On the epistemic role of hippocampal cells: the case of splitter cells

Naomi Chaix-Eichel[1,2,†], **Snigdha Dagar**[1,2,†], **Frédéric Alexandre**[1,2,‡], **Thomas Boraud**[1,‡], **Nicolas P. Rougier**[1,2,‡,*]

***For correspondence:** nicolas.rougier@inria.fr (FMS)

[†]These authors contributed equally to this work
[‡]These authors also contributed equally to this work

[1] Univ. Bordeaux, CNRS, IMN, UMR 5293, F-33000 Bordeaux, France; [2]Univ. Bordeaux, CNRS, Bordeaux INP, INRIA, LaBRI, UMR 5800, F-33400 Talence, France

Abstract Over the past decades, the hippocampal formation has undergone extensive study leading researchers to identify a vast array of cells with functional properties (place cells, splitter cells, etc). In the present work, we aim at investigating whether the activity of those cells derives from the anatomy and inner circuitry of the hippocampal formation or derives instead from the actual behavior of the animal. To do so, we simulated an agent navigating inside an 8-shaped track, making alternating choices (T-maze alternating task). We designed a random network, based on the reservoir computing paradigm, that processes distance-based sensors and outputs a direction change (constant speed). Despite its simplicity, the model successfully solved the task while bearing no structural similarity with the hippocampal formation. We subsequently followed the comprehensive and recent review on splitter cells by ***Duvelle et al.*** (***2023***), and applied the exact same analysis onto the activity on the cells composing our model. We were able to identify splitter cells (as well as place cells, head direction cells and decision cells) and confirm a significant portion of the theoretical hypotheses of Duvelle et al. regarding splitter cells. Beyond these results, this work strongly suggests that the activity of such cells originates from the actual behavior of the agent as opposed to any structural or anatomical origin: any model doing the same task might exhibit the same cell activity. From a broader point of view, this work questions the epistemic role of such cells in our understanding of the hippocampal formation.

Introduction

The hippocampal formation has been heavily studied in the past few decades (***Maguire et al., 1996***; ***Buzsáki and Moser, 2013***; ***Hartley et al., 2014***; ***Moser et al., 2017***) and researchers have since then established a huge repertoire of cells displaying very specific properties. In his book "The Brain from Inside out", ***György Buzsáki*** (***2019***) named a few of them in a footnote (page 356): *place cells, time cells, grid cells, head direction cells, boundary vector cells, border cells, object cells, object vector cells, distance cells, reward cells, concept cells, view cells, speed cells, memory cells, goal cells, goal direction cells, splitter cells, prospective cells, retrospective cells, island cells, ocean cells, band cells, object cells, pitch cells, trial order cells, etc*. Each and every cells have been characterized in terms of correlation between their activity and some combination of high level property involving space, time and internal state. This is the case for place (cell fires when animal is in a specific place), head direction (cell fires preferentially when head is pointing toward a specific direction), time cells (cell fires at successive moments), etc. Facing such huge repertoire, one may legitimately wonder how the brain orchestrates all these information and makes use of it to ensure survival. Part of the

answer is given by György Buzsáki himself in the same footnote where he further explained that *physiological attributes of neurons in the hippocampal-enthorinal system … might be explained by the apparent distinctiveness of a few sets of continuous distributions*. In other words, even though we can and we do observe these cells *in vivo*, they might be a simple epiphenomenon: their activity might be correlated to some unknown latent variables. In this work, we make the hypothesis that these latent variables are ascribed in the task itself rather than the structure (brain), hence supporting a behavioral signature for these cells. In order to investigate further, we propose to use a computational approach and check if we observe the same set of cells.

A classical method for the design of models in cognitive computational neuroscience is to take direct inspiration from the brain anatomy in order to specify the architecture of a model in terms of structures, populations and connectivity, with varying levels of accuracy and plausibility. The model can be subsequently fitted to the data such as to exhibit the same behavior as its biological counterpart: considering an animal performing a specific task, the model is expected to display a comparable level of performance when evaluated on the same task (possibly adapted). When successful, this classical method is able to provide explanations on the precise role of this or that structure / population / connection during the completion of the task. This approach corresponds actually to the implementational / physical level of the tri-level analysis of David Marr ***Marr*** (***1982***). A canonical example for such approach is the hippocampal formation whose precise anatomy has been known for a few decades in several species (mostly rodents and primates). Any biologically plausible model of the hippocampus is thus expected to be composed of the dentate gyrus (DG), CA3, CA1 and the subiculum, and can be possibly complemented with the enthorinal cortex (EC) since it represents the main input / output structure of the hippocampal formation (***Andersen, 2007***). If we now turn back to our initial question regarding the nature of this huge repertoire of cells, it comes as no surprise that any sufficiently detailed model of the hippocampal formation will (theoretically) exhibit the same repertoire of cells: same causes, same consequences. If the structure causes the function, then the model will possess such cells (even though things can be a bit more complicated because computational modeling can be tricky). Consequently, such kind of model cannot answer the question regarding the structural or behavioral nature of these cells.

An alternative approach is to target the behavior of the animal, without any regards to the actual structure of its brain, and observe *a posteriori* if such cells can be found. Ideally, we want to design a minimal model that is able to solve the task, without using any strong architectural hypothesis. This is exactly what an echo state network is: a recurrent neural network sparsely connected made of a single reservoir of neurons. This is incidentally very similar to the notion of a *huge reservoir of unique neuronal patterns* proposed by ***György Buzsáki*** (***2019***). In this work, we propose to study a reservoir model that is trained to perform a navigation task (alternating left right in a T-maze) and measure a posteriori for the existence of such cells. More precisely, we decided to focus our study on splitter cells about which ***Duvelle et al.*** (***2023***) published recently an extended review of the literature. Place cells are a type of hippocampal neurons that fire when an animal occupies a specific location in the environment, helping in building a cognitive representation of a specific location in space (***O'Keefe and Dostrovsky, 1971***). Within the population of place cells, there exists a subset known as "splitter cells" (***Wood et al., 2000***), that not only encode the animal's current location similar to conventional place cells, but also encode information about the past or the future trajectory. In other words, splitter cells demonstrate different firing patterns depending on the animal's origin or intended destination, even if the animal is in the same physical location. Our objective is to show that the internal state of some units within the model exhibits similar characteristics to biological splitter cells. By conducting a theoretical analysis similar to that performed by ***Duvelle et al.*** (***2023***), we examined various hypotheses and their implications within the reservoir computing model. The results of our study demonstrate that the majority of these hypotheses are indeed validated by the reservoir computing model and support the behavioral origin of such cells.

Methods

Task presentation

The class of tasks called spatial alternation has been widely used to study hippocampal and working memory functions (***Frank et al., 2000***). For the purpose of our investigation, we simulated two continuous versions of the task, wherein an agent must navigate through an 8-shaped track (see Figure 1). The first version corresponds to the classical alternating task, wherein the agent must alternate between right and left choices at a decision point, then returns to the central corridor, essentially following an 8-shape trace while moving. The second variation, termed the half-alternating task, deviates slightly. Here, the agent alternates its choices only half of the time, resulting in sequences of right-right and left-left alternations. This 8-maze environment offers an ideal setup for observing splitter cells when the agent enters the central corridor, because the agent can have various past trajectories leading to the central corridor and different future trajectories upon exiting it. Agent alternation tasks are widely believed to necessitate a form of working memory to recall the previous choices for effective alternations. Our findings show that the Echo State Network (ESN) (***Jaeger, 2001***) is capable of learning agent tasks without constructing any models based on structural hypotheses or explicit working memory.

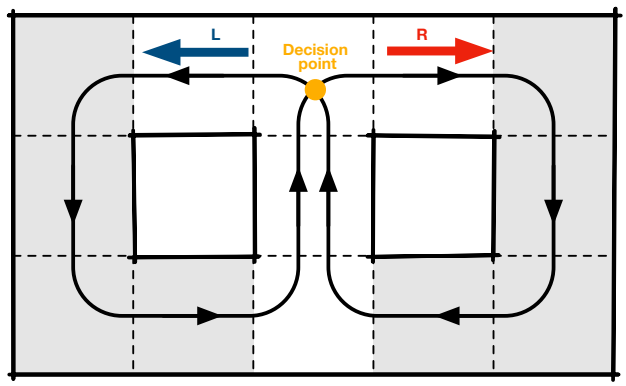

Figure 1. An expanded view of a T-Maze. At the decision point in yellow, the agent has to decide whether to go right (red) or left (blue).

