
A generative model of the hippocampal formation trained with theta driven local learning rules

Tom M George¹
Caswell Barry² Kimberly Stachenfeld^{3,4} Claudia Clopath^{5,1} Tomoki Fukai⁶

¹Sainsbury Wellcome Centre, UCL, UK ²Dept. of Cell and Developmental Biology, UCL, UK

³Google DeepMind, London, UK ⁴Columbia University, New York, NY

⁵Bioengineering Dept., Imperial College, UK ⁶Okinawa Institute of Science and Technology, Japan

tom.george.20@ucl.ac.uk

Abstract

Advances in generative models have recently revolutionised machine learning. Meanwhile, in neuroscience, generative models have long been thought fundamental to animal intelligence. Understanding the biological mechanisms that support these processes promises to shed light on the relationship between biological and artificial intelligence. In animals, the hippocampal formation is thought to learn and use a generative model to support its role in spatial and non-spatial memory. Here we introduce a biologically plausible model of the hippocampal formation tantamount to a Helmholtz machine that we apply to a temporal stream of inputs. A novel component of our model is that fast theta-band oscillations (5-10 Hz) gate the direction of information flow throughout the network, training it akin to a high-frequency wake-sleep algorithm. Our model accurately infers the latent state of high-dimensional sensory environments and generates realistic sensory predictions. Furthermore, it can learn to path integrate by developing a ring attractor connectivity structure matching previous theoretical proposals and flexibly transfer this structure between environments. Whereas many models trade-off biological plausibility with generality, our model captures a variety of hippocampal cognitive functions under one biologically plausible local learning rule.

1 Introduction

Generative models seek to create new data samples which are similar to those from the training set. To do so they must learn the probability distribution of the training data, comprising a rich, generalisable and accurate model of the world. Many of the recent advances in AI have involved types of generative models: VAEs [1], GANs [2], diffusion models [3] and autoregressive models [4] have seeded improvements in AI capabilities ranging from data compression [5] to image generation [6] and natural language [7]. In neuroscience, the animal brain has long been known to exploit generative models [8, 9]. The ability to generate representative sensory data samples can be used directly, for example during offline planning or memory recall. It can also be used indirectly to aid training of inference networks with the goal of processing rich, noisy and high dimensional streams of incoming sensory stimuli, as discussed in the predictive coding literature [10]. In a sentence: “What I cannot create [generate], I do not understand [inference]” (R. Feynman).




Figure 1: A biologically plausible generative model is trained with theta frequency wake-sleep cycles and a local learning rule. **a** Network schematic: high-D stimuli from an underlying environmental latent, z , arrive at the basal dendrites of the sensory layer, p , and map to the hidden layer, g (this is the inference model, weights in green). Simultaneously, top-down predictions from the hidden layer g arrive at the apical dendrites of p (this is the generative model, weights in blue). **b** Neurons in layers p and g have three compartments. A fast oscillation, $\theta(t)$, gates which dendritic compartment – basal (p_B , g_B) or apical (p_A , g_A) – drives the soma. A local learning rule adjusts input weights to minimise the prediction error between dendritic compartments and the soma. **c** This equates to rapidly switching “wake” and “sleep” cycles which train the generative and inference models. Panel c displays just two updates per theta-cycle, in reality there are many ($\delta t \ll T_\theta$).

The hippocampal-entorhinal system (aka. hippocampal formation) – a brain structure implicated in spatial [11] and non-spatial [12] memory – provides a pertinent example. Its primary role seems to be inference [13]: mapping sensory inputs into a robust and decodable representation of state (grid cells [14], place cells [11] etc. [15]). A generative model is thought to have a dual role in learning: supporting offline tasks such as route planning [16] and memory consolidation [17], and online during behaviour with path integration [18]. Path integration enables the hippocampal network to maintain an up-to-date and accurate estimate of its position in the absence of reliable sensory data by integrating self-motion cues. A recent flurry of computational [19, 20, 21] and theoretical [22, 21] work has highlighted the importance of path integration as a key objective explaining hippocampal function and representations.

Existing computational generative models of the hippocampal formation [23, 24] account for many of its cognitive functions and internal representations but require non-trivial learning rules and message passing protocols which don’t connect with known aspects of biology. Computational models of path integration [25, 26, 27] have mostly focussed on continuous attractor networks which, although experimentally supported [28], alone lack the complexity or expressivity required of a fully general model of the hippocampal memory system.

The primary contribution of this paper is to introduce a biologically plausible model of sequence learning in the hippocampus which unifies its capacities as a generative model of sensory stimuli and path integration under one schema. To do this we propose modeling the hippocampal formation as a Helmholtz machine [29] which learns to predict sensory stimuli given the current hidden state and action (e.g. velocity). We propose a deep connection between the hippocampal theta oscillation [30] and the unsupervised wake-sleep algorithm [31] for training Helmholtz machines. Though this class of generative models isn’t widely used, and lacks the scalability of the latest transformer-based sequence learners, it excels in this context since it has many natural points of contact with biology (both in terms of architecture and neural dynamics) yet still maintains the expressiveness afforded to models of the brain by deep neural networks.

In this paper we:

- introduce a new model of the hippocampal formation which learns the latent structure of an incoming stream of sensory stimuli analogous to a Helmholtz machine.
- describe a biologically plausible learning regime: Theta-oscillations gate information flow through multi-compartmental neurons which rapidly switches the system between “wake” and “sleep” phases. All plasticity is local.

- train our model on stimuli from a biologically relevant spatial exploration task and show it learns to path integrate by developing a ring attractor connectivity structure (comparable to theoretical predictions and empirical results in deep recurrent neural networks trained with gradient descent). Learning generalises: when the agent moves to a new environment, path integration capabilities recover without needing to relearn the path integration weights.

