UNIVERSITY OF CALIFORNIA, SAN DIEGO

Quantifying Information Storage in the Rat Hippocampus: Computational Analysis of
Electron Microscopy Data

A Dissertation submitted in partial satisfaction of the requirement for the degree
Doctor of Philosophy

in

Neurosciences

by

Cailey Thekla Bromer

Committee in charge:

        Terrence J Sejnowski, Chair
        Mark Ellisman
        David Kleinfeld
        Jill Leutgeb
        Charles F Stevens

2017

The dissertation of Cailey Thekla Bromer is approved, and is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____
Chair

University of California, San Diego

2017

DEDICATION

I dedicate my dissertation to my grandparents, who taught and inspired my parents, who taught and inspired me. In particular, I dedicate my dissertation to my maternal grandmother, Catherine Bradley Galvin, who is one of the strongest, kindest and smartest women I know, and to my paternal grandfather, Dr. Richard Bromer, whose "DR. BROMER" sign I have coveted for many years, and whom I regret I did not know better.

EPIGRAPH

'You', your joys and sorrows, your memories and your ambitions, your sense of personal identity and free will, are in fact no more than the behavior of a vast assembly of nerve cells
~Francis Crick

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

First and foremost, I acknowledge Dr. Terrence J Sejnowski for taking a chance on a second year student determined to learn to some programming; without that chance I would not be where I am. I acknowledge his diplomacy in the lab, and his intelligence, decisiveness and dedication to his field. I will be ever-grateful for that chance.

Second, I acknowledge Dr. Thomas M Bartol, for teaching me everything I know about computing, and much that I still do not. His upbeat attitude, passion for science and computing, and his friendship were an invaluable source of support and reassurance through moments of doubt, deadlines and new challenges. I will miss working with him immensely.

Third, I acknowledge Dr. Kristen M Harris for generously sharing the dentate data that comprises the bulk of the findings in Chapter 2, as well as for her insight and expertise in helping me to understand the technical and experimental details of my project.

I would additionally like to acknowledge Drs. Charles F Stevens, Jill Leutgeb, Mark Ellisman and David Kleinfeld for the guidance and feedback they have given me in their capacity as committee members.

Chapter 2, in full, is a reprint of the material as it appears in eLife, 2015. Bartol T.M., **Bromer C**., Kinney J., Chirillo M.A., Bourne J.N., Harris K.M. & Sejnowski T. J. The dissertation author was a primary investigator and co-author of this paper.

Chapter 3, in full, is currently being prepared for submission for publication of the material. Bromer C., Bartol T.M., Bowden J., Gonzalez P., Hanka D., Hubbard D., Kuwajimi M., Mendenhall J., Parker P.H., Abraham W.C., Sejnowski T.J., Harris, K.M. The dissertation author was the primary investigator and author of this material.

VITA

2011   Bachelor of Science, Brown University

2014   Master of Science, University of California, San Diego

2017   Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

Bartol T.M., **Bromer C**., Kinney J., Chirillo M.A., Bourne J.N., Harris K.M. & Sejnowski T. J. (2015). Nannoconnectomic upper bound on the variability of synaptic plasticity. *eLife*, 4:e10778.

RESEARCH EXPERIENCE

2009-2011   Department of Neuroscience, Brown University
              Advisors: Dr. Barry M Lester and Dr. Carmen J Marsit

2012-2017   Division of Biological Sciences, University of California, San Diego
              Advisor: Dr. Terrence J Sejnowski

ABSTRACT OF THE DISSERTATION

Quantifying Information Storage in the Rat Hippocampus: Computational Analysis of
Electron Microscopy Data

by

Cailey Thekla Bromer

Doctor of Philosophy in Neurosciences

University of California, San Diego, 2017

Professor Terrence J. Sejnowski, Chair

The importance of memory to consciousness, evolution and self-awareness is
unparalleled. Since early pivotal studies on hippocampal legions (Scoville, 1957), the
hippocampus has been recognized as an integral component of the memory formation.
Decades of research has refined and improved our understanding of this small, yet
complex, structure (Turner, 1969; Moser, 1998; Fortin, 2002). The hippocampus has been
implicated in not only in memory formation, but also recall of memory (Smith, 2006;
Foster, 2006), navigation (Thompson, 1989; Leutgeb, 2005) and time perception (Olton

1987; Mankin, 2012), among other things. Today we better understand how activity in this region and connections throughout the brain relate to the specific and parcellated functions it performs (Bliss, 1993; Leutgeb, 2005; Leutgeb 2007). The aim of this work is to understand the memory storage capacity of the brain, and, more specifically, synapses. In other words, it aims to quantify how much we can store, and in turn, how much we must forget.

Signal detection theory (Green & Swet, 1966) rose to popularity as a way to elucidate the "signal-to-noise" ratio (SNR) in Neuroscience (Schultz, 2007). At its simplest, the theory aims to discriminate two stimuli (Stanislaw, 1999). As such, signal detection theory has been the basis for discriminating a number of types of sensory stimuli (McGill, 1967; Foley, 1981). In our case, the two stimuli being distinguished are the synapse strength of two spines with a unique connectivity (Bartol, 2016). In the context of this thesis project we demonstrate that the dentate gyrus, a unique area of the hippocampus has the capacity to store less information at each of its synapses than does CA1.

Understanding how information storage capacity varies across species, brain region and age, can help further inform not only future research but also machine learning algorithms and other algorithms that aim to mimic and replicate the richness of the brain (Palm, 1989).

CHAPTER 1:

INTRODUCTION

The goal of this dissertation has been two-fold. Firstly, it aims to elucidate information storage capacity in the hippocampus. Secondly, and in pursuit of the first goal, it aims to develop a set of useful tools for the analysis of electron microscopy data with a goal of further exploring questions that might be answered by this type of data.

Hippocampus and Memory

The mammalian hippocampus is characterized by a "trisynaptic loop", namely referring to the transmission of information from the cortex received by the dentate gyrus (1), from dentate to CA3 (2) and from CA3 to CA1 (3) via synaptic connections. Information from CA1, the output of the tri-synaptic loop, is then relayed back to various parts of the cortex via widespread projections, many of which pass through the thalamus en route to their targets. This canonical connectivity model has been demonstrated and validated by innumerable experiments and is well-established in the field (Amaral 1989; van Strien 2009; Mizuseki 2009).

Spike Timing-Dependent Plasticity

In the hippocampus (as well as the cortex) the relative timing of axonal inputs and postsynaptic firing are instrumental in determining the strength of a synapse (Song, 2000; Dan 2004; Caporale 2008).The rules governing this phenomenon, or spike timing-

dependent plasticity (STDP), can be simplified as follows: when a presynaptic input precedes postsynaptic firing, synapses are strengthened; when postsynaptic firing precedes presynaptic input synapses weaken. The time that passes between these two events is critical for the induction of STDP. First demonstrated experimentally by Bi and Poo in 1998, the prevalence of STDP forms the basis for our understanding of the capacity for synapses to grow and shrink and is a prominent force in shaping long-lasting changes in synaptic strength in the hippocampus.

STDP, which has been extensively studied and characterized in brain regions involved in episodic memory formation, such as cortex and hippocampus, reflects the importance of coordinated changes on the presynaptic and postsynaptic sides of the synapse. For instance, the size of the active zone and the number of docked vesicles at a given synapse (presynaptic metrics) scale with the size of the postsynaptic density (PSD), the volume of the spine head and the relative ratio of some neurotransmitter receptors (postsynaptic metrics) and the strength of a synapse is related to a number of metrics such as spine head volume, postsynaptic density area, the number of docked vesicles and the ratio of various neurotransmitters (Harris and Stevens, 1989; Lisman and Harris, 1994; Schikorski and Stevens, 1997; Murthy, 2001; Branco, 2008; Bourne, 2013).

Though the basic premise of STDP is well-established, previous work has fallen short in characterizing how precisely STDP processes are controlled. In other words, while the demonstrated directional changes in synaptic strength resulting from STDP are well-established, the current techniques available in the field have made it impossible to characterize the precision in the control of such directional changes. One way to

characterize this kind of precision is able to take advantage of so-called "redundant synapses", or, pairs of spine heads sitting on the same dendritic branch, and contacted by a single axon. Such synapses offer a unique window to test the relative similarity of the strength of spines, which, if only controlled by STDP plasticity processes, should be identical in strength given their shared history of presynaptic input and postsynaptic firing.

Signal Detection Theory Applied

One way to measure the precision of STDP processes is to think of the strength of a spine as a "signal" and any unintended variability in strength as "noise". Using this framework, we can apply standards from Signal Detection Theory (described above) to quantify the strength of the signal (synaptic strength) relative to the noise (variability in spines with identical presynaptic and postsynaptic history). Measuring the strength and variability is a challenge in itself; because of the inherent stochasticity built into the brain, the observed metrics of synaptic strength are likely to change in the window in which they are measured. For instance, a patched cell may be undergoing ongoing transcription and translation even while current is being recorded. Similarly, a spine responding to Ca^{2+} or glutamate release changes artificially in the sense that it changes without its partner.

Electron microscopy data offers a unique opportunity to study STDP, counter-intuitively, given its static nature. By identifying pairs of spines with shared presynaptic and postsynaptic history and measuring their strength, we are able to STDP in 3D reconstructions of electron microscopy data. In order to effectively study such synapses,

it is necessary to accurately represent them in 3D space, with the goal of accurately

measuring metrics like spine head volume, which is most positively correlated with

synaptic strength. In the course of this thesis project, existing software was developed

and expanded to facilitate this and future studies of static electron microscopy image

data (Chapter 4). In this dissertation, the metrics of synaptic strength listed here are

measured and utilized collectively, or independently to represent synaptic strength.

CHAPTER 2:

NANOCONNECTOMIC UPPER BOUND ON THE VARIABILITY OF SYNAPTIC

PLASTICITY

Abstract

Information in a computer is quantified by the number of bits that can be stored and recovered. An important question about the brain is how much information can be stored at a synapse through synaptic plasticity, which depends on the history of probabilistic synaptic activity. The strong correlation between size and efficacy of a synapse allowed us to estimate the variability of synaptic plasticity. In an EM reconstruction of hippocampal neuropil we found single axons making two or more synaptic contacts onto the same dendrites, having shared histories of presynaptic and postsynaptic activity. The spine heads and neck diameters, but not neck lengths, of these pairs were nearly identical in size. We found that there is a minimum of 26 distinguishable synaptic strengths, corresponding to storing 4.7 bits of information at each synapse. Because of stochastic variability of synaptic activation the observed precision requires averaging activity over several minutes.

Introduction

Synapses between neurons control the flow of information in the brain and their strengths are regulated by experience. Synapses in the hippocampus are involved in the

formation of new declarative memories. Understanding how and why synaptic strengths undergo changes in the hippocampus is important for understanding how we remember facts about the world. A fundamental question is the degree of precision in the adjustment of synaptic strengths in view of the many sources of variability at synapses. In this study we provide an upper bound on the variability of synaptic plasticity and quantify a lower bound on the amount of information that can be stored at a single synapse.

Excitatory synapses on dendritic spines of hippocampal pyramidal neurons have a wide range of sizes. Anatomical measurements of the spine size, the area of the postsynaptic density (PSD), the number of AMPA receptors, the area of the presynaptic active zone and the number of docked vesicles in the presynaptic terminal are all highly correlated with each other and with physiological measurements of the release probability and the efficacy of the synapse (Harris and Stevens, 1989; Lisman and Harris, 1994; Harris and Sultan, 1995; Schikorski and Stevens, 1997; Murthy, 2001; Bronco, 2008; Bourne, 2013). Thus, each of these individual characteristics is a correlate of synaptic strength. The sizes and strengths of these synapses can increase or decrease according to the history of relative timing of presynaptic inputs and postsynaptic spikes (Bi and Poo, 1998).

If experience regulates synaptic strength then one might expect that synapses having the same pre- and postsynaptic histories would be adjusted to have the same strength. But what would be the inherent variability, or conversely the precision, of this process? Due to the high failure rate and other sources of stochastic variability at synapses one might expect that the precision of changes in the strengths of these synapses in vivo to be low. The failure rate at synapses depends inversely on the strength, and therefore the size, of the synapse. On this basis the strengths of weaker, and therefore smaller and less

reliable synapses, would be expected to be less precisely controlled than the larger and stronger synapses, which have a lower failure rate.

An ideal experiment to test for the precision of the changes in synaptic strength would be to stimulate in vivo the axonal inputs to two well-separated spines on the same dendrite to insure that they have the same presynaptic and postsynaptic history of stimulation. Nature has already done the experiment for us as pairs of spines on the same dendrite contacting the same axon satisfy this condition. Prior work suggests that such pairs of spines are more similar in size than those from the same axon on different dendrites (Sorra and Harris, 1993). Here we evaluated this axon-spine coupling in a complete nanoconnectomic three-dimensional reconstruction from serial electron microscopy (3DEM) (Harris, 2015) of hippocampal neuropil. We determined the similarity of synapses among pairs of spines and set an upper bound on the variability and the time window over which pre- and postsynaptic histories would need to be averaged to achieve the observed precision.

Results

In a 6 x 6 x 5 mm3 complete 3DEM from the middle of stratum radiatum in hippocampal area CA1 (Mishchenko, 2010; Kinney, 2013; see Materials and Methods). We identified 449 synapses, 446 axons and 149 dendrites, which except for one identified branch point, are likely to originate from different neurons based on the size of the reconstructed volume and the obtuse branching angles of dendrites from these neurons (Ishizuka, 1995; Megıas, 1997). We measured spine head volume and surface area, surface

area of the postsynaptic density (PSD) adjacent to the presynaptic active zone, and spine neck volume, neck length and neck diameter at the 287 spines that were fully contained within the volume. We also quantified the number of vesicles at the 236 spines and presynaptic boutons that were fully contained within the volume. The strong correlations between these metrics, the skewed shape of the frequency histograms, and the number of synapses per unit of volume (Figures 1–3), are consistent with previous observations (Harris and Stevens, 1989; Schikorski and Stevens, 1997; Sorra, 2006; Bourne and Harris, 2011; Bourne, 2013; Bell, 2014). To reduce error, we averaged over multiple independent spine volume measurements for each spine (Figure 3—figure supplement 1). We determined that the relationship between PSD area and spine head volume did not differ significantly across different dendrites (Figure 1—figure supplement 1). The correlation between spine head area and spine head volume accounted for 99% of the variance, despite the wide range in spine head shapes and dimensions (Figure 1A), which suggests that the accuracy of our measurements matched the precision of the spine. We also measured spine neck length, diameter, and volume and found a weak trend between the neck diameter (Figure 1D) and spine head volume but no correlations between neck length (Figure 1E) and spine head volume, consistent with previous studies (Harris and Stevens, 1989; Tønnesen, 2014).

Next, we analyzed spine volumes according to their axonal connectivity and dendrite origin. Pairs of spines on the same dendrite that received input from the same axon ('axon-coupled'), were of the same size and had nearly identical head volumes (Figure 4). We compared this sample of 10 axon coupled pairs on the same dendrite (Figure 4B, pairs a-j) to those identified on dendrites from the two additional animals (Bourne, 2013), for a

total of 17 axon-coupled spine pairs. When plotted against one another, the paired head volumes were highly correlated with slope 0.91, and despite the small sample size, were highly significantly different from random pairings of spines (Figure 4C and Figure 4—figure supplement 1A, KS test p=0.0002). Similarly, there was a strong positive correlation between their paired neck diameters (Figure 4D), PSD areas (Figure 4E), and number of presynaptic docked vesicles (Figure 4F). These features of axon-coupled spines from the same dendrite spanned the distribution of the overall spine population (Figure 3). In contrast, the spine neck lengths (Figure 4G), and neck volumes (Figure 4—figure supplement 1B) of the pairs were not well-correlated indicating that regulation of neck length and neck volume are not important for synaptic strength.

The coupled triplet of synapses (Figure 4C, gray points 'k, l, m') are on three different spines along a single dendrite and receive synaptic input from a single multi-synaptic bouton. A larger central spine between two similar in size (Figure 4B, 'k, l, m') produces one same size pair ('k') and two different size pairs (' l', 'm'). This unusual configuration is probably driven by processes, such as competition for available resources that differ from the other pairs (Sorra and Harris, 1993; Sorra, 1998). As one possibility, perhaps the size of the larger postsynaptic spine was influenced by the larger size of the available pool of presynaptic vesicles in close proximity to its active zone. Excluding this triple synapse, the median value of the coefficient of variation of volume differences between pairs was CV = 0.083 and was as precise for small synapses as it was for large ones (Figure 5). This precision (i.e. low CV) suggests that accurately maintaining the size of every synapse, regardless of size and strength, could be important for the function, flexibility and computational power of the hippocampus.

This near-identical size relationship does not hold for axon-coupled spines on different dendrites (Figure 6B, CV = 0.39, n = 127, example Figure 6A), nor for non-axon-coupled spines on the same or different dendrites (Figures 6E,F, example Figure 6D) – all cases which would have had different activation histories. The volumes of axon-coupled different-dendrite spines are no different from the volumes of random pairs when plotted against one another (KS test p=0.94, Figure 4—figure supplement 2A, and Figures 6B,C) and the distribution of their sizes was no different from the whole population (KS test p=0.41). The number of docked vesicles for pairs on different dendrites (Figure 6—figure supplement 1B) is not different from random pairings (KS test p=0.08), nor are the neck diameters (Figure 6—figure supplement 1C, KS test p=0.06), nor the neck lengths (Figure 6— figure supplement 1D, KS test p=0.75). The size difference of pairs of axon-coupled spines on the same or different dendrites shows a weak trend with separation distance along the axon or dendrite (Figure 6—figure supplement 2). The sizes of pairs of axon-coupled spines on the same or different dendrites is unaffected by proximity of glia processes to the synapses (Figure 7) (Ventura and Harris, 1999; Witcher, 2007), or location of mitochondria in the axon (Billups and Forsythe, 2002).

We found that spine head volumes ranged in size over a factor of 60 from smallest to largest while the CV of any given size was 0.083 and was constant across the range of sizes. Measurements of these of 20 pairs allowed us to estimate the number of distinct spine sizes, and by extension synaptic strengths, that can be reliably distinguished across this range. Signal detection theory holds that at a Signal-to-Noise Ratio (SNR) of 1, a common detection threshold used in psychophysical experiments, an ideal observer can correctly detect whether a signal is higher or lower than some threshold 69% of the time (Green and

Swets, 1966; Schultz, 2007). Put another way, if random samples are drawn from two Gaussian distributions whose areas overlap by 31%, an ideal observer will correctly assign a given sample to the correct distribution 69% of the time. Using this logic, we found that ~26 different mean synaptic strengths could span the entire range, assuming CV = 0.083 for each strength level, and a 69% discrimination threshold (Figure 8, see Materials and methods). These 26 distinct strength levels can be represented with 4.7 bits of information (i.e. 24.7 » 26) which means 4.7 bits of information that can be stored at each synapse as synaptic strength. At a discrimination threshold of 76% (corresponding to SNR = 2) there would be ~23 distinct strengths and 4.5 bits of information.

To explain the high precision observed in spine head volumes, we propose that time window averaging smooths out fluctuations due to plasticity and other sources of variability including differences in the age of the synapses. To set a lower bound on averaging time, we chose to examine neurotransmitter release probability as a single source of variability. Let us first consider release caused by single action potentials, ignoring short-term plasticity. Release of this type can be analyzed using a binomial model in which n presynaptic action potentials, each with a probability pr of releasing one or more vesicles, leads to a mean number of releases = n*pr having variance s2 = n*pr*(1pr). The coefficient of variation around the mean is CV = s/ = sqrt [(1pr)/(n*pr)] and can be compared with the measured values. Therefore, the number of spikes that are needed to reduce the variability to achieve a given CV is n=(1pr)/(pr*CV2). Table 1 gives averaging time windows T = n/R, where R is spiking rate of the presynaptic axon, for representative values of pr and a range of spiking rates.

Accounting for other known sources of variability at dendritic spines would require even longer time windows. In particular, the impact of short-term plasticity during bursts of action potentials on the length of the time-window is complicated by the interplay of facilitation and depression. Synapses with a low initial pr (and corresponding long time-window in Table 1) exhibit marked facilitation and slowly depress during bursts (Kandaswamy, 2010; Nadkarni, 2010) which would shorten the time-window. But synapses with a high initial pr (and short time-window) only weakly facilitate, if at all, and quickly depress, which would lengthen the time-window.

Discussion

Previous upper bounds on the variability of spine volume in the hippocampus, based on the whole spine volume (Sorra and Harris, 1993; O'Connor, 2005), underestimated the precision by including the spine neck volume (Figure 4—figure supplement 1A), which was not correlated between pairs of spines in our volume (Figure 4—figure supplement 1B). Our dense reconstruction included a complete inventory of every synapse in the reconstructed volume and in this respect was unbiased. Additional pairs of synapses from two other rats confirmed that this finding is not confined to a single brain. Of course, additional measurements in the hippocampus and other brain regions would be needed to confirm and extend this finding. The very high statistical significance of the finding (Figure 4C, KS test p=0.0002) despite a relatively small number of pairs in our sample implies a large effect magnitude, which would be much smaller if many more samples were needed to reach the same level of significance. To make this p value concrete,

if 17 random pairs were chosen from all 287 synapses in the reconstructed volume, there is only a one in 5000 chance that the spine heads would be as precisely matched as the 17 axon-coupled pairs discovered here.

Previous studies have shown that there is a high correlation of the size of the spine head with the PSD area and numbers of docked vesicles (Harris and Stevens, 1989; Lisman and Harris, 1994; Harris and Sultan, 1995; Schikorski and Stevens, 1997; Murthy, 2001; Branco, 2008; Bourne, 2013). Since the correlations between the head sizes of axon-coupled pairs of spines is high, the high correlation between the PSD areas and numbers of docked vesicles observed in axon-coupled spines is not surprising (Figures 4E and 4F). However, it was unexpected to find that the spine neck diameters were also highly correlated between axon-coupled pairs of spines (Figure 4D r2=0.70), since the correlation between spine head volumes and spine neck diameters is not statistically significant (Figure 1D). Thus, there are at least two geometric aspects of the spine geometry that are under tight control of synaptic plasticity, which may reflect different aspects of synaptic function. The diameter of the spine neck may reflect the need for trafficking of materials between the spine shaft and spine head, which is known to be regulated by LTP and LTD (Araki, 2015).