Modeling framework

Solving an alternating task necessitates the presence of a functional working memory, as the agent must retain information about its previous direction, such as turning right in the prior loop, to inform its decision to turn left at the subsequent decision point. Reservoir computing emerges as a promising solution for it, thanks to its fading memory (***Pascanu and Jaeger, 2011***). This attribute enables the system to retain past experiences, facilitating their recall for subsequent decisions. Consequently, our model consists of a reservoir computing network of type Echo State Network (ESN) (***Jaeger, 2007***) that controls the movement of the agent solving a continuous navigation task in the 8-maze of Figure 1. An ESN is a recurrent neural network (called reservoir) composed of randomly connected units, associated with an input and an output layer. Inputs are projected into the reservoir, a non-linear, high-dimensional space, allowing the integration of information over time and space. Only the output neurons, referred as the readout neurons, are trained (in red in Figure 2). The neurons have the following dynamics:

$$\frac{1}{\alpha}\frac{d\mathbf{x}}{dt} = -\mathbf{x} + \tanh(W.\mathbf{x} + W_{in}.\mathbf{u}) \quad \text{and} \quad \mathbf{y} = W_{out}.\mathbf{x} \tag{1}$$

where $\mathbf{x}$, $\mathbf{u}$ and $\mathbf{y}$ represent the reservoir states, input, and output. W, W_{in}, and W_{out} are weight matrices, while $tanh$ refers to the hyperbolic tangent function. α refers to the leak rate, a crucial parameter of the ESN that plays a central role in controlling the memory and timescale of the network's dynamics: a low leak rate indicates a longer memory and a slower dynamics, whereas a high leak rate leads to a shorter memory but a higher speed of update dynamics (***Lukoševičius, 2012***). Only the output weights W_{out} are trained, using a supervised learning with the ridge regression method (the Tikhonov regularization) on pre-generated data:

$$W_{out} = Y^{target} X^T (XX^T + \beta I)^{-1} \quad \text{with} \quad X = [1, \mathbf{u}, \mathbf{x}] \tag{2}$$

where Y^{target}, β and I are respectively the target signal to approximate, the regularization coefficient and the identity matrix. Although supervised learning is far from being biologically plausible, the reservoir computing approach has been shown to exhibit some biologically plausible characteristics. (***Lukoševičius and Jaeger, 2009***). The model was built thanks to the python library ReservoirPy (***Trouvain et al., 2020***).

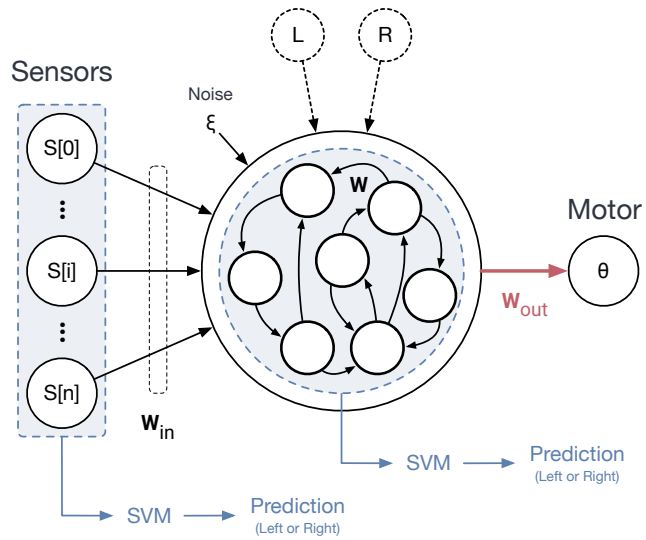

Figure 2. Model Architecture. Composed of 8 sensor inputs, a reservoir and a motor output (orientation). The black arrows are fixed connections while the red arrows are plastic connections and are trained. The reservoir states are used as the input to a SVM classifier which is trained to make a prediction about the decision (going left or right) of the agent. A left (L) and right (R) cue can be fed to the model depending on the experiment.

Tutor model

In order to generate data for learning, we implemented a simple Braintenberg vehicle (***Braitenberg, 1986***) where the agent moves automatically with a constant speed and changes its orientation according to the values of its sensors. This method has also been explored by ***Antonelo and Schrauwen*** (***2014***), where they also used an ESN to solve a navigation task with a robot. The model takes as input the sensor values $S_0, ..., S_n$ of the agent and outputs the next orientation θ (see Figure 2). At each time step the sensors measure the distance to the walls and the bot turns in order to avoid the walls. At each timestep, the position of the bot is updated as follows:

$$\theta(t) = \theta(t-1) + 0.01 \sum_i \alpha_i s_i \tag{3}$$

$$p(t) = p(t-1) + 2 * (cos(\theta(t)) + sin(\theta(t))) \tag{4}$$

where $p(t)$ and $p(t+1)$ are the positions of the agent at time step t and $t+1$, $\theta(t)$ is the orientation of the agent, calculated as the weighted (α_i) sum of the values of the sensors s_i. The norm of the

movement is kept constant and fixed at 2. Although the majority of this study focuses on analyzing this model, one variant will be employed in a specific scenario (see section Splitting effect at the population-level). This variant involves incorporating two additional inputs labeled as "contextual inputs" ("L" and "R" in Figure 2). These inputs indicate the direction the agent should choose at the next decision point. They are binary values, set to 1 only when the agent is expected to take the corresponding direction. One cue is associated with a right turn (R), and the other with a left turn (L).

Training data

The ESN is trained using supervised learning, containing samples from the desired 8-shaped trajectory. Since the Braitenberg algorithm only aims at avoiding obstacles, the agent is forced to follow the desired trajectory thanks to the timely addition of walls at the intersection points as shown on Figure 3. After generating the right pathway, the added walls are removed and the true sensor values are gathered as input. Gaussian noise is added to the position values of the robot at every time step in order to make the training more robust. Approximately 20,000 time steps were generated (equivalent to 28 complete alternating 8-loops) and separated into training and testing sets.

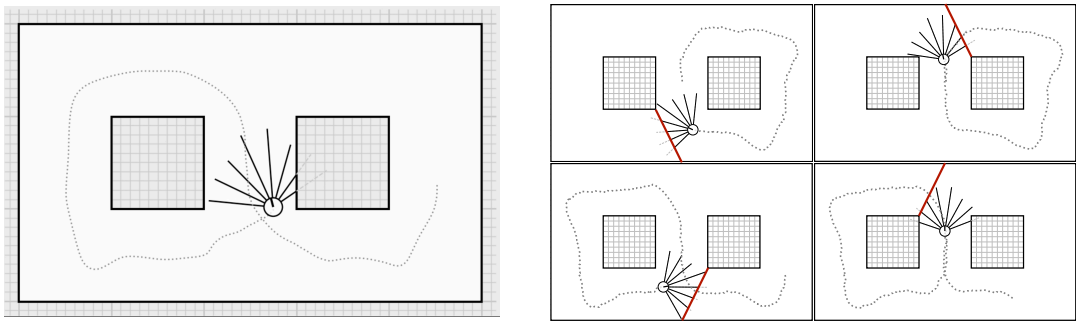

Figure 3. (Left) agent Simulation in 8-Maze. Snapshot from a simulation showcasing the agent's navigation within an 8-maze. The agent employs 8 sensors to measure wall distance in real-time. **(Right) Generation of the 8-shape trajectory** with the addition of walls at the intersection points

Hyper parameters tuning

Hyper-parameters were first tuned using the Hyperopt python library (***Bergstra et al., 2015***) in order to set the order of magnitude, and they were subsequently manually tuned (see Supplementary data for exact values). The different models possess different leak rates depending on the task since each task requires a different amount of working-memory: the alternating task requires to remember only the last decision in order to alternate at the next decision point, whereas the half-alternating task RR-LL requires to remember the last two decisions in order to know if alternation is needed at the next decision point. Consequently, it requires a larger memory, which explains the smaller leak rate. One other main difference is the maximum size of the sensors. Given the increased complexity of the RR-LL task, we also expanded the maximum sensory range to better account for the environment.

Analysis

After the training phase, models are able to solve respectively the alternating and the half-alternating task. When the agent navigates the maze, the reservoir's internal states are recorded and analyzed using two methods: single-cell level analysis and population-level analysis.

