SANDIA REPORT

SAND2016-9676
Unlimited Release
September 2016

Modeling Information Multiplexing in the Hippocampus

Frances S. Chance

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: http://www.ntis.gov/search

SAND2016-9676
Unlimited Release
September 2016

Modeling Information Multiplexing in the Hippocampus

Frances S. Chance
Data-Driven and Neural Computing (1462)
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1327

Abstract

Modern computers are constantly faced with the challenge of processing ever-growing quantities of data that span a wide range of modalities. Of particular relevance to national security interests is the ability to integrate multimodal data for the purpose of fast decision making. The brain is a biological system that is specialized for high performance at this task, suggesting that understanding the mechanisms by which neural circuits integrate multimodal data may lead to improved man-made detection systems. This research focused on understanding these neural algorithms, and specifically tested the hypothesis that hippocampal neurons multiplex information from two different input streams. Specifically, we compare the spiking behavior of a computational model of hippocampal circuitry with neurophysiological data recorded from rodent hippocampus.

ACKNOWLEDGMENTS

I am grateful to Lesley A. Schimanski and Carol Barnes for granting access to this data set.

Carol Barnes, Andrew Maurer, and Sara Burke were particularly helpful in getting the data analyses set up. I am grateful to them for helpful discussion and support throughout the lifetime of this project.

Kelsey Kalmbach developed and performed the clustering analysis of the data. Conrad James served as technical mentor on this project.

I am also grateful to the LDRD office for supporting the Early-Career LDRD program, without which this work would not have been possible.

CONTENTS

FIGURES

NOMENCLATURE

BIC	Bayesian Information Criterion
CA1	Cornu Ammonis 1 (a subdivision of the hippocampus brain area)
CA3	Cornu Ammonis 3
EC3	layer 3 of the entorhinal cortex
DOE	Department of Energy
LFP	local field potential
SNL	Sandia National Laboratories

1. INTRODUCTION

1.1. Multimodal Information in the Hippocampus

Like any complex signal processing structure, the brain is tasked with sending signals through a complex and highly interconnected web of pathways. Furthermore, biological neural networks must do so without the ability to label signals with specific addresses. This research was designed to test the hypothesis that the brain handles routing different types of information by multiplexing them such that each type is assigned to specific time segments relative to an oscillating brain signal, similar to time-division multiplexing used in telecommunications systems.

The hippocampus receives input ranging from sensory (sight, sound) to proprioceptive (how fast is the animal moving) to higher-level cognition-based signals (memory and perception of time). This part of the brain is therefore assumed to be involved with associating different types of information (the sights and sounds associated with a particular event) for memory formation. Of particular interest for this project is the long-standing theory that one function of the hippocampus is to detect novel combinations of otherwise familiar stimuli (O'Keefe & Nadel, 1978). Conveying either of these functionalities to a detection system would convey obvious advantages. New methods for rapidly integrating information from a range of modalities could potentially advance the state-of-the art, but the ability to rapidly detect novel combinations of otherwise familiar stimuli would be disruptive to the field.

The "output" area of the hippocampus is known as CA1 (Cornu Ammonis 1). This subdivision of the hippocampus served as the focus of our efforts to understand how multiple streams of information are represented in one output information stream. Aside from being the output area of the hippocampus, two other CA1 features made this specific component of the hippocampus an attractive target of study. First, the principal processing neurons in CA1 are known to receive inputs from two different areas of the brain: CA3, another subdivision of the hippocampus, and layer 3 of the entorhinal cortex (EC3). Second, CA3 is proposed to have an autoassociative memory function (Marr, 1971), while the entorhinal cortex is regarded as providing less-processed sensory input to the hippocampal system. Thus CA1 is known to receive inputs via two anatomically-distinct pathways and is thought integrate information of different modalities (carried by these two well-defined pathways) and then transmit this integrated information to cortex.

