Robustness of Several Versions of the Velocity-Controlled Oscillator Model of Grid Cells

Table of Contents

1 Introduction

In 2005, Edvard and May-Britt Moser found evidence of a kind of cells in rat dorsomedial entorhinal cortex that they called *grid cells* [1]. Similar to *place cells* [2], which have been known for a longer time, grid cells fire (generate action potentials) in specific preferred locations only. For a grid cell, the preferred firing locations form the vertices of an equilateral triangular grid, in contrast to place cells which fire when an animal is near a specific location. Neil Burgess [3] made the observation that an equilateral triangular grid of preferred locations would arise if a grid cell were driven by a superposition of *velocity-sensitive* oscillatory inputs.

First, Burgess' model of grid cell activity is presented, and some nuances of the model are discussed. Next, a method to determine the robustness of a model is provided. Finally, a comparison of the robustness of various models under the provided metric is discussed. The key results of the paper are that peaked oscillations can be much less robust than sinusoidal oscillations, and that the 6 input model is more robust than the 3 input model.

2 The simplest version of the V.C.O. model

To explain why and in which sense an equilateral triangular grid arises as Burgess shows to be the case, consider a cell that receives a baseline oscillatory input with a fixed frequency, and several other oscillatory inputs with frequencies dependent on the current velocity of the rat. We think of the *theta rhythm* (a pronounced 4–11 Hz rhythm in the hippocampus) as the origin of the baseline oscillatory input. Following [3], we assume it to be proportional to

$$\sin(2\pi(t + C)),$$

where t is time measured in theta periods. We take a *velocity-controlled oscillating input* with preferred direction $u \in \mathbb{R}^2$ (u is a vector of unit length) to be proportional to

$$\sin(2\pi(t + \beta u \cdot x(t) + C)), \tag{1}$$

where $\beta > 0$ and C are constants, and x denotes the position of the rat in the plane in which it moves. To motivate (1), note that the time rate of change of the phase $t + \beta u \cdot x(t) + C$ is $1 + \beta u \cdot v(t)$, where $v = dx/dt$ is the velocity of the rat. Thus the phase advances most rapidly when v points in the same direction as u, and least rapidly when v points in the direction opposite to that of u. The constant $\beta > 0$ is a measure of the degree to which the frequency of the oscillation (1) is velocity-sensitive. We always assume, without loss of generality, that

$$x(0) = 0. \tag{2}$$

Note that then all oscillations have phase C at time 0.

The total oscillatory input to the cell is taken to be, up to a constant of proportionality,

$$\epsilon \sin(2\pi(t+C)) + \sum_{i=1}^{n} \sin\left(2\pi\left(t + \beta u_i \cdot x(t) + C\right)\right), \tag{3}$$

where $\epsilon > 0$ is a constant. Here n is the number of velocity-controlled oscillations driving the cell, and the unit vectors $u_i \in \mathbb{R}^2$ are the preferred directions of those oscillations. (We assume $u_i \neq u_j$ for $i \neq j$.) Three assumptions have tacitly been made here:

1. Each velocity-controlled oscillator is associated with the same constant $\beta > 0$, i.e., all the velocity-sensitive oscillations are velocity-sensitive to the same degree. This is why β has no subscript i in (3).

2. All velocity-controlled oscillatory inputs are of equal strength. This is why there are no constants multiplying the sine terms in (3).

3. At time $t = 0$, all velocity-controlled oscillators are at the same phase. This is why C has no subscript i in (3).

We will discuss later what happens when these assumptions are violated.

At a location x, the cell receives oscillatory drive proportional to

$$I(x,t) = \epsilon \sin(2\pi(t+C)) + \sum_{i=1}^{n} \sin\left(2\pi\left(t + \beta u_i \cdot x + C\right)\right). \tag{4}$$

which has amplitude $n + \epsilon$ if and only if all oscillations have the same phase, that is,

$$\forall i \qquad \beta u_i \cdot x + C \equiv C \pmod{\mathbb{Z}}$$

The expressions (3) and (4) are the same, except that we have replaced $x(t)$ in (3) by x in (4), reflecting the assumption that the rat rests, or moves very little, during one theta cycle. This is a reasonable assumption, as one theta cycle is too short of a time span ($\sim$.175 seconds) to move very much. We assume that the cell fires where its oscillatory drive has maximal amplitude.

Changing β will result in changing the locations in which $\beta u_i \cdot x$ is an integer. Thus, this change effectively models changing the scale of the environment in which the rat moves. Consider each active oscillation to have a distinct β, call it β_i. If we increase only a single β_i, then one velocity-controlled oscillation is more velocity-sensitive than the other oscillations, causing the maxima of one of the velocity-sensitive oscillations in (4) to shift.

Thus, for a fixed β such that $\beta u_i \cdot x \in \mathbb{Z}$, and for any C, we can see that the amplitude of (4) as a function of time achieves periodic maxima of $n + \epsilon$, but those maxima will only arise in a regular grid of equilateral triangles when the u_i in (4) are

oriented in directions separated, or spaced, by $\frac{\pi}{3}$ or $\frac{2\pi}{3}$ such that for each unit vector in a given direction, a *family of lines* orthogonal to that vector will be generated with frequency proportional to the frequency of the oscillation whose argument contains the given unit vector. For example, for $n = 1$ where u_1 is the unit vector in the direction of the positive y axis, we see that (4) has maximal amplitude along the crest of a sinusoidal wave, which is an infinite set of lines–perpendicular to the u_i–that are separated pairwise by a uniform distance. A finite subsection of this family of lines looks like the following figure.

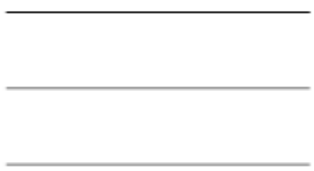

Figure 1: Family of lines

Then in the case of $n = 2$ with spacing $\frac{\pi}{3}$, meaning the family of lines as shown above superimposed with the family generated by a second active oscillation with u_2 being a unit vector in the direction of 60° above u_1–and by symmetry, the argument holds for a unit vector in the direction of 120° above u_1–we see the following section of the associated family of lines:

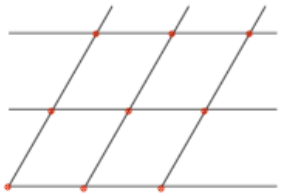

Figure 2: 2 families of lines

Then, for $n = 3$ with spacing $\frac{\pi}{3}$ we see the following section of the associated family of lines:

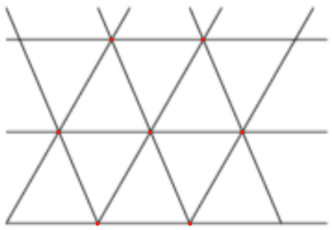

Figure 3: 3 families of lines

For $n = 6$ and equally spaced directions, the family of lines will not generate any distinct intersections from the $n = 3$ case, as each newly added direction will only repeat intersections. If we consider only the intersection points of the family of lines to be locations where a grid cell fires, we can see that an equilateral triangular grid of firing points will arise. Note that as ϵ tends to 0 the maximum of (4) tends to n, but the family of lines is preserved.

