



RESEARCH ARTICLE

WILEY

Partial coherence and frustration in self-organizing spherical grids

Federico Stella¹ | Eugenio Urdapilleta^{2,3} | Yifan Luo³ | Alessandro Treves^{3,4}

¹Institute of Science and Technology Austria, Klosterneuburg, Austria

²Centro Atómico Bariloche, Instituto Balseiro, Comisión Nacional de Energía Atómica (CNEA) and Universidad Nacional de Cuyo (UNCUYO), Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET), San Carlos de Bariloche, Argentina

³SISSA - Cognitive Neuroscience, Trieste, Italy

⁴NTNU - Centre for Neural Computation, Trondheim, Norway

Correspondence

Alessandro Treves, Cognitive Neuroscience, SISSA, via Bonomea 265, 34136 Trieste, Italy. Email: ale@sisssa.it

Funding information

FP7 Information and Communication Technologies, Grant/Award Number: Gridmap; Gordon and Betty Moore Foundation, Grant/Award Number: 2919.01; Human Frontier Science Program, Grant/Award Number: RGP0057/2016; National Institutes of Health, Grant/Award Number: R25GM067110; National Science Foundation, Grant/Award Number: PHY-1748958

Abstract

Nearby grid cells have been observed to express a remarkable degree of long-range order, which is often idealized as extending potentially to infinity. Yet their strict periodic firing and ensemble coherence are theoretically possible only in flat environments, much unlike the burrows which rodents usually live in. Are the symmetrical, coherent grid maps inferred in the lab relevant to chart their way in their natural habitat? We consider spheres as simple models of curved environments and waiting for the appropriate experiments to be performed, we use our adaptation model to predict what grid maps would emerge in a network with the same type of recurrent connections, which on the plane produce coherence among the units. We find that on the sphere such connections distort the maps that single grid units would express on their own, and aggregate them into clusters. When remapping to a different spherical environment, units in each cluster maintain only partial coherence, similar to what is observed in disordered materials, such as spin glasses.

KEYWORDS

continuous attractor, curved environment, ensemble coherence, grid cells, remapping

1 | INTRODUCTION

What are the defining properties of grid cells? In the 15 years since their discovery in the medial entorhinal cortex (mEC) (Fyhn, Molden, Witter, Moser, & Moser, 2004), two organizational principles appear to have emerged as the cornerstones of the phenomenon. The first one is, of course, the positioning of the fields of each individual cell at the vertices of a regular hexagonal tessellation of the environment (Hafting, Fyhn, Molden, Moser, & Moser, 2005). The second, a strong propensity of local ensembles of these cells to maintain their co-activation patterns across conditions and environments (Fyhn, Hafting, Treves, Moser, & Moser, 2007); in striking contrast to the behavior expressed by neighboring place cells, which make the swapping of activation partners (“remapping”) one of their defining features (Bostock, Muller, & Kubie,

1991). These two properties, the former expressed at the single-cell level, the latter constraining collective states of activity, have come to be regarded as quintessential to grid cells. They are often yoked together when discussing the “grid cell code” (Burak, 2014; Stemmler, Mathis, & Herz, 2015; Yoon et al., 2013), thus leaving it unclear whether such code is expressed more in the regularity of field arrangements or in the constancy of spatial phase relations, or in a necessary combination of both.

It should not be forgotten, however, that grid cells have been first described and mostly studied in flat, empty, bounded environments. Their entanglement could be possibly related to the Euclidean geometry of this very specific sort of environment, rather than being intrinsic to the cells. The question, then, is to what extent would single-cell and population properties still co-occur, when such a specific setting is abandoned, to reach for more naturalistic settings of complex,

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2019 The Authors. *Hippocampus* published by Wiley Periodicals, Inc.

curved, partially open environments, such as the burrows where rats live in the wild (Calhoun, 1963).

In two previous modeling studies, we argued that the notion of the hexagonal grid may need to be generalized in order to predict the behavior of such cells in environments of constant positive or negative curvature. With sufficient exposure to these environments, our model indicates how single grid cells may adapt by producing regular tessellations consistent with the underlying curvature (Stella, Si, Kropff, & Treves, 2013; Urdapilleta, Troiani, Stella, & Treves, 2015): tessellations with fivefold or lower symmetry for positive curvature; sevenfold or higher symmetry for negative curvature. An analysis based on Calhoun's (1963) study leads to the conclusion that the standard sixfold grid symmetry would arise at the single-cell level only when the curvature is near zero; while the actual range of curvature values of the natural Norway rat habitat extends further, both at the negative and at the positive ends of the spectrum (Figure 1).

What about the effect of interactions between grid cells? What was shown in (Stella et al., 2013) is only how a population of noninteracting grid-like units can self-organize a representation of the spherical surface, where each unit ends up displaying an independently "oriented," often quasi-regular grid. There, grid patterns emerged due to the progressive, unsupervised sculpting of the feed-forward connections through Hebbian plasticity induced by navigation-related activity. Contrary to

continuous attractor models (Burak & Fiete, 2009), interactions between mEC units were not needed for the emergence of individual grid maps—possibly, only for their coordination at the population level. Indeed, studies of the same model on *planar* environments have made clear how the introduction of lateral connectivity in the mEC population can induce an alignment among units, resulting in a common orientation of the fields emerging from the feed-forward self-organization process, while also reinforcing their symmetric arrangement (Si, Kropff, & Treves, 2012; Si & Treves, 2013). Urdapilleta et al. (2015) have observed that also in environments with constant negative curvature, the lateral interactions tend to favor a coherent arrangement of the fields across units, but such environments are dominated by their boundaries, which leads to arbitrary modeling choices, that in turn prevent reaching firm conclusions. A complete sphere, on the other hand, has no boundaries, and thus offers a conveniently simple model of a curved environment. Moreover, it has been effective as an experimental set-up, both on the outside (Harvey, Collman, Dombeck, & Tank, 2009) and on the inside (Krüge, 2016; Krüge, Wernle, Moser, & Moser, 2013), giving hope that once complemented with the appropriate sensory surround, it will allow developmental studies to investigate the formation of spherical representations in rodents. In the meantime, here we use the adaptation model to study the effect of collateral interactions among grid-like units self-organizing on a spherical world. We also briefly comment on the additional effects




FIGURE 1 Natural Norway rat environments span limited stretches with negligible curvature. Main graph: the symmetry expected at the single-unit level for different values of constant Gaussian curvature, measured by the ratio between grid spacing s and the radius of curvature λ . Blue curve: theoretical relation between the angle $\alpha = 2\pi/n$ of the n -fold symmetry and the ratio s/λ , $\cosh(s/\lambda) = \cos(\alpha)/[1 - \cos(\alpha)]$ (negative curvature) and $\cos(s/\lambda) = \cos(\alpha)/[1 + \cos(\alpha)]$ (positive curvature). Symmetric arrangements for $n = 4, 5$ (on a sphere, right), 6 (on the plane, center), and $7, 8$, and 9 (on a pseudo-sphere, left) are indicated, with possible curved environments to be used in the laboratory at the top of the left and right column. Green arrows from the flat sixfold grid example indicate the range where it may be relevant, $|s/\lambda| \leq 1$, before other symmetries prevail. Superimposed on the graph is a drawing of a Norway rat den, from Calhoun (1963). He estimated the inner radius of the tunnels, r , to be below 5 cm, which implies that for the radius of Gaussian curvature λ of a curved tunnel to be of order the grid spacing s , even for a small $s = 40$ cm, the outer radius R of the den has to be several meters. This means that the sixfold symmetry is relevant only to roughly straight tunnel segments, approximately indicated in green, while most of the den (example tunnel in red) does not admit symmetric grids. The chambers are too small to reveal spherical arrangements, and their representation may be more akin to that of the turning points in the hairpin maze (Derdikman et al., 2009) [Color figure can be viewed at wileyonlinelibrary.com]

expected of gravity, and of boundaries, when extending the analysis from our artificial spherical environment to ecologically plausible ones—an extension that we leave for future studies.