Complementing our observations and analysis in the hippocampus, highly correlated pr at multiple contacts in the neocortex between the axon of a given layer 2/3 pyramidal neuron and the same target cell has been reported (Koester and Johnston, 2005). Our estimate of synaptic variability, based on spine head volume, is an order of magnitude lower. In a recent connectomic reconstruction of the mouse cortex, the similarity in the volumes of axon-coupled pairs of dendritic spines were statistically significant (Kasthuri,

2015). This observation is further evidence for the high precision of synaptic plasticity and suggests that the same may be true in other brain areas.

The axon-coupled pairs of synapses that we studied were within a few microns of each other on the same dendrite, which raises the question of how far apart the two synapses can be and still converge to the same size. Related to this question, two synapses from the same axon on two different dendrites of the same neuron might not share the same postsynaptic history. These questions cannot be answered with our current data due to the small dimensions and the fact that the position in the neuropil from which our reconstruction was taken makes it highly unlikely that any of the dendrites, other than the one branch point captured in the volume, belong to the same neuron (Ishizuka, 1995). Synaptic tagging and capture, in which inputs that are too weak to trigger LTP or LTD can be 'rescued' by a stronger input to neighboring synapses if it occurs within an hour (Frey and Morris, 1997; O'Donnell and Sejnowski, 2014), is much less effective when the synapses are on different branches (Govindarajan, 2011), which would tend to make two synapses from the same axon on different dendritic branches less similar. Probing these questions will require reconstructing a larger extent of hippocampus when a single axon can contact multiple dendritic branches of the same neuron (Sorra and Harris, 1993) or of other cells, such as layer 5 pyramidal cells, which can have 4–8 connections between pairs of neurons (Markram, 1997).

An unusual triple synapse from a single axon (Figure 4B, 'k, l, m') was excluded from the analysis because the presynaptic terminal was a single large varicosity filled with vesicles (i.e. an MSB) shared by three synapses, unlike the other pairs that had isolated presynaptic specializations (n=9), or an MSB shared by two synapses (n=8). It is possible

that the large, central spine had an effectively larger pool of vesicles by virtue of proximity, whereas the two synapses on the outside had a more limited population to draw from, and the size of the postsynaptic spine was influenced by the size of the available pool. More examples are needed before we can reach any conclusions. Regardless of the explanation, our estimate of the variability would not be greatly affected by including these 3 additional pairs of synapses in the analysis.

How can the high precision in spine head volume be achieved despite the many sources of stochastic variability observed in synaptic responses? These include: 1) The low probability of release from the presynaptic axon in response to an action potential (Murthy, 2001); 2) Short-term plasticity of release of neurotransmitter (Dobrunz, 1997); 3) Stochastic fluctuations in the opening of postsynaptic NMDA receptors, with only a few of the 2–20 conducting at any time (Nimchinsky, 2004); 4) Location of release site relative to AMPA receptors (Franks, 2003; Ashby, 2006; Kusters, 2013) 5) Few voltage-dependent calcium channels (VDCCs) in spines that affect synaptic plasticity (smallest spines contain none) (Mills, 1994; Magee and Johnston, 1995); 6) VDCCs depress after back propagating action potentials (Yasuda, 2003); 7) Capacity for local ribosomal protein synthesis in some spines while others depend on transport of proteins from the dendrites (Ostroff, 2002; Sutton and Schuman, 2006; Bourne, 2007; Bourne and Harris, 2011); 8) Homeostatic mechanisms for synaptic scaling may vary (Turrigiano, 2008; Bourne and Harris, 2011); 9) Presence or absence of glia (Ventura and Harris, 1999; Witcher, 2007; Clarke and Barres, 2013); and 10) Frequency of axonal firing (Callaway and Ross, 1995).

One way that high precision can be achieved is through time averaging. Long-term changes in the structure of the synapse and the efficacy of synaptic transmission are

triggered by the entry of calcium into the spine. A strategy for identifying the time-averaging mechanism is to follow the calcium. Phosphorylation of calcium/calmodulin-dependent protein kinase II (CaMKII), required for spike-timing dependent plasticity processes, integrates calcium signals over minutes to hours and is a critical step in enzyme cascades leading to structural changes induced by long-term potentiation (LTP) and long-term depression (LTD) (Kennedy, 2005), including rearrangements of the cytoskeleton (Krama ́r, 2012). The time window over which CaMKII integrates calcium signals is within the range of time windows we predict would be needed to achieve the observed precision (Table 1). Similar time windows occur in synaptic tagging and capture, which also requires CaMKII (Redondo and Morris, 2011; de Carvalho Myskiw, 2014). These observations suggest that biochemical pathways within the postsynaptic spine have the long time scales required to record and maintain the history of activity patterns leading to structural changes in the size of the spine heads.

The information stored at a single synapse is encoded in the form of the synaptic strength, which reflects the pre- and postsynaptic history experienced by the synapse. But due to the many sources of variability, this information cannot be read out with a single input spike. This apparent limitation may have several advantages. First, the stochastic variability might reflect a sampling strategy designed for energetic efficiency since it is the physical substrate that must be stable for long-term memory retention, not the read out of individual spikes (Laughlin and Sejnowski, 2003). Second, some algorithms depend on stochastic sampling, such as the Markov Chain Monte Carlo algorithm that achieves estimates by sampling from probability distributions, and can be used for Bayesian inference (Gamerman and Lopes, 2006). Each synapse in essence samples from a

probability distribution with a highly accurate mean, which collectively produces a sample from the joint probability distribution across all synapses. A final advantage derives from the problem of overfitting, which occurs when the number of parameters in a model is very large. This problem can be ameliorated by 'drop out', a procedure in which only a random fraction of the elements in the model are used on any given trial (Wan, 2013; Srivastava, 2014). Drop out regularizes the learning since a different network is being used on every learning trial, which reduces co-adaptation and overfitting.

We are just beginning to appreciate the level of precision with which synapses are regulated and the wide range of time scales that govern the structural organization of synapses. The upper bound on the variability that we have found may be limited by errors in the reconstruction and could be even lower if a more accurate method could be devised to compute the volume of a spine head, neck diameter, PSD area, number of docked vesicles, or other salient features of dendritic spines. Much can be learned about the computational resources of synapses by exploring axon-coupled synaptic pairs in other brain regions and in other species.

Materials and Methods

Reconstruction of Neuropil

Three separate 3DEM data sets were used in this study. Each of these data sets has been used for other purposes in prior studies. Images were obtained from serial thin sections in the middle of stratum radiatum of hippocampal area CA1 from three adult male rats (55–65 days old) (Mishchenko, 2010; Bourne, 2013). One set of images was used to make a dense model of 6 x 6 x 5 mm3 of hippocampal neuropil and processed as previously

described in a study of the extracellular space (Kinney, 2013). In this data set, we identified 13 axon-coupled synaptic pairs on 11 dendrites (Figure 4—figure supplement 3 and Figure 4—figure supplement 4). The other two sets of images were part of a prior study (Bourne, 2013) in which subsets of dendrites and axons had been reconstructed. In this data set, we identified 7 axon-coupled synaptic pairs on 4 dendrites for a total of 20 axon-coupled dendrite-coupled spine pairs. To perform an accurate and robust geometric analysis of the dendrites, dendritic spines, axons, and glial processes, it was necessary to correct the reconstructed surface meshes for artifacts and make them into computational-quality meshes as described elsewhere (Kinney, 2013; Edwards, 2014).

The postsynaptic densities (PSDs) and presynaptic active zones (AZs) were identified in the serial section transmission electron microscopy (ssTEM) images by their electron density and presence of closely apposed presynaptic vesicles. We devised a method to segment the PSD-AZ features in the electron micrographs and mark their pre- and post-synaptic locations as subregions of the membrane in the final 3D mesh. To accomplish this, contours were hand-drawn on each serial section micrograph closely encompassing, as a single closed contour, the pre- and post-synaptic extent of the electron dense region. Taken together, the stack of contours for a given PSD-AZ forms a 3D capsule which encloses the entire feature. VolRoverN (Edwards, 2014) was used to reconstruct the 3D surface of the capsule enclosing each PSD-AZ pair in 3D. Because these capsules enclose the intracellular domain of both the PSD and AZ they also overlap with the pre- and post-synaptic membrane associated with these subcellular features. Each of these closed capsules was then used as a '3D lasso' to tag mesh triangles of the pre- and post-synaptic membrane contained within the lasso, marking the enclosed membrane area as a

synaptic contact region—PSD postsynaptically and AZ presynaptically. Figure 3—figure supplement 1A shows a postsynaptic contact area labeled in red on a dendritic spine.

The reconstructed neuropil models were then visualized and analyzed using Blender, a free, open-source tool for 3D computer graphics modeling (http://blender.org). A total of 449 synaptic contacts were found in the dense reconstructed volume of neuropil. We excluded a number of synapses from the analysis if they were partially clipped by the edge of the data set (142), or were shaft synapses (20) leaving 287 valid synapses on dendritic spines in the dense model. An additional 70 spines were excluded from the analysis of axon-coupled spines as the axon which contacted these spines did not contact any other spines within the reconstructed volume. Example visualizations of the spines and axons, generated using Blender, are shown in Figure 2A,B, Figure 3—figure supplement 1A, Figure 4A,B, Figure 4—figure supplement 3, Figure 4—figure supplement 4, Figure 6A,D, Figure 7B,D,F,H.

Segmentation of dendritic spines

Blender's functionality is user-extensible via a Python interface for creating add-ons. We created a Python add-on for Blender that enabled the selection of the mesh triangles of the dendrite corresponding to the spine head and whole spine of each individual spine. Our add-on tagged each selected set of triangles with metadata for the spine name and geometric attributes of the head, whole spine, and neck as described below. The selection of the spine head was made by hand based on a standardized procedure in which the junction between the head and neck was visually identified as half-way along the concave arc as the head narrows to form the neck (see Figure 3—figure supplement 1A).

To select the whole spine, a similar visual judgment was made to locate the junction where the neck widens as it joins the dendritic shaft.

Once the appropriate area was selected, the tool was designed to automatically create the convex hull of the selected region. The closed mesh formed by the Boolean intersection of the convex hull and the dendrite was used to determine the measured volume of the spine head or whole spine. The volume of the neck was calculated by taking the difference between these two measurements. Areas were computed from the selected regions for spine head and whole spine. Active zone and postsynaptic density areas were calculated using regions that had been determined during the hand-drawn reconstruction phase described above. Distances between spine heads along the axon were calculated as the Euclidean distance between the centroids of the PSD/AZ regions. Distances between whole spines along the dendritic shaft were calculated as the Euclidean distance between the spine necks to shaft junctions. Glial classification, mitochondria classification and shape classification were performed by hand using set criteria.

Estimation of measurement error of spine head volume

Some error in the measurement of spine head volume is expected to occur in the human judgment required to segment the dendritic spines into whole spine, head, and neck. To estimate this error, the valid spines in the dense model were segmented and measured a total of four times per spine (twice each by two people). The standard error of the mean in spine head volume decreases with volume and is less than 5% for the majority of spines with a median error of about 1% (Figure 3— figure supplement 1). The head volumes in the other two data sets were only measured once.

Segmentation of synaptic vesicles and estimation of docked vesicles

Synaptic vesicles in the presynaptic terminals, totaling 31,377 in number, were identified along with their 3D locations within the dense reconstruction. Of the 449 presynaptic terminals, we excluded 193 terminals from the analysis due to truncation at the edge of the volume, and 20 terminals at shaft synapses, leaving 236 valid terminals. A visualization of all the synaptic vesicles in the reconstruction is shown in Figure 2A.

Positive identification of docked vesicles in these ssTEM data sets is problematic due to the thickness of the sections and density of the staining. To estimate docked vesicles, we counted the number of vesicles whose centers were located within 100 nm of the presynaptic membrane across from the postsynaptic density of a given spine. Of the 31,377 vesicles, 3437 were labeled as docked according to this criterion which yielded estimates in good agreement with previous estimates (Harris and Sultan, 1995; Schikorski and Stevens, 1997; Figures 2B–D). An en face view of the docked vesicles at one synapse is shown in Figure 2B.

Statistical analysis

All statistical analysis and plots were generated using Python 2.7 (http://python.org) with NumPy, SciPy, and Matplotlib. The distributions of spine head volume, spine head area, spine neck volume, PSD area, and AZ area were highly skewed with a long tail at larger values (Figure 1). Consequently, all regression analysis was performed using Pearson's linear regression on the data after applying a log-normal transformation (r2 values shown in Figures 1–6).

The coefficient of variation (CV) of the population of spine pairings (Figures 4 and 6) was calculated as the median value of the CVs of each individual pair. The CV of each individual pair is simply the standard deviation of the volumes of the pair divided by the mean volume of the pair (Figure 5).

Population distributions were highly skewed making it necessary to make comparisons of distributions using non-parametric methods. We used the two-sample Kolmogorov-Smirnoff (KS) test to make these comparisons in Figures 4 and 6.

Estimation of number of distinguishable spine sizes and bits of precision in spine size

To estimate the number of distinguishable spine sizes and corresponding bits of precision we calculated the number of distinct Gaussian distributions of spine sizes, each with a certain mean size and standard deviation that together would cover and span the entire range of spine head sizes seen in Figure 4A. Figure 5 demonstrates that it is reasonable to assume that the CV of each these sub-distributions is a constant value of 0.083. From this CV, the spacing between the mean values of each sub-distribution can be chosen to achieve a total of 31% overlap with adjacent sub-distributions giving a 69% discrimination threshold. A 69% discrimination threshold is commonly used in the field of psychophysics and corresponds to a Signal-to-Noise Ratio (SNR) of 1 (Green and Swets, 1966; Schultz, 2007).

The 69% confidence interval, z, of a Gaussian distribution is given by:

$z = sqrt(2) * erf^{-1}(0.69)$

The spacing, s, of adjacent intervals of mean, , is given by:

$s = \mu * 2 * CV * z$

The number, N, of such distributions that would span the factor of 60 range of spine sizes is:

$N = log(60) / log(1 + 2 * CV * z)$

$N = 26.3$

The number of bits of precision implied by N distinguishable distributions is given by:

$bits = log_2(N)$

$bits = 4.72$

Figure 8 shows that ~26 distinguishable distributions can cover the entire range of spine sizes, implying that there are ~4.7 bits of precision in the spine size.

All data and software tools described here are available at:

http://www.mcell.cnl.salk.edu/models/hippocampus-spine-analysis-2015-1

Acknowledgments

Chapter 2, in full, is a reprint of the material as it appears in eLife, 2015. Bartol T.M., **Bromer C**., Kinney J., Chirillo M.A., Bourne J.N., Harris K.M. & Sejnowski T. J. The dissertation author was a primary investigator and co-author of this paper.

Figure 2.1: Correlations among metrics of dendritic spine morphology. Strong correlations were found between (A) Spine head area and spine head volume, (B) PSD area and spine head volume, and (C) Spine head area and PSD area. (D) Weak correlation was found between spine neck diameter and spine head volume. No correlation was found between (E) spine neck length and spine head volume and (F) spine neck diameter and spine neck length. Regression lines in red and error bars for each data point represent SEM based on multiple tracers who also edited each spine. Equations are based on the log-log distributions, with r2 values indicated, and n=287 complete spines.

