Environmental deformations dynamically shift the grid cell spatial metric

Alexandra T Keinath[1]*, Russell A Epstein[1], Vijay Balasubramanian[2]*

[1] Department of Psychology, University of Pennsylvania

[2] Department of Physics, University of Pennsylvania

* Corresponding authors. VB is lead contact.

Conflict of interest: The authors declare no competing conflicts of interest.

Summary
We propose a unified explanation for the diverse distortions in time-averaged activity of grid and place cells observed after environmental deformations. In our account, input from border cells resets the spatial phase but not the spatial scale of grid cells, maintaining learned relationships between grid phase and environmental boundaries. A computational model implementing this mechanism reproduced several commonly observed experimental effects, including scale-dependent distortions in time-averaged grid fields after environmental deformation, and stretched, duplicated, and fractured place fields. Furthermore, this model predicted a striking new effect: dynamic, history-dependent 'shifts' in grid phase. We reanalyzed two classic datasets on grid rescaling and found clear evidence for such shifts, which have not previously been reported. These results invite a reconceptualization of the effects of environmental deformations on spatial representations – rather than rescaling the spatial metric of the cognitive map, as previously believed, alterations in environmental geometric may dynamically shift the map.

The cognitive map is thought to be a metric representation of space that preserves distances between represented locations [1,2]. Entorhinal grid cells are hypothesized to generate this metric by maintaining an internally-generated, path-integrated representation of space [3–8]. Results of environmental deformation experiments have led to the belief that this metric is fundamentally malleable [9–11]. In these experiments, neural activity is recorded as a rat explores deformed versions of a familiar environment where chamber walls have been stretched, compressed, removed, or inserted. Such deformations induce a number of distortions in the time-averaged activity of both grid [9,11] and hippocampal place cells [12–16]. Often described as 'rescaling', these distortions have been taken to suggest that the spatial metric of the cognitive map can be reshaped by altering environmental geometry [9,17,18]. Crucially, however, this interpretation assumes that the distortions observed in the time-averaged rate maps of these cells directly reflect corresponding changes to the underlying spatial code. Here, we propose and test an alternative mechanism which challenges this interpretation and instead indicates that during environmental deformations the grid cell spatial metric does not rescale, but instead undergoes dynamical, history-dependent phase shifts.

We hypothesize, as others have [10,19–23], that border cells interact with grid cells in familiar environments to maintain learned relationships between grid phase and boundaries. Here we further propose that the same mechanism is responsible for inducing diverse and unexpected distortions observed in the time-averaged activity of grid and place cells during environmental deformations. Specifically, when an environment is deformed, border cell input continues to maintain learned boundary-grid phase relationships, inducing shifts in the grid pattern that depend on the most recent boundary encounter. When grid cell responses are averaged over time, these *boundary-tethered grid shifts* lead to apparent distortions of grid patterns, even though the underlying scale of the grid remains unchanged. The shifts are inherited by downstream place cells, producing heterogeneous place field changes. In this view, distortions observed in the time-averaged activity of place and grid cells are an artifact of averaging over paths originating from different boundaries; aligned by the most recently contacted boundary, the grid and place codes should appear shifted, not distorted.

To test this mechanism, we constructed a model (the *boundary-tethered* model) where the activity of a grid cell attractor network [24] is shaped by Hebbian-modified input from border cells [21]. Recent work has shown that such border cell input to grid cells can stabilize drift in the grid pattern [19,25], providing evidence for such interactions in familiar environments. The model also included a population of units corresponding to hippocampal place cells, whose responses were learned from grid unit output [26,27]. Our simulations show that during environmental deformations, modeled grid and place units reproduce a number of striking experimentally-observed behaviors: 1) when a familiar environment is rescaled, the firing patterns of large-scale grid units rescale to match the deformation, while the firing patterns of small-scale grid units do not [9,11]; 2) when a familiar environment is stretched, the fields of place units exhibit a mix of stretching, bifurcation, modulation by movement direction, and inhibition [12]; 3) when a familiar linear track is compressed, the place code is updated when a track end is encountered [13,28]; 4) when a new boundary is inserted in an open environment, place fields exhibit a mix of duplication, inhibition, and perseverance [14–16]. Our model further generated an additional unique prediction: a signature shift in grid and place field location, dependent on the most recently contacted boundary. To test this prediction, we reanalyzed the datasets from two previous environmental deformation experiments [9,11], and found compelling new evidence of *boundary-tethered grid shifts* in rodent grid cell activity. In sum, our results indicate that deformation of a familiar environment results in dynamic history-dependent shifts in grid and place field locations, rather than distortions of field shape. These results invite

a reconceptualization of how environmental deformations affect spatial representation and navigation more broadly.

Results

A model of border, grid, and place cell interactions

Previous recording studies in rodents have shown a number of striking effects of environmental deformations on the time-averaged firing maps of grid cells (e.g., scale-dependent rescaling [9,11]) and place cells (e.g., rescaling, bifurcation, modulation by movement direction, field duplication, and suppression [12–16]). Although numerous models of environmental deformations have been proposed to explain subsets of these results, including the influential boundary vector cell model [14,17,18,29–31], no single mechanism has been demonstrated to reproduce this set of results in its entirety. Thus, we sought a single circuit mechanism that could give a unified explanation for all these phenomena. To this end, we developed a spiking feedforward model of border, grid, and place cell interactions (Fig. 1; additional details in Methods).

The border population consisted of 36 units whose activity was designed to mimic the behavior of border cells [21]. (Throughout this paper, we use unit to refer to modeled data, and cell to refer to *in vivo* recorded data.) Each border unit was active only when a boundary was nearby, within 12 cm in a particular allocentric direction [19]. The preferred firing field of each border unit covered ~10% of the boundary length, and maintained proportional coverage if that boundary was deformed [21,25]. If a new boundary was inserted, the border unit was active at an allocentrically analogous location adjacent to the new boundary.

The grid population was subdivided into 5 modules, each consisting of a neural sheet of size 90 x 90 units. The internal connectivity and dynamics of each module was based on the attractor network model described in [24], and was identical across modules except for a single movement velocity gain parameter controlling the grid scale of each module. This parameter was adjusted to yield a geometric series of scales across modules (scale factor of 1.42), as observed experimentally [11] and explained theoretically [32,33]. In addition to these connections, each grid unit also received initially random excitatory input from all border units. These connections developed through experience via a Hebbian learning rule in which connections between coactive grid and border units were strengthened at the expense of connections from inactive border units [34].

The place population consisted of 64 units receiving initially random excitatory input from 500 random grid units. These connections also developed with experience via Hebbian learning [27,34]. In combination with uniform recurrent inhibition, these dynamics yield place-cell-like activity at the single unit level.

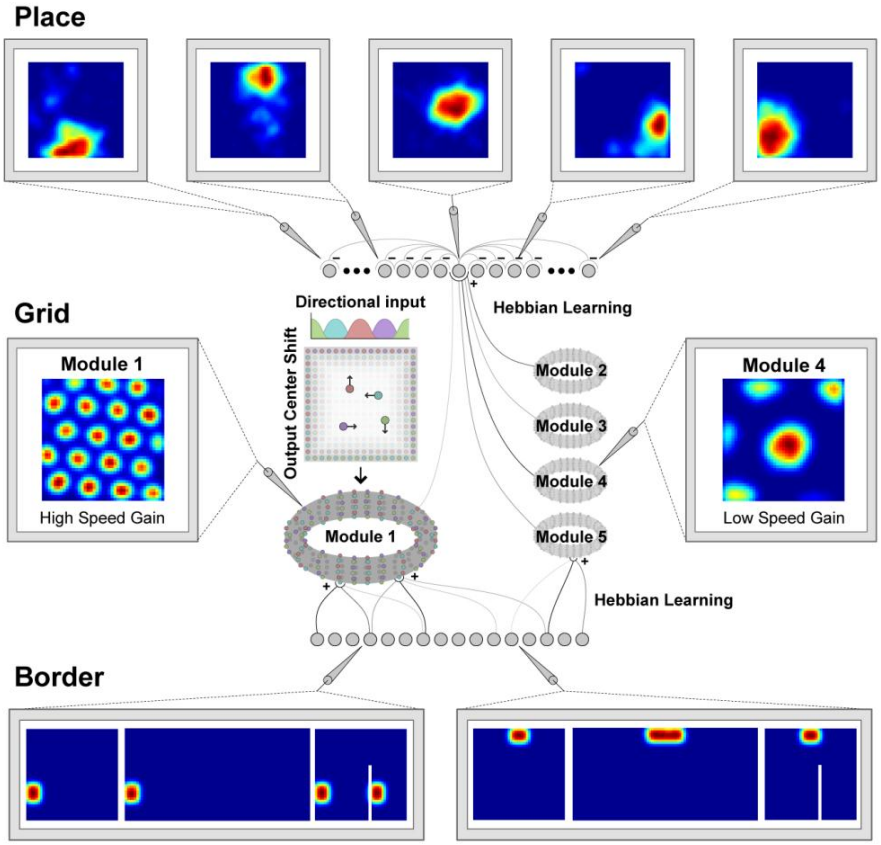

Figure 1. Schematic of the border, grid, and place cell network model. *See Methods for a complete description of model interactions. Note the behavior of border units during environmental deformations – stretching the environment stretches the border field and inserting a barrier duplicates the border field with the same allocentric relationship to both the original and inserted boundaries.*

Model grid units deform with the environment in a scale-dependent fashion.

Electrophysiological experiments have shown that rescaling a familiar environment can induce a corresponding rescaling of grid cell firing patterns, dependent on grid scale. To explore the effects of environmental rescaling on grid units, we first familiarized a naive virtual rat with a 150 cm x 150 cm square environment. During this familiarization period, the border-grid connectivity self-organized via Hebbian learning (see Methods). The rat then explored the familiar environment and deformed versions of this environment without new learning (chamber lengths between 75 cm to 225 cm in increments of 25 cm; chamber sizes chosen to match experiment [11]). Consistent with previous reports [9,11], we observed that these deformations induced rescaling of time-averaged rate maps in some grid modules (Fig. 2a). To quantify this module-dependent rescaling, we computed the grid rescaling factor required to stretch or compress the time-averaged rate maps in the deformed environment to best match the rate maps in the familiar environment, separately for each module. We found that the grid patterns of units in large-scale modules morphed with the environment, but pattern of units in small-scale modules tended not to (Fig. 2b). Precisely this behavior is observed experimentally [11]. These results demonstrate for the first time that connections from border cells are sufficient to induce

scale-dependent grid rescaling, without appealing to inputs that are module-specific or originate outside the medial entorhinal cortex as in [17,18,29,31,35].