Single-cell level analysis

Analyzing the mean firing activity of individual neurons within the reservoir can reveal various types of hippocampal cells, including head-direction cells, place cells, and splitter cells. Head-direction cells encode the agent's orientation, evident in their mean firing patterns. Place cells exhibit differential firing rates based on the agent's location, while splitter cells fire distinctively in overlapping segments of trajectories, depending on the animal's past or future trajectory. The differentiation in mean firing activity on intersecting trajectories serves as a qualitative indicator of splitter cell behavior. To quantitatively assess individual splitter cells, the trajectory selectivity index (SI) defined by ***Tang et al.*** (***2023***) and ***Kay et al.*** (***2020***), is employed. This index, previously utilized in identifying splitter cells in rat hippocampus, involves comparing the mean firing activity of splitter cells on two overlapping trajectories. For example, when evaluating the activity of the cells in the central corridor as the bot goes either from right to left (RL) or right to right (RR), the SI is calculated as follow:

$$SI_{RL-RR} = \frac{\bar{x}_i^{RL}(corridor) - \bar{x}_i^{RR}(corridor)}{\bar{x}_i^{RL}(corridor) + (\bar{x}_i^{RR}(corridor))} \tag{5}$$

Where $i \in [1, 1000]$, $\bar{x}_i^{RL}(corridor)$ and $\bar{x}_i^{RR}(corridor)$ represent the mean reservoir activity in the central corridor when the bot is respectively moving left and right at the decision point, and i is the neuron index of the reservoir. If $| SI |> 0.1$, the neuron is identified as a splitter cell. This metric facilitates the determination of the proportion of each type of splitter cell by separately considering trajectories with the same future but different past (retrospective) using SI_{RR-LR} and SI_{LL-RL}, as well as trajectories with the same past but different future (prospective) using SI_{RR-RL} and SI_{LL-LR}. The splitter cells are also being identified thanks to raster plot generation. The activity of individual neurons are first recorded in the central corridor of the maze, which is shared between two trajectories with a different past, and different future. To illustrate the difference in spiking activity of such splitter cells, the neuron activity was turned into a dynamic firing rate according to the Poisson process (***Heeger et al., 2000***).

$$P\{1 \text{ spike in } \delta t\} \approx r\delta t \quad \text{and} \quad r(t) = \alpha[V(t) - V_{th}] \tag{6}$$

where α is the maximum firing rate 150 Hz, $V(t)$ is the normalized activity of the reservoir neuron and V_{th} is the mean normalized activity of the neuron.

Population-level analysis

While executing the alternated task, the reservoir state of the ESN is analysed at the population level such as to inspect to which extent it forms an internal and hidden representation of the trajectory. Principal Component Analysis (PCA) is employed as a dimensionality reduction technique to identify patterns and significant features within the processed data. PCA is a fast, non-random linear mapping which is sufficient to capture global structure. Similarly to PCA, UMAP (Uniform Manifold Approximation and Projection) (***McInnes et al., 2018***) is another technique used to visualize high-dimensional data in lower dimensions while preserving its underlying structure. While PCA is a non-random linear mapping and tends to capture global structure, UMAP is a random projection that tries to capture local similarities. This method has also been used to study the activity of populations in the hippocampus (***Tang et al., 2023***; ***Yang et al., 2023***). In this study, UMAP is utilized to explore the organization and clustering of neuronal activity patterns, particularly to investigate two types of splitter cells: retrospective and prospective cells. Additionally, UMAP is employed to examine the relationship between hippocampal activity and decision-making processes. Based on the hypothesis that the future direction of the robot can be predicted by observing the internal states of the reservoir, a classification task is performed. In this task, the reservoir states are categorized into two groups: one associated with predicting a left turn and the other with predicting a right turn. The Support Vector Machine (SVM) classifier (***Cortes and Vapnik, 1995***) is used and

takes the reservoir state at each position of the agent during the 8-shaped trajectory as input and predicts the robot's decision at the next intersection.

Results

Single cells analysis reveals the presence of hippocampal cells

Several studies have provided evidence for the presence of hippocampal neurons crucial for navigation tasks, encoding essential information for the agent's spatial location. For example, place cells represent the agent's current position (***O'Keefe and Dostrovsky, 1971***), head-direction cells encode its orientation (***Taube, 1998***), and splitter cells, a subset of place cells, encode both current position as well as the recent past or the near future trajectories. As highlighted by ***György Buzsáki*** (***2019***), this list can be further extended and include for example time cells, boundary cells, distance cells, etc. This section aims to identify some of these hippocampal neurons by examining the reservoir activity while the agent is performing the navigation tasks.

Place cells and head-direction cells

Through the analysis of reservoir neuron activity during the agent's navigation, the activity of certain units similar to hippocampal cells has been uncovered, as illustrated in Figures reffig:hippocells and 5. Figure 4 displays the activity of various types of cells: a cell akin to a place cell, representing a distinct location within the maze, another resembling a corner cell encoding the maze's corners recently discovered in the mouse subiculum (***Sun et al., 2024***), and a cell termed here "loop cell", indicating which loop (right or left) the agent is navigating. The loop cells can be regarded as an extended form of place cells, as they also become active in a specific region, albeit with less localization, covering the entire loop in their instance. In Figure 5, the evolution of the agent's orientation over time is depicted in yellow. The figure showcases the activity of the five most correlated neurons with the agent's orientation, indicating that these neurons correspond to head-direction cells. These latter observations highlight the presence of a diverse hippocampal cells within the reservoir system.

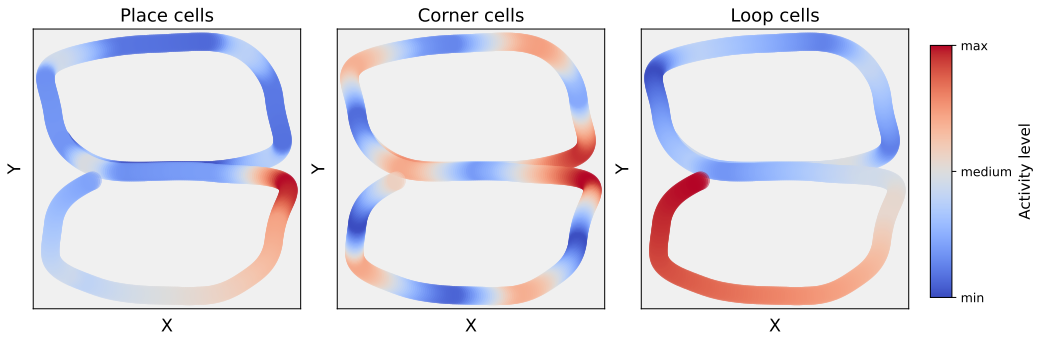

Figure 4. Discovering activity akin to place cells and corner cells, as well as cells that encode the loop location, denoted as "loop cells" as the agent navigates in the 8-shaped maze. Each diagram represents the activity of a single neuron while the agent is navigating into the 8 maze. Neuronal activity intensity is depicted by trace color. The right graph showcases loop cells, whose activity is influenced by loop locations. The middle graph illustrates corner cells, which exhibit distinct firing patterns in corners. The left graph displays place cells, which fire selectively at specific locations.

Splitter cells

In their recent review, ***Duvelle et al.*** (***2023***) conducted a comprehensive analysis of splitter cells, examining both experimental and theoretical aspects. The objective here is to replicate this analysis and verify whether their hypotheses are supported in the reservoir model. While recording the activity of individual neurons, splitter cells were identified through the computation of the selectivity

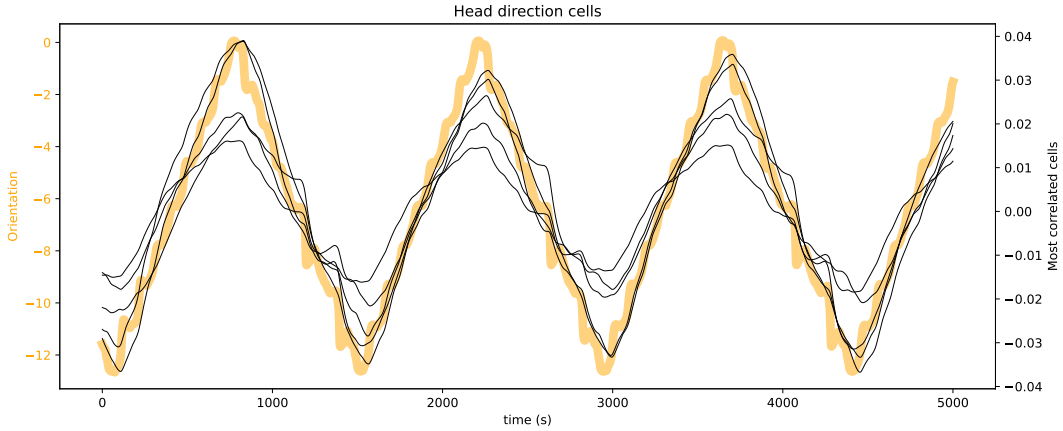

Figure 5. The emergence of Head-direction cells as the agent navigates through the 8 maze.The agent's orientation over time is represented in yellow, while the traces of the five neurons most closely correlated with this orientation are shown in black. These neurons accurately encode the agent's orientation.

