PREDICTION AND GENERALISATION OVER DIRECTED ACTIONS BY GRID CELLS

Anonymous authors
Paper under double-blind review

ABSTRACT

Knowing how the effects of directed actions generalise to new situations (e.g. moving North, South, East and West, or turning left, right, etc.) is key to rapid generalisation of action to new situations. Markovian tasks can be characterised by a state space and a transition matrix and recent work has proposed that neural grid codes provide an efficient representation of the state space, as eigenvectors of a transition matrix reflecting diffusion across states, that allows efficient prediction of future state distributions. Here we address the ability to predict the effects of directed actions across new tasks, irrespective of differences in local transition structure, to allow fast generalisation and form a "sense of direction". We extend the eigenbasis prediction model, utilising tools from Fourier analysis, to prediction over arbitrary translation-invariant directed transition structures (i.e. displacement and diffusion), showing that a single set of eigenvectors can support predictions over arbitrary directed actions via action-specific eigenvalues. We show the equivalence between the generalised prediction framework and traditional models of grid cells firing driven by self-motion to perform path integration using oscillatory interference models (via Fourier components as velocity-controlled oscillators) or continuous attractor networks (via analysis of the update dynamics). We thus provide a unifying framework for the role of the grid system in predictive planning, sense of direction and path integration. The resulting model allows a single grid-like representation to support efficient prediction over directed transitions in spatial and non-spatial tasks, supporting generalisable inference over directed actions across different tasks.

1 INTRODUCTION

A "cognitive map" encodes relations between objects and supports flexible planning (Tolman [32]), with hippocampal place cells and entorhinal cortical grid cells thought to instantiate such a map (O'Keefe and Dostrovsky [28]; Hafting et al. [20]). Each place cell fires when the animal is near a specific location, whereas each grid cell fires periodically when the animal enters a number of locations arranged in a triangular grid across the environment. Together, this system could support representation and flexible planning in state spaces where common representational structure is preserved across states and tasks, affording generalisation and inference, e.g., in spatial navigation where Euclidean transition rules are ubiquitous (Whittington et al. [35]).

Recent work suggests that place cell firing provides a local representation of state occupancy, while grid cells comprise an eigenbasis of place cell firing covariance (Dordek et al. [15]; Stachenfeld et al. [31]; Sorscher et al. [30]; Kropff and Treves [25]). Accordingly, grid cell firing patterns could be learned as eigenvectors of a symmetric (diffusive) transition matrix over state space, providing a basis set enabling prediction of occupancy distributions over future states. The "intuitive planning" operates by replacing multiplication of state representations by the transition matrix with multiplication of each basis vector by the corresponding eigenvalue (Baram et al. [2]; Corneil and Gerstner [12]). Thus a distribution over state space represented as a weighted sum of eigenvectors can be updated by re-weighting each eigenvector by its eigenvalue to predict future state occupancy.

Fast prediction and inference given different actions across environments is vital and economical for the survival of animals, and intuitive planning, in its original form, supports such ability under a single transition structures, most often corresponding to symmetrical diffusion. Here we generalise

prediction with grid codes to include the set of translation-invariant directed transitions by proposing an extended model of intuitive planning based on Fourier analysis, with the assumption that the underlying environment is periodically bounded. By defining a "sense of direction" quantity, the model is capable of predicting the effect of arbitrary directed transitions, irrespective of local details, e.g, obstacles, hence proposing a new computational role for the medial entorhinal grid cells. In the second half of the paper, we showed, on the theory-level, how the generalised prediction model corresponds to the mechanistic path-integration models of grid cells (where the grid firing patterns are formed via pattern formation).

2 "INTUITIVE PLANNING" WITH SYMMETRICAL TRANSITIONS

Intuitive planning represents distribution over the state space in terms of a weighted sum of the (real) eigenvectors of the (symmetric) transition matrix, so that the effect of one step of the transition dynamics on the distribution can be predicted by reweighting each of the eigenvectors by the corresponding eigenvalue. And this generalises to calculating the cumulative effect of discounted future transitions (Baram et al. [2]).

Specifically, consider a transition matrix, $T \in \mathbb{R}^{N\times N}$, $T_{ss'} = \mathbb{P}(s_{t+1} = s'|s_t = s)$ where s_t encodes the state at time t and N is the number of states. Then, T^n is the n-step transition matrix, and has the same set of eigenvectors as T. Specifically, the eigendecomposition of T and T^n are:

$$T = Q\Lambda Q^{-1}, \qquad T^n = Q\Lambda^n Q^{-1} \tag{1}$$

where each column of the matrix Q is an eigenvector of T and $\Lambda = diag(\sigma_P(T))$, where $\sigma_P(T)$ is the set of eigenvalues of T. Similarly any polynomial in T, $p(T)$, shares the same set of eigenvectors as T and the set of eigenvalues $\sigma_P(p(T)) = p(\sigma_P(T))$. Hence:

$$\sum_{k=0}^{\infty}(\gamma T)^k = (I - \gamma T)^{-1} = Q\text{diag}(\mathbf{w})Q^{-1}, \quad \text{where } \mathbf{w} = \left\{\frac{1}{1-\gamma\lambda}, \text{ for } \lambda \in \sigma_P(T)\right\} \tag{2}$$

The resolvent form (eq. 2) is an infinite discounted summation of transitions, which under a policy and ransition structure corresponding to diffusion, is equivalent to the successor representation (SR, Fig. 1E) with discounting factor γ (Dayan [14]; Stachenfeld et al. [31]). The SR has been shown to be useful for navigation via gradient ascent of the future probability of occupying the target state, and represent linear relationship with the true underlying Euclidean distances in spatial tasks (hence "intuitive planning", see Fig. 1 and Fig. 2D-E).

