

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/328078107>

A Possible Encoding of 3D Visual Space in Primates

Chapter in *Studies in Computational Intelligence* - January 2019

DOI: 10.1007/978-3-319-99283-9_14

CITATIONS

0

READS

143

2 authors:



Jochen Kerdels

FernUniversität in Hagen

36 PUBLICATIONS 131 CITATIONS

[SEE PROFILE](#)



Gabriele Peters

FernUniversität in Hagen

88 PUBLICATIONS 565 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



CManipulator [View project](#)

A Possible Encoding of 3D Visual Space in Primates

Jochen Kerdels and Gabriele Peters

FernUniversität in Hagen – University of Hagen,
Universitätsstrasse 1, D-58097 Hagen, Germany

Abstract. Killian et al. were the first to report on entorhinal neurons in primates that show grid-like firing patterns in response to eye movements. We recently demonstrated that these *visual grid cells* can be modeled with our RGNG-based grid cell model. Here we revisit our previous approach and develop a more comprehensive encoding of the presumed input signal that incorporates binocular movement information and fixation points that originate from a three-dimensional environment. The resulting *volumetric firing rate maps* exhibit a peculiar structure of regularly spaced *activity columns* and provide the first model-based prediction on the expected activity patterns of *visual grid cells* in primates if their activity were to be correlated with fixation points from a three-dimensional environment.

1 Introduction

Investigating the information processing that occurs on a neuronal level in deep stages of the mammalian brain is a challenging task. Neurons in these regions of the brain commonly process information that is highly aggregated and difficult to correlate with external variables that are observable in an experimental setting. A rare exception in this regard are so-called *grid cells* [8, 10]. The activity of these neurons correlates in a regular and periodic fashion with the organism’s location in its environment and thus facilitates experimental inquiry. Since their discovery in the entorhinal cortex of rats cells with grid-like activity patterns were found in several mammalian species (rats, mice, bats, and humans) and multiple regions of the brain (entorhinal cortex, pre- and parasubiculum, hippocampus, parahippocampal gyrus, amygdala, cingulate cortex, and frontal cortex) [8, 10, 4, 25, 6, 11]. In all reported cases the activity of the observed neurons correlated with the organism’s location supporting the early and predominant interpretation that grid cells are a functional component of a system for navigation and orientation. Particularly, the relatively stable and periodic firing patterns of grid cells are thought of to provide a kind of *metric for space* by means of path integration [19]. However, recent observations [3, 2, 17, 23] indicate that the firing patterns of grid cells are much more dynamic and adaptive than previously assumed questioning their utility as such a metric. Moreover, Killian et al. [16] observed neurons with grid-like firing patterns in the entorhinal cortex

of primates whose activity does not correlate with the animal’s location but instead with gaze-positions in the animal’s field of view. These new observations challenge the original notion of grid cells as a specialized component for path integration [19, 24, 9, 1, 5, 20] and suggest that the grid-like activity patterns may reflect a more general form of information processing. To explore this alternative hypothesis we developed a new computational model of grid cells based on the *recursive growing neural gas* (RGNG algorithm), which describes the behavior of grid cells in terms of principles of self-organization that utilize a form of competitive Hebbian learning [14, 12]. We could demonstrate [12] that our RGNG-based model can not only describe the basic properties of grid cell activity but can also account for recently observed phenomena of dynamic grid pattern adaptation in response to environmental changes [3, 2, 17, 23].

Furthermore, we recently showed [15] that our model can describe the behavior of “visual” grid cells in primates as they were reported by Killian et al. [16]. However, the data presented by the latter is not as clean as comparable data from, e.g., experiments with rats (see fig. 4). This difference may just be an artefact of challenging recording conditions in primates or caused by a limited amount of data, but it could also reflect a difference in the actual behavior of the observed cells. Here we investigate this question further by extending our previously used two-dimensional input model into a three-dimensional version. The following two sections summarize the RGNG-based grid cell model and revisit the results of our recent investigation into modelling and simulation of *visual grid cells* in primates [15]. Section 4 then introduces a three-dimensional input model that is based on the efference copy of motor neuron populations that putatively control the binocular gaze direction. Section 5 presents the simulation results we obtained using this new input model and shows how the results relate to the characteristics of the data published by Killian et al. [16]. Finally, section 6 draws conclusions and provides a set of testable predictions for future neuroscientific research.

2 RGNG-based Grid Cell Model

The predominant interpretation that grid cells are a functional component of a system for navigation and orientation is reflected by the majority of conventional grid cell models. Typically, they incorporate mechanisms that directly integrate information on the velocity and direction of the respective animal. Requiring this domain specific information as input renders these models incapable of describing the behavior of other neurons with similar grid-like firing patterns but different types of input signals like the *visual grid cells* observed by Killian et al. [16]. The RGNG-based grid cell model avoids such domain specific dependencies. It is a neuron-centric model in which neurons act in their “own interest” while being in local competition with each other. Biological neurons receive thousands of inputs from other neurons, and from the perspective of a single neuron these inputs are just electrochemical signals that carry no domain specific information. The entirety of these input signals and their possible values (i.e., states of activity)

constitute the input space of a neuron. Besides additional sources of information such as neuromodulators the *Umwelt* of a neuron is primarily defined by the structure of this input space. We hypothesize that neurons “want” to maximize their activity in response to this Umwelt as recent findings indicate that neuronal activity increases the direct glucose uptake of neurons [18]. Hence, from the perspective of a “selfish” neuron being more active means getting more energy. To this end, we assume that neurons are able to learn a limited number of input patterns or *prototypes* within their dendritic tree such that encountering any of these patterns will cause the neuron to become active. In that case, maximizing activity translates into learning those input patterns that occur most often while simultaneously trying to avoid learning near duplicate patterns which would waste “dendritic memory” capacity. As a consequence, the set of learned input patterns form a simple prototype-based representation of the input space structure residing in the dendritic tree of the respective neuron. In addition, if multiple neurons compete against each other via local inhibition they will form input space representations that are pairwise distinct from each other given that each competing neuron still wants to maximize it’s activity. In such a group of competing neurons the individual simple representations will interleave in such a way that a complex representation of the input space emerges that is distributed over the entire group. In our model we use a two layer recursive growing neural gas (RGNG) to describe both processes at once: the learning of prototypes within individual cells as well as the simultaneous competition among the cells in the group. Interestingly, in both cases the underlying dynamics follow the same principles of self-organization utilizing a form of competitive Hebbian learning.

