The Dissertation Committee for Wenke Li

certifies that this is the approved version of the following dissertation:

TIMING IN THE CEREBELLUM: A MATTER OF NETWORK INHIBITION

Committee:

Michael D. Mauk, Supervisor

Alexander C. Huk

Hiroshi Nishiyama

Eyal Seidemann

Peter H. Stone

TIMING IN THE CEREBELLUM: A MATTER OF NETWORK INHIBITION

by

Wenke Li, B.S.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

The University of Texas at Austin

May 2015

DEDICATION

This work is dedicated to the following people for their non-academic contributions:

To my parents, Chuan-fen and Jing-yan: I am fortunate to have you as my role model in science, and in life.

To friends and past members of the lab: Jenni, Randy, Chris, Andrei, Liz, Nick and Sam. Thank you for your discussions and inspirations, for making me think and grow. To Jenni, thank you for showing me how to walk the path of science.

To my advisor, Mike, for giving me the opportunity to do this work. I am grateful to be able to work on what I dreamed of since high school.

ACKNOWLEDGEMENTS

I would like to acknowledge the following people for their individual contributions to this work: Alex Huk, Eyal Seidemann, Hiroshi Nishiyama, Peter Stone, and Mike Mauk. Students and Post-docs: Matthew Hausknecht, Vyash Puliyadi, Randy Chitwood, and Jennifer Siegel.

TIMING IN THE CEREBELLUM: A MATTER OF NETWORK INHIBITION

Wen-Ke Li, Ph.D.

The University of Texas at Austin, 2015

Supervisor: Michael D. Mauk

The motor functions of an animal require precisely timed and coordinated sequences of movements. The cerebellum is crucial for performing these functions with precision. To investigate cerebellar computations involved in precise motor movements, behavioral paradigms such as delay eyelid conditioning have been used. Delay eyelid conditioning trains an animal to close its eye in response to a previously neutral stimulus. The timing of the eyelid closure responses suggests that the cerebellum is capable of keeping track of the elapsed time since the onset of the stimulus. This dissertation proposes a network mechanism for cerebellar timing based on biologically informed simulations of the cerebellum. In chapter 2, a simulation with over a million cells is described. This simulation approaches the observed cerebellar connectivity in several well studied mammals. Graphics processing units (GPUs) provide the computational power necessary to perform this simulation at a practical speed. This chapter describes simulation algorithms that efficiently utilize GPUs. In

chapter 3, the simulation is used to explore cerebellar timing mechanisms. The lateral inhibition among cerebellar Golgi cells is observed to be a potential mechanism for robust timing. Lateral Golgi inhibition enables the simulation to better replicate animal eyelid conditioning behavior for longer inter-stimulus intervals. In chapter 4, the emergent network mechanisms of lateral Golgi inhibition are analyzed by decomposing the network into its individual components. This component analysis demonstrates that nonreciprocal connectivity (where one Golgi cell inhibits another but does not receive inhibition in return) is useful for timing. Specifically, removing nonreciprocal connectivity greatly degrades the simulation's ability to keep track of time. This implies that the aforementioned component analyses are relevant to the emergent timing mechanisms of the network. Finally, in chapter 5, this dissertation discusses the relevance and limitations of the computational approach, biological predictions, and component analysis presented in previous chapters.

Table of Contents

LIST OF FIGURES x
CHAPTER 1: 1
GENERAL INTRODUCTION 1
Cerebellar network architecture 4
Cerebellar cortical input network 8
Cerebellar output network 11
Delay eyelid conditioning is an important tool for investigating cerebellar computation 14
Theoretical models for addressing computation in the cerebellum 20
Using bottom-up simulations and eyelid conditioning to investigate cerebellar timing 26
CHAPTER 2: 30
EXPANDING THE CEREBELLAR SIMULATION TO APPROCH THE OBSERVED CONNECTIVITY 30
Introduction 31
Simulation connectivity 40
Mossy fiber-Golgi-granule network 41
Purkinje-basket-stellate network 44
Purkinje-deep nucleus-inferior olive network 45
Neuron representation in the simulation 47
Granule cell calculations in the GPU 50
Updating inputs to granule cells in the GPU 52
Updating granule output to other cells in the GPU 53
Adjusting the weight of granule-Purkinje synapses 55
CPU calculations and communication to the GPU 57
Summary 58
CHAPTER 3: 60
MECHANISMS OF TIMING IN THE CEREBELLUM 60
Introduction 61
Methods 66
Simulating delay eyelid conditioning 66
Data recording 67
Data analysis 67
Results 69

