Intuitive planning: global navigation through cognitive maps based on grid-like codes

Alon B. Baram*, Timothy H. Muller*, James C.R. Whittington, Timothy E.J. Behrens
Wellcome Centre for Integrative Neuroimaging
University of Oxford

`alon.baram@merton.ox.ac.uk ; timothy.muller@magd.ox.ac.uk`

* denotes equal contribution

Abstract

It is proposed that a cognitive map encoding the relationships between objects supports the ability to flexibly navigate the world. Place cells and grid cells provide evidence for such a map in a spatial context. Emerging evidence suggests analogous cells code for non-spatial information. Further, it has been shown that grid cells resemble the eigenvectors of the relationship between place cells and can be learnt from local inputs. Here we show that these locally-learnt eigenvectors contain not only local information but also global knowledge that can provide both distributions over future states as well as a global distance measure encoding approximate distances between every object in the world. By simply changing the weights in the grid cell population, it is possible to switch between computing these different measures. We demonstrate a simple algorithm can use these measures to globally navigate arbitrary topologies without searching more than one step ahead. We refer to this as intuitive planning.

1 Introduction

In order to generate complex behaviour that is flexible to changes in circumstance and internal goals, it has been argued that animals maintain a model of the relationships between objects in the world, often termed a cognitive map [25]. In physical space the neural substrate for such a map is hypothesized to exist in the hippocampal formation in the form of place cells [20] and grid cells [15]. Place cells are found in the hippocampus and fire when an animal is in a specific location, whereas grid cells, found in the entorhinal cortex, amongst other regions, fire in a periodic hexagonal lattice over the environment.

It is not yet clear why it is advantageous to encode such a map both locally (using place cells that fire in specific locations) and globally (using grid cells that each fire in many different locations, but whose population can encode location precisely). However, one attractive recent suggestion is that these grid codes can be used to efficiently solve navigation problems as much of the necessary computation is already solved in the representation. Angles and distances between different points in the map can be extracted from grid codes using inexpensive computations that do not require searches through place cells [5, 24], endowing the agent with a sense of direction for navigating the map intuitively.

Tolman's original concept of a cognitive map was not, however, restricted to spatial relationships, but encompassed relationships that subserved flexible behaviour across multiple domains of life [25]. Indeed the hippocampal formation is critical for knowledge representation and generalisation in behavioural domains as distinct as social cognition [22], reward learning [14] and episodic memory [19]. Furthermore, neurons coding for non-spatial information in a fashion analogous to place cells have been found in human [21] and rodent [2] hippocampus, and signatures of grid-like activity have

been reported in rodents encoding time rather than space [16] and in humans navigating conceptual, rather than spatial, knowledge [6].

It has recently been suggested that these different properties of hippocampal neurons can be reconciled formally if the idea of a place cell is generalised to encode a state in a reinforcement-learning world [23]. This conceptual insight allows the machinery of reinforcement learning theory to be applied to understand neural coding in the hippocampal formation. Indeed it has been shown that global navigational strategies can be performed in place cells, if the scope of their coding is broadened to predict future states using successor representations [8] that relate them to the navigational goal [23]. Furthermore, signatures of predictive state representations have been observed in human entorhinal cortex in discrete state space environments [13], and there is evidence that humans use predictive representations for planning [18].

Here, we present a different technique for navigating arbitrary conceptual worlds. This technique allows place cells to retain local scope but instead, as in the spatial domain [5, 24], relies on grid-like codes to generate globally optimal (or nearly optimal) routes with efficient computations that do not require any sequential search. Similar to space, this technique uses grid-like codes to endow the agent with an intuitive sense of direction, but in arbitrary relational worlds. We therefore refer to it as intuitive planning.

2 Global distances from local learning

We rely on the recent suggestions that grid cells can be understood as the eigenvectors of the relationship between place cells. This has been demonstrated both using a PCA of place cell firing [10], and using the eigenvectors of successor representations [23]. This view is consistent with the recently acknowledged importance of the feedback connections from place cells to grid cells [26, 4, 17].