Most importantly, the RGNG-based neuron model can operate on arbitrary input spaces. For any input space the modeled group of neurons will try to learn the structure of the particular input space as well as possible. If the activity of a single modeled neuron is then correlated with a suitable external variable, the individual firing fields that can be observed correspond to the individual prototypes or input space patterns that the neuron has learned. The “locations” of the firing fields in relation to the external variable are a pointwise mapping between the learned input space structure and the value range of that particular external variable. For instance, to observe the typical grid-like firing pattern of grid cells the input patterns must originate from a two-dimensional, uniformly distributed manifold in the input space and have to correspond to the location of the organism in it’s environment. Notably, these basic requirements allow for a multitude of possible input spaces for grid cells, which can then be tested with respect to further observed properties like the dynamic adaptation of grid patterns to environmental changes. Likewise, by choosing a suitable input space it becomes possible to model observations of grid-like firing patterns in other domains like the *visual grid cells* reported by Killian et al. [16] (see next section). A formal description of the RGNG-based model is provided in the appendix. For an in-depth characterization of the model we refer to our prior work [15, 12].

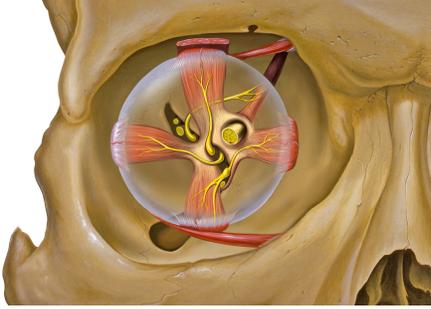


Fig. 1. Eye and orbit anatomy with motor nerves by Patrick J. Lynch, medical illustrator; C. Carl Jaffe, MD, cardiologist (CC BY 2.5). Extracted from [15].

3 Encoding of 2D Visual Space Revisited

Neurons with grid-like activity patterns were observed in several mammalian species and multiple regions of the brain [8, 10, 4, 25, 16, 6, 11]. Among these findings the observations of Killian et al. [16] stand out as they are the first to report grid-like activity patterns that are not correlated with the organism’s location but with gaze-positions in the field of view. We recently investigated if and how these *visual grid cells* can be described by our RGNG-based neuron model [15]. In this section we revisit and summarize our results from this investigation.

We hypothesized that the observed neurons may receive an input signal that originates as a so-called *efference copy* from the populations of motor neurons that control the four main muscles attached to the eye (Fig. 1). In such a population signal the number of active neurons corresponds to the degree with which the particular muscle contracts. Hence, the signal would provide an indirect measure of the gaze-position of an eye. We constructed a corresponding input signal $\xi := (v^{x_0}, v^{x_1}, v^{y_0}, v^{y_1})$ for a given normalized gaze position (x, y) by using four concatenated d -dimensional vectors v^{x_0} , v^{x_1} , v^{y_0} and v^{y_1} :

$$\begin{aligned}
 v_i^{x_0} &:= \max\left[\min\left[1 - \delta\left(\frac{i+1}{d} - x\right), 1\right], 0\right], \\
 v_i^{x_1} &:= \max\left[\min\left[1 - \delta\left(\frac{i+1}{d} - (1-x)\right), 1\right], 0\right], \\
 v_i^{y_0} &:= \max\left[\min\left[1 - \delta\left(\frac{i+1}{d} - y\right), 1\right], 0\right], \\
 v_i^{y_1} &:= \max\left[\min\left[1 - \delta\left(\frac{i+1}{d} - (1-y)\right), 1\right], 0\right],
 \end{aligned}$$

$$\forall i \in \{0 \dots d-1\},$$

with $\delta = 4$ defining the “steepness” of the population signal and the size d of each motor neuron population.

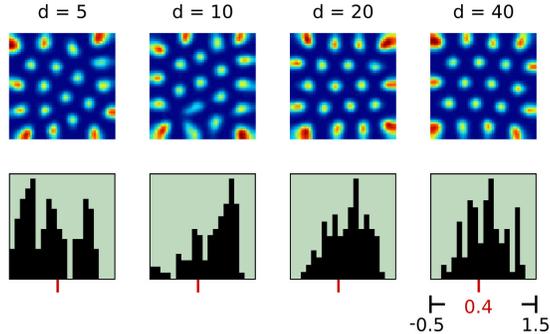


Fig. 2. Artificial rate maps and gridness distributions of simulation runs that processed input from a varying number d of presumed motor neurons per muscle (columns) that control the gaze-position. All simulation runs used a fixed set of parameters (table 1) and processed two million random gaze locations encoded as the activity of corresponding motor neuron populations. Each artificial rate map was chosen randomly from the particular set of rate maps and displays all of the respective neuron’s firing fields. The gridness distributions show the gridness values of all neurons from the particular neuron group. Gridness threshold of 0.4 indicated by red marks. Extracted from [15].

Using this input space encoding we conducted a series of simulation runs with varying sizes d of the presumed motor neuron populations. Each run simulated a group of 100 neurons with each neuron using 20 prototypes to learn an input space representation¹. The input consisted of two million, random gaze-positions per simulation run encoded as described above. Based on the resulting neuron activity and corresponding gaze-positions we constructed standard *firing rate maps* that represent a correlation of these variables. The maps were constructed according to the procedures described by Sargolini et al. [21] but using a 5×5 boxcar filter for smoothing instead of a Gaussian kernel as introduced by Stensola et al. [22]. This conforms to the de facto standard of rate map construction in the grid cell literature. The firing rate maps integrate activity and position information of 30000 time steps sampled from the end of a simulation run, which corresponds to typical experimental trial durations in grid cell research, i.e., 10 minutes recorded at 50 Hz. In addition, we calculated the *gridness score* for each simulated neuron based on the constructed firing rate maps. The gridness score is an established measure to assess how grid-like an observed activity pattern is [21]. It ranges from -2 to 2 with scores greater zero indicating a grid-like activity pattern. To classify an observed neuron as grid cell, recent publications choose more conservative thresholds between 0.3 and 0.4.

