Gateway Identity and Spatial Remapping in a Combined Grid and Place Cell Attractor.

Tristan Baumann[1] and Hanspeter A. Mallot[1*]

[1]Cognitive Neuroscience, Department of Biology, University of Tübingen, Auf der Morgenstelle 28, 72076 Tübingen, Germany
[*]hanspeter.mallot@uni-tuebingen.de

May 11, 2022

Abstract

The spatial specificities of hippocampal place cells, i.e., their firing fields, are subject to change if the rat enters a new compartment in the experimental maze. This effect is known as remapping. It cannot be explained from path integration (grid cell activity) and local sensory cues alone but requires additional knowledge of the different compartments, other types of "context", or the gateways that lead into each compartment. Here we present a model for the hippocampal-entorhinal interplay in which the activity of place and grid cells follows a joint attractor dynamic. Place cells depend on the current grid cell activity but can also reset the grid cell activity in the remapping process. Remapping is triggered by the recognition of a gateway, for example by visual cues (not modeled in this paper). When this happens, the previously stored pattern of place cell activity for this gateway is reactivated from a "gateway database". The joint attractor will then reinstate the grid cell pattern that was active when the gateway had first been learned and path integration can proceed from there. The model is tested with various mazes used in the experimental literature and reproduces the published results. We also make testable predictions for remapping in a new maze type. We hypothesize the existence of "gate cells" that drive the place cells and with them the joint hippocampal-entorhinal loop into the corresponding attractor whenever a gate is detected.

1 Introduction

The neural substrate of spatial representation, the *cognitive map*, is generally thought to be found in place cells located in the hippocampus (HC). First described by O'Keefe and Dostrovsky (1971), these cells fire maximally whenever the animal is within a small localized region in space, the cell's place field (O'Keefe and Nadel, 1978; Moser et al., 2017). The environment is covered by the overlapping place fields of different place cells and the population activity can be used to decode the animal's position and even future trajectories once the firing fields have been mapped (Wilson and McNaughton, 1993; Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015; Pfeiffer, 2020).

In the five decades since the discovery of place cells, many other cell types supporting spatial representation have been found, including (but not limited to) head direction cells (Taube et al., 1990a,b), border cells (Solstad et al., 2008), and grid cells (Hafting et al., 2005).

Especially grid cells have received much attention: Located in the medial entorhinal cortex (MEC) (Hafting et al., 2005) and pre- and parasubiculum (Boccara et al., 2010), these cells fire at regularly spaced intervals, tiling the environment into a regular hexagonal lattice. Grid cells are dorsoventrally organized into discrete modules of constant scale and orientation; between modules, scale is increasing in discrete steps (Hafting et al., 2005; Stensola et al., 2012). Within a module, grid phase and orientation between neighboring cells remains fixed (Fyhn et al., 2007). Due to their regular firing properties and the fact that the pattern persists in complete darkness, grid cells are generally thought to be involved in path integration (Hafting et al., 2005; Fyhn et al., 2007; Rowland et al., 2016; Grieves and Jeffery, 2017).

Grid cells are the most common spatially modulated cell in the MEC (Hafting et al., 2005) and form a major input to place cells (Brun et al., 2008; Jeffery, 2011; Rowland et al., 2016). However, silencing input from MEC to CA1 does not eliminate place cell activity (Brun et al., 2008; Bush et al., 2014; Brandon et al., 2014) and in rat pups, place cells form before grid cells (Langston et al., 2010; Wills et al., 2012; Bush et al., 2014). Place cells show relatively stable place fields before the hexagonal grid cell pattern is fully developed, suggesting a strong input from allothetic or external cues such as landmarks or boundaries (Bjerknes et al., 2014; Muessig et al., 2015; Moser et al., 2015; Grieves et al., 2018).

This external input causes context-dependent firing changes: When the animal moves or is moved from one compartment to another, place cell specificities may undergo complete changes to an independent pattern of firing fields. This process is known as "remapping" (Muller and Kubie, 1987; Lever et al., 2002; Leutgeb et al., 2005; Colgin et al., 2008; Julian et al., 2018): Place cells randomly rearrange their firing fields to new, uncorrelated locations, form additional firing fields or cease firing altogether. When the animal later returns or is returned to the original compartment, the place cells remap to their original pattern (Lever et al., 2002; Leutgeb et al., 2005). A given place cell therefore represents different places at different times or, more correctly, in different "spatial contexts". Remapping also occurs in grid cells: When the context changes, the firing field locations of different modules may shift but grid scale and relative orientation remain unchanged (Fyhn et al., 2007). Due to direct anatomical projections from MEC to the hippocampus and the finding that electrical stimulation of MEC grid cells causes place cell remapping (Kanter et al., 2017), grid cell realignment has previously been suggested as a possible candidate for triggering place cell remapping (Monaco and Abbott, 2011; Bush et al., 2014).

An important puzzle piece in understanding the remapping process is the nature of the "context change" by which it is triggered. Remapping can be observed when the animal moves to different compartments within a maze, or when the color and location of cues (Bostock et al., 1991) or the overall shape and texture of the compartment is changed (Lever et al., 2002; Leutgeb et al., 2005; Wills et al., 2005; Colgin et al., 2010; Jeffery, 2011). Non-spatial contexts may lead to a modulation of place cell firing rates ("rate remapping" (Latuske et al., 2018)), as has been shown for changes of the current task or goal (Allen et al., 2012) or in the dependence on recent history (Keinath et al., 2020). For a review of remapping types, see Latuske et al. (2018). Importantly, remapping seemingly does not occur if the context remains unchanged: For example, grid cell patterns are continuous within compartments in Fyhn et al. (2007); Derdikman et al. (2009) and Carpenter et al. (2015). If remapping did occur, discontinuities in the patterns would be apparent.

This suggests that the current compartment is equivalent to the current context, for as long as no remapping occurs, the rat must be within the same area. Therefore, remapping suggests a regionalization of the rat's cognitive map, although the rat may not have direct access to that information. In this sense, the pattern of active cells provides both information about the current context and the position of the animal within that context (Latuske et al., 2018).

Overall, the findings suggest that remapping is an active process that reflects the recruitment of a unique cognitive map from a "cognitive atlas" (Julian et al., 2018), a putative higher-level hierarchical structure of spatial representation.

How can the hippocampal-entorhinal circuit acquire, maintain and reactivate a multitude of unique, context-dependent maps in a noisy infinite-context world? Some hints to this question may be hidden in a number of particularities of the remapping process:

1. Similar contexts lead to similar representations.
   In visually identical but otherwise separate compartments, place and grid cell firing patterns have been observed to repeat (Derdikman et al., 2009; Spiers et al., 2015; Carpenter et al., 2015; Grieves et al., 2016; Harland et al., 2017). That is, the place and grid cell activity at corresponding positions in different rooms is highly correlated. This similarity is behaviorally relevant: In a spatial memory task by Grieves et al. (2016), rats had to memorize odor-coded buried rewards in each of four identical rooms connected by a lateral corridor. The rooms were either arranged in parallel or radially around a curved corridor (Fig. 1). The animals learned the task faster and performed significantly better in the distinguishable radial condition, where place cells remapped to different patterns in each room, but failed in the parallel condition, where the patterns repeated (Grieves et al., 2016).

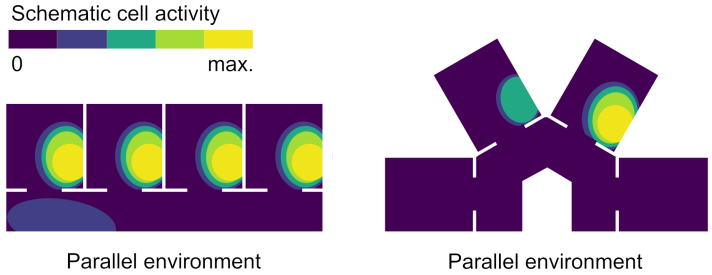

Figure 1: Schematic example of findings by Grieves et al. (2016). The place fields of the same place cell are shown in two environments. Firing repeats at corresponding positions in each room in the parallel environment, but not when the rooms are arranged radially.

2. Entrance orientation is sufficient to distinguish contexts.
   The superior performance in the radial condition is attributed to the head direction system which is assumed to maintain a global reference while the animal explores the different compartments (Harland et al., 2017). If the orientation of the compartments is the same and they are entered from the same geocentric direction, the reference is also the same and the rooms are confused. On the other hand, if the compartments extend into different directions, such as in the radial arrangement in Fig. 1, the head direction reference may be used to detect this difference and the compartments can be distinguished. If the head direction system is silenced, this ability disappears and the radially arranged rooms also show the same local pattern (Harland et al., 2017). In agreement, the inverse case, where locally and globally identical compartments show different cell patterns, has also been reported (Tanila, 1999; Grieves et al., 2016; Harland et al., 2017; Keinath et al., 2020). Here, the distinguishing cues are entry direction or entrance position relative to the compartment, i.e., also orientation references.

3. Local cues only play a limited role.
   Overall, the remapping dynamics suggest a strong bias towards local cues, i.e., the sensory

and geometrical information of the current surroundings. These findings led to the formulation of border or boundary vector cell models of place cell firing (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018): These cells fire whenever the animal is located at a preferred distance and angle from a nearby wall and thus reflect the local geometry. Based on the repeating patterns in identical compartments, the models suggest that geometric cues measured by boundary vector cells and mediated by the head direction system are the main determinants for place fields. However, boundary vector cell show no remapping. The cells have been reported to continue firing at the same fixed distance and allocentric direction from boundaries across multiple environments, regardless of context (Lever et al., 2009). Therefore, remapping dynamics must occur at least partially independent of boundary vector cells so that different patterns may arise in otherwise identical compartments.

4. Place and grid cells remap together.
   The exact moment of remapping is hard to pinpoint since firing maps are measured during extended running time of the animal. Still, both cell types seem to remap concurrently (Fyhn et al., 2007) and at the maxima of contextual difference, such as gateways between compartments or region transitions (Tanila, 1999; Leutgeb et al., 2005; Carpenter et al., 2015; Grieves et al., 2016). The firing patterns do not remap within a compartment, where the local context remains the same.

Overall, the findings imply that as long as compartments are distinguishable at their entrances, the context is different, and distinct place and grid cell patterns may arise. If, on the other hand, the compartments cannot be distinguished at the entrance, the context is the same and the cell patterns reflect that similarity. In this sense, the context of the current environment would be based on the most recent gateway rather than just local cues.

