



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/108622/>

Version: Accepted Version

Proceedings Paper:

Fox, C and Prescott, T (2010) Hippocampus as unitary coherent particle filter. In: The 2010 International Joint Conference on Neural Networks (IJCNN). The 2010 International Joint Conference on Neural Networks (IJCNN), 18-23 Jul 2010, Barcelona, Spain. IEEE, pp. 1-8. ISBN: 978-1-4244-6916-1.

<https://doi.org/10.1109/IJCNN.2010.5596681>

© 2010 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Hippocampus as unitary coherent particle filter

Charles Fox and Tony Prescott
Adaptive Behaviour Research Group
Department of Psychology, University of Sheffield
charles.fox@sheffield.ac.uk

Abstract—We present a mapping of the hippocampal formation onto a Temporal Restricted Boltzmann Machine [23] based architecture, running a deterministic version of Gibbs sampling, and extended with a lostness detection and recovery circuit modelled on subiculum and septal acetylcholine (ACh). The mapping approximates Bayesian filtering, which infers both auto-associative de-noised percepts and temporal sequences, the latter including sequences of places during navigation. Inference may be viewed as a neurally implemented particle filter with a single particle – as suggested previously [5] as a purely behavioural animal model.

I. INTRODUCTION

The classical view of hippocampus is as a single loop. The principal input structures of the hippocampus are the superficial layers of Entorhinal Cortex (ECs). ECs project to Dentate Gyrus (DG) which is believed to sparsify the encoding of ECs. Both ECs and DG project to CA3, which also receives strong recurrent connections that are disabled [10] by septal ACh. There is a recently-discovered back-projection from CA3 to DG [20]. CA3 and ECs project to CA1, which in turn projects to the deep layers of Entorhinal cortex, ECd. ECs, CA1 and ECd outputs appear to share a coding scheme, as evidenced by one-to-one topographic projections. In contrast, DG and CA3 outputs are thought to work in a second basis or latent space. There is also a second loop. ECs and CA1 project to Subiculum (Sub), which projects to the midbrain Septum (Sep) via fornix. Septal ACh and GABA fibres project back to all parts of hippocampus.

We make a distinction between two broad schools of hippocampal modelling: spatial sequences, and auto-associative. They differ in their representation of the world, and in the role attributed to the recurrent connections.

Spatial sequence models (e.g. [21], [2], [4]) are based on the existence of rodent cells responding to place. They use the activity of the i th CA3 (or sometimes CA1) cell to represent a strength of belief, $P(i)$, in the agent being at the i th discrete place,

$$CA3(i) = P(i). \quad (1)$$

Weights connect the place cells to sensory inputs, and possibly to features extracted from inputs by DG. Asymmetric recurrent CA3 weights may then specify probabilities of moving from one place to another. Assuming that the places are mutually exclusive and exhaustive, such models are similar to the Hidden Markov Model which we discuss in section III, inferring a complete distribution of place probabilities $\{P(i)\}_i$ from a vector of inputs and prior transitions.

In contrast, *auto-associative models* based on primate physiology [15], [12], [18], [22] use vector-coded discrete or continuous attractors to represent a single ('unitary') but complex ('coherent') state of the world. For example, each unit in a Hopfield network [12] may correspond to the presence of an object, and the network stores and recalls patterns of many objects which occur together. Thus an episode corresponds to a collection of facts experienced simultaneously, which may include the agent's location and its percepts. To construct stable attractors, the recurrent connections are generally assumed to be symmetric, $w_{ij} = w_{ji}$ (there is no biological evidence for this assumption), and trained so that an energy function,

$$E = \sum_{ij} w_{ij} CA3(i) CA3(j), \quad (2)$$

is optimised at stored episodes. By encoding the world state as a vector, we note that only a single world state,

$$\arg \max_{\{i\}} P(\{i\}_i), \quad (3)$$

can be represented at each point in time – not a probability distribution *over* world states.

A recent model in machine learning, the Temporal Restricted Boltzmann Machine (TRBM) [23], has been proposed [3] as a hippocampal model, and uses a rate-coded, mean-field assumption to approximate whole belief distributions over a vector-coded CA3. This combines aspects of both the above modelling approaches. Our model uses a similar architecture to the TRBM but uses a different inference algorithm, and extends the architecture with a lostness detection and recovery module, mapped to Sub and cholinergic Sep functions. Our model is also similar to that of Samu et al. [19] which uses a CA to correct rate-coded grid cells. As with the TRBM, this model treats CA as recognising and de-noising configurations of inputs. Our model differs from Samu et al. by specifying precise probabilistic semantics; adding recurrent CA3 connections; using sampling instead of rate coding; adding DG, CA1, Sub and Sep; and de-noising the whole input vector rather than just the inputs that are grid cells.

We do not model biological learning in the present study, rather we set weights by hand and with a machine learning algorithm. The aim is purely to illustrate inference and lostness detection and recovery in a pre-trained network. Previous experiments and models [10] suggest that septal ACh is involved with learning: we suggest an additional, complementary function for this signal in lostness recovery.

