

Can Grid Cell Ensembles Represent Multiple Spaces?

Davide Spalla

dspalla@sissa.it

SISSA, Cognitive Neuroscience, 34136 Trieste, Italy

Alexis Dubreuil

alexis.dubreuil@lpt.ens.fr

Laboratoire de Physique Théorique de l'ENS, 75231 Paris Cedex 05, France

Sophie Rosay

sophie.rosay@gmail.com

SISSA, Cognitive Neuroscience, 34136 Trieste, Italy

Remi Monasson

monasson@lpt.ens.fr

Laboratoire de Physique Théorique de l'ENS, 75231 Paris Cedex 05, France

Alessandro Treves

ale@sissa.it

SISSA, Cognitive Neuroscience, 34136 Trieste, Italy

The way grid cells represent space in the rodent brain has been a striking discovery, with theoretical implications still unclear. Unlike hippocampal place cells, which are known to encode multiple, environment-dependent spatial maps, grid cells have been widely believed to encode space through a single low-dimensional manifold, in which coactivity relations between different neurons are preserved when the environment is changed. Does it have to be so? Here, we compute, using two alternative mathematical models, the storage capacity of a population of grid-like units, embedded in a continuous attractor neural network, for multiple spatial maps. We show that distinct representations of multiple environments can coexist, as existing models for grid cells have the potential to express several sets of hexagonal grid patterns, challenging the view of a universal grid map. This suggests that a population of grid cells can encode multiple noncongruent metric relationships, a feature that could in principle allow a grid-like code to represent environments with a variety of different geometries and possibly conceptual and cognitive

D.S. and A.D. contributed equally to this work.

spaces, which may be expected to entail such context-dependent metric relationships.

1 Introduction

Grid cells appear to comprise an essential component of the cognitive representation of space in rodents (Hafting, Fyhn, Molden, Moser, & Moser, 2005) and other species (e.g., bats; Yartsev, Witter, & Ulanovsky, 2011). Located in the medial entorhinal cortex, these neurons are selectively active when the animal is in certain positions of the environment, the so-called fields, at the vertices of a remarkably regular hexagonal lattice. A study of the activity of grid cells in multiple environments (Fyhn, Hafting, Treves, Moser, & Moser, 2007) has shown that while the grid expressed by each neuron varies across environments in its spatial phase and orientation, between neurons the coactivity relations are largely preserved, at least for those recorded nearby in the tissue, with the same tetrad. In other words, the grids of different cells undergo a coherent rigid movement when a new environment is explored, as illustrated schematically in Figures 1A and 1B (bottom row). The subsequent discovery of quasi-discrete “modules” (Stensola et al., 2012) indicates that these relations are maintained at the local network level, presumably by recurrent collateral connections among grid cells. This finding has led to the hypothesis that local ensembles of grid cells each comprise a single continuous attractor network, expressing a universal, two-dimensional map, which encodes the metric of space independent of the environmental context. There is a crucial difference with the context-dependent spatial representations provided by hippocampal place cells, which display “global remapping” (Leutgeb et al., 2005) even between very similar rooms, in particular in the CA3 field (Alme et al., 2014): cells that were silent acquire one or more place fields, others lose theirs, and the fields that seem to have been maintained typically are in a different location (see Figure 1, top row). Global remapping has motivated the conceptual model of multiple charts (Samsonovich & McNaughton, 1997), in contrast with early and later models of continuous attractor grid cell networks, which envisage a single chart (Fuhs & Touretzky, 2006; Burak & Fiete, 2009; Couey et al., 2013). The dominant overall view, then, holds that the hippocampus encodes multiple, uncorrelated, context-dependent cognitive maps, while the grid system provides metric information that is independent of the environment. Recent evidence of context-dependent distortions in the grid pattern has begun to question the view that the collective map expressed by a grid module is universal, that is, it applies to any environment. Stensola, Stensola, Moser, and Moser (2015) have shown that when rats explore large environments, a single grid can exhibit multiple orientations, likely due to anchoring effects to the closest wall, which in any case amount to distortions of the hexagonal pattern. These effects have been

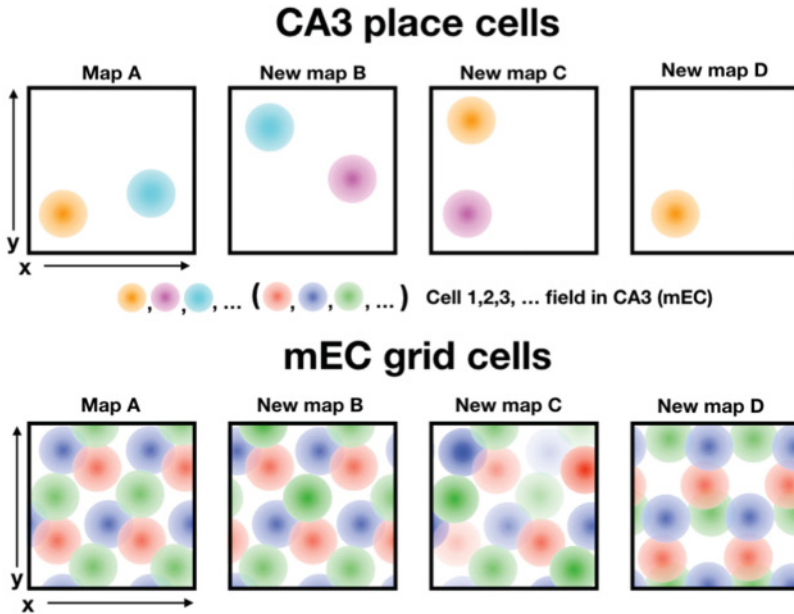


Figure 1: Types of change in grid cell activity in mEC (bottom) concurrent with global remapping in the CA3 field of the hippocampus (top). The universal grid map model, idealized from Fyhn et al. (2007), allows only for a coherent translation (and possibly a rotation) into a new map B, when changing environment. Under a manipulation that does not entail changing environment, the individual fields of each unit have been observed to independently vary their peak rates, keeping their relative position (Kanter et al., 2017; new map C). We assess the hypothesis that the same network may also express other maps, such as map D, with a complete repositioning of the grids of different units.

analyzed extensively in a more recent study (Häggglund, Mørreaunet, Moser, & Moser, 2019). Krupic et al. (Krupic, Bauza, Burton, Barry, & O'Keefe, 2015; Krupic, Bauza, Burton, & O'Keefe, 2018) have shown that the grid pattern deviates from perfect hexagonality, with both global and local distortions, in response to environmental features such as the geometry of the walls. Finally, a couple of recent studies (Boccarda, Nardin, Stella, O'Neill, & Csicsvari, 2019; Butler, Hardcastle, & Giocomo, 2019) have shown that the presence of salient features such as goals or rewards affects the entorhinal map, changing field locations and inducing remapping in other space-selective cells. These observations, moreover, refer solely to the position of the peaks of activity, that is, the place fields of each cell, and do not take into account the fact that they vary reliably in height, independently across peaks, from one environment to the other (Dunn,

Wennberg, Huang, & Roudi, 2017). Should we still regard grid cells as a sort of stack of millimeter paper, providing a universal metric for space?

