Toward a spiking model of the Hippocampus and Entorhinal Cortex

Emily Mankin

Abstract

A framework for a spiking model of Entorhinal Cortex, and the Dentate Gyrus and CA3 subfields of the hippocampus is developed. Taking the input to the Entorhinal Cortex as determined by influences outside the model, I define grid cells, head direction cells, and border cells in a simulated experiment in which a rat performs random foraging in a box. Cells are modeled with Izhikevich-style neurons and synapses are modified by spike-timing dependent plasticity. Resulting firing patterns in the dentate gyrus and CA3 are considered, and several questions are raised that this model – when realized fully – may be able to help answer.

Background

The entorhinal-hippocampal complex has been extensively studied in rodents. In one class of experiments, neural activity is monitored (typically through extracellular recordings) while a rodent performs a random foraging task, in which it runs around an environment (usually a box around 1 square meter) searching for food. In these tasks, the principal cells in the hippocampus demonstrate spatially-specific firing patterns within a single environment [8]. Typically a pyramidal cell in the CA1 or CA3 subregion has a single "place field," a connected area in which the maximum firing rate can be tens of Hz, compared to the area outside where baseline spiking is very low. A granule cell in the Dentate Gyrus (DG) may have several place fields within a single environment [6]. Many cells in layer II of the medial Entorhinal Cortex (mEC) also show spatially-modulated firing patterns. Approximately half of these cells are "head direction cells," which fire preferentially when the rodent's head is pointed in a specific direction. Another 35% are "grid cells," which show bumps of activity whose centers make the vertices of a triangular tesselation of space. Finally, approximately 10% are "border cells," which fire when the rodent is near one wall of the environment [10].

The primary input to the hippocampus comes via the perforant path from the entorhinal cortex. These collaterals synapse on the granule cells of the dentate gyrus and then continue on and synapse on the pyramidal cells of CA3. The dentate also sends its own collaterals, the mossy fibers, to CA3. CA3 also has its own recurrent connections, as well as sending axon collaterals to the CA1 region in the Schaffer collateral pathway. CA1 finally sends axons back to the entorhinal cortex, completing the "tri-synaptic loop" [11].

The hippocampus is known to be required for storing new spatial memories. The dentate gyrus has been proposed as a site for pattern separation [9]. The idea is that, in order to have distinct memories of similar but distinct events, there must be some way for the hippocampus to orthogonalize strongly overlapping input. Because the dentate gyrus has approximately five times as many excitatory cells as the entorhinal cortex, this is a place where there is room for the inputs to be orthogonalized. In fact, experimental studies have backed up this idea: A sparse set of dentate cells fires in each environment. The subset of active cells stays constant from one environment to the next, but the number, size, and location of place fields changes within each active cell. This is different from pattern separation in CA3, in which different cell populations are recruited in different environments [6].