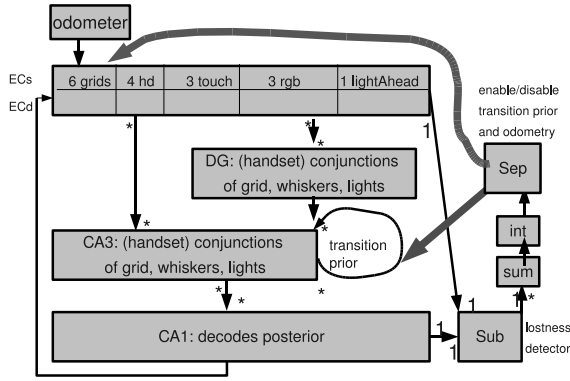


Fig. 1. Hippocampal network architecture used in the model. Connections are labelled with UML notation indicating many-to-many fully connected links ($* \rightarrow *$), one-to-one links ($1 \rightarrow 1$) and many-to-one links ($* \rightarrow 1$). Thick lines are ACh projections, thin lines are glutamate.

II. BAYESIAN FILTERING

The objectives of both auto-associative and spatial sequence memories are combined by a general Bayesian filter with noisy observations, which infers the hidden state of the world x_t at each time from a series of noisy observations $z_t = y_t + \epsilon_t$ of ideal, deterministic sensor states $y_t = f(x_t)$ through the update

$$P(x_t|z_{1:t}) = \frac{1}{Z} \left(\sum_{y_t} P(z_t|y_t)P(y_t|x_t) \right) \times \left(\sum_{x_{t-1}} P(x_{t-1}|z_{1:t-1})P(x_t|x_{t-1}) \right). \quad (4)$$

where Z is a normalising coefficient. Such a filter is further able to provide a de-noised version of the sensors,

$$P(y_t|z_{1:t}) = \frac{1}{Z} \sum_{x_t} P(y_t|x_t)P(x_t|z_{1:t}). \quad (5)$$

which is a form of auto-association, but one which also incorporates prior knowledge from the previous step in a sequence of inferences. If part of the sensor observation relates to place then spatial sequences can be captured, but also sequences of other senses.

We would like to map the Bayesian Filter to a high-level model of the hippocampus circuitry, to allow it to perform both auto-associative and spatial sequence inference. Let us first consider how a purely spatial sequence, place-cell based model may be mapped onto such a filter. In this case, the state of the world is a *single* variable representing the agent's location, rather than a complex coherent set of associations. The i th CA3 cell could each represent the i th location,

$$CA3_t(i) = P(x_t = i|z_{1:t}), \quad (6)$$

and each of the EC and DG cells $z_t(j) = (EC_t, DG_t)$ represents the presence or absence of a sensory feature. As the hidden states are exhaustive and mutual rivals, this special case of the Bayesian filter is a Hidden Markov Model

(HMM). Neural HMM mappings may be subsequently be constructed [17] so that the activation of the CA3 cells can represent the location posteriors. Such mappings allow optimal tracking during navigation, assuming a set of discrete places, but unfortunately cannot account for the presence of other CA3 cells in the hidden state such as those found in biology responding to odours [6] and objects [7].

III. UNITARY COHERENT PARTICLE FILTERING

Auto-associative models, on the other hand, can represent complex hidden states of the world, such as:

isAt(self,location1) AND isAt(edge,myRight) AND isVisible(light) AND isAt(light,location3),

which denotes the agent's own allocentric position at location 1, the egocentric presence of a boundary on the agent's right, the immediate visibility of light, and the position of a light in allocentric space. Each logical term can be represented by the activation of a CA3 cell, and thus the conjunction of terms by the activity of the whole CA3 vector. Such representation is *unitary* because it represents a single state of the world, rather than a pdf over many states. It is *coherent* because it comprises several facts which mutually cohere with one other.

Such a scheme can represent exponentially more unique states than the HMM mapping above, but by forfeiting the ability to represent a belief distribution over states. The presence of cells responding to particular facts is in agreement with the biological CA3. We would thus like to utilise such representations in a hippocampus model, as in auto-associative memories. But we would also like to retain an ability to track sequences of states, including locations, as in the HMM mapping.

We combine these ideas using a variant of the Temporal Restricted Boltzmann Machine (TRBM) architecture [23]. The TRBM assumes weight matrices W_{xz} and W_{xx} , Boolean vectors for the hidden state $x_t(i) \in \{0, 1\}$ and observations $z_t(j) \in \{0, 1\}$, and specifies joint probabilities,

$$P(x_t, x_{t-1}, z_t) =$$

$$\frac{1}{Z} \exp \sum_t (-x_t W_{xx} x_{t-1} - x_t W_{xz} z_t - b_x x_t - b_z z_t), \quad (7)$$

where b are biases that may be viewed as specifying the priors on each of x and z . If we extend each population vector v to a vector $v' = (v, 1)$, appending an additional node which is always on, then the biases may be moved inside the weight matrices:

$$P(x_t, x_{t-1}, z_t) = \frac{1}{Z} \exp \sum_t (-x'_t W'_{xx'} x'_{t-1} - x'_t W'_{x'z'} z'_t).$$