index (SI), selecting those with an absolute value greater than 0.1 ($|S1| > 0.1$). The upper section of Figure 6 illustrates some of these identified splitter cells. The bars represent the activity of neurons while the agent traverses the central corridor, indicating that all recorded activities correspond to the exact same location. Despite this identical spatial location, the recent past and near future of the agent are different. The red bars indicate neuron activity in the central stem when the agent comes from the left and will go to the right, while the blue bars indicate activity in the central stem when the agent comes from right and will go to the left. The selected neurons of the model exhibit significant distinctive firing rates in the central maze stem dependant on the journey (R-L in blue or L-R in red), illustrating the splitting. These selected splitter cells display a higher activity within the corridor (depicted in red and blue) compared to areas outside the corridor (depicted in green). While not all splitter cells exhibit this particular activity pattern, these figures serve as a noteworthy example of the firing selectivity observed in our model. The activity depicted in the top part of Figure 6 does not provide sufficient discrimination to determine whether the splitter cells are encoding for the past (retrospective cells) or for the future (prospective cells). This is because the recorded activity corresponds to trajectories that have different past and future directions at the same time (R-L and L-R). To find out the nature of the splitter cells, i.e. whether they encode the upcoming future or the recent past, the neuron's activity is recorded during trajectories that either share the same past but have different futures (L-L and L-R), or vice versa (L-L and R-L), as depicted in the bottom part of Figure 6. On the left part of the figure, the activity of six retrospective neurons is depicted, revealing a significant difference in their mean activity depending on the side of the maze the agent approaches from. Conversely, the right part of the figure displays an example of six prospective cells, demonstrating a significant difference in their mean activity depending on the future direction the bot is heading in the future. Raster plots generated from the Poisson process are depicted in Figure 8, illustrating the average firing activity of a subset of splitter neurons in the reservoir. Various patterns are evident: certain neurons, such as those with indices 752 and 350, exhibit differences in firing rates, others neurons fire exclusively for specific trajectories, resembling the behavior of the neuron indexed as 920. Additionally, some others such as those indexed as 723, 742, and 850, fire at different timings. Each individual neuron displays a distinct average activity pattern based on the trajectory, even when recorded at the same spatial location. This clear differentiation in activity patterns reinforces the idea of the existence of splitter cells within the reservoir, showcasing their selective responses based on the past and/or future trajectory of the agent. The splitter cells within the reservoir have the capacity to encode various types of information. This is clearly depicted in Figure 7 that shows the Representational Similarity Analysis (RSA).

A correlation matrix of three splitter cells is presented. Each row and column of each matrix corresponds to a trajectory taken by the agent, with the values representing the correlation between the mean activities of the two trajectories in the central stem. This framework enables to infer similarities between brain responses and the behavior of the agent. The correlations depicted in the matrix highlight the versatility of splitter cells in encoding different types of information. For example, neuron 38 exhibits a significant correlation when the bot chooses to remain within the same loop, even though the trajectories are different, while neuron 312 shows a high correlation when the bot decides to switch loops. Additionally, neuron 498 demonstrates a notable correlation among trajectories that share similar recent histories.

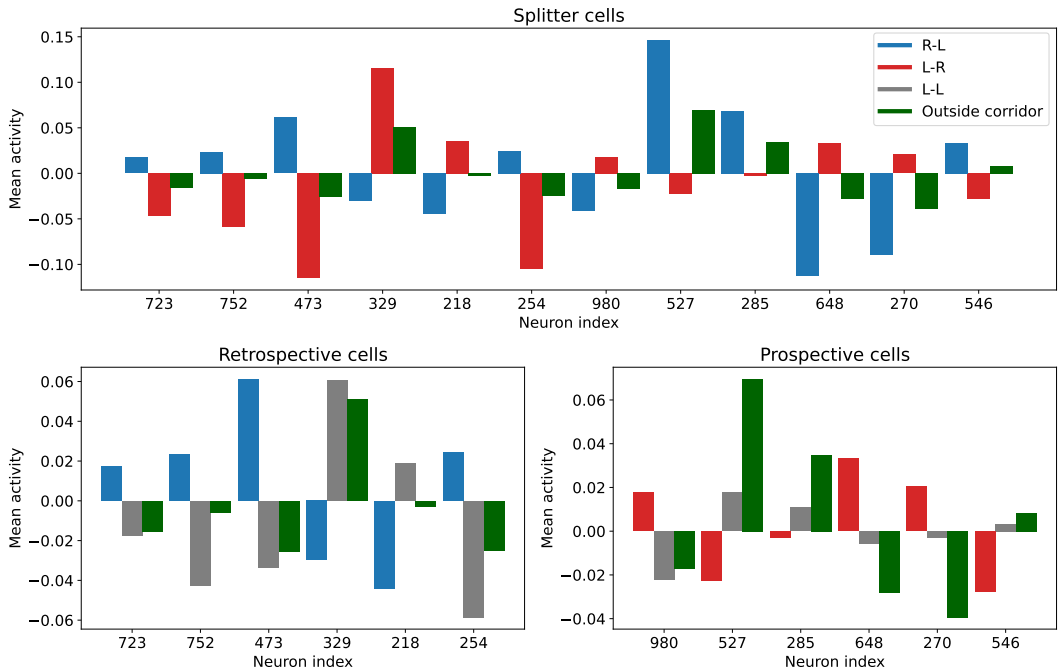

Figure 6. (Top) Comparing firing rates in the central stem. Distinct firing rates in the central maze stem confirm they are splitter cells. The cells exhibit high activity within the corridor compared to outside of the corridor, with different rates depending on the trajectory. (Bottom left) Retrospective cells. Comparing firing rates for shared futures (left) but different pasts (right and left). (Bottom right) Retrospective cells. Comparing firing rates for shared past (left) but different future (right and left).

Splitting effect at the population-level

While the initial section focused on analyzing the firing activity of individual neurons, this section aims to investigate neuron dynamics at the population level. Similar to the single-cell analysis, the reservoir neurons are recorded while the agent performs the navigation task.

Trajectory information encoded in the internal state

A 2D PCA analysis is conducted on the reservoir states captured during the agent's 8-shaped trajectory during 5000 time steps, as depicted in Figure 9. The figure shows two symmetric sub-attractors, which are linearly separable, that correspond to the two loops of the 8-shape trajectory. This result illustrates the ability of the ESN to implicitly approximate the state-space of the 8-maze environment within a continuous temporal dynamics. Essentially, the model achieves this by transforming the input data into a higher-dimensional space. The recurrent connections within the reservoir enable the network to incorporate and reverberate past information, allowing for the encoding of

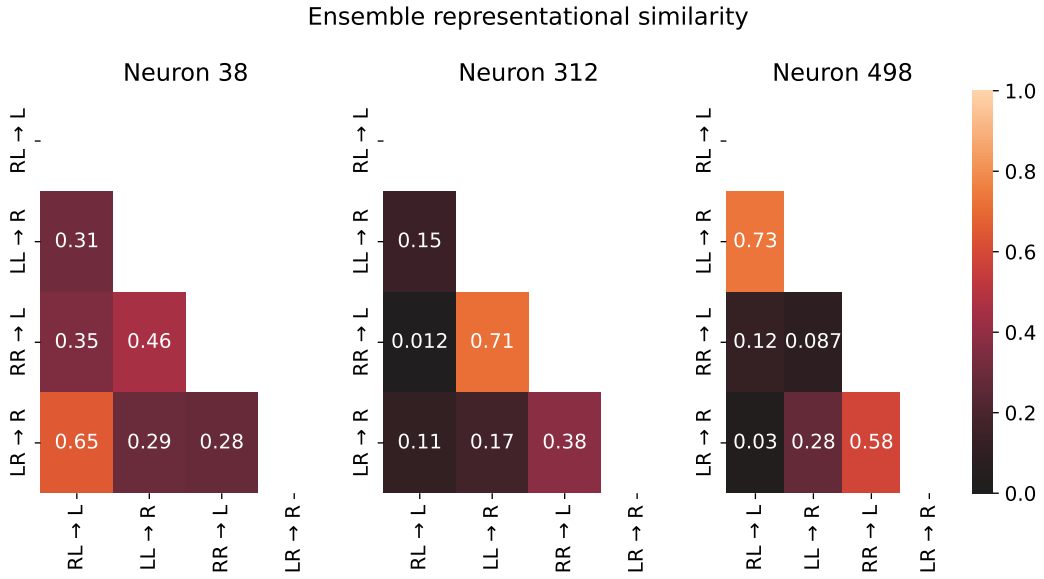

Figure 7. Representational Similarity Analysis (RSA): A correlation matrix depicting the average activity of three designated splitter cells as the bot navigates through the central corridor, considering various past and future scenarios. In the context provided, the notation 'RL- L' signifies that the bot previously completed a right loop (R), followed by a left loop (L), and now intends to turn left (L) after navigating the central corridor. For each trajectory, the activities of the neurons were recorded across 10 instances, and the resultant neuronal activity corresponds to the mean of these 10 trajectories. These correlations reveal that splitter cells possess the ability to encode diverse forms of information: for instance, neuron 38 exhibits a strong correlation when the bot opts to remain within the same loop, while neuron 312 demonstrates an high activity correlation when the bot decides to switch loops, and neuron 498 displays a notable correlation among trajectories that share similar recent histories.