Recent studies conducted in both rodents and humans, moreover, suggest that regular grids may not measure only physical space. Aronov, Nevers, and Tank (2017) find that both place cells and grid cells, in rats, are involved in the representation of a nonspatial but continuous, one-dimensional variable, such as the frequency of a sound. An fMRI study by Constantinescu, O'Reilly, and Behrens (2016) shows a hexagonal modulation of the BOLD signal in human entorhinal cortex, and elsewhere, in a task that requires subjects to navigate the 2D space spanned by the varying leg and neck lengths of a drawing of a bird. The representation of abstract or conceptual spaces, which could in principle be topologically and geometrically complex, would require of the grid cell system a flexibility that can hardly be reconciled with the universal grid hypothesis.

In a most interesting study (Kanter et al., 2017), a subset of grid units were depolarized in transgenic mice, leading to what appears to be global remapping in the hippocampus. What is so striking is that the manipulation induces extensive changes, up and down, in the peak firing rates of the different fields of individual grid units but not in their position. This elaborates the observation in Fyhn et al. (2007) and suggests that what might be universal in the grid representation expressed by an ensemble of units, if anything, are the relative positions of the fields, whereas their peak firing rates are variable (see Figure 1C, bottom row). On the other hand, a strict hexagonal periodicity of the field positions of individual units is possible only in flat 2D environments. The adaptation model of grid formation (Kropff & Treves, 2008) predicts instead, on surfaces with constant positive or negative gaussian curvature and appropriate radius, the emergence of grids with, for example, pentagonal (Stella, Si, Kropff, & Treves, 2013) or heptagonal (Urdapilleta, Troiani, Stella, & Treves, 2015) symmetry. In all other cases, including ecologically plausible natural environments, nonflat surfaces have varying curvature, making strictly periodic grids dubious and rigid phase coherence most unlikely. But what happens to the universality of the grid in natural environments?

To address these issues, the aim of the work we present here is to answer a first fundamental question: Is it at all possible to conceive of multiple, hence nonuniversal, ideal grid representations expressed in the same local network when the animal is placed in distinct, even if flat, environments? In other words, would the storage capacity of a recurrent network of grid cells be above unity, so that multiple continuous attractors can coexist, encoded in the same synaptic efficacies? We pose this question within two alternative mathematical models, both accepting the idealized assumptions that underlie the universal map hypothesis, that is, of strict periodicity and equal peak rates, depicted in Figure 1D (bottom row), but allowing for several uncorrelated grid representations. Under these assumptions, we analyze an ensemble of grid cells as a continuous attractor neural network, extending the

frameworks developed in Battaglia and Treves (1998) and Monasson and Rosay (2013, 2014) for the description of place cells. We emphasize that the storage capacity we are interested in quantifies the number of different, independent charts (or collective maps) that the network can store, and not the spatial resolution (which may be referred to as information capacity, i.e., the number of different positions that can be decoded from the ensemble activity), as considered, for example, in Hedrick and Zhang (2016) and Fiete, Burak, and Brookings (2008).

2 Complementary Network Models

We model the grid cell population as an ensemble of units interacting through recurrent connections, whose structure defines which activity states are robust—the dynamical attractors. We assume, however, that a separate process, based, for example, on adaptation (Kropff & Treves, 2008), has determined the emergence of a periodic grid, independently for each unit, during familiarization with each of p distinct environments; meanwhile, recurrent connections are shaped by a Hebbian learning process, such that neurons that happen to have nearby fields tend to fire together, strengthening their connections, while neurons with fields far apart remain weakly connected. The connection strength J_{ij} is therefore taken to be a sum of contributions from the exploration of p environments, with each contribution, once averaged across many trajectories, a function of the relative position of the fields in that environment. Exploiting the simplifying assumption that each grid is strictly periodic, we can focus on the elementary repetitive tile, which has only one field per unit and is, in the mathematical formulation, connected by periodic boundary conditions to adjacent tiles. The assumption of periodic boundary conditions is motivated by the remarkable regularity of the arrangement of the fields observed in the original experiments, and by the model being meant to describe interactions within a grid module, in which all cells share the same spacing and orientation. The contribution to the connection strength between two units i and j is then reduced to a function of their field centers \bar{x}_i^π and \bar{x}_j^π on the elementary tile in environment π ,

$$J_{ij} = \sum_{\pi=1}^p K(\bar{x}_i^\pi, \bar{x}_j^\pi), \quad (2.1)$$

where we refer to $K(\cdot)$ as the interaction kernel. The field peaks, or centers \bar{x}_i of N units, are taken to be randomly and uniformly distributed over the elementary tile. Our analysis focuses on two different models of neurons, binary and threshold-linear, and two types of attractor symmetry, square and hexagonal, which stem from the tile shape or the interaction kernel. Both neuron models allow, from complementary angles, a full statistical

analysis, leading to otherwise inaccessible results. The storage capacity turns out to depend more on how interference reverberates through loops (expressed by the parameter ψ ; see below) than on the type of units; interference, in the densely coded and densely connected regime, affects square much more than hexagonal grids.

2.1 Binary Units. The first model we consider is an extension of the model proposed by Monasson and Rosay (2013) for the modeling of place cells in CA3. Here the activity of neurons is described by binary variables, such that the pattern of activity of a network of N units is a vertex $\{\sigma\} \in \{0, 1\}^N$. For the binary model, the kernel $K(i, j)$ between units i and j relative to one environment is taken to be a step function of the distance between their field centers,

$$K(\cdot) = \frac{1}{N} \Theta(d_c - |\vec{x}_i - \vec{x}_j|), \quad (2.2)$$

where $\Theta(x) = 1$ for $x > 0$ and 0 otherwise; note that the distance $|\vec{x}_i - \vec{x}_j|$ is along the shortest path, considering the periodic boundary conditions. The periodic structure of the attractor depends on the shape of the rhomboid unitary tile in which the field center \vec{x}_i of each unit is located. The lattice symmetry is specified by the angle θ between its two primitive vectors. $\theta = 60^\circ$ corresponds to the standard case of hexagonal grids, while $\theta = 90^\circ$ describes a square grid pattern. These two cases and the resulting interaction kernel are depicted in Figures 2a and 2b. The cutoff distance d_c sets the number of nonzero connections each unit receives from the storage of a given environment, denoted by wN : $d_c = \sqrt{(w/\pi)} \sin \theta$. This measure of connectivity within one environment should not be confused with the global connectivity, taking into account all environments, $C = (N - 1)(1 - (1 - w)^p) \sim N$ for large p .

