-supervised’ manner in which internal representations are generated without an explicit reward, and later bootstrapped by other systems in order to perform some task. For example, once an internal representation of one’s location in two dimensional space is present, as well as the location of a water source, route planning from the current location to the goal is a relatively simple process.

Learning in biological systems has been experimentally shown to include synaptic modification dependent on the local state of neurons at the postsynaptic receptors at the time of an action potential (Gerstner & Kistler, 2002). In the simplest approaches this implements a form of associative learning between the presynaptic and postsynaptic action potential timing, suggesting that the role of time is critical for forming representations. In contrast, modern deep learning (DL) approaches, loosely inspired by neural systems, have largely moved away from systems involving temporal dynamics or learning over time, in favor of computationally optimized approaches. These methods have had great success in situations such as image classification or semantic compression, but have not generalized to other domains where target signals cannot be provided in a temporally concise manner,

such as dynamic motor control, rule-like manipulation, or continual learning. Inspired by this high performance, the central topic of this thesis is to investigate what specific architectural and dynamical constraints exist in neural systems, and how these enable learning of useful internal representations in an unsupervised manner. The remainder of this introduction will provide an overview of the problem space of choices to model learning in natural neural systems, and then review the high-level background information required to choose a specific avenue of investigation. The second chapter will then investigate the more specific architectural choices and performance. Chapter three then provides an extensive additional background on a specific behavioral paradigm which I model in the fourth chapter.

Problem Space When attempting to model a neural system which minimizes some error, there are three aspects of the system that can be manipulated (Richards et al., 2019). The first of these is the overall objective function which, as described above, may be a specific target signal, or may be some alternative signal such as overall activity (Clopath & Gerstner, 2010) or matching activity over time (O'Reilly, 1997). The second major aspect is the learning rule, which determines how the system adapts internal parameters in order to better meet the objective function. While the learning rule may include an overall global optimization (Rumelhart & McClelland, 1986), I will introduce several alternatives below in which, by choosing a corresponding objective function, the learning rule need only include local optimization. The final aspect that can be modified is the architecture and dynamics of units in the network. These form a set of inductive biases about the nature of stimuli and objective, such as spatial structure (Lecun et al., 2015) or temporal continuity (Rumelhart & McClelland, 1986). Such biases may inform our choice of

learning rule or otherwise alter the overall parameter adjustment in response to the objective function.

1.1 LEARNING RULES

Backpropagation In its typical form, as used in machine learning, backpropagation (BP) is a method for credit assignment to individual weights after the errors have been determined by an external objective function (Rumelhart & McLelland, 1986). These networks consist of a series of units with a piece-wise continuous output in response to inputs. Activity at each level of the network is calculated in a ‘feedforward’ manner, taken as the dot product of the activation of the previous layer by a set of learnable weights and passing this through the activation function of that layer. By combining the weights between these units appropriately, a deep neural network can function as a universal function approximator (Ablavsky & Sclaroff, 2011). However, because these networks are highly nonlinear and high dimensional, there is no simple mathematical algorithm for optimizing these weights, compared to, e.g: linear regression where the inverse matrix operation can find the optimal weights in a single iteration (Abbott, 2008). Instead BP utilizes the chain rule of derivatives to update weights in the state-space which most immediately decreases the error signal. While this approach is computationally powerful, it requires a separate phase for feedforward activity versus weight updates (temporal segregation) and also requires an exact method for calculating and carrying gradient information backwards in a network.

Contrastive Hebbian Learning Some of the first attempts to solve backprop in biologically plausible systems was the use of Contrastive Hebbian Learning (CHL),

which is based on the concept of Hebbian plasticity, but with a specifically temporal error component (Xie & Seung, 2003). In this approach, there is still a separation of the primary feedforward (termed negative) and weight updating (positive) phases. During the negative phase all units in the network evolve according to a nonlinear dynamic such as the leaky integrators used in Chapter 2. The final activation levels from this negative phase are saved, and then the network is run again, this time with the output layer activities clamped to the desired target value. After the clamped activity settles, an element-wise difference is taken in the activity from the negative and positive phases, resulting in a local error signal, and the weights are updated according to the difference. Critically, there are reciprocal connections (from layer N to $N-1$), and activity is calculated in multiple steps such that the network is now a complex dynamical system, meaning that the activities of output layers during the clamped phase can influence the activity and weight updates of lower levels. Feedback weights may be held constant (Detorakis et al., 2018) or updated according to the same error-driven rule as feedforward weights. In the latter case, CHL is shown to approximate the BP algorithm, though in practice the network fails to converge if there are more than 4-5 layers between the inputs and output layer (Xie & Seung, 2003). However, the explicit separation of a free phase and learning phase is a severe limitation to biological feasibility, though some authors have argued that these phases may be separated by intrinsic dynamics such as phasic changes in neurotransmitters (Hasselmo & Bower, 1993; Hasselmo, 1999), theta/alpha oscillations (Ketz et al., 2015; Hasselmo et al., 2002), or intrinsic bursting neurons (O'Reilly et al., 2021). Additionally, as originally introduced, Contrastive Hebbian Learning works only for continuous-valued activation functions, not for discrete action potentials (though see below).

Spike Timing Dependent Plasticity STDP is an extremely broad class of learning rules, dating back to the original theory of Hebb (Hebb, 1949) and including both early extracellular studies of synaptic modification (Levy & Steward, 1979; McNaughton et al., 1986) and intracellular studies (Gustafsson & Wigstrom, 1986; Markram et al., 1997). Whereas the learning rules discussed above rely on the activation value of the presynaptic and postsynaptic units, typically maximizing or minimizing their correlations, spike-based rules operate on some internal state of the postsynaptic unit at the point that a presynaptic action potential arrives. Often there is a model of a postsynaptic ‘trace’ which reflects some lingering intracellular changes after an action potential, such as NMDA receptor potentiation (Compte et al., 2000). The end result is a system in which synapses are strengthened when a presynaptic action potential shortly follows a postsynaptic one, and decreases otherwise. There exist a number of modifications of this basic setup which are sensitive to postsynaptic voltage (Clopath et al., 2010), calcium levels (Inglebert et al., 2020), or specific triplet firing patterns (Pfister, 2006), and which explain findings in different experimental protocols. Many of these modified learning rules can be captured by abstract learning rules such as the BCM rule (Bienenstock et al., 1982) which encapsulates a pre-post associative term with a postsynaptic homeostatic term to increase weights when the recent postsynaptic activity has been above a long-term running average. These STDP rules have been successfully combined with the clamping approach of Contrastive Hebbian Learning to implement accurate credit assignment in spiking networks (O’Reilly & Munakata, 2000), but is still restricted to a static target, due to the presence of a clamp signal.

Burst-Dependent Plasticity In order to avoid the requirement of a separate phase for forward activity and weight updates, we consider a recent class of learning

rules which utilize two multiplexed activities to simultaneously carry the signals required for both phases. These models utilize either the intracellular calcium levels (Larkum, 2013) or burst rates (Payeur et al., 2021) in pyramidal neurons to carry feedback information. In this approach, pyramidal neurons are modeled as two compartments. The main cell body is responsible for the generation of action potentials, while the distal dendritic tree integrates feedback signals to create an intracellular voltage or calcium plateau. When the distal component receives strong input, it results in a higher internal voltage which then causes a bursting, rather than single action potential, when the soma receives threshold input. The authors of this approach suggest that in a network of neurons modeled as such it is possible to multiplex a feedforward signal, through firing rates, and a feedback signal, through burst rates. The general architecture is that a feedforward chain is composed of pyramidal neurons projecting onto the proximal dendrites of higher up layers, with a feedback chain wherein pyramidal neurons project to the distal dendrites of the 'lower' layer. By chaining these neurons, burst rate can propagate downwards through multiple layers. With the transmission of these two signals, it is then possible for a synapse to locally compute a learning signal. Algorithmically, this means that the synapse increases if the presynaptic neuron simultaneously receives a top-down and feedforward signal, but decrease in strength if there is an 'error' of presynaptic firing without distal input, resulting in spikes which are not part of a burst. While the authors of these studies utilize this learning rule in the familiar contrastive manner, other recent work has adapted the firing rate versus burst rate approach to continuous time (Halvagal & Zenke, 2022) and self-supervised approaches (Illing et al., 2021).

1.2 OBJECTIVE FUNCTIONS

Unsupervised Learning The majority of the work in the learning rule literature discussed above focuses on how a system might align internal synaptic weights such that given a particular stimulus it elicits a desired activity at a different layer. It should be noted however that this is not necessarily required for all of the learning algorithms discussed. For instance, backpropagation requires only that an error be calculated, and some work has investigated how an error might be generated at higher regions based on stimuli that occur far apart from each other in time (Hinton, 2002). Errors may also be calculated by minimizing the difference in high level representation of two different processing streams, similar to multi-modal association (Chen et al., 2020). Finally, a loss could be directly calculated on the ability of a network to reconstruct a stimulus from internal states once the external target has been removed (Dayan et al., 1995).

Predictive Learning The approaches above may require separate observe and reconstruct phases, or a separate offline replay of temporally disjointed stimuli in the case of contrastive errors. An alternative approach, known as predictive learning, is to calculate error as the degree of instantaneous change, forcing stimuli which occur close together in time to have similar internal representations. In predictive learning, explicit teaching signals are replaced with a continuous input stream, and ‘error’ is the degree to which neural activity just before an observation does not match the activity after that observation. This is predictive in that the task of the network is then to calculate the future input as accurately as possible, resulting in a continuous and detail-rich target signal throughout the lifetime. The earliest studies of these types of systems focused on passive observation of the

world, such as how hierarchical visual cortex might use prediction not through time, but through space of visual scenes (Rao & Ballard, 1999). In these models simple cells in early visual cortex respond to their typical line and Gabor type receptive fields, while higher regions send back a prediction about a signal at a larger spatial level, such as environmental context. By analyzing this signal over time, we can expect that V1 cells will continue to code for their typical receptive fields, but with some top-down input (V2) that readies them (eg: by distal dendritic input) to fire or not, depending on what the state of the higher levels are at that time. The result is that V1 is responsible for the prediction of low-level detail (lines/pixels) that is it's input, while V2 is some estimation about a more stable latent signal in the world, similar to a Restricted Boltzmann Machine. More recent and computationally intensive implementations of predictive coding have specifically reconstructed videos over multiple predictive timesteps, and shown that the pixel-level representation is accurate for multiple steps into the future, and that overall activity levels increase when an unexpected frame is presented, consistent with experimental findings (Lotter et al., 2020; Halvagal & Zenke, 2022).

1.3 DYNAMICS & ARCHITECTURE

Learning Through Time In the previous section I suggested that in the realm of unsupervised learning, predictive learning is a powerful objective which we might fit. Some forms of predictive learning rely primarily on prediction of lower level activity patterns (Rao & Ballard, 1999). However, prediction over *time* allows the prediction of lower level activity, while also incorporating dynamics of stimuli and behavior, providing a richer error signal and providing the ability to control systems over time, such as in reinforcement learning (Baras & Meir, 2007; Sutton

& Barto, 1998). As with supervised learning, generation of temporal patterns for practical purposes is typically done utilizing backpropagation, in a specific configuration known as backpropagation through time (BPTT) (Williams & Zipser, 1989; Werbos, 1990). In this setup, recurrent neural networks (RNN) contain some recurrent connection from hidden layers back to the hidden layers, such that activation one point in time affects future activity. At the end of a trial, an error *vector* is then generated which calculates the difference in the target signal and output signal at each point in time. This error is backpropagated across virtual layers which represent each of the hidden layers at each point in time, and weights are updated accordingly, forming a kernel for temporal updates. While this approach can be extremely powerful for short temporal durations, the number of virtual units increases with each time step, and the final error signal must be divided among all of the virtual layers, resulting in a diminishing magnitude that interferes with long-term dependency learning (Bengio et al., 1994). Several attempts have been made to resolve these issues, such as modifying the architecture of the network (see below), avoiding a temporal component and instead treating a temporal sequence as a single observation (Vaswani et al., 2017), or incorporating recurrent weights but only learning from single frame errors (Sussillo & Abbott, 2009).

Practical Architectural Changes In order to solve the vanishing gradient problem of simple RNNs, an approach has been utilized which gates the input and output of temporal information based on current internal values (Sutskever et al., 2014). By preventing the feedforward signal from always being incorporated into the internal state, these approaches decrease the rate of gradient decay. This is an example in which a change in the modular structure of a network is introduced in order to solve a practical issue surround the learning rule. Other examples include

the use of convolutional connections which decrease the number of learnable parameters and introduce a bias towards certain spatial (Lecun et al., 2015), temporal (Cui et al., 2016), or both (Shi et al., 2015), regularities in the input. Additional constraints may be made on the macroscale organization of modules, such as deeper networks for image processing and shallow networks for temporal tasks.

Neural Dynamics and Learning Unlike the simplified feedforward networks originally designed to use BP, true neural systems inherently have spatiotemporal dynamics. The study of these dynamics is a field unto itself (Dayan & Abbott, 2003; Izhikevich, 2006), but less attention has been paid to how these dynamics may be improving or at least constraining weight updates (Clopath et al., 2010; Wunderlich & Pehle, 2021), objective functions (Boerlin et al., 2013), or entirely new functionality such as temporal coordination (Kramer et al., 2008), attention-based gating (Bastos et al., 2020), or explicitly temporal codes as opposed to firing rate based approaches (Skaggs et al., 1996; Frady & Sommer, 2019). Chapter 2 shows several steps towards harnessing these properties for low-pass interpolation, predictive bias, and temporal coincidence detection.

1.4 BEHAVIORAL TASKS

In studies one and three I will implement a biophysically inspired model which incorporates the principles outlined above. In order to evaluate how well these models learn both to predict external stimuli and create consistent internal representations, they will be trained in the context of behavioral tasks. A brief justification for the choice of tasks is provided in this section.

1.4.1 Sequence Prediction

The task utilized in the second chapter is a framework in which the (1-3 dimensional) state-space location of a stimulus moves according to an underlying deterministic set of ordinary differential equations. From a first-glance this task is highly inspired by tasks such as target tracking (Yoo et al., 2021; Alexander et al., 2022) in which an animal must predict the future location of a target object in order to intercept it. In such settings, it is beneficial for the animal to predict the future location of the target, rather than just the current location, in order to implement predictive pursuit, even when the target is briefly outside of the visual scene (Yoo et al., 2020). At a more abstract level, this task is also intended to represent how any cortical region may predict a low dimensional representation of afferent signals, and thus perform a general form of prediction, which has been shown to be possible in certain forms of neural models (Sussillo & Abbott, 2009), and gives rise to characteristic recurrent weights (Chalk et al., 2016) or error signals (Lotter et al., 2020; Gilra & Gerstner, 2017; O'Reilly et al., 2021).

Predicting such a signal is largely trivial under certain conditions for a simple RNN or LSTM network trained with backpropagation through time. In early proof-of-concept stages for this project, I trained such networks on relatively simple tasks such as sum of sinusoids, and more complex tasks such as the Lorenz equations. Consistent with current approaches in machine learning literature, I found that as long as the number of timesteps over which the ANN must predict is relatively small (<50), the trajectory has constant parameters (eg: frequency of oscillations is constant, though initial phase may differ), a narrow and shallow network is able to solve the task. However, as the number of time steps increases those networks are unable to correctly assign error (Hochreiter, 1998), and any stochastic

error signal tends to increase these difficulties (Bengio et al., 1994). However, these limitations are precisely the conditions under which neocortical predictions must occur.

Finally, because we are relatively agnostic as to the physical nature of the afferent signal, this task serves as a general-purpose hierarchical prediction mechanism (Gregor et al., 2014). This serves as an important contrast to the next task described in this thesis, which assumes some underlying latent structure of the signal.

1.4.2 Spatial Navigation

While the previous task operates on predictions over a fast time scale and early levels of neocortex, spatial navigation operates on the order of tens of seconds and longer (Bittner et al., 2017; Dannenberg et al., 2019), and involve several subcortical as well as cortical regions (Alexander et al., 2020). Many of the regions involved in this process are implicated not only in spatial reasoning, but more generally in sensory processing, motor control, reinforcement learning, and other functions. The earliest studies to find striking map-like (allocentric) representations in the brain found ‘place cells’ which respond when an animal is in a given location in space (Morris et al., 1982; McNaughton et al., 1984). Perhaps because these representations are so striking and intuitive to understand, they have been intensely studied for many decades, leading to an entire sub-field which has found other internal representations of allocentric space, such as head direction cells (Taube, 2007), cells coding running speed (Hinman et al., 2016; Kropff et al., 2015), boundary vector cells (Lever et al., 2009; Solstad et al., 2008) and grid cells (Hafting et al., 2005). These same regions have been shown to be necessary for spatial navigation tasks, giving them a causal role in behavior (Arolfo & Brioni, 1991).

However, such allocentric representations are clearly never directly provided as explicit input to the neurons which exhibit them, which brings into question how they are formed. Many early studies have investigated how the reciprocal connections in the spatial navigation system may lead to one functional phenotype such as grid cells giving rise to another such as place cells (Moser et al., 2013; Burgess et al., 2007), or how the intrinsic neural properties of regions such as entorhinal cortex may support spatially periodic firing fields (McNaughton et al., 2006). While such properties and connectivity patterns may provide internal consistency to such geometric representations, they must still be grounded to the more direct sensory observations at some level. The more behaviorally grounded of the frameworks (Byrne et al., 2007) have modeled how specific allocentric phenotypes connect to each other. This has led to predominant theories in which the regions involved in allocentric representations are providing some compression of the sensory stimuli (Recanatesi et al., 2021), providing a structured way in which those sensory signals may relate to each other over time (Stachenfeld et al., 2017; Whittington et al., 2022). More recent biological inspired models (Bicanski & Burgess, 2018) also incorporate sensory-driven egocentric representations found in cortical regions in the motor and retrosplenial cortex (Alexander et al., 2020), posterior parietal cortex (Alexander et al., 2022), postrhinal cortex (LaChance et al., 2019), and dorsal striatum (Hinman et al., 2019).

Overall spatial navigation is one of the most studied aspects of neuroscience from both a theoretical point and also in the number of neural responses known to exist in various brain regions. While the final subcortical representations are removed from the sensory space, there is a clear pathway to those responses from the directly observable environment. These properties make spatial navigation

an ideal case to study how self-supervised learning gives rise to a known set of responses.

1.5 TYING IT ALL TOGETHER & OUTLINE

In the sections above, I have introduced different biological learning rules, organizational principles, and objective functions that have been proposed in the space of understanding how neural systems learn. Based on the mutual constraints of the choices in each of these aspects, along with constraints such as continuous targets, I choose to investigate one possible combination of learning rule, architecture, and objective function. Based primarily on the constraint of unsupervised learning, I choose the objective of predictive coding, such that the network minimizes the difference between internal states and feedforward induced activity. Since the target signal in this objective is the same level of abstraction as the input signal, I choose an architecture in which we measure performance at the lowest levels of the network rather than the highest. This has the added advantage that the unit-specific error, the rate of change in activity, is local to each unit and should not require a global error signal. Of the learning rules discussed above, the burst-dependent plasticity rule is the one most suited to updating feedforward weights based on continuously varying hierarchical feedback signals. The remainder of this thesis works towards tying these choices of focus into a coherent model, and showing how the model learns internal representations in different behavioral situations.

Chapter 2 consists of a series of experiments and findings which aims to create a ‘predictive module’, which can be assembled to create a model of self-supervised learning of input stimuli. The model is inspired by much of the background material addressed above, and incorporates several biological inspirations and con-

straints to achieve performance above what is standard in back-propagation through time (BPTT) approaches. We introduce both a continuous version and a biophysical spiking model. This work has been significantly revised through peer review and is currently under re-review. These experiments were performed solely by myself, and the writing was done by myself and my primary advisor.

Chapter 3 consists of a previously published study (Alexander et al., 2019). This study investigates the neural responses of retrosplenial cortex (RSC) in freely behaving rats, and shows the presence of egocentric boundary cells (EBCs). It is chosen as an example of several experimental papers (Raudies et al., 2015; Shay et al., 2015; Tsuno et al., 2015; Ferrante et al., 2016; Hinman et al., 2016; Monaghan et al., 2017; Hinman et al., 2019; Alexander et al., 2019, 2022) where I have worked on the development of statistical analyses for detecting novel functional phenotypes involved in spatial navigation. This particular paper was chosen as it demonstrates the role of RSC, specifically the presence of egocentric boundary cells, which form a critical component of the transformation from egocentric observations to allocentric representations, as demonstrated in the following chapter.

Chapter 4 utilizes the basic architecture presented in Chapter 2 and incorporates it into a larger architecture to investigate the mechanisms of spatial representation learning. The overall connectivity is inspired by previous work (Bicanski & Burgess, 2018) which was hardwired to give specific functional cell types. In contrast, the work presented here begins with a *tabula rasa* connectivity and incorporates a self-supervised predictive learning rule to show the emergence of egocentric and allocentric representations purely from egocentric sensory information. This is the first released version of this work, performed exclusively by my advisor and myself, and will soon be submitted for peer review.

Taken together, these studies show that predictive learning is a biologically feasible mechanism whose implementation gives rise to internal representations of useful latent factors of inputs and mirrors experimental findings.

CHAPTER 2

Predictive Learning by a Burst-Dependent Learning Rule

ABSTRACT

Humans and other animals are able to quickly generalize latent dynamics of spatiotemporal sequences, often from a minimal number of previous experiences. Additionally, internal representations of external stimuli must remain stable, even in the presence of sensory noise, in order to be useful for informing behavior. In contrast, typical machine learning approaches require many thousands of samples, and generalize poorly to unexperienced examples, or fail completely to predict at long timescales. Here, we propose a novel neural network module which incorporates hierarchy and recurrent feedback terms, constituting a simplified model of neocortical microcircuits. This microcircuit predicts spatiotemporal trajectories at the input layer using a temporal error minimization algorithm. We show that this module is able to predict with higher accuracy into the future compared to traditional models. Investigating this model we find that successive predictive models learn representations which are increasingly removed from the raw sensory space, namely as successive temporal derivatives of the positional information. Next, we introduce a spiking neural network model which implements the rate-model through the use of a recently proposed biological learning rule utilizing dual-compartment neurons. We show that this network performs well on the same tasks as the mean-field models, by developing intrinsic dynamics that follow the dynamics of the external stimulus, while coordinating transmission of higher-order dynamics. Taken as a whole, these findings suggest that hierarchical temporal abstraction of sequences, rather than feed-forward reconstruction, may be

responsible for the ability of neural systems to quickly adapt to novel situations.

2.1 INTRODUCTION

Neocortical circuits mediate a broad variety of cognitive functions, including the extraction of rules in different behavioral tasks (Bhandari & Badre, 2018; Zhu et al., 2018; Hasselmo & Stern, 2018; Wallis et al., 2001; Buschman et al., 2012). One aspect of the extraction of rules involves the tracking of dynamics of sensory stimuli (Yoo et al., 2020) and self-location as an agent navigates in an environment (McNaughton et al., 2006; Byrne et al., 2007; Hasselmo, 2005; Bicanski & Burgess, 2018). A number of different cortical regions are implicated in these types of functions, including parietal cortex (Byrne et al., 2007; Bicanski & Burgess, 2018), retrosplenial cortex (Alexander et al., 2020), entorhinal cortex (Brandon et al., 2013), and prefrontal cortex. Simultaneously, anatomical evidence suggests that there may be common features of cortical circuitry throughout different cortical regions (Douglas et al., 1989; Bastos et al., 2012; Mountcastle, 1997; Rockland, 2010). Given the distributed nature of tracking, as well as this anatomical consistency across cortical circuits, we hypothesize that particular aspects of cortical organization may be responsible for building accurate internal representations of these external stimuli. Here, we work towards building a model of cortical microcircuits which replicates this ability to predicatively code for trajectories of stimuli.

Previous Work There have been many models of prediction of time series, both from a machine learning perspective and a neurally inspired framework. While our goal is to create a biologically realistic model of prediction, we discuss machine learning (ML) approaches as well. These ML approaches serve as a baseline to

which we can compare our model's performance, but also as extremely abstracted and mathematically optimized models of neural systems (Rumelhart & McClelland, 1986). The most common form of ML sequence prediction is sequence-to-sequence modeling, which utilizes backpropagation of errors through time (BPTT) (Williams & Zipser, 1989). While this approach has high success in areas such as natural language processing (Sutskever et al., 2014), they are essentially recurrent autoencoders, and typically fail when external teaching signals are inconsistent or sparse (Bengio et al., 1994). Alternative approaches based on echo-state (known as FORCE training) (Sussillo & Abbott, 2009), or liquid-state (Boerlin et al., 2013), networks are able to mimic external dynamics autonomously after a brief training period. FORCE and related methods however rely on a specific connectivity in which a decoded state is optimized by an external teacher and fed back into the network (Nicola & Clopath, 2017; Denève et al., 2017), and can not learn when the external teaching signal appears stochastically (see supplemental materials). In contrast, animals must perform in an environment where external stimuli may appear and disappear at random intervals, and building incorrect internal models of position is not sufficient for driving behavior. Prompted by this discrepancy, we consider three core aspects in which neural systems are thought to be organized, and consider how they may be essential for sequence prediction.

Hierarchy Proposals of cortical function have used hierarchical representations of information across different regions. These include many examples of function. For example, in progressing from caudal to rostral regions in the visual system different cortical regions mediate a hierarchical transition of coding level (Gilbert & Li, 2013; DiCarlo et al., 2012) of individual points of an image with the extraction of edges (V1), to the coding of higher order features such as movement (MT), color

(V4) and ultimately the identity of an object such as a face (IT) (Desimone & Schein, 1987; Hasselmo et al., 1989; DiCarlo et al., 2012). Another example concerns position information in which the elements of self-movement can be extracted separately in terms of position (O’Keefe & Dostrovsky, 1971; O’Keefe & Burgess, 2005), or separately as velocity (Kropff et al., 2015; Hinman et al., 2016) or as acceleration (Kropff et al., 2021). This also applies to higher level cognitive control modeled on reinforcement learning in which sub-policies in posterior frontal cortex are controlled by higher level hierarchical control policies in more anterior regions (Koechlin & Summerfield, 2007; Badre & Frank, 2012; Badre & D’Esposito, 2009). In modeling, these types of hierarchical representations are used in hierarchical reinforcement learning (Sutton et al., 1999) and in hierarchical Bayesian coding (Kingma & Welling, 2019). Here we build on the idea of hierarchical representations of cortical regions to form a predictive representation of dynamics, where the higher cortical regions track successively derived features.

Supervised vs Predictive Coding Previous approaches tend to rely on supervised learning, in which a whole or portion of a sequence is provided at one ‘end’ of a network, and a prediction is formed at the other end of the network, which is then used to update synaptic weights. In contrast, a number of cortical models involve generative-predictive coding (Rao & Ballard, 1999). In this framework, a given cortical region receives a feedforward (eg: sensory) signal, and a feedback signal consisting of the expected feedforward signal. Each region then calculates the difference between the feedforward (true) and feedback (expected) activity, which is then passed to the next cortical region (Bastos et al., 2012). In the example of early visual cortex, these feedback signals may be a line segment, while feedforward signals then carry the spatial error, giving rise to the end-stopping

phenomenon and other experimental findings (Lotter et al., 2020). While these forms of predictive coding hypothesize that the feedforward activity represents prediction errors, other studies have hypothesized that there is no such explicit computation of error. These alternative forms of predictive coding pose a feedforward signal which is, in each cortical region, simply a transformed version of its own input, and feedback signals represent some latent feature which may be informative for improving these lower-level predictions (O’Reilly et al., 2021). A common aspect to both the feedforward error and feedforward prediction frameworks is that reconstruction or prediction of the external stimulus occurs at the lowest cortical regions. This is similar to Helmholtz Machines, in which external stimuli drive the formation of self organizing maps, and feedback weights create generative biases that can reconstruct stimuli from scratch or distorted signals (Dayan et al., 1995). More recent work has also begun to investigate how the lower-level reconstruction approach can improve or simplify temporal prediction in machine-learning contexts (Gregor et al., 2014; Sutskever et al., 2009). In addition to *hierarchical* prediction, explicit prediction over time has also been shown to create compressed representations of stimuli, whereas non-predictive autoencoders do not (Recanatesi et al., 2021). Inspired by the success of hierarchical predictive coding, in both biological plausibility and successful stimulus prediction, we look for a model which incorporates a predictive and generative framework. Specifically, we expect that some temporal features of stimuli, such as temporal derivatives, may be derived in a hierarchy, thus forming the basis for a temporally changing context which can improve lower-level predictions.

Learning Rules Most of the articles cited up until this point rely either on some form of backpropagation through time (BPTT), or an abstracted form of statistical

optimization (eg: (Bastos et al., 2012)). However backpropagation is biologically implausible, in its generally presented form. There have been several proposed learning rules which approximate backpropagation in a more biologically plausible form, such as Contrastive Hebbian Learning (O’Reilly, 1997), Feedback Alignment (Lillicrap et al., 2016), or burst propagation networks (Payeur et al., 2021). However, these studies tend to focus on supervised learning problems, in which a series of feedforward regions attempt to match a given input to a desired output label. Even if the architecture of the network is modified such that reconstruction is performed at the lowest level, BPTT is not guaranteed to converge on optimal weights (as we show in our section ‘Intermediate Models’ below). In contrast, a more biologically plausible and local learning rule may converge on more optimal weights, despite being worse for generating universal function approximations, if it is more suited towards directly minimizing temporal error. We therefore propose to investigate a learning rule which incorporates both feedforward and feedback weights to minimize errors over time, and then connect it to a spike-based learning rule from experimental and computational literature (Payeur et al., 2021).

Scope of Work In the current work we focus on training networks to predict deterministic dynamical systems for long periods after stimulus has stopped being presented. We begin by illustrating how training by backpropagation through time fails in the presence of unreliable external input. We then present three sequential modifications to the baseline architecture, each of which is based on the biological principles discussed above. We then present a final modification in the introduction of a local learning rule, inspired by theories of error-driven learning in neocortex. We show that this biologically inspired “predictive module” performs significantly better at predicting dynamics over long timescales. Investigating the

learned representations in the predictive module, we find that successive modules learn to encode successive temporal derivatives. Finally, in the third portion of the results section, we present a spiking adaptation of the predictive module. This spiking model utilizes a set of biologically plausible learning rules, and performs with similar accuracy and feature learning as the continuous approximation.

2.2 METHODS

2.2.1 Task

In order to test the degree to which our models can predict underlying dynamics of an external stimulus, we utilize the commonly used sum of sinusoids task (Fig 2.1A) which evolves according to equation 2.1 (Duggins & Eliasmith, 2022; Sussillo & Abbott, 2009), with two modifications. First, in comparison to earlier papers, we generate the task in a procedural manner, where the relative phase and frequency of the underlying task change on each presentation (see table 2.5 for parameter ranges). By altering these parameters on each trial, we test the ability of the networks to learn general dynamics, rather than forecasting from previously seen histories. Secondly, we utilize a paradigm known as teacher forcing in which, over the course of training, the interval at which the network observed the ground-truth stimulus decreases. This frequency of observation of ground-truth, termed ‘teaching ratio’, is decreased over training according to a hyperbolic schedule (Fig 2.1E). On frames where the network is able to observe the ground-truth, the stimulus is presented to the lowest level of the network (Fig 2.1B, top). On frames where the network is not able to view the stimulus (“rollout”), the lowest region is instead presented with a decoded input from the previous time-step’s activity (Fig 2.1B, Middle), or no input at all (Fig 2.1B Bottom). During the initial phases of learn-

ing, having a high teaching ratio ensures that the ground truth and input follow similar statistical distributions, and is necessary to ensure learning during early phases of training. As training progresses, predictions a few time steps out are likely to follow the same distribution as the ground truth, as the network predicts each next-frame with high accuracy. On any given trial, the network is guaranteed to receive the true stimulus input for the first half (150) of frames, in order to allow appropriate historical observations to propagate through the network. On each epoch, we test the network on an unseen set of trials, and measure performance using the teaching ratio according to the schedule described above (henceforth the ‘local ratio’), and also using a teaching ratio of zero (Fig 2.1C). This allows investigation of how the networks learn to respond to the task which they are being trained to (local teaching ratio), while simultaneously investigating how they respond to long periods without external feedback.

$$P(t) = a_1 \sin(f_1 t + \phi_1) + a_2 \sin(f_2 t + \phi_2) \quad (2.1)$$

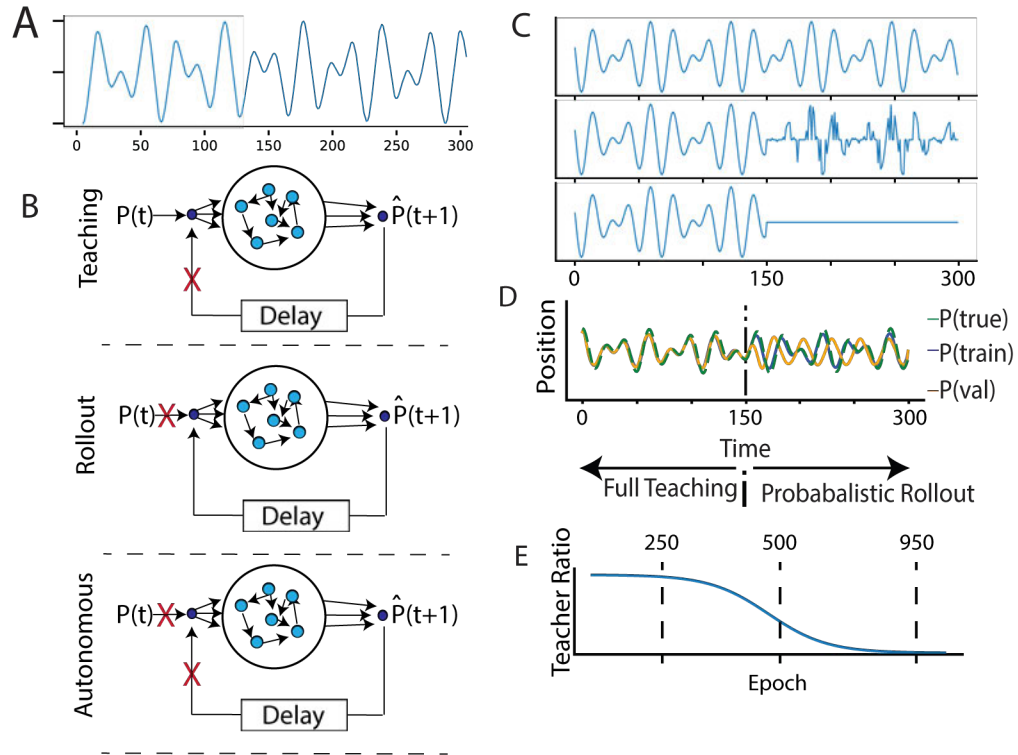


Figure 2.1: Caption on next page

2.2.2 Simulation Tools and Model Training

In the following sections we introduce a number of models which progress from a simple baseline model to a final novel architecture and learning rule. In each section equations which model the dynamics and weight updates are introduced, and parameters and initial distributions are set in tables. All models, with the exception of the spiking model, are implemented in PyTorch (Paszke et al., 2019), and evolve according to a simple forward Euler method, utilizing a timestep of 1. In all cases, error is calculated as the sum of mean-square errors across the entirety of a trial. All models before the predictive module are optimized by backpropagation through time (BPTT), using the ADAM optimizer (Kingma & Dhariwal,

A Example of the task used, showing the simulated trajectory as a function of time. Trajectories are parametrically generated according to equation 1. **B** Different approaches to time-series prediction. (Top) During a teaching frame, a network receives the ground truth position, and estimates the next frame. No information about the decoded prediction is returned to the network. (Middle) On a rollout frame, the network does not receive the ground truth, and instead receives its previous decoded prediction. (Bottom) In alternative frameworks, used later in the paper, a network may evolve autonomously, receiving neither the ground truth nor the decoded prediction. **C** Example of a target signal (top), and the corresponding *external* signal to the networks. (Middle) Throughout training, the network is presented with a training signal which is equal to the target signal for the first half of the trial, but randomly set to zero during the probabilistic rollout period ($p=0.50$ in example). On frames where the external signal is zero, the network evolves according to the corresponding rules from B. (Bottom) Throughout training an additional ‘validation’ signal is presented, in which the rollout teacher ratio is set to zero, but no learning occurs; this shows how the network would perform at the fully autonomous task at any point in training. **D** Example of a trial for a backpropagation through time (BPTT) network partly through training. During the full teaching period the network closely matches the target signal (green). During probabilistic rollout (blue), the network continues to closely approximate the target signal; the validation configuration (orange) diverges from the target. Note that while we use a signal frozen example in this and future traces, the actual signal given to the network is randomly generated on each trial. **E** The teacher ratio decreases over the course of training. Dashed lines indicate points where examples are drawn from in subsequent figures.