2 | THE MODEL

The basic details of the model are identical to (Stella et al., 2013), and are described in Appendix A, with the critical addition of a set of recurrent collaterals connecting units of the EC layer.

Time is discretized in steps of length $t = 0.01$ s. The total length of a simulation is of 100 million steps (corresponding to nearly 12 days of continuous running, a very long time, to ensure that the self-organization process has approached its asymptote). The virtual rat moves on the surface of a sphere of radius 52.6 cm with a constant speed of $v = 40$ cm/s. To obtain *smooth* random trajectories, resembling those observed in experiments, running direction changes gradually after each step, resulting in an extended correlation over time. For simplicity, the change in running direction between two consecutive steps of the virtual rat is sampled from a Gaussian distribution with zero mean and standard deviation $h = 0.2$ rad. The virtual rat always runs along the great circle determined by its running direction. Our model is comprised of two layers. The input array represents, for example, the CA1 region of the hippocampus and includes $N_{\text{hipp}} = 1,400$ units with their fields regularly arranged to evenly tile the spherical surface. The output network is comprised of a population of $N_{\text{mEC}} = 250$ would-be grid units, all with the same adaptation parameters—hence they represent a single mEC module, in relation to the modules discovered by Stensola et al. (2012).

In a limited set of simulations intended to explore the effects of boundaries, and of gravity, we used hemispheres instead of full spheres. The effects of a boundary, corresponding to the equator, was assessed both with isolated hemispheres, in which case the boundary was reflecting the trajectory of the virtual rat, and with hemispheres embedded in a flat surround, where the boundary amounted only to a sudden change of intrinsic curvature. Without gravity, whether the hemisphere is concave or convex makes no difference. We also simulated trajectories on concave and convex hemispheres with gravity, and the latter was modeled by an additional force dragging the trajectory toward the equator (in the convex hemisphere) or the bottom of the bowl (in the concave one). This force was parametrized by the change in speed when moving downward and the strength of the downward pull applied to the vector expressing the current direction of motion (see Appendix A).

Similarly to the planar case, we introduce collateral weights between mEC units to induce the coordinated development of their firing maps. The weights are set in the following way: each unit is temporarily assigned a preferred position, an “auxiliary field” at coordinates (ϕ, θ) on the sphere, and a “preferred direction” (angle relative to the meridian, with 0 pointing toward the North Pole). The coordinates as well as the angle are randomly chosen. These auxiliary fields are introduced solely to define the collateral weights, and not to position the subsequently developing grid fields, nor do they play any role in other parts of the simulations. They are only used, in other words, to induce a notion of similarity among output units. The collateral weight between unit i and unit k is then calculated as

$$J_{ik} = \left[f_{\theta_i}(\omega_{ik}) f_{\theta_k}(\omega_{ik}) \exp\left(-\frac{d_{ki}^2}{2\sigma_f^2}\right) - \kappa \right]^+ \quad (1)$$

where $[\]^+$ denotes the Heaviside step function, $\kappa = 0.05$ is an inhibition factor to favor sparse weights, f is a tuning function described in the Appendix A and ω_{ik} is the direction, with respect to the North Pole, of the line connecting the auxiliary fields of unit i and k , along the great circle. $\sigma_f = 10$ cm denotes how broad the connectivity is, and d_{ki} is defined as

$$d_{ki} = R \cos^{-1}(X_i X_{\text{exp}} + Y_i Y_{\text{exp}} + Z_i Z_{\text{exp}}) \quad (2)$$

that is, it is the distance between the coordinates of the auxiliary field (X_i, Y_i, Z_i) and the *expected* position of a virtual rat that had started at the auxiliary field of unit k and had moved 10 cm along the geodesics joining both fields, corresponding to 250 ms of reverberatory delayed activity of movement along this direction. The definition of the weights leads to a localized connectivity pattern, such that strong positive interactions are only generated between units with similar preferred head direction and activation fields appropriately shifted along the same head direction (Kropff & Treves, 2008; Si et al., 2012). The resulting connectivity is rather sparse, with only about 8% of the possible pairs sharing a nonzero weight. As with the feed-forward connectivity, normalization on this set of connections is performed by setting a unitary L^2 norm on the presynaptic strengths for each mEC unit. Moreover, the relative strength of the recurrent input with respect to the feed-forward input was reduced to 0.2 for most of the simulations (see Appendix A). Notice that our model does not include plasticity on the recurrent set of connections: their value is defined once, at the beginning of each simulation, and then kept unmodified throughout. Note also that with a radius $R = 52.6$ cm and the adaptation parameters we use, most units tend to have 12 fields in the simulations with recurrent connections, but 13 or 14 fields without (see Figure S1, top). Still, we choose to use the same value of R in the two cases for ease of comparison. When separately varying R in order to optimize the proportion of units evolving 12 fields in each condition, and taking into account all units, the simulation without collaterals yields maps much closer to the “soccer ball” ideal (Figure S1, bottom).

3 | RESULTS

Simulations produce units with different number of fields, as shown in Figure S1. In the following, we set the radius to the same value $R = 52.6$ cm and consider only units which have developed 12 fields. We compare simulations with and without lateral connections.

For each simulation, the same set of grids are left to self-organize in two distinct, independent environments.

3.1 | Grid map distortion

At the single-map level, the introduction of lateral connections has the immediate result of interfering with the development of a regular grid structure. We can quantify this phenomenon by computing the spherical correlation (as the best match over any possible rotation) of the activity

map developed by each unit with the template of a perfect 12-field “soccer ball” rate map (Stella et al., 2013). One sees a marked effect of recurrent interactions as a general increase in the distance from the best-matching template, even though the radius of the sphere can be adjusted so that even with the recurrent collaterals most units produce 12 fields (Figures 2c,d and S1). The simulations without collaterals produce fairly good exemplars of an ideal spherical grid. This is not the case when recurrent collaterals are introduced: the interactions among EC units lead to a disruption of the regular arrangement of their fields (see the examples in Figure 2a,b; and the quantitative measures in Figure S2).

This illustrates how quickly the spherical case departs from what is observed on a planar surface. There, it has been shown (Si et al., 2012; Si & Treves, 2013), the presence of lateral connections has the crucial role of inducing a common orientation in the grid population and does not hinder, in fact enhances, the quality of the grids developed by the system. The same process does not occur here, where a similar attempt to induce coordination in grid-evolving units appears at odds with the regularity of the grids.

3.2 | Grid units tend to cluster

One can then ask, what are the effects of connections on the whole ensemble of grid units, and what does the “common orientation” that

they should induce look like, on the sphere? We answer these questions by analyzing the spatial (spherical) similarity in the structure of the maps developed by different units. To do so, for each unit, we computed the activity autocorrelation after 373,248 different rotations. Rotations were randomly drawn to evenly span the space of possible Euler rotations ($2\pi \times \pi \times 2\pi$, considering also the cosine factor). We then sorted the rotations according to their autocorrelation score, from highest to lowest. We used the first 5,000 best rotations to compute their distribution density in the three-dimensional space of Euler rotations $\Gamma_i(\phi, \theta, \psi)$ (where i denotes the unit). For each pair of units (i, j) we then computed the overlap between $\Gamma_i(\phi, \theta, \psi)$ and $\Gamma_j(\phi, \theta, \psi)$ (as a correlation). This correlation (or effectively, distance) matrix was used to identify clusters of units sharing a similar rotation distribution. The clustering was performed with the “ward” algorithm and the number of clusters was optimized over the range (4–15).