The key insight in the current work is that grid cells representing eigenvectors learnt from only local inputs can be used to construct distributions of distant future states as well as a global distance matrix that encodes the approximate distances between every pair of points in the world. We are proposing this eigenstructure contains the information necessary to compute all possible distances without any search.

To see why this is true, consider the adjacency matrix, A , that defines relationships between neighbouring states (or place cells) by encoding the weight of the edge between i and j , as element A_{ij} (figure 1). To compute the distribution of states after n steps, it is necessary to compute the matrix A^n . Notably, element A^n_{ij} is the number of walks of length n from states i to j . As it is simply a power of A , all matrices A^n share the same eigenvectors of A (see figure 2 for example eigenvectors of A of a simple two-dimensional world):

$$A^n = \varphi \lambda^n \varphi^{-1} \quad (1)$$

where the matrices φ and λ are the eigenvectors and eigenvalues of A , respectively.

Furthermore, any summation of $C_i A^{n_i}$, (where n_i is an integer, and C_i is an arbitrary scalar) will share these eigenvectors, so if we choose an approximate distance metric that is a weighted sum of future states, it is easily computable from the eigenvectors of A .

For example, consider the following two weighted sums:

$$\sum_{k=0}^{\infty} \frac{A^k}{k!} = e^A = \varphi e^{\lambda} \varphi^{-1} \quad (2)$$

and

$$\sum_{k=0}^{\infty} \gamma^k A^k = (I - \gamma A)^{-1} = \varphi (I - \gamma \lambda)^{-1} \varphi^{-1} \quad (3)$$

where $\gamma < \frac{1}{\lambda_{max}}$ and λ_{max} is the largest eigenvalue of A in modulus. Throughout the paper, γ is set to the commonly used [3, 1] $\gamma = \frac{0.85}{\lambda_{max}}$. The infinite sums in equations 2 and 3 converge to the matrix exponential (e^A) and resolvent ($(I - \gamma A)^{-1}$), respectively.

Figure 1: Construction of local and global knowledge from locally-learned eigenvectors. a) A two-way linear track. For the three columns on the right, bottom is a matrix and top is a slice through that matrix (red line) at state four. From left to right: Schematic of the environment with a mouse occupying state four; the associated adjacency matrix A ; A^n , where $n = 5$, and therefore A_{ij} is the number of possible walks of length n from state i to state j ; $-\log(C)$, where C_{ij} is the communicability from state i to state j . **b)** The same format as **a** for a two-way circular track. Red circles in the adjacency matrix mark the new adjacency between the ends of the linear track in **a** turning it in to a circular track. **Bottom:** Crucially, not only the matrix A containing local knowledge, but also A^n and $-\log(C)$ containing global knowledge can be constructed from the eigenvectors φ and eigenvalues λ of the matrix A . The local knowledge in state four is identical across the two environments (compare the slices through the adjacency matrices in **a** and **b**) but the global knowledge, obtained from the eigenvectors of A , is very different (compare slices in A^n and $-\log(C)$ across **a** and **b**).

In each of these sums the dominant contribution to element (i, j) is given by the shortest path between i and j . Longer paths are dramatically downweighted (particularly with the matrix exponent, equation 2). Hence element (i, j) contains information that closely relates to the shortest distance between i and j (this is shown empirically below, including some exceptions). In graph theory, these two summations are used to compute the communicability between nodes of a graph [11, 1, 12]. The second of these two summations (equation 3) is similar to the successor representation in reinforcement learning [8], which is computed from the transition rather than adjacency matrix.

Notably, in order to compute these functions it is simply necessary to recombine the eigenvectors of A (grid cells) with different weights (figure 1). Instead of using the original eigenvalues λ , λ^n will give the distribution of future states after n steps, e^λ will give the matrix exponent, and $(I - \gamma\lambda)^{-1}$ will give the resolvent (or successor representation). Further, these reweightings are inexpensive to implement because the number of eigenvalues is equal to the number of grid cells (eigenvectors), not the number of distances (relationships) that can be reconstructed.