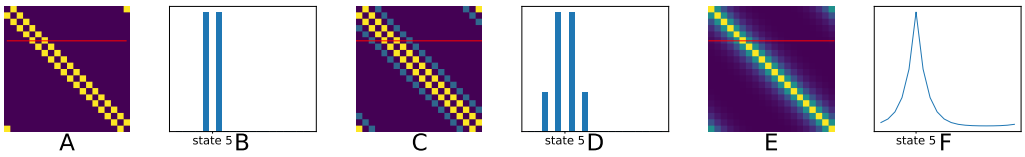

Figure 1: **Demonstration of intuitive planning on a diffusive transition task on a ring track with 20 states**. The transition matrix allows prediction of future state occupancy. E.g., if an agent diffuses equally one step left and right on a ring, the transition matrix T is shown in (A), and the distribution of next states from a randomly chosen state (state 5), $\mathbb{P}(s_{t+1} = s'|s_t = 5)$, equals to the row of T is shown in (B), and similar for T^3 (C-D) showing predicted distribution over the next three time steps. (E) shows the the resolvent form/SR (eq. 2) computed from the eigenstructure of local transition matrix, and (F) shows that the SR values can be used for gradient-based navigation.

The eigenvectors of the diffusion transition matrix generally show grid-like patterns, suggesting a close relationship to grid cells. However, intuitive planning cannot predict the effects of directed actions or drift (i.e., variable asymmetric transition structure), which is useful in itself and also corresponds to the presumed role of grid cells in path integration. Moreover, predictions over different directed actions would require separate realisations of eigendecomposition, hence incurring high computational costs that undermines its biological plausibility. In Section 3 we unify the prediction and path integration approaches by exploiting translation invariant symmetries to generalise across actions, using a single common eigenbasis and cheaply calculated updates via action-dependent eigenvalue.

3 FLEXIBLE PLANNING WITH DIRECTED TRANSITIONS

Updating state representations to predict the consequences of arbitrary directed actions is an important ability of mobile animals, known as path integration (McNaughton et al. [27]). To generalise to directed transition structures, we consider the transition dynamics corresponding to translation (drift) and Gaussian diffusion with arbitrary variance (including 0, equivalent to plain translation). We impose an assumption of periodic boundary conditions, so that the transition structure is translation invariant, hence leading to circulant transition matrices. We note, however, that a task without periodic boundary conditions could be implemented in a subset of the full state space to be unaffected by this assumption (Fig. 3 F-G).

Consider a 2D rectangular environment with length L and width W where each state is a node of the unit square grid, then the transition matrix can be represented by $\mathbf{T} \in \mathbb{R}^{LW \times LW}$, with each row the vectorisation ($\mathbf{vec}(\cdot)$) of the matrix of transition probabilities starting from the specified location, i.e., $\mathbf{T}[i,:] = \mathbf{vec}[\mathbb{P}(s(t+1)|s(t)=i)]$, where $\mathbf{T}$ is constructed by considering the 2D state space as a 1D vector and concatenating the rows ($i = xL + y$ for $(x,y) \in [0, W-1] \times [0, L-1]$), see Fig. 2A.

The transition matrix is circulant due to the translation invariance of the transition structure (see Appendix Prop. A.1), and takes the following form:

$$\mathbf{T} = \begin{bmatrix} T_0 & T_{LW-1} & \cdots & T_2 & T_1 \\ T_1 & T_0 & T_{LW-1} & \cdots & T_2 \\ \vdots & T_1 & T_0 & \ddots & \vdots \\ T_{LW-2} & \cdots & \ddots & \ddots & T_{LW-1} \\ T_{LW-1} & T_{LW-2} & \cdots & T_1 & T_0 \end{bmatrix} \tag{3}$$

