

On the optimality of grid cells *

Christos Papadimitriou
EECS, UC Berkeley

June 16, 2016

Abstract

Grid cells, discovered more than a decade ago [5], are neurons in the brain of mammals that fire when the animal is located near certain specific points in its familiar terrain. Intriguingly, these points form, for a single cell, a two-dimensional triangular grid, not unlike our Figure 3. Grid cells are widely believed to be involved in path integration, that is, the maintenance of a location state through the summation of small displacements. We provide theoretical evidence for this assertion by showing that cells with grid-like tuning curves are indeed well adapted for the path integration task. In particular we prove that, in one dimension under Gaussian noise, the sensitivity of measuring small displacements is maximized by a population of neurons whose tuning curves are near-sinusoids — that is to say, with peaks forming a one-dimensional grid. We also show that effective computation of the displacement is possible through a second population of cells whose sinusoid tuning curves are in phase difference from the first. In two dimensions, under additional assumptions it can be shown that measurement sensitivity is optimized by the product of two sinusoids, again yielding a grid-like pattern. We discuss the connection of our results to the triangular grid pattern observed in animals.

1 Introduction

Grid cells [5] are neurons in the dorsocodal medial entorhinal cortex of mammals that fire when the animal is near specific locations in its familiar environment; intriguingly, these locations form, for a single cell, a two-dimensional regular triangular grid [5]. Ever since their discovery, grid cells have been hypothesized to be involved in space representation [5, 4], and in particular in neural algorithms “*that integrate information about place, distance, and direction*” [5], a task usually referred to as *path integration* [5, 8]. But why are neurons with grid-like tuning curves well adapted for the task of path integration? This is the question we address in this paper.

Path integration presumably entails the measurement of small displacements. Therefore, for path integration to be effective, measurement of small displacements has to be as accurate

*Research supported by NSF grant CCF - 1408635.




Figure 1: The optimum tuning curve

as possible. What is the tuning curve, for neurons measuring small displacements, that has the highest possible sensitivity, that is, the smallest possible variance? We show that, in one dimension, optimal measurement sensitivity is achieved through a one-dimensional grid.

In particular, we consider a population of neurons measuring small displacements on the circle. Working on the circle instead of the line segment, or the infinite line, simplifies the analysis by avoiding edge effects. We assume that the tuning curves of these neurons are cyclical shifts of one another, and that the noise of the measurement is Gaussian. We seek the tuning curve maximizing the accuracy of the measurement. A useful surrogate of accuracy is the *Fisher information* [6], which upper-bounds the accuracy of any estimator. We establish that the tuning curve maximizing Fisher information is a sinusoidal-like wave (see Figure 1) — that is to say, a tuning curve whose peaks form a grid. The frequency of the wave is that of the eigenvector corresponding to the smallest positive eigenvalue of the noise correlation matrix. We say “sinusoidal-like” because, mathematically, the optimal solutions form a family of near-sinusoidal functions parametrized by any function $\psi : [0, 1] \mapsto [0, 1]$, with the common sinusoid corresponding to the identity function (notice the differences between the sinusoid-like waves in Figure 2).

But how can the displacement be read out from the change in spiking rates in such a population? We notice that the actual displacement can be computed with the help of a second cell population which is identical to the first, albeit with tuning curves phase-shifted by 90° .

This last observation about the computation of displacement is relevant when exploring how our analysis can be extended to two dimensions. Again, in order to avoid edge effects we are working on the torus (the unit square with its opposite sides identified). Under quite hefty assumptions (of independence of both correlation and tuning curves, as well as of independent optimization of sensitivity in each direction) it can be shown that the optimum tuning curve in two dimensions is indeed a grid. We speculate that the triangular grid may be the result of optimization of measurement sensitivity under an additional constraint, and that constraint may be computational: In two dimensions, the problem of inferring the displacement from tuning curve change requires not one additional phase-shifted population of cells, but *three* such populations. The interface of the four phase-shifted populations is best achieved through a triangular grid (Figure 3).