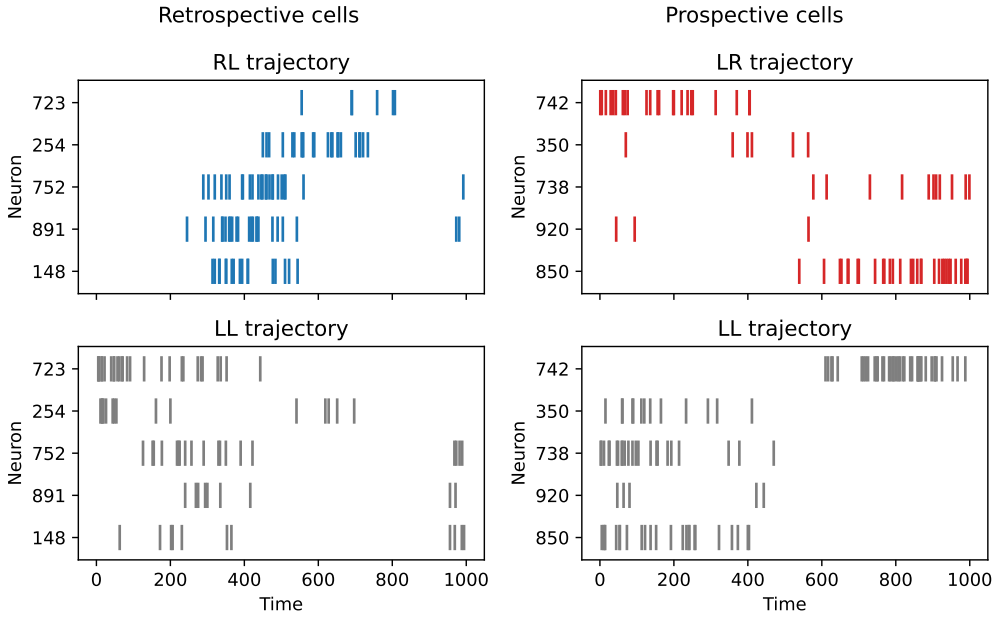

Figure 8. Raster plots. Different types of pattern are depicted: the same neurons can have a difference in firing rate, or fire for only one of the trajectories. The different colors correspond to the trajectories: R-L (blue) vs L-L (grey), L-R (red).

temporal context and the integration of the past trajectory. As mentioned by ***Duvelle et al.*** (***2023***), incorporating temporal context by augmenting sensory input is a crucial aspect for modeling splitter cells. A classification task is conducted on the reservoir states to check whether they can predict the next decision of the agent. The accuracy score obtained is of 1. The prediction result depicted in Figure 9 demonstrates that if the agent recently turned right, the reservoir alters its internal state to favor a left turn within just a few dozen time steps. This shows that by simply observing the internal states of the reservoir, it is possible to predict its next prediction. Consequently, it can be inferred that the internal neurons of the robot encapsulate information about future predictions. In contrast, applying the same classification to sensor values failed in the predictions, as anticipated.

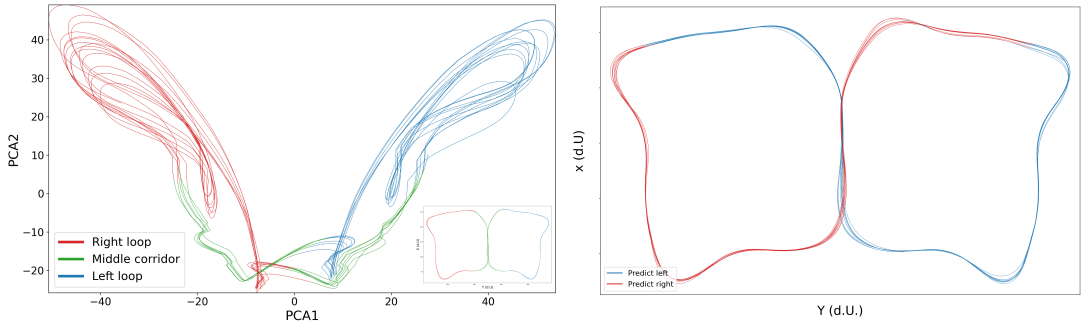

Figure 9. (Left) 2D PCA on the reservoir states, applied during 7 loops of the robot. **(Right) SVM predictions** Agent's Trajectory and SVM Classification. Colors represent SVM predictions: red for right, blue for left directions. By simply observing the internal states of the reservoir, it is possible to predict the next prediction: the neurons of the agent encapsulate information about future predictions

Temporal Context Model versus Latent State Inference Model

At the population level, splitter activity is defined as the presence of different pattern of activity at the same spatial location (***Duvelle et al., 2023***). This phenomenon is depicted in the 3D PCAs of Figure 10. The top-left figure corresponds to the PCA of the reservoir during a single 8-loop. The green trace followed by the blue trace indicates the agent moving from the right and heading left, while the green trace followed by the red trace indicates the agent moving from the left and heading right. Remarkably, these two trajectories occupy distinct regions in the PCA space, despite the robot passing through the same location in the main corridor (green trace). The three markers (C1, C2, C3) indicate specific positions within the main corridor. If there were no splitter effect, these three markers would overlap since they represent the same spatial location. However, they are distinctly separated with a measurable PCA distance. This observation further highlights the presence of splitter cells and their role in generating distinct patterns of activity within the reservoir and align with the splitter cells concepts proposed by ***Duvelle et al.*** (***2023***). The PCA analysis has been conducted on two different model variants. The first one (top part of Figure 10) corresponds to the PCA of the classical model being used so far: only the sensor value are fed as input. In the second one (bottom part of Figure 10), two additional inputs are fed in the reservoir, called "contextual input" ("L" and "R" in the figure), indicating which direction the agent should take at the next decision point. They are binary values activated only in the central corridor of the maze, switched to a value of 1 when the bot is known to take the corresponding direction. The main goal of employing these two versions of the model is to investigate the impact of cues, which provide explicit information about the subsequent state. The impact of the leak rate on the splitting effect is also investigated. The leak rate is a crucial parameter of ESN models, since it plays a major role in controlling the memory and timescale of the network's dynamics. It governs the rate at which the internal states of the reservoir decay over time, thus determining the extent to which past in-

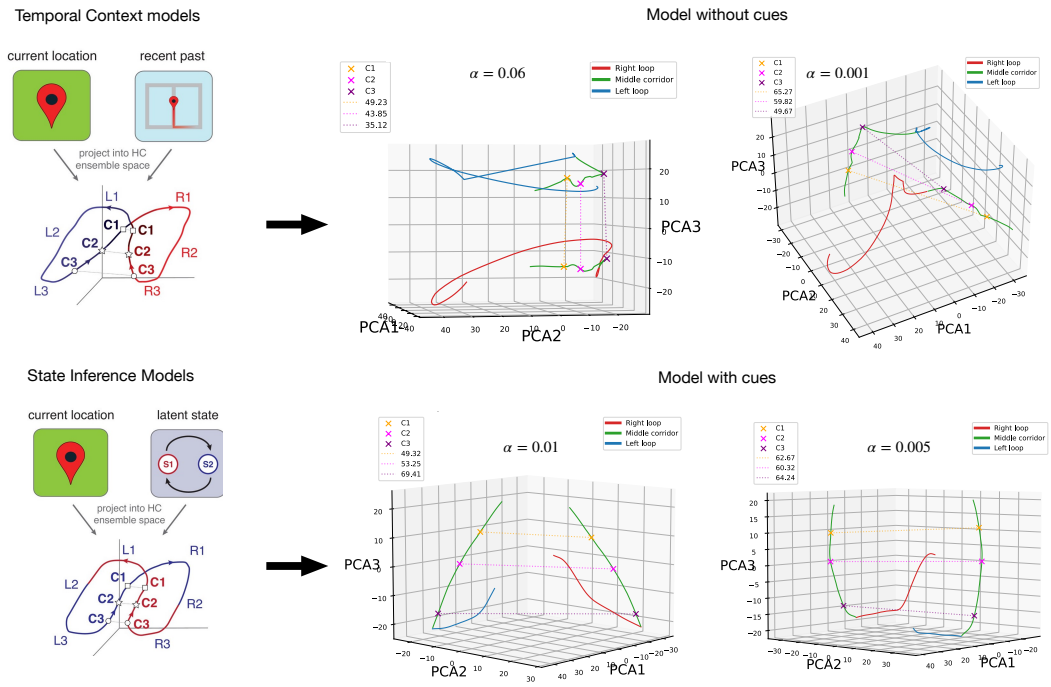

Figure 10. (Top) Temporal Context Model. Left: Taken from Duvelle et al.***Duvelle et al.*** (***2023***). Hypothesized neural Activity Trajectories in PCA Space illustrating the temporal context model (TCM). The population activity at a specific location exhibits closer proximity in PCA space when the recent past is similar and greater distance when it is less similar. This is demonstrated in the distances between markers C1, C2, and C3. Right: 3D PCA of the cue-less model for varying leak rate (α). With a smaller leak rate, the reservoir model contains a larger memory capacity, the distances between C1, C2, C3 markers increase, indicating amplified impact of past trajectory on neuron activity. Consequently as the leak rate decreases, the impact of the agent's past trajectory is more pronounced, which is in line with Duvelle's TCM hypothesis. **(Bottom) Latent State Inference Model** Left: Taken from ***Duvelle et al.*** (***2023***) Hypothesized neural activity trajectories for the latent-state inference (LSI) hypothesis. The distances between the markers C1, C2, and C3 don't vary. Right: 3D PCA analysis with the model with cues for varying leak rate (α). For $\alpha = 0.01$, the further the agent gets into the central corridor, the more distinct its neural activity becomes (C3 distance is larger than C1 distance, contrary to the TCM case), indicating a stronger association with one of the two discrete states. As the leak rate decreases, the distances between C1, C2 and C3 markers converge, which aligns with (LSI).

