

# A hierarchical active inference model of spatial alternation tasks and the hippocampal-prefrontal circuit

Received: 2 September 2023

Accepted: 4 November 2024

Published online: 15 November 2024

 Check for updatesToon Van de Maele<sup>1,2</sup>, Bart Dhoedt<sup>1</sup>, Tim Verbelen<sup>2,4</sup> & Giovanni Pezzulo<sup>1,3,4</sup> 

Cognitive problem-solving benefits from cognitive maps aiding navigation and planning. Physical space navigation involves hippocampal (HC) allocentric codes, while abstract task space engages medial prefrontal cortex (mPFC) task-specific codes. Previous studies show that challenging tasks, like spatial alternation, require integrating these two types of maps. The disruption of the HC-mPFC circuit impairs performance. We propose a hierarchical active inference model clarifying how this circuit solves spatial interaction tasks by bridging physical and task-space maps. Simulations demonstrate that the model's dual layers develop effective cognitive maps for physical and task space. The model solves spatial alternation tasks through reciprocal interactions between the two layers. Disrupting its communication impairs decision-making, which is consistent with empirical evidence. Additionally, the model adapts to switching between multiple alternation rules, providing a mechanistic explanation of how the HC-mPFC circuit supports spatial alternation tasks and the effects of disruption.

To solve cognitive problems, such as finding the shortest route to a goal destination in a busy city, we can use so-called cognitive maps of the situation that affords flexible planning. The concept of the cognitive map has been initially popularized by Tolman, especially in the context of spatial navigation<sup>1</sup>. An extensive body of literature has investigated the neuronal underpinnings of cognitive maps for spatial navigation in humans, rodents, and other animals. This research has established the crucial importance of structures of the medial temporal lobe and especially of the hippocampal formation in the creation of codes for allocentric space, such as place cells (in the hippocampus)<sup>2,3</sup> and grid cells (in the entorhinal cortex)<sup>4</sup>. An emerging perspective to understand these findings is that cognitive mapping in the hippocampal formation might be described at the computational level as a sequence learning problem and that place cells, grid cells and other specialized codes might emerge from such sequence learning<sup>5–13</sup>. Furthermore, it has been suggested that the same hippocampal codes that support spatial navigation might also support navigation in cognitive domains<sup>14–16</sup>.

In parallel, the concept of a cognitive map can apply to the abstract state spaces that describe the sequential stages of a task at hand (for example, buying a gift in one's preferred shop and then bringing it to a friend's house) as opposed to the more fine-grained spatial codes found in the hippocampal formation. Recent studies reported that prefrontal structures, such as the orbitofrontal cortex, might host such coarse task codes, which might permit representing (for example) the current and the next stages of the task<sup>17,18</sup> as well as future navigational goals<sup>19</sup>, as opposed to a physical location.

Crucially, during goal-directed spatial navigation (and other cognitive problems), it is often necessary to combine cognitive maps of task space (to find a sequence of coarse-grained actions to progress in task space, such as ensuring the correct sequence of destinations to buy and deliver a gift) and physical space (to find a path that reaches the next intended destination)<sup>20–22</sup>. However, we still lack a comprehensive theory that explains the ways cognitive maps of physical and task space interact when executing cognitive tasks.

<sup>1</sup>IDLab, Department of Information Technology, Ghent University - imec, Ghent, Belgium. <sup>2</sup>VERSES Research Lab, Los Angeles, USA. <sup>3</sup>Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy. <sup>4</sup>These authors contributed equally: Tim Verbelen, Giovanni Pezzulo.

✉ e-mail: [giovanni.pezzulo@istc.cnr.it](mailto:giovanni.pezzulo@istc.cnr.it)

A useful starting point to understand the interaction of task-related and spatial codes is rodent memory-guided, spatial alternation tasks. A prominent example is the spatial alternation task in the W-maze studied in ref. 23. As visualized in Fig. 1, the W-maze comprises three corridors. To collect rewards, the animal has to visit the corridors in the correct order, according to an alternation rule (e.g., left, center, right, center, left, etc) that is initially unknown and has to be learned by trial and error.

A series of experiments showed that solving the spatial alternation task engages coordinated patterns of neural activity in the hippocampus (HC), which putatively encodes the spatial aspects of the task, and the prefrontal cortex (mPFC), which putatively learns the alternation rule and uses it to guide behavior<sup>24–28</sup>. For example, the HC - mPFC interaction is selectively enhanced during epochs requiring spatial working memory<sup>29</sup>. Furthermore, the coordination of neural activity in both structures spans multiple timescales, from theta sequences during navigation to reactivation (replay) activity during inter-trial periods prior to navigation—and is considered crucial for the information exchange between the two areas during the task<sup>30</sup>.

Notably, disrupting awake hippocampal reactivations (sharp wave ripples, SWR) at decision points impaired the animal's performance, but only for the aspects of the alternation task that require spatial working memory<sup>23</sup>. Specifically, disrupting awake SWRs impaired outbound decisions (from the central corridor to the left or right corridor), which require memory for immediate past outer arm location. Rather, the same SWR disruption did not impair inbound decisions (from the left or right corridor to the central corridor) that do not require memory of the past corridor; nor did it impair self-localization<sup>23</sup>. The study concludes that impairing SWRs prevents the hippocampus from providing information about past locations (and/or future options) to the prefrontal cortex, which is responsible for learning the alternation rule and to use it to guide behavior. In other words, the SWR disruption impairs communication between the hippocampus and the prefrontal cortex, hence preventing the latter from correctly inferring the current stage in task space.

The importance of hippocampal-prefrontal communication for memory-guided decisions is confirmed by another study in which animals had to learn and subsequently switch between three spatial alternation rules in the same three-arm maze<sup>31</sup>. This study showed that disrupting neural activity in the mPFC directly following hippocampal sharp-wave ripples (SWRs) (but not after a random delay) impairs the animal's ability to switch between learned rules.

In this work, we advance a computational theory that casts the interactions between the hippocampus (HC) and the medial prefrontal cortex (mPFC) during spatial alternation tasks, in terms of hierarchical active inference (HAI)<sup>32–38</sup>. Our theory has two main tenets. First, the HC and the mPFC learn cognitive maps for navigation in physical and task space, respectively, using the same statistical computations modeled as Partially Observable Markov Decision Processes,

(POMDPs<sup>39</sup>), but based on different inputs (see Section “Methods”). Second, the neural circuit formed by the HC and the mPFC realizes a HAI architecture, which can solve spatial alternation tasks. HAI rests on reciprocal, bottom-up, and top-down message passing. In the bottom-up pathway, the lower hierarchical level (encoding the HC map) infers the current location based on sensory information and communicates it to the higher level (encoding the mPFC map). In turn, the higher level infers the next goal location (in the mPFC map) and sets it as a goal for spatial navigation (in the HC map), in a top-down manner. Impairments of the message passing prevent the architecture from correctly solving spatial alternation tasks. The separation of timescales automatically emerges from the structure of HAI. The lower level passes the bottom-up message when a free-energy threshold is met, indicating it is confident it has reached its goal.


We exemplify the new theory by presenting two simulations of rodent spatial alternation tasks, with intact and impaired HC-mPFC interactions<sup>23,31</sup>. Our first simulation, presented in Section “Experiment 1: solving the spatial alternation task”, brings three main conclusions. First, the model is able to learn the spatial structure of the maze (HC map) and spatial alternation rules (mPFC) by navigating in the environment. Second, when the HC - mPFC circuit is configured as a HAI system, it effectively solves spatial alternation tasks, through bottom-up and top-down message passing. Third, disrupting the interaction from the HC to the mPFC breaks the spatial memory and impairs the animal's ability to make outbound but not inbound decisions, as observed empirically<sup>23</sup>. Our second simulation, presented in Section “Experiment 1: solving the spatial alternation task”, shows two additional features of the model. First, in a task in which three alternative spatial alternation rules are in play, the HC - mPFC circuit permits inferring the current rule. Second, selectively inhibiting the mPFC impairs this ability and provokes distinct behavioral patterns, as observed empirically in the rodent study reported in<sup>31</sup>.

## Results

### The W-maze setup and the hierarchical active inference (HAI) agent

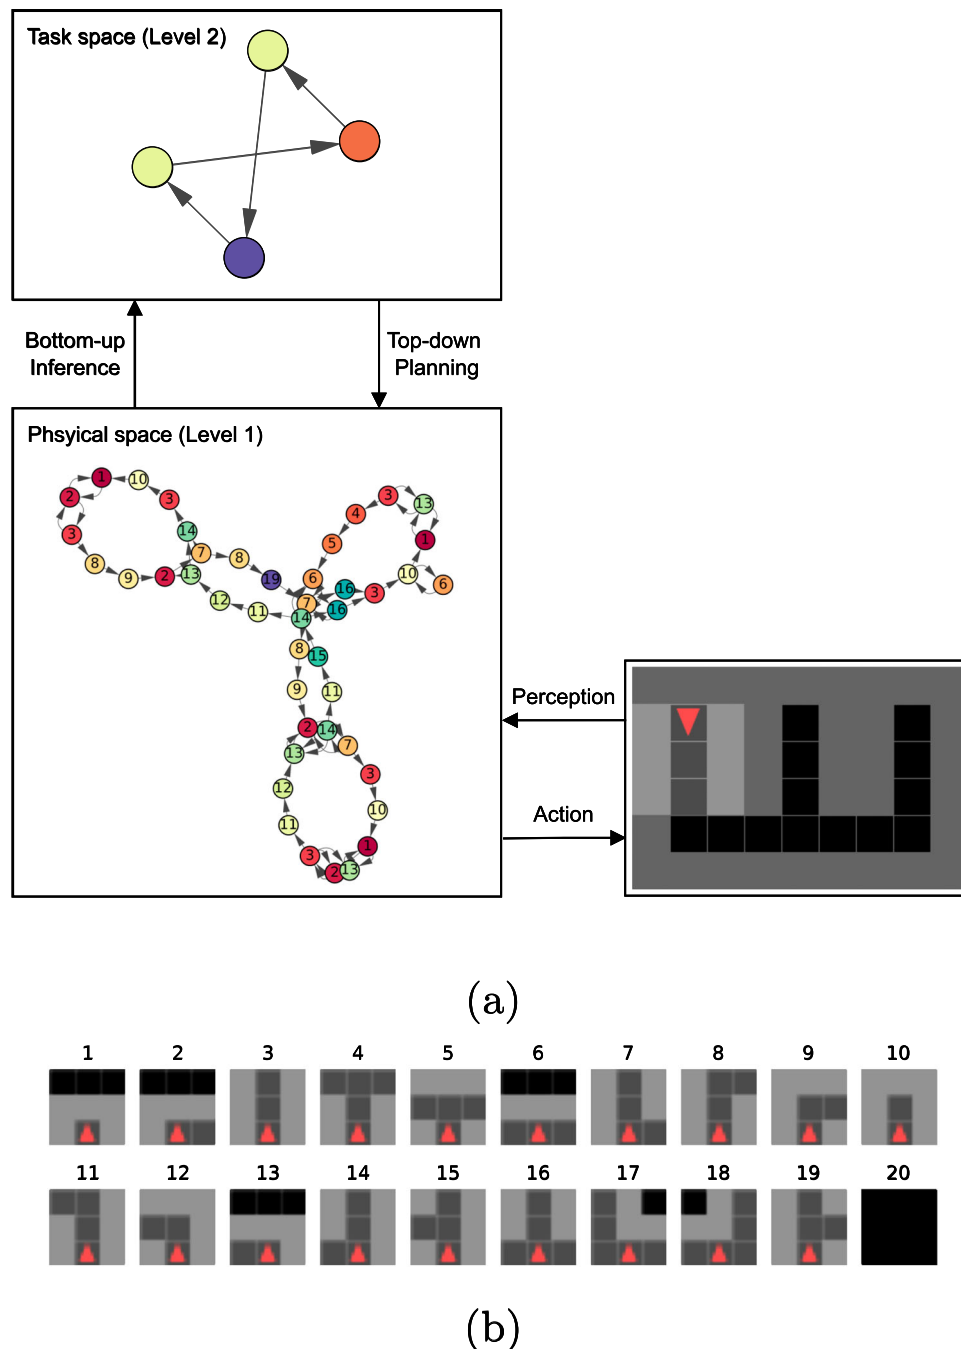
In this work, we consider the spatial alternation task in the W-maze of ref. 23, see Fig. 1. In this task, an animal (or an artificial agent) can acquire rewards at the end of each corridor, providing that the corridors are visited in the correct order (e.g., left, center, right, center, left).

