Predictive Learning for Self-Supervised Mapping and Localization

G. William Chapman (wchapman@bu.edu)
Boston University

Andrew S. Alexander (andyalexander@ucsb.edu)
University of California, Santa Barbara

Frances S. Chance (fschanc@sandia.gov)
Sandia National Laboratories

Michael E. Hasselmo (hasselmo@bu.edu)
Boston University

Abstract

Spatial navigation involves the formation of coherent representations of a map-like space, while simultaneously tracking current location in a primarily unsupervised manner. Despite a plethora of neurophysiological experiments revealing spatially tuned neurons across the mammalian neocortex and subcortical structures, it remains unclear how such representations are acquired in the absence of explicit allocentric targets. Drawing upon the concept of predictive learning, we utilize a biologically plausible learning rule which utilizes sensory-driven observations with internally driven expectations and learns through a contrastive manner to better predict sensory information. We implement this learning rule in a network with the feedforward and feedback pathways known to be necessary for spatial navigation. After training, we find that the receptive fields of the modeled units resemble experimental findings, with allocentric and egocentric representations in the expected order along processing streams. These findings suggest that a self-supervised prediction of sensory information can extract latent structure from the environment.

Introduction

A sense of allocentric, or map like space, is essential for navigating to goal which are outside of the immediate sensory space. Neurophysiological studies have revealed a series of brain regions with egocentric and allocentric representations, allowing insight to the steps of transformation from egocentric signals to allocentric representations (Alexander et al., 2020), and have been formalized in models of spatial navigation (Bicanski & Burgess, 2018). It is still unclear however, how these representations are learned. While modern machine learning methods can utilize supervised learning to create allocentric representations which match experimental findings (Banino et al., 2018), it is unclear how such representations may be learned in the absence of explicit supervision. Recent work has shown that self-supervised *prediction* of sensory information in shallow recurrent networks can learn internal representations of allocentric space (Recanatesi et al., 2021). Here, we utilize expand this framework of predictive coding to utilize a biologically inspired learning rule, and embed the current networks in a topology matching brain regions known to contain spatial representations, to investigate how predictive coding may generate additional intermediate transformations from sensory information to spatial knowledge.

Methods

Task & Setup We created a virtual agent which follows an algorithmically generated trajectory in a virtual 1m X 1m environment, mimicking the behavior of rodents (Raudies & Hasselmo, 2012; George, de Cothi, Clopath, Stachenfeld, & Barry, 2022), and generate a 3-dimensional rendering of visual information (Chevalier-Boisvert, 2018) to create RGB-Depth inputs to provide to the agent inputs. The RGB, depth, rotational velocity and linear running speed are provided to the agent through separate input streams discussed below. We then run the simulation, with dynamics and learning rules described below, for one-thousand trials of 60 simulated minutes, after which learning is disabled. The tuning curves are analyzed on a separate 60-minute trial after weights have been frozen. We then introduce an impassable 50cm boundary in the middle of the environment to analyze responses to novel environments after learning.

Model Architecture The overall topology of the network is illustrated in Figure 1, where white blocks indicate a linear encoder and the grey blocks indicate a model of canonical microcircuit as described in the next subsection. Feedforward inputs onto granular neurons are indicated in blue, and feedback inputs onto distal dendrites are shown in red. The network receives a visual input which is decomposed into identity and distance information before entering the network. A separate input pathway supplies self-motion information in the form of rotational velocity and running speed.

