Project Title: Cellular automata models and neuroscience

Name: Alana Firl

Department: Vision Science

Email: ajfirl@berkeley.edu

Abstract:

CAs are a simple and therefore nice example of how a network of weakly coupled units with their own internal state dynamics can give rise to interesting spatial and temporal patterns of activity, without requiring external stimuli or supervision. My thesis work studies a period of retinal development that includes these same properties. Here I looked into whether CA models provide insight into the internal dynamics of the retina during this period of development.

Introduction

1. Motivation

Theory and reality often become hilariously nonoverlapping as soon as one changes the scale of study, a classic example being the different assumptions required for Newtonian versus quantum mechanics. We generally keep our realities in a comfortable space where the models still accurately predict things. We also have the considerably less lucrative option of investigating systems for which models consistently fail, with the intent of finding inspiration to build better models. Naturally occurring biological systems are really good at exposing the inflexibility and cluelessness of existing models and algorithms. I do not know if CA models will contribute to explaining things about the retinal system that I study, but the motivation is to see if they can help.

2. Why it is interesting

The system of study is a self-organizing circuit of retinal neurons. When mature, retinal circuits can borrow from each other, adapt within themselves, and voluntarily inactivate. Ultimately, the circuits are able to convert visible light into a code that the brain uses to form a visual percept. Within a few days of development, the necessary cells find each other and form the appropriate synapses. This is all done without any top-down instruction, although genetics instruct some of the cell differentiation.

Any and all insights into this process are interesting to me. I can invent any number of irrelevant, ultimately economic-driven incentives for why this is process is still interesting, but am feeling inclined, at this time, to save that energy for times of grant writing.

3. Synopsis of project and results

The objective here was to implement a CA model and check whether it displays any of the behaviors seen in the developing retina. I found the activity patterns do have similarities, but were not compelling enough for me to pursue this directly. However I think they may be useful once I figure out a few of the physical rules that govern retinal dynamics. Namely, I would like to identify how many distinct populations of cells are necessary for giving rise to the observed spontaneous dynamics before continuing with CAs.

Background

The rules governing the spontaneous retinal dynamics (referred to as "waves") are as follows: the retina receives no stimulation. It is loaded with a dye that brightens when a cell fires a spike, which is the only indicator we have that a neuron did something. Data is based on passive recordings of cell activity. I will focus on imaging data, which is comprised of 1Hz sampled image frames of cells that are either active or not active. Activity is defined by the cell's brightness compared to baseline. Here is an image of the tiling pattern of retinal waves over 30 seconds. The viewing perspective is looking down on the surface of the retina, and covers roughly 200 cells:

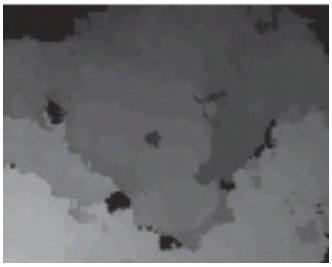

Gray indicates that a wave passed over that region. The grayscale corresponds to amount of time elapsed, with lightest gray being the most recent event, and darkest being the earliest. Black is baseline. Note that waves respect boundaries, and are for the most part spatially nonoverlapping across time.

The objective is to see if a CA model can approximate this behavior. If it does, I could use it to make predictions about the system.

Dynamical system

1. Particular system is a 2D lattice of cells. Each cell has its own internal state equation, and each cell adds some fraction of its internal state to its nearest neighbors. Cells on edges wrapped cylindrically.

2. Equation governing each cell internal state is a linear circle map:

$$y = r + 0.95x;$$

$$y \%= 1.0$$

where x is the temporally previous state value. The resulting value of y is added to all nearest neighbors, with coupling constant c.

Parameters to be varied:

r- affects internal state

c- weights contribution from neighbors

N- number of cells, doesn't make much difference once you exceed a 10x10 lattice size.

3. How the terms model various aspects of the system

The coupling constant describes the connectivity of the cells. In retina, the coupling constant is fairly weak, at least between neighboring cells, which are actually not directly connected. Roughly speaking, r governs the frequency of the internal state. Below are time plots of a linear circle map with no outside inputs, for values of r at different orders of magnitude.

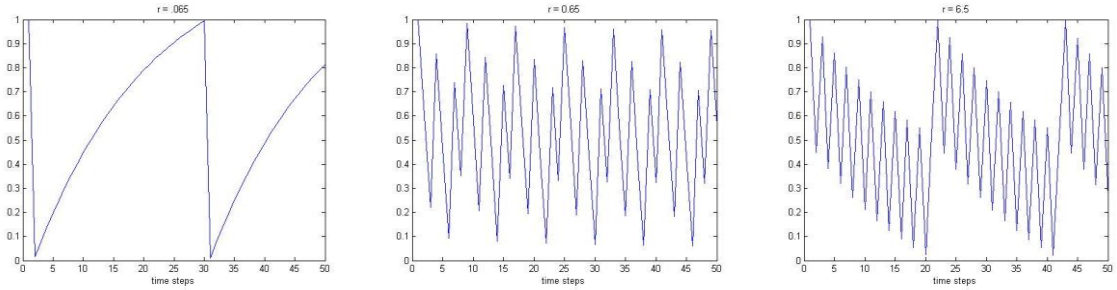

However it should be noted that a linear circle map is by no means the only possible internal state dynamic. I used it because it provides more stability in a network than other functions that do not have inherent periodicity.