Our model of the hippocampal formation simultaneously (i) accounts for its role as a generative model of sensory stimuli, (ii) can learn to path integrate and (iii) can transfer structural knowledge between environments. The model, though here applied to the hippocampus, can be viewed as a step towards a general solution for how biological neural networks in many brain regions (for example visual cortex [10]) can learn generative models of the world.¹

1.1 Related work

A recent generative model of the hippocampus, the Tolman-Eichenbaum Machine [23], proposed that the hippocampal formation be thought of as a hierarchical network performing latent state inference. Medial entorhinal cortex (MEC) sits atop the hierarchy and learns an abstract representation of space which is mapped to the hippocampus (HPC) where it is bound onto incoming sensory stimuli. Once trained the system can act in a generative fashion by updating the hidden representation with idiothetic action signals and then predicting the upcoming sensory experience. The drawback of this model, and others which share a similar philosophical approach [32, 24], is that it requires training via backpropagation through time (or equivalent end-to-end optimisation schemes, as in [24]) without clear biological correlates. Related hierarchical network architectures have also been studied in the context of reinforcement learning [33] and hippocampal associative memory [34].

Historically, hippocampal models of path integration have focused on continuous attractor networks (CANs) [25, 26, 27, 21] in entorhinal cortex. A bump of activity representing location is pushed around the CAN by speed and/or head-direction selective inputs, thus integrating self-motion. CANs have received substantial experimental support [28] but few studies adequately account for *how* this structure is learned by the brain in the first place. One exception exists outside the hippocampal literature: Vafidis et al. [35] built a model of path integration in the fly head-direction system which uses local learning rules. Here we go further by embedding our path integrator inside a hierarchical generative model. Doing so additionally relaxes the assumption (made by Vafidis et al. [35] and others [36]) that sensory inputs into the path integrator are predefined and fixed. Instead, by allowing all incoming and outgoing synapses to be learned from random initialisations, we achieve a more generalisable model capable of transferring structure between environments (see section 3.3).

Hippocampal theta oscillations have been linked to predictive sequence learning before [37, 38, 39, 40] where research has focused on the compressive effects of theta *sequences* and how these interplay with short timescale synaptic plasticity. Instead of compression, here we hypothesize the role of theta is to control the direction information flows through the hierarchical network.

Finally, a recent theoretical work by Bredenberg et al. [41] derived, starting from principles of Bayesian variational inference, a biologically plausible learning algorithm for approximate Bayesian inference of a hierarchical network model built from multi-compartmental neurons and trained with local learning rules using wake-sleep cycles. Here we build a similar network to theirs (i) extending it to a spatial exploration task and mapping the hidden layers onto those in the hippocampal formation, (ii) simplifying the learning rules and relaxing a discrete-time assumption – instead, opting for a temporally continuous formulation more applicable to biological tasks such as navigation – and (iii) adapting the hidden layer to allow idiothetic action signals to guide updates (aka. path integration). Their work provides a theoretical foundation for our own, helping to explaining *why* learning converges on accurate generative models.

2 A biologically plausible generative model trained with rapidly switching wake-sleep cycles and local learning rules

In sections 2 and 3 we give concise, intuitive descriptions of the model and experiments; expanded details can be found in the supplementary material.

¹Code provided at <https://github.com/TomGeorge1234/HelmholtzHippocampus>

2.1 Basic model summary

We consider learning in an environment defined by a latent state, $z(t)$, which updates according to stochastic dynamics initially unknown to the network,

$$\frac{dz}{dt} = f_z(t). \quad (1)$$

These dynamics depends on the task; first we consider $z(t)$ to be a set of mutually independent random variables and later we consider the more realistic task of an agent moving on a 1D track.

The network receives sensory input which is a function of the latent state into a sensory layer, $\mathbf{p}(t)$, and communicates this to a hidden layer (aka “internal state”), $\mathbf{g}(t)$. The network contains both an *inference* (aka. *recognition*) model which infers the hidden state from the sensory input (green arrows, Fig. 1a) and a *generative* model which updates the hidden state with recurrent synapses and maps this back to the sensory layer (blue arrows). As we will soon identify these processes with Basal and Apical dendritic compartments of pyramidal neurons we label activations sampled from the inference model with the subscript B and those from the generative model with the subscript A .² In summary

$$\left. \begin{aligned} \mathbf{p}_B(t + \delta t) &= \bar{\mathbf{p}}(z(t)) \\ \mathbf{g}_B(t + \delta t) &= \sigma_{g_B}(\mathbf{w}_{g_B} \mathbf{p}(t)) \end{aligned} \right\} \quad \text{Inference model} \quad (2)$$

$$\left. \begin{aligned} \mathbf{g}_A(t + \delta t) &= \sigma_{g_A}(\mathbf{w}_{g_A} \mathbf{g}(t)) \\ \mathbf{p}_A(t + \delta t) &= \sigma_{p_A}(\mathbf{w}_{p_A} \mathbf{g}(t)) \end{aligned} \right\} \quad \text{Generative model.} \quad (3)$$

$\mathbf{w}_{g_B}, \mathbf{w}_{p_A}, \mathbf{w}_{g_A}$ are matrices of randomly initialised and plastic synaptic weights. $\bar{\mathbf{p}}$ maps the environmental latent into a vector of neural inputs. σ ’s denote activation functions applied to the dendritic pre-activations – either the identity ($\sigma(x) = x$) or rectified tanh functions ($\sigma(x) = \max(0, \tanh(x))$). A small amount of noise is added to the dendritic activations to simulate realistic biological learning.

We believe that the widely adopted convention of modelling neurons as single-compartment perceptrons is limiting. By considering, in a minimal extension, the distributed dendritic structure of real neurons we can tap into significant potential for explaining hippocampal learning. Theoretical [42, 43, 44, 45] and experimental [46, 47, 48] research into credit assignment in biological neurons has identified different roles for basal and apical dendrites: basal dendrites are thought to receive bottom-up drive from sensory inputs whereas apical dendrites receive top-down drive from higher layers in the sensory hierarchy [49]. Following this line of research — and matching an equivalent theoretical model of latent state inference described by [41] — we identify the inference process with synaptic inputs into a basal dendritic compartment of pyramidal neurons and the generative process with synaptic inputs into an apical dendritic compartment. In summary, each $\mathbf{p}$ and $\mathbf{g}$ neuron in our model has three compartments: a somatic compartment, a basal dendritic compartment and an apical dendritic compartment (Fig. 1b). Only the somatic activation is used for communication between layers (right hand side of Eqns. (2) and (3)) while dendritic compartment activations are variables affecting internal neuronal dynamics and learning as described below (Eqns. (4) and (6)).

