3D grid cells in human entorhinal cortex: Theoretical and methodological considerations and fMRI findings

Misun Kim, Eleanor A. Maguire

Wellcome Centre for Human Neuroimaging, Institute of Neurology,
University College London, London WC1N 3AR, UK

Corresponding author: Eleanor A. Maguire
Wellcome Centre for Human Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3AR, UK; e.maguire@ucl.ac.uk
Tel: +44 020 3448 4362
Fax: +44 020 7813 1445

Abstract

Recent human functional magnetic resonance imaging (fMRI) and animal electrophysiology studies suggest that grid cells in entorhinal cortex are an efficient neural mechanism for encoding knowledge about the world, not only for spatial location but also for more abstract cognitive information. The world, be it physical or abstract, is often high-dimensional. However, grid cells have been mainly studied on a simple two-dimensional (2D) plane and little is known about how grid cells encode the three-dimensional (3D) physical space in which we live. This dearth of knowledge is partly due to the technical difficulty of recording freely moving animals in 3D. Here, we first developed interactive software to help researchers visualize 3D grid cells and predict their activity. This relies on the known property of grid cells where their activity is modulated by the movement direction of animals relative to the grid axis. We then searched for 3D grid cells in human entorhinal cortex using a novel 3D virtual reality paradigm and a new fMRI analysis method. We found that signals in the left entorhinal cortex were best explained by one particular 3D grid cell model - a face-centred cubic lattice model. This is the first empirical evidence for 3D grid cells in the human brain, notwithstanding the inherent methodological limitations of fMRI. We believe that our findings and software serve as a useful initial stepping-stone for studying grid cells in realistic 3D worlds and also, potentially, for interrogating abstract high-dimensional cognitive processes.

Keywords

3D, grid cell, fMRI, virtual navigation, entorhinal cortex

Highlights

We present software and an analysis method to probe 3D grid cells using human fMRI

Based on an alignment score between 3D movement direction and grid orientation

We then tested this using a 3D virtual environment and fMRI

Signals in entorhinal cortex were best explained by a face-centred cubic lattice model

Introduction

Grid cells in entorhinal cortex (EC) have received much attention from researchers in the field of spatial navigation because of their unique firing pattern. A grid cell fires at multiple periodic locations resembling a hexagon when an animal explores an environment (Hafting et al., 2005; Fig. 1A). Importantly, different grid cells have different spatial scales and phases, the combination of which enables efficient encoding of an entire space using relatively few cells, compared to when each individual cell fires at unique locations, as is the case with hippocampal place cells. More recently, human fMRI and animal electrophysiology studies found that grid cells do not only encode the physical location of animals but also encode a continuously changing auditory tone (Aronov et al. 2017), locations during mental imagery (Bellmund et al., 2016; Horner et al., 2016), features of abstract visual stimuli (Constantinescu et al., 2016) and eye position during 2D visual search (Nau et al., 2018; Julian et al., 2018). This suggests that grid cells may be suitable for more abstract "cognitive mapping" (Tolman, 1948). If grid cells are indeed involved in abstract cognitive mapping, the 'space' cannot be limited to simple 2D physical space on which most grid cell research has to date been conducted. Grid cells should also be able to efficiently encode 3D and higher dimensional space.

Mathematical studies predict that 3D grid cells would fire at locations which correspond to the position of the centre of spheres tightly packed in 3D space, known as a face-centred cubic lattice (FCC) and hexagonal close packing (HCP) arrangements (Mathis et al., 2015; Fig. 2). There are only a few mammals that naturally explore 3D volumetric space upon which experimenters can test such theoretical predictions, for example, bats or primates in an arboretum. Although at least one research group is currently testing grid cells in flying bats (Ginosar et al., 2016), there is not yet clear empirical evidence of grid cells showing 3D lattice structures. Technical difficulties associated with recording animals freely moving in 3D space might be one reason for the dearth of empirical findings relating to 3D grid cells.

This is where studying the human brain becomes advantageous. With recent advances in affordable and easy-to-use virtual reality (VR) technologies, experimenters can create a virtual 3D world where human participants feel like they are floating or flying. A complex tethered or wireless recording device is not required for human participants in such a virtual environment. Although virtual environments have inherent limitations, such as lacking some primary sensory inputs, it is nevertheless a valuable starting point for studying the human brain in a volumetric 3D space and with which 3D grid cells can be potentially probed. fMRI is the first choice to study grid cells in the human brain because it is non-invasive and can detect signals in a deep structure like EC. An obvious limitation of fMRI is that it does not have cellular resolution. Nevertheless, previous fMRI studies have successfully investigated grid cells in 2D space using a known property of grid cells whereby their activity is modulated by movement direction (Doeller et al., 2010). This property is central to the approach we describe here.

As discussed by Stangl et al. (2017), grid code analysis is non-trivial, even in 2D, with various analysis strategies employed by different researchers. Extending grid analysis into 3D requires additional theoretical and methodological considerations. Although this article does not aim to provide a complete review of grid analysis in the literature, we outline key methodological issues that are either unique to 3D or relevant to both 2D and 3D. Furthermore, we showcase graphical user interface (GUI) software that we developed to visualize 3D grid structure and direction-modulated grid signals.

In the Methods section we first describe the proposed firing patterns of 3D grid cells, namely FCC and HCP. We then explain the principles of detecting 3D grid cells using fMRI with the software we developed. A detailed description of the analysis method then follows. Next, we present our experimental paradigm that involved VR in zero gravity during fMRI where we test the grid cell

models empirically in EC. Following this, in the Results section we report our findings from the fMRI study.

Methods

Grid analysis

Response profiles of 3D grid cells

Optimal encoding of 3D volumetric space with grid cells has been extensively discussed in previous theoretical and modelling studies (Mathis et al., 2015; Stella and Treves, 2015), and here we briefly describe two arrangements of 3D grid cells: FCC and HCP. Both FCC and HCP arrangements are analogous to the spatial arrangement of tightly stacked spheres inside a box with a minimum gap (e.g. oranges in a crate). These 3D arrangements can be viewed as 2D hexagonal lattices stacked on top of each other with a translational shift between the layers. FCC is composed of three repeating layers (blue, yellow and red spheres, see Fig. 2A left panel) and HCP is composed of two repeating layers (blue and red spheres, see Fig. 2B left panel). The grid axis, which is the key property used for our analysis, is the direction linking one grid field to its neighbouring grid fields (the green lines linking a centre black sphere to neighbouring coloured spheres in Fig. 2 middle panel). When we consider one centre sphere and its neighbouring 12 spheres, the positions of 3 spheres differ between the FCC and HCP arrangements (the yellow spheres in Fig. 2A middle panel). This difference leads to different direction-modulated fMRI signal between the FCC and HCP which is described in the next section. To enable ourselves and other researchers to visualize the 3D grid cells, we developed interactive web-based software where users can zoom, pan, rotate, and cross-sect a 3D grid cell's structure (the software, including a manual, can be accessed here: www.fil.ion.ucl.ac.uk/Maguire/grid3D_gui). Example screenshots are shown in Fig. 3. The software was implemented using Unity 5.4 (Unity Technologies, CA, United States).

A principle for detecting 3D grid cells using fMRI

fMRI measures the gross activity of thousands of neurons via complex neural-hemodynamic coupling. When the thousands of grid cells that fire at different locations are summed up, the gross activity is no longer expected to respond to fixed periodic locations in the environment. However, there is another important property of grid cells that enables their detection at a macroscopic level like fMRI. The activity of grid cells is known to be modulated by the alignment between the movement direction of an animal and the grid axis (Doeller et al., 2010). This means that a grid cell shows greater activity when an animal moves along the grid axis (Fig. 1B). As the majority of grid cells share a common grid axis, the summed response of thousands of grid cells can be systematically modulated by the movement direction of a participant. In previous studies that investigated grid cells in 2D, fMRI activity was modelled as a cosine function of movement direction relative to the grid axis with a period of 60 degrees to account for hexagonal symmetry.