The normalised eigenvectors of the circulant matrix $\mathbf{T} \in \mathbb{R}^{N \times N}$ ($N = LW$) are the vectors of powers of Nth roots of unity (the Fourier modes):

$$\mathbf{q}_k = \frac{1}{\sqrt{N}} \left[1, \omega_k, \omega_k^2, \cdots, \omega_k^{N-1}\right]^T \tag{4}$$

where $\omega_k = \exp(\frac{2\pi i}{N} k)$, for $k = 0, \ldots, N-1$, and $i = \sqrt{-1}$. Hence the matrix of eigenvectors (as the columns), $F = (\mathbf{q}_0, \mathbf{q}_1, \ldots, \mathbf{q}_{N-1})$, is just the (inverse) discrete Fourier transform matrix (Bracewell [3]), where $F_{kj} = \omega_j^k$ for $0 \le k, j \le N-1$. The Fourier modes projected back onto the $L \times W$ 2D spatial domain are plane waves, as shown in Fig. 2G, with wavevector determined by the value of k that specifies the direction and spatial frequency of each plane wave (see Appendix B). We can immediately compute the corresponding eigenvalues for the eigenvectors in eq. 4 (equivalent to taking the discrete Fourier transform (DFT) of the first row (or column) of T, see Bracewell [3]):

$$\lambda_m = \sum_{j=0}^{N-1} T_j \omega_j^m, \quad \text{for } m = 0, \ldots, N-1 \tag{5}$$

where $\{T_0, \ldots, T_{N-1}\}$ are the N unique elements that fully specifies the circulant matrix T (eq. 3).

We can then utilise tools from Fourier analysis for efficient updating of the eigenvalues whilst leaving the universal eigenbasis unaffected. For a transition matrix $\mathbf{T}^{\mathbf{v}}$ corresponding to an arbitrary action (translation velocity) $\mathbf{v} = (v_x, v_y)$, each row of $\mathbf{T}^{\mathbf{v}}$ is again a circulant, but shifted version of the corresponding row vector of the symmetric transition matrix corresponding to zero drift velocity, $\mathbf{T}^0$. Specifically, the first rows of the two matrices are related as follows:

$$\mathbf{T}^0(k + v_x L + v_y) = \mathbf{T}^v(k), \quad \text{for } k = 0, \ldots, N-1 \tag{6}$$

Given the eigenvalues for $\mathbf{T}^0$, $\Lambda^0 = \left[\lambda_0^0, \lambda_1^0, \ldots, \lambda_{N-1}^0\right] \in \mathbb{C}^N$ (via the DFT of the first row of $\mathbf{T}^0$, eq. 5), we can immediately derive the eigenvalues of $\mathbf{T}^{\mathbf{v}}$, $\Lambda^{\mathbf{v}}$, via a one-step update based on the Fourier shift theorem (Bracewell [3]) without recomputing the eigendecomposition:

$$\Lambda^{\mathbf{v}}[k] = \exp\left(\frac{2\pi i}{N}(v_x L + v_y)k\right) \Lambda^0[k] \quad \text{for } k = 0, \ldots, N-1 \quad \text{for arbitrary } \mathbf{v} \tag{7}$$

$$\text{i.e., } \Lambda^{\mathbf{v}} = \Phi_{\delta^{\mathbf{v}}} \Lambda^0, \quad \Phi_{\delta(\mathbf{v})} = \left[1, \omega_{\delta(\mathbf{v})}, \omega_{\delta(\mathbf{v})}^2, \ldots, \omega_{\delta(\mathbf{v})}^{N-1}\right], \quad \text{where } \delta(\mathbf{v}) = v_x L + v_y \tag{8}$$

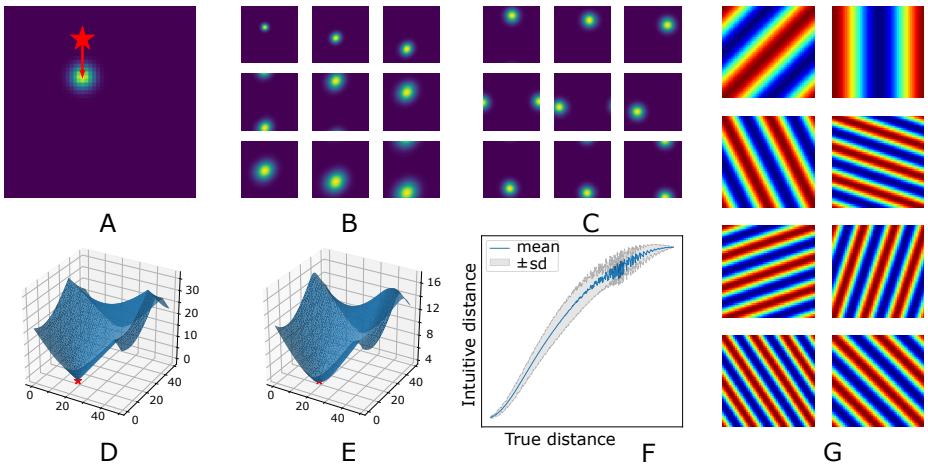

Figure 2: **Demonstration of our method in a 2D 50×50 environment with periodic boundary conditions**. **A**: Transition matrix $\mathbb{P}(x_{t+1} = (i', j')|x_t = (i, j))$ starting from a randomly chosen state $((8, 10)$; red star) with drift velocity 10 units southward and Gaussian diffusion (red arrow); **B**: Usage of the eigenbasis (eq. 4) and our analysis (eq. 8) for predicting the distribution over future states given the transition structure given in A, showing successive changes to state occupancy; **C**: Application of our model to translation-only transition structures (drift velocity $v = (3, 3)$); **(D-E)**: Ground-truth shortest distance to state $(8, 20)$ (red star; D) as a function of of starting location, and distance estimated under the intuitive planning framework (E) using plain diffusion; **F**: Estimated distance measure (E) v.s. the corresponding ground-truth distance (D) over all pairs of states; **G**: Examples of 2D Fourier modes (real parts shown, see Appendix Fig. 6 for the top 100 eigenvectors).

