

Accurate representation for spatial cognition using grid cells

Nicole Sandra-Yaffa Dumont (ns2dumont@uwaterloo.ca) & Chris Eliasmith (celiasmith@uwaterloo.ca)

Computational Neuroscience Research Group, University of Waterloo
200 University Ave W, Waterloo, ON N2L3G1, Canada

Abstract

Spatial cognition relies on an internal map-like representation of space provided by hippocampal place cells, which in turn are thought to rely on grid cells as a basis. Spatial Semantic Pointers (SSP) have been introduced as a way to represent continuous spaces and positions via the activity of a spiking neural network. In this work, we further develop SSP representation to replicate the firing patterns of grid cells. This adds biological realism to the SSP representation and links biological findings with a larger theoretical framework for representing concepts. Furthermore, replicating grid cell activity with SSPs results in greater accuracy when constructing place cells. Improved accuracy is a result of grid cells forming the optimal basis for decoding positions and place cell output. Our results have implications for modelling spatial cognition and more general cognitive representations over continuous variables.

Keywords: Neural Engineering Framework; Semantic Pointer Architecture; spatial semantic pointer; spatial representation; spiking neural networks; cognitive maps; grid cells; place cells

Introduction

In the 1940's, Tolmann observed that rats learning maze layouts discovered shortcuts, and proposed that they navigate using a "cognitive map" (Tolman, 1948). Evidence for the idea that animals have an internal allocentric representation of space has accumulated, specifically from electrophysiological recordings in the hippocampal-entorhinal circuit. Two main types of spatially responsive cells have been observed. Place cells that become active at specific regions in an environment were discovered in the hippocampus (O'Keefe & Nadel, 1978). Grid cells in the medial entorhinal cortex (MEC) activate at hexagonally tiled points in an environment (Hafting, Fyhn, Molden, Moser, & Moser, 2005). The resolution and orientation of the pattern varies among grid cells. These cells are believed to provide a basis for hippocampal cognitive maps.

Recent experiments have indicated that the neural mechanisms behind spatial navigation may be a general mechanism for encoding continuous variables (Aronov, Nevers, & Tank, 2017; Constantinescu, O'Reilly, & Behrens, 2016). This evidence has been used to support the idea of cognitive spaces – a generalization of cognitive maps applied to concepts (Bellmund, Gärdenfors, Moser, & Doeller, 2018). Cognitive representation in a geometric fashion is central to the Semantic Pointer Architecture (SPA); (Eliasmith, 2013). In this framework, compressed vectors – termed semantic point-




Figure 1: Similarity plots of SSPs decoded from a 361-dimensional memory vector storing 4 different spatially-bound objects. The ground truth locations of the items are marked by symbols.

ers – are used to represent anything from low level visual features to high level concepts. Cognitive semantic pointers can be bound together to build syntactic structures which live in a continuous vector space where distance represents concept similarity. These vectors can be encoded, decoded, and computed with the collective activity of spiking neurons using the Neural Engineering Framework (NEF); (Eliasmith & Anderson, 2003).

A core feature of semantic pointers is their ability to be composed in hierarchies using binding and bundling operations. Until recently, cognitive semantic pointers, like most vector symbolic architectures (VSAs), have been restricted to representing discrete structures (e.g., trees, graphs, lists, etc.). However, continuous structures are needed for spatial representation, and many other tasks. The hippocampus plays a key role in spatial representation. As a result, representing continuous spatial structures with SPA in a manner consistent with hippocampal data may improve our understanding of the mechanisms underlying spatial cognition.

Background

Spatial Semantic Pointers

In Komer, Stewart, Voelker, and Eliasmith (2019), Spatial Semantic Pointers (SSPs) were introduced as a way of encoding spatial coordinates, or, more generally, representing any continuous variables within the SPA framework. Let us begin by considering discrete VSA representations using circular con-

volution (Plate, 1994). In this case, an integer index, $n \in \mathbb{N}$, of a structure (e.g., this could be the ordered position in a list) can be encoded by binding a semantic pointer to itself n times.

$$B^n = \underbrace{B \circledast B \circledast \dots \circledast B}_{n \text{ times}} \quad (1)$$

To ensure that the magnitude of the pointer does not change with repeated binding, a unitary B vector is used. A unitary vector is defined to be a vector whose discrete Fourier transform has entries with an absolute value of one. With such a representation in hand, we could bind objects in an ordered list to these index vectors, and sum across them to generate a distributed representation of the full list (Choo & Eliasmith, 2010). Notably, any such structure has discrete ‘slots’.