We assume that the same principle of direction-modulation will hold in 3D grid cells, so that fMRI activity can be modelled as the degree of alignment between 3D movement direction and the grid axis. We defined the alignment score as a cosine of the angle between the 3D movement direction and the nearest grid axis. Thus, a larger signal is expected when the angle is smaller (i.e. the movement is aligned to the grid axis). The nearest grid axis forms the minimum angle with the direction vector. We can predict the grid cell's response when a participant is moving in a particular 3D direction, defined by azimuth (horizontal angle) and pitch (vertical angle) using our interactive software (example screenshots are shown in Fig. 4A,B).

According to our model, FCC and HCP grid cells will show complex response patterns as a function of the vertical and horizontal components of a participant's movement direction relative to the grid axis (Fig. 2 right panels). The response pattern of FCC and HCP are largely similar except for the difference in vertical symmetry. As the HCP arrangement is symmetric across the horizontal plane

(e.g. the blue spheres on a layer above the red spheres are located at the identical position as the blue spheres on a layer below, Fig. 2B middle panel), two movement directions that only differ in the sign of vertical pitch (e.g. (azimuth, pitch)=(45°, 30°) and (45°, -30°)) result in the same grid activity (symmetric across θ=0 plane in Fig. 2B right panel). In contrast, the FCC arrangement is symmetric across the origin (e.g. the blue and yellow spheres are facing each other, Fig. 2A middle panel) and therefore the two directions opposing across the origin will result in the same grid response.

The orientation of the grid axis relative to the environment

Crucially, the activity of a grid cell, or a grid voxel in fMRI, depends upon both the movement direction of a participant (which is known to experimenters) and the orientation of grid axis relative to the 3D environment (which is unknown to experimenters). Fig. 5 describes two hypothetical cases where a participant moves in the same direction but the grid axis is oriented differently. In Fig. 5A, a participant's movement direction (black arrow) is relatively closely aligned to the grid axis with an angular deviation of 14°, resulting in high activity (the yellow bar graph). In Fig. 5B, due to a different orientation of the grid axis, the same movement direction is farther away from the grid axis with an angle of 35°, resulting in low activity. The orientation of the grid axis should be numerically estimated by iteratively fitting the experimental data - a process we describe in a later section. Of note, the orientation of the grid axis can be analytically estimated in a simpler 2D case.

Unlike in 2D where the orientation of the grid axis can be specified by one polar angle from a reference direction (e.g. 20° from the north-south axis), the grid axis in 3D can be rotated along any three arbitrary axes and the order of applying each rotation also matters (known as the non-commutative property of 3D rotation). Although users can explore these 3D rotation options in our software, we restricted the rotation of the 3D grid axis to only one axis so that six hexagonal grid fields (the red spheres in Fig. 5) remain parallel to the ground of the environment when we analyse our fMRI data. This restriction in the rotation axis can be justified by the fact that grid cells on a 2D

horizontal surface show the corresponding grid fields. This restriction also simplifies the modelling process and reduces the computational cost.

The relationship between the grid alignment score and fMRI activity

A grid voxel's activity is expected to be modulated by the degree of alignment between the movement direction and the grid axis, and we defined the grid alignment score as the cosine of the angle between the movement direction and the nearest grid axis. This is equivalent to the previous grid analysis in 2D which used parametric regressors of cosine and sine functions (Doeller et al., 2010; Horner et al., 2016), except for a minor difference in that our method predicts more abrupt changes in grid alignment scores when the movement direction passes the midpoint of two neighbouring grid axes (Supplementary Fig. S1B,C). Of note, the precise form of the direction-modulated firing rate of grid cells is not known either in 2D or 3D. There is also additional complexity in measuring the grid cell's signal via neural-hemodynamic coupling. Therefore, it is also possible to model the grid voxel's activity with a non-sinusoidal function, like a linear or binary function (Supplementary Fig. S1D,E). The exact relationship between direction-modulated grid activity and the fMRI response should be examined in future studies.

Estimating 3D grid orientation in fMRI data

In this section, we describe how to estimate 3D grid orientation from the fMRI time series. As explained earlier, the grid alignment score can be calculated as the cosine of the angle between the participant's movement direction (known to experimenters) and the nearest grid axis. The nearest grid axis is determined by the orientation of the grid axis relative to the environment, which is unknown to the experimenters (Fig. 5). We, therefore, first assume that the grid orientation is aligned at 0 degrees from a reference direction (e.g. parallel to the side wall of the environment) and then calculated the grid alignment score. This grid alignment vector is then convolved with the hemodynamic response function (SPM canonical hemodynamic response function). The resulting

vector serves as a hypothetical grid voxel signal. We create a general linear model (GLM) which contains this predictive 3D grid signal and nuisance regressors that include six head motion realignment parameters and experiment-specific conditions like an occasional question and response period. The fMRI time series (after standard preprocessing) in each voxel and in each scanning session is then fitted with the GLM, and the outcomes - beta (regression coefficients) and mean square residual - are saved for each voxel.

We then repeat the whole procedure with newly calculated grid alignment scores with different assumptions, namely that the grid orientation is aligned at 15, 30, 45, …, 120 degrees relative to the environment (cf. we only need to sample the grid orientation up to 120 degrees, as the geometry of the 3D lattice structure of both FCC and HCP is symmetric for the 120 degrees rotations on a plane (we can check this property using our 3D visualization software). For each voxel and each scanning session, we select the orientation of the grid axis that gives the best fit by comparing the mean square residual of these multiple GLMs. A GLM with the smallest mean square residual and a positive regression coefficient for the grid signal regressor is selected. The reason we select the GLM with a positive regression coefficient is to avoid the inverted relationship between the hypothetical 3D grid cell's signal and the fMRI response (e.g. when movement is more aligned to the grid axis, the fMRI signal is lower). In rare case (<10% of voxels in our empirical data – see later sections) where all grid orientation models yield a negative regression coefficient, we simply select the GLM with the smallest mean square residual. To summarise, this iterative fitting process identifies which grid orientation best describes the fMRI signal in each voxel and in each scanning session.

Testing for a grid signal in the fMRI data

We then test whether each voxel shows a consistent 3D grid signal across different scanning sessions by quantifying the regression coefficient of the grid signal model. For instance, if the fMRI data in one scanning session (e.g. run 1) is best fitted with a grid model that aligns at 15 degrees, we

measure the grid score as the beta of the same grid orientation (15 degrees) model in the another scanning session (e.g. run 3). The beta values of voxels in the brain region of interest (ROI) are averaged for each participant, and a t-test is used to test whether the beta is positive at the group level (excluding outliers, participants with more than a standard deviation of 3 in our empirical data). This approach is similar to previous 2D grid analyses where the grid orientation is estimated from one half of the dataset and is tested on the other half of the dataset, and the regression coefficient is analysed at the group level (e.g. Doeller et al., 2010; a standard group level inference for fMRI experiments).

However, there is a difference between our study and some of the previous studies in terms of grid orientation averaging. In Doeller at al. (2010) and Horner et al. (2016), the estimated grid orientation of each voxel within the EC ROI was averaged, and this averaged grid orientation model was tested in the other half of the data. Here, we estimate and test the grid orientation model within each voxel, then we later summarise the grid score of voxels within the ROI. Neighbouring grid cells share a common grid orientation which is the essential property of grid cells that allows the detection of the direction-modulated signal at the fMRI voxel level, and earlier fMRI studies assumed one unique grid orientation for the entire EC. However, there is also evidence of multiple grid modules in the EC that have different grid orientations and scales (Stensola et al., 2012), and we believe that estimating and testing grid orientation at the voxel level, instead of the whole ROI, can maximise the sensitivity of analysis. This voxel-by-voxel estimation and test approach was used in more recent 2D grid cell study (Nau et al., 2018).