1.2. A Dual Input-Component Model of CA1 neurons

Historically, CA3 has been considered the primary input to CA1, although in studies where CA3 transmission to CA1 is disrupted, CA1 neurons continue to represent spatial information in their spiking outputs, suggesting CA3 is not the only driver of CA1 spiking (Brun et al, 2002; Nakashiba et al, 2008; Middleton & McHugh, 2016), and that there is at least one other input pathway of spatial information that is capable of independently driving CA1 output. A recent model of CA1 information integration (Chance, 2012) proposed that information from CA3 and EC3 is multiplexed in the CA1 output, with CA3 input driving spiking during a specific phases of an oscillating brain signal known as the theta oscillation and EC3 input driving spiking over a different range of theta phases. Supporting this model of hippocampal activity are studies examining gamma oscillations in the rodent hippocampus suggest that lower-frequency gamma oscillations (associated with input from CA3) and higher-frequency gamma oscillations (associated

with input from EC3) tend to occur at different theta phases (Colgin et al 2009; Schomburg et al 2014).

Validating the Chance (2012) model of CA1 information integration requires a manipulation of the input strength from either EC3 or CA3 followed by an assessment of the impact of the manipulation on CA1 spike theta phase distributions. Because EC3 and CA3 are thought to carry different types of information, it is possible to indirectly impact the input strength by varying task condition. For example, Cabral et al (2014) compared CA1 spike theta phases in mouse hippocampi while the mouse was navigating a maze based on spatial cues vs memory and found shifts in the mean spike theta phase distributions. This research took advantage of a decrease in the functional strength of CA3 input that occurs during aging (Rosenzweig & Barnes, 2003; Burke & Barnes, 2010) as an indirect method of manipulating the strength of CA3 input and validate the Chance model of CA1 information integration by examining CA1 spike patterns in young and aged rats.

2. RESULTS/ANALYSIS

2.1. A Model Neuron with Multimodal Information Multiplexing

The purpose of this research was to both validate and extend a model of CA1 neuron integration and multiplexing of two inputs streams. A certain class of neurons in the hippocampus and entorhinal cortex are referred to as "place cells" because they fire action potentials (also referred to as spikes) only within a specific region in a particular environment, referred to as the "place field" of that neuron. Another interesting feature of hippocampal place cells is that the relationship between the theta phase of a fired action potential and the animal's location within the place field of that neuron follows a particular relationship (referred to as phase precession - see Figure 4 for examples). Specifically, as an animal enters the place field of a neuron, spikes are fired relatively late in the theta cycle (we define the peaks of the theta oscillation to be the boundaries of one theta cycle). As the animal traverses the place field, the spikes systematically shift earlier in theta phase. The mechanism by which theta phase precession occurs is unknown, yet this work was inspired by a model of this phenomenon in this report. Full details of the model have been previously published (Chance, 2012). A brief description of this CA1 neuron model is given here to clarify the motivation underlying this research project.

2.1.1. Model

The dynamics of the CA1 neuron membrane potential (V) are described by

$$C\frac{dV}{dt} = g_L(E_L - V_L) + g_E(E_E - V),$$

where C is the capacitance of the cell membrane, g_L is the neuronal membrane conductance in the absence of any excitatory synaptic input, g_E is the total excitatory synaptic conductance (described below), E_L is the resting neuronal membrane potential, and E_E is the reversal potential of the excitatory synaptic current. If V depolarizes above a pre-determined threshold, a spike (or action potential) is fired and the membrane potential is reset to E_L.

The model neuron receives two distinct components of synaptic input (one from CA3 input and one from EC3), denoted by i in the equations below. The synaptic input rate of each input component is

$$r_i = A_i(x)\lfloor(\cos(2\pi ft + \phi_i) + b_i)\rfloor_+.$$

$\lfloor\ \rfloor_+$ denotes rectification (firing rates cannot be negative), f is the theta frequency, ϕ is the theta phase of the input component, and b is a parameter that the degree of theta modulation of the input. Input events are generated as Poisson input trains with underlying rate r_i.

With each input event, g_E is increased by a given Δg such that the resulting deflection in V is approximately 1 mV. In the absence of any input events, g_E exponentially decays to zero with a time constant of 2 ms.