Hence allows path integration by reweighting the common set of eigenvectors at each timestep by the updated eigenvalues corresponding to the current drift velocity (eq. 8). Note that additionally, T^0 can include diffusion, thus reweighting by the eigenvalues of the diffusive transition matrix also allows tracking of increasing uncertainty.

Utilising the fixed eigenbasis (eq. 4) and the respective eigenvalues (eq. 8) for arbitrary transition structures, we can make efficient prediction for the distribution of future state occupancy with respect to arbitrary action (see Figs. 2B-C).

Adding translation to the translation-invariant transition matrix does not change the set of eigenvectors - allowing one set of eigenvectors (Fourier modes) to support prediction for actions in all directions (or plain diffusion), hence prediction of effects of directed actions can be efficiently generalised across environments. We define a "sense of direction", as the angle of the transitions that maximise the future probability of reaching the target state given an initial state, which is modelled by the SR matrix.

$$\theta^* = \arg\max_{\theta} \sum_j \frac{\exp[2\pi i(\mathbf{s}_G - \mathbf{s}_0) \cdot \mathbf{k}_j]}{1 - \gamma D_j \exp[2\pi i \mathbf{v}_\theta \cdot \mathbf{k}_j]} \tag{9}$$

where γ is the discounting factor, $D_j, j = 1, \cdots, LW$ are the eigenvalues for the symmetric diffusion transition matrix, $\mathbf{k}_j, j = 1, \ldots, LW$ are the wavevectors for the j-th Fourier components, $\mathbf{s}_0, \mathbf{s}_G$ are the coordinates of the start and goal states, and $\mathbf{v}_\theta = (v\cos(\theta), v\sin(\theta))$ represents the velocity (with speed v and head direction θ). We see that the "sense of direction" supports generalisation of predictions of effects of actions across all environments with the same translation-invariant transition structure, i.e., such predicted effects ignore any differences in local structure. See Appendix B for the derivation of eq. 9.

The proposed prediction framework can be applied to flexible planning under arbitrary drift velocity as demonstrated in Fig. 3 (A-E). An agent is trying to navigate towards a goal state in a windy grid world. The navigation is based on following the ascending "gradient" of the SR for occupancy of the target state (the resolvent metric, eq. 2). The SR computed from the (asymmetric) transition matrix including the effects of diffusion and wind (Fig. 3 B) based on our analysis (eq. 8) leads straight

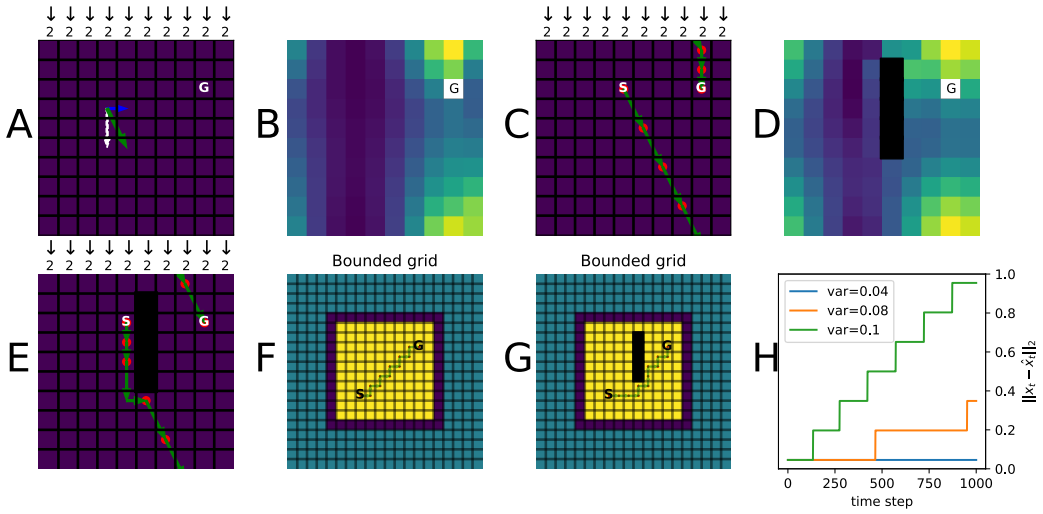

Figure 3: **Application to spatial navigation in grid worlds.** **A**: The 10×10 windy grid world task environment, with toroidal boundary conditions and a constant external force causing two units of displacement southward acting on every state (white arrow: wind; blue arrow: one-step action rightward; green arrow: actual displacement, G: goal state);**B**: Estimated SR (using eq. 2 and eq. 8) under diffusion plus the wind effect (color indicates the strength of future probabilities of occupying the target states); **C**: Path following the diffusion SR plus the wind effect; **D**: Updated diffusion SR plus wind given the insertion of a barrier (dark blocks); **E**: Path following the updated SR; **F**: Navigation in the task space (S_0, yellow) with boundaries (magenta) embedded in a larger pseudo state space (S_p, blue), utilising the Fourier modes computed from S_p; **G**: Navigation in S_0 with inserted local obstacles (black); **H**: Path integration performance as a function of time step under different levels of variance (uncertainty).