Assumption 1. Unlike the standard TRBM – which uses rate coding – we will assume when inferring x_t that both the sensors and previous hidden state are *observed*. That is, we assume that the previous inference step produced an exact, correct estimate $\hat{x}_{t-1} = x_{t-1}$. Under this assumption, all

links in the TRBM graphical model [23] become effectively directed, and we may write

$$P(x_t|x_{t-1}, y_t) = \frac{1}{Z} \exp \sum_t (-x'_t W_{x'x'} x'_{t-1} - x'_t W_{x'z'} z'_t).$$

A previous mapping of TRBM to hippocampus [3] used the rate-coded, variational approximation to the full posterior,

$$Q(x_t) = \prod_i Q(x_t(i)) \approx P(x_t|z_{1:t}), \quad (8)$$

as in the original TRBM [23]. In contrast, we suggest the following update as a model of CA3 function,

$$\begin{aligned} CA3_t &= \hat{x}_t \leftarrow \arg \max P(x_t|\hat{x}_{t-1}, z_t) \\ &= \{\hat{x}_t(i) = (P(x_t(i)|\hat{x}_{t-1}, z_t) > \frac{1}{2})\}_i \end{aligned} \quad (9)$$

This is a deterministic update which may be viewed as the zero-temperature limit of an annealed Gibbs sampler. As we update to a unitary coherent state at each time step, all nodes are Boolean valued and the conditioning on \hat{x}_{t-1} becomes definite, satisfying Assumption 1. Therefore the cells $\hat{x}_t(i)$ become independent over i and may be updated locally and individually, using sigmoidal threshold-and-fire units with

$$\begin{aligned} P(x_t(i)|\hat{x}_{t-1}, z_t) &= \frac{P(x_t(i)|\hat{x}_{t-1}, z_t)}{P(x_t(i)|\hat{x}_{t-1}, z_t) + P(\neg x_t(i)|\hat{x}_{t-1}, z_t)} \\ &= \text{sig}(W_{x'x'}(i)\hat{x}_{t-1} + W_{x'z'}(i)z'_t), \end{aligned} \quad (10)$$

where $W(i)$ denotes the i th row of matrix W , $\text{sig}(x) = (1 + \exp(-x))^{-1}$ is the sigmoid function, and the output of eqn. 10 is used with a threshold of $\frac{1}{2}$ to set the Boolean vector \hat{x}_t as in eqn. 9. This tracking method may be viewed as a particle filter (PF) with a single particle, and the zero-temperature annealed Gibbs distribution as its proposal distribution. The particle has a complex, vector-coded coherent state, hence we call the algorithm a *unitary coherent particle filter (UCPF)*. Unitary particle filters have previously been proposed as a model of external behaviour [5] – here we give an internal neural implementation which could give rise to such behaviours.

We map the noisy inputs to the combined ECs and DG, where the DG activations are functions of the ECs activations, $z_t = (ECs_t, DG_t(ECs_t))$. Finally we map the estimated de-noised output to ECd, $\hat{y}_t = ECd_t$. Each neural population is a Boolean vector at each discrete time step t , which may be viewed as an abstracted spike or absence of spike within a theta cycle.¹

We will later describe subpopulations within each region. For example, ECs contains a subpopulation of grid cells, which in turn is indexed by 2D row and column coordinates r, c . We use chains of functions to index such populations, for example, ECs denotes the whole ECs population, $ECs(\text{grid})$ denotes the subpopulation of grid cells, and $ECs(\text{grid})(r, c)$ denotes a particular grid cell at row r

¹The precise timing of theta oscillations and phase-precession within them in hippocampus have attracted much modelling interest but we do not consider them in the present model.

and column c . We may also write $z_t(ECs)(\text{grid})(r, c)$ to emphasise that ECs is itself a subpopulation of the input z_t .

IV. LOSTNESS DETECTION AND RECOVERY

A major problem with particle filters having small numbers of particles – and especially a unitary particle – is getting lost. PFs approximate the posterior at each step by a small number of samples, and the UCPF approximates it with a single sample. If the sample set drifts away from the true state, then it becomes difficult or impossible to regain tracking. An approach to dealing this problem used in robotics [13] is to monitor the performance of the filter and heuristically detect when tracking is lost, for example by thresholding a moving average,

$$e_t = \alpha \sum_{i \in -o} (z(i)_t - \hat{y}(i)_t)^2 + (1 - \alpha)e_{t-1}, e_0 = 0, \quad (11)$$

where the sum ranges over the non-odometric inputs only (discussed further in section V). If tracking is lost, then priors (and odometry-dependent sensors, discussed in section V) should be disabled and an alternative proposal distribution used, based only on immediate sensor likelihoods.

We hypothesise that the Subiculum-Septum circuit performs such monitoring. We model CA1 as performing partial decoding of the CA3 hidden state into the de-noised, posterior beliefs about the sensors, which are then relayed or fully decoded in ECd. Sub is then well-placed to compare the de-noised CA1 information against the original ECs input, receiving one-to-one connections from both regions. If the posterior CA1 decodings are sufficiently different from the ECs likelihoods for a period of time, this indicates loss of tracking. The cholinergic projections from Sep, activated via Sub, are especially well-placed to disable the CA3 priors as described above, as they are known [10] to disable the recurrent connections in CA3.