Figure S2.1-1: Area of postsynaptic density plotted against spine head volume. Nine individual dendrites all have similar slopes that are not significantly different sho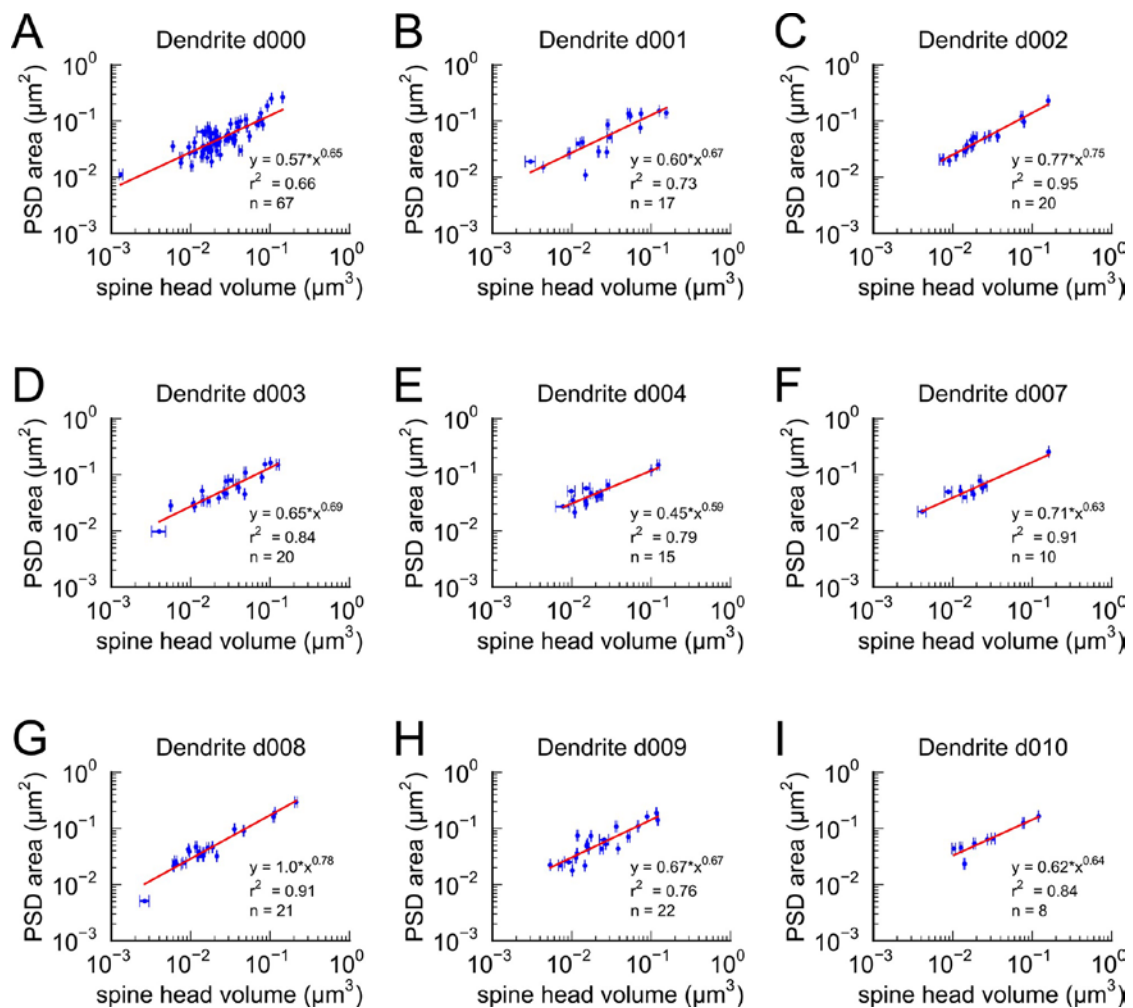wing the uniformity of this comparison across dendrites. Error bars, regression lines and equations as described in Figure 1.
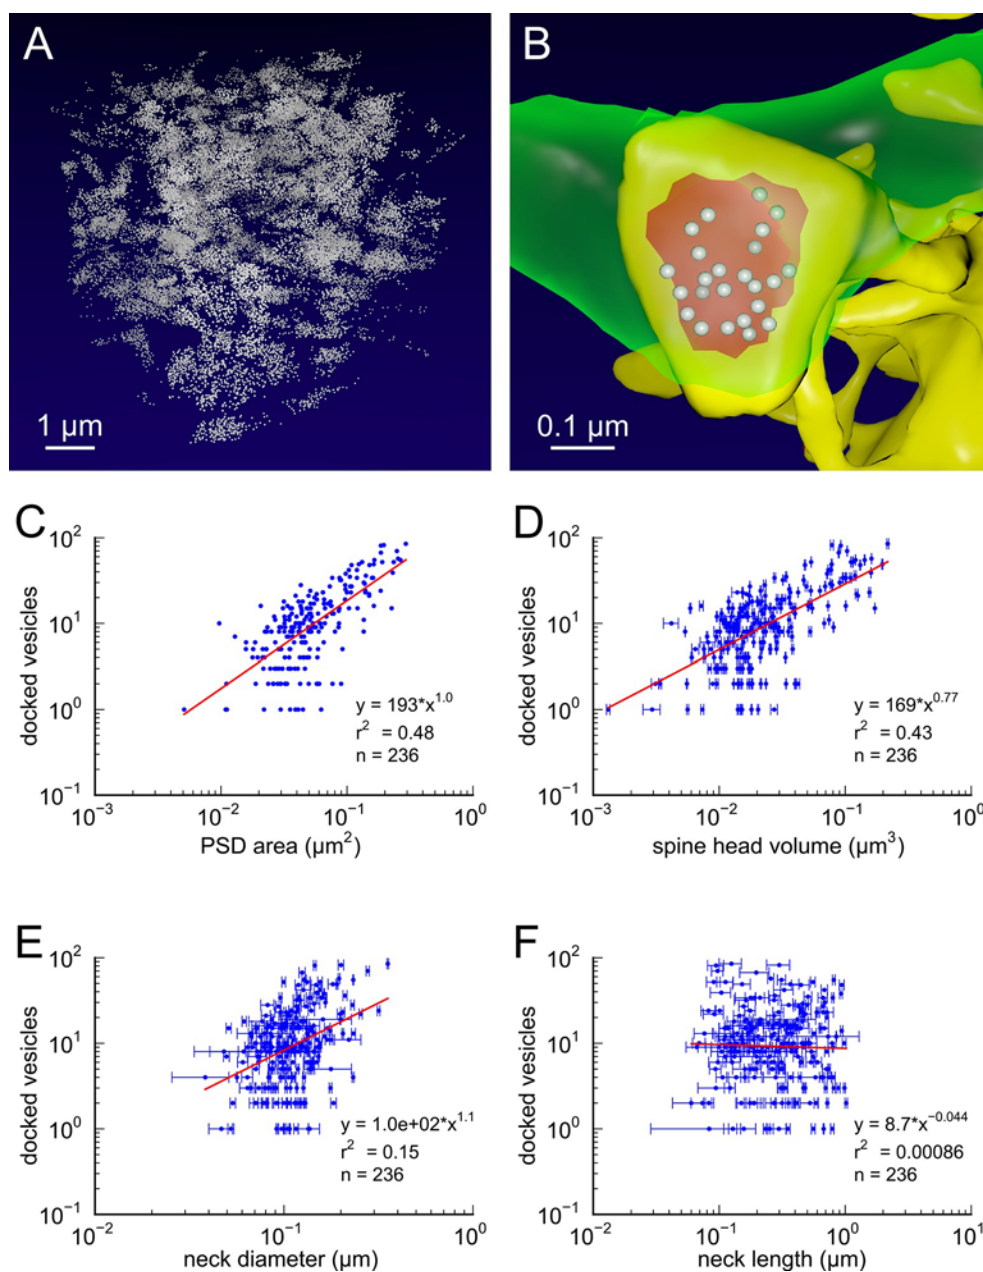
Figure 2.2: Presynaptic docked vesicle numbers are correlated with PSD areas, spine head volumes, and neck diameter, but not with neck length. (A) All 31,377 presynaptic vesicles. (B) En face view of the 24 docked vesicles (gray spheres) viewed through an axon (green) onto the PSD (red) of example spine (yellow). (C) Number of docked vesicles is correlated strongly with both PSD area and (D) spine head volume, weakly with (E) neck diameter, but is not correlated with (F) spine neck length. Regression lines, SEM (from multiple tracers), and r2 are as in Figure 1 n = 236 complete axonal boutons, each associated with one of the 287 complete spines. One human tracer marked PSDs and vesicles, hence no SEM for these two metrics
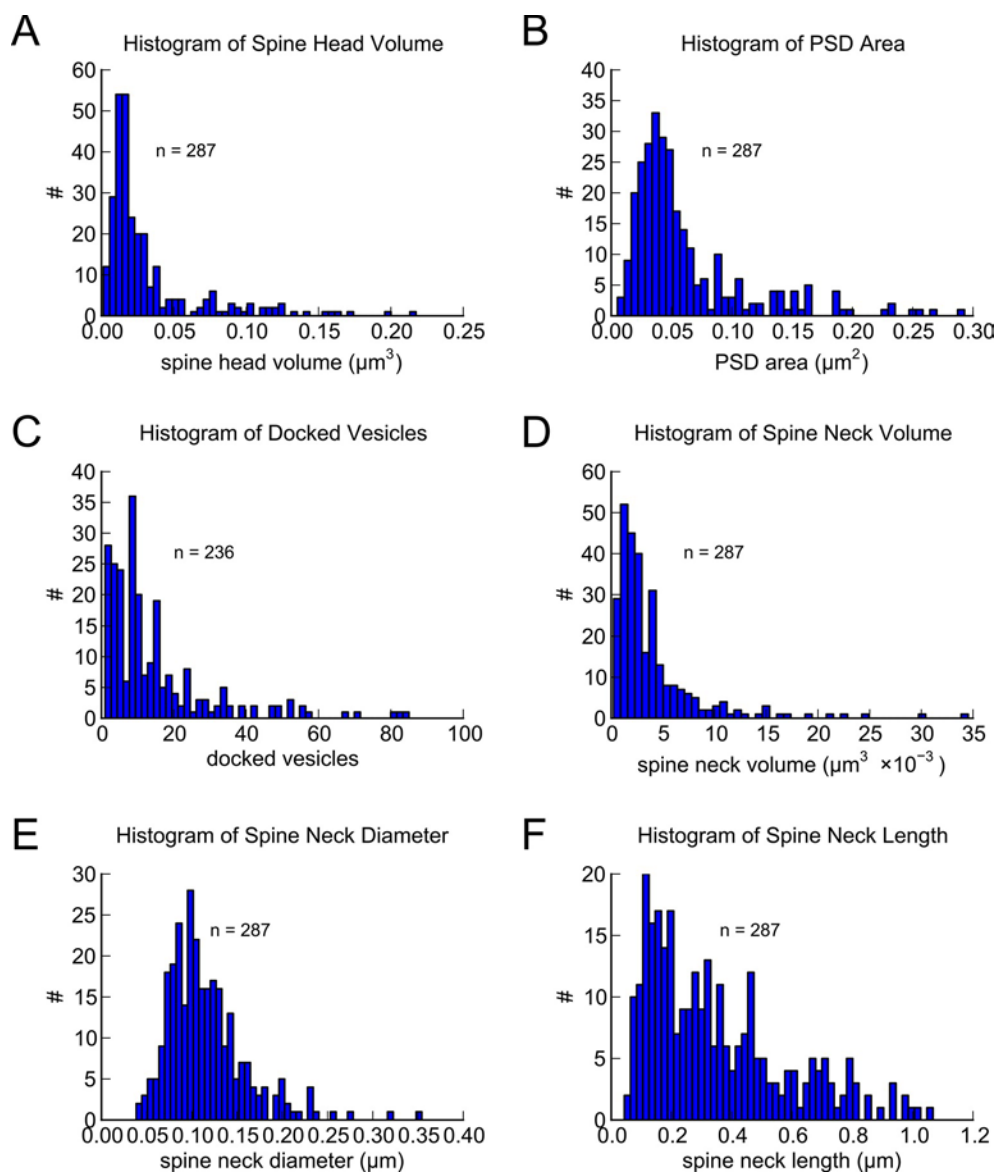
Figure 2.3: Morphometric analysis of 287 complete spines in reconstruction. Distributions of (A) spine head volumes, (B) PSD areas, (C) docked vesicles, (D) spine neck volumes, (E) spine neck diameters, and (F) spine neck lengths are highly skewed with a long tail.
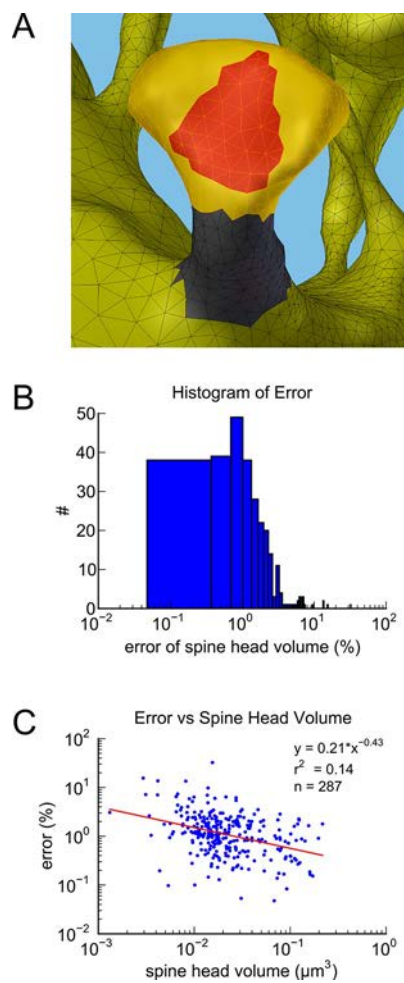
Figure S2.1-1: Spine measurement and estimation of measurement error. (**A**) Example segmentation of spine head (yellow), neck (gray), and PSD area (red). (**B**) Histogram of the measurement error across all spines measured. (**C**) Measurement error plotted against spine head volume.
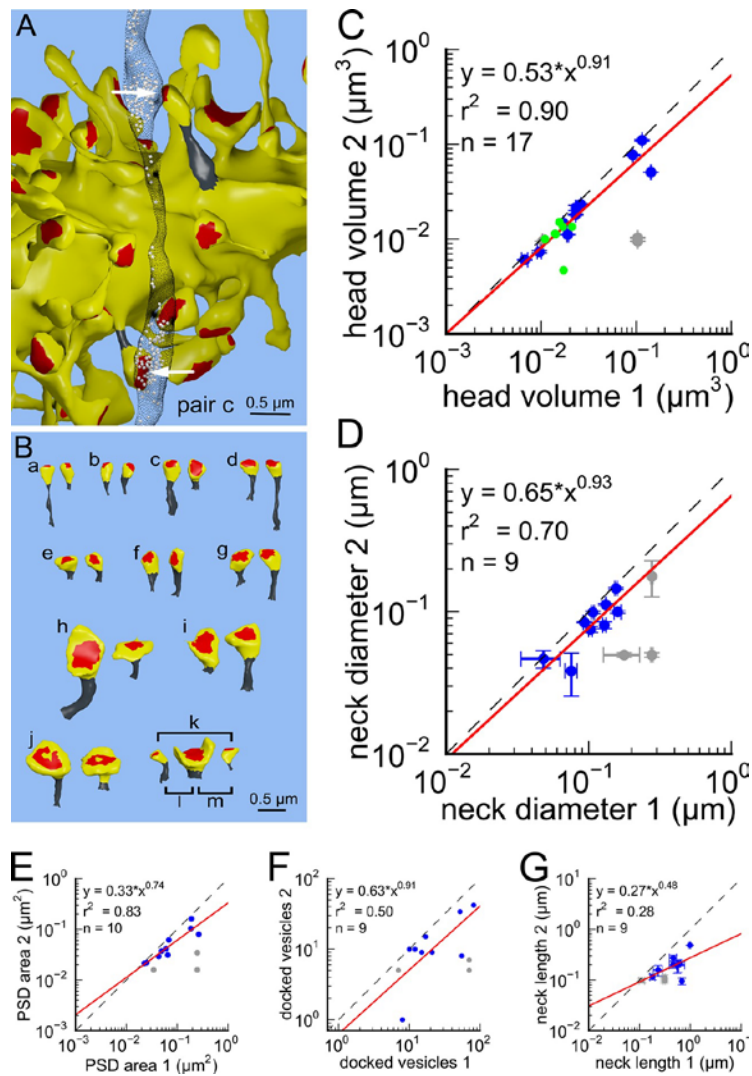
Figure 2.4: Spine head volumes, PSD areas and neck diameters, but not neck lengths, are highly correlated between pairs of axon-coupled same-dendrite spines. (A) Visualization of a pair of spines (gray necks) from the same dendrite (yellow) with synapses (red, indicated by white arrows) on the same axon (black stippling) with presynaptic vesicles (white spheres). (B) All axon-coupled same-dendrite spine pairs (colors as in A, pair c is elaborated in A). Strong correlations with slopes near 1 (dashed diagonal line) occur between paired (C) spine head volumes (slope = 0.91), (D) neck diameters (slope = 0.93), (E) PSD areas (slope = 0.74), and (F) docked vesicles (slope = 0.91); but not (G) spine neck lengths (slope = 0.48). Larger values from each pairing are plotted on the X axis. Regression lines (red) include the 10 a-j pairings (blue points) and 7 pairs from 2 additional animals (green points in (C)), but do not include triplet bouton pairings (k-m, gray points).
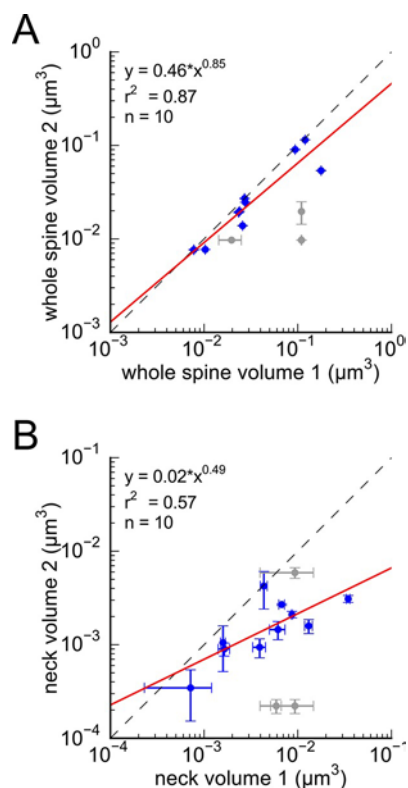
Figure S2.4-1: Analysis of whole spine volume and spine neck volume of axon-coupled same dendrite spines.(A) Whole spine volumes of pairs of axon-coupled spines on the same dendrite are highly correlated and significantly different from random pairs (KS test $p = 0.018$). (B) Correlation of neck volumes of pairs of axon-coupled spines on the same dendrite are not significantly different from random pairs (KS test $p = 0.74$). Larger value in each pair is plotted on the X axis. Regression lines shown in red. Equations are based on regression of log-log distributions, with r^2 values indicated.
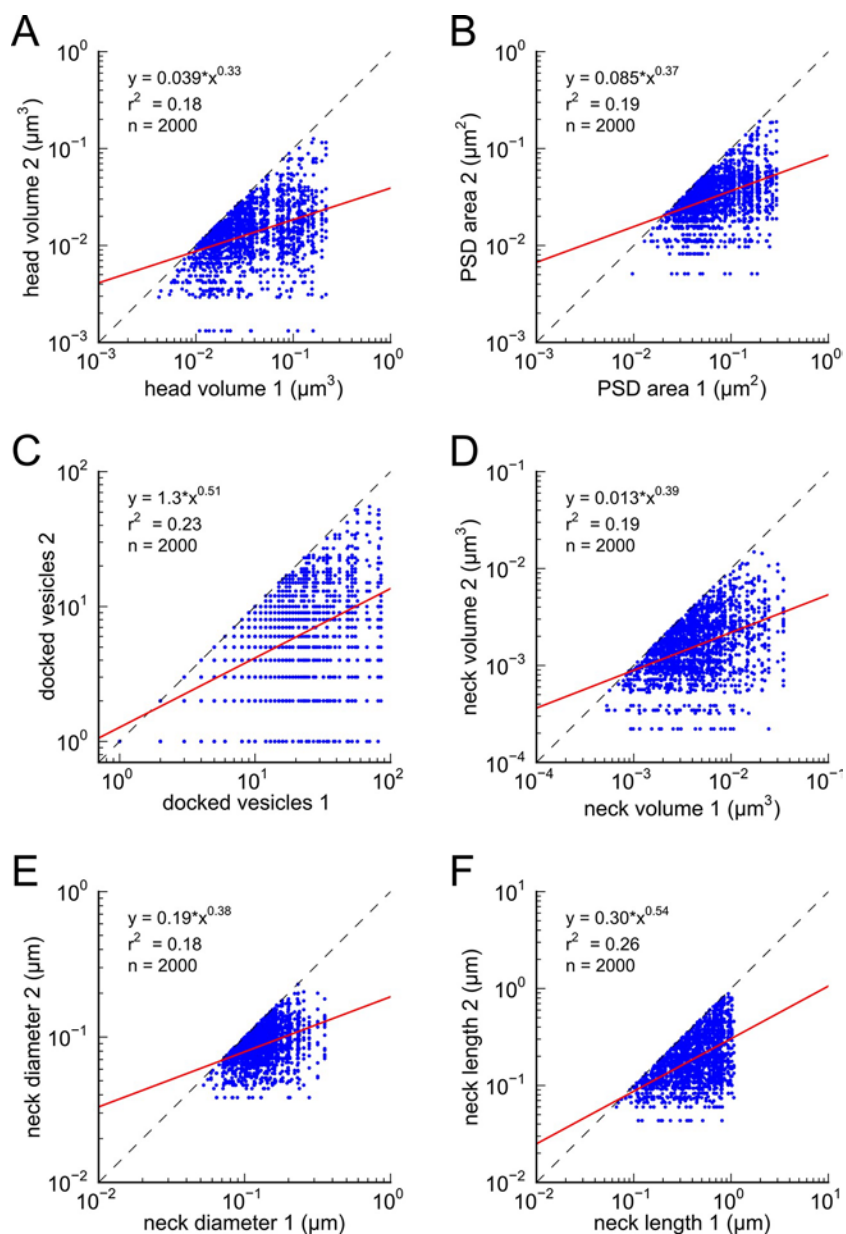
Figure S2.4-2: Analysis of spines paired randomly. Distributions represent random pairings of (A) spine head volumes, (B) PSD areas, (C) docked vesicles, (D) neck volumes, (E) neck diameters, and (F) neck lengths, from the population of complete spines in the reconstruction. Larger value in each pair is plotted on the X axis. Regression lines shown in red. Error bars for each data point are not shown for clarity. Equations are based on regression of log-log distributions, with r^2 values indicated.
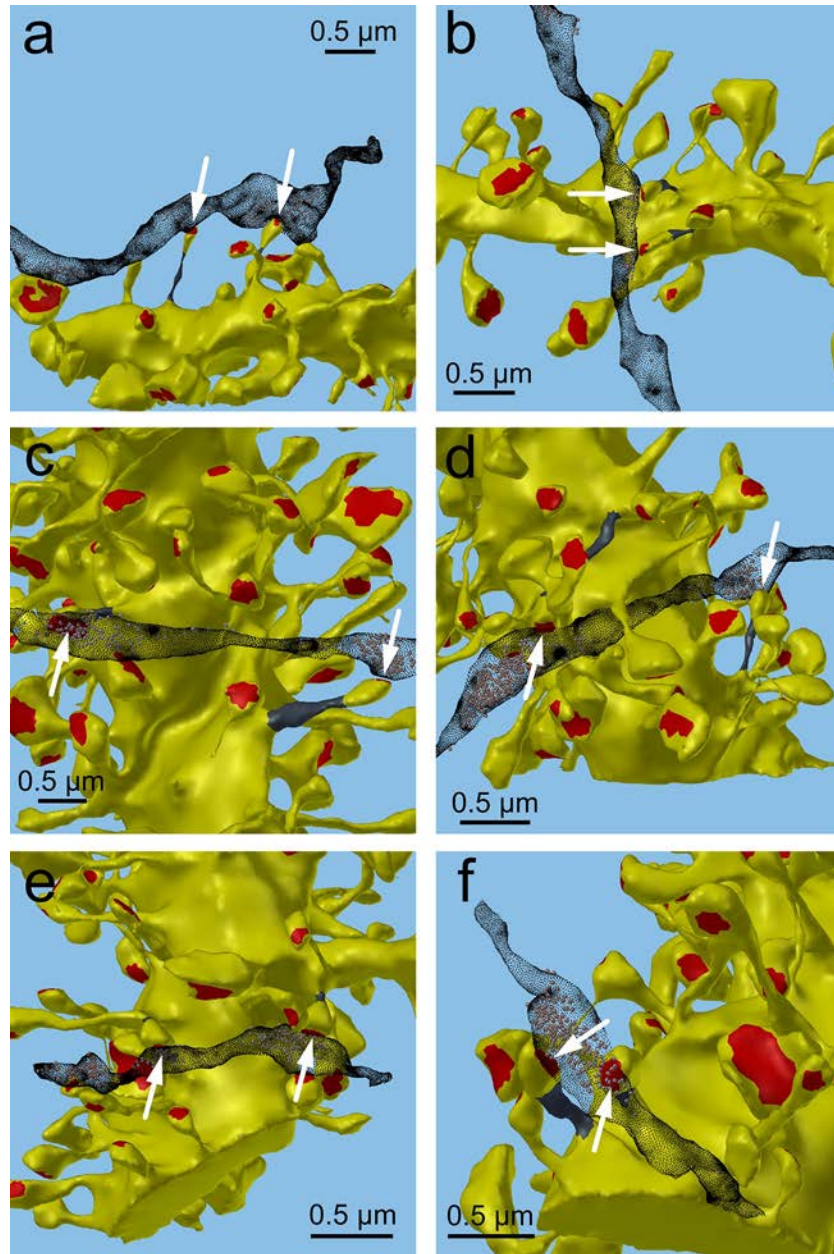
Figure S2.4-3: Axon-coupled same dendrite pairs a–f**.** Large white arrows indicate the red PSDs of the spine pairs, the edited necks are dark gray, and the axons are stippled black with vesicles inside. These illustrate how the axon weaves through the neuropil, synapses with two spines yet passes by others
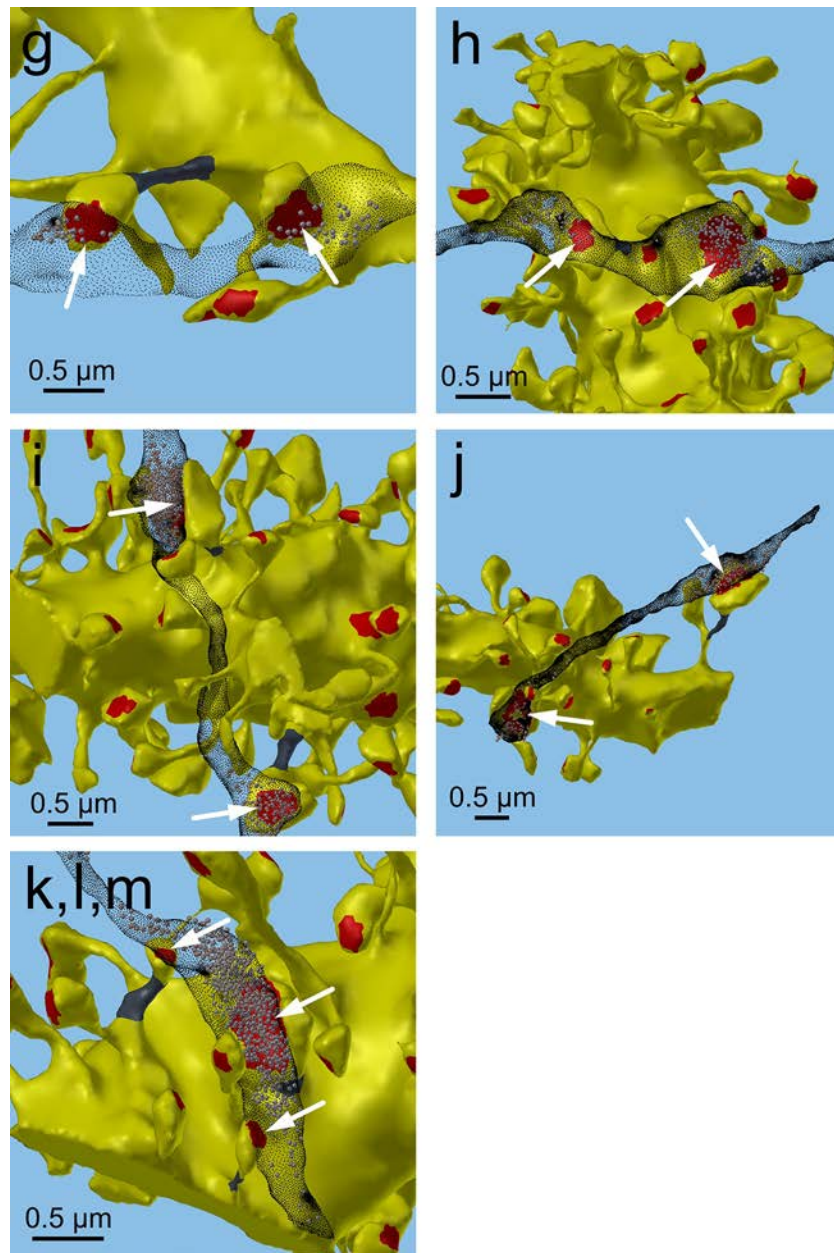
Figure S2.4-4: Axon-coupled same-dendrite pairs g-m, illustrated in same way as in Figure 4 (S3)
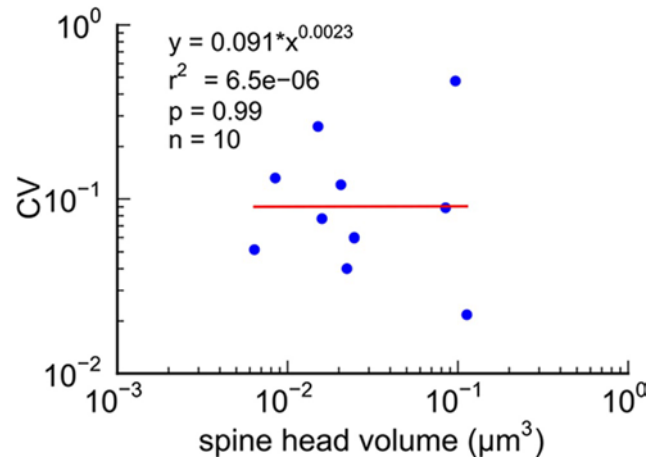
Figure 2.5: CV of axon-coupled spines on the same dendrite does not vary with spine size. There is no significant correlation, which implies that paired small synapses are as precisely matched as paired large synapses.
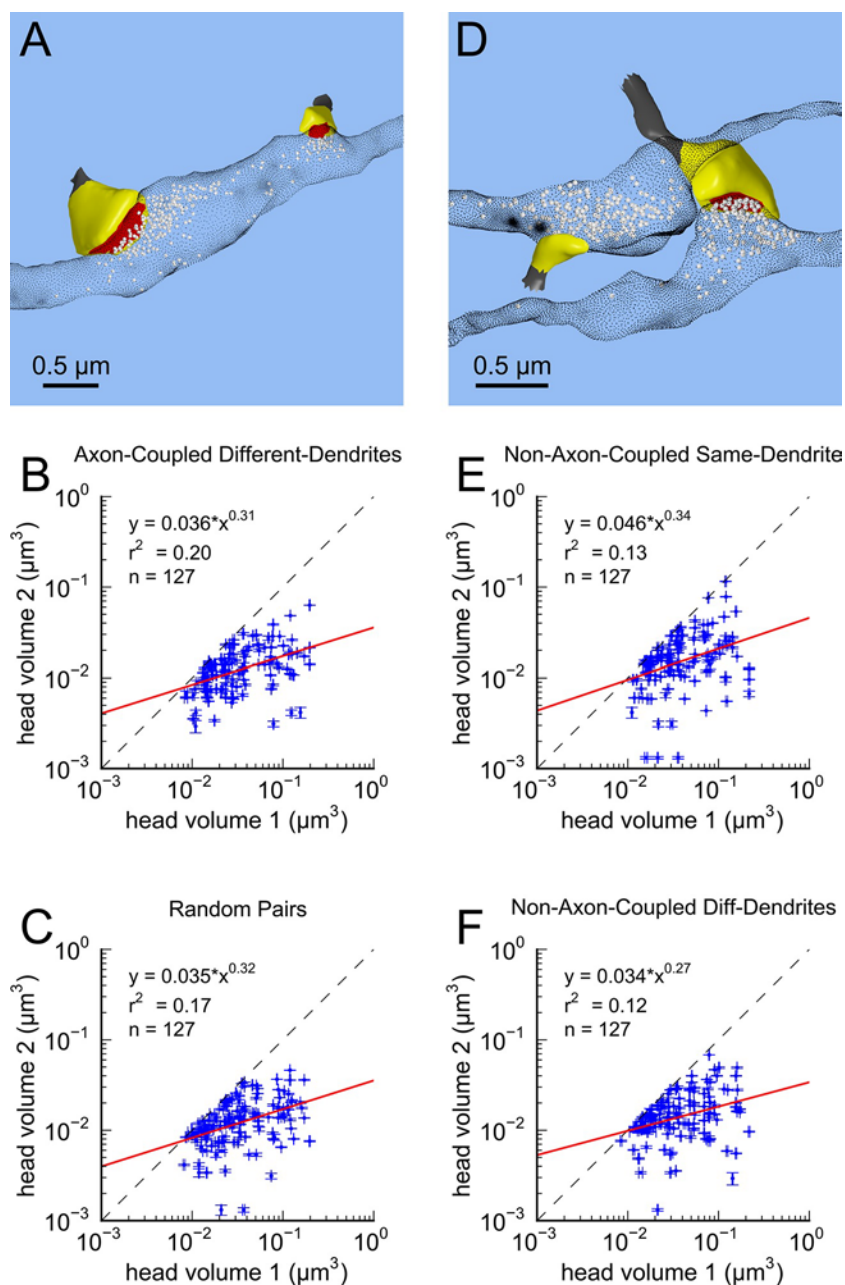
Figure 2.6: Paired spine head volumes are not correlated when they are not both axon and dendrite coupled. (A) Representative visualization and (B) plot showing lack of correlation between spine head volumes of all pairs of axon-coupled spines on different dendrites (n=127). (C) Similarly, randomly associated pairs of spine head volumes were not correlated. (D) Representative visualization and plots show lack of correlation between spine head volumes from randomly selected pairs (n=127) of non-axon-coupled spines (E) on the same or (F) different dendrites. Color scheme and regression analyses as in Figure 4.
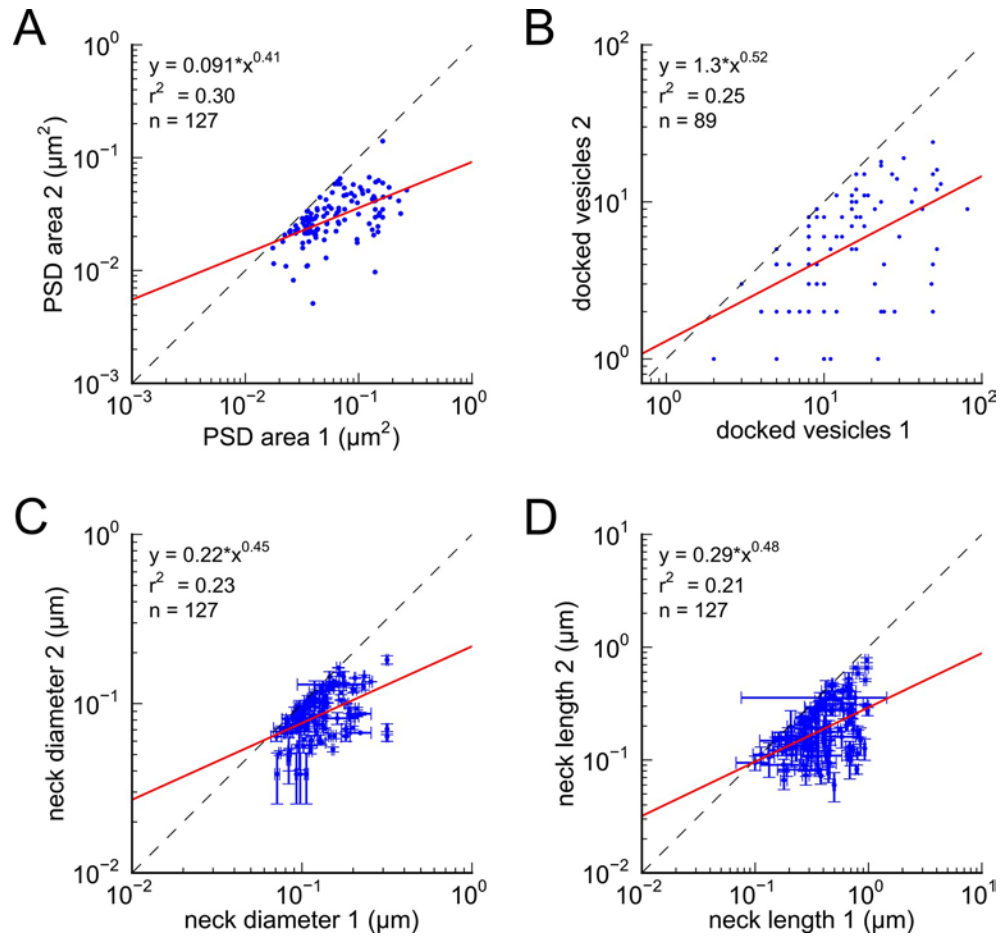
Figure S2.6-1: Morphologies of PSD, docked vesicles, and necks are not correlated when spines are not both axon and dendrite coupled. There is no correlation between (**A**) PSD areas, (**B**) docked vesicles, (**C**) neck diameters, and (**D**) neck lengths, in pairs of axon-coupled spines on different dendrites.
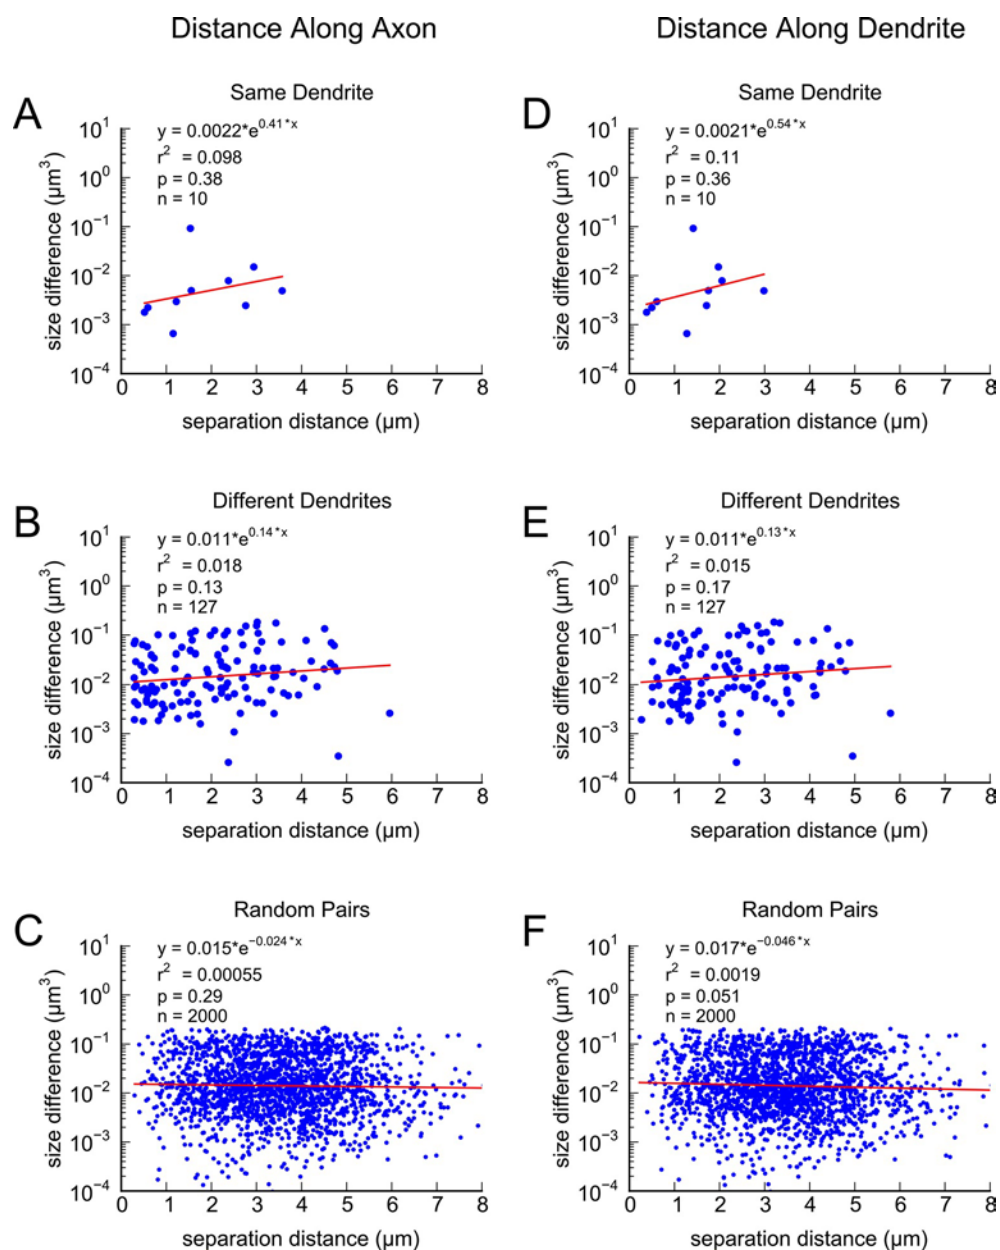
Figure S2.6-2: Difference in volume between pairs of axon-coupled spines exhibits a weak trend with separation distance. Differences in spine head volumes plotted against: **A**) Distance along the axon for axon-coupled spines on the same dendrite; **B**) Distance along the axon for axon-coupled spines on different dendrites; **C**) Distance along the axon for randomly paired spines; **D**) Distance along the dendrite for axon-coupled spines on the same dendrite; **E**) Distance along the dendrite for axon-coupled spines on different dendrites; **F**) Distance along the dendrite for randomly paired spines
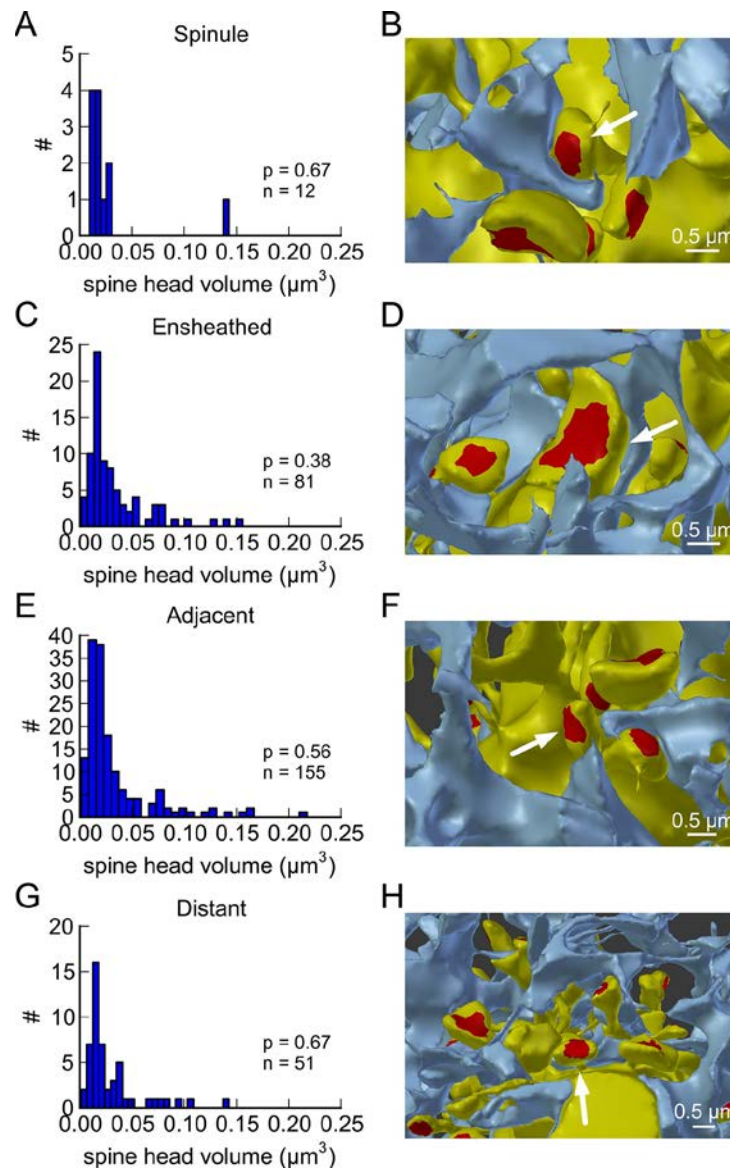
Figure 2.7: Proximity of the glial cell to axon-coupled dendritic spines on either the same or different dendrites. Proximity of astrocytic glial processes is not significantly correlated with spine head volumes of axon coupled pairs. (A) Histogram of spine head volume for spines that contain a spinule that is engulfed within the glial process (' spinule'). (B) Representation of an engulfed spinule. (C) Histogram of spine head volume for spines that are surrounded by and making contact with a glial process ('ensheathed'). (D) Representation of 'ensheathed' spine. (E) Histogram of spine head volume for spines that are proximal but not contacting a glial process ("' adjacent"'). (F) Representation of "'adjacent"' spine. (G) Histogram of spine head volume for spines that are distant from any glial process. (H) Representation of a spine "'distant"' from the glial process. The KS p value is shown on each inset and indicates that none of these distributions differ from the distribution for the whole population of spines.
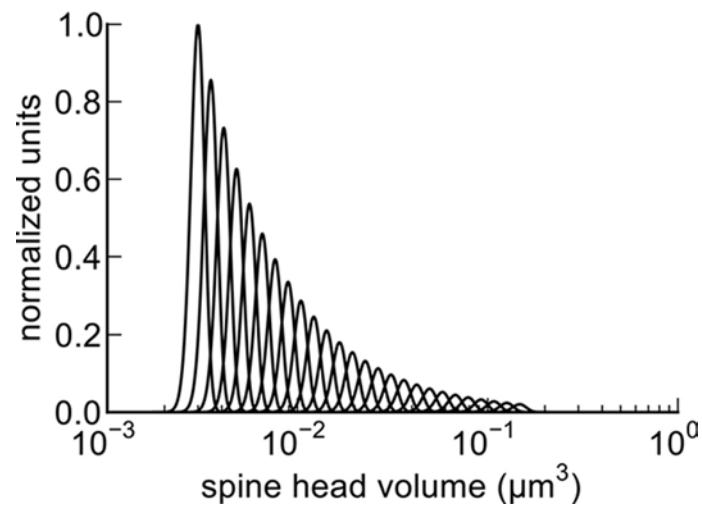
Figure 2.8: Distinguishable spine sizes. Over the factor of 60 range in spine head volumes from the data set there are 26 distinguishable intervals of spine sizes with a discrimination probability of 69% for each interval based on signal detection theory (Green and Swets, 1966; Schultz, 2007). The graph illustrates how distinct Gaussian distributions of spine sizes, each with a certain mean size and standard deviation, covers the entire range of spine head sizes on a log scale. The CV of each distribution is a constant value of 0.083 (Figure 5) and the intervals are spaced to achieve a total of 31% overlap with adjacent intervals giving a 69% discrimination threshold (see Materials and Methods). Note that the constant CV observed in the data set (Figure 5) means that the intervals appear uniform in width and spacing on a logarithmic scale. This is a form of non-uniform quantization which efficiently encodes the dynamic range of synaptic strengths at constant precision.