Figure 2 summarizes the results of these simulation runs for motor neuron populations of size $d \in \{5, 10, 20, 40\}$ per eye muscle. Both the shown firing rate maps as well as the gridness score distributions show that a significant proportion of the simulated neurons form grid-like activity patterns in response to the

¹ The full set of parameters used is given in the appendix (table 1).

input signal described above. Furthermore, the response of the neurons appears to be robust over the entire range of population sizes that were investigated. The fact that the activity of the simulated neurons forms a grid-like firing pattern when correlated with a two-dimensional, external variable (gaze-position) indicates, that the inputs themselves originate from a two-dimensional manifold lying within the high-dimensional ($d \times 4$) input space – which is unsurprising in this case since we constructed the encoding of the input signal to have this property (but see section 4 why this is only an *indication* in general).

The gaze-positions that were randomly chosen as input signals in the simulation runs cover the entire field of view and the resulting firing rate maps show the entire input space representation that was learned by a single neuron based on these inputs. As a consequence, a strong alignment of firing fields, i.e., prototypes at the borders of the firing rate maps can be observed since the borders of the firing rate maps coincide with the outer limits of the underlying, low-dimensional input space manifold in this case. As it is unlikely that experimental observations of natural neurons will cover the entire extent of an underlying input space, we investigated how the partial observation of firing fields may influence rate map appearance and gridness score distributions. To this end we conducted a second series of simulation runs using the same set of parameters as before with the exception of using 80 prototypes per neuron instead of 20. Furthermore, we constructed additional firing rate maps for each neuron that contain only one-quarter or one-sixteenth of the respective neuron’s firing fields emulating a partial observation of natural neurons. Figure 3 shows the results of this second series of simulation runs. The additional firing rate maps that contain only a subset of firing fields display firing patterns that are much more regular and periodic compared to the firing patterns shown in figure 2. The corresponding gridness score distributions support this assessment by showing a larger proportion of simulated neurons with gridness scores above a threshold of 0.4. This indicates that the distortions introduced by the alignment of firing fields at the outer limits of the underlying input space manifold remain local with respect to their influence on the grid-like structure of other regions. Thus, natural neurons may receive signals from input space manifolds that are only partially two-dimensional and evenly distributed. In such a case grid-like firing patterns would only be observed if the experimental conditions happen to restrict the input signals to these regions of input space. Any shift towards regions with different properties would result in a distortion or loss of the grid-like structure.

4 Encoding of 3D Visual Space

In our recent investigation on modeling *visual grid cells* [15] we argued, as summarized above, that the firing rate map derived from the observation of a biological neuron would likely show only a subset of this neuron’s firing fields since experimental conditions would typically not allow to probe the entire input space of that neuron. Analyzing only a fraction of the firing fields of our simulated neu-

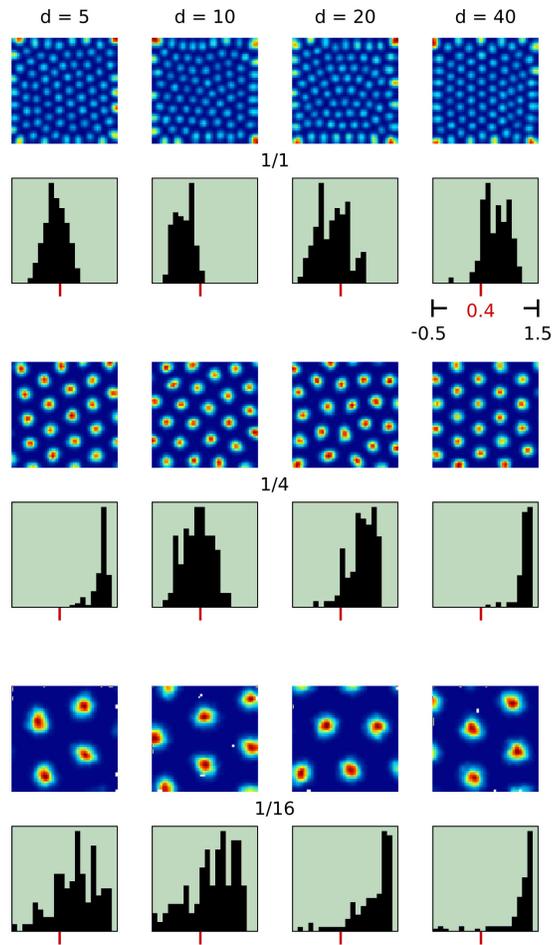


Fig. 3. Artificial rate maps and gridness distributions of simulation runs presented like in figure 2, but with 80 prototypes per neuron and containing either all, one-quarter, or one-sixteenth (rows) of the respective neuron's firing fields. Extracted from [15].

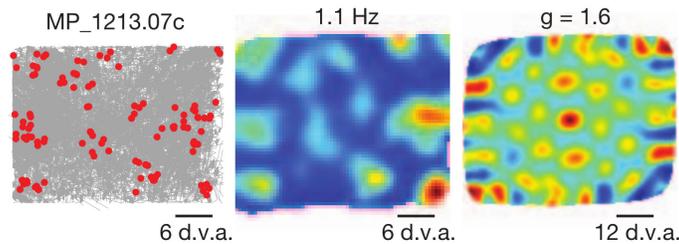


Fig. 4. Example of a neuron with a grid-like activity pattern in the monkey as observed by Killian et al. [16]. Left, plots of eye position (grey) and spikes (red) reveal non-uniform spatial density of spiking. For clarity, only spikes corresponding to locations of firing rate above half of the mean rate were plotted. The monkey's name and unit number are indicated at the top. Middle, spatial firing-rate maps show multiple distinct firing fields. Maps are colour coded from low (blue) to high (red) firing rates. The maximum firing rate of the map is indicated at the top. Right, the spatial periodicity of the firing fields shown against spatial autocorrelations. The colour scale limits are ± 1 (blue to red), with green being 0 correlation. d.v.a., degrees of visual angle; g, gridness. Reprinted by permission from Macmillan Publishers Ltd: Nature, Killian et al. [16], copyright 2012.

rons showed an improvement in the overall regularity of the firing patterns and a reduction of the influence of alignment artifacts located at the outer limits of the underlying input space manifold. However, the activity patterns observed by Killian et al. [16] differ from these simulation results. They appear to be even less regular and “clean” in their grid-like structure (Fig. 4) as our simulated rate maps of the entire visual field (Fig. 2). It is possible that the distortions in the observed activity patterns are just artifacts that result from challenging recording conditions in primates, or potential inaccuracies in the measurement of the external variable (eye tracking); and in our previous investigation we assumed that this was indeed the case. Yet, it is also possible that the distortions are no artifacts and may reflect a more complex underlying input space structure.