3 The phase of the oscillators at time zero

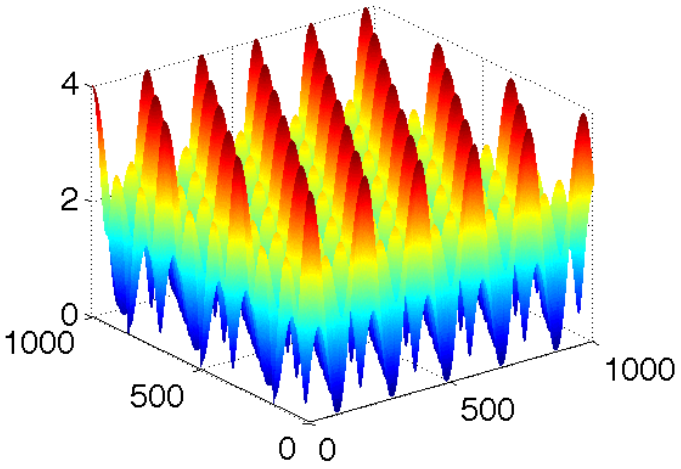

Figure 4: Oscillatory surface with $n = 3$

For a fixed x, the oscillation $I(x, t)$ defined by eq. (4) can at most have amplitude $n + \epsilon$ in special cases–the argument of each oscillation has to be the same, which is sufficient for a maximal amplitude as all oscillations have period 2π. Now, we can see that taking the maximum of eq. (4) and plotting it such that x is a 2-dimensional vector, denoting location in a plane as coordinates, generates a surface in $\mathbb{R}^3$ similar to the one in Fig. 4, where the maxima represent the firing of a grid cell. The maxima will be generated along a family of lines which form a grid, along directions decided by the u_i. Unless all directions and the baseline have the same $C \pmod 1$, the family of lines generated by the oscillatory drive will not each have a line that passes through the origin, and thus no precisely regular grid will arise.

Not having uniform β across the active oscillations will distort the grid. Effectively, increasing β in a single active oscillation will increase the velocity sensitivity of that oscillation. Note that increasing β will not necessarily increase $\beta u_i \cdot x$ in (4) by an integer, so this increase may take the active oscillations out of phase with other oscillations. For the purposes of this section, let $\beta = 6$. Upon adding a random

number between 0 and 1 a single β in a single active oscillation, we see that the oscillatory surface becomes distorted, as the maximum amplitude of the individual altered oscillation now occur in different places than the maximum amplitude of the other oscillations, which have uniform β. In other words, the family of lines generated by the maxima of the active oscillations will not have intersections that include a line from every possible family derived from the oscillations. This alters the surface from looking like it does in Fig. 4, and instead results in something like the following figure.

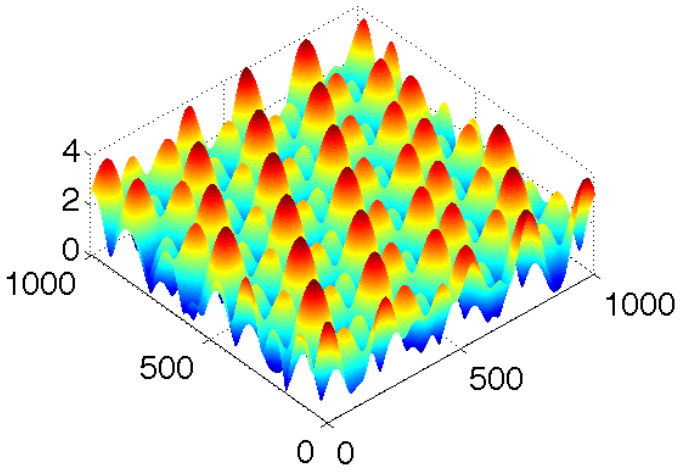

Figure 5: Oscillatory surface with $n = 3$, but one β altered

The maxima of this surface form a regular, equilateral triangular grid, so though the surface is distorted, it does not affect where we consider the grid cell to fire.

Now, if we were to alter each β by some random number in the interval $(0, 1)$, we would see an even more dramatic distribution of the maximal amplitude of the surface. The following figure is such an example, in which the maxima will not form a regular grid.

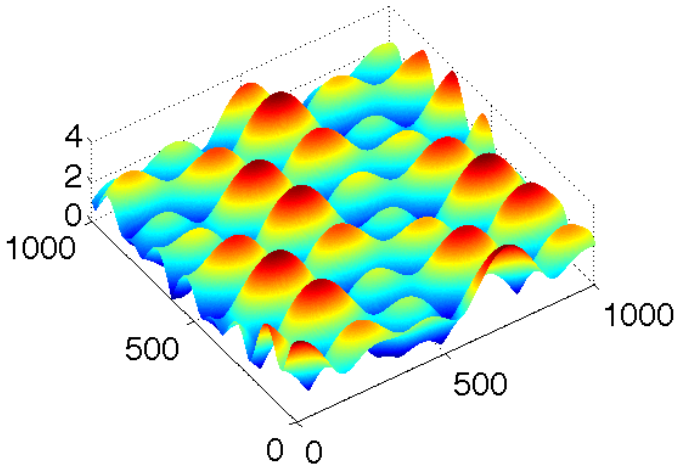

Figure 6: Oscillatory surface with $n = 3$, but all β altered

Upon a direction u_i becoming more velocity sensitive than the other directions, the grid generated by (4) becomes sheared perpendicular to u_i in the altered oscillation proportionally to the change in β.

Interestingly, varying C across the oscillations will not greatly distort the grid that arises from the intersection of the active oscillations' family of lines. Here, let all C initially be 0. Denoting individual, distinct constants as C_i, altering a single C_i by some amount in the interval $(0, 1)$ will take a single active oscillation out of phase with the others, and thus shift the family of lines for that active oscillation whose C_i has shifted. This, for u_i comprised of the vector in the direction of $0°, 120°$, and $240°$, looks like the following figure for a random change in C_3, i.e. the C_i in the active oscillation whose unit vector is $240°$ above the horizontal.