The dynamics of the network is governed by the energy function

$$E_j[\{\sigma\}] = - \sum_{i < j} J_{ij} \sigma_i \sigma_j \quad (2.3)$$

and constrained by the requirement that a fixed fraction f of units be in the active state, $\sum_i \sigma_i = fN$, at any time. We call f the coding level, or sparsity of the representation. This constraint is taken to reflect some form of global inhibition. Later we focus, given w , only on the optimal coding level in terms of storage capacity, hence on a specific value $f^*(w)$, which turns out to be a monotonic function of w (see Figure 3). This model then allows an explicit focus on the dependence of the storage capacity on the width of the kernel and the resulting optimal sparsity of the representation.




Figure 2: Interaction kernels for the binary (a, b) and rate (c, d) models. The white lines show the elementary tile of each lattice.

2.2 Threshold-Linear Units. We extend our analysis to firing rate units, whose activity is described by a continuous positive value corresponding to their instantaneous firing rate. This second model allows us to capture the graded nature of neural activity, which is salient when it represents space, which is itself continuous. The activity of the network is given by a configuration $\{V_i\} \in (\mathbb{R}^+)^N$, and each unit integrates the inputs it receives through a threshold-linear transfer function (Treves, 1990b)

$$V_i = \begin{cases} g(h_i - h_0) & \text{if } h_i \geq h_0 \\ 0 & \text{if } h_i \leq h_0 \end{cases}, \quad (2.4)$$

where g (the linear gain) and h_0 (the activation threshold) are global parameters of the network, and the local field h_i is a real-valued variable summarizing the input influence on unit i from the rest of the network, which we take to come from a random but fixed set of C among the $N - 1$ other units, as well as from external sources. The interaction kernel $K(\cdot)$ is given by the special sum-of-cosines form,

$$K(\cdot) = \frac{1}{C} \sum_{l=1}^d (\cos[\phi^l(\vec{x}_i) - \phi^l(\vec{x}_j)] + 1), \quad (2.5)$$

which had been considered a toy case by Battaglia and Treves (1998) before the discovery of grid cells. The field center of each unit on the elementary tile is expressed by a set of angles $\phi^l(\vec{x})$. We shall see that $d = 2$ and 3 are equally valid choices on the plane, as well as $d = 1$, which leads to band solutions (see below). This model therefore allows decoupling the form of the kernel, which is extended, from interactions among units far away on the elementary tile (and the resulting coding level is correspondingly nonsparse) from the connectivity, which can be made arbitrarily sparse if $C/N \rightarrow 0$. As a superposition of d cosine functions, the kernel can also be conveniently written as a sum of dot products. The $+1$ term is added to enforce excitatory connections. While not circularly symmetric like the radial kernel used in the binary model, this cosine kernel allows for the analytical study of periodic patterns that are spread out on the plane, with a large fraction of the units active at any given time. The solutions for the hexagonal kernel (see Figure 2d), in particular, given by three cosine functions at a 60° angle from one another, have been considered a reasonable model for experimentally observed grid cells. In the figure, the hexagonal elementary tile extends in the range $x = \pm 1/2$ and $y = \pm 1/\sqrt{3}$, and the three angles span the directions $\phi^1 = 2\pi x$, $\phi^{2,3} = \pi(x \pm \sqrt{3}y)$. The square kernel is obtained for $d = 2$ and the two cosines at 90° from each other (Figure 2c). Note that as with the binary model, N units are concentrated on an elementary tile that in the hexagonal case is $\sqrt{3}/2$ of the area of the square case.

An energy function would look similar to the one in equation 2.3 but now expressed in terms of the continuous variables $\{V\}$. When $C < N - 1$ and connections are not symmetric, the energy formalism does not apply, but we can still analyze the model (see below and in appendix B), and again we take global inhibition, which can now also act through a modulation of the common gain g to depend on the average activity of the network and to be such as to optimize storage capacity.

3 Storage Capacity

Both models can store a single population map, as in the bottom panels of Figure 1, A and B, and the equations for such a map admit periodic bump solutions that reproduce the shape of the tile or kernel (as well as potentially other solutions, such as stripes, discussed later). We are interested, however, in their capacity to store several distinct maps, as in panels A and D of Figure 1, and in the possibility of calculating such storage capacity analytically, in the mean field approximation. The general strategy involves formulating and resolving a set of self-consistent equations relating the activity of the units in the network. When the model admits an energy function, these are the saddle point equations derived from the computation of the free energy of the system with the replica trick, which allows taking into account the statistics of the field centers in each environment. Without an energy

function (e.g., when the connections are sparse and not symmetric), equivalent equations can be derived through so-called self-consistent signal-to-noise analysis (Roudi & Treves, 2004). The solutions to these equations, which describe the activity in one map, disappear sharply at a critical value α_c of the storage load $\alpha = (p/C)$, which measures the ratio of the number of maps to the number of connections to each unit. α_c then gives the maximum number of maps that the network can store and retrieve or express, normalized by the connectivity. Crucially, we have developed a novel method to assess whether below α_c , these solutions are indeed stable and prevail on others (see Figures 6 and 7). The details of these methods, which build on Monasson and Rosay (2013, 2014) for the binary model and on Treves (1990a) and Battaglia and Treves (1998) for the rate model, can be found in appendix B. We focus, in the calculation of the storage capacity, on so-called bump states, in which activity is localized along each of the two dimensions of the elementary tile (anywhere on the tile, given the translation invariance of the interaction kernel). Other solutions, however, do exist, as discussed in section 4.

3.1 Binary Units. The statistical analysis of the minima of the free energy leads to the patterns of activity $\{\sigma\}$ that are likely to be observed given the connectivity. More precisely, we have derived self-consistent equations for the average activity $\rho(\vec{x}) = \langle \sigma_i \rangle$ of unit i having its grid field centered in \vec{x} (in the elementary tile),

$$\rho(\vec{x}) = \int dz \frac{e^{-z^2/(2\alpha r)}}{\sqrt{2\pi \alpha r}} \Theta[\mu(\vec{x}) + z + \lambda], \quad (3.1)$$

where

$$\mu(\vec{x}) = \int d\vec{y} K(\vec{x}, \vec{y}) \rho(\vec{y}) \quad (3.2)$$

is the signal input received by the unit through the interactions corresponding to the environment in which the bump is localized, say, $\pi = 1$, and z is the noisy, gaussian input due to the interference from the other environments, say, $\pi = 2, \dots, p$ (see equation 2.1). The variance αr of these gaussian inputs is self-consistently derived from the knowledge of the activity profile ρ (see appendix A). The uniform (inhibitory) input λ enforces the constraint $\int d\vec{x} \rho(\vec{x}) = f$. We have considered the limit case of neurons responding deterministically to their inputs, although the analysis extends naturally to stochastic noise.

We calculate, from the saddle point equations, the storage capacity $\alpha_c(w, f)$ as the maximal value of the load α for which a bump-like solution to equation 3.1 exists. Then, for a given value of w , we find the coding level $f^*(w)$ that maximizes the storage capacity. Over a broad range




Figure 3: Optimal coding level for the binary model.




Figure 4: Storage capacity as a function of w for square and hexagonal grids in the binary model, given an optimal coding level $f^* \simeq w/2$.