Here, we propose a combined grid-and-place cell attractor model that is able to capture these remapping dynamics. The attractor is based on the observed connectivity loop of direct projections from MEC to HC (Brun et al., 2008; Jeffery, 2011; Monaco and Abbott, 2011; Rowland et al., 2016) and recurrent connections back from HC to MEC (Hafting et al., 2005; Boccara et al., 2010; Bush et al., 2014; Rennó-Costa and Tort, 2017): As in other models (de Almeida et al., 2009; Monaco and Abbott, 2011; Lyttle et al., 2013; Li et al., 2020), the activity of multiple grid cell modules is summed to form realistic place fields; in turn, place cells also connect to grid cells and may thus support the current attractor or drive the system into a new one.

Such bifurcation events happen at region transitions and exploit the ability of place cells to remap to arbitrary positions when directly stimulated (Diamantaki et al., 2018). Recurrent input from the place cells to grid cells then causes a shift in the grid cell firing fields until a new attractor is reached. To obtain repeating place and grid cell patterns in similar environments, the model has access to a "gateway database" that stores place cell patterns at gateways between rooms. In the database, the patterns are associated with a "gateway identifier", used to index local contextual cues at the gateway. If two gateways in different regions have the same local context, they share the same gateway identifier.

New gateway identifiers are added when the local position information provided by visual or other sensory cues abruptly changes to an unknown pattern. This is generally the case when the agent walks through a narrow aperture into a novel room. Whenever a known identifier is encountered again later, the associated pattern is reactivated and the place cell activity changes to it.

For simplicity, the place cell activity within a region is assumed to be based on path integration only. As such, the patterns are directly determined by the context at the last gateway

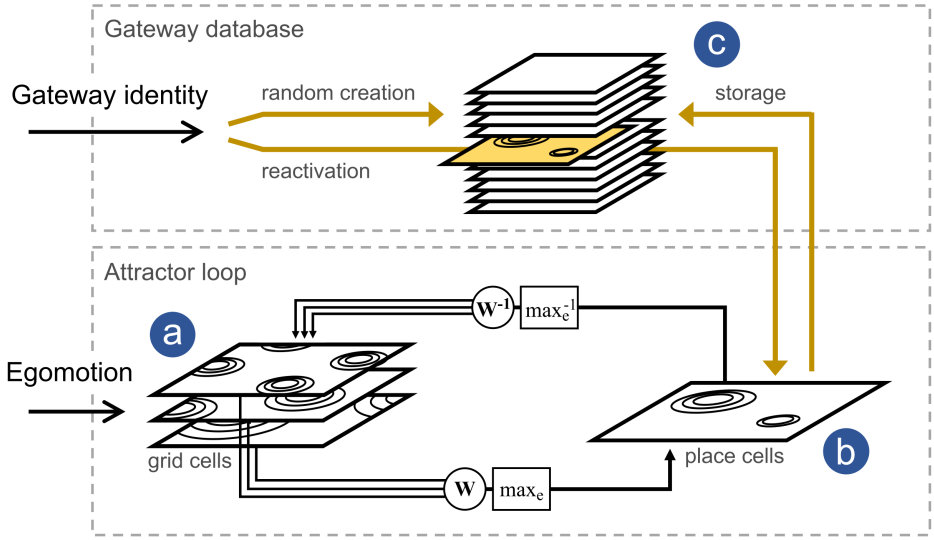

Figure 2: Model overview. **(a)** 3×100 grid cells are modeled by moving a stable activity peak over a toroidal surface. The cells form three layers with different input gains, resulting in hexagonal patterns with different scale and orientation. The circles on the layers sketch the firing fields, not the attractor peaks (which are the same for every layer). **(b)** Grid cell activity is combined in a weighted sum and filtered by an (invertible) winner-take-all scheme to form the place cell activity. The place cells in turn activate the grid cell module with the inverse operation, completing the larger attractor loop. **(c)** At region transitions such as room entrances, place cell patterns are memorized so that they may be recalled in the future (yellow arrows). When the virtual rat enters another region, place cell activity is either randomized or a corresponding place cell pattern is reactivated, followed by a shift of the grid cell firing peaks to the best matching attractor.

and the experimentally observed repetition of firing fields in similar multicompartment environments becomes a necessary consequence.

Our model is related to models of working memory by Bouchacourt and Buschman (2019), where representations are maintained through recurring connections between a structured sensory layer (the grid cells) and a random, unstructured layer (the place cells). As a model of working memory, the model is inherently limited in its capacity once representations in the unstructured layer start to overlap. In this sense, the gateway database is a long-term memory that binds working memory patterns to context and is able to reinstantiate these patterns in a top-down manner.

In the following, we show that the model is able to replicate the remapping dynamics observed in multicompartment environments (e.g., Tanila, 1999; Carpenter et al., 2015; Grieves et al., 2016) and we propose a novel multicompartment setup (L-shaped rooms, Virtual environments) to confirm the role of region transitions as the main determinants for local context.

2 Materials and Methods

2.1 Model overview

The model consists of an attractor neural network with fixed weights, simulating the cognitive map of a rat exploring a series of virtual environments. The cell activity is iteratively simulated in real-time; we denote the discrete time step of a single computational sequence as t.

The core of the model is a grid cell attractor based on (Guanella et al., 2007). Input arrives in the form of an egomotion velocity vector, which shifts grid cell activity peaks to form the well-known hexagonal firing fields. Grid cells are arranged in multiple layers that differ in orientation and gain; from these, place cell activity is obtained by summing over the grid cells in a winner-take-all scheme. We adopted this modeling approach to place cell firing fields from previous studies (de Almeida et al., 2009; Monaco and Abbott, 2011; Lyttle et al., 2013; Li et al., 2020) in order to keep the system simple. However, local position information can easily be added as another input to the place cell layer. The place cells then feed back into the grid cell module to complete a larger attractor loop (Fig. 2).

To explore remapping dynamics, an external remapping signal is triggered whenever the virtual rat enters another compartment. This signal may be generated from abrupt changes of visual or other sensory input as they would occur when crossing a narrow gateway. Initially, place cell activity is randomized and the system is forced into the nearest attractor. By associating region transitions with specific place cell patterns, the same patterns may be reactivated upon later revisits and the environments become regionalized: Each region in the environment is associated with a specific place cell pattern, and crucially, this pattern remains active for as long as the rat does not move to another region.

The head direction system is not modeled, but simply assumed to be perfect, meaning the virtual rat is always aware of "true north". It is therefore always able to distinguish rooms with different relative orientations. Throughout the paper, we use symbolic cardinal directions, i.e., north, east, south and west means up, right, down and left in the figures.

2.2 Grid Cells

Egomotion is processed by the grid cell module, consisting of 300 grid cells divided into three independent granularity layers with increasing scale and varying orientation. Three layers are required for realistic remapping based on independent shifts, as described in (Monaco and Abbott, 2011). In accordance with measurements by (Stensola et al., 2012), the grid scale, i.e., the shortest distance between two activity peaks, was set to 40cm for the most fine-grained granularity layer and increased by a factor of 1.42 per layer to 57cm and 80cm for the coarser layers.

In the rodent brain, grid cells in the same granularity layer have been reported to share the same orientation (Stensola et al., 2012), are aligned to nearby walls, and the orientation remains constant across different contexts (Krupic et al., 2015). In our model, grid orientation was fixed at 8.8° for layers 1 and 2 and 98.8° for layer 3, in accordance with the measurements by Krupic et al. (2015).

Within a layer, a stable activity peak is maintained by attractor dynamics (Guanella et al., 2007). $n^2 = 100$ cells form a topologically organized periodic sheet where neighboring cells have neighboring firing fields (Fig. 3a). The cells are connected to all others by short-range activation and long-range inhibition. In order to achieve the hexagonal geometry, the cells are arranged in a rhomboid with angles 60° and 120°, i.e., composed of two equilateral triangles. The rhomboid is the primitive cell of the hexagonal grid. The Cartesian coordinates $\mathbf{x}_{ij}$ of grid

cell ij within a rhomboid with lattice constant $1/n$ are

$$\mathbf{x}_{ij} = \frac{1}{2n}\begin{pmatrix}(i+j)\\ \sqrt{3}(j-i)\end{pmatrix} = \frac{1}{2n}\begin{pmatrix}1 & 1\\ -\sqrt{3} & \sqrt{3}\end{pmatrix}\begin{pmatrix}i\\ j\end{pmatrix} =: \mathbf{H}\begin{pmatrix}i\\ j\end{pmatrix}. \tag{1}$$

A 5×5 cell example is depicted in Fig. 3a. To derive an expression for the hexagonal periodic distance function, we pad a set of eight shifted rhomboids around the standard rhomboid by considering the shifts $\mathbf{s}_{pq} = \mathbf{H}\ (p,q)^\top$ for $(p,q) \in \{-n, 0, n\}^2$. For given Cartesian points $\mathbf{x}, \mathbf{x}'$ the periodic distance function is

$$d(\mathbf{x}, \mathbf{x}') = \min_{(p,q)\in\{-n,0,n\}^2} ||\mathbf{x} - \mathbf{x}' + \mathbf{s}_{pq}||, \tag{2}$$

where $||\cdot||$ is the Euclidean norm. This description of periodic boundary conditions in a rhomboid is equivalent to the twisted torus approach of Guanella et al. (2007), but can easily be generalized to an arbitrary number of grid cells.

We can now describe the activity of grid cell ij at time step t,

$$g_{ij}(t) = \sum_{k=1}^{n}\sum_{l=1}^{n} g_{kl}(t-1)\ w_{kl}^{ij}, \tag{3}$$

where the weights w_{kl}^{ij} follow a Gaussian function of the periodic distance defined above,

$$w_{kl}^{ij} = I \exp\left(-\frac{d(\mathbf{x}_{ij}, \mathbf{x}_{kl} + \alpha \mathbf{R}_\beta \hat{\mathbf{v}}(t))^2}{\sigma^2}\right) - T. \tag{4}$$

The intensity $I = 0.3$ defines the overall synaptic strength, $\sigma = 0.24$ is the width of the Gaussian and $T = 0.05$ is a global inhibition parameter. The scale of the grid is defined by the gain parameter $\alpha \in \{0.05, 0.035, 0.025\}$ corresponding to 40cm, 57cm and 80cm for the three layers. The orientation is given by the rotation matrix $\mathbf{R}_\beta$ with $\beta \in \{8.8°, 8.8°, 98.8°\}$, respectively. Initially, the cell activity $g_{ij}(t_0)$ is set to a random value drawn from a uniform distribution $\mathcal{U}(0, 0.1)$, but a stable peak is formed in the next time steps (initial attractor formation, Fig. 3b).