Methods

Model parameters were set to various values and the model ran for 500-1000 iterations. Each pixel value was logged over iterations to obtain one timeseries per pixel. For each timeseries, baseline level of brightness was computed by averaging, and each element in the time series was compared to its baseline. Any pixel value exceeding its baseline by 15% was considered "active" and assigned a value of 1. All other values were assigned a 0. This is the same heavy-handed abuse that my retinal data endures.

Results

Roughly speaking, the CA system exhibited spatiotemporally structured patterns that are similar to those observed in retina.

Binarized raster plot of retina activity (white ticks indicate active cells at that time point):

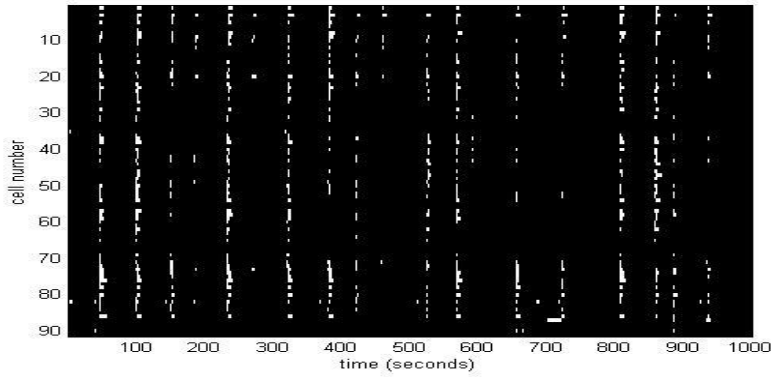

Binarized raster plot of CA activity with coupling 0 and r of .065 (top). Cell count is the number of active cells at each time point (bottom).

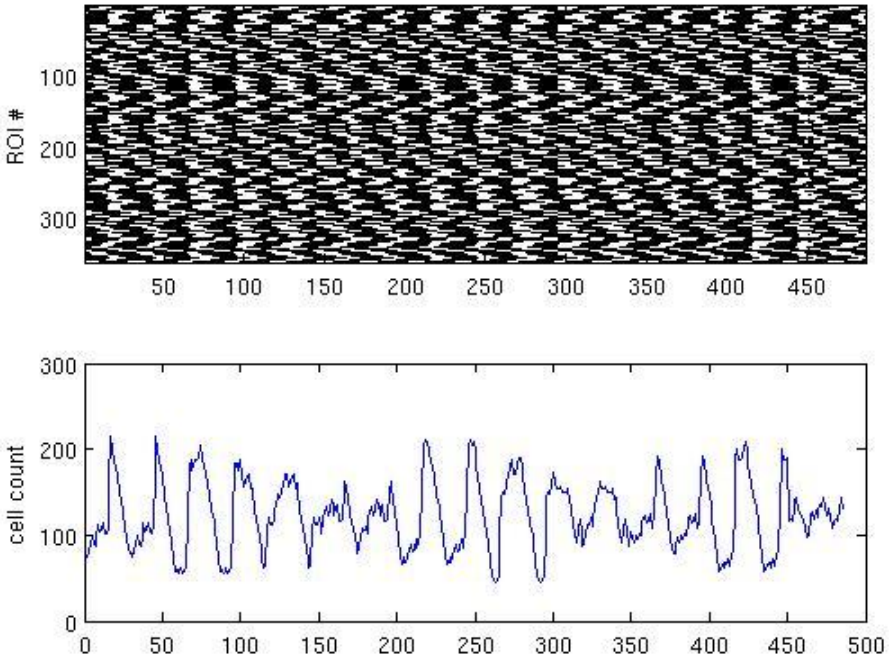

Binarized raster plot of CA activity with coupling of .20 and r of .065:

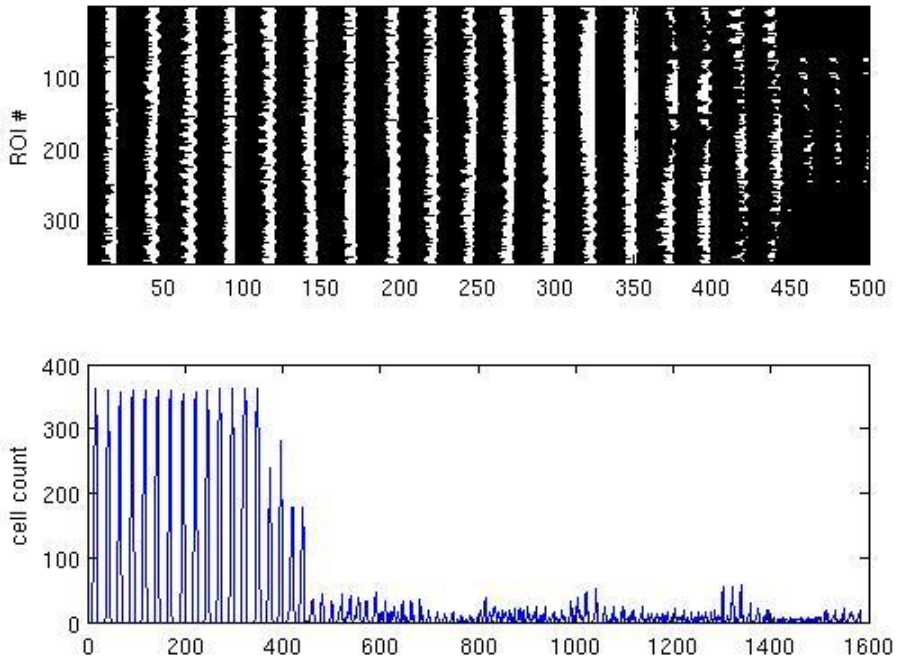

In terms of using information theoretic measures to assay the network, I have a few scantily informed, spurious opinions. Various people have successfully used mutual information versus spike timing differences to look for any interesting curve shapes. Here "interesting" means something with minima and maxima so the researcher can point out optimal parameter ranges and then argue that their biological system exploits that range. This may rapidly become a cliché; perhaps someday it will be akin to saying that a system must balance excitation with inhibition to display behavior that neither blows up to infinity nor dies to zero. People no longer dispute this in neuroscience or even in philosophy for the most part. However I maintain some faith that using mutual information or, conversely, entropy as a qualifier for determining functional system parameters can yield more subtle hints for tuning a model than assays of how many timesteps the system needed before it exploded or died.

We could do this here nonetheless, although the utility of any results should be made clear before taken up as a central endeavor. As applied to neuroscience, information is usually based on distributions of distances between pairs of neurons and the distribution of relative spike timing between said pairs given their distances. When papers mention mutual information in neuroscience it often is in the context of how much time can elapse between one neuron's spiking and another's, with the largest contributions coming from whatever the majority of pair distance and spike timing combinations are. In other words it is not safely dissimilar from taking

an average. It doesn't tell us about specificity of circuits in terms of sequences of neural activation, which is more central to my project since I care about circuit formation and refinement. As a result I have come to view measures of mutual information, when conventionally applied for the purpose of characterizing a system, to be a more appropriate conclusive step than an explorative one. It is something one can use to greatest effect after one already knows what is important for the system.

Other people have looked at information content of binarized spike/no spike words. This is closer to what one might use for measuring entropy rate and building a word distribution library that cares about sequences of events, rather than simply defining a measure over which all sequences are pooled into a distribution. However, people often fix the length of allowable words to make the calculations tractable, which reduces a lot of otherwise good debates into squabbles over what's an appropriate time window. Meanwhile I'm unsure the retina thinks about time windows. To my knowledge it has provided no evidence of such preferences. A more compelling use of timescale of spiking, to me, is that it could be used to convey different aspects of a stimulus, such as distinct spatial frequency ranges (a lame example but hopefully conveys the idea). It is my belief that retinal neurons can use time to expand channel capacity of their axons, through which various independent or complementary aspects of a stimulus can be sent using the same cable provided it is somehow kept distinct in time. There is evidence that neurons adjust their own membrane timescales for spiking depending on the type and temporal variance of inputs received, and receptive field fits improve to near-competence when the model incorporates spatial and temporal aspects of a stimulus.

Conclusion

To conclude I think CAs are a good model for constructing a self-sustaining system of cells. I am not convinced of their far-reaching utility as an explorative tool in neuroscience, but I think they may prove more useful to retinal wave research once a few more facts about the physiological basis of waves is filled in. Currently I have very little insight into what the internal state dynamics are, or what the coupling parameters should be, or if I would benefit from combining two or more populations of CAs with different internal state equations. The search space is unfeasibly large.

I think a better approach would be to find eMs for calcium imaging data, and look for states that are predictive of waves. I'd first thought MEA-based spike trains, with their ultra precise sampling of 10000Hz, would be the best option. However after trying a few things, the imaging sampling of 1Hz is actually starting to look like a blessing rather than a curse. Never thought I'd say that, and may retract that statement soon, and that's as much as I can say conclusively at this time.

References

Berkes P, Orban G, Lengyel M, Fiser J (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science 331(83-87).

Bialek W, Owen WG (1990). Temporal filtering in retinal bipolar cells. Biophysical J. 58(5):1227-1233.

Blankenship AG, Ford KJ, Johnson J, Seal RP, Edwards RH, Copenhagen DR, Feller MB (2011). Synaptic and extrasynaptic factors governing glutamatergic retinal waves. Neuron 62:230-241.

Butts DA, Rokhsar DS (2001). The information content of spontaneous retinal waves. J Neurosci 21(3):961-973.

Padmanabhan K, Urban NN (2010). Intrinsic biophysical diversity decorrelates neuronal firing while increasing information content. Nat Neurosci 13(10):1276-82.