Related work: Over the past decade, there has been much theoretical investigation of grid cells, their origin, and their role. It has been noted that a grid-like pattern can result from the interference of two, or possibly three, sinusoidal waves [2], while periodic tuning curves on the circle [11] and the torus [8] can be generated by neural networks (see Figure 3A and B in [4]); however, these models were not proposed in the context of optimizing the accuracy of measurement. It has also been shown that continuous attractor models can generate triangular grid-like responses [1], while experimental data are consistent with a 2-dimensional response of the population [15]. Grid cells were interpreted in [12] and, in a different way, in [3], as very efficient novel neural codes for encoding position and velocity, and in [7] it is shown that in this arena grid cells are more apt than place cells; these works are methodologically close to ours in that they also employ Fisher information for their analysis and comparisons — without, however, seeking the tuning curve design that maximizes it. Cells with one-dimensional grid-like firing patterns have also very recently come up in the analysis of the responses of animal grid cells to one-dimensional environments [14], characterized as projections (“slices”) of a two-dimensional lattice to the one-dimensional circumstances of the experiment. Finally, recently it was claimed in [10] that the tuning curve in one dimension with maximum Fisher information is a sinusoid curve, while the product of two such curves is optimum in two dimensions, results very similar to ours; unfortunately (as pointed out in Sections 2 and 3) the mathematical development in that manuscript contains significant gaps.

2 The optimal tuning curve is periodic

Consider a population of N neurons measuring a small angular displacement at a point on the circle. We assume that the tuning curves of the N neurons are identical, albeit shifted by multiples of the angle $\frac{2\pi}{N}$. We seek the tuning curve that maximizes the sensitivity of measuring small displacements.

The neurons respond to an angular stimulus $\theta \in [0, 2\pi]$, and the tuning curve of the i th neuron is denoted by $f_i(\theta)$; we assume that the tuning curves are identical but shifted by multiples of $\frac{2\pi}{N}$, that is, $f_{i+1 \bmod N}(\theta) = f_i(\theta + \frac{2\pi}{N})$. The average population activity caused by a stimulus is thus the vector $f(\theta) = (f_1(\theta), f_2(\theta), \dots, f_N(\theta))$. The derivative $\frac{df_i}{d\theta}$ is denoted \dot{f}_i , and we denote by $\dot{f}(\theta)$ the vector $(\dot{f}_1(\theta), \dot{f}_2(\theta), \dots, \dot{f}_N(\theta))$.

A stimulus θ results in the response $r_i(\theta) = f_i(\theta) + \eta_i$, for $i = 1, \dots, N$, where η_i is Gaussian noise. The values of noise at different neurons, η_i and η_j , are correlated, and this correlation is assumed to be independent of θ , and denoted $C_{i,j}$, a quantity that depends on the distance on the ring of the neurons i and j . It follows that the noise correlation matrix C is both circulant and symmetric.

Importantly, we also assume that the total signal power — the sum of the squares of the slopes of the tuning curves — is bounded from above by a constant, which we take to be one. Thus $\dot{f}(\theta)^T \dot{f}(\theta) \leq 1$ for all θ .

If the stimulus changes from θ to $\theta + \Delta\theta$, this results in population activity $r(\theta + \Delta\theta)$. The goal of the decoding system is to estimate $\Delta\theta$ from the change in the population response — that is, from $r(\theta) - r(\theta + \Delta\theta)$. To do this as effectively as possible, the overall variance of the measurement must be as small as possible. Instead of this “overall variance”, it is convenient

in this context to work with the *Fisher information* of the population, a function of the tuning curves and the correlation matrix which is known, by the Cramer-Rao theorem [6], to bound from above the accuracy of any unbiased estimator. We seek the tuning curves $f_i(\theta)$ with the largest Fisher information under correlation matrix C .

Under Gaussian noise, and the assumption that C is nonsingular, it is well known [6, 13] that the Fisher information can be written as follows:

$$I(\theta) = \dot{f}(\theta)^T C^{-1} \dot{f}(\theta) \quad (1)$$

Thus, we seek the vector f satisfying $f^T f \leq 1$ that maximizes the right-hand side of (1).

Furthermore, C^{-1} is also symmetric, and its eigenvalues are the inverses of the eigenvalues of C , while its eigenvectors are the same as those of C .

Recall now the Courant-Fischer theorem [9] (stated below for the case of real symmetric matrices and the largest eigenvalue only):

Theorem 1 (Courant-Fischer, 1953) *If A is symmetric, then the vector x in the unit ball $x^T x \leq 1$ that maximizes $x^T A x$ is the eigenvector corresponding to the largest eigenvalue of A .*

Comparing with equation (1), we conclude that the optimum tuning curve vector f has derivative \dot{f} equal to the eigenvector corresponding to the *smallest* positive eigenvalue of C (the inverse of the largest eigenvalue of C^{-1}). What is this eigenvector?