If the virtual rat moves with velocity $\mathbf{v}(t) = (v_1, v_2)^\top$, synaptic weights in direction of the egomotion estimate $\hat{\mathbf{v}}(t)$ increase, causing the activity peak on the grid cell layer to move by a corresponding amount. The egomotion estimate is obtained from the true velocity by adding multiplicative noise to both components,

$$\hat{\mathbf{v}}(t) = ((1+\epsilon_1)v_1, (1+\epsilon_2)v_2)^\top,\ \epsilon_1, \epsilon_2 \sim \mathcal{U}(-0.25, 0.25). \tag{5}$$

In the sequel, rather than describing the dynamics of individual cells, we combine the activities of all three layers into a single 1×300 grid cell activity matrix

$$\mathbf{G}(t) = [g_1(t), ..., g_{300}(t)]\,. \tag{6}$$

For example, the grid cell activity g_{221} now refers to g_{21} in granularity layer 3. This allows us to describe cell and layer-layer interactions in the form of matrix multiplications.

2.3 Place Cells

Place cells form a single layer of 300 cells. In most situations, i.e., as long as no remapping occurs, the place cell activity $\mathbf{P}(t)$ is obtained by summing over grid cell activity with random weights:

$$\mathbf{P}(t) = \max_e \left(\mathbf{P}(t)\ \mathbf{W}_g^p\right), \tag{7}$$

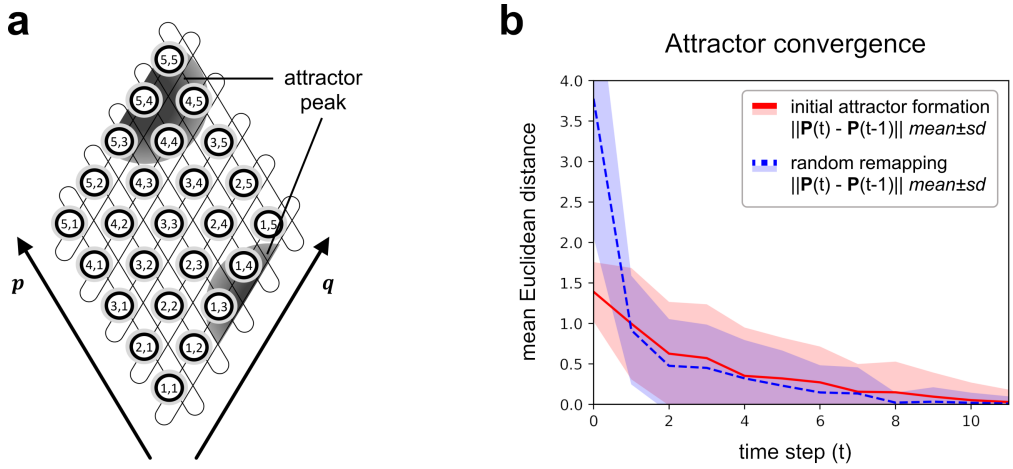

Figure 3: **(a)** Grid cell rhomboid (with only 5×5 cells for clarity). An activity peak (dark gray circle) is centered on cell $g_{5,4}$. The thin lines show periodic boundary conditions along grid cell indexes i, j. The vectors p, q correspond to the side length of the rhomboid and give the shift distance s_{pq} for eq. 2. **(b)** Convergence to an attractor when the model is initialized with zero place cell activity and random grid cell activity (red), or when the model remaps to a random place cell pattern (blue). Average of 100 repetitions.

where $\mathbf{W}_g^p$ is a 300×300 synaptic weight matrix drawn from $\mathcal{U}(-1, 1)$. The weight matrix describes the excitatory and inhibitory input from grid to place cells. As in previous studies describing place fields as originating from random summation of grid cell inputs (de Almeida et al., 2009; Monaco and Abbott, 2011; Lyttle et al., 2013), a winner-take-all nonlinearity $\max_e$ was used to sparsify place cell activity:

$$y = \max_e(x) = \frac{1}{x_{\max}} \begin{cases} x & \text{if } x > (1-e)x_{\max}, \\ 0.01\ x & \text{otherwise.} \end{cases} \tag{8}$$

$x_{\max}$ is the highest x in the input set and $e = 0.1$ is the percentage of most active cells whose activity is maintained, while all others are attenuated by a factor 0.01.

The winner-take-all scheme restricts cell activity to local maxima, resulting in place fields qualitatively similar to real place cell recordings. Importantly, by only attenuating cells below threshold rather than setting their activity to zero, the entire operation becomes invertible, i.e., it enables place cells to predict grid cells in turn. This property is required for the combined grid-and-place cell attractor loop and remapping dynamics described in the following.

2.4 Combined attractor and inverse place cell - grid cell activation

By linking the place cell activity output back to the grid cells, an attractor is formed. It ensures that a specific place code will always match a specific grid cell arrangement. However, during normal exploration, the place cell input must not hinder the movement of the grid cell peak by the velocity vector. That is, the estimated activity of grid cells from place cell input, $\hat{\mathbf{G}}(t)$ should not differ substantially from the grid cell activity $\mathbf{G}(t)$. The required feedback connections from the place cell layer to the grid cell layers have been shown to exist in the form of recurrent MEC-HC connections (Hafting et al., 2005; Boccara et al., 2010; Bush et al.,

2014; Rennó-Costa and Tort, 2017) and could be learned by standard neural network techniques (Zhang et al., 1997; Linsker, 1997). For the purpose of this study we implemented a straightforward inversion,

$$\hat{\mathbf{G}}(t) = \max_e^{-1}\left(\mathbf{P}(t)\right)\ \left(\mathbf{W}_g^p\right)^{-1}, \tag{9}$$

where $(\mathbf{W}_g^p)^{-1}$ is the inverse or Moore-Penrose-pseudoinverse of $\mathbf{W}_g^p$ from Eq. 7. In our simulations, the random matrices came out with full rank in about 99% of the trials. The nonlinearity $\max_e^{-1}(x)$ is the inverse operation of eq. 8,

$$x = \max_e^{-1}(y) = y_{\max} \begin{cases} y & \text{if } y > 1 - e, \\ 100\ y & \text{otherwise.} \end{cases} \tag{10}$$

From eqs. 7 and 9, it is easy to see that $\hat{\mathbf{G}}(t) \approx \mathbf{G}(t)$ except for a scaling factor. We can therefore stop the intrinsic dynamic of the grid cell module and restart it with the image of the place cell activity $\mathbf{G}(t) := \hat{\mathbf{G}}(t)$. This will lead to a normal continuation of the grid cell dynamics, as long as $\hat{\mathbf{G}}(t)$ falls in the basin of attraction of the current grid cell attractor. The intrinsic dynamics of the grid cell module will take care of remaining errors in the inversion process in the next few time steps. The convergence to the attractor after remapping is shown in Fig. 3b.

2.5 Remapping and gateway database

Remapping occurs whenever the virtual rat moves from one compartment to another. At the transition, the model is provided a "gateway identifier" Γ, an integer corresponding to the local context of the entered compartment. In the virtual environments described below, identifiers were predefined based on visual similarity: if two places look the same and are oriented in the same direction, they are considered to be identical (Fig. 4). In a biological implementation, this information could for example be provided by visual or geometrical information and references to local cues and landmarks (Trullier et al., 1997; Cheng et al., 2013). Gate information could be modeled by an additional "gate cell" (see Discussion) input to the place-cell layer; the hypothetical gate cell would be a cell that fires whenever the rat crosses a specific gate.

In the current implementation of the model, the newly encountered gateway identifier Γ is used to index the gateway database $D = \mathbf{P}_0, ..., \mathbf{P}_n$. If the database has an entry $\mathbf{P}_\Gamma \in D$, the stored place cell pattern is reactivated by setting $\mathbf{P}(t) := \mathbf{P}_\Gamma$. With the attractor rules described in eq. 9, the grid cell attractor will shift to the attractor matching the reactivated place cell pattern. As a side effect, accumulated path integrator noise is also reset when a known gateway is passed.

This mechanism may generate errors at wide gateways that can be passed with different sideways offsets. This problem could be avoided by assuming that the rat shows strict centering behavior at gateways. Here, we measure the sideways offset and add it to the displacement vector $\hat{\mathbf{v}}$ in eq. 4. In biology, offset information could for example be provided by border or boundary vector cells, or depth perception.

Resetting the place cell activity and the subsequent convergence of the network to the new grid-and-place cell attractor concludes the remapping process. Grid and place cells have completely changed their firing fields, but the change is consistent in that the same pattern will be activated whenever a given gateway is encountered. Therefore, similar compartments that share the same gateway identifiers will show the same place and grid cell patterns.

The gateway database models the rat's overall knowledge of the compartments; it is a cognitive atlas in the sense of Julian et al. (2018) that encompasses all context-specific maps. Initially, the gateway database is empty and the model is initialized with random activity.

When the rat learns a novel gateway identifier, the current place cell activity pattern is stored. This may happen in three different cases:

1. When entering an unknown room, a novel view will appear. A new gateway identifier is encountered and a random pattern of place cell activity is generated and stored in the database together with the identifier.

2. When entering a known room, the gateway identifier is recognized and the associated place cell activity pattern is loaded into the place cell layer.

3. New gateway identifiers can also be stored when leaving a room such that the room can be recognized when it is later reentered through the same gate. In order to model direction-specific firing fields observed in linear track environments (McNaughton et al., 1983; Derdikman et al., 2009), this feature of the model is applied only in wide rooms in which the rat can turn around, but not in narrow corridors.

2.6 Simulation and exploration procedures

For fast parallel computing of the cell matrices, the model was implemented in the machine learning interface TensorFlow (Abadi et al., 2015) running in real-time on a normal notebook PC (CPU: Intel i7-10750H, GPU: NVIDIA GeForce RTX 2070). The model was evaluated in three repetitive multicompartment environments, realized as virtual 3D models in Unity (Unity Technologies, 2021). The environments are described in detail below.

An agent, the "virtual rat", represented by a 4.8 cm diameter cylinder, explored the environments either by random walking (parallel rooms and L-shaped rooms) or guided by the experimenter (Hairpin maze). The virtual rat moved at a constant speed of 14 cm/s. For the random walk, at each time step t, the movement direction was rotated by a random angle drawn from $\mathcal{U}(-10°, 10°)$. To avoid exploration deadlocks such as continued movement into a corner, the random walk trajectory reflected off of walls. The hairpin maze environment was explored by hand to obtain complete traversals which are unlikely to result from random walking. Position data and cell activities were sampled once per time step, which had a duration of approximately 0.1 s.

2.7 Virtual environments

Parallel rooms (Fig. 4a). The layout is based on the two- or four-room mazes by (Skaggs and McNaughton, 1998; Fuhs et al., 2005; Spiers et al., 2015; Grieves et al., 2016), where identical rooms are arranged in parallel and are either entered from the same direction, or from different directions (Tanila, 1999; Grieves et al., 2016). We combined both layouts in a single two-by-two rooms environment where two sets of parallel rooms are placed on the opposite sides of a central corridor. The rooms are locally indistinguishable except for the position and direction of the entrance. In the figure, the two rooms entered from the north have gateway identifier 1 and the rooms entered from the south gateway identifier 6. The four identifiers of the corridor are all distinguishable by the combination of view and heading.