The outcome of this clustering algorithm shows how the population effectively breaks down into subgroups of segregated units, each developing an internal degree of coherence that is higher than that shared by the entire population. In Figure 3, the spatial position (on a 2D projection of the sphere) of the field centers of all units with 12 fields in the population are shown colored according to their cluster membership. To a large extent, each of these clusters expresses a

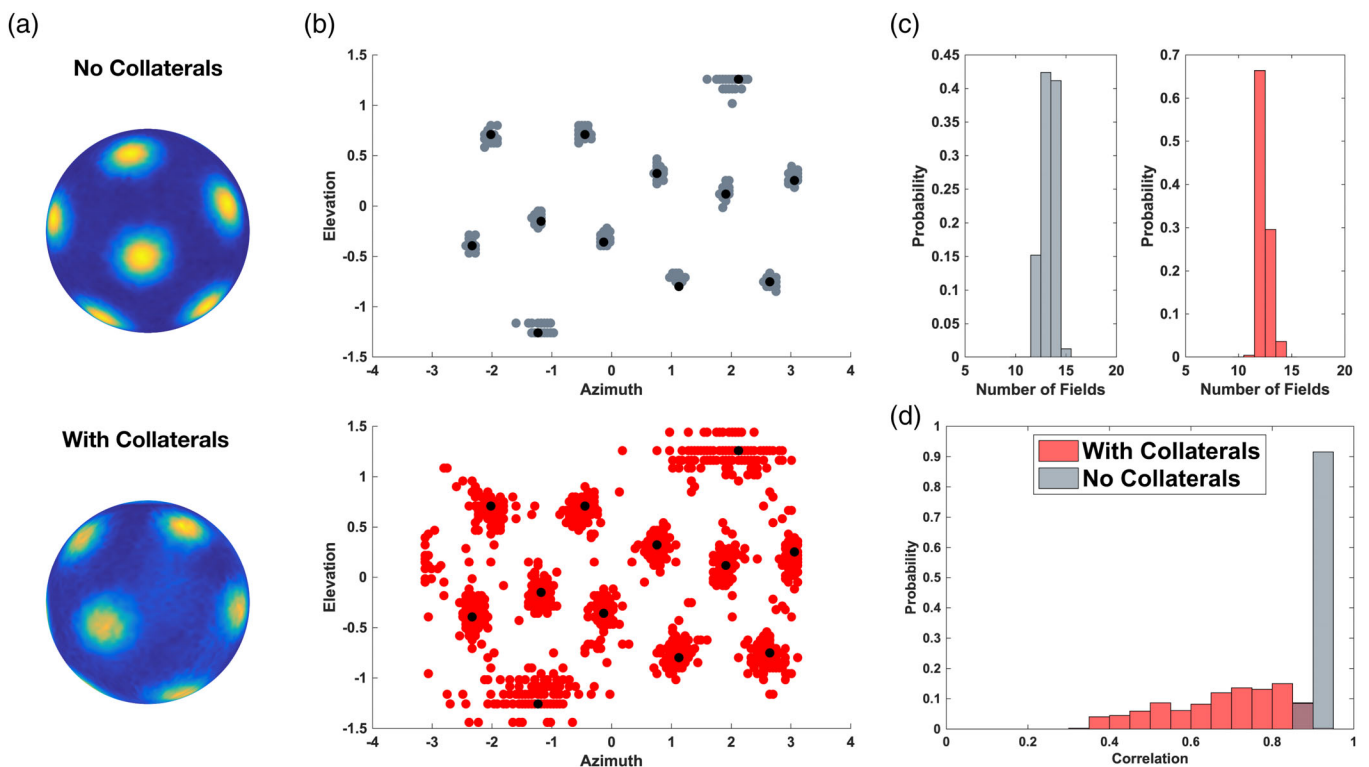
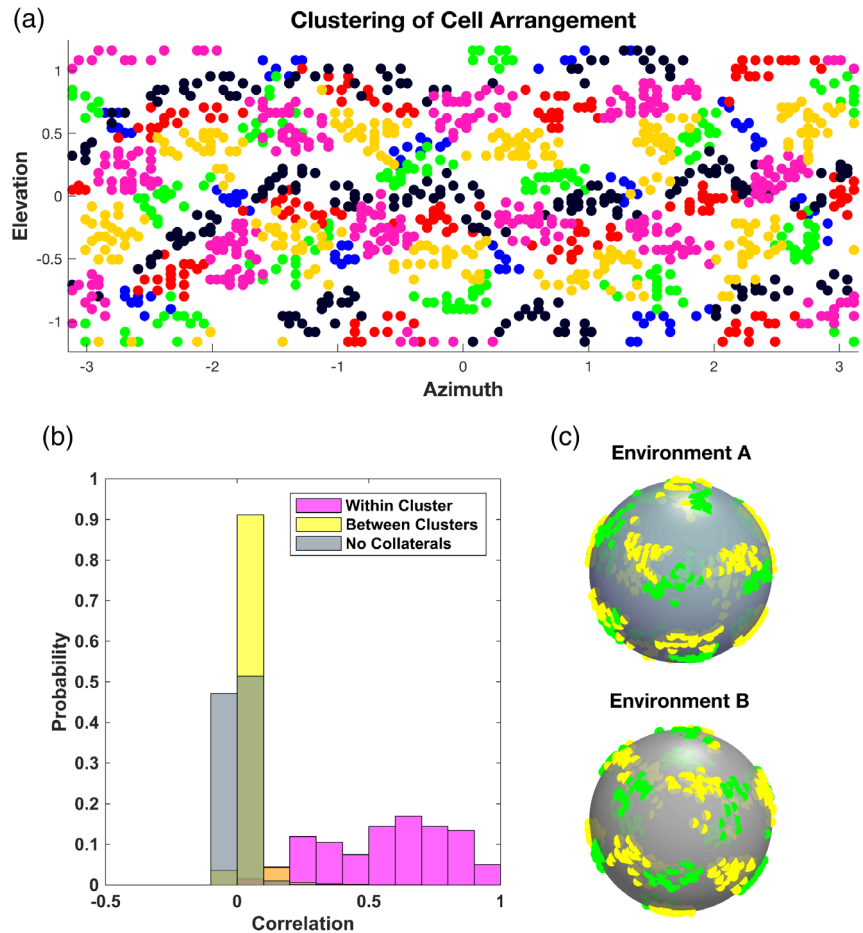


FIGURE 2 Collateral interaction distorts the grid pattern on a sphere. (a) Two representative examples of activity developed by grid units on the sphere. Top: unit from a simulation without collateral connectivity. Bottom: unit from a simulation in which units interact through collaterals. (b) Distribution of the position of all the fields from a population after the rate map of each unit has been rotated to maximize its overlap with a common “soccer ball” template (the field centers of this perfect grid are shown in black). Top: fields of a population of noninteracting units (mean distance from perfect center: 3.52°); Bottom: fields of a population of interacting units (mean distance from perfect center: 7.96°). (c) Distribution across the population of the number of fields developed by units, with $R = 52.6$ cm. Left: no collaterals; right: with collaterals. Sample sessions. Values across sessions: mean fraction with 12 fields, for no interaction 0.13, with interactions 0.68. (d) Correlation of all units with 12 fields with a best-rotated “soccer ball.” Aggregate from all sessions [Color figure can be viewed at wileyonlinelibrary.com]

FIGURE 3 Interacting grid units tend to cluster. (a) Spatial distribution of the fields of a population of grid units over the surface of a sphere. Each field center of each unit is assigned a color according to which cluster the unit belongs to. Sample session. Average number of clusters across sessions: 7.6. (b) Measure of clustering quality: distribution of pairwise unit correlations between the 3D density of their 5,000 best rotations in the space of Euler angles (the measure used to define clusters). Sample session. (c) Example of the arrangement of the fields from the units of two clusters (green and yellow, as in a). The positions of their fields are shown in two different spherical environments. Clusters were defined solely in the first environment [Color figure can be viewed at wileyonlinelibrary.com]



common orientation, meaning that each of the 12 fields of a unit tends to appear grouped with those of every other unit in the cluster. Different clusters appear instead unrelated, roughly to the same extent that individual units are in the nonconnected simulations (Figure 3b).