to the target (Fig. 3 C). Given the definition of the SR, we could efficiently adjust the SR matrix to accommodate local changes in the state space, e.g., insertion of a barrier, using the Woodbury inversion formulae (Piray and Daw [29]); and again in this case, the agent correctly adjusts for the wind as well as taking the shortest path around the inserted wall (Fig. 3 D-E).

Our generalised prediction model requires translation-invariance, consistent with periodic boundary conditions. Here we address the applicability of our model to tasks without periodic boundary conditions. Consider the task state space, S_0, being embedded into a larger, periodically bounded pseudo state space S_p, at least twice as large in each dimension as S_0 (Fig. 3 F). We again follow the previous procedures, utilising the Fourier modes, this time computed on S_p, to perform predictions in S_0 (Fig. 3 F-G), and the performance is unaffected. We can also use our model for path integration (Fig. 3 H, see Section 4) in S_0, by taking velocity inputs (from a randomly generated path in the grid world) to update the state occupancy distribution utilising our previous analysis (eq. 8). We observe that the path integration performance is strongly sensitive with respect to the noise, with perfect path integration when the uncertainty is low up to 1000 time steps (empirical results shows the perfect path integration is maintained up to the uncertainty level reaches 0.075), and monotonically non-decreasing path integration error when the uncertainty is higher, where the uncertainty could represent the degree of familiarity of the agent in the environment. We thus propose that a computational role for the neural grid codes: generating a "sense of direction" (eq. 9) even in new or bounded environments, via utilising a Fourier basis for a larger toroidal pseudo space.

4 A UNIFYING FRAMEWORK FOR MODELS OF GRID CELL FIRING

Computational models of grid cells can be roughly categorised into two classes: normative models, which specify a neural architecture, learning rule and training tasks, from which grid-like patterns emerge (Dordek et al. [15]; Cueva and Wei [13]; Stachenfeld et al. [31]; Banino et al. [1]; Baram et al. [2]); and mechanistic models, which suggest that the grid-like patterns arise from specifically

designed architecture (e.g., recurrently connected networks) via pattern formation while performing path integration without the necessity of learning (Fuhs and Touretzky [18]; Burak and Fiete [4]; Welday et al. [34]; Burgess et al. [7]; Burgess [6]; Corneil and Gerstner [12]). Our focus so far has been on proposing a flexible and efficient extension of the normative prediction model of grid cells (Baram et al. [2]; Dordek et al. [15]; Stachenfeld et al. [31]) to arbitrary directed transitions. In this section we discuss the second aim: unifying the predictive framework with the path integration models - both the CAN models and the OI models.

4.1 Relation to continuous attractor network models of grid cells

Several unified frameworks for grid cell models have been proposed, one of the most prominent by Sorscher et al. [30], where they proved the equivalence between the maximisation of a spatial representation objective function under the normative models and the pattern formation dynamics of the CANs under the models of path integration. Their analysis, however, is implicit about the equivalence between the two classes of models when there is non-zero velocity inputs, corresponding to non-trivial asymmetric velocity-dependent connectivities in the CANs (Fuhs and Touretzky [18]; Burak and Fiete [4]), which is the key element for performing path integration. We explicitly address such equivalence utilising our framework based one Fourier analysis. Assuming that a grid cell is formed via linear combination of selected Fourier modes (e.g., 6 Fourier modes at $\frac{\pi}{3}$ degrees increments with same frequencies):

$$g = \sum_{j=1}^{G} w_j f_j \tag{10}$$

where f_j's are the selected Fourier modes with weights w_j. Inserting the single-cell representation into eqs. (14, 15) from [30] we have the equations governing the dynamics of the grid cells under the normative and mechanistic path integration models respectively as following.

$$\frac{dg}{dt} = \begin{cases} -\alpha g + \gamma(\sum_{j=1}^{G} \lambda_j w_j f_j) + \mu, & g > 0 \\ -\alpha g + \sigma(\gamma(\sum_{j=1}^{G} \lambda_j w_j f_j) + \mu), & g = 0 \end{cases} \tag{11}$$

$$\tau \frac{dg}{dt} = -g + \sigma(\mathbf{W} g + b(v)) \tag{12}$$

where λ_j is the corresponding eigenvalue of f_j with respect to the (symmetric) transition matrix, $\mathbf{T}^0$, $\sigma(\cdot)$ is the Heaviside function, $\mathbf{W}$ is the recurrent connectivity matrix, $b(v)$ is the velocity-dependent feedforward input to the grid cell under continuous attractor model which involves a constant baseline term and a velocity dependent term (Burak and Fiete [4]), α is the multiplicative constant, γ and μ are the multiplicative constant corresponding to the additive penalty terms in the (singel-neuron) Lagrangian under the normative models:

$$\mathcal{L} = g^T \mathbf{T}^0 g - \gamma g^T g + \mu \mathbb{1}^T g \tag{13}$$