Hairpin maze (Fig. 4b). *Hairpin maze.* We created a virtual hairpin maze similar to Derdikman et al. (2009). In hairpin mazes, cell activity has been observed to repeat every second arm and to depend on running direction (Derdikman et al., 2009). Accordingly, we defined the turns between arms as region transitions, with every second arm sharing the same gateway identifier. As described above, place cell patterns in corridors are only memorized when the corridor is entered, and not when it is left. That is, entrances and exits depend on the traversal direction, which will lead to direction-selective firing fields.

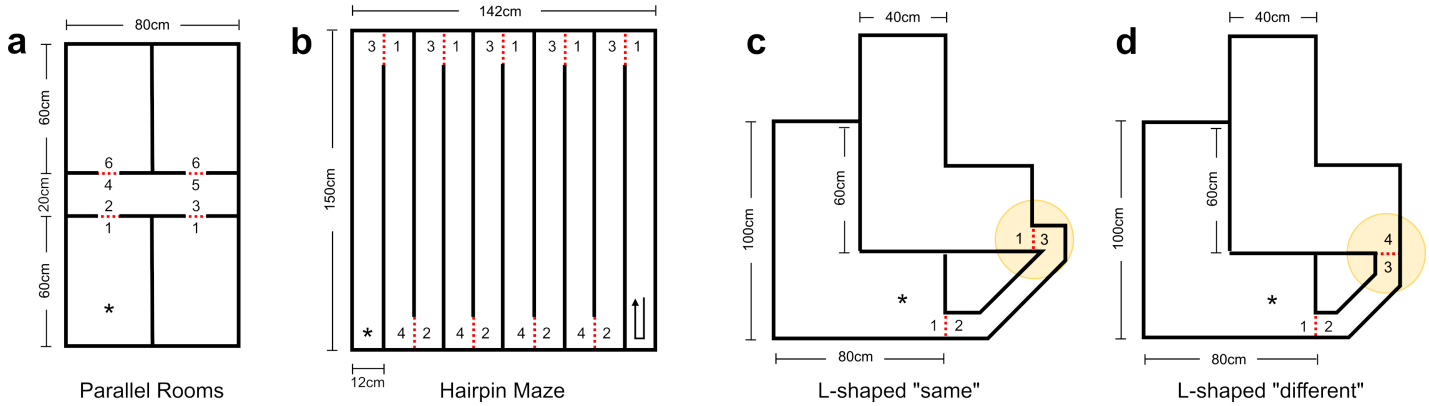

Figure 4: **(a) - (d)** Virtual environment floor plans used in the simulations. The red dotted lines mark region transitions and the adjacent integers the gateway identifier Γ: Equal numbers indicate the same local context and therefore cause remapping to the same patterns. The asterisk **(*)** marks the starting position. Note that in **(c)**, the gateway identifier of both rooms is the same , $\Gamma = 1$, but in **(d)**, it is different: $\Gamma = 1$ for the southwest and $\Gamma = 4$ for the northeast room.

L-shaped rooms (Fig. 4c,d). This novel layout consisted of two adjacent identical L-shaped rooms connected by a single diagonal hallway. In the "same context" condition (Fig. 4c), both rooms are entered at the same position (south-east corner) and with the same heading. The rooms are therefore locally indistinguishable and share the same gateway identifier. In the "different context" condition (Fig. 4d), one entrance was rotated by 90° so that the rooms are entered with a different heading. Previous studies have suggested that entry direction alone may be sufficient for rats to distinguish otherwise identical contexts (Tanila, 1999; Grieves et al., 2016). However, in these experiments, the entrances are visible from the entire room and could have been used as landmarks. In the L-shaped rooms, the entrance is not visible from the far part of the room. Models based on local cues would therefore predict the same activity in the far part of the room in both conditions.

2.8 Data analysis

For each environment, place and grid cell activity maps were produced by sorting the localized activity of each cell into 2×2cm bins, averaged over all time steps. Activity maps were smoothed by a 3×3 bin Gaussian kernel. To assess firing field repetition, Pearson correlations were calculated between different regions of interest on the same cell activity map. Only same-size regions were compared; we did not attempt to correlate scaled, rotated or mirrored cell activities. To account for random variations, each environment was explored ten times with different random number seeds, and the correlations were averaged. Place and grid cell results are reported separately.

Additionally, as a means of confirming the regions of interest as correlation maxima, we created autocorrelograms by correlating activity maps with bin-wise shifted versions of themselves.

3 Results

In the following, we report model predictions for a series of environments, in the form of binned place and grid cell activities. In the model, the cells form a combined attractor where the activity of one cell type directly influences the other. Therefore, correlation results for the different cell types, although reported separately, are highly similar.

3.1 Parallel rooms

We compared place and grid cell activity of matching bins for the four rooms (numbers 1 to 4 in fig. 5a) combined into parallel (west: rooms 1 and 3 versus east: 2 and 4) and opposed (north: 3 and 4 versus south: 1 and 2) sets. All place cells in all repetitions showed space-dependent activity patterns. Consistent with findings from animal research (Tanila, 1999; Fuhs et al., 2005; Spiers et al., 2015; Carpenter et al., 2015; Grieves et al., 2016), we observed highly repetitive activity for the parallel condition (mean east/west correlation 0.75, 95% confidence interval (CI) [0.69, 0.79]), but not in the opposed condition (north/south 0.01, 95% CI [-0.10,0.12]). Parallel room correlations were significantly higher than those of opposed rooms ($u = 0.0, p < 0.0001$, Mann-Whitney-U-Test, MWU).

Grid cell correlations match the findings from place cells. Activity similarly repeated in the parallel condition (east/west 0.84, 95% CI [0.81, 0.87]) and was uncorrelated in the opposed condition (north/south 0.01, 95% CI [-0.10,0.12]). The difference was highly significant ($u = 0.0, p < 0.0001$, MWU). The results suggest that remapping at room entrances alone is sufficient to decorrelate place and grid cell patterns for different contexts while still producing repetitive activity under similar contexts.

3.2 Hairpin maze

In the hairpin maze, firing fields have been shown to depend both on the direction of movement and arm identity (Derdikman et al., 2009). Trajectories (Fig. 6a) started in the first arm and traversed the maze *eastbound* until the end of the last arm, where the agent turned around and returned to the start, *westbound*. For each of the ten random initializations of the model, the maze was traversed twice in either direction, resulting in 40 traversals.

Results depended on the overall traversal directions (*eastbound* vs. *westbound*), as well as the local directions within each arm. Remapping occurred at the apices between the arms (Fig. 4b). In agreement with animal studies (Derdikman et al., 2009), model place and grid cell firing fields repeated in every second arm.

When correlating place fields in different arms, a clear checkerboard pattern emerged (Fig. 6b). Except for the initial arms (arm 1 for eastbound and arm 10 for westbound traversals), arms traversed in the same direction, i.e., odd- or even-numbered arms, showed significantly higher correlation (eastbound 0.88, 95% CI [0.86, 0.91]; westbound 0.89, 95% CI [0.86, 0.91]) than arms traversed in a different direction (eastbound 0.00, 95% CI [-0.11, 0.12]; westbound 0.08, 95% CI [-0.03, 0.19]; Same direction vs. alternating: eastbound $u = 0.0, p < 0.0001$, westbound $u = 0.0, p < 0.0001$, MWU). Firing fields in the eastbound and westbound traversals were uncorrelated (0.10, 95% CI [-0.02, 0.21]).

Grid cell results are generally similar to place cells. Grid cell firing fields are also highly correlated in same-direction arms (eastbound 0.91, 95% CI [0.89, 0.93]; westbound 0.91, 95% CI [0.89, 0.93]) but not in alternating arms (eastbound -0.01, 95% CI [-0.13, 0.10]; westbound 0.09, 95% CI [-0.02,0.2]; Same direction vs. alternating: eastbound $u = 0.0, p < 0.0001$, westbound $u = 0.0, p < 0.0001$, MWU). Again, overall eastbound and westbound traversal correlations are low (0.10, 95% CI [-0.01, 0.21]).