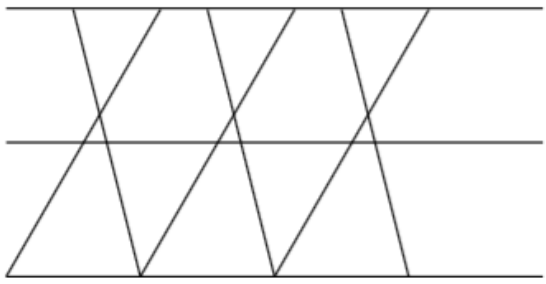

Figure 7: Family of lines with C_3 altered

Something subtle occurs here–Namely, a surface with randomly altered C_i, i.e. each C_i has a random number between 0 and 1 added to it, with a surprisingly large probability does not have a set of maxima that differ very greatly from the case

wherein all C_i are equivalent, modulo 1. Here is a figure to illustrate that occurrence that looks similar to a surface with non-random C_i, and one that looks nothing like it.

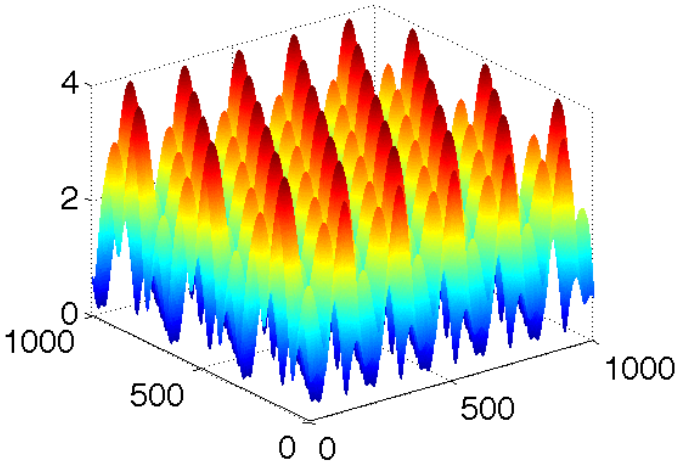

Figure 8: Surface with random C_i, similar to Fig. 4

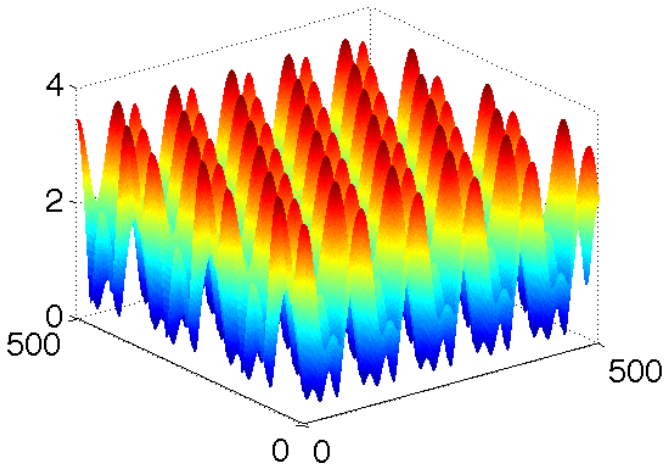

Figure 9: Surface with random C_i, dissimilar to Fig. 4

Now, the reason that randomly choosing C_i in the case of $n = 2$ in (4) doesn't mess up the grid of maxima is that a shift in C_i results in shifting the family of lines generated by u_i in the direction orthogonal to the u_i by the amount C_i has been shifted. For two directions, it is impossible to destroy an already equilateral grid by shifting one or both of the families of lines that generate it in a way that does not affect the angle that separates them–it will only result in the maxima shifting location, not that each is the vertex of some equilateral triangle.

Then, for $n = 3$ with u_i spaced evenly by 120°, we must be more careful. Let a *maxima neighborhood* mean the circular area whose center is a maxima of an oscillatory surface, i.e. an intersection point in some family of lines. This looks like the following:

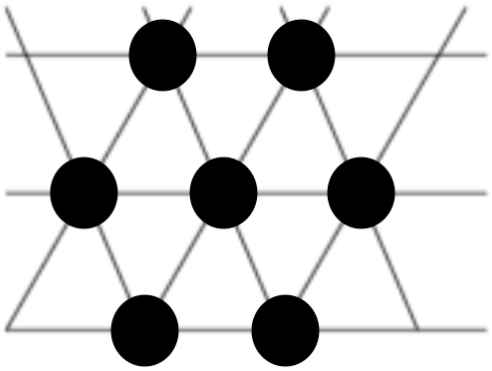

Figure 10: Family of lines with maxima neighborhoods

These maxima neighborhoods correspond to considering an oscillatory surface above some intersecting plane in $\mathbb{R}^3$. That is, for maxima neighborhoods of radius $\delta \leq \frac{1}{2}$, picking a δ corresponds to picking a height for a plane in $\mathbb{R}^3$ such that the plane intersects an oscillatory surface in circles of radius δ. Note that this height can at be at minimum 0, and at most $n + \epsilon$. This is illustrated by the following figure.

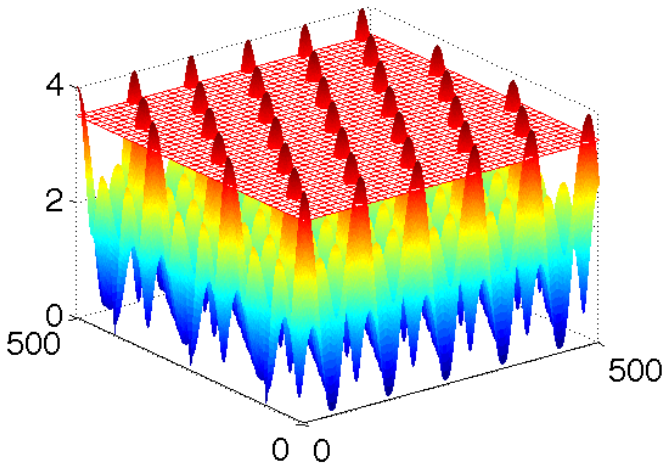

Figure 11: Oscillatory surface and intersecting plane at $z = 3.5$

Then, if each active oscillation is given a random C_i, the first two active oscillations do not require a careful choice of C_i will result in a regular equilateral grid,

as described above. Then the probability that a family of lines generated from the active oscillation with u_3 passes through the maxima neighborhoods is 2δ, as the only requirement for the lines is to avoid the band of width $1 - 2\delta$ between maxima neighborhoods that resulted from the previous families of lines. Note that adding new families of lines result in an additional factor of 2δ when calculating the probability that a set of families of lines result in intersections in some maxima neighborhood. Thus, the probability for n active oscillations with randomly chosen C_i to meet that criterion is $(2\delta)^{n-2}$. A computed example for $n = 3$ is, for $z = 3.0$, $\delta \sim .15$.