Crucially, the membership to one cluster is a feature that is carried over to new environments almost entirely unaltered (we describe the remapping procedure in the next paragraph). In Figure 3c we show a similar plot of the field centers, this time for only two of the clusters, to highlight their correspondence in two different environments. The grouping is conserved, as is the mutual avoidance of the fields in the two clusters.

Thus we observe how on the sphere, the interaction between units results into a break-down of the population, with different subsets acquiring a coordinated arrangement, while at the same time it forces each unit to distort its firing pattern away from that of an ideal grid. These features are consistently reproduced across environments.

3.3 | Remapping

In order to gain a better understanding of the mechanisms underlying grid formation and of the spatial code they can generate, we need to address the properties of remapping. To that end, maps were developed independently in two environments: the set of place cell inputs

and the associated feed-forward connectivity was randomly initialized in each of the two environments. Only the strength of the recurrent collaterals (when present), and thus of the grid cell interactions, was maintained after remapping. For each unit, we compare the maps developed in the two environments, maps A and B. Taking map A, we again apply a large random set of rotations spanning the entire Euler rotations range, and for each rotated version of map A we compute the resulting overlap (correlation) with map B. We first consider only the rotation associated with the highest similarity score—the “best” rotation—and its associated rotated version of the map, A' .

Results are shown in Figure 4a: in red one can see the distribution of correlation values between B and A' , and for comparison the results for the best rotation of a perfect grid (in light brown). The two distributions lie in the same range of high correlation, although the perfect grid can usually be rotated to achieve a higher correlation with the map in B, suggesting that the distortion observed in A is independent of that in B. We can compare both distributions with that obtained by correlating for each unit the best-rotated map A' with the map of each other unit in the same cluster, in B (magenta). In this case, correlation values are somewhat lower but, consistent with the partially coherent behavior expressed by units in the same clusters, they are still significantly higher than those obtained using either the maps in B of units in other clusters (yellow) or of all the other units (not shown).

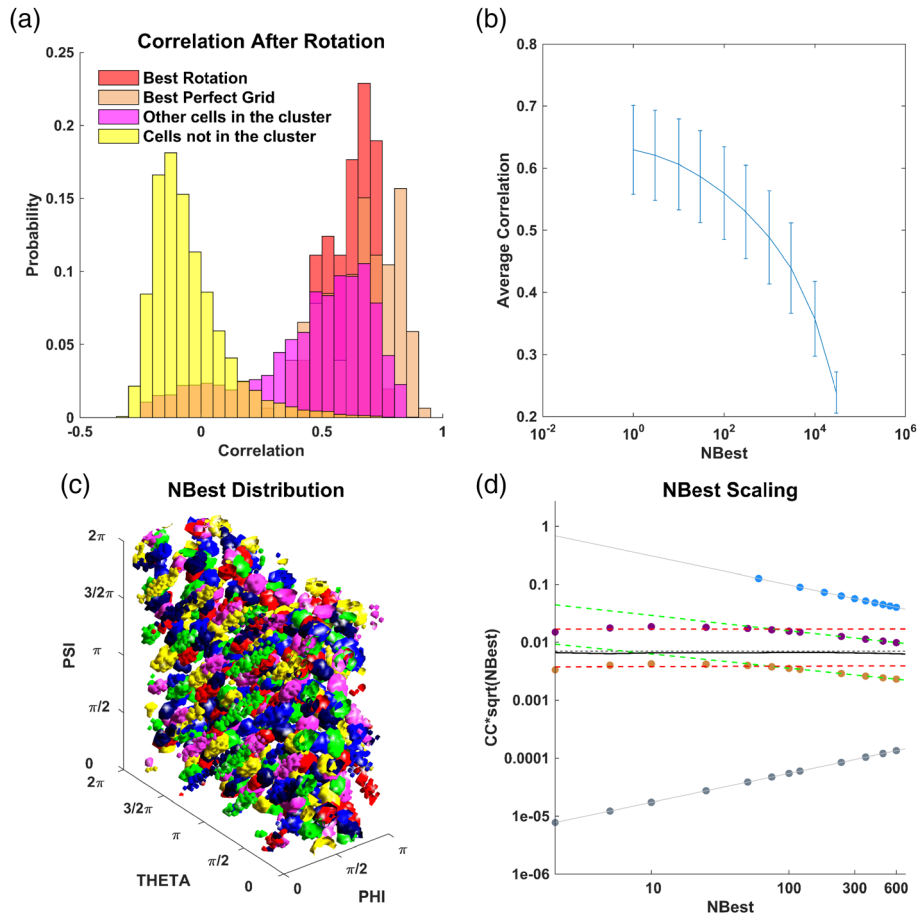


FIGURE 4 Remapping preserves the clusters. (a) Distribution of overlaps after the best rotation (out of 373,248 randomly drawn rotations). All sessions. Mean overlap: with the map in B of the same unit, 0.615; with those of other units in the cluster, 0.455; with those of units not in the same cluster, -0.050 ; with those of all units, 0.006 (not shown). Mean overlap of the best rotated perfect grid with the maps in B, 0.659. (b) Correlation- N_{Best} dependence. Average value of the correlation between the map in B and the map of the same unit in A rotated N_{Best} times (bars denote SDs). (c) Example of the spatial density distribution of best rotations for cells belonging to different clusters. Each color represents the distribution of the $N_{\text{Best}} = 500$ angles for five random units for each cluster. Sample session. (d) N_{Best} scaling and partial coherence. Different scaling $CC = 1/N_{\text{Best}}^\beta$ of the Clustering Coefficient (computed over distributions like the one in c) for different types of remapping. Logarithmic scales. For clarity, the quantity on the y-axis is $CC \cdot N_{\text{Best}}^{1/2}$. Gray: random remapping of noninteracting units, $\beta = 0$. Blue: coherent remapping, $\beta = 1$. Purple: CC computed using interacting grid units from the same cluster. Brown: CC computed using all units in an interacting population. Dashed black: $\beta = 0.5$; dashed red: CC computed using N_{Best} up to 100 (average over simulations: $\beta = 0.53$ all units, $\beta = 0.52$ units in same cluster); dashed green: CC computed using $100 < N_{\text{Best}} \leq 700$ (average over simulations: $\beta = 0.74$ all units, $\beta = 0.71$ units in same cluster); Black: computed using numerical estimation of the analytical formulation (see Appendix B) [Color figure can be viewed at wileyonlinelibrary.com]

Next, we observe that, if the two maps were perfect soccer balls, there would be $12 \times 5 = 60$ equivalent ways to rotate one into the other. Because of the distortions from the most symmetric configuration, the degeneracy is only approximate, but still massive: there are many different rotations, spanning a diverse set of Euler angle triplets, that lead to almost the same correlation values as the best rotation, for each unit. In fact, if we take for each unit its N_{Best} rotations of the map in A, we see that their average correlation with the corresponding map in B is a smooth function of N_{Best} , which averaged over the population shows a significant decrease only when N_{Best} reaches into the thousands (among an arbitrarily set range of 373,248 randomly chosen triplets (ϕ, θ, ψ) of Euler angles; Figure 4b).