In this section, we present a HAI agent<sup>36</sup> for solving the spatial alternation task. Central to this framework is the fact that agents are endowed with a generative model, describing how observed outcomes are generated from a hidden, unobserved state, which is influenced by the agent's actions. The agent infers these hidden states by minimizing the free energy functional with respect to its belief over the state and



**Fig. 1 | Spatial alternation tasks (LCRC).** This figure exemplifies the W-maze used in ref. 23 to study spatial alternation and the role of the hippocampal - prefrontal dynamics. The W-maze is characterized by three separate corridors. The animal can acquire the reward at the end of each corridor, however, to receive the reward it must do so in the correct order (e.g., Left (L), Center (C), Right (R), Center (C), Left (L), etc). The spatial alternation task implies that when the animal is in the center

position, it has to make a (difficult) outbound decision: it has to go either to the left or the right, depending on where it comes from. This outbound decision, therefore, requires a spatial memory component. In contrast, when the animal is in one of the side corridors (left or right), the only correct action is to move to the center. This inbound decision is considered simpler since it does not require a memory component. Trajectory order is indicated by the shade where lighter is earlier.



**Fig. 2 | Illustration of the hierarchical active inference (HAI) agent and the W-maze setup. a** The hierarchical generative model of the HAI agent. The figure shows that the two layers of the model include cognitive maps of physical space (level 1) and task space (level 2). These cognitive maps essentially represent transitions between learned locations and stages of the task, respectively. Both maps are learned using a probabilistic sequence learning algorithm: the clone structured cognitive graph (CSCG<sup>7</sup>). Within each block, a subset of states from the CSCG is visualized, namely, the active states when pursuing the alternation rule (LCRC). The dynamics between the physical (level 1) and task space (level 2) CSCG in the hierarchical model are regulated by bottom-up message passing (that supports inference) and top-down message passing (that supports planning). Bottom-up

inference messages from the physical space model are passed to the task space model, while top-down actions drive planning from the task-space model to the physical space model. An essential aspect of these interactions is that the two levels operate on a different temporal scale. Transitions at the physical level, through the agent's movement, occur at every time step. Rather, the task level only transitions when the intermediate goal (of reaching a certain position in the maze) is achieved. This abstraction allows for hierarchical planning. Finally, the interaction with the environment using an action-perception loop with the world, the W-maze, is shown on the right. **b** The 20 distinct observations the agent can encounter in the W-maze created in the Minigrid environment<sup>40</sup>.

acts to minimize its expected free energy over time<sup>36</sup>. For more details on these mechanisms, the reader is referred to Section “Hierarchical active inference”.

In our simulations, the HAI agent is endowed with a hierarchical generative model of two layers, which is schematically illustrated in

Fig. 2 (see Section “Hierarchical active inference” for a formal introduction). In this generative model, the lowest layer takes up the functional role of the hippocampus, namely encoding the spatial structure of the maze<sup>2</sup>. In contrast, the highest layer takes the functional role of the prefrontal cortex, namely encoding the structure of

task space<sup>18</sup>. The agent learns both maps using the same sequence learning algorithm: the clone structured cognitive graphs (CSCG), see Section “Methods” for details. The latent states learned by the CSCG at the two levels are schematically illustrated within the two boxes of Fig. 2.

Crucially, the HAI agent performs inference (of where it is in physical space, i.e., its pose, and in task space) and planning (of the next goal in physical and task space) through message-passing between the two hierarchical levels. As shown in Fig. 2a, level 1 observes the structural aspects of the environment through its bottom-up observations, i.e., the three-by-three grid around the agent as shown in Fig. 2b and selects local movements in physical space such as turn left, turn right or move forward to reach any goal location (which is fed by the highest level, see below).

Once level 1 has navigated towards its local goal, it passes a bottom-up message containing the current state of the agent to level 2. In turn, level 2 processes this message and infers the next stage of the task to achieve a reward (as it is endowed with a prior preference to achieve reward at each step, see Section “Hierarchical active inference”). Actions at level 2 are abstract and are matched to the bottom-up received observable states from level 1. Since level 1 states encode the agent’s pose, these actions essentially mean move to target location. We specifically encode only actions corresponding to target locations at the end of a corridor (i.e., observation 1 in Fig. 2b). The selected action is sent to level 1 in a top-down manner and corresponds to the goal location for level 1. Once level 1 reaches the goal location, this process repeats.

Finally, the right part of Fig. 2 illustrates the fact that the lowest level of the model is responsible for the action-perception loop with the environment—here, an implementation of the W-maze in the Minigrid<sup>40</sup> simulator. Observations in the Minigrid<sup>40</sup> are acquired as a square around the agent. In this work, we chose a range of three, yielding three-by-three observations which are then mapped to a unique one-hot index. For each step, the agent also observes a binary value indicating the reward presence. The actions the agent can perform are turn left, turn right and move forward, which are also represented as a one-hot index. See Section “Hierarchical active inference” for more details on the model implementation

### Experiment 1: solving the spatial alternation task

We replicate the task and experiment from a study on rodents<sup>23</sup> where the animal is taught this specific alternation task. In this study, disruption of the hippocampal SWR reduces the performance on outbound trajectories, but not on inbound trajectories.

We first allow the HAI agent to learn the two cognitive maps for solving the spatial alternation task, using two CSCGs (see<sup>7</sup> and Section “Methods” for details on the learning procedure). The learned maps for the W-maze are shown within the two boxes of Fig. 2. For ease of visualization, within each box, only a subset of states from the CSCG is shown, namely, the states that are active when correctly pursuing the alternation rule. We denote the sequence of this alternation rule as LCRC (for Left, Center, Right, Center).

The hippocampal map on the lowest level organizes the 20 observations that the agent encounters during navigation (illustrated in Fig. 2b) into a coherent graph, which nicely reflects the 3 corridors of the W-maze. In this map, the circles reflect the states learned by the CSCG, while the transitions reflect the agent’s spatial actions (e.g., turn left, turn right, or move forward). The nodes are color-coded (and numbered) according to the observations that they encode. Note that while the agent encounters the same observations multiple times during navigation (i.e., the observations are aliased), the CSCG correctly reflects the fact that they are part of different behavioral sequences. For example, observation 1 (that corresponds to the location in which a reward can be collected) appears three times, at the

apex of each of the latent sequences that represent each of the three corridors.

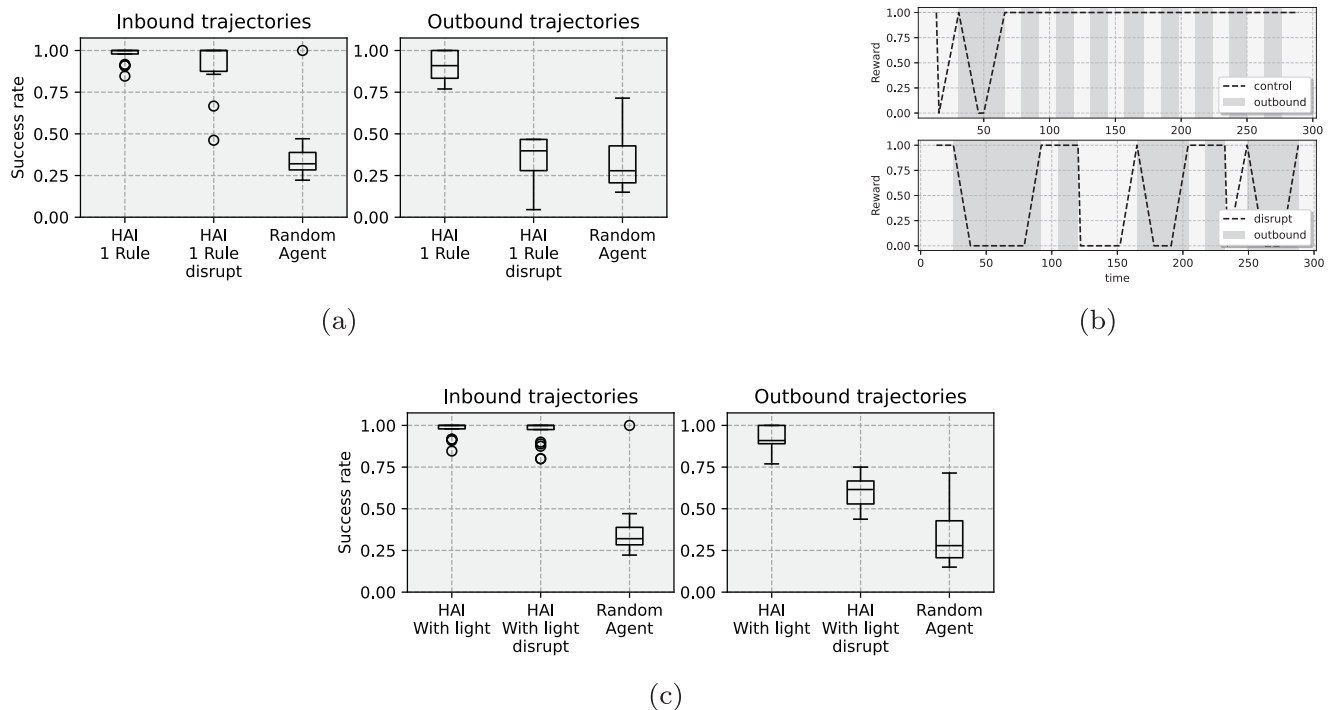
The prefrontal map on the highest level corresponds to a much smaller graph that encodes the task-specific sequence of corridors that the agent has to visit to secure rewards. In this graph, the nodes correspond to (color-coded) corridors in which rewards have been collected and the edges correspond to higher-order actions (go to corridor). Note that there are two nodes of the same color: both correspond to the central corridor but in different phases in task space (i.e., after the left and the right corridors, respectively). This indicates that the agent has successfully learned a low-dimensional representation of (or a finite state machine for) the task rule LCRC.

We then investigate whether the HAI agent equipped with the learned cognitive maps for physical and task space is able to solve spatial alternation tasks. For this, we endow the HAI agent with a prior preference for maximizing the reward and place it in a random pose in the W-maze. We then let the agent run for 300 steps, and record whether it selected the correct or the incorrect corridor to find the reward. We repeat this experiment for 20 trials, to collect statistics. Figure 3a shows that the HAI agent has near-perfect performance during the easiest, inbound decisions and very high (above 80%) performance during the harder, outbound decisions, closely approximating the empirical results reported in<sup>23</sup>. Furthermore, as expected, it significantly outperforms a Random agent that selects the next corridor randomly ( $t$ -test  $p = 0$ ). Figure 3b shows that after a few time steps, once the HAI agent has inferred where it is (in both physical and task space), it consistently follows the rule and is able to collect rewards at every step.

We performed two control experiments, which reveal that the HAI model is robust to a wide choice of parameters and to greater levels of uncertainty (see Supplementary Materials). Our control experiments also illustrate that in scenarios characterized by greater uncertainty, the expected free energy functional used by the HAI model to select policies is more advantageous compared to the sole objective of reward maximization that is common in economic and reinforcement learning settings. The expected free energy functional effectively balances two components: a pragmatic imperative to maximize reward (e.g., to reach the next subgoal to secure reward) and an epistemic imperative to gain information (e.g., to reduce uncertainty about the current pose, by going to places where unambiguous observations could be found, e.g., the T-junction), see Section “Hierarchical active inference”. By comparing HAI agents with and without the epistemic imperative, we found that the pragmatic imperative to secure rewards is sufficient to address a W-maze with low uncertainty, but adding the epistemic imperative to reduce uncertainty about one’s hidden state (here, the pose) significantly increases performance in a W-maze with greater uncertainty (see Supplementary Materials). This is because an agent endowed with epistemic uncertainty can explicitly plan to self-localize (e.g., by visiting unambiguous locations, such as the bottom of corridors), which is especially advantageous with greater uncertainty<sup>36,41</sup>.

Having established that the HAI agent correctly solves the spatial alternation task, we next ask whether an impairment of spatial memory disrupts its performance in outbound decisions, as shown experimentally by ref. 23. In this study, the experimenters disrupted hippocampal SWR at decision points and observed that the disruption prevented access to spatial memory—rendering the rodents unable to make correct outbound decisions—but left hippocampal spatial codes intact. The experimenters, therefore, concluded that the disruption was caused by a failure of communication or updating at the level of the PFC.