The boundary-tethered model predicts that whether the rate map of a grid cell appears to rescale is not just an inherent property of its grid scale or module, but also of the size and type (stretch vs. compression) of the deformation. Preliminary evidence for this prediction comes from the observation of small-scale grid rescaling during less extreme [9] but not during more extreme compression deformations [11]. A detailed experimental exploration of the dependence of grid rescaling on deformation size and type would provide an additional strong test of our boundary-tethered model.

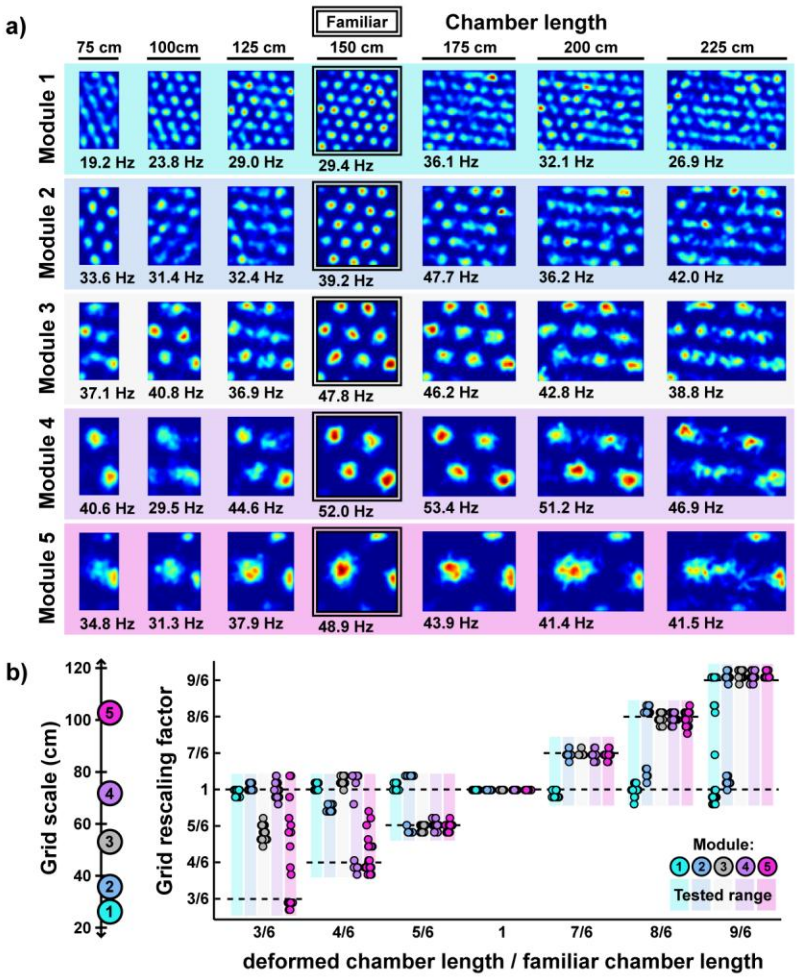

Figure 2. Grid unit response during deformation of a familiar open field environment. *a) Rate maps from five example grid units, one from each module, across all tested deformations. Colors normalized to the maximum peak for each set of rate maps. Peak firing rate noted below the lower left corner of each map. b) Grid rescaling factors that best match the familiar enclosure rate maps for 30 randomly chosen grid units from each Module when the familiar open (150 cm x 150 cm) environment is rescaled to various chamber lengths (see Methods). Each dot denotes the grid rescaling factor of a single grid unit, and the color indexes the module (grid scale of each module is indicated at left).*

Model place units distort heterogeneously during environmental deformations.

Electrophysiological experiments have shown that stretching a familiar environment induces a heterogeneous mix of responses in the time-averaged activity of place cells [12]. To explore the effects of stretching deformations on model place units, we began by familiarizing the naive virtual rat with a 61 cm x 61 cm square open environment, during which period the border-grid connectivity and grid-place connectivity self-organized via Hebbian learning. Following this familiarization, the virtual rat again explored the familiar environment, as well as a number of deformed environments without new learning (various chamber lengths between 61 cm and 122 cm, chamber widths 61 cm or 122 cm; chamber sizes chosen to match experiment [12]). During these deformations, we observed heterogeneous changes in the time-averaged rate maps of place units. A number of place units exhibited place field stretching in proportion to the rescaling deformation (Fig. 3ai). Other units exhibited place field bifurcations accompanied by progressively lower peak firing rates during more extreme deformations (Fig. 3aii). Finally, some units exhibited emergent modulation by movement direction (Fig. 3aiii), with place fields shifting 'upstream' of the movement direction. Precisely this mix of place field distortions is observed experimentally [12].

Electrophysiological experiments have also demonstrated that when a familiar linear track is compressed, the place code is updated when track ends are encountered [13]. We therefore examined the effects of compressing a familiar linear track on model place units. We first familiarized the naive virtual rat with running laps on a 161 cm long linear track, during which period the border-grid connectivity and grid-place connectivity self-organized via Hebbian learning. Following this familiarization, the virtual rat ran laps along both the familiar track and a number of compressed tracks without new learning (track lengths between 53 cm to 161 cm; lengths chosen to match experiment [13]). During laps on compressed tracks, place unit activity unfolded as if unaffected by the compression, no matter how extreme, until the opposing track end was reached. Once encountered, the place code previously active at this track end during familiarization reemerged (Fig. 3b), as observed experimentally [13]. In recording experiments, similar boundary-tethered updating persists in darkness indicating that such dynamics arise even in the absence of visual cues [28], a result consistent with the absence of explicitly visual input in our model. However, we note that in these recording experiments the particular transition point does differ depending on the availability of visual input, which may reflect changes to border cell activity [36] or the influence of additional mechanisms outside the scope of the boundary-tethered model [17,18].

Finally, electrophysiological experiments have shown that when a boundary is inserted in a familiar open environment, place fields exhibit a mix of duplication, suppression, and perseverance [14–16]. We explored the effects of inserting a new boundary on model place units. We first familiarized the naive virtual rat with a 65 cm x 65 cm square open environment, during which period the border-grid connectivity and grid-place connectivity self-organized via Hebbian learning. Following this familiarization, the rat explored, without new learning, the familiar environment and a deformed version of this environment containing an additional 40 cm long boundary adjacent to one wall and evenly dividing the space (chosen to match experiment [14]). Again, we observed heterogeneous changes in the time-averaged rate maps of place units. Some units exhibited place field duplication during boundary insertion (Fig. 3ci) while other units exhibited place field inhibition (Fig. 3cii). Still others persevered largely unaffected (Fig. 3ciii). Precisely this mix of responses is observed experimentally during boundary insertions [14–16].

Together, these results demonstrate that the boundary-tethered model can give rise to the heterogeneous behaviors of place cells observed across three unique environmental deformation paradigms.