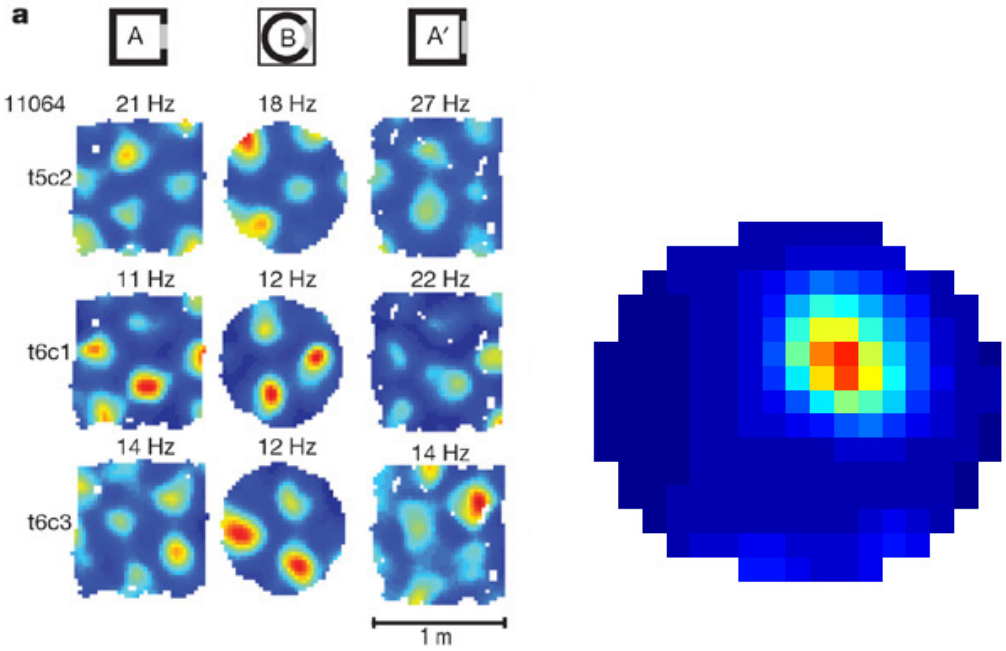

Figure 1: Sample grid cells taken from [4]. Notice that their activity bumps tesselate the plane. The cells are sampled in environment A and then in B, and then in A again. Note that all cells have the same grid orientation as each other in each environment, but not the same orientation across environments. b. A sample CA3 place field [7].

Objectives

While a fair amount of theoretical work on the dentate gyrus (and EC/Hippocampal complex as a whole) has been done, to date no spiking model of the DG has been published. My goal for this project was to build a framework for a spiking model of the dentate gyrus and CA3 subregions of the hippocampus that could be used to explore topics such as

- The effects of plasticity on pattern separation
- The extent to which dentate input can be degraded and pattern separation in CA3 can be maintained
- The role of neurogenesis in hippocampal function

Originally, I followed the example of Cerasti and Treves and modeled dentate input as "given." That is, for each environment and each cell, I used a random processes to determine the number of activity bumps the cell would have in the environment and the location, size, and peak firing rate for each. I then created functions that would determine the appropriate amount of "current" that cell should receive in that location, in order for it to fire at the pre-determined rate. So, while the idealized behavior of the neuron was pre-determined for an environment, the actual spiking behavior was modeled with differential-equation-based Izhikevich-style neurons. From there, I modeled synaptic connections between dentate and CA3 and recurrent CA3-CA3 connections and used the spiking of pre-synaptic neurons to determine the current received by post-synaptic neurons.

While this produced interesting preliminary results, further reading convinced me that it was important to model entorhinal input to the dentate and CA3 as well. This will allow exploration of how the pattern-separating activity in the DG comes about, as well as issues such as the individual contributions of grid cells, head direction cells, and border cells. This was achieved by using the same approach for EC cells as previously used for DG cells, and then modeling current into DG and CA3 with synaptic connections.

Methods

Exploring the Environment

In typical live experiments, the rodent is placed in an environment and allowed to run around. Its location is recorded along with spiking activity, so that the two may be correlated. The running of the rat could be modeled with a modified version of a random walk. To save some computational and programming effort, for this preliminary model, I used trajectory data from a live experimental session. Unfortunately, Matlab's memory restrictions have made it such that so far I have only been able to model up to two minutes of exploration time. So the entire environmental space is not completely sampled in the results below.

Spiking Neurons

In the rat brain, there are approximately 300,000 CA3 pyramidal cells, 1,000,000 dentate gyrus cells, and 200,000 principal cells in the entorhinal cortex. These are reduced by a factor of 1000 in my model to 300 CA3 cells, 1000 DG cells, and 200 EC cells. Approximately 2-5% of DG cells are active at a time [3], so in the first instantiation of the model (in which DG input was defined by the environment), DG was reduced to 30 active cells.

Every cell was modeled using the framework developed by Izhikevich [5], a 2-dimensional system of differential equations with variables: v, the voltage, and u, a recovery variable. The system is given as

$$\begin{aligned} \frac{dv}{dt} &= \frac{k}{C}(v - v_r)(v - v_t) - u + I \\ \frac{du}{dt} &= a(b(v - v_r) - u) \end{aligned}$$

where v_r is the resting potential and v_t is the threshold. I is the input current. At any step when v exceeds a given value, v_{peak}, the cell is said to have fired, the voltage is reset to v_{reset}, typically somewhere between v_r and v_t. Although there are many parameters that should be adjusted for different neuron types, as a starting place I simply used the parameters given in Izhikevich's book for regular spiking neurons.