2018). In the predictive module section we replace BPTT with a novel learning rule which utilizes only local variables, and update weights in the same manner and frequency that activities are updated. Hyperparameter optimizations are found in table 2.4 and were chosen by an exhaustive grid search, representing 36-48 configurations for each model. For each hyperparameter configuration, five separate randomly models were run and their minimum local error or validation error were averaged together. The minimum error for each models best hyperparameter configuration are given in table 2.4 The spiking model was implemented in BindsNET (Hazan et al., 2018), a spiking neural network simulator built in Python, and utilize

a timestep of 0.1ms, and is described in more depth in the corresponding section below.

2.2.3 Baseline Model

As a comparison for our biologically-inspired network described later, we implement a standard sequence-to-sequence model. This baseline model consists of a linear encoder, a number of hidden Elman RNN layers, and a linear decoding layer. The activity of each layer, denoted 'R(t)' evolved according to the equation:

$$\begin{aligned}
 R_0(t) &= \tanh(W_{I0}I(t) + b_0) \\
 R_i(t) &= \tanh(W_{ii}R_i(t-1) + W_{(i-1)i}R_{i-1}(t) + b_i) \\
 R_N(t) &= W_{(i-1)i}R_{N-1}(t) + b_N
 \end{aligned}
 \tag{2.2}$$

Where subscripts denote either layer or weight (PrePost). On each time-step, the external input I is either the ground-truth stimulus (Fig 1B, top), or previous decoded output (Fig 1B, middle), according to the teaching ratio (Fig 1C and Fig 1E). We also investigated the performance of a similar architecture but instead utilizing LSTM units.

2.2.4 Intermediate Models




Figure 2.2: Intermediate models, showing stepwise changes towards biologically inspired architecture. **A** Stacked RNN: In the baseline model activity entered a stack of Simple RNNs at the base and propagates to a top layer where readout occurred. In the stacked RNN case, there are bidirectional connections between each hidden layer, and readout occurs either at the top layer (top) or input layer (bottom). **B** Leaky integrators (LI): Instantaneous responses in the simple RNNs are now replaced with leaky integrators which have a finite-time response to inputs. Here we show the impulse response of a single (blue) LI, two LIs in series (orange) and three LIs in series (green). Each successive layer introduces an additional delay (offset) and smoothing of the input. **C** We now introduce separate feedforward (FF) and feedback (FB) pathways. Each colored block represents a set of leaky integrators, now arranged in a manner similar to neocortical motifs (see text), allowing separate feedforward (black) and feedback (red) pathways. Feedback connections project to L2/3 and L5/6 of lower regions.

We next introduce a series of intermediate models, each of which incorporates one of the principles outline in the introduction. The purpose of these models is to introduce aspects of the final proposed model, while allowing investigation into how each of these architectural or dynamical principles influence the ability of a standard learning rule (BPTT) to learn in the presence of our task. For each of these changes, we provide a biological rationale, and a normative explanation for how the change may improve performance.

Recurrence & Readout Layer Compared to the feed-forward networks of the baseline approach, cortical regions tend to be bidirectionally connected. For this

reason our first intermediate model utilizes a series of stacked Elman RNN layers but adds an additional weight from each layer back to its preceding layer (Figure 2.2 A, top). Next, we address the issue of where signals are reconstructed. In a hierarchical circuit the further we move up a hierarchy, the less information about the external stimulus is directly available. A common approach in machine learning settings is to introduce a ‘U’ like structure or provide skip connections, which bypass intermediate transformations and provide direct routes for lower-level information to be integrated into higher-level (deeper) regions. Here, we take a slightly different approach, inspired by biological models of predictive coding. Instead of providing skip connections that bring low-level information to higher regions, we provide backwards-projections, such that the representation from a deeper region influences the activity of earlier layers. This achieves the same result that less processed and more processed information integrate in a given layer, but creates several patterns of activity that have been observed in neural data (Lotter et al., 2020) and satisfies the general framework of predictive coding in which deeper layers represent some information not present in earlier layers and project that backwards (Rao & Ballard, 1999). The result is a stacked hierarchy of RNN cells, but reconstruction now occurs at the lowest level of the model (Fig 2.2 A, bottom). Mathematically, both of these scenarios can be expressed as:

$$\begin{aligned}
 R_0(t) &= \tanh(W_{00}R_0(t-1) + W_{10}R_1(t-1) + W_{I0}I(t) + b_0) \\
 R_i(t) &= \tanh(W_{ii}R_i(t-1) + W_{(i-1)i}R_{i-1}(t-1) + W_{(i+1)i}R_{i+1}(t-1) + b_i) \\
 R_N(t) &= \tanh(W_{NN}R_N(t-1) + W_{(N-1)N}R_{N-1}(t-1) + b_N) \\
 R_{out}(t) &= W_{X-out}R_X(t) + b_{out}
 \end{aligned} \tag{2.3}$$

Where “X” denotes either the bottom-most or top-post layer, depending on the condition, and R_{out} consists only of a single unit. In order for each layer to update synchronously, inputs to all layers are collected before any layer’s activity is updated. This results in a 1-frame delay between each layer (2 frames from bottom layer to top layer, 4 frames for bottom layer activity to propagate all the way to the deepest layer and then back to the lowest layer).

Leaky Integrator Units From a biological perspective, the activity of a population of neurons can not change instantaneously, and is often modeled as a leaky-integrator (LI), in which the potential of a single unit decays to zero in the absence of any input, according to the equation:

$$\tau \frac{dv_L}{dt} = -v_L(t) + \sum_{N \in A} W_{NL} R_N(t) \quad (2.4)$$

$$R_L(t) = \tanh(v_L(t))$$

Where τ represents a slow-leak time constant (set to 10 frames), v represents the ‘membrane potential’, r represents the outgoing activity, and A is the set of all other layers which project to layer L , again following the same bidirectional connectivity pattern from above. Leaky integrators have also been shown to be useful in generation of complex trajectories by providing a temporal reservoir of past activity (see supplemental material). Functionally, LI units act as a low-pass filter, which can filter out transiently absent inputs to a certain degree, and may be useful in smoothing out small errors in rollout dynamics. An additional consequence of the low-pass function is that input to a given unit has a *delay* before causing maximal changes in a downstream unit (Fig 2.2 B). This delay means that in our stacked RNN setting, changes that occur in a given layer and propagate into hidden layers

and back have a $2\tau\Delta N$ delay before they can be fully propagated and integrated into the layer. This means that in order for a signal to be useful by the time it propagates back, there must be some explicit bias towards prediction at each synaptic step.

Separate Feedforward and Feedback Pathways For the final intermediate model, we increase the connective complexity within each region. In neocortex there is a distinct laminar organization in which there is a separate feedforward pathway (lower regions \rightarrow granular \rightarrow superficial \rightarrow higher regions) and feedback pathway (higher regions \rightarrow infragranular \rightarrow lower regions) (Haeusler & Maass, 2007). Compared to a simple stacked approach, this branching connectivity allows infragranular neurons to integrate feedforward and feedback signals before passing them back down the cortical hierarchy, and has been suggested to be important for predictive coding approaches (Lotter et al., 2017). When combined with leaky-integrator units, this integration of paths also occurs over time, since feedback pathways represent older representations than feedforward ones. In order to test whether this separation of pathways can be utilized for prediction of external stimuli we modify the previous intermediate stacked model such that each region contains these separate sub-modules. The overall connectivity pattern is in Table 2.1, or illustrated in figure 2.2 C, and each sub-module continues to follow the dynamics of equation 2.4.

	$W_{PrePost}$
Recur	W_{GG}
	W_{SS}
	W_{II}
Intra	W_{GS}
	W_{SI}
Inter	W_{SG}
	W_{IS}
	W_{II}

Table 2.1: Organization of connections in the laminar intermediate model. Weights are indicated as PrePos. G is “granular” (lamina 4), S is “superficial” (Lamina 2/3) and I is “infragranular” (Lamina 5/6). Interegional SG is the primary feedforward pathway (projects to higher regions), while IS and II project from higher regions to lower ones, in accordance with general neuroanatomical studies.

2.2.5 Predictive Module

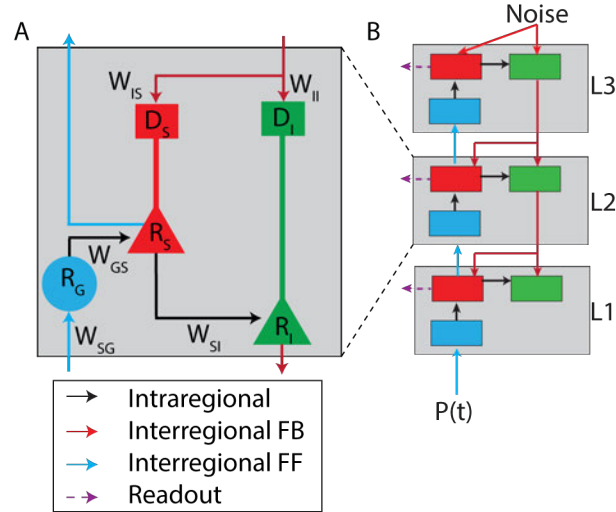


Figure 2.3: Architecture of the final predictive module. **A** The structure of each module, where rates ('R') of each submodule represent the firing rate of a given population of neurons, and weights are labeled according to their source and target submodules. Units are coloured by their laminar location, also denoted by subscript (G)ranular, (S)uperficial, (I)nfragranular, and (D)istal. Weights are coloured by according to their primary function in the feedforward (cyan), feedback (maroon) or local (black) pathways. **B** Zooming out to show connectivity among multiple predictive modules. For the lowest module feedforward activity comes from the external stimulus $P(t)$, while for higher regions it comes from lower level superficial neural activity (R_S of A). For the highest level, feedback activity comes from an OU process representing background fluctuations, while lower regions receive it from higher level infragranular activity (R_I of A).

The goal here is to introduce a learning rule which can effectively utilize the separate feedback pathway to self-supervise learning in lower cortical regions. In order to reach this we perform three key modifications

Dual-Compartment Units Results from the intermediate models (see results) indicate that in the BPTT approach having a branching feedback pathway decreases the accuracy of backpropagation based credit assignment. However, several lines

of research have suggested that the feedback component of cortical microcircuits is critical for guiding learning in biological systems (Magee & Grienberger, 2020; Larkum, 2013; Greedy et al., 2022). A common thread in each of these approaches is that the feedback pathway terminates on the distal dendrites of pyramidal neurons, and provides a mechanism through which to gate traditional Hebbian-like learning on the feedforward path synapses. Thus, in order to implement a similar rule, we modify the superficial and deep units of our model to implement a mean-rate version of the dual-compartment spiking pyramidal model. Each pyramidal neuron now contains an additional term in their dynamics:

$$\begin{aligned}
 D_L(t) &= \tanh\left(\sum W_{ID_L} R_I(t)\right) \\
 \tau \frac{dv_L}{dt} &= -v_L(t) + D_L(t) + \sum_{N \in A} W_{NL} R_N(t) \\
 R_L(t) &= \tanh(v_L(t))
 \end{aligned} \tag{2.5}$$

Here, the first term indicates the distal dendritic potential for a given unit, which follows its own activation function based on the sum of feedback activities R_I . The second term specifies that the ‘soma’ of the dual compartment model follows the same leaky-integrator dynamics of equation 2.4, but also receives a 1:1 input from its own distal component, representing passive intracellular coupling of these compartments.

Learning Rule Previous models of biological learning have suggested that the distal dendritic potentials (D_L) can guide learning by changing either burst rate (Payeur et al., 2021) or intracellular calcium levels (Bittner et al., 2017; Larkum, 2013). We will model the burst-based rule explicitly in the spiking model, but for the rate model implement this simply as a unit-by-unit gating of learning rate.

We then identify the overall goal of the network, which is to minimize temporal discrepancies in the sum of inputs. Given the temporal error, and a feedback signal for guiding learning, we propose a three-factor learning rule:

$$E_L(t) = (v_L(t) - v_L(t - 1)) \tag{2.6}$$

$$\Delta W_{PrePost} = \eta r_{pre}(t) \otimes D_L(t) \odot E_{post}(t)$$

Where \odot represents the Hadamard product, \otimes the outer product, and η is a learning-rate (set at 0.01). The outer product of the presynaptic and postsynaptic terms is identical to the associative term found in many learning rules (Kohonen, 1982), but now contains an additional ‘pushing’ term in the minimization of temporal change (Greedy et al., 2022), and further modulated by the degree of feedback prediction. Weights were initialized as shown in table 2.2, and learning occurred on every time step, regardless of whether the frame was forced or rollout, with the exception of validation trials.

Dynamic Vs Static Connections The learning rule proposed above is reliant on a feedback gating signal as well and has an explicit term to modify weights in order to minimize feedforward prediction errors (Brea et al., 2016), posing two issues. First, because the learning rule is locally greedy, there is a global minimum error of all synapses reach zero, resulting in a constant zero error-term. Secondly, the gating term is reliant on the distal compartment of pyramidal units, indicating that the rule can not apply to the granular units. However, experimental evidence has suggested that error-driven gating primarily occurs in the pyramidal neurons and other physiologically similar populations. We therefore implement the granular layer as a reservoir, and do not update the weights terminating onto these layers (see Table 2.2). When combined with random initial conditions (see below)

the reservoir dynamics allow the granular layers to avoid the global minimum of zero activity, and push other layers from the same region. Additionally, we leave the feedback weights, terminating on the distal compartments at their initial conditions, as the calcium or burst-dependent inspiration for the gating term is not valid for these synapses. Pragmatically, this means the predictive module implements a form of feedback alignment (Lillicrap et al., 2016). Overall, it is only the feedforward synapses onto pyramidal units that are expected to undergo the *error-driven* learning we investigate here.

	Connection _{pre-post}	Updates?	Initial Distribution
Recur	W_{XG-XG}	No	$Gauss(0, \frac{1}{2\sqrt{N}})$
	W_{XS-XS}	Yes	$Gauss(0, \frac{1}{2\sqrt{N}})$
	W_{XI-XI}	Yes	$Gauss(0, \frac{1}{2\sqrt{N}})$
Intra	W_{XG-XS}	Yes	$Unif(-\frac{1}{2\sqrt{N}}, \frac{1}{2\sqrt{N}})$
	W_{XS-XI}	Yes	$Unif(-\frac{1}{2\sqrt{N}}, \frac{1}{2\sqrt{N}})$
Inter	W_{XS-YG}	No	$Unif(-\frac{1}{2\sqrt{N}}, \frac{1}{2\sqrt{N}})$
	W_{YI-XS}	No	$Gauss(0, \frac{1}{2\sqrt{N}})$
	W_{YI-XI}	No	$Gauss(0, \frac{1}{2\sqrt{N}})$

Table 2.2: Initialization of weights in the predictive module. Weights are denoted as PreregionPreLamina-PostregionPostLamina, where X indicates a given region, and Y is X+1; G is “granular” (lamina 4), S is “superficial” (Lamina 2/3) and I is “infragranular” (Lamina 5/6). Feedback weights (going from Y to X) indicate connections targeting the postsynaptic learning-gate, rather than the strong somatic potential.

Noise There are two sources of noise in the model we implement here. First, on each trial the activity of all units (and distal compartments) is randomly initialized. Secondly, for the upper-most module there is no higher level to provide feedback activity to the distal dendrites, and we replace this instead with a small OrnsteinUhlenbeck (OU) process of mean zero, τ of 2 timesteps, and σ of 0.05.

Readout Because the predictive module does not directly optimize for readout of the target signal by a readout weight, we instead optimize readout weights (W) for each layer of the predictive module by constrained least-squares ('Ridge' in sklearn) such that the error term:

$$E = \sum (Y(t) - WR(t))^2 + \lambda \|W\|_2^2 \quad (2.7)$$

was minimal, where $Y(t)$ is the signal we are attempting to decode, and λ is a constant (0.01) which penalizes large readout weights. Weights were optimized by a 5-fold cross-validation from the full forcing period, and reused to create the decoded signal.

2.2.6 Spiking Model

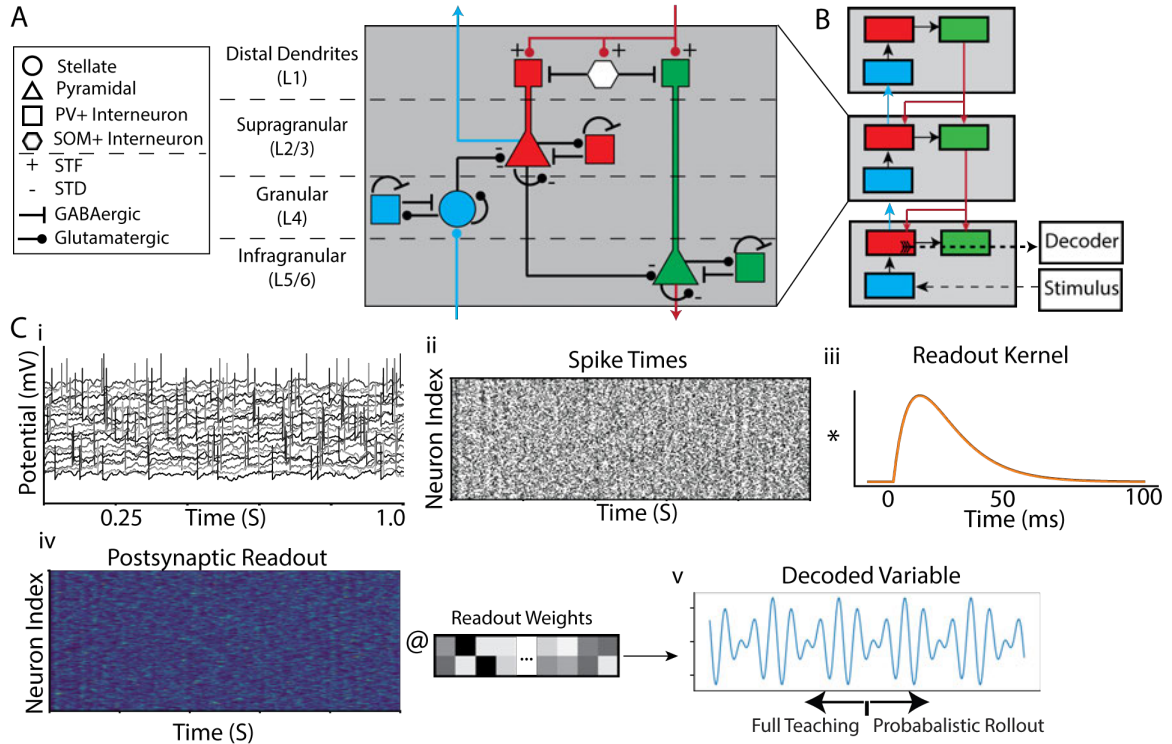


Figure 2.4: Spiking Model Architecture. Each box represents a cortical region, and coloured blocks represent lamina (matching those in Fig 2.3). L2/3, L4, and L5/6 each contain an excitatory population (Pyramidal and stellate, as triangles/circles), along with a loosely balanced inhibitory population (parvalbumin interneurons, squares). Layer L1 consists only of an inhibitory population of somatostatin neurons, and the distal dendrites of layer L2/3 and L5/6 pyramidal neurons. Feed-forward activity begins in layer L4, propagates to layer L2/3, and then diverges to intra-columnar L5/6, and inter-columnar L4. Feedback activity travels from L5/6 to inter-columnar L1, where it guides burst rate of upstream pyramidal neurons. **B** The full spiking model, where each grey block is as in A. External stimuli are encoded according to a Poisson process and randomly projected into L4 of the lowest cortical region. The activity of each region is then read out with an independently trained decoder. **C** Structure of the decoder in. For a given lamina we record the membrane potential from all excitatory neurons (i), from which we extract spike trains (ii). The spike train is convolved by an idealized postsynaptic potential(iii) to give an output signal (iv) which is read out by optimized weights (W) to give the reconstructed signal (v).

Next, we implemented a spiking neural network, trained to perform the same sum-of-sinusoids task set as above. As with the predictive module, the repeating anatomical connective motifs of this model are inspired by the canonical cortical microcircuit (Douglas et al., 1989). Now however, we introduce several additional components such as inhibition and short-term potentiation, that are necessary for translation to a spiking setup. Each introduced set of parameters follows general findings from neuroanatomical studies. All spiking simulations are performed in BindsNET and utilize a simple forward method (timestep of 0.1ms) to implement the dynamics described in the equations below. On each timestep all of the incoming activities to each neural population are collected before any updates are made. Each population of neurons then updates according to its local dynamics before weights are updated according to short-term or long-term updates.

Overall structure As shown in 2.4A. Each region of the spiking model consisted of 3 lamina: granular, supragranular, and infragranular, similar to the predictive module. Now however, each of these three lamina contained a population of excitatory neurons, modeled as stellate cells for granular, and pyramidal neurons for supragranular and infragranular, along with a loosely coupled population of parvalbumin (PV) interneurons. All excitatory populations consisted of 4000 units, and inhibitory populations consisted of 1000 units, reflecting proportions found *in vivo*. An incoming signal first passes through the granular layer, to supragranular, where it then diverges to the local infragranular neurons and the granular neurons of the next cortical region. Feedback activity travels from infragranular neurons to the distal dendrites of superficial and deep layers of the previous cortical region (Larkum, 2013). Connection patterns were initialized based on experimental findings (Haeusler & Maass, 2007), summarized in Table 2.3. Connection weights were

uniformly distributed at +/- 10% of the mean weight given in the table, and 1-p of those weights, were set to zero.

Short Term Plasticity Short term plasticity (STP) was modeled in all synapses according to the Markram Tsodyks model (Markram et al., 1998), modeling depletion of neurotransmitters and presynaptic calcium potentiation:

$$\begin{aligned}\frac{dR}{dt} &= \frac{1 - R(t)}{D} - u(t)R(t)\delta(t - t_{spike}) \\ \frac{du}{dt} &= \frac{U - u}{F} - f[1 - u(t)]\delta(t - t_{spike})\end{aligned}\tag{2.8}$$

STP parameters for each type of connection are summarized in table 2.8, and are chosen to match short-term facilitation (STF) or short-term depression (STD). The purpose of the STP dynamics are two-fold. First, they regulate the overall level of activity in the network, even as long-term weights are modified by the learning rule described below. Secondly, they allow a filtering of quick (STD) or slower (STF) spiking activities, as described in-depth in (Naud & Sprekeler, 2018). This second attribute is important for the learning rule described below.

Pre	Post	$\ W\ $ (pA)	p	Delay (ms)	STP	LTP
L4E	L4E	29	0.17	2	EE	None
L4E	L4I	60	0.19	1	EI	None
L4I	L4I	-41	0.5	1	II	None
L4I	L4E	-23	0.1	1	IE	None
L2/3E	L2/3E	45	0.26	2	EE	BDP
L2/3E	L2/3I	50	0.21	1	EI	None
L2/3I	L2/3I	-36	0.25	1	II	None
L2/3I	L2/3E	-17	0.16	1	IE	None
L5/6E	L5/6E	45	0.09	2	EE	BDP
L5/6E	L5/6I	24	0.1	1	EI	None
L5/6I	L5/6I	-32	0.6	1	II	None
L5/6I	L5/6E	-32	0.12	1	IE	None
L4E	L2/3E	60	0.28	2	EE	BDP
L2/3E	L5/6E	37	0.55	2	EE	BDP
L5/6E	L1D	15	.2	1	F	None
L5/6E	L1I	15	.2	1	F	None
L1I	L1D	-15	.2	1	D	None

Table 2.3: Connectivity between lamina and regions in the spiking model. Weights were initialized to the mean weight given in this table, with probability p, then allowed to evolve. STP column refers to the parameters given in table 2.8. For pre and post populations, E refers to excitatory population somatic compartments, I refers to inhibitory somatic compartments, and D refers to distal dendritic components (pyramidal neurons only).

Neural Dynamics As with other recent studies, pyramidal neurons were implemented with two compartments, representing the somatic and distal dendrites:

$$\begin{aligned}
C_s \frac{dV_s}{dt} &= \frac{C_s}{\tau_s} (V_s - E_L) + g_s f(V_d) + I_s - w_s \\
\frac{dw_s}{dt} &= \frac{-w_s}{\tau_s}
\end{aligned} \tag{2.9}$$

$$V \geq V_{Th} \rightarrow \begin{cases} V(t_+) &= V_{reset} \\ w(t_+) &= w(t_-) + b \end{cases} \tag{2.10}$$

Where $f(V_d)$ is represents the voltage/calcium gated channels in the dendritic compartment

$$f(V_d) = \frac{1}{1 + e^{-(V_d - E_d)/D_d}} \quad (2.11)$$

g_s is then the passive coupling parameter from the distal dendrites to the soma.

And for the dendritic compartment

$$\begin{aligned} C_d \frac{dV_d}{dt} &= -\frac{C_d}{\tau_d} (V_d - E_L) + g_d f(V_d) + c(S(t)) + I_d - w_d \\ \frac{dw_s}{dt} &= -\frac{w_d}{\tau_d} + a_w (V_d - E_L) \end{aligned} \quad (2.12)$$

Here, $S(t)$ represents the backpropagating action potential from the somatic compartment, which takes the form of a 2ms long pulse, delayed by 0.5ms from the time of the somatic spike($\delta(t)$):

$$S(t) = \int_{t-2.5ms}^{t-0.5} \delta(t) dt (\max 1) \quad (2.13)$$

This dual compartment approach allows separate control of firing rate and burst-rate, enabling multiplexing of feedforward and feedback activity. Similar to the D_L term from the predictive module, the distal dendritic compartment is receives feedback activity, and is responsible for controlling the learning rate of the somatic compartment, through rules described below (Payeur et al., 2021). In order to maintain the overall level of bursting in these neurons, an additional somato-statin (SOM) population of neurons was added which receives the same feedback activity as the distal dendrites, and provides a level of normalization keeping the burst rate approximately linear to the overall level of feedback activity (Vercruysse et al., 2021). The spiking properties of neurons in the network are illustrated in

part B and C of Fig 2.4. All other neurons were implemented with adaptive exponential integrate and fire models (see section 2.5.3.1), with parameters to match their physiology (Naud et al., 2008).

2.2.6.1 Inputs and Readout

Each compartment received a multicomponent input current, consisting of an external stimulus-dependent component, a synaptic component, and a stochastic background noisy input.

$$I_i = I_i^{ext} + I_i^{BG} + I_i^{syn} \quad (2.14)$$

The synaptic current I^{syn} is calculated as:

$$\frac{dI_{ij}^{syn}}{dt} = -I_{ij}^{syn} / \tau_{C_{ij}} + W_{ij} R \delta_j(t - d_{ij}) \quad (2.15)$$

Here, the spike train of presynaptic action potentials is delayed by a connection specific period d_{ij} before being modified by the short term plasticity term R and scaled by the synaptic weight matrix W . This modified impulse train is added at each point to a low-pass filter of the postsynaptic current, τ (5ms for inhibitory connections and 1ms for excitatory connections) allowing the postsynaptic synaptic potential (PSP) to follow the typical double-exponential pattern. The background input is necessary to keep units in the same lamina in a decorrelated state. This background was modeled separately for each compartment as an OU process, with parameters specific to each compartment type (see table 2.9). The background input was separately modeled as an OU process for each compartment:

$$\frac{d}{dt} I_i^{BG} = \frac{\mu - I_i^{BG}}{\tau_{OU}} + \sigma \sqrt{2/\tau_{OU}} \epsilon_i \quad (2.16)$$

Finally, the external input was zero for all units, with the exception of granular units in the lowest cortical region. For this region we encoded the external position as a Poisson random variable and processed this in the same manner as other synaptic inputs:

$$\begin{aligned} \delta(t) &\sim Pois(P(t)) \\ \frac{dI^{ext}}{dt} &= -I^{ext}/\tau_{C_{ij}} + W_{Input-4}\delta(t) \end{aligned} \quad (2.17)$$

Where $W_{Input-4}$ is initialized according to the distribution patterns for intraregional weights onto the granular lamina in table 2.3

2.2.6.2 Learning Rule

We next translate the predictive module learning rule to a spiking version, which reflects the recently proposed Burst-propagation rule (Payeur et al., 2021). As with the predictive module, we are only modeling the ‘prediction-driven’ aspects of learning, that is the portions of learning which are driven by top-down expectations which travel along the activity of infragranular activity to distal dendrites. The feedback pathway (L5/6-to-Distal dendrites) remained static, and this model therefore implements a version of feedback alignment (Lillicrap et al., 2016). With these structural changes in place, we now replace the predictive-error-driven learning rule from above with a direct implementation of Burst-Dependent Plasticity (BDP, Eq 2.18) (Payeur et al., 2021).

Feedforward excitatory synaptic weights onto pyramidal neurons were updated using the recently proposed burst-propagation rule, implemented at every time step except for validation trials: (Payeur et al., 2021)

$$\begin{aligned}
\frac{dw_{ij}}{dt} &= \eta \left\{ \underbrace{B_i}_{\text{Burst?}} - \underbrace{(\bar{P}_i)}_{\text{P(Burst)}} \underbrace{E_i}_{\text{event?}} \right\} \underbrace{\tilde{E}_j}_{\text{pre eligibility trace}} \\
\bar{E}_i(t) &= \frac{1}{\tau_{avg}} \int_0^{\infty} E_i(t - \tau) e^{-\tau/\tau_{avg}} d\tau \\
\bar{B}_i(t) &= \frac{1}{\tau_{avg}} \int_0^{\infty} B_i(t - \tau) e^{-\tau/\tau_{avg}} d\tau \\
\bar{P}_i(t) &= \frac{\bar{B}_i(t)}{\bar{E}_i(t)}
\end{aligned} \tag{2.18}$$

Where E_i is 1 when the neuron emits either a single spike, or the second spike within a 16ms window, and B_i is 1 only for the second spike during a 16ms window. Thus $\bar{E}_i(t)$ and $\bar{B}_i(t)$ are time-weighted averages of the ‘event’ and ‘burst’ rates. $\bar{P}_i(t)$ is therefore the probability that an event is the second spike of a burst.

Interpretation Intuitively, this learning rule is a modified pre-post product rule, where postsynaptic term (typically a Boolean indicating a spike) is replaced by $B_i - \hat{P}_i E_i$. This term means that weights will increase if the postsynaptic neuron fires the second spike of a burst, and decrease when the postsynaptic neuron fires a singlet or first spike of a burst. Because \hat{P} is primarily driven by the distal dendritic inputs (Naud & Sprekeler, 2018), this means that the feedback term is responsible for driving learning. We can directly compare this spike-based model to our gated-and-error-driven associative rule from the continuous model. Both equations contain a short-timescale representation of presynaptic activity, either as the low-passed activation ($R_{pre}(t)$ in the PM), or as a synaptic spike-trace ($E_j(t)$ in BDP). The feedback gating term from PM ($D(t)$) is analogous to the Burst Rate (B) in BDP. The most difficult analogy is how the error driven term from the predictive module ($D(t) * (v(t) - v(t-1))$) relates to the error-driven aspect of BDP ($B_i - P_i * E_i$).