$0 \leq w \leq 0.5$, the optimal f^* turns out to be approximately half the value of w (see Figure 3). That the optimal value for the coding level is proportional to w can be understood intuitively by considering the spatial profile of the signal $\mu(\vec{x})$: if too few cells are allowed to be active, the connections to the cells that are forced to be silent, within the connectivity range of the active cells, will be frustrated. But if too many cells are active, those outside the connectivity range will contribute more to the noise than to the signal. This optimal storage capacity is plotted in Figure 4 for the square and hexagonal grids as a function of w . At low w , the two values are similar, but when w increases, their trends diverge, leading to a substantially higher-capacity value in the hexagonal case, of order 10^{-2} for $w \simeq 0.5$. This value would definitely allow, in a real cortical network with order thousands (or tens of

thousands) of neurons, the storage and retrieval of multiple independent grid maps. Again, considering the spatial profiles of the signal $\mu(\vec{x})$ allows gaining intuition about this divergence. At very low w , that is, short-range interactions, what happens in other tiles can be neglected, and the two grids behave similarly. When the range is wider, the location of the fields in the immediately neighboring tiles starts to be relevant. In the square case, there are four first neighbors that contribute to exciting silent neurons in between the fields. For a hexagonal arrangement of the fields, there are six neighboring tiles that each contribute relatively less excitation in between fields. Intuitively, this last geometrical arrangement makes the structure more rigid and reduces the influence of the noise due to the storage of other charts.

3.2 Threshold-Linear Units. In this model, the coding level and the connectivity range are both fixed by the shape of $K(\cdot)$. The mean field approach, however, can be extended to the case of arbitrary values of the connectivity density C/N , with the self-consistent signal-to-noise analysis (Roudi & Treves, 2004). The storage capacity is given by the α for which the solution to the equation

$$\bar{\mu}^2 - d \left[1 + \frac{C}{N} \left(\frac{(2 - \psi)\psi}{(1 - \psi)^2} \right) \right] \alpha r = 0 \quad (3.3)$$

disappears. In fact, the disappearance of the solution gives only an upper bound on α_c ; one has to check its stability as well. The details of the derivation and the expression of the average signal $\bar{\mu}$ and of the interference noise r are reported in appendix B. We plot critical values for square and hexagonal grids with the respective kernels, as a function of the inverse density N/C , in Figure 5 (full curves, blue and red). In the fully connected regime, we find a result, also corroborated by computer simulations, similar to the one obtained with the binary model; however, there is a huge difference in capacity between square and hexagonal grids and a value $\sim 10^{-2}$ only for the latter. Moreover, it turns out that for the square kernel, the stripe or band solutions of the next section are the global minima, and the square solutions are only marginally stable. In all cases, the capacity increases as the connectivity density decreases, reaching an asymptotic value as $N/C \rightarrow \infty$. The quantitative results for hexagonal grids have implications consistent with those of the binary model. They suggest that, again, a network of grid cells, for which a plausible number of synapses per neuron may be in the order of thousands and with a connectivity, say, of order $C/N \simeq 0.1$, would have the capacity to encode perhaps 100 different environments.

3.3 Sparsity and Noise Reverberation. The binary model shows that the difference in capacity between hexagonal and square grids results from the effective interactions among the fields in different tiles, as it emerges




Figure 5: Storage capacity in the threshold-linear model as a function of the inverse connectivity density N/C , on a log-log scale. Solid lines give α_c for the three different interaction kernels (bands in green, square grids in red, and hexagonal grids in blue). Dashed lines indicate what the capacity would be without noise reverberation. The crosses on the left show the capacity obtained with numerical simulations for a fully connected network.

only with wide kernels and dense coding. When both are sparse, hexagonal and square grids are roughly equivalent. The $w \rightarrow 0$ limit can be worked out analytically and $\alpha_c \rightarrow 0$ in both cases, but only after having reached a maximum around $\alpha_c \simeq 0.02$ for quite sparse codes, $w \simeq 0.03$ and $f \simeq 0.015$. Sparse coding is known to suppress noise reverberation (leading to small ψ), but remarkably, this relatively large capacity is approximately preserved for hexagonal grids with dense coding, $w \simeq 0.5$ and $f \simeq 0.25$, illustrating the efficiency with which this compact arrangement minimizes interference.

The threshold-linear model affords complementary insight into how the hexagonal/square capacity difference depends on the units active in each attractor reverberating their activity. Mathematically, this is expressed explicitly by the dependence of equation 3.3 on the order of parameter ψ , which quantifies the amount of reverberation through the loops in the networks. The physical meaning of ψ can be inferred from the expression derived in appendixes B and C:

$$\psi = g' \frac{T_0}{d} f. \tag{3.4}$$

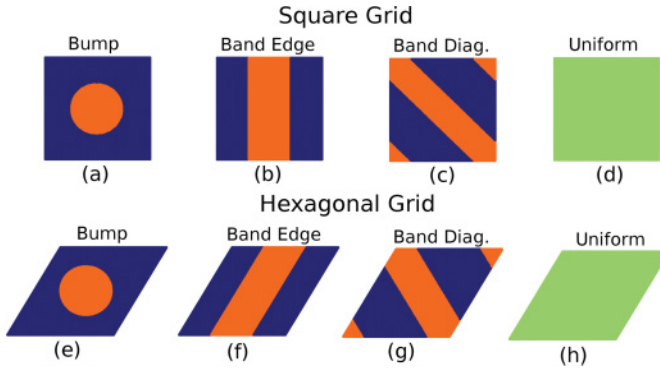


Figure 6: Different solutions to the saddle point equations in the binary model. Bumps (a, e) are stable at low f ($f = 0.2$ in the figure). Edge-oriented and diagonal bands are stable solutions for the $\theta = 60^\circ$ tile at higher f (e.g., $f = 0.4$, panels f and g) but only the former (panel b) are stable for $\theta = 90^\circ$. Uniform solutions (panels d and h) are always unstable below the critical capacity.

The factor $g'T_0/d$ is in fact the typical noise T_0/d amplified by the renormalized gain g' and multiplied by the average fraction of active units, the f parameter as in the binary model. ψ is then the one-step loop term in the reverberation of the noise; its effect on the capacity is illustrated by the dashed lines in Figure 5, in which such contribution is factored out. For densely connected networks, storage capacity would massively increase and relative differences would decrease without noise reverberation. The optimal capacity for the hexagonal kernel is then (mainly) the result of a reduced reverberation of the noise due to the shape of the activity distribution of its attractors: the average fraction of active units ($f \sim 0.46$) in the attractive state of the hexagonal kernel model is considerably lower than the same fraction in the square kernel, where it would be $f \sim 0.79$ for the square grids, and is only somewhat reduced to $f \sim 0.68$ for the stripes, which replace them as the stable solutions for this kernel.