Environment-Defined Input

Dentate Gyrus

In the first model version, idealized dentate firing was defined for each active cell in each environment. The number of activity bumps a cell had was modeled by a Poisson process with $\lambda = 1.7$ [3]. Each bump was modeled as proportional to a two-dimensional Gaussian distribution. The center of each bump (mean of the Gaussian) was chosen uniformly from the area inside the environment. The size of the field (standard deviation of Gaussian) and peak firing rate were chosen independently from normal distributions. The idealized firing rate could be modeled, thus, as a function of x and y.

Simple experiments showed that, with the parameters as defined for the spiking neurons, there was an almost linear relationship between input current and firing rate. Taking the inverse of this function allowed me to define the amount of input current at (x, y) that would (if the rat stayed in the same location) yield the idealized firing rate at (x, y).

Entorhinal Cortex

In the expanded model that included the Entorhinal Cortex, all EC cells were defined to have either spatially-specific idealized firing rates (grid cells and border cells) or movement-angle-specific idealized firing rates (head direction cells). The method of converting from idealized firing rate to input current was as described for dentate gyrus cells in the first model, so here we describe only the method for defining the idealized rate maps.

Grid Cells

The functions

$$\begin{aligned} f_1(x,y) &= \frac{\beta}{2}(1+\sin\left(\frac{1}{radius}(x-\tan(\frac{\pi}{3})y\right) \\ f_2(x,y) &= \frac{\beta}{2}(1+\sin\left(\frac{1}{radius}(x-\tan(\frac{\pi}{6})y\right) \\ f_3(x,y) &= \frac{\beta}{2}(1+\sin\left(\frac{1}{radius}(x+\tan(\frac{\pi}{6})y\right) \end{aligned}$$

when multiplied together form a pattern of activity akin to typical grid cells: bumps of activity whose centers form the vertices of equilateral triangles that tesselate the plane (loosely based on ideas in [2]). The lines that form the base of these triangles is at some angle relative to the environment that is constant across cells. Thus for each environment, the angle, θ, is chosen from a uniform distribution between 0 and $\pi/3$. Then for each cell, a radius size is chosen. Because the grids are known to expand in size as one moves from dorsal to ventral hippocampus [1], I used deterministic radii, spaced logarithmically between a radius of 8 cm and 100 cm. Finally, some amount of offset for the center, (x_o, y_o) is chosen. This leads to the final idealized firing rate funtion:

$$f(x,y) = \frac{1}{4\beta^2}\prod_{i=1}^{3} f_i\left(\cos(\theta)(x-x_o) - \sin(\theta)(y-y_o), \sin(\theta)(x-x_o) + \cos(\theta)(y-y_o)\right)$$

Border Cells

Typically, border cells fire only along one portion of the border of the environment [10]. In the case of a square environment, this means that most cells fire along the entire length of a single wall. I model this by randomly choosing one side for each cell and defining a gradient that takes a maximum value along the side and extendings perpendicularly away. For example, if the environment has vertices $(\pm w, \pm w)$, β is the maximum firing rate, and σ is the standard deviation that defines the speed of decay from the wall, the idealized firing rate of a "west wall" cell is given by:

$$f(x,y) = \frac{\beta}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{(x-w)^2}{2\sigma^2}\right) I_{\{x\leq w\}}(x)$$