Since C is circulant and symmetric, it is well known [9] that each eigenvalue λ_k , for $k = 0, \dots, \frac{N}{2} - 1$, has multiplicity two, and the two corresponding eigenvectors are the two sinusoidal waves v_k and w_k :

$$v_k = [\cos(0), \cos(k\delta), \cos(2k\delta), \dots, \cos((N-1)k\delta)] \text{ and} \\ w_k = [\sin(0), \sin(k\delta), \sin(2k\delta), \dots, \sin((N-1)k\delta)],$$

where $\delta = \frac{2\pi}{N}$ and $k = 0, 1, \dots, \frac{N}{2} - 1$.

We conclude that the optimum tuning curve vector has derivative of the form

$$\dot{f}(\theta) = \alpha_k(\theta)v_k + \beta_k(\theta)w_k,$$

where $\alpha_k^2(\theta) + \beta_k^2(\theta) = 1$ and the smallest positive eigenvalue of C is λ_k . We now apply the change of variables ¹ $\alpha_k(\theta) = \sin(\phi(\theta))$, $\beta_k(\theta) = \cos(\phi(\theta))$ to obtain

$$\dot{f}_i(\theta) = \sin(\phi(\theta)) \cos(ik\delta) + \cos(\phi(\theta)) \sin(ik\delta) = \sin(\phi(\theta) + ik\delta). \quad (2)$$

¹The otherwise similar argument in [10] does not contain this step, and as a result it is incomplete and the full spectrum of optimal solutions (see Figure 2) is missed.

Recall that we are assuming that the tuning curves of the N neurons are identical, albeit shifted by $\delta = \frac{2\pi}{N}$; that is, for all i and θ , $f_{i+1}(\theta) = f_i(\theta + \delta)$, and thus $\dot{f}_{i+1}(\theta) = \dot{f}_i(\theta + \delta)$. Substituting into (2) we conclude that $\sin(\phi(\theta) + (i+1)k\delta) = \sin(\phi(\theta + \delta) + ik\delta)$ for all θ , or equivalently

$$\phi(\theta + \delta) = \phi(\theta) + k\delta + 2n\pi, \quad (3)$$

for some integer n . The simplest solution of (3) is $\phi(\theta) = (k + nN)\theta + c$ for all θ and for some constant c (which we take zero without loss of generality) and integer n . It follows that $\dot{f}_i(\theta) = \sin(K\theta + ik\delta)$, and thus

$$f_i(\theta) = \frac{1}{K} \cos(K\theta + ik\delta) + \left| \frac{1}{K} \right|, \quad (4)$$

where $K = k + nN$ for some integer n (positive, zero, or negative), and we took the constant of the integration to be $\left| \frac{1}{K} \right|$ so the tuning curve takes only positive values.

However, (3) has many more solutions. Let $\psi(\theta)$ be any function mapping $[0, \delta]$ to the reals, and, for any $\theta \in [0, 2\pi]$, define $\theta' \in [0, \delta]$ and $\theta'' \in \{0, \delta, 2\delta, \dots, 2\pi - \delta\}$ by the equation $\theta = \theta' + \theta''$. Then the function

$$\dot{f}(\theta) = \sin(\psi(\theta') + \theta'')$$

is also the derivative of an optimal tuning curve. Indicatively, in Figure 2 we show the tuning curves resulting from six simple functions ψ , including $\psi(x) = x$ (the true sinusoid).




Figure 2: Six optimum tuning curves

3 Two Dimensions

In view of the one-dimensional result, one may suspect that grid-like structures may also be optimal in two dimensions. Intuitively, it is tempting to try and reduce the two-dimensional

case to the one-dimensional case just solved, and show that the optimum two-dimensional tuning curve must be the product of two sinusoids, and therefore a grid. Unfortunately, this matter turns out to be quite a bit more complicated.

Assume a population of neurons measuring displacements on the torus (the product of two circles, or, equivalently, a square with opposite edges identified in the parallel way), whose tuning curves are shifts of one another along some lattice on the torus defined by the unit vectors along two different directions x and y (not necessarily orthogonal). We make some additional assumptions:

- The noise correlations in the x and y directions have the same form and are independent of each other. That is, the correlation tensor decomposes into the product of two identical correlation matrices.
- We further assume that the tuning curve of the neurons can also be decomposed as the product of identical one-dimensional tuning curves in the x and y dimensions. It then follows that the Fisher information in the x dimension has the form

$$I_x(x) = \left[\dot{f}_x(x)^T C^{-1} \dot{f}_x(x) \right] \cdot \left[(f_y(y)^T C^{-1} f_y(y)) \right], \quad (5)$$

and similarly for the y direction.