Similarly, assigning different weights to the individual active oscillations will greatly distort the surface, but not the grid of maximal points–a regular triangulation is preserved, though the maximum of (4) will be shifted proportional the assignment of weights, i.e. if one oscillation is weighted double, the maximum of (4) will increase by 1. To illustrate this, here is a figure wherein each active oscillation is given a randomly assigned weight in the interval (1,2).

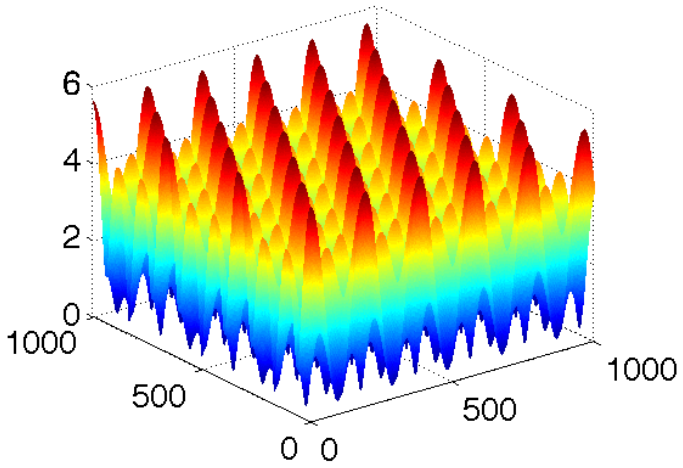

Figure 12: Surface with randomly assigned oscillation weights

4 The possible direction choices

In this section, we return to the assumption that C and β are uniform across all oscillations. The only scenarios in which a regular triangular grid will arise is with 2 or 3 directions, discounting sets of directions where opposites, or additive inverses, such as u and $-u$ occur. However, any other number of distinct directions up to additive inverse will not yield a faithful representation of a grid cell's tiling of the plane. For example, it is clear that a single direction will not yield any intersections in the family of lines generated by the direction set. For one direction along the

horizontal and one direction above the horizontal by 60°, i.e.

$$u_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad u_2 = \begin{bmatrix} \frac{1}{2} \\ \frac{\sqrt{3}}{2} \end{bmatrix}$$

we have the following subsection of the family of lines:

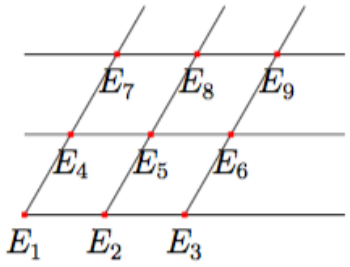

Figure 13: 2 families of lines with intersections

From this, we can see that the intersections will generate a regular grid of equilateral triangles. By reflection symmetry, the same is true for a two directional set separated by 120°. Now, for three directions, we can see that the only possible family of lines that generate a regular grid of equilateral triangles will be one of the two-directional family of lines coupled with a new family of lines such that for all i, E_i is preserved as an intersection. Without loss of generality, consider the two-directional case as illustrated above. Then the only possibility for the third direction must pass through all intersection points in order to preserve their status as intersections, so there are only two options–the family of lines that is generated by including the set of lines in the direction of 120° above the horizontal, and the one 30° above the horizontal.

The reason that 30° will not generate the grid is that, although it results in a set of maxima that lie along the family of lines in the 60° direction and thereby one set of sides of the triangles will be of length k, the other two sides of triangles resulting from connecting maximal vertices of the surface in $\mathbb{R}^3$ will yield a grid of isosceles triangles as follows.

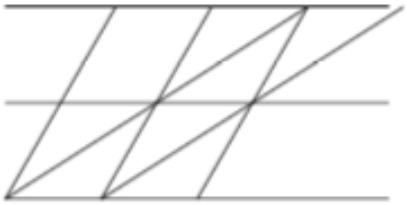

Figure 14: $0°, 30°$, and $60°$ families of lines

Thus, our only remaining option is $120°$ above the horizontal to generate an equilateral triangulation with three directions. Then, any other additional direction must not create a different set of intersecting points among the existing families of lines. Any other line will not pass through all E_i, thus the triangulation will be imperfect. Here, a family of vertical lines will pass through all the maximal points of the surface, but will again yield an isosceles triangulation. Thus, the only viable addition directions are negations of preexisting directions. Note that for 6 distinct directions, the only 6 that work by this argument are spaced evenly by $60°$.

5 Computational Method

The manner in which this paper assumes a computational representation of a model is as follows: first, we pick an appropriate n, and some dimensions in which a grid may be generated. Then, we generate the surface according to our model–for every point in the dimensions we've selected to hold our grid, we generate a surface point as the maximum of (4) at that point in $\mathbb{R}^3$.

We proceed to smooth the surface by comparing each point with neighbors, and averaging the position to have less conflict between points over claims to being a maximum or a minimum. This allows us to define a threshold z value, where points above that threshold represent candidate points where the grid cell may fire. We collect the set of points above that threshold, and continue to formatting them as a possible regular, equilateral triangulation.

6 Measuring the quality of the equilateral triangulation

It is hard to imagine that directions occur precisely in some equally spaced manner in nature so it is important to ask what happens when we perturb the directions by some random noise to approximate what might happen in nature. The question of how "regular" a grid is arises when computing a grid from the input of n directions,

which is to say how close to a regular equilateral triangular grid the output is.

Here, a discussion of practice is merited. When generating a surface in $\mathbb{R}^3$ from a model, it is necessary to approximate the continuous surface with a discrete mesh of points which arises by computing a function such as (4) pointwise. This computation is achieved by iterating over an $n \times n$ set of points to define the function surface, where the tradeoff for low n is faster computation but a *coarser* mesh of points with which we may experiment to explore the properties of a model. This will accumulate numerical error.