Timing in the expanded simulation with low granule-Golgi convergence ratios....69
Lateral Golgi inhibition is a mechanism for timing in the expanded simulation....75
Summary and discussion....82
CHAPTER 4:....85
COMPUTATIONAL MECHANISMS OF LATERAL GOLGI INHIBITION....85
Introduction....86
Methods....88
Manipulating Golgi network connectivity....88
Producing isolated Golgi cell networks to test the contributions of components of the network....89
Results....90
Early Golgi cells are important for generating a stimulus-temporal code....90
Component analysis of the inhibition to early cells....94
Testing the necessity of the nonreciprocal inhibition to early Golgi cells for the emergence of stimulus-temporal code....98
Nonreciprocal inhibition enhances stimulus-temporal code....99
Summary and discussion....104
CHAPTER 5:....108
LIMITATIONS AND SIGNIFICANCE....108
Summary....109
Implications for future scaling the simulated cerebellar network....113
Utilizing graphics processing units for high performance computing....115
Simulation predictions of the timing mechanism in the cerebellum....116
Network connectivity properties illustrated by the simulation....122
REFERENCES....127

LIST OF FIGURES

Figure 1.1. Cerebellar connectivity..... 6
Figure 1.2. The delay eyelid conditioning paradigm..... 16
Table 2.1. Connectivity ratios of the 12000 granule cell simulation, the one million granule cell simulation, and observed connectivity.....32
Figure 3.1. A mechanism with which the cerebellum can generate timed responses with eyelid conditioning..... 65
Figure 3.2. The expanded simulation with 4096 granule inputs per Golgi cell does not exhibit well timed responses.....71
Figure 3.3. Decreasing the number of granule inputs per Golgi cell improves stimulus-temporal code for Golgi and granule cell population.....74
Figure 3.4. Lateral Golgi inhibition improves stimulus-temporal code in the expanded simulation and produced well timed responses.....78
Figure 3.5. The activity of Golgi cells does not reflect the inputs from granule cells for the expanded simulation with lateral Golgi inhibition.....80
Figure 3.6. Lateral Golgi inhibition is necessary and sufficient to produce Golgi stimulus-temporal code.....81
Figure 4.1. Early Golgi cells are specifically important for producing stimulus-temporal code..... 93
Figure 4.2. Early cell activity in isolated simulations shows that nonreciprocal 1st order cells are the primary factor for decreasing early cell activity.....97
Figure 4.3. Nonreciprocal inhibition to early cells is important for producing stimulus-temporal code.....101
Figure 4.4. Exclusive reciprocal lateral Golgi inhibition does not produce stimulus-temporal code.....102
Figure 4.5. Networks containing only Golgi cells show that nonreciprocal inhibition generated stimulus-temporal code in a wide range of inhibitory synaptic strengths..... 103

CHAPTER 1:

GENERAL INTRODUCTION

The motor functions of an animal require precisely timed and coordinated sequences of contractions of different muscles. The cerebellum has been shown to be necessary to perform these functions with precision. Animals and humans with lesions to the cerebellum exhibit impaired motor functions (Bastian, Zackowski, & Thach, 2000; Flament, Vilis, & Hore, 1984; Manto et al., 2012; Palliyath & Hallett, 1998; Topka, Konczak, & Dichgans, 1998). Existing evidence suggests that while the cerebral motor cortex is necessary for initiating voluntary motor movement (Arezzo & Vaughan, 1975; Davey & Romaiguere, 1994; Deecke, Scheid, & Kornhuber, 1969; Roland & Larsen, 1980), the cerebellum provides the tuning signals to make such movements precise (Flanagan & Wing, 1993; Manto et al., 2012; Nowak, Topka, Timmann, Boecker, & Hermsdörfer, 2007). The computations performed by the cerebellum to generate these tuning signals have been investigated by several behavioral paradigms, such as smooth pursuit (Lisberger & Fuchs, 1978; Medina & Lisberger, 2007; Stone & Lisberger, 1990), vestibulo-ocular reflex (DuLac, Raymond, Sejnowski, & Lisberger, 1995; Ito, 1982; Miles & Lisberger, 1981), and delay eyelid conditioning (Garcia, Steele, & Mauk, 1999; Mauk, Steinmetz, & Thompson, 1986; Mauk & Thompson, 1987; McCormick & Thompson, 1984; Steinmetz et al., 1987; Steinmetz, Lavond, & Thompson, 1989). These paradigms have revealed the timing (Jirenhed & Hesslow, 2011a; Li & Lisberger, 2011; Medina, Garcia, Nores, Taylor, & Mauk, 2000), amplitude (DuLac et al., 1995; Kreider & Mauk, 2010), and adaptability

(Boyden, Katoh, & Raymond, 2004; DuLac et al., 1995; Medina, Nores, & Mauk, 2002; Perrett & Mauk, 1995) properties of cerebellar output. Experimental results based on these paradigms suggest that the cerebellum is capable of keeping track of time internally (Ivry, Spencer, Zelaznik, & Diedrichsen, 2002), specifying the amount of output to produce (Flanagan & Wing, 1993; Kreider & Mauk, 2010; MacKay & Murphy, 1979; Nowak et al., 2007), and adapting to new conditions with new output (Bastian, 2006; Carey & Lisberger, 2002; Contreras-Vidal, Grossberg, & Bullock, 1997; Morton & Bastian, 2006; Ohyama, Nores, Murphy, & Mauk, 2003). However, these claims are not universally accepted. Harrington et al. (2004) conclude that impaired sensory and cognitive information transfer can explain the impact on timing from cerebellar lesions. There are also alternate theories of cerebellar computation that do not involve learning (Llinas, Lang, & Welsh, 1997; Llinás & Welsh, 1993; Pellionisz & Llinas, 1979; Welsh et al., 2005).

In addition to well defined behaviors, investigations into cerebellar computation also benefit from detailed observations of the network architecture of the cerebellum (Eccles, Ito, & Szentágothai, 1967; Ito, 1984, 2006a). The connectivity and physiology of the neurons in the cerebellum have been studied for over a century (Sotelo, 2003). These studies (Eccles et al., 1967; Ito, 1984, 2006b) provide important information for proposing biologically constrained and relevant models of cerebellar computation (Buonomano & Mauk, 1994; Marr,

1969; Medina et al., 2000; Pellionisz, Llinas, & Perkel, 1977; Pellionisz, 1973).

Cerebellar network architecture

The physiology and connectivity of the cerebellar architecture have been extensively studied (Eccles et al., 1967; Ito, 1984). The microzone hypothesis of the cerebellum postulates that the cerebellum is divided into different functional areas (Balaban, Schuerger, & Porter, 2000; Ito, 1984; Oscarsson, 1979; Sugihara, 2006). Each functional area (microzone) is responsible for driving a muscle or set of related muscles (Gibson, Robinson, Alam, & Houk, 1987). However, the cerebellum has been observed to be involved in extra-motor functions such as cognition (Schmahmann, 2004; Strick, Dum, & Fiez, 2009; Van Overwalle, Baetens, Mariën, & Vandekerckhove, 2013). The connectivity of each microzone is relatively uniform (Ito, 1984) in the sense that the connectivity of one microzone is similar to every other microzone (but see (DiÑO, Willard, & Mugnaini, 1999) for differences). Classical observations (Eccles et al., 1967) suggest that there are a limited number of cell types and input pathways. Much of the connectivity has been observed in detail (Ito, 2006b; Palkovits, Magyar, & Szentágothai, 1972). However, in recent years additional connectivity has been observed (Hull & Regehr, 2012; Vervaeke, Lorincz, Nusser, & Silver, 2012; Xu & Edgley, 2008), suggesting the cerebellum network architecture is more complex and interconnected than previously thought.