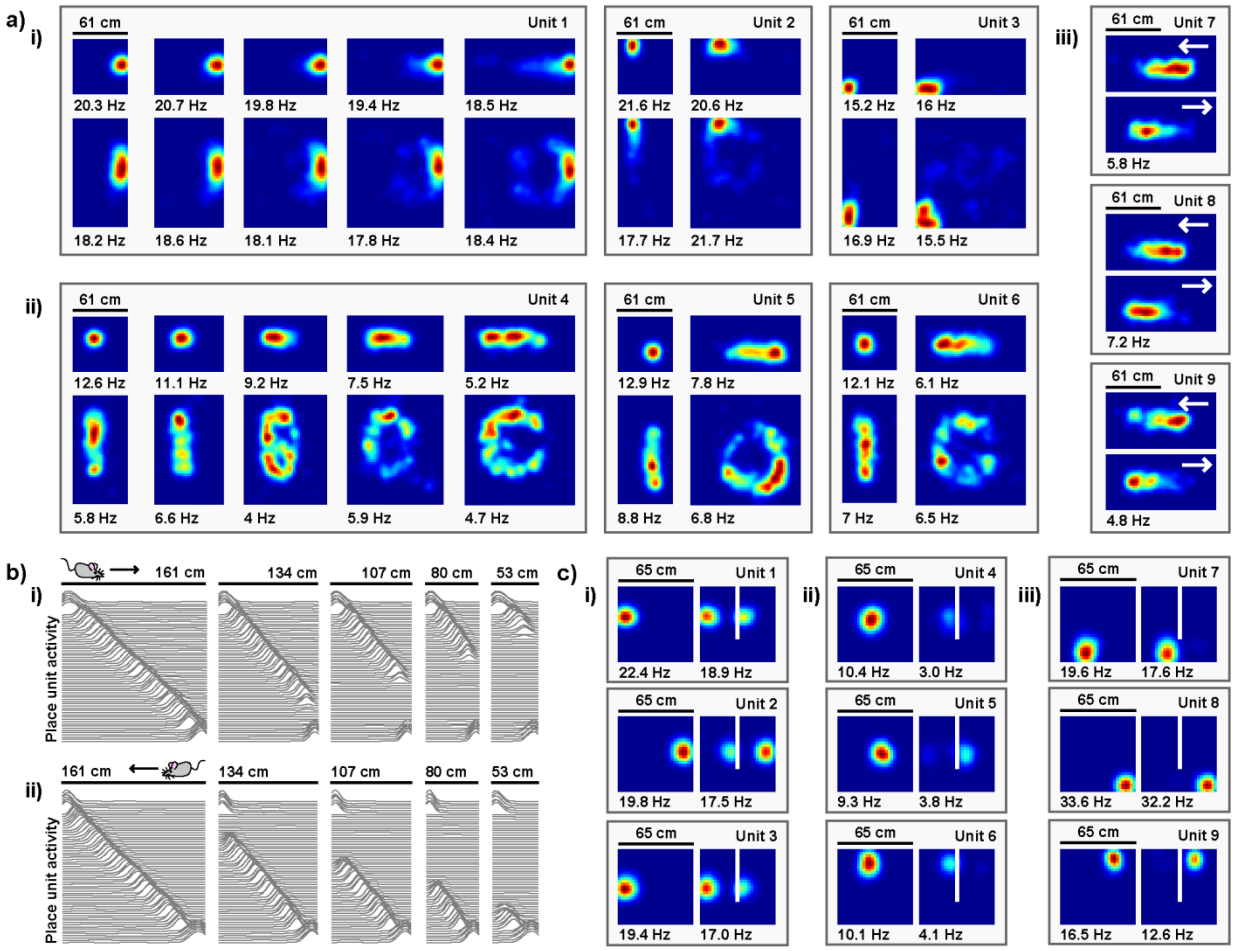

Figure 3. Place unit response during environmental deformations. *a) Place unit rate maps when a familiar open (61 cm x 61 cm) environment is stretched. Place fields exhibit (i) stretching, (ii) bifurcation, and (iii) emergent modulation by movement direction (indicated by white arrow). Colors normalized to the peak for each rate map. Peak firing rate noted below the lower left corner of each map. Note that peak firing rate tends to decrease with more extreme deformations for cells with place fields near the center of the environment. b) Place unit activity for all 64 place units during compressions of a familiar (161 cm) linear track, separated by (i) eastward and (ii) westward laps. Each line indicates the firing rate of a single place unit at each location across the entire track during movement in the specified direction, normalized to the familiar track peak rate. Units sorted by place field location on the familiar track. Note that, during compressions, the place code unfolds as if anchored to the beginning of the track until the end of the track is encountered, at which point the familiar end-of-track place units are reactivated. c) Place unit rate maps demonstrating a mix of place field (i) duplication, (ii) inhibition, and (iii) perseverance when a new boundary (white line) is inserted in a familiar open (65 cm x 65 cm) environment. Peak firing rate noted below the lower left corner of each map.*

Boundary-tethered grid shift underlies model grid and place unit distortions.

How do model interactions give rise to these grid and place unit distortions? During familiarization, Hebbian learning strengthens the connections from active border units to active grid units at the expense of connections from inactive border units (Fig. 4a; see Methods). Once familiarized, border unit activity reinstates the grid network state associated with the same pattern of border unit responses during familiarization. This grid reinstatement occurs even when border inputs are activated at a new location, such as when a new or displaced boundary is encountered. In a rescaled open environment, grid reinstatement leads to 'shifts' in the spatial phase of the grid pattern, such that the phase relative to the most recent border input matches the phase entrained during familiarization in the undeformed environment (Fig. 4b,c). Averaged over time, these shifts can resemble a rescaling of the grid pattern. Note that the boundary-tethered model thus implies that rescaling of grid maps is an epiphenomenon that results from time-averaging over dynamical shifts in deformed environments. This view departs from previous accounts that view grid rescaling itself as the fundamental phenomenon and propose mechanisms to directly reproduce this effect [18,35].

Why does apparent grid rescaling depend on grid scale and module identity in the boundary-tethered model? The grid representation is periodic, so border input can only reset the grid network state to within one grid period. If the deformation extent is less than the grid scale, then the different boundaries will reinstate different grid phases, yielding an apparently rescaled time-averaged grid pattern. When the deformation extent matches the grid scale, different boundaries will reinstate the same grid phase, yielding an apparently undistorted time-averaged grid pattern. When the deformation extent exceeds the grid scale, different boundaries will again reinstate different grid phases; thus the time-averaged grid pattern will again appear distorted. However, in the latter case, additional fields will appear (during stretches) or previously-observed fields will disappear (during compressions). Thus the resulting time-averaged grid pattern, although distorted, will not resemble a complete rescaling (Fig. S1). Additionally, because module identity is assigned primarily according to grid scale the appearance or absence of time-averaged rescaling in the model will appear to be module-specific as well. Note that the boundary-tethered model implies that environmental deformations affect the grid patterns of rescaling and non-rescaling cells alike, in contrast with other accounts in which the appearance or absence of rescaling is hypothesized to reflect a fundamental difference in function [11,18]. This model further implies that the appearance or absence of rescaling in simultaneously recorded modules is not evidence of a functional dissociation between modules, as previously claimed [11].

As for place cells, in this model place unit activity is constructed as a normalized, thresholded sum of grid unit input following [26]. Because of the boundary-tethered grid shifts induced during environmental deformations, the location of each place field will also shift, maintaining its spatial relationship to the most recently contacted boundary (Fig.4d). Critically, the likelihood of having most recently encountered a given boundary differs throughout an open environment: locations near the center of the chamber are more likely to have had an even distribution of previously encountered boundaries, while locations near a boundary are more likely to be visited following an encounter with that boundary (Fig. 4e). When averaged across time, these most recent boundary biases result in a mix of place field stretching (closer to boundaries) and bifurcation distortions (further from boundaries). Furthermore, the most recently encountered boundary is correlated with the direction of movement: the rat is more likely to have most recently encountered a given boundary when moving away from that boundary (Fig. 4f). Thus, if the rat is traveling eastward in a stretched environment, then the place field will

typically be tethered to the west wall; if the rat is traveling westward, then the field will typically be tethered to the east wall. Because the environment has been stretched, west wall-tethered fields will be shifted westward of east wall-tethered fields. Thus, boundary-tethered place field shift causes place fields to be displaced 'upstream' along the direction of movement. Finally, more extreme deformations of an enclosure lead to more extreme boundary-tethered shifts and less frequent convergence of grid inputs at the same location, and, thus, overall systematic decreases in the peak firing rate of grid and place units.

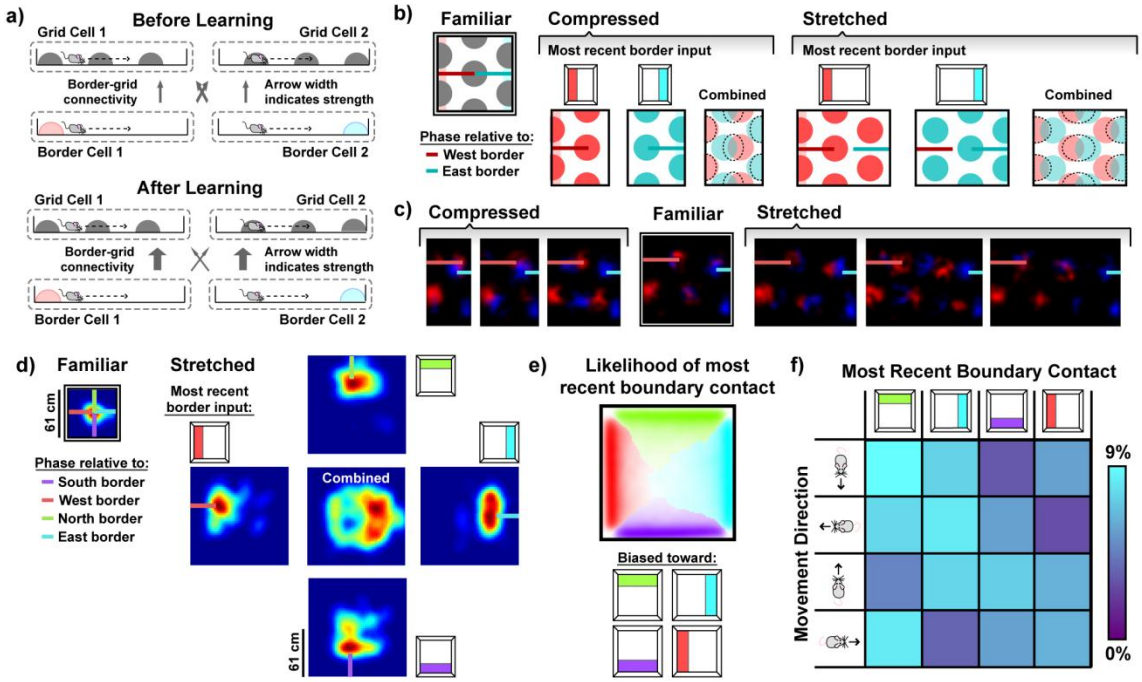

Figure 4. Boundary-tethered grid shift underlies model grid and place unit distortions. *a) During familiarization, Hebbian learning strengthens the connections between coactive border and grid cells, at the expense of non-coactive connections. b) During deformations, border input acts to maintain the previously learned relationship between grid phase and the most recent border input. c) Rate map of a Module 4 grid unit when the west border was the most recently contacted boundary (red), overlaid with the rate map of the same unit when the east border was the most recently contacted boundary (blue). The spatial phase relative to the most recent border input (indicated by red/blue bars) is preserved during all deformations. Thus the grid pattern is undistorted when separated by the most recent border input. d) Place fields shift to maintain their previously learned relationships relative to the most recent border input. e) Likelihood of having most recently contacted each border as a function of location in the square environment. Color saturation denotes the strength of the bias. Data from [11]. f) Each square in the grid depicts the joint probability distribution in a square environment of the indicated most recently contacted boundary and movement direction. Data from [11].*

When the rat is trained to run laps on a linear track, movement and the most recently contacted track end are constrained. Thus linear track compressions provide an especially clear view of boundary-tethered updating. Until a track end is encountered, grid and place unit activity unfolds according to path integration alone. When a track end is encountered, border input reinstates the grid network state and, in turn, the place network state that coincided with that

track end on the familiar track. This produces the upstream place field shifts observed in Fig. 3b.

Inserting a boundary in an open environment elicits identical border unit activity when either the old boundary or new boundary is nearby in the preferred allocentric direction, inducing boundary-tethered reinstatement of the grid network state at both locations (Fig. S2). This grid shift translates to a duplication of the place unit representation adjacent to the old and inserted boundaries. Because a new grid and thus place representation are now active around the inserted boundaries, the old representations previously active at this location in the familiar environment are no longer activated. This leads to an apparent inhibition of place units participating in the old representation. However, grid and place units that were active at locations distant from the duplicated boundaries will generally persevere unaffected.

Thus, in this model, boundary-tethered grid shift drives the diverse grid and place field distortions observed during geometric deformations.