Figure 2: **Example eigenvectors of the adjacency matrix of a simple two-dimensional world are periodic.** Eigenvectors of the adjacency matrix A of a discrete-state two-dimensional box projected back in to the space of the box have periodic waveforms. These are similar to the eigenvectors reported in [10] before a non-negativity constraint was applied that rendered them hexagonal. The permitted transitions between states encoded in A were in the up/down and left/right directions.

Figure 3: **Communicability measures in arbitrary topologies.** **a)** A one-way linear track. Left: The adjacency matrix A . Middle: the matrix $-\log(C)$ where element C_{ij} is the communicability between states i and j . Red line marks the column of a target state (note this is different to figure 1 where it is the row of the current state). Right: This column represented on the space defined by A . Red circle indicates target state. **b)** A two dimensional 10x10 box. The same format as **a** except the target column of $-\log(C)$ is represented as a heatmap on a two-dimensional environment with warmer colours corresponding to higher values. **c)** A box with barriers and a wormhole connecting states which are not adjacent in a two-dimensional space. Same format as **b**. States separated by a magenta line are not adjacent. White squares and circles mark entrance and exit gates of a wormhole, respectively: the states are adjacent yet distant in the regular two-dimensional topology. The conventions used in this figure will be used throughout the paper.

We emphasise that these measures work for arbitrary topology environments. They accurately reflect communicability in environments that cannot be easily mapped on to normal maps, such as those containing wormholes linking disparate parts of a map (figure 3), as well as environments with one-way links between states, which result in complex eigenvectors and eigenvalues (figure 3). Such arbitrary topologies are more likely those found in non-spatial cognitive maps.

Figure 4: Approximately linear relationship between intuitive distance and true distance. Intuitive distances between pairs of states are approximately linearly related to the true distance, as defined by the length of the shortest path. **a)** Intuitive distance reliably increases with true distance on a two way, one-dimensional track (the same as that reported in figure 1b). **b)** Spread of intuitive distances for a given true distance in a 10x10 two-dimensional box (the same as that reported in figure 3b), because communicability is influenced by the number of routes between states.

Both the matrix exponential and resolvent measures can be approximately linearly related to the true distance, as defined by the length of the shortest path between states, by taking their $-\log$ (figure 4), a measure we refer to as intuitive distance. This relationship is dependent on the sets of walks between pairs of states with the same distance. In some environments, the set of walks between a pair of states can be obtained by a permutation from the set of walks between another pair of states with the same true distance. In such a case there is no variance in the relationship between intuitive and true distance (figure 4a). However, in environments where there are many short, yet not necessarily shortest, paths between states, the communicability might differ between pairs of states with the same true distance, and hence so will the intuitive distance (figure 4b). This predicts failures to navigate the shortest path between states (see below).

3 Global navigation from local computation

Given current and target states, as well as the eigenvectors and eigenvalues, all calculations necessary for navigation are local: The agent needs to calculate which of the available next states has the shortest distance to the target. This can be done in the case of the matrix exponential by calculating for target state T and each neighbouring state, N :

$$C_{NT} = \sum_{v=1}^n \varphi_v(N) \varphi_v(T) e^{\lambda_v} \quad (4)$$

and in the case of the resolvent by calculating:

$$C_{NT} = \sum_{v=1}^n \varphi_v(N) \varphi_v(T) (1 - \gamma \lambda_v)^{-1} \quad (5)$$

where $\varphi_v(N)$ is the N th element of the v th eigenvector of the adjacency matrix associated with the eigenvalue λ_v and n is the number of eigenvectors. The agent can then proceed to the neighbouring state with the maximal C_{NT} , or equivalently minimal intuitive distance $-\log(C_{NT})$. This procedure finds trajectories that are approximately globally optimal without any global (tree) search. Because it stores the relational structure, the algorithm can flexibly reroute when the values of candidate target states change, for example due to changing internal goals. Further, parallel neural circuitry can naturally implement the summations in equations 4 and 5.