$A(x)$ defines the place field of each input component, where the amplitude of the oscillating input component rate is

$$A_i(x) = \alpha_i \frac{\exp(x - x_i)^2}{2\sigma_i^2}.$$

Natural variations in the relative amplitudes of the two input components and the distances between peak input locations are thought to underlie the diversity of place fields profiles and diversity of phase precession "shapes" observed in the hippocampus (Chance, 2012). While both input components oscillate at the approximately the same frequency (f), their inherent theta phases (ϕ) are separated by approximately half of one theta cycle, with CA3 inputs driving model neuron spiking at later theta phases and EC3 inputs driving spiking at earlier theta phases.

The model proposes that the CA3 and EC3 input components to a given place cell are arranged in space such that as an animal enters the place field, the neuron is first driven almost exclusively by CA3 input. In the center of the place field the neuron receives input via both pathways. As the animal exits the place field, EC3 input dominates. Details of how this produces biologically-realistic phase precession are given in Chance (2012).

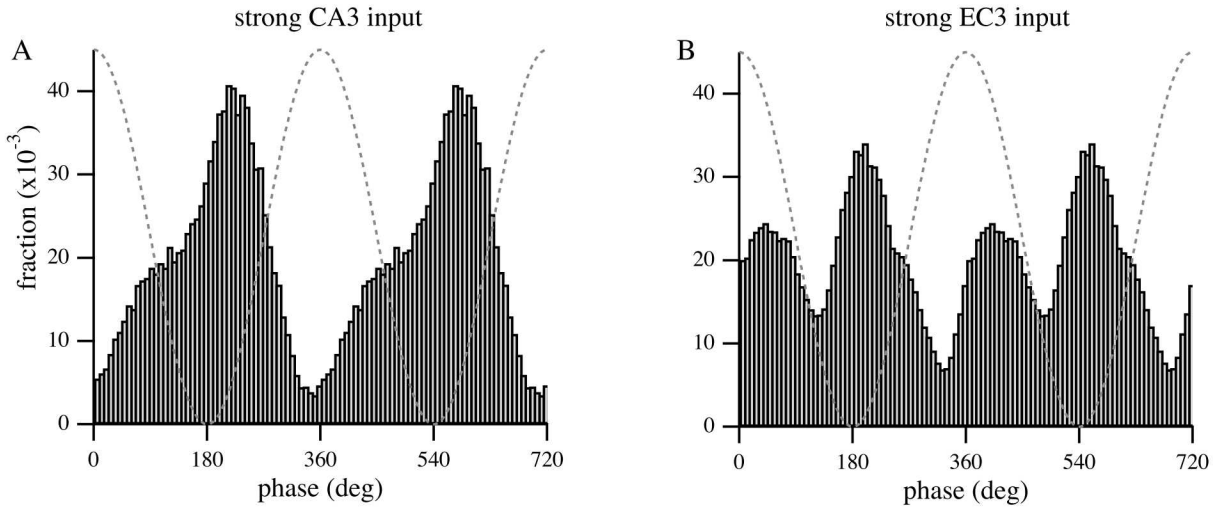

Figure 1. Theta phase distributions of spikes fired by the model CA1 cell in the presence of (A) strong CA3 input and weaker EC3 input or (B) strong EC3 input and weaker CA3 input. The y-axis is the fraction of spikes that were fired at each theta phase (x-axis). Two theta cycles are repeated on the x-axis for clarity. The dashed line is the model theta oscillation.

2.1.2. Variations in Input Component Strength Results in Different Theta Phase Distributions in the Model

Because the two input components of the model described above drive CA1 spiking at different theta phases in the model, the relative strengths of the two input components may be inferred from the distribution of CA1 spike theta phases. Figure 1 demonstrates

how theta-phase distributions vary as the ratio of CA3 to EC3 input strength is modified in the model. In (A) CA3 input is strong relative to EC3 input. The resulting theta-phase distribution appears unimodal with a peak during the ascending phase of the theta oscillation (slightly later than 180°). By contrast, in (B) the EC input is much more effective at driving CA1 spiking by itself. With this configuration of inputs, a secondary peak in the theta-phase distribution, appearing during the descending phase of the theta oscillation, becomes more pronounced.

