



Does the entorhinal cortex use the Fourier transform?

Jeff Orchard^{1,2*}, Hao Yang² and Xiang Ji^{1,2}

¹ Centre for Theoretical Neuroscience, University of Waterloo, Waterloo, ON, Canada

² David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada

Edited by:

Si Wu, Beijing Normal University, China

Reviewed by:

Abdelmalik Moujahid, University of the Basque Country UPV/EHU, Spain

Kechen Zhang, Johns Hopkins University, USA

*Correspondence:

Jeff Orchard, David R. Cheriton School of Computer Science, University of Waterloo, 200 University Ave. West, Waterloo, ON N2L 3G1, Canada
e-mail: jorchard@uwaterloo.ca

Some neurons in the entorhinal cortex (EC) fire bursts when the animal occupies locations organized in a hexagonal grid pattern in their spatial environment. Place cells have also been observed, firing bursts only when the animal occupies a particular region of the environment. Both of these types of cells exhibit theta-cycle modulation, firing bursts in the 4–12 Hz range. Grid cells fire bursts of action potentials that precess with respect to the theta cycle, a phenomenon dubbed “theta precession.” Various models have been proposed to explain these phenomena, and how they relate to navigation. Among the most promising are the oscillator interference models. The bank-of-oscillators model proposed by Welday et al. (2011) exhibits all these features. However, their simulations are based on theoretical oscillators, and not implemented entirely with spiking neurons. We extend their work in a number of ways. First, we place the oscillators in a frequency domain and reformulate the model in terms of Fourier theory. Second, this perspective suggests a division of labor for implementing spatial maps: position vs. map layout. The animal’s position is encoded in the phases of the oscillators, while the spatial map shape is encoded implicitly in the weights of the connections between the oscillators and the read-out nodes. Third, it reveals that the oscillator phases all need to conform to a linear relationship across the frequency domain. Fourth, we implement a partial model of the EC using spiking leaky integrate-and-fire (LIF) neurons. Fifth, we devise new coupling mechanisms, enlightened by the global phase constraint, and show they are capable of keeping spiking neural oscillators in consistent formation. Our model demonstrates place cells, grid cells, and phase precession. The Fourier model also gives direction for future investigations, such as integrating sensory feedback to combat drift, or explaining why grid cells exist at all.

Keywords: entorhinal cortex, fourier transform, path integration, oscillators, neural engineering framework, phase precession, grid cells, place cells

1. INTRODUCTION

Some neurons in the entorhinal cortex (EC) exhibit spatial firing patterns (Hafting et al., 2005). These neurons, called “grid cells,” spike preferentially when the animal is at points arranged in a hexagonal grid pattern. Before that, neurons in the hippocampus were found to activate when the animal was in a particular location in the environment (O’Keefe and Dostrovsky, 1971; Muller et al., 1987). These neurons are called “place cells.”

Both types of cells, place cells and grid cells, are modulated by the theta rhythm, a pattern of activity that oscillates at between 4 and 12 Hz. Moreover, the frequency at which grid cells oscillate is influenced by the animal’s movement. As the animal moves, the frequencies increase slightly. The amount of increase depends on what direction the animal is moving, and how fast (Sargolini et al., 2006). If the animal moves in a cell’s preferred direction, the frequency increases more, whereas in the opposite direction, the frequency increases only marginally. The term “velocity-controlled oscillator,” or VCO, denotes a neuron or population of neurons whose activity oscillates, but at a frequency that is modulated by velocity (Welday et al., 2011).

The result of increased frequencies causes the timing of spike bursts to precess in phase relative to the baseline theta cycle (O’Keefe and Recce, 1993; Tsodyks et al., 1996; Geisler et al., 2007). This phenomenon, first described by O’Keefe and Recce (1993), is called “phase precession.”

Combining the ideas of VCOs and grid cells, researchers proposed that the grid patterns might arise from an interference pattern between VCOs (Burgess et al., 2007; Blair et al., 2008). As the animal moves, these VCOs take on slightly different frequencies, and hence their relative phases change.

Phase is the time integral of frequency, just like position is the time integral of velocity. Hence, if a VCO’s frequency is proportional to the animal’s velocity, then the VCO’s phase is proportional to the animal’s position. This is the basis for oscillator interference models.

By combining (adding) two VCOs with different frequencies, the result is a beat interference pattern that generates periods of constructive and destructive interference as their phase difference evolves (Blair et al., 2008). Since phase and position are linked, this interference pattern overlays the animal’s spatial environment. Combining three VCOs (that differ in preferred direction

by multiples of 60°) tends to create a hexagonal grid interference pattern (Burgess et al., 2007; Krupic et al., 2012).

Despite this progress in understanding grid cells, a satisfactory explanation of the relationship between grid cells and place cells remained unclear. As recently as 2008, researchers had only a handful of ideas of how grid cells might combine to produce place cells (Moser et al., 2008). Some have experimented with combining a random selection of grid cells to produce place-cell like behavior (Fuhs and Touretzky, 2006; Solstad et al., 2006). Others suggested that a sum of grid cells could create place cells, but offered only vague justification (O'Keefe and Burgess, 2005; McNaughton et al., 2006). A more detailed proposal argued that place cells resulted from the Moiré interference patterns between small-scale grid patterns (Blair et al., 2007). However, their method involves intricate resealing of so-called “theta cell” grids, which the authors point out as “potentially a serious limitation” (Blair et al., 2007).

A recent study concluded that distributed encoding using grid cells formed a more efficient representation of spatial location than the same number of place cells (Mathis and Herz, 2012). That work is interesting, but does not discuss the mechanisms underpinning these various cell types. A comprehensive review of the various proposed models can be found in Zilli (2012).

A spiking-neuron based model of path integration used Gaussian surfaces to represent place cells, but encoded these Gaussians by their Fourier coefficients (Conklin and Eliasmith, 2005). This implementation takes advantage of the Fourier Shift Theorem (discussed later), moving the Gaussian pattern of excitation around by applying phase shifts to the Fourier coefficients. However, their model does not address grid cells. Can the Fourier Shift Theorem be used in conjunction with grid cells?

In 2011, Wolday et al. showed that “Theta cell burst frequencies varied as the cosine of the rat’s movement direction.” In other words, the frequency of each oscillator includes a component proportional to the cosine between the preferred direction vector, and the animal’s velocity vector (Wolday et al., 2011). They formed a bank of VCOs arranged into a 2-dimensional (2-D) array, where one dimension spans a variety of preferred directions, and the other dimension represents the degree to which frequency is increased by movement. **Figure 1A** is a recreation of a portion of Figure 7 from their paper.

In their firing-rate model, each VCO is represented as a ring oscillator with a wave of activity that cycles at theta frequency. Hence, each neuron in a ring activates at a particular phase. A read-out node connected to a number of neurons in these ring oscillators will act as a coincidence detector, and fire only when the incoming spikes are sufficiently synchronized. In this sense, a read-out node detects the phase convergence of its inputs.

According to their paper, a read-out node connected to all the phase-matched neurons from rings in the same row could behave like a place cell, and fire only when all the input neurons are in phase with each other. Similarly, choosing only three phase-matched neurons from a row, but with preferred directions separated by 120° , yields a grid cell. Finally, choosing all the phase-matched neurons from rings with the same preferred direction vector can generate a “border cell.”

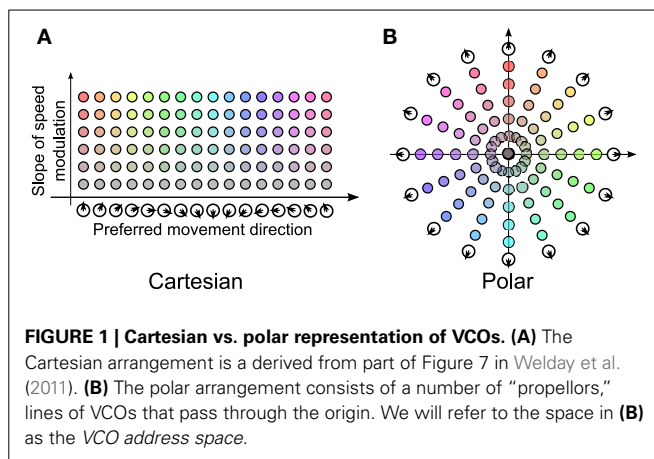


FIGURE 1 | Cartesian vs. polar representation of VCOs. (A) The Cartesian arrangement is derived from part of Figure 7 in Wolday et al. (2011). **(B)** The polar arrangement consists of a number of “propellers,” lines of VCOs that pass through the origin. We will refer to the space in **(B)** as the VCO address space.

The authors point out that their model can produce not only grid fields and place fields, but “a sufficiently large number of VCOs should be able to approximate the spatial frequency spectrum of almost any desired spatial function.” However, their justification for this assertion is mathematically unconventional.

Furthermore, they point out that such a system of oscillators is susceptible to drift; before long, these ring oscillators will start to drift out of phase and thereafter fail to produce the desired phase convergence. A coupling mechanism is needed to keep the bank of oscillators in a coherent phase state. Progress has been made to incorporate phase-resetting mechanisms into such models (Monaco et al., 2011). However, the mechanism does not improve the stability of path integration, but rather resets the phases to a previously known state as the animal approaches a landmark.

While the authors hint at the use of Fourier theory, they do not take full advantage of the theory and its interpretations. For example, their arrangement of the VCOs into a rectangular array obscures phase patterns that are revealed by a more conventional frequency-domain layout. They also miss the intuitive link between this frequency domain and hexagonal grid orientation and frequency, and the link between “border cells” and the Fourier Projection-Slice Theorem (Natterer and Wübbeling, 2001).