4 Navigation using intuitive planning

We constructed several discrete-state mazes and tested whether intuitive planning could successfully navigate between start and target locations. We demonstrate the algorithm successfully navigates the

shortest route to the target in several two-dimensional, undirected mazes (figure 5). Cognitive maps encoding non-spatial relationships, however, will likely have arbitrary topologies. We show that the algorithm generalises to arbitrary topologies as its performance is not affected by the insertion of a wormhole in to the maze, rendering it no longer topologically two-dimensional (figure 5). Successful navigation of the shortest route to the target, as defined by the minimum number of transitions to reach it, is achieved in all of the above mazes when using both the matrix exponential and resolvent measures. The mazes in figure 5 can be mapped on to two-dimensional spaces for visualisation purposes. The algorithm also achieved successful navigation to the target in randomly generated, unweighted graphs, provided there was a route to the target (figure 5).

Since the relationship between intuitive distance and true distance is only approximately linear (figure 4), it is possible to construct mazes where the algorithm successfully navigates to the target, but does not take the shortest route. This is achieved by increasing the communicability of initial states not associated with the shortest route, but with multiple longer routes to the target (figure 6).

Figure 5: Navigation in arbitrary topology worlds. Successful navigation from start (red square) to target (red circle) states in two-dimensional worlds with barriers and differing start and target states (a-c). Introducing wormholes that connect disparate parts of the map and therefore warp it rendering it no longer described by a simple map does not affect the ability to successfully navigate to the target (d-f). Nor does the introduction of a one-way wormhole (e-f), resulting in complex eigenvalues and eigenvectors. In all navigation simulations, state transitions were permitted in the up/down and left/right directions (unless, for example, blocked by a barrier), and black circles denote the agent's path. g) Successful navigation in a randomly generated graph.

Figure 6: Failure to navigate the shortest path. Introducing multiple wormholes that cause multiple longer routes to the target to pass through a nearby state (magenta star) that is not in the direction of the shortest route (green star) causes the agent to take a longer route to the target (a). With fewer such wormholes, the shortest route is taken (b)

5 Experimental predictions

This theoretical work makes some experimental predictions:

- Unlike explicit planning [9], such intuitive planning will not require a precession through place cells leading to a goal.
- In intuitive planning the distance to a goal does not influence the difficulty of the next decision. The opposite is true in explicit planning (tree-search). We therefore predict that for local goals, distance to goal (where explicit planning is possible) will predict reaction time, but for distant goals it will not.
- The intuitive distance measures are most likely to differ from shortest-path measure when there are many possible routes to the goal that pass through one nearby state, but where the shortest route passes through a different state (figure 6). This predicts a precise pattern of planning errors.

6 Discussion

We believe that the observation that grid cell patterns can be related to the eigenspectrum of local place cell transitions [10, 23] is an important one. We have shown that one consequence of this observation is that grid cell patterns contain information not only about the current location but also about all possible future locations. It is therefore possible to compute distance metrics between every pair of locations in the state space using grid patterns that have been learnt from only one-step transitions. With one choice of weighting of future states, this distance metric allows rapid computation of successor-like representations [8, 23] but with states (place cells) that retain local scope. Other distance metrics can also be computed. Like successor representations, this provides only a partial solution to the model-based reinforcement-learning problem leaving open the problem of how values are assigned to different target locations. The eigenvector-based computations suggested here are easily implemented in parallel, and therefore suggest natural implementations in recurrent neural circuits that support grid codes. Indeed, Corneil & Gerstner [7] have elegantly implemented a planning algorithm based on the eigendecomposition of the SR of the environment in a recurrent attractor neural network. The global metrics we have presented here allow global navigation through arbitrary topology maps without searching more than one step ahead. We refer to this as intuitive planning. Such a strategy could be combined, for example, with hierarchical representations of state space to further reduce computational complexity. Together, such strategies demonstrate how the expensive computations that underlie state-space navigation can be finessed by efficient choices of neuronal representation.