Head Direction Cells

Head direction cells fire preferentially when the head of the rat is facing in a specific direction. Because head direction data was not available for the trajectories I used, I used the angle of the direction of movement as a proxy for head direction. For each head direction cell, I defined a preferred angle between $-\pi$ and π, and used a scaled one-dimensional Gaussian centered at the preferred angle, and cut off outside 3 standard deviations. Finally, I took any piece that was non-zero on $(-2\pi, -\pi)$ and added it to the function between 0 and π, and similarly, the non-zero portion from $(\pi, 2\pi)$ I added to the portion between $-\pi$ and 0. This allowed the Gaussian to cross smoothly over the branch cut at $-\pi$. The final result is

$$f(\theta) = f_1(\theta) + f_1(\theta - 2\pi) + f_1(\theta + 2\pi)$$

where

$$f_1(\theta) = \frac{\beta}{\sqrt{2\pi\sigma^2}}\exp\left(-\frac{(\theta-\theta_0)^2}{2\sigma^2}\right) I_{\{|\theta-\theta_0|<3\sigma^2\}}(\theta)$$

Synapses

For each connected area, the probability of there existing synapse between a cell in area 1 and a cell in area 2 was defined. The simplifying assumption was made (but could easily be relaxed) that there is at most one synapse between each pair of cells. The existence or non-existence of a synapse was defined for each pair of potentially-connected neurons with a Bernoulli distribution. Then the initial weight of each synapse was chosen from a uniform distribution between 0 and 1. Because synapses between some pairs of areas are stronger than others, a weight-impact was defined for each pair of areas and multiplied by the synaptic weight, which always remained between 0 and 1.

Synaptic Currents

The input current to areas that receive synaptic currents instead of environmentally defined current is defined as the sum of the excitatory inputs (as modulated by synaptic weights) minus inhibition times the driving force. More precisely, at time t, the current into cell j is given by

$$I_j(t) = (v - v_{rest}) \left(\left(\sum_i \eta_i \sum_{k_i} w_{k_i j} f_{k_i}(t) \right) - \eta_{inhibition} g(t) \right)$$

where i ranges over areas that synapse onto the area of interest (e.g. for CA3, i will range over EC, DG, and CA3) and η_i is the weight-impact of that area onto the area of interest. k_i ranges over the cells in area i, and $w_{k_i j}$ is the weight of the synapse between cells k_i and j. $f_{k_i}(t) = 1$ if cell k_i has fired recently (where the definition of recent may depend on i) and is 0 otherwise. $\eta_{inhibition}$ is the weight-impact of inhibition onto the area of interest, and $g(t)$ is the amount of inhibition the cell receives at time t, currently chosen from a normal distribution.

Plasticity

For some runs of the model, I allowed learning through spike-timing dependent plasticity. Every time a cell fired, cells that had fired recently and were connected to that cell had their synapses updated. Synapses in which the pre-synaptic cell fired before the post-synaptic cell were strengthened, while synapses in which the post-synaptic cell fired before the pre-synaptic were weakened. In order to keep the synaptic weights of the cells between 0 and 1, a sigmoidal function was used to translate between synapse-weight and an incrementable step number. Specifically,

$$g(w) = \frac{2w - 1}{2\sqrt{\frac{\lambda}{(w - w^2)}}}$$

That is, given a synaptic weight, w, we may find $s = g(w)$. The step, s, is then increased by 1 for synaptic strengthening or decreased by 1 for weakening. To find the new weight, we apply $w_{new} = g^{-1}(s_{new})$.

Results

Actual firing reflects idealized firing

The first test of whether the model works appropriately is to make sure that the input currents that are defined by the environment actually translate into spiking that has the same overall properties. This, indeed holds true, as may be seen in Figure 2.

In Model version 1, Learning has little effect

I wanted to explore the effects of spike-timing dependent plasticity on the firing of CA3 cells. In the first model version, I ran simulations both with and without learning. The raster plots of the first 100 CA3 and all active dentate gyrus cells are displayed in Figure 3,