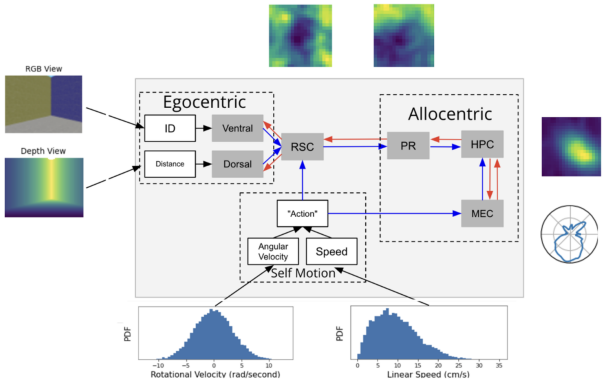

Figure 1: Macroscale model architecture: in which the agent receives only egocentric information and self motion information. Each module indicated by a grey block implements a form of predictive learning between feedforward (red arrows) and feedback (blue arrows). Tuning curves in the outer portions of the figure show the firing rates (higher rates in yellow) for example individual units modeled in each region, across a simulated 1m virtual environment. The ratemaps show the primary responses in each region, beginning with egocentric boundary cells in the retrosplenial cortex, boundary vector cells in perirhinal, place cells in HPC, and head direction tuning in the entorhinal cortex.

Microscale Architecture We implement a rate-approximation of a model of canonical cortical microcircuit as illustrated in Figure 2. Each region received a combination of feedforward inputs onto granular neurons and feedback information onto distal dendrites of pyramidal units. Such networks have been shown to learn in a self-supervised manner for simple tasks when arranged in a hierarchical manner (Chapman & Hasselmo, 2023). Granular neurons follow a rate-based approximation of point neurons, where each unit is a leaky-integrator with a sigmoidal activation function. Pyramidal neurons were designed to replicate the control of firing rate by somatic in-

put and burst rate by distal dendritic input (Naud & Sprekeler, 2018; Payeur, Guerguiev, Zenke, Richards, & Naud, 2021), by incorporating additional dynamics:

$$\begin{aligned}
\tau_v \frac{dv_L}{dt} &= -v_L + \sigma(u_L/D_d) + \sum_{N \in FF} W_{NL} r_N \\
r_L &= \sigma(v_L) \\
\tau_u \frac{du_L}{dt} &= -u_L + 2\frac{dr_L}{dt} + \sum_{N \in FB} W_{Nu} b_N \\
b_L &= r_L \odot \sigma(u_L)
\end{aligned} \tag{1}$$

The added state u_L is the potential of the distal dendrites, while v_L is specifically the *somatic* potential. FB ("feedback") is the collection of all regions which project onto the distal dendrites, while FF ("feedforward") is the local inputs from either the granular neurons or superficial pyramidal neurons. The burst rate, b_L, is controlled through the activation of distal dendrites

Learning To replicate the local feedback gated learning rules demonstrated in spiking networks (Chapman & Hasselmo, 2023), feedforward synapses onto the pyramidal neurons evolve according to the learning rule:

$$\Delta W_{XY} = \eta((p_Y - \bar{p_Y}) \odot r_Y) \otimes r_X \tag{2}$$

Where X and Y are the pre and post synaptic neurons, respectively. This rule contains the standard 'pre-post' associational learning, with an additional gating by the instantaneous burst-probability compared to a long-term running average, which implements a weak version of an error signal. This equation can also be shown to approximate a temporal error term by the coupling of the dendritic derivative to the derivative of firing rate (equation 1, line 3). All other weights, including feedback connections and terminals onto granular neurons, are static. The untrained granular layer is necessary to sustain activity when utilizing the temporal error learning rule.

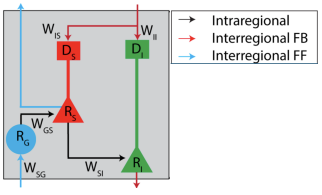

Figure 2: Architecture of the predictive module: The structure of each module, where rates ('R') of each submodule represent the firing rate of a given population of neurons, and weights are labeled according to their source and target submodules. Units are coloured by their laminar location, also denoted 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.

Results

Tuning Curves After training, we analyze the tuning curves of each region utilizing a parametric point-process regression, specifically testing for place fields, head direction tuning, running speed, egocentric boundary tuning, and grid cells. Similar to previous approaches (Alexander et al., 2022) significance for each of these tuning curves was determined by an iterative leave-one-out approach from the point-process model, and statistical significance was determined by comparison to an F-test. The tuning curves on the outer-ring of Figure 1 indicate the primary feedforward receptive fields identified, with a transition from sensory-driven representations in outer regions, to allocentric representations in innermost regions. Intermediate areas contained tuning to a combination of allocentric and egocentric information, such as egocentric boundary cells (shown in plots at the top of Figure 1).