References

- [1] M. Aprahamian, D. J. Higham, and N. J. Higham. Matching exponential-based and resolvent-based centrality measures. *Journal of Complex Networks*, 4(2):157–176, 2016.
- [2] D. Aronov, R. Nevers, and D. W. Tank. Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature*, 543(7647):719–722, 2017.
- [3] M. Benzi and C. Klymko. Total communicability as a centrality measure. pages 1–27, 2013.
- [4] T. Bonnevie, B. Dunn, M. Fyhn, T. Hafting, D. Derdikman, J. L. Kubie, Y. Roudi, E. I. Moser, and M. B. Moser. Grid cells require excitatory drive from the hippocampus. *Nature Neuroscience*, 16(3):309–317, 2013.
- [5] D. Bush, C. Barry, D. Manson, and N. Burgess. Using Grid Cells for Navigation. *Neuron*, 87(3):507–520, aug 2015.
- [6] A. O. Constantinescu, J. X. O’Reilly, and T. E. J. Behrens. Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292):1464–1468, 2016.
- [7] D. S. Corneil and W. Gerstner. Attractor Network Dynamics Enable Preplay and Rapid Path Planning in Maze-like Environments. *Advances in Neural Information Processing Systems*, pages 1675–1683, 2015.
- [8] P. Dayan. Improving Generalisation for Temporal Difference Learning: The Successor Representation. *Neural Computation*, 5(4):613–624, 1993.
- [9] K. Diba and G. Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10):1241–1242, 2007.

- [10] Y. Dordek, D. Soudry, R. Meir, and D. Derdikman. Extracting grid cell characteristics from place cell inputs using non-negative principal component analysis. *eLife*, 5(MARCH2016):1–36, 2016.
- [11] E. Estrada and N. Hatano. Communicability in complex networks. pages 1–12, 2008.
- [12] E. Estrada and D. J. Higham. Network properties revealed through Matrix Functions. 52(4):696–714, 2010.
- [13] M. M. Garvert, R. J. Dolan, and T. E. Behrens. A map of abstract relational knowledge in the human hippocampal–entorhinal cortex. *eLife*, 6:1–20, 2017.
- [14] A. Gilboa, M. Sekeres, M. Moscovitch, and G. Winocur. Higher-order conditioning is impaired by hippocampal lesions. *Current Biology*, 24(18):2202–2207, 2014.
- [15] T. Hafting, M. Fyhn, S. Molden, M. B. Moser, and E. I. Moser. Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806, 2005.
- [16] B. J. Kraus, M. P. Brandon, R. J. Robinson, M. A. Connerney, M. E. Hasselmo, and H. Eichenbaum. During Running in Place, Grid Cells Integrate Elapsed Time and Distance Run. *Neuron*, 88(3):578–589, 2015.
- [17] R. F. Langston, J. A. Ainge, J. J. Couey, C. B. Canto, T. L. Bjerknes, M. P. Witter, E. I. Moser, and M.-B. Moser. Development of the Spatial Representation System in the Rat. *Science*, 328(5985):1576–1580, jun 2010.
- [18] I. Momennejad, E. M. Russek, J. H. Cheong, M. M. Botvinick, N. D. Daw, and S. J. Gershman. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9):680–692, 2017.
- [19] L. Nadel and M. Moscovitch. Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology*, 7:217–227, 1997.
- [20] J. O’Keefe and L. Nadel. *The hippocampus as a cognitive map*. 1978.
- [21] R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045):1102–1107, 2005.
- [22] F. Sams-Dodd, B. K. Lipska, and D. R. Weinberger. Neonatal lesions in the rat ventral hippocampus result in hyperlocomotion and deficits in social behaviour in adulthood. *Psychopharmacology*, 132(3):303–310, 1997.
- [23] K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman. The hippocampus as a predictive map. *Nature Neuroscience*, 20(11):1643–1653, 2017.
- [24] M. Stemmler, A. Mathis, and A. V. M. Herz. Connecting multiple spatial scales to decode the population activity of grid cells. *Science Advances*, 1(11):e1500816–e1500816, dec 2015.
- [25] E. C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189–208, 1948.
- [26] T. J. Wills, F. Cacucci, N. Burgess, and J. O’Keefe. Development of the Hippocampal Cognitive Map in Prewanling Rats. *Science*, 328(5985):1573–1576, jun 2010.