A control model

Our 3D grid cell analysis (as well as the 2D grid cell analyses in the literature) relies on the dependency of the neural signal on movement direction, and one concern is whether a neural signal that is responsive to one particular direction (or the view associated with a direction) could be

weakly correlated with a grid cell model and so identified as a grid cell. This was why in our empirical experiment (described later) we used a direction (or view) encoding model as a control analysis. We created a direction-sensitive model signal which was sensitive to one of nine 3D directions that were visited by participants in a virtual environment. The nine directions were regularly sampled both horizontally and vertically: (azimuth, pitch in degrees) = (-60, -60), (-60, 0), (-60, 60), (0, -60), (0, 0), (0, 60), (60, -60), (60, 0), (60, 60). Following Bellmund et al. (2016), we assumed that each direction-sensitive neural response had a margin of 30 degrees. This meant that neurons or voxels that responded strongly to (0, 0) direction would also respond strongly to (±30, ±30), and would respond weakly to the rest of the movement directions. We convolved the binary direction response vector with the hemodynamic response function. We created a GLM similar to the grid cell model described in the previous section but now the grid signal was replaced by the direction encoding signal. Again, the best direction-encoding model was selected for each voxel from one scanning session and then tested on the other scanning session. If voxels in our ROIs (left and right EC) responded to unique directions, we would see a significantly positive regression coefficient for a direction model at the group level.

Experimental protocol

In this section we describe an empirical fMRI experiment where data were acquired while participants were exploring a virtual 3D environment. The experimental paradigm included a pre-scanning session with VR goggles. We applied our proposed 3D grid analysis to this dataset.

Participants

Thirty healthy adults took part in the experiment (16 females; mean age 25.9±4.8 years; range 19-36 years; all right-handed). All had normal or corrected to normal vision and gave informed written consent to participation in accordance with the local research ethics committee.

The virtual environment

Participants were instructed that they were exploring a virtual zero-gravity environment, called "spaceship", where they could move freely up, down, forwards and backwards. This spaceship had two rectangular compartments linked by a corridor (Fig. 6A, the environment can also be viewed in our online software). Participants could orient themselves in each compartment because the visual appearance of the walls differed (e.g. a window on the west side and a grey wall on the east side; the ceilings and floors also had different textures or hues). A snapshot of this virtual spaceship as seen from a participant's perspective is shown in Fig. 6B,C. The virtual environment was implemented using Unity 5.4 (Unity Technologies, CA, United States).

The virtual spaceship was rendered on two different mediums for pre-scanning tasks and scanning tasks respectively: a head-mounted VR display (Samsung Gear VR, model: SM-R322 with Samsung Galaxy S6 phone) and a standard computer screen (Dell Optiplex 980 with an integrated graphic chipset). The VR display was used because it has been suggested that prior vestibular experience with a VR environment can later reinstate the relevant body-based cues during fMRI scanning, when only visual input is available due to head immobilisation (Shine et al., 2016). The VR head-mounted display provided participants with a fully immersive sensation of 3D space via its head motion tracking system, stereoscopic vision and a wide field-of-view (96°). A rotation movement in the VR display was made by a participant's physical head rotation and a translational movement was made by a button press on the Bluetooth controller (SteelSeries Stratus XL, Denmark). For example, a participant could move up to the ceiling in the virtual spaceship by physically looking above and press the forward button. To rotate to the right, they physically rotated their head to the right or rotated their whole body when the required rotation was beyond the range of motion for neck rotation. For ease of rotation, participants were seated on a swivel chair.

During fMRI scanning, participants watched a video rendered on a standard computer screen (aspect ratio=4:3, Fig. 6C). The video was a first-person perspective that gave the participants the feeling of moving in a virtual spaceship. The stimuli were projected on the screen using a projector at the back of the MRI scanner bore (Epson EH-TW5900 projector), and participants saw the screen through a mirror attached to the head coil. The screen covered a field of view of ~19° horizontally and ~14° vertically.

Pre-scan: VR memory task

Wearing the VR display, participants freely explored the environment (5 minutes) and then performed a spatial memory test (15 minutes) where they had to recall the location of balls by physically directing their head to the remembered locations (Fig. 6B). The purpose of this test was to ensure that participants were sufficiently exposed to the environment prior to the scanning task.

fMRI scan: direction judgment task during passive viewing

During scanning, participants watched a video rendered on a standard display and performed a direction judgment task. The video provided participants with the feeling that they were flying in a controlled 3D trajectory within the spaceship (Fig. 6C). The pre-programmed video allowed tight control of location, direction and timing for all participants. The trajectory consisted of multiple short linear movements (each of 3 seconds) followed by rotation (2/2.6 seconds). We restricted the range of movement directions (-60 to 60 degree with 30 degree steps, both vertically and horizontally) to increase the frequency of each movement direction visited within the limited scanning duration. A smooth trajectory was used without abrupt rotations. A constant linear and angular velocity was applied in order to control velocity, which can also influence grid cells' activity. If a participant reached the boundary of the spaceship, a blank screen appeared for 2 seconds and then a trajectory started again from the other end of the spaceship. For 25% of the time, a question screen appeared immediately after a linear movement and participants indicated the direction of

their last movement by pressing an MR-compatible button pad (a 5-alternative forced choice question with a time limit of 5 seconds; Fig. 6C). Vertical or horizontal questions were randomly presented. This direction judgment task was used to ensure participants kept track of their movements during scanning. The two compartments of the spaceship were visited alternatively for each of 4 scanning sessions. Half of the participants started in one compartment and half started in the other compartment. Each scanning session lasted ~11 minutes with a short break between the sessions, making a total functional scanning time of ~50 min.

Post-scan debriefing

After scanning, participants were debriefed about how much they felt immersed in the virtual environment during the pre-scan session with VR goggles and during scanning. Participants chose from multiple options: “I felt like I was really in the spaceship”; “I occasionally thought about the environment as being computer-generated, but overall the environment was convincing and I felt I was moving around in the spaceship”; “I was often distracted by the feeling that I was not in a real environment”.

Behavioural analyses

The overall accuracy of the direction judgment task during scanning was measured (chance=20%) to check whether participants knew in which direction they were moving in the 3D environment. For the debriefing question, we counted the number of responses for each option.

Scanning and pre-processing

T2*-weighted echo planar images (EPI) were acquired using a 3T Siemens Trio scanner (Siemens, Erlangen, Germany) with a 32-channel head coil. Scanning parameters optimised for reducing susceptibility-induced signal loss in areas near the orbitofrontal cortex and medial temporal lobe were used: 44 transverse slices angled at -30°, TR=3.08 s, TE=30 ms, resolution=3×3x3mm, matrix

size=64x74, z-shim gradient moment of -0.4mT/m ms (Weiskopf et al. 2006). Fieldmaps were acquired with a standard manufacturer's double echo gradient echo field map sequence (short TE=10 ms, long TE=12.46 ms, 64 axial slices with 2 mm thickness and 1 mm gap yielding whole brain coverage; in-plane resolution 3 x 3 mm). After the functional scans, a 3D MDEFT structural scan was obtained with 1mm isotropic resolution.

Preprocessing of data was performed using SPM12 (www.fil.ion.ucl.ac.uk/spm). The first 5 volumes from each functional session were discarded to allow for T1 equilibration effects. The remaining functional images were realigned to the first volume of each run and geometric distortion was corrected by the SPM unwarp function using the fieldmaps. Each participant's anatomical image was then coregistered to the distortion corrected mean functional images. Functional images were normalised to MNI space and smoothed with 6mm kernel.