In this paper, we extend their work in a number of critical ways:

- We express the oscillator states and their interrelations in terms of Fourier theory by arranging the VCOs in a frequency domain.
- Our formulation exposes a linear constraint on the phase relationships between the VCOs.
- We suggest the division of two different components involved in forming activation maps: position is encoded in VCO phases, and map layout is encoded in connection weights.
- We propose and implement novel coupling mechanisms that stabilizes VCO phase relationships.
- We implement a partial model using spiking leaky integrate-and-fire neurons.

2. MATERIALS AND METHODS

2.1. FOURIER MODEL

In this section, we describe how Fourier theory leads to a deeper understanding of the interactions between oscillators. For a primer on the relevant theory, see Appendix A.

The bank-of-oscillators model (Welday et al., 2011) states that a VCO's frequency depends on two parameters: the speed of the animal, and the cosine of the animal's velocity vector with the VCO's preferred direction. In Welday et al. (2011), these VCOs were organized into a 2-D space in which one axis enumerates preferred directions, and the other axis represents the influence of speed (labeled "slope of speed modulation" in Figure 1A).

Another, perhaps more intuitive way of presenting the same 2-D parameter space is to use polar coordinates, as shown in Figure 1B. In this view, each VCO's preferred direction is indicated by its polar angle, while its speed modulation is indicated by its distance from the origin. We will refer to this polar arrangement as the VCO address space, and refer to a VCO's address in this space as \mathbf{d} , using either Cartesian or polar coordinates.

Instead of thinking of a VCO as a ring of neurons, we will think of a VCO as a population of neurons that represents a 2-D unit vector, which we will call a *phase vector* (For more on how a population of neurons can represent a vector, see Appendix B.). Each VCO's phase vector simply rotates around the unit circle at the VCO's specified frequency.

Consider the VCO with address \mathbf{d}_A in Figure 2A, four units from the origin, in the direction of 30° . If the animal's velocity, \mathbf{v} , is in that direction, the VCO will exhibit a higher frequency than the VCO at the origin. More precisely, the VCO's frequency will increase by an amount proportional to $\mathbf{d}_A \cdot \mathbf{v}$, the dot-product of

\mathbf{d}_A and \mathbf{v} . After a time t , the difference in phase between the VCOs at \mathbf{d}_A and the origin will be

$$\phi(t) = \int_0^t \mathbf{d}_A \cdot \mathbf{v}(\tau) d\tau = \mathbf{d}_A \cdot \int_0^t \mathbf{v}(\tau) d\tau = \mathbf{d}_A \cdot \mathbf{x}(t).$$

That is, the phase difference depends linearly on the animal's position, $\mathbf{x}(t)$. Even though ϕ is an angle, we will also represent it as a vector on the unit circle. We can write this phase vector as a complex number using Euler's formula,

$$(\cos \phi, \sin \phi) \equiv \cos \phi + i \sin \phi = \exp(i\phi),$$

where $i = \sqrt{-1}$. Thus, we can represent our phase difference as

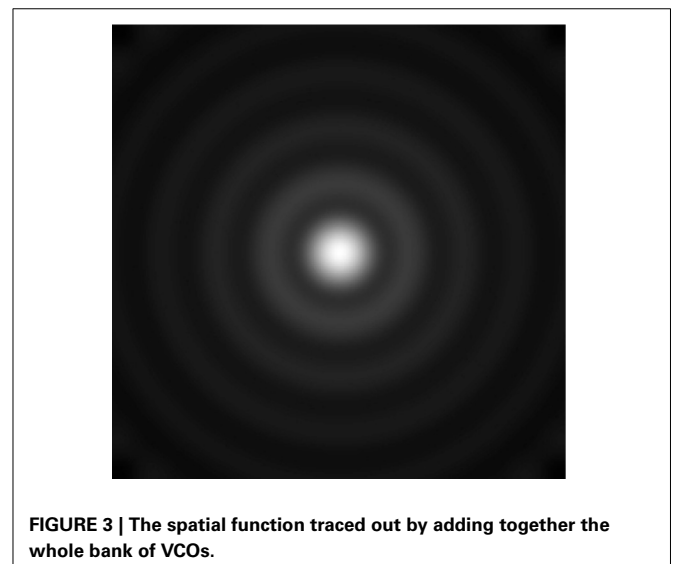
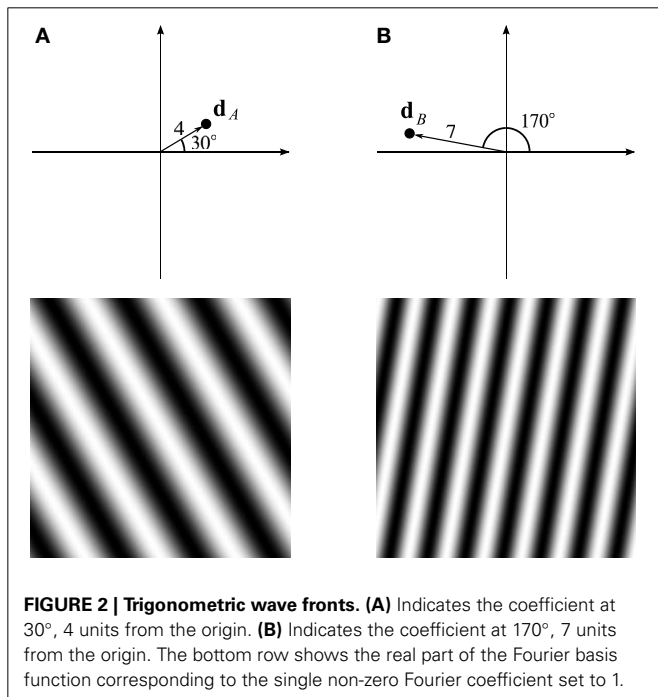
$$(\cos \phi, \sin \phi) = \exp(i\mathbf{d}_A \cdot \mathbf{x}).$$

Thus, the components of the phase vector trace out sine and cosine wave fronts that are fixed in the animal's environment. The bottom row of Figure 2A shows the wave front corresponding to the VCO located at \mathbf{d}_A . A different VCO at location \mathbf{d}_B traces out a different wave front, as shown in Figure 2B.

Considering that the animal has many such VCOs, what happens if we combine them all into a sum, as if a single read-out node was adding together all the phase vectors from all the VCOs? The value of the read-out node would be

$$p(\mathbf{x}) = \sum_{\rho, \theta} \exp(i\mathbf{d}_{\rho, \theta} \cdot \mathbf{x}), \quad (1)$$

where $\mathbf{d}_{\rho, \theta}$ is the address of a VCO, and the subscripts ρ and θ index distance from the origin and orientation, respectively. An image created using this simple method is shown in Figure 3. The activity of this read-out node corresponds to a spatial map akin to a place cell. Why is that? The answer has to do with the fact that Equation (1) almost looks like an inverse Discrete Fourier



Transform¹ (Oppenheim and Schaffer, 1999). In the following sections, we outline the benefits of thinking about the EC in terms of the Fourier transform, and use it to extend the bank-of-oscillators model.

2.2. ENTORRHINAL CORTEX

Here we outline our model of how the Fourier transform relates to the navigational function of the EC.

We can think of the VCO address space as a 2-D frequency spectrum (frequency domain). In the frequency domain, each location corresponds to a Fourier basis function (wave front), with orientation and frequency dictated by its location, as illustrated in **Figure 2**. Analogously, each VCO's phase vector traces out a Fourier basis function as the animal moves around, with orientation and frequency dictated by the VCO's address in this space.

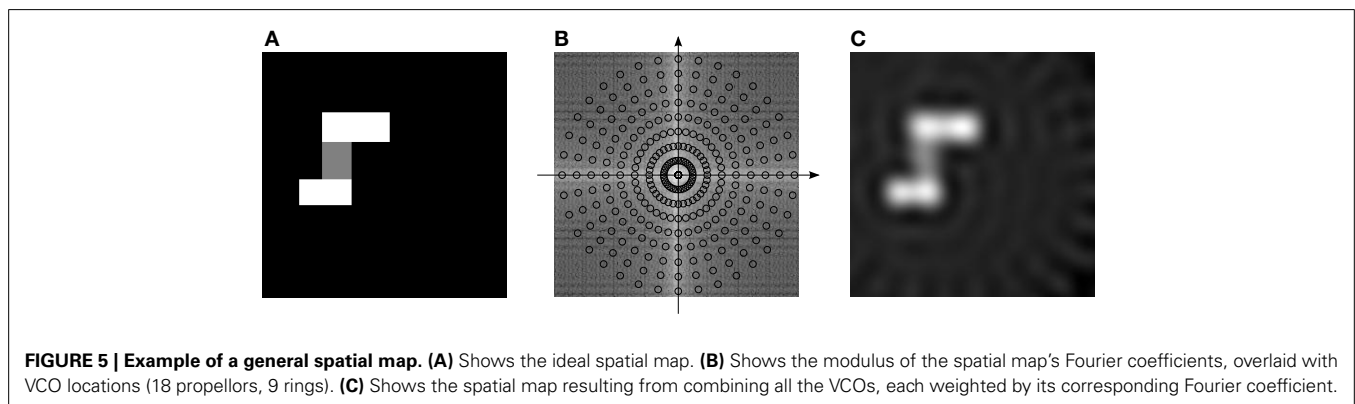
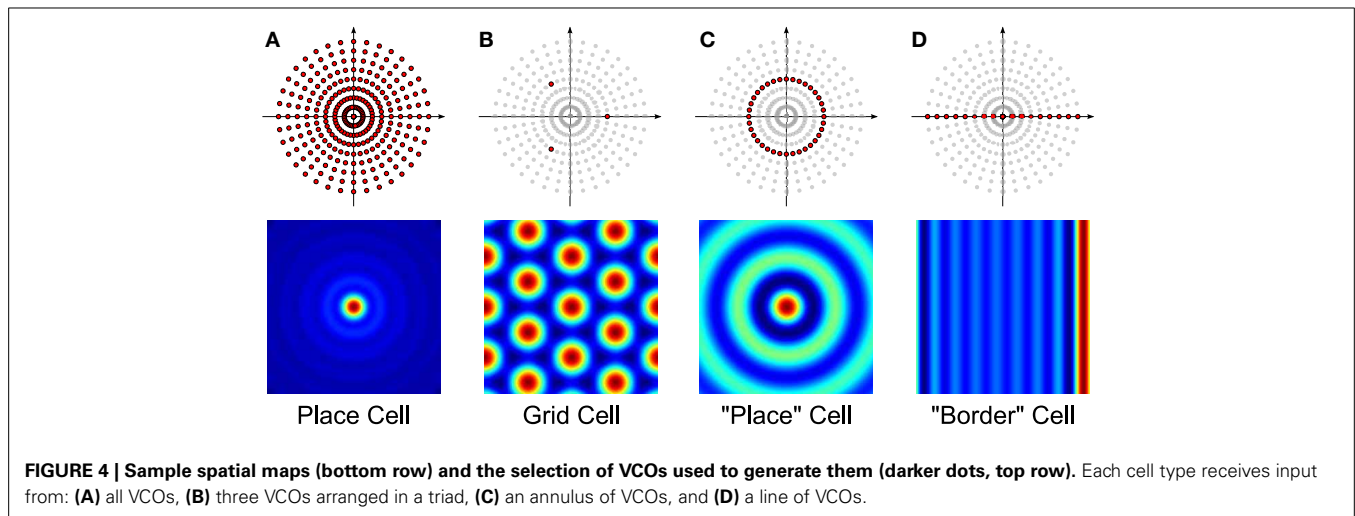
A read-out node that adds together the phase vectors of two VCOs receives the sum of those two Fourier basis functions. (Recall that the read-out node is a population of neurons that can encode the resulting vector). Hence, adding these phase vectors together is like adding Fourier basis functions.