4 Band Solutions

In the previous analysis, we focused on bump states, in which activity is localized in a grid pattern. Another possibility are partially localized solutions: band states, where activity is localized along a single direction in the elementary tile and extends along a stripe in the orthogonal direction.

In the binary model, these band states can be oriented along an edge of the tile (see Figures 6b and 6f) or along the diagonal of the tile (see Figures 6c and 6g) or in a discrete multiplicity of other orientations. Individual units fire along stripes of the same orientation, with relative offsets. We can study

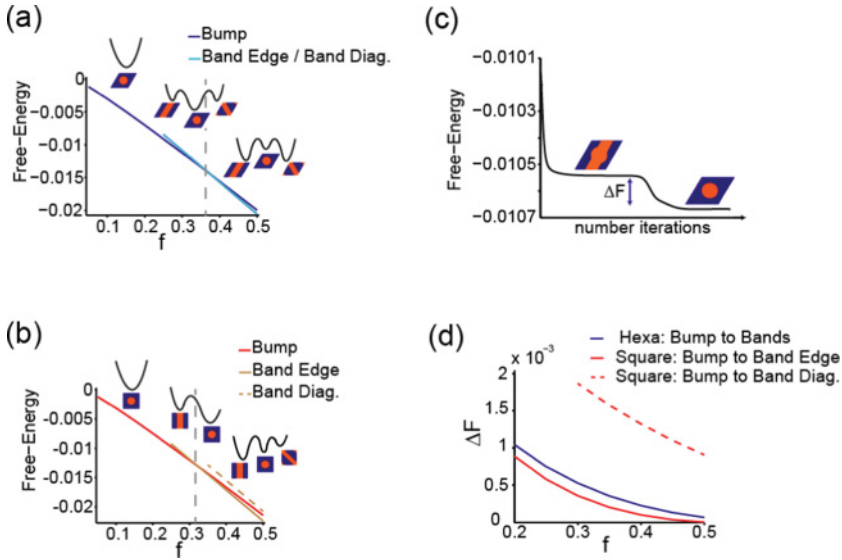


Figure 7: Bump and band states in the binary model. Free energies of the bump and band states for hexagonal grids (a) and square grids (b). (c) Free energy barriers are given by the difference in free energies between an unstable mixed state (band edge plus bump shown here) and a metastable state (bump state shown here). (d) Size of the free energy barriers to cross to go from the bump state to band states. $w = 0.1, \alpha \rightarrow 0$.

the properties of some of these band states in the $w - f$ parameter space, to find that they are particularly favored in regions of high coding level. Given the connectivity range set by w , bump states are the global minima of the free energy for low f , and one of the band states (which one depends on θ) becomes the minimum for higher f . For example, for both square and hexagonal grids, at small connectivity range $w = 0.1$, band states have lower free energy than the bump state for coding levels beyond 0.35, while for the larger connectivity range $w = 0.5$, this happens for coding levels beyond 0.4. This is intuitive, since for sufficiently large f , a band state has a shorter boundary between active and quiescent units than a bump, and it is the length of the boundary that raises the free energy above its minimum value. Moreover, we can study how these different states are separated by computing the size of the free energy barrier to cross to go from one state to another. The method to compute this barrier is sketched in Figure 7c and explained in more detail in appendix D. In Figure 7d, we show the size of the barriers to cross to go from a bump state to band states. On the range of coding levels where these two kinds of states coexist, the bump state is always more robust for a hexagonal grid compared to a square grid, as shown by

the higher barrier size in a hexagonal grid (blue curve, from bump to band edge or band diagonal state) compared to square grid (full red curve, from bump to band edge state).

Different behavior is observed in the threshold-linear network. In this case, the rigid symmetry imposed by the three-cosine interaction kernel makes the bump pattern a global minimum. In the two-cosine case instead, band state are stable solutions, corresponding to a macroscopic overlap with only one of the two cosines. We can describe bands also with a 1D interaction kernel, with a single cosine, and compare the storage capacity for band patterns with the one for square and hexagonal grids. In Figure 5, the green line shows the capacity for band patterns as a function of the connectivity. For a densely connected network, it is above that for square grids, and the barrier method indicates that these are only marginally stable to collapsing into stripes. This is in line with the reduction of the capacity from one to two dimensions shown in Battaglia and Treves (1998). Interestingly, as soon as the emergence of a third cosine is allowed, the capacity is instead enhanced, surpassing the 1D kernel except for very low values of connectivity density.

5 Discussion

Our results indicate that given appropriate conditions, a neural population with recurrent connectivity can effectively store and retrieve many hexagonally periodic continuous attractors. This possibility suggests that a regular grid code may not be restricted to represent only physical space; it could also express continuous abstract relations between arbitrary features, at least if they can be mapped to a two-dimensional space. This would, however, require a system flexible enough to store and retrieve uncorrelated grid representations. Our results show that this flexibility does not need, in principle, separate neural populations for separate representations, but can be achieved by a single local ensemble, provided it can effectively learn orthogonal representations.

Given the recent observation of nonspatial coding—a consistently tuned response to the position along a 1D nonspatial variable, sound frequency, during a sound manipulation task—by neurons that qualify as grid cells in a 2D spatial exploration task (Aronov et al., 2017), it would be interesting to know whether a similar selectivity can be observed for a 2D nonspatial variable, as suggested by indirect observations of hexagonal modulation (Constantinescu et al., 2016).

Several important questions are left open for future investigation. First, if global remapping is possible within a grid cell population, why has it not been observed experimentally? Possibly a remapping capacity of grid cells may have been hidden by the fact that multiple mappings were studied only in simple, empty, flat environments—and then they turned out to be the same, modulo translations (Fyhn et al., 2007). The hypothesis of a universal

grid, which shifts without deformation across an environment and from one environment to the other, faces severe difficulties as soon as curvature is taken into consideration. In curved environments, rigid translations are not possible, and the geodesic transformations that partially substitute for them do not leave field-to-field relations unchanged, making a universal grid a priori impossible. Nevertheless, natural environments show a wide range of both positive and negative curvature, which does not seem to pose any problem to the navigational skills of rodents or other species. It is then conceivable that the apparent universality of the grid pattern comes from the experimental restriction to flat environments, which all belong to the same, rather special, class of 2D spaces with zero curvature, and that a richer grid behavior is required in order to code for position in more general spaces.

The emergence of grid representations in curved environments has been investigated with a model based on single cell adaptation (Stella et al., 2013; Urdapilleta et al., 2015), which illustrates the emergence of different regular patterns for distinct ranges of curvature. Estimating the storage capacity of recurrent networks expressing curved grids, however, poses some challenges. Since shifting the grid pattern along a curved surface moves individual fields by a different amount, the relationships between grid units cannot be reduced to relationships between a single pair of their fields. Long-range translational coherence becomes impossible. Curved grids can be only partially coherent, and whether this partial coherence is sufficient to build stable attractors is an open problem (Stella et al., 2019).