V. TREATING ODOMETRY AS NOISY GPS OBSERVATIONS

Including odometry information in the sensory input z_t is somewhat problematic for an HMM-like architecture. In the Bayesian filter, inputs must be independent of one another conditioned on the hidden states. One allowable type of place sensor would be a *noisy GPS* system, which gives noisy estimates $(\theta_t + \epsilon_t)$ of pose θ_t , being a vector comprised of 2D position and angular heading. However animals do not possess such a sensor, rather they must integrate a series of noisy differential odometry measurements $\Delta\theta_t + \epsilon_t$ to obtain a sense of location. If summed naively, a sequence of heavily error-correlated estimates of pose is obtained. Such correlation violates the Bayesian filter requirements.

To force the place and heading observation inputs to respect HMM semantics, our ECs assumes that the previous hippocampal output estimate of pose is so accurate as to be perfect, so when added to the latest differential odometry gives an input equivalent to a noisy GPS observation, appropriate for the Bayesian filter. This is a similar assumption as used in the UCPF update: assuming that the previous inference was perfect. Of course – as with the UCPF




Fig. 2. Plus-maze environment used in the demonstration.

update – the noisy-GPS assumption will break when tracking is lost. We already have a system for lostness detection described above. So its output can be reused to completely disable the noisy GPS input until tracking is regained, in addition to its role in disabling the prior in the UCPF itself.

To provide the capability of switching between full inference (using odometry, other senses, and CA3 priors) and inference from odometry-independent sensors only, care must be taken in setting up the TRBM and its biases in particular. We partition the input vector z_t into a pure sensory component $z_t(-o)$ which is independent of all odometry, and a component dependent on odometry, $z_t(o)$. These vectors are extended to $z'_t(-o)$ and $z'_t(o)$ respectively, by appending a 1 at the end, to provide separate bias terms. We partition the de-noised version y_t similarly. (This is an orthogonal partition to $(z_t(EC), z_t(DG))$, as both ECs and DG will contain both odometry dependent and independent variables.) We then define a global bias on the hidden TRBM nodes,

$$b_x = \text{sig}^{-1}(P(x)), \quad (12)$$

yielding the first-order priors on the hidden states under the Boltzmann distribution,

$$P(x) = \frac{1}{Z} \exp b_x. \quad (13)$$

Next we define separate sets of weights for the three information sources, conditioning on the global bias, so that the joint posterior is

$$\log ZP(x_t, x_{t-1}, z_t) = \sum_t (b_x + x'_t W_{x'x'} x'_{t-1} + x'_t W_{x'z'(-o)} z'(-o)'_t + x'_t W_{x'y(o)} y(o)'_t)$$

As addition in the weight domain corresponds to distribution fusion in the probability domain, the terms in this sum correspond exactly to the factors in

$$P(x_t, x_{t-1}, z_t) = \frac{1}{Z} P(x) P(x_{t-1}|x_t) P(z(o)_t|x_t) P(z(-o)_t|x_t). \quad (14)$$




Fig. 3. The first 500 steps around the maze in the simulation. The x and y locations shown here correspond to the top two graphs in figs. 5 and 6.

It is necessary to define the global bias first before learning the other weights, so that the bias terms in the other weight matrices exclude the contribution from the global prior and thus form likelihood terms which may be fused correctly with it. In particular this then allows the odometry-dependent and recurrent prior information to be removed using

$$P(x_t, z(-o)_t) = \frac{1}{Z} \exp \sum_t (b_x + x'_t W_{x'z(-o)} z'(-o)'_t). \quad (15)$$

The terms in this exponent then correspond, as desired, to

$$P(x_t, z(-o)_t) = \frac{1}{Z} P(x) P(z(-o)_t|x_t). \quad (16)$$

We may write equation 15 as

$$P(x_t, z(-o)_t) = \frac{1}{Z} \exp \sum_t (b_x + x'_t (0 \times W_{x'x'}) x'_{t-1} + x'_t W_{x'z'(-o)} z'(-o)'_t + x'_t W_{x'z(o)} (0 \times z(o)'_t)), \quad (17)$$

where $(0 \times W_{x'x'})$ is equivalent to disabling the recurrent CA3 connections, and $(0 \times z(o)'_t)$ is equivalent to inhibiting all activity in ECs and DG areas receiving odometry-dependent input. Importantly, both of these disabling assume that the source population includes an always-on bias node. In contrast, the effect of the global bias may easily be moved into the firing threshold for CA3 cells.

VI. PLUS-MAZE EXAMPLE

We illustrate the UCPF with Subicular lostness detection and recovery using a simple microworld based on a standard environment used in hippocampal research [1]. Fig. 2 shows a simulated plus-maze, consisting of 13 discrete locations. An agent can move between these locations, and rotate to face the four discrete compass headings. The north and east arms of the maze have coloured posters placed some distance behind them, to provide visual cues when facing in those directions. Two components of the state of the world are the agent's own place, $place \in [0 : 13]$, and heading, $hd \in \{N, E, S, W\}$. In order to illustrate more complex coherent world states, lights are placed at the ends of each arm. One light is on at each time step. When the agent reaches the end of the arm with the active light, it is




Fig. 4. Walsh-like basis functions used as simplified grid cell receptive fields. Each of the six cells shown here has a Boolean valued receptive field ranging over the 2D space occupied by the maze.

switched off and another light, from a fixed 4-light sequence, is activated. So a third component of the world state is, $light \in \{N, E, S, W\}$, describing which of the four lights is currently on. The complete world state is thus the triple, $(place, hd, light)$.