There are two major input pathways into the cerebellum: mossy fibers from the pontine nuclei (Ito, 1984), and climbing fibers from the inferior olivary complex (Desclin, 1974; Shinoda, Sugihara, Wu, & Sugiuchi, 2000). The mossy fiber activity has been observed to correlate with changes in limb positions (van Kan, Gibson, & Houk, 1993), and the presence of various sensory stimuli (Aitkin & Boyd, 1978; K Maekawa & Takeda, 1975). In addition, the primary motor cortex has been observed to project to the pontine nuclei (Kelly & Strick, 2003). These projections have been suggested to carry motor commands to the cerebellum (Ito, 2005). Based on these observations, it is inferred that the mossy fibers carry information about the state of the world to the cerebellum (Ito, 1984; Marr, 1969). In contrast, the climbing fiber activity has been observed to correlate with the onset of unexpected stimuli (Andersson & Armstrong, 1987; Gellman, Gibson, & Houk, 1985; Gilbert & Thach, 1977; Simpson, Wylie, & De Zeeuw, 1996). The climbing fiber activity is observed through the complex spike activity of Purkinje cells (Bell & Grimm, 1969; Eccles, Llinas, & Sasaki, 1966; Ito & Simpson, 1971). The climbing fibers have been hypothesized to carry error signals to the cerebellum that instruct the cerebellum to modify its output (Marr, 1969; Simpson et al., 1996).

The mossy fiber inputs to the cerebellum connect to a large number of granule cells in the cerebellar cortex (Palkovits, Magyar, & Szentágothai, 1971b; Shinoda et al., 2000), and the granule cells in turn provide inputs to the Purkinje

cells (Eccles et al., 1967). The Purkinje cells also receive inputs from climbing fibers (Eccles et al., 1967; Ito & Simpson, 1971). The Purkinje cell axons are the only output of the cerebellar cortex. They provide inhibition onto cerebellar deep nucleus cells (M. Ito, Yoshida, Obata, Kawai, & Udo, 1970; Palkovits, Mezey, Hámori, & Szentágothai, 1977). The deep nucleus cell axons are the output of the cerebellum (Dum & Strick, 2003; Ito, 1984; Middleton & Strick, 1998).

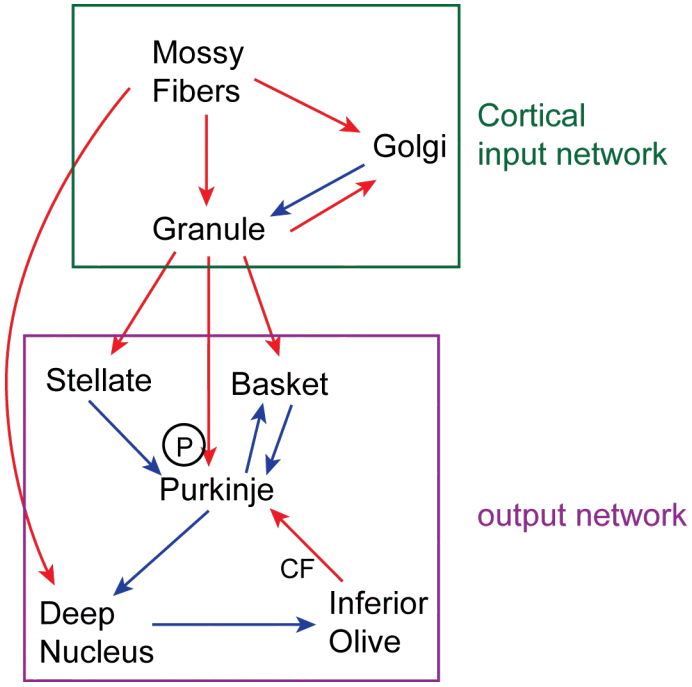

Figure 1.1. Cerebellar connectivity. Red arrows: excitatory connections. Blue arrows: inhibitory connections. P: plasticity at granule-Purkinje synapses. CF: climbing fibers. The granule cells from the cortical input network (green) provide input to the basket, stellate, and Purkinje cells of the output network (purple). The Purkinje cells inhibit the deep nucleus cells in the output network that are the output of the cerebellum.