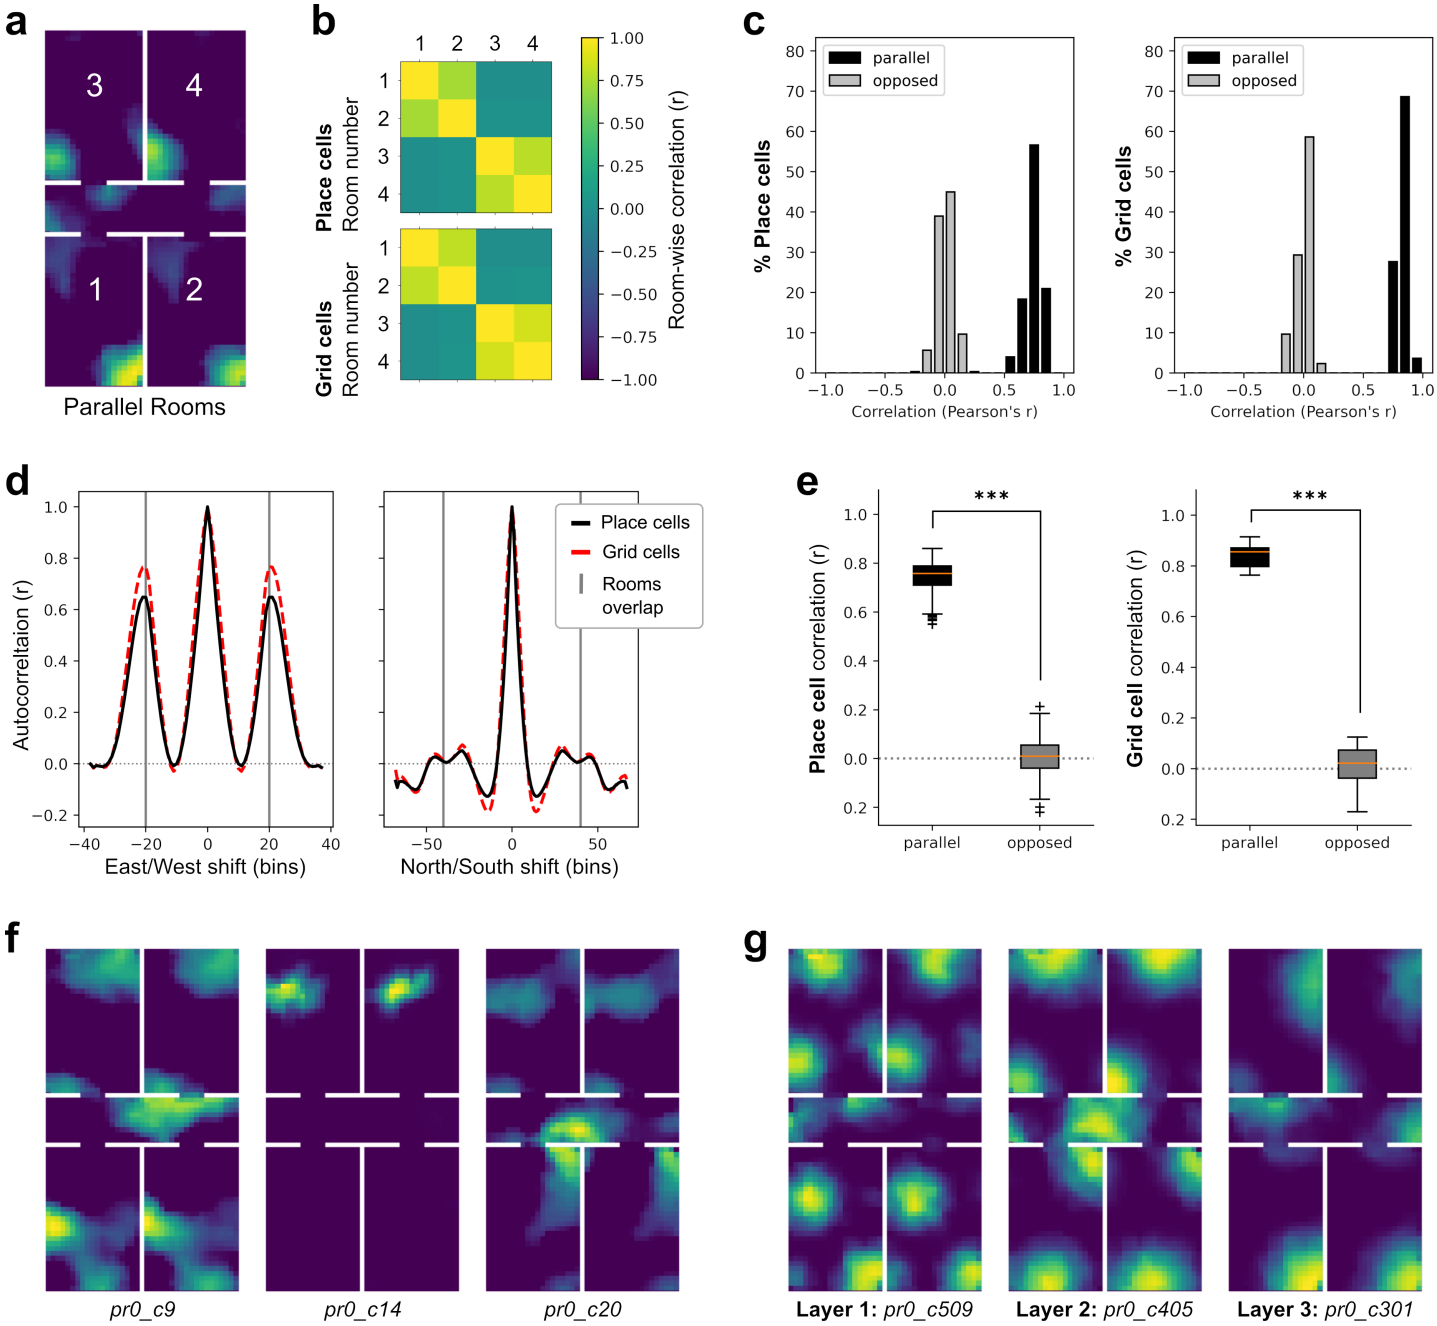

Figure 5: Parallel rooms. **(a)** Example of place cell activity in the different rooms (1-4) and the corridor (middle). **(b)** Average room-wise correlations for place and grid cells. As expected, the correlation between parallel rooms is high and the correlation between opposed rooms low. **(c, d)** Place cell (left) and grid cell (right) correlation distribution between parallel and opposed rooms. In agreement with the average room-wise correlation matrix in **(a)**, cell correlations cluster around 0 for the opposed condition. The distributions are relatively narrow and don't overlap. **(e)** Linear autocorrelogram of all cells shifted in bin-wise increments, including activity in the corridor. The vertical gray lines signal shifts where the rooms line up. **(f)** Examples of 3 place cells, expressing similar place fields in parallel rooms but different place fields in opposed rooms. **(g)** The same can be seen in grid cell firing fields. The grid structure in the larger layers is not visible because the rooms are shorter than the grid scale. The codes refer to the model repetition (pr0-pr9) and cell identity (0-599).

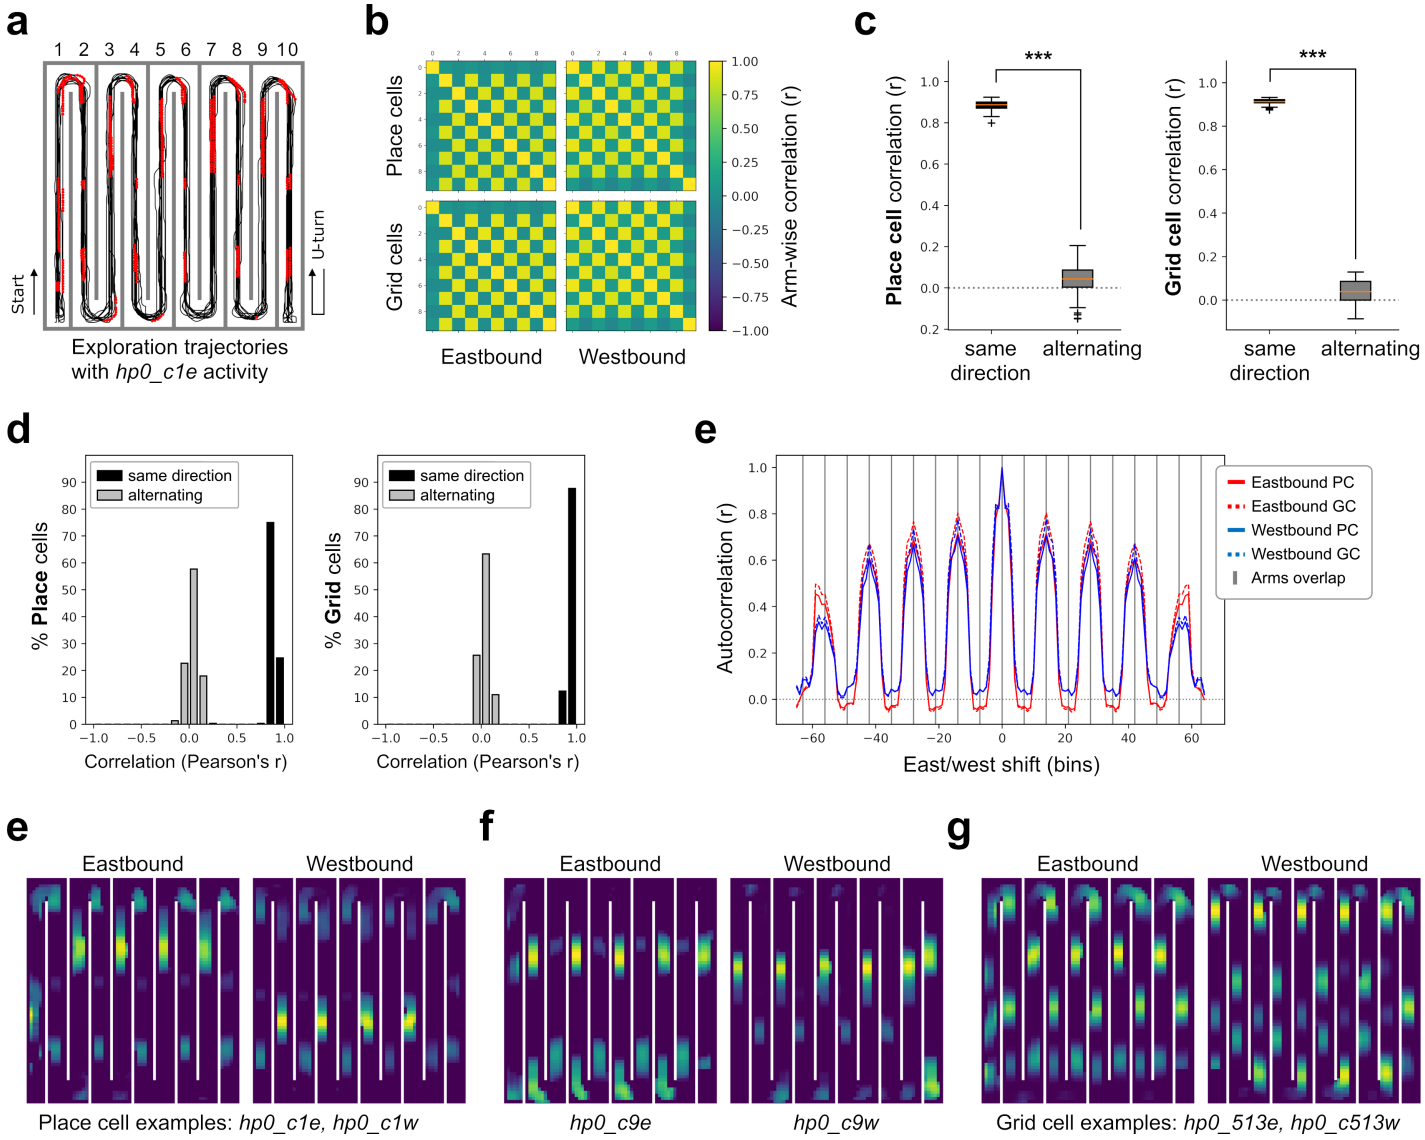

Figure 6: Hairpin Maze. **(a)** Exploration of the maze. The red dots mark eastbound cell activity of the cell in (f), where activity > 10% of max. **(b)** Arm-wise place cell (top) and grid cell (bottom) correlation for eastbound (left) and westbound (right) traversals. The low correlations for the outer arms result from the initial random model state for arm one, and the still active *eastbound* context for arm ten. **(c, d)** Place cell (left) and grid cell (right) correlation distribution for arms traversed in the same or different directions. The plots show the combined eastbound and westbound data. **(e)** Linear autocorrelogram of all cells, shifted in one-bin increments. The autocorrelation peaks whenever the activity is shifted by two arms. **(f,g)** Example place fields of the same cell for eastbound and westbound traversals. The cells show the characteristic repeating patterns. **(h)** Same as (f,g) but for a grid cell from the smallest layer. The grid pattern is obscured by the narrow maze arms.