ROI selection

We used anatomical ROIs, left and right EC masks (Fig. 7A). The ROIs were manually delineated on the group-averaged MRI scans from a previous independent study on 3D space representation (Kim et al., 2017) following the protocol in Pruessner et al. (2002). The number of functional voxels (3 x 3 x 3 mm) within the ROI masks was 47 (left) and 49 (right). Of note, EC was further divided into posterior medial and anterior lateral parts in one previous fMRI study (Bellmund et al., 2016), based on the finding in rodents that grid cells are typically reported in the medial EC. However, our study used standard resolution fMRI and further segmentation of this kind was not feasible. Functional specialisation within the EC is an interesting topic that needs to be further addressed in future studies with high-resolution scanning sequences.

Another important consideration for selecting EC ROIs is that EC is notoriously difficult to image because of fMRI susceptibility artefact in this vicinity. Although sequence development continues in

this regard, EC still has inherently low raw BOLD signal compared to other cortical regions. Crucially, standard fMRI analysis software like SPM excludes voxels of low signal by default. The "global masking threshold" parameter in the first-level model specification in SPM determines which voxels are to be included in the analysis based on the raw intensity, and EC voxels can often be excluded. It can also result in a different number of voxels in an EC ROI for each participant. We caution researchers about excluding voxels for two reasons. First, an exclusion criterion based on the mean BOLD intensity is arbitrary. Depending on the version of the software, some voxels can be included or excluded from the analysis. Second, raw BOLD intensity alone does not predict whether a voxel shows functional modulation. For instance, whereas the raw signal intensity of cerebrospinal fluid (CSF) is higher than most other cortical areas, we rarely observe meaningful signal in the CSF in typical cognitive experimental paradigms. In our study, we defined the EC ROIs purely anatomically, without excluding any voxels based on raw signal intensity.

fMRI grid analysis

We applied the 3D grid analysis that we described in detail above to the preprocessed fMRI dataset. In essence, we estimated the orientation of the 3D grid (separately for FCC and HCP grid types) in each voxel and scanning session by iteratively fitting the fMRI time series to the predicted grid alignment score defined as the cosine of movement direction and the nearest grid axis. The grid model was tested on another scanning session. Because our virtual spaceship had two compartments, we trained and tested the grid cell models within each compartment, and averaged the regression coefficient of the two compartments.

fMRI control analysis

As outlined above, we tested the unique direction (or associated view) encoding model as a control. In addition, we compared the orientation of the estimated 3D grid axis across participants. If the visual features of the environment had spatial regularity (e.g. the view when looking 60 degrees is

similar to the view when looking 0 degrees) that resembled grid signals, we would expect every participant to show a similar grid orientation given they were in the same environment. We calculated the average grid orientation across all voxels within the ROI for each subject and each compartment using a circular mean. We then applied a Rayleigh test for non-uniformity of circular data. The circular mean and non-uniformity test was computed using CircStat2012a toolbox (Berens, 2009).

Empirical results

Behavioural

The mean accuracy of direction judgments during scanning was well above chance level (mean 74%, SD 16%; chance 20%), suggesting that participants knew their 3D movement direction in the virtual environment during scanning.

The rating data showed that participants felt immersed in the virtual environment, with the vast majority choosing either "I felt like I was really in the spaceship" (57% for the pre-scan VR goggle task, 10% for the scanning task) or "I occasionally thought about the environment as being on a computer screen, but overall the environment was convincing and I felt I was moving around in the spaceship" (43% for the pre-scan VR goggle task, 80% for the scanning task). This result implies that our virtual environment effectively conveyed a sense of being in 3D space.

fMRI

We tested whether fMRI signal in the left and right EC was modulated by participants' 3D direction as predicted by different 3D grid cell models: FCC and HCP. The FCC grid model was significant in the left EC ($t(28)=2.9$, $p=0.007$), and showed a trend in the right EC ($t(28)=1.8$, $p=0.09$) (Fig. 7B). The HCP grid model did not explain the response of either left or right EC (Fig. 7C). To exclude the possibility of a neural signal sensitive to one particular direction (or associated view) being identified as a grid

cell, we tested a unique direction encoding model as a control. The direction encoding model was not significant in either EC (Fig. 7D), suggesting that the FCC grid-like signal found in the left EC was not driven by one particular direction. Furthermore, the grid orientation was not significantly clustered across participants (Rayleigh's test for non-uniformity was $p>0.1$), making it even less likely that particular visual features of the environment caused grid-like signals. This non-clustered grid orientation across participants is consistent with previous studies (Doeller et al., 2010; Nau et al., 2018)

Discussion

Here we presented a VR experimental paradigm and a novel analysis method to investigate 3D grid cells non-invasively in humans. We also developed associated software to help researchers visualize 3D grid cells and predict their responses. Using these methods, in our empirical data we observed an fMRI signal in the EC which was in line with one particular 3D grid cell model – an FCC lattice model.

A prerequisite for probing 3D grid cell is to use an appropriate 3D environment. 3D lattice structures, like FCC or HCP, are optimal for encoding volumetric space (Mathis et al., 2015), and the existence of substructures like 2D walls or a 1D track could affect the response of grid cells. For instance, grid cells recorded in rats moving on a sloped terrain showed firing pattern similar to a 2D horizontal plane rather than a 3D lattice pattern (Hayman et al., 2015). Therefore, we built a fully volumetric virtual "zero-gravity" environment where participants could move freely along all directions. Furthermore, we complemented the fMRI scanning, where only visual input was available because of in-scanner head immobilisation, by using VR goggles during pre-scan tasks as in Shine et al. (2016). We believe this pre-scan experience of physical head rotation when using the VR goggles was particularly helpful to participants in building the mental and neural representation of a 3D space, and this was supported by participants' reports in the debriefing session at the end of the experiment.

The main finding of this study relates to our probing of putative 3D grid cells using fMRI by predicting the neural signal as a function of 3D movement direction and the grid axis. The principle of measuring direction-modulated grid signals has been widely used in 2D (Doeller et al., 2010; Constantinescu et al., 2016; Bellmund et al., 2016). However, in our study, we extended, for the first time, this principle into 3D volumetric space, thereby opening up the possibility of empirically studying grid cells in high-dimensional space. Predicting the activity of a 3D grid cell was inevitably more complicated than a 2D grid cell due to the added dimension, as we described in earlier. Consequently, we estimated the 3D grid orientation by iteratively fitting the neural data with signals predicted by different possible grid orientations. We successfully demonstrated the feasibility of this analysis approach by finding our data was concordant with the FCC model in the left EC, the candidate brain structure for 3D grid encoding.

Our result is suggestive of 3D grid cells in the human EC when they explore a volumetric space, but we remain cautious until this finding is further corroborated by future studies that directly measure cellular responses. It is not yet proven that the direction modulation principle also holds in 3D because there is no clear evidence of grid cells showing a regular 3D lattice pattern (cf. a recent study in bats found multiple firing fields that were not perfectly regular; Ginosar et al., 2016). Invasive recordings from the human brain in a similar 3D virtual reality experiment to the one we used here may be able to provide more direct evidence of 3D grid cells. Once there is a fuller understanding of the cellular physiology of 3D grid cells, it will be possible to determine the optimal fMRI analysis protocol by considering the multiple factors we have described here, such as whether the orientation of the 3D grid axis is parallel to the ground or not, the precise model between the grid alignment and the fMRI signal (e.g. cosine, linear, binary), and the distribution of the grid orientation across different voxels within EC.