However, one can see the mapping if we consider the BDP framework with a small population of neurons all of which have nearly-balanced distal dendritic inputs. This BDP error term can then be approximated as $F(D(t), v(t)) - D(t)_{lowpass} * v(t)$, where F incorporates the slight nonlinear effect of distal and basal potentials on burst probability. This revised BDP now resembles the error term from the predictive module, by incorporating the difference between an instantaneous and low-passed activity rate.

Readouts were performed by convolution of the recorded spike-train of interest with a synaptic kernel (Fig 2.4C). A linear readout for this convolved signal was then created by minimal least-squares linear weighting (Eq 2.21). We note that this readout term, and corresponding error terms, are never provided to the model, and are only used for interpreting the function of the network.

2.3 RESULTS & DISCUSSION

2.3.1 Overall Model Comparisons

We begin with a high level summary of means-square error performance across all models and identification of hyperparameters. Table 2.4 shows a summary of best hyperparameter sets for each model type, identified by minimum validation (full rollout/autonomous) error at any point in training. Trends for each model-types best parameter set are described in their corresponding sections below. Overall, the LSTM and baseline Elman RNN models show the lowest local error, supporting the general claim that these machine learning models, which utilize backpropagation through time, are highly accurate at predicting dynamics in the short term. However, when evaluating the full rollout condition, the leaky-integrator approach outperforms the LSTM, and this performance is further improved by

incorporating the novel learning rule of the predictive module.

Model Base	Hyperparameters	Train Error	Val Error
Baseline RNN	Units: 32, 64 , 128 Depth: 1 , 2, 3, 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.004	0.144
LSTM	Units: 32, 64 , 128 Depth: 1 , 2, 3, 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.002	0.119
Stacked RNN	Units: 32, 64 , 128 Depth: 2 , 3, 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.063	0.125
Leaky Integrator	Units: 32, 64 , 128 Depth: 2, 3 , 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.060	0.082
Separate FF & FB	Units: 32, 64 , 128 Depth: 2, 3 , 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.108	0.109
Predictive Module	Units: 32, 64 , 128 Depth: 2, 3 , 4 Learning Rate: 10^{-2} , 10^{-3} , 10^{-4}	0.008	0.027
Spiking Model	N/A	0.067	0.070

Table 2.4: Summary of performance across all tested models, including hyperparameters searched. Each row shows a model class, as separated out in the methods section. The bolded parameters in the second column indicate the best hyperparameter set, chosen by minimum validation error. The best models for local error (LSTM) and validation error (predictive module) are identified by bolding of their MSE.

2.3.2 Baseline Model

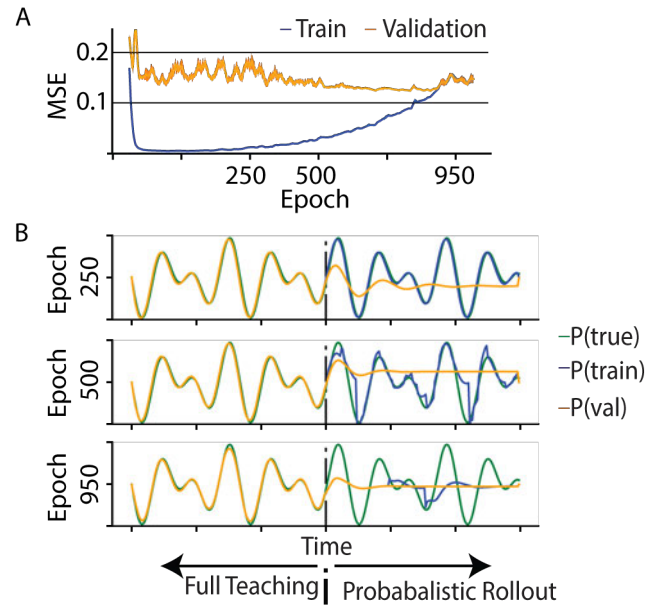


Figure 2.5: Performance of the baseline RNN. **A** Performance over training for both the local training ratio (blue) and validation ratio of zero (orange) as teaching ratio decreases. Initially, local loss approaches zero and the model is able to very accurately predict the next input, given the ground truth. As training progresses, both the partially driven and rollout system approach the same performance. **B** Showing example outputs for the partially driven (blue), and full rollout (orange) systems compared to the target signal (green) at three points in training. (Top) Early in training, the network predicts the next input accurately, given the ground truth, but decays to zero when run in full rollout mode. (Middle) Midway through training, the network can predict a few timesteps into the future accruing noise (blue) and snap back to the correct trajectory when a teaching frame occurs. At this point, the full rollout configuration will tend to maintain an output close to the last observed value. (Bottom) Late in training, the network is severely dampened, and will quickly create a constant-zero output, changing briefly when teaching frames occur (blue shifts away from orange for periods when teaching appears).

This network initially learns to predict the next input with a high degree of accuracy, but fails to predict in an undriven state (Fig 2.5 A). As training progresses, the network continues to learn short-term predictions (Fig 2.5 B, middle). During this stage of training the BPTT network has learned some of the longer term dynamics,

and will poorly fit 1-2 cycles before decaying to a constant output (orange trace). At this point small sequences of frames without the ground truth tend to cause network activity to diverge from the target, but quickly return when the external stimulus is presented again. In the final stages of learning the BPTT network performs very similarly to the full-rollout case from partway through training. From these examples we conclude that while the baseline model is able to accurately act as a predictive auto-encoder, the network is never able to predict beyond a handful of frames faithfully.

2.3.3 Intermediate Models

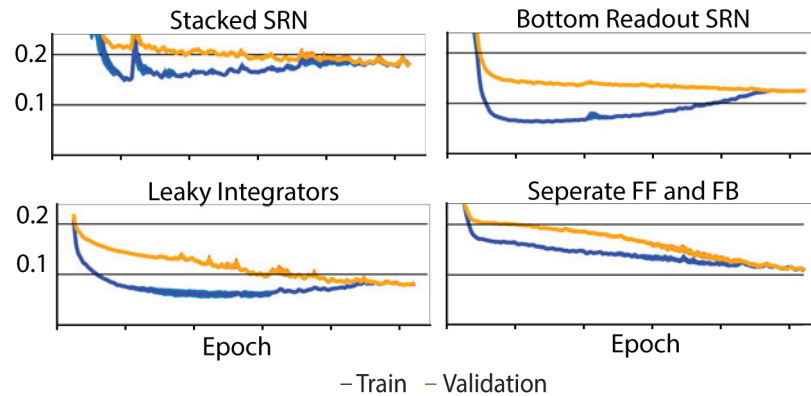


Figure 2.6: Intermediate Models, showing stepwise changes in overall performance. **Top Left** Beginning with a stacked bidirectionally connected RNN, but reconstructing from the ‘top’ layer, BPTT is unable to converge even for largely externally driven (early blue) cases. **Top Right** Continuing with a stacked RNN, but now reconstructing from the activity of the lowest layer, the network is more closely able to match the performance of the baseline model. This suggests that having sensory-driven feedforward inputs is critical to utilizing the hidden representations derived by other regions. **Bottom Left** By replacing the RNN dynamics with leaky integrators both the probabilistic and full rollout conditions improve. The training performance (blue) still does not perform as well as baseline models, but the full rollout condition (orange) now more closely matches the external dynamics. **Bottom Right** The introduction of separate feedforward and feedback pathways, while anatomically realistic, decreases performance for both training and validation forcing.

Investigating the learning curves of the intermediate models gives us a high level understanding of what architectural changes the backpropogation approach is able to make use of. In the stacked RNN (Fig 2.6) performance on both the teacher and full rollout conditions is significantly lower than the performance of the baseline RNN. This suggests that the reciprocal connections interfere with BPTT over long periods, consistent with the observation that the baseline model performed best with a lower number of layers. When utilizing the same bidirectional connectiv-

ity and reading out from the lowest layers ('Bottom Readout RNN') however, the short-term prediction increases significantly, though still higher than the baseline. The fully autonomous mode though now has a minimum error that is lower than the baseline, suggesting the the bottom readout approach allows useful information from the hidden layers to integrate into the primarily externally driven signal without disrupting the overall dynamics.

Continuing to utilize a bottom-readout approach and replacing the RNN dynamics with a leaky integrator further decrease the validation-mode error (MSE 0.082), far below the minimum achieved in the baseline models (MSE 0.119). However, the speed with which teacher learning occurs is significantly lower, as indicated by the slower decent of the blue line in Figure 2.6. When introducing separate pathways for feedforward and feedback information, performance was significantly worse for both the local forcing (MSE 0.108) and full rollout (MSE 0.109) cases, where the system primarily learned only to replicate the first frequency component of the target signal. Similar to the initial findings of the top-readout stacked RNN, this suggests that the BP based method is not able to utilize the additional pathway.

2.3.4 Predictive Module

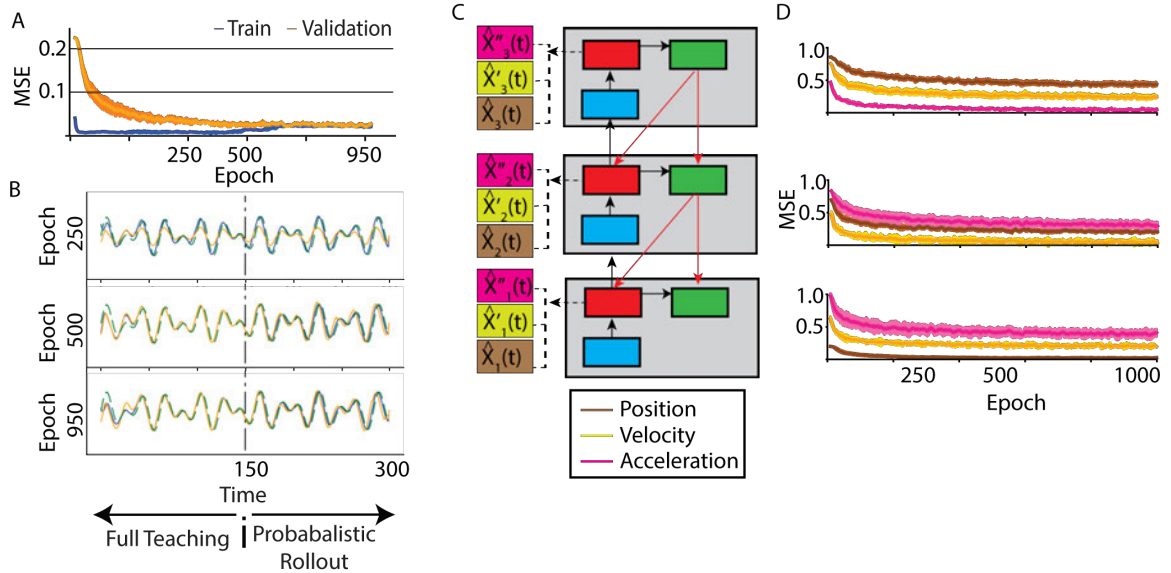


Figure 2.7: Performance of the predictive module approach. **A** Overall performance over training. The training-forcing condition (blue) quickly approaches performance similar to the backprop-based methods, showing that the lowest region quickly encodes the next-frame location. Compared to the baseline and intermediate models however, the validation (fully autonomous) mode also increases in accuracy quickly, showing that the network is learning to follow the external dynamics, rather than a weighted history of recent activity. **B** Example traces at early, mid, and late stages of training. In the early condition (top) the conditions match closely, though the autonomous mode has a tendency to undershoot local peaks. By mid training (middle) the autonomous mode now closely follows the ground truth. In late training (bottom) the predictions continue to match the target. **C** Decoding additional signals from each region of the model. (Left) Brown, yellow, and pink boxes indicate separate decoders trained for position, acceleration, and velocity for each of the regions, for a total of nine separately trained decoders. **D** Right Shows the accuracy for each of these decoders over the course of training. The lowest region most strongly codes for position (brown), with decreasing accuracy for decoding higher derivatives. The intermediate region can most strongly decode velocity (yellow), while the highest has the strongest tuning for acceleration (pink).

Under the predictive module framework, the network is able to quickly learn to reproduce the next observed frame (Fig 2.7 A, blue), but initially does not pre-

dict well without external drive. However, within the first 200 epochs, the network begins to faithfully predict trajectories in both the partially-driven and undriven states, with a slight undershooting of local peaks in the undriven state. Over the course of training the partially-driven error continues to stay low, despite a decreasing rate of external stimulus, while autonomous error approaches similar performance (down to 0.027 minimum autonomous error, down from 0.082 in the leaky-integrator intermediate model). In addition to investigating how well the predictive module network learned to predict long temporal delays, we were interested in what role each of those regions played in the overall learned representations. We attempted to decode not only the position, but also the temporal derivatives of this variable, utilizing three separate decoders for each module. Each decoder was optimized by cross-validated least-squares minimization between the ground-truth signal and the supragranular activity on each trial (Eq 2.21). We found that each module most strongly represented increasing temporal derivatives of position (i.e. distance, velocity and acceleration). This suggests that the later modules are encoding temporal derivatives of lower-level variables, which in turn provide contextual information to lower regions, further increasing predictive power. The multiple transformations between each module's output and the feedback signal allows for learned coefficients and non-linearities in how that feedback signal affects the current prediction.




Figure 2.8: Predictive Module activity on a more complicated task. Performance of the predictive module approach. **A** Showing decoded position along the first dimension at various points in training. At all points the local forcing ratio closely (blue) closely matches the ground truth (orange). Early in training the autonomous dynamics quickly diverge from ground truth, are lower magnitude. In mid training (middle) the autonomous teacher-forcing (blue) will tend to diverge from ground-truth, but comes back when a sample is presented. The autonomous trajectory appears to match the general switching behavior of the Lorenz equations. Later in training (bottom) the autonomous mode stays close to ground truth for long periods before diverging due to chaotic behavior. However, the general behavior still matches the dual-mode oscillations of the ground truth. **B** Showing the first two dimensions at various points in training, only for the second half of the trial. During early training (top) the autonomous trajectory is off center and lower magnitude than ground truth and semi-forcing case. During mid training (middle) the fully autonomous trajectory fills regions that the ground truth does not. Later in training (bottom) the autonomous trajectory more fully matches the 2D regions that the ground truth visits.

In order to test how this model generalizes to more complex tasks we next tested the model on the Lorenz equations, a set of deterministic differential equations known for their sensitivity to initial conditions. As with the sum of sinusoids task, decoding is performed by training on the forced (first half) of a trial, and using the learned weights to decode for the second half of the trial. We find, that by the end of training the network is able to accurately recreate the general attractor dynamics of the task (Fig 2.8A, bottom). Consistent with the nature of the task, the exact trajectories diverge when the internal representation is near the origin.

At the midpoint in training (Fig 2.8A, middle) the autonomous trajectory similarly matches the general dynamics of the task, but tends to visit locations that the underlying dynamics do not (Fig 2.8B, middle). However when partial teacher forcing is present, the external stimulus is sufficient to keep the decoded position near the ground truth. For the fully trained model (epoch 950) we attempted to decode temporal derivatives of the *internal* position generated by the position decoding. Consistent with the above results, we found that the magnitude of instantaneous position change was strongly represented in the second region but not the first or third. The third region did not strongly encode position or speed (MSE greater than 0.14 in both cases). a

2.3.5 Spiking Model

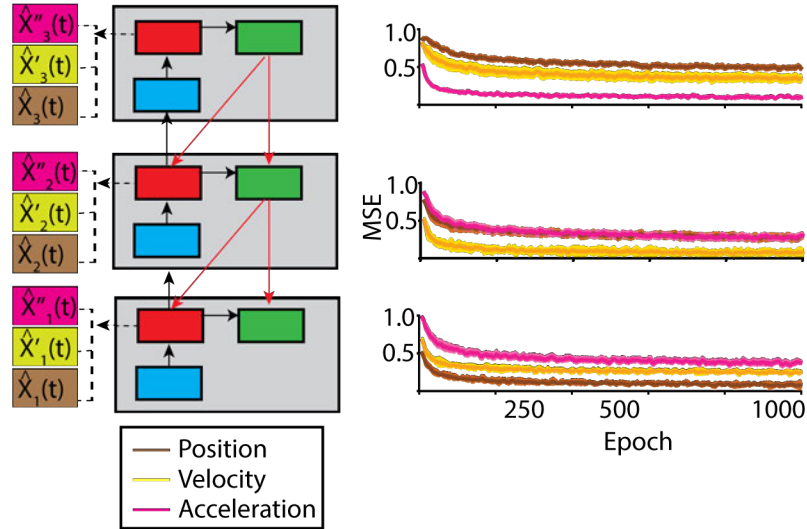


Figure 2.9: Summary of spiking model decodability. (Left) Showing the read-out architecture, identical to the setup for the predictive module. (Right) Showing decodability of position and temporal derivatives over training in superficial pyramidal neurons after convolving by a postsynaptic readout kernel (see methods). The optimized decoders show that the lowest region again encodes position, with velocity being most strongly encoded in the second region, and region 3 most strongly encoding acceleration terms.

As shown in Fig 2.9, the spiking network showed similar capabilities for function to the predictive module presented in Fig 2.7. The spiking network is able to accurately predict at long time-scales. The L2/3 neurons in region 1 (bottom) most accurately encodes position, the L2/3 neurons in region 2 (middle) most accurately encode velocity, and the L2/3 neurons in region 3 (top) encode acceleration (Fig 2.9). Thus, the region most directly connected to the external input mimics the dynamics of spatiotemporal patterns, while the more distal regions generate abstracted signals which can help guide lower-level activity, but do not directly mimic the inputs.

2.4 CONCLUSIONS

We examined how an architecture inspired by connective motifs from the neocortex might predict the spatiotemporal dynamics of external stimuli. Using a continuous firing rate approximation, we found that incorporating a leak-term with a hierarchical network in which reconstruction occurs at the input layers (Fig 2.6D) significantly improves long-term prediction compared to standard backpropagation approaches. By then incorporating a laminar structure, in which deeper regions gate learning in lower regions, we were able to further increase these predictions (Fig 2.7). Investigating the activity patterns in various layers of this network, we found that successive regions were predicting temporal derivatives of activity in their inputs, such that deeper layers represent further temporal derivatives of the external stimulus. This finding highlights general patterns of cortical hierarchy which suggest that low-level sensory cortices give rise to higher level cortex which provides less physically grounded representations.

Utilizing a spiking model based on previous work (Naud & Sprekeler, 2018; Payeur et al., 2021), we showed that a biologically plausible learning rule can likewise result in a network that develops intrinsic connectivity that enables prediction of external stimuli, without an extrinsic teaching signal. Instead, the dynamics within each cortical region appear to follow the temporal derivative of their inputs, and must align the dynamics with more abstracted activity from higher cortical regions.

Previous work has investigated how reservoir networks may learn to predict complex time series (Denève et al., 2017) with external negative feedback systems, or by optimized networks (Sussillo & Abbott, 2009) akin to the sequence-to-sequence approach (Boerlin et al., 2013). However, the present work focused on

how a potential interaction of multiple cortical regions forming a hierarchy might develop a representation of network dynamics. The arrangement of a hierarchical system not only resulted in an improved performance, but also resulted in a structured organization of temporal information along the hierarchy. Essential to this performance was feedback along the hierarchy, which allowed higher order information to be integrated in superficial neurons of lower cortical regions (Haeusler & Maass, 2007). The different roles of somatostatin and parvalbumin allowed activity to continue stably when external input is absent, as proposed in recent work (Hertäg & Sprekeler, 2020).

Future Directions There were several purposeful simplifications in this study, that can be expanded upon in future studies. Because the focus here was on translating biologically inspired learning rules to a self-supervised temporal prediction setting, we utilized a low-dimensional and deterministic stimulus. Without additional architectural changes, such as convolution or topologically organized connections which are essential for receptive field generation (Gilbert & Li, 2013), the degree to which this network can be scaled up to behaviorally relevant tasks such as object recognition are limited. We also implemented long-term plasticity only on a minority of connections. This is done in order to stay within the realm of self-supervised signals, but neglects the potentiation of inhibitory synapses (Vogels et al., 2011) and granular-terminating connections which may be important for the formation of purely associative subnetworks. Despite these simplifications, the current model shows how a neural system can utilize a hierarchical organization to closely track the dynamics of external stimuli. Future studies may investigate uses of this model in more complex scenarios, which may lead to predictions about the functional phenotypes of cells that arise throughout the cortex.

2.5 ADDITIONAL METHODS

2.5.1 Task Parameters

Param	Sin1 Lower	Sin1 Upper	Sin2 Lower	Sin2 Upper
a	1	1	0.5	2.0
f	0.15	0.30	1.5 * f1	2.0 * f1
p	-pi	pi	-pi + p1	pi + p1

Table 2.5: Task Parameters

2.5.2 Predictive Module

2.5.3 Spiking Model

2.5.3.1 AdEx Dynamics

$$\begin{aligned}
 C \frac{dV}{dt} &= -g_L(V - E_L) - g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) + I - w \\
 \tau_w \frac{dw}{dt} &= a(V - E_L) - w
 \end{aligned} \tag{2.19}$$

$$V \geq V_{Th} \rightarrow \begin{cases} V(t_+) = V_{reset} \\ w(t_+) = w(t_-) + b \end{cases} \tag{2.20}$$

Parameter	Somatostatin Neurons	Parvalbumin Neurons	L4 Stellate
C (pF)	100	100	200
g_L (nS)	5	5	12
E_l (mV)	-70	-70	-70
V_T (mV)	-62	-48	-50
Δ_T (mV)	4	2	2
a (nS)	-.5	0	2
τ_w (ms)	100	1	300
b	10	0	60
V_r (mv)	-65	-55	-58
τ_s (ms)	10	10	5

Table 2.6: Stellate and interneuron parameters

Parameter	Soma	Dendrite	Description
C	370pF	170pF	Membrane Capacitance
τ	7ms	16ms	Membrane time constant
E_L	-70mV	-70mV	Reversal potential
g	1200pA	1300pA	Passive coupling from gated dendritic channels
τ_w	30ms	100ms	Recovery variable time constant
c	N/A	2600pA	Strength of backpropogating action potential
a_w	N/A	-13pA	Strength of passive leak current in dendrites

Table 2.7: Pyramidal Neuron parameters. Soma column corresponds to Eq 2.9, and dendrite column to values in Eq 2.12

Type	D (s)	F (s)	U	f
EE	1.1	.05	.5	.1
EI	.125	1.2	.05	.1
II	.144	.06	.32	.1
IE	.7	.02	.25	.1
F	.1	.1	.02	.1
D	.02	1	0.9	.1

Table 2.8: Parameter sets used for the short-term dependent plasticity.

Compartment	$\mu(\text{pA})$	$\tau_{ou}(\text{ms})$	$\sigma(\text{pA})$
Somatostatin	0	2	450
Parvalbumin	0	2	450
Stellate	0	2	450
Pyramidal Soma	70	2	450
Pyramidal Dend	-270	2	450
Region 3 Dend	-23	2	450

Table 2.9: Parameters for the OU noisy input. The final region distal dendrites received different input in order to compensate for the lack of feedback activity and maintain a burst rate similar to other regions.

2.5.3.2 Synaptic Readout

Spiking models were decoded by convolving a given spike train with an idealized dual-exponential post-synaptic kernel:

$$I_{readout}(t) = \frac{(t - \delta)}{\tau_r} * e^{-(t-\delta)/\tau_r} \quad (2.21)$$

Where τ_r is the synaptic readout time constant of 10ms. On each trial, the readout weights (W) were optimized by constrained least-squares ('Ridge' in sklearn) such that the error term:

$$E = \sum (Y(t) - WI_{readout}(t))^2 + \lambda \|W\|_2^2 \quad (2.22)$$

was minimal, where $Y(t)$ is the signal we are attempting to decode (eg: $P(t)$, $P'(t), \dots$), and λ is a constant (0.01) which penalizes large readout weights. These same readout weights are then used during the probabilistic rollout to create the decoded rollout signal. Readout weights are optimized for each epoch.

2.5.4 FORCE

As a second type of baseline, we implemented FORCE training. This type of training, and its resulting outputs, differ greatly from BPTT approaches, and provide additional insight as to how a network may fail at the presented tasks. This network consists of a single recurrent network of leaky integrators, and a linear readout. At each point in time, the value from the readout node is fed back into the network according to a random weight matrix, similar to the rollout mode in the BPTT models. Notably, the only weights that adjust are the weights from the readout node back into the network, thus reducing the number of learned parameters compared to the BPTT approach. These readout weights are adjusted quickly according to the recursive-least-squares (RLS) learning rule. During the rollout period this network performs similarly for both local and teacher ratios. This is because the tight control between the readout term and network activity is directly optimizing the ‘autonomous’ mode of function. However, as learning progresses and the network must operate in the presence of unreliable teaching signals, both the local and validation errors increase significantly. This change occurs because the network is able to drift slightly during the autonomous periods, resulting in an inaccurate correlation of presynaptic activity (Eq 2.24) which results in RLS driving learned readout weights incorrectly. This illustrates how a network that may perform well initially, even at full rollout, degrades in the presence of unreliable external signals.

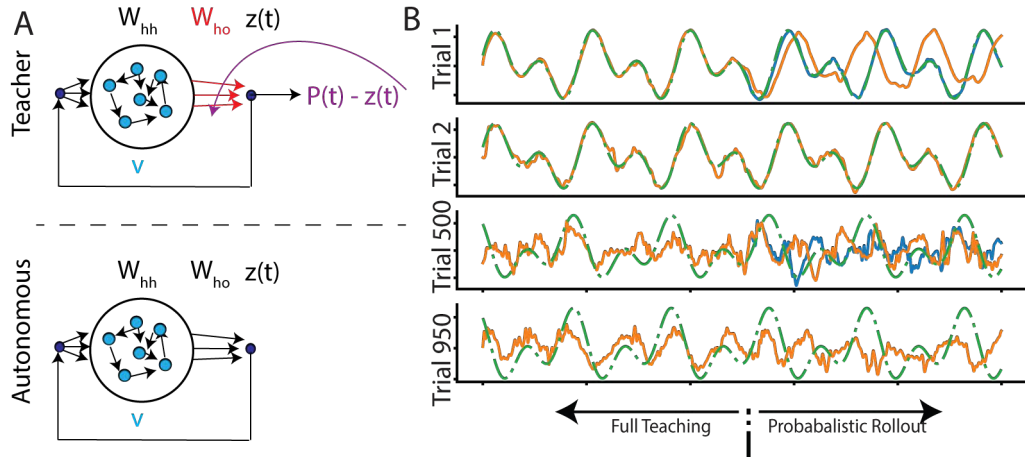


Figure 2.10: FORCE Training. **A** Setup for training of FORCE. During teacher frames (top) the readout value $z(t)$ is subtracted from the goal readout $P(t)$. This error signal is used to adjust the readout weights W_{ho} . During autonomous frames, the system evolves the same as during teaching, but the readout weights are not updated. **B** Example of training FORCE networks. (Top) During the first trial, the network is able to closely follow the inputs signal, due to the quick adjustment of readout weights. During full autonomous (orange) however, the network begins to diverge, indicating that the 150 frames presented during full teacher are not sufficient to train the readout weights. During the second trial (Second row) the network is able to perform both the probabilistic autonomous and full autonomous conditions. However, as training progresses (third and bottom row) the network produces a noisy and chaotic output. The unreliable teaching signal has driven the readout weights and correlation matrix from their initially tuned regime, and even full teaching frames can not bring the network back to reliable readout.

Units in the FORCE network are simple recurrent networks with a leak term, and a linear readout. The update of the system was defined as:

$$\begin{aligned} \tau \frac{dv}{dt} &= -v(t) + W_{hh}r(t) + W_{ih}I(t) + W_{zh}Z(t-1) \\ r &= \tanh(v(t)) \\ z(t) &= W_{ho}r(t) \end{aligned} \tag{2.23}$$

Where $r(t)$ represents the “firing rate” of the hidden units, $z(t)$ the activity of the readout units, and $I(t)$ represents the task-dependent input signal. At every time

step a running average of presynaptic correlations was updated:

$$Q(t) = Q(t-1) - \frac{Q(t-1)r(t)r^T(t)Q(t-1)}{1 + r^T Q(t-1)r(t)} \quad (2.24)$$

On time steps where learning occurred, readout weights were updated according to the recurrent least squares (RLS) algorithm:

$$\begin{aligned} e(t) &= z(t) - y(t) \\ W_{ho}(t) &= W_{ho}(t-1) - e(t)Q(t)r(t) \end{aligned} \quad (2.25)$$

Weights were initialized and parameters set as standard (Sussillo & Abbott, 2009) and reported in table 2.10, Weights were initialized according to a Gaussian distribution, and clamped to zero with a probability 1-p (see table for values).

τ	10ms
N_I	2
N_h	1000
N_z	2
W_{ih}	$Gauss(0, \frac{1}{\sqrt{N_I}})$, p=0.1
W_{hh}	$1.5 * Gauss(0, \frac{1}{\sqrt{N_h}})$, p=0.1
W_{hz}	0, p=1
W_{zh}	$Unif(-1, 1)$, p=1

Table 2.10: Parameters used for FORCE network. Weight parameters indicate the probability distribution that initial values were drawn from.

CHAPTER 3

Egocentric boundary vector tuning of the retrosplenial cortex

This chapter has appeared as a published article: Alexander, A. S., Carstensen, L. C., Hinman, J. R., Raudies, F., Chapman, G. W., & Hasselmo, M. E. (2019). Egocentric boundary vector tuning of the retrosplenial cortex. *Science Advances* 6(8), eaaz2322

3.1 INTRODUCTION

Spatial cognition is a critical component of intelligent behavior. The ability to effectively recall and navigate between known goals relies on stored representations of spatial interrelationships. Further, episodic experiences can be thought of as situated within a stored mental map indicating the places in which events occurred. Spatial representations that support both navigation and episodic memory are observed in many brain regions, including the hippocampus and medial entorhinal cortex, where neurons exhibit receptive fields that are correlated with the position or orientation of the animal relative to the array of locations and cues that define the structure of the outside world. This viewpoint-invariant coordinate system is commonly referred to as the allocentric reference frame (Hafting et al., 2005; O'Keefe & Dostrovsky, 1971; Taube et al., 1990). Although it has been repeatedly shown that intact function of allocentric spatial circuits is critical for spatial memory and navigation (Morris et al., 1982; Steffenach et al., 2005), it is important to consider that all spatial information enters the brain via sensory organs and their corresponding processing streams. Accordingly, knowledge of the position of a prominent landmark and a neighboring goal location would be, at least initially, incorporated into a stored spatial map in egocentric coordinates relative to

the animal itself (Bicanski & Burgess, 2018; Byrne et al., 2007). Further, enacting navigational plans can be based on stored allocentric representations but would ultimately require translation into sequences of actions anchored in an egocentric reference frame (e.g., one turns clockwise relative to their own previous orientation position) (Bicanski & Burgess, 2018; Byrne et al., 2007; Whitlock et al., 2012). Neural mechanisms by which egocentric and allocentric coordinate systems are interrelated are still the subject of intense examination. Computational models have predicted that cortical networks capable of integrating allocentric and egocentric information for either constructing or using stored spatial representations require neurons with egocentric sensitivity to external locations (Bicanski & Burgess, 2018; Byrne et al., 2007). Most investigations into egocentric representations in unconstrained animals have focused on the neural substrates of path integration, a navigational computation wherein self-location is approximated via continuous integration of angular and linear displacement (Mittelstaedt, 1980). Neural correlates of these movement variables have been reported in several structures (Whitlock et al., 2012; Cho & Sharp, 2001; Hinman et al., 2016; McNaughton et al., 1994; Wilber et al., 2014). Only recently have externally anchored egocentric representations that extend beyond self-motion been reported (Gofman et al., 2019; Hinman et al., 2019; LaChance et al., 2019; Peyrache et al., 2017; Wang et al., 2018a; Wilber et al., 2014). Egocentric representations of this nature may anchor to environmental boundaries. Boundaries present a unique intersection between egocentric and allocentric coordinate systems, as they have fixed positions that define the navigable allocentric space and simultaneously restrict the egocentric affordances of the agent such as what can be viewed or what motor plans can be executed. Environmental bounds or walls extend along large regions of an environment and

thus enable extended interaction from multiple allocentric or egocentric perspectives. Egocentric neural responses have now been reported in multiple areas such as lateral entorhinal cortex (Wang et al., 2018a), dorsal striatum (Hinman et al., 2019), and postrhinal cortices (Gofman et al., 2019; Wang et al., 2018a). However, none of these regions have the reciprocal inter-connectivity between egocentric and allocentric spatial circuitry that might mediate bidirectional reference frame transformations. From a connectivity standpoint, the retrosplenial cortex (RSC) is an excellent candidate to examine egocentric representations during navigation. Further, theoretical work has posited that RSC forms a computational hub for supporting coordinate transformations (Bicanski & Burgess, 2018; Byrne et al., 2007; Clark et al., 2018; Rounds et al., 2018). RSC is composed of two interconnected subregions, dysgranular and granular, which have slightly different connectivity with cortical and subcortical regions. Dysgranular RSC (in mice agranular RSC) is positioned along the dorsal surface of the brain and has biased interconnectivity with association, sensory, and motor processing regions that code in egocentric coordinates (van Groen & Wyss, 1992; Vogt & Miller, 1983). In contrast, granular RSC has strong reciprocal innervation with the hippocampal formation and associated structures that are primarily sensitive to the allocentric coordinate system (van Groen & Wyss, 1992). Despite having dense reciprocal connectivity with numerous regions known to support spatial cognition, few reports have examined spatial response properties of neurons within the RSC. Most assessments of functional properties of RSC neurons have occurred in rodents performing track running tasks (Alexander & Nitz, 2017, 2015; Mao et al., 2017; Miller et al., 2019; Smith et al., 2012). Track running experiments have revealed that RSC neurons exhibit spatial correlates with conjunctive sensitivity to allocentric and egocentric

coordinate systems (among others) (Alexander & Nitz, 2015). Conjunctive tuning of this type has been shown in modeling work to facilitate spatial coordinate transformations, further supporting a role for RSC in the required transformation between these two spatial reference frames (Bicanski & Burgess, 2018; Pouget & Sejnowski, 1997). However, grid cells, head direction cells, place cells, and other forms of well-characterized spatial receptive fields have primarily been examined in two-dimensional (2D) environments. Only a few experiments have studied RSC in similar conditions, and all such reports have focused on head direction encoding (Cho & Sharp, 2001; Chen et al., 1994; Jacob et al., 2017). To examine externally referenced egocentric representations in RSC capable of supporting both navigation and reference frame transformations, we recorded from both RSC subregions while rats freely explored familiar 2D environments. We report that subsets of RSC neurons exhibit a variety of spatially stable activation patterns in egocentric and allocentric coordinate systems. These findings support predictions from computational modeling related to translation between spatial reference frames and highlight important navigation-related variables encoded in association cortex (Bicanski & Burgess, 2018; Byrne et al., 2007).