Now suppose the agent is moving with velocity v, which causes changes in the transition structure, given our previous formulation based on Fourier analysis (eq. 8) and using the first-order Taylor expansion around 0 for approximating the exponential, eq. 11 and eq. 12 becomes equivalent, by identifying the corresponding terms τ, W and $b(v)$, see Appendix B for detailed proof of equivalence. The resulting expression for $b(v)$ is directly related to the Fourier modes inputs, hence it correlates with the underlying grid codes with the baseline μ, hence is equivalent to the asymmetric velocity-dependent parts of the recurrent connectivity matrix in the CAN.

However, we further note that the the unifying framework in Sorscher et al. [30] only addressed the unifying framework between the normative models and the pattern formation path integration models that utilise recurrent neural networks, while leaving out the discussion of another class of path integration model that yields grid firing patterns via coincidence detection the oscillatory interference models (Burgess [6]; Burgess et al. [7]; Welday et al. [34]). In the following subsection we aim to draw qualitative equivalence between the prediction framework based on our formulation and the oscillatory interference models, in terms of path integration and phase precession in grid cells.

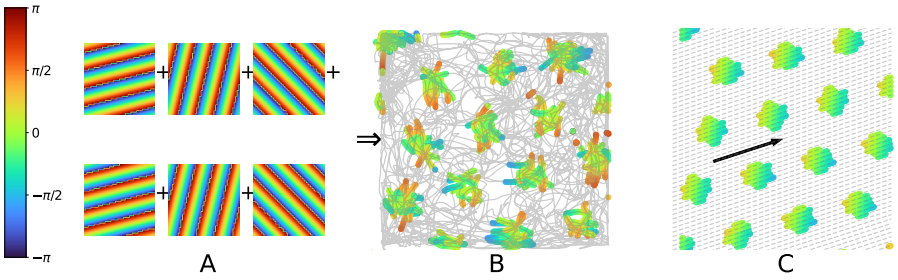

Figure 4: **Generated grid-like patterns given our model**. **A**: 6 input Fourier modes with the same wavelength (complex phases shown); **B**: Simulated grid cell firing given a real rat trajectory (gray line) using coincidence detection and baseline modulation given the 6 Fourier modes inputs in A, each spike is represented by a colored scatter with the color indicates the corresponding "theta" phase at the spiking time; **C**: Simulated grid cell firing fields given multiple runs in the same direction (black arrow) showing theta phase precession.

4.2 Relation to oscillatory interference models of grid cells

For the oscillatory interference models of grid cells, path integration is achieved by tracking the phases of the velocity controlled oscillators (VCOs), which encode movement speed and direction by variations in burst firing frequency (Burgess et al. [7]; Burgess [6]). The VCOs generate grid-like firing patterns via coincidence detection by the grid cells (Hasselmo [22]; Welday et al. [34]). The variation of frequency with velocity produces a phase code for displacement, enabling the modelled grid cells to perform path integration. Namely, VCOs change their frequencies relative to the baseline according to the projection of current velocity, $v(t)$, onto the VCO's "preferred direction", $\mathbf{d}$:

$$f_a(t) = f_b(t) + \beta \mathbf{v}(t) \cdot \mathbf{d} \tag{14}$$

where β is a positive constant, and $f_b(t)$ is the baseline frequency (the $4 - 11$Hz EEG theta rhythm). It follows that the VCOs perform linear path integration since the relative phase between VCO's frequency and baseline at time t, $\phi_{ab}(t) = \phi_a(t) - \phi_b(t)$, is proportional to the agent's displacement in the preferred direction:

$$\phi_{ab}(t) - \phi_{ab}(0) = \int_0^t 2\pi[f_a(\tau) - f_b(\tau)]d\tau = 2\pi\beta[\mathbf{x}_t - \mathbf{x}_0] \cdot \mathbf{d} \tag{15}$$

where $\mathbf{x}(t)$ is the agent's location at time t. The interference of VCOs whose preferred directions differ by multiples of $\pi/3$ generate triangular grid-like patterns, provides an explanation of theta phase precession in grid cells (Hafting et al. [21]; Burgess [6]) and complements the attractor dynamics given by symmetrical connections between grid cells (Bush and Burgess [8]).

We simulated the firing of a grid cell synthesised from 6 Fourier modes inputs (Fig. 4 A-B), each firing a spike at its complex phase in the current state, as a leaky integrate and fire neuron performing coincidence detection using a real trajectory of a rat exploring a 50cm $\times$ 50cm box. At each time step (corresponding to one theta cycle), the phase of each Fourier mode is updated according to eq. 8 given the current velocity, and fires a spike at this phase if it is within the interval $[-\pi/4, \pi/4]$ (modelling modulation by the baseline theta oscillation). The grid cell fires a spike at the current location if the integrated inputs reach a threshold. Note that we could also simulate a set of grid cells, with different offsets (depending on the initial phases of the Fourier modes) and different scales and orientations (depending on the choice of Fourier modes), such that the set of grid cells, like the Fourier modes, comprises a basis for the state space and do so on the basis of path integration (for which environmental inputs are also required to prevent error accumulation [6; 5]).