CHAPTER 3:

QUANTIFYING INFORMATION STORAGE IN THE HIPPOCAMPUS: LESSONS

FROM DENTATE GYRUS

Abstract

Reconstruction of synapses on dendritic spines of hippocampal CA1 pyramidal

cells revealed a capacity to store ~4.7 bits of information at each synapse, across the

population range of synapses. This prediction was based on the properties of a set of

dendritic spines found to share the same presynaptic axon and postsynaptic dendrite.

The degree to which these spine pairs were similar in size was assumed to reflect how

tightly controlled synaptic strength can be under conditions where spike-timing

dependent plasticity (STDP) dominates. Here, spine pairs with the same pre-post

connectivity were explored for another cell type in the hippocampus, the dentate

granule cell. Under control conditions, the similarity in size between the axon-coupled

spine pairs was less in dentate MML than in CA1 stratum radiatum. The resulting

information storage capacity per synapse in MML was calculated to be ~3.2 bits (30%

less than CA1). In the dentate MML we also assessed the effect of long-term

potentiation (LTP) induced by delta-burst stimulation delivered to the medial perforant

pathway in the angular bundle *in vivo*. At 30 minutes following induction of LTP, the

distribution of synapse sizes shifted in both directions, producing a higher frequency of smaller than average, as well as larger than average spine head and synapse sizes. These findings suggest that the growth of potentiated synapses was homeostatically counterbalanced by heterosynaptic long-term depression at neighboring synapses. Thus, the shift in distribution increases the available number of strengths, while having negligible impact on the precision of control over size for paired spines in MML. The results of this study elucidate the significance of functionally distinct brain regions and states in terms of information storage capacity.

Introduction

Evidence for Hebbian plasticity, and in particular, spike-timing-dependent plasticity (STDP) in the hippocampus and cortex is abundant (Bi and Poo, 1998; Caporale and Dan, 2008). The STDP literature highlights the importance of the timing of axonal input to spines relative to postsynaptic cell depolarization. Notably, the shifts in synapse size that occur in response to STDP are correlated across a number of metrics including spine head volume, postsynaptic density area (PSD), AMPAR complement and docked vesicle number (Harris and Stevens, 1989; Lisman and Harris, 1994; Schikorski and Stevens, 1997; Murthy, 2001; Branco, 2008; Bourne, 2013).