In analogy with this procedure, we realized a variant of the HAI agent (HAI disrupt) in which we implement the SWR disruption by preventing the belief updates (through the transition model) at the second level of the hierarchical model. By doing this, we remove the



**Fig. 3 | Experiment 1: solving the spatial alternation task with intact and disrupted spatial memory.** **a** The success rate in the spatial alternation task for the HAI agent without disruption, the HAI agent with disruption (HAI disrupted), and the Random agent, separated for both in- and outbound trajectories. The success rate is computed over 20 trials per agent, with 300 steps per trial, only recording the presence of reward at the end of the selected corridors. Data is represented as the mean success rate, box limits indicate the central 50% of the data, with the horizontal line representing the median value. The whiskers represent the minimum and maximum values and outliers are visualized as circles. **b** The reward over time for the HAI agent and a disrupted HAI agent. The shaded areas in the plot indicate outbound trajectories. **c** The success rate for the HAI agent with (HAI

disrupted) and without (HAI) disruption in the alternative light environment. Data is represented as the mean success rate, box limits indicate the central 50% of the data, with the horizontal line representing the median value. The whiskers represent the minimum and maximum values and outliers are visualized as circles. The success rate is computed over 20 trials per agent, with 300 steps per trial, only recording the presence of reward at the end of the selected corridors. The agent is disrupted by disabling the connection from the transition model of the POMDP, separated for in- and outbound trajectories. We observe that the performance on inbound trajectories does not change, while for outbound trajectories this drops to random chance.

agent's spatial memory about task space. However, the agent is still able to receive the bottom-up message from level 1 and thus knows where it currently is in physical space—in keeping with the finding that spatial codes are intact after SWR disruptions<sup>23</sup>.

When we compare the HAI agent with and without disruption in the spatial alternation task, we observe the same behavior as the experimental study: the HAI disrupted agent correctly addresses inbound decisions but makes outbound decisions at a chance level, see Fig. 3a. The simulation shows that the disruption only impairs the outbound trajectories (please see the Supplementary Materials for additional simulations exploring how the choice of the gamma parameter influences the results). A more detailed example of this pattern of results can be appreciated in Fig. 3b, which shows that the HAI disrupted agent tends to miss rewards, but only during outbound decisions.

Finally, we performed a control simulation to identify more clearly why the disruption impairs the ability to make the outbound decisions. It is worth noting that in the higher-level (prefrontal) map of the HAI model, the task variables (color-coded nodes in Fig. 2) are hidden (i.e., they do not have corresponding observations), unlike the spatial states in the hippocampal map that have corresponding observations. Hence, the agent's estimate of the hidden state at the level of task space depends entirely on spatial memory, as encoded in the model's transition function – which we disrupt.

To assess whether the partial observability of task space is key to explain the effects of the disruption, we considered an alternative (control) environment where the task variable is observable, however,

the component required for spatial memory is unobserved. We consider the following variant of the W-maze environment: the maze is augmented with a light signal. Depending on where the agent just collected reward, a different color is shown. We implement this by adding an additional observation modality containing a categorical variable, indicating the previous rewarding corridor.

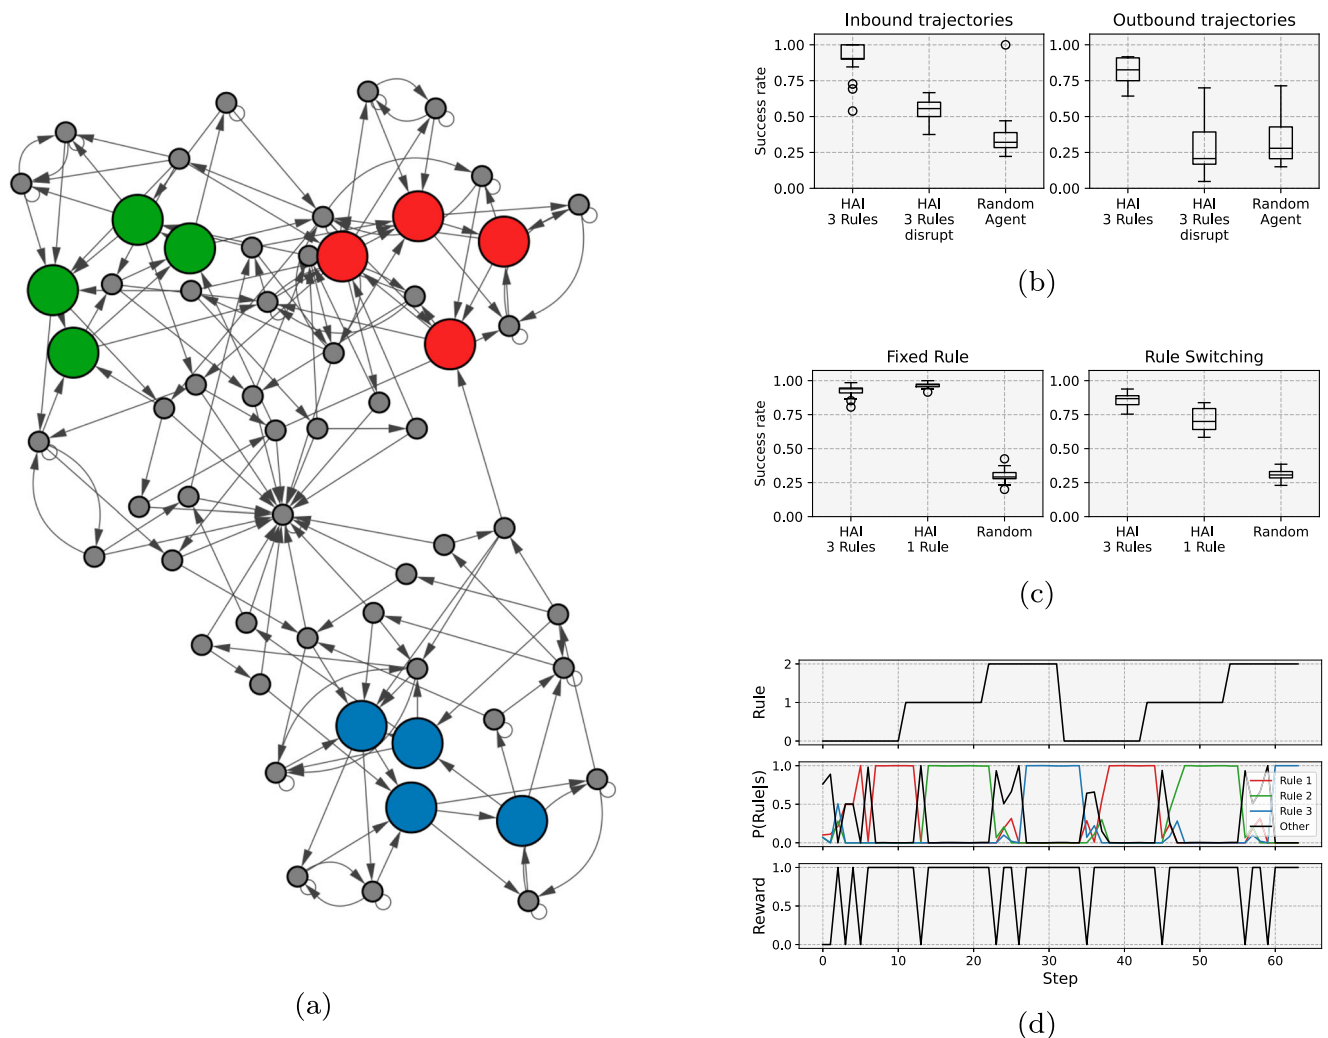
We repeated the same disruption experiment as before and found that the disrupted HAI agent maintains a high success rate for inbound decisions, but makes outbound decisions at chance level (Fig. 3c). This control experiment therefore indicates that disruption of spatial memory impairs the ability to make the outbound decisions, even when task variables are observable.

## Experiment 2: learning multiple rules and flexible switching between them

In the next experiment, we investigate whether the HAI agent is able to learn multiple alternation rules as proposed in<sup>31</sup>. We consider three spatial alternation rules, where for each rule, a different corridor serves as the alternation point (LCLR, LCRC, and RCRL). Note that for simplicity, we use the same structure and notation as in the previous experiment with the W-maze, despite the study of multiple alternation rules of ref. 31 has been conducted on a three-arm radial maze, not a W-maze.

In this experiment, the HAI agent is endowed with a cognitive map of physical space which is the same as Experiment 1 (since the maze is always the same). However, it has to learn a new cognitive map of task space, which now comprises the dynamics of the three rules. For this,





**Fig. 4 | Experiment 2: Multiple spatial alternation rules (LCLR, LCRC, and RCRL).** **a** Representation of the states used in the CSCG for the three distinct rules. States again map to the conjunction of the visited corridor and the presence of reward. Colors in the graph indicate the rule for which the state is used (red for rule 1, green for rule 2, and blue for rule 3). The gray nodes represent other states, either without reward or in transition between multiple rules. **b** Disruption experiment for the model that learned three rules, evaluated on a single rule. The success rate is computed over 20 trials per agent, with 300 steps per trial. Data is represented as the mean success rate, box limits indicate the central 50% of the data, with the horizontal line representing the median value. The whiskers represent the minimum and maximum values and outliers are visualized as circles. **c** Average reward for the agents trained on 3 rules, 1 rules and random selected. The

case where the rule is constant (left), and the rule switches every 300 steps to a randomly selected next rule (right) is considered. This experiment was repeated over 20 trials of 900 steps. Data is represented as the mean success rate, box limits indicate the central 50% of the data, with the horizontal line representing the median value. The whiskers represent the minimum and maximum values and outliers are visualized as circles. **d** Performance during a single representative trial of 1000 steps, during which the rule switches randomly, every 150 steps. The top plot shows the rule that is currently in play. The center plot shows the belief over rule, which is computed by first extracting states belonging to a single rule, and then measuring the likelihood of being in one of these states. The bottom plot shows the rewards collected over time.

we let the agent explore the maze for 8000 steps, with alternating task rules after blocks of 1000 steps (in task space, i.e., visits to corridors). To aid the learning process, we guide the sequence and select the correct action 75% of the time, and random otherwise.

The learned cognitive map of task space is visualized in Fig. 4a. For ease of visualization, we extract the states that are active when the agent has correctly inferred the rule, after a warmup period (15 steps) in each block. Similar to Section “Experiment 1: solving the spatial alternation task”, the states of the task space correspond to the distinct colors, however since the different rules have distinct transition dynamics, these are mapped to unique states. Figure 4a uses different colors to visualize the states of the three different rules (red for rule 1, green for rule 2, and blue for rule 3). The smaller dark gray states are states used for transitioning between rules, or encode the corridors where no reward is found. We observe that for each of the individual

rules, four states are active. This is an indication of correct learning since four is the optimal number required for each alternation problem.

Then, we evaluate whether the HAI agent is able to infer which rule is currently in place—in order to continue collecting rewards when the rule switches. For this, we test the HAI agent (endowed with a preference to maximize reward) in a task in which the rule switches every 300 steps.

Figure 4b shows the success rate of various agents for a fixed rule over 20 trials. From the success rate of the HAI agent trained on the three rules (HAI 3 rules), we can conclude that the agent properly encodes the rule. When we apply the same disruption as in Experiment 1 (Section “Experiment 1: solving the spatial alternation task”), the success rate for inbound trajectories is now also affected (HAI 3 rules disrupt). This drop in performance is caused by the fact that when

considering multiple rules, there is also ambiguity over which rule is currently in play, for which the spatial memory component is crucial.

Figure 4c compares the performance of various agents in two scenarios: in the former, there is a single rule, which remains fixed throughout the experiment (left plot), whereas in the latter, there are three rules, which switch randomly every 300 steps, over 20 trials. The left plot shows that both the Active Inference agents trained on three rules (HAI 3 rules) and on one rule (HAI 1 rule) achieve a very high success rate in the first scenario and both outperform a random agent (random). This result indicates that learning multiple rules does not decrease performance when a single rule is in place. The right plot shows that the Active Inference agent trained on multiple rules (HAI 3 rules) achieves a high success rate in the second scenario, too, indicating it correctly learned all the rules and successfully alternates them. The Active Inference agent trained on a single rule (HAI 1 rule) achieves a lower success rate, as expected, but still it outperforms a random agent. This result shows that learning only a subset of the task rules could lead to a relatively good performance in the W-maze, but still it is possible to discriminate between agents that learn the task completely, partially or select randomly from their behavior.

Figure 4d illustrates the behavior of the agent during a single, representative trial. As shown in the figure, shortly after the rule switches (top panel), the HAI agent is able to correctly update its belief about the current rule in its plan (center panel) and secure rewards (bottom panel). The inference of the rule currently in play follows standard Bayesian approach: when the rule changes, the agent receives unexpected observations (since expected rewards are not delivered) and correctly makes a transition to the subspace of the task rule map that encodes the most likely rule. In the center panel of Fig. 4d, we visualize the belief over each rule as the probability of being in one of the four states used in a particular rule (the rules are color-coded as Fig. 4a) in or in another phase if the state does not occur in either of the rules (grey line). When the agent has a high probability of a rule, it no longer misses the reward. In some cases, it misses the reward because there is a potential overlap between the two rules. For example, around step 200 it could be either of the rules, as it only picks the right one once, and the wrong one after—which could happen for each corridor in each of the rules. Note that in this architecture, the belief about the currently active rule is implicit in the model, but it could be made explicit by adding a hierarchical layer that maintains a probability distribution over the map or the task one is currently in, as in<sup>10</sup>.