This research was driven by the hypothesis that the CA1 neuron model could be verified by comparing theta-phase distributions from young and aged rats. Because CA3 functional input weakens with age (see Burke & Barnes for a review), it was hypothesized that the theta-phase distribution of CA1 spiking in aged rats would look more like the distribution of spike phases illustrated in Figure 1B (strong EC3 input), whereas the spike theta-phase distribution recorded in CA1 of younger rats would look more like the theta-phase distribution illustrated in Figure 1A (produced by a version of the neuron model that received strong CA3 input).

2.1.3. Experimental Data

We were graciously offered a data set from a previously published study (Schimanski et al 2013) to test this prediction of the model. In this study, cohorts of six young (9-12 months old) and six aged (25-28 months old) rats were trained to run on a circular maze in exchange for food rewards. A barrier prevented the animal from completing the full circle, so each rat essentially followed a horseshoe-shaped pattern on the circular maze. All experiments were performed according to protocols previously approved by the University of Arizona Institutional Animal Care and Use Committee and followed the guidelines of the United States National Institutes of Health Guide for the Care and Use of Laboratory Animals. All data collection was completed prior to 2013 (the time of proposal of this LDRD). Full details of the methods associated with conducting electrophysiological recordings in rat hippocampi and isolating spikes fired by individual neurons are given in Schimanski et al (2013).

Each day of the experiment, rats ran on the maze for two consecutive running sessions separated by a short break. This protocol was repeated daily for 31 consecutive days. For five of the rats (four of the young rats and one of the aged rats), all recorded place fields remained stable over the course of each day (that is, all recorded place cells encoded the same location on the maze across the two consecutive running sessions). For six of the rats (two young rats and four aged rats), "remapping" place fields (where the same place cell did not encode the same location on the maze during both sessions) were observed (but only after day 13). One aged rat became sick and did not complete all 31 days of the experiment.

Spike theta phases were identified by first applying a band-pass (1-80 Hz) filter to the recorded local field potential (LFP), followed by identification of local maxima and minima of the filtered LFP. Theta phases were assigned by linear interpolation between the identified maxima and minima such that the time interval between any maximum and minimum always comprised half of one theta cycle. The theta oscillation is known to be present only during exploration (when the animal is moving) and REM sleep. We therefore identified time epochs within which the animal was moving faster than 10 cm/s. Only spikes from within these epochs were used for our analyses.