3.2 RESULTS

3.2.1 RSC neurons exhibit stable spatial activity during free exploration

We recorded 555 neurons extracellularly in RSC in either hemisphere from male Long-Evans rats ($n = 7$) during free exploration. To enable comparisons between functional properties of neurons recorded in dysgranular versus granular subregions of RSC, we estimated tetrode placement and depth for each session (Fig. 3.1A, $n = 130$ sessions). Of the total population, 41.5% ($n = 230/555$) were recorded

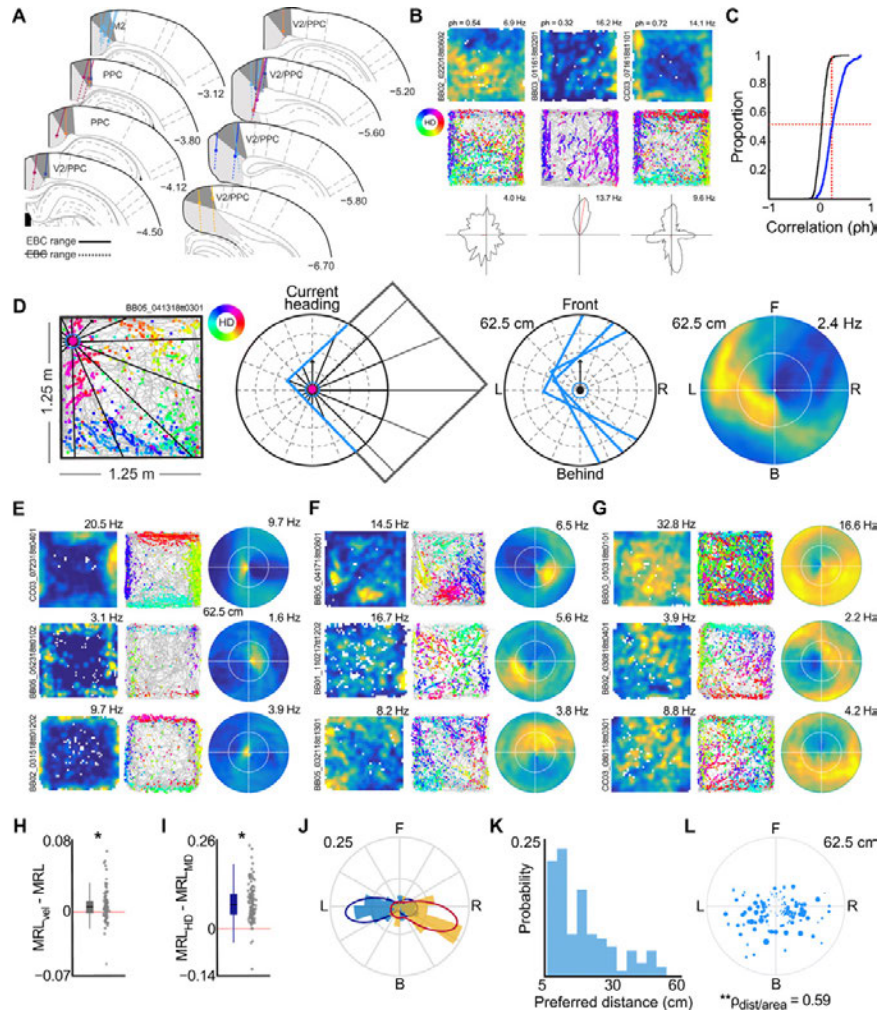


Figure 3.1: Caption on next page

from dysgranular RSC, 15.1% ($n = 84/555$) from the border between dysgranular RSC and granular RSC, and 43.4% ($n = 241/555$) within granular RSC. For baseline sessions, rats foraged for scattered reward in 1.25-m^2 arenas with observable fixed distal cues.

RSC neurons exhibited complex firing rate fluctuations, as rats randomly foraged within open arenas (Fig. 3.1B). To assess the spatial stability of these representations for each neuron individually, we began by examining correlations between 2D spatial firing ratemaps constructed from first and second halves of each exper-

Egocentric boundary vector representations of RSC neurons during free exploration. (A) Locations of RSC tetrode tracts where neurons with egocentric boundary sensitivity were observed. (B) Example 2D ratemaps (top), trajectory plots (middle), and head direction tuning plots (bottom) for three RSC neurons with significant stability in spatial firing. The location of individual spikes is shown with colored circles, which indicate the corresponding head direction of the animal according to the legend on the left. (D) Schematic for construction of EBRs. Left and middle panels: An example spike is mapped with respect to egocentric boundary locations in polar coordinates. (E) 2D ratemaps, trajectory plots, and EBRs for three example RSC EBCs with animal-proximal receptive fields. (F) Same as in (E) but for three RSC EBCs with animal-distal receptive fields. (H) Difference in strength of EBC tuning when a speed threshold was applied (MRL_{vel}) versus no speed threshold (MRL). (I) Difference in strength of EBC tuning when egocentric bearing was referenced to head direction (MRL_HD) rather than movement direction (MRL_MD). (J) Polar histogram of preferred orientation of receptive field across all RSC EBCs. Yellow and blue bars correspond to EBCs recorded in the left and right hemisphere, respectively. (K) Distribution of preferred distance of all RSC EBC receptive fields. (L) Polar scatter plot of preferred orientation versus preferred distance. Circle size indicates the area of the egocentric boundary vector receptive field.

imental session (Fig. 3.1C). Across the full population of RSC neurons, 45.8% of cells ($n = 254/555$) had spatial correlations greater than the 99th percentile ($\rho = 0.23$) of the distribution of correlations observed following 100 random shifts of the complete spike train for each neuron relative to spatial position (Fig. 3.1C). In some cases, RSC neurons with spatially anchored responses had slight differences in basic firing properties from those that were not spatially stable. Of particular interest and consistent with the presence of spatial receptive fields, RSC neurons with spatially reliable activity had significantly more spatial coherence and broader peak firing rate dispersion than nonstable cells. Spatially anchored firing patterns were also observed at more ventral recording sites, where it was difficult to resolve whether the recording tetrode was in RSC or the cingulum bundle. Recordings from these sites were not included in the pool of RSC neurons for

analysis.

3.2.2 Egocentric boundary vector responsivity of RSC

Of neurons with stable spatial firing in the open field, several had receptive fields that were qualitatively proximal to environmental boundaries (Fig. 3.1B, right). Inspection of the relationship between each spike and the head direction or movement direction of the animal revealed that these responses manifested when the animal was oriented in a similar manner relative to any wall, suggesting that the receptive field was defined in an egocentric manner. As such, these responses were reminiscent of egocentric boundary cells (EBCs) recently reported in the dorsal striatum (Hinman et al., 2019), lateral entorhinal cortex (Wang et al., 2018a), and postrhinal cortex (Gofman et al., 2019; LaChance et al., 2019). To test this explicitly, we constructed egocentric boundary ratemaps (EBRs) using procedures previously described (Fig. 3.1D) (Hinman et al., 2019). Briefly, for each behavioral frame, the distance to the nearest wall in each 3° offset from the animals head direction or movement direction was calculated (Fig. 3.1D). The same process is repeated for the position of each spike from each neuron, and then ratemaps in polar coordinates were constructed by dividing the number of spikes by the total behavioral occupancy in seconds. From each EBR, we computed the mean resultant length (MRL) of angular tuning for the full session as well as the absolute difference in angular tuning direction and distance between first and second halves of the baseline session. RSC cells were determined to exhibit significant egocentric boundary sensitivity if they met the following criteria: (i) they had an MRL for the session that was greater than the 99th percentile of the distribution of resultants computed following repeated shifted spike train randomizations, (ii) they had an absolute

difference of mean directional tuning between halves of the baseline session that was less than 45° , and (iii) the change in preferred distance relative to the full session was less than 50% for both halves. Using these metrics, 17.3% ($n = 96/555$) of RSC neurons were determined to be EBCs (Fig. 3.1, E to G). Application of a speed threshold (>5 cm/s) modestly increased the MRL and size of the EBC population (Fig. 3.1H, $n = 106/555$, 19.1%; median MRL difference = 0.006, interquartile range (IQR) = 0.0012 to 0.013; Wilcoxon signed-rank, $z = 4.70$, $P = 2.65 \times 10^{-6}$). From this result, we hypothesized that egocentric receptive fields of EBCs were defined by the movement direction of the animal rather than head direction, which can be computed even when the animal is motionless. RSC neurons (21.4%, $n = 119/555$) were determined to be EBCs when referenced to movement direction, but egocentric boundary vectors were overall more strongly tuned to head direction (Fig. 3.1I, EBC MRL with head direction = 0.22, IQR = 0.17 to 0.29; EBC MRL with movement direction = 0.14, IQR = 0.11 to 0.20, Wilcoxon signed-rank test for zero difference, $z = 9.61$, $P = 7.60 \times 10^{-22}$). A subpopulation of the recorded population, 16.4% ($n = 91/555$), were detected as EBCs when using either head direction or movement direction as the angular reference. However, a subpopulation of neurons were only detected as EBCs when using head direction ($n = 15/555$) or movement direction ($n = 28/555$). Because the directional estimates are largely dependent on one another and there were no differences in the stability of receptive fields for EBCs detected using either, we pooled these populations ($n = 134/555$, 24.1%).

3.2.3 Properties of RSC egocentric boundary vector receptive fields

Subpopulations of RSC EBCs exhibited either increased or decreased activation when the animal occupied a particular orientation and distance relative to envi-

ronmental boundaries (Fig. 3.1, E to G). In accordance with previous literature, we refer to those neurons that were inhibited as inverse EBCs (Fig. 3.1G) and neurons with excitatory receptive fields as EBCs (Fig. 3.1, E and F). Identification of the center of mass of EBC receptive fields (either excitatory or inhibitory) revealed a bimodal distribution of preferred orientations that was best fit by a two-component Von Mises mixture model with means of $108.1^\circ(\text{L})$ and $249.3^\circ(\text{R})$ relative to directly in front of the animal (Fig. 3.1J; number of Von Mises mixture components determined via sequential improvement in log-likelihood ratio and cross-validation). EBCs were recorded in both hemispheres, and there was a contralateral relationship between the preferred orientation and the hemisphere in which the neuron was recorded, with larger numbers of rightward EBC receptive fields recorded in the left hemisphere and vice versa (Fig. 3.1J, right-hemi EBCs preferring left side = 32/54, 59.3%; left-hemi EBCs preferring right side = 55/80, 68.8%; Watson-Williams test, $F = 41.93$, $P = 1.7 \times 10^{-109}$). The distribution of preferred distances was skewed toward animal-proximal receptive fields, with numerous EBCs having receptive fields extending beyond range of direct somatosensory stimulation of whiskers (Fig. 3.1K). Further, the size of EBC receptive fields increased as a function of the preferred distance of the egocentric vector, indicating that the resolution of the representation was dependent on proximity to boundaries (Fig. 3.1L, Spearman's $\rho = 0.59$, $P = 5.1 \times 10^{-14}$). The presence of EBCs with preferred distances distal to the animal suggested that the EBC response property was not dependent on physical interaction with arena borders.

3.2.4 EBC responses are localized within dysgranular RSC but lack topographic organization

Egocentric boundary vector sensitivity was primarily observed in dysgranular RSC, where 38.7% ($n = 89/230$) of neurons recorded were classified as EBCs. In contrast, EBCs were observed in 9.9% ($n = 24/241$) of granular RSC and 25.0% ($n = 21/84$) of intermediary area cells between the two subregions. By and large, the distribution of EBCs among RSC subregions was consistent across animals. The EBC response property was observed across a wide range of anterior/posterior (A/P) coordinates spanning most of the RSC but had no further anatomical organization beyond subregion specificity (range = 2.9 to 6.8 mm relative to bregma). The distribution of spike waveform widths across all RSC neurons was bimodal, with identified EBCs primarily found in the cluster of neurons with longer duration waveforms (k-means clustering on waveform width, cluster 1 median = 0.15 s, IQR = 0.15 to 0.18 s, EBCs in cluster 1, $n = 9/175$, 5.1%; cluster 2 median = 0.31 s, IQR = 0.31 to 0.34 s; EBCs in cluster 2, $n = 125/380$, 32.9%). Further, EBCs had overall low mean firing rates (EBCs = 1.62 Hz, IQR = 0.90 to 2.75 Hz, not-EBCs = 3.58 Hz, IQR = 1.15 to 8.51 Hz). Together, the EBC subpopulation was determined to be primarily composed of putative principal neurons, suggesting that the EBC signal is propagated across RSC subregions or into other brain regions. EBCs could often be simultaneously recorded, which enabled an analysis of potential topography in the distribution of preferred distance and orientation of the egocentric boundary vector. Overall, 142 pairs of RSC EBCs were co-recorded across 29 sessions. Of these pairs, 31.7% ($n = 45/142$) were recorded on the same tetrode, while the remaining 68.3% of EBC pairs ($n = 97/142$) were concurrently recorded on different tetrodes. To assess whether there was organization to preferred orientation

and distance as a function of proximity of two EBCs (i.e., observed on the same or different tetrodes), we next calculated the difference in receptive field center of mass for both angular and distance components for all pairs. Neither preferred distance nor orientation was statistically different for EBCs recorded on the same versus different tetrodes (absolute difference in preferred distance same tetrode = 7.5 cm, IQR = 5 to 12.5 cm; absolute difference in preferred distance different tetrode = 10 cm, IQR = 2.5 to 17.5 cm; Wilcoxon rank sum test, $z = 1.44$, $P = 0.15$; absolute difference in preferred orientation same tetrode = 63° , IQR = 20.25 to 93° ; different tetrodes = 66.0° , IQR = 24.0 to 114.1° ; Wilcoxon rank sum test, $z = 1.08$, $P = 0.28$). Accordingly, we conclude that there is a lack of topographic organization of egocentric boundary vector tuning in the RSC.

3.2.5 Egocentric boundary vector tuning in secondary motor cortex and posterior parietal cortex but not medial entorhinal cortex

In three animals, a subset of more anterior recording tetrodes were positioned in secondary motor cortex (M2, from bregma: A/P: 1.1 to 2.9 mm, medial/lateral (M/L): 0.8 to 1.2 mm), and 56 neurons were recorded there. Of M2 neurons, 37.5% reached EBC criterion ($n = 21/56$). Similarly, 95 neurons across five rats were recorded in posterior parietal cortex (from bregma: A/P: 3.7 to 5.9 mm, M/L: 1.5 to 2.4 mm) and a subpopulation of 21.1% ($n = 20/95$) reached EBC criterion. EBCs and inverse EBCs were observed in both structures, and receptive fields had similar angular and distance distributions as those observed in RSC. In contrast, only 3.0% ($n = 9/297$) of medial entorhinal cortex neurons recorded in similar conditions reached EBC criterion, indicating that the egocentric vector signal was generally not present within the region (Wang et al., 2018a).

3.2.6 EBC responsivity is not explained by self-motion correlates

In free exploration, spatial locations near environment boundaries uniquely restrict the behavioral affordances of the animal. Many observed EBC receptive fields were proximal to the rat, firing only when the animal was close to boundaries and thus most limited in its possible actions. We next tested whether the manifestation of egocentrically referenced boundary vector tuning was in actuality reflective of self-motion-related firing that was stereotyped near borders. We began by constructing self-motion-referenced firing ratemaps during open-field sessions (Whitlock et al., 2012). The angular difference between movement direction ($\Delta\theta$) and the Euclidian distance in 2D location was calculated across a sliding 100-ms window for every position of the animal throughout a free exploration session (Fig. 3.2A, left). These displacement values were converted to Cartesian coordinates referenced to the previous location of the animal at each step, thus producing a map of the distance and direction of movement of the animal for all position samples within the exploration session (Fig. 3.2A, middle and right).

Firing rate as a function of these displacement values is presented for representative RSC neurons in Fig. 3.2 (C to F). The zero-line intersection indicates the position of the animal at the beginning of each 100-ms window, and the x and y axes reflect displacement in lateral and longitudinal dimensions, respectively. Thus, values to the right of the vertical zero line reflect the activity of the neuron when the animal moved to the right relative to the previous position and direction of its body axis, and the distance that the action took the animal is reflected in the position of the value along the y axis. To quantify the stability of self-motion tuning, we correlated self-motion ratemaps for each neuron that were individually computed from interleaved temporal epochs (1 s in duration) within the free exploration ses-

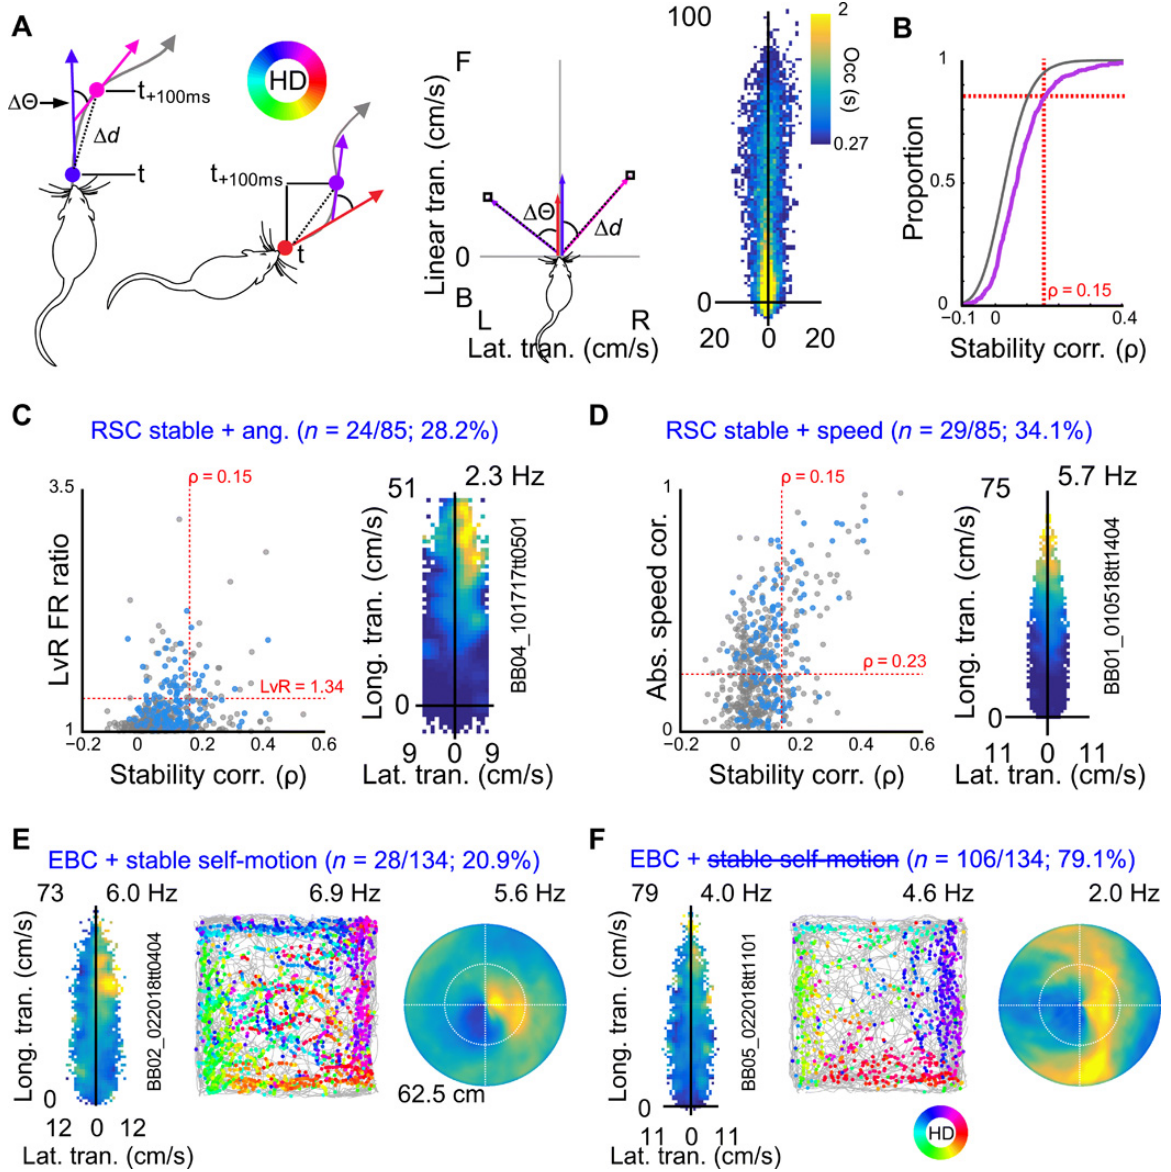


Figure 3.2: Caption on next page

RSC egocentric boundary vector representations cannot be explained purely by self-motion correlates. (A) Schematic of generation of self-motion referenced ratemaps. Left: Example angular and distance translation across 100-ms temporal windows for two hypothetical position samples. Middle: Corresponding lateral and longitudinal translation for left examples in self-motion referenced coordinates. Right: Heat map depicting mean occupancy in seconds for lateral and longitudinal translation combinations across a complete experimental session. d , distance displacement; θ , angular displacement; colored arrows depict individual position samples 100 ms apart; gray arrow, hypothetical trajectory. (B) In pink, cumulative density functions for self-motion ratemap stability values (Spearman's ρ) for all RSC neurons (randomization in gray). Red vertical line shows 95th percentile of randomized distribution and its intersection with the real distribution of spatial stability. Percentage of neurons above red horizontal line have significant self-motion stability. (C) Left: Self-motion stability score (x axis) versus absolute ratio of activity on left versus right halves of self-motion ratemaps (y axis) for all RSC neurons. Blue dots correspond to identified RSC EBCs. Red lines and corresponding values correspond to 95th percentiles of randomized distributions for both metrics. Neurons with values in upper right region were determined to have significant angular displacement tuning. Right: Example RSC neuron with significant firing rate modulation for clockwise movements. (D) On the left, same as in (C) but for self-motion stability score versus absolute correlation between mean firing rate and longitudinal displacement (speed). Right: Example RSC neuron with significant firing rate modulation as a function of longitudinal displacement. (E) Example RSC EBC with stable self-motion correlates. (F) Example RSC EBC without stable self-motion correlates.

sion. A subset of RSC neurons (15.3% ($n = 85/555$)) exhibited self-motion-related activity that had greater stability than the 95th percentile of the distribution of stability correlation values calculated following permutation tests (Fig. 3.2B). Of this subpopulation, 28.2% ($n = 24/85$) had firing rate modulation that was biased for leftward or rightward movements (Fig. 3.2C), while 34.1% ($n = 29/85$) were sensitive to longitudinal movements consistent with speed tuning (Fig. 3.2D). Of the EBC population, 20.9% ($n = 28/134$) met the stability criteria, indicating that a small subpopulation of neurons exhibiting egocentric boundary vector tuning had stable self-motion correlates (Fig. 3.2E). However, the vast majority of RSC EBCs did not exhibit self-motion correlates, confirming that egocentric boundary vector

tuning was primarily not an epiphenomenon of movement-related activity near borders (Fig. 3.2F, $n = 106/134$, 79.1%).

Beyond EBCs, the present analysis demonstrated overall limited self-motion tuning in RSC during free exploration. This observation shines new light on previously reported turn-sensitive neurons in RSC during track running tasks (Alexander & Nitz, 2015). In previous work, the magnitude of clockwise or counterclockwise activation during track running was demonstrated to be generally insensitive to the magnitude of angular velocity on a trial-by-trial basis. In combination with the lack of self-motion tuning during free foraging observed here, the results collectively suggest that reported egocentric correlates in RSC are externally referenced and unrelated to the speed of angular movement.

3.2.7 Generalized linear models demonstrate robust egocentric vector tuning of RSC EBCs

Self-motion is necessarily conflated with egocentric boundary vector tuning because the response primarily manifested during movement. An EBC may exhibit stable firing rate fluctuations as a function of self-motion that are driven by the egocentric boundary vector receptive field, not the action state of the animal. For example, an EBC with a receptive field to the animals left may also show self-motion tuning for clockwise movements as a result of the animal being more likely to turn clockwise when there is a wall occupying the egocentric receptive field. Yet, the same neuron may not be activated when the animal turns clockwise in other locations within the environment that do not satisfy the egocentric boundary vector. Thus, although informative about the prevalence of self-motion sensitivity in RSC, a different approach was required to tease out the influence of self-motion and

other potential spatial covariates on EBC activity patterns. We next implemented a generalized linear model (GLM) framework to predict the probability of spiking at each time point as a function of the relative influence of classes of allocentric, self-motion, or EBC-related predictors (Fig. 3.3A). Allocentric predictors included the head direction of the animal and x and y positions within the arena. Self-motion-related predictors included linear speed and angular displacement (i.e., the differential of animal movement direction in 100-ms windows)

EBC-related predictors were more complicated as a single position sample or spike had relationships to multiple locations along boundaries simultaneously. Accordingly, the EBC predictor could take many forms. To minimize the number of subpredictors, EBC predictors were composed of the animal's distance and egocentric bearing to the center of the arena. Unlike arena boundaries, the center of the arena is a single coordinate that can be described as a function of individual angular and distance components or their conjunction for each position sample. Critically, EBCs were found to exhibit robust egocentric bearing and distance tuning to the center of the arena, making the predictor a reasonable counterpart to referencing single unit activity to arena walls.

We assessed the overall influence of each predictor class (allocentric, self-motion, and EBC-related) on model fit by constructing a nested GLM, dropping each predictor class, and then making comparisons between resulting model fits. Figure 3B depicts the difference of Akaike information criteria (dAIC), a metric of the decrement to model fit, for both EBCs and non-EBCs between the full model and reduced models, with all allocentric, self-motion, or egocentric boundary predictors removed. Larger dAIC values indicate greater impact of the predictor class within the full model. Models without EBC and allocentric predictors had significant dif-

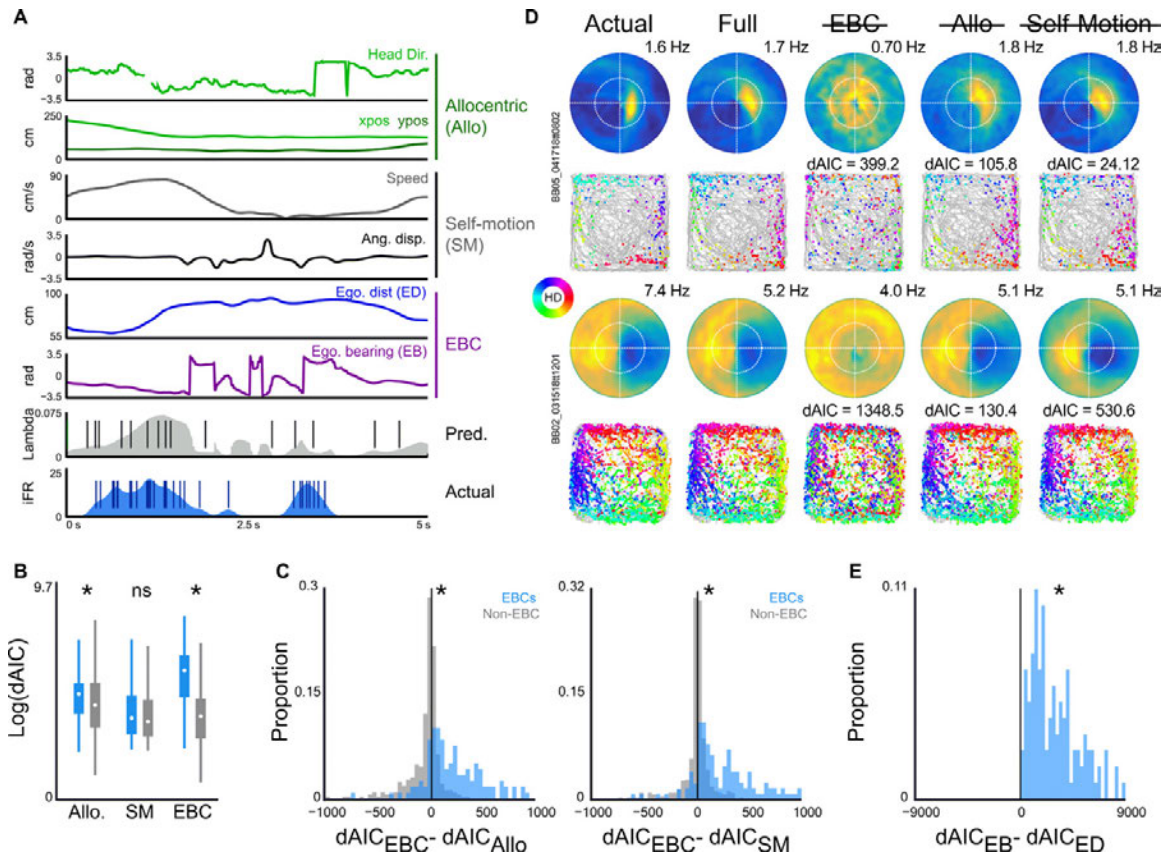


Figure 3.3: Egocentric vector tuning is more robust than allocentric or self-motion correlates using a generalized linear modeling framework. (A) Example GLM predictors composing allocentric, self-motion, and egocentric vector classes with corresponding actual and predicted firing rates and spike trains over a 5-s window. (B) Boxplots depicting median and quartiles of log-transformed dAIC scores for models, with all allocentric, self-motion, or egocentric vector predictors removed (blue bars, EBCs; gray bars, non-EBCs). Larger dAICs indicate greater error in model fit with removal of a predictor class. (C) Comparison of dAIC scores for models with egocentric vector versus allocentric predictors removed (left) or egocentric vector versus self-motion predictors removed (right) for EBCs (blue) and non-EBCs (gray). Rightward shifts indicate greater error in model fit for models with removed egocentric vector predictors. (D) For two example RSC EBCs, predicted GLM spike trains from all models were used to construct EBRs and trajectory plots. Left column: Actual EBR and corresponding trajectory plot below. Second column: For the same cell, an EBR and corresponding trajectory plot for the GLM constructed using all egocentric vector, allocentric, and self-motion predictors. Final three columns: EBRs and trajectory plots for each reduced model and corresponding dAIC scores. (E) Comparison of dAIC scores for models with the egocentric bearing versus the egocentric distance removed reveal greater impact of egocentric bearing for EBCs

ferences in fit between EBCs and non-EBCs (Fig. 3.3B, Kruskal-Wallis, $\chi^2 = 270.73$, $P = 1.95 * 10^{-56}$, post-hoc Scheffe test EBC predictors, $P = 1.73 * 10^{-39}$, post-hoc Scheffe test allocentric predictors, $P = 0.03$). There was no difference between these subpopulations of RSC neurons for the removal of self-motion covariates from the model, further supporting that EBCs were not more sensitive to speed or angular displacement than the remainder of the RSC population (Fig. 3.3B, post hoc Scheffe test self-motion predictors, $P = 0.17$).

A clear divergence emerged in the importance of EBC-related predictors for the EBC and non-EBC subpopulations. As reflected in the difference in magnitude of dAIC, EBC predictors had greater impact than either allocentric or self-motion predictors for the EBC population (Fig. 3.3C, blue; $dAIC_{EBC} - dAIC_{Allo}$ for EBCs = 186.8, IQR = 33.8 to 450.9; $dAIC_{EBC} - dAIC_{SM}$ for EBCs = 238.8, IQR = 44.1 to 507.1) in contrast to the non-EBC population, which had similar dAIC scores (near 0) for models lacking EBC predictors and other predictor classes (Fig. 3.3C, gray; $dAIC_{EBC} - dAIC_{Allo}$ for non-EBCs = 15.3, IQR = 113.3 to 9.8; $dAIC_{EBC} - dAIC_{SM}$ for non-EBCs = 0.05, IQR = 22.5 to 29.8). Overall, the impact of EBC-related predictors relative to other predictor classes was significantly greater for EBC versus non-EBC sub-populations ($dAIC_{EBC} - dAIC_{Allo}$ for EBCs versus non-EBCs, Wilcoxon rank sum, $z = 12.0$, $P = 1.68 * 10^{-33}$; $dAIC_{EBC} - dAIC_{SM}$ for EBCs versus non-EBCs, Wilcoxon rank sum, $z = 11.2$, $P = 2.60 * 10^{-29}$). EBCs observed in both parietal cortex and M2 were also substantially more influenced by egocentric vector predictors than either allocentric or self-motion predictor classes (parietal cortex: $dAIC_{EBC} - dAIC_{Allo}$ for EBCs versus non-EBCs, Wilcoxon rank sum, $z = 3.01$, $P = 0.003$; $dAIC_{EBC} - dAIC_{SM}$ for EBCs versus non-EBCs, Wilcoxon rank sum, $z = 4.55$, $P = 5.35 * 10^{-6}$; M2: $dAIC_{EBC} - dAIC_{Allo}$ for EBCs versus non-EBCs,

Wilcoxon rank sum, $z = 3.72$, $P = 1.96 * 10^{-4}$; $dAIC_{EBC} - dAIC_{SM}$ for EBCs versus non-EBCs, Wilcoxon rank sum, $z = 3.42$, $P = 6.29 * 10^{-4}$).