The lack of correlation in the initial arms was also reported in (Derdikman et al., 2009). In our model, it is due to the fact that these arms have different gateway identifiers, as can be seen in Fig. 4b. Initially, no gateway identifier is stored for arm 1, and the pattern for further odd-numbered arms (Gateway identifier 2, Fig. 4b) is created only when arm 3 is entered for the first time. On the way back, arm 1 has gateway identifier 3 and is therefore correlated with the other odd-numbered arms. Arm 10 has the same gateway identifier before and after the u-turn, while the identifiers of the other even-numbered arms change on the way back, which explains the shifted correlation patterns between eastbound and westbound.

In general, the matching correlations in the hairpin maze are higher and the confidence interval is narrower compared to the other two environments. This is likely a result of the manual exploration with frequent region transitions, which caused accumulated path integrator noise to be reset.

3.3 L-shaped rooms

The L-shaped rooms environment was explored in two entrance conditions, *same*, where both rooms were entered from the same direction, and *different*, where both rooms were entered from different directions. Note that the conditions were explored by separate, unpaired models, i.e., cells are not matched between conditions.

As in the other environments with repeating activity, place cell correlation was significantly higher in the *same* condition (0.63, 95% CI [0.56, 0.69]) compared to the *different* condition (0.06, 95% CI [-0.05, 0.17], Place cell *same* vs. *different*: $u = 0.0, p < 0.0001$, MWU).

For grid cells, correlations in *same* (0.74, 95% CI [0.69, 0.79]) were higher than between place cells but similarly uncorrelated in the *different* condition (0.10, 95% CI [-0.01, 0.21]) (Statistic: $u = 0.0, p < 0.0001$, MWU).

Cell activity repeated when the rooms were indistinguishable at their entrances but differed otherwise. Compared to repeating activity in the other two environments, the correlations in the *same* condition were lower. This may be a result of a less frequent path integrator reset, since the L-shaped rooms contained only two region transitions. Therefore, input noise played a higher role.

The existence of the bend in the rooms plays no role for the model. It does not matter for the attractor whether the entrance is visible from a position or not. The L-shaped rooms only serve as an illustration of what sort of activation to expect in a case that has not previously been covered.

4 Discussion

Place and grid cells have been observed to express uncorrelated firing patterns in different contexts (Muller and Kubie, 1987; Lever et al., 2002; Leutgeb et al., 2005; Fyhn et al., 2007; Colgin et al., 2008; Julian et al., 2018). The phenomenon, called "remapping", has long been understood as a means of spatial pattern separation (Leutgeb et al., 2007) and may reflect a hierarchical organization of spatial representation (Julian et al., 2018). Remapping does not reflect purely local sensory differences, as both separate and identical patterns have been observed in similar contexts (Tanila, 1999; Carpenter et al., 2015; Grieves et al., 2016). However, the dynamics that give rise to the different firing patterns remain unknown.

In this study, we explore the idea that remapping is triggered by context change, defined as a transition between spatial regions such as the compartments of a maze. The mechanism requires a memory of contexts, implemented in our model as the gateway database. Within this memory, each context is represented as a pair consisting of (i) an identifier that can be detected

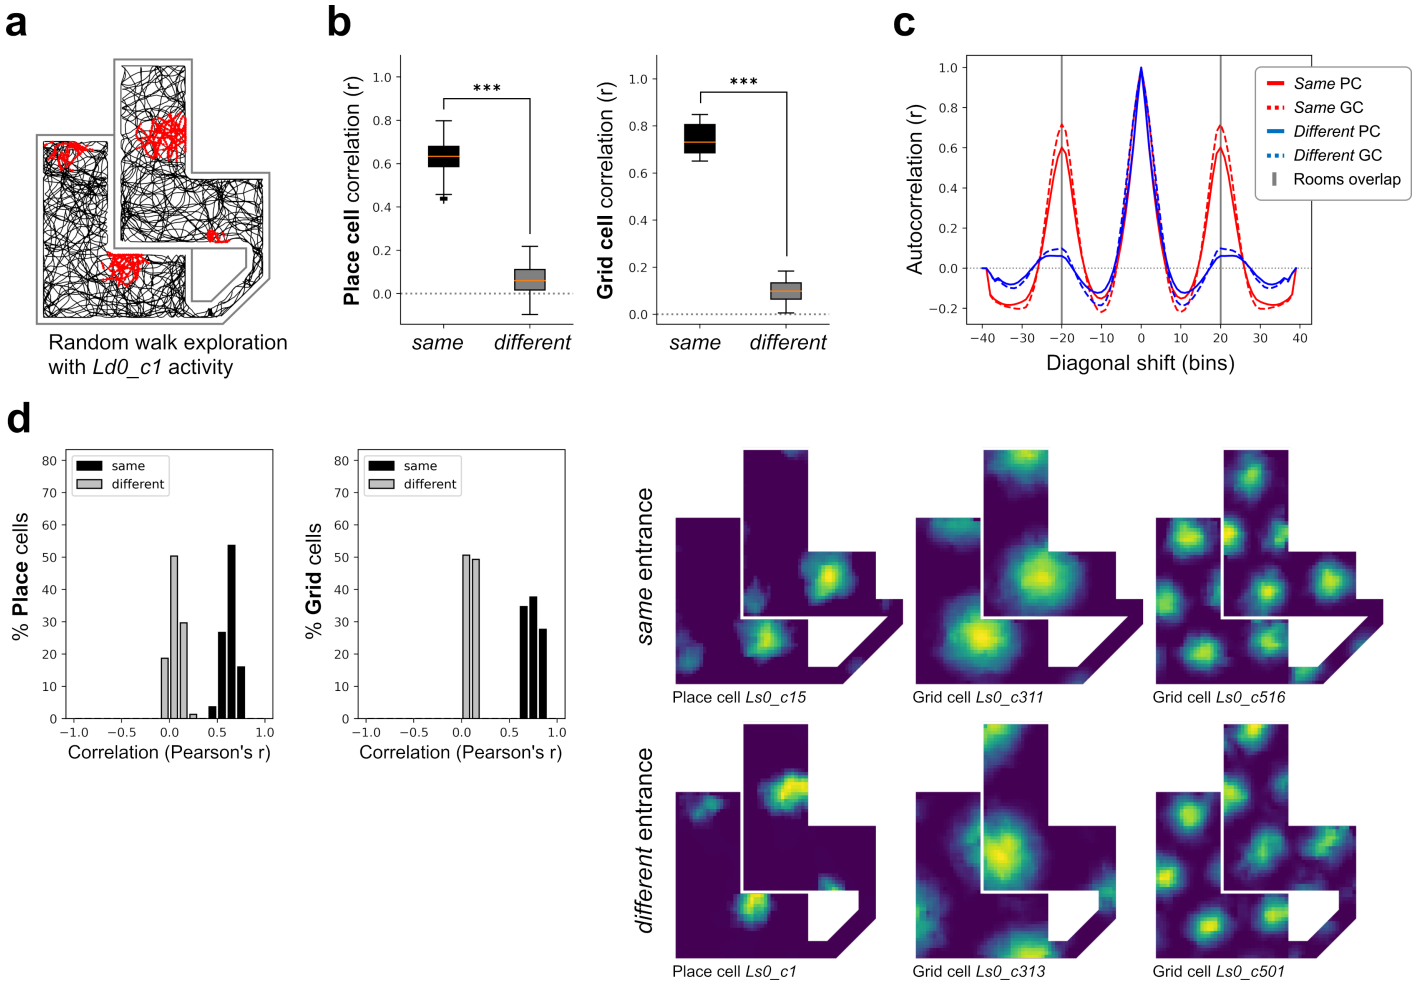

Figure 7: L-shaped rooms. **(a)** Random walk exploration in the *different* entrance condition. The red dots mark eastbound cell activity of the lower left cell in (e), where activity > 10% of max. **(b)** Place cell (left) and grid cell (right) correlation for *same* and *different* entrances. **(c)** Linear autocorrelogram of all cells, shifted diagonally in one-bin increments. The autocorrelation peaks in the *same* condition where the rooms overlap. The autocorrelogram includes firing information from the corridor, which was excluded from the other comparisons. **(d)** Cell correlations as in (b). **(e)** Firing field examples, from left to right in both conditions: Place cell, grid cell from the largest granularity layer, grid cell from the smallest granularity layer.

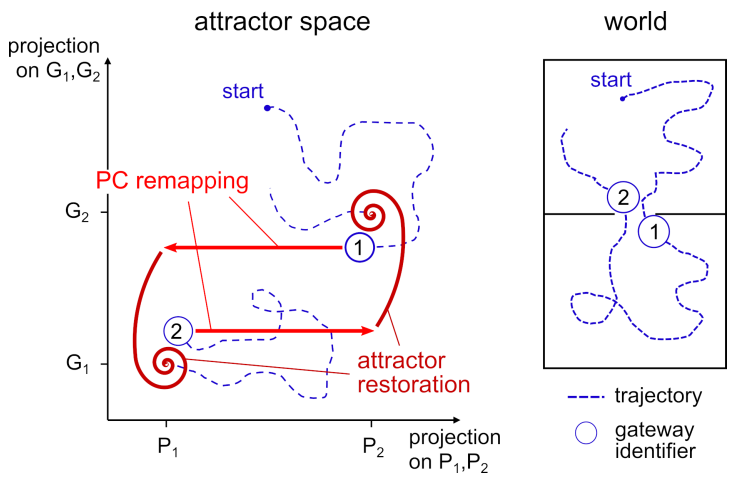

Figure 8: Phase diagram of the remapping process. Left: 2-dimensional projection of the joint state space of place and grid cells for two attractors P_1, G_1 and P_2, G_2. The projection is spanned by the straight lines trough P_1, P_2 (x-axis) and trough G_1, G_2 (y-axis). Right: Agent trajectory in a two-room maze. For further explanation, see text.

from sensory input such as a snapshot and (ii) the pattern of place cell activity prevailing when the context was first encountered. Remapping occurs when a known gateway identifier is detected, e.g. by snapshot matching. In this case, the ongoing place cell activities are overridden by the stored place cell activity pattern of the recognized gateway. Projections into entorhinal cortex drive the grid cell attractor into its corresponding state and the normal path integration dynamics proceed from there.

The process is schematically illustrated in fig. 8: In a two-room environment, the agent starts by exploring the northern room. The accompanying attractor trajectory is shown as a dashed blue line in the left part of the figure. As soon as the agent enters the southern room, the gateway identifier 1 is recognized, the place cells remap to the stored pattern, and the system converges to the closest attractor, (P_1, G_1). Exploration of the southern room results in continuous attractor shifts until the gateway is passed again (2) and another remapping event occurs.

Importantly, similar regions are assumed to share the same gateway identifier and entering them causes the same pattern to be activated. Thereby, the model was able to generate repetitive place and grid cell firing patterns qualitatively similar to those observed in multi-compartment environments. In the hairpin maze, the model also replicated the directionally selective firing commonly observed in corridors.