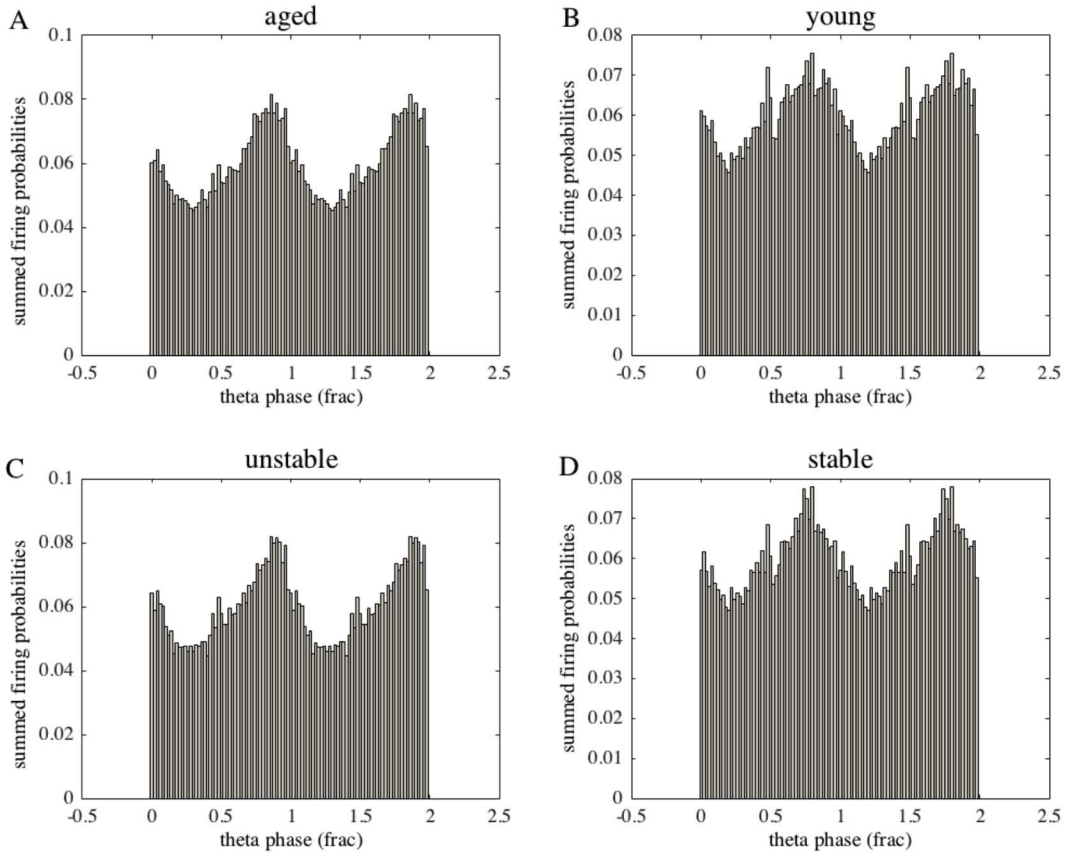

Figure 2. Experimental Day 1: Theta-Phase Distributions of Place Cells from Different Experimental Groups. Each panel is the probability of firing (summed across all cells in all rats of the indicated experimental group) as a function of theta phase (plotted as fraction of one theta-cycle). Experimental groups are A) Aged rats, B) Young rats, C) Rats showing unstable place fields (defined in text), and D) Rats showing only stable place fields over the course of the experiment. Two identical theta cycles are shown.

2.1.4. Analyses of Experimental Data: Theta-Phase Distributions

2.1.4.1. Examination of Theta-Phase Distributions

Once we had assigned theta-phases to spikes fired during active exploration/locomotion, we next examined theta-phase distributions as a function of experimental group. Initially we compared theta-phase distributions of spikes fired by place cells in aged rats to the theta-phase distributions from place cells in young rats. Figure 2A shows the theta-phase distribution of spikes collected from aged rats during the first day of recording, expressed as the summed firing probability across all aged rats. Figure 2B is the spike theta-phase distribution collected from young rats. Although visual inspection suggests the presence of a secondary theta peak, indications of this second peak were not present in subsequent days (see Figure 3B, for example),

although it is possible that the second peak is present but less-defined (spread out over a wider range of theta phases, for example).

We next compared spike theta-phase distributions collected during Experimental Day 1 from rats with unstable place fields (place cells that were observed to remap at least once over the course of one day) and rats with place fields that remained stable throughout the course of the experiment. As with the young and aged groups, the stable group seemed to show a more pronounced secondary peak in the theta phase distribution (partially expected because the stable group was primarily comprised of young rats and the majority of the unstable group was aged rats). As with the young and aged groups, this secondary peak is less pronounced in Experimental Day 21, although the possibility exists that the peak is broadening rather than disappearing (see Figure 3D).

Figure 3 shows theta phase distributions from the same rats, from the indicated experimental groups, but for spikes collected on Day 21. Unlike Day 1, spiking from place cells in young or stable groups did not show a pronounced secondary peak. However, there does appear to be significantly less theta-modulation (the spikes do not appear to be as synchronized with the theta rhythm). One interpretation of this phenomenon is that there is a shift of spiking to being driven by the secondary component in young and stable rats on this day as well, but the secondary component is not as strongly theta-modulated as on Day 1. We also examined data from Experimental Day 7 and Experimental Day 14, but did not observe any consistent trends within groups or across days.

2.1.4.1. Theta Modulation of Spike Trains

In Figure 3, one interpretation of the theta phase distributions in Figure 3B and Figure 3D is that the secondary component of spiking is growing in strength but is less coherent than on Day 1. We therefore fit spike trains of individual place cells to cosine functions, and compared average amplitude of theta modulation for place cells from a given rat across experimental groups on a given day, and also across experimental days. We observed a trend that, in general, place cells from the unstable group were more theta-modulated than place cells from the stable group (indicating a more unimodal distribution of spike phases in the unstable group). We also observed a trend that for both groups, degree of theta-modulation tended to decrease over time. In neither case was this trend statistically significant, however.