These results suggested that although models without allocentric or self-motion predictors could yield significantly decreased model fit, the vast majority of EBC neurons were significantly more affected by EBC predictors. Two example EBCs in Fig. 3.3D illustrate this point, wherein a spike train generated from the output of each model was used to construct trajectory plots and EBRs. In both cases, the model lacking egocentric orientation and distance information yields a trajectory plot and EBR that is substantially poorer at reconstructing the actual data than any other reduced model. Although egocentric predictors were the dominant influence on EBC activation, all EBCs were statistically affected by the removal of more than one predictor category (assessed via chi-square tests of log-likelihood ratios, $P < 0.001$). In this manner, the GLM analyses also revealed that RSC EBCs were conjunctively sensitive to the position of arena boundaries in egocentric coordinates and allocentric heading or location simultaneously. This feature of EBC responsiveness is consistent with theoretical work proposing a transformation between egocentric and allocentric spatial representations within RSC (Bicanski & Burgess, 2018).

3.2.8 GLM confirms vectorial representation

The use of the GLM framework provided an opportunity to verify that RSC neurons with egocentric boundary sensitivity actually formed vector representations of the relationships between environmental boundaries and the animal. By dropping out egocentric bearing and egocentric distance from the model individually, we were able to investigate the relative influence of the individual components

of the egocentric boundary vector in isolation for each neuron. Significant model decrements were observed in 93.3% ($n = 125/134$) of EBCs following removal of the egocentric bearing component and 55.9% ($75/134$) of EBCs were affected by the removal of egocentric distance predictors. Overall, the magnitude of error to model fit was substantially greater when egocentric bearing was removed, indicating that, although both distance and orientation components are critical for egocentric boundary vector responsiveness, the directional component more robustly drives neurons exhibiting this tuning preference (Fig. 3.3E, difference in dAIC for egocentric bearing versus egocentric distance = 2269.7, IQR = 1283.3 to 3767.5; Wilcoxon signed-rank test, $z = 9.7$, $P = 2.96 \times 10^{-22}$). Consistent with observations from the GLM, the number of detected EBCs decreased as the allowed variability in preferred distance across the session was systematically decreased (i.e., detected EBCs were required to have more reliable distance tuning across the session). These results indicate that although a large proportion of RSC neurons can be described as having egocentric boundary vectors, a number of neurons are primarily responsive to egocentric boundary bearing as observed in the lateral entorhinal and postrhinal cortices (Gofman et al., 2019; LaChance et al., 2019; Wang et al., 2018a).

3.2.9 EBCs respond to local, not distal, environmental features

Characterization of EBC properties and self-motion correlates were conducted in baseline sessions in which the open arena remained in a fixed location relative to the experimental room and fixed distal cues therein. We next conducted a series of experimental manipulations of the relationship between the familiar arena and the testing room to confirm that EBC response properties were defined by the

relationships between environmental boundaries and the animal itself. First, we rotated the open field 45° to maximally disrupt correspondence between arena and distal walls or cues present within the recording environment to verify that EBC responses were anchored to local boundaries and not the broader recording room. Under these conditions, we recorded a total of 65 RSC neurons (across four rats and 14 sessions), of which 44.6% ($n = 29/65$) had EBC sensitivity. Consistent with EBC responses being referenced to the rat, receptive fields in rotated arenas maintained the same orientation and distance with respect to the animal, although arena boundaries now fell along completely different allocentric axes (Fig. 3.4, A to C; difference between baseline and rotated preferred orientation = 9°, IQR = 11.25 to 33.0°; Wilcoxon signed-rank test, $z = 1.17$, $P = 0.24$; difference between baseline and rotated preferred distance = 3.0 cm, IQR = 3.0 to 5 cm; Wilcoxon signed-rank test, $z = 1.44$, $P = 0.15$). Although vector tuning remained intact, there were slight but significant changes to ratemap coherence between baseline and rotation sessions, which suggested that the quality of the egocentric boundary receptive field was decremented across conditions (Fig. 3.4D; difference between baseline and rotated ratemap coherence = 0.02, IQR = 0.01 to 0.04; Wilcoxon signed-rank test, $z = 2.12$, $P = 0.03$).

Consistency in tuning could emerge if the allocentric map anchored to local boundaries rather than distal cues. This was not the case, as a population of simultaneously recorded head direction cells ($n = 4$; Fig. 3.4A, right) exhibited similar mean tuning across the rotated and nonrotated conditions ($n =$ absolute median tuning difference = 11.9°, absolute maximum tuning difference = 20.9°). Accordingly, arena rotation experiments dissociated the directional component of EBCs from the allocentric reference frame of head direction cells.

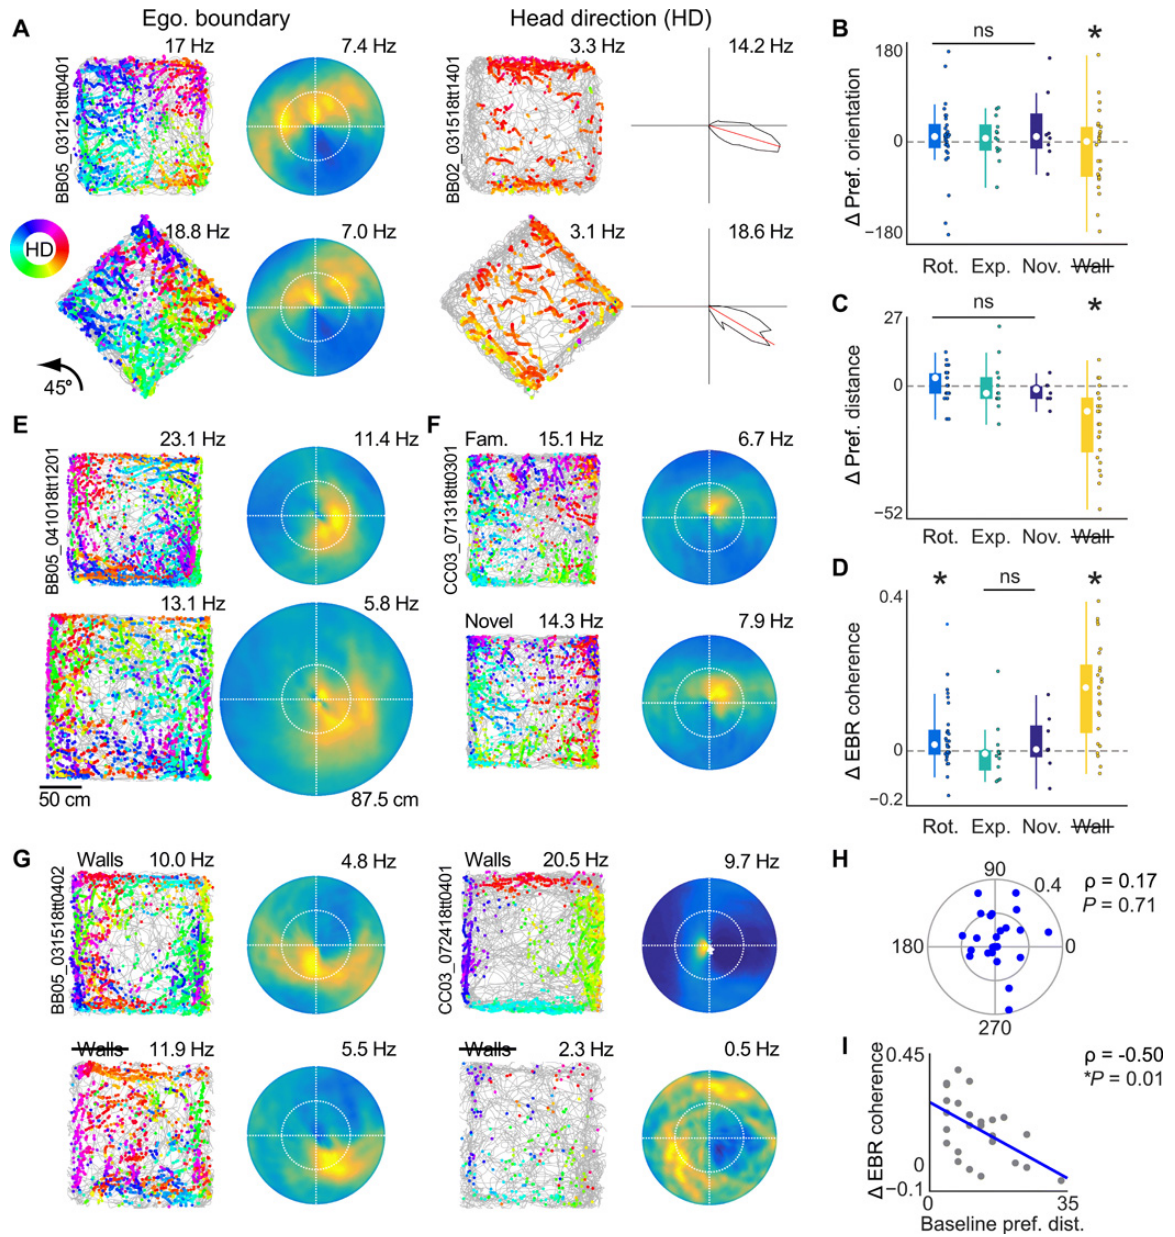


Figure 3.4: Caption on next page

EBCs are anchored to local boundaries, respond in novel environments, and lose sensitivity in arenas without explicit borders (A) Left: Trajectory plot and EBR for an example EBC with similar egocentric boundary vector tuning in baseline experimental session (top) and a second session in an environment rotated 45° (bottom). Right: Example head direction neuron sustains directional tuning across both conditions. (B) Preferred orientation of EBC receptive fields in all arena manipulation sessions subtracted from the preferred orientation in baseline sessions. (C) Preferred distance of EBC receptive fields in all arena manipulation sessions subtracted from the preferred distance in baseline sessions. (D) EBC receptive field coherence in all arena manipulation sessions subtracted from receptive field coherence in baseline sessions. (E) Trajectory plot and EBR for an example EBC with similar egocentric boundary vector tuning in baseline experimental session (top) and a second session in an expanded arena (bottom). (F) Trajectory plot and EBR for an example EBC with similar egocentric boundary vector tuning in baseline experimental session (top) and a second session in a novel arena (bottom). (G) Trajectory plot and EBR for two example EBCs between baseline session (top) and session with walls removed (bottom). Left EBC has a more distal receptive field and exhibits similar egocentric boundary vector tuning. Right EBC has a more proximal receptive field and has disrupted tuning in arena with no walls. (H) For EBCs recorded in arenas without walls, the preferred orientation at baseline plotted against the change in EBC receptive field coherence between the two sessions is shown. (I) Same as (H) but for change in coherence as a function of baseline preferred distance.

3.2.10 EBC responsivity is anchored to boundaries, not the center of the environment

RSC EBCs exhibited egocentric vector sensitivity to both arena boundaries and the center of the environment, which we used to our advantage in GLM analyses. This occurs because arena boundaries have a fixed relationship relative to the center of the environment. Accordingly, an obvious question is whether the egocentric boundary response is in actuality defined as an egocentric vector to the center of the arena. We addressed this possibility by comparing preferred orientation and distance for 13 RSC EBCs (from four rats across 11 sessions) between baseline arenas and open fields expanded up to 1.75 m^2 (Fig. 3.4E). If EBC responses

were anchored to boundaries, we anticipated that the orientation and preferred distance would remain consistent across both conditions. Conversely, if the receptive field was defined by a vector to the center of the arena, then the distance component of the egocentric boundary vector would remain fixed to this point. In this scenario, the preferred distance would either move away from the animal in expanded arenas or potentially scale with the arena expansion. We observed that the preferred orientation, preferred distance, and ratemap coherence were not altered between baseline and expanded field sessions, confirming that EBCs were anchored to boundaries and not the center of the arena (Fig. 3.4, B to D, difference between baseline and expanded preferred orientation = 6° , IQR = 15.75 to 32.25°; Wilcoxon signed-rank test, signed-rank = 55.5, $P = 0.51$; difference between baseline and expanded preferred distance = 3 cm, IQR = 5.0 to 3.5 cm; Wilcoxon signed-rank test, signed-rank = 30, $P = 0.51$; difference between baseline and rotated ratemap coherence = 0.01, IQR = 0.04 to 0.01; Wilcoxon signed-rank test, signed-rank = 31, $P = 0.34$).

3.2.11 EBC responsivity is stable in novel environments

Neurons within the broader neural spatial circuitry such as grid cells, head direction cells, and place cells exhibit consistent, albeit remapped, spatial receptive fields in novel environments. We next questioned whether egocentric boundary vectortuned neurons of RSC would exhibit similar stability in their selectivity. We recorded 17 RSC cells including 8 EBCs in familiar then novel environment sessions (Fig. 3.4F, four rats across five sessions). Neither distance nor orientation components of the egocentric boundary vector were altered in the novel environment relative to baseline, illustrating that EBCs are not experience dependent and

do not remap between environments (Fig. 3.4, B and C, difference between baseline and novel preferred orientation = 9° , IQR = 12 to 52.5° ; Wilcoxon signed-rank test, signed-rank = 23.5, $P = 0.48$; difference between baseline and novel preferred distance = 1.5 cm, IQR = 5 to 0 cm; Wilcoxon signed-rank test, signed-rank = 3, $P = 0.31$). Coherence of EBC receptive fields was unchanged between environments, providing evidence that the resolution of the egocentric location signal was robust in both familiar and novel arenas (Fig. 3.4D; difference between baseline and novel ratemap coherence = 0.01, IQR = 0.02 to 0.06; Wilcoxon signed-rank test, signed-rank = 22, $P = 0.64$).

3.2.12 Stability of EBC subpopulations requires physical boundaries

Sensory information originating from multiple modalities likely underlies the egocentric nature of the RSC boundary vector responses. There are two reasons to believe that somatosensation may inform the preferred orientation and distance of a subset of EBCs. First, many RSC neurons with egocentric boundary vector tuning had preferred distances that were proximal to the animal and within or near whisker range (the preferred distance was less than 10 cm for a subset of 24.7% of EBCs; Fig. 3.1K). Second, the preferred orientation of EBCs spanned all egocentric bearing angles but was biased laterally (in a contralateral manner), perhaps reflecting whisker interaction with borders (Fig. 3.1J). As such, we questioned whether the presence of a physical boundary was required for EBC spatial tuning and/or particular subsets of EBC receptive fields. To this end, baseline sessions were compared to recordings in environments that were bordered by drop-offs with no arena walls ($n = 35$ neurons from seven sessions across three rats). Twenty-five neurons recorded under these conditions exhibited EBC sensitivity in the baseline session

(Fig. 3.4G). EBCs detected in the baseline session had similar preferred orientations but more distal preferred distances in sessions with no physical walls (Fig. 3.4, B and C, difference between baseline and no walls preferred orientation = 0° , IQR = 63.75 to 27.75° ; Wilcoxon signed-rank test, $z = 1.02$, $P = 0.31$; difference between baseline and no walls preferred distance = 10.0 cm, IQR = 25.75 to 4.5 cm; Wilcoxon signed-rank test, $z = 3.86$, $P = 1.12 \times 10^{-4}$). In addition, the overall coherence of the egocentric receptive field was significantly decreased in the absence of physical walls and fewer EBCs were detected in these sessions (Fig. 3.4D; difference between baseline and no walls ratemap coherence = 0.16 , IQR = 0.05 to 0.22 ; Wilcoxon signed-rank test, $z = 3.94$, $P = 8.1 \times 10^{-5}$; EBCs with no walls = 48.6% , $n = 17/35$; EBCs with walls = 68.6% , $n = 24/35$). Collectively, these results suggest that the EBC population signal is degraded in the absence of explicit borders. Despite this fact, numerous EBCs sustained their preferred egocentric vector across conditions. To investigate why some neurons were disrupted and not others, we next examined the difference in receptive field coherence as a function of baseline preferred orientation and distance. There was no relationship between the baseline preferred orientation of the neuron and the magnitude of degradation of the spatial signal with no physical walls (Fig. 3.4H, circular-linear correlation, $r = 0.17$, $P = 0.71$). In contrast, the more proximal the egocentric boundary receptive field was to the animal at baseline, the more decreased the tuning was in an arena with no physical walls (Fig. 3.4I, Spearman's correlation, $r = 0.50$, $P = 0.01$). These results support the idea that the subset of animal-proximal EBCs (Fig. 3.1E) may rely on somatosensory interaction with borders, while EBCs with more animal-distal receptive fields (Fig. 3.1F) are preserved in environments with no physical walls because they rely on other sensory modalities.

3.2.13 RSC EBCs are insensitive to environmental geometry, which yields a directional representation of environment shape

Boundaries are unique environmental features in that they both restrict navigational affordances and define the spatial structure of the broader environment. Accordingly, the presence of boundary-sensitive neurons within RSC indicates that the region is capable of detecting features of environmental geometry. In a square open field like the one used for baseline experimental sessions, there are two primary defining features of environmental geometry: (i) conjunctions of walls forming 90° corners and (ii) boundaries that are orientated along two axes of allocentric environmental directions. As such, we questioned whether EBCs would maintain their preferred tuning in circular environments that excluded both of these geometric features. We recorded 23 RSC EBCs as animals free foraged in square and circular environments across two experimental sessions each day (Fig. 3.5A, total RSC neurons recorded under these conditions = 32 across four rats and 10 sessions). As with most other environmental manipulations, EBC boundary vectors were unchanged when the geometry of the environment was altered (Fig. 3.5B, difference between square and circle preferred orientation = 9.0°, IQR = 28.5 to 15.75°; Wilcoxon signed-rank test, $z = 1.14$, $P = 0.25$; difference between square and circle preferred distance = 3.0 cm, IQR = 8.0 to 2.75 cm; Wilcoxon signed-rank test, $z = 1.60$, $P = 0.11$; difference between square and circle ratemap coherence = 0.01, IQR = 0.03 to 0.03; Wilcoxon signed-rank test, $z = 0.76$, $P = 0.44$).

A notable feature of many EBCs (Fig. 3.5A, left), but not all (Fig. 3.5A, right), was the structure of head direction tuning between square and circular environments. EBCs would typically have four-pronged directional tuning that aligned with the orientation of the walls in square environments (Fig. 3.5A, left top). How-

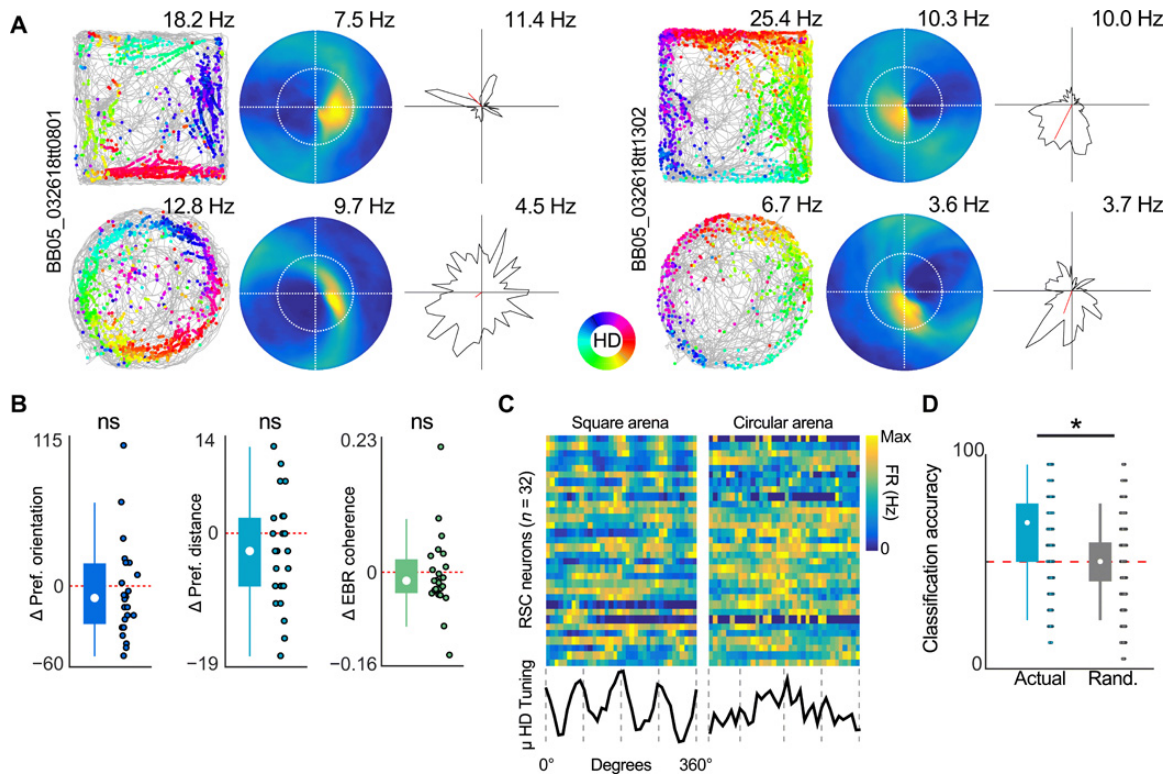


Figure 3.5: RSC EBCs are insensitive to environmental geometry, which generates a directional representation of environment shape. (A) Trajectory plots, EBRs, and head direction tuning plots for two example RSC EBCs for experimental sessions in a square (top) and circular environment (bottom). (B) Preferred orientation, preferred distance, and EBC receptive field coherence from recording sessions in the circular arena subtracted from the corresponding metrics in square arena sessions. (C) Head direction tuning plots for all RSC neurons (in the same order) in the square arena (left) and circular arena (right). Color depicts intensity of activation (blue is zero firing rate; yellow is maximum firing rate). Bottom in black: The average head direction tuning across the full population of RSC neurons for the square and circular environments. Gray dashed lines depict 90° axes. (D) Arena classification accuracy for linear discriminant analysis on head direction tuning from (C). Teal, actual classification; gray, classification after randomizing arena identity. Red dashed line is statistical chance.

ever, as a consequence of consistent egocentric boundary vector tuning in environments of different shapes, EBC tuning yielded directionally uniform tuning in circular environments (Fig. 3.5A, left bottom). Figure 5C depicts head direction tuning plots for the full population of RSC neurons recorded in square and circular arenas. When the mean population head direction tuning was examined, distinct peaks fell at the four cardinal directions in square arenas, but no such peaks were observed in their circular counterparts (Fig. 3.5C, bottom plots). We hypothesized that differences between directional tuning, as a consequence of the presence of EBCs, would allow downstream regions to disambiguate environments of different geometries. To test this, we trained a linear classifier on a random 80% of the directional tuning curves from both environments and attempted to predict which environment the other 20% of head direction tuning curves were recorded within (Fig. 3.5D, linear discriminant classifier, $n = 100,000$ iterations). Consistent with the hypothesis that geometry could be decoded from a population with EBC tuning, the arena could be identified correctly with 67.7% accuracy (IQR = 50 to 75%), which was statistically significant from both statistical chance (Wilcoxon signed-rank with 50% accuracy median, $z = 218.2$, $P = 0$) and a classifier ran with arena identity randomized (randomized arena identity = 50%; IQR = 41.7 to 58.3%; Wilcoxon rank sum test, $z = 188.4$, $P = 0$). We conclude that regions having egocentric boundary vector tuning may provide punctate directional signals to downstream regions such as the medial entorhinal cortex that can be compared to other directional inputs to inform circuits about environment geometry.

3.2.14 A subpopulation of RSC EBCs are theta-modulated

In building off of geometry detection in RSC EBC ensembles, a natural next question is how might these egocentric positional signals be integrated within the broader spatial circuitry. Previous work has demonstrated that RSC local-field potentials feature a prominent theta oscillation during active movement that is strongly coherent with theta rhythms observed in the dorsal hippocampal formation (Alexander et al., 2018; Borst & Anderson, 2017). Spatial representations in regions with strong theta rhythmicity, such as medial entorhinal cortex or hippocampus, are strongly influenced by boundaries and environmental geometry (Muller & Kubie, 1987; O'Keefe & Burgess, 1996; Keinath et al., 2017; Kinsky et al., 2018; Krupic et al., 2015). We next questioned whether RSC neurons exhibiting egocentric boundary vector sensitivity were potentially synchronized with these areas via theta oscillations. Consistent with previous work, we observed a strong RSC theta oscillation and that individual RSC neurons engaged with the theta oscillation (Fig. 3.6, A and B). A small subpopulation of RSC neurons exhibited theta rhythmic spiking (as revealed by autocorrelations of their spike trains) and were phase-locked to the locally recorded theta oscillation (Fig. 3.6B; 4%, $n = 22/555$). A larger subset of RSC neurons did not have detectable theta rhythmic spiking but were significantly phase-locked to the theta oscillation (Fig. 3.6C; 27.6%, $n = 153/555$, see Materials and Methods).

Virtually no EBCs exhibited intrinsically theta rhythmic spiking (0.08%, $n = 1/134$), but 25.4% of EBCs ($n = 34/134$) were phase-locked to RSC theta oscillations (Fig. 3.6D). The strength of theta modulation (MRL) was similar for theta-modulated non-EBCs and EBCs (Fig. 3.6E, non-EBC MRL = 0.13, IQR 0.10 to 0.18; EBC MRL = 0.11, IQR = 0.09 to 0.14, Wilcoxon rank sum test, $z = 1.59$, $P = 0.11$).

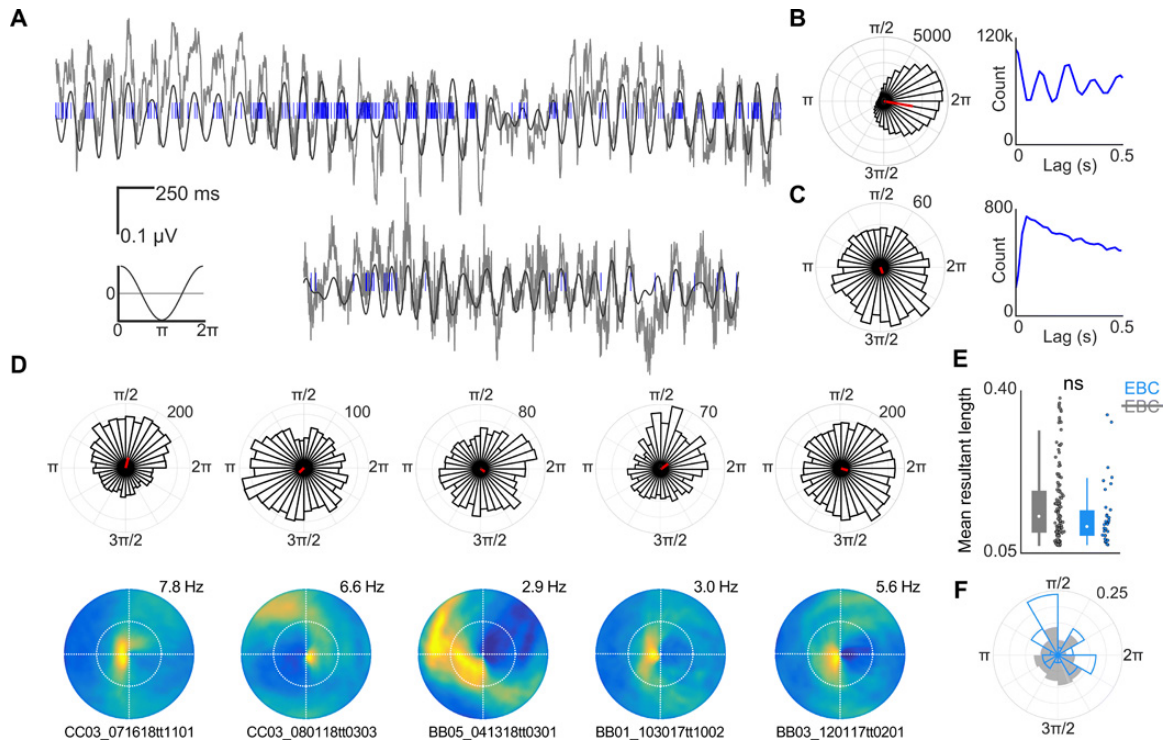


Figure 3.6: A subset of RSC EBCs is theta-modulated. (A) Two examples of RSC theta oscillation (gray) and spike train of simultaneously recorded neurons (blue). Bottom left: Scale bar and schematic depicting correspondence between oscillation and theta phase. (B) Left: Circular histogram depicting spike counts as a function of theta phase for the neuron in the top row of (A). Density of spikes near 2π indicates that the neuron is locked to the peak of the theta phase. Right: Spike train autocorrelogram for the same neuron shows theta rhythmic spiking. (C) Same as in (B) but for the neuron depicted in the bottom row of (A). This neuron is significantly theta phase-modulated but does not exhibit theta rhythmic spiking. (D) Example theta phase-modulated EBCs. Top row: Circular histogram of spike counts versus theta phase. Bottom row: Corresponding EBRs. (E) Strength of theta phase modulation as measured by the MRL for non-EBCs (gray) and EBCs (blue). EBCs have theta modulation significantly similar to non-EBCs with significant phase relationships. (F) Preferred theta phase for all EBCs (blue) and non-EBCs (gray). EBCs tended to prefer the falling phase of the theta oscillation, while non-EBCs preferred the rising phase.

Non-EBC theta-locked RSC neurons were biased to firing during the rising phase of the theta rhythm, whereas theta-locked EBCs preferred the falling phase (Fig. 3.6F, non-EBCs phase = 3.55 rad, IQR = 1.94 to 4.87 rad; EBCs phase = 2.31 rad, IQR = 1.64 to 5.07 rad; Watson-Williams test, $F = 45.4$, $P = 3.14 \times 10^{-10}$). These results confirm that a subpopulation of RSC EBCs are phase-locked to theta oscillations present in RSC, consistent with recent modeling work suggesting periodic modulation as a mechanism for comparing current sensory input about the environment against stored spatial representations (Bicanski & Burgess, 2018).

3.3 DISCUSSION

3.3.1 RSC spatial representations facilitate reference frame transformations

The current data support and extend the functional role of RSC in reference frame transformations. Specifically, the RSC population exhibits sensitivity to multiple spatial coordinate systems, an essential characteristic of circuitry capable of generating such translations. In the current work, we report a large subset of spatially reliable neurons that encode the position of boundaries in egocentric coordinates. Referred to as EBCs, these neurons robustly encoded a vectorial representation of the distance and orientation of any boundary relative to the animal itself (i.e., in an egocentric reference frame) (Gofman et al., 2019; Hinman et al., 2019; LaChance et al., 2019; Wang et al., 2018b). Egocentric boundary representations are predicted to form a critical component of the coordinate transformation circuit, as the response property could function to inform the broader spatial circuitry about the position of external landmarks in a viewpoint-dependent manner (Bicanski & Burgess, 2018; Byrne et al., 2007). RSC neurons also exhibited multiple forms of allocentric modulation. Nearly half of RSC neurons exhibited reliable and

spatially anchored responses during free foraging behavior. Spatially stable cells had complex 2D spatial representations that, in some cases, were reminiscent or possibly descended from spatial non-grid cells observed in medial entorhinal cortex (Diehl et al., 2017), allocentric boundary vector cells and axis-tuned neurons of dorsal subiculum (Hartley et al., 2000; Lever et al., 2009; Olson et al., 2017), and/or location-modulated head direction cells of post-subiculum (Peyrache et al., 2017). A second form of allocentric response was observed in a subset of RSC neurons that exhibit allocentric head direction sensitivity. These forms of allocentric spatial information may be processed or compared with egocentric boundary vector information within theta time scales. Both subsets of neurons exhibited theta phase modulation, which is well known to synchronize information processing throughout the broader allocentric spatial circuit. When paired with the unique anatomical connectivity of RSC with both egocentric and allocentric processing regions, the presence of neurons, such as EBCs, that are sensitive to one or more spatial coordinate systems signifies that the region is capable of interrelating external and internal spatial information for the initial construction and use of stored spatial representations. This fact may explain the diversity of impairments observed in spatial navigation, learning, and memory that occur following damage, lesion, or inactivation of the area (Vann et al., 2009).

3.3.2 The RSC egocentric boundary vector code is context independent, which generates a directional code that reflects environment geometry

EBC spatial receptive fields were activated when the animal was positioned with both a specific orientation and distance from an environmental boundary. EBCs maintained their preferred vector tuning preference in rotated, expanded, and

novel arenas (Fig. 3.4). Accordingly, the EBC signal does not remap across environments, thus providing a stable, context invariant, positional metric. This stability can be contrasted to the vast majority of allocentric representations, such as place cells, grid cells, or head direction cells, that are known to either show global or rate remapping, translations, or rotations between environments (Muller & Kubie, 1987; Fyhn et al., 2007; Yoganarasimha et al., 2006). In contrast, border cells of medial entorhinal cortex and boundary vector cells of dorsal subiculum maintain similar tuning preferences in a context invariant manner analogous to that observed in the EBCs shown here (Solstad et al., 2008; Lever et al., 2009). It remains to be seen what interactions exist between cells having these different types of boundary-anchored receptive fields; however, the current data suggest that boundary-sensitive neurons may provide a foundational map upon which other spatial representations can be situated. Like border and boundary vector cells, RSC EBC vector representations did not remap in environments of different geometries (Solstad et al., 2008; Lever et al., 2009). However, because EBCs respond in a directionally dependent manner along every environmental border, the mean directional tuning of the RSC population reflected the shape of the environment (Fig. 3.5). Here, we demonstrated that this directional signal could be used to distinguish arena shape. Relative positions of boundaries have repeatedly been shown to alter or anchor allocentric spatial representations, especially in medial entorhinal grid and hippocampal place cells ((Muller & Kubie, 1987; O'Keefe & Burgess, 1996; Kinsky et al., 2018; Krupic et al., 2015; Barry et al., 2007). Medial entorhinal cortex receives excitatory projections, both directly and indirectly, from RSC and projects into hippocampus. We hypothesize that the RSC arena geometry-related directional signal may serve to provide excitatory drive at specific allocentric

head directions to inform the circuit about the relative angles among borders. The invariance of EBC response fields and their relationship to environmental geometry may also support or drive route-referenced spatial representations observed in RSC and parietal cortex, which are anchored to space as defined by the route itself (Alexander & Nitz, 2015; Miller et al., 2019; Nitz, 2006). These route-centric activity patterns are often modulated by the geometric shape of the trajectory, which could be partially explained by EBCs responding to specific environmental features such as corners, repeating egocentric views, or some conjunction of EBC responsivity and other spatial covariates (Whitlock et al., 2012; Alexander & Nitz, 2017; Nitz, 2006). Accordingly, future work is needed to assess the relationship between EBC sensitivity in the open field and responses of RSC and parietal cortex in linearized environments.