Grid cells show "precession" in their firing phase relative to theta as the animal moves through the firing field (signalling distance travelled within the field) (Hafting et al. [21]; Jeewajee et al. [23]; Climer et al. [11]), which is captured by our model, similar to the OI model. The Fourier modes whose wavevectors that are aligned with the current direction of translation advance in phase as the movement progresses. Phase precession results from assuming that the Fourier modes aligned with movement direction are dominant influence on grid cell firing (c.f. those aligned to the reverse

direction). The baseline "theta frequency" corresponds to the mean rate of change of phase of all Fourier components and so could vary (for noise reduction, see Burgess and Burgess [5]; Burgess [6], without precluding phase coding (Eliav et al. [17]; Bush and Burgess [9])). By simulating random straight runs, we can see clear late-to-early phase precession (Fig. 4 C), analogous to what was observed in OI models (Burgess [6]).

We thus conclude that the OI model and our model perform path integration or prediction in the same way: the OI model changes the phase of each VCO relative to baseline corresponding to the component of translation along the VCO's preferred direction, and this is exactly analogous to multiplying the Fourier mode by its eigenvalue: shifting its complex phase corresponding to the component of translation along its wavevector (eq. 8); hence showing the qualitative equivalence between the prediction framework and OI models.

5 Discussion

Understanding how different actions affect the agent's state across environments is essential for generalisation. Existing models are capable of such prediction when the underlying transition structure of the task is symmetric. Here we propose a spectral model based on neural grid codes that is able to generalise prediction to arbitrary directed transitions, by exploiting translation invariant symmetries to find a single common eigenbasis (Fourier modes) that generalises prediction across actions via translation-specific eigenvalues. The resulting prediction framework is computationally efficient and enables predictions beyond translation-invariant state spaces (e.g., bounded state space with local obstacles). We thus propose a new computational role for the grid codes in addition to path integration function: generating generalisable "sense of direction" for navigation, irrespective of environmental boundaries and local structures, while other (e.g., fronto-parietal) areas of the brain are responsible for detecting and avoiding obstacles (Edvardsen et al. [16]; Maguire et al. [26]), or that direct translations occur in clear regions of state space between "bottlenecks" that are also identified by spectral methods (Stachenfeld et al. [31]). Our model also offers potential generalisation to non-spatial tasks such as transitive inference (Von Fersen et al. [33]), we provide a self-contained trivial example in Appendix D.

The current model applies to translation-invariant transition structures, and use of the Fourier shift theorem to calculaet eigenvalues also assumes a Euclidean state space. However, we demonstrated ways of generalising planning to bounded of locally non-translation invariant transition structures in Section 3. We note that machine learning methods based on a similar premise (creating a single representation to support planning via multiple different actions) might work even when the transitions are not strictly translation invariant (e.g., family trees, see Whittington et al. [35]).

We build upon the previous work of unifying the normative and path integration models of grid cells firing by Sorscher et al. [30], extending their discussions and showed the equivalence of our generalised prediction framrework with models of path integration by grid cells based on continuous attractor dynamics or oscillatory interference.

The current model predicts future state occupancy from the transition matrix, future works could look at the reversed direction: vector navigation, i.e., inferring the translation between two locations given the phase codes for each (see Bush et al. [10]; corresponding to discrinminative and generative models). It might be that the grid cells/Fourier modes provide navigational direction, while other processes deal with obstacles (Edvardsen et al. [16]; Maguire et al. [26]), or that direct translations occur in clear regions of state space between "bottlenecks" that are also identified by spectral methods (Stachenfeld et al. [31]).

An important extension of the current work is the generalisation of the spectral model based on Fourier analysis to non-spatial tasks, one promising direction is to consider Fourier analysis on groups of operators utilising group-theoretic knowlegde (Kondor and Trivedi [24]; Gao et al. [19]). Such generalisation might enable application to domains far from physical space, including aspects of logical or linguistic processing.

ACKNOWLEDGEMENTS

The authors thanks support from various funding resources. The authors declare no competing financial interests.

REFERENCES

[1] A. Banino, C. Barry, B. Uria, C. Blundell, T. Lillicrap, P. Mirowski, A. Pritzel, M. J. Chadwick, T. Degris, J. Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.

[2] A. B. Baram, T. H. Muller, J. C. Whittington, and T. E. Behrens. Intuitive planning: global navigation through cognitive maps based on grid-like codes. *bioRxiv*, page 421461, 2018.

[3] R. N. Bracewell. *The Fourier transform and its applications*, volume 31999. McGraw-Hill New York, 1986.