To explore this alternative explanation we considered a more realistic input space model that includes both eyes and stereoscopic vision. Instead of encoding two-dimensional gaze-positions of a single eye we took random fixation points p_f in a three-dimensional view box (width: 400 cm \times height: 250 cm \times depth: 300 cm) and derived the corresponding horizontal and vertical viewing angles for both eyes at their respective locations. Allowed eye movement was restricted to a typical range of 50° adduction (towards the nose), 50° abduction (towards the temple), 60° infraduction (downwards), and 45° supraduction (upwards). Eyes were positioned 6 cm apart at 200 cm \times 125 cm \times 0 cm. Fixation points outside of the field of view of either eye as well as fixation points closer than 30 cm were rejected and not processed. After determining the horizontal and vertical viewing angles for each eye the angles were normalized to the interval $[0, 1]$ and encoded as presumed population signal from the respective eight motor neuron

populations (four per eye) as in our previous simulation experiments. A fixed population size of 20 motor neurons was used resulting in a 160-dimensional input space. We conducted a series of simulation runs with varying numbers of prototypes $\{20, 40, 80, 160\}$ per neuron. Each run simulated a group of 100 neurons that processed 6 million input patterns derived from randomly chosen fixation points within the view box as described above. The full set of parameters used is provided in table 2 in the appendix.

5 Results

To assess the outcome of the simulation runs we constructed for each neuron a *volumetric firing rate map* that correlates the simulated neural activity with the corresponding fixation points within the view box. Figure 5 shows one exemplary firing rate map for each simulation run. The firing rate maps have a resolution of $3\text{ cm} \times 3\text{ cm} \times 3\text{ cm}$, and a $5 \times 5 \times 5$ boxcar filter adapted from the two-dimensional version introduced by Stensola et al. [22] was applied for smoothing. The normalized neuronal activity is color-coded from blue-transparent (low activity) to red-opaque (high activity). The maps are oriented such that the eyes are located at the center of the horizontal plane (width \times height) looking upwards in the positive direction of the depth dimension.

The most prominent feature visible in these *volumetric firing rate maps* are columns of activity that radiate outwards from the location of the eyes. The columns appear evenly spaced, having similar diameters, and approaching a hexagonal arrangement with increasing depth. Columns at the outer limits of the input space manifold appear thicker and elongated outwards. The activity distribution within each column is non-homogeneous, but seems to never fall below a level that remains significantly larger than the activity outside of the column. In case of neurons with a number of 160 prototypes this continuity of activity appears to weaken as some columns show stretches of very low activity. This observation matches the expected increase in competition among individual neurons caused by the increasing number of prototypes per neuron, i.e., multiple neurons will compete for regions of input space within single columns.

In order to study the alignment of the activity columns with progressing depth we extracted a series of horizontal slices through the *volumetric firing rate maps* shown in figure 5. The resulting slices are shown in figure 6 with each slice integrating a depth range of 30 cm. The increase in regularity of the column alignment with increasing depth is clearly visible. The resulting patterns at a depth of 240 cm (Fig. 6, last row) vary between more diagonally aligned rectangular patterns in neurons with 20 and 40 prototypes to more hexagonal patterns in neurons with 80 and 160 prototypes. At depths of about 180 cm or less the outer limits of the underlying input space manifold determined by the maximum viewing angles of the modeled eyes become visible in the slices. Similarly to our previous results (Fig. 2) an alignment of the firing fields with respect to these limits can be observed. In addition, the outer firing fields appear larger and elongated outwards. This change in size and shape may in part be

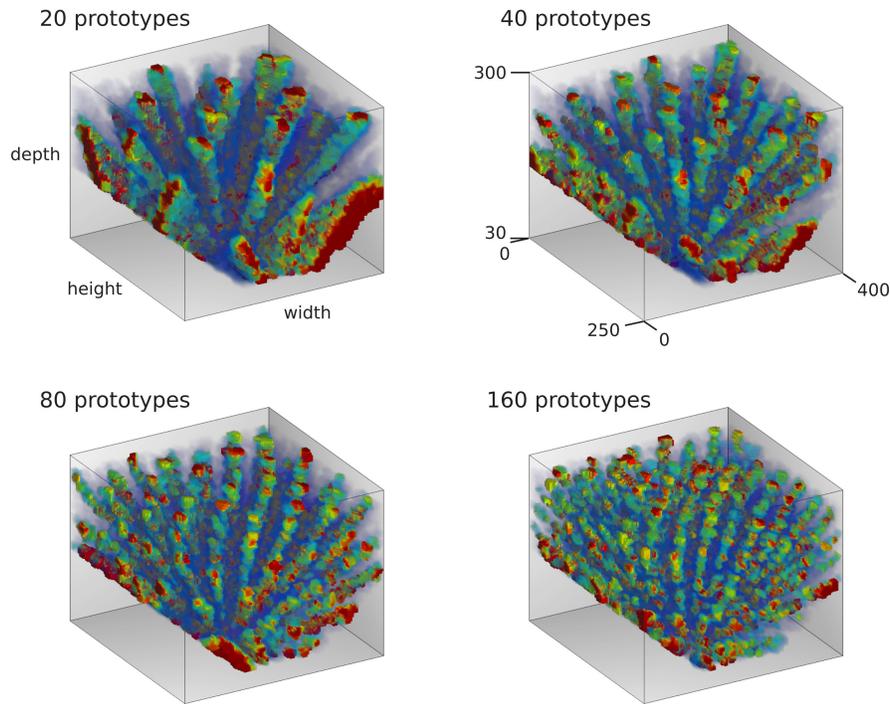


Fig. 5. Volumetric firing rate maps derived from simulation runs with varying numbers of prototypes $\{20, 40, 80, 160\}$ per neuron. Each firing rate map was chosen randomly from the respective set of neurons. The rate maps correlate the simulated neuronal activity of a single neuron with the corresponding fixation point within the view box. Eyes are located at the center of the width and height dimensions looking upwards in the direction of the depth dimension. Normalized neuronal activity is color-coded from blue-transparent (low activity) to red-opaque (high activity). Resolution of the firing rate maps is $3\text{ cm} \times 3\text{ cm} \times 3\text{ cm}$.

caused by cutting horizontally through columns that are angled outwards from the center. However, if this would be the sole cause for the deformation we would expect to see a gradual increase in the deformation of firing fields with respect to their distance from the center, which is not observable in the slices.