Different firing-rate spatial maps can be constructed by picking and choosing which VCOs to combine. For example, the bottom row of **Figure 4** shows a number of spatial maps similar to those elicited by place cells, grid cells, and "border" cells. The VCOs used to create the maps are colored darker in the top row. For the sake of comparison, we included **Figure 4C** to show the "place" cell that was proposed in Welday et al. (2011), using only the VCOs from a given spatial frequency (but all orientations).

There is more flexibility though. A read-out node can also combine the VCO phase vectors using different weights, and even apply a phase rotation to individual VCOs. Remember, each VCO and read-out node is a population of neurons that encodes a 2-D phase vector. Using the Neural Engineering Framework (NEF, see Appendix B.), we can scale and rotate a VCO's phase vector by choosing the appropriate connection strengths. Scaling and phase rotation is exactly what a Fourier coefficient does. Hence, combining the VCO phase vectors in this way performs an inverse Fourier transform, where the connection weights are determined by the desired Fourier coefficients.

Consider the following example, illustrated in **Figure 5**. Part A in the figure shows an idealized spatial activation map, where brightness indicates the desired activity. We wish to use a combination of VCOs (wave fronts) to duplicate—as accurately as possible—that activation map. **Figure 5B** shows the Fourier transform of the ideal map in A, overlaid with sampling locations

¹It would be the IDFT, except that the summation indices correspond to polar coordinates, whereas the IDFT indices use rectangular spacing.



that correspond to the VCO addresses. By sampling the Fourier transform at those locations, we get the Fourier coefficients that we should use as weights for the VCO phase vectors. That is, each VCO's phase vector gets transformed by the Fourier coefficient as it is transmitted to the read-out node. Part C of the figure shows the theoretical activity of the read-out node, attained by combining the VCO phase vectors using these Fourier coefficients. Note that a cone filter was also used to compensate for the polar sampling, as is common practice in computed tomography reconstruction (Natterer and Wübbeling, 2001).

Why is this Fourier interpretation helpful? The answer comes from how it decomposes the construction of spatial maps into two parts. The animal's position in its environment is encoded by the phases of the VCOs, while the shape of the spatial map is encoded by the connection weights.

We have already discussed how the connection weights can produce almost any spatial map. Now we will look at how the animal's location is represented by the collective phases of the VCOs.

The Fourier Shift Theorem (described in Appendix A.) states that any function f can be shifted by multiplying each of its Fourier coefficients, F_k , by a phase ramp, as shown in Equation (A2) in Appendix A. In 2-D, we have two spatial dimensions, (x, y) , and two frequency dimensions, (k, ℓ) . The Fourier Shift Theorem states that the Fourier coefficients of our function, shifted by (x, y) , can be written,

$$G_{(k, \ell)} = \exp\left(-2\pi i \frac{(k, \ell) \cdot (x, y)}{N}\right) F_{(k, \ell)}.$$

Or, more concisely,

$$G_{\mathbf{d}} = \exp\left(-2\pi i \frac{\mathbf{d} \cdot \mathbf{x}}{N}\right) F_{\mathbf{d}}, \quad (2)$$

where \mathbf{d} is the address of the VCO, and \mathbf{x} is the shift. We shift our function simply by multiplying its Fourier coefficients by a phase ramp. It is called a phase ramp because the expression inside the exponential function is a ramp (or linear function) in 2-D. The slope of the ramp is controlled by \mathbf{x} . If \mathbf{x} is the animal's position, then the slope of the ramp changes as the animal moves, as illustrated in Figure 6.

This view of the EC makes clear a strict set of conditions on the VCO phases. They all must change in a way consistent with a linear phase ramp. We can take advantage of this condition. As

will be seen, spiking neurons do not make perfect oscillators. In the next section, we will use this constraint to generate coupling mechanisms that stabilize the process.

The results shown thus far are from an idealized implementation of the EC and its bank of VCOs. However, we also implemented a partial version of the EC Fourier model using spiking leaky integrate-and-fire (LIF) neurons (Koch, 1999). Here we describe our implementation of the model, outline the challenges, and display results from simulation experiments.

2.3. DYNAMICS OF COUPLED OSCILLATORS

2.3.1. Neural simple harmonic oscillator

The simple harmonic oscillator is governed by the system of differential equations

$$\begin{aligned} \frac{dx}{dt} &= cy \\ \frac{dy}{dt} &= -cx \end{aligned}$$

where c is a scalar. Solutions to this dynamical system include all circular orbits around the origin in the (x, y) -plane. The frequency of oscillation is proportional to c . To implement this behavior in neurons, we compute the decoders D that decode

$$f(x, y) = [x + \tau_s cy, y - \tau_s cx]$$

(see Appendix B.). The decoded state is immediately fed back into the population, leading to the trajectory shown in Figure 7.

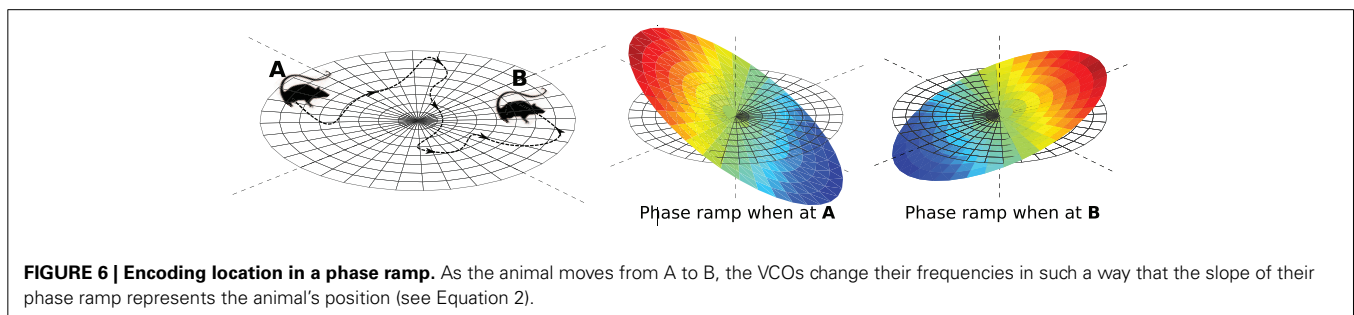
A VCO can be constructed by making the population encode a 3-D vector of the form (x, y, θ) , where the x and y components oscillate at frequency (in radians per second) specified by θ . In this case, the decoder would be designed to decode

$$f(x, y, \theta) = [x + \tau_s \theta y, y - \tau_s \theta x].$$

We used a stabilized version of the simple harmonic oscillator by incorporating a unit-vector constraint into the decoder,

$$f(x, y, \theta) = \frac{[x + \tau_s \theta y, y - \tau_s \theta x]}{\|[x + \tau_s \theta y, y - \tau_s \theta x]\|_2}.$$

In our model, the VCOs were modeled using populations of 300 LIF neurons. The degree to which the animal's velocity vector influences the frequency of the oscillators depends on where the



VCO sits in the plane. Similar to the frequency law stated in Equation (27) of Welday et al. (2011), we derive the frequency of the VCO at location \mathbf{d}_n from the animal's velocity vector, \mathbf{v} , using

$$\theta_n = 8 + 1.6 \|\mathbf{v}\|_2 + 1.273 \mathbf{d}_n \cdot \mathbf{v} \quad (3)$$

where the vectors \mathbf{v} and \mathbf{d}_n are assumed to be in the unit circle. The coefficient of 1.6 comes from using Equation (27) from Welday et al. (2011) with a maximum speed of 25 cm/s, while the coefficient of 1.273 comes from $4/\pi$, a factor that scales from radians to radius and increases the influence of that term by a factor of 4. **Figure 8** shows how the frequencies vary with the VCO's distance from the origin, and that the frequencies are always above the baseline theta-rhythm of 8 Hz no matter which direction the rat is moving.

Our model is composed of 1-D arrays of equally-spaced VCO nodes, each one centered on the origin of the address space, and each one containing 17 VCOs. We will refer to each linear array as a "propellor." The partial model includes three such propellers.