2.2 Theta oscillations gate the direction of information flow through the network

The dynamics of the somatic activations $\mathbf{p}(t)$ and $\mathbf{g}(t)$ are as follows: the voltage in each soma is either equal to the voltage in the basal compartment *or* the voltage in the apical compartment depending on the phase of an underlying theta oscillation. This is achieved by a simple theta-gating mechanism (Fig. 1b):

$$\begin{aligned} \mathbf{p}(t) &= \theta(t) \mathbf{p}_B(t) + (1 - \theta(t)) \mathbf{p}_A(t) \\ \mathbf{g}(t) &= \theta(t) \mathbf{g}_B(t) + (1 - \theta(t)) \mathbf{g}_A(t). \end{aligned} \quad (4)$$

where $\theta(t)$ is a 5 Hz global theta oscillation variable defined by the square wave function:

$$\theta(t) = \begin{cases} 1, & \text{if } t/T \bmod 1 \leq 0.5 \\ 0, & \text{if } t/T \bmod 1 > 0.5 \end{cases} \quad (5)$$

²These labellings conveniently match the notion that inferences are made from layers Below in the sensory hierarchy (bottom-up) whereas generative predictions arrive from Above (top-down).

for $T = 1/f_\theta$ and $f_\theta = 5$ Hz, matching the hippocampal theta frequency (5-10 Hz) [50]. According to this model theta-band oscillations in the hippocampal local field potential gate which dendritic compartment drives the soma. Experimental [47, 51, 52] and modelling work [53] gives provisional support for this assumption.

These local theta-dynamics have global consequences: the early phase ($\theta(t) = 1$) of each theta cycle can be thought of as a “wake” phase where information flows upwards through the network from the environment to the hidden layer, sampling the inference model. The latter phase ($\theta(t) = 0$) of each theta cycle is a “sleep” phase where information flows down from the hidden layer to the sensory units, sampling the generative model. These dynamics are displayed in Fig. 1.

2.3 Hebbian-style learning rules train synapses to minimise local prediction errors

In contrast to comparable models which are optimised end-to-end using backpropagation through time our model learns synaptic weights according to a local plasticity rule which is a simplified variant of a rule proposed by Urbanczik and Senn [43]. Incoming synaptic projections are continually adjusted in order to minimize the discrepancy between the somatic activation and the dendritic activation. The full learning rules are described in the supplement but simplified versions are given here:

$$\begin{aligned}\frac{d\mathbf{w}_{g_B}}{dt} &\propto (\mathbf{g}(t) - \mathbf{g}_B(t))\mathbf{p}(t)^\top \\ \frac{d\mathbf{w}_{p_A}}{dt} &\propto (\mathbf{p}(t) - \mathbf{p}_A(t))\mathbf{g}(t)^\top \\ \frac{d\mathbf{w}_{g_A}}{dt} &\propto (\mathbf{g}(t) - \mathbf{g}_A(t))\mathbf{g}(t)^\top\end{aligned}\tag{6}$$

Notably this learning rule is equivalent for *all* plastic synapses in the model: $\mathbf{p}$ to $\mathbf{g}$, $\mathbf{g}$ to $\mathbf{p}$ and the recurrent $\mathbf{g}$ to $\mathbf{g}$ synapses (see Fig. 1b). If a local prediction error is detected, for example the somatic activation is larger than the dendritic activation, then the synaptic strength of inputs into that dendritic compartment which are positive/negative are strengthened/weakened to reduce the error. This model can equivalently be viewed as a type of Hebbian learning – weight change is proportional to the correlation of pre- and post-synaptic activity (the first term) – regularised (by the second term) to prevent unbounded growth.

During the wake phase the weights of the generative model ($\mathbf{w}_{p_A}$ and $\mathbf{w}_{g_A}$) are trained and plasticity on the inference weights ($\mathbf{w}_{g_B}$) falls to zero. This occurs naturally because $\mathbf{p} = \mathbf{p}_B$ so there will be no basal prediction errors to correct. During sleep the reverse occurs; the weights of the inference model are trained and plasticity on the generative model falls to zero. Experimentally, apical activity is known to guide plasticity at basal synapses in CA1 [46]. This alternating, coordinated regime of sampling and learning (sample-inference-train-generative, then sample-generative-train-inference) is a hallmark of the wake-sleep algorithm. It fundamentally differs from the forward and backward sweeps of backpropagation since neurons remain provisionally active at all times so the process of learning minimally perturbs perception. Also, whereas backpropagation sends error signals down through the network to train synaptic weights, here only predictions are sent between layers and error signals are calculated locally at each dendrite.

As discussed in section 1, Bredenberg et al. [41] mathematically derive learning rules similar to these starting from a loss function closely related to the evidence lower bound (ELBO). As such our identification of early- and late-theta phases as “wake” and “sleep” cycles can be considered precise: from a Bayesian perspective our hippocampal model is minimising a modified ELBO loss (see supplement) thus learns to find approximately optimal inference and generative models accounting from the temporally varying stimulus stream it is presented.

2.4 Velocity inputs into the hidden layer

For path integration, the hidden state needs access to an idiothetic (internally generated) velocity signal. To satisfy this we endow the hidden layer, $\mathbf{g}$, with conjunctive velocity inputs, henceforth “conjunctive cells”, as shown in Fig. 3a & b. Conjunctive cells are organised into two groups: $\mathbf{g}_{v_L}$ is responsible for leftward motion and $\mathbf{g}_{v_R}$ for rightward motion. Each conjunctive cell receives input from the hidden units and either the leftward ($v_L = \max(0, -\dot{x})$) or rightward ($v_R = \max(0, \dot{x})$) component of the velocity. For the results shown this connectivity is one-to-one $[\mathbf{w}_{g_{v_L}}]_{ij} = [\mathbf{w}_{g_{v_R}}]_{ij} = \delta_{ij}$ but




Figure 2: Learning in an environment of temporally varying latents. **a** In this artificial task the latent space comprises of $N_z = 5$ independent random variables with an autocorrelation decay timescale of 1 s. **b** Prediction errors (difference between apical and basal activations) in sensory and hidden layers reduce over training time. **c** Tested in wake mode ($\theta = 1$) after training, the ground truth stimulus matches apical prediction for all stimulus dimensions (one shown) implying the network is efficiently “autoencoding” the sensory inputs into and back out of the compressed hidden layer. **d** Tested in sleep mode ($\theta = 0$, no environmental inputs), generated data from the hidden units, g , have an autocorrelation curve which matches that of the true latents implying a statistically accurate generative model has been learned. More extensive samples from this model, before and after training, can be found in Fig. ??

random connectivity works too, see supplement. Finally, conjunctive cells send return connections back to the apical dendritic compartment of the hidden units via a randomly initialised plastic synaptic weight matrix. This inputs are what drive the hidden units to path integrate.