The distinction between the FCC and HCP models also deserves further investigation. Although the two models have slightly different grid axes, and only the FCC model was significant in our empirical data, a theoretical study showed that both have identical encoding efficiency (Mathis et al., 2015). One simulation study found that a mixture of the FCC and HCP models emerged (Stella and Treves al., 2016). It may be that factors such as the anisotropy of the environment (all animals are influenced by gravity, thus the vertical axis is distinguished from the other two axes) might influence the optimal encoding of space. It is currently unknown how this might be related to either the FCC or HCP models. Furthermore, grid cells are known to be influenced by the geometry of environment (Krupic et al., 2014). Here, we used a spaceship with a rectangular shape, therefore FCC and HCP models should be tested in the environment of different shapes (e.g. spherical or asymmetric shapes) in the future.

In the process of extending grid analysis from 2D to 3D, we felt there was a pressing need for 3D visualization software, because it was not trivial to imagine the lattice structures of 3D grid cells and the relationship between the grid axis and 3D movement vectors. Therefore, we developed simple interactive software that helps us, and we hope other researchers, to visualize and understand the direction-modulated signal of 3D grid cells. This web-based software is easy to use without requiring an additional third-party CAD (computer-aided design) program. Of note, this software was developed for visualization purposes, and it does not have a data analysis function. There is already grid analysis software for 2D cases (Stangl et al., 2017), and researchers who are familiar with 2D grid analysis can readily implement our 3D analysis with their own protocols, once they grasp the geometry of 3D grid cells using our visualization software.

In summary, we believe that our experimental paradigm, analysis method and software serve as a useful initial stepping-stone for studying grid cells in realistic 3D worlds. Grid cells have been reported to encode not only physical space, but also more abstract knowledge (Aronov et al., 2017;

Constantinescu et al., 2016; Nau et al., 2018; Julian et al., 2018), and we hope our approach will in due course also promote interrogation of abstract high-dimensional cognitive processes.

Acknowledgements

This work was supported by the Wellcome Trust (101759/Z/13/Z to E.A.M.; 203147/Z/16/Z to the Centre; 102263/Z/13/Z to M.K.) and a Samsung Scholarship (to M.K). We thank Kate Jeffery, Tim Behrens, Alexa Constantinescu and Neil Burgess for helpful discussions.

The authors declare no competing financial interests.

References

Aronov, D., Nevers, R., Tank, D.W., 2017. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. Nature. 543, 719–722.

Bellmund, J.L., Deuker, L., Navarro Schröder, T., Doeller, CF., 2016. Grid-cell representations in mental simulation. eLife, 5, e17089.

Berens, P., 2009. CircStat : A *MATLAB* Toolbox for Circular Statistics. J. Stat. Softw. 31.

Constantinescu, A.O., O'Reilly, J.X., Behrens, T.E.J., 2016. Organizing conceptual knowledge in humans with a gridlike code. Science. 352, 1464–1468.

Doeller, C.F., Barry, C., Burgess, N., 2010. Evidence for grid cells in a human memory network. Nature. 463, 657–661.

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

Horner, A.J., Bisby, J.A., Zotow, E., Bush, D., Burgess, N., 2016. Grid-like processing of imagined navigation. Curr. Biol. 26, 842–847.

Ginosar, G., Finkelstein, A., Las, L., Ulanovsky, N., 2016. Representation of 3D space in the entorhinal cortex of flying bats, Society for Neuroscience abstract, San Diego, USA.

Julian, J.B., Keinath, A.T., Frazzetta, G., Epstein, R.A., 2018. Human entorhinal cortex represents visual space using a boundary-anchored grid. Nat. Neurosci. 21, 191-194

Krupic, J., Bauza, M., Burton, S., Lever, C., O'Keefe, J., 2014. How environment geometry affects grid cell symmetry and what we can learn from it. Philos. Trans. R. Soc. B Biol. Sci. 369, 20130188.

Mathis, A., Stemmier, M.B., Herz, A.V.M., 2015. Probable nature of higher-dimensional symmetries underlying mammalian grid-cell activity patterns. eLife, 4, e05979

Nau, M., Navarro Schröder, T., Bellmund, J.L.S., Doeller, C.F., 2018. Hexadirectional coding of visual space in human entorhinal cortex. Nat. Neurosci. 21, 188-190.

Pruessner, J.C., Köhler, S., Crane, J., Pruessner, M., Lord, C., Byrne, A., Kabani, N., Collins, D.L., Evans, A.C., 2002. Volumetry of temporopolar, perirhinal, entorhinal and parahippocampal cortex from high-resolution MR images: considering the variability of the collateral sulcus. Cereb. Cortex. 12, 1342–1353.

Stangl, M., Shine, J., Wolbers, T., 2017. The GridCAT: A toolbox for automated analysis of human grid cell codes in fMRI. Front. Neuroinform. 11, 1–12.

Shine, J.P., Valdes-Herrera, J.P., Hegarty, M., Wolbers T., 2016. The human retrosplenial cortex and thalamus code head direction in a global reference frame. J. Neurosci. 36, 6371–6381.

Stella, F., Treves, A., 2015. The self-organization of grid cells in 3D. eLife, 4, e05913.

Stensola, H., Stensola, T., Solstad, T., FrØland, K., Moser, M.-B., Moser, E.I., 2012. The entorhinal grid map is discretized. Nature. 492, 72–78.

Tolman, E.C., 1948. Cognitive maps in rats and men. Psychol. Rev. 55, 189–208.

Weiskopf, N., Hutton, C., Josephs, O., Deichmann, R., 2006. Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. Neuroimage. 33, 493–504.

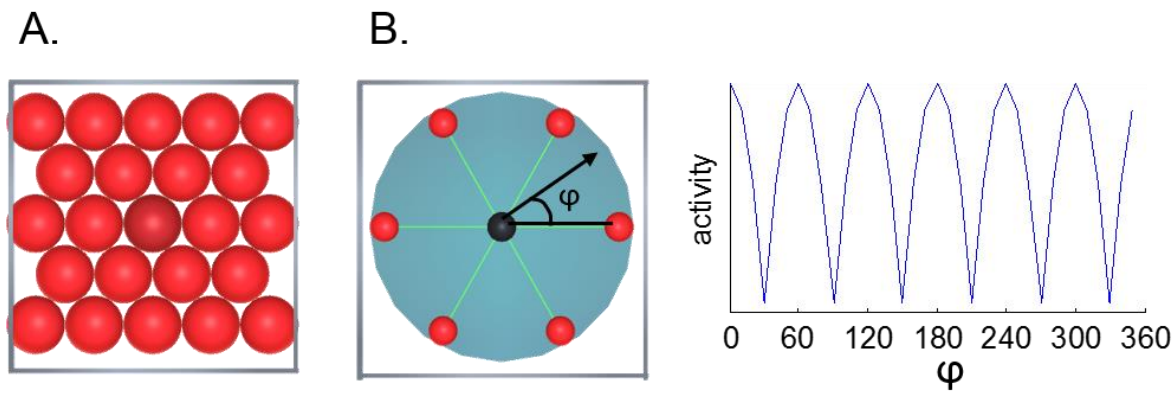

Fig 1. 2D grid cells. **A.** A grid cell fires at multiple locations (called grid fields, red circles) which correspond to the centre of circles closely fitted in a 2D box. **B.** A grid cell's activity is modulated by the animal's movement direction (black arrow) relative to the grid axis (green lines linking one grid field to its neighbouring six grid fields). φ denotes the angle between the movement direction and the grid axis. The grid cell fires more when an animal's moving direction is aligned to one of the grid axes. A grid axis is regularly displaced with 60° periodicity (right panel), therefore a grid cell's activity also shows a periodic response pattern depending on the animal's movement direction.