2.1.5. Analyses of Experimental Data - K-means Clustering of Place Cell Spikes

Because there were indications of a difference in spike patterns between the different experimental groups, and also across experimental days, we sought to find a more sensitive measure of comparing spike patterns. We determined that a clustering algorithm would facilitate us examining the relationship between spike timing and animal location when the spike was fired (as hypothesized by the model). Thus our next step was to adapt a K-means clustering algorithm to analyze this data set.

2.1.5.1. Data clustering

Each spike is associated with a location x_i (calculated as a one-dimensional linearization of the circular track) and a theta phase ϕ_i, defined over the interval $[0,2\pi)$. Because the

outcome of various clustering algorithms can be strongly impacted by the relative variance along different dimensions, we chose to scale both theta phase and location such that the standard deviation in the theta-dimension was the same as the standard deviation in the x-dimension.

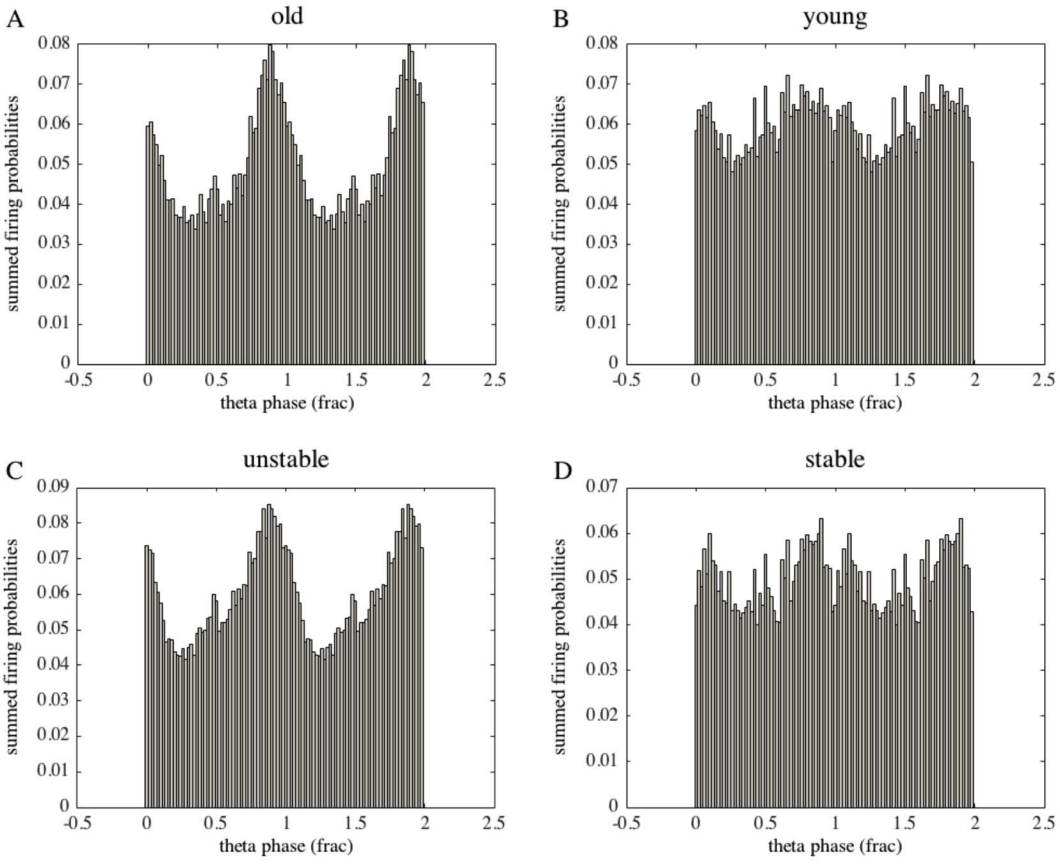

Figure 3. Experimental Day 21: Theta-Phase Distributions of Place Cells in Different Experimental Groups. Each panel is the probability of firing (summed across all cells in all rats in the indicated experimental group) as a function of theta phase (plotted as fraction of one theta cycle). Experimental groups are A) Aged rats, B) Young rats, C) Rats showing unstable place fields, and D) Rats showing only stable place fields. Two complete theta cycles are shown.