This model takes inspiration from so-called conjunctive grid cells [54] found in the medial entorhinal cortex (MEC). These cells, though to be an integral component of the mammalian path integration system [27], are jointly tuned to head direction and location much like the conjunctive cells in our model. An important and novel aspect of our model is that synaptic weights between or into the hidden units are *learned*. This deviates from other models for example that by Burak and Fiete [27] (where all connectivity is predefined and fixed) or Vafidis et al. [35] and Widloski and Fiete [36] (where sensory inputs to the hidden units are pre-defined and fixed). This is not only more realistic but affords the model flexibility to translate path integration abilities between environments without having to relearn them, a form of transfer learning which we demonstrate in section 3.3.

3 Results

3.1 Validation on an artificial latent learning task

We begin by testing the basic model (i.e. without conjunctive inputs, Fig. 1a) on an artificial task. $N_z = 5$ latents, $z_i(t)$, are independently sampled from a smooth, random process with an autocorrelation timescale of 1 second (Fig. 2a). The sensory layer, $N_p = 50$, then receives a high-dimensional random linear mixture of the latents into the basal compartments:

$$p_B(t) = \mathbf{A}z(t), \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{50 \times 5}$ and $[\mathbf{A}]_{ij} \sim \mathcal{N}(0, \frac{1}{\sqrt{N_z}})$. The hidden layer, $g(t)$, is matched in size to the latent process, $N_g = N_z = 5$, and all dendritic activation functions are linear. We train the model for 30 minutes of simulated time and track prediction errors, the difference between the basal and apical activations in the sensory and hidden layers, which reliably decreased throughout training (Fig. 2b). We then perform two tests designed to confirm whether the model has learnt accurate inference and generative models.




Figure 3: The hippocampal model learns to path integrate on a 1D track using a ring attractor. **a** Position selective (place cell) inputs drive basal dendrites of the sensory layer **p** (HPC). **b** Hidden units (MEC) are connected to two sets of “conjunctive cells” which each connect back to one of the hidden neurons (**g**) and either the leftward (for g_{vL}) or rightward (for g_{vR}) velocity of the agent allowing velocity information to enter the network. Synaptic strengths of the return connections from the conjunctive cells to the MEC hidden units, as well as those for the MEC recurrent connectivity (collective denoted w_{gA}), are randomly initialised and plastic. **c** After training, reordering the hidden units by the position of peak activity reveals a ring attractor in the synaptic weight matrices. Centre-surround recurrent connectivity stabilises an activity bump which is then “pushed” around the attractor manifold by asymmetric connections from the conjunctive cells, integrating velocity. Bands of zero weights show MEC neurons which have become perpetually inactive (aka “died”). The bottom panel displays the matrix row-averages, utilizing the circular symmetry of the environment to align rows before averaging. **d** Learning plateaus after 15 mins of simulated time. **e** Path integration ability is demonstrated in a lesion study: after 10 seconds in the normal oscillatory mode the network is placed into sleep mode (aka generative mode), lesioning the position-dependent sensory inputs. Despite this HPC continues to accurately encode position, evidence that the MEC ring attractor is path integrating the velocity inputs and sending predictions back to HPC. Lower panel shows the accumulated decoding error as well as the mean $\pm$ SEM over 50 trials.

First, we set the dynamics of the model to “wake” mode ($\theta = 1$) and measure the basal and apical activations of one of the sensory neurons for 60 seconds. Close correspondence (Fig. 2c) confirms that the network accurately “autoencodes” the high-dimensional sensory inputs through the compressed hidden layer. Since all activation functions are linear this implies that w_{gB} and w_{pA} are pseudoinverses. Next, we place the network in “sleep” mode ($\theta = 0$) and allow the generative model to run freely. The autocorrelation of the generated hidden states ($g(t|\theta = 0)$, displayed fully in the supplement) match that of the true environmental latents ($z(t)$), Fig. 2d, implying the generative model has statistics closely matching those of the true underlying generative process.

3.2 Learnable path integration with a hidden ring attractor

Next we turn our attention to the hippocampal formation’s role in spatial navigation, and our central result. The environment consists of an agent randomly moving around a 1 m 1D circular track (motion and cell data is generated using the RatInABox package [55]). The basal compartment of each HPC neuron is spatially tuned to a single different Gaussian input however non-Gaussian

randomly spatially tuned inputs work as well (see supplement Fig. ??b):

$$[\mathbf{p}_B(t)]_i = \exp \left[- \frac{(x(t) - x_i)^2}{2\sigma^2} \right]. \quad (8)$$

$x(t)$ is the position of the agent and $\{x_i\}_{i=1}^{N_p}$ are the centres of the Gaussian inputs ($\sigma = 6$ cm), intended to simulate hippocampal place fields, evenly spaced at 1 cm intervals along the track. MEC (i.e. the hidden layer, $\mathbf{g}(t)$) is matched in size $N_g = N_p = 100$ with rectified tanh activation functions on both dendritic compartments ($\sigma_{g_B}(x) = \sigma_{g_A}(x) = \max(0, \tanh(x))$) and HPC (the sensory layer $\mathbf{p}(t)$) is linear ($\sigma_{p_A}(x) = x$). Two populations of conjunctive cells (Fig. 3a & b) feed into the apical compartments of the MEC recurrent units. Random initialisation of $\mathbf{w}_{g_B}$ means that MEC neurons start off with random non-Gaussian spatial tunings. $\mathbf{w}_{g_A}$ and $\mathbf{w}_{p_A}$ are also randomly initialised.