Experimental observation of predicted boundary-tethered grid shifts

Above we showed that many previously-known grid and place cell distortions can be explained as arising from underlying grid shifts, providing an indirect test of the boundary-tethered model. Here, we directly test the novel prediction that the locations of both grid and place fields should shift to preserve their spatial relationship relative to the most recently contacted boundary. To this end, we reanalyzed data from two classic environmental deformation studies ([9] and [11]). In [9], rats were familiarized with either a 100 cm x 100 cm square or a 100 cm x 70 cm rectangular open environment, and then reintroduced to deformed and undeformed versions of these environments (i.e. all combinations of chamber lengths and widths of 70 cm or 100 cm), while the activity of grid cells was recorded (familiar square: 42 grid cells; familiar rectangle: 23 grid cells). In [11], rats were familiarized with a 150 cm x 150 cm square open environment, and then reintroduced to deformed (100 cm x 150 cm rectangular) and undeformed versions of this environment, while data were recorded from 51 grid cells.

To test for the predicted boundary-tethered shifts, we first separated the spiking data of each cell by the most recently contacted boundary, either the north, south, east or west, with contact defined as coming within 12 cm of the boundary [19]. From these data, we created four *boundary rate maps* which summarized the spatial firing pattern of the grid cell after contacting each boundary. Comparison of such rate maps, conditioned on contact with opposing boundaries (north-south vs. east-west), revealed clear examples of grid shift along deformed dimensions (Fig. 5a). To quantify shift separately for each dimension, we cross-correlated the corresponding boundary rate maps (i.e., north-south or east-west boundary rate map pairs; see Methods). Only pixels sampled after contacting both opposing boundaries were included. Next, we computed the distance from center of the cross-correlogram to the peak nearest the center (see Methods). This distance measures the relative shift between the opposing boundary rate maps. Even in a familiar environment, finite sampling noise will cause this measure of shift to be nonzero. Relative to this baseline, grid shift increased along deformed, but not undeformed, dimensions across deformation trials (Fig. 5b; Fig. S3a). Moreover, an increase in shift was observed even in specifically small-scale and non-rescaling grid cells (Fig. S4). This suggests, as predicted by our model, that grid shifts affect all grid cells alike regardless of whether the time-averaged rate maps appear to show rescaling.

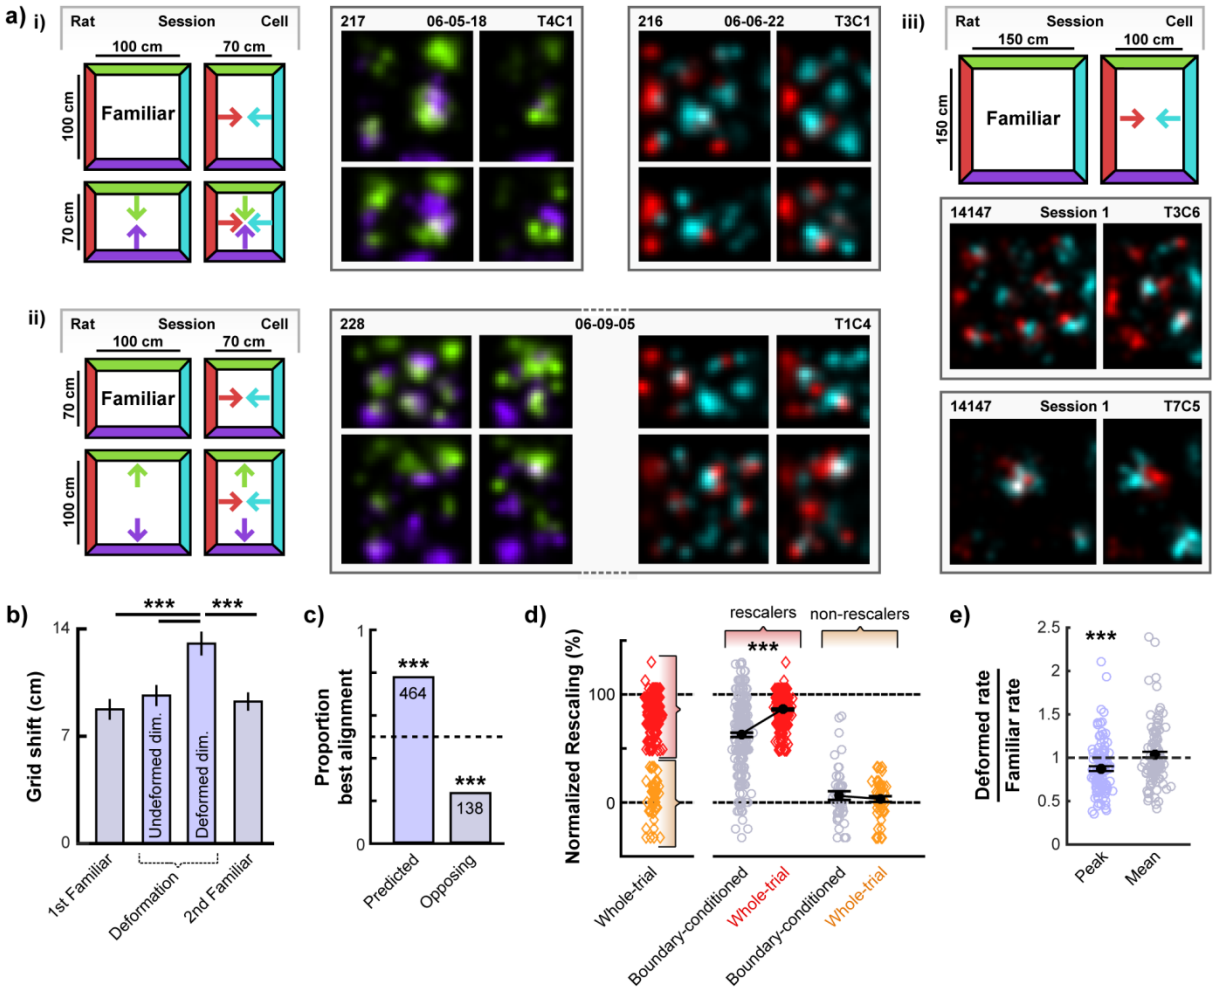

Figure 5. Boundary-tethered grid shifts are observed in the activity of recorded grid cells during environmental deformations. *a) Examples of grid shifts visible in 'boundary rate maps' (see text) created by dividing spiking data according to the most recently contacted boundary (indicated by rate map color). Maps organized by opposing north-south (green—purple) and east-west (blue—red) boundary pairs. Colored arrows indicate the directions of shifts predicted by our model during each deformation. b) Grid shift as measured by the relative phase between opposing boundary rate maps (see text), separately for deformed and undeformed dimensions (errors bars ±1 SEM). Significance markers denote the outcome of a 2-tailed paired t-test between the indicated conditions (1st familiar vs. deformed: $t(115)=4.9$, $p < 0.001$; undeformed vs. deformed: $t(115) = 3.9$, $p < 0.001$; 2nd familiar vs. deformed: $t(111) = 4.8$, $p < 0.001$; all other comparisons: $t < 1.24$, $p > 0.22$, n.s.). c) Proportion of trials for which each boundary rate map was best matched with its familiar environment rate map when aligned by the most recently contacted boundary (as predicted by the boundary-tethered model) vs. the opposing boundary (counts shown within the bars). Significance markers denote the outcome of a 2-tailed sign test versus 50%. d) Boundary-conditioned grid rescaling factors and classic whole-trial grid rescaling factors normalized to range from no rescaling (0%) to complete rescaling (100%), split by the extent of whole-trial grid rescaling (split at 45% rescaling). Each dot denotes rescaling along one dimension (north-south or east-west) during one deformation trial (errors bars ±1 SEM). Significance markers denote 2-tailed paired t-test between indicated conditions (Boundary-conditioned versus whole-trial, rescalers: $t(256) = 12.4$, $p < 0.001$; non-rescalers: $t(44) = 0.8$, $p = 0.423$, n.s.). e) Ratio of average peak and mean firing rates in deformed environments divided*

*by average rate in the familiar environment (errors bars ±1 SEM). Significance marker denotes the outcome of a 2-tailed t-test versus 1 (peak rate ratio: $t(115) = 4.5$, $p < 0.001$; mean rate ratio: $t(115) = 1.1$, $p = 0.257$). ***$p < 0.001$*

Next we asked whether the grid pattern in each boundary rate map maintained its spatial phase with the corresponding boundary, as the boundary-tethered model predicts. To address this question, we compared each of the boundary rate maps to the whole-trial familiar environment rate map, while varying the alignment of the two compared maps along the deformed dimension. If the spatial relationship relative to the most recently contacted boundary is preserved, then each boundary rate map should be most similar to the familiar environment rate map when the two maps are aligned by the corresponding boundary. If, on the other hand, reshaping a familiar environment rescales the grid pattern uniformly, then the familiar and boundary rate maps should be equally well aligned by the corresponding or opposite boundary. Consistent with the boundary-tethered model, we found that a deformed environment boundary rate map and the familiar environment map were usually best aligned by the corresponding boundary, rather than the opposite boundary (Fig. 5c; Fig. S3b).

The boundary-tethered model further predicts that the appearance of rescaling is an artifact of averaging over trajectories originating from different boundaries. Thus, the appearance of rescaling should be minimized when the data are divided according to the most recently contacted boundary. In contrast, if grid response rescaled, similar rescaling should be observed regardless of whether or not data are divided according to the most recently contacted boundary. To test these predictions, we computed the grid rescaling factor between the familiar rate map and each deformed-dimension boundary rate map, aligned by the corresponding boundary, and then averaged across these factors to yield a boundary-conditioned grid rescaling factor. Next, we computed the classic grid rescaling factor between the familiar rate map and the whole-trial rate map, separately for each deformed dimension. This analysis revealed a significant reduction in the appearance of rescaling when the data are first divided according to the most recently contacted boundary (Fig. 5d; Fig. S3c). The reduction in the appearance of rescaling upon boundary-conditioning was specific to grid cells which exhibited rescaling in their whole-trial rate maps. Thus, boundary rate map grid patterns exhibited significantly less rescaling than whole-trial rate maps, consistent with the boundary-tethered model.

We next tested whether environmental deformations affect grid field size. The boundary-tethered model predicts that deformations induce shifts in the spatial phase of the grid pattern. Averaged over the entire trial, these shifts should yield an increase in field size, regardless of whether the environment was compressed or stretched. On the other hand, a pure rescaling account predicts an increase in field size during stretching, but a decrease in field size during compressions. Because both accounts predict an increase in field size during stretching deformations, we focused on only on compression trials. For each grid cell we computed the ratio of the average field size during compression trials to the average field size averaged during familiar trials from the whole trial rate map. This ratio was significantly larger than one (1.21 ± 0.07 mean ± SEM; 2-tailed t-test versus 1: $t(114) = 2.9$, $p = 0.004$), indicating that field size increased during compressions deformations as predicted by the boundary-tethered model prediction but not a pure rescaling account.