The major cell types present in all areas of the cerebellum are granule cells (Chadderton, Margrie, & Häusser, 2004; Gabbiani, Midtgaard, & Knöpfel, 1994), Golgi cells (Galliano, Mazzarello, & D'Angelo, 2010), Purkinje cells (Eccles et al., 1967), stellate cells (V Chan-Palay & Palay, 1972), basket cells (O'Donoghue, 1989; Palkovits, Magyar, & Szentágothai, 1971c), and deep cerebellar nucleus cells (Jahnsen, 1986; Llinás & Mühlethaler, 1988; Ohyama, Nores, Medina, Riusech, & Mauk, 2006; Palkovits et al., 1977), in addition to the mossy fiber and climbing fiber input pathways. The architecture of the network (figure 1.1) can be divided into two subnetworks: 1. the cortical input network (D'Angelo & De Zeeuw, 2009; Kanichay & Silver, 2008; Mapelli & D'Angelo, 2007), which consists of mossy fibers, Golgi cells, and granule cells, for which the granule cells produce the primary output, (Eccles et al., 1967; Palkovits et al., 1971c) and 2. the output network, which consists of Purkinje cells, basket cells, and stellate cells, all of which receive input from granule cells (Palkovits et al., 1971c). The output network also contains inferior olivary cells (that provide climbing fiber input), mossy fibers, and deep nucleus cells. The deep nucleus cells provide the only output of the cerebellum. The predominant connectivity between the input and output network is the granule cell output to Purkinje, basket, and stellate cells. However, there is evidence that cerebellar Lugaro cells (Melik-Musyan & Fanardzhyan, 2004) receive inputs from Purkinje cells and can potentially provide input to Golgi cells. This would provide a feedback connection

from the output network (Lainé & Axelrad, 1996, 1998).

Cerebellar cortical input network

The cortical input network consists of a large number of granule cells (Lange, 1975), and a small number of mossy fibers and Golgi cells in the granule layer (Palkovits et al., 1971b). The total granule cell population in the cerebellum has been reported to account for over 50% of the total number of neurons in the mammalian central nervous system (Ito, 1984). These cells vastly outnumber the Golgi cells (Palkovits et al., 1971b) (5000 granule cells per Golgi cell) and mossy fibers. The Golgi cells have been reported to be distributed in a two dimensional grid in the granule layer (Palkovits et al., 1971b), while the spatial distribution of mossy fibers terminals (part of the glomeruli) has been suggested to follow the microzones on a gross scale (Ji & Hawkes, 1994). However, the significance of this possible arrangement remains unclear. With the exception of granule input to Golgi cells, the connectivity between these cell types primarily occurs in the granule layer through the glomeruli.

The glomeruli are synaptic structures in the granule layer of the cerebellar cortex that connect the mossy fibers, Golgi cells, and granule cells together (Eccles et al., 1967; Ito, 1984; Jakab & Hámori, 1988; Palkovits et al., 1971b; Spacek, Parízek, & Lieberman, 1973). Each glomerulus is composed of multiple granule cell dendrites (between 20-110), (Eccles et al., 1967; Palkovits et al.,