In populations of spines where Hebbian plasticity (STDP) dominates and a skewed distribution of spine dimensions is present, pairs of spines sharing a dendritic branch and an axonal input tend to be similar in size (Chicurel and Harris, 1992). The tendency for spines with this shared activation history to have similar size and strength is

most likely not an accident, but rather a reflection of the shared Hebbian processes at work. In the cortex, spines that share a presynaptic and postsynaptic partner appear similar in size in a number of studies (Sorra and Harris, 1993; Kasthuri, 2015; Markram, 1997). Our recent work quantified the similarity in spine size for spines with a shared presynaptic and postsynaptic history in rat CA1 s. radiatum and demonstrated that the information storage capacity at these synapses is much higher than previous estimates (Bartol, 2016).

We applied information theory (or signal detection theory, more specifically) to estimate the number of possible states that a dendritic spine can assume on hippocampal pyramidal cells in CA1 (s. radiatum). Information theory is often used in computing. We know that a binary bit in a computer has 2 possible values (0 or 1), but it can only have one of these states at a given time. To assess how many unique states a dendritic spine can assume, we calculated the coefficient of variation for the differences between two spines that share the same activation history, based on sharing the same pre and postsynaptic partners (same pre-post connectivity). For the hippocampal area CA1 population, we found that axon-dendrite coupled spines presumed to have the same activation history had a CV of 0.10 (N=3). We applied the principles of information theory to the full range of observed spine sizes (60 in CA1 s. radiatum) to calculate the number of distinguishable synaptic states, keeping in mind that these calculations represent the lower limit. The outcome of the information theory calculation yielded 26 distinct states for spines in hippocampal area CA1 (Bartol, 2016). If we assume that each spine in the population can grow or shrink to occupy any size in the observed range, then this calculation states that every spine has the capacity to store 4.7 bits of information in

its size (a metric for strength). In other words, under these assumptions, the calculated information storage capacity is assigned to each spine in the population.

Our current work explores the information storage capacity at dentate granule cell synapses, in the middle molecular layer (MML). The dentate gyrus is a unique region of the hippocampus. It is one of the few brain regions capable of generating newborn neurons in adulthood in mammals (i.e. neurogenesis) and the plasticity of neurons here is influenced by their relative age (Snyder, 2001; Saxe, 2006). Learning paradigms implemented in the dentate gyrus have revealed a low threshold for intrinsic plasticity in addition to any synaptic plasticity that is induced (Lopez-Rojas, 2016). In addition to these functional differences to neurons in other regions of the brain, dentate granule cells have a very different dendritic arbor compared to CA1 or cortical pyramidal cells, most notably the absence of an apical trunk dendrite and basilar dendrites. The chalice-shaped arbor of dentate granule cells is associated with poor back-propagation of action potentials (Krueppel, 2011). For these various reasons, the degree of correlation in the strengths of the synapse pairs sharing a presynaptic and postsynaptic history may be less than in other brain regions. By exploring and quantifying the capacity for information storage across the hippocampus, we can begin to explore how the diverse functions of hippocampal regions intersect with plasticity processes in these regions.

Information theory was applied here to new data from the dentate gyrus MML to assess whether the information storage capacity at each synapse was the same or different than in CA1 s. radiatum. Toward this end, the dynamic range in spine sizes was measured in MML, as well as the CV for spines with shared activation history (same pre-post connectivity). Further, we tested whether the observed shift in spine head sizes in

LTP had a functionally significant influence on the precision of axon-spine size coupling or dynamic range in spine sizes that affected the number of distinct functional states, and thus capacity to store information at the synapse (as measured in bits). We find that by using the larger dynamic range observed in LTP we see a shift information storage capacity between control and LTP hemispheres. Moreover, we presume that LTP in the CA1 region would shift the information storage capacity in a similar manner to what we observe in the dentate.

Results

Spine head distributions and information stored at the synapse differ between CA1 and MML

In data from the middle molecular layer of dentate gyrus (MML), we sampled 205 spine heads from the control hemispheres and 189 spine heads from LTP hemispheres. We demonstrated in our previous work that spine head volume, PSD area and vesicle number are highly correlated. For the purposes of this study we chose to focus on spine head volume because it is the most accurately measured of these correlated metrics. This accuracy of measurement is due to our ability to reconstruct spine heads in 3D space in Blender. While the PSD area and vesicle number are easily estimated from the Reconstruct (.ser) files, the ability to take multiple measurements of the spine head volume in a more computationally accurate representation enables improved quality of measurements. We found that the distribution of spine head volume in MML was

significantly different from that in CA1 as measured by the Kolmogorov-Smirnoff (KS) test (Fig. 1).

Repeating our analysis of pairs of spines sharing the same dendritic branch and same axonal input, or SDSA pairs (example pairs from CA1 and dentate are shown in Fig. 2), we identified, measured and plotted the volumes of such pairs (n=16) from the MML (Fig. 3). We found that when plotted against one another, the slope of the paired spine volumes was lower in the MML (0.66) than in CA1 (0.91). A greater variability in paired spine sizes was also evidenced by the higher coefficient of variation for SDSA pairs in this area (CV = 0.42 +/- 0.066) compared to CA1 (CV = 0.10 +/- 0.037). These data indicate that there is less concordance in spine head size, and thus synaptic efficacy, for dentate SDSA pairs than for CA1 SDSA pairs. The measured CV for the pairs of spine head volumes is an important component in estimating distinguishability of spine sizes using signal-detection theory (Bartol, 2016; Green and Swets, 1966; Schultz, 2007). It allows us answer the question of how well two spine heads with different sizes can be distinguished if they are drawn from two distributions with different means but the same CV

We set the threshold for an ideal observer at 69%, corresponding to a signal-to-noise ratio of 1. This gives us an estimate for the the minimum spacing between distinguishable spine head volumes; and knowing the range of volumes we can then calculate the total number of distinguishable volumes. The number of bits of information that can be stored is the logarithm base 2 of this number. We calculated ~3.2 bits of information capacity in control synapses in MML based on the observed CV (0.44) and range of available sizes (270). Compared to CA1 (4.7 bits, Bartol 2016), this outcome

represents an almost 30% lower information storage capacity in the dentate gyrus granule cell.

LTP stimulation in the MML shifts the distribution of spine heads but not information stored at the synapse

In the LTP hemispheres, which showed an average of 41.0% potentiation (Fig. 4A,4B), we sampled 189 spines and demonstrated that the distribution of these spine heads was significantly different from spine heads in the same region in the control hemispheres. Using the spine head volume measurements, we found that the shift in the distribution of spine head volume in the LTP condition in MML was significantly different as measured by the Kolmogorov-Smirnoff (KS) test (S. Fig. 4), with an increase in both tails of the distribution. In other words, following LTP induction, some spines grew much larger than in the control hemisphere, and appear to have been counterbalanced by an increase in the number of small spines (Fig. 4E, 4F). We observed additional differences in the LTP and control spine populations. First and foremost, we observed variation in the number of synapses per micron (Fig. 4D). We also noted that the discrepancy in measurements (4 measurements taken by two people) was only notable in LTP condition, which we believe is suggestive of the fluctuating nature of these spines at the 30 minute time point (Fig. 4C).

In the MML control hemispheres we observed 10 SDSA pairs, with a regression slope for the paired volumes of 0.72. The CV for the observed pairs was 0.44, and the range of spine sizes (observed only in the control hemispheres) was 78. Using this CV and range, specific to the control condition, we calculated 7. 4 distinguishable spine sizes

for the control condition, slightly lower than our calculation across both hemispheres (above) (Fig. 5A, 5C). This results in 2.9 bits of information storage capacity in the control hemisphere, given the smaller range of observed spine sizes.

In the LTP hemispheres, we observed 6 SDSA pairs and plotted the paired spine head volumes with a regression slope of 0.35 (Fig. 5B, 5D). A nearly identical coefficient of variation was observed under the LTP condition as compared to control condition in MML (CV =.40 +/- 0.106). However, the shift in the distribution of spine sizes increased the range of available sizes (270). Applying signal detection theory as above, we calculated that the information stored at the synapse in the MML was affected by the shift in the underlying distribution of spine head volumes under LTP conditions and calculated 9.4 distinguishable sizes for 3.2 bits per synapse. Importantly, this increase in information is enabled by the shift in the distribution rather than a shift in the similarity in the strength of paired spines.

Does spine shape matter?

In CA1 s. radiatum data, SDSA spine pairs are not only similar in size, but also in shape. In order to quantify the relative similarity of SDSA spine pairs to one another, as compared to other spines of similar volume, we devised a new method to quantitatively evaluate spine shape. Spines were centered, rotated and scaled in 3D space. The Jaccard Distance was calculated by taking the intersection of the two spines from the union. This numerical measure from 0 to 1, normalized for size and rotation, allows for a quantitative comparison of spines' shape. In Fig. 6 we plotted the Jaccard distance for random spine

comparisons (Fig. 6A), as well as those spines sharing axonal input in CA1 (Fig. 6B). Interestingly, the slope of the line was not influenced by connectivity in CA1, where the slope does not differ significantly from random as a result of connectivity. We also plotted random spine comparisons in MML control and LTP conditions (Fig. 6C, 6D). The relationship between spine head shape and Jaccard distance is stronger in MML, and varies slightly in response to condition. The stronger relationship between spine head shape and Jaccard distance in the MML of dentate gyrus is in keeping with our hypothesis that there is indeed a relationship between spine head volume and spine head shape, and that the precision of control over this relationship is related to the degree of STDP precision.

Discussion

In the current study, we provide evidence that the information storage capacity at the synapse in dentate gyrus MML, the site where the predominant stream of information from cortex arrives in the hippocampus is about 3.2 bits per synapse, compared to 4.7 bits per synapse in CA1 s. radiaturm. This difference could be functionally accounted for by a number of factors. The first factor is the variable age of neurons in the dentate that results from neurogenesis, including the hypothesized "retirement" of older neurons in this region from participation in Hebbian learning processes (Snyder, 2001; Mongiat, 2011). The second is the lower basal rate of activity in dentate gyrus as compared to CA1 (Chawala, 2005; Neunuebel, 2012). The third is that dendritic response properties and the efficiency of action potential back-propagation differ in the dentate gyrus granule cells

compared to CA1 pyramidal cells (Schmidt-Hieber, 2007; Krueppel, 2011; Brunner, 2016).

We hypothesize that the second and third factors contribute most to the discrepancy in our measurement of information storage capacity across the hippocampal regions. The combined effect of these two factors is to vastly decrease the impact and importance of the back-propagating action potential (bAP) in the dentate granule cell relative to the pyramidal cells of CA1. Lower activity rates in granule cells are believed to shape the response to the deluge of information received from the entorhinal cortex and to aid in pattern separation (Neunuebel, 2012). It may be that this relatively low firing rate also makes STDP processes less effective or relevant in the dentate gyrus compared to CA1 and other hippocampal regions. Similarly, response properties of granule cells in the dentate gyrus privilege dendritic computation over somatic bAP (Krueppel, 2011). These various factors decrease the impact of the bAP at the synapse as the spine receives presynaptic input. The lower information per synapse observed in this brain region might suggest simply that our metric for comparison, namely SDSA pairs, does not sufficiently capture the plasticity processes at play in this region where the influence STDP processes is reduced except where the effect of the bAP is enhanced by neuromodulators (Yang and Dani, 2014). In other words, granule cell SDSA synapses may not share presynaptic and postsynaptic histories as much as CA1 pairs as a result of privileging of dendritic computation in the granule cells.

We believe that neurogenesis plays less of a role in the observed difference in information storage capacity at the synapse between CA1 and dentate than the above factors. Although we are not able to know definitely the age of the neurons in a static

image representation, we do know that the technique used for the selection dendritic spines should ensure that the dendrites are similar in age across animal and condition (see Methods). Mature granule cells are believed to be involved in STDP processes, while newborn cells are not. Our analysis presumes our selected dendrites fall into the former category, and are able to participate in STDP processes.

We found that while the precision of control over spine size was not influenced by condition, the amount of information stored by synapses in the dentate gyrus MML region was was greater in the LTP hemispheres alone than under control hemispheres in isolation. In particular, our study demonstrates that LTP induction dramatically shifts the distribution of spines head sizes, thereby increasing capacity for information storage in the region temporarily. This increase is indeed likely temporary, as the spines in these series may still be in flux 30 minutes following LTP stimulation. The degree to which spines are in flux is evidenced by higher rates of variability in measurements of spines in the LTP series compared to control; the boundaries between spine head and neck are more difficult to identify in these two LTP series compared to control series in the same region. The very similar CV for spine pairs with a shared activation history in MML across control and LTP hemispheres suggests that the mechanism of plasticity is not altered by stimulation. Rather, the effective decrease in information storage observed in the control hemispheres is a result of the increased range of observed sizes in the LTP condition, not observed in control hemispheres. We chose to use this range for our overall calculation of information storage capacity per synapse in MML, as we believe the limits of the range of possible sizes provides the best predictor of information storage in the brain. We believe that in CA1, a stimulation paradigm paradigm would also likely

increase the observed range of sizes, and consequently the distinguishable spine sizes for a given CV. Even using the maximum range of spine sizes in the dentate gyrus MML did not bring the calculated information storage capacity close to that of CA1, the difference between regions lies in the CV of paired spines (CA1 = 0.10, dentate = 0.42).

With respect to spine shape, we observe a difference in correlation with spine head volume across our two regions of interest: CA1 s. radiatum and dentate MML. The flatter relationship between size and shape in CA1 may suggest that the precision of spine head size dominates in this region, where dendritic computations are less important relative to bAP as compared to dentate. Perhaps then, in dentate, the computations performed by spines are influenced by their shape, and the relationship between the volume and shape is more important. Further exploration of this relationship will be enabled by the tools developed in the context of this project.

The implications of our findings are twofold. Firstly, our findings suggest the need for further exploration of information storage capacity at synapses throughout the brain. Is there variability across cortical regions and between cortex and hippocampus? Is the dentate an aberration from the precision observed in CA1 or does the capacity for information storage at the synapse fall across a range? And, if STDP processes are not dominant in the dentate, how can we better capture the precision of control over spine size in this area, and others like it? Secondly, the shift in spine head volumes following LTP induction bears further exploration. The "lognormal" distribution of spine head sizes, and other metrics throughout the brain have been well documented (Buzsaki, 2014). How this long-tail distribution of spine sizes responds to activity, such as LTP induction, and how the observed shift in distribution supports the underlying

computations is less understood. How might this temporary shift in the distribution of spine sizes (increase in both larger and smaller spines) and information storage capacity serve to accommodate the underlying computational processes occurring at individual spines? The answers to these questions will further inform our understanding of computation in the brain, and its relationship to a well-characterized form of plasticity: STDP.

Methods

Surgery and electrophysiology

Data were collected from two young-adult male Long-Evans rats aged 179 and 121 days at the time of LTP induction and perfusion. They had been surgically implanted as previously described (Bowden, 2012) with wire stimulating electrodes separately into the medial and lateral perforant pathways running in the angular bundle in the LTP hemisphere, and in the medial perforant pathway only in the control hemisphere (only medial path data are described in this paper). Wire field EPSP recording electrodes were implanted bilaterally in the dentate hilus. Two weeks after surgery, baseline recording sessions (30 min) commenced, with animals being in a quiet alert state during the animals' dark cycle. Test pulse stimuli were administered to each pathway 1/30 s as constant current biphasic square-wave pulses (150 μs half-wave duration), alternating between the three stimulating electrodes. The test pulse stimulation intensity was set to evoke medial path waveforms with fEPSP slopes >3.5 mV/ms in association with population spike amplitudes between 2-4 mV, at a stimulation current ≤ 500 μA. On the

day of LTP induction, after stable baseline recordings were achieved, animals received 30 min of test pulses followed by delta-burst stimulation (DBS) delivered to the ipsilateral medial perforant path, while the contralateral hippocampus served as a control. The LTP-inducing DBS protocol consisted of five trains of 10 pulses (250 µs half-wave duration) delivered at 400 Hz at a 1 Hz interburst frequency, repeated 10 times at 1 min intervals (Bowden, 2012). Test pulse stimulation then resumed until the animal was sacrificed. The initial slope of the medial path fEPSP (mV/ms) was measured for each waveform and expressed as a percentage of the average of the last 15 min of recording before DBS.

Perfusion and fixation

At 30 minutes after the commencement of DBS, animals were perfusion fixed under halothane anesthesia and tracheal supply of oxygen. The perfusion involved brief (~20 sec) wash with oxygenated Krebs-Ringer Carbicarb buffer (concentration in mM; 2.0 $CaCl_2$, 11.0 D-glucose, 4.7 KCl, 4.0 $MgSO_4$, 118 NaCl, 12.5 Na_2CO_3, 12.5 $NaHCO_3$, pH = 7.35–7.4, osmolality = 300–330 mmol/kg), followed by 2% formaldehyde and 2.5% glutaraldehyde (both aldehydes from Ladd Research, Williston, VT) in 0.1M cacodylate buffer (pH = 7.35-7.4) containing 2 mM $CaCl_2$ and 4 mM $MgSO_4$ for ~1 hr (~1900 ml of fixative was used per animal). The brains were removed from the skull at about 1 hour after end of perfusion, wrapped in several layers of cotton gauze and shipped on ice in the same fixative from the Abraham Laboratory in Dunedin, New Zealand, to the Harris Laboratory in Austin, Texas by overnight delivery (TNT Holdings B.V., Amsterdam, The Netherlands).

Tissue processing and serial sectioning

The fixed tissue was then cut into parasagittal slices (70 µm thickness) with a vibrating blade microtome (Leica Microsystems, Buffalo Grove, IL) and processed for electron microscopy as described previously (Harris, 2006; Kuwajima, 2013 Methods Mol Biol). Briefly, the tissue was treated with reduced osmium (1% osmium tetroxide and 1.5% potassium ferrocyanide in 0.1M cacodylate buffer), followed by microwave-assisted incubation in 1% osmium tetroxide under vacuum. Then the tissue underwent microwave-assisted dehydration and *en bloc* staining with 1% uranyl acetate in ascending concentrations of ethanol. The tissue was embedded into LX-112 epoxy resin (Ladd Research) at 60°C for 48 hr before being cut into series of ultrathin sections at the nominal thickness of 45 nm with a 35° diamond knife (DiATOME, Biel, Switzerland) on an ultramicrotome (Leica Microsystems). The serial ultrathin sections from MML (region of molecular layer ~125 µm from top of granule cell layer in dorsal blade of the hippocampal dentate gyrus) were collected onto Synaptek® Be-Cu slot grids (Electron Microscopy Sciences, Hatfield, PA, or Ted Pella, Redding, CA) coated with Pioloform (Ted Pella), and stained with a saturated aqueous solution of uranyl acetate followed by lead citrate (Reynolds, 1963).

Imaging and alignment

The serial ultrathin sections were imaged, blind as to condition, with either a JEOL JEM-1230 TEM (Tokyo, Japan) or a transmission-mode scanning EM (tSEM; Zeiss SUPRA® 40 field-emission SEM with a retractable multi-mode transmitted

electron detector and ATLAS package for large-field image acquisition; Oberkochen, Germany; Kuwajima, 2013 PLoS). On the TEM, sections were imaged in two-field mosaics at 5,000× magnification with a Gatan UltraScan 4000 CCD camera (4,080 pixels × 4,080 pixels; Pleasanton, CA), controlled by DigitalMicrograph software (Gatan). Mosaics were then stitched with "Photomerge" function in Adobe Photoshop (San Jose, CA). The serial TEM images were first manually aligned in RECONSTRUCT™ (http://synapseweb.clm.utexas.edu/software-0; Fiala, 2005) and later with Fiji with the TrakEM2 plugin (http://fiji.sc; Saalfeld, 2012; Cardona, 2012; Schindelin, 2012). On the tSEM, each section was imaged with the transmitted electron detector from a single field encompassing 32.768 µm × 32.768 µm (16,384 pixels × 16,384 pixels at 2 nm/pixel). The scan beam was set for a dwell time of 1.3-1.4 ms, with the accelerating voltage of 28 kV in high-current mode. Serial tSEM images were aligned automatically using Fiji with the TrakEM2 plugin. The images were aligned rigidly first, followed by application of affine and then elastic alignment. Images from a series were given a five-letter code to mask the identity of experimental conditions in subsequent analyses with RECONSTRUCT™. Pixel size was calibrated for each series using the grating replica image that was acquired along with serial sections. The section thickness was estimated using the cylindrical mitochondria method (Fiala and Harris, 2001).

Reconstructions

Spine heads and postsynaptic densities (PSDs) were identified in by their electron density and presence of closely apposed presynaptic vesicles. Contours were drawn using

Reconstruct software on serial images for each spine head. For spines making a synaptic

contact a second contour was drawn, closely encompassing, as a single closed contour,

the pre- and postsynaptic extent of the electron dense region. For each of the 4 series

presented here (N = 2 x condition), all spines on three dendrites of similar caliber were

traced. Caliber has previously been shown to scale with dendrite cross-section and

microtubule count, which is a more reliable feature (Fiala, 2001; Bourne and Harris,

2011; Bowden, 2008). The microtubule count ranged from 30-35 for this sample and

represents the average for dendrites found in this region.