A second open problem is the ability of a network encoding multiple charts to support path integration, since the noise introduced by other charts is likely to introduce discontinuities in the dynamics shifting the activity bump, affecting the accuracy of the integrator. It has recently been suggested (Mosheiff & Burak, 2019) that interactions between different grid modules (each encoding a single chart or coherent ensemble of maps) can enhance the robustness to noise during path integration. The possibility that this result generalizes to modules encoding multiple charts, and the analysis of the capacity deriving from interactions between modules, is beyond the scope of this letter but deserves future investigation.

Finally, a third issue concerns the learning dynamics that sculpts the grid attractors. What is the mechanism that leads to the attractors of the recurrent network? Does a single grid dominate it, in the case of flat environments? Can self-organization be unleashed by the interplay between the neural populations of medial entorhinal cortex (mEC), including non-grid units, and hippocampal place cells, aided by the dentate gyrus (Treves, Tashiro, Witter, & Moser, 2008)? Including the hippocampus may be needed also to understand the distortion of the grid pattern, reported in several experimental studies (Stensola et al., 2012; Stensola et al., 2015; Krupic et al., 2015), that disrupting long-range order also weakens coherence. At the system level, a finite storage capacity for the grid cell network implies the possibility that mEC, or any other area in the brain (Constantinescu et al., 2016)

that is observed to include grid-like units, can serve context memory. This would turn upside down the widely shared notion that memory for the specific spatial features of each environment is available only downstream, in the hippocampus, and conceptually reunites mEC with other regions of the mammalian temporal lobe, known to be dedicated to their own flavor of memory function (Eichenbaum, 2000). Moreover, the possibility of multiple uncorrelated continuous attractors in flat environments, combined with the discovery of transitions between (highly correlated) states in which the grid is the same but the peak firing rate of each field is different (Kanter et al., 2017), and with a new understanding of the disorder and frustration inherently associated with the grid representation of the curved environment (Stella et al., 2019), puts to rest the rigid order that appeared as the most salient characteristic of the newly discovered grid cells. It suggests instead a sort of spin glass at intermediate temperature—that in order to code densely and efficiently for position on (many) continuous manifolds, grid cells have to be equipped with the flexibility and the ability to compromise characteristic of a self-organized disordered system.

Appendix A: Mean Field Equations: Binary Model

The free-energy can be written, in the large N limit, in terms of macroscopic quantities:

$$\begin{aligned} \mathcal{F} = & \frac{\alpha\beta}{2}r(f-q) - \frac{\alpha}{\beta}\Omega(q, \beta) + \int \mu(\vec{x})\rho(\vec{x}) \\ & - \frac{1}{2} \int d\vec{x}d\vec{y} \rho(\vec{x})K(|\vec{x}-\vec{y}|)\rho(\vec{y}) \\ & - \frac{1}{\beta} \int d\vec{x} \int Dz \ln[1 + e^{\beta z \sqrt{ar} + \beta\mu(\vec{x})}], \end{aligned} \quad (\text{A.1})$$

where β is an inverse temperature or noise level and the function $\Omega(q, \beta)$ is given by

$$\begin{aligned} \Omega(q, \beta) = & 2 \sum_{k_1=1}^N \sum_{k_2=1}^N \left\{ \frac{\beta(q-f^2)}{1/\lambda_{k_1, k_2} - \beta(f-q)} - \ln[1 - \lambda_{k_1, k_2}\beta(f-q)] \right\} \\ & + \sum_{k=1}^N \left\{ \frac{\beta(q-f^2)}{1/\lambda_{0, k} - \beta(f-q)} - \ln[1 - \lambda_{0, k}\beta(f-q)] \right. \\ & \left. + \frac{\beta(q-f^2)}{1/\lambda_{k, 0} - \beta(f-q)} - \ln[1 - \lambda_{k, 0}\beta(f-q)] \right\}. \end{aligned} \quad (\text{A.2})$$

The order parameters minimizing the free energy functional are the average activity $\rho(\vec{x})$ (see the main text) and

$$q = \int d\vec{x} \int Dz \left[1 + e^{-\beta z \sqrt{\alpha r} - \beta \mu(\vec{x})} \right]^{-2} \quad (\text{A.3})$$

$$r = 4(q - f^2) \sum_{k_1=1}^N \sum_{k_2=1}^N \left[\frac{1}{\lambda_{k_1, k_2}} - \beta(f - q) \right]^{-2} \\ + 2(q - f^2) \sum_{k=1}^N \left\{ \left[\frac{1}{\lambda_{0, k}} - \beta(f - q) \right]^{-2} + \left[\frac{1}{\lambda_{k, 0}} - \beta(f - q) \right]^{-2} \right\}, \quad (\text{A.4})$$

where λ enforces the constraint $\int d\vec{x} \rho(\vec{x}) = f$ and λ_{k_1, k_2} are the eigenvalues of the kernel K and are given by

$$\lambda_{k_1, k_2} = w \frac{J_1(2\sqrt{z_{k_1, k_2}})}{\sqrt{z_{k_1, k_2}}}, \\ z_{k_1, k_2} = w\pi \left(k_1^2 + \left(\frac{k_2 - k_1 \cos \theta}{\sin \theta} \right)^2 \right), \quad (\text{A.5})$$

where J_1 is the Bessel function of the first kind of order 1.

In the text, we focus on the limit of vanishing stochastic noise $\beta \rightarrow \infty$, and the term $\beta(q - f)$, which remains finite in such limit, can be identified with the parameter ψ of the threshold-linear model, quantifying the reverberation through the loops of the network of the quenched noise, which is due to the interference of the other maps.

Appendix B: Mean Field Equations: Threshold-Linear Model

When an energy function can be defined (with full or, in any case, symmetric connectivity), the thermodynamics of the system is dominated by the minima of the free energy density,

$$\mathcal{F} = -T \left\langle \left\langle \int Dz \ln \text{Tr}(h, h2) \right\rangle \right\rangle - \frac{1}{2} \sum_{\sigma, l} |m^{\sigma, l}|^2 \\ - B(m) - \sum_{\sigma, l} (m^{\sigma, l})^2 + mB'(m) - r_0 y_0 + r_1 y_1 \\ + \frac{\alpha d}{2\beta} \left(\ln[1 - T_0 \beta(y_0 - y_1)] - \frac{\beta y_1}{1 - T_0 \beta(y_0 - y_1)} \right) \quad (\text{B.1})$$

where we have maintained a notation consistent with Treves (1990a) and Battaglia and Treves (1998)—for example,

$$Tr(h, h_2) = k + \left(\frac{\pi g'}{2\beta}\right)^{1/2} \exp\left(\frac{\beta g'}{2}(h_0 - h)^2\right) \left\{ 1 + \operatorname{erf}\left[\frac{\beta g'^{1/2}}{2}(h_0 - h)\right] \right\}, \tag{B.2}$$

$$h = \sum_{\sigma l} m^{l\sigma} \cdot \eta^{l\sigma} + B'(m) - z\sqrt{-2T_0 r_1}, \tag{B.3}$$

$$h_2 = r_1 - r_0, \tag{B.4}$$

$$1/g' = 1/g - 2h_2, \tag{B.5}$$

$$Dz = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz, \tag{B.6}$$

while $\langle\langle \cdot \rangle\rangle$ denotes an average over the quenched noise (the field centers in all other stored maps, distinct from the one currently expressed); and $B(x)$, together with the gain g , can be used to constrain the mean activity and the sparsity of the activity pattern (Treves, 1990a), analogous to the parameter λ in the binary model.