[4] Y. Burak and I. R. Fiete. Accurate path integration in continuous attractor network models of grid cells. *PLoS computational biology*, 5(2), 2009.

[5] C. P. Burgess and N. Burgess. Controlling phase noise in oscillatory interference models of grid cell firing. *Journal of Neuroscience*, 34(18):6224–6232, 2014.

[6] N. Burgess. Grid cells and theta as oscillatory interference: theory and predictions. *Hippocampus*, 18(12):1157–1174, 2008.

[7] N. Burgess, C. Barry, and J. O'keefe. An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9):801–812, 2007.

[8] D. Bush and N. Burgess. A hybrid oscillatory interference/continuous attractor network model of grid cell firing. *Journal of Neuroscience*, 34(14):5065–5079, 2014.

[9] D. Bush and N. Burgess. Advantages and detection of phase coding in the absence of rhythmicity. *Hippocampus*, 2020.

[10] D. Bush, C. Barry, D. Manson, and N. Burgess. Using grid cells for navigation. *Neuron*, 87(3): 507–520, 2015.

[11] J. R. Climer, E. L. Newman, and H. M. E. Phase coding by grid cells in unconstrained environments: two-dimensional phase precession. *European Journal of Neuroscience*, 38(4): 2526–2541, 2013.

[12] D. S. Corneil and W. Gerstner. Attractor network dynamics enable preplay and rapid path planning in maze–like environments. In *Advances in neural information processing systems*, pages 1684–1692, 2015.

[13] C. J. Cueva and X.-X. Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. *arXiv preprint arXiv:1803.07770*, 2018.

[14] P. Dayan. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5(4):613–624, 1993.

[15] 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:e10094, 2016.

[16] V. Edvardsen, A. Bicanski, and N. Burgess. Navigating with grid and place cells in cluttered environments. *Hippocampus*, 2019.

[17] T. Eliav, M. Geva-Sagiv, M. M. Yartsev, A. Finkelstein, A. Rubin, L. Las, and N. Ulanovsky. Nonoscillatory phase coding and synchronization in the bat hippocampal formation. *Cell*, 175 (4):1119–1130, 2018.

[18] M. C. Fuhs and D. S. Touretzky. A spin glass model of path integration in rat medial entorhinal cortex. *Journal of Neuroscience*, 26(16):4266–4276, 2006.

[19] R. Gao, J. Xie, S.-C. Zhu, and Y. N. Wu. A representational model of grid cells based on matrix lie algebras. *arXiv preprint arXiv:2006.10259*, 2020.

[20] 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.

[21] T. Hafting, M. Fyhn, T. Bonnevie, M.-B. Moser, and E. I. Moser. Hippocampus-independent phase precession in entorhinal grid cells. *Nature*, 453(7199):1248–1252, 2008.

[22] M. E. Hasselmo. Grid cell mechanisms and function: contributions of entorhinal persistent spiking and phase resetting. *Hippocampus*, 18(12):1213–1229, 2008.

[23] A. Jeewajee, C. Barry, V. Douchamps, D. Manson, C. Lever, and B. N. Theta phase precession of grid and place cell firing in open environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635):20120532, 2014.

[24] R. Kondor and S. Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. *arXiv preprint arXiv:1802.03690*, 2018.

[25] E. Kropff and A. Treves. The emergence of grid cells: Intelligent design or just adaptation? *Hippocampus*, 18(12):1256–1269, 2008.

[26] E. A. Maguire, N. Burgess, J. G. Donnett, R. S. Frackowiak, C. D. Frith, and J. O'Keefe. Knowing where and getting there: a human navigation network. *Science*, 280(5365):921–924, 1998.

[27] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser. Path integration and the neural basis of the'cognitive map'. *Nature Reviews Neuroscience*, 7(8):663–678, 2006.

[28] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 1971.

[29] P. Piray and N. D. Daw. A common model explaining flexible decision making, grid fields and cognitive control. *bioRxiv*, page 856849, 2019.

[30] B. Sorscher, G. Mel, S. Ganguli, and S. Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. In *Advances in Neural Information Processing Systems*, pages 10003–10013, 2019.

[31] K. L. Stachenfeld, M. M. Botvinick, and S. J. Gershman. The hippocampus as a predictive map. *Nature neuroscience*, 20(11):1643, 2017.

[32] E. C. Tolman. Cognitive maps in rats and men. *Psychological review*, 55(4):189, 1948.

[33] L. Von Fersen, C. D. Wynne, J. D. Delius, and J. E. Staddon. Transitive inference formation in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, 17(3):334, 1991.

[34] A. C. Welday, I. G. Shlifer, M. L. Bloom, K. Zhang, and H. T. Blair. Cosine directional tuning of theta cell burst frequencies: evidence for spatial coding by oscillatory interference. *Journal of Neuroscience*, 31(45):16157–16176, 2011.

[35] J. C. Whittington, T. H. Muller, S. Mark, G. Chen, C. Barry, N. Burgess, and T. E. Behrens. The tolman-eichenbaum machine: Unifying space and relational memory through generalisation in the hippocampal formation. *bioRxiv*, page 770495, 2019.