2.3.2. Phase coupling

The stochastic nature of spiking neurons causes imperfect behavior of the oscillators. If set to the same frequency and started in phase, perfect oscillators will remain in phase. In reality, however, slight errors in frequencies will accumulate and cause the oscillators to drift out of phase. Then, the VCO phases will no longer accurately represent the location of the animal.

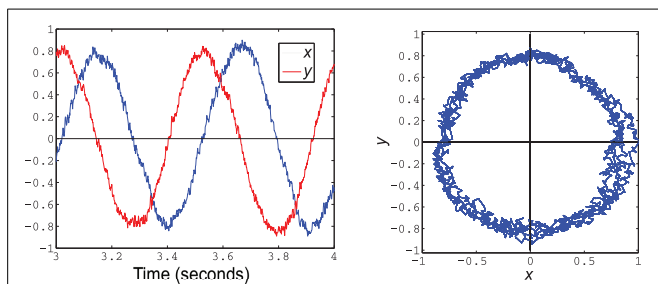


FIGURE 7 | Decoded state of a simple harmonic oscillator implemented using 200 spiking LIF neurons. The plot on the left shows the x and y components over 1 s of time, while the graph on the right plots the phase portrait of x vs. y over the same time frame.

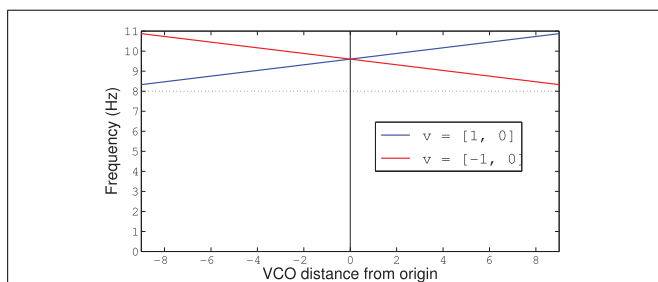


FIGURE 8 | Frequency modulation for VCOs with preferred direction $\mathbf{d}_n = [1, 0]$ and \mathbf{v} either $[1, 0]$ or $[-1, 0]$. The dotted line shows the baseline theta-frequency of 8 Hz. Notice that all VCOs have a frequency above the baseline 8 Hz, no matter which way the rat is running.

Phase drift in VCOs can be corrected, as demonstrated in Monaco et al. (2011). In that work, the animal records a snapshot of the VCO phase offsets for a number of spatial landmarks. Later, as the animal passes near one of the landmarks, the VCOs are smoothly brought into their corresponding (recalled) phases using error feedback. This approach addresses the incorporation of sensory feedback to correct absolute phase errors, but does not intrinsically stabilize the path integration process.

In this paper, we address the internal consistency of the VCOs, rather than absolute drift. We use phase coupling to maintain an internal state consistent with a phase ramp. Such coupling can be implemented in a number of ways. For brevity, we outline our solution in detail for part of the network (coupling within a single propellor), and include less detail for the other forms of coupling (between propellers).

The absolute phase of the oscillators does not matter. What matters is the phase differences between VCOs. In particular, the phases should follow a linear phase ramp along the propellor, as described by Equation (2). The phases can drift, as long as their relative phases maintain a linear trend across the array.

One way to stabilize the relative phases is to couple the oscillators to each other. We reasoned that the system should be free to allow any linear slope in phase, but discourage other phase differences.

Since the phase should change linearly and the VCO nodes are equally-spaced, then each pair of adjacent VCOs should have the same phase difference, or phase step. For each propellor, we introduce an array of neural populations to couple the oscillators within that propellor. We call these nodes "phase step" nodes. **Figure 9** shows the array of VCOs (the propellor), and their connections to an array of phase-step coupling nodes. Each adjacent pair of oscillators is connected to the same phase-step node. Each phase-step node contains 500 LIF neurons, and represents a 6-D vector of the form $(a, b, \alpha, \beta, c, s)$, where (a, b) and (α, β) are the phase vectors from the two connecting oscillators, and (c, s) represents a phase difference of ϕ , where $c = \cos \phi$ and $s = \sin \phi$.

Ideally, all the phase-step nodes for a given propellor should report the same phase difference. In reality, the phase-step array creates a consensus for this phase difference. Each coupling node decodes the phase difference, (\tilde{c}, \tilde{s}) , between its afferent

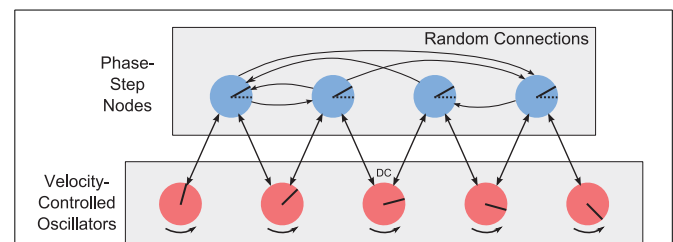


FIGURE 9 | Array of velocity-controlled oscillators (propellor). Each adjacent pair of VCOs is coupled by a phase-step node. The phase-step nodes are randomly connected to each other to arrive at a consensus for what the phase difference should be between adjacent VCOs.

VCOs using

$$(\tilde{c}, \tilde{s}) = (a, b) \cdot \overline{(\alpha, \beta)},$$

where (a, b) and (α, β) are the oscillator states, and $\overline{(\alpha, \beta)}$ is the conjugate of (α, β) . If $a + ib = \exp(i\phi_k)$ and $\alpha + i\beta = \exp(i\phi_{k+1})$, then their phase difference can be represented by another phase vector,

$$\begin{aligned} \tilde{c} + i\tilde{s} &= \exp(i\phi_k) \exp(-i\phi_{k+1}), \\ &= \exp(i(\phi_k - \phi_{k+1})). \end{aligned}$$

Note, however, that the computation is done by the network entirely in Cartesian coordinates, using

$$\begin{aligned} \tilde{c} + i\tilde{s} &= (a + ib) \cdot (\alpha - i\beta) \\ &= (a\alpha + b\beta) + i(-a\beta + b\alpha). \end{aligned}$$

Each phase-step node decodes (\tilde{c}, \tilde{s}) and projects it out to a random subset of other phase-step nodes (including itself). Thus, each phase-step node receives (\tilde{c}, \tilde{s}) from a number of other nodes, each equally weighted with all the weights adding to 1. This weighted-average consensus gets stored in the (c, s) -components of each phase-step node's state.

Recall that our phase-step nodes store vectors of the form $(a, b, \alpha, \beta, c, s)$. In a perfect world, the VCO states (a, b) and (α, β) would differ in phase by exactly (c, s) . However, there is always some error. To reduce the error, each phase-step node projects phase adjustments back to their connected VCOs. Given the consensus phase difference (c, s) , we can generate corrected estimates of (a, b) and (α, β) using

$$\begin{aligned} (\tilde{a}, \tilde{b}) &\approx (\alpha, \beta) \cdot \overline{(c, s)}, \\ (\tilde{\alpha}, \tilde{\beta}) &\approx (a, b) \cdot (c, s). \end{aligned}$$

In other words, we rotate (α, β) clockwise to get (\tilde{a}, \tilde{b}) , an approximation of (a, b) . Likewise, we rotate (a, b) counter-clockwise to get $(\tilde{\alpha}, \tilde{\beta})$, an approximation of (α, β) . Then we can compute phase correction vectors,

$$\begin{aligned} (\Delta\alpha, \Delta\beta) &= (\tilde{\alpha}, \tilde{\beta}) - (\alpha, \beta), \\ (\Delta a, \Delta b) &= (\tilde{a}, \tilde{b}) - (a, b). \end{aligned}$$

This process is illustrated in **Figure 10**.

Only half of each correction needs to be incorporated to bring the two oscillators into the correct phase relationship. In our implementation, we divide the corrections by 5 and feed them back into the (x, y) components of the corresponding VCOs. The divisor of 5 was chosen instead of 2 for stability reasons. In our experience, a wide range of divisors work equally well.

This phase-coupling method maintains a linear progression in phase across each propellor array of VCOs. However, there is nothing keeping the propellers in phase with each other. There

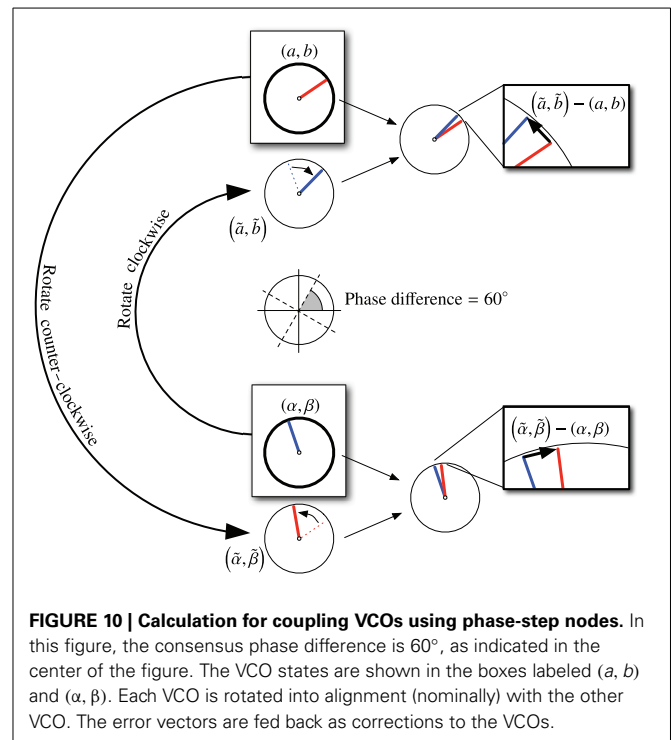


FIGURE 10 | Calculation for coupling VCOs using phase-step nodes. In this figure, the consensus phase difference is 60° , as indicated in the center of the figure. The VCO states are shown in the boxes labeled (a, b) and (α, β) . Each VCO is rotated into alignment (nominally) with the other VCO. The error vectors are fed back as corrections to the VCOs.

are two forms of phase locking required to keep all the propellers consistent.