3.3.3 EBCs are primarily restricted to the dysgranular RSC

A notable anatomical feature of the EBC population was that it was primarily localized to the dysgranular subregion of RSC. Dysgranular RSC has connectivity weighted toward egocentric coordinate systems, as it is reciprocally innervated by cortical regions important for processing sensory and motor information as well as association areas such as parietal cortex wherein egocentrically referenced spatial responses have been observed (9, 13, 14, 20). Further, the concentration of EBCs in dysgranular RSC is consistent with theoretical work proposing a circuit for translating between egocentric and allocentric coordinate systems that includes parietal cortex, RSC, and the extended hippocampal formation as primary hubs (Bicanski & Burgess, 2018; Byrne et al., 2007; Clark et al., 2018; Rounds et al., 2018). Of note, dysgranular RSC was shown by Jacob et al. to have bidirectional head di-

rection cells that respond to local reference frames in multicompartment environments with distinct contextual cues (Jacob et al., 2017). This sensitivity ultimately yields allocentric directional tuning plots that are bimodal. In the current work, strongly tuned EBCs commonly exhibited quad-modal allocentric directional tuning that was aligned with the four walls of square environments. This similarity in directional tuning response of EBCs and bidirectional head direction cells and their colocalization in dysgranular RSC raises questions as to the nature of the relationship or interactions between these functional subpopulations. One possibility is that neurons in dysgranular RSC are prone to represent the locations of spatial landmarks using egocentric vectors and that EBCs and bidirectional head direction cells are both special cases constrained by their respective experimental setups. In the case of EBCs reported here, the vector may anchor to boundaries because borders are the only landmarks present in the open field that can cause activation of the receptive field. In the work of Jacob et al., the egocentric vector may respond to borders as well as local visual landmarks or doorways between two compartments. The bimodal directional tuning in the latter experiment may arise from constrained egocentric sampling along two axes as a consequence of the multicompartment environment segmenting two opposing walls. This proposed egocentric vector encoding of environment features in RSC may underlie functional correlates of local heading orientation, scene processing, or goal location in RSC in humans (Chrastil et al., 2015; Vass & Epstein, 2013; Marchette et al., 2014).

3.3.4 A network of vector-based egocentric spatial representation

In addition to RSC, EBCs were also observed in both parietal and secondary motor cortices (M2) but not medial entorhinal cortex, which is commonly thought to

represent space in allocentric coordinates. The presence of EBCs in parietal cortex converges nicely with previous work, demonstrating egocentric bearing sensitivity of neurons within the region to visual cues positioned along boundaries (Wilber et al., 2014). Computational models exploring circuitry for reference frame transformations and spatial imagery initially predicted EBCs to exist in parietal cortex (Byrne et al., 2007). However, egocentric responses were initially reported in lateral entorhinal cortex, dorsal striatum, and postrhinal cortex (Gofman et al., 2019; Hinman et al., 2019; Wang et al., 2018a) and now, here in RSC, parietal cortex and M2. Accordingly, a picture of a distributed network of interconnected regions with egocentric vector representations is beginning to emerge. Given the presence of EBCs in several midline structures, it is possible that EBCs are also present in the anterior cingulate cortex as well as thalamic structures that innervate midline associative cortex. In the current report, we observe lateralization of the preferred orientation of EBC receptive fields that is contralateral with the hemisphere that the neuron is within. This may offer clues as to the origin of the egocentric signal. Specifically, it suggests that EBCs are likely to be driven directly by contralaterally projecting thalamic afferents or early cortical processing regions, which also have this form of lateralization. Regardless, future investigations should focus on dependencies among the regions currently implicated, as the EBC network may have functional and anatomical connectivity resembling the well-characterized extended head direction cell network.

3.4 MATERIALS AND METHODS

3.4.1 Subjects

Male Long-Evans rats ($n = 7$; Charles River Laboratories, Wilmington, MA) were housed individually in plexiglass cages and kept on a 12-hour light/dark cycle. Rats had continuous access to food during a habituation period lasting approximately 1 week. After this period, animals were food-restricted until they reached 85 to 95% of their weight during free feeding. Water was available continuously. All procedures were approved by the Institutional Animal Care and Use Committee at Boston University.

3.4.2 Data Analysis

3.4.2.1 *2D spatial ratemaps and spatial stability*

Animal positional occupancy within an open field was discretized into 3 cm \times 3 cm spatial bins. For each neuron, the raw firing rate for each spatial bin was calculated by dividing the number of spikes that occurred in a given bin by the amount of time the animal occupied that bin. Raw firing ratemaps were smoothed with a 2D Gaussian kernel spanning 3 cm to generate final ratemaps for visualization. Individual raw firing ratemaps were also computed after dividing the session into halves. To assess spatial stability of an individual RSC neuron, the similarity of the two raw firing ratemaps from nonoverlapping halves of the recording session was calculated using the nonparametric Spearman's rank correlation coefficient. To determine whether a given spatial stability value was greater than expected by chance, we next conducted randomization tests wherein the spike train for each RSC neuron was circularly shifted relative to spatial position 100 times, and indi-

vidual firing ratemaps were constructed for nonoverlapping halves that were then correlated. The spatial stability correlation values following randomizations were collapsed into a single distribution for all neurons and randomizations, and the 99th percentile of all values was calculated. RSC neurons with real spatial stability correlations greater than this threshold were determined to have robust 2D spatial stability.

3.4.2.2 *Construction of EBRs*

EBRs were computed in a similar manner as 2D spatial ratemaps but referenced relative to the animal rather than the spatial environment. The position of the boundaries relative to the animal was calculated for each position sample (i.e., frame). For each frame, we found the distance, in 2.5-cm bins, between arena boundaries and angles radiating from 0° to 360° in 3° bins relative to the rats position. Critically, angular bins were referenced to the head direction or movement direction of the animal such that 0°/360° was always directly in front of the animal, 90° to its left, 180° directly behind it, and 270° to its right. Intersections between each angle and environmental boundaries were only considered if the distance to intersection was less than or equal to one-half the length of the most distant possible boundary (in most cases, this threshold was set at 62.5 cm or half the width of the arena). In any frame, the animal occupied a specific distance and angle relative to multiple locations along the arena boundaries, and accordingly, for each frame, the presence of multiple boundary locations was added to multiple 3°x2.5 cm bins in the egocentric boundary occupancy map. The same process was completed with the locations of individual spikes from each neuron, and an EBR was constructed by dividing the number of spikes in each 3°x2.5 cm bin by the amount of time that

bin was occupied in seconds. Smoothed EBRs were calculated by convolving each raw EBR with a 2D Gaussian kernel (5 bin width, 5 bin SD). For EBR construction and other analyses in the current work, head direction was the instantaneous angle calculated from the location of two position tracking diodes, and movement direction was defined as the instantaneous derivative of the position signal. Head direction was used as the primary directional variable in EBR construction, but a comparison to movement direction determined the latter to be a less robust signal for egocentric boundary vector tuning. However, some neurons showed significant egocentric boundary tuning in relationship to movement direction rather than head direction.

3.4.2.3 *Head direction cell identification*

For each neuron, the MRL of the firing rate as a function of head direction was calculated as

$$R_m = \frac{\cos(\bar{\theta}) \sum_{i=1}^n F_i \cos(\theta_i) + \sin(\bar{\theta}) \sum_{i=1}^n F_i \sin(\theta_i)}{\sum_{i=1}^n F_i} \quad (3.1)$$

where $\bar{\theta}$ was the head direction of firing and F_i and θ_i were the firing rate and head direction for bin i . Head direction cells were identified as those cells with R_m greater than 0.20. Head direction cells ($n = 27/555$; 4.9% of all RSC neurons) were removed from the possible pool of RSC EBCs.

3.4.2.4 *Identification of neurons with egocentric boundary vector tuning*

To identify neurons with significant egocentric boundary vector sensitivity, we began by calculating the mean resultant (MR) of the cells egocentric boundary directional firing collapsed across distance to the boundary. The mean resultant was

calculated as

$$MR = \left(\sum_{\theta=1}^n \sum_{D=1}^m F_{\theta,D} * e^{i*\theta} \right) / (n * m) \quad (3.2)$$

where θ is the orientation relative to the rat, D is the distance from the rat, $F_{\theta,D}$ is the firing rate in a given orientation-by-distance bin, n is the number of orientation bins, m is the number of distance bins, e is the Euler number, and i is the imaginary constant. The MRL is defined as the absolute value of the mean resultant and characterized the strength of egocentric bearing tuning to environment boundaries. We next computed the preferred orientation of the EBR as the mean resultant angle (MRA)

$$MRA = \arctan2\left(\frac{\text{imag}(MR)}{\text{real}(MR)}\right) \quad (3.3)$$

We estimated the preferred distance by fitting a Weibull distribution to the firing rate vector corresponding to the MRA and finding the distance bin with the maximum firing rate. The MRL, MRA, and preferred distance were calculated for each neuron for the two halves of the experimental session independently. Next, the MRL was computed for each neuron following 100 random, unrestricted, circular shifts of the spike train relative to position. The 99th percentile of the MRL distribution across all neurons was determined.

A neuron was characterized as having egocentric boundary vector tuning (i.e., an EBC) if it reached the following criteria: (i) the MRL from the baseline session was greater than the 99th percentile of the randomized distribution, (ii) the absolute circular distance in preferred angle between the first and second halves of the baseline session was less than 45° , and (iii) the change in preferred distance for both the first and second halves relative to the full session was less than 50%.

To refine our estimate of the preferred orientation and preferred distance of each neuron, we calculated the center of mass of the receptive field defined after thresholding the entire EBR at 75% of the peak firing and finding the largest continuous contour (contour in MATLAB). We repeated the same process for the inverse EBR for all cells to identify both an excitatory and inhibitory receptive field and corresponding preferred orientation and distance for each neuron.

3.4.2.5 *Ratemap coherence, dispersion, and receptive field size*

For either EBRs or 2D spatial ratemaps, coherence was defined as the Spearman's correlation between each spatial bin and the mean firing rate of all adjacent bins. Dispersion was calculated as the mean within ratemap distance of the top 10% of firing rate bins. Receptive field size was only calculated for EBRs (described below) and was defined by the area (percentage of all EBR degree \times cm bins) of the largest single contour detected after selecting for bins with firing rates greater than 75% of the maximum firing rate.

3.4.2.6 *Self-motion ratemaps and assessment of self-motion sensitivity*

Angular displacement (θ) was calculated by determining the circular difference in movement direction between two position samples (frames) separated by a 100-ms temporal window. The total distance (d) traveled between these two frames was also calculated. The process was repeated for the full recording by sliding a 100-ms temporal window across all position frames and calculating these values. Angular displacement and distance traveled were converted to Cartesian coordinates to generate x- and y-displacement values in centimeters, which corresponded to lateral and longitudinal displacements for each frame across the full recording. 2D

displacements were binned (1 cm) and convolved with a 2D Gaussian spanning 3 cm. For each neuron, the same process was repeated for displacement values that co-occurred with spike times to generate a spike occupancy map as a function of displacement. Self-motion ratemaps were constructed by dividing the spike occupancy map for each neuron by the total time in each displacement bin. Bins occupied for less than 267 ms were removed from analyses, as they typically were observed at extreme displacement values. Self-motion ratemaps for each neuron were additionally constructed independently from interleaved, nonoverlapping, 1-s periods throughout the entire session to assess stability of self-motion tuning. For quantification of self-motion tuning relative to a randomized distribution, all aforementioned ratemaps were generated for each neuron 100 times after randomly shifting the spike train relative to position. Self-motion ratemaps were quantified for their stability, left- versus right-turning preference (LvR), and speed modulation. First, stability of self-motion tuning was quantified by correlating self-motion ratemaps generated from nonoverlapping periods for each neuron (Fig. 3.2B). Next, turning biases for clockwise versus counterclockwise movements (LvR FR ratio, Fig. 3.2C) were quantified by computing the ratio of summated firing for all similarly occupied displacement bins on the right and left sides of the zero vertical line, respectively. Last, speed modulation was approximated by finding the mean firing rate as a function of longitudinal displacement (i.e., averaging over columns of self-motion ratemaps) and correlating with the y-displacement value (Fig. 3.2D). All self-motion ratemap quantification was repeated for all randomized self-motion ratemaps as mentioned above. Stability and turn-bias quantification were computed from displacement bins that were occupied in both self-motion ratemaps or both halves of an individual self-motion ratemap. All quan-

tification was completed on nonsmoothed self-motion ratemaps.

3.4.2.7 Generalized linear models

To more directly test to what degree neurons represented egocentric compared to allocentric variables, we adopted a GLM framework. The probability of spiking in a given behavioral frame (33 Hz) is described by an inhomogeneous Poisson process, where the probability of spiking in a given frame is described by the time varying variable λ

$$\begin{aligned}
 P(\text{spike}|t) &= e^{-\lambda(t)} \\
 \lambda(t) &= \lambda_{FR} * \lambda(t)_{\text{SelfMovement}} * \lambda(t)_{\text{Egocentric}} \\
 \lambda_{FR} &= \beta_0 \\
 \lambda_{\text{SelfMovement}}(t) &= v_1 S + v_2 A \\
 \lambda_{\text{Allocentric}}(t) &= \rho_1 x + \rho_2 y + \rho_3 x^2 + \rho_4 y^2 + \rho_5 xy + \rho_6 \cos(\phi) + \rho_7 \sin(\phi) \\
 \lambda_{\text{Egocentric}}(t) &= \epsilon_1 d + \epsilon_2 d^2 + \epsilon_3 \sin(\theta) + \epsilon_4 \cos(\theta) + \epsilon_5 d * \sin(\theta) + \epsilon_6 d * \cos(\theta)
 \end{aligned} \tag{3.4}$$

where β_θ defines the baseline firing rate of the neuron. All subscripted variables are fit coefficients weighting the other (time-varying) variables. S is the running speed of the animal, and A is the angular displacement of the animal, as described above. x and y are measurements of the animals position in the environment in pixels, and ϕ is the head direction. Last, d is the animals distance from the center of the environment, and θ is the egocentric angle to the center of the environment. Coefficients were determined by fitting to maximize log-likelihood (MATLAB function glmfit) of the experimental spike train given the behavioral variables. For statistical tests, some numbers of the coefficients were set to zero, giving a log-likelihood

for the reduced model. The difference in likelihood for the full versus reduced model was compared to a chi-square distribution (degrees of freedom equal to the number of coefficients set to zero) to generate an analytic P value. While theoretically the change in log-likelihood should follow a chi-square distribution, this is only the case when the spike train has been fit very well (e.g., including all neuron-neuron coupling terms). In line with previous approaches, we therefore also compared the change in log-likelihood in two additional manners. First, we compared the change in log-likelihood to that from 1000 randomly shuffled spike train, giving an empirical null distribution. Second, when comparing the relative effects of two variables (that is, comparing two reduced models to each other), we can derive the dAIC for each of the reduced models and compare their magnitudes. Representative spike trains for each model were generated by evaluating lambda for each behavioral time point (`glmeval` in MATLAB) and using this as the input to a random Poisson generator (`poissrnd` in MATLAB).

3.4.2.8 *Classification of arena shape condition from directional tuning*

Head direction tuning curves were calculated for each neuron ($n = 32$) by normalizing the number of spikes occurring in a given 10° directional bin by the amount of time that bin was occupied. Head direction tuning curves were independently calculated for sessions in square and circular environments and normalized by peak firing rate. Directional tuning curves were labeled with the arena condition in which they were sampled and classified. Tuning curves were down-sampled to include alternating directional bins (every 20° , to ensure the covariance matrix would be positive definite) and then classified using linear discriminant analysis 100,000 times for arena condition with cross-validation (80% train, 20% test;

classify in MATLAB) and uniform priors. Classification accuracy was assessed by finding the percentage of correct labels on the test dataset. A randomized distribution of classification accuracy was calculated in parallel by randomizing the true arena condition label of the training set.

3.4.2.9 *Assessment of theta phase modulation*

For each experimental session, an local field potential (LFP) channel was identified that was qualitatively noise free. The LFP signal was filtered in the theta frequency range (6 to 10 Hz), and the phase for each spike from each neuron was estimated as the instantaneous phase angle of the corresponding Hilbert transform (`hilbert` in MATLAB). For each neuron, the MRL and MRA were calculated on the full spike phase distribution using the circular statistics toolbox (MRL, `circr`; MRA, `circmean`). We next randomly shifted the spike train relative to theta phase 100 times for each neuron to generate a null distribution of MRL values. RSC neurons with MRLs greater than the 95th percentile of the full distribution of randomized MRL values were determined to be theta phaselocked.

3.4.2.10 *Assessment of theta rhythmic spiking*

Spike train autocorrelograms were estimated by generating a histogram of temporal lags between spikes in a 400-ms temporal window discretized into 20-ms bins. For each neuron, the power spectrum of the autocorrelogram was computed using the Fourier transform (`fft` in MATLAB), and the peak in the theta frequency range was identified (if it existed). If the mean power within 1 Hz of this theta peak was 50% greater than the mean power for the full power spectrum, the neuron was determined to exhibit intrinsic theta spiking.

3.4.2.11 *Von Mises mixture models*

Distribution of preferred orientation estimates was modeled as mixtures of Von Mises distributions using orders from 1 to 5. Optimal models were identified as the model increasing model fit over the one-component model with the next complex model yielding less improvement. Models were cross-validated using randomly selected halves of the preferred orientation distribution. Theta of each Von Mises component is reported, and a distribution function of the optimal model was generated to visualize mixture model fit.

CHAPTER 4

Allocentric Representation and Transformation Learning from Egocentric Sensory Prediction

4.1 ABSTRACT

Spatial navigation is a fundamental ability that is conserved across a wide range of species. It involves the formation of coherent representations of one's location in map-like space. Despite a plethora of neurophysiological experiments revealing spatially tuned neurons across the mammalian neocortex and subcortical structures such as the hippocampus, it remains largely unclear how such representations are acquired in the absence of explicit allocentric inputs. Theoretical neuroscience and machine learning studies have proposed that a type of self-supervised learning may be responsible for extracting the latent structure of navigational space by contrasting sensory stimuli over time.

In this study, we employ a modular simulation of neocortex/hippocampus that integrates feedforward pathways onto pyramidal cell soma with feedback pathways onto distal dendritic arborizations. Our learning rule aims to minimize the temporal difference in feedforward signal changes by modifying feedforward synapses to align with feedback activity, thereby implementing a form of contrastive and associative learning. We organize these modules into the major cyclic connectivity pattern of the visual cortex, the retrosplenial cortex, the postrhinal cortex, the hippocampus and the entorhinal cortex, and train during an open field foraging-like task in which the network receives sensory signals and self-motion signals.

After training, we compare the overall tuning activity of units in the trained

model to known functional phenotypes such as egocentric boundary cells (that respond to sensory-driven egocentric bearing and distance of boundaries), as well as allocentric head direction and place codes. Moreover, we make a series of modifications on the environment after learning. Given that the feedback and feedforward signals in our network have distinct physiological origins, these modifications result in incongruities in sensory signals and top-down expectations, which create unique signatures of network dynamics that can be investigated in future experimental studies.

4.2 INTRODUCTION

A sense of allocentric, or map like space, is essential for navigating to goals which are outside of the immediate sensory space. The simplest approach to path planning is known as dead reckoning, or path integration (McNaughton et al., 2006). In this approach, given an initial starting location, an agent begins to move, and afferent motor and sensory information, such as running speed (Hinman et al., 2016; Kropff et al., 2015), rotational velocity (Alexander & Nitz, 2015), and optic flow (Raudies & Hasselmo, 2012) is integrated over time. Integrating these afferent signals over time yields a net displacement, and could be used as a homing signal back to the origin location (Mittelstaedt, 1980). However, such an approach is extremely sensitive in the presence of noise (Hardcastle et al., 2015), would not persist across multiple experiences, and do not provide a full representation of space for more difficult navigation involving obstacles. In contrast, animal studies have shown that animals will mentally simulate planned trajectories in a combination of forward and backward search, similar to the A* algorithm used in artificial route planning (Ferbinteanu & Shapiro, 2003; Erdem & Hasselmo, 2012). However,

such planning requires that there first be an internal representation of location in space and that such a representation can be indexed by sensory information.

Egocentric Representations While view-invariant allocentric codes are required for longer-term navigation, animals rarely, if ever, receive direct allocentric information. Instead, the sensory information is encoded in the neocortex, where the location of stimuli will be relative to an animal, such as an object being distant and ahead in the visual field, or an object stimulating whiskers on the left side of an animal. Therefore, even if path integration alone is insufficient for formation of stable global maps, some degree of processing of sensory and motor information must occur, creating a transformation from the egocentric reference frame to a global map-like space. This transformation must also have an inverse transformation, so that allocentric codes can inform egocentrically centered behavior, though such transformations may occur in other brain regions (Hinman et al., 2019). Neocortical regions are known to encode the location and identity of objects relative to an animal, beginning with primary sensory regions and convergent to association areas such as the parietal cortex. There have recently been numerous reports of neurons that respond to environmental boundaries in the postrhinal cortex (LaChance et al., 2019) and retrosplenial cortex (Alexander et al., 2019), as well as other egocentric responses, such as the relative location of a visual target (Alexander et al., 2022) in posterior parietal cortex. The majority of these regions also receive afferent self-movement information and feedback information from the allocentric regions discussed below. Therefore, these regions are thought to be essential to transforming sensory information into allocentric codes.

Allocentric Representations Extensive experimental evidence has shown that the medial temporal lobe (MTL) and associated structures in mammals contain allocentric codes for space. Place cells, found primarily in the hippocampus (O'Keefe & Dostrovsky, 1971), respond to a single location (or a few locations) in an environment but will remap between environments, creating a reusable fabric for allocentric codes (Muller & Kubie, 1987). Grid cells in the medial entorhinal cortex (Colgin et al., 2009) respond in a regular array of locations in the environment and project bidirectionally to place cells. They may form a basis for performing mental operations between place cell representations (Erdem & Hasselmo, 2012; Bush et al., 2014). Head direction cells that respond to the current compass direction of the head are found in many of these regions Taube (2007). These MTL regions receive connections from the processed egocentric signals discussed above, but it is not clear how they directly form from those signals.

Learning Allocentric & Egocentric Predictions Temporal prediction of spatial signals is known to create different latent representations depending on the specific input and target signals. Computational models that learn to encode future sensory states from current sensory states and motor information directly lead to allocentric representations, but do not form the same representations in an auto-encoder approach without a predictive component (Recanatesi et al., 2021). These findings suggest that the prediction of sensory inputs extracts the transitions of sensory information in a reduced two-dimensional latent space, similar to successor representation (Stachenfeld et al., 2017). Other work has shown that given past allocentric locations and motor commands and tasked to predict future allocentric locations, in essentially a form of path integration, hidden units will form even more compressed representations in the form of regular grid-like activity (Banino

et al., 2018). More recent work has also shown that simple RNN models will exhibit a representational state in which principal components code separately for spatial and contextual information (Low et al., 2023), similar to splitter cells (Ferbinteanu & Shapiro, 2003) or remapping (Ocko et al., 2018). However, allocentric predictions of the medial temporal lobe are also responsible for guiding learning in the neocortex, suggesting that stable map-like representations may guide sensory driven representations in uncertain environments (Doron et al., 2019). These previous results show that in abstract networks temporal prediction of egocentric sensory signals gives rise to allocentric codes, and prediction of allocentric codes can lead to more compressed allocentric codes.

Current Approach: Learned Egocentric-Allocentric Associations While the computational studies discussed above have shown the learning of allocentric spatial representations from temporal prediction of sensory representations, it remains unclear how such a representation would be learned in a biological substrate. Here we make use of the connectivity structure of a previous systems-level model of spatial cognition (Bicanski & Burgess, 2018; Byrne et al., 2007) (the ‘BBB models’ below). The BBB models hand-coded several functional cell types such as egocentric responses and place cells, and showed how they could be trained to drive each other in either a sensory driven or memory driven manner. We utilize the same overall connectivity of this previous model, but instead of pre-determining the functional cell types and then training the weights, we implement a simplified model of learning. We replace each of the abstract regions from the BBB model with a biophysically inspired model of canonical microcircuit (chapter 2), coupled with a learning rule which alters feedforward weights in order to predict future states.

4.3 METHODS

4.3.1 General Approach

We utilize a virtual agent moving in a 1m x 1m open field, sampled at 500Hz (dt=2ms) in order to give appropriate temporal resolution for the neural modeling described below. The agent moves according to a parametrically driven approach described elsewhere (Raudies & Hasselmo, 2012; George et al., 2022) with equations and parameters chosen (Table 4.1) to realistically replicate the behavior of rats moving in an environment for 5 hours. Briefly, on each frame the rotational velocity is modeled as an OrnsteinUhlenbeck (OU) process and running speed is modeled as a Rayleigh transformed OU process. The agent location and head direction are then generated by integrating these terms over time. Near boundary locations there was an anti-parallel repulsive force proportional to the distance from the wall, as well as a perpendicular force. Together, these two forces create a slowing down as the agent approaches a wall and causes a thigmotaxis-like following of the boundaries. We then utilize a three-dimensional rendering software (Chevalier-Boisvert, 2018) to generate agent-based visual scenes in colour and depth (RGB-D). In this rendering each wall is given a different colour in order to provide some grounding information as to the structure of the environment, and could be a stand in for sensory cues.

Variable	Description (Units)	Range x Shape
dt	Simulation timestep (s)	0.001
t	Time (s)	[0, 1200]
X	X location (m)	[0, 1]
Y	Y location (m)	[0, 1]
θ	Head direction (radians)	$[-\pi, \pi]$
S	Running Speed (m/s)	[0, 0.3]
ω	Rotational Velocity (radians/sec)	$[-\pi, \pi]$
I	Agent RGB View (1)	$[0, 1] \times [60, 80, 3]$
D	Agent Depth View (m)	$[0, 1.4] \times [60, 80]$

Table 4.1: Behavioral Variables for Simulation

Macroscale Connectivity The overall connection between modules is arranged to reflect the macroscale connections among brain regions strongly implicated in spatial navigation of rodents (Fig. 4.1). External stimuli come from two major sources, the egocentric sensory (vision) and the motor / proprioceptive pathway. The vision pathway begins with the identity, and distance of the visual field array projecting to a visual modules termed ‘ventral’ and ‘dorsal’ respectively. Similarly, the linear velocity and rotational velocity from the behavioral simulation project to a singular ‘action’ module. The two visual streams and the action module have convergent feedforward pathways onto the retrosplenial cortex, which in turn projects to the perirhinal cortex, and hippocampus, representing the primary feedforward visual to allocentric pathway. The entorhinal module also receives direct feedforward activity from the ‘action’ module and projects to the feedforward compartment of the hippocampal module. Feedback connections travel in the opposite order and along the same route. Thus, the dorsal, ventral, and action modules are the only ones which directly relate to the environment and provide a basis for interpretation later.

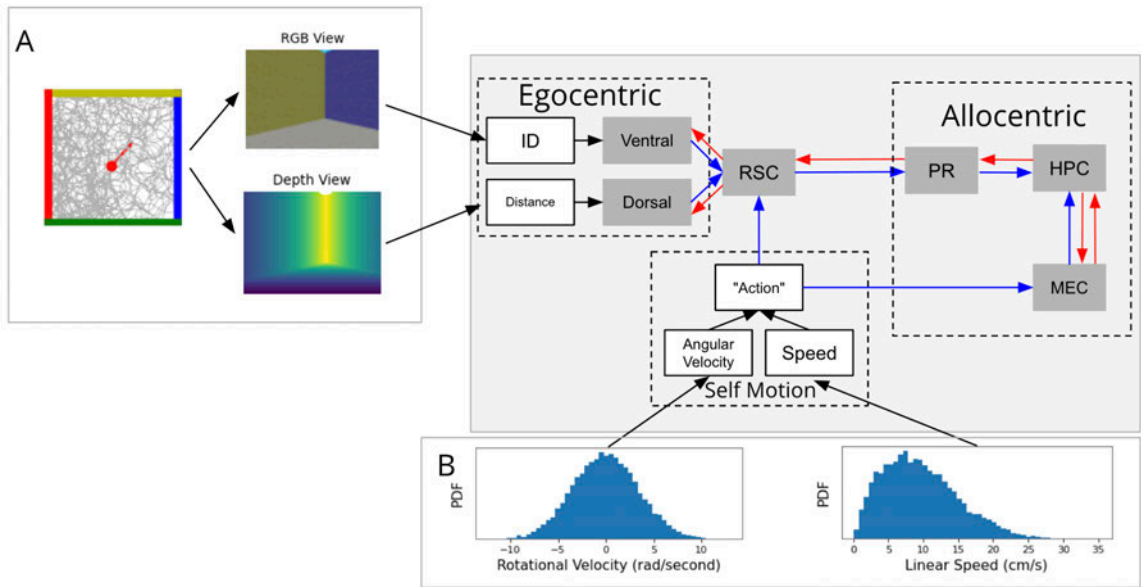


Figure 4.1: Overall simulation paradigm. **A** (Left) As an agent (red dot, facing towards arrow) navigates around an environment (grey line), behavioral frames are generated. (Top) RGB images are encoded with a one-hot paradigm to provide object identity. (Bottom) Depth images are also generated, representing early processing along the ventral visual stream. **B** Behaviorally realistic rotational velocity (left) and running speed (right) distributions are used to generate the overall trajectory from A. These variables are separately encoded on each frame into a self-motion representation. **C** Show the macroscale connectivity between modules. Black lines indicate static weights which encode information from the behavioral simulation. Blue lines indicate inter-area feedforward synapses which target onto the granular layer of neocortex or dentate gyrus populations and carry overall activation. Red lines indicate feedback weights which originate from the infragranular layers of neocortex or the CA1 region of hippocampus and terminate on distal dendrites, carrying burst rates.

4.3.2 Neural Dynamics

Microcircuit Connectivity The model that we utilize here relies on separate feed-forward and feedback pathways, both in the neocortex and hippocampus. Large scale summaries of cortical microcircuit connectivity (Felleman & Van Essen, 1991; Haeusler & Maass, 2007) suggests that in neocortex there is a primarily three-layer separation of granular, superficial, and infragranular layers, which are organized into a primarily feedforward (granular to superficial) and feedback (infragranular to distal dendrites) pathways. Hippocampus, with a distinctly different architecture consists of a different nomenclature. We model hippocampus as dentate gyrus, consisting of granule cells, and CA3 and CA1 consisting of pyramidal cells. DG projects to CA3 somatic regions through the mossy fibers which then projects to CA1 through Schaffer Collaterals (O'Mara, 2005). CA3 and CA1 also receive input to distal dendrites from entorhinal cortex, and exhibit burst transfer functions similar to that observed in pyramidal neurons of neocortex in response to feedback (Takahashi & Magee, 2009; Naud et al., 2023). The internal connectivity pattern of hippocampus is primarily the same as in neocortex. The primary difference that we address here is that hippocampal formation receives feedforward signal into the dentate gyrus from both the entorhinal cortex and perirhinal cortex, but also the distal dendrites by the temporammonic pathway, creating an associative loop that is not present in neocortex (Anderson et al., 2007).

Within each simulated 'region' there exist three distinct populations of 128 continuous activation units, each of which reflects the mean-field activity of a population of neurons which lie in different lamina. In the cortical regions, these regions correspond to the granular ('G') layer stellate neurons, superficial ('S') pyramidal neurons, and infragranular ('I') pyramidal neurons. In the hippocampal module,

these instead correspond to dentate gyrus, CA3, and CA1 respectively.

Following prior work, granular neurons follow the general recurrent neural network (RNN) formulation reflecting their compact physiology:

$$\begin{aligned}\tau \frac{dv_L}{dt} &= -v_L(t) + \sum_{N \in A} W_{NL} R_N(t) \\ r_L(t) &= \sigma(v_L(t))\end{aligned}\tag{4.1}$$

Where τ is the population specific time constant (10ms throughout), v represents the membrane potential. W represents the weights of afferent connections from presynaptic populations A (in the summation index), including recurrent connections among a population. For the case of the granular neurons, A is the set of all external inputs and feedforward afferents (black lines in figure 4.1). σ is the non-linear activation function converting membrane potential v to action potential r , and we utilize the sigmoidal activation function.

The pyramidal units contain the same leaky-integrator format for their somatic potential, but incorporate an additional term representing the potential of the distal dendrites and overall burst probability:

$$\begin{aligned}\tau \frac{dv_L}{dt} &= -v_L + \sigma(u_L/D_d) + \sum_{N \in A} W_{NL} r_N \\ r_L &= \sigma(v_L) \\ \tau_u \frac{du_L}{dt} &= -u_L + 2 \frac{dr_L}{dt} + \sum_{N \in B} W_{Nu} b_N \\ p &= \sigma(u) \\ b_L &= r_L \odot p_L\end{aligned}\tag{4.2}$$

The additional term u represents the distal dendritic compartment, and the sum-

mation index B represents the population of neurons projecting to that compartment, compared to those projecting to the soma (summation index A). For superficial units A is the recurrent connections and activity of the local granular units. For infragranular units, A is the recurrent connections and activity of the local superficial units. For both sets of pyramidal neurons, B is the set of feedback connections (red lines in figure 4.1). p represents the probability of a burst event and follows a sigmoidal function of the distal dendrites (Naud & Sprekeler, 2018). b represents the burst rate of the pyramidal neurons and is the direct product of the overall firing rate and the burst probability. The middle terms in both the somatic potential and dendritic potential equations represent subthreshold coupling and backpropagating action potentials (BAP) respectively, and allow mixing of feedforward and feedback signals, and increasing the burst rate (b) when these inputs are coincident in time.