The network is trained for 30 minutes with learning plateauing after 15 (Fig. 3d). A lesion study, designed to test path integration, is then performed as follows: First, the network is run for 10 seconds normally (i.e. with theta-oscillating periods of wake and sleep). Since the simulated HPC neurons receive place-tuned inputs uniformly ordered along the track (i.e. $x_j > x_i \forall i, j > i$) an activity heatmap of HPC reveals a bump of activity accurately tracking agent’s position (Fig. 3e, left). The network is then placed into a sleep phase ($\theta = 0$) for 20 seconds. This amounts to a full sensory lesion since top-down MEC inputs, not bottom-up place-tuned sensory inputs, drive HPC. Despite the full sensory lesion, hippocampal activity remains approximately unperturbed and the activity bump continues to accurately track position, slowly accumulating errors (Fig. 3e right). Since our HPC layer has no recurrent connectivity it cannot support this post-lesion activity on its own. Instead feed-forward drive from an MEC ring attractor, which we turn our attention to now, is responsible for maintaining the HPC code.

To find the ring attractor we must first reorder the MEC cells. We do this according to the position of the peak of their receptive fields (defined in the supplement). After reordering, the recurrent connectivity matrix can be seen to have acquired a centre-surround connectivity profile. Nearby MEC cells were, on average, strongly and positively recurrently connected to one another. Those far apart weakly inhibit one another (Fig. 3c, left; band of strong positive weights along diagonal flanked by weak negative weights). This profile matches that of a quasi-continuous ring attractor: local excitatory and long-range inhibitory connections stabilise a bump of activity on the attractor manifold in the absence of sensory input [56]. Weights from the conjunctive cells acquired asymmetric connectivity (Fig. 3c, middle & right) skewed towards the velocity direction for which they are selective. These asymmetric connections enable conjunctive cells to “push” the activity bump around the manifold, integrating velocity (see supplement for a visualisation of the MEC bump attractor). Theoretical work on ring attractors has demonstrated that for accurate path integration the asymmetric weights must be proportional to the derivative of the symmetric weights [56], approximately observed here. A noteworthy observation is that some MEC neurons become perpetually inactive; this is a consequence of the fact that *both* top-down and bottom-up synapses into the hidden layer are plastic and can fall to zero (Fig. 3c bands of zero-weights) satisfying a trivial $g_A = g_B = 0$ solution for minimising the prediction error. Despite this, not all MEC neurons die and the surviving subset are sufficient for path integration. In supplementary section ?? we discuss additional results showing when the network learns robust path integrate under a variety of plasticity, initialisation and noise manipulations.

Crucially, what sets this model apart from others [19, 20, 21, 22] is that the network is not optimized using a conventional path-integration objective and backpropagation. Instead, it has been demonstrated how path integration can naturally arise in a biologically constrained network subject to a much simpler (yet more broadly applicable) local objective, in cases where idiothetic velocity signals are available to the hidden layers.

3.3 Remapping: transfer of structural knowledge between environments

Finally, we demonstrate how our trained network can transfer structural knowledge – which here means the ring attractor and thereby path integration – between environments. We start by training the network as in section 3.2; the only difference is that for simplicity we choose to fix $\mathbf{w}_{g_B} = \delta_{ij}$ giving rise to MEC representations which, like HPC, are unimodal (this constraint can be relaxed and, in the more general case, MEC units typically have multiple receptive fields, Fig ??d, reminiscent of grid cells). We then simulate a hippocampal “remapping” event by shuffling the sensory inputs to the HPC layer (Fig. 4a & b, top panel) and retraining the network for a further 30 minutes but this time




Figure 4: Remapping and transfer of structural knowledge between environments. **a** After training (as in Fig. 2) place cell inputs are shuffled to simulate a “remapping” event observed when an agent moves to a new environment. The agent then retrain for an additional 30 minutes; during this period internal MEC weights, and weights from the conjunctive cells to MEC are held fixed while MEC $\leftrightarrow$ HPC weights remain plastic. **b** Receptive fields of the HPC and MEC neuronal populations at different stages in the experiment: Initially after remapping HPC and MEC inputs are randomised. MEC relearns rate maps as they were before remapping but with a constant phase shift. Note: neurons are ordered by the position of their peak activity on the track *before* remapping and this ordering is maintained in subsequent panels. **c** The error ($\pm$ SEM over 50 trials) after 1 second of path integration is shown at different stages of the experiment. Although path integration is initially disrupted after remapping it recovers despite no relearning of the MEC synapses where the ring attractor is stored.

holding weights in the hidden layer, w_{gA} . Only the HPC $\leftrightarrow$ MEC synapses (w_{gB} & w_{pA}) remain plastic during retraining. Biologically this may be accounted for by the observation that cortical plasticity is substantially slower than hippocampal plasticity [57].

During biological remapping events place cells remap independently whereas grid cells remap *en masse* with entire modules shifting by the same constant phase [58]. This observation is reproduced in our model: after retraining MEC units regroup with receptive fields as they were before remapping but with a constant phase shift along the track. This re-emergence of structure occurs because the ring attractor seeds a bump of activity on the attractor manifold (during the “sleep” phases of retraining) onto which the shuffled HPC inputs then bind. Since nothing constrains *where* on the circularly symmetric attractor manifold this regrouping can initiate, only relative correlations, modulo a phase shift, are preserved.

Decoding error one second after a sensory lesion is tested just *before* remapping, just *after* remapping and after retraining (Fig. 4c). After the remapping path integration abilities temporarily disappear because the MEC ring attractor is still tuned to the old and invalid HPC receptive fields. After relearning – and despite *no adjustments to the MEC weights, w_{gA} , where the ring attractor is stored* – path integration recovers to almost the level before remapping. This differs substantially from other local models of path integration learning [35, 36] which don’t consider plasticity on the ring attractor inputs. In these models, adaptation to a new environment necessarily requires complete relearning of the ring attractor. Instead our model exploits the basic fact that movement (path integration) in one environment is fundamentally the same as in another, one must simply learn a new mapping to/from the ring attractor, “translating” it to fit the new sensory stimuli.