Interestingly, this repeated self-binding can be generalized to encode a continuous variable $k \in \mathbb{R}$. In particular, we can define “fractional binding” by recalling the relationship between circular convolution and the discrete Fourier transform.

$$B^k = \mathcal{F}^{-1} \{ \mathcal{F} \{ B \}^k \}, \quad k \in \mathbb{R} \quad (2)$$

where the exponentiation of $\mathcal{F} \{ B \}^k$ is element-wise. We now have a means of generating continuous ‘slots’ in our representation, where the distance between such slots is determined by the difference between k . To represent multiple continuous variables, such as an x, y coordinate in 2-dimensional space, we can bind the continuous representation of each dimension together,

$$S(x, y) = X^x \circledast Y^y = \mathcal{F}^{-1} \{ \mathcal{F} \{ X \}^x \odot \mathcal{F} \{ Y \}^y \}, \quad (3)$$

where $\odot$ is the Hadamard (element-wise) product and $X, Y \in \mathbb{R}^d$ are randomly chosen unitary vectors. This vector $S(x, y)$ is a Spatial Semantic Pointer (SSP). The vectors X, Y are called the bases vectors of the representation. The vector space in which a SSP lives can be thought of as a cognitive map.

This spatial representation has several useful properties. Binding can be used to shift the SSP,

$$S(x_1, y_1) \circledast S(x_2, y_2) = S(x_1 + x_2, y_1 + y_2) \quad (4)$$

Spatial representations can be attached to other semantic pointers (e.g., vector representations of objects, landmarks, shapes, colours, etc.). For example, a semantic pointer representing an object can be bound with the SSP that represents its location, $OBJ \circledast S(x, y)$. A set of such location-tagged object representations can be stored in a memory vector.

$$M = \sum_i OBJ_i \circledast S(x_i, y_i) \quad (5)$$

To obtain the location of an object, the semantic pointer of the object can be unbound from the memory vector.

$$M \circledast (OBJ_i)^{-1} \approx S(x_i, y_i) \quad (6)$$

The high-dimensional memory vector M can be visually represented by computing the dot product similarity of it with




Figure 2: Example wave vectors (top) and the resulting interference patterns (bottom). Combining multiple sets of such wave vectors at different orientations produces a bump with surrounding rings (center). Adding more rotations and wave vectors of different magnitudes reduces the magnitude of the surrounding rings and localizes the representation (right).

SSPs representing evenly tiled points in x, y space. The similarity can be plotted as a heat map over that space. Fig. 1 provides an example of this visualization. It shows the similarity plot of $M \circledast (OBJ_i)^{-1}$ for four different objects stored in the same 361 dimensional memory vector. Several examples of encoding, decoding, and manipulating such vectors in a spiking neural network can be found in Komer et al. (2019).

In this paper, we extend the earlier results of Komer et al. (2019) to show that SSPs can be modified to naturally generate representations at the neuron level that resemble those found in the brain (both grid cells and place cells). Furthermore, we show that these brain-like SSPs are significantly more accurate for encoding cognitive maps than the randomly generated SSPs used in that previous work.

Grid cells

Sorscher, Mel, Ganguli, and Ocko (2019) examined several different neural network architectures trained on path integration tasks. Interestingly, hidden layers with grid cell-like activity arose. The optimization problem common to all of the models was reconstructing a matrix of place cell output using a lower dimensional matrix of output from a hidden layer.