4.3.3 Learning Rule

Consistent with previous studies (Payeur et al., 2021) and the methods reported in Section 2 we implement a learning rule on the feedforward synapses of pyramidal units. The learning rule combines a classical “pre-post” term with a top-down gain based on the difference in instantaneous and long term burst probabilities:

$$\Delta W_{XY} = \eta((p_Y - \bar{p}_Y) \odot r_Y) \otimes r_X \quad (4.3)$$

Where X is the presynaptic population and Y is the postsynaptic, and \bar{p}_Y represents the low-pass burst probability on the temporal order of behavioral variables (equation 2.18, $\tau_{avg} = 1s$). The learning rate η is held at 10^{-4} throughout.

In the implementation here we utilize the numerically evaluated values for p_Y

and \bar{p}_y . However, for the purposes of this chapter, we provide a relationship between equation 2.6 and 4.3. In the case of a smoothly varying input the difference between instantaneous burst probability and low pass average is approximately the temporal derivative. Then, letting x and y denote individual units for notational simplicity, and following the chain rule for the derivative of p_Y :

$$\begin{aligned}
\Delta W_{xy} &= \eta((p_y - \bar{p}_y)r_y)r_x \\
&= (\eta r_x r_y)(p_y - \bar{p}_y) \\
&\approx (\eta r_x r_y) \frac{dp}{dt} \\
&\approx (\eta r_x r_y) \frac{d\sigma(u)}{du} \frac{du(t)}{dt} \\
&\approx (\eta r_x r_y) (\sigma(u)(1 - \sigma(u)) \frac{1}{\tau_u} (-u + 2 \frac{dr_y}{dt} + \sum Wb)
\end{aligned} \tag{4.4}$$

Evaluating at a given value of u :

$$\Delta W_{xy}(u) \propto (\eta r_x r_y) \frac{dr_y}{dt} F(u) \tag{4.5}$$

From this formulation the instantaneous change in feedforward synaptic weight is proportional to the presynaptic activity times the postsynaptic activity (Hebbian associative learning) times the instantaneous change in unit rate (temporal error) and multiplied by the distal dendritic potential (feedback gain). Thus, the learning rule given in 4.3 contains the same terms as those in 2.6, with the additional presynaptic activity term and non-linearity on u . Weights are initialized using random distributions, identical to those in table 2.2.

4.3.4 Tuning Curves

When attempting to classify the receptive fields of units we utilize a general framework:

$$R(t) = \sigma(\beta_0 + \beta_{pc}\lambda_{PC}(t) + \beta_{hd}\lambda_{HD}(t) + \beta_{ls}\lambda_{LS}(t) + \beta_{ebc}\lambda_{EBC}(t) + \beta_{grid}\lambda_{grid}(t)) \quad (4.6)$$

Where $R(t)$ is the activation of a given unit at a single point in time and σ is the non-linear activation function. In this equation, PC=place cell, HD=head direction cell, LS=running speed cell, EBC=egocentric boundary cell and grid=grid cell. β_0 Is the baseline bias in activity, and the remaining β terms are strengths of behavioral variable influence on activation. The λ terms represent parametric tuning curves that match the general shapes of experimental findings and are individually optimized for each unit:

$$\begin{aligned} \lambda_{pc} &= e^{-\frac{1}{2(1-\rho^2)}\left(\left(\frac{x-\mu_x}{\sigma_x}\right)^2 - 2\rho\left(\frac{x-\mu_x}{\sigma_x}\frac{y-\mu_y}{\sigma_y}\right) + \left(\frac{y-\mu_y}{\sigma_y}\right)^2\right)} \\ \lambda_{hd} &= (\cos(\theta - \mu_{hd}))^k \\ \lambda_{ls} &= \lfloor \beta_s s + (1 - \beta_s)e^{-\frac{(s-\mu_s)^2}{2\sigma_s^2}} \rfloor \\ \lambda_{ebc} &= \max(I \odot |_{l \in L, r \in R} e^{-\left(\left(\frac{l-\mu_l}{\sigma_l}\right)^2 + \left(\frac{r-\mu_r}{\sigma_r}\right)^2\right)}) \\ \lambda_{grid} &= \frac{1}{3} \left[\cos\left(2\pi \frac{x e_\theta}{\Lambda} + \phi\right) + \cos\left(2\pi \frac{x e_{\theta+\pi/3}}{\Lambda} + \phi\right) + \cos\left(2\pi \frac{x e_{\theta+2\pi/3}}{\Lambda} + \phi\right) \right] \end{aligned} \quad (4.7)$$

Within each subequation behavioral variables are as in table 4.1 and valid ranges for tuning parameters and their description is in table 4.2. The place-cell term is a two dimensional Gaussian centered at any point in the environment. The head direction model is a cosine function with mean μ_{hd} and tuning parameter k which decreases with width of the tuning curve. Running speed tuning is modeled as the sum of a linear slope and a superimposed Gaussian function, enabling a di-

verse range of speed-curves including linear increase, saturating responses, and maximum response models (Alexander et al., 2022). The egocentric boundary cell predictor is a two dimensional Gaussian displaced from the current location of the agent by μ_{ex} and μ_{ey} relative to the direction of the animal. I is an indicator function evaluated at 0.1cm increments in the x and y direction relative to the animal for the presence of a boundary, and then rotated to the current movement direction. The second term of the EBC response is a 2D Gaussian intensity function evaluated at Rostral-Caudal (r) and Lateral-Medial (l) offsets from the agent location. The grid cell model is the sum of three spatial sinusoids with frequency Λ , spaced 60 degrees apart and overall orientation to the environment θ (Burgess et al., 2007). Finally, we do not explicitly provide an equation for boundary vector cells (BVCs), since a place field with unconstrained location and covariance is degenerate with the typical equations used to model BVCs. In the results however when a unit is significantly modulated by the ‘place cell’ predictor but is elongated along a boundary we identify it as a BVC instead of a place cell.

Parameter	Description	Range
$\sigma_{x/y}$	X/Y width of place cell (cm)	
ρ	X/Y Covariance (angle) of place cell (1)	
$\mu_{x/y}$	X/Y center of place cell (cm)	
k	Tuning width of HD cell (1)	[1/3, 3]
μ_{hd}	Center of HD cell (rad)	$[-\pi, \pi]$
β_{s_1}	Slope of speed cell (s/cm)	[0, 30]
β_{s_2}	Strength of Gaussian speed cell (s/cm)	[0, 30]
L	Lateral offset from the agent (cm)	[-50, 50]
R	Rostral offset from the agent (cm)	[-50, 50]
$\mu_{R/L}$	R/L center of egocentric boundary cell (cm)	[-50, 50]
θ	Angle offset of grid cells (rad)	$[0, \pi/3]$
ϕ	Phase offset of grid cells (rad)	$[-\pi, \pi]$
Λ	Spatial scale of grid cells (cm)	[20, 50]

Table 4.2: Tuning curve parameters fitting or explicitly setting functional cell types

Measuring Representation Learning After every simulated three minute interval ('epoch') we evaluate the tuning of all learned units. This is done by, without enabling learning, simulating the network on a separately generated 60 minute session and then down-sampling the behavior and unit activity to 10Hz (36K frames). Down-sampling is performed for computational efficiency of the tuning curve fitting procedure. The activity from this 'test' session is then fit to the explicit tuning curve given in equation 4.6 by gradient descent (SGD, minimum error over 10 random initializations of 5-fold verification), as well as submodels that each drop one of the predictors. β values were constrained to be non-negative during the fitting process, to match experimental methods in which place, grid, and head direction cells are positively modulated compared to the mean firing rate. The format of the λ_{ls} predictor allows negatively modulated speed cells, as shown in experimental findings (Hinman et al., 2016; Dannenberg et al., 2019; Alexander et al., 2022). The non-negativity constraint precludes any explicit 'negative-EBCs' (Hinman et al., 2016), which are instead be detected as having a wide and distal receptive field. For each sub model we compare the mean-square error in predicted activity compared to the true activity and generate an F-statistic, which is analogous to the log-likelihood approaches from Chapter 3 but now for a continuous activation function instead of spike trains. The proportion of units that are statistically significantly tuned ($p < 0.001$) is measured on each epoch and is reported below. In figures 4.2 and 4.3 all tuning curves are represented as normalized activity in 2D (spatial), polar (head direction), or 1D (running speed) bins.

4.4 RESULTS

4.4.1 Learned Responses

We first investigate the final somatic event rate tuning curves and burst-rate tuning curves, as summarized in Figure 4.2. Speed tuning was present in all regions, with a total of 87% (890/1024) units being significantly modulated by running speed. Consistent with expected findings, the feedforward signals of the RSC module contains a number (80/128, 63%) of units which respond to the egocentric boundary predictor. These can be recognized in figure 4.2 by the strong responses around the edges of the environment. Another 41 (32%) were classified as head-direction sensitive, though notably no units had significant tuning to both head direction and EBC tuning. In the feedback layer of RSC, one unit (unit 0 in Figure 4.2) showed significant EBC response. Half (64/128) of FB RSC units responded to head direction, and 108 (84%) to movement speed.

Superficial perirhinal units expressed several BVC-like responses, identified as a response to the place cell predictor, but with a variance along one axis at least three times greater than along the other (52/128 41% BVC, 23/128 18% PC). These were largely conjunctive (67/75, 90%) with head direction tuning. Deep PR units generally expressed the same variables as their superficial counterparts, but with a more pronounced tendency for place cell (48/128, 38%) than BVC (12/128, 9%) response.

Hippocampal feedforward units were by large place-cell like (115/128, 89%), with the remaining portion being non-spatially responsive. The majority of units were conjunctive with head direction (102/128, 80%). Deep hippocampal units were tuned to both place (27/128, 21%) and BVC (47/128, 37%), with a slightly

less tendency to be head direction tuned (45/128, 35%). Two deep hippocampal units were significantly tuned to the grid predictor, but upon visual inspection primarily tended to fire at three corners of the environment.

Entorhinal units showed a combination of speed (108/125, 84%) and head direction (42/128, 33%) tuning. A subset of units (18/128, 14%) showed tuning to location and exhibited a mixture of near-boundary and near-center responses.

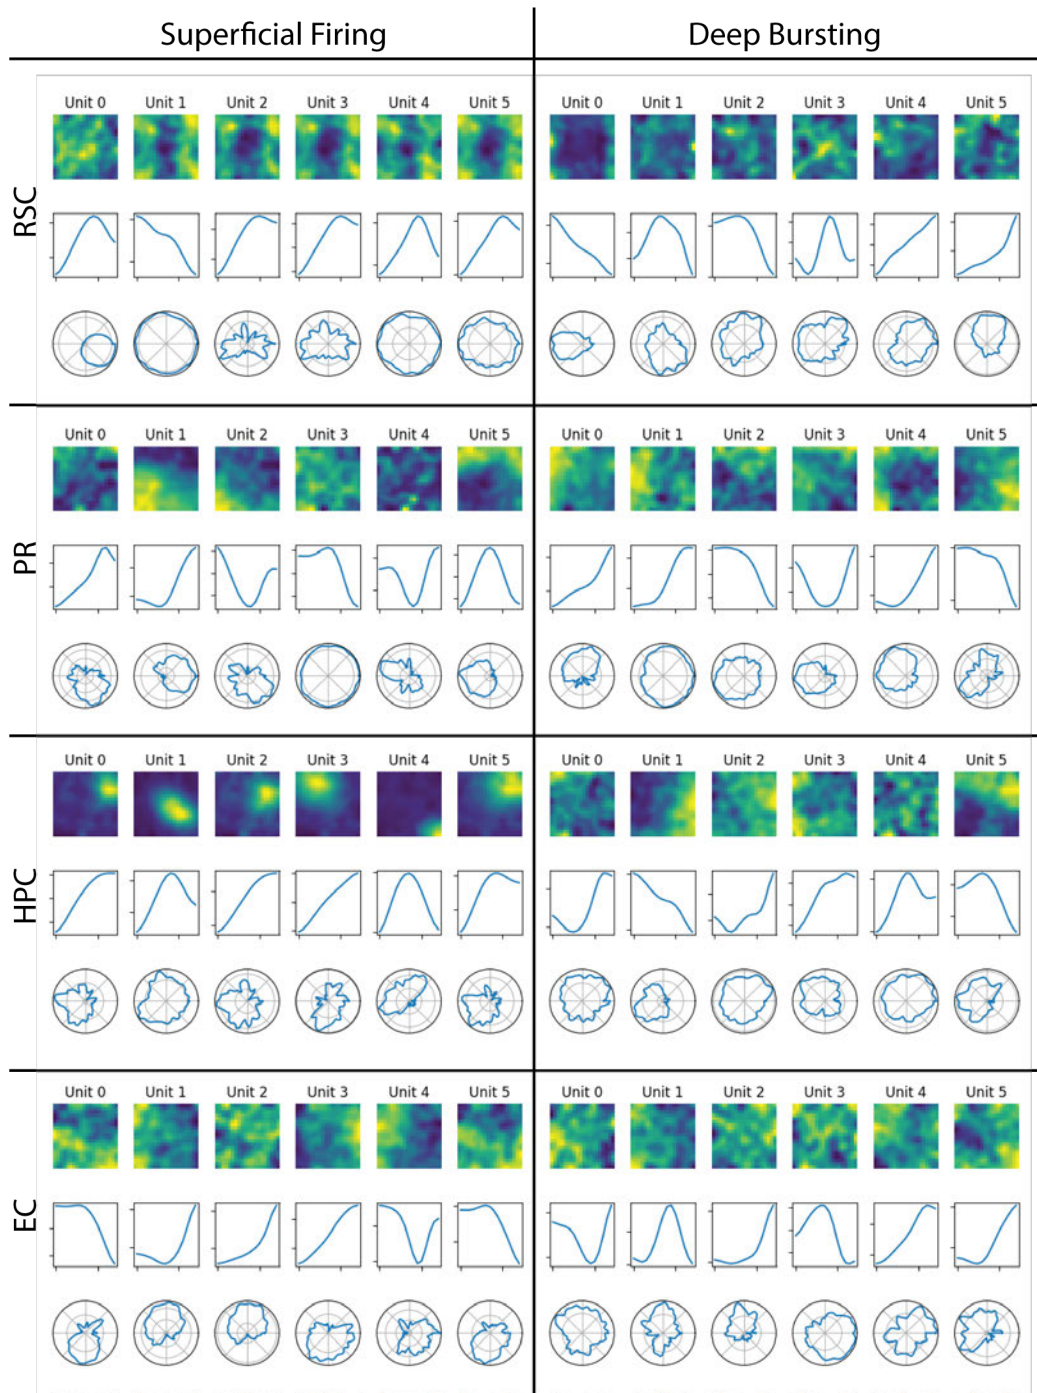


Figure 4.2: Caption on next page

For each box the feedforward (left) or feedback (right) activity is shown for six randomly selected units at the end of training. The top row is the occupancy-normalized mean activity, the middle row is mean response at each running speed, and the bottom is a polar directional tuning curve at each movement angle (which is equivalent to head direction in this simulation). **RSC-FF** Many units exhibit egocentric-boundary responses, firing near the boundaries of the environment, regardless of head direction (unit 3). Others respond to head direction (unit 0). Nearly all respond to movement speed. **RSC-FB** Compared to FF, there is relatively little spatial modulation, though head direction and speed tuning remain prominent. **PR-FF** Spatial tuning (top) tended to be strongly locked to one boundary (unit 3) or corner (unit 1). Head direction and speed tuning are present. **PR-FB** Tuning is similar to the FF case. **HPC-FF** (Top) Place-cell like representations are found, with centers distributed throughout the environment. (Bottom) A subset of units are conjunctive with head direction (unit 0), while others are not (unit 1). **HPC-FB** Spatial representation is more PR-like, with receptive fields on the boundaries of the environment and higher dispersion. **EC-FF** Spatial modulation exists for a subset of units, and tends to be anchored to an environmental boundary and conjunctive with head direction. **EC-FB** Feedback spatial activity is dispersed, but head direction tuning is still present.

Order of Learning We also evaluated the rate at which representations converged to their final value over the course of training. On each epoch we found the mean correlation coefficient between the holdout session activity at that point in training and the holdout activity at the end of training (Figure 4.2). Contrary to expectations, it was neither the most (RSC) or least (HPC) driven representations which converged most quickly. Instead the perirhinal, which was predominantly responsible for encoding distance to a singular boundary, converged within the first tens of epochs. This region was then followed by hippocampus and retrosplenial, suggesting that the feedforward weights of perirhinal aligned to their inputs before the hippocampal module received a consistent enough signal to create place-like representations. Feedback bursting signals converged more slowly than any of the feedforward activities (Figure 4.2).




Figure 4.2: Correlation coefficient of feedforward (top) and feedback (bottom) representational similarity on each epoch compared to at the end of training.

4.4.2 Response to Novel Environment

We now consider a case where the feedforward activity violates the trained expectations. At the end of training we introduce an additional session in which an unexpected barrier with a unique identity encoding has been inserted into the middle of the environment ($y=0.48 - 0.52$, $x = 0.25 - 0.75$). For clarity of analysis, we disable learning during these experiments, and measure mean responses across a simulated 60 minute exposure. The tuning responses of units in this condition are demonstrated in Figure 4.3, and fitting of the curves is performed as in the baseline condition, and the chosen example units are the exact ones shown in Figure 4.2.

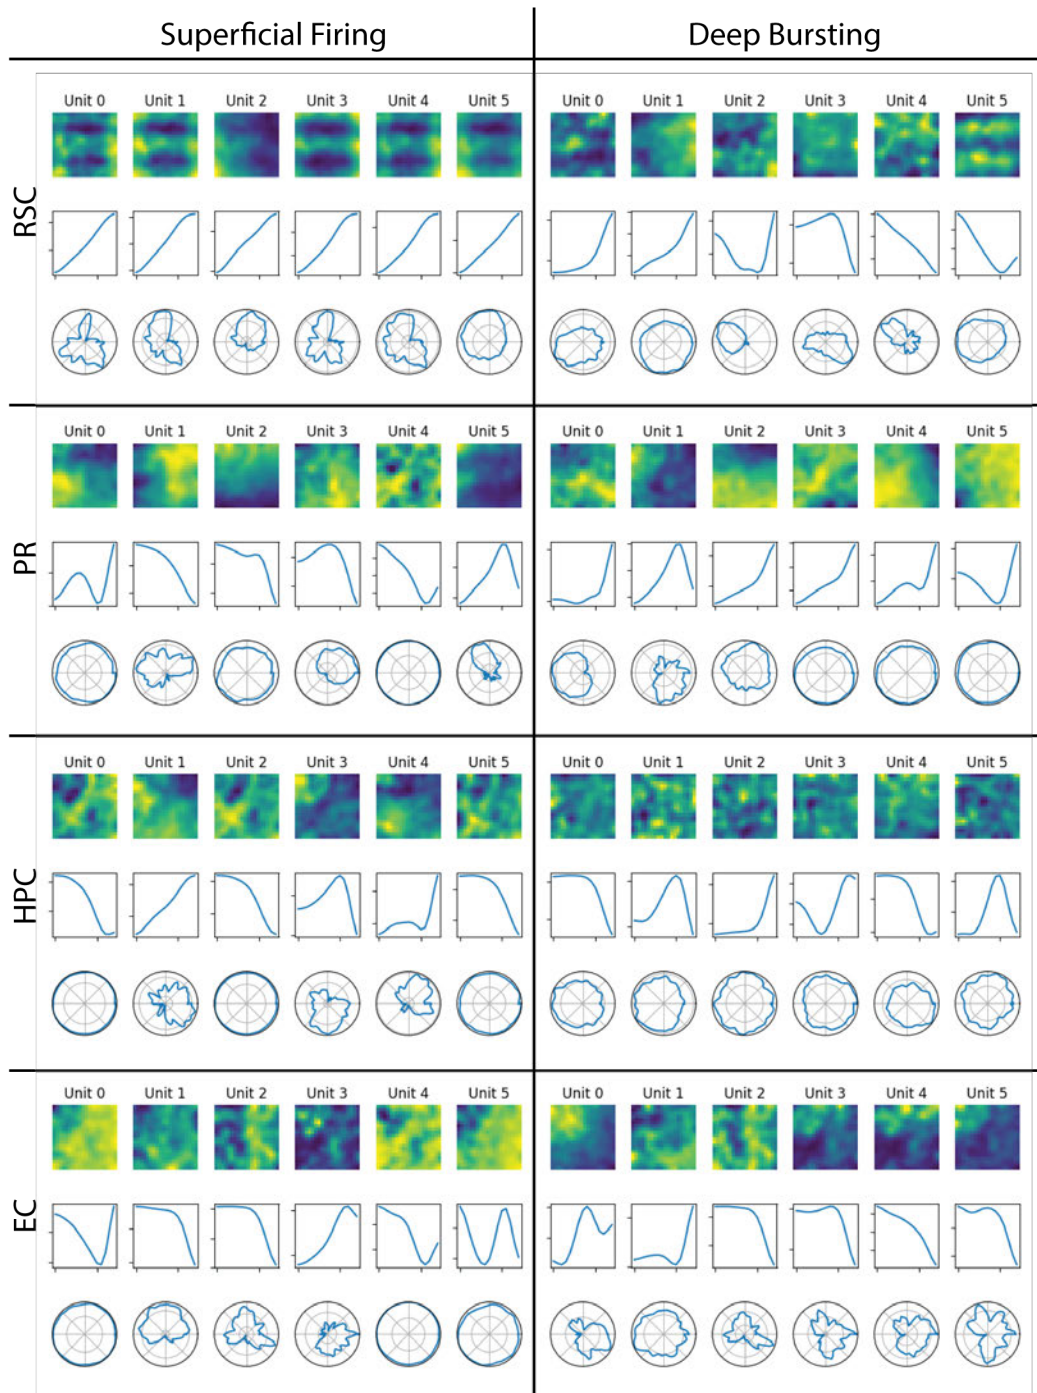


Figure 4.3: Figure on next page

In general, individual feedforward allocentric tuning curves were disrupted in the novel environment, while speed tuning and EBC responses were consistent. In

Novel Environment Tuning **RSC** Many units (eg: FF unit 1) continue to exhibit EBC-like tuning, showing a strong response near the familiar and novel boundary. Head direction tuning is largely disrupted. **PR-FF** A number of PR units continue to show HD tuning. Many spatial tunings have shifted compared to the baseline, but still tend to fire near an edge of the outer boundary. **PR-FB** Deep PR activations are more likely to be tuned to BVC responses than during baseline, but these responses tend to be wide. **HPC** Place-cell like representations are largely disrupted. Units that remain HD tuned tend to have a different preferred direction. **EC** Some units remain tuned to head direction, but spatial modulation is largely gone.

superficial RSC, 75 (93%) of the units which were EBC tuned in the baseline continued to be so in the novel environment, while only 10 (24%) remained modulated by head direction. In superficial PR all units previously classified as BVCs continued to be so, but their preferred boundary and x/y standard deviations changed. Hippocampal place-cell tuning was only present in a fraction of the units from baseline (2/115, 2%). Feedback activity was largely disrupted as much as the corresponding feedforward activity, with the exception of deep RSC where head direction tuning was preserved in 93% (60/64) of the original units. Of the feedback PR activity which was tuned to place-cell predictors in the familiar environment, a significant portion changed to being BVC tuned (20/48, 42%). Across all regions, feedforward and feedback, speed tuning was maintained in 793 of the original 890 (89%).

While Figure 4.3 clearly shows that individual unit representations are disrupted in the novel environment, we next asked whether it was necessarily the case that the information coded by distributed representations was disrupted. To address this, for each region and each behavioral variables (location, head direction, and speed) we train a linear decoder to produce the minimum mean-square error readout both during the familiar and novel environment. For head direction, the decoder was trained on the mean sum error of sin and cosine components of the angle. This was done using increments of ten units (1, 11, 21, ... 128) and five

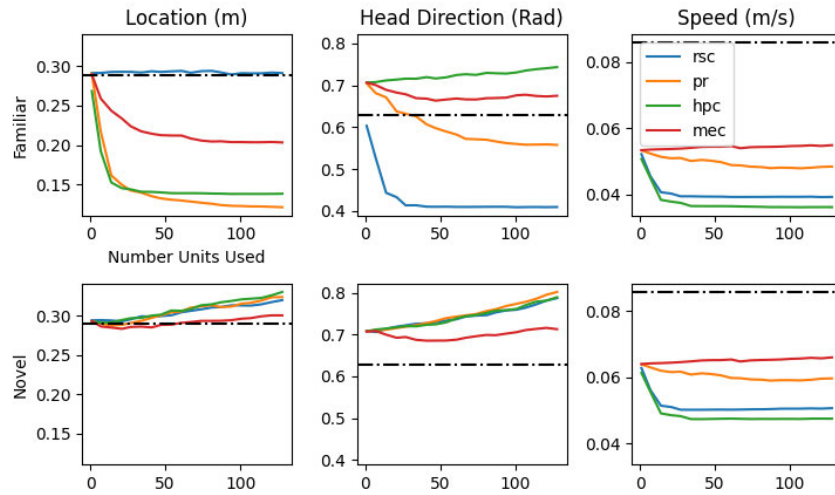


Figure 4.4: *Behavioral Decoding* Mean-square error of linear readout of behavioral variables in each modeled region during foraging in the familiar environment. Dashed black line indicates chance. **Top** Location can be significantly read out in all regions except RSC. Conversely, head direction is most strongly *linearly* read out from RSC and PR. Speed is significantly represented everywhere, but is more redundant in MEC And PR than RSC and HPC. **Bottom** In the novel environment with inserted barrier, the decoding is greatly altered. Using the same error scales as the baseline, location and head direction are not significantly read out from any regions. Speed continues to be represented in all regions, with a slight decrease in performance compared to the familiar environment.

repetitions of 5-fold cross validation. Overall, we find that in the baseline (familiar) condition this distributed readout was able to decode location from all regions except RSC, head direction from PR And RSC, and movement speed from all regions (Figure 4.4, Top). In contrast, position and head direction information was not significantly represented, even in a distributed manner, in any regions during the novel environment (Figure 4.4). Running speed was still coded in all regions, only slightly less prominently than in the familiar session.

4.5 DISCUSSION

Above we introduced preliminary results of a biophysically inspired model of spatial representation learning. Each module receives a combination of feedforward event rates and feedback burst rates which helps to guide a synapse update to minimize the moment-to-moment changes in somatic potential. After simulating this system in a passive navigation setting for an extended period, we see the emergence of internal representations which resemble those found in animal studies. In particular we find the emergence of units similar to egocentric boundary cells (Alexander et al., 2020) which respond to the presence of a barrier based on its egocentric displacement relative to the current direction and position of the agent. As they receive a combination of sensory distance signal ('dorsal'), self-motion information ('action'), and feedback allocentric location ('PR' boundary cells), we hypothesize that these implement a form of transformation from egocentric to allocentric signal (Wilber et al., 2014). The next region in the hierarchy then provides a response based on distance from environmental boundary signal (an allocentric boundary response) that the final region ('HPC') integrates over time to provide allocentric information even when far from the boundary locations. This place-

centric code can then provide additional contextual information to the perirhinal to resolve any ambiguity in feedforward signal, such as two barriers with similar appearance. The general internal representations we find are in line with both experimental findings and theoretical (Bicanski & Burgess, 2018) expectations. In particular, consistent with experimental data, the model shows formation of neural responses that progress from regions showing egocentric representation of barriers (egocentric boundary cells in retrosplenial) to allocentric representation of barriers (boundary vector cells) to allocentric representation of the spatial location of the agent (place cells in hippocampus). An interesting comparison can be drawn between the responses learned here and those learned in Recanatesi et al. (2021), where a single layer RNN trained in a similar manner produces place-like representations but not the types of responses we see in the intermediate RSC and PR modules. This suggests that the restriction of weights between subsets of RNN units, via hierarchy, leads to the intermediate representations being beneficial to overall performance.

When placed in a novel situation (Figure 4.3) the majority of the allocentric codes are diminished, although self-motion and egocentric related signals persist. While allocentric codes do not spontaneously adjust to the novel environment in our model, experimental evidence suggests that remapping may occur within the first few minutes of exposure to a novel environment (Lee et al., 2012), suggesting that some global modulation of learning rate may be necessary to more fully replicate experimental findings. We also saw, compared to experimental settings, a higher portion of units which responded conjunctively to place, head direction, and movement speed. Introducing some regularization constraints on either the weights or overall activity level within a layer may help to address these discrep-

ancies (Clopath et al., 2010).

A notable omission from the current work is that the hippocampal module exhibits all of the same dynamics and learning rules as the cortical modules. While the general DG \rightarrow CA3 \rightarrow CA1 connectivity does match the connectivity pattern of neocortex, the different specifics of short-term facilitation and depression, and the pattern of inhibitory interneurons is omitted from this model. In particular, the proximal FF into CA3 is thought to be facilitating rather than depressive. This means that the input to this region may actually be more responsive to bursting from CA1 than to event rate, of which the later is assumed in the simplified activation model presented here. A more fully biophysical model, such as that in 2, may incorporate these differences. A notable difference from experimental findings is that none of the learned responses in this model were tuned to grid-like activation, with the exception of a small number of feedback units in hippocampus. We expect this is largely due to the absence certain constraints in the behavior of the EC module. Recent computational work has shown that grid units only arise in specific constraints to weight matrices and place field structure (Schaeffer et al., 2022). Additionally, grid cells are thought to arise at least in part due to the presence of certain rhythmic-inducing components such as medial septal input (Brandon et al., 2011) or rebound current (Hasselmo & Shay, 2014). Incorporating one or both of these components in future work could result in grid-like representations.

CHAPTER 5

Discussion

5.1 SUMMARY OF FINDINGS

In the introductory section, I outlined the various aspects of neural models that might be manipulated, and how various biological constraints may guide these changes. In the following chapters I then introduced models which incorporate these changes in order to predict external dynamics (Chapter 2), convert an ego-centric stimulus to a latent two dimensional space (Chapter 4), and express interpretable internal representations.

Chapter 2 introduced several stepwise changes in architecture, which progressively incorporate neural hierarchy, simplified dynamics, and separate feedforward and feedback pathways. While hierarchy and dynamics increased the generalizability and autonomous performance of the overall system, the increased modular complexity did not increase this performance further. However, when combined with a learning rule which specifically utilizes the feedback connection to gate the rate and direction of learning (eq. 2.6) the performance was increased beyond what can be achieved with artificial neural networks and backpropagation through time approaches. Furthermore, because the dynamics of these modules and the corresponding components of the learning rule have concrete inspirations from biology, we were able to introduce a biophysical spiking version of this model, which has similar performance to the rate approximation. As important as the overall performance however, was the finding that the modules in this network strongly encode the temporal derivative of their inputs, suggesting that the predictive learning rule is able to extract some underlying structure of the stim-

ulus. While the tasks used in this study are one or three dimensional location trajectories, scaling up the predictive module to incorporate earlier regions such as visual cortex may allow predictive learning purely from video stimuli.

Chapter 3 demonstrated a recently discovered functional phenotype in the retrosplenial cortex of behaving rats. This region lies at the intersection of sensory driven egocentric representations of cortex and the allocentric representations of medial temporal lobe, and has therefore been postulated to perform a transformation between these coordinate systems. The novel cell type (egocentric boundary cells) may perform the substrate for this transformation, and have been found in other regions implicated in translating two dimensional space to sensory space, such as posterior parietal cortex, motor cortex, and striatum (LaChance et al., 2019; Alexander et al., 2019; Hinman et al., 2019).

Inspired by the success of the previous two chapters, Chapter 4 incorporates the dynamics and learning rules of Chapter 2 into the behavior and macroscale connectivity of Chapter 3, with the goal of investigating the representations learned and comparing them with experimental findings. Overall, the results show the expected trend of earlier regions encoding an egocentric (primarily sensory driven) response, with representations in later regions that are allocentrically informed (such as the response to distance from a single barrier) and purely allocentric representations (place cells). Consistent with experimental findings and expectations, self motion information was encoded in all regions, allowing the overall movement speed to provide a gain on rate of representational change. Also consistent with findings, though not directly in line with the egocentric to allocentric trend, head direction modulation was found in most regions.

5.2 CURRENT WORK

Immediate research directions center on the preliminary findings presented in Chapter 4. The large scale connectivity of this model was based on cognitive neuroscience models, but may require modification in order to more directly align with physiology. Of particular note, the module mediating responses between retrosplenial cortex and hippocampus, termed “perirhinal”, may more appropriately be referred to as subiculum or entorhinal cortex. While there are few studies of perirhinal responses in behaving animals (Burwell et al., 1998), several studies report boundary vector responses, similar to the observed “perirhinal” module responses, in subiculum and entorhinal cortex (Lever et al., 2009). More closely investigating the interconnectivity of these candidate regions with the hippocampus may provide slight modifications in module connectivity, and provide more realistic results in the modeled entorhinal cortex. Additional analyses of low dimensional representations in each module, compared to baseline RNN approaches, may further elucidate how the self-supervised learned representations differ from those learned with backpropagation (Recanatesi et al., 2021; Low et al., 2023).

5.3 FUTURE DIRECTIONS

The overall goal of computational neuroscience is to understand, and possibly utilize, how the dynamics and structure of neural systems enable a wide range of behavior. However, the work presented in chapters 2-4 focuses primarily on representation learning and utilization. While the performance of the network in chapter 2 is extremely high, and chapter 4 provides a concrete tie to a more behaviorally relevant scenario, there are several aspects of cognition beyond repre-

sensation learning that have not been addressed by these models. One clear application of the predictive learning utilized here is the case of control systems, such as motor cortex, where any differences in expected proprioceptive state over time can directly inform improved motor control. Another avenue of research would be how to integrate the columnar dynamics of pyramidal neurons to implement approximations of cognitive attention. The feedback dynamics of chapter 2 are primarily used to guide the learning of feedforward signals, but other researchers have utilized the same feedforward-feedback contrast to implement forms of routing (Bastos et al., 2018), consistent with experimental findings that oscillatory frequencies associated with feedback are implicated in attention (Friese et al., 2016; Richter et al., 2017). While small scale models have been able to incorporate such dynamics for gating of low-dimensional signals, it has proven difficult to scale these models to cognitive tasks as removing biological fidelity is often a requirement for such scaling, and loss of fidelity quickly approaches BP like algorithms. As the burst-dependent learning rule presented in Chapter 2 has shown, incorporating some degree of fidelity can be important for finding new mechanisms of learning. Therefore, careful consideration must be made when scaling biophysical models of attention, so that the essential underlying components are not lost.