Remarkably, the clusters of units defined on sphere A maintain a *partial* coherence once remapped onto sphere B, as already suggested

by the example in Figure 3c. If we randomly choose five sample units per cluster, and consider their $N_{\text{Best}} = 500$ individually most correlated rotations, we find that they cluster into distinct “islands” in Euler space, with each cluster contributing 60 regularly arranged islands, as shown in Figure 4c. Inside each island, however, chaos prevails. Some single units contribute many more of their best rotations to particular islands, and the shape of each island appears randomly distorted.

After failing to observe any further geometrical structure within the islands, we resorted to a quantitative measure of the extent to which the best rotations are concentrated across units. We define a clustering coefficient, CC (Cerasti & Treves, 2013), that measures, starting from the 373,248 randomly chosen Euler triplets and taking the N_{Best} distinct rotations for each of N units, the probability that two such triplets coincide (see Appendix A for the definition). For a

perfectly coherent rotation the N_{Best} rotations would be the very same triplets across units, hence $CC = 1/N_{\text{Best}}$. For a totally incoherent remapping, triplets would coincide at chance level, hence $CC = 1/373,248$. Figure 4d shows that, whether we take only units within the same cluster or in the entire population, the clustering coefficient has intermediate values, scaling approximately as $1/\sqrt{N_{\text{Best}}}$, for N_{Best} small—corresponding to a horizontal line in Figure 4d. For larger N_{Best} , a steeper decrease prevails, presumably because the population remains less than fully coherent also when allowing for looser remapping correspondence. In Appendix B we show that an intermediate scaling can be expected from a simple analytical model. A direct numerical evaluation of the mean field formulation for the clustering coefficient (see Equation B4 in Appendix B) yields a behavior in very good agreement with the one observed in simulations (Figure 4d, black continuous line). Here the deflection from the $1/\sqrt{N_{\text{Best}}}$ scaling appears to happen for larger values of N_{Best} (lying outside the right margin of the plot), presumably as an effect of the mean field approximation. The model does not seem to predict, however, an exact square root scaling, and it remains unclear to us whether the exponent $\beta \approx 1/2$ that we find to characterize $CC \approx 1/N_{\text{Best}}^\beta$ (for N_{Best} small) is fundamental or a mere coincidence.

An extensive set of simulations with varying overall interaction strength among the units (through a prefactor) indicates that the absolute value of the clustering coefficient depends mildly on the prefactor, but its scaling exponent $\beta \approx 1/2$ for N_{Best} small is the same (not shown), and the clustering coefficient is constant at $1/373,248$ only when the prefactor is strictly zero.

3.4 | Toward ecological plausibility

The full spherical environment may be approached with an appropriate experimental set-up (M. Mehta, personal communication (June 2019); by using spherical virtual reality, see also Aghajan et al., 2015), but is far from those in which rodents have evolved in the wild. The sphere does not include several features of, for example, the systems of burrows rodents dig as their homes. To begin considering the relevance of such features to grid maps, we start here with two: the presence of boundaries and the pull of gravity. By altering exploration and navigation behavior, both these features are expected to have at least an indirect influence on spatial codes, also in curved environments.

The effect of a *boundary* can be appreciated already by simply slicing a sphere in two halves, and running simulations of noninteracting units in one hemisphere (Figure S4a,b). If the cut were to be randomly oriented with respect to a perfect soccer-ball grid, one would expect definite proportions of the two hemisphere patterns in Figure S4a. In particular, the bottom arrangement with three fields around the pole should occur with $q = 28.6\%$ of the units, as can be calculated from the exact formula

$$q = 10 \left[1 - \left(\frac{3}{\pi} \right) \arctan(\varphi) \right] \quad (3)$$

here $\varphi = 1.618$ is the golden ratio. In simulations, however, it occurs with about half that frequency, $q = 14 \pm 5\%$ (Figure S4b), as the fields

of individual units tend to form away from the border. This also distorts the “pentagonal gridness” of each map on the hemisphere. Further, the presence of a hemispheric bump or cavity in an otherwise flat environment distort the hexagonal gridness of the fields near the boundary with the hemisphere. We have quantified this effect by running simulations on a flat ring which may contain either a circular hole or a hemisphere, and measuring the standard grid score on the flat part of would be grid units, this time interacting through recurrent collaterals. The presence of the (curved) hemisphere halved the average score, with respect to simulations run around a hole, from 0.39 ± 0.16 to 0.19 ± 0.07 (Figure S4c,d).

Further grid distortions appear if gravity is present. In the adaptation model, they are due to the unequal exploration of different latitudes on the hemisphere, which obviously deviates in opposite directions from an even sampling, depending on whether the hemisphere is set as a hill or as a valley. In Figure S5 we show examples of simulations run on a hill, in which trajectories were determined by adding to the standard algorithm, generating a random movement vector at each time step, a downward bias in speed (40% faster) and in turn selection, to model the downward gravity pull (see Appendix A for details). As can be seen, grid fields tend to cluster around the equator (Figure S5a,b). Interestingly, the resulting decrease in correlation with the perfect soccer ball field distribution is similar for non-interacting grids (Figure S5c) and for interacting ones (Figure S5d), which as discussed above already deviate more from the perfect arrangement also in the absence of gravity.

4 | DISCUSSION

The simulation of our adaptation model, allowing for collateral interactions among the units, indicates a radically different nature of the grid code on the sphere. The same interactions which on a plane suppress fluctuations and lead to a collapse into a smooth continuous attractor recruiting all units with the same or similar grid spacing (Si et al., 2012; Urdapilleta, Si, & Treves, 2017), on a sphere lead to a hierarchy of effects on single-unit maps that

1. are distorted from the available symmetric “soccer ball” field arrangement
2. are forced into clusters of units with approximately overlapping fields
3. remap to a different sphere with only partial coherence, even within clusters.

We argue that such effects are due to the unresolved conflict between regular tessellation at the single-unit level (due to adaptation dynamics) and global coherence at the population level (induced by recurrent connections). Unlike the planar case, where these two aspects can coexist in the same grid code, spherical geometry only allows for a compromise solution, where both regular tessellation and population coherence are only partially attained. We regard these as predictions that could be validated or falsified by experiments which

are doable in rodents, even though they may require ad hoc arrangements to allow for the slow emergence, possibly only during a 2-week developmental period, in rats (Langston et al., 2010), of a stable set of 2D maps.

Alternative models of grid map formation may lead to different predictions, but we would not know how, and are not currently aware of attempts by others, to extend existing models, for example, the oscillator interference model (Burgess, Barry, & O'Keefe, 2007) or the continuous attractor model (Burak & Fiete, 2009), to work on a sphere. Crucially, the continuous attractor model is based on the compatibility, contingent to Euclidean (in 2D, planar) spaces, of a single-unit hexagonal pattern, potentially extended to tile environments of any size, with congruent phase-offsets in different units. While this model can account for several grid cell properties and for their rapid manifestation, as observed in laboratory experiments, it appears that its theoretical premises make it inapplicable to environments of nonzero curvature. The regular phase offset that allows to project the activity from the "cell layer" envisaged by the model onto real space is just not possible with spherical geometry, leading to a loss of coherence in the activity of different units. It is also unclear to us how the oscillator interference model could be applied to curved environments.