In Model version 2, DG has very sparse activity

After adding the Entorhinal Cortex to the model, I returned the number of Dentate cells to 1000, and allowed them to "self-select" for activity level. In environment 1, of the 1000 cells, 68 cells had activity. In environment 2, 63 were active. This is higher than the 2-5% activity estimate of Cerasti and Treves [3], but not by too much. Of greater concern, the place fields for all 68 cells in environment 1 were largely overlapping. This was mostly true in environment 2, as well, though to a lesser degree. I expect that this is due to the fact that with my current strategy for defining grid cells, I am not able to tile the environment evenly. That is, the average activity of all the grid cells is not spatially consistent, but has hot spots and cold spots. Since it is input from the entorhinal cortex that drives dentate gyrus cells, it is likely that cells that receive enough inputs to cross a threshold and fire do so only where the average grid activity is hot. I originally thought the high amplitude of the spatial modulation of average activity was due to the small number of cells I was using and

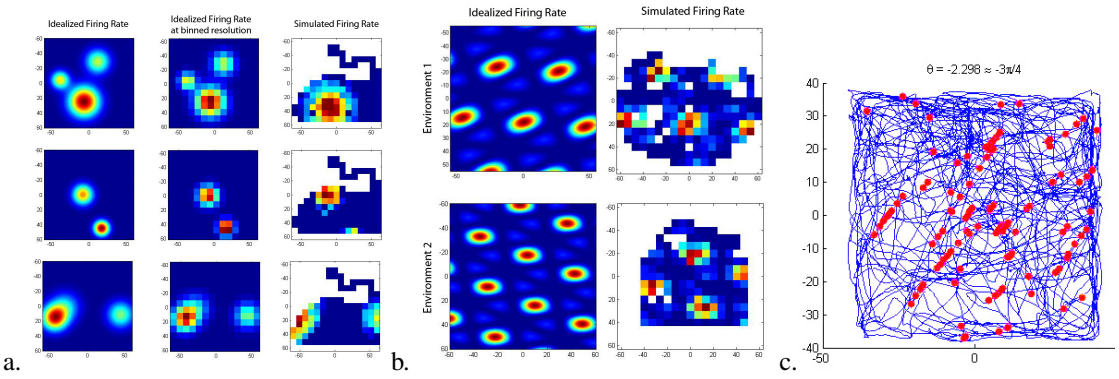

Figure 2: a. Three example cells from the Dentate Gyrus in the original version of the model. On the left is the idealized firing rate, as defined by the sum of Gaussians. In the middle column, this rate has been binned at the same resolution as binning during the simulation, so that direct comparison can be made to the simulation. On the right is the firing rate from the simulation. Unsampled bins are represented in white. b. Two example grid cells. Note that the orientation of the grid is rotated from one environment to the next (but is constant from one cell to the next within an environment (not shown)). c. A head direction cell, spikes are shown as red dots superimposed on the path of the rat. Notice that the spikes tend to fall along a line at an angle of approximately $-\frac{3\pi}{4}$, which is close to the preferred head direction of the cell of -2.298.

that increasing the number of grid cells to something more biologically realistic would even out grid activity, and then changing the sparsity and/or the weight-impact of EC-DG synapses would then allow for sparse activity with a wider range of possible firing locations. However, the modulation is almost identical for 100 cells vs. 5000 cells, so I think there is something about my generation of grid cells that precludes even tiling, but I don't know what it is.

In model version 1, CA3 has sparse activity and fires in Place Fields (mostly), while in model version 2, all CA3 cells are active and only some fire in a spatially restricted manner

In the first version of the model (with DG defined environmentally and no EC), 103 of 300 CA3 cells were active (fired more than 25 spikes during 120 seconds). On the other hand, in the second version of the model, all 300 cells were active. In model 1, the majority of active CA3 cells displayed a place-field-like firing pattern. In model 2, some of the CA3 cells had place-field like firing patterns, but many more had multiple peaks or even a lack of spatial modulation. With my original parameters of model 2, I actually had fewer than 100% of CA3 cells active, but I felt that it was not sparse enough. I decreased the synaptic impact of EC-CA3 synapses, thinking that less excitation would lead to sparser activity, but the opposite happened. I had experienced similar results in model 1, when increasing inhibition decreased sparseness. I would like to explore this issue further; I suspect that something about the recurrent connectivity of CA3 is contributing to these findings, but I haven't been able to sort that out yet.