The slices shown in figure 6 are also an approximation of the two-dimensional firing rate maps we would expect to observe if an animal with forward-facing binocular vision, like the macaque monkeys studied by Killian et al. [16], would be given the task of watching images presented on a two-dimensional monitor in front of them. In case of the experiments performed by Killian et al. the distance between the monkeys and the monitor that displayed the images was 60 cm,

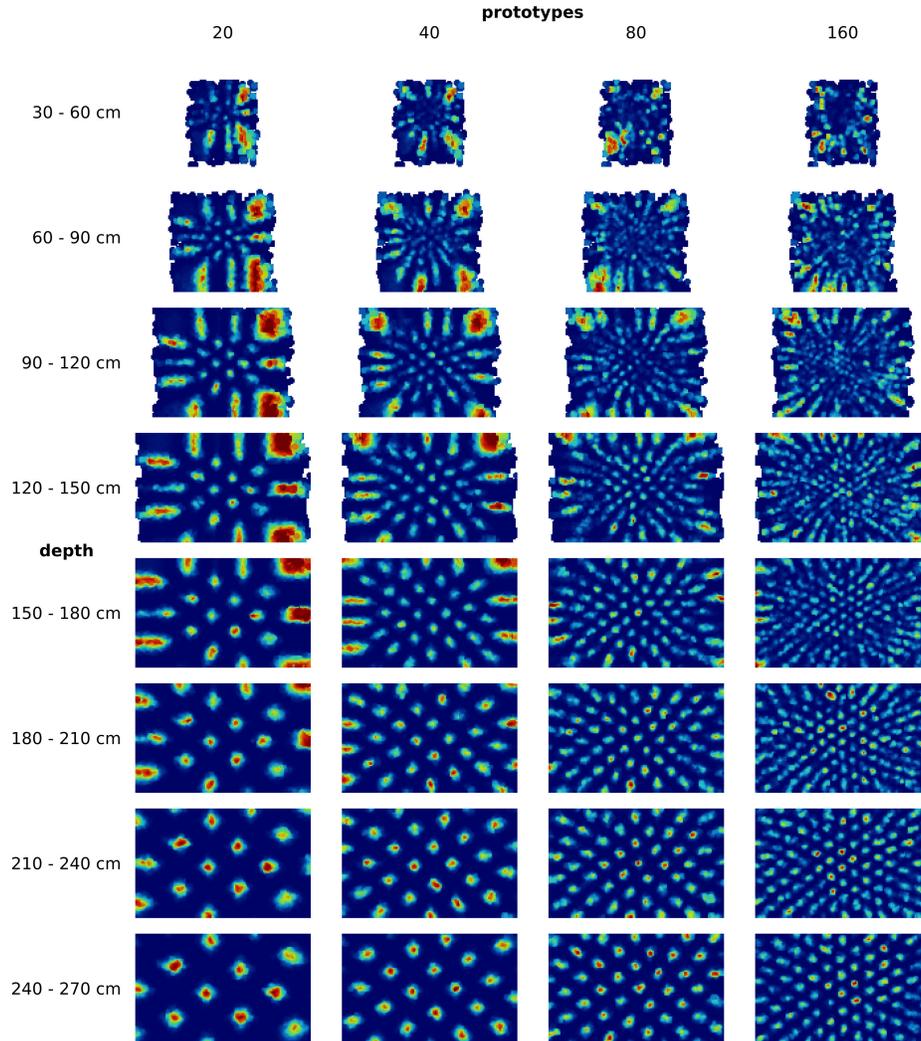


Fig. 6. Horizontal slices through the *volumetric firing rate maps* shown in figure 5. Each slice integrates a depth range of 30 cm. Normalized neuronal activity is color-coded from blue (low activity) to red (high activity). Regions with no data points are drawn in white.

and the constructed firing rate maps covered $33^\circ \times 25^\circ$ of visual angle. Values for adduction, abduction, infraduction, and supraduction as well as distance between the eyes in macaque monkeys was not provided by Killian et al., and we were unable to acquire this information reliably from other sources. Thus, we used common values for human eyes in our simulation experiments. Despite this shortcoming, the comparison of the firing rate map provided by Killian et

al. [16] (Fig. 4) and the slices of *volumetric firing rate maps* shown in the first rows of figure 6 shows some interesting structural similarities: In both cases the inner firing fields are weaker in activity, less regular, and smaller in size, whereas the outer firing fields are more pronounced, stronger in activity, and appear elongated outwards.

6 Conclusions

Killian et al. [16] were the first to report on entorhinal neurons in primates that show grid-like firing patterns in response to eye movements. We recently investigated how such a behavior could be modeled with our RGNG-based neuron model and we were able to demonstrate that it is feasible in principle [15]. In this paper we revisited our approach to modeling these neurons and provided a more comprehensive encoding of the presumed input signal that incorporates movement information from both eyes and fixation points that originate from a three-dimensional environment. The resulting *volumetric firing rate maps* of the simulated neurons exhibit a peculiar structure of regularly spaced *activity columns* that are angled outwards from the presumed eye locations. To the best of our knowledge these *volumetric firing rate maps* are the first model-based prediction on the expected activity patterns of *visual grid cells* in primates if their activity were to be correlated with fixation points from a three-dimensional environment.

In addition, horizontal slices through these *volumetric rate maps* show that the alignment of the activity columns increases in regularity with increasing distance from the eyes. Thus, for an experimental setup as it was used by Killian et al. [16] we predict that the distance between subject and the presented stimuli has a significant influence on the gridness of the observed activity patterns. A possible modification of the experimental design to test this prediction would be to use a video projector while placing the projection screen at various distances keeping the horizontal and vertical viewing angles constant.