The proposed measure of robustness for a given model is given by taking the set of points above the threshold, mentioned in the previous section, which represent the maximal vertices of the oscillatory surface. Then, we perform a Delaunay triangulation [4], which gives us a set of edges of triangles. Next, we remove border triangles until we only have proper candidate triangles, and not triangles that include edges as a side-effect of code, i.e. connecting vertices that are not adjacent.

We compile a measure of how equilateral the triangles are by taking each edge of each triangle of the form:

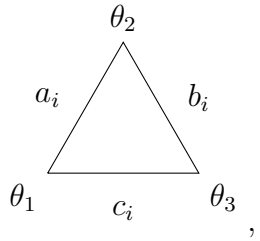

,

namely, this is a_i, b_i, and c_i . Then, where

$$v_i = \begin{bmatrix} a_i \\ b_i \\ c_i \end{bmatrix},$$

we take the 2-norm of the difference between each edge of a triangle and the mean edge length of the triangle divided by the mean edge length, given by $\psi = \frac{||v_i - \overline{v_i}||}{\overline{v_i}}$. Then, we examine the angles of a given triangle by taking the three individual angles θ_1, θ_2, and θ_3 which gives us

$$R = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix},$$

so we take the 2-norm of the difference between each angle of a triangle and the mean angle divided by the value of mean angle, given by $\rho = \frac{||R - \overline{R}||}{\overline{R}}$, which tells us how close the angles of a given triangle are to 60 degrees. We average both of these measures over all candidate triangles in a model's output grid.

7 Perturbing the directions

Now, when we perturb a single direction we can compare the resulting grid to the unperturbed case. Especially relevant here is the coarseness mesh with which we approximate our input model. For coarse mesh, an unperturbed model will have nonzero ψ and ρ due to discretization error. This error can be reduced by increasing the dimensional fineness of the mesh, and smoothing the result by comparing each point in the mesh to adjacent points and only keeping the maximal point in that comparison. We proceed with the three-directional case where the u_i are direction vectors in the direction of the horizontal, 120° above the horizontal, and 120° below the horizontal.

Then, we may compare these arguments to the set of arguments subject to perturbation by 1° in a single direction. Principally, our interest is in seeing which model is most robust when perturbed in the leading order, analogous to perturbing a partial derivative of a function of multiple variables, as it can be considered without loss of generality to be for any individual direction. We may then measure the quality of the perturbed triangulation relative to the unperturbed case with $\rho = 0$ and $\psi = 0$, which is otherwise identical, to provide a method of comparison between models. Hereafter, $\beta = 6$ for all models for the sake of consistency. Note that some numerical error may arise due to integration techniques incorporated in generating the surface, depending on the choice of model. This is visible in the unperturbed case, which is ideally 0. The results from using a 1000×1000 mesh and perturbing one direction by $0°, 1°$, and 5° can be found in Table 1.

We can see that perturbing the three-directional case in a direction of leading order has a far more drastic effect than that of the six-directional case. This is because the position of the maxima of the generated surface in the 3-directional case are determined by fewer terms in the sum in (4) than the 6-directional case, and thus are a greater distance from the "true" maxima that an unperturbed model will yield.

8 Non-sinusoidal oscillations

We can replace the use of explicit sine functions by other oscillatory functions. The motivation for this comes from evidence [5] that the signal obtained from a neuron is periodic, but experiences more severe peaks than a regular sinusoidal function. An example is the following:

$$s(t) = \frac{e^{\alpha \sin^2(\pi t)}}{\int_0^1 e^{\alpha \sin^2(\pi u)} du} - 1.$$

Now, as α tends to 0, $s(t)$ becomes sinusoidal. To demonstrate this, consider that by the local linear approximation $e^x - 1 \sim x$, we have that

$$s(t) \sim \frac{1 + \alpha \sin^2(\pi t)}{\int_0^1 1 + \alpha \sin^2(\pi u) du} - 1.$$

Now, consider the denominator:

$$\int_0^1 (1 + \alpha \sin^2(\pi u)) du = 1 + \alpha \int_0^1 \sin^2(\pi u) du = 1 + \frac{\alpha}{2},$$

so we have that $s(t) \sim \frac{1+\alpha \sin^2(\pi t)}{1+\frac{\alpha}{2}} - 1$. Now, this is equivalent to the following:

$$\frac{\alpha \sin^2(\pi t) - \frac{1}{2}}{1 + \frac{\alpha}{2}} = \alpha \frac{\frac{1}{2} - \frac{1}{2}\cos(2\pi t) - \frac{1}{2}}{1 + \frac{\alpha}{2}},$$

as we have that $\cos(2\pi t) = \cos^2(\pi t) - \sin^2(\pi t) = 1 - 2\sin^2(\pi t)$, so $\sin^2(\pi t) = \frac{1-\cos(2\pi t)}{2}$. Continuing, we have that

$$\alpha \frac{\frac{1}{2} - \frac{1}{2}\cos(2\pi t) - \frac{1}{2}}{1 + \frac{\alpha}{2}} \sim -\frac{\alpha}{2}\cos(2\pi t).$$

Then as $\cos(t) = -\sin(t - \frac{\pi}{2})$, we have that $s(t) \sim \frac{\alpha}{2}\sin(2\pi(t - \frac{\pi}{4}))$ and thus is sinusoidal for low α. As α increases, s then becomes increasingly "peaked." That is, as $\alpha \to \infty, s(t) \to \sum_{k \in \mathbb{Z}} \delta(t - kT)$, where δ denotes the *Dirac delta function* and T denotes the period of $s(t)$.

Here, error may accumulate due to the choice of method in evaluating the denominator. The numbers in this paper use Romberg Integration to evaluate the denominator, an algorithm which approximates integrals to within some tolerance, which is 10^{-14} in this paper.

Now, the maxima specified by the amplitude of $s(t)$ will map to the triangular grid much like those of the purely sinusoidal model, but will result in a more peaked output grid for significantly large α values. The following figures show examples of $s(t)$ with various values of α.