4 Discussion

We propose that the hippocampal formation resembles a Helmholtz machine, simultaneously learning an inference and generative model of sensory stimuli. Like previous models [23] medial entorhinal

cortex (MEC) sits hierarchically above the hippocampus (HPC) to which it sends generative predictions. Our model differs in the learning rules and neural dynamics: local prediction errors are minimised between distinct dendritic compartments receiving bottom-up and top-down signals. Theta oscillations regulate internal neural dynamics, switching the network between wake and sleep phases. In a navigation task our MEC model forms a ring attractor capable of path integration. Despite simple learning rules and dynamics our model retains key cognitive capabilities of the hippocampal formation including the ability to transfer knowledge across different sensory environments.

Local learning rules are commonly recognised as essential in biologically plausible learning algorithms [43]. However, the importance of learning *scheduling* – how neural systems coordinate or multiplex distinct phases of forward and backward information flow – is often overlooked [59]. Neural oscillations such as theta, hypothesized to temporally coordinate communication between neuronal populations [60], likely play an underexplored role in this regard (neural “bursting” has also been pointed out as a potential solution to multiplexing [61]). One advantage of the wake-sleep algorithm, which this study suggests neural oscillations can support, compared to forward and backward sweeps is that, during convergence, the two phases become highly similar, allowing learning to proceed without affecting perception.

While our discussion has primarily focused on theta oscillations as a mechanism for learning, they have also been proposed as a mechanism for short-range future prediction via so-called “mind-travel” [62]. During the latter phase of each theta cycle (i.e. the sleep phase) gain amplified velocity signals might rapidly drive the MEC activity bump along the manifold allowing the agent to assess nearby upcoming locations. This complimentary proposition could neatly integrate into the framework proposed here and emphasizes the need for further investigation into the multifaceted functions of neural rhythms within the hippocampal/entorhinal system.

Beyond theta oscillations, both faster gamma cycles [63] and the slower physiological states of sleep and wake [64] have been associated with learning. Based on our model we suggest a tentative hypothesis that theta oscillations may be favored due to an optimality criterion; whilst faster oscillations could be a mechanism to prevent extreme drift during sleep that might disrupt learning their frequency might be upper bounded biophysically by the neural time constants associated with the biophysical processes supporting dendritic gating the soma. These ideas, their relevance to other brain regions involved in generative learning, 2D spatial dynamics, and offline memory consolidation/replay remain exciting questions for future theoretical and experimental investigation.

References

- [1] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.
- [2] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.
- [3] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- [4] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017. URL <http://arxiv.org/abs/1706.03762>.
- [5] Tom M George and Pietro Lio. Unsupervised machine learning for data encoding applied to ovarian cancer transcriptomes. November 2019. doi: 10.1101/855593. URL <https://doi.org/10.1101/855593>.
- [6] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *CoRR*, abs/2102.12092, 2021. URL <https://arxiv.org/abs/2102.12092>.
- [7] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with gpt-4, 2023.
- [8] Karl Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, January 2010. doi: 10.1038/nrn2787. URL <https://doi.org/10.1038/nrn2787>.
- [9] Samuel J. Gershman. The generative adversarial brain. *Frontiers in Artificial Intelligence*, 2, September 2019. doi: 10.3389/frai.2019.00018. URL <https://doi.org/10.3389/frai.2019.00018>.
- [10] Rajesh P. N. Rao and Dana H. Ballard. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1):79–87, January 1999. doi: 10.1038/4580. URL <https://doi.org/10.1038/4580>.
- [11] John OKeefe. Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51(1):78–109, January 1976. doi: 10.1016/0014-4886(76)90055-8. URL [https://doi.org/10.1016/0014-4886\(76\)90055-8](https://doi.org/10.1016/0014-4886(76)90055-8).
- [12] Larry R. Squire. "memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans": Correction. *Psychological Review*, 99(3):582–582, July 1992. doi: 10.1037/0033-295x.99.3.582. URL <https://doi.org/10.1037/0033-295x.99.3.582>.
- [13] Honi Sanders, Matthew A Wilson, and Samuel J Gershman. Hippocampal remapping as hidden state inference. *eLife*, 9, June 2020. doi: 10.7554/elife.51140. URL <https://doi.org/10.7554/elife.51140>.
- [14] Torkel Hafting, Marianne Fyhn, Sturla Molden, May-Britt Moser, and Edvard I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, June 2005. doi: 10.1038/nature03721. URL <https://doi.org/10.1038/nature03721>.
- [15] Edvard I Moser, May-Britt Moser, and Bruce L McNaughton. Spatial representation in the hippocampal formation: a history. *Nature Neuroscience*, 20(11):1448–1464, November 2017. doi: 10.1038/nn.4653. URL <https://doi.org/10.1038/nn.4653>.
- [16] Hugo J. Spiers and Eleanor A. Maguire. Thoughts, behaviour, and brain dynamics during navigation in the real world. *NeuroImage*, 31(4):1826–1840, July 2006. doi: 10.1016/j.neuroimage.2006.01.037. URL <https://doi.org/10.1016/j.neuroimage.2006.01.037>.
- [17] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2):147–153, January 2011. doi: 10.1038/nn.2732. URL <https://doi.org/10.1038/nn.2732>.