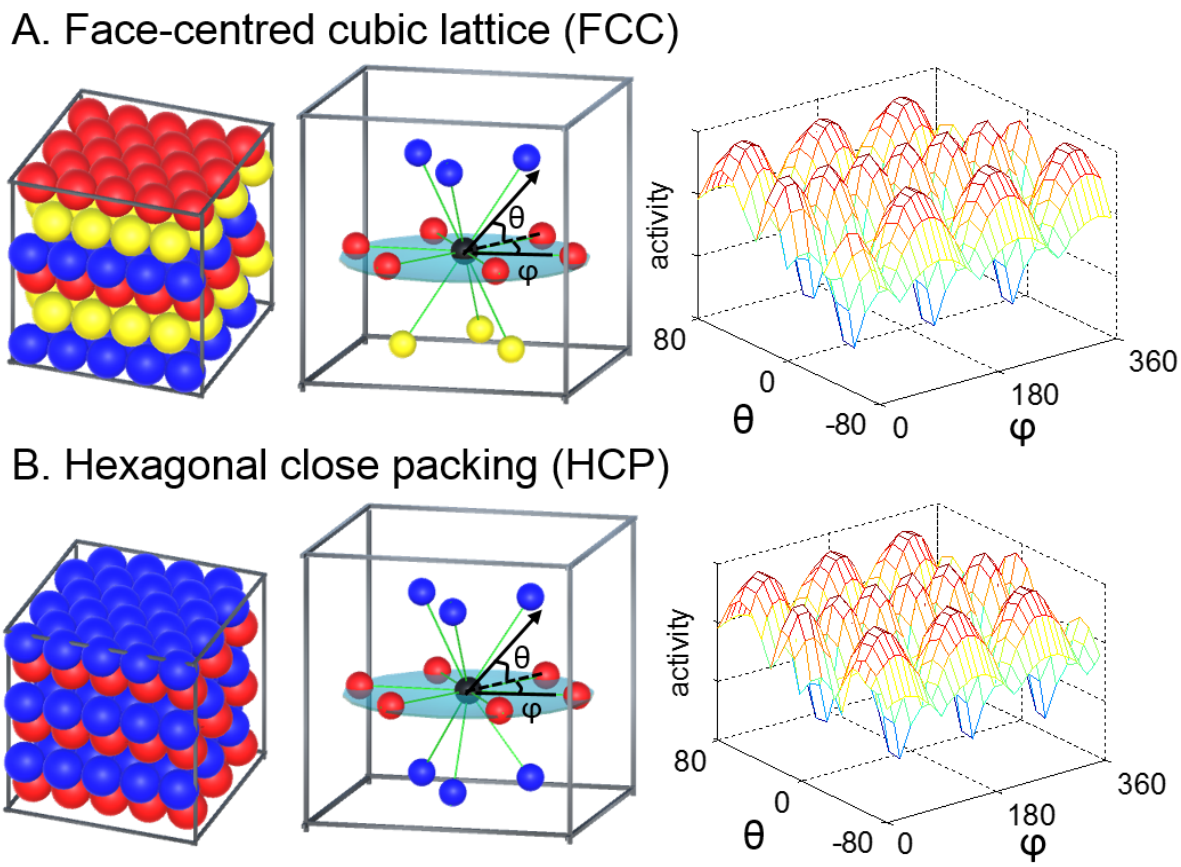

Fig. 2. 3D grid cells. **A.** FCC, **B.** HCP. Left panel, a 3D grid cell is proposed to fire at multiple locations (called grid fields, coloured spheres) which correspond to the centre of spheres closely fitted in a box. This is analogous to the spatial arrangement of oranges or cannonballs stacked as many as possible in the limited space. The highest packing density is known to enable the most efficient encoding of 3D space. The FCC and HCP arrangements have equally high packing density. The 3D arrangements can be viewed as 2D hexagonal lattices stacked on top of each other with a translational shift between the layers. FCC is composed of three repeating layers (blue, yellow and red spheres) and HCP is composed of two repeating layers (blue and red). FCC is a pure lattice in the mathematical sense, that is symmetric along the origin (blue and yellow spheres are symmetric) and can be described by three basis vectors. HCP lacks such symmetry (blue spheres on the top layer and bottom layer are not symmetric across the origin). Middle panels, grid axes are shown as green lines linking one grid field to the neighbouring 12 grid fields (coloured spheres). Similar to a 2D grid cell, it is expected that a grid cell's activity is modulated by the alignment between the animal's movement direction (black arrow) and the grid axis (green lines). Depending on the vertical (θ) and horizontal (φ) component of the movement direction relative to the grid axis, a grid cell's activity is expected to show a complex pattern of response (right panel).

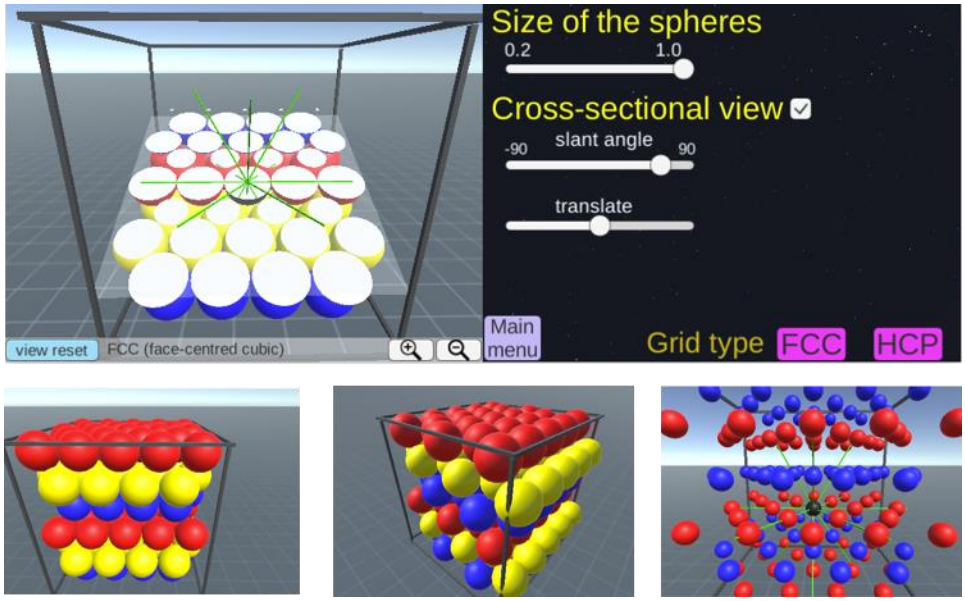

Fig. 3. Screenshots of our 3D grid cell visualization software. The software, including a manual, can be accessed here: www.fil.ion.ucl.ac.uk/Maguire/grid3D_gui. Users can change the viewpoint, the size of the spheres, and select cross-sectional views of grid fields. Users can also switch between the FCC and HCP arrangements.