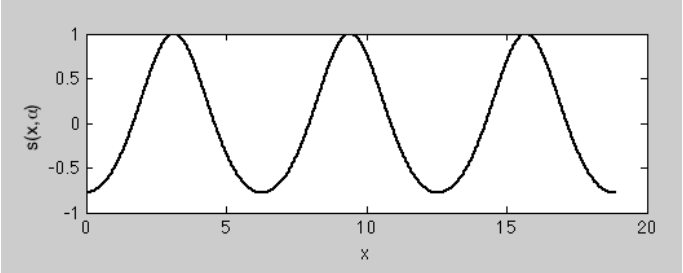

Figure 15: s(x) with $\alpha = 1$

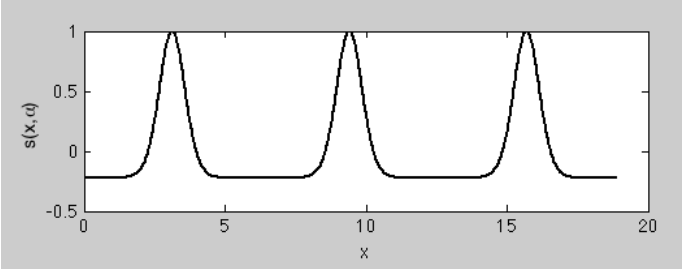

Figure 16: s(x) with $\alpha = 10$

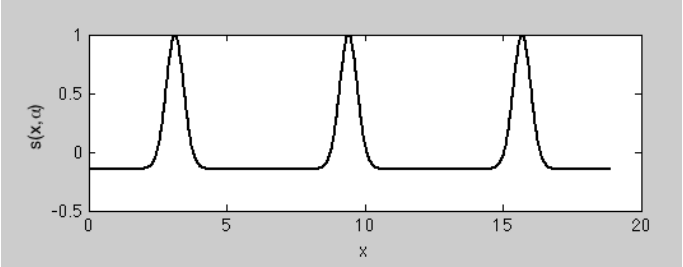

Figure 17: s(x) with $\alpha = 20$

Now, we can subject this model to the same tests as the sinusoidal case. As α tends to 0, the maxima of the amplitude of $s(t)$ become negligibly different from those of sinusoidal model. Thus, for small α, the oscillatory surface is extremely similar to that of the sinusoidal model. To illustrate this, here is a figure of an oscillatory surface for $\alpha = 0.00001$, and $\alpha = 20$. Note that because this surface has very peaked "hills," the maxima neighborhoods will be very narrow circles, thus random C_i will less likely to result in a nearly-equilateral grid than a sinusoidal model.

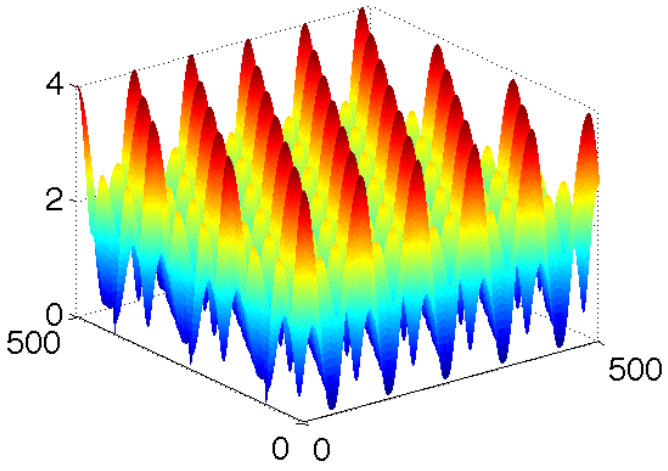

Figure 18: $s(t)$ with $\alpha = 0.00001$

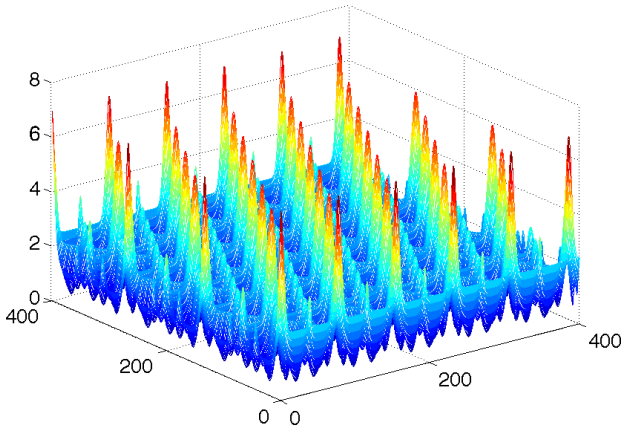

Figure 19: $s(t)$ with $\alpha = 20$

For $\alpha = 20$, the results can be found in Table 1.

9 Results

Table 1: ρ and ψ values for perturbing directions

3-directional sinusoidal	ρ	ψ
unperturbed	0.0095	0.0057
1°-perturbed	0.0408	0.247
5°-perturbed	0.0511	0.0309
3-directional peaked		
unperturbed	0.0095	0.0057
1°-perturbed	0.0392	0.0237
5°-perturbed	0.1279	0.0773
6-directional sinusoidal		
unperturbed	0.0095	0.0057
1°-perturbed	0.0123	0.0074
5°-perturbed	0.0570	0.0344
6-directional peaked		
unperturbed	0.0115	0.0069
1°-perturbed	0.0123	0.0074
5°-perturbed	0.0554	0.0334

Ultimately, the difference between resulting grids from the purely sinusoidal case and $s(t)$ for low α is extremely small, while high α yields a more robust model in some cases. The 3-directional peaked model is more robust under 1° perturbation, but falls apart for more severe perturbation. Interestingly, the 6-directional model is no more

robust under 1° perturbation, but is more robust for more severe perturbation. Following this section are some pictures of the output grids, labelled with the model that generated them (red triangles are candidate triangles for measurement). Note that the grids associated with 5° perturbation in peaked models seem less complete than their sinusoidal counterparts. This is because the oscillatory surface for the peaked models have maxima that "fall off" in some sense. That is, while some of the maxima from the peaked surface are grouped at a similar height to their sinusoidal maxima counterparts, many are at a significantly lower height. This is best illustrated in the following side views of oscillatory surfaces.