Let $P \in \mathbb{R}^{n_x \times n_p}$ (where n_p is the number of place cells and n_x is the number of training points in space) be a given matrix of place cell responses sampled across space. We will use Gaussian bumps as the ideal place cell responses.

$$P_{i,n} = P_n(\mathbf{x}_i) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2\sigma^2}(\mathbf{x}_i - \mu_n)^T(\mathbf{x}_i - \mu_n)} \quad (7)$$

The task is to find the hidden layer activations $G \in \mathbb{R}^{n_x \times n_g}$ (where n_g is the number of hidden neurons and $n_g < n_p$) and the matrix of read-out weights $W \in \mathbb{R}^{n_g \times n_p}$ that minimize the

reconstruction error of the place cell responses.

$$\min_{G,W} \|P - \hat{P}\|_F^2, \quad (8)$$

$$\hat{P} = GW \quad (9)$$

The optimal W for a fixed G is given by

$$W^* = (G^T G)^{-1} G^T P. \quad (10)$$

This W should be thought of as the connection weights between the final two layers of some deep neural network. The input to the full network would be low level sensory information and the output would be the place cell activity, P . The hidden layer with activations G is the last layer before the place cells, and, since $n_g < n_p$, it creates an information bottleneck. We are interested in finding the optimal G - a compressed representation of spatial position that is optimal for reconstructing P in a single layer.

This is a low-rank approximation problem, i.e., the problem of fitting a data matrix using an approximating matrix that has a reduced rank. By the matrix approximation lemma, the columns of the optimal G will span the top n_g eigenvectors of PP^T - the correlation matrix of place cell responses. As stated in Sorscher et al. (2019), if the number of place cells is large and their receptive fields uniformly cover space (and space has periodic boundary conditions) then PP^T will approximately be a circulant matrix and its eigenvectors will be Fourier modes.

Thus, the optimal responses of hidden neurons will be linear combinations of plane waves. This will produce hidden neurons with grid-like spatial responses. Adding a non-negativity constraint to this optimization problem will result in the activity of an individual hidden neuron being proportional to a sum of three plane waves whose wave vectors are 120° degrees apart. Specifically, a column of G will have entries,

$$\sum_{j=1}^3 e^{i\mathbf{k}_j \cdot \mathbf{x}_n} + e^{-i\mathbf{k}_j \cdot \mathbf{x}_n} \quad (11)$$

$$\text{where } |\mathbf{k}_j| = |\mathbf{k}_i| \quad \forall i, j \quad (12)$$

$$\sum_{j=1}^3 \mathbf{k}_j = 0 \quad (13)$$

The interference pattern of these waves will have a hexagonal grid pattern, like grid cells. The magnitude of the wave vectors, $|\mathbf{k}_j|$, will determine the resolution of the grid: the distance between nearest peaks in the grid pattern will be $\frac{4\pi}{\sqrt{3}|\mathbf{k}_j|}$. Fig. 2 shows this interference pattern, along with the patterns resulting from adding multiple sets of such plane waves together.¹ A linear combination of many sets of these waves with correct weighting provides a finite approximation to a

¹While the bump in Fig. 2 appears isolated, it will repeat over a large enough scale. The size of space we want to represent place cells over and the width of their place fields will determine the optimal resolution scales of the grid cells (i.e., $|k_j|$).

Gaussian bump in a plane wave basis. This is how P is reconstructed. In short, this provides a reason to think that hexagonal grid cells are optimal for representing place cells.

Methods

Forming grid cells with SSPs

The Neural Engineering Framework (NEF) provides a set of principles for performing computations with spiking neural networks. The first principle of the NEF is representation. The time varying SSP representing a point travelling through space, $S(x(t), y(t)) \in \mathbb{R}^d$, can be represented by a population of n_g neurons. Each neuron encodes the SSP via the formula,

$$a_i(t) = \mathcal{G}_i[\alpha_i \mathbf{e}_i \cdot S(x(t), y(t)) + \beta_i], \quad (14)$$

where $a_i(t)$ is the activity of neuron i , $\mathbf{e}_i \in \mathbb{R}^d$ is the encoder of the neuron, $\alpha_i > 0$ is its gain, β_i is its bias, and $\mathcal{G}_i$ is a nonlinear, non-negative activation function. Depending on the activation function used, $a_i(t)$ will either be a spike train or firing rate. For simplicity, a rate approximation of the leaky integrate-and-fire model will be used here unless stated otherwise. This is given by

$$\mathcal{G}_i[J] = \begin{cases} \frac{1}{\tau_{ref} - \tau_{RC} \ln\left(1 - \frac{J}{J_i^{thres}}\right)}, & \text{if } J > J_i^{thres} \\ 0, & \text{else} \end{cases} \quad (15)$$

Neuron model parameters are chosen so that the maximum firing rate of the neurons is uniformly distributed from 20Hz to 40Hz.