A sphere is of course an even more artificial rearing environment than a flat box, but we believe that it may help capture a fundamental trait of grid cell coding, by pointing at those properties of grid cells often assumed to be universal but in fact stemming from the use of flat, bounded environments. The qualitative characteristics 1, 2, and 3 may be general to any curved environment, and they can be contrasted with the character of grid cell activity in rodents reared in standard laboratory conditions. In this sense, a sphere may be closer than a plane to the ecological condition of a Norway rat system of burrows (Calhoun, 1963). Grid cell representations may be presumed to have evolved to be relevant to rodents living in the wild.

In humans, the same fMRI hexagonal signature that has been hypothesized to reflect grid cell activity in a virtual reality navigation task (Doeller, Barry, & Burgess, 2010) has later been reported when subjects "move" in a 2D space of drawings (Constantinescu, O'Reilly, & Behrens, 2016), raising the issue of whether hexagonal symmetry may characterize even abstract conceptual spaces, when described by assigning two dimensions (Bellmund, Gärdenfors, Moser, & Doeller, 2018). Our model suggests that this may occur only around locations that are either flat a priori, or where curvature has been ironed out, perhaps by extensive training.

From a complex systems point of view, it is remarkable how curvature opens up a scenario different from that of a strictly regular, periodic 2D tessellation. In a separate study, we have already argued how the extension of such 2D tessellation to a 3D crystal, a scenario potentially relevant to bats and other animals navigating through 3D volumes, is in fact implausible, because of the time scales involved (Stella & Treves, 2015). Here, we make the case that also navigation on 2D manifolds embedded in 3D Euclidean space (such as tree-branches or multi-store buildings) might be associated with a "broken" grid cell representation, retaining only part of the planar symmetry.

In the new scenario, a network of grid units "behaves" more like a disordered system than like a crystal. The approximate inverse square root scaling of the clustering coefficient of the rotations, under remapping, reminds us of the partial coherence of a physical system with impurities, where some interactions are perforce "frustrated" (Mézard, Parisi, & Virasoro, 1987). When interactions are short-range, local coherence may survive, avoiding the impurities, somewhat like grid maps away from objects placed in a flat environment (Boccaro, Nardin, Stella, O'Neill, & Csicsvari, 2019; Hoydal, Skytoen, Moser, & Moser, 2019). A prevailing nonzero curvature is more akin, however, to a system with impurities and long-range interactions, where disorder affects even the shortest organizational scale. Such systems, not unlike human society, can offer only partial coherence.

Partial coherence, together with the realization that a network of grid cells may be endowed with a significant storage capacity (Spalla, Dubreuil, Rosay, Monasson, & Treves, 2019), in a sense brings back grid cells to the fold of memory systems, next to the place cells, with their multiple charts (Battaglia & Treves, 1998; Samsonovich & McNaughton, 1997). For years, it has been thought that grid cells may afford long-distance path integration (Fuhs & Touretzky, 2006; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006). Partial coherence, however, limits accurate path integration to short distances. Together with the emerging observation that mEC outputs to the cortex, mainly from layer Va, include virtually no grid-cell signal (A. Egorov and D. Rowland, personal communication), this weakens the theory that mEC operates as a sort of spatial computer, and suggests instead that grid maps are one input that helps set up the spatial component of hippocampal memory representations. Alternative sets of coactivity relations stored on the same synaptic connections, as well as curvature, act on the currently active grid representation as "quenched" disorder, and coexisting with such spin-glass-like disorder appears to be the ultimate challenge for memory systems in the brain (Treves, 2009).

ACKNOWLEDGMENTS

We are grateful to Matteo Toso, who modeled the habitat of the Norway rat, to YuQiao Gu, who carried out some of the early simulations, and to colleagues in the EU "Gridmap" and HFSP RGP0057/2016 collaborations, which also partially funded this study. This report was in part drafted during a 2018 visit by AT to the KITP, thereby also supported in part by NSF Grant No. PHY-1748958, NIH Grant No. R25GM067110, and the Gordon and Betty Moore Foundation Grant No. 2919.01.

ORCID

Federico Stella  <https://orcid.org/0000-0001-9439-3148>

Eugenio Urdapilleta  <https://orcid.org/0000-0002-1656-3188>

Yifan Luo  <https://orcid.org/0000-0001-6137-6536>

Alessandro Treves  <https://orcid.org/0000-0001-7246-5673>

REFERENCES

- Aghajan, Z. M., Acharya, L., Moore, J. J., Cushman, J. D., Vuong, C., & Mehta, M. R. (2015). Impaired spatial selectivity and intact phase precession in two-dimensional virtual reality. *Nature Neuroscience*, 18(1), 121–128.
- Battaglia, F. P., & Treves, A. (1998). Attractor neural networks storing multiple space representations: A model for hippocampal place fields. *Physical Review E*, 58(6), 7738–7753.
- Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415), eaat6766.
- 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.
- Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2), 193–205.
- Burak, Y. (2014). Spatial coding and attractor dynamics of grid cells in the entorhinal cortex. *Current Opinion in Neurobiology*, 25, 169–175.
- Burak, Y., & Fiete, I. R. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Computational Biology*, 5(2), e1000291.
- Burgess, N., Barry, C., & O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9), 801–812.
- Calhoun, J. B. (1963). Ecology and sociology of the Norway rat (US Public Health Service Publication No. 1008).
- Cerasti, E., & Treves, A. (2013). The spatial representations acquired in CA3 by self-organizing recurrent connections. *Frontiers in cellular neuroscience*, 7, 112.
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. (2016). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292), 1464–1468.
- Derdikman, D., Whitlock, J. R., Tsao, A., Fyhn, M., Hafting, T., Moser, M. B., & Moser, E. I. (2009). Fragmentation of grid cell maps in a multicompartiment environment. *Nature Neuroscience*, 12(10), 1325–1332.
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657–661.
- 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., Molden, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688), 1258–1264.
- Fyhn, M., Hafting, T., Treves, A., Moser, M. B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(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), 801–806.
- Harvey, C. D., Collman, F., Dombeck, D. A., & Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, 461(7266), 941–946.
- Kropff, E., & Treves, A. (2008). The emergence of grid cells: Intelligent design or just adaptation? *Hippocampus*, 18(12), 1256–1269.
- Kruege, I. U. (2016). *Ready, steady, go! Development of spatial representations in the rat*. (Doctoral dissertation 320). Trondheim: NTNU. Retrieved from <http://hdl.handle.net/11250/2423621>
- Kruege, I. U., Wernle, T., Moser, E. I., & Moser, M. B. (2013). Grid cells of animals raised in spherical environments. *Society for Neuroscience Abstracts*, 39, 769.14.
- Langston, R. F., Ainge, J. A., Couey, J. J., Canto, C. B., Bjerknes, T. L., Witter, M. P., ... Moser, M. B. (2010). Development of the spatial representation system in the rat. *Science*, 328(5985), 1576–1580.
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, 7(8), 663–678.
- Mézard, M., Parisi, G., & Virasoro, M. (1987). *Spin glass theory and beyond: An introduction to the replica method and its applications* (Vol. 9). Singapore: World Scientific Publishing Company.
- Høydal, Ø. A., Skytøen, E. R., Andersson, S. O., Moser, M. B., & Moser, E. I. (2019). Object-vector coding in the medial entorhinal cortex. *Nature*, 568(7752), 400.
- 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.
- Si, B., Kropff, E., & Treves, A. (2012). Grid alignment in entorhinal cortex. *Biological Cybernetics*, 106(8–9), 483–506.
- Si, B., & Treves, A. (2013). A model for the differentiation between grid and conjunctive units in medial entorhinal cortex. *Hippocampus*, 23(12), 1410–1424.
- Spalla, D., Dubreuil, A., Rosay, S., Monasson, R., & Treves, A. (2019). Can grid cell ensembles represent multiple spaces? *BioRxiv*, 527192.
- Stella, F., Si, B., Kropff, E., & Treves, A. (2013). Grid cells on the ball. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(03), P03013.
- Stella, F., & Treves, A. (2015). The self-organization of grid cells in 3D. *eLife*, 4, e05913.
- Stemmler, M., Mathis, A., & Herz, A. V. (2015). Connecting multiple spatial scales to decode the population activity of grid cells. *Science Advances*, 1(11), e1500816.
- Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M. B., & Moser, E. I. (2012). The entorhinal grid map is discretized. *Nature*, 492(7427), 72–78.
- Treves, A. (2009). Spatial cognition, memory capacity and the evolution of mammalian hippocampal networks. In L. Tommasi, M. A. Peterson, & L. Nadel (Eds.), *Cognitive biology: Evolutionary and developmental perspectives on mind, brain, and behavior Vienna Series in Theoretical Biology* (pp. 41–59). Cambridge, MA: MIT Press.
- Urdapilleta, E., Troiani, F., Stella, F., & Treves, A. (2015). Can rodents conceive hyperbolic spaces? *Journal of the Royal Society Interface*, 12(107), 20141214.
- Urdapilleta, E., Si, B., & Treves, A. (2017). Selforganization of modular activity of grid cells. *Hippocampus*, 27(11), 1204–1213.
- Yoon, K., Buice, M. A., Barry, C., Hayman, R., Burgess, N., & Fiete, I. R. (2013). Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nature Neuroscience*, 16(8), 1077–1084.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Stella F, Urdapilleta E, Luo Y, Treves A. Partial coherence and frustration in self-organizing spherical grids. *Hippocampus*. 2019;1–12. <https://doi.org/10.1002/hipo.23144>

