Rat on a sphere: Exploring alternate grid cell representations on a hemisphere

Muhammad Sahal Goolam (muhammadgoolam1@students.wits.ac.za)
School of Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg, South Africa

Clémentine Carla Juliette Dominé (clementine.domine98@gmail.com)
Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom

Luke Hollingsworth (l.hollingsworth@ucl.ac.uk)
Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom

Rodrigo Carrasco-Davis (Rodrigo.cd.20@ucl.ac.uk)
Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom

Geraud Nangue Tasse (geraud.nanguetasse1@wits.ac.za)
School of Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg, South Africa
Machine Intelligence and Neural Discovery Institute, University of the Witwatersrand, Johannesburg, South Africa

Devon Jarvis (devon.jarvis@wits.ac.za)
School of Computer Science and Applied Mathematics, University of the Witwatersrand, Johannesburg, South Africa
Machine Intelligence and Neural Discovery Institute, University of the Witwatersrand, Johannesburg, South Africa

Abstract

Grid cells in the medial entorhinal cortex are fundamental to mammalian spatial navigation, forming periodic, hexagonal firing patterns that serve as an internal map. While these cells exhibit isotropic firing in flat two-dimensional environments, their behaviour in three dimensions is more irregular, with the mechanisms governing spatial cognition beyond planar surfaces remaining largely unexplored. This study examines how spatial encoding adapts to a curved two-dimensional surface embedded in three dimensions. Specifically, we simulate a rat's movement on a hemispherical surface under varying gravitational strengths and consider three spatial encoding strategies based on the Successor Representation framework: (i) aligned with the sphere's curvature (logarithmic projection); (ii) relative to the horizontal plane tangent to the hemisphere at its origin (orthogonal projection) and (iii) using polar coordinates. Our results show that orthogonal projection does not elicit particularly salient or regular grid patterns, which compromises localisation. In contrast, logarithmic projection improves localisation while encoding in polar coordinates enhances grid regularity at a cost to localisation. These findings highlight the need to account for third-dimensional effects on grid cell organisation even on two-dimensional surfaces. Our model establishes a theoretical foundation for future empirical studies and offers extensions to computational models exploring grid cell formation on flat two-dimensional surfaces (embedded in two-dimensional space).

Keywords: Successor Representations; Grid Cells; Spherical Geometry

Introduction

While grid cells in the MEC have long been implicated in encoding odometric spatial information, emerging evidence points to their role in forming predictive spatial representations that integrate both geometric and topological features of the environment (Hafting et al., 2005; Stachenfeld et al., 2017). Beyond 2D space, there is growing evidence that grid cells exhibit highly anisotropic behaviour, diverging from the regular, grid-like patterns seen in planar environments (R. Hayman et al., 2011; R. M. Hayman et al., 2015; Grieves et al., 2021). Given that mammals are subject to gravitational forces and face more restricted movement vertically, it seems natural that such behaviour would arise in non-planar contexts (Sosa & Giocomo, 2021). However, this loss of global regularity is observed even in bats (Ginosar et al., 2021) with no clear generalisation of such ordered arrangement to the third dimension (Stella & Treves, 2015). In this work, we generalise the theoretical study of grid cell activity on a flat 2D plane, to a 2D surface embedded in 3D space and provide three specific predictions for spatial encoding in this setting using the Successor Representations (SR) framework (Dayan, 1993; Stachenfeld et al., 2017; Carvalho et al., 2024). We study the impact of gravity on the representations of space in the controlled setting of a hemisphere, reminiscent of more complex landscapes a rodent may navigate.

Method

We introduce a novel simulation environment that enables a rat agent to traverse the inner surface of a lower hemisphere, under the influence of gravitational pull (described in Figure 1). Using this simulation environment, we model three rodents randomly navigating the sphere over time, each distinctly employing one of the three spatial representations to highlight differences in their effects (described in Figure 2).