Blender, a free, open-source, user-extensible animation tool was used in

conjunction with 3D models generated in Reconstruct. We enhanced our Python add-on

to Blender, "Neuropil Tools" (Bartol, 2016) to better facilitate the processing of the 3D

reconstruction and evaluation of spines. The additions encompassed in *Processor Tool*

were as follows:

a.   The software allows for the selection of traced objects from Reconstruct (.ser)

files by filter, allowing the user to select only desired contour traces (in this case spine

head and PSD contours for three dendrites per series)

b.   At the press of a button, the tool generates 3D representations of selected contours

in Blender. This step calls software from VolRoverN (Edwards, 2014) from Blender,

which generates "mesh objects" by the addition of triangle faces between contour traces

c.   Smoothing and evening of the surface of spine objects is accomplished by calling

the GAMer (http://fetk.org/codes/gamer/) software

d.   In a few cases, the formation of triangles was uneven and required additional

manipulation by Blender tools and repeating of step C before proceeding to step E

e.   Lastly, postsynaptic density (PSD) areas are assigned as metadata (represented by red triangles) on reconstructed spine heads; the assignment is performed based on the overlap of PSD and spine head contours (described above) in 3D space

Segmentation and Evaluation of Spines

Dendritic spines were segmented as previously described (Bartol, 2016) using the "Neuropil Tools" analyzer tool. The selection of spine head from spine neck, and from spine neck to dendritic shaft, were made using the same standardized criterion as before (visually identified as half-way along the concave arc as the head narrows to form the neck). Spines were excluded if they were clipped by the edge of the dataset. In order to ensure the accuracy of the measurements, segmentation and spine head volume evaluation were completed four times (twice each by two people) and averaged. A further check was added at this step, whereby spine heads with an SEM>=0.2 microns for all four measurements were visually evaluated by an expert, and any discrepancy in the segmentation was corrected. Interestingly, the only spines with an SEM larger than 0.2 microns were in the LTP condition. We believe this occurs because of the spines undergoing LTP are likely to be in transition at the 30 minute time point, and as such, the delineation between spine head and spine neck are more difficult for the human eye to see. In the two control condition series, further evaluation by an expert was performed, adjustments were made accordingly to better represent this view (**Table 1**).

Identification of Same-Dendrite-Same-Axon (SDSA) Pairs

Identification of spines sharing an axonal input and dendritic branch was performed in Reconstruct. Beginning in the center of each of three dendrites,axons were traced bidirectionally to the next synaptic contact from about 10 initial contact points. If the second contact was a spine from the same dendritic branch, the two spine heads were labeled an SDSA pair. This process was repeated in all four series (N=2 x condition).

Statistical Analysis

Statistical analysis and plots were generated using Python 2.7 with NumPy, SciPy and Matplotlib. Cumulative distributions (CDFs) were generated and plotted using spine head volume measurements output from Blender. Kolmogorov-Smirnoff (KS) test, a nonparametric test, was used for comparing distributions due to the skewed (non-normal) distribution of the data. The coefficient of variation (CV) for SDSA pairs was calculated from the standard deviation of the pair divided by the mean volume of the pair. The average CV for each series or condition was the average of the CV of included pairs.

Estimation of number of distinguishable spine sizes and bits of precision in spine size

To estimate the number of distinguishable spine size and bits of precision we calculated the number of distinct Gaussian distributions of spine sizes, each with a certain mean size and standard deviation that together would cover and span the entire range of spine head sizes for each series or condition. Given the CV, the spacing between the mean values of each sub-distribution can be chosen to achieve a total of 31% overlap with adjacent sub-distributions giving a 69% discrimination threshold. A 69%

discrimination threshold is commonly used in the field of psychophysics and corresponds to a Signal-to-Noise (SNR) of 1.

The 69% confidence interval, z, of a Gaussian distribution is given by:

$z = sqrt(2)*erf−1(0.69)$

The spacing, s, of adjacent intervals of mean, μ, is given by:

$s = \mu*2*CV*z$

The number, N, of such distributions that would span the factor of 50 range of spine sizes is:

$N = log(50)/log(1+2*CV*z))$

Where the CV = 0.44 (control) and 0.40 (LTP)

$N = 7.4, N = 9.4$

The number of bits of precision implied by N distinguishable distributions is given by:

$bits = log2(N)$

$bits = 2.887 (control), bits = 3.23 (LTP)$

Chapter 3, in full, is currently being prepared for submission for publication of the material. Bromer C., Bartol T.M., Bowden J., Gonzalez P., Hanka D., Hubbard D., Kuwajimi M., Mendenhall J., Parker P.H., Abraham W.C., Sejnowski T.J., Harris, K.M. The dissertation author was the primary investigator and author of this material.
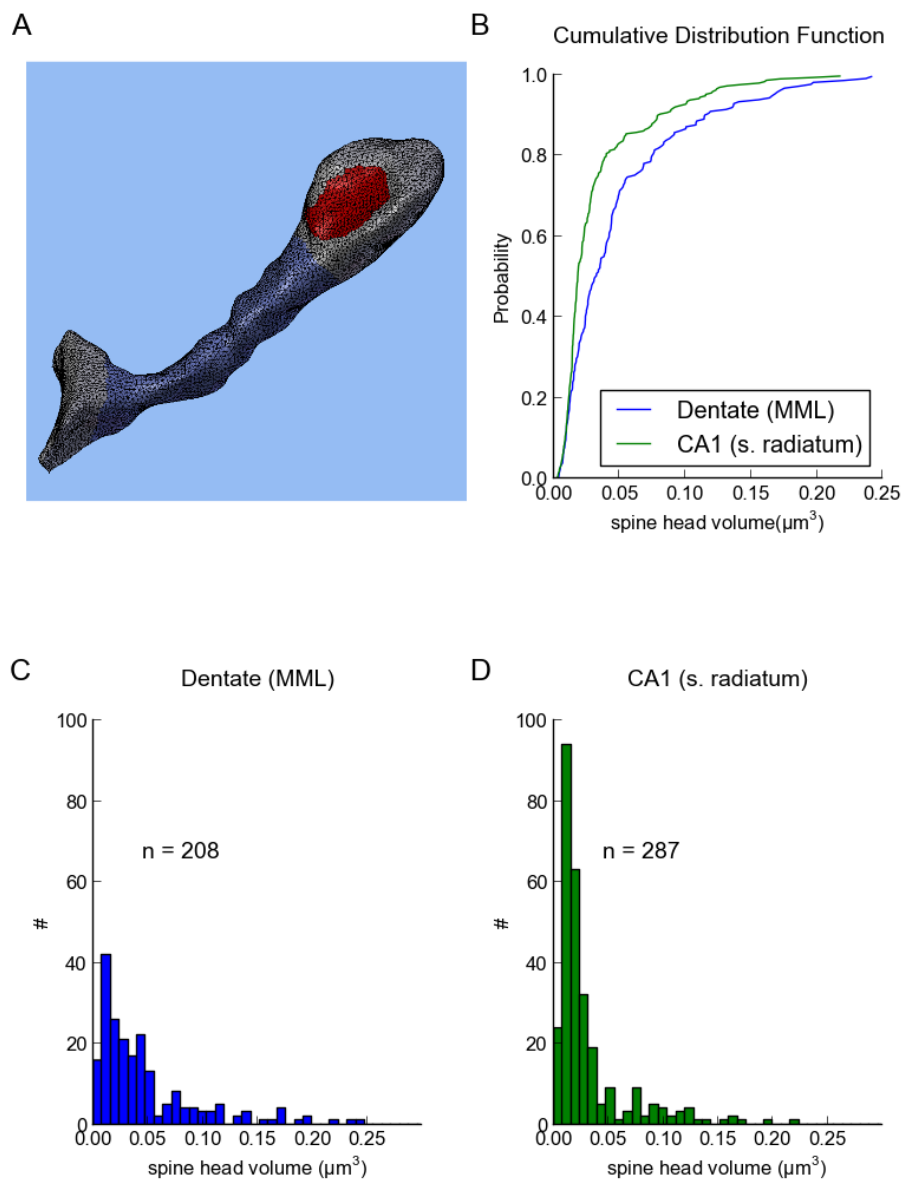
Figure 3.1: Distributions of spine head volumes differ across hippocampal regions under control condition. A) Example of spine head volume measurement using Neuropil Tools (spine head represented in gray, spine neck in purple, PSD area in red) B) Cumulative distribution function plot showing that the two spine populations are significantly different as measured by head volume (KS test p-value = 4.20e-07) C) Histogram of spine head volume in dentate gyrus MML (n = 205 synapses) D) Histogram of spine head volume in CA1 (n = 287 synapses).
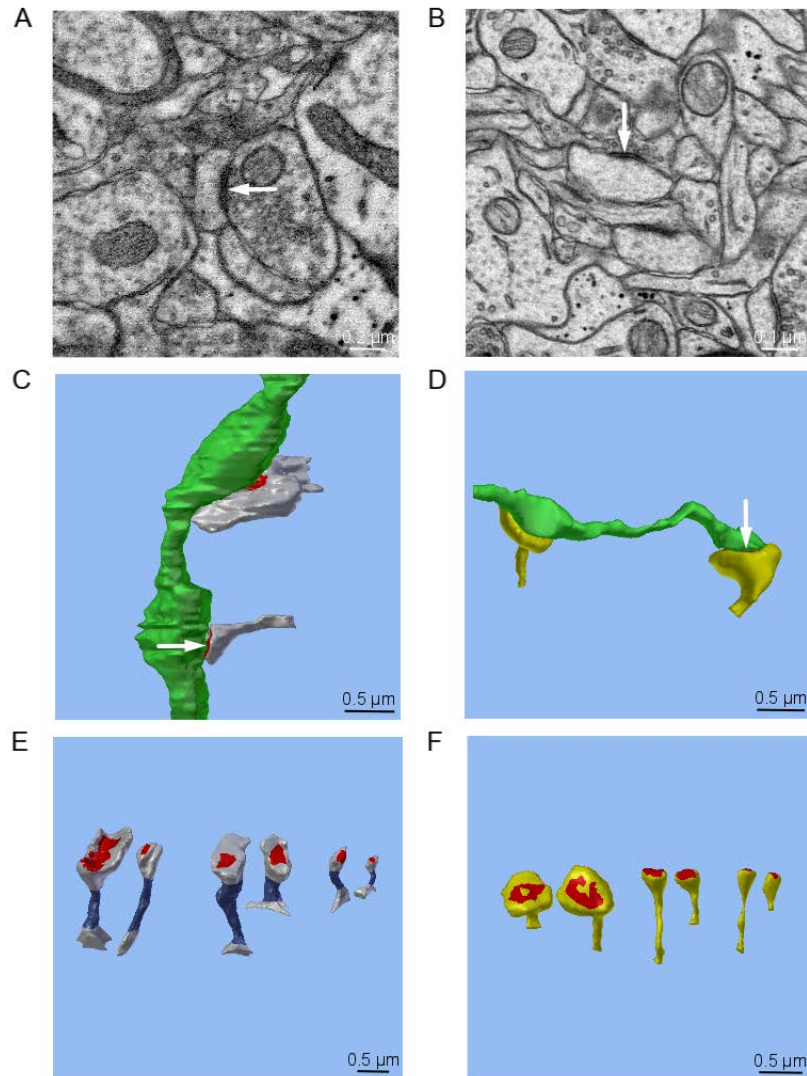
Figure 3.2: Depiction of "Same Dendrite Same Axon" (SDSA) spines in dentate gyrus MML (A,C,E) and CA1 s. radiatum (B,D,F). A) Representative EM of one spine in a pair of SDSA spines in MML B) Representative EM of one spine in a pair of SDSA spines in CA1 C) Example SDSA pair represented in Blender (MML) D) Example SDSA pair represented in Blender (CA1)   E) Largest, median and smallest SDSA pairs in MML F) Largest, median and smallest SDSA pairs in CA1