Lastly, we examined firing rate predictions of the boundary-tethered model. If, during deformations, grid vertices are shifted to different locations when different boundaries are encountered, then averaging across trajectories originating from different boundaries will

necessarily reduce the peak values of the whole trial rate map. Thus the boundary-tethered model predicts a reduction in the *peak firing rate* during environmental deformations, as measured by the peak value of the whole trial rate map. On the other hand, because the density of grid fields within the environment remains unchanged on average, grid shift does not predict a change in *mean firing rate*, as measured by the total number of spikes across the entire trial divided by the trial duration. To address this prediction, we computed for each grid cell the ratio of its average peak rate across all deformation trials to the average peak rate across all familiar trials. This ratio was significantly below 1 (Fig. 5e), indicating that peak firing rates were lower during environmental deformations, as predicted by the boundary-tethered model. Conversely, the ratio of the mean firing rates did not significantly differ from 1 (Fig. 5e), suggesting that the mean firing rate did not change during environmental deformations, also as predicted. In sum, these results provide convergent evidence of boundary-tethered grid shifts during environmental deformations.

Discussion

We proposed and tested the hypothesis that a single mechanism, boundary-tethered grid shift, underlies the effects of environmental deformations on grid and place cells. We first demonstrated that a biologically plausible model of interactions between border cells and grid cells instantiating this mechanism can account for key results from a number of environmental deformation studies: 1) scale-dependent grid rescaling [9,11]; 2) place field stretching, bifurcation, modulation by movement direction, and inhibition during stretching deformations [12]; 3) end-of-track place code updating during compressions of a linear track [13,28]; and 4) place field duplication, inhibition, and perseverance during insertion of a new boundary [14–16]. Although previous computational models can explain subsets of these experimental phenomena [14,17,18,29–31], this is the first time that a single mechanism has been demonstrated to reproduced all of them. Next, we reanalyzed two classic environmental deformation datasets, and observed striking evidence of the history-dependent grid shifts that are the hallmark of our model. Together, these results provide new and compelling evidence that a phenomenon that was previously unrecognized in the data, boundary-tethered grid shift, underlies the distortions observed in grid and place cell activity during environmental deformations.

The apparent rescaling of grid patterns during environmental deformations has been taken as evidence that the mental metric for space maintained by grid cells can be reshaped by altering spatial geometry [9,37]. Our results instead suggest that the grid pattern continues to reflect unbiased metric path integration during deformations, at least relative to the most recently contacted boundary. However, our results further imply that in these conditions a particular grid network state no longer coincides with a unique spatial location within the environment. This dissociation between grid network state and spatial location may make environmental deformation paradigms particularly well-suited for distinguishing the contributions of grid cells to navigation from other factors [38,39].

Previous work has revealed striking parallels between deformation-induced distortions of spatial representations in the rat brain and the spatial memory of humans in deformed environments [12,38,40,41], leading to the suggestion that a common mechanism might underlie these effects. Consistent with this idea, recent evidence suggests that rescaling can be observed in the time-averaged activity of human grid cells [42]. In light of our results, we suggest that boundary-tethered grid shift might be the common mechanism underlying these cross-species similarities, and predict that boundary-anchored shifts in human spatial memory should be observable during environmental deformations.

The phenomenon of boundary-tethered grid shift could have various neural circuit realizations. Here we implemented a particular model of interactions between border, grid and place cells that gave rise to these shifts. This model was feedforward between layers [27], included an attractor network of grid cells [24], and generated place cells from grid cell output alone [26]. Although each of these components was motivated by prior work, this model is not intended as a complete recreation of entorhinal-hippocampal connectivity, but rather demonstrates how border cell input can give rise to the complex dynamics we describe, even in a relatively simple network. As such, this model excludes known connections that are not necessary for these dynamics. For example, this model lacks visual inputs [36], input to place cells from sources other than grid cells [43], as well as reciprocal connections from place to grid cells [44], all of which play important roles in developing and maintaining a functional spatial code. Thus while our mechanism is one of many proposed models of the effects of environmental deformations [14,17,18,29–31], our new analysis of the experimental data demonstrates that any alternative circuit model of the hippocampal formation should also incorporate a mechanism to produce boundary-tethered grid shifts, as a key phenomenon underlying the experimentally observed deformation-induced grid and place field distortions.

The prevailing interpretation of these distortions has been that a rescaled environment induces a rescaled version of the cognitive map [10,12,14,16,18,22,29,37,39,41,45]. Our results suggest an alternative interpretation. Grid and place representations are not rescaled or otherwise distorted during environmental deformations. Rather, these representations are dynamically updated when boundaries are nearby. Our results add to a growing body of literature highlighting the unique ways in which environmental boundaries shape spatial representations in the brain [9–14,16,19,20,46–48], and invite a reconceptualization of how these representations dynamically adapt to a changing world.

Acknowledgements. We are grateful to the laboratories of Edvard Moser, Kate Jeffery and Dori Derdikman for making the data from [11] and [9] available for our reanalysis. We also thank Eli Pollock, Niral Desai and Xuexin Wei for advice on implementing spiking border-grid connections. Finally, we gratefully acknowledge support from NSF grant PHY-1607611 (VB), NIH grants EY022350 and EY022350 (RAE), and NSF IGERT grant 0966142 (ATK). VB thanks the Aspen Center for Physics (Aspen, Colorado) and the International Center for Theoretical Physics (Trieste, Italy) for support and hospitality while this paper was being completed.

References

1. O'Keefe, J., and Nadel, L. (1978). The Hippocampus as a Cognitive Map (Oxford University Press).
2. O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. *34*, 171–5.
3. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature *436*, 801–806.
4. McNaughton, B.L., Battaglia, F.P., Jensen, O., Moser, E.I., and Moser, M.-B. (2006). Path integration and the neural basis of the "cognitive map." Nat. Rev. Neurosci. *7*, 663–678.
5. Moser, E.I., and Moser, M.-B. (2008). A metric for space. Hippocampus *18*, 1142–56.
6. Buzsáki, G., and Moser, E.I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. Nat. Neurosci. *16*, 130–8.
7. Moser, E.I., Roudi, Y., Witter, M.P., Kentros, C., Bonhoeffer, T., and Moser, M.-B. (2014). Grid cells and cortical representation. Nat. Rev. Neurosci. *15*, 466–81.
8. Fuhs, M.C. (2006). A Spin Glass Model of Path Integration in Rat Medial Entorhinal Cortex. J. Neurosci. *26*, 4266–4276.
9. Barry, C., Hayman, R., Burgess, N., and Jeffery, K.J. (2007). Experience-dependent rescaling of entorhinal grids. Nat. Neurosci. *10*, 682–4.
10. Krupic, J., Bauza, M., Burton, S., and O'Keefe, J. (2016). Framing the grid: effect of boundaries on grid cells and navigation. J. Physiol. *594*, 6489–6499.
11. Stensola, H., Stensola, T., Solstad, T., Frøland, K., Moser, M.-B., and Moser, E.I. (2012). The entorhinal grid map is discretized. Nature *492*, 72–8.
12. O'Keefe, J., and Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. Nature *381*, 425–8.
13. Gothard, K.M., Skaggs, W.E., and McNaughton, B.L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J. Neurosci. *16*, 8027–40.
14. Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K., and Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. Rev. Neurosci. *17*, 71–97.
15. Muller, R.U., and Kubie, J.L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. *7*, 1951–68.
16. Lever, C., Burgess, N., Cacucci, F., Hartley, T., and O'Keefe, J. (2002). What can the hippocampal representation of environmental geometry tell us about Hebbian learning? Biol. Cybern. *87*, 356–72.
17. Sheynikhovich, D., Chavarriaga, R., Strösslin, T., Arleo, A., and Gerstner, W. (2009). Is there a geometric module for spatial orientation? Insights from a rodent navigation model. Psychol. Rev. *116*, 540–66.
18. Raudies, F., Hinman, J.R., and Hasselmo, M.E. (2016). Modelling effects on grid cells of sensory input during self-motion. J. Physiol. *594*, 6513–6526.
19. Hardcastle, K., Ganguli, S., and Giocomo, L.M. (2015). Environmental boundaries as an error correction mechanism for grid cells. Neuron *86*, 827–39.

20. Giocomo, L.M. (2016). Environmental boundaries as a mechanism for correcting and anchoring spatial maps. J. Physiol. *594*, 6501–6511.

21. Solstad, T., Boccara, C.N., Kropff, E., Moser, M.-B., and Moser, E.I. (2008). Representation of geometric borders in the entorhinal cortex. Science *322*, 1865–8.

22. Bush, D., Barry, C., and Burgess, N. (2014). What do grid cells contribute to place cell firing? Trends Neurosci. *37*, 136–45.

23. Cheung, A. (2014). Estimating location without external cues. PLoS Comput. Biol. *10*, e1003927.

24. Burak, Y., and Fiete, I.R. (2009). Accurate path integration in continuous attractor network models of grid cells. PLoS Comput. Biol. *5*.

25. Pollock, E., Desai, N., Wei, X., and Balasubramanian, V. (2017). A mechanism for self-organized error-correction of grid cells by border cells. In COSYNE Abstracts 2017 (Salt Lake City).

26. Solstad, T., Moser, E.I., and Einevoll, G.T. (2006). From grid cells to place cells: a mathematical model. Hippocampus *16*, 1026–31.

27. Pilly, P.K., and Grossberg, S. (2013). Spiking neurons in a hierarchical self-organizing map model can learn to develop spatial and temporal properties of entorhinal grid cells and hippocampal place cells. PLoS One *8*, e60599.

28. Gothard, K.M., Hoffman, K.L., Battaglia, F.P., and McNaughton, B.L. (2001). Dentate gyrus and ca1 ensemble activity during spatial reference frame shifts in the presence and absence of visual input. J. Neurosci. *21*, 7284–92.