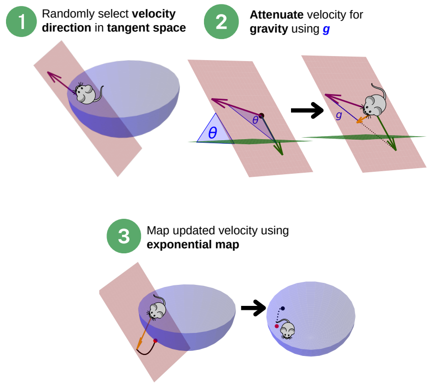

Figure 1: Single transition dynamic step of simulated motion of a rat in the hemisphere: *(1)* Randomly sample velocity direction for the current timestep from the instantaneous tangent plane (red shade) to the surface at the rat's location. *(2)* Adjust the velocity direction based on the degree of alignment θ of the tangent plane to the velocity direction using gravity parameter *g*. *(3)* Geometrically project the rat with the gravity-attenuated velocity to the sphere to obtain the new location of the rat.

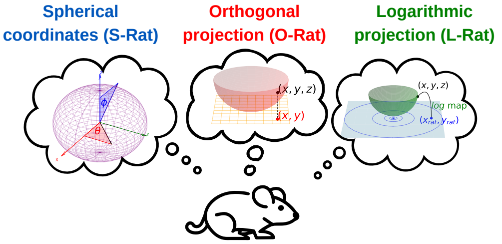

Figure 2: Three encoding strategies employed by the rat to encode the spatial location (x, y, z): *Spherical coordinates (S-Rat)* encodes the location in spherical coordinates (ϕ, θ), *Orthogonal projection (O-Rat)* discards the vertical dimension and encodes the location uniquely as (x, y) and *Logarithmic projection (L-Rat)* scales the orthogonal projection nonlinearly based on z to obtain (x', y').

Each rodent is simulated under varying values of gravitational strength $g \in [0, 1]$ using a standard reinforcement learning framework, in which the agent takes actions and observes state transitions, to estimate SR matrices for each. Each resulting SR matrix is then decomposed into its eigenvectors to obtain the grid cell predictions, as done by Stachenfeld et al. (2017).

Results

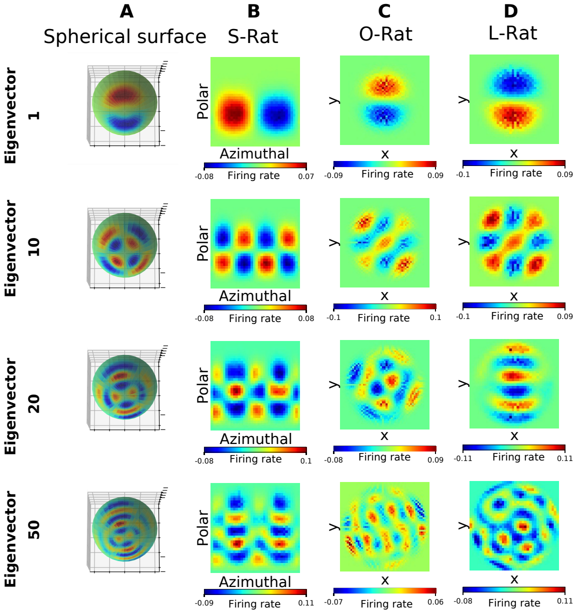

Figure 3: Grid cell predictions across differing spatial representations for $g = 0.15$. **A** illustrates grid cell firing on the spherical manifold, obtained from the grid cell activity of S-Rat mapped to 3D space. **B**, **C** and **D** illustrate the grid cell predictions for S-Rat, O-Rat and L-Rat respectively across various representative grid cells.

Figure 3 illustrates grid cell firing patterns on the spherical manifold and grid cell predictions under each of the three representations for fixed gravity. We observe varying periodicity in the grid cells within each representation (aligned with characteristic grid modules). As exhibited in Figure 3B, grid cells derived from S-Rat are observed to form generally stable and grid-like firing patterns that align very rigidly to the polar and azimuthal axes. Although the grid cell patterns maintain a regular structure in this coordinate frame, they become highly elongated when mapped onto the spherical surface, as seen in Figure 3A. The alignment of the firing patterns with the 2D manifold of the sphere is an expected result, as this parametrisation is isomorphic to a rat navigating a square room, where movement occurs in two basis directions which bound the environment. This suggests that such an encoding scheme inherently supports stable spatial representations by aligning the agent's movement with the underlying geometry. In contrast to S-Rat, both O-Rat and L-Rat display varying orientation across different eigenvectors. This variation fundamentally arises from the inherent symmetry of the environment when considered in Euclidean space, resulting in grid cells that align in a circular pattern with no strict alignment to any axes. Although less stable, this property reflects findings from the experiment by Stensola et al. (2015), where the absence of unique orientation cues in a circular room produced similar effects, highlighting a more ecologically-grounded contrast with the rigid alignment seen in S-Rat.