Notice that all DC nodes should be in phase with each other since there is no direction-dependent frequency modulation on the DC nodes; only the speed affects their frequency. Hence, we need to make sure the phase of the DC components match across all (three) propellers. We achieved this by adding a single 6-D node to perform a coupling similar in nature to the phase-step coupling described above. This DC coupling node acts as a mechanism for finding a consensus phase among the DC nodes. This function could also be accomplished using random connections between DC nodes, similar to how the phase-step nodes arrive at a consensus.

A more complex form of coupling is required to keep the 1-D phase ramps from the individual propellers coplanar with each other. While the phase-step nodes keep the phase linear within a 1-D propellor, we still need a way to ensure that the VCO phases form a linear function (a plane) in 2-D that passes through the origin. For example, drift could cause one propellor to attain a disproportionately steep slope that makes it tilt out of the plane delineated by the other propellers.

In order to ensure that the 1-D phase ramps are coplanar, we chose to couple together three phase-step nodes (from three different propellers). If the propellor orientations are offset by multiples of 60° , the coupling constraint becomes quite simple (omitted here). The resulting phase adjustments are fed back to the phase-step nodes.

2.4. SIMULATION OF RAT MOTION

We created our network model to test some specific aspects of the Fourier model. In particular, we wanted to see if we would find

grid cells that fired spikes on a hexagonal grid of locations. We also wanted to see if these grid cells would exhibit phase precession compared to a global theta cycle. We added a 2-D VCO node that oscillates at approximately 8 Hz, and used this node's state as the definitive theta cycle.

To simulate the movement of a rat in a circular environment, we added a random-walk function that adjusts the velocity vector smoothly. One could predict the rat's location by numerically integrating the rat's velocity. However, the rat's own perceived location (as encoded in the phase ramp of the EC VCOs) soon drifted away from the computed position. This drift phenomenon has been observed before (Zilli and Hasselmo, 2010) and is probably due to temporal delays in network activity, network transients caused by sudden input changes, and inaccuracies in the frequencies of the VCOs². A real rat seems to avoid this problem by updating its perceived location with sensory information (Burgess et al., 2007). Our model has no sensory input. In terms of assessing the spatial maps of the rat's EC neurons, what is important is where the rat believes it is, not necessarily where the rat actually is (Barry et al. (2007) and Blair et al. (2007)). We determine the rat's perceived location from the slopes of the phase ramps of the three propellers. In particular, the phase-step nodes encode the slope that we need. Each propeller gives us a projection of the rat's position onto the propeller³. Combining the three projections gives us an over-determined system; we find the least-squares solution to get a good estimate of the rat's perceived location.

2.5. NETWORK ARCHITECTURE

As shown in **Figure 11**, the network consists of three “wheels” of nodes, along with a velocity node, DC phase-coupling node, a

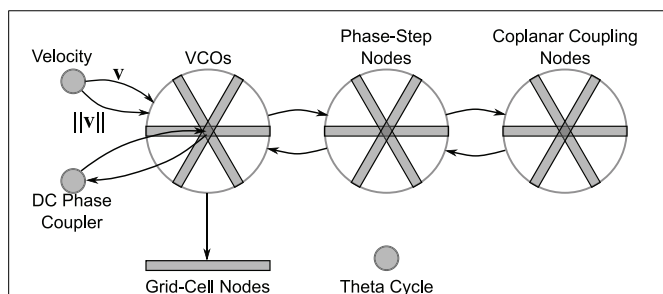


FIGURE 11 | Network overview. The velocity node modulates the frequency of the VCOs (see Equation 3). The phase-step nodes couple the VCOs to maintain a 1-D phase ramp within each propeller. The coplanar coupling nodes further keep the phase slopes of the different propellers linearly consistent (so that they all rest in a common plane). The DC phase coupler node keeps the absolute phase of the propellers synchronized. The grid cells sum triads of VCOs. The theta cycle node is a stand-alone oscillator that maintains a frequency of approximately 8 Hz.

²We did not assess the accuracy of the VCO frequencies after implementing the various couplings, though it would be an interesting study.

³Proof that the rat's location is projected onto each propeller comes from yet another fascinating and useful property of the Fourier transform, the Fourier Projection Theorem. This property is used heavily in computed tomography reconstruction (Natterer and Wübbeling, 2001).

theta-cycle node, and an array of grid-cell nodes. Each wheel has three propellers at angles 0°, 120°, and 240° (though a full model would include more propellers per wheel). The first wheel contains 17 VCO populations per propeller. Each population has 300 LIF neurons and encodes a 3-D vector. The recurrent connections of these oscillating populations have a synaptic time constant [τ_s , in Equation (B2) of Appendix B] of 10 ms.

The phase-step wheel also has three propellers, but with 16 nodes per propeller (since they model the phase differences between the VCO nodes). Each phase-step population has 500 LIF neurons and encodes a 6-D vector as described in the previous section. The coplanar coupling wheel mirrors the phase-step wheel, with each coplanar coupling node having 500 LIF neurons and encoding a 6-D vector.

In hopes of observing grid-cell behavior in our model, we created an array of grid-cell nodes (see **Figure 11**). The grid-cell array contains 17 nodes, mirroring the 17 nodes in each of the VCO propellers. Each grid-cell node (containing 200 LIF neurons) encodes a 2-D vector sum of the phase states from the three corresponding VCOs. That is, each grid-cell node receives the phase state from a triad of VCOs (VCOs at the same radius, but oriented 120° apart) and simply adds them together.

The DC phase-coupler node has 500 LIF neurons and encodes a 6-D vector that duplicates the phases of the three DC nodes. The velocity node has 100 LIF neurons and encodes a 2-D vector. Finally, the theta-cycle node contains 500 LIF neurons and encodes a 2-D vector that oscillates at approximately 8 Hz. The recurrent connections of the theta-cycle population use a synaptic time constant of $\tau_s = 5$ ms.

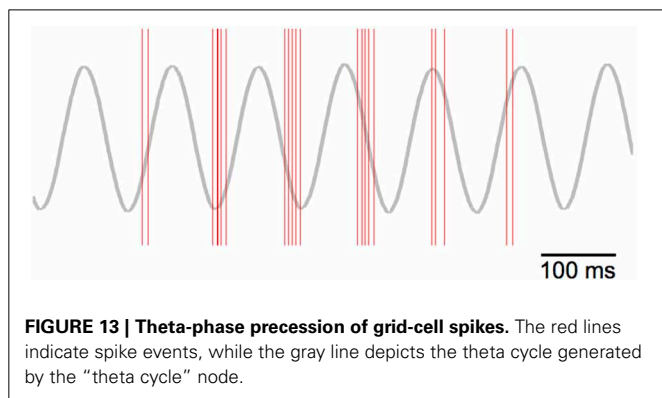
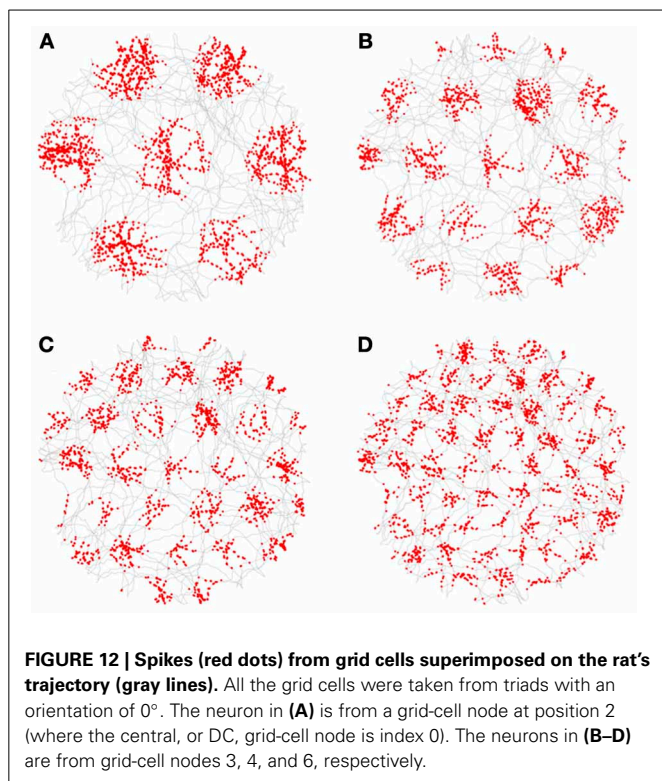
Unless otherwise specified, we used the following parameter values for all neurons: synaptic time constant $\tau_s = 5$ ms, refractory period $\tau_{ref} = 2$ ms, membrane time constant $\tau_m = 20$ ms, spiking threshold $J_{th} = 1$, encoding vectors (\mathbf{e}_n from Equation B1) selected randomly (uniformly) from the unit hypersphere, neural gain and bias (α_n and β_n from Equation B1) chosen to randomly (uniformly) sample the unit hypersphere of the representational space, with a maximum firing rate in the range 200–400 Hz.

3. RESULTS

The simulations were run using the Nengo software package (nengo.ca). The whole model includes 119 nodes, for a total of 68,700 LIF neurons. We ran the model for 300 s simulation time. The execution of the model took about 110 min to run on a laptop with a 2.9 GHz Intel Core i7 processor and 8 GB of RAM.

3.1. GRID CELLS

Figure 12 shows a sampling of activity from neurons in the grid-cell nodes, with their spikes (shown as red dots) superimposed over of the rat's trajectory (shown as gray lines). In the figure, the spatial frequency of the grid-cell triad increases from A to D. Recall that each triad combines the output from the three corresponding VCO nodes, one taken from each propeller. The further from the center (DC), the higher the spatial frequency. The red dots of spikes clearly occur on a hexagonal grid with different scales. Not all neurons in the grid nodes exhibited grid firing patterns. However, about 10% did.