The present model is concerned only with hippocampal inference, so no learning is modelled, and the agent moves in a random walk. The task is to infer the world state from its noisy sensors. The model network is shown in fig. 1 and weight types summarised in table I.

A. Entorhinal inputs

The agent has *touch* sensors detecting the presence of boundaries on its *left*, *right* and *forward*; a *lightAhead* sensor which is active when facing the active light; and *red*, *green* and *blue colour* sensors, active when facing a coloured poster. The values of these sensors are all Boolean and are placed into ECs. For example, we have cells $ECs(touch)(left)$ and $ECs(colour)(red)$.

The other two senses are history-dependent, based on odometry information under the noisy GPS assumption. First, medial ECs is known to encode position using grid cells [9] rather than unique place cells – we model a population of six simplified grid cells, $ECs(grid)$, having Boolean valued receptive field functions of 2D location (denoted by r, c , for row and column),

$$EC(grid)(n, 1) = (r \bmod 2^{n+1} < 2^n)(c \bmod 2^n < 2^{n-1})$$

$$EC(grid)(n, 2) = (c \bmod 2^{n+1} < 2^n)(r \bmod 2^n < 2^{n-1}),$$

for $n \in \{1, 2, 3\}$, as shown in fig. 4. Secondly, there are four head direction cells $EC(hd)(dir)$ responding to $dir \in \{N, S, E, W\}$ headings.

The effects of odometry on the ECs grid and head direction cells are modelled only functionally. The previous location estimate is read from ECd, odometry is added, and the resulting locations is encoded and placed in the ECs population,

$$\begin{aligned} EC(grid)_t &= EC(grid)_{t-1} + odom_t \\ EC(hd)_t &= EC(hd)_{t-1} + gyro_t, \end{aligned} \quad (19)$$

where the addition operator is here defined as decoding the grid or head direction cells, adding the Cartesian odometry (*odom*) or angular gyroscopy (*gyro*), then re-encoding new grid or head direction activations. (See [16], [8] for neural-level models of this process.)

B. Dentate Gyrus

DG has handset weights, $W_{EC \rightarrow DG}$, which form a sparse encoding of entorhinal feature combinations including: combinations of head direction and light ahead; left and right touches together; left, right and center touches together; and the 13 discrete places in the maze (given by combinations of grid cell activity). We denote these by, for example $DG(hd \otimes lightAhead)(N)$ for the head direction and lightAhead combination responding to north facing and the light ahead; $DG(touch \otimes touch)(left)(right)$ for the combination of two boundaries on the left and right. We write $DG(place) = DG(grid \otimes grid \otimes grid \otimes grid \otimes grid)$, defining place as a combination of grid activations to reduce notation.

C. CA3

The *semantics* (not the weights) of the CA3 cells are set by hand. We specify CA3 cells responding to each of the 13 discrete locations, $CA3(place)$; combinations of place and head direction, $CA3(place \otimes hd)$; the state of the light sequence $CA3(light)$; and combinations of location and light sequence, $CA3(place \otimes light)$ (Thus the micro-domain illustrates CA3 cells responding to simple facts about the world, and to combinations of those facts, cf. [7]).

The weights to CA3 are partitioned into $(W_{x'z'(-o)}, W_{x'z'(o)}, W_{x'x'})$ as in equation 15, where $\hat{x} = CA3$, $z(o)$ is the collection of EC and DG cells dependent on odometry, and $z(-o)$ is the collection of EC and DG cells that are independent of odometry.

The weights are fit using a version of the wake-sleep algorithm [11] as follows, used for convenience rather than as a biological model of learning. World states $(place, hd, light)$ and noisy sensors (including a genuine noisy GPS training-only sensor) from a random walk of 30,000 steps are recorded. Using the prescribed CA3 semantics, ideal CA3 activations are determined directly from the world states, and global biases b_x computed from their occurrence frequencies. Populations in EC and DG are pooled to form $z(o)$ and $z(-o)$, according to whether they include odometry information. In wake steps, ECs_t , $CA3_t$ and the previous $CA3_{t-1}$ are all thus directly observed, and Hebbian learning is performed, conditioned on b_x . In sleep steps, predictions of the next CA3 and next EC are made from the ideal CA3, and anti-Hebbian learning is performed, conditioned on b_x .²

D. CA1 and ECd

ECd consists of the same population types as ECs (being the deep layers of the same cortical columns as the ECs units), but stores de-noised sensor values \hat{y}_t .

²We found that it was important to train all three weight matrices together, in the same wake-sleep cycles. Theoretically they could be trained separately, all conditioned on the same global bias, to learn the likelihoods in equation 14. However in practise the likelihoods are never learned exactly, and variations in weight strength can occur during separated training, which combine to give inaccurate predictions when fused together.