Figure 3.3: Spine head volume for pairs of spines sharing a presynaptic and postsynaptic history is more similar in CA1 s. radiatum compared to dentate gyrus MML A) spine head volumes dentate gyrus MML (slope = 0.66, CV = 0.42), gray points represent a quadruplet of spines sharing a dendritic branch and a single axonal input (N=2) B) spine head volumes CA1 s. Radiatum (slope = 0.91, CV = 0.10), gray points represent a triplet of spines sharing a dendritic branch and single axonal input (N=3) C) Distinguishable spine sizes in dentate gyrus MML (9.4) D) Distinguishable spine sizes in CA1 s. radiatum (27.4)
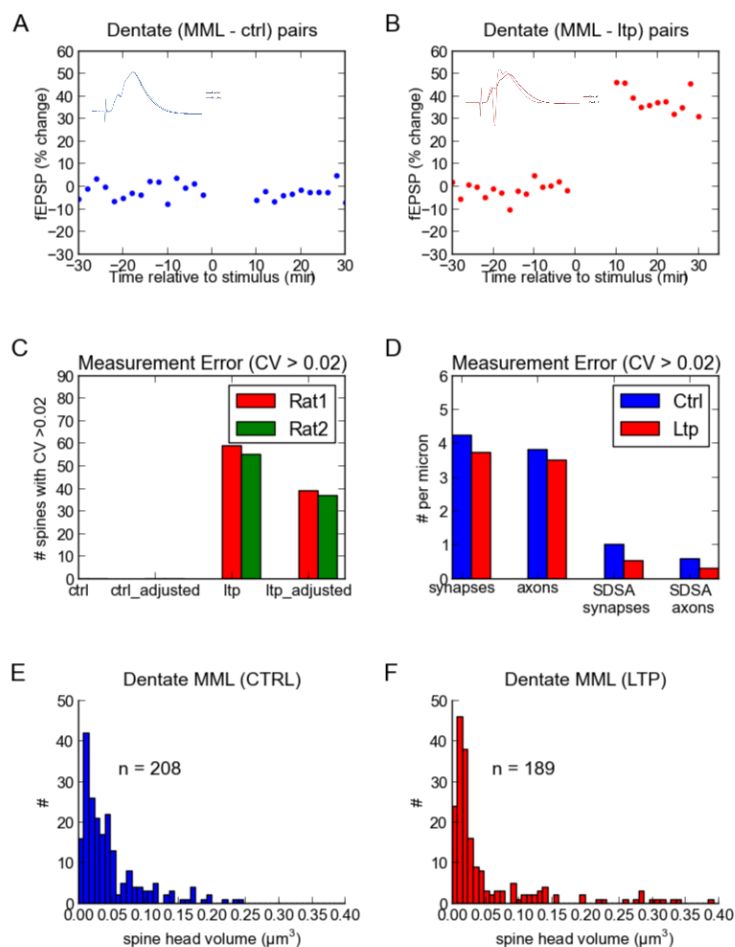
Figure 3.4: Acute induction of LTP in the middle molecular layer (MML) of the dentate gyrus of adult rat by stimulation of the angular bundle *in vivo* shifts the distribution of spine head volumes. Graphs illustrate change relative to baseline stimulation prior to induction of LTP by delta-burst stimulation in A) Control and B) LTP conditions. Control animals received test pulses only to the contralateral medial perforant pathway throughout the experiment and Representative waveforms [INSET] during baseline stimulation and following delta-burst stimulation. Scale is for both LTP and Control waveforms C) To test the accuracy of spine head volume metrics, SEM was calculated across 4 measurements for each spine. The number of spines where the SEM across measurements was greater than 0.02 microns was much higher in LTP condition (59 and 55) even after expert evaluation (39 and 37) compared to control D) Bar plot showing the frequency of SDSA synapses per micron relative to the total number of synapses E) Histogram of spine head volume for spines in the dentate gyrus MML under E) control condition and F) LTP condition
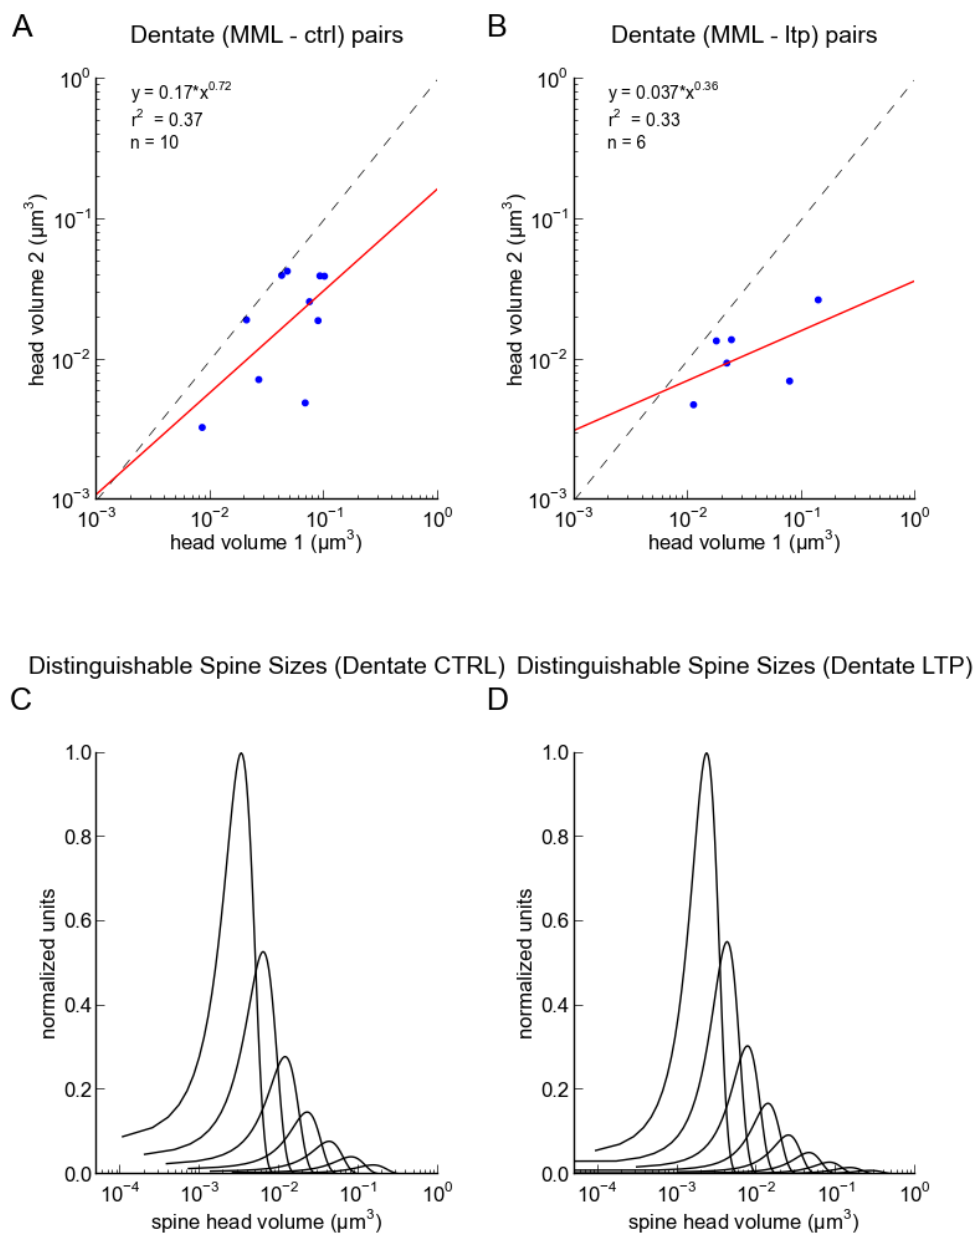
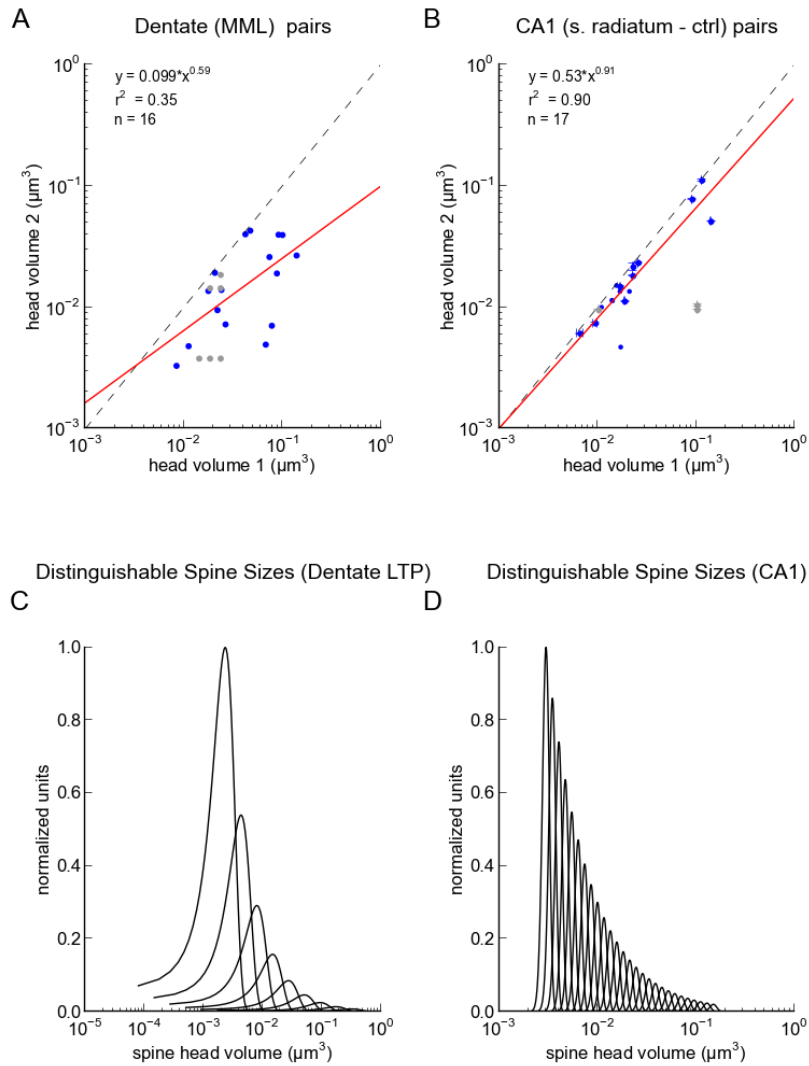**Figure 3.5: Spine head volume for pairs of spines sharing a presynaptic and postsynaptic history is similar in dentate gyrus MML under control and LTP conditions.** A) spine head volumes dentate gyrus MML under control conditions (slope = 0.723, CV = 0.44) B) spine head volumes dentate gyrus MML under LTP conditions (slope = 0.355, CV=0.40) C) Distinguishable spine sizes in dentate gyrus MML under control condition (7.4) D) Distinguishable spine sizes in dentate gyrus MML under LTP condition (9.4)
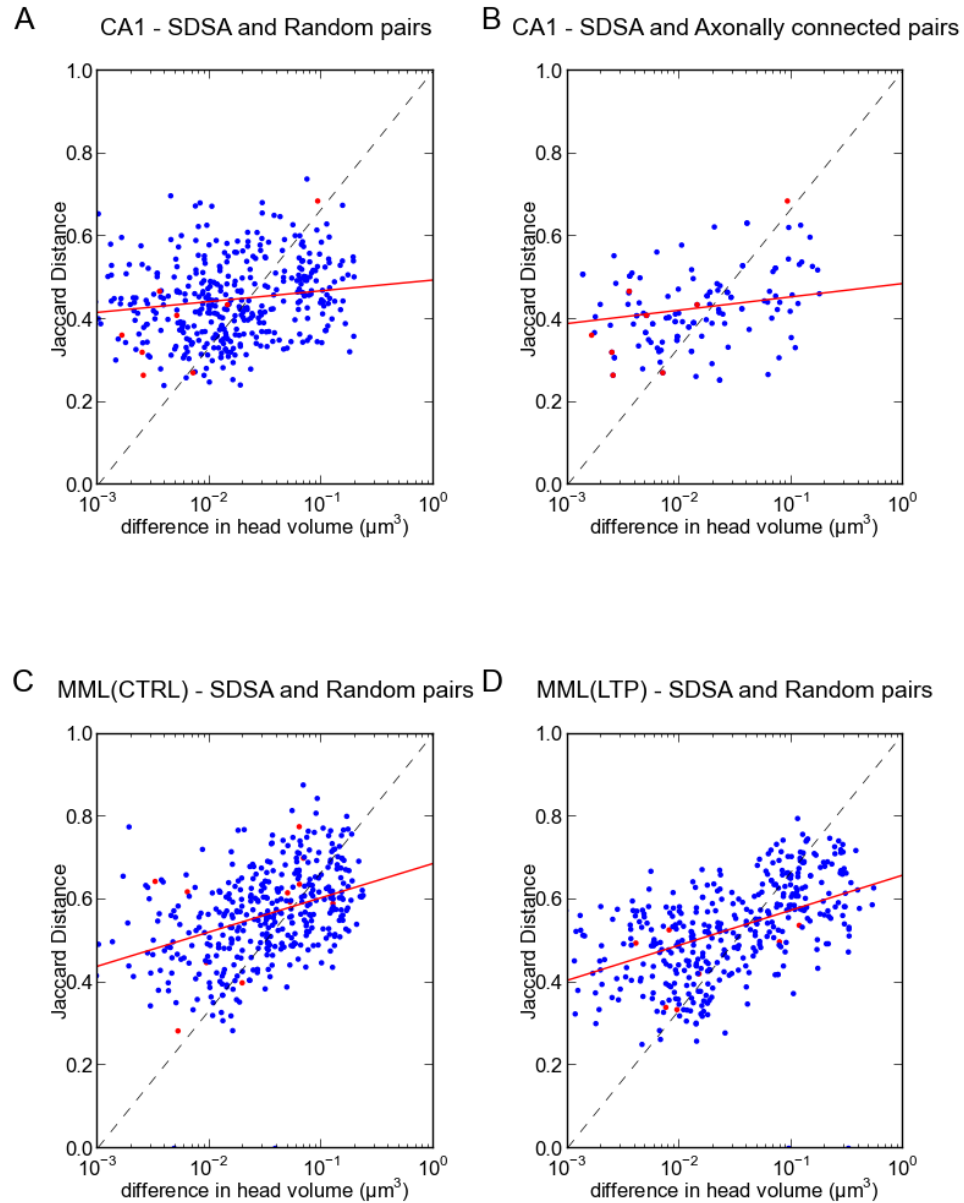
Figure 3.6: The relationship between spine head volume and shape varies across the hippocampus. Plots of Jaccard distance as a function of volume (SDSA pairs are represented in red in all plots) for A) random spine pairs in CA1 s. radiatum (slope = 0.0113) and B) axonally connected spine pairs in CA1 s. radiatum (slope =0.0140). C) random spine pairs in dentate MML control condition (slope = 0.0359) and D) random spine pairs in dentate MML LTP condition (slope = 0.0367)
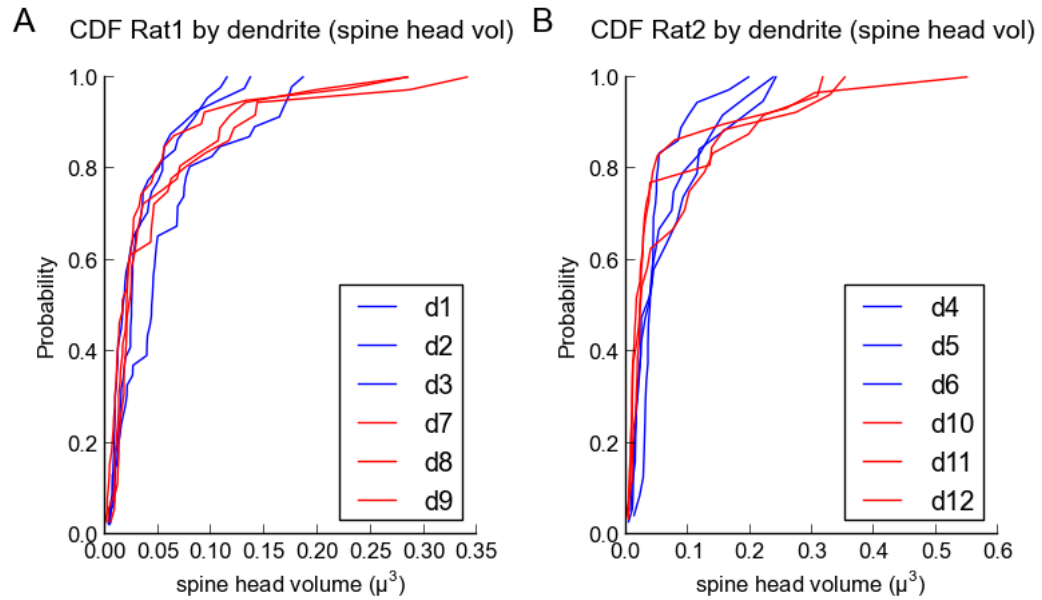
Figure S3-1: No significant difference in distribution of spine head volume occurs between dendrites within animal A) blue and red lines represent the distribution of spine head volume in three dendrites in the control and LTP condition respectively in Rat1 B) blue and red lines represent the distribution of spine head volume in three dendrites in the control and LTP condition respectively in Rat2
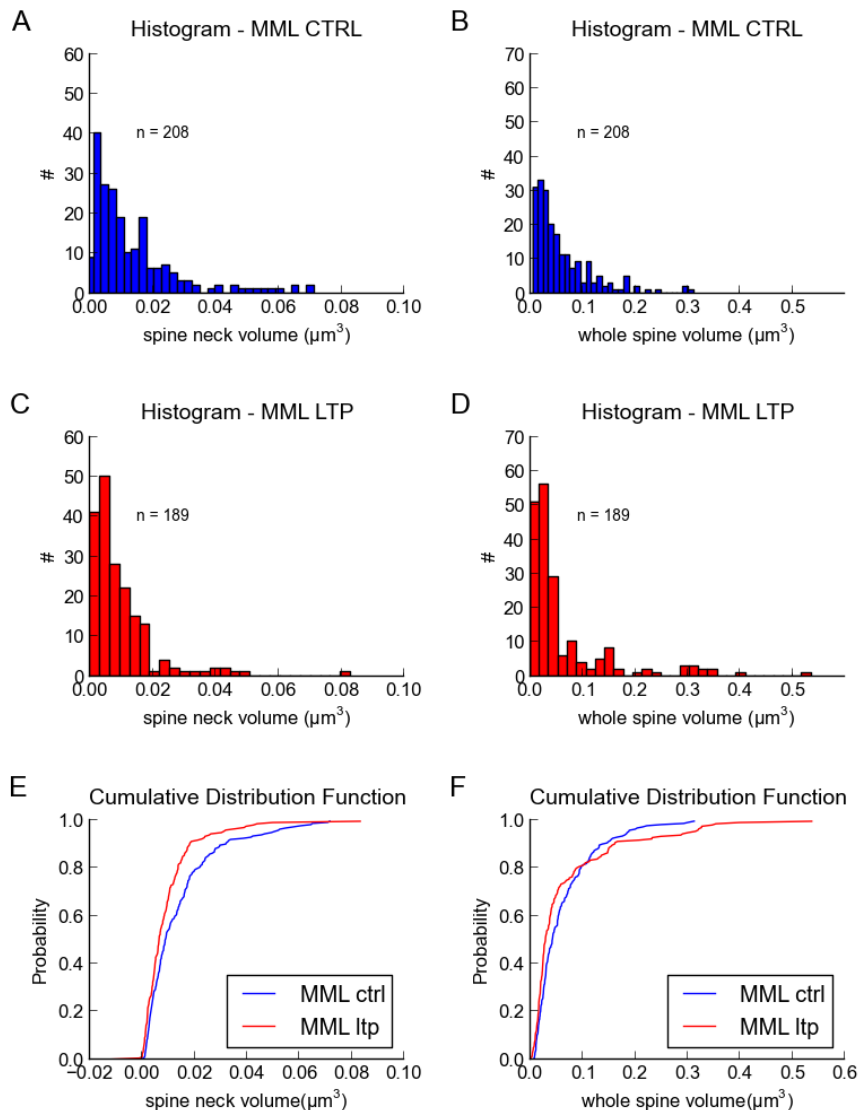
Figure S3-2: No significant difference exists in whole spine or spine neck populations in the control or LTP condition Histograms of A) spine neck volume and B) whole spine volume for the population of spines in dentate gyrus MML in the control condition. Histograms of C) spine neck volume and D) whole spine volume for the population of spines in dentate gyrus MML LTP condition. E) Cumulative distribution function plot showing that the two spine populations are significantly different as measured by spine neck volume (**KS test p-value = 0.003**) F) Cumulative distribution function plot showing that the two spine populations are significantly different as measured by whole spine volume (**KS test p-value = 0.019**)
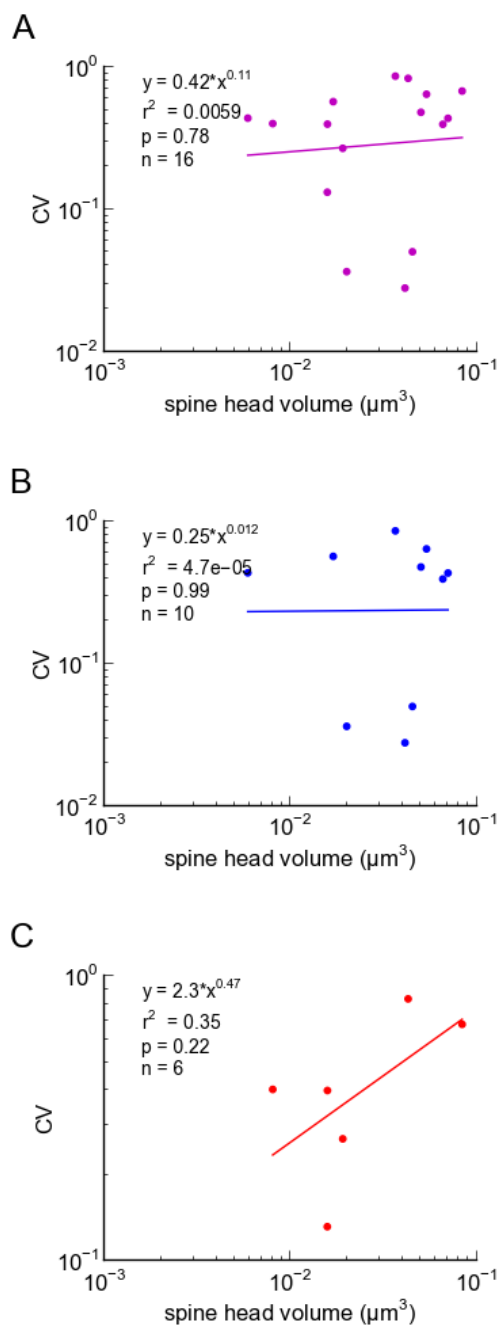
Figure S3-3: CV varies with spine size in the LTP condition but not the control condition. A) The combined SDSA pairs plotted against spine head volume, a slight increase in CV is seen with size when control and LTP pairs are aggregated B) CV for control SDSA pairs does not increase with spine head size C) CV for LTP SDSA pairs shows some relationship with size
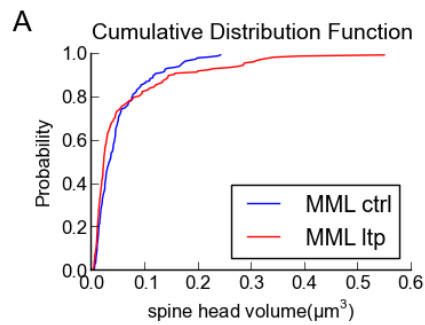
Figure S3-4: Control and LTP spine population distributions are significantly different from one another. Cumulative distribution function plot showing that the two spine populations are significantly different as measured by head volume (KS test p-value = 4.20e-07)

CHAPTER 4:

DEVELOPING TOOLS FOR THE ANALYSIS OF ELECTRON

MICROSCOPY DATA

Abstract

Improvements in imaging technology and computation in the last two decades have facilitated collection of large amounts of microscopy data and lowered barriers to collecting and storing such data. However, tools for evaluation and analysis of microscopy data have lagged behind, creating a bottleneck in utilization of existing datasets to address key questions in biology. As part of a larger pipeline, we have developed a set of tools for rigorous quantitative analysis of morphometric data generated from microscopy images, specialized for analysis of brain tissue. This software package, Neuropil Tools, is an add-on to Blender (https://www.blender.org/; an open access animation tool) with dependencies on existing software packages (CellBlender and VolRoverN, and GAMer) and is available for public use and distribution according to the GNU public license.

Introduction

The problem of accurately representing electron microscopy data in 3D has long plagued the field of Neuroscience and microscopy. A number of tools are available to

facilitate the tracing and naming of individual spines, axons and intracellular components, including Reconstruct, IMod, TrakEM2, SWIFT-IR, FIJI, NeuroMorph and iLastik (Cardona 2012; Jorstad 2015). Such tools have built in capabilities to visualize traced objects in 3D as well as tools to measure spine head volume, PSD area and other metrics of interest. However, existing tools lack the ability to measure volume and PSD area of reconstructed 3D data. The need for morphometric analysis of this nature has remained an unmet need for researchers.

Beyond the need for improved quality of measurements, expediting analysis of available electron microscopy data can greatly enhance use of the existing resources. While work on automatic segmentation of 2D images is ongoing, and currently unable to match human tracing work, facilitating generation of 3D analysis can make use of painstaking work on traces, and can identify and correct human errors made during the segmentation process.

Neuropil Tools is an open-source software package designed to aid in conversion of 2D electron microscopy images to 3D "mesh objects" for analysis of area, diameter and volume of neuronal components (github.com/mcellteam/neuropil_tools). The tool is under development and is encompassed in the larger CellBlender software package (github.com/mcellteam/cellblender), which utilizes Blender (blender.org), a state-of-the-art open-source 3D content creation and visualization tool. Its current capabilities include construction of computational quality mesh objects from labelled trace files (currently accepts .ser file format), and a set of tools to accurately, and in a semi-automated manner, measure volume, area, diameter, and length of spine heads and spine necks. The tool will be expanded to include analysis of various cellular components such as ER and

mitochondria. Such a tool greatly facilitates the speed of data processing and increases the accuracy of neuromorphic measurements. Moreover, the mesh objects generated using this tool are of computational quality and are automatically annotated with necessary metadata making them ready for use in simulations of biochemical reaction dynamics using MCell (mcell.org) and other 3D reaction/diffusion simulators such as VCell (vcell.org). Neuropil Tools is composed of several subcomponents, namely the "Processor Tool", "Spine Head Analyzer Tool", and "Connectivity Tool".

Methods and Materials

The tool assembled for this work relies on existing software packages and algorithms as well as new and original code. It aims to make complex functions available in a user-friendly and open-access interface. Ideally, such a tool will be utilized by numerous researchers and present a standard in the field for quality of 3D representations of neurons. Outlined below is information on the tools and algorithms utilized in this software package.

CellBlender

Neuropil Tools is a powerful package that draws on previously developed software to go above and beyond the capabilities of other morphometric analysis packages such as Neuromorph. Many of the functionalities in Neuropil Tools depend on another add-on built by our group, which is designed to visualize molecular scale simulations created by MCell. CellBlender creates the capacity for labelled regions and

other data to be stored on the triangles associated with each mesh, diminishing the clutter in the Blender workspace and ensuring that information is not lost or misplaced.

VolRoverN

VolRoverN surpasses existing methods for mesh generation, using an algorithm for the aggregation of traces into 3D objects that ensures manifoldness, water-tightness and lack of self- and object-object intersections, as well as geometric accuracy. Incorporating the functionalities in VolRoverN into Neuropil Tools enhances the user experience and ease of use, but Neuropil Tools relies fundamentally on the algorithms employed by VolRoverN

Results

Processor Tool

This tool handles the initial intake of segmented traces from RECONSTRUCT and preparation of quality meshes from these objects. Figure 1 depicts a TEM image of hippocampal Neuropil in which dendritic spines have been hand-traced in Reconstruct. On the right are a number of dendritic spines represented in 3D space in Neuropil Tools, with postsynaptic density (PSD) regions in red. At the click of a button, the user can navigate to and select a Reconstruct file (.ser) from a menu, and select any number of traced objects to be imported, reconstructed, and analyzed in batch. The series of steps performed automatically with no further human intervention includes mesh generation,

smoothing, and tagging of synaptic contact regions (PSD post-synaptically, and active-zone pre-synaptically). A first pass of mesh generation is performed by creating a series of connected triangular tiles spanning the original traces in 3D space. The section thickness (distance between traces) is read automatically by the tool and meshes are generated according to this spacing. Next, computational quality meshes are generated by remeshing and smoothing this first pass meshes. In the final step, the patches of mesh triangles on an object corresponding to synaptic contact regions are automatically identified and tagged with metadata obtained from the RECONSTRUCT trace files (Fig 13).

Spine Head Analyzer Tool

The Spine Head Analyzer component of Neuropil Tools is specialized for the analysis of neuromorphic reconstructions. The user is able to select a spine of interest from a list, and to select the border of the spine head (L) and and whole spine (R) as shown (the eventual pipeline will incorporate an algorithm that identifies inflection points like those below automatically, reducing the need for human input). Following the selection of these regions (Fig 14), the morphometric analysis of the whole spine, spine head, and spine neck is performed automatically by the tool. The morphometric data are stored as metadata attached to the object in Blender and can be output as a text file for convenient post-processing.

Connectivity Tool:

This tool facilitates exploration, visualization, and analysis of the connectivity among the axons and dendrites within a reconstructed block of neuropil. With this tool one can quickly find, for example, the axon that makes contact with a particular spine, or all the dendrites contacted by a particular axon. Reports of the morphometric data for all objects in the dataset can be generated, collated by their connectivity.

Discussion

Neuropil Tools (Figure 3) differs from other available software packages, such as Neuromorph, in the several ways (Jorstad, 2015). Firstly, Neuropil Tools generates high quality meshes directly from tracing software output (.ser file). To our knowledge, other existing tools are for 3D representations of 2D electron microscopy images are limited to importing existing 3D mesh objects. Secondly, Neuropil Tools uses user-given trace names in its representations of mesh objects and their connectivity, which it integrates automatically from the .ser trace files. In other existing packages like Neuromorph, all naming of objects is manual. Thirdly, morphometric calculations of volume, diameter and length are calculated at the press of a button, achieved by behind the scenes, consistently performing algorithms rather than by human intervention, as is the case in other software packages. Lastly, anatomically relevant regions, such as the PSD of spine, can be automatically tagged using metadata from the trace files, whereas in other packages, region labeling is done by hand in the software package and is thus far less accurate. For all of these reasons the Neuropil Tools software package is a useful and necessary

addition to the field, and should be improved and updated to remain consistent with

advances in microscopy and automatic segmentation (tracing) of 2D images.
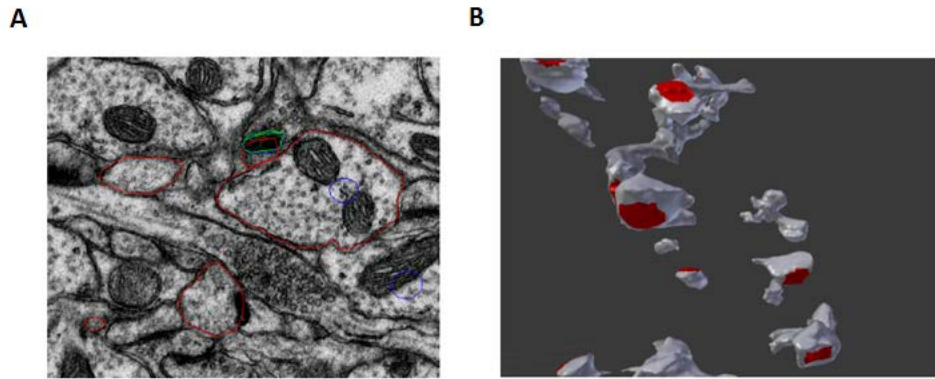
Figure 4.1: Electron microscopy image traces are used to generate 3D objects A) representative spine head traces from Reconstruct series in dentate B) Spine heads reconstructed using Neuropil Tools are pictured in 3D space in Blender.
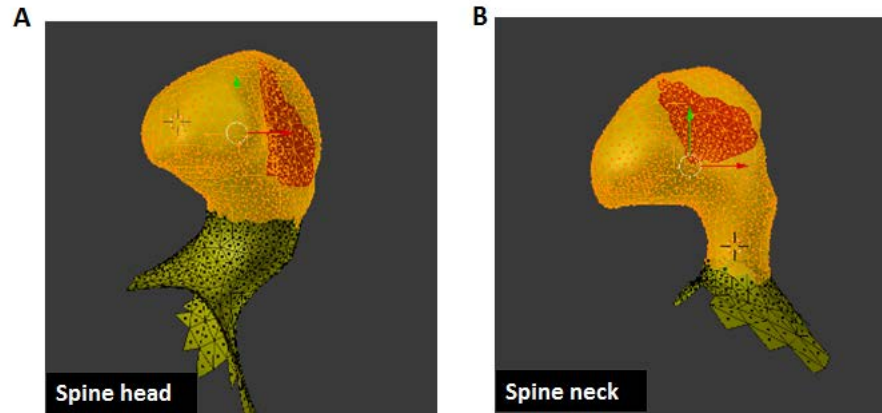
Figure 4.2: The boundary between the A) spine head and B) spine neck are selected by the user in Neuropil Tools. Volume, surface area and diameter are then automatically calculated by the click of a button.