The signal $S(x(t), y(t))$ can be decoded from the neural activity by

$$\hat{S}(x(t), y(t)) = \sum_{i=1}^{n_g} a_i(t) \mathbf{d}_i, \quad (16)$$

where $\mathbf{d}_i \in \mathbb{R}^d$ are the decoders of the population. The optimal set of decoders - i.e., the set that minimizes the error between the true signal and the signal reconstructed from neural activity - can be found exactly. Additionally, optimal decoders for computing some function of the input signal from the activity can also be computed. The optimal decoders for computing $f(S(x(t), y(t))) \in \mathbb{R}^{n_p}$ are

$$D = (G^T G)^{-1} G^T f(S(\mathbf{x}, \mathbf{y})), \quad (17)$$

here $\mathbf{x}, \mathbf{y}$ are n_x dimensional vectors of points sampled in space, $G \in \mathbb{R}^{n_x \times n_g}$ is the matrix of activities of the n_g neurons given input $S(x, y)$ at the sample points, $D \in \mathbb{R}^{n_g \times n_p}$ is the matrix of decoders, and $f(S(\mathbf{x}, \mathbf{y})) \in \mathbb{R}^{n_x \times n_p}$ contains samples of the function we wish to compute. This is the same as Eq. 10 for computing W^* with $D = W$ and $f(S(\mathbf{x}, \mathbf{y})) = P$. These decoders can be used to compute the weights connecting the population of n_g neurons to a population of n_p neurons so that they have the desired place cell activity.

To optimally produce the place cell responses, the activity matrix G must have columns proportional to linear combinations of plane waves as in Eq. 11. Consider the activation of




Figure 3: Example firing rates of neurons used to represent a 361 dimensional SSP with random unitary bases vectors.




Figure 4: Example firing rates of neurons from a population of grid cells used to represent the SSP with bases vectors X_{total} and Y_{total} .




Figure 5: Examples of the spiking activity of neurons from a grid cell population. The line in grey is the trajectory of points that the neuron population represents as a SSP over time. Each circle shows an example neuron from the population. The dots show where on the trajectory the neuron fired.



Figure 6: Examples of the spiking activity of neurons from a place cell population whose input is comes from a grid cell population. The line in grey is the trajectory of points that the grid cell population represents as a SSP over time. Each circle shows an example neuron from the population. The dots show where on the trajectory the neuron fired.

a single neuron with a SSP of a fixed location (a preferred firing location of the neuron) as its encoder. The $\mathbf{e}_n \cdot S(x(t), y(t))$ part of the activity of neuron n will be

$$\mathbf{e}_n \cdot S(x(t), y(t)) = (X^x \otimes Y^y) \cdot (X^{x_n} \otimes Y^{y_n}) \quad (18)$$

$$\propto (\mathcal{F}\{X\}^x \odot \mathcal{F}\{Y\}^y) \cdot (\mathcal{F}\{X\}^{x_n} \odot \mathcal{F}\{Y\}^{y_n}) \quad (19)$$

$$= \frac{1}{d} \sum_{j=0}^{d-1} \mathcal{F}\{X\}_j^{x+x_n} \mathcal{F}\{Y\}_j^{y+y_n} \quad (20)$$

$$= \frac{1}{d} \sum_{j=0}^{d-1} r_{x,j}^x r_{y,j}^y e^{i(\theta_{x,j}(x+x_n) + \theta_{y,j}(y+y_n))} \quad (21)$$

where $\mathcal{F}\{X\}_j = r_{x,j} e^{i\theta_{x,j}}$ and $\mathcal{F}\{Y\}_j = r_{y,j} e^{i\theta_{y,j}}$. We want to set these $r_{x,j}, \theta_{x,j}, r_{y,j}, \theta_{y,j}$ parameters so that this equation is equal to Eq. 11 so let

$$\bar{X}_j = \begin{cases} e^{iu_j}, & \text{for } j = 0, 1, 2 \\ e^{-iu_{d-j}}, & \text{for } j = d-2, d-1, d \\ 1, & \text{else} \end{cases} \quad (22)$$