Fig. 5. Results with septum lesioned. Top two graphs show the actual (dark) and inferred (light) x and y locations in the plus maze respectively, during 3000 steps of a random walk. These are the same x, y coordinates shown in 2D in fig. VI. The walk consists roughly of journeys to arms N,E,W,N,S,N,W,S. Tracking performs well until $t = 500$ and then becomes lost – i.e. the actual and inferred lines diverge – until regained by chance alone at $t = 800$. Tracking is lost again at $t = 1000$ and regained by chance at $t = 1300$. It is lost a third time at $t = 1700$ and regained by chance around $t = 1900$. It is lost a fourth time at $t = 2600$ and regained at $t = 2700$. The third, heading, graph shows actual (dark) and inferred (light) head directions, which have four discrete values for N,E,S,W and also diverge during lostness. The lower graph shows Subicular immediate and averaged activation. This activation is seen to be highest at times when the agent is lost, i.e. when there is divergence between the actual and inferred locations. With Sep lesioned the Sub output has no effect on the inference.

A naive decoding method would be to map the unitary TRBM posteriors $\hat{y}(EC)_t$ directly to ECd, using

$$ECd = \hat{y}'_t(EC) = (\text{sig}(W_{x'y'}^T \hat{x}'_t) > \frac{1}{2}), \quad (20)$$

and ignoring the $\hat{y}(DG)_t$ components. This can be done, but the results are poor because the TRBM assumes the sensor units to be independent of one another. However we know that some of the y_t subpopulations – place, heading and whisker combinations – in fact always have exactly one active unit, as they represent mutually rival and exhaustive hypotheses. Much of this structure is contained in the $\hat{y}(DG)_t$ components that are discarded in the naive decoding.

We can exploit the structure by using a two stage decoding processes, mirroring the two-stage (EC,DG) encoding. Our CA1 contains all the same population types as our DG, representing combinations of EC features. For example, $CA1(\text{place})$ is the de-noised version of $DG(\text{place})$. (It also contains populations mirroring the EC populations that do not benefit from mutual rivalry, namely *lightAhead* and *colour*.) The mutual rival features in each population s are then sampled using the winner-take-all rule,

$$CA1(s)(i) \leftarrow (P(CA1(s)(i)) > P(CA1(s)(j)); \forall j \neq i)$$

(such an update may be implemented neurally using winner-take-all feedback.)

The composite features in DG and CA1 are all conjunctions of EC features, so it is simple to handset weights




Fig. 6. Results with septum intact. When the moving average subicular output reaches a threshold, it causes Sep to fire (shown as spikes on the lower plot), releasing ACh into CA3 and ECs. This disables the recurrent prior connections in CA3 and inhibits all grid cell activity in ECs, causing inference to proceed using sensory likelihoods only for the next step. This sometimes allows the network to jump to the correct solution. If recovery fails, then subiculum continuous to fire until recovery succeeds. For example around $t = 1100$ there are two long Sep bursts before recovery occurs; in contrast at about $t = 2500$ there are two errors which are corrected almost immediately. Comparing with fig. 5 it can be seen that periods of lostness begin at the around the same times, but last for shorter periods due to the corrections.

$W_{CA1 \rightarrow ECd}$ to decode them back into the EC basis in ECd. (An ECd cell is on if any of its conjunctive parents is on.) In the case of non-rival populations, CA1 activations are simply relayed to the corresponding ECd populations.

E. Subiculum and Septum

We postulate that Subiculum consists of four parts: a decoder, comparator, an accumulator and an integrator. They are modelled functionally, not neurally. The decoder performs the same computation as ECd, to obtain a copy of the de-noised signals from the partially-decoded CA1. The comparator has cells in one-to-one correspondence with the sense cells in ECs and ECd (excluding place/grid and head direction cells), and which fire when the corresponding ECs and ECd cells differ. (It receives one-to-one input from ECs and decodes input from CA1 in the same way as ECd.) The accumulator computes the integer sum of these differences. The integrator tracks the exponentially weighted moving average of the accumulator to give an overall indication of tracking error. (Accumulation and integration could perhaps take place along the fornix.) When the integrator exceeds a threshold, a single-unit cholinergic Septum unit is activated. The ACh modulation then projects from Sep to ECs and to CA3, having the two effects described by equation 17.

F. Results


Results of the model running in the plus maze simulation are shown in figs. 5 and fig. 6, with Septum lesioned and intact respectively. Without the septum, there is no recovery from loss of tracking (except when the agent moves by

TABLE I

CONNECTIVITY IN THE MODEL. H=HAND SET WEIGHTS, L=LEARNED WEIGHTS, T=TRANSPOSE OF LEARNED WEIGHTS, R=ONE-TO-ONE RELAY, I=IMMEDIATE INPUTS, O=ODOMETRIC INPUTS. CONNECTIVITY IS SHOWN IN FIG. 1. SUB RECEIVES INPUTS FROM (ECs,CA1).