formation influences the current state of the network. Various values are tested, represented by α, with a higher leak rate on the left part of Figure 10) and a smaller one on the right part of Figure 10). As the green trace does not overlap in all of the graphs, despite corresponding to the same location, the splitting effect is evident in both models and all leak rate values. All graphs show different shape of activity pattern in the PCA space. In the model without cues (top part of Figure 10), the PCAs are in line with the "Temporal context model" (TCM) mentioned by Duvelle et al. in ***Duvelle et al.*** (***2023***). The population activity at a given location has a closer proximity in the PCA space when the recent past is similar, and a greater distance when the recent past differs. This effect is demonstrated by the varying distances between the markers C1, C2, and C3, which correspond to specific positions within the central corridor as depicted in Figure 10. Specifically, the distance between the C1 markers is larger than the distance between the C2 markers, and the distance between the C2 markers is larger than the distance between the C3 markers. Notably, when the bot reaches C3 in the maze, it shares the entire main corridor as a common past trajectory, whereas at C1, the agent's past experience differs. The increase in distance in the PCA space is consistent with the decrease in the leak rate. A smaller leak rate enhances sensitivity to the reservoir's past activity, resulting in a larger PCA distance, as confirmed in the figure, where the distances between the markers increase when the leak rate is smaller. Therefore, these findings provide support for considering the model without cues as a Temporal Context Model. In this model, not only the current inputs are considered, but also recently experienced events, such as the past trajectory of the bot stored through the reverberation of recurrently connected neurons, enabling a context-dependent response to current events. The model with the contextual cues (bottom part of Figure 10) deviates from a TCM. Instead, it aligns more closely with the concept of Latent State Inference introduced by Duvelle et al., where experiences are organized into discrete states. In the left Figure (bottom part of Figure 10), the further the agent gets into the central corridor, the more distinct its neural activity becomes (C3 distance is bigger than C1 distance, which is the contrary effect of the TCM case), indicating a stronger association with one of the two discrete states. As the leak rate α decreases (right figure), the distances between C1, C2 and C3 markers converge, until not varying for $\alpha = 0.005$, which aligns with the LSI model. The presence of contextual cues informs at every time step about the next direction to take, consequently, the model does not require as much memory of the past trajectory to complete the navigation task. Thus, the apparition of the cues in the reservoir model is enough to go from a TCM model to an LSI model, with the first needing additional contextual information thanks to the past trace, and the latter reflecting the inference of hidden states (***Duvelle et al., 2023***).

These results confirm Duvelle et al.'s theoretical framework introducing TCM and LSI in their comprehensive review. These two theories represent leading perspectives in the study of hippocampal functions. Our findings are in alignment and consistent with all the hypothesized predictions based on these two models, thereby supporting these theoretical concepts. As Duvelle et al. mention, these two theories aren't mutually exclusive, leaving possibility for hybrid models that contain elements of both. While reservoir computing primarily relies on maintaining working memory within the reverberation of reservoir neurons, it's natural to view this model as having a temporal gradedness comparable to TCM. This forms the basis of reservoir computing. By incorporating contextual cues as input, one can seamlessly transition from a TCM to an LSI model, where the weighting of contextual cues determines the balance between TCM and LSI components. Adjusting the scaling of these cues allows for a flexible hybrid model, combining temporally decaying representations of the past with discrete state representations. For instance, in tasks like the continuation alternation task, where task-state switches occur with changes in cue values, the internal states shift between distinct regions of activity space while still permitting continuous neural activity, as noted by Duvelle et al.

Population analysis reveals the apparition of retrospective and prospective cells

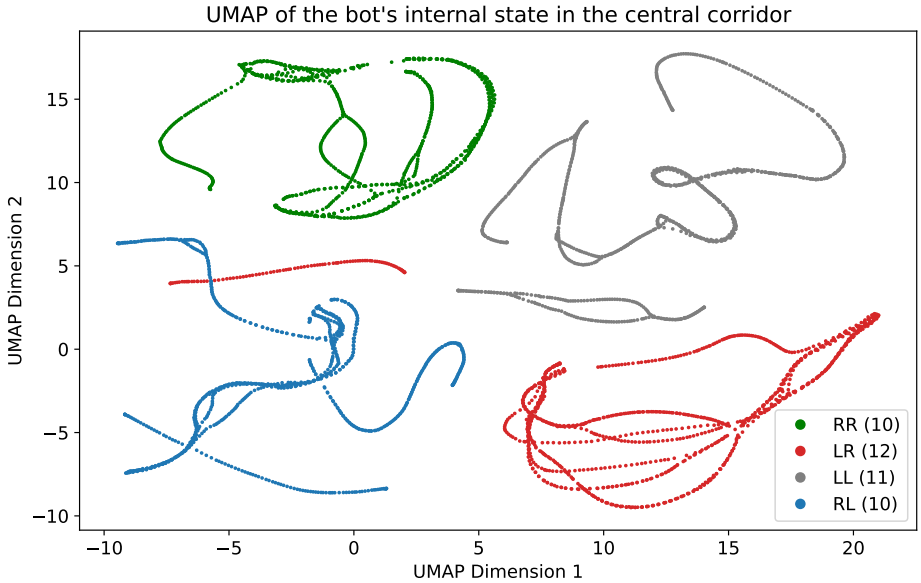

Figure 11. UMAP analysis. The numbers in the labels correspond to the quantity of recorded trajectories in the central corridor for each scenario. Divergent pasts and identical futures show distinct traces, highlighting the influence of past experiences (retrospective cells).

In the same way as the single-cell analysis, the reservoir activity is recorded during trajectories with the same past but different future (R-L and L-L), or with the same future but different past. The corresponding UMAP analysis is shown in Figure 11 where each trace corresponds to the internal state of the agent when going through the central corridor. A distinct separation is observed between the traces associated with a different future (LL and LR, or RL and RR), as well as between traces associated with a different past (LL and RL, or LR and RR). This suggest that the agent's past and future can significantly influence its activity, confirming the presence of prospective and retrospective cells whose influence can be observed at the population-level. This result is also found in the model with the contextual cues (See Supplementary data). Just like splitter

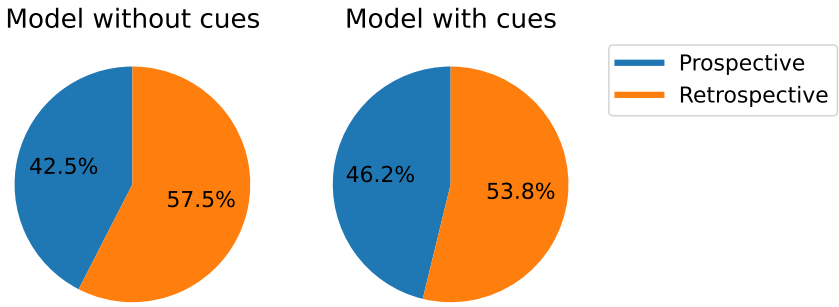

Figure 12. Prospective and retrospective splitter cells proportions

cells, Prospective and Retrospective cells are distinguished through the computation of the selectivity index (SI) for each neuron. In the case of retrospective cells, the SI assesses two overlapping trajectories with distinct pasts, while for prospective cells, it evaluates two overlapping trajectories with different futures. In both models, a higher proportion of retrospective cells is observed

compared to prospective cells, with a distribution of 32%-68% for prospective-retrospective in the cue-less model and 46%-54% for prospective-retrospective in the model with cues, as illustrated in Figure 12. This result is in line with all biological surveys as it is highlighted in ***Duvelle et al.*** (***2023***): for instance, 22%-72% prospective-retrospective in the CA1 region of the rat hippocampus were found in ***Catanese et al.*** (***2014***) and 40%-60% in ***Bahar and Shapiro*** (***2012***). The higher amount of retrospective cells may be attributed to the fact that prospective cells require training and do not emerge naturally, unlike retrospective cells, which are encoded directly from past experiences. Prospective cells necessitate the model to first learn a representation of the task to adequately encode the near future as highlight ***Duvelle et al.*** (***2023***). This idea is supported by the observation that the model with cues, which is closer to representing latent states thanks to the cues, exhibits a more balanced proportion of retrospective-prospective cells. In this model, the importance of knowing the past for making decisions about the future is diminished since contextual cues provide information on time.