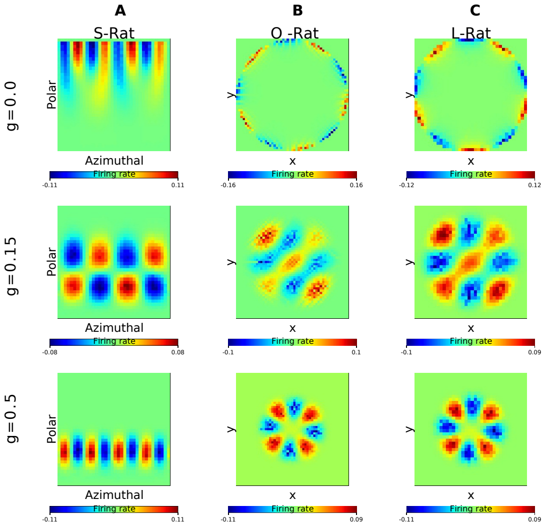

Figure 4: Grid cell predictions from the SR across different strengths of gravity g for a single eigenvector. **A**, **B** and **C** exhibit elicited grid cell predictions for S-Rat, O-Rat and L-Rat.

Figure 4 illustrates the effect of varying gravitational strength on grid cell patterns across each representation. In S-Rat, increasing gravity induces a clear translation of grid patterns along the polar axis, with firing patterns remaining highly periodic as exhibited in Figure 4A. This stability again stems from the geometric alignment of this representation to the surface itself, shifting in the direction of gravity. In contrast, this translation of firing fields in S-Rat manifests as a radial contraction of grid cell patterns in O-Rat and L-Rat as gravity increases (Figure 4B and C). There is a remarkable bias of the grid fields which point towards the base of the sphere, as seen for $g = 0.5$. Such a tendency suggests a learned environmental adaptation, where grid cells appear to encode a directional bias aligned with gravitational pull, anchoring spatial representations toward the base of the sphere.

These results conjunctively highlight how the bases of spatial representation and gravitational embedding fundamentally shape grid pattern emergence and localisation in 3D space, with different encodings yielding distinct spatial adaptations. Our work lays a theoretical foundation for empirical testing of these predictions in modulated environments, while also inviting exploration of alternative encoding schemes and refinements to models of grid cells.

References

Carvalho, W., Tomov, M. S., de Cothi, W., Barry, C., & Gershman, S. J. (2024). Predictive representations: Building blocks of intelligence. *Neural Computation*, *36*.

Dayan, P. (1993). Improving generalization for temporal difference learning: The successor representation. *Neural computation*, *5*.

Ginosar, G., Aljadeff, J., Burak, Y., Sompolinsky, H., Las, L., & Ulanovsky, N. (2021). Locally ordered representation of 3d space in the entorhinal cortex. *Nature*, *596*.

Grieves, R. M., Jedidi-Ayoub, S., Mishchanchuk, K., Liu, A., Renaudineau, S., Duvelle, É., & Jeffery, K. J. (2021). Irregular distribution of grid cell firing fields in rats exploring a 3d volumetric space. *Nature neuroscience*, *24*.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*.

Hayman, R., Verriotis, M. A., Jovalekic, A., Fenton, A. A., & Jeffery, K. J. (2011). Anisotropic encoding of three-dimensional space by place cells and grid cells. *Nature neuroscience*, *14*.

Hayman, R. M., Casali, G., Wilson, J. J., & Jeffery, K. J. (2015). Grid cells on steeply sloping terrain: evidence for planar rather than volumetric encoding. *Frontiers in psychology*, *6*.

Sosa, M., & Giocomo, L. M. (2021). Navigating for reward. *Nature Reviews Neuroscience*, *22*.

Stachenfeld, K. L., Botvinick, M. M., & Gershman, S. J. (2017). The hippocampus as a predictive map. *Nature neuroscience*, *20*.

Stella, F., & Treves, A. (2015). The self-organization of grid cells in 3d. *Elife*, *4*.

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