5.4 CONCLUSION & BROADER IMPLICATIONS

Many recent computational neuroscience studies have attempted to identify biologically plausible methods of backpropagation of errors (O'Reilly, 2000; Bourdoukan & Denève, 2017; Naud & Sprekeler, 2018; Payeur et al., 2021). While these approaches have had some degree of success, they are still fundamentally biologically implausible and limited in that they rely on supervised learning, target

signals, and explicit errors (though see Sussillo & Abbott (2009); Brea et al. (2016); Lotter et al. (2020)). Our approach here has instead been to take the same biologically plausible learning rules and organize them into a modular design which encourages alignment of feedforward and feedback signals, even when the feedback signals arise from internal dynamics rather than explicit targets. The results presented here indicate this approach can produce highly interpretable models, and may provide a basis for further research into such “internal alignment” approaches to predictive learning.

BIBLIOGRAPHY

- Abbott, L. F. (2008). Theoretical neuroscience rising. *Neuron*, 60(3), 489–95.
- Ablavsky, V., & Sclaroff, S. (2011). Layered Graphical Models for Tracking Partially Occluded Objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(9), 1758–1775.
- Alexander, A., & Nitz, D. a. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature neuroscience*, 18(July), 1–12.
- Alexander, A., & Nitz, D. A. (2017). Spatially Periodic Activation Patterns of Retrosplenial Cortex Encode Route Sub-spaces and Distance Traveled. *Current Biology*, 27(11), 1551–1560.e4.
- Alexander, A. S., Carstensen, L. C., Hinman, J. R., Raudies, F., Chapman, G. W., & Hasselmo, M. E. (2019). Egocentric boundary vector tuning of the retrosplenial cortex. *Science Advances*.
- Alexander, A. S., Rangel, L. M., Tingley, D., & Nitz, D. A. (2018). Neurophysiological signatures of temporal coordination between retrosplenial cortex and the hippocampal formation. *Behavioral Neuroscience*, 132(5), 453–468.
- Alexander, A. S., Robinson, J. C., Dannenberg, H., Kinsky, N. R., Levy, S. J., Mau, W., Chapman, G. W., Sullivan, D. W., & Hasselmo, M. E. (2020). Neurophysiological coding of space and time in the hippocampus, entorhinal cortex, and retrosplenial cortex. *Brain and Neuroscience Advances*, 4, 239821282097287.
- Alexander, A. S., Tung, J. C., Chapman, G. W., Conner, A. M., Shelley, L. E., Hasselmo, M. E., & Nitz, D. A. (2022). Adaptive integration of self-motion and goals in posterior parietal cortex. *Cell Reports*, 38(10), 110504.
- Anderson, P., Morris, R., Amaral, D., Bliss, T., & O'Keefe, J. (Eds.) (2007). *The Hippocampus book*. New York, NY: Oxford University Press, first ed.
- Arolfo, M., & Brioni, D. (1991). Diazepam impairs place learning in the Morris water maze. *Behavioral and Neural Biology*.
- Badre, D., & D'Esposito, M. (2009). Is the rostro-caudal axis of the frontal lobe hierarchical? *Nature reviews. Neuroscience*, 10(9), 659–669.
- Badre, D., & Frank, M. J. (2012). Mechanisms of hierarchical reinforcement learning in cortico-striatal circuits 2: Evidence from fMRI. *Cerebral Cortex*, 22(3), 527–536.

- Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., Pritzel, A., Chadwick, M. J., Degris, T., Modayil, J., Wayne, G., Soyer, H., Viola, F., Zhang, B., Goroshin, R., Rabinowitz, N., Pascanu, R., Beattie, C., Petersen, S., Sadik, A., Gaffney, S., King, H., Kavukcuoglu, K., Hassabis, D., Hadsell, R., & Kumaran, D. (2018). Vector-based navigation using grid-like representations in artificial agents. *Nature*, *557*(7705), 429–433. `Tex.ids= banino_kumaran_2018a`.
- Baras, D., & Meir, R. (2007). Reinforcement learning, spike-time-dependent plasticity, and the BCM rule. *Neural Computation*, *19*(8), 2245–2279.
- Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature neuroscience*, *10*(6), 682–684. Publisher: Nature Publishing Group US New York.
- Bastos, A. M., Loonis, R., Kornblith, S., Lundqvist, M., & Miller, E. K. (2018). Laminar recordings in frontal cortex suggest distinct layers for maintenance and control of working memory. *Proceedings of the National Academy of Sciences*, (p. 201710323).
- Bastos, A. M., Lundqvist, M., Waite, A. S., Kopell, N., & Miller, E. K. (2020). Layer and rhythm specificity for predictive routing. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/2020.01.27.921783> (Accessed:2020-01-31).
- Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical Microcircuits for Predictive Coding. *Neuron*, *76*(4), 695–711.
- Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, *5*(2), 157–166.
- Bhandari, A., & Badre, D. (2018). Learning and transfer of working memory gating policies. *Cognition*, *172*, 89–100.
- Bicanski, A., & Burgess, N. (2018). A Neural Level Model of Spatial Memory and Imagery. *eLife*, *7*, e33752. `Tex.ids= bicanski_burgess_2018a`.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *2*(1), 32–48.
- Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S., & Magee, J. C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science*, *357*(6355), 1033–1036.

- Boerlin, M., Machens, C. K., & Denève, S. (2013). Predictive Coding of Dynamical Variables in Balanced Spiking Networks. *PLoS Computational Biology*, 9(11), e1003258.
- Borst, J. P., & Anderson, J. R. (2017). A step-by-step tutorial on using the cognitive architecture ACT-R in combination with fMRI data. *Journal of Mathematical Psychology*, 76, 94–103.
- Bourdoukan, R., & Denève, S. (2017). Enforcing balance allows local supervised learning in spiking recurrent networks. (p. 9).
- Brandon, M. P., Bogaard, A. R., Libby, C. P., Connerney, M. A., Gupta, K., & Hasselmo, M. E. (2011). Reduction of Theta Rhythm Dissociates Grid Cell Spatial Periodicity from Directional Tuning. *Science*, 332(6029), 595–599.
- Brandon, M. P., Bogaard, A. R., Schultheiss, N. W., & Hasselmo, M. E. (2013). Segregation of cortical head direction cell assemblies on alternating theta cycles. *Nature Neuroscience*, 16(6), 739–748.
- Brea, J., Gaál, A. T., Urbanczik, R., & Senn, W. (2016). Prospective Coding by Spiking Neurons. *PLoS Comput Biol*, 12(6), 1005003.
- Burgess, N., Barry, C., Keefe, J. O., Al, B. E. T., & O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 000, 1–3.
- Burwell, R. D., Shapiro, M. L., OMalley, M. T., & Eichenbaum, H. (1998). Positional firing properties of perirhinal cortex neurons. *NeuroReport*, 9(13), 3013–3018.
- Buschman, T. J., Denovellis, E. L., Diogo, C., Bullock, D., & Miller, E. K. (2012). Synchronous oscillatory neural ensembles for rules in the prefrontal cortex. *Neuron*, 76(4), 838–46.
- Bush, D., Barry, C., & Burgess, N. (2014). What do grid cells contribute to place cell firing? *Trends in Neurosciences*, 37(3), 136–145.
- Byrne, P., Becker, S., & Burgess, N. (2007). Remembering the Past and Imagining the Future: A Neural Model of Spatial Memory and Imagery. *Psychological Review*, (p. 36).
- Chalk, M., Gutkin, B., & Deneve, S. (2016). Neural oscillations as a signature of efficient coding in the presence of synaptic delays. *ELife*, (p. 23).
- Chen, L. L., Lin, L.-H., Green, E. J., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. *Behavioral Neuroscience*.

- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A Simple Framework for Contrastive Learning of Visual Representations. ArXiv:2002.05709 [cs, stat], <http://arxiv.org/abs/2002.05709> (Accessed:2023-02-28).
- Chevalier-Boisvert, M. (2018). Miniworld: Minimalistic 3D Environment for RL & Robotics Research. <https://github.com/maximecb/gym-miniworld>.
- Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience*, *115*(1), 3–25.
- Chrastil, E. R., Sherrill, K. R., Hasselmo, M. E., & Stern, C. E. (2015). There and Back Again: Hippocampus and Retrosplenial Cortex Track Homing Distance during Human Path Integration. *Journal of Neuroscience*, *35*(46), 15442–15452.
- Clark, B. J., Simmons, C. M., Berkowitz, L. E., & Wilber, A. A. (2018). The retrosplenial-parietal network and reference frame coordination for spatial navigation. *Behavioral Neuroscience*, *132*(5), 416–429.
- Clopath, C., Büsing, L., Vasilaki, E., & Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nature Neuroscience*, *13*(3), 344–352.
- Clopath, C., & Gerstner, W. (2010). Voltage and spike timing interact in STDP - a unified model. *Frontiers in Synaptic Neuroscience*, *2*(JUL), 1–11.
- Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., & Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357.
- Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cerebral Cortex*, *10*(9), 910–923.
- Cui, Z., Chen, W., & Chen, Y. (2016). Multi-Scale Convolutional Neural Networks for Time Series Classification General Terms. Tech. rep. <https://arxiv.org/pdf/1603.06995v4.pdf>.
- Dannenberg, H., Alexander, A. S., Robinson, J. C., & Hasselmo, M. E. (2019). The Role of Hierarchical Dynamical Functions in Coding for Episodic Memory and Cognition. *Journal of Cognitive Neuroscience*, (pp. 1–19).
- Dayan, P., & Abbott, L. F. (2003). *Theoretical neuroscience: Computational and Mathematical Modeling of Neural Systems*.
- Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz Machine. *Neural Computation*, *7*(5), 889–904.

- Denève, S., Alemi, A., & Bourdoukan, R. (2017). The Brain as an Efficient and Robust Adaptive Learner. *Neuron*, *94*(5), 969–977.
- Desimone, R., & Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. *Journal of Neurophysiology*, *57*(3), 835–868.
- Detorakis, G., Bartley, T., & Neftci, E. (2018). Contrastive Hebbian Learning with Random Feedback Weights. *arXiv:1806.07406 [cs, q-bio, stat]*. ArXiv: 1806.07406.
- DiCarlo, J., Zoccolan, D., & Rust, N. (2012). How Does the Brain Solve Visual Object Recognition? *Neuron*, *73*(3), 415–434.
- Diehl, G. W., Hon, O. J., Leutgeb, S., & 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. Publisher: Elsevier.
- Doron, G., Shin, J. N., Takahashi, N., Bocklisch, C., Skenderi, S., Drüke, M., de Mont, L., Toumazo, M., von Heimendahl, M., Brecht, M., Naud, R., & Larkum, M. E. (2019). Perirhinal input to neocortical layer 1 controls learning. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/713883> (Accessed:2020-12-18).
- Douglas, R. J., Martin, K. A., & Whitteridge, D. (1989). A Canonical Microcircuit for Neocortex. *Neural Computation*, *1*(4), 480–488.
- Duggins, P., & Eliasmith, C. (2022). Constructing functional models from biophysically-detailed neurons. *PLOS Computational Biology*, *18*(9), e1010461.
- Erdem, U. M., & Hasselmo, M. E. (2012). A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *European Journal of Neuroscience*, *35*(6), 916–931.
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex*, *1*(1), 1–47.
- Ferbinteanu, J., & Shapiro, M. L. (2003). Prospective and Retrospective Memory Coding in the Hippocampus. *Neuron*, *40*(6), 1227–1239.
- Ferrante, M., Shay, C. F., Tsuno, Y., Chapman, G. W., & Hasselmo, M. E. (2016). Post-Inhibitory Rebound Spikes in Rat Medial Entorhinal Layer II/III Principal Cells: In-Vivo, In-Vitro, and Computational Modeling Characterization. *Cerebral Cortex*, *27*(March).
- Frady, E. P., & Sommer, F. T. (2019). Robust computation with rhythmic spike patterns. *Proceedings of the National Academy of Sciences*, *116*(36), 18050–18059.

- Friese, U., Daume, J., Göschl, F., König, P., Wang, P., & Engel, A. K. (2016). Oscillatory brain activity during multisensory attention reflects activation, disinhibition, and cognitive control. *Scientific reports*, 6(July), 32775.
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132), 190–194. Publisher: Nature Publishing Group UK London.
- George, T. M., de Cothi, W., Clopath, C., Stachenfeld, K. L., & Barry, C. (2022). RatInABox: A toolkit for modelling locomotion and neuronal activity in complex continuous environments.
- Gerstner, W., & Kistler, W. M. (2002). Mathematical formulations of Hebbian learning. *Biological Cybernetics*, 87(5-6), 404–415.
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, 14(5), 350–363.
- Gilra, A., & Gerstner, W. (2017). Predicting non-linear dynamics by stable local learning in a recurrent spiking neural network. *eLife*, 6, e28295.
- Gofman, X., Tocker, G., Weiss, S., Boccara, C. N., Lu, L., Moser, M.-B., Moser, E. I., Morris, G., & Derdikman, D. (2019). Dissociation between Postrhinal Cortex and Downstream Parahippocampal Regions in the Representation of Egocentric Boundaries. *Current Biology*, (p. S0960982219308528).
- Greedy, W., Zhu, H. W., Pemberton, J., Mellor, J., & Costa, R. P. (2022). Single-phase deep learning in cortico-cortical networks. *Advances in Neural Information Processing Systems*. Publisher: arXiv Version Number: 1.
- Gregor, K., Danihelka, I., Mnih, A., Blundell, C., & Wierstra, D. (2014). Deep AutoRegressive Networks. *arXiv:1310.8499 [cs, stat]*. ArXiv: 1310.8499.
- Gustafsson, B., & Wigstrom, H. (1986). Hippocampal long-lasting potentiation produced by pairing single volleys and brief conditioning tetani evoked in separate afferents. *Journal of Neuroscience*, 6(6), 1575–1582. Publisher: Soc Neuroscience.
- Haeusler, S., & Maass, W. (2007). A statistical analysis of information-processing properties of lamina-specific cortical microcircuit models. *Cerebral Cortex*, 17(1), 149–162.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806.

- Halvagal, M. S., & Zenke, F. (2022). The combination of Hebbian and predictive plasticity learns invariant object representations in deep sensory networks. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/2022.03.17.484712> (Accessed:2022-03-27).
- Hardcastle, K., Ganguli, S., & Giocomo, L. M. (2015). Environmental Boundaries as an Error Correction Mechanism for Grid Cells. *Neuron*, 86(3), 827–839.
- Hartley, T., Burgess, N., Lever, C., Cacucci, F., & O'Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus*, 10(4), 369–379.
- Hasselmo, M. E. (1999). Neuromodulation: acetylcholine and memory consolidation. *Trends in cognitive sciences*, 3(9), 351–359.
- Hasselmo, M. E. (2005). What is the function of hippocampal theta rhythm? Linking behavioral data to phasic properties of field potential and unit recording data. *Hippocampus*, 15(7), 936–949.
- Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural computation*, 14(4), 793–817. Publisher: MIT Press One Rogers Street, Cambridge, MA 02142-1209, USA journals-info .
- Hasselmo, M. E., & Bower, J. M. (1993). Acetylcholine and memory. *Trends in neurosciences*, 16(6), 218–222. Publisher: Elsevier.
- Hasselmo, M. E., Rolls, E. T., & Baylis, G. C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behavioural Brain Research*, 32(3), 203–218.
- Hasselmo, M. E., & Shay, C. F. (2014). Grid cell firing patterns may arise from feedback interaction between intrinsic rebound spiking and transverse traveling waves with multiple heading angles. *Frontiers in systems neuroscience*, 8(October), 201.
- Hasselmo, M. E., & Stern, C. E. (2018). A network model of behavioural performance in a rule learning task. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373(1744).
- Hazan, H., Saunders, D. J., Khan, H., Patel, D., Sanghavi, D. T., Siegelmann, H. T., & Kozma, R. (2018). BindsNET: A Machine Learning-Oriented Spiking Neural Networks Library in Python. *Frontiers in Neuroinformatics*, 12.
- Hebb, D. O. (1949). *The organization of behavior: a neuropsychological theory*. Mahwah, N.J: L. Erlbaum Associates.

- Hertäg, L., & Sprekeler, H. (2020). Learning prediction error neurons in a canonical interneuron circuit. *eLife*, 9, e57541.
- Hinman, J. R., Brandon, M. P., Climer, J. R., Chapman, G. W., & Hasselmo, M. E. (2016). Multiple Running Speed Signals in Medial Entorhinal Cortex. *Neuron*, 91(3), 666–679.
- Hinman, J. R., Chapman, G. W., & Hasselmo, M. E. (2019). Neuronal representation of environmental boundaries in egocentric coordinates. *Nature Communications*, 10(1), 2772.
- Hinton, G. E. (2002). Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation*, 14(8), 1771–1800.
- Hochreiter, S. (1998). The Vanishing Gradient Problem During Learning Recurrent Neural Nets and Problem Solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 06(02), 107–116.
- Illing, B., Ventura, J., Bellec, G., & Gerstner, W. (2021). Local plasticity rules can learn deep representations using self-supervised contrastive predictions. *arXiv:2010.08262 [cs]*. Tex.ids= illing_gerstner_2021a arXiv: 2010.08262.
- Inglebert, Y., Aljadeff, J., Brunel, N., & Debanne, D. (2020). Altered spike timing-dependent plasticity rules in physiological calcium. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/2020.03.16.993675> (Accessed:2020-03-18).
- Izhikevich, E. M. (2006). Polychronization: Computation with Spikes. *Neural Computation*, 18(2), 245–282.
- Jacob, P.-Y., Casali, G., Spieser, L., Page, H., Overington, D., & Jeffery, K. (2017). An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nature Neuroscience*, 20(2), 173–175.
- Keinath, A. T., Julian, J. B., Epstein, R. A., & Muzzio, I. A. (2017). Environmental geometry aligns the hippocampal map during spatial reorientation. *Current Biology*, 27(3), 309–317. Publisher: Elsevier.
- Ketz, N. A., Jensen, O., & O’reilly, R. C. (2015). Thalamic pathways underlying prefrontal cortex-medial temporal lobe oscillatory interactions. *Trends in Neurosciences*, 38(1).
- Kingma, D. P., & Dhariwal, P. (2018). Glow: Generative Flow with Invertible 1x1 Convolutions. *arXiv*.

- Kingma, D. P., & Welling, M. (2019). An Introduction to Variational Autoencoders. *Foundations and Trends in Machine Learning*, 12(4), 307–392. ArXiv: 1906.02691.
- Kinsky, N. R., Sullivan, D. W., Mau, W., Hasselmo, M. E., & Correspondence, H. B. E. (2018). Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. *Current Biology*.
- Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences*, 11(6), 229–235.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 69, 59–69.
- Kramer, M. A., Roopun, A. K., Carracedo, L. M., Traub, R. D., Whittington, M. A., & Kopell, N. J. (2008). Rhythm Generation through Period Concatenation in Rat Somatosensory Cortex. *PLoS Computational Biology*, 4(9), e1000169.
- Kropff, E., Carmichael, J. E., Moser, E. I., & Moser, M.-B. (2021). Frequency of theta rhythm is controlled by acceleration, but not speed, in running rats. *Neuron*, 109(6), 1029–1039.e8.
- Kropff, E., Carmichael, J. E., Moser, M.-B., & Moser, E. I. (2015). Speed cells in the medial entorhinal cortex. *Nature*.
- Krupic, J., Bauza, M., Burton, S., Barry, C., & O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538), 232–235.
- LaChance, P. A., Todd, T. P., & Taube, J. S. (2019). A sense of space in postrhinal cortex. *Science Advances*, (p. 12).
- Larkum, M. (2013). A cellular mechanism for cortical associations: An organizing principle for the cerebral cortex. *Trends in Neurosciences*, 36.
- Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Lee, D., Lin, B.-J., & Lee, A. K. (2012). Hippocampal place fields emerge upon single-cell manipulation of excitability during behavior. *Science (New York, N.Y.)*, 337(6096), 849–853.
- Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(31), 9771–9777.
- Levy, W. B., & Steward, O. (1979). Synapses as associative memory elements in the hippocampal formation. *Brain research*, 175(2), 233–245. Publisher: Elsevier.

- Lillicrap, T. P., Cownden, D., Tweed, D. B., & Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nature Publishing Group*, 7.
- Lotter, W., Kreiman, G., & Cox, D. (2017). Deep Predictive Coding Networks for Video Prediction and Unsupervised Learning. *arXiv:1605.08104 [cs, q-bio]*. ArXiv: 1605.08104.
- Lotter, W., Kreiman, G., & Cox, D. (2020). A neural network trained for prediction mimics diverse features of biological neurons and perception. *Nature Machine Intelligence*, 2(4), 210–219.
- Low, I. I., Giocomo, L. M., & Williams, A. H. (2023). Remapping in a recurrent neural network model of navigation and context inference. <http://biorxiv.org/lookup/doi/10.1101/2023.01.25.525596> (Accessed:2023-02-06).
- Magee, J. C., & Grienberger, C. (2020). Synaptic Plasticity Forms and Functions. *Annual Review of Neuroscience*, 43(1), annurev-neuro-090919-022842.
- Mao, D., Kandler, S., McNaughton, B. L., & Bonin, V. (2017). Sparse orthogonal population representation of spatial context in the retrosplenial cortex. *Nature communications*, 8(1), 243. Publisher: Nature Publishing Group UK London.
- Marchette, S. A., Vass, L. K., Ryan, J., & Epstein, R. A. (2014). Anchoring the neural compass: coding of local spatial reference frames in human medial parietal lobe. *Nature Neuroscience*, 17(11), 1598–1606.
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science (New York, N.Y.)*, 275(5297), 213–215. Publisher: American Association for the Advancement of Science.
- Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences*, 95(9), 5323–5328.
- McNaughton, B., Barnes, C., Rao, G., Baldwin, J., & Rasmussen, M. (1986). Long-term enhancement of hippocampal synaptic transmission and the acquisition of spatial information. *Journal of Neuroscience*, 6(2), 563–571. Publisher: Soc Neuroscience.
- McNaughton, B. L., Barnes, C. a., & O'Keefe, J. (1984). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 54(1), 195.

- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the cognitive map. *Nature Reviews Neuroscience*, (p. 16).
- McNaughton, B. L., Mizumori, S., Barnes, C., Leonard, B., Marquis, M., & Green, E. (1994). Cortical representation of motion during unrestrained spatial navigation in the rat. *Cerebral Cortex*, 4(1), 27–39. Publisher: Oxford University Press.
- Miller, A. M., Mau, W., & Smith, D. M. (2019). Retrosplenial cortical representations of space and future goal locations develop with learning. *Current Biology*, 29(12), 2083–2090. Publisher: Elsevier.
- Mittelstaedt, M. L. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, 68, 566.
- Monaghan, C. K., Chapman, G. W., & Hasselmo, M. E. (2017). Systemic administration of two different anxiolytic drugs decreases local field potential theta frequency in the medial entorhinal cortex without affecting grid cell firing fields. *Neuroscience*, 364, 60–70.
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, 297(5868), 681–683.
- Moser, E. I., Moser, M.-B., & Roudi, Y. (2013). Network mechanisms of grid cells. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635), 20120511.
- Mountcastle, V. (1997). The columnar organization of the neocortex. *Brain*, 120(4), 701–722.
- Muller, R., & Kubie, J. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7(7), 1951–1968.
- Naud, R., Friedenberger, Z., & Toth, K. (2023). Silences, Spikes and Bursts: Three-Part Knot of the Neural Code. ArXiv:2302.07206 [q-bio], <http://arxiv.org/abs/2302.07206> (Accessed:2023-02-22).
- Naud, R., Marcille, N., Clopath, C., & Gerstner, W. (2008). Firing patterns in the adaptive exponential integrate-and-fire model. *Biological Cybernetics*, 99(4-5), 335–347.
- Naud, R., & Sprekeler, H. (2018). Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences*, 115(27), E6329–E6338. Tex.ids: naud_sprekeler_2018a.

- Nicola, W., & Clopath, C. (2017). Supervised learning in spiking neural networks with FORCE training. *Nature Communications*, 8(1).
- Nitz, D. A. (2006). Tracking route progression in the posterior parietal cortex. *Neuron*, 49(5), 747–756. Publisher: Elsevier.
- Ocko, S. A., Mallory, C. S., Campbell, M. G., Low, I. I. C., Ganguli, S., & Giocomo, L. M. (2018). Principles governing the integration of landmark and self-motion cues in entorhinal cortical codes for navigation. *Nature Neuroscience*, 21, 1096.
- O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381.
- O'Keefe, J., & Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: Theoretical significance and relationship to entorhinal grid cells. *Hippocampus*, 15(7), 853–866.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175.
- Olson, J. M., Tongprasearth, K., & Nitz, D. A. (2017). Subiculum neurons map the current axis of travel. *Nature neuroscience*, 20(2), 170–172. Publisher: Nature Publishing Group US New York.
- O'Mara, S. (2005). The subiculum: what it does, what it might do, and what neuroanatomy has yet to tell us. *Journal of Anatomy*, 207(3), 271–282.
- O'Reilly, R. C. (1997). The LEABRA model of neural interactions and learning in the neocortex. *Dissertation Abstracts International: Section B: The Sciences and Engineering*, 57, 6792.
- O'Reilly, R. C. (2000). Generalization in Interactive Networks: The Benefits of Inhibitory Competition and Hebbian Learning. 1241, 1199–1241.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press, 2nd ed.
- O'Reilly, R. C., Russin, J. L., Zolfaghar, M., & Rohrlich, J. (2021). Deep Predictive Learning in Neocortex and Pulvinar. *Journal of Cognitive Neuroscience*, 33(6), 1158–1196.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Rai-son, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., & Chintala, S.

- (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. ArXiv:1912.01703 [cs, stat], <http://arxiv.org/abs/1912.01703> (Accessed:2023-01-13).
- Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., & Naud, R. (2021). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience*.
- Peyrache, A., Schieferstein, N., & Buzsáki, G. (2017). Transformation of the head-direction signal into a spatial code. *Nature Communications*, 8(1), 1752.
- Pfister, J.-P. (2006). Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *Journal of Neuroscience*, 26(38), 9673–9682.
- Pouget, A., & Sejnowski, T. J. (1997). Spatial Transformations in the Parietal Cortex Using Basis Functions. *Journal of Cognitive Neuroscience*, 9(2), 222–237.
- Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87.
- Raudies, F., Brandon, M. P., Chapman, G. W., & Hasselmo, M. E. (2015). Head direction is coded more strongly than movement direction in a population of entorhinal neurons. *Brain Research*, 1621, 355–367.
- Raudies, F., & Hasselmo, M. E. (2012). Modeling Boundary Vector Cell Firing Given Optic Flow as a Cue. *PLoS Computational Biology*, 8(6), 1–17.
- Recanatesi, S., Farrell, M., Lajoie, G., Deneve, S., Rigotti, M., & Shea-Brown, E. (2021). Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nature Communications*, 12(1), 1417.
- Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., de Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Thérien, D., & Kording, K. P. (2019). A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11), 1761–1770.
- Richter, C. G., Thompson, W. H., Conrado, X., Bosman, A., & Fries, X. P. (2017). Top-Down Beta Enhances Bottom-Up Gamma.
- Rockland (2010). Five points on columns. *Frontiers in Neuroanatomy*.

- Rounds, E. L., Alexander, A. S., Nitz, D. A., & Krichmar, J. L. (2018). Conjunctive coding in an evolved spiking model of retrosplenial cortex. *Behavioral Neuroscience*, 132(5), 430–452.
- Rumelhart, D. E., & McLelland, D. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press.
- Schaeffer, R., Khona, M., & Fiete, I. R. (2022). No Free Lunch from Deep Learning in Neuroscience: A Case Study through Models of the Entorhinal-Hippocampal Circuit. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/2022.08.07.503109> (Accessed:2022-10-02).
- Shay, C. F., Ferrante, M., Chapman, G. W., & Hasselmo, M. E. (2015). Rebound spiking in layer II medial entorhinal cortex stellate cells: Possible mechanism of grid cell function. *Neurobiology of Learning and Memory*.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-k., & Woo, W.-c. (2015). Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. (p. 9). Tex.ids: shi.woo.2015a arXiv: 1506.04214.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. a., & Barnes, C. a. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2), 149–172.
- Smith, D. M., Barredo, J., & Mizumori, S. J. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus*, 22(5), 1121–1133. Publisher: Wiley Online Library.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M. B., & Moser, E. I. (2008). Representation of Geometric Borders in the Entorhinal Cortex. *Science (New York, N.Y.)*, 322(December), 1865–1868.
- Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature Neuroscience*, 20(11), 1643–1653.
- Steffenach, H.-A., Witter, M., Moser, M.-B., & Moser, E. I. (2005). Spatial Memory in the Rat Requires the Dorsolateral Band of the Entorhinal Cortex. *Neuron*, 45(2), 301–313.
- Sussillo, D., & Abbott, L. (2009). Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4), 544–557. Tex.ids= sussillo_abbott_2009a.
- Sutskever, I., Hinton, G., & Taylor, G. (2009). The Recurrent Temporal Restricted Boltzmann Machine. *Advances in Neural Information Processing Systems*, (p. 8).

- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to Sequence Learning with Neural Networks. *Neural Information Processing*, (p. 9).
- Sutton, R. S., & Barto, A. G. (1998). Introduction. *Reinforcement Learning*.
- Sutton, R. S., Precup, D., & Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2), 181–211.
- Takahashi, H., & Magee, J. C. (2009). Pathway Interactions and Synaptic Plasticity in the Dendritic Tuft Regions of CA1 Pyramidal Neurons. *Neuron*, 62(1), 102–111.
- Taube, J., Muller, R., & Ranck, J. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, 10(2), 420–435.
- Taube, J. S. (2007). The head direction signal: origins and sensory-motor integration. *Annual review of neuroscience*, 30, 181–207.
- Tsuno, Y., Chapman, G. W., & Hasselmo, M. E. (2015). Rebound spiking properties of mouse medial entorhinal cortex neurons in vivo. *The European journal of neuroscience*, 42, 2974–2984.
- van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *Journal of Comparative Neurology*, 315(2), 200–216. Publisher: Wiley Online Library.
- Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, 10(11), 792–802.
- Vass, L. K., & Epstein, R. a. (2013). Abstract representations of location and facing direction in the human brain. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 33(14), 6133–6142.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention Is All You Need.
- Vercruyse, F., Naud, R., & Sprekeler, H. (2021). Self-organization of a doubly asynchronous irregular network state for spikes and bursts. preprint, Neuroscience. <http://biorxiv.org/lookup/doi/10.1101/2021.03.29.437548> (Accessed:2021-08-10).
- Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C., & Gerstner, W. (2011). Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science*, 334(6062), 1569–1573.

- Vogt, B. A., & Miller, M. W. (1983). Cortical connections between rat cingulate cortex and visual, motor, and postsubicular cortices. *The Journal of Comparative Neurology*, *216*(2), 192–210.
- Wallis, J. D., Anderson, K. C., & Miller, E. K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature*, *411*(6840), 953–956.
- Wang, C., Chen, X., Lee, H., Deshmukh, S. S., Yoganarasimha, D., Savelli, F., & Knierim, J. J. (2018a). Egocentric coding of external items in the lateral entorhinal cortex. *Science*, *362*(6417), 945–949.
- Wang, L., Hagoort, P., & Jensen, O. (2018b). Gamma Oscillatory Activity Related to Language Prediction. *Journal of Cognitive Neuroscience*, *30*(8), 1075–1085.
- Werbos, P. (1990). Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, *78*(10), 1550–1560.
- Whitlock, J. R., Pfuhl, G., Dagslott, N., Moser, M.-B., & Moser, E. I. (2012). Functional split between parietal and entorhinal cortices in the rat. *Neuron*, *73*(4), 789–802.
- Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W., & Behrens, T. E. J. (2022). How to build a cognitive map. *Nature Neuroscience*, *25*(10), 1257–1272.
- Wilber, A. A., Clark, B. J., Forster, T. C., Tatsuno, M., & McNaughton, B. L. (2014). Interaction of Egocentric and World-Centered Reference Frames in the Rat Posterior Parietal Cortex. *Journal of Neuroscience*, *34*(16), 5431–5446.
- Williams, R. J., & Zipser, D. (1989). A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation*, *1*(2), 270–280.
- Wunderlich, T. C., & Pehle, C. (2021). Event-based backpropagation can compute exact gradients for spiking neural networks. *Scientific Reports*, *11*(1), 12829.
- Xie, X., & Seung, H. S. (2003). Equivalence of Backpropagation and Contrastive Hebbian Learning in a Layered Network. *Neural Computation*, *15*(2), 441–454.
- Yoganarasimha, D., Yu, X., & Knierim, J. J. (2006). Head direction cell representations maintain internal coherence during conflicting proximal and distal cue rotations: comparison with hippocampal place cells. *Journal of Neuroscience*, *26*(2), 622–631. Publisher: Soc Neuroscience.
- Yoo, S. B. M., Tu, J. C., & Hayden, B. Y. (2021). Multicentric tracking of multiple agents by anterior cingulate cortex during pursuit and evasion. *Nature Communications*, *12*(1), 1985.

- Yoo, S. B. M., Tu, J. C., Piantadosi, S. T., & Hayden, B. Y. (2020). The neural basis of predictive pursuit. *Nature Neuroscience*.
- Zhu, H., Paschalidis, I. C., & Hasselmo, M. E. (2018). Neural circuits for learning context-dependent associations of stimuli. *Neural Networks*.

CURRICULUM VITAE