- [18] B. L. McNaughton, C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Qin, W. E. Skaggs, M. Suster, and K. L. Weaver. Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *Journal of Experimental Biology*, 199(1):173–185, January 1996. doi: 10.1242/jeb.199.1.173. URL <https://doi.org/10.1242/jeb.199.1.173>.
- [19] Christopher J. Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization, 2018.
- [20] Andrea Banino, Caswell Barry, Benigno Uribe, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J. Chadwick, Thomas Degris, Joseph Modayil, Greg Wayne, Hubert Soyer, Fabio Viola, Brian Zhang, Ross Goroshin, Neil Rabinowitz, Razvan Pascanu, Charlie Beattie, Stig Petersen, Amir Sadik, Stephen Gaffney, Helen King, Koray Kavukcuoglu, Demis Hassabis, Raia Hadsell, and Dharshan Kumaran. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, May 2018. doi: 10.1038/s41586-018-0102-6. URL <https://doi.org/10.1038/s41586-018-0102-6>.
- [21] Ben Sorscher, Gabriel C. Mel, Samuel A. Ocko, Lisa M. Giocomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137.e13, 2023. ISSN 0896-6273. doi: <https://doi.org/10.1016/j.neuron.2022.10.003>. URL <https://www.sciencedirect.com/science/article/pii/S0896627322009072>.
- [22] William Dorrell, Peter E. Latham, Timothy E. J. Behrens, and James C. R. Whittington. Actionable neural representations: Grid cells from minimal constraints, 2023.
- [23] James C.R. Whittington, Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E.J. Behrens. The tolmán-eichenbaum machine: Unifying space and relational memory through generalization in the hippocampal formation. *Cell*, 183(5):1249–1263.e23, November 2020. doi: 10.1016/j.cell.2020.10.024. URL <https://doi.org/10.1016/j.cell.2020.10.024>.
- [24] Dileep George, Rajeev V. Rikhye, Nishad Gothoskar, J. Swaroop Guntupalli, Antoine Dedieu, and Miguel Lázaro-Gredilla. Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nature Communications*, 12(1), April 2021. doi: 10.1038/s41467-021-22559-5. URL <https://doi.org/10.1038/s41467-021-22559-5>.
- [25] W. E. Skaggs, J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton. A model of the neural basis of the rat’s sense of direction. *Advances in neural information processing systems*, 7:173–180, 1995. ISSN 1049-5258.
- [26] Alexei Samsonovich and Bruce L. McNaughton. Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17(15):5900–5920, 1997. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.17-15-05900.1997. URL <https://www.jneurosci.org/content/17/15/5900>.
- [27] Yoram Burak and Ila R. Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS Computational Biology*, 5(2):e1000291, February 2009. doi: 10.1371/journal.pcbi.1000291. URL <https://doi.org/10.1371/journal.pcbi.1000291>.
- [28] Mikail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *arXiv*, 2021. doi: 10.48550/arxiv.2112.03978.
- [29] Peter Dayan, Geoffrey E. Hinton, Radford M. Neal, and Richard S. Zemel. The helmholtz machine. *Neural Computation*, 7(5):889–904, September 1995. doi: 10.1162/neco.1995.7.5.889. URL <https://doi.org/10.1162/neco.1995.7.5.889>.
- [30] Gyorgy Buzsaki. Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340, January 2002. doi: 10.1016/s0896-6273(02)00586-x. URL [https://doi.org/10.1016/s0896-6273\(02\)00586-x](https://doi.org/10.1016/s0896-6273(02)00586-x).
- [31] Geoffrey E. Hinton, Peter Dayan, Brendan J. Frey, and Radford M. Neal. The "wake-sleep" algorithm for unsupervised neural networks. *Science*, 268(5214):1158–1161, May 1995. doi: 10.1126/science.7761831. URL <https://doi.org/10.1126/science.7761831>.
- [32] Benigno Uribe, Borja Ibarz, Andrea Banino, Vinicius Zambaldi, Dharshan Kumaran, Demis Hassabis, Caswell Barry, and Charles Blundell. A model of egocentric to allocentric understanding in mammalian brains. November 2020. doi: 10.1101/2020.11.11.378141. URL <https://doi.org/10.1101/2020.11.11.378141>.

- [33] Dongqi Han, Kenji Doya, and Jun Tani. Self-organization of action hierarchy and compositionality by reinforcement learning with recurrent neural networks. *Neural Networks*, 129:149–162, September 2020. doi: 10.1016/j.neunet.2020.06.002. URL <https://doi.org/10.1016/j.neunet.2020.06.002>.
- [34] Sugandha Sharma, Sarthak Chandra, and Ila Fiete. Content addressable memory without catastrophic forgetting by heteroassociation with a fixed scaffold. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 19658–19682. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/sharma22b.html>.
- [35] Pantelis Vafidis, David Oswald, Tiziano D’Albis, and Richard Kempster. Learning accurate path integration in ring attractor models of the head direction system. *eLife*, 11:e69841, jun 2022. ISSN 2050-084X. doi: 10.7554/eLife.69841. URL <https://doi.org/10.7554/eLife.69841>.
- [36] John Widloski and Ila R. Fiete. A Model of Grid Cell Development through Spatial Exploration and Spike Time-Dependent Plasticity. *Neuron*, 83(2):481–495, 2014. ISSN 0896-6273. doi: 10.1016/j.neuron.2014.06.018.
- [37] William E. Skaggs, Bruce L. McNaughton, Matthew A. Wilson, and Carol A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–172, 1996. doi: 10.1002/(sici)1098-1063(1996)6:2<149::aid-hipo6>3.0.co;2-k. URL [https://doi.org/10.1002/\(sici\)1098-1063\(1996\)6:2<149::aid-hipo6>3.0.co;2-k](https://doi.org/10.1002/(sici)1098-1063(1996)6:2<149::aid-hipo6>3.0.co;2-k).
- [38] M.R. Mehta, M.C. Quirk, and M.A. Wilson. Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields. *Neuron*, 25:707–715, 2000.
- [39] Tom M George, William de Cothi, Kimberly L Stachenfeld, and Caswell Barry. Rapid learning of predictive maps with STDP and theta phase precession. *eLife*, 12, March 2023. doi: 10.7554/elif.80663. URL <https://doi.org/10.7554/elif.80663>.
- [40] Tom M George. Theta sequences as eligibility traces: A biological solution to credit assignment. In *International Conference on Learning Representations 2023 (TinyPapers track)*, 2023. doi: <https://doi.org/10.48550/arXiv.2305.08124>. URL <https://openreview.net/forum?id=vd16AYbem3Z>.
- [41] Colin Bredenberg, Eero P Simoncelli, Benjamin S H Lyo, and Cristina Savin. Impression learning: Online representation learning with synaptic plasticity. page 13, 2021.
- [42] Konrad P. Kording and Peter König. Supervised and unsupervised learning with two sites of synaptic integration. *Journal of Computational Neuroscience*, 11(3):207–215, 2001. doi: 10.1023/a:1013776130161. URL <https://doi.org/10.1023/a:1013776130161>.
- [43] Robert Urbanczik and Walter Senn. Learning by the Dendritic Prediction of Somatic Spiking. *Neuron*, 81(3):521–528, 2014. ISSN 08966273. doi: 10.1016/j.neuron.2013.11.030. URL <https://linkinghub.elsevier.com/retrieve/pii/S0896627313011276>.
- [44] Joao Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems*, pages 8721–8732, 2018.
- [45] Blake A Richards and Timothy P Lillicrap. Dendritic solutions to the credit assignment problem. *Current Opinion in Neurobiology*, 54:28–36, February 2019. doi: 10.1016/j.conb.2018.08.003. URL <https://doi.org/10.1016/j.conb.2018.08.003>.
- [46] Katie C Bittner, Christine Grienberger, Sachin P Vaidya, Aaron D Milstein, John J Macklin, Junghyup Suh, Susumu Tonegawa, and Jeffrey C Magee. Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nature Neuroscience*, 18(8):1133–1142, July 2015. doi: 10.1038/nn.4062. URL <https://doi.org/10.1038/nn.4062>.
- [47] Jurij Brankack, Mark Stewart, and Steven E. Fox. Current source density analysis of the hippocampal theta rhythm: associated sustained potentials and candidate synaptic generators. *Brain Research*, 615(2):310–327, July 1993. doi: 10.1016/0006-8993(93)90043-m. URL [https://doi.org/10.1016/0006-8993\(93\)90043-m](https://doi.org/10.1016/0006-8993(93)90043-m).