pop.region	ECs	DG	CA3	CA1	ECd	Sub
grid	o				h	r,h
hd	o			t	r	r,r
touch	i				h	r,h
colour	i			t	r	r,r
lightAhead	i			t	r	r,r
place		h		t		
hd \otimes lightAhead		h				
touchCombis		h		t		
place			1			
place \otimes hd			1			
light			1			
light \otimes hd			1			

Fig. 7. Mean errors made about places. (Error bars are vanishingly small as a large amount of simulation data is used.)




chance to the location that it believes in). For example, at $t = 500$, $t = 1000$, $t = 1700$ and $t = 2700$ the ground truth and estimate lines diverge and do not rejoin for several hundred steps. With Septum intact, tracking is able to recover following the detection of failure and deactivation of CA3 recurrent connections via septal ACh. At each of the time above, the Sub-Sept circuit becomes active and the priors are removed. This often leads to burst of activation, as the likelihood-driven estimates may also be wrong, but eventually the correct location is found, the error falls, then tracking returns to normal.

The error introduced by the Sub-Sept circuit is less than the tracking error that it corrects, as shown in fig. 7. We also show the error in a purely likelihood based model, which has both CA3 priors and odometry permanently removed. The transitions priors alone do improve the place error over the likelihood-only model; but a further significant improvement is gained by using the Sub-Sept circuit. Fig. 8 shows the decreased error in the odometry independent sensor denoising when the Sub-Sept circuit is active.

VII. DISCUSSION

The UCPF model combines and extends ideas from Samu et al. [19] and the TRBM hippocampal model of Becker and Hinton [3]. Samu et al. use a hidden CA population (assumed to encompass all of DG, CA3 and CA1) to modulate and

Fig. 8. Discrepancies between input and de-noised sensors. These may be due to genuine de-noising, or to lostness producing incorrect inferences. (Error bars are vanishingly small as a large amount of simulation data is used.)



correct activity in a single EC population of grid cells. Our model also uses hippocampus to modulate the grid cells. However it makes a more detailed distinction between the ECs input and ECd output populations, and adds recurrent temporal priors. Under the UCPF's probabilistic semantics, ECs and ECd are well-defined representations of the noisy input z_t variables and the estimated denoised \hat{y}_t respectively. This allows us to interpret both grid and sensor deep layers as de-noised versions of the superficial layers. These semantics are drawn from the TRBM model, which the present work extends with a more detailed mapping to hippocampus, and the new sampling and lostness inference algorithms. We have used the temporal sequence semantics of the TRBM (i.e. recurrent CA3 connections as priors), for example to learn the places and sensations that typically follow a given place or sensation. This draws together the notion of asymmetric weights as spatial transitions between place cells – used in purely spatial models – with the notion of the weights as compatibilities between variables – used in auto-associative models.

The wake-sleep algorithm in particular learns whatever weights give the best model of the temporal sequence of vector-valued data, so will incorporate both temporal and associative factors. (We note that if each input vector was held constant for a time tending to infinity, then the recurrent weights would tend towards symmetry, thus the symmetric auto-associative models are a limiting case of temporal sequence models.) The wake-sleep algorithm was used here merely to set the computational weights, and not as part of the biological model. However it has previously been suggested as a theory of on-line learning, with learning and anti-learning triggered by tonic ACh in the theta cycle. (In contrast, the septal signal here is the phasic ACh level.)

Three biological connections were not modelled: ECs \rightarrow CA1, the CA3 \rightarrow DG back-projection, and Sub \rightarrow EC. We speculate that ECs \rightarrow CA1 could play a role in learning: CA1 must learn to decode CA3 back to the EC basis, so this connection could act as the supervised training signal. The CA3 \rightarrow DG back-projection might likewise play a role in learning, perhaps part of a wake-sleep-like algorithm where DG sleep states are sampled given CA3 (though full

wake-sleep would also require sampling from ECs).

The mapping would have been more elegant if the Sub comparator received input from ECs and ECd, instead of ECs and CA1. The present model requires Sub to perform the same decoding from the CA1 basis to the EC basis, as well as the comparison itself. An early version of the model assigned the full decoding task to CA1, which then simply relayed all information to ECd, and allowed Sub to act as a simple one-to-one comparator. While this is a possible alternative mapping, we found that using an extra layer of partial decoding produced better results, as it allowed mutual rivalry to be expressed and inferred in CA1 during the partial decoding step. It may be possible to move all such rivalry functionality into CA3, allowing CA1 to act as a pure decoder and relay, though this would make the function of ECd rather redundant, and would also require the presence of grid cells in CA1, rather than the CA1 place cells known in the biology.

The biological mechanism for resetting grid cells is unclear – the model merely requires that lostness detection causes reset. We modelled the Sep→ECs connections as performing this role, but it is possible that the Sub→ECs projection could perform a similar function instead. There is some evidence for septal ACh activating inhibitory cells in EC [24], but it is also known to make ECd cells ignore their inputs and maintain their firing patterns [?].

ACh is well-known to play an important role in learning in CA3, as well as switching CA1 from responding mostly to CA3, to responding mostly to ECs. Learning and lostness go together in many circumstances, namely those where the agent is in a novel environment, so is both lost and needing to learn. During novelty, ACh would both disable the priors and trigger learning. However we might not always want to disable odometry during novelty. In contrast, in cases of being lost in a known environment, the removal of both priors and odometry is desired but without learning. Whilst lostness and novelty often go together, they are different concepts, and might be represented by combinations of modulators, for example dopamine is also known to play a role in hippocampal learning [14], but its interactions with ACh are not yet understood.