$$\bar{Y}_j = \begin{cases} e^{iv_j}, & \text{for } j = 0, 1, 2 \\ e^{-iv_{d-j}}, & \text{for } j = d-2, d-1, d \\ 1, & \text{else} \end{cases} \quad (23)$$

$$\text{where } \sum_{j=0}^2 u_j = \sum_{j=0}^2 v_j = 0 \quad (24)$$

$$\sqrt{u_j^2 + v_j^2} = \sqrt{u_i^2 + v_i^2}, \quad i, j \in \{0, 1, 2\} \quad (25)$$

and take $X = \mathcal{F}^{-1}(\bar{X})$ and $Y = \mathcal{F}^{-1}(\bar{Y})$. This gives the desired bases vectors. The choice of (u_j, v_j) for $j \in \{0, 1, 2\}$ that satisfy the above equations will determine the orientation and resolution of the hexagonal grid activity patterns exhibited by the neurons representing $S(x(t), y(t))$.² The dimension of these bases vectors must be $d \geq 7$ to capture the 3 plane waves and a zero-frequency term. Since $r_{x,j} = r_{y,j} = 1$, the vectors X and Y will be unitary and, thus, binding will not affect the magnitude of the SSP. The placement of the complex conjugate pairs will ensure that X and Y are real.

Multi-scale grid cells

While Eqs. 22 and 23 show how to obtain bases vectors that will produce hexagonally gridded neural activity, we need

²With these bases vectors, even using random encoders will produce neurons with grid cell-like activity patterns but with added distortion.




Figure 7: The grid cell population representing $S(x,y) = X_{total}^x \otimes Y_{total}^y$ consists of N groups of neurons. Here two neurons from three of such groups are shown as rows in box labelled grid cells. The grids vary in scale, much like in the MEC where grid scales increase along dorsal to ventral axis (Brun et al., 2008). This population is connected to a place cell population using connection weights given in 10.

grids of different orientations and resolutions to accurately compute place cell responses. This can be done by using higher dimensional bases vectors that contain the vectors of Eqs. 22 and 23 as sub-vectors in the Fourier domain.

Consider the following basis vector in the Fourier domain

$$\bar{X}_{total} = [e^{iu_0^{(0)}}, \dots, e^{iu_2^{(N)}}, 1, e^{-iu_2^{(N)}}, \dots, e^{-iu_0^{(0)}}]^T \in \mathbb{C}^{6N+1} \quad (26)$$

This vector contains N sets of the 6 Fourier modes that give rise to hexagonal grid patterns. Each of the N sets will have wave vectors that lie on an equilateral triangle. The difference between the sets will be the orientation of the wave vectors and their magnitude.

We will define the n^{th} sub-vector as sub-vector

$$\bar{X}_n \equiv [e^{iu_0^{(n)}}, e^{iu_1^{(n)}}, e^{iu_2^{(n)}}, 1, e^{-iu_2^{(n)}}, e^{-iu_1^{(n)}}, e^{-iu_0^{(n)}}]^T. \quad (27)$$

We can construct a matrix $\bar{B}_n$ that will take the n^{th} sub vector $\bar{X}_n \in \mathbb{C}^7$ and project it to a $6N+1$ vector that contains the sub-vector in the same position as in $\bar{X}_{total}$ and zeros elsewhere. The sum of all such projections will produce the com-

plete vector.

$$\bar{X}_{total} = \sum_{n=0}^{N-1} \bar{B}_n \bar{X}_n \quad (28)$$

Converting to the time domain,

$$X_{total} = \sum_{n=0}^{N-1} W_{6N+1}^{-1} \bar{B}_n W_7 X_n. \quad (29)$$

where W_7 is the 7×7 DFT matrix and W_{6N+1}^{-1} is the $6N+1 \times 6N+1$ IDFT matrix.