A. Movement direction: (azimuth, pitch)=(-61°, 49°)

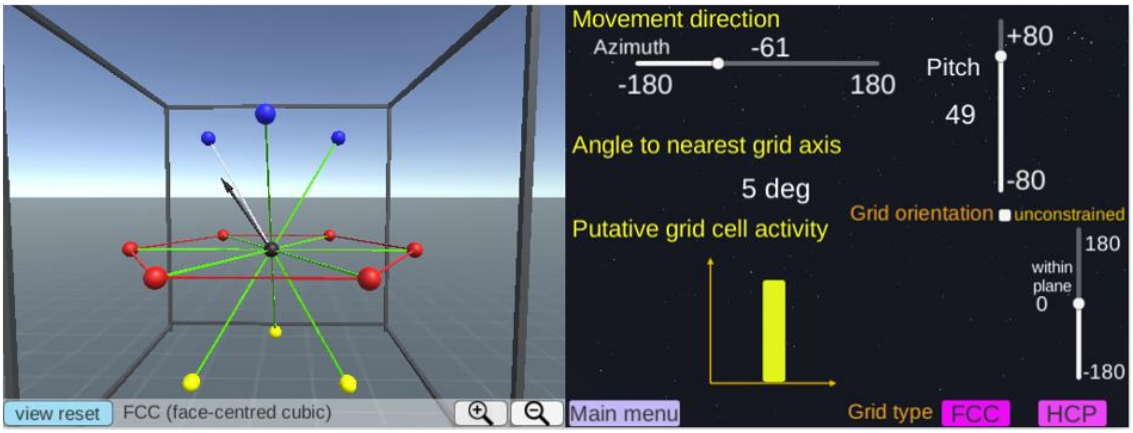

B. Movement direction: (azimuth, pitch)=(-61°, 24°)

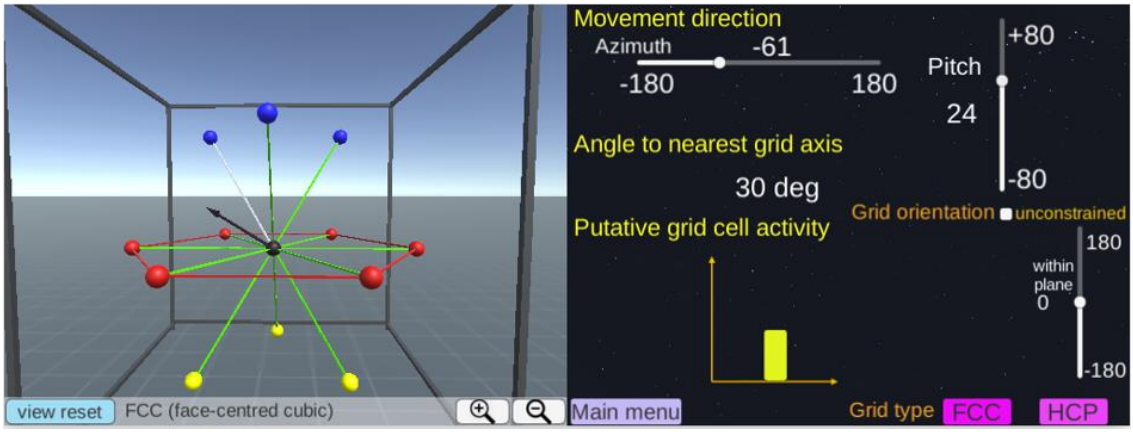

Fig. 4. A grid cell's activity is modulated by the movement direction relative to the grid axis. **A.** A participant's 3D movement direction (black arrow, left panel) is close to the grid axis (white line, left panel) with 5° deviation. Thus, the activity of grid cells is expected to be high (yellow bar graph, right panel). **B.** A participant's movement direction (black arrow, left panel) is far away from the grid axis (white line, left panel) with 30° deviation. Consequently, less activity is expected (yellow bar graph, right panel). Users can change the movement direction using the sliders on the right panel. Users can also switch between the FCC and HCP models and change the viewpoint.

A. Grid axis orientation=17°

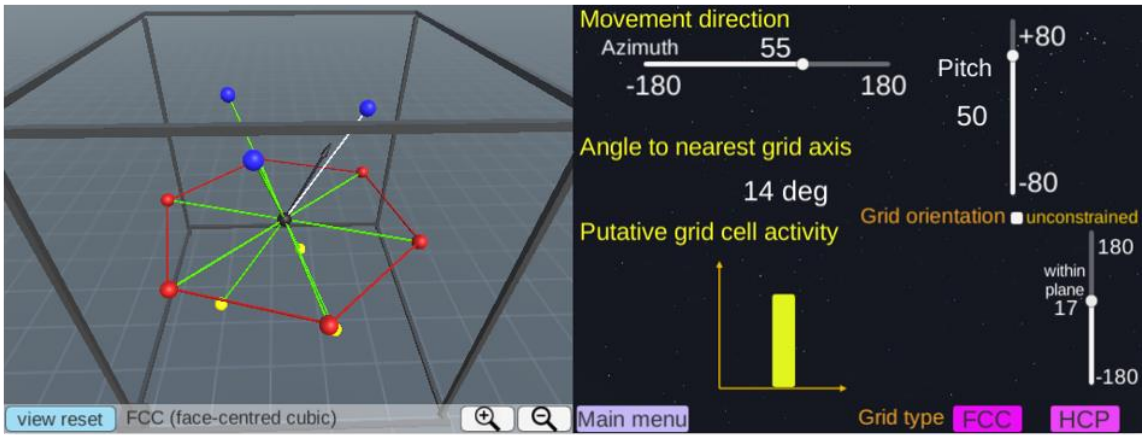

B. Grid axis orientation=56°

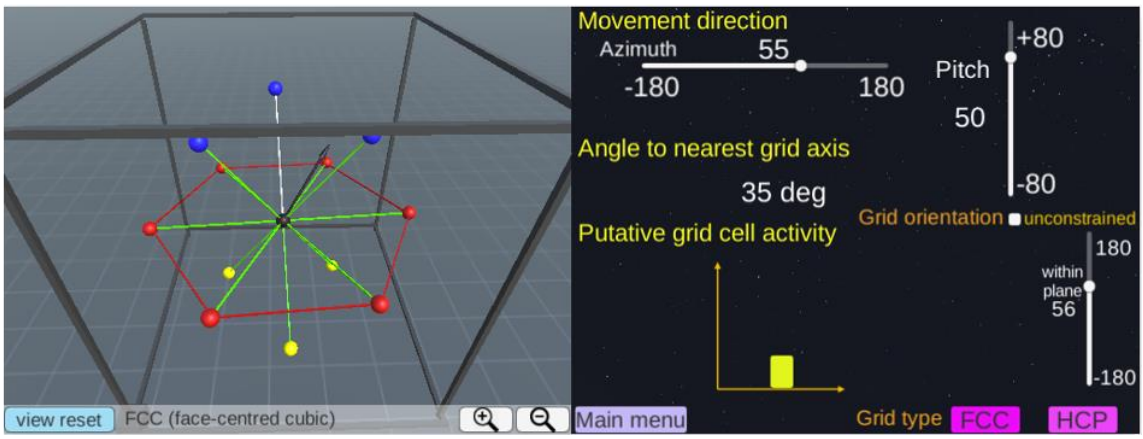

Fig. 5. The orientation of the grid axis relative to the environment. The movement direction (black arrow, left panel) is identical in **A** and **B** (azimuth=55°, pitch=50° in this example). However, grid cells are aligned differently relative to the environment (rectangular frame, left panels in **A** and **B**), meaning that the grid axes are rotated from each other (left panel). Thus, the grid alignment scores measured as the angle between the movement direction and the nearest grid axis differ (14° versus 35°), resulting in different amounts of grid activity (yellow bars graph, right panels in **A** and **B**). Users can change the orientation of the grid axis using the sliders on the right panel.