We defined the angular standard deviation of ϕ to be

$$s_\phi = \sqrt{2(1-R)},$$

where

$$R = \frac{1}{N}\sqrt{(\sum_{i=1}^{N}\sin(\phi_i))^2 + (\sum_{i=1}^{N}\cos(\phi_i))^2}.$$

We then applied a modified version of the K-means clustering algorithm to cluster the spikes from each place cell. Because spikes fall in a cylindrical space, modifications to the standard K-means algorithm were required to adapt it to our purposes. While the distance between two locations is calculated using standard Euclidean distance, distance between spike phases is

$$\min(|\phi_i - \phi_j|, 2\pi - |\phi_i - \phi_j|),$$

and mean phase is

$$\overline{\phi} = \arctan \frac{\sum_i \sin(\phi_i)}{\sum_i \cos(\phi_i)},$$

where arctan is the quadrant-specific arctangent function.

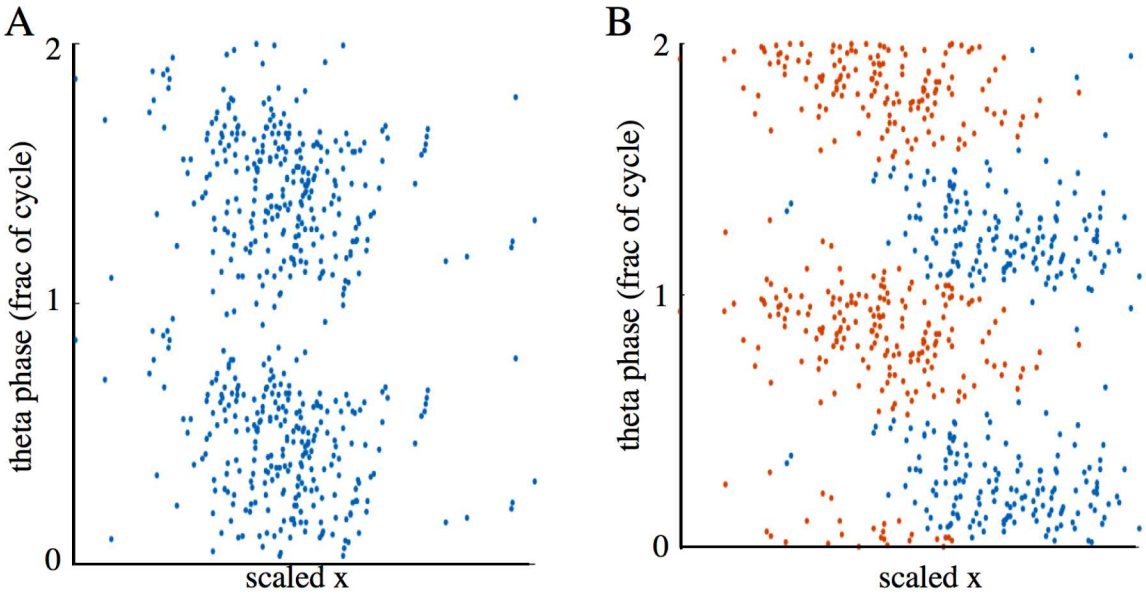

Figure 4. Examples of Spike Clustering. For the two example cells above, each point represents one spike, plotted with the theta phase of the spike on the y-axis and the scaled position plotted on the x-axis. In both examples, the rat runs right to left. Left: Example of a one-cluster place cell. Right: Example of a two-cluster place cell (spikes from each cluster are colored in either red or blue).

With the K-means algorithm modified to accommodate the circular nature of the theta phases, we then ran the clustering algorithm on the spikes recorded from each place cell. We calculated the Bayesian Information Criterion (BIC) for clustering done with the number of clusters k = 1, ... , 6. In general, the spiking patterns of most place cells in this data set are best described with one (see Figure 4A) or two clusters (see Figure 4B), consistent with previous analyses of place field spiking (Yamaguchi et al, 2002). If the spike patterns are well described by two clusters, the center-of-mass of each of the

clusters tends nearly antiphase to each other. Across rats, the average center-of-mass of the first cluster (the spikes fired first as the rat traverses the place field) is 4.73 radians and the center-of-mass of the second cluster is 1.69 radians. We did not observe any consistent trends when comparing cells from young, aged, stable, and unstable experimental groups (data not shown). Nor did we observe consistent trends when comparing spiking across days, regardless of the experimental group the spikes were drawn from.