APPENDIX

Recursive Growing Neural Gas

The recursive growing neural gas (RGNG) has essentially the same structure as the regular growing neural gas (GNG) proposed by Fritzke [7]. Like a GNG an RGNG g can be described by a tuple²:

$$g := (U, C, \theta) \in G,$$

with a set U of units, a set C of edges, and a set θ of parameters. Each unit u is described by a tuple:

$$u := (w, e) \in U, \quad w \in W := \mathbb{R}^n \cup G, \quad e \in \mathbb{R},$$

² The notation $g.\alpha$ is used to reference the element α within the tuple.

with the *prototype* w , and the *accumulated error* e . Note that in contrast to the regular GNG the prototype w of an RGNG unit can either be a n -dimensional vector or another RGNG. Each edge c is described by a tuple:

$$c := (V, t) \in C, \quad V \subseteq U \wedge |V| = 2, \quad t \in \mathbb{N},$$

with the units $v \in V$ connected by the edge and the *age* t of the edge. The *direct neighborhood* E_u of a unit $u \in U$ is defined as:

$$E_u := \{k | \exists (V, t) \in C, \quad V = \{u, k\}, \quad t \in \mathbb{N}\}.$$

The set θ of parameters consists of:

$$\theta := \{\epsilon_b, \epsilon_n, \epsilon_r, \lambda, \tau, \alpha, \beta, M\}.$$

Compared to the regular GNG the set of parameters has grown by $\theta.\epsilon_r$ and $\theta.M$. The former parameter is a third learning rate used in the adaptation function A (see below). The latter parameter is the maximum number of units in an RGNG. This number refers only to the number of “direct” units in a particular RGNG and does not include potential units present in RGNGs that are prototypes of these direct units.

Like its structure the behavior of the RGNG is basically identical to that of a regular GNG. However, since the prototypes of the units can either be vectors or RGNGs themselves, the behavior is now defined by four functions. The distance function

$$D(x, y) : W \times W \rightarrow \mathbb{R}$$

determines the distance either between two vectors, two RGNGs, or a vector and an RGNG. The interpolation function

$$I(x, y) : (\mathbb{R}^n \times \mathbb{R}^n) \cup (G \times G) \rightarrow W$$

generates a new vector or new RGNG by interpolating between two vectors or two RGNGs, respectively. The adaptation function

$$A(x, \xi, r) : W \times \mathbb{R}^n \times \mathbb{R} \rightarrow W$$

adapts either a vector or RGNG towards the input vector ξ by a given fraction r . Finally, the input function

$$F(g, \xi) : G \times \mathbb{R}^n \rightarrow G \times \mathbb{R}$$

feeds an input vector ξ into the RGNG g and returns the modified RGNG as well as the distance between ξ and the best matching unit (BMU, see below) of g . The input function F contains the core of the RGNG’s behavior and utilizes the other three functions, but is also used, in turn, by those functions introducing several recursive paths to the program flow.

$F(g, \xi)$: The input function F is a generalized version of the original GNG algorithm that facilitates the use of prototypes other than vectors. In particular, it allows to use RGNGs themselves as prototypes resulting in a recursive structure. An input $\xi \in \mathbb{R}^n$ to the RGNG g is processed by the input function F as follows:

- Find the two units s_1 and s_2 with the smallest distance to the input ξ according to the distance function D :

$$s_1 := \arg \min_{u \in g.U} D(u.w, \xi),$$

$$s_2 := \arg \min_{u \in g.U \setminus \{s_1\}} D(u.w, \xi).$$

- Increment the age of all edges connected to s_1 :

$$\Delta c.t = 1, \quad c \in g.C \wedge s_1 \in c.V.$$

- If no edge between s_1 and s_2 exists, create one:

$$g.C \Leftarrow g.C \cup \{(\{s_1, s_2\}, 0)\}.$$

- Reset the age of the edge between s_1 and s_2 to zero:

$$c.t \Leftarrow 0, \quad c \in g.C \wedge s_1, s_2 \in c.V.$$

- Add the squared distance between ξ and the prototype of s_1 to the accumulated error of s_1 :

$$\Delta s_1.e = D(s_1.w, \xi)^2.$$

- Adapt the prototype of s_1 and all prototypes of its direct neighbors:

$$s_1.w \Leftarrow A(s_1.w, \xi, g.\theta.\epsilon_b),$$

$$s_n.w \Leftarrow A(s_n.w, \xi, g.\theta.\epsilon_n), \quad \forall s_n \in E_{s_1}.$$

- Remove all edges with an age above a given threshold τ and remove all units that no longer have any edges connected to them:

$$g.C \Leftarrow g.C \setminus \{c | c \in g.C \wedge c.t > g.\theta.\tau\},$$

$$g.U \Leftarrow g.U \setminus \{u | u \in g.U \wedge E_u = \emptyset\}.$$

- If an integer-multiple of $g.\theta.\lambda$ inputs was presented to the RGNG g and $|g.U| < g.\theta.M$, add a new unit u . The new unit is inserted “between” the unit j with the largest accumulated error and the unit k with the largest accumulated error among the direct neighbors of j . Thus, the prototype $u.w$ of the new unit is initialized as:

$$u.w := I(j.w, k.w), \quad j = \arg \max_{l \in g.U} (l.e),$$

$$k = \arg \max_{l \in E_j} (l.e).$$

The existing edge between units j and k is removed and edges between units j and u as well as units u and k are added:

$$\begin{aligned} g \bullet C &\Leftarrow g \bullet C \setminus \{c \mid c \in g \bullet C \wedge j, k \in c \bullet V\}, \\ g \bullet C &\Leftarrow g \bullet C \cup \{(\{j, u\}, 0), (\{u, k\}, 0)\}. \end{aligned}$$

The accumulated errors of units j and k are decreased and the accumulated error $u \bullet e$ of the new unit is set to the decreased accumulated error of unit j :

$$\begin{aligned} \Delta j \bullet e &= -g \bullet \theta \bullet \alpha \cdot j \bullet e, & \Delta k \bullet e &= -g \bullet \theta \bullet \alpha \cdot k \bullet e, \\ u \bullet e &:= j \bullet e. \end{aligned}$$

– Finally, decrease the accumulated error of all units:

$$\Delta u \bullet e = -g \bullet \theta \bullet \beta \cdot u \bullet e, \quad \forall u \in g \bullet U.$$

The function F returns the tuple (g, d_{\min}) containing the now updated RGNG g and the distance $d_{\min} := D(s_1 \bullet w, \xi)$ between the prototype of unit s_1 and input ξ . Note that in contrast to the regular GNG there is no stopping criterion any more, i.e., the RGNG operates explicitly in an online fashion by continuously integrating new inputs. To prevent unbounded growth of the RGNG the maximum number of units $\theta \bullet M$ was introduced to the set of parameters.