1972), a single Golgi cell axon terminal, occasionally a Golgi cell descending dendrite (Hámori & Szentágothai, 1966), and a single mossy fiber terminal (Eccles, Llinás, & Sasaki, 1966). Through these structures, the mossy fibers provide excitatory output to the granule cells and Golgi cells (Hámori & Szentágothai, 1966), and Golgi cells provide inhibitory output to granule cells (Hámori & Szentágothai, 1966; S. Mitchell & Silver, 2000). However, see Victoria Chan-Palay and Palay (1971) for observations that mossy fibers direct contact Golgi soma. The number of gromeruli terminals per mossy fiber has been estimated to be between 16 (Eccles et al., 1967) and 44 (Palkovits et al., 1971b). This, combined with the number of granule dendrites per glomerulus, results in 320-4400 granule cells dendrites per mossy fiber. The Golgi cell axons have been reported to connect to 60-100 gromeruli near the cell body (Eccles et al., 1967; Palkovits et al., 1971b), and result in a divergence of 1200-11,000 granule cell dendrites per Golgi axon. Golgi cells also have descending dendrites (Hámori & Szentágothai, 1966) that connect to the glomeruli and receive mossy fiber excitatory inputs (Ito, 1984). Each granule cell dendrite has been observed to only connect to a single glomerulus (Eccles et al., 1967), and most granule cells have been observed to have 4 dendrites (Eccles et al., 1967). The length of dendrites has been reported to be 10-25um (Ito, 1984), which constrains a granule cell to only receive inputs from the glomeruli that are close to the cell.

The granule inputs to Golgi cells are located in the molecular layer of the

cerebellar cortex (Palay, 1974), which contains the axons and dendrites of various neurons (Palkovits et al., 1971c). The granule cell axons first ascend to the molecular layer and then bifurcate in the same direction to produce parallel fibers (Eccles et al., 1967). The parallel fibers provide excitatory inputs to Golgi cells and other neurons in the output network. The length of the parallel fibers has been reported to be between 1-10 mm (Brand, Dahl, & Mugnaini, 1976; Ito, 1984), depending on the region of the cerebellum and the animal species. The Golgi ascending dendrites have been observed to extend throughout the entire molecular layer (Palay, 1974). The dendrites have been observed to not have many branches and are spatially sparse (Eccles et al., 1967). Given the orientation of the parallel fibers, a Golgi cell can potentially receive inputs from granule cells that are far away from the Golgi cell but are located in the same direction as the parallel fibers (Volny-Luraghi, Maex, Vos, & De Schutter, 2002). In contrast, a granule cell that is located close to the Golgi cell but perpendicular to the parallel fiber orientation would not provide input for that Golgi cell.

In addition to mossy fibers, Golgi cells, and granule cells, uni-polar brush cells (Mugnaini, Sekerková, & Martina, 2011) have also been observed in the granule layer of the cerebellar cortex. These cells receive inputs from mossy fibers and Golgi cells and provide excitatory output to granule cells (Dino, Schuerger, Liu, Slater, & Mugnaini, 2000; Nunzi & Birnstiel, 2001). These cells have been observed to be concentrated in the medial vestibular regions of the

cerebellum (Diño, Nunzi, Anelli, & Mugnaini, 2000), and are less common in the lateral regions (DiÑO et al., 1999).

Cerebellar output network

The output network consists of basket cells, stellate cells, Purkinje cells, deep nucleus cells and inferior olivary cells. The basket, stellate, and Purkinje cells are in the cerebellar cortex, whereas the deep nucleus cells are in the deep nuclei of the cerebellum, and inferior olivary cells are in the inferior olivary nuclei in the brainstem. The Purkinje, basket, and stellate cells are thought to be arranged into microzones (Balaban et al., 2000; Oscarsson, 1979; Ozden, Sullivan, Lee, & Wang, 2009; Pijpers, Voogd, & Ruigrok, 2005), where each functional unit is aligned in the direction perpendicular to the parallel fibers. All three types of cells receive excitatory input from the parallel fibers. The Purkinje cell dendrites are sheet-like structures that permeate throughout the height of the molecular layer (Eccles et al., 1967; Sotelo, 2003). The orientation of each dendritic sheet is perpendicular to the direction of the parallel fibers (Fox & Barnard, 1957). It has been hypothesized that this arrangement maximizes the number granule inputs to Purkinje cells with respect to a fixed amount of space (Ito, 1984). Each Purkinje cell has been observed to receive between 80,000 (Palkovits et al., 1971c) to 200,000 (Eccles et al., 1967) granule cell inputs. The granule-Purkinje synapses have been suggested as a site of plasticity that

mediate learning in the cerebellar cortex (Albus, 1975; Marr, 1969), which is further discussed below. The dendritic structures of basket and stellate cells are in the same general orientation as Purkinje dendrites (Mertz, Koscheck, & Schilling, 2000), but less regular (V Chan-Palay & Palay, 1972).