3. CONCLUSIONS

Although it is known that there are two anatomically distinct input pathways converging onto principal neurons in CA1, the function of these two pathways (what types of information they carry, for example) is unknown. The fact that the K-means algorithm tends to cluster CA1 neuron spiking into one or two clusters is interesting, as it has not (to our knowledge) been demonstrated that the input pathway from EC3 is capable of directly driving spiking in CA1. The results of the K-means clustering presented here suggest that there at least two separable pathways or mechanisms by which CA1 spiking is driven.

Direct validation of the model proposed by Chance (2012) would require acute and targeted manipulation of either EC3 or CA3 activity (in contrast, the manipulations performed by Brun et al and Nakashiba et al required at least 6 weeks between manipulation and measurement of effect). To our knowledge, the neuroscience field still lacks the technology to successfully complete such an experimental design while maintaining quality recording in CA1.

Multiple methods of impacting input from either EC3 or CA3 exist, however (e.g. Brun et al, 2002; Nakashiba et al, 2008; Middleton and McHugh, 2016). Moreoever, the functional strength of CA3 input to CA1 is known to decrease with aging, and certain cognitive loads (for example navigation by memory) are thought to enhance the effectiveness of CA3 input at driving CA1 spiking (Cabral et al 2014).

Although we did not measure a significant impact of aging on CA1 spike theta-phase distributions, the lack of evidence does not invalidate the Chance model of CA1 input integration. The impact of aging on EC3 input is unknown, for example. It is possible that both CA3 input and EC3 input weaken with aging, in which case the model predicts no observable effect of aging on CA1 spiking (assuming the CA3 and EC3 input do not drop below transmission threshold during aging).

4. REFERENCES

Brun VH, Otnæss MK, Molden S, Steffenach H-A, Witter MP, Moser M-B, Moser EI (2002) Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science 296:2243-2246.

Cabral HO, Vinck M, Fouquet C, Pennartz CMA, Rondi-Reig L, Battaglia FP (2014) Oscillatory Dynamics and Place Field Maps Reflect Hippocampal Ensemble Processing of Sequence and Place Memory under NMDA Receptor Control. Neuron 81: 402-415.

Chance FS (2012) Hippocampal phase precession from dual input components. J. Neurosci. 32: 16693-16703.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser M-B, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. Nature 462: 353-357.

Marr, D. (1971). Simple memory: A theory for archicortex. Phil Trans Roy Soc B B262: 23-81.

Middleton SJ, McHugh TJ (2016) Silencing CA3 disrupts temporal coding in the CA1 ensemble. Nature Neuroscience 19: 945-951.

Nakashiba T, Young JZ, McHugh TJ, Buhl DL, Tonegawa S (2008) Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. Science 319:1260 –1264.

O'Keefe J, Nadel L (1978) The Hippocampus as a Cognitive Map. Oxford University Press, Oxford, England

Schimanski LA, Lipa P, Barnes CA (2013) Tracking the Course of Hippocampal Representations during Learning: When Is the Map Required? J. Neuroscience 33: 3094-3106.

Schomburg EW, Fernández-Ruiz A, Mizuseki K, Barényi A, Anastassiou CA, Koch C, Buzsáki G (2014) Theta Phase Segregation of Input-Specific Gamma Patterns in Entorhinal-Hippocampal Networks. Neuron 84: 470-485.

Yamaguchi Y, Aota Y, McNaughton BL, Lipa P (2002) Bimodality of theta phase precession in hippocampus place cells in freely running rats. J. Neurophysiology 87: 2629-2642.

DISTRIBUTION

1	MS0359	D. Chavez, LDRD Office	1911
1	MS0899	Technical Library	9536 (electronic copy)
1	MS1327	Frances Chance	1462


Sandia National Laboratories