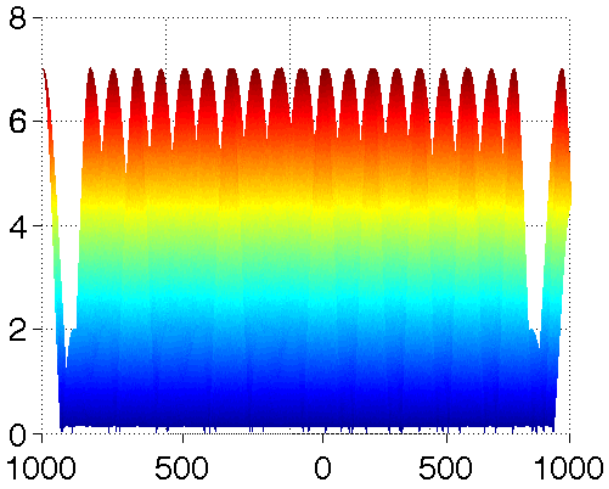

Figure 20: Unperturbed side view of peaked oscillatory surface

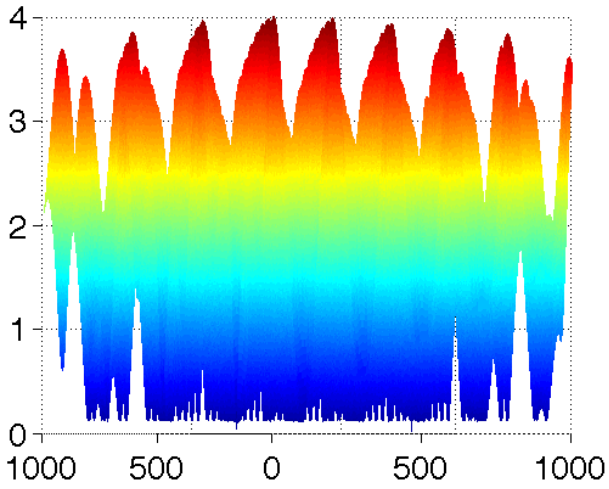

Figure 21: 3-directional peaked, perturbed 5°

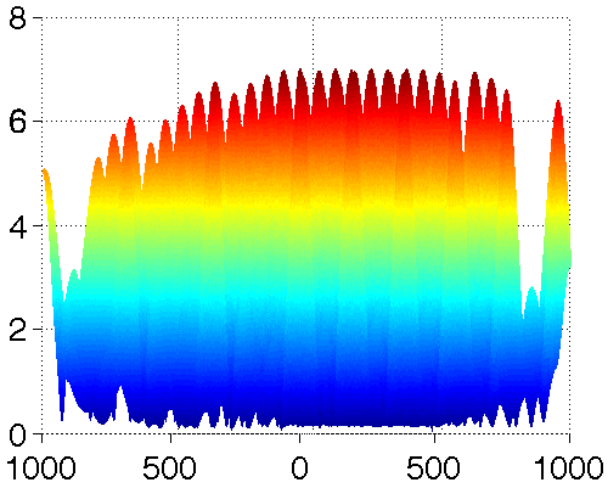

Figure 22: 6-directional peaked, perturbed 5°

Here are some pictures of the grids.