29. Cheung, A. (2016). Probabilistic Learning by Rodent Grid Cells. PLoS Comput. Biol. *12*, e1005165.

30. Cheung, A., Ball, D., Milford, M., Wyeth, G., and Wiles, J. (2012). Maintaining a cognitive map in darkness: the need to fuse boundary knowledge with path integration. PLoS Comput. Biol. *8*, e1002651.

31. Raudies, F., and Hasselmo, M.E. (2015). Differences in Visual-Spatial Input May Underlie Different Compression Properties of Firing Fields for Grid Cell Modules in Medial Entorhinal Cortex. PLOS Comput. Biol. *11*, e1004596.

32. Wei, X.-X., Prentice, J., and Balasubramanian, V. (2015). A principle of economy predicts the functional architecture of grid cells. Elife *4*, e08362.

33. Mathis, A., Herz, A.V.M., and Stemmler, M.B. (2013). Multiscale codes in the nervous system: the problem of noise correlations and the ambiguity of periodic scales. Phys. Rev. E. Stat. Nonlin. Soft Matter Phys. *88*, 22713.

34. Grossberg, S. (1980). How does a brain build a cognitive code? Psychol. Rev. *87*, 1–51.

35. Bush, D., and Burgess, N. (2014). A hybrid oscillatory interference/continuous attractor network model of grid cell firing. J Neurosci *34*, 5065–5079.

36. Pérez-Escobar, J.A., Kornienko, O., Latuske, P., Kohler, L., and Allen, K. (2016). Visual landmarks sharpen grid cell metric and confer context specificity to neurons of the medial entorhinal cortex. Elife *5*.

37. Evans, T., Bicanski, A., Bush, D., and Burgess, N. (2016). How environment and self-motion combine in neural representations of space. J. Physiol. *594*, 6535–6546.

38. Chen, X., He, Q., Kelly, J.W., Fiete, I.R., and McNamara, T.P. (2015). Bias in Human Path Integration Is Predicted by Properties of Grid Cells. Curr. Biol. *25*, 1771–6.

39. Bush, D., Barry, C., Manson, D., and Burgess, N. (2015). Using Grid Cells for Navigation. Neuron *87*, 507–20.

40. Hartley, T., Trinkler, I., and Burgess, N. (2004). Geometric determinants of human spatial memory. Cognition *94*, 39–75.

41. Hartley, T., Burgess, N., Lever, C., Cacucci, F., and O'Keefe, J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. Hippocampus *10*, 369–79.

42. Nadasdy, Z., Nguyen, T.P., Török, Á., Shen, J.Y., Briggs, D.E., Modur, P.N., and Buchanan, R.J. (2017). Context-dependent spatially periodic activity in the human entorhinal cortex. Proc. Natl. Acad. Sci. *114*, E3516–E3525.

43. Wills, T.J., Cacucci, F., Burgess, N., and O'Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. Science *328*, 1573–6.

44. Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J.L., Roudi, Y., Moser, E.I., and Moser, M.-B. (2013). Grid cells require excitatory drive from the hippocampus. Nat. Neurosci. *16*, 309–17.

45. Hasselmo, M.E., Hinman, J.R., Dannenberg, H., and Stern, C.E. (2017). Models of spatial and temporal dimensions of memory. Curr. Opin. Behav. Sci. *17*, 27–33.

46. Stensola, T., Stensola, H., Moser, M.-B., and Moser, E.I. (2015). Shearing-induced asymmetry in entorhinal grid cells. Nature *518*, 207–12.

47. Krupic, J., Bauza, M., Burton, S., Barry, C., and O'Keefe, J. (2015). Grid cell symmetry is shaped by environmental geometry. Nature *518*, 232–5.

48. Savelli, F., Luck, J., and Knierim, J.J. (2017). Framing of grid cells within and beyond navigation boundaries. Elife *6*.

49. Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.-B., and Moser, E.I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science *312*, 758–62.

Methods

Model

Border layer. The border layer consisted of 36 units. First, the area near each wall in 4 allocentric directions (North, South, East, West) was divided into 'bricks' (see [25] for a similar treatment). Each brick extended 12 cm from the wall and covered 11.1% of the total environment length along that dimension. The jth unit received a uniform input $b_j = 0.1$ whenever the simulated rat was within one of three adjacent bricks. This input was converted to stochastic spiking activity (see below). Each brick was assigned to a single border unit.

Grid layer. The grid layer, derived from the model of [24], consisted of 5 grid 'modules'. Each module consisted of a neural sheet with periodic boundary conditions, visualized as a torus. This neural sheet was composed of 45^2 identical 2 unit x 2 unit tiles (90^2 units per module). Each unit in a tile was associated with a particular direction (North, South, East, West), which determined both the movement-direction-specific excitatory input received, as well as its local connectivity. Movement-direction-specific excitatory input v_j to grid unit j was determined by

$$\mathrm{v_j} = \gamma + g_m\big(d\cos(\theta - \phi_\mathrm{j})\big)$$

where d is the distance moved since the previous timestep, θ is the direction of movement, ϕ_j is the preferred direction of unit j, g_m is a gain factor specific to the module m to which to unit j belongs, and $\gamma = 0.6$ is a constant. Local connections within each module consisted of shifted radial inhibition, in which each unit inhibited all units within a 12 unit radius by a uniform weight of -0.02. The center of this radial inhibition output for each unit was shifted by 2 units away from that unit in a direction consistent with each units preferred direction. In the absence of other inputs, each grid module yields a hexagonal grid-like pattern of activation on the neural sheet, which is translated during movement at a rate proportional to the gain factor. Thus, to model modules with varying grid scales, the gain factor g_m of the mth module was set by

$$g_m = \frac{g_1}{2^{\left(\frac{m-1}{2}\right)}}$$

where $g_1 = 0.92$ is the gain of the smallest-scale module, module 1. This results in a geometric series of biologically-plausible [11] grid scales for each module.

Place layer. The place layer consisted of 64 units, subject to uniform recurrent inhibition from all place layer units with a weight of -0.15.

Border-to-grid connectivity. All grid units received additional excitatory feed-forward projections from all border units. These connections were initialized with random weights uniformly sampled from the range 0 to 0.00222, and developed through experience via Hebbian learning (see below and [25]).

Grid-to-place connectivity. Each place unit received additional excitatory feed-forward projections from 500 random grid units. These connections were initialized with random weights

uniformly sampled from the range 0 to 0.022, and developed through experience via Hebbian learning (see below).

Model dynamics

Activation. The dynamics of the network was developed following the methods in [24]. The activation a_j of unit j was determined by first computing the total input b_j to unit j according to

$$b_j = \begin{cases} v_j + \sum_i^I a_i w_{ij}\,, & \text{grid units} \\ \sum_i^I a_i w_{ij}\,, & \text{place units} \end{cases}$$

where a_i is a variable quantifying activation of unit i, w_{ij} is the weight from unit i to unit j, and I enumerates all the units. (Note that some weights w_{ij} can be zero.) Also recall from above that a border unit receives a constant input when the rat is in a boundary region associated with that unit. The total input b_j was used to stochastically determine the spiking s_j of each unit j during the current timestep, according to

$$s_j = \begin{cases} 1, & \kappa(b_j - \beta_j)\mathrm{d}t > \mathrm{unif}(0,1) \\ 0, & \kappa(b_j - \beta_j)\mathrm{d}t \leq \mathrm{unif}(0,1) \end{cases}$$

where κ = 500 is a scale factor, β_j (border units: β_j = 0; grid units: β_j = 0.1; place units:

β_j = 0.05) is the spike threshold for unit j, $\mathrm{unif}(0,1)$ is a single draw from a random uniform distribution ranging from 0 to 1, and $\mathrm{d}t$ = 0.001 sec is the length of each timestep. Finally, this spiking activity was integrated to update the activation variable a_j of unit j after each timestep according to

$$a_j = a_j - a_j \frac{\mathrm{d}t}{c} + \alpha s_j$$

Where α = 0.5 is a scale factor and c = 0.02 is the time constant of integration.

Hebbian learning. All Hebbian weights were updated by the competitive learning rule

$$w_{ij} = w_{ij} + \lambda a_j \left(\left((\xi_j - w_{ij}) a_i \right) - \left(w_{ij} \sum_{n \neq i} a_n \right) \right)$$

where the sum is only over the set of units with nonzero Hebbian weights to unit j, λ = 0.00001 is the learning rate, ξ_j is a constant specific to the connection type (border-to-grid: ξ = 0.4; grid-to-place: ξ = 0.5) [27,34]. This rule results in competitive activity-dependent weight changes

among incoming Hebbian connections, and leads over time to a total weight of ξ_j across incoming synapses.

Simulation details

Generating simulated rat paths. Because some of the deformed environments that we tested have not been experimentally studied, it was necessary to generate simulated rat paths, rather than using experimentally recorded paths. Open field paths were generated via a bounded random walk model, parameterized by speed and movement direction. At each timestep, unbiased normally-distributed random noise was added to both speed (σ = 0.001 cm/msec) and movement direction (σ = 1 °/msec). To approximate actual rat exploration, speed was bounded to the range [0, 40] cm/sec. If a step would result in the rat path crossing a boundary, random noise was again added repeatedly to the movement direction until the next step would no longer cross the boundary. Open field paths always began in the center of the environment, with the simulated rat stationary and facing a random direction. Linear track paths were generated as straight end to end laps at a constant speed of 20 cm/sec.

Familiarization. In all simulations, familiarization with the environment was mimicked by allowing the naive simulated rat to explore the environment for 60 min. Prior to familiarization, grid layer activity was allowed to settle into its grid-like attractor state for 2 sec without learning. Initialization of the grid layer was biased so that an axis of the settled grid network state would lie at an angle of -7.5° relative to east, consistent with experiments [46,47]. Following familiarization, the model weights were saved so that all post-familiarization simulations could begin with the familiarized model.

Post-familiarization testing simulations. The model weights were reset to the state saved after familiarization, and the experienced virtual rat was allowed to explore each tested environment for 30 min. Grid layer activity was also initially reset to the familiar environment state corresponding to the rat's start location. Learning was turned off during the testing phase.

Analysis

Unit sampling. Due to computational constraints and the redundant nature of grid unit activity, only the spikes from 30 randomly chosen grid units in each module were recorded and analyzed during all simulations. All place units were recorded and analyzed.