APPENDIX A: MODEL

Additional aspects of the model, besides those reported in the main text:

The input to mEC unit i at time t is given by

$$h_i(t) = \sum_j W_{ij}(t)r_j(t) \quad (\text{A1})$$

The weight W_{ij} connects input unit j to mEC unit i . We assume that at the time the mEC units develop their maps, spatially modulated or place cell-like activity is already present, either in parahippocampal cortex or in the hippocampus. The network model works in the same way with any kind of spatially modulated input, but the place-cell assumption reduces the averaging necessary for learning. Each input unit activity in space is modeled as a Gaussian place field centered at preferred position x_{j0}

$$r_j(t) = \exp\left[-\frac{\|x(t) - x_{j0}\|^2}{2\sigma_p^2}\right] \quad (\text{A2})$$

where $x(t)$ is the position at time t of the simulated rodent, $\sigma_p = 0.05$ m is the width of the field and $\|a-b\|$ is the great-circle distance on the sphere.

Single-unit dynamics

The firing rate $\Psi_i(t)$ of mEC unit i is determined by a nonlinear transfer function

$$\Psi_i(t) = (\pi/2)\arctan[g(t)(\alpha_i(t) - \mu(t))]\Theta(\alpha_i(t) - \mu(t)), \quad (\text{A3})$$

which is normalized to have maximal firing rate equal to 1 (in arbitrary units), while $\Theta(\cdot)$ is the Heaviside function. The variable $\mu(t)$ is a threshold while $\alpha_i(t)$ represents the adaptation-mediated input to unit i . It is related to $h_i(t)$ as follows:

$$\begin{aligned} \alpha_i(t) &= \alpha_i(t-1) + b_1[h_i(t-1) - \beta_i(t-1) - \alpha_i(t-1)] \\ \beta_i(t) &= \beta_i(t-1) + b_2[h_i(t-1) - \beta_i(t-1)], \end{aligned} \quad (\text{A4})$$

where β_i has slower dynamics than α_i , with $b_2 = b_1/3$, $b_1 = 0.1$ (in a continuous formulation, the b coefficients become rates, in units of $[\Delta t]^{-1}$). These adaptive dynamics make it more difficult for a neuron to fire for prolonged periods of time, and correspond to the kernel K considered in the analytical treatment (Kropff & Treves, 2008). The gain $g(t)$ and threshold $\mu(t)$ are iteratively adjusted at every time step to fix the mean activity $a = \sum_i \Psi_i(t)/N_{\text{mEC}}$ and the sparsity $s = (\sum_i \Psi_i(t))^2 / (N_{\text{mEC}} \sum_i \Psi_i(t)^2)$ within a 10% relative error bound from prespecified values, $a_0 = 0.1$ and $s_0 = 0.3$, respectively. If k is indexing the iteration process:

$$\begin{aligned} \mu_t(k+1) &= \mu_t(k) + b_3[a(t) - a_0] \\ g_t(k+1) &= g_t(k) + b_4[s(t) - s_0], \end{aligned} \quad (\text{A5})$$

$b_3 = 0.01$ and $b_4 = 0.1$ are also rates, but in terms of intermediate iteration steps. a^k and s^k are the values of mean activity and sparsity determined by $\mu_t(k)$ and $g_t(k)$ in the intermediate iteration steps. The iteration stops once the gain and threshold have been brought within

the 10% error range, and the activity of mEC units are determined by the final values of the gain and threshold.

Synaptic plasticity

The learning process modifies the strength of the feed-forward connections according to a Hebbian rule

$$W'_{ij}(t) = W_{ij}(t) + \varepsilon [\Psi_i(t)r_j(t) - \langle \Psi_i(t) \rangle \langle r_j(t) \rangle] \quad (\text{A6})$$

with a rate $\varepsilon = 0.002$. $\langle \Psi_i(t) \rangle$ and $\langle r_j(t) \rangle$ are estimated mean firing rates of mEC unit i and place unit j that are adjusted at each time step of the simulation

$$\begin{aligned} \langle \Psi_i(t) \rangle &= \langle \Psi_i(t-1) \rangle + \eta [\Psi_i(t) - \langle \Psi_i(t-1) \rangle] \\ \langle r_j(t) \rangle &= \langle r_j(t-1) \rangle + \eta [r_j(t) - \langle r_j(t-1) \rangle] \end{aligned} \quad (\text{A7})$$

With $\eta = 0.05$ a time averaging factor. After each learning step, the $W'_{ij}(t)$ weights are normalized into unitary norm

$$\sum_j (W_{ij}(t))^2 = 1. \quad (\text{A8})$$

Head direction input

Head direction (HD) on the sphere is defined as the angle between a vector and the vector pointing toward the north pole. With the addition of HD modulation and collateral connections, the overall input to unit i for the interacting case is:

$$h_i(t) = f_{\theta_i}(\omega(t)) \left[\sum_j W_{ij}(t)r_j(t-1) + \rho \sum_k J_{ik} \Psi_k(t-\tau) \right] \quad (\text{A9})$$

with $\rho = 0.2$ a factor setting the relative strength of feed-forward $W_{ij}(t)$ and collateral weights J_{ik} , and $\tau = 25$ steps a delay in signal transmission, as discussed by Si et al. (2012). The multiplicative factor $f_{\theta_i}(\omega(t))$ is a tuning function which is maximal when the current direction of the animal movement $\omega(t)$ is along the preferred direction θ_i assigned to unit i

$$f_{\theta}(\omega) = c + (1-c)\exp[\nu(\cos(\theta - \omega) - 1)] \quad (\text{A10})$$

where $c = 0.2$ and $\nu = 0.8$ are parameters determining the minimum value and the width of the cell tuning curve. Preferred head directions are randomly assigned to mEC units and they uniformly span the 2π angle.