Discussion

The goal of this project was to assemble a framework that could be used for a spiking model of the dentate gyrus and surrounding areas. I feel that overall I have accomplished this goal. However, there remain many parameters that need to be adjusted in order for this to be an accurate model that can be used to make predictions about dentate and hippocampal function. For exmaple, all neurons are currently modeled with identical parameters that make them behave as regular spiking neurons. However, the shape of spikes varies considerably between cell types, so the parameters for DG neurons should be different than those of EC or CA3 neurons.

With more time, I will be able to fit parameters much more carefully. Some parameters can be fit on a small scale. For instance, the spike shape of the principal cells in DG, EC, and CA3 is probably known. I can use that information to fit parameters for the spiking neuron model. Some parameters,

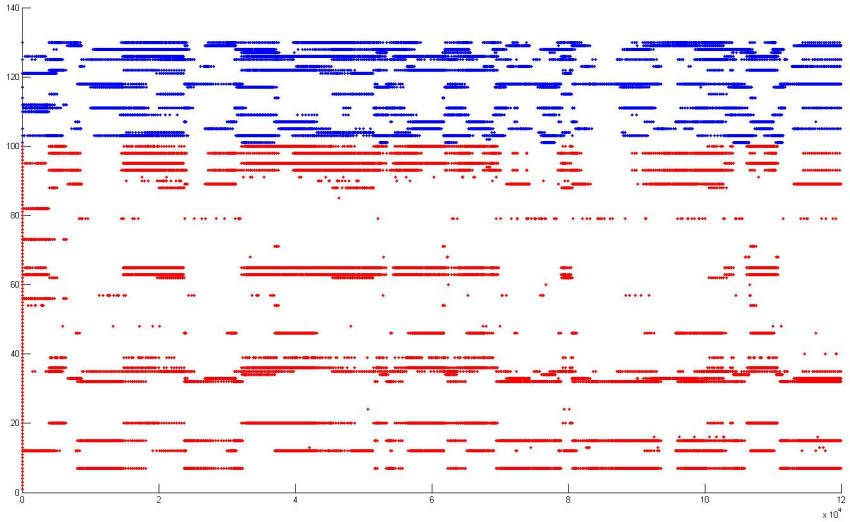

a.

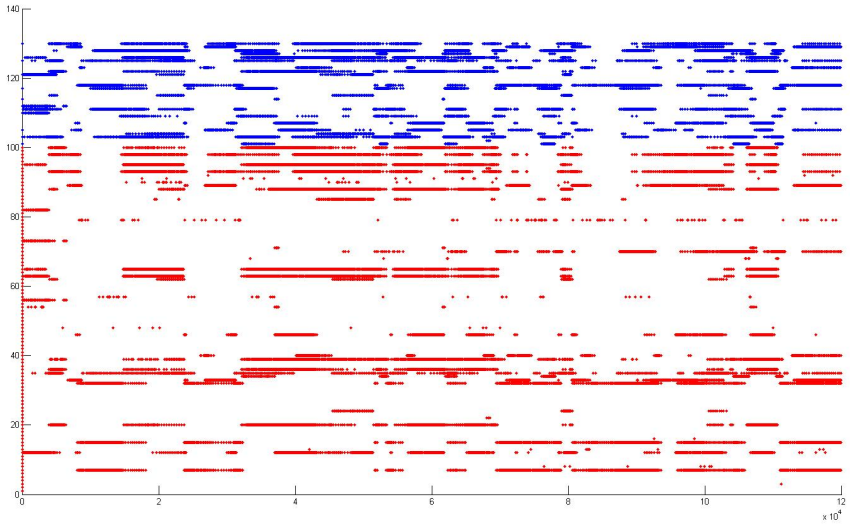

b.

Figure 3: Raster plots of the same cells in the same environment. Time (between 0 and 120 seconds) is displayed on the x-axis; cell identity is on the y-axis. At any point when a cell fires, a dot is drawn. CA3 cells are in red; DG cells are in blue. In a, there is no plasticity. In b, synapses are updated via STDP. No difference in firing patterns is discernible between the two cases.

however, will probably need to be fit on a larger scale. For example, what learning rate should be used to see effects of STDP? Should the learning rate be the same for DG-CA3 synapses as for CA3-CA3 synapses? Why does decreasing the synaptic impact from EC to CA3 cause an increase in the number of active CA3 cells? These are questions that will need to be answered by working with the entire simulation.