The minima are given, in the limit $T \rightarrow 0$, by the saddle point equations:

$$m^{l\sigma} = g' \left\langle\left\langle \eta^{l\sigma} \int_{h>Th} Dz(h - Th) \right\rangle\right\rangle, \tag{B.7}$$

$$m = g' \left\langle\left\langle \int_{h>Th} Dz(h - Th) \right\rangle\right\rangle, \tag{B.8}$$

$$y_0 = g'^2 \left\langle\left\langle \int_{h>Th} Dz(h - Th)^2 \right\rangle\right\rangle, \tag{B.9}$$

$$r_0 = \frac{\alpha T_0}{2} \frac{1 - T_0\beta(y_0 - 2y_1)/d}{(1 - T_0\beta(y_0 - y_1)/d)^2}, \tag{B.10}$$

$$y_1 = g'^2 \left\langle\left\langle \int_{h>Th} Dz(h - Th)^2 \right\rangle\right\rangle - T_0 g' \left\langle\left\langle \int_{h>Th} Dz \right\rangle\right\rangle, \tag{B.11}$$

$$r_1 = \frac{\alpha T_0}{2d} \frac{T_0\beta y_1}{(1 - T_0\beta(y_0 - y_1)/d)^2}. \tag{B.12}$$

Introducing the variables

$$\rho^2 = \frac{\alpha T_0^2 y_0}{d(1 - \psi)^2} \tag{B.13}$$

$$\psi = g' \frac{T_0}{d} \left\langle \left\langle \int^+ Dz \right\rangle \right\rangle, \quad (\text{B.14})$$

we can write the free energy as a function of macroscopic quantities,

$$\begin{aligned} \mathcal{F} = & -\frac{g'}{2} \left\langle \left\langle \int_{h>Th} Dz (h - Th)^2 \right\rangle \right\rangle + \sum_{\sigma l} \frac{(m^{l\sigma})^2}{2} \\ & + mb(m) - B(m) + \frac{T_0 \rho^2 \psi d}{2}, \end{aligned} \quad (\text{B.15})$$

now with

$$g' = \frac{1}{\frac{1}{g} - \frac{\alpha T_0}{1-\psi}}. \quad (\text{B.16})$$

To calculate the storage capacity, we focus on the case in which a single environment is retrieved by the network,

$$\begin{aligned} m^{ll} & > 0, \\ m^{\pi l} & = 0, \quad \forall \pi \neq l, \end{aligned}$$

although the analysis can be extended to the retrieval of bump states that are localized in multiple environments. Without loss of generality, we assume, therefore, that environment $\pi = 1$ is retrieved. With this assumption and introducing the two signal-to-noise ratios,

$$v^l = \frac{m^l}{\rho}, \quad (\text{B.17})$$

$$w = \frac{b(m) - Th}{\rho}, \quad (\text{B.18})$$

which represent, respectively, the environment-specific component of the signal and the uniform background inhibition acting on each unit, the saddle point equations can then be reduced to a system of two equations in two variables:

$$E_1(v, w) = A_1^2(v, w) - d\alpha A_3(v, w) = 0, \quad (\text{B.19})$$

$$E_2(v, w) = A_1(v, w) \left(\frac{d}{gT_0} - A_2(v, w) \right) - d\alpha A_2(v, w) = 0, \quad (\text{B.20})$$

where $A_1(w, v)$, $A_2(w, v)$, and $A_3(w, v)$ are the averages:

$$A_1(w, v) = \frac{1}{v^2 T_0} \left\langle \left\langle \sum_l \mathbf{v}^l \cdot \boldsymbol{\eta}^l \int^+ Dz \left(w + \sum_l \mathbf{v}^l \cdot \boldsymbol{\eta}^l - z \right) \right\rangle \right\rangle - \left\langle \left\langle \int^+ Dz \right\rangle \right\rangle, \quad (\text{B.21})$$

$$A_2(w, v) = \frac{1}{v^2 T_0} \left\langle \left\langle \sum_l \mathbf{v}^l \cdot \boldsymbol{\eta}^l \int^+ Dz \left(w + \sum_l \mathbf{v}^l \cdot \boldsymbol{\eta}^l - z \right) \right\rangle \right\rangle \quad (\text{B.22})$$

$$A_3(w, v) = \left\langle \left\langle \int^+ Dz \left(w + \sum_l \mathbf{v}^l \cdot \boldsymbol{\eta}^l - z \right)^2 \right\rangle \right\rangle. \quad (\text{B.23})$$

Solutions to equations B.19 and B.20 give the minima of the free energy that correspond to the retrieval of one of the stored environments. $E_1(v, w) = 0$ describes a closed curve in the $w - v$ plane, and these solutions are the intersections with $E_2(v, w) = 0$, which depends on the gain g .

As the storage load $\alpha = p/C$ increases, this closed curve shrinks and eventually disappears. The value $\alpha = \alpha_c$ at which the curve vanishes marks a phase transition: for $\alpha > \alpha_c$, retrieval solutions do not exist. The storage capacity α_c can therefore be calculated by finding the vanishing point of $E_1 = 0$, and in this way, one automatically selects the optimal value of the gain g as the one for which the curve $E_2 = 0$ intersects the single point $E_1 = 0$.

Appendix C: Finite Connectivity and Noise Reverberation

Equations B.19 and B.20 can be extended to arbitrary value of connectivity density C/N following the self-consistent signal-to-noise analysis developed in Roudi and Treves (2004). This gives

$$E_1 = A_2^2 - \left(1 + \frac{C}{N} \left(\frac{(2 - \psi)\psi}{(1 - \psi)^2} \right) \right) d\alpha A_3 = 0, \quad (\text{C.1})$$

$$E_2 = \left(\frac{d}{gT_0} - d\alpha \frac{C\psi}{N(1 - \psi)} \right) - A_2 = 0. \quad (\text{C.2})$$

These equations interpolate, as the free parameter C/N varies, between the two limiting cases of a fully connected network ($C/N = 1$) and the extremely diluted case ($C/N \rightarrow 0$) studied in Treves (1991). We see that the reverberation factor ψ enters in the equation for the storage capacity as a correction on the loopless equation $A_2^2 - d\alpha A_3 = 0$, modulated by the connectivity density C/N , and that the lower the ψ , the higher the storage capacity.

For the fully connected network, this correction gives

$$\frac{\psi}{1 - \psi} = \sum_{k=1}^N \psi^k, \quad (\text{C.3})$$

which is the sum over all the k-loops' contributions to the reverberation of the noise.