Another important finding of the study of ref. 31 is that disrupting optogenetically the medial prefrontal cortex (mPFC) immediately after a hippocampal SWR is detected significantly impairs performance, by increasing the occurrence of three distinct maladaptive patterns: (i) rotated alternation, where the agent follows a different rule, (ii) back and forth, where the agent alternates between two corridors, and (iii) circling behavior, where the agent iterates over all corridors in a cyclic pattern (see the Supplementary Materials for a graphical illustration of

these patterns). The authors reported a proportion of 35% alternation, 20% rotated alternation, 15% back and forth, and 10% circling.

The behavioral patterns that are generated from the same type of disruption as conducted in experiment 1 - where the transition model of the POMDP is lesioned - are investigated. We match the patterns observed in the biological counterpart of this experiment<sup>31</sup>, and evaluate the trajectories generated by the agent trained on the three rules, for an agent with and without the spatial memory disruption from Section “Experiment 1: solving the spatial alternation task”, as well as a random agent. The results are reported in Table 1. This shows that, when lesioned, our model reproduces all four behavioral patterns seen in lesioned animals<sup>31</sup> (alternation, rotated alternation, back and forth, and circling) although in slightly different proportions compared to the empirical study (please see the Supplementary Materials for a simulation how the choice of  $\gamma$  affects these proportions). We observe that for the disrupted agent, a large part of the behaviors can be classified as back-and-forth behavior. The specific proportions of categorization are a function of the agents’ habit policy  $E$ . This habit acts as a prior over the selected action (of level 2). When no temporal information is considered through the lesion, this prior will dominate, and the agent will act according to the statistics of the training data. This can be observed by looking at HAI disrupt (1), where the agent behavior falls back to either the alternation or the back-and-forth case. Finally, we also note that the patterns generated by the random agent are generally classified as the other category, indicating that the behavior generated by our disruption is not random.

Furthermore, we performed the same experiment for an agent trained on a single rule. We observe that both rotated alternation and circling behavior drops to near 0% as shown in the table indicated by HAI control (1) and HAI disrupt (1). This is due to the agent not considering other rules, when falling back to its habit policy. To the best of our knowledge, the pattern of errors in animals trained with a single rule have not been systematically studied and hence our results could be considered as a prediction of the HAI model.


The previous simulations illustrate that disruptions of the HAI model permit reproducing empirical findings of rodent studies<sup>23,31</sup>. However, our model permits realizing synthetic experiments—for example, other types of disruptions—that have not been studied empirically so far, but which could be interpreted as predictions of the model.

Here, we investigate the potential effects of a disruption of hippocampal-prefrontal communication that completely impairs the communication between the two levels of the HAI model. Specifically, in this simulation, we prevent level 1 from sending any bottom-up information to level 2. This is a severe impairment, since hierarchical inference rests upon the reciprocal messages passing between the two levels—and level 2 only receives observations from level 1, not from the environment. Figure 5 shows the differences between the physical trajectory of the intact HAI agent (left) and the disrupted HAI agent (right) in which communication between the two levels is impaired. The shade of the trajectory indicates the order of visits (lighter is earlier). The figure shows that while the intact HAI agent follows the

**Table 1 | Behavioral patterns during disruption**

	HAI control (1)	HAI disrupt (1)	HAI control (3)	HAI disrupt (3)	Random
Alternation	86.33%	49.07%	76.00%	16.80%	3.47%
Rotated alternation	0.00%	0.80%	1.07 %	13.55%	8.24%
Back and forth	2.68%	38.40%	2.40 %	38.48%	8.86%
Circling	1.61%	1.60%	8.80 %	17.34%	5.86%
Other	9.38%	10.13%	11.73%	13.82%	73.75%

Classified observed behavioral patterns according to<sup>31</sup> in different scenarios for agents trained on the three rules indicated by (3) and a single rule indicated by (1).



**Fig. 5 | Disruption experiment for rule switching.** when communication between the level 1 and level 2 model is suppressed, perseveration behavior is observed. Trajectory order is indicated by the shade where lighter is earlier, noise was added to the trajectory to visualize overlapping visits.

correct rule, the disrupted HAI agent perseveres in choosing the center corridor.

Interestingly, the agent does check at the unambiguous T-junction to self-localize and ensure that it is still in the correct corridor (which is a feature of the epistemic dynamics of active inference, see Section “Hierarchical active inference”), but never selects the next goal in the task rule. This is because, without the bottom-up message from level 1, level 2 is unable to correctly recognize that one phase of the task has been achieved and hence never updates the top-down plan. This result illustrates the possibility to use the model to generate predictions—in this case, about how disrupting communication between the two levels produces perseveration—which could be subsequently tested empirically.

## Discussion

Here, we advanced a computational theory of hippocampal (HC) - prefrontal (PFC) interactions during cognitive tasks that require navigating in both physical and task space—such as spatial alternation tasks. Empirical studies of spatial alternation have assessed that they depend on the animal’s spatial memory, which is at least in part maintained by HC-PFC communication<sup>25,29,30,42–44</sup> and that disruption of this communication can impair difficult (outbound) decisions<sup>23</sup> and the ability to switch between multiple rules<sup>31</sup>.

Our computational model is based on—and unites—two established theories. The former is a theory of cognitive map formation, based on a statistical sequence learning algorithm: the clone-structured cognitive graph (CSCG)<sup>7</sup>. Previous studies assessed that CSCG are computationally effective and have biological validity since they are able to successfully reproduce a number of empirical observations about cognitive map formation in the hippocampus, including for example the emergence of place cell coding and splitter cells<sup>8,45</sup> and orthogonalized state representations<sup>45</sup> (see also<sup>5,6</sup> for alternative proposals about the computational principles that underlie spatial map formation). Here we extended this body of work by showing that CSCGs can be used in a hierarchical scheme to learn not just cognitive maps for physical space (putatively linked to hippocampal computations), but also more abstract cognitive maps for task space (putatively linked to prefrontal computations). Furthermore, we used learned CSCGs as components of a hierarchical generative model for active inference<sup>36</sup>. Active inference is a normative framework to model sentient behavior in terms of free energy minimization and (approximate Bayesian) inference over a generative model, which is gaining popularity in cognitive neuroscience. We have shown that by combining two learned CSCG maps (for physical and task space), it is possible to design an effective HAI agent able to solve spatial alternation tasks. Interestingly, this scheme affords (hierarchical) planning by only using local computations—that is, top-down and bottom-up message-passing between the two hierarchical levels.

Furthermore, by simulating the interruption of HC - PFC communication in our model, we were able to correctly reproduce impairments of difficult (outbound) decisions<sup>23</sup> (Experiment 1) and of correct rule switching, unveiling the same kind of maladaptive behavior—rotated alternation, back and forth, and circling behavior (Experiment 2)—observed empirically in rodents<sup>31</sup>.

Our model suggests that the selective impairment of outbound decisions provoked by hippocampal SWR disruptions<sup>23</sup> is due to the fact that the SWRs convey messages to higher structures, like the PFC, which are used to update a belief about the current stage of the task (specifically, this message is key to propagate the belief about task state at level 2 over time). This interpretation is in keeping with the idea of<sup>23</sup> that the impairment is at the level of spatial memory, not of hippocampal place coding, but to our knowledge, ours is the first work that provides a mechanistic model of this theoretical proposal. Furthermore, our model suggests that the maladaptive behavior found in ref. 31 could be due to the impossibility for the higher, prefrontal

component (level 2) to correctly update its belief, based on bottom-up message passing from the hippocampal component (level 1). This perspective is coherent with the finding of ref. 31 that the only disruption of mPFC that prevents spatial rule switching is one that directly follows hippocampal SWRs—hence highlighting the importance of coherent HC-PFC reactivations to solve spatial alternation tasks. Finally, our simulations suggest that other, more severe interruptions of HC-PFC communication could produce a specific pattern of maladaptive behavior—namely, perseverative behavior—that differs from those observed in the above disruption experiments. This is a prediction that remains to be tested empirically.

Besides helping understand HC-PFC interactions, our simulations suggest that looking at the animal’s epistemic behavior (e.g., the selection of actions to self-localize and reduce uncertainty about the current pose) during uncertain tasks could be important to elucidate its strategy; in particular, whether it aims to maximize reward or also to minimize its uncertainty – as predicted by active inference<sup>36,41,46</sup> (see the Supplementary Materials). Future studies might assess whether epistemic imperatives are important drivers of behavior during spatial alternation or similar tasks, especially in conditions of high uncertainty, such as when the animal is randomly placed in a maze (as in our simulations) or when spatial contingencies or task rules change unexpectedly.

It could, in principle, be possible to solve the navigation tasks studied in this article using a non-hierarchical generative model with a single map (and a single CSCG) that encompasses both spatial and task-related components. However, the hierarchical structure of the HAI generative model used in this study better reflects the implicit division of labor between HC and PFC circuits, which is well-established empirically in rodent studies. For example, inactivating prefrontal structures during navigation tasks tends to disrupt rule-related contingencies and deliberative behavior, while sparing spatial representation<sup>31,47</sup>. Therefore, our model reflects the widespread perspective that goal-directed navigation depends on the coordinated interplay between (inferential) processes at two levels, which could be associated with HC and PFC structures<sup>20,25,30,43</sup>. The division of labor between HC and PFC is also central to other prominent accounts, such as the Complementary Learning Systems framework, which highlights that faster learning in HC facilitates slower learning in neocortex—with the latter integrating across episodes to extract semantic structure<sup>48</sup>. Furthermore, from a computational perspective, the hierarchical structure of the HAI model affords a useful factorization: learning a new rule in a known maze, as we did in Experiment 2, only requires retrain the higher-level CSCG, while leaving the lower-level CSCG unchanged. This might be more problematic when using a single CSCG that combines spatial and task-related information.

The current study has several limitations that will need to be addressed in future studies. First, for efficiency reasons, we learned the CSCGs offline (before embedding them into the HAI), using a simplified procedure: we used predefined trajectories that exhaustively covered the W-maze as inputs for the cognitive map of physical space and 75% correct trajectories as inputs for the cognitive map of task space. In the future, it would be interesting to train CSCG online, similar to ref. 49, by guiding the exploration through active inference dynamics<sup>33,36,41</sup>. A second research avenue is to relax the separation of the timescales between the two levels, by selecting their inputs (e.g., level 1 takes all sensory observations as inputs, whereas level 2 only considers observations that could be rewarding—and in particular, observation 1 in Fig. 2b). In the future, it would be interesting to explore methods to learn hierarchical models with multiple timescales<sup>50,51</sup> and effective state spaces for navigation and for task rules in self-supervised (and/or reward-guided) ways, as shown in prior work<sup>10,17,52,53</sup>. This might also help understand the reciprocal influences between cognitive map learning at different levels and in different (e.g., prefrontal versus hippocampal) brain structures. A third



challenge is to avoid having the agent learn from scratch each new maze or rule. Recent work in transfer learning shows that it is possible to reuse existing cognitive maps or latent task representations to learn novel and similar tasks much faster<sup>10,53–55</sup>. Extending our architecture with transfer learning abilities would be important to provide more accurate models of how animals learn cognitive maps, especially given the strong evidence for the reuse of existing neural sequences and cognitive maps in the hippocampus<sup>56,57</sup>.

Another limitation of the model presented in this study is that in the CSCG of the task layer, we made explicit the fact that the relevant task states are the ends of corridors because the animals can (only) acquire rewards in such states. This design choice reflects the animals' knowledge, given that in the spatial alternation task<sup>23</sup>, the three arms had reward food wells at their endpoints. However, future work should consider generative models that learn autonomously what the relevant task states are, even if they are not explicitly cued as in<sup>23</sup>. There is an established literature showing that latent task states relevant for cognitive control can be learned autonomously using Bayesian (nonparametric) methods<sup>52,53,58,59</sup> and that task rules can be learned using deep reinforcement learning<sup>60</sup>. However, it remains to be assessed in future research how to integrate these methods in the HAI model.

Future studies might consider how to adapt the HAI agent to robot navigation and planning. There is increasing interest in applying active inference to robotic domains, given the appeal of its unified approach to control, state estimation, and world model learning<sup>61–63</sup>. HAI is especially appealing since it facilitates planning over long time horizons. Indeed, the computational burden required to calculate the expected free energy of long policies is high, but it can be substantially reduced by splitting work between (higher-level) policies that select (a sequence of) subgoal(s) along with (lower-level) policies that select the specific actions to reach each subgoal<sup>64</sup>. For example, one study built a hierarchical model for robot navigation, using multiple layers of recurrent state space models<sup>65</sup>. Another study realized a hierarchical model for next-frame video prediction, using a subjective timescale for the predictions<sup>66</sup>. In another study<sup>67</sup>, a single layer CSCG was embedded within the active inference framework and was able to support navigation in different mazes. An interesting direction of future work could be adding a (learned) hierarchical layer that maps the high-dimensional observations space to a discrete state space and stack the proposed HAI on top with multiple layers of CSCGs. This would potentially afford abstract reasoning and planning for complex tasks, directly from sensory observations such as pixels.