The primary output of basket and stellate cells has been observed to inhibit Purkinje cells. The stellate cell axons have been reported to inhibit parts of Purkinje cell dendrites (V Chan-Palay & Palay, 1972; Eccles et al., 1967; Midtgaard, 1992), whereas the basket cells inhibit the Purkinje cell body (O'Donoghue, 1989; Palkovits et al., 1971c). The basket cell axons are arranged in the direction perpendicular to the parallel fibers (Palkovits et al., 1971c), which is consistent with the orientation of the functional units. The Purkinje cells have also been observed to inhibit basket cells (O'Donoghue, 1989). The Purkinje cell axons are the only output from the cerebellar cortex (Eccles et al., 1967). They provide inhibition to the deep cerebellar nucleus (Ito et al., 1970; Zheng & Raman, 2010).

The deep cerebellar nucleus cells have been observed to provide the only output of the cerebellum (Palkovits et al., 1977), and receive excitatory input from mossy fiber collaterals (Shinoda, Sugiuchi, Futami, & Izawa, 1992) (however, see (Brodal, Dietrichs, & Walberg, 1986) for different observations) and inhibition from Purkinje cells (Ito et al., 1970). The mossy fiber-deep nucleus synapses have been observed to be plastic and mediated by inhibition from the Purkinje

cells (Ohyama et al., 2006; Pugh & Raman, 2006, 2008). In the cat, each nucleus cell has been estimated to receive inhibition from up to 200 Purkinje cells (Palkovits et al., 1977). Deep nucleus cells have been observed to provide inhibition to inferior olivary neurons (Best & Regehr, 2009; Lang, Sugihara, & Llinás, 1996) and excitatory output to the downstream areas such as the red nucleus (Asanuma, Thach, & Jones, 1983; Flumerfelt, Otabe, & Courville, 1973). The inferior olivary neurons provide the climbing fiber inputs into the cerebellum (Desclin, 1974) that are thought to carry teaching signals for the cerebellum to modify its output (Marr, 1969; Simpson et al., 1996; Türker & Miles, 1986). These neurons appear to have extensive electrical synapses among their dendrites (Placantonakis, Bukovsky, Aicher, Kiem, & Welsh, 2006). The axon of each inferior olivary neuron provides climbing fibers to multiple Purkinje cells (Desclin, 1974; H. Fujita & Sugihara, 2013), whereas each Purkinje cell receives a single climbing fiber input (Eccles et al., 1967). It has been suggested that the Purkinje cells that receive climbing fiber input from the same inferior olivary cell all perform the same functions since they receive the same teaching signals (Ito, 2000).

In summary, the network architecture of the cerebellum has been studied for over a century, and has been characterized in considerable detail. These characterizations enable the construction of biologically informed models, and enable these models to provide relevant and testable hypotheses of the

computational properties and mechanisms of the cerebellum.

Delay eyelid conditioning is an important tool for investigating cerebellar computation

Investigations of the computational properties of the cerebellum also enjoy the advantages of well defined behaviors that directly engage the input and output of the cerebellum (DuLac et al., 1995; Li & Lisberger, 2011; McCormick & Thompson, 1984). Such a behavior needs well defined and controllable inputs to the cerebellar network, and the network needs to produce clearly characterizable behavioral output to provide the experimental foundations to test hypotheses about the cerebellum. Ideally, the behavioral output should capture the richness of cerebellar functions such that the behavior can be used to investigate all aspects of cerebellar computation. The delay eyelid conditioning paradigm (figure 1.2) is one such behavior that provides the experimental basis for testing hypotheses about cerebellar computation (Medina et al., 2000).