Within each region, grid cell activity is solely controlled by noisy path integration. This causes drift over time, which is also reset by remapping. Local cues, landmarks and geometry do not influence cell activity and were indeed not modeled at all. We thus present an alternative to boundary vector cell models of place cell remapping, where geometric cues measured by boundary vector cells are the main determinants for place fields (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018). In these models, repetitive patterns in similar compartments are a result of locally similar geometry.

With the L-shaped rooms environment, we propose a novel multicompartment setup where the two model types may be compared: When looking from the far part of the room. the bend in the L-shape obscures the room entrance, which is manipulated in two different contexts.

Our model predicts that the cell patterns selected at the entrance persist over the entire region, even in parts of the rooms that are locally indistinguishable. On the other hand, a model based on local geometry alone should find minimal differences in either context.

Note that we do not wish to imply that boundary vector cells play no role in place field generation. They provide position encoding independent from path integration (Bicanski and Burgess, 2020) and are able to explain place field deformation when the environment changes shape (O'Keefe and Burgess, 1996; Grieves et al., 2018) or barriers are inserted (Barry et al., 2006; Grieves et al., 2018). However, as mentioned in the introduction, boundary vector cells do not remap (Lever et al., 2009) and would need additional contextual gating to maintain distinct place cell patterns in different contexts. For this reason, we did not include the cells in our model.

4.1 Implications for biology

The underlying assumption behind the context-dependent cognitive map is one of hierarchy. This is directly implied by the vaguely defined term "context" itself, but also by the lack of remapping *within* a context; the spatial representation can therefore be at least divided in similar and dissimilar parts. The re-use of place cells in different population codes implies that, within a specific context, information about other contexts is unavailable. However, rats are clearly able to execute trajectories across region boundaries, so some sort of higher-level representation likely exists. Possible candidates may be cell ensembles in the medial prefrontal cortex (Hyman et al., 2012) or neurons in the perirhinal cortex (Bos et al., 2017): The former show different firing patterns depending on region, time and task and could therefore encode context (Hyman et al., 2012). The latter are neurons that fire continuously while the rat is within a specific region in space, such as an arm of a maze.

Of higher importance, however, is the representation of gateways. When another region is entered, the correct place cell pattern is selected based on information available at the entrance. Thus, a failure to recognize the gateway could result in catastrophic mislocalization. Clearly, rats are able to detect these context changes, but evidence for a representation is scarce. Spiers et al. (2015) report the clustering of place fields around doorways to compartments, indicating that the doorways are at least salient landmarks; on the other hand, Duvelle et al. (2021) explicitly did not find such an overrepresentation in an experiment where the passage between compartments was blocked by lockable doors.

After detection, the rat must also be able to store gateways and associated place cell patterns in long-term memory. A possible candidate may be the retrosplenial cortex (RSC): The rat RSC is reciprocally connected to the hippocampal formation and involved in a variety of spatial tasks (Vann et al., 2009). Lesions of the RSC only cause modest spatial deficits within a context, but impair spatial memory (Vann et al., 2009) and the ability to account for context changes, such as distal cue rotation or switching from light to darkness (Wesierska et al., 2009; Pothuizen et al., 2008).

Around 9% of RSC cells are head direction cells, some of which are additionally tuned to specific movements and locations (Chen et al., 1994; Jacob et al., 2017). Here, of particular interest may be a rare type of RSC cell described in the supplementary material of Jacob et al. (2017): The cells show localized firing at the doorway between two compartments and are tuned to a particular direction. This is precisely the behavior predicted for the gateway identifier units in our model.

Mathematically speaking, the definition and detection of gateways is non-trivial. A biologically plausible solution would be the prediction or tracking of changes geometric and visual information over time; at region transitions, this change is especially high (*surprise*, Butz et al.

(2004); Klukas et al. (2022)). In mobile robotics, the detection of doorways and narrow passages is a well-known problem with a variety of approaches based on cameras and laser range finders (e.g., Anguelov et al. (2004)) and posterior map segmentation (Thrun, 1998).

In humans, the cognitive map is generally thought to be hierarchically organized: in spatial memory and path planning tasks, differences arise in subjective distance estimation, recall speed (McNamara, 1986; McNamara et al., 1989) and preferred routes depending on region partition and transitions (Schick et al., 2019). In rodents, on the other hand, path planning is relatively unexplored. The best candidates are "preplays", high-frequency place cell sequences observed in periods of inactivity. The preplays depict future trajectories, sometimes even through unexplored parts of the environment, indicating a planning component (Dragoi and Tonegawa, 2011; Pfeiffer and Foster, 2013; Ólafsdóttir et al., 2015). However, whether preplays can cross region boundaries into other sub-maps is yet unknown.

4.2 Remapping at positions besides gateways

Remapping at gateways is sufficient to replicate the place and grid cell dynamics observed in rodent experiments, and it is the minimal requirement for distinct patterns in different contexts. However, in principle, remapping could also occur within a region. Key positions like feeding or resting places or salient landmarks could likewise be associated with stored place cell patterns. The key positions could serve as additional anchors for the path integrator or head direction system, or correct mislocalization, if the animal initially fails to recognize the region. In these cases, additional reactivation of the stored pattern would shift the active pattern to the correct attractor. In the model, this type of remapping would be straight-forward to implement, but would hinge on the correct recognition and identification of the key positions. As such, this type of remapping requires uniquely identifiable places and should be rarer than remapping at gateways to avoid aliasing.

4.3 Remapping without the MEC

Evidence points towards grid cells in the MEC as an important driving force behind hippocampal place cell remapping (Monaco and Abbott, 2011; Bush et al., 2014; Kanter et al., 2017): Grid cell remap concurrently with place cells, the MEC directly projects to the hippocampus (Fyhn et al., 2007), and electrical stimulation of the area leads to hippocampal remapping (Kanter et al., 2017). In accordance, place cell activity and remapping in our model directly arise from grid cell input. However, place cells must be at least partially independent from MEC input: Their activity persists and the cells continue to remap in different contexts, even if grid cell firing is disrupted (Brandon et al., 2014) or the entire MEC is lesioned (Schlesiger et al., 2018).

In our model, place cell activity is sustained by grid cell input and would therefore immediately cease if the input was severed. To maintain activity in the absence of grid cell input, the model place cells would either need to be augmented with recurrent weights or input from another source. This presents an opportunity to reconcile our model with boundary vector cell models (Hartley et al., 2000; Barry et al., 2006; Bush et al., 2014; Grieves et al., 2018), where local geometry, i.e., the distance to walls, determines place fields. A combined model could continue to function with interrupted path integration. Perhaps, a combined model with context-dependent grid cells and context-independent boundary vector cells could also explain the phenomenon of *partial* remapping Latuske et al. (2018), where only some firing fields remap, while others stay at the same location.

Interestingly, place fields that are newly formed while grid cells are interrupted, remain the

same when grid cell activity is later restored (Brandon et al., 2014). This is consistent with the proposed feedback connections from place cells to grid cells in our model; once grid cell activity is restored, the newly formed place fields should move the grid cell activity to the closest attractor.

In conclusion, we suggest that remapping is a mechanism subserving the integration of the systems for path integration and place recognition, i.e., O'Keefe and Nadel's (1978) Locale and Taxon systems. The cognitive map is thereby structured into a set of contexts or local charts, memorized by gateway identifiers. Within the charts, navigation is based on local metrics. On a coarser scale, the map is likely organized by the adjacencies of the charts, i.e., some sort of graph structure.

Acknowledgement

This work was carried out at the Department of Biology of the University of Tübingen, Tübingen, Germany. This research did not receive any specific grant from funding agencies in the public, commercial, or non-profit sectors.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org.

Allen, K., Rawlins, J. N. P., Bannerman, D. M., and Csicsvari, J. (2012). Hippocampal Place Cells Can Encode Multiple Trial-Dependent Features through Rate Remapping. *Journal of Neuroscience*, 32(42):14752–14766.

Anguelov, D., Koller, D., Parker, E., and Thrun, S. (2004). Detecting and modeling doors with mobile robots. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, volume 4, pages 3777–3784 Vol.4.

Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K., and Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, 17(1-2):71–97.

Bicanski, A. and Burgess, N. (2020). Neuronal vector coding in spatial cognition. *Nature Reviews Neuroscience*, 21(9):453–470.

Bjerknes, T. L., Moser, E. I., and Moser, M.-B. (2014). Representation of Geometric Borders in the Developing Rat. *Neuron*, 82(1):71–78.

Boccara, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., and Moser, M.-B. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience*, 13(8):987–994.

Bos, J. J., Vinck, M., van Mourik-Donga, L. A., Jackson, J. C., Witter, M. P., and Pennartz, C. M. A. (2017). Perirhinal firing patterns are sustained across large spatial segments of the task environment. *Nature Communications*, 8(1):15602.

Bostock, E., Muller, R. U., and Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2):193–205.

Bouchacourt, F. and Buschman, T. J. (2019). A Flexible Model of Working Memory. *Neuron*, 103(1):147–160.e8.

Brandon, M. P., Koenig, J., Leutgeb, J. K., and Leutgeb, S. (2014). New and Distinct Hippocampal Place Codes Are Generated in a New Environment during Septal Inactivation. *Neuron*, 82(4):789–796.

Brun, V. H., Leutgeb, S., Wu, H.-Q., Schwarcz, R., Witter, M. P., Moser, E. I., and Moser, M.-B. (2008). Impaired Spatial Representation in CA1 after Lesion of Direct Input from Entorhinal Cortex. *Neuron*, 57(2):290–302.

Bush, D., Barry, C., and Burgess, N. (2014). What do grid cells contribute to place cell firing? *Trends in Neurosciences*, 37(3):136–145.

Butz, M. V., Swarup, S., and Goldberg, D. E. (2004). Effective online detection of task-independent landmarks. *Online Proceedings for the ICML'04 Workshop on Predictive Representations of World Knowledge*.

Carpenter, F., Manson, D., Jeffery, K., Burgess, N., and Barry, C. (2015). Grid Cells Form a Global Representation of Connected Environments. *Current Biology*, 25(9):1176–1182.

Chen, L. L., Lin, L.-H., Green, E. J., Barnes, C. A., and McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. *Experimental Brain Research*, 101(1):8–23.

Cheng, K., Huttenlocher, J., and Newcombe, N. S. (2013). 25 years of research on the use of geometry in spatial reorientation: A current theoretical perspective. *Psychonomic Bulletin & Review*, 20(6):1033–1054.

Colgin, L. L., Leutgeb, S., Jezek, K., Leutgeb, J. K., Moser, E. I., McNaughton, B. L., and Moser, M.-B. (2010). Attractor-Map Versus Autoassociation Based Attractor Dynamics in the Hippocampal Network. *Journal of Neurophysiology*, 104(1):35–50.