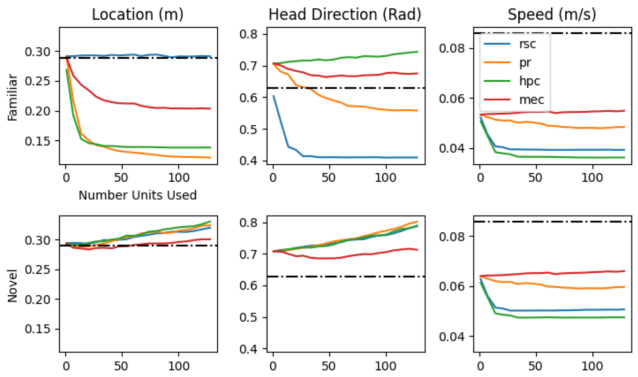

Figure 3: Decoding Results: Behavioral variables across regions, as a function of the number of units used in the familiar (top) and novel (bottom) environments. Mean squared error is measured, with lower values indicating higher performance. (Top) Behavioral variables are easily decoded in familiar environments in the various regions, even with a small proportion of units. (rsc=retrosplenial, pr=perirhinal, hpc=hippocampus, mec=medial entorhinal cortex). (Bottom) In the novel environment, allocentric codes are performing at or worse than chance, while self-motion decoding is conserved.

Behavioral Decoding We next attempt to *linearly* decode the behavioral variables (position, head direction, and movement speed) from each region in both the baseline and novel settings using post-training weights. The novel setting allocentric, but not egocentric, decoding is largely disrupted compared to the familiar environment (see Figure 3). This suggests that additional mechanisms may be required to reconcile large differences between sensory inputs and learned expectations.

Discussion Overall, these results validate the potential for prediction of egocentric information to generate internal representations of allocentric space. While a single layer of a predictive network was shown in previous work to directly extract allocentric representations, we show here that a hierarchy of predictive regions forms intermediate representations that match experimental findings.

References

Alexander, A. S., Robinson, J. C., Dannenberg, H., Kinsky, N. R., Levy, S. J., Mau, W., ... Hasselmo, M. E. (2020, January). Neurophysiological coding of space and time in the hippocampus, entorhinal cortex, and retrosplenial cortex. *Brain and Neuroscience Advances*, *4*, 239821282097287. doi: 10.1177/2398212820972871

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.

Banino, A., Barry, C., Uria, B., Blundell, C., Lillicrap, T., Mirowski, P., ... Kumaran, D. (2018). Vector-based navigation using grid-like representations in artificial agents. *Nature*, *557*(7705), 429–433. doi: 10.1038/s41586-018-0102-6

Bicanski, A., & Burgess, N. (2018). A Neural Level Model of Spatial Memory and Imagery. *eLife*, *7, e33752*. doi: 10.7554/eLife.33752

Chapman, G. W., & Hasselmo, M. E. (2023). Predictive learning by a burst-dependent learning rule. *Neurobiology of Learning and Memory*, 107826.

Chevalier-Boisvert, M. (2018). *Miniworld: Minimalistic 3D Environment for RL & Robotics Research.*

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.* bioRxiv.

Naud, R., & Sprekeler, H. (2018, July). Sparse bursts optimize information transmission in a multiplexed neural code. *Proceedings of the National Academy of Sciences*, *115*(27), E6329-E6338. doi: 10.1073/pnas.1720995115

Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A., & Naud, R. (2021, July). Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience*, *24*(7), 1010–1019. doi: 10.1038/s41593-021-00857-x

Raudies, F., & Hasselmo, M. E. (2012). Modeling Boundary Vector Cell Firing Given Optic Flow as a Cue. *PLoS Computational Biology*, *8*(6), 1–17. doi: 10.1371/journal.pcbi.1002553

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.