Clustering coefficient

Given q cells and taking the first N_{best} Euler rotations from each, the clustering coefficient (CC) was defined as:

$$\text{CC} = \left[\frac{1}{(qN_{\text{best}})^2 - qN_{\text{best}}} \sum_{ij} \sum_{mn} \exp[-d(\theta_{i,m}, \theta_{j,n})/\xi] \right] \quad (\text{A11})$$

with $i \neq j$ and where $d(\Theta_i, m, \Theta_j, n)$ is the distance between two three-dimensional rotations; and ξ is set to 5° so that only nearly coinciding rotations contribute to the sum.

Grid field definition and properties

Individual fields for each developing unit were identified as continuous portions of the spherical surface where the unit firing rate was above two times the average firing rate computed over the entire environment. Field size was defined as the number of bins passing the threshold in each continuous region, field height as the maximum firing rate within the continuous region, and the field ellipticity as the ratio between the radii of a circle circumscribed to the field and a circle inscribed in the field.

Effects of gravity on the animal trajectory

To simulate the change in movement statistics due to the presence of gravity, we modulated the generation of random trajectories in two ways.

1. We assumed that downward movement is executed at greater speed than upward movement. Therefore the animal speed was modulated depending on its running direction as $v_g = v(1 - \zeta \cos(\theta_z))$ where θ_z is the angle between the direction of motion and the vertical axis (0 when pointing upward), and ζ is a parameter regulating the strength of the gravitational effect.
2. We also applied a constant, downward pull to the direction of motion by applying a bias in the step-wise choice of a new running direction. Downward turns were favored by implementing a Metropolis Markov Chain that only accepted upward turns with a certain probability. Namely, a turn was rejected when $u < \alpha$ where u is a uniform random number on $[0,1]$ and $\alpha = \exp[\zeta(\theta_z^{t+1} - \theta_z^t)]$. θ_z is the angle with respect to the z -axis (0 when pointing completely upward), ζ is a parameter regulating the strength of the bias, taken here to numerically coincide with that of point 1).

APPENDIX B: STATISTICAL ANALYSIS

Let us consider N units that have developed grid representations on a sphere A, and now develop also grid representations on another sphere B. For every triplet of Euler angles (ϕ, θ, ψ) one can define the overlap (or spatial Pearson correlation) C_i between the representation of unit i in A and that in B rotated by (ϕ, θ, ψ) . Assume that $-1 < C_i < 1$. Then define $C_{\text{mean}}(\phi, \theta, \psi)$ as the mean $C_i(\phi, \theta, \psi)$ across all units, or the units in a cluster. Of course, $-1 < C_{\text{mean}}(\phi, \theta, \psi) < 1$ as well (in practice its range is much more restricted, if different units do not coincide in their "best rotations"). Now position all Euler triplets along the x -axis given by their C_{mean} value, and define $f(x)$ as their density (density of angles) along the axis. That is, if one considers a total of N_{angles} , there are $N_{\text{angles}} f(x) dx$ of them between x and $x + dx$. In Figure S3a we show the C_{mean} distribution for an entire population (dashed line) and separately for each of eight grid unit clusters (colored lines).

Assume now that among all N_{angles} angles, we pick for each unit N_{best} of them, those that have the highest C_i value. How will all the $N_{\text{best}} \times N_{\text{units}}$ angles be distributed, on average, in terms of $f(x)$, the C_{mean} -ordered histogram? On average, they will concentrate more at higher C_{mean} values, at the very least because each unit gives a $1/N$ contribution to C_{mean} ; but possibly more concentrated than that. How much they concentrate is critical in order to determine the clustering coefficient CC, which measures simply how many of the N_{best} angles (what fraction) coincide among pairs of distinct units. We assume then that, at least within a cluster,

1. one can write:

$$C_i = C_{\text{mean}} + \eta_i \quad (\text{B1})$$

where η_i is a form of "noise," that is, the combined effect of all other factors independent and unrelated to C_{mean} . Note that this decomposition is a strong assumption.

2. this "noise" is normally distributed, with a width $\sigma(N_{\text{best}}, \text{cluster})$ that is the same across units in a cluster.

If we denote with $b(x; N_{\text{best}})$ the average fraction of N_{best} angles, among the N_{angles} $f(x)$ present at a given C_{mean} -value x , such that, within a cluster,

$$N_{\text{best}} = N_{\text{angles}} \int f(x) b(x; N_{\text{best}}) dx \quad (\text{B2})$$

and with $x_b(N_{\text{best}}, \text{cluster})$ the value of x such that this average fraction is $1/2$, with these assumptions one has that $b(x)$ can be expressed as the complementary error function

$$b(x) \approx (1/2) \text{erfc}\left((x_b - x)/\sigma\sqrt{2}\right) \quad (\text{B3})$$

which Figure S3b shows is not a bad approximation, if one allows x_b and hence $b(x)$ to depend on both the cluster and N_{best} .

We make now the additional, critical (mean-field) assumption that

3. the clustering coefficient, CC, at least within each cluster, is only determined by the average density $b(x; N_{\text{best}})$. Therefore, considering a generic pair of units,

$$CC = \left[N_{\text{angles}} / (N_{\text{best}})^2 \right] \int f(x) b^2(x; N_{\text{best}}) dx \quad (\text{B4})$$

One may observe that CC is given by the extent of the overlap between the two distributions $f(x)$ and $b(x)$ (dropping for ease of notation its argument N_{best}). Equation (B2) and Figure S3 show that for small N_{best} the overlap is limited to the opposing tails of the two distributions. We can evaluate the goodness of the assumptions made so far by comparing the values of the clustering coefficient obtained

from the last equation to those obtained from the full analysis of simulation results. Equation B4 can be numerically evaluated by making use of the $f(x)$ computed from simulations (Figure S3a) and of the parameters for $b(x)$ obtained from Gaussian fits (Figure S3b-d). The resulting mean field approximation curve shows a remarkable similarity with simulation results (Figure 4d). If we proceed and make the final assumption that

4. $f(x)$ has a quasi-Gaussian upper tail

$$f(x) \approx k \exp\left[-(x-x_0)^2/2\xi^2\right] \quad (\text{B5})$$

with k a suitable factor and ξ the effective width of the tail, we can obtain an analytical estimate of the CC. The result is

$$\ln(\text{CC}(N_{\text{best}})) \approx -\ln(N_{\text{best}}) - (x_0 - x_b)^2 \sigma^2 / [2(\sigma^2 + \xi^2)(\sigma^2 + 2\xi^2)] \quad (\text{B6})$$

(keeping only the leading exponent). Both x_b and σ may depend on N_{best} , but Figure S3 indicates that within each cluster the dependence of σ is weak, while x_b shifts leftward as N_{best} increases: from Equation (B2) one can derive

$$\ln(N_{\text{best}}) \approx \ln(N_{\text{angles}}) - (x_0 - x_b)^2 / [2(\sigma^2 + \xi^2)] \quad (\text{B7})$$

This yields a scaling of $\ln(\text{CC}(N_{\text{best}})) \approx -\beta \ln(N_{\text{best}})$ with $0 < \beta < 1$, and the particular value $\beta = 1/2$ is obtained for $\sigma^2 \approx 2\xi^2$. We have no explanation for why this last relation appears to hold, approximately, for N_{best} small, as shown in Figure 4d.