Figure 4.3: Example of the Neuropil Tools interface with CA1 dendrite visible.

CHAPTER 5:

CONCLUSIONS

The primary focus of this dissertation work is to establish a framework for thinking about information storage capacity at the synapse, and to set benchmarks by quantifying information stored in each synapse in the hippocampus. This research lays the groundwork for additional experimental studies regarding the precision of control of STDP processes in the hippocampus and cortex. Also following from this work is a need to experimentally probe whether the variability in information stored at the synapse between CA1 and dentate gyrus is the result of neurogenesis and other unique properties of the dentate, or if some other form of plasticity is dominant of STDP in the region. Understanding how such differences might relate to the specified functions of these two regions is also in order. All of the work described above (Chapter 2 and 3) was enabled by the concurrent development of software to aid in the accurate, timely and consistent analysis of electron microscopy data (Chapter 4).

The work in Chapter 2 utilized data cleaned and represented in 3D space by Dr. Justin Kinney over the course of his work in Dr. Sejnowski's lab, and an observation by Dr. Thomas Bartol regarding the notable similarity in size of spines with shared presynaptic and postsynaptic history in CA1. The author of the dissertation came on board the project and assisted in developing the "Spine Head Analyzer" component of Neuropil Tools. The dissertation author was also responsible for the measurement of spine head volumes, and assisted with developing and carrying out the analysis strategy and generation of the manuscript. Demonstrated consistency in synaptic strength in

spines sharing a presynaptic and postsynaptic input, presumed to arise after averaging over sources of stochasticity built into the system, was notable. Using a signal detection framework to assess just how consistent the control over synaptic size is (using spine head volume as a proxy for synaptic strength) enabled us a means of comparison across brain regions.

Chapter 3 consists of a follow up study to the work presented in Chapter 2. However, rather than beginning with a dense reconstruction, we began working from four sparsely traced series from the dentate gyrus MML layer. Further axon tracing was done in the context of this project for the purposes of identifying pairs on the same axonal branch with a shared axonal connection. Many of the additions to Neuropil Tools, including the "Processor Tool" module described in Chapter 4 were made during this phase of the dissertation research. Using two animals we were able to test the how the distribution of the spine population differed by head volume, as well as calculate how many bits of information were stored at the synapse in this region (using the same signal detection framework applied in Chapter 2), under both control and LTP stimulation conditions. This work provides impetus for additional replication of the signal detection application in diverse brain regions, to understand whether the dentate gyrus is unique in holding fewer bits of information at its synapses in MML, or if there is widespread variation across the cortex and hippocampus. A secondary analysis was included as part of this chapter, namely the "Jaccard Distance" analysis. The justification for this analysis was to test for the relationship between shape and volume, or shape and connectivity, as we noticed our original 10 pairs in the CA1 data were similar to one another in shape.

Chapter 4 is an explanation of the Neuropil Tools software package add-on to Blender that was developed and utilized as part of this work.

In total, this work bridges multiple specialties within the Neuroscience world, from molecular level imaging, to application of computational algorithms. It provides a baseline for comparatively, and quantitatively, evaluating information storage within and across brain regions, and furthers our understanding of the brain as a computational unit.

REFERENCES

Amaral D, Witter M. 1989. The three-dimensional organization of the hippocampal formation: A review of anatomical data. Neuroscience 31(3): 571-591.

Araki Y, Zeng M, Zhang M, Huganir RL. 2015. Rapid dispersion of SynGAP from synaptic spines triggers AMPA receptor insertion and spine enlargement during LTP. Neuron 85:173–189.

Ashby MC, Maier SR, Nishimune A, Henley JM. 2006. Lateral diffusion drives constitutive exchange of AMPA receptors at dendritic spines and is regulated by spine morphology. Journal of Neuroscience 26:7046–7055.

Bell ME, Bourne JN, Chirillo MA, Mendenhall JM, Kuwajima M, Harris KM. 2014. Dynamics of nascent and active zone ultrastructure as synapses enlarge during long-term potentiation in mature hippocampus. Journal of Comparative Neurology 522:3861–3884.

Bi GQ, Poo MM. 1998. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. The Journal of Neuroscience 18:10464–10472.

Billups B, Forsythe ID. 2002. Presynaptic mitochondrial calcium sequestration influences transmission at mammalian central synapses. The Journal of Neuroscience 22:5840–5847.

Bliss TV, Collingridge GL. 1993. A synaptic model of memory: long-term potentiation in the hippocampus. Nature 361 (

Bliss TV, Lomo T. 1973. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the performant path. Journal of Physiology 232: 331-356.

Bourne JN, Sorra KE, Hurlburt J, Harris KM. 2007. Polyribosomes are increased in spines of CA1 dendrites 2 h after the induction of LTP in mature rat hippocampal slices. Hippocampus 17:1–4.

Bourne JN, Harris KM. 2011. Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. Hippocampus 21:354–373.

Bourne JN, Chirillo MA, Harris KM. 2013. Presynaptic ultrastructural plasticity along CA3fiCA1 axons during long-term potentiation in mature hippocampus. The Journal of Comparative Neurology 521:3898–3912.

Bowden JB, Abrahan WC, Harris KM. 2012. Differential effects of train, circadian cycle and stimulation pattern on LTP and concurrent LTD in the dentate gyrus of freely moving rats. Hippocampus 22(6): 1363-70.

Branco T, Staras K, Darcy KJ, Goda Y. 2008. Local dendritic activity sets release probability at hippocampal synapses. Neuron 59:475–485.

Burel-Jungerman E, Davis S, Rampon C, Laroche S. 2006. Long-term potentiation enhances neurogenesis in the adult dentate gyrus. Journal of Neuroscience 26(22): 5888-5893.

Buzsáki, G, and Kenji M. 2014. The Log-Dynamic Brain: How Skewed Distributions Affect Network Operations. Nature Reviews Neuroscience 15 (4): 264–78.

Callaway JC, Ross WN. 1995. Frequency-dependent propagation of sodium action potentials in dendrites of hippocampal CA1 pyramidal neurons. Journal of Neurophysiology 74:1395–1403.

Cardona A, Saalfeld S, Schindelin J, Arganda-Carreras I, Preibisch S, Longair M, Tomancak P, Harstenstein V, Doublas RJ. 2012. TrakEM2 Software for Neural Circuit Reconstruction. PLoS ONE 7(6) e38011.

Caporale, N and Yang D. 2008. Spike Timing–Dependent Plasticity: A Hebbian Learning Rule. Annual Review of Neuroscience 31 (1): 25–46.

Carvalho Myskiw J, Furini CRG, Benetti F, Izquierdo I. 2014. Hippocampal molecular mechanisms involved in the enhancement of fear extinction caused by exposure to novelty. Proceedings of the National Academy of Sciences of the United States of America 111:4572–4577.

Chawla, M.K., J.F. Guzowski, V. Ramirez-Amaya, P. Lipa, K.L. Hoffman, L.K. Marriott, P.F. Worley, B.L. McNaughton, and C.A. Barnes. 2005. Sparse, Environmentally Selective Expression of Arc RNA in the Upper Blade of the Rodent Fascia Dentata by Brief Spatial Experience. Hippocampus 15 (5): 579–86.

Chicurel, Marina E., and Kristen M. Harris. 1992. Three-Dimensional Analysis of the Structure and Composition of CA3 Branched Dendritic Spines and Their Synaptic Relationships with Mossy Fiber Boutons in the Rat Hippocampus. The Journal of Comparative *Neurology* 325 (2): 169–82.

Claiborne, Brenda J., David G. Amaral, and W. Maxwell Cowan. 1990. Quantitative, Three-Dimensional Analysis of Granule Cell Dendrites in the Rat Dentate Gyrus. The Journal of Comparative Neurology 302 (2). Wiley-Liss, Inc.: 206–19.

Clarke LE, Barres BA. 2013. Emerging roles of astrocytes in neural circuit development. Nature Reviews Neuroscience 14:311–321.

Dan Y, Poo M. 2004. Spike Timing-Dependent Plasticity of Neural Circuits. Neuron 44(1):23-30.

Dobrunz LE, Huang EP, Stevens CF. 1997. Very short-term plasticity in hippocampal synapses. Proceedings of the National Academy of Sciences of the United States of America 94:14843–14847.

Edwards J, Daniel E, Kinney J, Bartol T, Sejnowski T, Johnston D, Harris K, Bajaj C. 2014. VolRoverN: enhancing surface and volumetric reconstruction for realistic dynamical simulation of cellular and subcellular function. Neuroinformatics 12:277–289.

Foley J, Legge G. 1981. Contrast detection and near-threshold discrimination in human vision. Vision Research 21(7): 1041-1053.

Fortin N, Agster K, Eichenbaum H. 2002. Critical role of the hippocampus in memory for sequences of events. Nature Neuroscience 5(5):458.

Foster D, Wilson M. 2006. Reverse replay of behavioral sequences in hippocampal place cells during the awake state. Nature 440 (7084):690-683.

Franks KM, Stevens CF, Sejnowski TJ. 2003. Independent sources of quantal variability at single glutamatergic synapses. The Journal of Neuroscience 23:3186–3195.

Frey U, Morris RGM. 1997. Synaptic tagging and long-term potentiation. Nature 385:533–536.

Gamerman D, Lopes HF. 2006. Markov Chain Monte Carlo. CRC Press.

Govindarajan A, Israely I, Huang S-Y, Tonegawa S. 2011. The dendritic branch is the preferred integrative unit for protein synthesis-dependent LTP. Neuron 69:132–146.

Green DM, Swets JA. 1966. Signal Detection Theory and Psychophysics. ed. Los Altos, CA: Peninsula Publishing.

Harris KM, Spacek J, Bell ME, Parker PH, Lindsey LF, Baden AD, Vogelstein JT, Burns R. 2015. A resource from 3D electron microscopy of hippocampal neuropil for user training and tool development. Scientific Data 2: 150046.

Harris, K M, and J K Stevens. 1989. Dendritic Spines of CA 1 Pyramidal Cells in the Rat Hippocampus: Serial Electron Microscopy with Reference to Their Biophysical Characteristics. The Journal of Neuroscience : The Official Journal of the Society for Neuroscience 9 (8): 2982–97.

Harris KM, Sultan P. 1995. Variation in the number, location and size of synaptic vesicles provides an anatomical basis for the nonuniform probability of release at hippocampal CA1 synapses. Neuropharmacology 34:1387– 1395.

Hiber-Schmidt C, Jonas P, Bischofberge J. 2004. Enhanced synaptic plasticity in newly generated cells in the adult hippocampus. Letters to Nature 429: 184-187

Ishizuka N, Cowan WM, Amaral DG. 1995. A quantitative analysis of the dendritic organization of pyramidal cells in the rat hippocampus. The Journal of Comparative Neurology 362:17–45.

Jorstad A, Nigro B, Cali C, Wawrzyniak M, Fua P, Knott G. 2015. NeuroMorph: A toolset for the morphometric analysis and visualization of 3D models derived from electron microscopy image stacks. Neuroinformatics 13: 83-93.

Kandaswamy U, Deng P-Y, Stevens CF, Klyachko VA. 2010. The role of presynaptic dynamics in processing of natural spike trains in hippocampal synapses. Journal of Neuroscience 30:15904–15914.

Kasthuri N, Hayworth KJ, Berger DR, Schalek RL, Conchello JA, Knowles-Barley S, Lee D, Va´zquez-Reina A, Kaynig V, Jones TR, Roberts M, Morgan JL, Tapia JC, Seung HS, Roncal WG, Vogelstein JT, Burns R, Sussman DL, Priebe CE, Pfister H, Lichtman JW, Schalek RL, Conchello JA, Knowles-Barley S, Lee D. 2015. Saturated reconstruction of a volume of neocortex. Cell 162:648–661.

Kennedy MB, Beale HC, Carlisle HJ, Washburn LR. 2005. Integration of biochemical signalling in spines. Nature Reviews Neuroscience 6:423–434.

Kinney JP, Spacek J, Bartol TM, Bajaj CL, Harris KM, Sejnowski TJ. 2013. Extracellular sheets and tunnels modulate glutamate diffusion in hippocampal neuropil. Journal of Comparative Neurology 521:448–464.

Koester HJ, Johnston D. 2005. Target cell-dependent normalization of transmitter release at neocortical synapses. Science 308:863–866.

Kramar EA, Babayan AH, Gavin CF, Cox CD, Jafari M, Gall CM, Rumbaugh G, Lynch G. 2012. Synaptic evidence for the efficacy of spaced learning. Proceedings of the National Academy of Sciences of the United States of America 109:5121–5126.

Krueppel, Roland, Stefan Remy, Heinz Beck. 2011. Dendritic Integration in Hippocampal Dentate Granule Cells. Neuron 71 (3). Elsevier: 512–28.

Kusters R, Kapitein LC, Hoogenraad CC, Storm C. 2013. Shape-induced asymmetric diffusion in dendritic spines allows efficient synaptic AMPA receptor trapping. Biophysical Journal 105:2743–2750.

Laughlin SB, Sejnowski TJ. 2003. Communication in neuronal networks. Science 301:1870–1874.

Leutgeb JK, Leutgeb S, Moser MB, Moser EI. 2007. Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. Science 315(5814): 961-966.

Leutgeb S, Leutgeb JK, Barnes C, Moser EI, McNaughton B, Moser MB. 2005. Independent Codes for Spatial and Episodic Memory in Hippocampal Neuronal Ensembles. Science 309 (5734):619-623

Lisman J, Harris K. 1994. Who's been nibbling on my PSD: is it LTD? Journal of Physiology-Paris 88:193–195.

Lopez-Rojas J, Kreutz MR. 2016. Mature granule cells of the dentate gyrus – Passive bystanders of principal performers in hippocampal function? Neuroscience and Biobehavioral Reviews 64: 167-174.

Lopez-Rojas J, Heine M Kreutz MR. 2016. Plasticity of intrinsic excitability in mature granule cells of the dentate gyrus. Scientific Reports 6(21615).

Magee JC, Johnston D. 1995. Characterization of single voltage-gated Na and Ca2 channels in apical dendrites of rat CA1 pyramidal neurons. The Journal of Physiology 487:67–90.

Mankin E, Sparks F, Slayyeh B, Sutherland R, Leutgeb S, Leutgeb JK. 2012. Neuronal code for extended time in the hippocampus. PNAS 109(47): 19462-7

Markram H, Lu ̈bke J, Frotscher M, Roth A, Sakmann B. 1997. Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. The Journal of Physiology 500:409– 440.

Megías M, Verduga R, Ferna´ndez-Viadero C, Crespo D. 1997. Neurons co-localizing calretinin immunoreactivity and reduced nicotinamide adenine dinucleotide phosphate diaphorase (nADPH-d) activity in the hippocampus and dentate gyrus of the rat. Brain Research 744:112–120.

Mills LR, Niesen CE, So AP, Carlen PL, Spigelman I, Jones OT. 1994. N-type Ca2 channels are located on somata, dendrites, and a subpopulation of dendritic spines on live hippocampal pyramidal neurons. The Journal of Neuroscience 14:6815–6824.

Mishchenko Y, Hu T, Spacek J, Mendenhall J, Harris KM, Chklovskii DB. 2010. Ultrastructural analysis of hippocampal neuropil from the connectomics perspective. Neuron 67:1009–1020.

Mizuseki K, Sirota A, Pastalkova E, Buzsaki G. 2009. Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. Neuron 64(2): 267-280.

Mongiat, Lucas A., and Alejandro F. Schinder. 2011. Adult Neurogenesis and the Plasticity of the Dentate Gyrus Network. European Journal of Neuroscience 33 (6): 1055–61.

Moser MB, Moser EI. 1998. Functional differentiation in the hippocampus. Hippocampus 8(6): 608-619.

Murthy VN, Schikorski T, Stevens CF, Zhu Y. 2001. Inactivity produces increases in neurotransmitter release and synapse size. Neuron 32:673–682.

Nadkarni S, Bartol TM, Sejnowski TJ, Levine H, McCulloch AD. 2010. Modelling vesicular release at hippocampal synapses. PLoS Computational Biology 6:e1000983.

Neunuebel, Joshua P., and James J. Knierim. 2012. Spatial Firing Correlates of Physiologically Distinct Cell Types of the Rat Dentate Gyrus. Journal of Neuroscience 32 (11).

Nimchinsky EA. 2004. The number of glutamate receptors opened by synaptic stimulation in single hippocampal spines. Journal of Neuroscience 24:2054–2064.

O'Connor DH, Wittenberg GM, Wang SS-H. 2005. Graded bidirectional synaptic plasticity is composed of switch-like unitary events. Proceedings of the National Academy of Sciences of the United States of America 102:9679–968

O'Donnell C, Sejnowski TJ. 2014. Selective memory generalization by spatial patterning of protein synthesis. Neuron 82:398–412.

Olton D, Meck W, Church R. 1987. Separation of hippocampal and amygdaloid involvement in temporal memory dysfuntions. Brain Research 404(1): 180-188

Ostroff LE, Fiala JC, Allwardt B, Harris KM. 2002. Polyribosomes redistribute from dendritic shafts into spines with enlarged synapses during LTP in developing rat hippocampal slices. Neuron 35:535–545.

Palm G. 1989. On the asymptotic information storage capacity of neural networks. Neural Computers: 271-280.

Redondo RL, Morris RGM. 2011. Making memories last: the synaptic tagging and capture hypothesis. Nature Reviews Neuroscience 12:17–30.

Saxe, Michael D, Fortunato Battaglia, Jing-Wen Wang, Gael Malleret, Denis J David, James E Monckton, A Denise R Garcia, Sofroniew MV, Kandel ER, Santarelli L, Hen R, Drew MR. 2006. Ablation of Hippocampal Neurogenesis Impairs Contextual Fear Conditioning and Synaptic Plasticity in the Dentate Gyrus. Proceedings of the National Academy of Sciences of the United States of America 103 (46). National Academy of Sciences: 17501–6.

Schikorski T, Stevens CF. 1997. Quantitative ultrastructural analysis of hippocampal excitatory synapses. The Journal of Neuroscience 17:5858–5867. Schultz S. 2007. Signal-to-noise ratio in neuroscience. Scholarpedia 2:2046.

Schmidt-Hieber, Christoph, Peter Jonas, and Josef Bischofberger. 2007. Subthreshold Dendritic Signal Processing and Coincidence Detection in Dentate Gyrus Granule Cells. Journal of Neuroscience 27 (31).

Scoville, W. B., & Milner, B. 1957. Loss of recent memory after bilateral hippocampal lesions. Journal of Neurology, Neurosurgery, and Psychiatry 20(1): 11.

Smith D, Mizumori S. 2006. Hippocampal place cells, context, and episodic memory. Hippocampus 16(9): 716-729.

Snyder, J S, N Kee, and J M Wojtowicz. 2001. Effects of Adult Neurogenesis on Synaptic Plasticity in the Rat Dentate Gyrus. Journal of Neurophysiology 85 (6): 2423–31.

Song S, Miller K, Abbott L. 2000. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. Nature Neuroscience 3(9):919-926.

Sorra KE, Harris KM. 1993. Occurrence and three-dimensional structure of multiple synapses between individual radiatum axons and their target pyramidal cells in hippocampal area CA1. The Journal of Neuroscience 13: 3736–3748.

Sorra KE, Fiala JC, Harris KM. 1998. Critical assessment of the involvement of perforations, spinules, and spine branching in hippocampal synapse formation. The Journal of Comparative Neurology 398:225–240.

Sorra KE, Mishra A, Kirov SA, Harris KM. 2006. Dense core vesicles resemble active-zone transport vesicles and are diminished following synaptogenesis in mature hippocampal slices. Neuroscience 141:2097–2106.

Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. 2014. Dropout: a simple way to prevent neural networks from overfitting. The Journal of Machine Learning Research 15:1929–1958.

Stanislaw H and Todorov N. 1999. Calculation of signal detection theory measures. Behavior Research Methods, Instruments & Computers 3:37-149

Sutton MA, Schuman EM. 2006. Dendritic protein synthesis, synaptic plasticity, and memory. Cell 127:49–58.

Thompson L, Best P. 1989. Place cells and silent cells in the hippocampus of freely-behaving rats. Journal of Neuroscience 9(7): 2382-2390.

Turner E. 1969. Hippocampus and memory. The Lancet 294(7630): 1123-1126.

Turrigiano GG. 2008. The self-tuning neuron: synaptic scaling of excitatory synapses. Cell 135:422–435.

Tønnesen J, Katona G, Ro´zsa B, Na ¨gerl UV. 2014. Spine neck plasticity regulates compartmentalization of synapses. Nature Neuroscience 17:678–685.

Van Strien N, Cappaert N, Witter M. 2009. The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network

Ventura R, Harris KM. 1999. Three-dimensional relationships between hippocampal synapses and astrocytes. The Journal of Neuroscience 19:6897–6906.

Wan L, Zeiler M, Zhang S, Cun YL, Fergus R. 2013. Regularization of Neural Networks Using DropConnect:1058– 1066.

Witcher MR, Kirov SA, Harris KM. 2007. Plasticity of perisynaptic astroglia during synaptogenesis in the mature rat hippocampus. Glia 55:13–23.

Yang, Kechun, and John A Dani. 2014. Dopamine D1 and D5 Receptors Modulate Spike Timing-Dependent Plasticity at Medial Perforant Path to Dentate Granule Cell Synapses. The Journal of Neuroscience 34 (48): 15888–97.

Yasuda R, Sabatini BL, Svoboda K. 2003. Plasticity of calcium channels in dendritic spines. Nature Neuroscience 6: 948–955.