Finally, future studies might consider in more detail the functional role and content of hippocampal reactivations and replay in the hippocampus and other brain structures, such as the prefrontal cortex and the ventral striatum<sup>68–81</sup>. Previous studies suggested that replay might play different roles, which range from memory functions to planning, compositional computation, and the optimization of the brain's generative models<sup>10,82–87</sup>. However, these works have mostly considered hippocampal replay, not coordinated replay in the hippocampus and the prefrontal cortex (and other brain areas). It might be interesting to combine the insights of these studies with the hierarchical architecture of the HAI agent, to test whether (for example) planning and model optimization benefit from the combined replay of multiple brain structures, as opposed to local replay in the hippocampus.

## Methods

In this paper, we develop a hierarchical planning agent by combining two components: we use CSCG to learn cognitive maps of physical and task space; and HAI to form hierarchical plans. In the next sections, we explain the two key components of our approach: CSCG (Section “Clone-structured cognitive graphs (CSCG) for learning cognitive maps of physical and task space”) and HAI (Section “Hierarchical active

inference”). Finally, we explain how we combined these (Section “Casting CSCGs as partially observable Markov decision processes”).

## Clone-structured cognitive graphs (CSCG) for learning cognitive maps of physical and task space

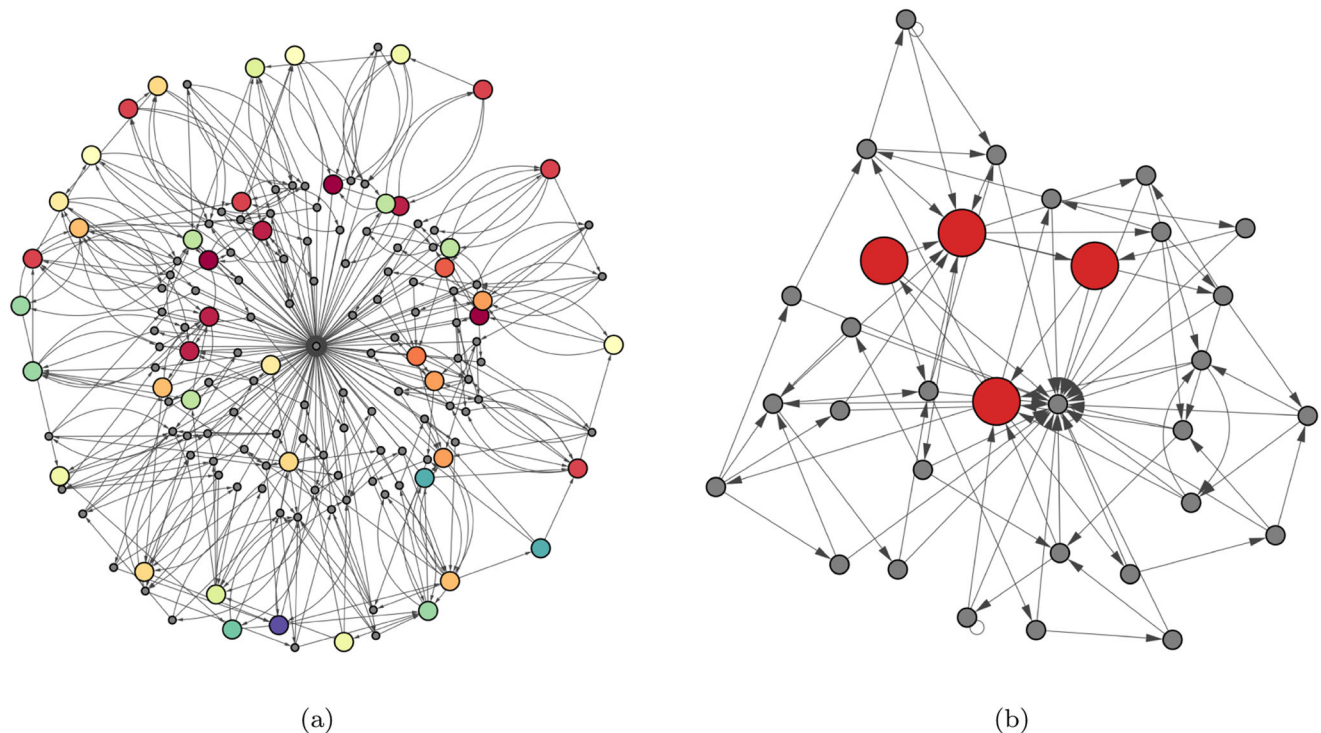
Clone-structured cognitive graphs (CSCGs) are a probabilistic model for representing sequences of data, e.g., a sequence of action-observation pairs<sup>7,88</sup>. They are a special case of Hidden Markov Models (HMM), where each observation maps to a subset of the hidden states: the so-called clones of this observation. While these states have the same observation likelihood, they differ in the implied dynamics encoded in their transition model. Through the sequence of action-observation pairs, specific clones will have a higher likelihood and can therefore disambiguate the aliased observations.

The CSCGs are learned using the expectation-maximization (EM) algorithm for HMMs (the Baum–Welch algorithm) which maximizes the evidence lower bound (ELBO)<sup>7</sup>. During the E-step, the posteriors over states are estimated through smoothing, i.e., the forward-backward algorithm. The M-step then selects the optimal parameters for the transition model, given this sequence of visited states. For the update equations, the reader is referred to ref. 7.

We consider two cognitive maps, that represent the W-maze task at two distinct levels. The first level considers the structure of the environment (i.e., where can the agent walk and where are the walls), and the second level encodes the task rule (i.e., in which order the corridors yields the reward). We learn these two maps in a sequential fashion, using two distinct CSCGs.

We first collect a sequence of data for learning the spatial structure by exploring the maze. We designed the exploration sequence to exhaustively cover the W-maze such that a path from each pose to each possible other pose is present. This sequence was used to learn the parameters of a CSCG for the first (physical) level, using the expectation-maximization mechanism described in ref. 7. We initialize the level 1 CSCG with 20 clone states per observation. After learning, the model is pruned using a Viterbi-decoding algorithm<sup>7</sup> and mapped to the POMDP of the first level of the hierarchical generative model (as will be explained in Section “Casting CSCGs as partially observable Markov decision processes”) for engaging in active inference. The learned graph is visualized in Fig. 6a. The nodes in the graph correspond to the learned nodes and their colors represent the observations encountered in the states, during the correct execution of the spatial alternation task. The gray nodes represent other states required for transitioning between states or trajectories outside the correct spatial alternation. Note that the graph shown in Fig. 2 only shows the colored nodes, but not the gray nodes.

We use the learned cognitive map of physical space to learn the cognitive map of task space. We first extract the states that are distinctly representing the end of each corridor (i.e., matching with observation 1 in Fig. 2b) and use them as the actions for the task level. When an action is selected, this means that the agent is moving toward one of these states, and thus the end of one of the corridors. The observations at this level are the presence of reward (“1” if the agent is following the rule, “0” otherwise), and the reached level 1 state (physical level). To aid the learning process, we guide the agent to follow the rule 75% of the time. We now learn a CSCG with 10 clones per observation using the same expectation-maximization scheme for the single rule scheme, and 32 clones for the scenario with three rules. After learning, the model is pruned using a Viterbi algorithm<sup>7</sup> and mapped to the POMDP of the second level of the hierarchical generative model to engage in active inference (see Section “Casting CSCGs as partially observable Markov decision processes”). For a robustness analysis on the model capacity with respect to amount of rules, rule length and amount of clones per state, the reader is referred to the Supplementary Materials. We visualize the learned cognitive graph in Fig. 6b, where the states used when pursuing the rule are



**Fig. 6 | Learned cognitive graphs.** The arrows indicate possible transitions from one state to the next. The node in the center where all nodes point to, is the added and absorbing dummy state. **a** The learned CSCG for the physical space model (level 1). The color in the graph represents distinct observations, only the states that are active when pursuing the spatial alternation rule (LCRC) are colored, the other

states are marked in gray. **b** The learned CSCG for the task space model (level 2). The color in the graph indicates the states that were active when pursuing the spatial alternation rule (LCLR). While only four states are necessary to encode the rule, it can be observed that many more states are added by learning the other paths that might have been visited during training.

colored in red. It can be observed that the spatial alternation rule is encoded in this graph, using four states, as this is the optimal amount required for the spatial alternation rule (two states for the center, one for the left, and one for the right corridor). Note that the graph shown in Fig. 2 only shows the colored nodes, but not the gray nodes.

### Hierarchical active inference

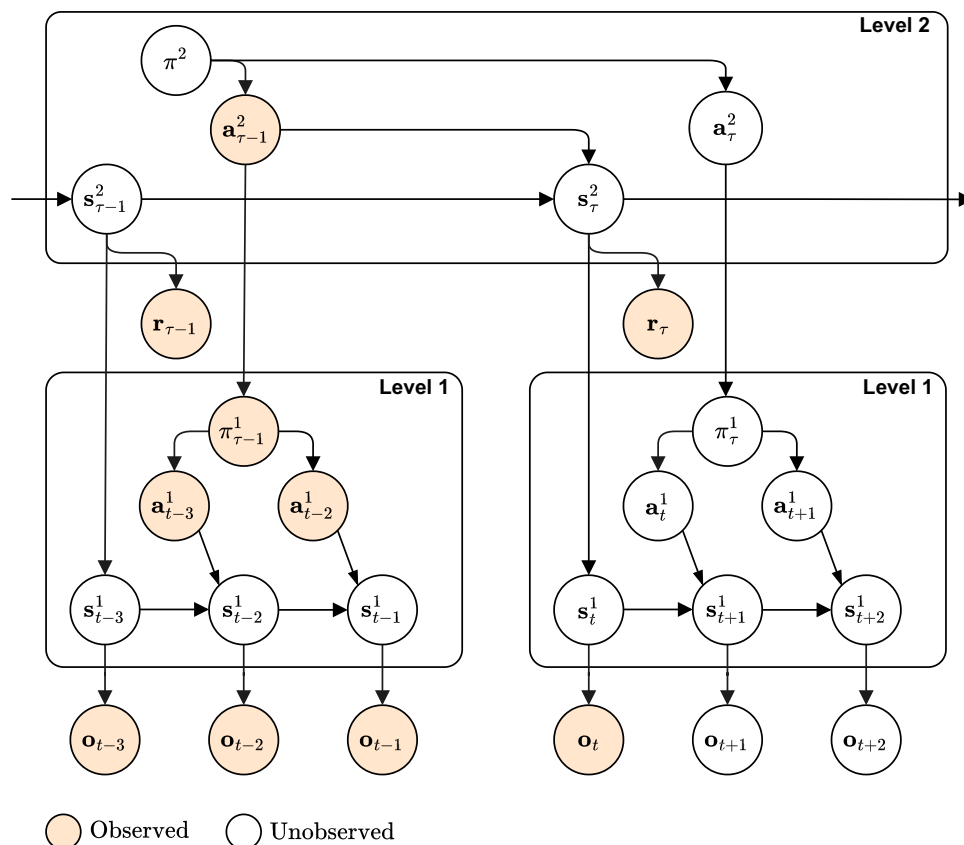
Active inference is a normative framework to describe cognitive processing and brain dynamics in living organisms<sup>36</sup>. It assumes that action and perception both minimize a common functional: the minimization of variational free energy (which is an upper bound to the organisms' surprise). Central to active inference is the idea that living organisms are equipped with generative models, capturing the causal relations between observable outcomes, the agent's actions, and hidden states. It is important to note the difference between the agent's generative model and the generative process. The former represents the internal generative model that the agent uses to attribute consequences to its actions, while the latter represents the true process from which outcomes are generated in the world. Crucially, the agent is unable to observe the hidden state of the generative process, as they are separated through a Markov blanket given the observation and action hidden variables. However, the agent is able to perform actions and observe the generated outcomes<sup>36,89,90</sup>.

**The hierarchical generative model.** We endow the active inference agent with the hierarchical generative model depicted in Fig. 7. In this section, we provide a high-level overview of the generative model, whereas in Section "Casting CSCGs as partially observable Markov decision processes" we discuss implementation details.