3.2. THETA-PHASE PRECESSION

If we focus on the timing of the grid-cell spike bursts, we can see that the start of the bursts precess through the theta cycle. **Figure 13** plots the spikes as red lines over the theta cycle produced by the “theta” node. The frequency of oscillation for the VCOs—and hence the grid cells—is higher than the nominal 8 Hz theta cycle (as shown in **Figure 8**). Thus, we see the bursts of grid-cell activity precess through the lower-frequency theta cycle.

4. DISCUSSION

The model proposed in this paper extends the model proposed by Welday et al. (2011). We have re-formulated their bank-of-oscillators model using Fourier theory. The VCOs oscillate, so various combinations of them can result in complex interference patterns. Fourier theory is a convenient and powerful tool for understanding these interference patterns.

Our Fourier interpretation points out a global constraint on the VCOs, that they must maintain a phase ramp in the address space. The Fourier Shift Theorem shows us that the location of the animal in its 2-D environment is encoded in the slope of this phase ramp.

Moreover, this phase constraint can be used to generate (or interpret) coupling mechanisms that keep the phases in the proper formation. The VCO phases tend to drift, destroying the location information. Coupling mechanisms that allow phases to change, but only in a manner consistent with a linear phase ramp, maintains the location information.

Each VCO generates a Fourier basis function on the animal's environment. Spatial maps, such as place fields and grid fields, result from the interference patterns between these basis functions. Our Fourier-theory interpretation suggests that the connection weights from the VCOs to the read-out nodes constitute Fourier coefficients. That is, the connections alter the amplitude and phase of each Fourier basis function. Applying a different weight to each VCO is like applying a different Fourier coefficient to each Fourier basis function. Hence, each read-out node performs an inverse Fourier transform.

The inverse Fourier transform implemented by the connections is modulated by the phase state of the VCOs. Since the VCOs are constrained to maintain a linear phase ramp, the spatial maps get shifted (thanks to the Fourier Shift Theorem). The slope of the phase ramp encodes the animal's position in its environment, so the spatial maps are all shifted in concert with the animal's motion. This framework makes it easy to understand the relationship between the VCOs and cells that exhibit spatial maps.

Interestingly, grid cells might be a byproduct of the computation of place cells. All the VCOs in **Figure 4A** are added together to get the place cell. However, that sum could be done in stages. As an intermediate step to adding all the VCOs together, all possible triads could be added together to form a bank of grid-cell nodes. Then, all those grid cells could be added together to form the place cell. This two-part process is equivalent to adding all of the VCOs together. In addition, variants of grid-like cells have been observed in the EC (Krupic et al., 2012). Just like traditional grid cells, these cells share the property that their spatial activation maps are composed of a small number of Fourier components. That is, their Fourier transforms consist of only a few non-zero Fourier coefficients. Perhaps there is something intrinsically beneficial in combining small sets of Fourier components.

Additionally, grid cells might emerge as a by-product of a phase coupling mechanism. Some research has shown that the distributed nature of grid-cell encoding offers better accuracy than the same number of sparse place cells (Mathis and Herz, 2012). But this theory still does not address *why* grid cells appear, since the bank of VCOs already offers a distributed representation of location. Another theory, and one that we plan to investigate, is that grid cells are a by-product of the coupling mechanisms that maintain the phase relationships within the bank of VCOs. It seems intuitive that place cells could offer a stable and accurate representation of location as long as the underlying network that feeds into the place cells encodes location in a stable and accurate manner. Coupling between nodes harnesses the redundancy in the network and enables wide-spread resources to be focused on

lower-dimensional data, such as location. The coplanar-coupling nodes assess the linear consistency among three or more other nodes. In general, a linearity constraint in 2-D will always require input from at least three VCOs (in addition to the implicitly-included DC node). We plan to investigate more general implementations of the coplanar constraint and observe whether these mechanisms inherently generate grid-cell behaviors.

Couching the behavior of the EC in terms of Fourier theory opens up new vistas of interpretations and predictions. It gives us the mathematical machinery to contemplate other neurophysiological observations. For example, how might sensory feedback be incorporated into the EC? If an animal is given a visual cue of its location, that sensory data might excite the corresponding place cell, which—in turn—could feed back through to the phase-step nodes to adjust the slope of the phase ramp so that it matches the animal's location. It would seem that this feature would be accompanied by a phase-resetting mechanism that allows the VCOs to rapidly realign their phases (or take on some other phase-ramp state). A similar phase-resetting method was able to correct phase drift using error feedback (Monaco et al., 2011). In that work, each landmark has its own phase signature for the VCOs, and as the animal approaches the landmark, its VCOs receive corrective adjustments to bring them into the corresponding phase state. While this approach uses phase adjustments similar to our phase-coupling mechanisms, their method corrects the absolute phase and requires that the entire set of VCO phases be recorded. Their method does not improve the stability of the path-integration process itself. Future work on our Fourier EC model could incorporate these sensory-based error correction mechanisms.

It is well known that different environments elicit different spatial representations in the hippocampus and EC (Barry et al., 2007; Fyhn et al., 2007). While place cells can show entirely different activation maps, grid cells all seem to alter their orientation, scale and phase in concert (Fyhn et al., 2007). Orientation and scale can be accounted for in our model by rotating and scaling the input velocity vector. A phase shift can be implemented by setting the bank of VCOs with an initial phase ramp. As a future investigation, we could consider allowing arbitrary initial VCO phases, similar to the method outlined by Monaco et al. (2011), in which they assign random initial phases to generate independent phase codes for different environments.

The network we have built involves 119 neural populations, and contains a total of 67,800 LIF neurons. Our implementation is an important step in demonstrating the capabilities and behaviors of our model. However, an obvious question remains, how might such a system get established? What self-organizing principles might apply, and where? Spatial maps of place cells have been learned using Hebbian learning (Rolls et al., 2006). Grid cells can emerge spontaneously in a topographically connected network with local excitation and lateral inhibition (Fuhs and Touretzky, 2006; McNaughton et al., 2006). However, these “Turing grids” are not found in adults, leading researchers to suggest that they form during a developmental stage and are used to guide the formation of grid cells in the non-topographical, adult EC network. Even a random selection of grid cells can produce place cells (Solstad et al., 2006; de Almeida et al., 2009). We plan

to investigate unsupervised and supervised learning algorithms to derive neural oscillators. We also plan to investigate learning mechanisms that could result in the proper phase-coupling between two or more VCOs.

A recent study found that the dynamics of grid-cell subthreshold membrane potentials not only exhibit theta oscillations, but also systematic depolarizations (so called “ramps”) that coincide with grid fields (Domnisoru et al., 2013). That paper proposes that these depolarizations explain grid fields better than theta interference patterns, posing a challenge to simple oscillator-interference models. The version of our model presented here does not attempt to address those observations; the LIF neuron model we use does not have a facility to change the resting (reset) potential. However, grid cells are not an essential part of our model, and we merely show that they can easily be constructed from the model's raw materials. In the future, it will be interesting to try to incorporate these new observations into an expanded Fourier model of the EC.

5. CONCLUSION

Our Fourier model of the EC path-integration system extends the bank-of-oscillators model proposed by Welday et al. (2011). Using the rich field of Fourier theory, our model suggests the separation of navigation into two components: location, and spatial maps. An animal's location is encoded by the collective phases of the VCOs, which Fourier theory tells us should be arranged in a linear ramp. The slope of the ramp indicates position. The spatial maps that stem from these oscillators can be thought of as inverse Fourier transforms, where the connection weights dictate the corresponding Fourier coefficients. This division of labor allows the EC to maintain a single, dynamic representation of its location, while enabling the construction of many different spatial maps.

The fact that the VCOs all have to maintain a linear ramp in the address space means that this large set of oscillators is actually a distributed encoding of a small piece of data, location. Coupling can keep the VCOs in ramp formation, but allow the slope of the ramp to change freely. We have demonstrated this coupling feature in a system model that uses populations of spiking LIF neurons.

Knowledge about other properties of the Fourier transform can help to guide further development of the model, and assess how it may (or may not) be extended to explain or predict other observations.

ACKNOWLEDGMENTS

We are thankful for the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Foundation for Innovation (CFI), and the Ontario Innovation Trust.

REFERENCES

- Barry, C., Hayman, R., Burgess, N., and Jeffery, K. (2007). Experience-dependent rescaling of entorhinal grids. *Nat. Neurosci.* 10, 682–684. doi: 10.1038/nn1905
- Blair, H., Gupta, K., and Zhang, K. (2008). Conversion of a phase- to a rate-coded position signal by a three-stage model of theta cells, grid cells, and place cells. *Hippocampus* 18, 1239–1255. doi: 10.1002/hipo.20509