- Identifying the optimum tuning curve, even under these assumptions, is still ill-defined, because of the possibly unbounded second factor in (5): f_y must be obtained by integrating \dot{f}_y , a step that introduces an unbounded integration constant. We could of course impose an upper bound on f_y — a justified assumption since neurons cannot fire at arbitrarily high rates, — but then the optimization problem becomes an intractable one, involving integral inequality constraints². We can obtain a meaningful solution only under one additional assumption: That the Fisher information in each of the x and y directions is maximized *independently of the other direction*. That is, the overall sensitivity is not maximized, and instead the sensitivities along the two directions x and y are maximized separately, yielding an overall suboptimal solution. This is not implausible, if one considers the independent evolution of two separate modules, each measuring displacement in one of the two directions.

Under these assumptions, the result for the one-dimensional case does generalize immediately, and the tuning curve maximizing the Fisher information at all stimuli turns out to be the outer product of near-sinusoidal waves in the x and y directions. If the two directions form an angle of 120° , the familiar triangular grid results. The idea that the triangular firing field structure can result from the interference of two oscillations has been suggested before [2]; however, the advantage of this structure was unclear.

But why should the two directions x and y be at an angle of 120° to form the familiar triangular grid observed in [5]? One possible answer comes from algorithmic considerations, discussed next.

²This difficulty is ignored in [10].

4 Computing the displacement

Consider a population of cells around the circle as in the previous section with the sinusoid tuning curve $f(\theta)$ in (4) above, measuring (under noise) the change in firing rate $\Delta f = f(\theta) - f(\theta + \Delta\theta)$. What is the mechanism whereby the displacement $\Delta\theta$ is inferred from the measurement of Δf ? This seems problematic, since the value of the stimulus θ appears to be needed, and keeping track of θ is the purpose of path integration... But upon closer consideration, we note that the decoding mechanism does not quite need the stimulus θ , but just its cosine. Recall that, as $\Delta\theta$ goes to zero, $\Delta f = \dot{f}\Delta\theta$, where $f = \sin(K\theta + \text{const})$. Hence,

$$\Delta\theta = \Delta f \frac{1}{K \cos(K\theta)}. \quad (6)$$

In the one-dimensional case, this additional information can be obtained with a very simple architecture: Suppose that there is a *second* population of neurons, with identical tuning curve $g(\theta)$ with the primary population, albeit shifted by an angle $\alpha \neq 0, \pi$; say $\alpha = 90^\circ$. Then this new population yields a similar equation, with sin replacing cos because of the phase shift:

$$\Delta\theta = \Delta g \frac{1}{K \sin(K\theta)}. \quad (7)$$

Thus we have two trigonometric equations for the two unknowns $\Delta\theta$ and $\cos(K\theta)$, which can be easily solved: By dividing (6) by (7) we note that $\tan(\theta) = \frac{\Delta g}{\Delta f}$, and hence

$$\Delta\theta = \Delta f \frac{1 + \left(\frac{\Delta g}{\Delta f}\right)^2}{K}.$$

That is, the displacement measurement can be computed from the two populations. We conclude that, in one dimension, a second identical population of neurons shifted by 90° suffices for an effective readout.

Now in two dimensions, the equivalent of (6) is

$$\Delta\theta_x K \cos(K\theta_x) + \Delta\theta_y K \cos(K\theta_y) = \Delta f.$$

Notice that now there are *four* unknowns ($\Delta\theta_x$, $\cos(K\theta_x)$, $\Delta\theta_y$, and $\cos(K\theta_y)$). We conclude that *three* additional populations of neurons seem to be required, with different two-dimensional shifts from the original one. The most natural way to implement such a scheme is through shifts in three directions, forming equal angles of 120° with each other (see Figure 3). Hence the familiar triangular grid may be the most natural way to implement this mechanism. Further analytical articulation of this point is the subject of future work.