The hierarchical generative model illustrated in Fig. 7 is split into two distinct levels. The highest hierarchical level (level 2) reasons at a

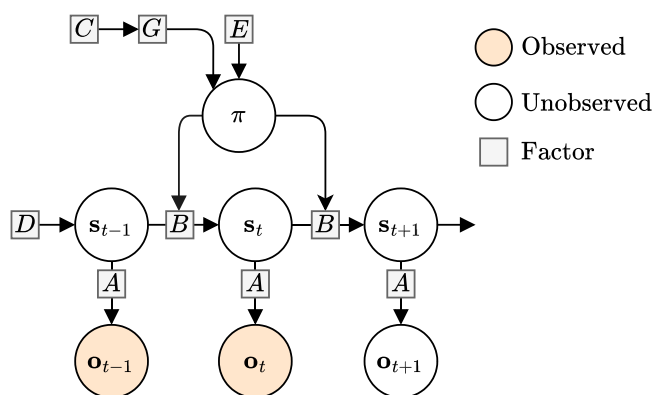
more abstract level, e.g., about which corridor to visit, while the lowest level (level 1) considers the step-by-step navigation in the environment. Each of these levels can be interpreted as an individual POMDPs<sup>39</sup>, operating at different timescales: faster for level 1 and slower for level 2. However, crucially, the two levels interact reciprocally by exchanging messages, as described below. The hierarchical model supports the hierarchical selection of abstract policies that move the agent from one goal to the other, following the task rule(s) (at level 2) and from one spatial location to the other, based on the spatial map (at level 1). Conceptually, this hierarchical selection starts when the agent's preferred observation is set to find the reward (technically, in active inference, this is done by assuming a strong prior for the reward observation) and it involves both levels.

The most fundamental goal of Level 2 is to select a policy  $\pi^2$  that moves the agent through task space, by steering abstract actions  $\mathbf{a}_t^2$  that follow the rule learned by the agent (or in the case of multiple rules, as in Experiment 2, the rule currently inferred to be in place). These actions indicate the location that the agent will then try to reach in the spatial map (level 1). In this way, the level 2-action conditions the level 1 policy  $\pi^1$ . The level 2 state  $\mathbf{s}_t^2$  is a latent representation of the corridor in which the agent is currently located in, and where it comes from (i.e., where the agent is in rule space). This abstraction is enforced by the choice of considering the reachable states to be matched to observation 1 (Fig. 2b), reflecting the fact that the agent knows the potential reward locations (which, in the animal experiments, are baited with food wells). The agent only considers reward on level 2 and thus generates the presence of reward directly from the level 2 state. At level 1, the policy  $\pi^1$  is conditioned by level 2, by setting the preferred state as reaching one of the spatial locations encoded in the cognitive graph. The policy  $\pi^1$  generates the low-level actions that navigate the agent from one pose (position and orientation, encoded in hidden



**Fig. 7 | The hierarchical generative model.** This generative model consists of two hierarchical levels, where the top level operates at a slower timescale than the lower level. The policy at the highest level sets the actions for the top level, which determines temporal transitions between states at level 2. The level 2 state then generates the presence of reward and the level 1 state. At level 1, the states generate the observations and the policy generates the actions, depending upon the selected level 2 action. Finally, when a policy is pursued, a message is passed to the upper

level, transitioning the upper level to the next state. This highlights the different timescales, whereas level 1 operates in the timescale of the agent's movement, the upper level operates at the timescale of the reached subgoals (visited corridors). The policies and actions at future timesteps are unobserved for both levels, and are inferred using the expected free energy. Orange circles denote observed variables, while white circles denote unobserved variables.



**Fig. 8 | Factor graph of a POMDP.** The conditional dependencies in Hidden Markov Models are parameterized by a set of matrices. The **A** matrix parameterizes the likelihood model, i.e., how states  $s_t$  map to observations  $o_t$ . The **B** matrix parameterizes the transition model, i.e., how state  $s_{t+1}$  changes at each timestep, dependent on policy  $\pi$ . The **C** vector denotes the preferred observation or state, and directly influences the **G** factor which conditions  $\pi$ . The policy  $\pi$  depends on the habit **E**. Orange circles denote that the variable is observed.

state  $s_t^1$  to the next. Then, the agent receives observations  $o_t$  from the environment, which reflect its current pose in the maze and which permit updating the hidden state estimate. Note that reaching the preferred spatial location requires making multiple transitions at level

1, but level 2 only makes a transition between states when the preference at level 1 is reached (or equivalently when the level 2 action is performed). This creates a separation of timescales because multiple transitions at level 1 are nested between two subsequent actions at level 2.

**Active inference.** As shown in Fig. 7, the generative model considered in this paper is a hierarchically stacked (two levels) POMDP. In order to reduce complexity, we first discuss active inference for a single layer POMDP, as depicted in Fig. 8. At time step  $t$ , we have observation  $o_t$ , action  $a_t$  generated from policy  $\pi$ , and state  $s_t$ . The generative model is factorized as:

$$P(\tilde{s}, \tilde{a}, \tilde{o}) = P(s_0) \prod_{t=1} P(o_t | s_t) P(s_t | s_{t-1}, a_{t-1}) P(a_{t-1}), \quad (1)$$

where the tilde represents a sequence over time. As computing the posterior distribution over the state is intractable for large state spaces, we resort to variational inference and introduce the approximate posterior  $Q(\tilde{s} | \tilde{o}, \tilde{a})$ , the variational free energy measures the discrepancy between the joint distribution and the approximate posterior:

$$F = \mathbb{E}_{Q(\tilde{s} | \tilde{o}, \tilde{a})} [\ln Q(\tilde{s} | \tilde{o}, \tilde{a}) - \ln P(\tilde{s}, \tilde{a}, \tilde{o})] \quad (2)$$

Active inference agents minimize this quantity through learning (by optimizing the model parameters), perception (by estimating the most



likely state), and planning (by selecting the policy or action sequence that results in the lowest expected free energy).

The expected free energy  $G$  is a quantity that is only used during planning and represents the agent's variational free energy expected after pursuing a policy  $\pi$ . It is distinct from the variational free energy, because it requires considering future observations generated from the selected policy, as opposed to the current (and past) observations.

Active inference realizes goal-directed behavior by selecting policies that minimize expected free energy—and that are expected to yield observations that are closer to preferred observations (or prior preferences). This is done by setting the approximate posterior of the policy proportional to this quantity:

$$Q(\pi) = \sigma(\gamma G(\pi) + \ln \mathbf{E}), \quad (3)$$

where  $\sigma$  is the softmax function,  $\gamma$  is a temperature variable, and  $\mathbf{E}$  is a prior distribution over actions, or habit. Actions are sampled according to this posterior distribution, where low temperatures yield more deterministic behavior. At a given timestep  $\tau$ ,  $G$  is computed for a given policy according to the following quantity as formalized in<sup>91</sup>:

$$G(\pi, \tau) = -\mathbb{E}_{Q(\mathbf{o}_\tau|\pi)} [\mathbb{D}_{\text{KL}}[Q(\mathbf{s}_\tau|\mathbf{o}_\tau, \pi) \| Q(\mathbf{s}_\tau|\pi)]] - \mathbb{E}_{Q(\mathbf{o}_\tau|\pi)} [\ln P(\mathbf{o})] \quad (4)$$

$$+ \underbrace{\mathbb{E}_{Q(\mathbf{o}_\tau|\pi)} [\mathbb{D}_{\text{KL}}[Q(\mathbf{s}_\tau|\mathbf{o}_\tau, \pi) \| P(\mathbf{s}_\tau|\mathbf{o}_\tau, \pi)]]}_{\text{Expected Approximation Error} \geq 0} \quad (5)$$

$$\geq -\underbrace{\mathbb{E}_{Q(\mathbf{o}_\tau|\pi)} [\mathbb{D}_{\text{KL}}[Q(\mathbf{s}_\tau|\mathbf{o}_\tau, \pi) \| Q(\mathbf{s}_\tau|\pi)]]}_{\text{Epistemic value}} - \underbrace{\mathbb{E}_{Q(\mathbf{o}_\tau|\pi)} [\ln P(\mathbf{o})]}_{\text{Pragmatic Value}} \quad (6)$$

This quantity is decomposed into an epistemic value and a pragmatic value<sup>33</sup>. When evaluating  $G$ , we compute the upper bound shown in Eq. (6). We thus make the assumption that the expected approximation error (Eq. (5)) is zero<sup>91</sup>. The epistemic value computes the expected information gain (Bayesian surprise) between the prior  $Q(\mathbf{s}_\tau|\pi)$  and posterior  $Q(\mathbf{s}_\tau|\mathbf{o}_\tau, \pi)$ . Intuitively, this quantity represents how much the agent expects the belief over the state to shift when pursuing this policy. The pragmatic value then measures the expected log probability of observing the preferred observation under the selected policy, intuitively computing how likely it is that this policy will drive the agent to its prior preferences. The full expected free energy for a finite time horizon  $T$  is computed as  $\sum_{\tau=1}^T G(\pi, \tau)$ .

**Hierarchical active inference.** Here, we discuss how the active inference scheme introduced above is extended to realize hierarchical perception and planning, through bottom-up and top-down message passing between the two levels of the hierarchical generative model shown in Fig. 7.

During perception, first, the low-level state  $\mathbf{s}^1$  is inferred given observations  $\mathbf{o}_\tau$  and level 1 actions  $\mathbf{a}_\tau^1$  using the inference mechanism implemented in PyMDP<sup>91</sup>. When the free energy of the lowest level reaches a pre-specified threshold, a message containing the most likely level 1 state  $\mathbf{s}_\tau^1$  the agent is in is passed to the level above (level 2). This threshold is chosen to be reached when the level 1-preference is reached, and uncertainty is below a fixed level. When level 2 gets this bottom-up message, it queries the environment for the presence of reward and observes  $\mathbf{r}_\tau$ . This conjunction of observations, together with high-level action  $\mathbf{a}^2$  is used to infer the belief over the level 2 state  $Q(\mathbf{s}^2)$ , in the same way as done for level 1.

During planning, i.e., generating a sequence of actions that (should) drive the agent toward its goal, policies are generated in a top-down manner, in the sense that goals set at level 2 determine the prior preference (and then in turn the policy) at level 1. First, the level 2 policy is selected by sampling from the approximate posterior over  $\pi^2$

proportional to its expected free energy. The action selected at this level then conditions the policy at level 1  $\pi^1$ , for which in turn the approximate posterior  $Q(\pi^1)$  is computed. The low-level action  $\mathbf{a}_\tau^1$  is then sampled according to this distribution, similar to the hierarchical generative model described by ref. 92.

As depicted in Fig. 7, the policy of the lower level sets a limited amount of steps. In particular, when the preference of the lower level is reached, a message is passed to the upper level, transitioning it into the next state. In this way, inferring the policy at the lower level only considers one action of the upper level at a time.

The posterior over the level 2 policy  $Q(\pi^2)$  is first inferred, as it has no dependencies from above. To do this, the expected free energy is computed, for which we look one time step ahead in our implementation as this could predict the next corridor in which the agent encounters reward. We consider Eq. (7), for which the observations are now a conjunction of reward  $\mathbf{r}_\tau$  and level 1-states  $\mathbf{s}_\tau^1$ :

$$G^2(\pi^2, \tau) \geq -\mathbb{E}_{Q(\mathbf{r}_\tau, \mathbf{s}_{t=\tau}^1|\pi^2)} [\mathbb{D}_{\text{KL}}[Q(\mathbf{s}_\tau^2|\mathbf{r}_\tau, \mathbf{s}_\tau^1, \pi^2) \| Q(\mathbf{s}_\tau^2|\pi^2)]] - \mathbb{E}_{Q(\mathbf{r}_\tau, \mathbf{s}_{t=\tau}^1|\pi^2)} [\ln P(\mathbf{r}_\tau, \mathbf{s}_{t=\tau}^1)], \quad (7)$$

Note that the level 1 state is synchronized with the level 2 state, and only the state observed at this synchronized time is considered, hence we denoted the time index for the level 1 state by  $t = \tau$ . In practice, we used a temperature ( $\gamma$ ) value of 0.5 for this level. For the habit policy  $\mathbf{E}$ , we use a categorical distribution over action, that is conditioned on the current state of the agent. The action  $\mathbf{a}_\tau^2$  is then sampled according to  $Q(\pi^2)$ .

The agent, however, is not able to act upon the physical world yet with this level 2 action. The agent now has to infer the posterior distribution over the policy at level 1  $Q(\pi^1)$ . Because we now want to move towards a specific state, i.e., a disambiguated observation, we set the preference in state space for computing the expected free energy at level 1. In practice, Eq. (6), for level 1 then boils down to:

$$G^1(\pi^1, \tau) \geq -\mathbb{E}_{Q(\mathbf{o}_\tau|\pi^1)} [\mathbb{D}_{\text{KL}}[Q(\mathbf{s}_\tau^1|\mathbf{o}_\tau, \pi^1) \| Q(\mathbf{s}_\tau^1|\pi^1)]] - \mathbb{E}_{Q(\mathbf{s}_\tau^1|\pi^1)} [\ln P(\mathbf{s}_\tau^1)], \quad (8)$$

where the first term yields the expected information gain after pursuing policy  $\pi^1$ , and the latter the expected utility. In other words, how close the expected state is from the preferred state, set by the action of the level above. This quantity is then used to approximate the posterior over the policy at level 1 from which action  $\mathbf{a}_\tau^1$  is sampled. At this level, we set a temperature ( $\gamma$ ) value of 0.5 and a uniform habit distribution. Crucially, in order to evaluate  $G^1$  at planning depth 1, we set the prior preference ( $\mathbf{C}$  matrix) as such that the preference for each state is proportional to the distance to the goal, i.e., intuitively this means that the agent can follow a breadcrumb trail towards the goal, given that it properly inferred the current state.