Rate maps. Rate maps were created by first dividing the environment into 2.5 cm x 2.5 cm pixels. Then the mean firing rate within each pixel was calculated. Finally, this map was smoothed with an isotropic Gaussian kernel with a standard deviation of 1.5 pixels (3.75 cm) and square extent of 9 pixels x 9 pixels (22.5 cm x 22.5 cm). Pixels which were never visited were ignored during further analyses.

Autocorrelations and cross-correlations. Autocorrelations of rate maps were computed similar to previous reports [49]. Briefly, the correlation r between overlapping pixels of the original rate map and a shifted version of itself was computed as

$$r = \frac{\sum_{i=1}^{I}\sum_{j=1}^{J}\left(f_{ij} - \bar{f}\right)\left(f'_{ij} - \bar{f'}\right)}{\sqrt{\sum_{i=1}^{I}\sum_{j=1}^{J}\left(f_{ij} - \bar{f}\right)^2}\sqrt{\sum_{i=1}^{I}\sum_{j=1}^{J}\left(f'_{ij} - \bar{f'}\right)^2}}$$

where f is the rescaled rate map, f' is the familiar rate map, i and j run over pixels in the overlapping regions of these maps, and $\bar{f}$ and $\bar{f'}$ indicate the mean firing rate across overlapping pixels, at a series of single pixel (2.5 cm) step lags. *Cross-correlations* were computed similarly, except that two different rate maps, rather than two copies of the same rate map, were correlated. Autocorrelations and cross-correlations were only estimated for spatial lags with at least 20 overlapping pixels.

Grid scale. To compute grid scale we first averaged the autocorrelations of all grid units within a module. Next, we computed the mean distance from the center of the autocorrelation to the center of mass of the six closest surrounding peaks. In cases where the grid period was larger than the size of the environment thus obscuring the periodicity, grid scale was instead estimated by multiplying the scale of the next smaller module by $\sqrt{2}$, reflecting the parameters set in the attractor model creating the grid.

Grid rescaling factor. The grid rescaling factor during each deformation trial was computed separately for each unit by comparing rescaled versions of the deformed environment rate map to the familiar environment rate map. Following [11], the deformed rate map was uniformly rescaled to a series of chamber lengths, ranging from 10 cm below the smaller of the deformed and familiar chamber lengths, through 10 cm above the larger of these chamber lengths in 5 cm (2 pixel) increments. This yielded a set of rescaled rate maps for each unit, which were aligned with the familiar environment rate map by each of the two displaced boundaries. For each rescaled map and alignment, we computed the correlation r (defined above) between the rescaled and familiar rate maps. The grid rescaling factor was then defined as the ratio between the rescaled chamber length that yielded the highest correlation and the familiar chamber length, across either alignment.

Reanalysis of experimental data. A complete description of the experiments was provided in [9,11]. To test these data for the presence of grid shifts during environmental deformations, we first divided the spiking activity of each cell according to the most recent boundary contact (North, South, East, or West). Boundary contact was defined as the rat being within 12 cm of a boundary. Spiking activity prior to boundary contact at the beginning of the trial was ignored. Next, four separate rate maps were created, one for each most recently contacted boundary. To quantify grid shift along a particular dimension for each cell, the rate maps of opposing boundaries perpendicular to the chosen dimension were cross-correlated at a series of lags in single pixel steps (see above) within the range of ±20 pixels (±50 cm). The distance from the center to the nearest peak of this cross-correlation was computed as the measure of grid shift. The nearest peak was defined by first partitioning the cross-correlogram into 'blobs' of contiguous pixels which had correlations of at least 30% of the maximum value. Then, the location with the maximum correlation value within the blob nearest to the center was taken as the nearest peak. For alignment and rescaling analyses, the rate maps from the first familiar trials were used for comparison; in the few cases where no rate map was recording during the first familiar trial, the rate map from the second familiar trial was used instead. Grid field size for each rate map was measured as the median area of all patches of at least 9 contiguous pixels with rates >50% of the peak firing rate.

Data and code availability. All simulations were conducted with custom-written Matlab scripts. These scripts, the simulation results presented here, and similar results from an additional simulation of each experiment (to confirm reliability), are available from the authors upon reasonable request. All reanalyzed data are available upon request from the corresponding authors of the relevant papers.

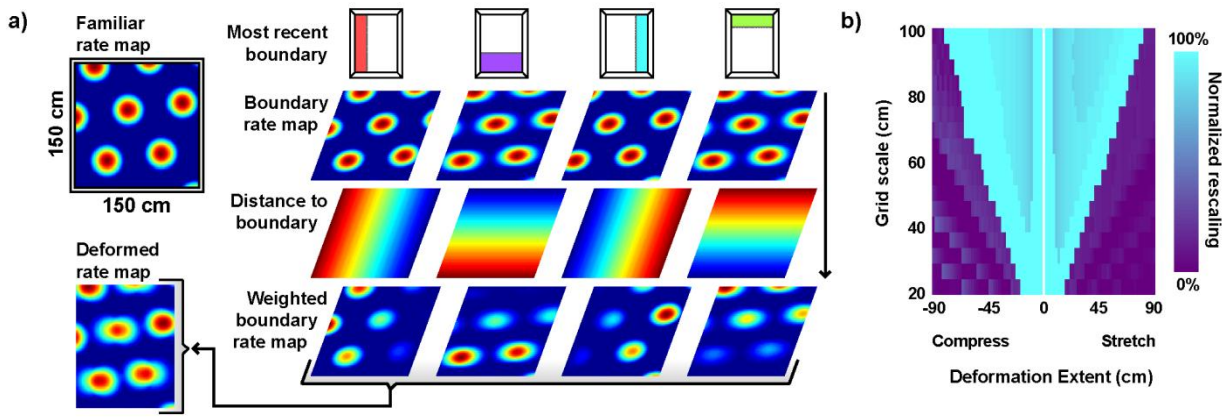

Supplementary Figure 1. Exploring the dependence of grid rescaling on grid scale and deformation extent with a sinusoidal interference model of grids. a) In the main text simulations were carried out over a limited range of grid scales and deformation extents because of computational constraints. To extend this analysis to a much wider range of conditions, and to verify that our results are independent of the specific implementation of grid cells, we approximated the effect of boundary-tethered grid shift without a network model. To do so, we first created familiar environment grid cell rate maps by overlapping and rectifying three 60°-offset sine waves [26]. Then, we shifted the spatial phase of these sine waves to reflect the most recently contacted boundary in the deformed environment. Finally, to approximate sampling biases we weighted each boundary rate map by both the inverse distance from the corresponding boundary and the relative length of the corresponding boundary, before adding together all four boundary rate maps. b) Grid rescaling when a familiar 150 cm x 150 cm square environment is deformed as a function of grid scale and deformation extent. Grid rescaling normalized to range from no rescaling (0%) to full rescaling (100%). Note that the best fitting rescaling is a full rescaling when the deformation extent is less than the grid scale and no rescaling when the deformation extent is larger than the grid scale. Each pixel represents the average best fitting rescaling of 30 cells with random spatial phases.

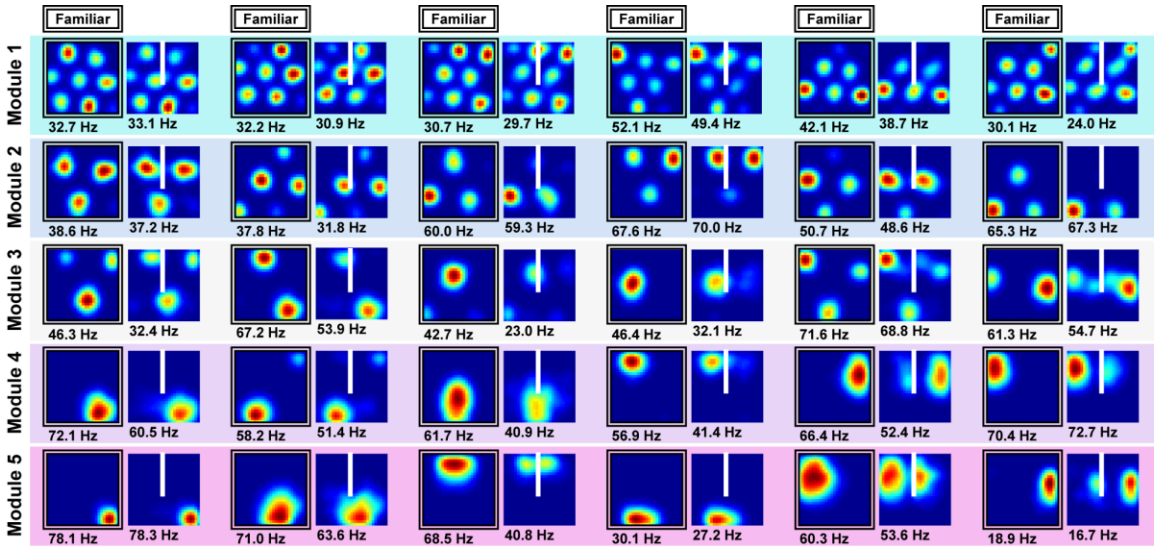

Supplementary Figure 2. Grid unit activity during insertion of a new boundary in an open environment. Examples of grid unit activity during exploration of a familiar chamber and boundary insertion (white line) -- 6 units shown from each module. Distortions are minimal in the time-averaged rate maps of small-scale grid units (matching experiment [21]), but become apparent in the activity of large-scale grid units. Peak firing rate noted below the lower left corner of each map.