$D(x, y)$: The distance function D determines the distance between two prototypes x and y . The calculation of the actual distance depends on whether x and y are both vectors, a combination of vector and RGNG, or both RGNGs:

$$D(x, y) := \begin{cases} D_{RR}(x, y) & \text{if } x, y \in \mathbb{R}^n, \\ D_{GR}(x, y) & \text{if } x \in G \wedge y \in \mathbb{R}^n, \\ D_{RG}(x, y) & \text{if } x \in \mathbb{R}^n \wedge y \in G, \\ D_{GG}(x, y) & \text{if } x, y \in G. \end{cases}$$

In case the arguments of D are both vectors, the Minkowski distance is used:

$$\begin{aligned} D_{RR}(x, y) &:= (\sum_{i=1}^n |x_i - y_i|^p)^{\frac{1}{p}}, \quad x = (x_1, \dots, x_n), \\ & \quad y = (y_1, \dots, y_n), \\ & \quad p \in \mathbb{N}. \end{aligned}$$

Using the Minkowski distance instead of the Euclidean distance allows to adjust the distance measure with respect to certain types of inputs via the parameter p . For example, setting p to higher values results in an emphasis of large changes in individual dimensions of the input vector versus changes that are distributed over many dimensions [13]. However, in the case of modeling the behavior of grid cells the parameter is set to a fixed value of 2 which makes the Minkowski distance equivalent to the Euclidean distance. The latter is required in this

context as only the Euclidean distance allows the GNG to form an induced Delaunay triangulation of its input space.

In case the arguments of D are a combination of vector and RGNG, the vector is fed into the RGNG using function F and the returned minimum distance is taken as distance value:

$$\begin{aligned} D_{GR}(x, y) &:= F(x, y) \cdot d_{\min}, \\ D_{RG}(x, y) &:= D_{GR}(y, x). \end{aligned}$$

In case the arguments of D are both RGNGs, the distance is defined to be the pairwise minimum distance between the prototypes of the RGNGs' units, i.e., *single linkage* distance between the sets of units is used:

$$D_{GG}(x, y) := \min_{u \in x \cdot U, k \in y \cdot U} D(u \cdot w, k \cdot w).$$

The latter case is used by the interpolation function if the recursive depth of an RGNG is at least 2. As the RGNG-based grid cell model has only a recursive depth of 1 (see next section), the case is considered for reasons of completeness rather than necessity. Alternative measures to consider could be, e.g., *average* or *complete* linkage.

$I(x, y)$: The interpolation function I returns a new prototype as a result from interpolating between the prototypes x and y . The type of interpolation depends on whether the arguments are both vectors or both RGNGs:

$$I(x, y) := \begin{cases} I_{RR}(x, y) & \text{if } x, y \in \mathbb{R}^n, \\ I_{GG}(x, y) & \text{if } x, y \in G. \end{cases}$$

In case the arguments of I are both vectors, the resulting prototype is the arithmetic mean of the arguments:

$$I_{RR}(x, y) := \frac{x + y}{2}.$$

In case the arguments of I are both RGNGs, the resulting prototype is a new RGNG a . Assuming w.l.o.g. that $|x \cdot U| \geq |y \cdot U|$ the components of the interpolated RGNG a are defined as follows:

$$\begin{aligned} a &:= I(x, y), \\ a \cdot U &:= \left\{ (w, 0) \left| \begin{array}{l} w = I(u \cdot w, k \cdot w), \\ \forall u \in x \cdot U, \\ k = \arg \min_{l \in y \cdot U} D(u \cdot w, l \cdot w) \end{array} \right. \right\}, \\ a \cdot C &:= \left\{ (\{l, m\}, 0) \left| \begin{array}{l} \exists c \in x \cdot C \\ \wedge u, k \in c \cdot V \\ \wedge l \cdot w = I(u \cdot w, \cdot) \\ \wedge m \cdot w = I(k \cdot w, \cdot) \end{array} \right. \right\}, \\ a \cdot \theta &:= x \cdot \theta. \end{aligned}$$

The resulting RGNG a has the same number of units as RGNG x . Each unit of a has a prototype that was interpolated between the prototype of the corresponding unit in x and the nearest prototype found in the units of y . The edges and parameters of a correspond to the edges and parameters of x .

$\mathbf{A}(x, \xi, r)$: The adaptation function A adapts a prototype x towards a vector ξ by a given fraction r . The type of adaptation depends on whether the given prototype is a vector or an RGNG:

$$A(x, \xi, r) := \begin{cases} A_R(x, \xi, r) & \text{if } x \in \mathbb{R}^n, \\ A_G(x, \xi, r) & \text{if } x \in G. \end{cases}$$

In case prototype x is a vector, the adaptation is performed as linear interpolation:

$$A_R(x, \xi, r) := (1 - r)x + r\xi.$$

In case prototype x is an RGNG, the adaptation is performed by feeding ξ into the RGNG. Importantly, the parameters ϵ_b and ϵ_n of the RGNG are temporarily changed to take the fraction r into account:

$$\theta^* := (r, r \cdot x.\theta.\epsilon_r, x.\theta.\epsilon_r, x.\theta.\lambda, x.\theta.\tau, \\ x.\theta.\alpha, x.\theta.\beta, x.\theta.M),$$

$$x^* := (x.U, x.C, \theta^*),$$

$$A_G(x, \xi, r) := F(x^*, \xi).x.$$

Note that in this case the new parameter $\theta.\epsilon_r$ is used to derive a temporary ϵ_n from the fraction r .

This concludes the formal definition of the RGNG algorithm.

Activity Approximation

The RGNG-based model describes a group of neurons for which we would like to derive their “activity” for any given input as a scalar that represents the momentary firing rate of the particular neuron. Yet, the RGNG algorithm itself does not provide a direct measure that could be used to this end. Therefore, we derive the activity a_u of a modelled neuron u based on the neuron’s best and second best matching BL units s_1 and s_2 with respect to a given input ξ as:

$$a_u := e^{-\frac{(1-r)^2}{2\sigma^2}},$$

with $\sigma = 0.2$ and ratio r :

$$r := \frac{D(s_2.w, \xi) - D(s_1.w, \xi)}{D(s_1.w, s_2.w)}, \quad s_1, s_2 \in u.w.U,$$

using a distance function D . This measure of activity allows to correlate the response of a neuron to a given input with further variables.

Table 1. Parameters of the RGNG-based model used for the simulation runs in our previous work [15]. Parameters θ_1 control the top layer RGNG while parameters θ_2 control all bottom layer RGNGs of the model.