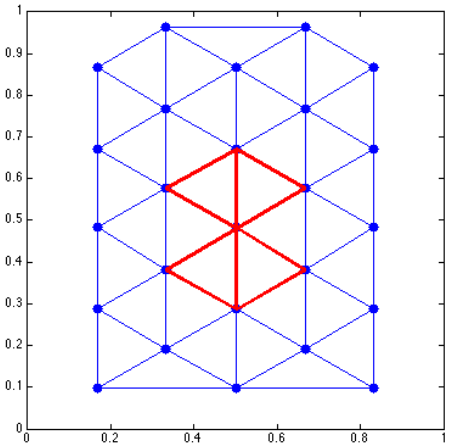

Figure 23: 3-directional unperturbed sinusoidal

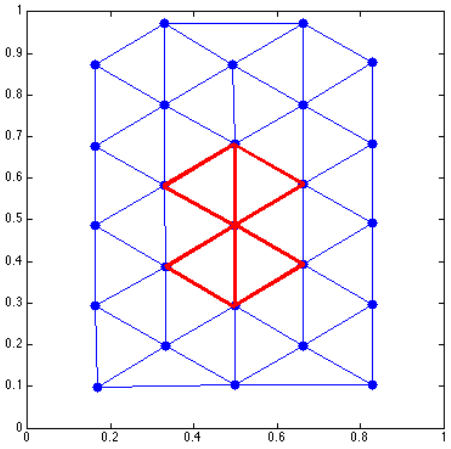

Figure 24: 3-directional 1°-perturbed sinusoidal

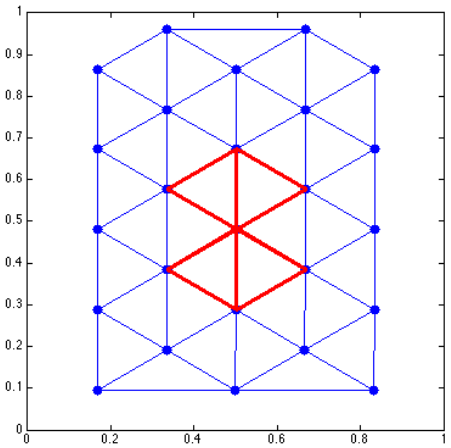

Figure 25: 3-directional unperturbed peaked

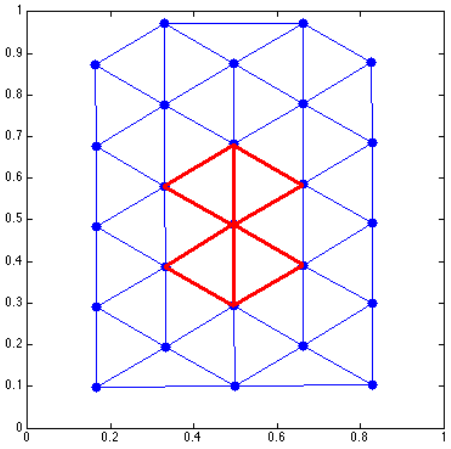

Figure 26: 3-directional 1°-perturbed peaked

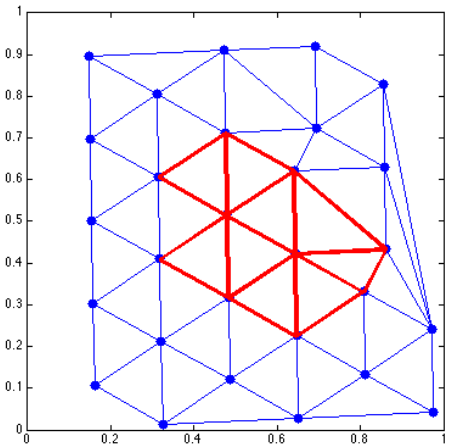

Figure 27: 3-directional 5°-perturbed sinusoidal

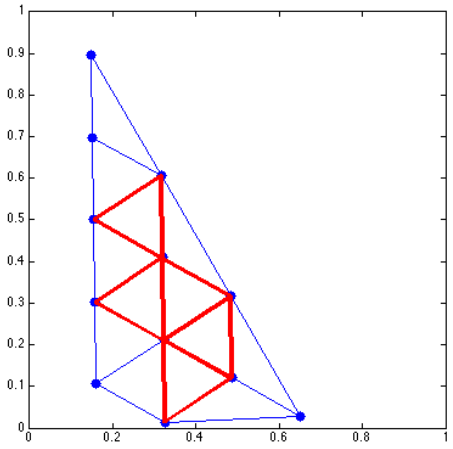

Figure 28: 3-directional 5°-perturbed peaked

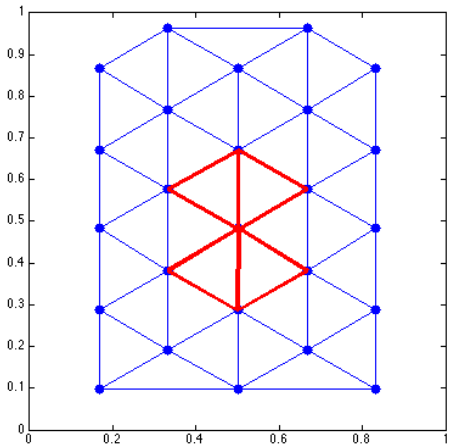

Figure 29: 6-directional unperturbed sinusoidal

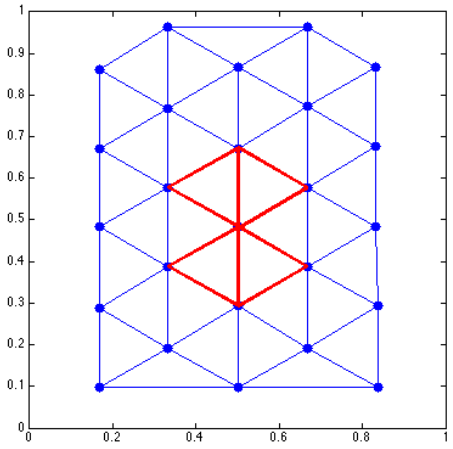

Figure 30: 6-directional 1°-perturbed sinusoidal

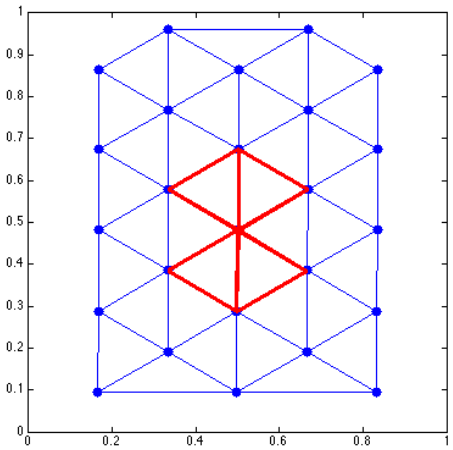

Figure 31: 6-directional unperturbed peaked

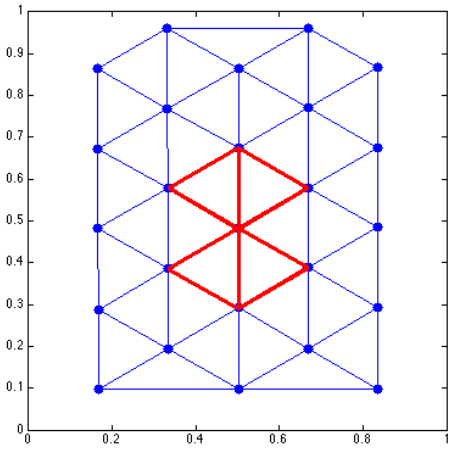

Figure 32: 6-directional 1°-perturbed peaked

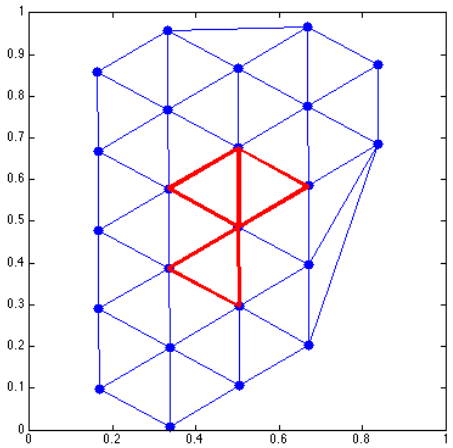

Figure 33: 6-directional 5°-perturbed sinusoidal

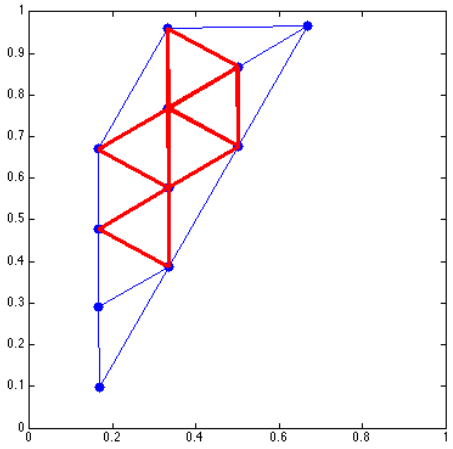

Figure 34: 6-directional 5°-perturbed peaked

10 Discussion

It seems that it's difficult to do better than a sinusoidal model, especially for a low number of directions. It's interesting that it's possible to keep an equilateral grid with random C_i, as it leads me to believe it's not necessary to reset phase with the hippocampal theta rhythm to obtain a serviceable grid. It also isn't too surprising that the 6-directional model is more robust under small perturbations than the 3-directional model, but the 3-directional model's superiority in the larger perturbation is somewhat surprising.

Randomly altering velocity sensitivity of inputs seems the most effective way to ruin a grid, which makes sense because it corresponds to altering perception of spatial environment scale in different directions, which would make it very hard to reason about that environment.

References

[1] T. Hafting, M. Fyhn, S. Molden, M.-B. Moser, and E. I. Moser.
Microstructure of a spatial map in the entorhinal cortex.
Nature, 436:801–806, 2005.

[2] John O'Keefe.
Place units in the hippocampus of the freely moving rat.
Experimental Neurology, 51:78–109, 1976.

[3] N. Burgess.
Grid Cells and Theta as Oscillatory Interference: Theory and Predictions.
Hippocampus, 18:1158–1172, 2008.

[4] Steven Fortune.
Handbook of discrete and computational geometry.
chapter Voronoi Diagrams and Delaunay Triangulations, pages 377–388. CRC Press, Inc., Boca Raton, FL, USA, 1997.

[5] György Buzsáki.
Theta oscillations in the hippocampus.
Neuron, 33(3):325–340, January 2002.