5 Discussion

What is the origin and utility of the grid cells' distinctive firing field, and what does it have to do with path integration? We have shown that, in one dimension, the tuning curve that optimizes the accuracy of displacement measurements is a near-sinusoid wave, whose peaks




Figure 3: Four interlaced populations of grid cells

naturally form a one-dimensional grid. In two dimensions, we needed several further assumptions in order to show that the optimum tuning curve is the product of two sinusoidal waves in two non-parallel directions. If these two directions form an angle of 120° , the familiar triangular grid results. We have also presented ideas about a possible mechanism for computing the change in position from the change in the response. For two dimensions, our proposed mechanism predicts the existence of four populations of grid cells with regularly displaced firing fields, and suggests that the triangular architecture may be the optimum solution of the joint problem of maximizing both sensitivity and accuracy of decoding.

One may further speculate about grid cells for *three dimensions* (relevant for animals such as bats and sea mammals). Here, our analysis predicts a total of *six* populations, and these must be interlaced in a way analogous to that in Figure 3. The arrangement in Figure 3 works because the simple triangular lattice of one population, say population A, if copied three times and appropriately shifted, has the property that each point has at least one neighbor from each of the other three populations. This suggests the following technical question in 3-dimensional geometry: Is there a lattice in three dimensions with the property that, if it is duplicated six times, each point of each copy of the lattice is adjacent to at least one point from each of the other five copies? It turns out that the lattice generated by the vectors $(0, 3, 0)$, $(0, 0, 2)$, $(1, 1, 1)$ has this property: Notice that the corresponding matrix has determinant 6, suggesting that six copies of the lattice can be arranged in space, and it is easy to check that each point of each copy is at distance one from a point from each of the other copies. Therefore, intriguing compromises do exist between the design of a neural architecture for path integration in three dimensions and the realities of three-dimensional geometry.

Acknowledgment: Many thanks to Reza Moazzezi for fruitful collaboration during the early stages of this work, to Umesh Vazirani for illuminating discussions, and to Christos-Alexandros Psomas for help with the figures.

References

- [1] Y. Burak, I. Fiete “Accurate path integration in continuous attractor network models of grid cells,” *PLOS Computational Biology* 5, 2, 2009.
- [2] N. Burgess “Grid cells and theta as oscillatory interference: theory and predictions,” *Hippocampus*, 2008.
- [3] I. Fiete, Y. Burak, T. Brookings “What grid cells convey about rat location,” *The Journal of Neuroscience* 28, 6858–6871, 2011.
- [4] L. Giacomo, M. Moser, E. Moser “Computational models of grid cells,” *Neuron* 71, 589–603, 2011.
- [5] T. Hafting, M. Fyhn, S. Molden, M. Moser, and E. I. Moser “Microstructure of a spatial map in the entorhinal cortex,” *Nature*, 436, 801–806, 2005.
- [6] S. M. Kay *Fundamentals of statistical signal processing, volume I: estimation theory*, Prentice Hall, 1993.
- [7] A. Mathis, A. Herz, M. Stemmler “Optimal population codes for space: Grid cells outperform place cells,” *Neural Computation* 24, 2280–2317, 2012.
- [8] B. L. McNaughton, F. P. Battaglia, O. Jensen, E. I. Moser, and M.-B. Moser “Path integration and the neural basis of the ‘cognitive map’ ” *Nature Reviews Neuroscience*, 7, 8, 663–678, 2006.
- [9] C. D. Meyer, *Matrix analysis and applied linear algebra*, SIAM, 2000.
- [10] R. Moazzezi “Grid-Like structure is optimal for path integration,” Arxiv, [1606.01239], June 2016. Also, unpublished manuscript February 2016.
- [11] Z. Navratilova, L. Giacomo, L. Fellous, J. Hasselmo, and B. L. McNaughton “Phase precession and variable spatial scaling in a periodic attractor map model of medial entorhinal grid cells with realistic after-spike dynamics,” *Hippocampus* 22, 4, 2011.
- [12] S. Sreenivasan, I. Fiete “Grid cells generate an analog error-correcting code for singularly precise neural computation,” *Nature Neuroscience* 14, 1330–1337, 2011
- [13] H. Yoon and H. Sompolinsky “The effect of correlations on the Fisher information of population codes,” *NIPS* 1999.
- [14] K. Yoon, S. Lewallen, A. Kinkhabwala, D. Tank, I. Fiete “Grid cells in 1D environments assessed as slices through a 2D lattice,” *Neuron* 89, 1085–1099, 2016.

- [15] K. Yoon, M. Bruce, C. Barry, R. Hayman, N. Burgess, I. Fiete “Specific evidence of low-dimensional continuous attractor dynamics in grid cells,” *Nature Neuroscience* 16, 8, 1077–1084, 2014.