- Blair, H., Welday, A. C., and Zhang, K. (2007). Scale-invariant memory representations emerge from Moiré interference between grid fields that produce theta oscillations: a computational model. *J. Neurosci.* 27, 3211–3229. doi: 10.1523/JNEUROSCI.4724-06.2007
- Burgess, N., Barry, C., and O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus* 17, 801–812. doi: 10.1002/hipo.20327
- Conklin, J., and Eliasmith, C. (2005). A controlled attractor network model of path integration in the rat. *J. Comput. Neurosci.* 18, 183–203. doi: 10.1007/s10827-005-6558-z
- de Almeida, L., Idiart, M., and Lisman, J. E. (2009). The input–output transformation of the hippocampal granule cells: from grid cells to place fields. *J. Neurosci.* 29, 7504–7512. doi: 10.1523/JNEUROSCI.6048-08.2009
- Domnisoru, C., Kinkhabwala, A. A., and Tank, D. W. (2013). Membrane potential dynamics of grid cells. *Nature* 495, 199–204. doi: 10.1038/nature11973
- Eliasmith, C., and Anderson, C. H. (2003). *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. Cambridge, MA: MIT Press.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, C., et al. (2012). A large-scale model of the functioning brain. *Science* 338, 1202–1205. doi: 10.1126/science.1225266
- Fuhs, M. C., and Touretzky, D. S. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *J. Neurosci.* 26, 4266–4276. doi: 10.1523/JNEUROSCI.4353-05.2006
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., and Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446, 190–194. doi: 10.1038/nature05601
- Geisler, C., Robbe, D., Zugaro, M., Sirota, A., and Buzsáki, G. (2007). Hippocampal place cell assemblies are speed-controlled oscillators. *PNAS* 104, 8149–8154. doi: 10.1073/pnas.0610121104
- Hafting, T., Fyhn, M., Molden, S., Moser, M., and Moser, E. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* 436, 801–806. doi: 10.1038/nature03721
- Koch, C. (1999). *Biophysics of Computation: Information Processing in Single Neurons*. New York, NY: Oxford University Press.
- Krupic, J., Burgess, N., and O'Keefe, J. (2012). Neural representations of location composed of spatially periodic bands. *Science* 337, 853–857. doi: 10.1126/science
- Mathis, A., and Herz, A. (2012). Optimal population codes for space: grid cells outperform place cells. *Neural Comput.* 24, 2280–2317. doi: 10.1162/NECO_a_00319
- McNaughton, B., Battaglia, F. P., Jensen, O., Moser, E., and Moser, M. (2006). Path integration and the neural basis of the 'cognitive map'. *Nat. Rev. Neurosci.* 7, 663–678. doi: 10.1038/nrn1932
- Monaco, J. D., Knierim, J. J., and Zhang, K. (2011). Sensory feedback, error correction, and remapping in a multiple oscillator model of place-cell activity. *Front. Comput. Neurosci.* 5:39. doi: 10.3389/fncom.2011.00039
- Moser, E., Kropff, E., and Moser, M. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.* 31, 69–89. doi: 10.1146/annurev.neuro.31.061307.090723
- Muller, R. U., Kubie, J. L., and Ranck, J. B. Jr. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J. Neurosci.* 7, 1935–1950.
- Natterer, F., and Wübbeling, F. (2001). *Mathematical Methods in Image Reconstruction*. Philadelphia: SIAM. doi: 10.1137/1.9780898718324
- O'Keefe, J., and Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15, 853–866. doi: 10.1002/hipo.20115
- O'Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34, 171–175. doi: 10.1016/0006-8993(71)90358-1
- O'Keefe, J., and Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317–330. doi: 10.1002/hipo.450030307
- Oppenheim, A. V., and Schaffer, R. W. (1999). *Discrete-Time Signal Processing*. 2nd Edn. Upper Saddle River, NJ: Prentice Hall.
- Rolls, E. T., Stringer, S. M., and Elliot, T. (2006). Entorhinal cortex grid cells can map to hippocampal place cells by competitive learning. *Netw. Comput. Neural Syst.* 17, 447–465. doi: 10.1080/09548980601064846
- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B., Witter, M. P., Moser, M., et al. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312, 758–762. doi: 10.1126/science.1125572
- Solstad, T., Moser, E., and Einevoll, G. T. (2006). From grid cells to place cells: a mathematical model. *Hippocampus* 16, 1026–1031. doi: 10.1002/hipo.20244
- Tsodyks, M., Skaggs, W., Sejnowski, T., and McNaughton, B. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. *Hippocampus* 6, 271–280. doi: 10.1002/(SICI)1098-1063(1996)6:3<271::AID-HIPO5>3.3.CO;2-Q
- Welday, A. C., Shlifer, I. G., Bloom, M. L., Zhang, K., and Blair, H. T. (2011). Cosine directional tuning of theta cell burst frequencies: evidence for spatial coding by oscillatory interference. *J. Neurosci.* 31, 16157–16176. doi: 10.1523/JNEUROSCI.0712-11.2011
- Zilli, E. A. (2012). Models of grid cell spatial firing published 2005–2011. *Front. Neural Circ.* 6, 1–17. doi: 10.3389/fnirc.2012.00016
- Zilli, E. A., and Hasselmo, M. E. (2010). Coupled noisy spiking neurons as velocity-controlled oscillators in a model of grid cell spatial firing. *J. Neurosci.* 30, 13850–13860. doi: 10.1523/JNEUROSCI.0547-10.2010

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 05 November 2012; accepted: 25 November 2013; published online: 11 December 2013.

Citation: Orchard J, Yang H and Ji X (2013) Does the entorhinal cortex use the Fourier transform? *Front. Comput. Neurosci.* 7:179. doi: 10.3389/fncom.2013.00179
This article was submitted to the journal *Frontiers in Computational Neuroscience*. Copyright © 2013 Orchard, Yang and Ji. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

APPENDIX

A. FOURIER THEORY BASICS

We will develop the theory using the Discrete Fourier Transform (DFT), but point out that analogous properties exist for the continuous-domain Fourier transform (Oppenheim and Schaffer, 1999).

An efficient way to write the Fourier transform is to use complex numbers. Recall that we have two ways of denoting a complex number: using the Cartesian form $a + ib$, or the polar form $r \exp(i\phi)$. In the Cartesian form, we refer to a as the *real part*, and b as the *imaginary part*. In the polar (or exponential) form, we call r the *modulus*, and ϕ the *phase angle*. The polar notation carries with it several advantages. Multiplying $r_1 \exp(i\phi_1)$ and $r_2 \exp(i\phi_2)$ is consistent with the algebraic rules for multiplying exponential functions, giving $r_1 r_2 \exp(i(\phi_1 + \phi_2))$. In words, you multiply their moduli, and add their phases.

We will start our review of the Fourier transform with a 1-D function. Consider a sampled function f with N samples indexed $n = 0, \dots, N - 1$,

$$f = [f_0, f_1, \dots, f_{N-2}, f_{N-1}].$$

The DFT (Oppenheim and Schaffer, 1999) of f is

$$F_k = \sum_{n=0}^{N-1} f_n \exp\left(-2\pi i \frac{nk}{N}\right), \quad k = 0, \dots, N - 1.$$

Each complex number F_k is called a Fourier coefficient. We can also denote the transform using $F = \text{DFT}(f)$. In essence, the DFT is a frequency decomposition; it takes a spatial signal and represents it as a sum of wave fronts of various frequencies (and orientations, in higher dimensions). Each Fourier coefficient corresponds to a different location in the frequency domain. Each location in the frequency domain represents a different wave front. The value of a Fourier coefficient, F_k , represents the contribution of the corresponding wave front. The coefficient F_0 has a special name; it is called the DC, and it always corresponds to the zero frequency and is located at the origin of the frequency domain.

The Fourier basis functions are N -periodic. That is,

$$\exp\left(-2\pi i \frac{(n+N)k}{N}\right) = \exp\left(-2\pi i \frac{nk}{N}\right).$$

In other words, it is perfectly valid to refer to F_{-1} , since $F_{-1} = F_{N-1}$. Because of this periodicity, we can shift our array of Fourier coefficients and visualize them as being centered on the DC. For example, if N is 33, then we list our Fourier coefficients using $[F_{-16}, \dots, F_{-1}, F_0, F_1, \dots, F_{16}]$. Likewise, use of the DFT implicitly assumes that f is also periodic and can be indexed in the same manner. Thus, an equivalent formulation of the DFT is

$$F_k = \sum_{n=-\tilde{N}}^{\tilde{N}} f_n \exp\left(-2\pi i \frac{nk}{N}\right), \quad k = -\tilde{N}, \dots, \tilde{N},$$

where we assume for simplicity that N is odd⁴ and use the symbol \tilde{N} to represent $\lfloor \frac{N}{2} \rfloor$, where the delimiters $\lfloor \cdot \rfloor$ denote rounding toward zero. We will use this equivalent, centered version of the DFT throughout this paper.

As an example relevant to our purposes, suppose we have a rat that moves along a small corridor, and that we have broken the corridor into N blocks, indexed $n = -\tilde{N}, \dots, \tilde{N}$. We can represent the location of the rat using an array, δ , indicating where the rat is by placing a 1 in the element corresponding to the rat's location, and setting all the other elements to zero. As the rat moves along the corridor, the array δ changes so that the location of the 1 reflects which block the rat occupies. By this definition, δ is the Kronecker Delta,

$$\delta_{xn} = \begin{cases} 1 & \text{if } n = x \\ 0 & \text{if } n \neq x \end{cases}$$

where the subscript x represents the index of the block containing the rat, and the subscript n indexes the elements of the array δ_x .

The DFT of δ has special properties. The Fourier coefficients of δ_0 are $F_k = 1$, all unit-modulus with zero phase. However, when the rat is at position x , the DFT of δ_x is,

$$\begin{aligned} G_k &= \sum_{n=-\tilde{N}}^{\tilde{N}} \delta_{xn} \exp\left(-2\pi i \frac{nk}{N}\right) \\ &= \exp\left(-2\pi i \frac{xk}{N}\right). \end{aligned}$$

These Fourier coefficients all have a modulus of 1, and their phases vary linearly with k (the frequency index). That is, if $G_k = r_k \exp(i\phi_k)$, then $r_k = 1$ and

$$\phi_k = -2\pi \frac{xk}{N}. \quad (\text{A1})$$

The modulus and phase of δ_2 are shown in **Figure A1**. Notice that the phases in **Figure A1** form a line (or ramp) as you move along the frequency axis, although the phases are wrapped into the range $[-\pi, \pi]$. The slope of the line is $(-2\pi x/N)$, so the larger x is, the steeper the slope. The sign of the slope reflects whether x is positive or negative (which direction the rat moved along the corridor).