Relationship with decision-making

A central question that emerges in several studies on splitter cells focuses on understanding the relationship between these cells and the behavioral choices of animals (***Ferbinteanu and Shapiro, 2003***; ***Allen et al., 2012***; ***Bahar and Shapiro, 2012***). When the animal makes an incorrect decision, the investigation extends to whether the splitter cells encode this error or only capture what has been learned (i.e., the correct choice), regardless of the animal's final decision. Furthermore, researchers are exploring whether this splitting activity has an impact on the overall activity of splitter cells. We proceed to a similar analysis within our models by introducing an error during the trial execution. The objective is to determine whether this predictive capability remains consistent when the bot makes an error. For example, if the bot has been trained to go right but ultimately goes left, the focus is on understanding what the splitter cells encode - the intention to go right or the actual action of going left. To address this question, the model without cues is initially built with a minor adjustment introduced to their spectral radius, rendering it suboptimal in executing the task and prone to errors. After learning, the models successfully follow the 8-shaped path until a point where they make a mistake, such as going left instead of right. During this error, the activity of the selected splitter cells are recorded and their firing rates are being compared with the correct case. In Figure 13, we depict a dissociation between splitting activity and decision-making. Even though the bot deviates from the correct direction (e.g., choosing right instead of left), the splitter cells continue to encode the accurate decision rather than the observed behavior. This is clearly shown in the left part of Figure 13, where the activity of the splitter cells is similar to the firing pattern associated with the bot going left (the error case in yellow is similar as the activity of the R-L trajectory in blue), despite the bot actually moving to the right. If the splitter cells were to encode actual behavior, the firing activity during an error would be similar to the R-R trajectory in green. The dissociation at the population level is further illustrated in the right part of Figure 13, where the error case is positioned in the R-L region of the UMAP, while the bot's behavior corresponds to R-R. Comparable results have been shown in the hippocampus of rats, in studies such as ***Catanese et al.*** (***2012***). In the study presented in ***Catanese et al.*** (***2012***), rats were required to alter their trajectory initially set as an alternating task upon the appearance of a cue. Interestingly, even when the cue instructed the rat to take the opposite direction from what it had learned, the splitter cells continued to encode the initial direction. These studies are in line with the behavior of our model providing conclusive evidence of a dissociation between splitter activity and behavioral choices. However, not all studies align with these findings. The opposite effect, where splitter activity is influenced by error trials, has been demonstrated in studies such as ***Ferbinteanu and Shapiro*** (***2003***); ***Bahar and Shapiro*** (***2012***); ***Robitsek et al.*** (***2013***).

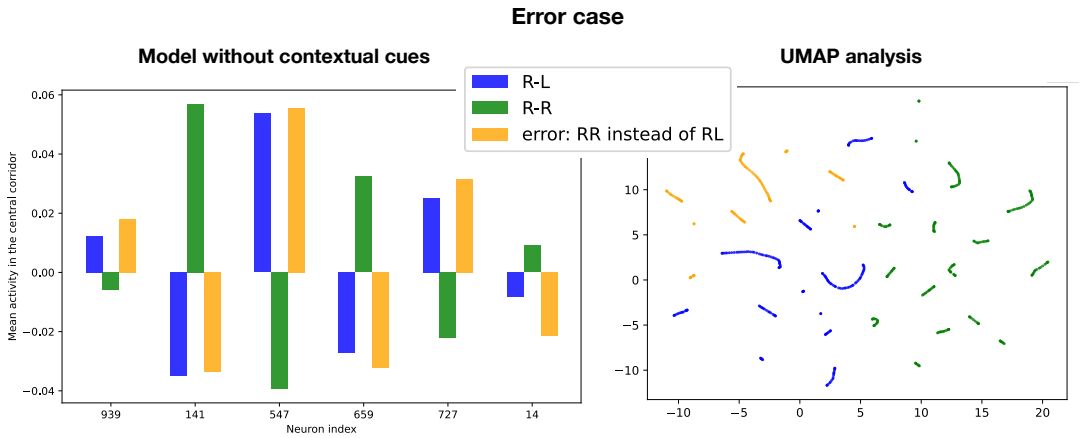

Figure 13. Splitting activity during correct and error trials. The activity during the error case corresponding to the 'R-R' trajectory (yellow) is similar as the R-L trajectory (blue) as if the bot were not doing any mistakes. In other words, the splitter cells are not encoding the real choice of the bot, but only what it has initially learnt. The error is not significantly encoded in the splitter activity.

Discussion

The model we proposed to study is as simple as it can be: a reservoir of neurons is fed by 8 sensors and outputs a simple change in direction that is fed back to the reservoir. We did not claim any architectural plausibility with the hippocampal formation even though the reservoir could be compared to some extents to CA3 due to its highly recurrent nature. Despite this apparent simplicity, the model is able to solve the alternating navigation task and we've been able to demonstrate for the existence of splitter cells as described by ***Duvelle et al.*** (***2023***) in their review. By redoing all their analyses on the reservoir, we've found activities that characterize splitter cells (at least in rodents). Furthemore, we've also validated the proposed theoretical framework by developing a model that incorporates concepts from both TCM (Temporal Context Model, ***Howard and Kahana*** (***2002***)) and LSI (Latent State Inference, ***Gershman and Niv*** (***2010***)). Even though these concepts appear to be two distinct notions, they are both needed for the interpretation of experimental data on splitter cells and the model we proposed suggests a simple mechanism where the transition from TCM to LSI is driven by contextual cues.

One noticeable aspect of our work is the ability to examine prospective and retrospective cells separately, based on the half-alternation task, which is not commonly addressed in existing computational models. For example, in ***Cone and Clopath*** (***2024***), the splitting effect emerges transiently during task learning while in our model, the reservoir (that contains the splitter cells) is random and remains unmodified. Only the output set of weights, that drive behavior, is learned. Consequently, the splitter-cell-like activity observed inside the reservoir results solely from internal structure and the input whose activity derives from actual behavior which is in turn driven by the output set of weights that control behavior. In our case, learning process is a simple identification of a random pattern of activity inside the reservoir to the motor command to be issued. This process is indeed well aligned with ***György Buzsáki*** (***2019***)'s inside-out view of the brain, where he suggests that brain activity is not generated from scratch during exploration and learning. It is rather a vast reservoir capable of generating an extensive repertoire of neuronal patterns, initially independent of experience. In other words, these patterns have initially no meaning and acquire it through exploration and learning. In this sense, experience is primarily a process of matching preexisting neuronal dynamics to outside events, and learning does not create new brain activity from scratch, but rather selects preexisting pattern of activity responding to the external stimuli. Although our model with random weights may not capture the full complexity of the brain, any modification of weights may

have only transient effects on network dynamics. This idea is further supported by the work of ***Golub et al.*** (***2018***) where it was demonstrated that during learning, the population activity of the motor cortex of rhesus macaques followed a reassociation strategy, in which the same activity patterns observed before learning were associated with different movements. This suggests again that learning does not necessarily create new activity patterns from scratch, but rather modifies the associations between existing patterns and behavior. This inside-out view of the brain is supported by evidence that spatial memory cells and their features were already present when young pups made their first movement (***Langston et al., 2010***), or emerge early in postnatal development before significant spatial exploration (***Muessig et al., 2016***; ***Wills et al., 2014***). These findings suggest that the hippocampus possesses many possible neuronal trajectories even before exploration has begun.

Beyond splitter cells, we've also suggested the presence of place cells, head direction cells and decision cells. Even though we did not look for each and every other type of cells, we're confident that we would probably find them provided they're correlated with the T-maze task we've been using. This strongly suggests that the nature of those cells is behavioral rather than structural: their activities do not derive from structure nor learning (we have no learning inside the reservoir) but are rather indicative of a correlation with some pre-existing dynamics that are linked to the actual behavior. Considering for example decision cells (indicating decision to go left or right), it comes as no surprise that we can identify them inside the reservoir since the agent is really going left or right at the end of the central corridor, at least for an external observer. For the agent however, the sensory world is fully described by a set of eight sensors that do not convey the notions of *left* or *right*. Ultimately, this means that these decision cells may only exist in the eye of the observer. In other words, if a reservoir can solve a task with a minimal set of inputs, any variable that can be decoded from the activity of the reservoir derived from the resolution of the task itself as opposed to the structures involved in the resolution or even learning. The notion that these cells are behaviorally defined rather than structurally localized suggests a disconnection between these neural entities and the hippocampal formation. This idea is supported by numerous studies demonstrating the existence of place cells, head direction cells, boundary vector/border cells, and grid cells in the somatosensory cortex of rats (***Long and Zhang, 2021***) while grid cells have also been identified in the neocortex of rats (***Hawkins et al., 2019***; ***Constantinescu et al., 2016***) and humans (***Epstein et al., 2017***; ***Doeller et al., 2010***).From a broader perspective, we think this work questions the methodology in both Neurosciences and Computational Neuroscience: the huge repertoire of cells that has been established for the hippocampal formation over several decades might be only indicative of a bias in the method: we only find what we're looking for.