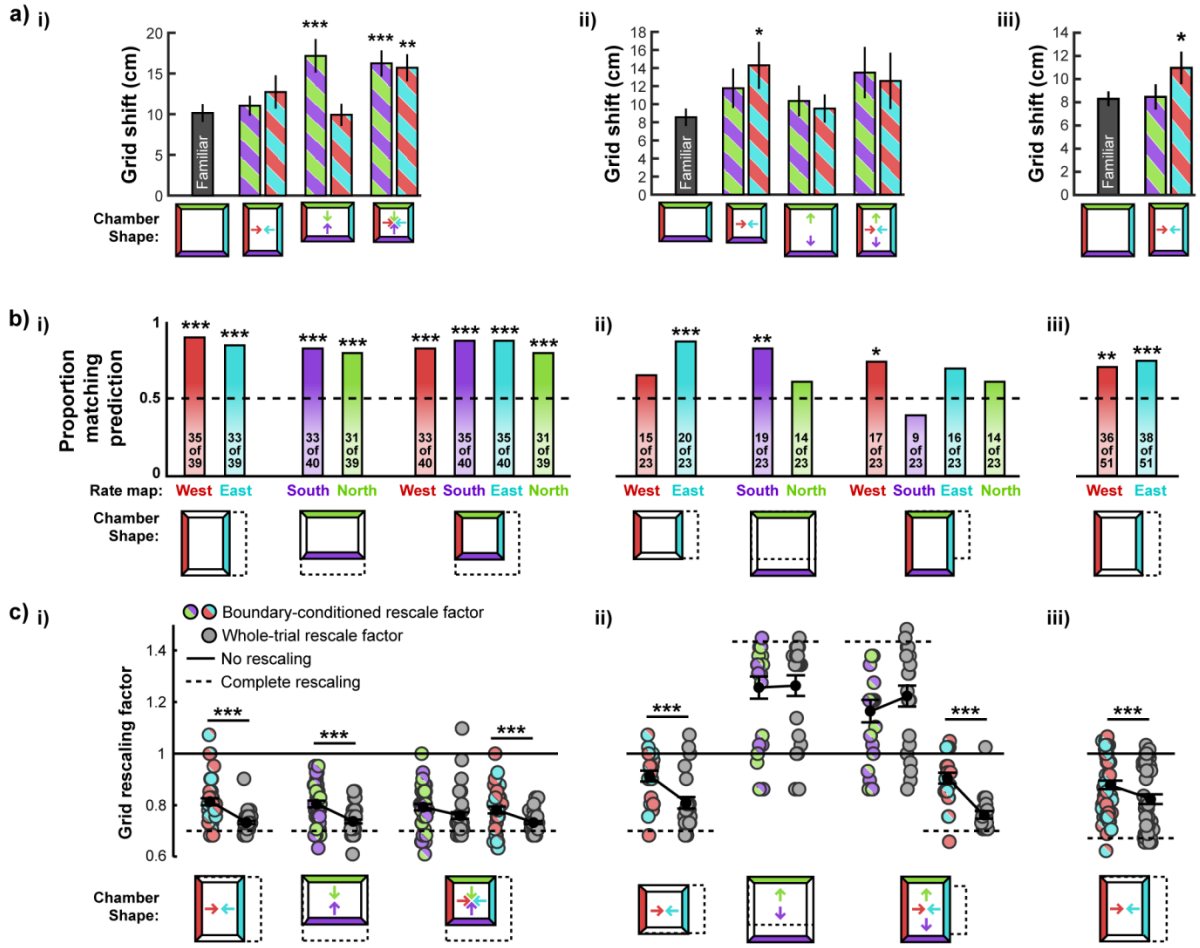

Supplementary Figure 3. Grid shift, alignment, and rescaling separated by condition. In each case rats trained in (i) a familiar square (data from [9]), (ii) a familiar rectangle (data from [9]), and (iii) a familiar square (data from [11]). a) Grid shift computed for each condition separately (see Text). Significance markers indicated 2-tailed paired t-tests versus familiar environment shift (full statistics in Table 1). Colored arrows indicate the dimensions along which our model predicts positive grid shift. b) Proportion of trials for which each boundary rate map was best matched with its familiar environment rate map when aligned by the most recently contacted boundary (as predicted by the boundary-tethered model) vs. the opposing boundary (counts shown within the bars). Familiar environment (dashed box), deformed environment (solid walls), and boundary (colored walls) shown in lower insets. Significance markers denote the outcome of a 2-tailed sign test versus 50% (i.e. equal probability of best alignment by either boundary; full statistics in Table 2). c) Boundary-conditioned and whole-trial grid rescaling factors for all deformation trials (errors bars ±1 SEM). Boundary-conditioned grid rescaling factors are closer to one than whole-trial grid rescaling factors. Rats trained in (i) a familiar square (data from [9]), (ii) a familiar rectangle (data from [9]), and (iii) a familiar square (data from [11]). Significance markers denote the outcome of a 2-tailed paired t-test between indicated conditions (full statistics in Table 3). *p<0.05, **p<0.01, ***p<0.001

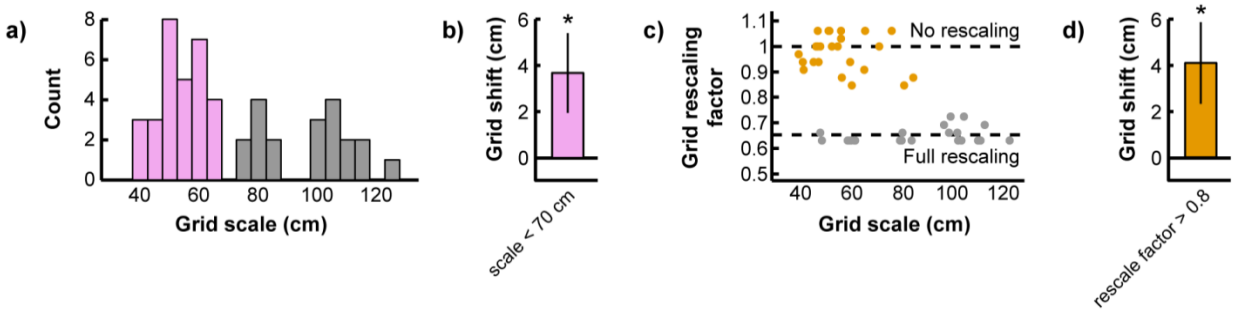

Supplementary Figure 4. Grid shift split by grid scale and rescaling. Data from [11]. a) Histogram of grid scale computed as the average grid scale across both familiar trials. b) Grid shift along the deformed dimension for cells split by grid scale, with average shift in the familiar environment subtracted. Note that shift is induced even in small-scale (scale < 70 cm) grid cells (1-sample t-test vs. 0: $t(30) = 2.13$, $p = 0.042$). c) Grid scale versus grid rescale factor for a two-thirds compression of the familiar environment. As reported previously [11], large-scale grid cells tend to exhibit rescaling while small-scale grid cells do not. d) Grid shift along the deformed dimension for cells split by grid rescaling factor, with average shift in the familiar environment subtracted. Note that shift is induced in grid cells which do not exhibit rescaling (rescale factor > 0.8, 1-sample t-test vs. 0: $t(27) = 2.3$, $p = 0.0288$). * $p < 0.05$

Dataset	Trained in	Deformed Dimension(s)	Measured Dimension	df	t	p
Barry et al., 2007	Square	east-west	north-south	41	0.65	0.520
Barry et al., 2007	Square	east-west	east-west	41	1.86	0.070
Barry et al., 2007	Square	north-south	north-south	40	3.35	0.001
Barry et al., 2007	Square	north-south	east-west	39	0.04	0.971
Barry et al., 2007	Square	both	north-south	41	3.85	0.000
Barry et al., 2007	Square	both	east-west	41	2.89	0.006
Barry et al., 2007	Rectangle	east-west	north-south	22	1.33	0.196
Barry et al., 2007	Rectangle	east-west	east-west	22	2.17	0.041
Barry et al., 2007	Rectangle	north-south	north-south	22	0.99	0.332
Barry et al., 2007	Rectangle	north-south	east-west	22	0.64	0.528
Barry et al., 2007	Rectangle	both	north-south	22	1.68	0.101
Barry et al., 2007	Rectangle	both	east-west	22	1.44	0.164
Stensola et al., 2012	Square	east-west	north-south	49	0.38	0.706
Stensola et al., 2012	Square	east-west	east-west	49	2.20	0.033

Supplementary Table 1. Statistical information for the grid shift tests depicted in Fig. S3a. All tests 2-tailed t-test vs. 0.

Dataset	Trained in	Deformed Dimension(s)	Boundary Rate Map	n	sign stat	p
Barry et al., 2007	Square	east-west	West	42	4	<0.0001
Barry et al., 2007	Square	east-west	East	42	6	<0.0001
Barry et al., 2007	Square	north-south	South	42	7	<0.0001
Barry et al., 2007	Square	north-south	North	42	8	0.0003
Barry et al., 2007	Square	both	West	42	7	<0.0001
Barry et al., 2007	Square	both	South	42	5	<0.0001
Barry et al., 2007	Square	both	East	42	5	<0.0001
Barry et al., 2007	Square	both	North	42	8	0.0003
Barry et al., 2007	Rectangle	east-west	West	23	8	0.2100
Barry et al., 2007	Rectangle	east-west	East	23	3	0.0005
Barry et al., 2007	Rectangle	north-south	South	23	4	0.0026
Barry et al., 2007	Rectangle	north-south	North	23	9	0.4049
Barry et al., 2007	Rectangle	both	West	23	6	0.0347
Barry et al., 2007	Rectangle	both	South	23	9	0.4049
Barry et al., 2007	Rectangle	both	East	23	7	0.0931
Barry et al., 2007	Rectangle	both	North	23	10	0.5413
Stensola et al., 2012	Square	east-west	West	51	15	0.0046
Stensola et al., 2012	Square	east-west	East	51	13	0.0006

Supplementary Table 2. Statistical information for the grid alignment tests depicted in Fig. S3b. All tests 2-tailed sign test vs. 0.5.

Dataset	Trained in	Deformed Dimension(s)	Measured Dimension	df	t	p
Barry, et al. 2007	Square	east-west	east-west	38	7.0	0.00000002
Barry, et al. 2007	Square	north-south	north-south	39	6.2	0.0000003
Barry, et al. 2007	Square	both	north-south	39	4.3	0.0001
Barry, et al. 2007	Square	both	east-west	39	1.9	0.06
Barry, et al. 2007	Rectangle	east-west	east-west	22	4.3	0.0003
Barry, et al. 2007	Rectangle	north-south	north-south	22	0.3	0.75
Barry, et al. 2007	Rectangle	both	north-south	22	1.8	0.086
Barry, et al. 2007	Rectangle	both	east-west	22	6.5	0.000002
Stensola, et al., 2012	Square	east-west	east-west	50	3.9	0.0003

Supplementary Table 3. Statistical information for the boundary-conditioned rescaling tests depicted in Fig. S3c. All tests 2-tailed paired t-test between boundary-conditioned rescaling factors and whole trial rescaling factors.