This property of the DFT is similar to a more general concept known as the Fourier Shift Theorem. The Fourier Shift Theorem tells us how shifting (translating) a signal influences its Fourier coefficients. Suppose that F_k are the Fourier coefficients of a signal f_n . Consider a shifted version, f_{n-x} , and its Fourier coefficients, G_k . What is the relationship between G_k and F_k ? The Fourier coefficients of f_{n-x} can be written

$$G_k = \sum_{n=-\tilde{N}}^{\tilde{N}} f_{n-x} \exp\left(-2\pi i \frac{nk}{N}\right).$$

⁴If N is even, then we sum over $n = -\tilde{N}, \dots, \tilde{N} - 1$.

Using the change of variables $m = n - x$, we get

$$G_k = \sum_{m=-\tilde{N}-x}^{\tilde{N}-x} f_m \exp\left(-2\pi i \frac{(m+x)k}{N}\right).$$

Both $\exp(\cdot)$ and f are periodic (by assumption, for f), so we can shift the summation range without changing its value. Moreover, the exponential function can be split into two components, one containing m , while the other can be pulled out of the summation,

$$G_k = \exp\left(-2\pi i \frac{xk}{N}\right) \sum_{m=-\tilde{N}}^{\tilde{N}} f_m \exp\left(-2\pi i \frac{mk}{N}\right)$$

$$G_k = \exp\left(-2\pi i \frac{xk}{N}\right) F_k, \quad k = -\tilde{N}, \dots, \tilde{N}. \quad (\text{A2})$$

This is the Fourier Shift Theorem, and it tells us that the Fourier coefficients of the shifted signal, f_{n-x} , can be derived from the coefficients of the original signal by simply multiplying each coefficient by a phase-shift, where the amount of the phase-shift is a linear function of the frequency index k .

Looking back at our example, our rat started at location 0 and the Fourier coefficients of δ_0 were $F_k = 1$, $k = -\tilde{N}, \dots, \tilde{N}$. However, when the rat moved to location x , the Fourier coefficients still had unit modulus, but their phases formed a ramp, $\exp(-2\pi i xk/N)$. That is, $G_k = \exp(-2\pi i xk/N) F_k$. Applying such a phase ramp can shift any function, not just our δ functions.

The DFT is an invertible transform, and the inverse DFT (IDFT) yields back the original function. For example, IDFT(DFT(δ_x)) = δ_x . With this in mind, we can represent any shifted version of δ_0 by multiplying its Fourier coefficients by a phase ramp of the desired slope. In other words, the slope of the phase ramp in the Fourier coefficients indicates the shift applied to δ_0 . In this way, the phases of the Fourier coefficients encode the position of the rat.

What if the slope ($-2\pi x/N$) of the phase ramp along k corresponds to a value of x that is not an integer? We cannot expect the IDFT to give us δ_x exactly, since that is only defined for integer values of x . As it turns out, non-integer values of x yield δ -like signals. The operation is called “Fourier interpolation,” and corresponds to summing the continuous-domain wave fronts

and sampling the resulting function. **Figure A2** illustrates a non-integer shift, and how the samples computed by the IDFT relate to the underlying continuous-domain Fourier reconstruction.

All of this theory extends trivially to 2-D domains (and higher). A Kronecker delta (or any function, for that matter) on a 2-D domain can be shifted by multiplying its Fourier coefficients by a phase ramp in 2-D.

B. NEURAL ENGINEERING FRAMEWORK

To build our neural network, we used the Neural Engineering Framework (NEF) (Eliasmith and Anderson, 2003), a powerful and versatile platform that has proven useful for large-scale cognitive modeling (Eliasmith et al., 2012).

The NEF is comprised of three main principles: (1) neural encoding and decoding, (2) neural transformations, and (3) neural dynamics.

B.1. NEURAL ENCODING AND DECODING

In this framework, data is stored in the collective activities of a population of LIF neurons. The neurons have a range of parameters so that their tuning curves span a wide range of possibilities. A population of N neurons (a “node”) can encode a value \mathbf{x} in its neural activities using

$$a_n(t) = G_n(\mathbf{x}(t)) \cdot \mathbf{e}_n \alpha_n + \beta_n \quad (\text{B1})$$

where \mathbf{e}_n is the encoding vector (preferred direction vector), and α_n and β_n are scalar gain and bias terms that account for the neural climate of the neuron. The input to the function $G_n(\cdot)$ can be thought of as the input current driving neuron n . The function G_n translates the input current to neural activity, either in the form of a firing rate, or a series of spikes. In the case of firing rate, G_n is the steady-state LIF tuning curve,

$$G_n(J) = \begin{cases} \frac{1}{\tau_{\text{ref}} - \tau_m \ln\left(1 - \frac{J_{th}}{J}\right)} & \text{for } J > J_{th} \\ 0 & \text{otherwise} \end{cases}$$

where τ_{ref} is the refractory period, τ_m is the membrane time constant, and J_{th} is the threshold current, below which the neuron has

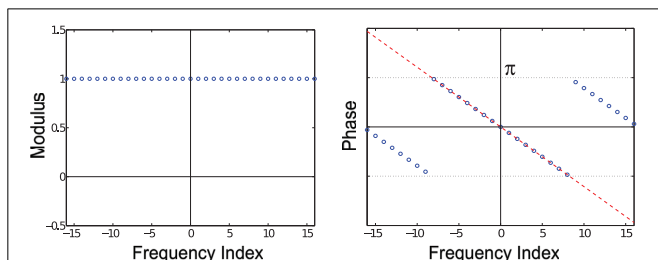


FIGURE A1 | Fourier transform of δ_2 . The moduli (left) of the 33 Fourier coefficients are all 1. The phase (right) is linear; the dashed red line shows a linear function, while the blue dots show the same linear function wrapped into the range $[-\pi, \pi]$.

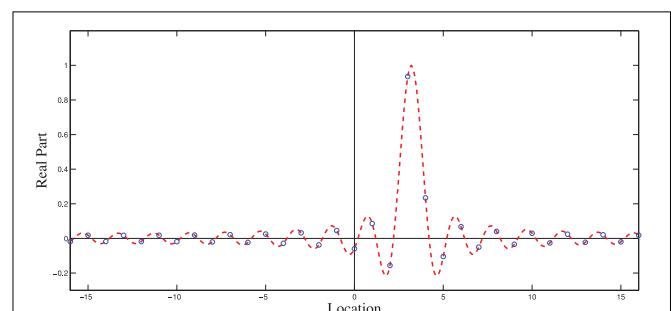


FIGURE A2 | Fourier interpolation of δ_0 shifted to the right by 3.2 units. The dotted red line shows the continuous-domain reconstruction from the phase-shifted Fourier coefficients. The blue circles show the samples produced using the IDFT.

a firing rate of zero. On the other hand, if using spikes, then the output of G_n is represented as a sum of time-shifted Dirac delta functions (Oppenheim and Schaffer, 1999),

$$G_n(J_n(t)) = \sum_p \delta(t - t_{np}),$$

where t_{np} is the time of the p th spike from neuron n . The membrane potential, v , is governed by the differential equation,

$$\tau_m \frac{dv}{dt} = RJ_n(t) - v,$$

where $J_n(t)$ is the input current, and R is the membrane resistance. Once the membrane potential reaches its threshold, the neuron spikes, the timing of the spike is recorded, and the membrane potential is reset to zero.

If we wish to decode the neural activities of a population of neurons, we can compute the optimal linear decoders. We do this by collecting a sampling of inputs, X , and corresponding neural activities, A . That is, each row of matrix X stores a sample input, and each row of matrix A stores the corresponding neural activities for all N neurons (usually stored as firing rates). To decode from our population, we seek the linear weights D that solve

$$\min_D \|AD - X\|_2^2.$$

Thus, the weights in D perform a linear transformation from the space of neural activities to the space of input values. This is a linear least-squares problem, and there are multiple ways to compute D numerically. Once we have D , we can decode neural activities to get an estimate of the value being encoded. Moreover, we can decode arbitrary functions of our encoded data by finding the decoders that solve

$$\min_D \|AD - f(X)\|_2^2,$$

where $f(X)$ is a function of the encoded values.

B.2. NEURAL TRANSFORMATIONS

We can use our neural encoders and decoders to transform data from one population P , to another population Q . To do this, we essentially decode the desired function from P and re-encode the

result into Q . Collapsing those processes together gives the $N \times M$ weight matrix

$$W = D_p E_Q \alpha_Q,$$

where D_p is the matrix that decodes the neural activities from the N neurons in P , and $E_Q \alpha_Q$ is the matrix of M encoders for the neurons in Q . The weight matrix simply combines the linear decoders and encoders into a single matrix.

B.3. NEURAL DYNAMICS

The neural coding and transformation principles above can be built into a dynamic framework by including the temporal action of synapses. When a spike arrives at a synapse, it induces a current on the post-synaptic neuron. We model this post-synaptic current using exponential decay⁵. That is, we convolve an incoming spike train with the post-synaptic filter, $h(t)$,

$$h(t) = \frac{1}{\tau_s} \exp\left(\frac{-t}{\tau_s}\right), \quad (\text{B2})$$

where τ_s is the decay time constant. The total current entering a neuron is a weighted sum over all incoming spike trains so that the total input current arriving at neuron m is

$$J_m(t) = h(t) \star \left[\sum_n w_{nm} \sum_p \delta(t - t_{np}) \right],$$

where \star represents convolution, and w_{nm} is the connection weight from neuron n to neuron m .

We can use a recurrently-connected population of neurons to implement a dynamic model of the form

$$\frac{dx}{dt} = f(\mathbf{x})$$

by choosing the recurrent connection weights so that they decode and feed back $\tau_s f(\mathbf{x}) + \mathbf{x}$, where τ_s is the post-synaptic time constant (Eliasmith and Anderson, 2003).

With these three principles in place, one can implement a dynamical system using spiking LIF neurons by assigning populations to state variables, and connecting them with the appropriate transformations and time constants. A good example of such a dynamical system is the neural simple harmonic oscillator, described in section 2.3.

⁵The NEF offers other post-synaptic filters.