Note, finally, that for ease of comparison with the binary model, we have written in the main text,

$$\begin{aligned} \bar{\mu} &= A_2, \\ r &= A_3. \end{aligned} \quad (\text{C.4})$$

Appendix D: Free-Energy Barriers in the Binary Model

Free-energy values for the different metastable states are calculated using equation A.4 after order parameters have been computed by solving the saddle-point equations. These equations are solved iteratively, starting from an initial condition for order parameters and iterating the values of the order parameters until convergence to fixed values. The free-energy values of the different metastable states are obtained by initializing $\rho(\vec{x})$ as $\rho_{Bp}(\vec{x})$ for bump states (see Figures 6a and 6e) and $\rho_{BaE}(\vec{x})$ for band edge states (see Figures 6b and 6f) or $\rho_{BaD}(\vec{x})$ for band diagonal states (see Figures 6g and 6f). In order to estimate the size of the barrier that must be jumped over in order to go from one state X to another state Y , we proceed as follows. The activity profile is initialized as $\rho^{k=0,z}(\vec{x}) = z\rho_X(\vec{x}) + (1-z)\rho_Y(\vec{x})$, with z chosen such that $\rho^{k \rightarrow +\infty, z}(\vec{x}) = \rho_X(\vec{x})$ and $\rho^{k \rightarrow +\infty, z-\epsilon}(\vec{x}) = \rho_Y(\vec{x})$ for $\epsilon \ll z$. When solving equations from such an initial condition, the network state goes close to a saddle point lying at the boundary between the two basins of attraction associated with states X and Y , before sliding into state X , as shown in Figure 7c. The size of the barrier is then given by the difference between the free energy of the saddle point and that of the metastable state X .

Acknowledgments

This work was supported by Human Frontier grant RGP0057/2016.

References

- Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M. B. (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences*, 111(52), 18428–18435.

- Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature*, *543*(7647), 719.
- Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, *58*(6), 7738.
- Boccaro, C. N., Nardin, M., Stella, F., O'Neill, J., & Csicsvari, J. (2019). The entorhinal cognitive map is attracted to goals. *Science*, *363*(6434), 1443–1447.
- Burak, Y., & Fiete, I. R. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Computational Biology*, *5*(2), e1000291.
- Butler, W. N., Hardcastle, K., & Giocomo, L. M. (2019). Remembered reward locations restructure entorhinal spatial maps. *Science*, *363*(6434), 1447–1452.
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, *352*(6292).
- Couey, J. J., Witoelar, A., Zhang, S. J., Zheng, K., Ye, J., Dunn, B., . . . Witter, M. P. (2013). Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature Neuroscience*, *16*(3), 318.
- Dunn, B., Wennberg, D., Huang, Z., & Roudi, Y. (2017). *Grid cells show field-to-field variability and this explains the aperiodic response of inhibitory interneurons*. arXiv:1701.04893.
- Eichenbaum, H. (2000). A cortical/hippocampal system for declarative memory. *Nature Reviews Neuroscience*, *1*(1), 41–50.
- Fiete, I. R., Burak, Y., & Brookings, T. (2008). What grid cells convey about rat location. *Journal of Neuroscience*, *28*(27), 6858–6871.
- Fuhs, M. C., & Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, *26*(16), 4266–4276.
- Fyhn, M., Hafting, T., Treves, A., Moser, M. B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* (7132), 190–194.
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*(7052), 801.
- Häggglund, M., Mørreaunet, M., Moser, M. B., & Moser, E. I. (2019). Grid-cell distortion along geometric borders. *Current Biology*, *29*, 1047–1054.
- Hedrick, K. R., & Zhang, K. (2016). Megamap: Flexible representation of a large space embedded with nonspatial information by a hippocampal attractor network. *Journal of Neurophysiology*, *116*(2), 868–891.
- Kanter, B. R., Lykken, C. M., Avesar, D., Weible, A., Dickinson, J., Dunn, B., . . . 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.
- Kropff, E., & Treves, A. (2008). The emergence of grid cells: Intelligent design or just adaptation? *Hippocampus*, *18*(12), 1256–1269.
- Krupic, J., Bauza, M., Burton, S., Barry, C., & O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature*, *538*, 232–235.
- Krupic, J., Bauza, M., Burton, S., & O'Keefe, J. (2018). Local transformations of the hippocampal cognitive map. *Science*, *359*, (6380), 1143–1146.
- Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, *309*(5734), 619–623.

- Monasson, R., & Rosay, S. (2013). Cross-talk and transitions between multiple spatial maps in an attractor neural network model of the hippocampus: Phase diagram (I). *BMC Neuroscience*, *14*(1), 015.
- Monasson, R., & Rosay, S. (2014). Crosstalk and transitions between multiple spatial maps in an attractor neural network model of the hippocampus: Collective motion of the activity. *Phys. Rev. E*, *89*, 032803.
- Mosheiff, N., & Burak, Y. (2019). *Velocity coupling of grid cell modules: Stable embedding of a low dimensional variable in a high dimensional neural attractor*. bioRxiv, 651513.
- Roudi, Y., & Treves, A. (2004). An associative network with spatially organized connectivity. *Journal of Statistical Mechanics: Theory and Experiment*, *2004*(7), P07010.
- Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, *17*(15), 5900–5920.
- Stella, F., Si, B., Kropff, E., & Treves, A. (2013). Grid cells on the ball. *Journal of Statistical Mechanics: Theory and Experiment*, *2013*(3), P03013.
- Stella, F., Urdapilleta, E., Luo, Y., & Treves, A. (2019). *Partial coherence and frustration in self-organizing spherical grids*. BioRxiv:560318 [Preprint].
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M. B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, *492*(7427), 72.
- Stensola, T., Stensola, H., Moser, M. B., & Moser, E. I. (2015). Shearing-induced asymmetry in entorhinal grid cells. *Nature* (7538), 207–212.
- Treves, A. (1990a). Graded-response neurons and information encodings in auto-associative memories. *Phys. Rev. A*, *42*, 2418–2430.
- Treves, A. (1990b). Threshold-linear formal neurons in auto-associative nets. *Journal of Physics A: Mathematical and General*, *23*(12), 2631.
- Treves, A. (1991). Dilution and sparse coding in threshold-linear nets. *Journal of Physics A: Mathematical and General*, *24*(1), 327–335.
- Treves, A., Tashiro, A., Witter, M. P., & Moser, E. I. (2008). What is the mammalian dentate gyrus good for? *Neuroscience*, *154*(4), 1155–1172.
- Urdapilleta, E., Troiani, F., Stella, F., & Treves, A. (2015). Can rodents conceive hyperbolic spaces? *Journal of the Royal Society Interface*, *12*(107), 20141214.
- Yartsev, M. M., Witter, M. P., & Ulanovsky, N. (2011). Grid cells without theta oscillations in the entorhinal cortex of bats. *Nature*, *479*(7371), 103.

Received April 18, 2019; accepted July 20, 2019.