Thus, a population of n_g neurons that encodes a SSP constructed with such bases vectors will have hexagonal grid activity patterns of different orientations and resolutions. The total population will represent the $6N+1$ dimensional SSP, $S(x,y) = X_{total}^x \otimes Y_{total}^y$. Neurons can be sorted into N groups with different orientations and resolutions. Neuron i from group n will have an encoder given by

$$\mathbf{e}_i = N W_{6N+1}^{-1} \bar{B}_n W_7 (X_n^{x_i} \otimes Y_n^{y_i}). \quad (30)$$

Neurons within a group will have hexagonal activity patterns of the same orientation and resolution but with different shifts and receptive field widths (from variations in threshold current and bias parameters). Populations of neurons constructed in this way will be referred to as grid cell populations. Examples of the activity patterns of select neurons within a grid cell population are shown in Fig 4. In contrast, Fig. 3 shows the activity patterns of neurons from a population that represents a SSP constructed using *random* unitary bases vectors of dimension $6N+1$. In the random SSP population, d -dimensional SSPs representing uniform locations over space are used as encoders. As can be seen, the random SSPs do not replicate a hexagonal grid structure.

To increase biological realism, populations can be simulated using the spiking version of the leaky-integrate-and-fire model. Examples of the hexagonal firing patterns of neurons in a grid cell population are given in Fig. 5.

Neurons with place cell-like firing patterns can be simulated by connecting this population of spiking grid cells to population of n_p neurons using connection weights:

$$\omega_{ij} = \alpha_j \mathbb{1}_j \cdot \mathbf{d}_i \quad (31)$$

where ω_{ij} is the connection weight from grid cell i to place cell j , $\mathbf{d}_i$ are the decoders of the grid cell population given in Eq. 17, α_j is the gain of place cell j , and $\mathbb{1}_j$ is a one-hot vector with a one in the j^{th} position – this is the encoder of the place cell. Example firing patterns of neurons in such a place cell population are given in Fig. 6.

Accurate place cell representation

Replicating grid cell patterns is useful for biological fidelity, but does not address the question of why such patterns arise in the first place. As mentioned, Sorscher et al. (2019), suggested that such patterns should be optimal under weak assumptions. Here we test that hypothesis using SSP-based representations.




Figure 8: The squared Frobenius norm between the ideal place cell output and the model place cell output versus the dimension of the SSP used in computing the approximation. The dashed line shows the mean values obtained over 10 trials using random bases vectors. The error bands around this line are the min and max values.

Using the decoders of Eq. 10, place cell output was reconstructed using the activity of a population of grid cells and, for comparison, a population of neurons representing SSPs with random unitary bases vectors of dimension d . The dimensions of the SSPs and the population used were $N = 60$ (from using 5 different sets of grid orientations and 12 sets of different resolutions), $d = 6N + 1 = 361$, $n_g = 600$, and $n_p = 3000$. The activity was recorded over $x, y \in [-10, 10]$ with $n_x = 10000$. The grid spacing varied from 9 to 3.6 across the neurons. The error in the reconstruction of place cell output was computed with the squared Frobenius norm.

$$\|P - GW^*\|_F^2 \quad (32)$$

where the ground truth P matrix of place cell firing rates is as given in Eq. 7 with field centers randomly distributed over the space.

In addition, the average distance from the true center of the place cells fields (μ_n from Eq. 7) and the center of the reconstructed place cells was computed. The results are given in Table 1. In both measures, the grid cell population outperformed the random SSP population.

We also investigated the effects of using higher dimensional SSPs. Figs 8 and 9 show the matrix norm and center distance error measures plotted versus the SSP dimension d . The grid cell population consistently performed better on both measures for all values of d greater than 37. Increasing d beyond 400 did not improve the performance with grid cells. As low a value of $d = 169$ can be used with grid cells while maintaining accurate place cell output.

Lastly, the effect of the number of place cells, n_p , on the reconstruction accuracy was explored. As the number of place cells increases, the ratio of hidden layer neurons to place cells decreases. The hidden layer becomes a greater information




Figure 9: The average distance between the ideal place cell centers and the centers of the model place cell output versus the dimension of the SSP used in computing the approximation. The dashed line shows the mean values obtained over 10 trials using random bases vectors. The error bands around this line are the min and max values.

bottleneck. Fig. 10 shows the matrix norm error measure plotted versus n_p . The error increases with n_p faster in the random SSP population than in the grid cell population. The grid cell population scales better as it provides a better basis for decoding place fields.