A further limitation of the model right now is that Matlab runs out of memory after about 120 seconds of simulation time. In a typical experiment in my lab, the animal will run for 600 seconds. I think that with some effort I can fix the code so that it will allow for building one simulation onto another. That is, I can run for 120 seconds, save crucial results, and then run for another 120

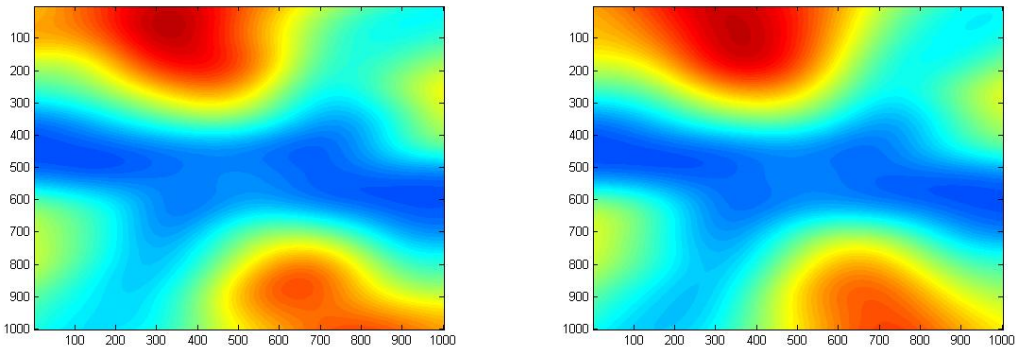

a. b.

Figure 4: The spatial distribution of average firing activity for 100 (a) vs 5000 (b) grid cells. Having more grid cells does not allow for a more even spatial distribution.

seconds and add these results to the model. This, however, did not prove as trivial to implement as I had hoped, so I haven't done that yet.

Once I have fixed these issues, I would like to explore the questions listed in the Objectives section of this paper and compare the model with experimental data.

References

[1] Vegard Heimly Brun, Trygve Solstad, Kirsten Brun Kjelstrup, Marianne Fyhn, Menno P Witter, Edvard I Moser, and May-Britt Moser. Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, 18(12):1200–1212, Dec 2008.

[2] Neil Burgess, Caswell Barry, and John O'Keefe. An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9):801–12, Jan 2007.

[3] Erika Cerasti and Alessandro Treves. How informative are spatial ca3 representations established by the dentate gyrus? *arXiv*, q-bio.NC, Jan 2009.

[4] Marianne Fyhn, Torkel Hafting, Alessandro Treves, May-Britt Moser, and Edvard I Moser. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, 446(7132):190–4, Mar 2007.

[5] E Izhikevich. *Dynamical Systems in Neuroscience*. MIT Press, 2007.

[6] Jill K Leutgeb, Stefan Leutgeb, May-Britt Moser, and Edvard I Moser. Pattern separation in the dentate gyrus and ca3 of the hippocampus. *Science*, 315(5814):961–6, Feb 2007.

[7] E Mankin. CA3 place field. 2010.

[8] J O'Keefe. Place units in the hippocampus of the freely moving rat. *Exp Neurol*, 51(1):78–109, Apr 1976.

[9] R C O'Reilly and J L McClelland. Hippocampal conjunctive encoding, storage, and recall: avoiding a trade-off. *Hippocampus*, 4(6):661–82, Dec 1994.

[10] T Solstad, C. N Boccara, E Kropff, M.-B Moser, and E. I Moser. Representation of geometric borders in the entorhinal cortex. *Science*, 322(5909):1865–1868, Dec 2008.

[11] Menno P Witter. Intrinsic and extrinsic wiring of ca3: indications for connectional heterogeneity. *Learn Mem*, 14(11):705–13, Nov 2007.