Supplementary data

Code availability

The code associated with this work is publicly available on GitHub at https://github.com/naomichx/splitter-cells

Models parameters

Alternating task R-L		
Parameter	No cues	With cues
Input size	8	10
Number of units	1000	1000
Input connectivity	0.2	0.09
Reservoir connectivity	0.19	0.19
Input scaling	1	1 (sensors), 10 (cues)
Spectral Radius	1.4	1.5
Leak Rate	0.02	0.1
Regularization	4e-6	6e-3
Maximum distance sensors	60	60

Half-alternating task RR-LL		
Parameter	No cues	With cues
Input size	8	10
Number of units	1000	1000
Reservoir connectivity	0.2	0.2
Input scaling	1	1 (sensors), 10 (cues)
Spectral Radius	0.9	1.4
Leak Rate	0.005	0.005
Regularization	1.2e-1	3e-2
Maximum distance sensors	80	80

Table 1. Parameter configuration for both cued and cue-less ESN models. Top: optimized parameters for the classical continuous alternating task (the bot is alternating between right and left direction). Bottom: optimized parameters for the half alternating task (the bot is alternating every two times).

References

Allen K, Rawlins JNP, Bannerman DM, Csicsvari J. Hippocampal place cells can encode multiple trial-dependent features through rate remapping. Journal of Neuroscience. 2012; 32(42):14752–14766.

Andersen P. The hippocampus book. Oxford university press; 2007.

Antonelo EA, Schrauwen B. On learning navigation behaviors for small mobile robots with reservoir computing architectures. IEEE transactions on neural networks and learning systems. 2014; 26(4):763–780.

Bahar AS, Shapiro ML. Remembering to learn: independent place and journey coding mechanisms contribute to memory transfer. Journal of Neuroscience. 2012; 32(6):2191–2203.

Bergstra J, Komer B, Eliasmith C, Yamins D, Cox DD. Hyperopt: a python library for model selection and hyperparameter optimization. Computational Science & Discovery. 2015; 8(1):014008.

Braitenberg V. Vehicles: Experiments in synthetic psychology. MIT press; 1986.

Buzsáki G, Moser EI. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nature neuroscience. 2013; 16(2):130–138.

Catanese J, Cerasti E, Zugaro M, Viggiano A, Wiener SI. Dynamics of decision-related activity in hippocampus. Hippocampus. 2012; 22(9):1901–1911.

Catanese J, Viggiano A, Cerasti E, Zugaro MB, Wiener SI. Retrospectively and prospectively modulated hippocampal place responses are differentially distributed along a common path in a continuous T-maze. Journal of Neuroscience. 2014; 34(39):13163–13169.

Cone I, Clopath C. Latent representations in hippocampal network model co-evolve with behavioral exploration of task structure. Nature Communications. 2024; 15(1):687.

Constantinescu AO, O'Reilly JX, Behrens TE. Organizing conceptual knowledge in humans with a gridlike code. Science. 2016; 352(6292):1464–1468.

Cortes C, Vapnik V. Support-vector networks. Machine learning. 1995; 20:273–297.

Doeller CF, Barry C, Burgess N. Evidence for grid cells in a human memory network. Nature. 2010; 463(7281):657–661.

Duvelle É, Grieves RM, Van der Meer MA. Temporal context and latent state inference in the hippocampal splitter signal. ELife. 2023; 12:e82357.

Epstein RA, Patai EZ, Julian JB, Spiers HJ. The cognitive map in humans: spatial navigation and beyond. Nature neuroscience. 2017; 20(11):1504–1513.

Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. Neuron. 2003; 40(6):1227–1239.

Frank LM, Brown EN, Wilson M. Trajectory encoding in the hippocampus and entorhinal cortex. Neuron. 2000; 27(1):169–178.

Gershman SJ, Niv Y. Learning latent structure: carving nature at its joints. Current Opinion in Neurobiology. 2010 Apr; 20(2):251–256. http://dx.doi.org/10.1016/j.conb.2010.02.008, doi: 10.1016/j.conb.2010.02.008.

Golub MD, Sadtler PT, Oby ER, Quick KM, Ryu SI, Tyler-Kabara EC, Batista AP, Chase SM, Yu BM. Learning by neural reassociation. Nature neuroscience. 2018; 21(4):607–616.

György Buzsáki M. The brain from inside out. Oxford University Press; 2019.

Hartley T, Lever C, Burgess N, O'Keefe J. Space in the brain: how the hippocampal formation supports spatial cognition. Philosophical Transactions of the Royal Society B: Biological Sciences. 2014; 369(1635):20120510.

Hawkins J, Lewis M, Klukas M, Purdy S, Ahmad S. A framework for intelligence and cortical function based on grid cells in the neocortex. Frontiers in neural circuits. 2019; 12:431889.

Heeger D, et al. Poisson model of spike generation. Handout, University of Standford. 2000; 5(1-13):76.

Howard MW, Kahana MJ. A Distributed Representation of Temporal Context. Journal of Mathematical Psychology. 2002 Jun; 46(3):269–299. http://dx.doi.org/10.1006/jmps.2001.1388, doi: 10.1006/jmps.2001.1388.

Jaeger H. The "echo state" approach to analysing and training recurrent neural networks-with an erratum note. Bonn, Germany: German National Research Center for Information Technology GMD Technical Report. 2001; 148(34):13.

Jaeger H. Echo state network. scholarpedia. 2007; 2(9):2330.

Kay K, Chung JE, Sosa M, Schor JS, Karlsson MP, Larkin MC, Liu DF, Frank LM. Constant sub-second cycling between representations of possible futures in the hippocampus. Cell. 2020; 180(3):552–567.

Langston RF, Ainge JA, Couey JJ, Canto CB, Bjerknes TL, Witter MP, Moser EI, Moser MB. Development of the spatial representation system in the rat. Science. 2010; 328(5985):1576–1580.

Long X, Zhang SJ. A novel somatosensory spatial navigation system outside the hippocampal formation. Cell research. 2021; 31(6):649–663.

Lukoševičius M. A practical guide to applying echo state networks. In: *Neural Networks: Tricks of the Trade: Second Edition* Springer; 2012.p. 659–686.

Lukoševičius M, Jaeger H. Reservoir computing approaches to recurrent neural network training. Computer science review. 2009; 3(3):127–149.

Maguire E, Frackowiak RS, Frith CD. Learning to find your way: a role for the human hippocampal formation. Proceedings of the Royal Society of London Series B: Biological Sciences. 1996; 263(1377):1745–1750.

Marr D. Vision: A Computational Investigation into the Human Representation and Processing of Visual Information. USA: Henry Holt and Co., Inc.; 1982.

McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:180203426. 2018; .

Moser EI, Moser MB, McNaughton BL. Spatial representation in the hippocampal formation: a history. Nature neuroscience. 2017; 20(11):1448–1464.

Muessig L, Hauser J, Wills T, Cacucci F. Place cell networks in pre-weanling rats show associative memory properties from the onset of exploratory behavior. Cerebral Cortex. 2016; 26(8):3627–3636.

O'Keefe J, Dostrovsky J. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. Brain research. 1971; .

Pascanu R, Jaeger H. A neurodynamical model for working memory. Neural networks. 2011; 24(2):199–207.

Robitsek RJ, White JA, Eichenbaum H. Place cell activation predicts subsequent memory. Behavioural brain research. 2013; 254:65–72.

Sun Y, Nitz DA, Xu X, Giocomo LM. Subicular neurons encode concave and convex geometries. Nature. 2024; p. 1–9.

Tang W, Shin JD, Jadhav SP. Geometric transformation of cognitive maps for generalization across hippocampal-prefrontal circuits. Cell reports. 2023; 42(3).

Taube JS. Head direction cells and the neurophysiological basis for a sense of direction. Progress in neurobiology. 1998; 55(3):225–256.

Trouvain N, Pedrelli L, Dinh TT, Hinaut X. ReservoirPy: An Efficient and User-Friendly Library to Design Echo State Networks. In: *Artificial Neural Networks and Machine Learning – ICANN 2020* Springer International Publishing; 2020.p. 494–505. https://doi.org/10.1007/978-3-030-61616-8_40, doi: 10.1007/978-3-030-61616-8_40.

Wills TJ, Muessig L, Cacucci F. The development of spatial behaviour and the hippocampal neural representation of space. Philosophical Transactions of the Royal Society B: Biological Sciences. 2014; 369(1635):20130409.

Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron. 2000; 27(3):623–633.

Yang W, Sun C, Huszar R, Hainmueller T, Buzsaki G. Selection of experience for memory by hippocampal sharp wave ripples. bioRxiv. 2023; p. 2023–11.