- [48] Kenji Mizuseki, Anton Sirota, Eva Pastalkova, and György Buzsáki. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron*, 64(2):267–280, October 2009. doi: 10.1016/j.neuron.2009.08.037. URL <https://doi.org/10.1016/j.neuron.2009.08.037>.
- [49] Matthew E. Larkum. Are dendrites conceptually useful? *Neuroscience*, 489:4–14, May 2022. doi: 10.1016/j.neuroscience.2022.03.008. URL <https://doi.org/10.1016/j.neuroscience.2022.03.008>.
- [50] David J. Foster and Matthew A. Wilson. Hippocampal theta sequences. *Hippocampus*, 17(11):1093–1099, 2007. doi: 10.1002/hipo.20345. URL <https://doi.org/10.1002/hipo.20345>.
- [51] Christian Holscher, Roger Anwyl, and Michael J. Rowan. Stimulation on the positive phase of hippocampal theta rhythm induces long-term potentiation that can be depotentiated by stimulation on the negative phase in area ca1. *The Journal of Neuroscience*, 17(16):6470–6477, August 1997. doi: 10.1523/jneurosci.17-16-06470.1997. URL <https://doi.org/10.1523/jneurosci.17-16-06470.1997>.
- [52] Yoko Yamaguchi, Yoshito Aota, Bruce L. McNaughton, and Peter Lipa. Bimodality of theta phase precession in hippocampal place cells in freely running rats. *Journal of Neurophysiology*, 87(6):2629–2642, June 2002. doi: 10.1152/jn.2002.87.6.2629. URL <https://doi.org/10.1152/jn.2002.87.6.2629>.
- [53] Michael E. Hasselmo, Clara Bodelón, and Bradley P. Wyble. A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, 14(4):793–817, April 2002. doi: 10.1162/089976602317318965. URL <https://doi.org/10.1162/089976602317318965>.
- [54] Francesca Sargolini, Marianne Fyhn, Torkel Hafting, Bruce L. McNaughton, Menno P. Witter, May-Britt Moser, and Edvard I. Moser. Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762, May 2006. doi: 10.1126/science.1125572. URL <https://doi.org/10.1126/science.1125572>.
- [55] Tom M George, William de Cothi, Claudia Clopath, Kimberly Stachenfeld, and Caswell Barry. RatInABox: A toolkit for modelling locomotion and neuronal activity in continuous environments. aug 2022. doi: 10.1101/2022.08.10.503541. URL <https://doi.org/10.1101/2022.08.10.503541>.
- [56] K Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *The Journal of Neuroscience*, 16(6):2112–2126, March 1996. doi: 10.1523/jneurosci.16-06-02112.1996. URL <https://doi.org/10.1523/jneurosci.16-06-02112.1996>.
- [57] Ceren Ergorul and Howard Eichenbaum. Essential role of the hippocampal formation in rapid learning of higher-order sequential associations. *The Journal of Neuroscience*, 26(15):4111–4117, April 2006. doi: 10.1523/jneurosci.0441-06.2006. URL <https://doi.org/10.1523/jneurosci.0441-06.2006>.
- [58] Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I. Moser. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194, February 2007. doi: 10.1038/nature05601. URL <https://doi.org/10.1038/nature05601>.
- [59] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. *eLife*, 6, December 2017. doi: 10.7554/elife.22901. URL <https://doi.org/10.7554/elife.22901>.
- [60] Pascal Fries. Rhythms for cognition: Communication through coherence. *Neuron*, 88(1):220–235, October 2015. doi: 10.1016/j.neuron.2015.09.034. URL <https://doi.org/10.1016/j.neuron.2015.09.034>.
- [61] Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake A. Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience*, 24(7):1010–1019, May 2021. doi: 10.1038/s41593-021-00857-x. URL <https://doi.org/10.1038/s41593-021-00857-x>.
- [62] Honi Sanders, César Rennó-Costa, Marco Idiart, and John Lisman. Grid cells and place cells: An integrated view of their navigational and memory function. *Trends in Neurosciences*, 38

- (12):763–775, December 2015. doi: 10.1016/j.tins.2015.10.004. URL <https://doi.org/10.1016/j.tins.2015.10.004>.
- [63] Kwan Tung Li, Junhao Liang, and Changsong Zhou. Gamma oscillations facilitate effective learning in excitatory-inhibitory balanced neural circuits. *Neural Plasticity*, 2021:1–18, January 2021. doi: 10.1155/2021/6668175. URL <https://doi.org/10.1155/2021/6668175>.
- [64] William E. Skaggs and Bruce L. McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271(5257):1870–1873, March 1996. doi: 10.1126/science.271.5257.1870. URL <https://doi.org/10.1126/science.271.5257.1870>.