From these results, we now have evidence that the SSP characterization of spatial representation, which supports the representation of objects over continuous slots with fractional binding, can be implemented using bases vectors that are efficient for generating place cell activity. This suggests that SSPs can act as a unifying representation from low-level grid cell activity to cognitive representations of objects (with hundreds of features) at continuous spatial locations.

Conclusion

In summary, Spatial Semantic Pointers can be represented by the activity of a population of grid cells organized into modules of different orientation and spatial resolutions – much like the real modular organization of grid cells in the entorhinal cortex (Fig. 7). This method of representing space can be modelled by spiking neural networks and reproduces the neural activity patterns seen in the brain.

The method presented for constructing SSPs produces a basis for accurately decoding place cell output. It outperforms SSPs constructed from random unitary bases vectors at both decoding place cell responses across space and matching the correct place field centers. This methodology could be extended as a general way to accurately represent continuous variables. Recent experiments have found that grid cell-like activity occurs during tasks involving navigation over continuous non-spatial dimensions such as sound and visual features. Additionally, some hippocampal cells have been found to be preferentially receptive to particular concepts, just as

Table 1: The results of the place cell output reconstruction using the activity of neuron population of grid cells and populations representing random SSPs with $d = 361$. The results using random SSPs were averaged over 10 trials. The squared Frobenius norm is a measure of the matrix reconstruction error (see Eq. 32). The center distance is the average distance from the true place field centers and the centers of the reconstructions. The numbers in brackets are the 95% confidence intervals for the 10 random trials.

Bases Vectors	Frobenius norm	Center Distance
Grid Cells	1.621	0.089
Random SSPs	15.085 (12.844, 17.326)	0.338 (0.324, 0.353)




Figure 10: The squared Frobenius norm between the ideal place cell output and the model place cell output versus the number of place cells being approximated. The dashed line shows the mean values obtained over 10 trials using random bases vectors. The error bands around this line are the min and max values.

place cells are to specific regions in space.

Since grid cells are an optimal basis for representing places then they may also be an optimal basis for representing concepts that reside in continuous feature spaces. SSPs are well suited for such representations - they can be used to represent any continuous variable and can be bound together with semantic pointers representing discrete variables/concepts as well. Semantic pointers that reside in different cognitive spaces can be bound together in hierarchical structures, allowing for rich and complex representations all modelled in spiking neural networks.

Future work will include exploring the generalization of SSPs to concept representation and advancing the realism of spatial cognition by including path integration - integrating over velocity input to update the SSP over time and using visual feedback for correction.

Acknowledgements

The authors would like to thank Terry Stewart and Peter Dugins for discussions that helped improve this paper. This work was supported by CFI and OIT infrastructure funding as well as the Canada Research Chairs program and NSERC Discov-

ery grant 261453, and AFOSR grant FA9550-17-1-0026.

References

- Aronov, D., Nevers, R., & Tank, D. W. (2017, Mar). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature*, 543(7647).
- Bellmund, J. L., Gärdenfors, P., Moser, E. I., & Doeller, C. F. (2018). Navigating cognition: Spatial codes for human thinking. *Science*, 362(6415).
- Brun, V. H., Solstad, T., Kjelstrup, K. B., Fyhn, M., Witter, M. P., Moser, E. I., & Moser, M.-B. (2008). Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, 18(12), 1200–1212.
- Choo, F.-X., & Eliasmith, C. (2010). A spiking neuron model of serial-order recall. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 32).
- Constantinescu, A. O., O'Reilly, J. X., & Behrens, T. E. J. (2016, Jun). Organizing conceptual knowledge in humans with a gridlike code. *Science*, 352(6292).
- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering*. Massachusetts Institute of Technology.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052).
- Komer, B., Stewart, T. C., Voelker, A. R., & Eliasmith, C. (2019). A neural representation of continuous space using fractional binding. In *41st annual meeting of the cognitive science society*. Montreal, QC: Cognitive Science Society.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.
- Plate, T. A. (1994). *Distributed representations and nested compositional structure*. Citeseer.
- Sorscher, B., Mel, G., Ganguli, S., & Ocko, S. (2019). A unified theory for the origin of grid cells through the lens of pattern formation. In *Advances in neural information processing systems*.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, 55(4).