We have seen that the TRBM uses vector coding, but can learn weights modelling both the temporal and auto-associative aspects of the input vectors; we then constructed a unitary coherent particle filter algorithm operating on the TRBM structure and illustrated how tracking failure may be detected and corrected in a UCPF by monitoring input/output difference and removing the effects of transition and priors and odometry when lost. We have showed how to map the UCPF onto the hippocampal EC-DG-CA3-CA1-EC loop, and how tracking correction could be implemented by the Subiculum and Septum, using ACh to deactivate CA3 recurrent connections.

REFERENCES

- [1] E. Tabuchi A.B. Mulder and S.I. Wiener. Neurons in hippocampal afferent zones of rat striatum parse routes into multi-pace segments during maze navigation. *European Journal of Neuroscience*, 19(7):1923–32, 2004.
- [2] A. Arleo and W. Gerstner. Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. *Biological Cybernetics*, 83:287–299, 2000.
- [3] S. Becker and G. Hinton. Caching and replay of place sequences in a temporal restricted Boltzmann machine model of the hippocampus. In *Cosyne (abstract only, Poster II-56)*, 2007.
- [4] Ricardo Chavarriaga, Thomas Strasslin, Denis Sheynikhovich, and Wolfram Gerstner. A computational model of parallel navigation systems in rodents. *Neuroinformatics*, 3:223–241, 2005.
- [5] Nathaniel Daw and Aaron Courville. The rat as particle filter. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 369–376. MIT Press, Cambridge, MA, 2008.
- [6] E.R. Wood, P.A. Dudchenko, and H. Eichenbaum. The global record of memory in hippocampal activity. *Nature*, 397:613–616, 1999.
- [7] E.T. Rolls, J. Xiang, and L. Franco. Object, space and object-space representations in the primate hippocampus. *Journal of Neurophysiology*, 94:833–844, 2005.
- [8] Alexis Guanella and Paul F. M. J. Verschure. A model of grid cells based on a path integration mechanism. In Stefanos D. Kollias, Andreas Stafylopatis, Włodzisław Duch, and Erkki Oja, editors, *ICANN (1)*, volume 4131 of *Lecture Notes in Computer Science*, pages 740–749. Springer, 2006.
- [9] T. Hafting, M. Fyhn, S. Molden, M.B. Moser, and E.I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436:801–6, 2005.
- [10] ME Hasselmo, E Schnell, and E Barkai. Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region ca3. *Journal of Neuroscience*, 15:5249–62, 1995.
- [11] Geoffrey E. Hinton, Peter Dayan, Brendan J. Frey, and Radford M. Neal. The wake-sleep algorithm for unsupervised neural networks. *Science*, 268:1158–1161, 1995.
- [12] J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc. National Academy of Science*, 79:2554–2558, 1982.
- [13] S. Lenser and M. Veloso. Sensor resetting localization for poorly modelled mobile robots. In *Proceedings of the International Conference on Robotics and Automation*, 2000.
- [14] J. Lisman and A.A. Grace. The hippocampal-VTA loop: Controlling the entry of information into long-term memory. *Neuron*, 46:703–13, 2005.
- [15] D. Marr. Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society B*, 262:23–81, 1971.
- [16] Bruce L. McNaughton, Francesco P. Battaglia, Ole Jensen, Edvard I Moser, and May-Britt Moser. Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, pages 663–678, 2006.
- [17] Rajesh P. N. Rao. Bayesian computation in recurrent neural circuits. *Neural Comput.*, 16(1):1–38, 2004.
- [18] E.T. Rolls, S.M. Stringer, and T.P. Trappenberg. A unified model of spatial and episodic memory. *Philosophical Transactions of the Royal Society B*, 269:1097–1093, 2002.
- [19] D. Samu, P. Eros, B. Ujfalussy, and T. Kiss. Robust path integration in the entorhinal grid cell system with hippocampal feedback. *Biological Cybernetics*, 101(1):19–34, 2009.
- [20] HE Scharfman. The CA3 backprojection to the dentate gyrus. *Progress in Brain Research*, 163:627–37, 2007.
- [21] P.E. Sharp. Computer simulation of hippocampal place cells. *Psychobiology*, 19:103–115, 1991.
- [22] S. Kali and P. Dayan. The involvement of recurrent connections in area ca3 in establishing the properties of place fields: a model. *Journal of Neuroscience*, 20(19):7463–7477, 2000.
- [23] Graham W. Taylor, Geoffrey E. Hinton, and Sam T. Roweis. Modeling human motion using binary latent variables. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1345–1352. MIT Press, Cambridge, MA, 2007.
- [24] Zhaoyang Xiao, Pan-Yue Deng, Chuanxiu Yang, and Saobo Lei. Modulation of gabaergic transmission by muscarinic receptors in the entorhinal cortex of juvenile rats. *J Neurophysiol*, 109:659–669, 2009.