θ_1	θ_2
$\epsilon_b = 0.04$	$\epsilon_b = 0.01$
$\epsilon_n = 0.04$	$\epsilon_n = 0.0001$
$\epsilon_r = 0.01$	$\epsilon_r = 0.01$
$\lambda = 1000$	$\lambda = 1000$
$\tau = 300$	$\tau = 300$
$\alpha = 0.5$	$\alpha = 0.5$
$\beta = 0.0005$	$\beta = 0.0005$
$M = 100$	$M = \{20, 80\}$

Table 2. Parameters of the RGNG-based model used for the simulation runs presented in this paper (section 4).

θ_1	θ_2
$\epsilon_b = 0.004$	$\epsilon_b = 0.001$
$\epsilon_n = 0.004$	$\epsilon_n = 0.00001$
$\epsilon_r = 0.01$	$\epsilon_r = 0.01$
$\lambda = 1000$	$\lambda = 1000$
$\tau = 300$	$\tau = 300$
$\alpha = 0.5$	$\alpha = 0.5$
$\beta = 0.0005$	$\beta = 0.0005$
$M = 100$	$M = \{20, 40, 80, 160\}$

Parameterization

Each layer of an RGNG requires its own set of parameters. In case of our two-layered grid cell model we use the sets of parameters θ_1 and θ_2 , respectively. Parameter set θ_1 controls the main top layer RGNG while parameter set θ_2 controls all bottom layer RGNGs. Table 1 summarizes the parameter values used for the simulation runs presented in our previous work [15], while table 2 contains the parameters of the simulation runs presented in this paper. For a detailed characterization of these parameters we refer to Kerdels [12].

References

1. Barry, C., Burgess, N.: Neural mechanisms of self-location. *Current Biology* 24(8), R330 – R339 (2014)
2. Barry, C., Ginzberg, L.L., OKeefe, J., Burgess, N.: Grid cell firing patterns signal environmental novelty by expansion. *Proceedings of the National Academy of Sciences* 109(43), 17687–17692 (2012)
3. Barry, C., Hayman, R., Burgess, N., Jeffery, K.J.: Experience-dependent rescaling of entorhinal grids. *Nat Neurosci* 10(6), 682–684 (jun 2007)

4. Boccara, C.N., Sargolini, F., Thoresen, V.H., Solstad, T., Witter, M.P., Moser, E.I., Moser, M.B.: Grid cells in pre- and parasubiculum. *Nat Neurosci* 13(8), 987–994 (Aug 2010)
5. Burak, Y.: Spatial coding and attractor dynamics of grid cells in the entorhinal cortex. *Current Opinion in Neurobiology* 25(0), 169 – 175 (2014), theoretical and computational neuroscience
6. Domnisoru, C., Kinkhabwala, A.A., Tank, D.W.: Membrane potential dynamics of grid cells. *Nature* 495(7440), 199–204 (Mar 2013)
7. Fritzke, B.: A growing neural gas network learns topologies. In: *Advances in Neural Information Processing Systems* 7. pp. 625–632. MIT Press (1995)
8. Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., Moser, M.B.: Spatial representation in the entorhinal cortex. *Science* 305(5688), 1258–1264 (2004)
9. Giocomo, L., Moser, M.B., Moser, E.: Computational models of grid cells. *Neuron* 71(4), 589 – 603 (2011)
10. Hafting, T., Fyhn, M., Molden, S., Moser, M.B., Moser, E.I.: Microstructure of a spatial map in the entorhinal cortex. *Nature* 436(7052), 801–806 (Aug 2005)
11. Jacobs, J., Weidemann, C.T., Miller, J.F., Solway, A., Burke, J.F., Wei, X.X., Suthana, N., Sperling, M.R., Sharan, A.D., Fried, I., Kahana, M.J.: Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat Neurosci* 16(9), 1188–1190 (Sep 2013)
12. Kerdels, J.: A Computational Model of Grid Cells based on a Recursive Growing Neural Gas. Ph.D. thesis, FernUniversität in Hagen, Hagen (2016)
13. Kerdels, J., Peters, G.: Analysis of high-dimensional data using local input space histograms. *Neurocomputing* 169, 272 – 280 (2015)
14. Kerdels, J., Peters, G.: A new view on grid cells beyond the cognitive map hypothesis. In: 8th Conference on Artificial General Intelligence (AGI 2015) (July 2015)
15. Kerdels, J., Peters, G.: Modelling the grid-like encoding of visual space in primates. In: *Proceedings of the 8th International Joint Conference on Computational Intelligence, IJCCI 2016, Volume 3: NCTA, Porto, Portugal, November 9–11, 2016*. pp. 42–49 (2016)
16. Killian, N.J., Jutras, M.J., Buffalo, E.A.: A map of visual space in the primate entorhinal cortex. *Nature* 491(7426), 761–764 (Nov 2012)
17. Krupic, J., Bauza, M., Burton, S., Barry, C., O’Keefe, J.: Grid cell symmetry is shaped by environmental geometry. *Nature* 518(7538), 232–235 (Feb 2015), <http://dx.doi.org/10.1038/nature14153>
18. Lundgaard, I., Li, B., Xie, L., Kang, H., Sanggaard, S., Haswell, J.D.R., Sun, W., Goldman, S., Blekot, S., Nielsen, M., Takano, T., Deane, R., Nedergaard, M.: Direct neuronal glucose uptake heralds activity-dependent increases in cerebral metabolism. *Nature Communications* 6, 6807 (Apr 2015)
19. Moser, E.I., Moser, M.B.: A metric for space. *Hippocampus* 18(12), 1142–1156 (2008)
20. Moser, E.I., Moser, M.B., Roudi, Y.: Network mechanisms of grid cells. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369(1635) (2014)
21. Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.B., Moser, E.I.: Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312(5774), 758–762 (2006)
22. Stensola, H., Stensola, T., Solstad, T., Froland, K., Moser, M.B., Moser, E.I.: The entorhinal grid map is discretized. *Nature* 492(7427), 72–78 (Dec 2012)
23. Stensola, T., Stensola, H., Moser, M.B., Moser, E.I.: Shearing-induced asymmetry in entorhinal grid cells. *Nature* 518(7538), 207–212 (Feb 2015)

24. Welinder, P.E., Burak, Y., Fiete, I.R.: Grid cells: The position code, neural network models of activity, and the problem of learning. *Hippocampus* 18(12), 1283–1300 (2008)
25. Yartsev, M.M., Witter, M.P., Ulanovsky, N.: Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature* 479(7371), 103–107 (Nov 2011)