### Casting CSCGs as partially observable Markov decision processes

CSCGs can be used directly to plan<sup>7</sup>. However, in this work, we use the two learned CSCGs to create a hierarchical generative model for active inference—and then use HAI to solve the spatial alternation tasks.

Using the two CSCGs to create a hierarchical generative model for active inference requires mapping them into two POMDPs<sup>39</sup>, the mathematical framework used in discrete-time active inference.

Practically, a POMDP is described by a set of four arrays as shown in Fig. 8. We describe these arrays using the symbol notation typically adopted in discrete time active inference<sup>36</sup>. The  $\mathbf{A}$  matrix encompasses the likelihood model  $P(\mathbf{o}|\mathbf{s})$ , or how states are mapped to observations. The  $\mathbf{B}$  tensor entails the transition model  $P(\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t)$ , or how states change over time, conditioned on the selected action. The  $\mathbf{C}$  vector sets the prior over future observations or states, depending on the



implementation. This vector is used within active inference to embed the preference or goal state/observation into the agent. Finally, the **D** vector describes the initial belief over the state  $P(\mathbf{s})$ .

As mentioned earlier, clone-structured cognitive graphs are a special case of HMM and can be therefore easily mapped to POMDPs (which extend HMMs, too). First, we consider the mapping of the likelihood. This **A** matrix represents the mapping of observation to state, specifically  $\mathbf{A}_{ij} = P(\mathbf{o}_i | \mathbf{s}_j)$ . The matrix can be constructed through the definition of the CSCG, a deterministic mapping of clone state to its observation is set:  $P(\mathbf{o}_i | \mathbf{s}_j) = 1 \forall \mathbf{s}_j \in C(\mathbf{o}_i)$  and  $P(\mathbf{o}_i | \mathbf{s}_j) = 0 \forall \mathbf{s}_j \notin C(\mathbf{o}_i)$ , where  $C(\mathbf{o}_i)$  denotes all clone states for observation  $\mathbf{o}_i$ . When constructing the model we only consider states for which the marginalized probability  $P(\mathbf{s}_j)$  surpasses a threshold of 0.0001. Finally, we add an additional dummy state to the additional dummy observation, to which unlikely actions are mapped (see below).

To define the transition model, or **B** tensor, we use the learned parameters of the model. Through the learning process described in Section “Clone-structured cognitive graphs (CSCG) for learning cognitive maps of physical and task space”, the CSCG has learned the transition probability  $P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ . This is parameterized as a tensor for which  $\mathbf{B}_{ijk} = P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$ . We construct this tensor with the learned probabilities, for the same states that are considered in the likelihood matrix. Note that for ease of implementation, the active inference routines that we used require that any action could be executed from any state. However, in our task, some actions are not available or highly unlikely in some states (e.g., turning actions in a corridor). To handle this, we created a dummy state in the POMDP, which corresponds to observation 20 in Fig. 2b to which we map the “highly unlikely” actions and to which we assign a large negative value (see below) to ensure that it is never selected during planning. In practice, we set a transition probability of  $10^{-12}$  from every state to the dummy state (after which we re-normalize the tensor to sum to one for each action and state). We also set a self-transition of 1 for the dummy state, therefore ensuring that it is “absorbing”.

We endow the active inference agent with prior preferences to secure rewards in the spatial alternation tasks in the **C** matrix. For ease of implementation, we set these prior preferences manually, rather than extracting them from the learned parameters of the CSCG. For level 2, we set a preference over the conjunction of observations where the reward is 1 to encourage the agent to follow the learned rule. As described above, when the level 2 action is selected, it sets the **C** matrix of level 1 to a large preference for the preferred state and to a value proportional to the distance (in the learned state space of level 1) from said given state to the preferred state (except for the dummy state, for which the **C** vector has a fixed value. This is set to the lowest value, making the dummy state the least preferred state).

The prior over the state is parameterized by the **D** matrix. We parameterize this as a uniform distribution over all the states, except for the dummy state. This reflects the fact that in our simulations, the active inference agent is placed in the W-maze with a random pose and has to self-localize.

Finally, the habit, or the prior over action is parameterized by the **E** tensor. For level 1, we parameterize this as a uniform distribution. For level 2, we model this habit as a Dirichlet distribution, that is conditioned on the current state of the agent and is proportional to actions that were rewarding during learning.

Using these tensors, inference at each single level can be implemented as a Bayes filter, iteratively computing the posterior over state at each timestep using the following update:

$$\mathbf{q}\mathbf{s}_t = \sigma(\mathbf{A} \cdot \mathbf{o}_t + \mathbf{B}_{\pi_{t-1}} \cdot \mathbf{q}\mathbf{s}_{t-1}), \quad (9)$$

where  $\mathbf{q}\mathbf{s}_t$  denotes the parameters of the categorical distribution over state, i.e.,  $Q(\mathbf{s}_t | \mathbf{o}_t, \pi) = \text{Cat}(\mathbf{q}\mathbf{s}_t^1)$ ,  $\mathbf{o}_t$  represents the observation as a one-hot vector,  $\sigma$  denotes the softmax function, and the **B**-tensor is

sliced by the policy  $\pi_{t-1}$ <sup>33</sup>. In the first timestep,  $\mathbf{q}\mathbf{s}_0$  is initialized as the prior matrix **D**.

The expected free energy  $G$  for a considered policy  $\pi_t$  can then be evaluated using these tensors for each future timestep  $\tau$ . As we specify the preference over state for the first level, and over observation for the second level, the distinction is made explicitly. For the first level, this boils down to<sup>36</sup>:

$$G_t^1 = \text{diag}(\mathbf{A}^{1\top} \cdot \ln \mathbf{A}^1) \cdot \mathbf{q}\mathbf{s}_t^1 - \mathbf{q}\mathbf{o}_t^1 \cdot \ln \mathbf{q}\mathbf{o}_t^1 + \mathbf{q}\mathbf{s}_t^1 \cdot \ln \mathbf{C}^1, \quad (10)$$

where the superscript <sup>1</sup> denotes the first level,  $\mathbf{q}\mathbf{s}_t^1$  denotes the parameters of the categorical distribution over state, i.e.,  $Q(\mathbf{s}_t^1 | \mathbf{o}_t^1, \pi_t^1) = \text{Cat}(\mathbf{q}\mathbf{s}_t^1)$ , which is computed as  $\mathbf{q}\mathbf{s}_t^1 = \mathbf{B}_{1,\pi_t^1} \cdot \mathbf{q}\mathbf{s}_t^1$ . The parameters of the categorical over the observation is denoted by  $\mathbf{q}\mathbf{o}_t^1$ , i.e.,  $Q(\mathbf{o}_t^1) = \text{Cat}(\mathbf{q}\mathbf{o}_t^1)$ , and is computed as  $\mathbf{q}\mathbf{o}_t^1 = \mathbf{A}^1 \cdot \mathbf{q}\mathbf{s}_t^1$ . The same equation can be used for the expected free energy at the second level, except now the preference is specified as an expected observation:

$$G_t^2 = \text{diag}(\mathbf{A}^{2\top} \cdot \ln \mathbf{A}^2) \cdot \mathbf{q}\mathbf{s}_t^2 - \mathbf{q}\mathbf{o}_t^2 \cdot \ln \mathbf{q}\mathbf{o}_t^2 + \mathbf{q}\mathbf{o}_t^2 \cdot \ln \mathbf{C}^2, \quad (11)$$

where the superscript<sup>2</sup> denotes the second level.  $Q(\mathbf{s}_t^2 | \mathbf{o}_t^2, \pi_t^2) = \text{Cat}(\mathbf{q}\mathbf{s}_t^2)$ , which is computed as  $\mathbf{q}\mathbf{s}_t^2 = \mathbf{B}_{2,\pi_t^2} \cdot \mathbf{q}\mathbf{s}_t^2$ . The parameters of the categorical over the observation at level 2 are denoted by  $\mathbf{q}\mathbf{o}_t^2$ , i.e.,  $Q(\mathbf{o}_t^2) = \text{Cat}(\mathbf{q}\mathbf{o}_t^2)$ , and is computed as  $\mathbf{q}\mathbf{o}_t^2 = \mathbf{A}^2 \cdot \mathbf{q}\mathbf{s}_t^2$ . The expected free energy  $G$  for a policy  $\pi$  can then be acquired by summing over time horizon  $T$ :  $G = \sum_{t=1}^T G_t$  and the posterior over the policies is achieved through  $Q(\pi) = \sigma(\gamma \cdot \mathbf{G} + \ln \mathbf{E})$ , where **G** is a vector with the expected free energy for all considered policies.

## Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

## Code availability

We provide the code for replicating the simulations at <https://github.com/toonvdm/bridging-cognitive-maps><sup>93</sup>.

## References

- Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189 (1948).
- O’Keefe, J. & Nadel, L. Précis of O’Keefe & Nadel’s The hippocampus as a cognitive map. *Behav. Brain Sci.* **2**, 487–494 (1979).
- O’Keefe, J. & Dostrovsky, J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* **34**, 171–175 (1971).
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).
- Whittington, J. C. et al. The Tolman-Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249–1263.e23 (2020).
- Whittington, J. C. R., McCaffary, D., Bakermans, J. J. W. & Behrens, T. E. J. How to build a cognitive map. *Nat. Neurosci.* **25**, 1257–1272 (2022).
- George, D. et al. Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nat. Commun.* **12**, 2392 (2021).
- Raju, R. V., Guntupalli, J. S., Zhou, G., Lazaro-Gredilla, M. & George, D. Space is a latent sequence: a theory of the hippocampus. *Sci. Adv.* **10**, 31 (2024).
- Stachenfeld, K. L., Botvinick, M. M. & Gershman, S. J. The hippocampus as a predictive map. *Nat. Neurosci.* **20**, 1643–1653 (2017).
- Stoianov, I., Maisto, D. & Pezzulo, G. The hippocampal formation as a hierarchical generative model supporting generative replay and continual learning. *Progr. Neurobiol.* **217**, 102329 (2022).