Colgin, L. L., Moser, E. I., and Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, 31(9):469–477.

de Almeida, L., Idiart, M., and Lisman, J. E. (2009). The Input–Output Transformation of the Hippocampal Granule Cells: From Grid Cells to Place Fields. *Journal of Neuroscience*, 29(23):7504–7512.

Derdikman, D., Whitlock, J. R., Tsao, A., Fyhn, M., Hafting, T., Moser, M.-B., and Moser, E. I. (2009). Fragmentation of grid cell maps in a multicompartment environment. *Nature Neuroscience*, 12(10):1325–1332.

Diamantaki, M., Coletta, S., Nasr, K., Zeraati, R., Laturnus, S., Berens, P., Preston-Ferrer, P., and Burgalossi, A. (2018). Manipulating Hippocampal Place Cell Activity by Single-Cell Stimulation in Freely Moving Mice. *Cell Reports*, 23(1):32–38.

Dragoi, G. and Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401.

Duvelle, É., Grieves, R. M., Liu, A., Jedidi-Ayoub, S., Holeniewska, J., Harris, A., Nyberg, N., Donnarumma, F., Lefort, J. M., Jeffery, K. J., Summerfield, C., Pezzulo, G., and Spiers, H. J. (2021). Hippocampal place cells encode global location but not connectivity in a complex space. *Current Biology*, 31(6):1221–1233.e9.

Fuhs, M. C., VanRhoads, S. R., Casale, A. E., McNaughton, B., and Touretzky, D. S. (2005). Influence of Path Integration Versus Environmental Orientation on Place Cell Remapping Between Visually Identical Environments. *Journal of Neurophysiology*, 94(4):2603–2616.

Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., and Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–194.

Grieves, R. M., Duvelle, É., and Dudchenko, P. A. (2018). A boundary vector cell model of place field repetition. *Spatial Cognition & Computation*, 18(3):217–256.

Grieves, R. M. and Jeffery, K. J. (2017). The representation of space in the brain. *Behavioural Processes*, 135:113–131.

Grieves, R. M., Jenkins, B. W., Harland, B. C., Wood, E. R., and Dudchenko, P. A. (2016). Place field repetition and spatial learning in a multicompartment environment. *Hippocampus*, 26(1):118–134.

Guanella, A., Kiper, D., and Verschure, P. (2007). A model of grid cells based on a twisted torus topology. *International Journal of Neural Systems*, 17(04):231–240.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806.

Harland, B., Grieves, R. M., Bett, D., Stentiford, R., Wood, E. R., and Dudchenko, P. A. (2017). Lesions of the Head Direction Cell System Increase Hippocampal Place Field Repetition. *Current Biology*, 27(17):2706–2712.e2.

Hartley, T., Burgess, N., Lever, C., Cacucci, F., and O'Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus*, 10(4):369–379.

Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., and Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proceedings of the National Academy of Sciences*, 109(13):5086–5091.

Jacob, P.-Y., Casali, G., Spieser, L., Page, H., Overington, D., and Jeffery, K. (2017). An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nature Neuroscience*, 20(2):173–175.

Jeffery, K. J. (2011). Place Cells, Grid Cells, Attractors, and Remapping. *Neural Plasticity*, 2011:e182602.

Julian, J. B., Keinath, A. T., Marchette, S. A., and Epstein, R. A. (2018). The Neurocognitive Basis of Spatial Reorientation. *Current Biology*, 28(17):R1059–R1073.

Kanter, B. R., Lykken, C. M., Avesar, D., Weible, A., Dickinson, J., Dunn, B., Borgesius, N. Z., Roudi, Y., and Kentros, C. G. (2017). A Novel Mechanism for the Grid-to-Place Cell Transformation Revealed by Transgenic Depolarization of Medial Entorhinal Cortex Layer II. *Neuron*, 93(6):1480–1492.e6.

Keinath, A. T., Nieto-Posadas, A., Robinson, J. C., and Brandon, M. P. (2020). DG–CA3 circuitry mediates hippocampal representations of latent information. *Nature Communications*, 11(1):3026.

Klukas, M., Sharma, S., Du, Y., Lozano-Perez, T., Kaelbling, L., and Fiete, I. (2022). Fragmented Spatial Maps from Surprisal: State Abstraction and Efficient Planning.

Krupic, J., Bauza, M., Burton, S., Barry, C., and O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, 518(7538):232–235.

Langston, R. F., Ainge, J. A., Couey, J. J., Canto, C. B., Bjerknes, T. L., Witter, M. P., Moser, E. I., and Moser, M.-B. (2010). Development of the Spatial Representation System in the Rat. *Science*, 328(5985):1576–1580.

Latuske, P., Kornienko, O., Kohler, L., and Allen, K. (2018). Hippocampal Remapping and Its Entorhinal Origin. *Frontiers in Behavioral Neuroscience*, 11.

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., and Moser, E. I. (2007). Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. *Science*, 315(5814):961–966.

Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., and Moser, M.-B. (2005). Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. *Science*, 309(5734):619–623.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary Vector Cells in the Subiculum of the Hippocampal Formation. *Journal of Neuroscience*, 29(31):9771–9777.

Lever, C., Wills, T., Cacucci, F., Burgess, N., and O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, 416(6876):90–94.

Li, T., Arleo, A., and Sheynikhovich, D. (2020). Modeling place cells and grid cells in multicompartment environments: Entorhinal–hippocampal loop as a multisensory integration circuit. *Neural Networks*, 121:37–51.

Linsker, R. (1997). A Local Learning Rule That Enables Information Maximization for Arbitrary Input Distributions. *Neural Computation*, 9(8):1661–1665.

Lyttle, D., Gereke, B., Lin, K. K., and Fellous, J.-M. (2013). Spatial scale and place field stability in a grid-to-place cell model of the dorsoventral axis of the hippocampus. *Hippocampus*, 23(8):729–744.

McNamara, T. P. (1986). Mental representations of spatial relations. *Cognitive Psychology*, 18(1):87–121.

McNamara, T. P., Hardy, J. K., and Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(2):211–227.

McNaughton, B. L., Barnes, C. A., and O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, 52(1):41–49.

Monaco, J. D. and Abbott, L. F. (2011). Modular Realignment of Entorhinal Grid Cell Activity as a Basis for Hippocampal Remapping. *Journal of Neuroscience*, 31(25):9414–9425.

Moser, E. I., Moser, M.-B., and McNaughton, B. L. (2017). Spatial representation in the hippocampal formation: A history. *Nature Neuroscience*, 20(11):1448–1464.

Moser, M.-B., Rowland, D. C., and Moser, E. I. (2015). Place Cells, Grid Cells, and Memory. *Cold Spring Harbor Perspectives in Biology*, 7(2):a021808.

Muessig, L., Hauser, J., Wills, T. J., and Cacucci, F. (2015). A Developmental Switch in Place Cell Accuracy Coincides with Grid Cell Maturation. *Neuron*, 86(5):1167–1173.

Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *Journal of Neuroscience*, 7(7):1951–1968.

O'Keefe, J. and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381(6581):425–428.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1):171–175.

O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Clarendon Press ; Oxford University Press, Oxford : New York.

Ólafsdóttir, H. F., Barry, C., Saleem, A. B., Hassabis, D., and Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, 4:e06063.

Pfeiffer, B. E. (2020). The content of hippocampal "replay". *Hippocampus*, 30(1):6–18.

Pfeiffer, B. E. and Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79.

Pothuizen, H. H. J., Aggleton, J. P., and Vann, S. D. (2008). Do rats with retrosplenial cortex lesions lack direction? *European Journal of Neuroscience*, 28(12):2486–2498.

Rennó-Costa, C. and Tort, A. B. L. (2017). Place and Grid Cells in a Loop: Implications for Memory Function and Spatial Coding. *Journal of Neuroscience*, 37(34):8062–8076.

Rowland, D. C., Roudi, Y., Moser, M.-B., and Moser, E. I. (2016). Ten Years of Grid Cells. *Annual Review of Neuroscience*, 39(1):19–40.

Schick, W., Halfmann, M., Hardiess, G., Hamm, F., and Mallot, H. A. (2019). Language cues in the formation of hierarchical representations of space. *Spatial Cognition & Computation*, 19(3):252–281.

Schlesiger, M. I., Boublil, B. L., Hales, J. B., Leutgeb, J. K., and Leutgeb, S. (2018). Hippocampal Global Remapping Can Occur without Input from the Medial Entorhinal Cortex. *Cell Reports*, 22(12):3152–3159.

Skaggs, W. E. and McNaughton, B. L. (1998). Spatial Firing Properties of Hippocampal CA1 Populations in an Environment Containing Two Visually Identical Regions. *Journal of Neuroscience*, 18(20):8455–8466.

Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., and Moser, E. I. (2008). Representation of Geometric Borders in the Entorhinal Cortex. *Science*, 322(5909):1865–1868.

Spiers, H. J., Hayman, R. M. A., Jovalekic, A., Marozzi, E., and Jeffery, K. J. (2015). Place Field Repetition and Purely Local Remapping in a Multicompartment Environment. *Cerebral Cortex*, 25(1):10–25.

Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., and Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492(7427):72–78.

Tanila, H. (1999). Hippocampal place cells can develop distinct representations of two visually identical environments. *Hippocampus*, 9(3):235–246.

Taube, J. S., Muller, R. U., and Ranck, J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10(2):420–435.

Taube, J. S., Muller, R. U., and Ranck, J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience*, 10(2):436–447.

Thrun, S. (1998). Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71.

Trullier, O., Wiener, S. I., Berthoz, A., and Meyer, J.-A. (1997). BIOLOGICALLY BASED ARTIFICIAL NAVIGATION SYSTEMS: REVIEW AND PROSPECTS. *Progress in Neurobiology*, 51(5):483–544.

Unity Technologies (2021). Unity. Version 2020.3.4f1.

Vann, S. D., Aggleton, J. P., and Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, 10(11):792–802.

Wesierska, M., Adamska, I., and Malinowska, M. (2009). Retrosplenial cortex lesion affected segregation of spatial information in place avoidance task in the rat. *Neurobiology of Learning and Memory*, 91(1):41–49.

Wills, T., Barry, C., and Cacucci, F. (2012). The abrupt development of adult-like grid cell firing in the medial entorhinal cortex. *Frontiers in Neural Circuits*, 6.

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science*, 308(5723):873–876.

Wilson, M. A. and McNaughton, B. L. (1993). Dynamics of the Hippocampal Ensemble Code for Space. *Science*, 261(5124):1055–1058.

Zhang, K., Ganis, G., and Sereno, M. (1997). Anti-Hebbian synapses as a linear equation solver. In *Proceedings of International Conference on Neural Networks (ICNN'97)*, volume 1, pages 387–389 vol.1.