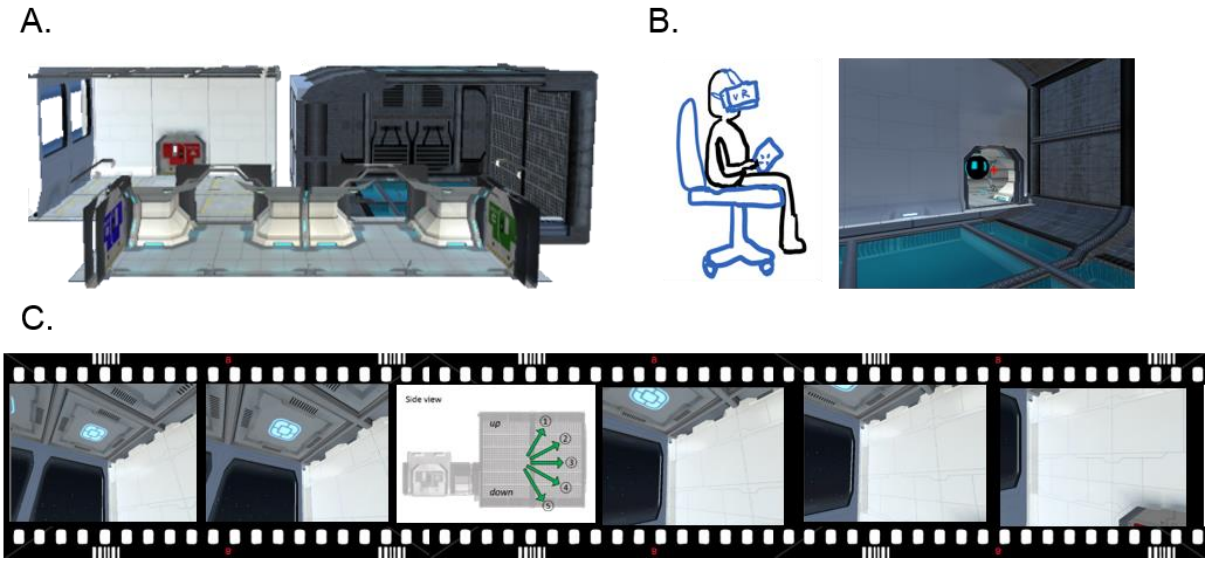

Fig. 6. The empirical experiment. **A.** A side view showing the layout of the virtual 3D spaceship. **B.** Prior to scanning, participants explored the 3D environment while wearing virtual reality goggles. **C.** During scanning, participants viewed a video rendered on a standard display. The video provided the feeling that participants were flying in a 3D trajectory within the spaceship. Participants occasionally indicated, via a keypad response, their vertical and horizontal movement direction when questioned.

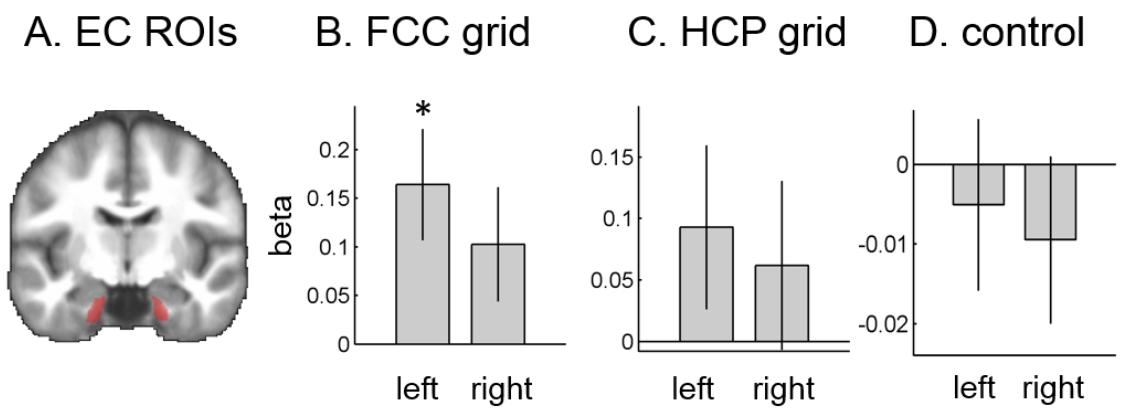

Fig. 7. 3D grid cell-like representations in the left entorhinal cortex. **A.** The anatomical ROIs where left and right EC are shown on the group-averaged structural MRI scan. **B.** The mean beta of the FCC grid cell model was significantly positive in the left EC. There was a similar trend in the right EC (p=0.09). **C.** The HCP model was not significant in either EC. **D.** The control direction (or associated view) encoding model was also not significant in either EC. *p<0.01

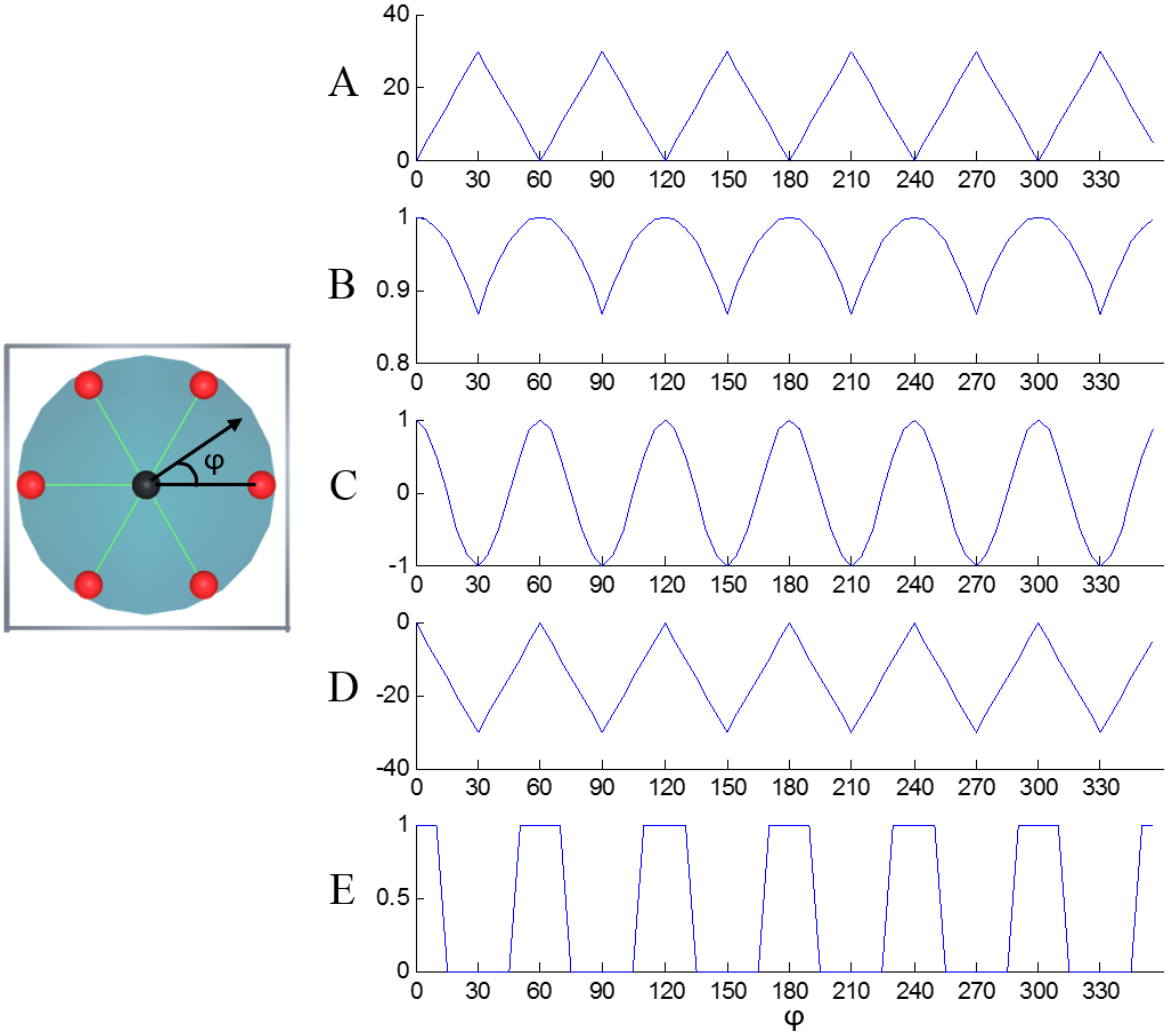

Supplementary Fig. S1. The grid alignment score and the fMRI signal. **A.** An angle between the movement direction (black arrow) and the nearest grid axis (green lines). **B.** A cosine of the nearest angle. **C.** A sinusoidal signal with a periodicity of 60 degrees. **D.** A linear signal (the negative of the nearest angle). **E.** A binary signal.