11. Chen, Y., Zhang, H., Cameron, M. & Sejnowski, T. Predictive sequence learning in the hippocampal formation. *Neuron* **112**, 2645–2658 (2024).
12. Levenstein, D., Efremov, A., Eyono, R. H., Peyrache, A. & Richards, B. A. Sequential predictive learning is a unifying theory for hippocampal representation and replay. Preprint at <https://www.biorxiv.org/content/10.1101/2024.04.28.591528v1> (2024).
13. Recanatesi, S. et al. Predictive learning as a network mechanism for extracting low-dimensional latent space representations. *Nat. Commun.* **12**, 1417 (2021).
14. Bellmund, J. L., Gärdenfors, P., Moser, E. I. & Doeller, C. F. Navigating cognition: spatial codes for human thinking. *Science* **362**, eaat6766 (2018).
15. Epstein, R. A., Patai, E. Z., Julian, J. B. & Spiers, H. J. The cognitive map in humans: spatial navigation and beyond. *Nat. Neurosci.* **20**, 1504–1513 (2017).
16. Buzsáki, G. & Moser, E. I. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* **16**, 130–138 (2013).
17. Niv, Y. Learning task-state representations. *Nat. Neurosci.* **22**, 1544–1553 (2019).
18. Schuck, N., Cai, M., Wilson, R. & Niv, Y. Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* **91**, 1402–1412 (2016).
19. Basu, R. et al. The orbitofrontal cortex maps future navigational goals. *Nature* **599**, 449–452 (2021).
20. Ito, H. T., Zhang, S.-J., Witter, M. P., Moser, E. I. & Moser, M.-B. A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature* **522**, 50–55 (2015).
21. Pezzulo, G., Verschure, P. F., Balkenius, C. & Pennartz, C. M. The principles of goal-directed decision-making: from neural mechanisms to computation and robotics. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20130470 (2014).
22. Verschure, P. F., Pennartz, C. M. & Pezzulo, G. The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20130483 (2014).
23. Jadhav, S. P., Kemere, C., German, P. W. & Frank, L. M. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
24. Benchenane, K., Tiesinga, P. H. & Battaglia, F. P. Oscillations in the prefrontal cortex: a gateway to memory and attention. *Curr. Opin. Neurobiol.* **21**, 475–485 (2011).
25. Shin, J. D. & Jadhav, S. P. Multiple modes of hippocampal-prefrontal interactions in memory-guided behavior. *Curr. Opin. Neurobiol.* **40**, 161–169 (2016).
26. Colgin, L. L. Oscillations and hippocampal-prefrontal synchrony. *Curr. Opin. Neurobiol.* **21**, 467–474 (2011).
27. Siapas, A. G., Lubenov, E. V. & Wilson, M. A. Prefrontal phase locking to hippocampal theta oscillations. *Neuron* **46**, 141–151 (2005).
28. Khodagholy, D., Gelineas, J. N. & Buzsáki, G. Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science* **358**, 369–372 (2017).
29. Jones, M. W. & Wilson, M. A. Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol.* **3**, e402 (2005).
30. Tang, W., Shin, J. D. & Jadhav, S. P. Multiple time-scales of decision-making in the hippocampus and prefrontal cortex. *eLife* **10**, e66227 (2021).
31. Den Bakker, H., Van Dijck, M., Sun, J.-J. & Kloosterman, F. Sharp-wave ripple associated activity in the medial prefrontal cortex supports spatial rule switching. *Neuroscience* **42**, 112959 (2022).
32. Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
33. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. & Pezzulo, G. Active inference: a process theory. *Neural Comput.* **29**, 1–49 (2017).
34. Bogacz, R. A tutorial on the free-energy framework for modelling perception and learning. *J. Math. Psychol.* **76**, 198–211 (2017).
35. Buckley, C. L., Kim, C. S., McGregor, S. & Seth, A. K. The free energy principle for action and perception: a mathematical review. *J. Math. Psychol.* **81**, 55–79 (2017).
36. Parr, T., Pezzulo, G. & Friston, K. J. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (The MIT Press, 2022).
37. Smith, R., Friston, K. J. & Whyte, C. J. A step-by-step tutorial on active inference and its application to empirical data. *J. Math. Psychol.* **107**, 102632 (2022).
38. Isomura, T., Kotani, K., Jimbo, Y. & Friston, K. J. Experimental validation of the free-energy principle with in vitro neural networks. *Nat. Commun.* **14**, 4547 (2023).
39. Kaelbling, L. P., Littman, M. L. & Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99–134 (1998).
40. Chevalier-Boisvert, M., Willems, L. & Pal, S. Minimalistic gridworld environment for gymnasium. *Adv. Neural Inf. Process. Syst.* **101**, 8024–8035 (2018).
41. Schwartenbeck, P. et al. Computational mechanisms of curiosity and goal-directed exploration. *eLife* **8**, e41703 (2019).
42. Spiers, H. J. Keeping the goal in mind: prefrontal contributions to spatial navigation. *Neuropsychologia* **46**, 2106–2108 (2008).
43. Patai, E. Z. & Spiers, H. J. The versatile wayfinder: prefrontal contributions to spatial navigation. *Trends Cogn. Sci.* **25**, 520–533 (2021).
44. Simons, J. S. & Spiers, H. J. Prefrontal and medial temporal lobe interactions in long-term memory. *Nat. Rev. Neurosci.* **4**, 637–648 (2003).
45. Sun, W. et al. Learning produces a hippocampal cognitive map in the form of an orthogonalized state machine. Preprint at <https://www.biorxiv.org/content/10.1101/2023.08.03.551900v2> (2023).
46. Rens, N. et al. Evidence for entropy maximisation in human free choice behaviour. *Cognition* **232**, 105328 (2023).
47. Schmidt, B., Duin, A. A. & Redish, A. D. Disrupting the medial prefrontal cortex alters hippocampal sequences during deliberative decision making. *J. Neurophysiol.* **121**, 1981–2000 (2019).
48. McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol. Rev.* **102**, 419 (1995).
49. Lazaro-Gredilla, M., Deshpande, I., Swaminathan, S., Dave, M. & George, D. Fast exploration and learning of latent graphs with aliased observations. Preprint at <https://arxiv.org/abs/2303.07397> (2023).
50. Yamashita, Y. & Tani, J. Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS Comput. Biol.* **4**, e1000220 (2008).
51. Hinton, G. E., Osindero, S. & Teh, Y.-W. A fast learning algorithm for deep belief nets. *Neural Comput.* **18**, 1527–1554 (2006).
52. Stoianov, I. P., Pennartz, C. M., Lansink, C. S. & Pezzulo, G. Model-based spatial navigation in the hippocampus-ventral striatum circuit: A computational analysis. *PLoS Comput. Biol.* **14**, e1006316 (2018).
53. Stoianov, I., Genovesio, A. & Pezzulo, G. Prefrontal goal codes emerge as latent states in probabilistic value learning. *J. Cogn. Neurosci.* **28**, 140–157 (2016).
54. Guntupalli, J. S. et al. Graph schemas as abstractions for transfer learning, inference, and planning. Preprint at <https://arxiv.org/abs/2302.07350> (2023).
55. Swaminathan, S. et al. Schema-learning and rebinding as mechanisms of in-context learning and emergence. Preprint at <https://arxiv.org/abs/2307.01201> (2023).

56. Liu, K., Sibille, J. & Dragoi, G. Preconfigured patterns are the primary driver of offline multi-neuronal sequence replay. *Hippocampus* **29**, 275–283 (2019).
57. Farzanfar, D., Spiers, H. J., Moscovitch, M. & Rosenbaum, R. S. From cognitive maps to spatial schemas. *Nat. Rev. Neurosci.* **24**, 63–79 (2023).
58. Collins, A. G. & Frank, M. J. Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychol. Rev.* **120**, 190 (2013).
59. Gershman, S. J. & Blei, D. M. A tutorial on Bayesian nonparametric models. *J. Math. Psychol.* **56**, 1–12 (2012).
60. Wang, J. X. et al. Prefrontal cortex as a meta-reinforcement learning system. *Nat. Neurosci.* **21**, 860–868 (2018).
61. Lanillos, P. et al. Active inference in robotics and artificial agents: survey and challenges. Preprint at <https://arxiv.org/abs/2112.01871> (2021).
62. Da Costa, L., Lanillos, P., Sajid, N., Friston, K. & Khan, S. How active inference could help revolutionise robotics. *Entropy* **24**, 361 (2022).
63. Taniguchi, T. et al. World models and predictive coding for cognitive and developmental robotics: frontiers and challenges. *Adv. Robot.* **37**, 780–806 (2023).
64. Donnarumma, F., Maisto, D. & Pezzulo, G. Problem solving as probabilistic inference with subgoal: explaining human successes and pitfalls in the tower of hanoi. *PLoS Comput. Biol.* **12**, e1004864 (2016).
65. Çatal, O., Verbelen, T., Van de Maele, T., Dhoedt, B. & Safron, A. Robot navigation as hierarchical active inference. *Neural Netw.* **142**, 192–204 (2021).
66. Zakharov, A., Guo, Q. & Fountas, Z. Variational predictive routing with nested subjective timescales. Preprint at <https://arxiv.org/abs/2110.11236> (2021).
67. Van de Maele, T., Dhoedt, B., Verbelen, T. & Pezzulo, G. Integrating cognitive map learning and active inference for planning in ambiguous environments. In *International Workshop on Active Inference*, 204–217 (Springer Nature Switzerland, Cham, 2023).
68. Wilson, M. A. & McNaughton, B. L. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679 (1994).
69. Foster, D. J. Replay comes of age. *Annu. Rev. Neurosci.* **40**, 581–602 (2017).
70. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).
71. Liu, Y., Dolan, R. J., Kurth-Nelson, Z. & Behrens, T. E. Human replay spontaneously reorganizes experience. *Cell* **178**, 640–652.e14 (2019).
72. Lansink, C. S., Goltstein, P. M., Lankelma, J. V., McNaughton, B. L. & Pennartz, C. M. Hippocampus leads ventral striatum in replay of place-reward information. *PLoS Biol.* **7**, e1000173 (2009).
73. Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S. I. & Battaglia, F. P. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* **12**, 919–926 (2009).
74. Wittkuhn, L. & Schuck, N. W. Dynamics of fmri patterns reflect sub-second activation sequences and reveal replay in human visual cortex. *Nat. Commun.* **12**, 1795 (2021).
75. Liu, K., Sibille, J. & Dragoi, G. Generative predictive codes by multiplexed hippocampal neuronal tuples. *Neuron* **99**, 1329–1341 (2018).
76. Buzsáki, G. Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. *Hippocampus* **25**, 1073–1188 (2015).
77. Buzsáki, G. *The Brain from Inside Out* (Oxford University Press, 2019).
78. Gupta, A. S., Van Der Meer, M. A., Touretzky, D. S. & Redish, A. D. Hippocampal replay is not a simple function of experience. *Neuron* **65**, 695–705 (2010).
79. Nour, M. M., Liu, Y., Arumham, A., Kurth-Nelson, Z. & Dolan, R. J. Impaired neural replay of inferred relationships in schizophrenia. *Cell* **184**, 4315–4328 (2021).
80. Pezzulo, G., Van der Meer, M. A., Lansink, C. S. & Pennartz, C. M. Internally generated sequences in learning and executing goal-directed behavior. *Trends Cogn. Sci.* **18**, 647–657 (2014).
81. Pezzulo, G., Kemere, C. & Van Der Meer, M. A. Internally generated hippocampal sequences as a vantage point to probe future-oriented cognition. *Ann. N. Y. Acad. Sci.* **1396**, 144–165 (2017).
82. Shin, H., Lee, J. K., Kim, J. & Kim, J. Continual learning with deep generative replay. *Advances in neural information processing systems* 30 (2017).
83. Mattar, M. G. & Daw, N. D. Prioritized memory access explains planning and hippocampal replay. *Nat. Neurosci.* **21**, 1609–1617 (2018).
84. Pezzulo, G., Donnarumma, F., Maisto, D. & Stoianov, I. Planning at decision time and in the background during spatial navigation. *Curr. Opin. Behav. Sci.* **29**, 69–76 (2019).
85. Pezzulo, G., Zorzi, M. & Corbetta, M. The secret life of predictive brains: what's spontaneous activity for? *Trends Cogn. Sci.* **25**, 730–743 (2021).
86. Wittkuhn, L., Chien, S., Hall-McMaster, S. & Schuck, N. W. Replay in minds and machines. *Neurosci. Biobehav. Rev.* **129**, 367–388 (2021).
87. Kurth-Nelson, Z. et al. Replay and compositional computation. *Neuron* **111**, 454–469 (2023).
88. Gothoskar, N., Guntupalli, J. S., Rikhye, R. V., Lázaro-Gredilla, M. & George, D. Different clones for different contexts: hippocampal cognitive maps as higher-order graphs of a cloned hmm. Preprint at <https://www.biorxiv.org/content/10.1101/745950v1> (2019).
89. Pezzulo, G., Rigoli, F. & Friston, K. J. Hierarchical active inference: a theory of motivated control. *Trends Cogn. Sci.* **22**, 294–306 (2018).
90. Pezzulo, G., Rigoli, F. & Friston, K. Active inference, homeostatic regulation and adaptive behavioural control. *Prog. Neurobiol.* **134**, 17–35 (2015).
91. Heins, C. et al. pymdp: A Python library for active inference in discrete state spaces. Preprint at <https://arxiv.org/abs/2201.03904> (2022).
92. Friston, K. J., Rosch, R., Parr, T., Price, C. & Bowman, H. Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* **77**, 388–402 (2017).
93. Van de Maele, T. & Verbelen, T. toonvdm/bridging-cognitive-maps: v1.0.0 (2024). <https://zenodo.org/doi/10.5281/zenodo.13769225>.

## Acknowledgements

G.P. received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreements No. 945539 (Human Brain Project SGA3) and No. 952215 (TAILOR); the European Research Council under the Grant Agreement No. 820213 (ThinkAhead), the Italian National Recovery and Resilience Plan (NRRP), M4C2, funded by the European Union - Next-GenerationEU (Project IR0000011, CUP B51E22000150006, "EBRAINS-Italy"; Project PE0000013, CUP B53C22003630006, "FAIR"; Project PE0000006, CUP J33C22002970002 "MNESYS"), and the PRIN PNRR P20224FESY. The GEFORCE Quadro RTX6000 and Titan GPU cards used for this research were donated by the NVIDIA Corporation. T.VdM was supported by a grant for a research stay abroad by the Flanders Research Foundation (FWO), and from the Flemish Government (AI Research Program). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

## Author contributions

The study was conceptualized by T.VdM., T.V., and G.P., with contributions from B.D. T.VdM. wrote the software for the simulations, conducted the experiments, did the formal analysis, and created the visualizations.

T.VdM., T.V., and G.P. designed the experiment methodology. T.V. and G.P. guided the study. T.VdM. and G.P. wrote the manuscript, T.V. and B.D. reviewed and commented on the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at

<https://doi.org/10.1038/s41467-024-54257-3>.

**Correspondence** and requests for materials should be addressed to Giovanni Pezzulo.

**Peer review information** *Nature Communications* thanks Sho Aoki, and the other, anonymous, reviewers for their contribution to the peer review of this work. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024