The delay eyelid conditioning paradigm has been shown to directly engage the cerebellum (McCormick & Thompson, 1984). Animals with the cerebral cortex removed could still learn this paradigm (Mauk & Thompson, 1987), whereas animals with lesions in the cerebellum exhibited impaired learning (Garcia et al., 1999; Lavond, Hembree, & Thompson, 1985; Steinmetz, Logue, & Steinmetz, 1992). While the behavioral output of this paradigm does

not reflect *all* of the functions in the cerebellum (such as coordination between different muscles (Ramnani, Toni, Passingham, & Haggard, 2001; Thach, Goodkin, & Keating, 1992; Thach, 1998)), the output does capture three essential features thought to be useful for tuning signals for motor control: adaptability (Kehoe & Holt, 1984; Medina, Garcia, & Mauk, 2001), timing (White, Kehoe, Choi, & Moore, 2000), and amplitude (Kreider & Mauk, 2010). Each trial of this paradigm presents the animal with a conditioned stimulus (CS) such as an auditory tone. After a fixed delay to the onset of the CS, the unconditioned stimulus (US) such as an air puff to the eye is presented to the animal. Initially, the animal reflexively closes its eyelid after the onset of the US. However, after many trials, the animal learns to associate the CS with the US by closing to its eyelid to the CS prior to the onset of the US, even if the US is absent.

The delay eyelid conditioning paradigm is a useful tool for investigating the cerebellum because the inputs to the cerebellum during the behavior have been studied and can be manipulated (Steinmetz et al., 1989). Auditory sensory stimuli (for the CS) have been observed with *in vivo* recording of mossy fibers (Aitkin & Boyd, 1978), and the US inputs have been characterized by the Purkinje cells' responses to climbing fiber inputs (Rasmussen, Jirenhed, & Hesslow, 2008). Furthermore, the CS and US inputs can be replaced by stimulating mossy fibers

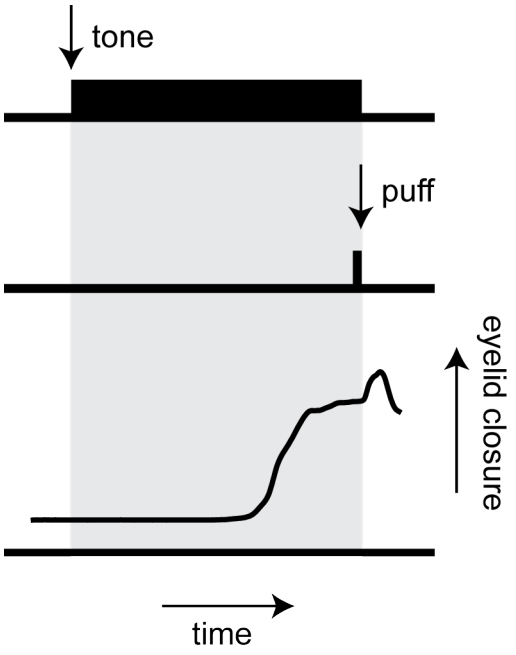

Figure 1.2. The delay eyelid conditioning paradigm. A single trial of the paradigm is shown. The conditioned stimulus (CS, tone), is a tonic stimulus that is presented to an animal and persists at least until the onset of the unconditioned stimulus (US, puff to the eye). After training, the animal learns to close its eyelid to the CS prior to the onset of the US.

(Hesslow, Svensson, & Ivarsson, 1999; Steinmetz, Rosen, Woodruff-Pak, Lavond, & Thompson, 1986; Steinmetz, 1990) and climbing fibers (Mauk et al., 1986) respectively, which provide effective tools for manipulating the inputs to the cerebellum. For example, when the auditory stimuli are replaced by mossy fiber stimulations, the stimulation (CS input) must persist until the onset of the US in order for the animal to learn to respond (Kalmbach, Voicu, Ohyama, & Mauk, 2011). In contrast, when using an auditory tone as the CS input, the CS could terminate well before the US onset and the animal can still learn to respond (Trace eyelid conditioning) (Kalmbach, Ohyama, Kreider, Riusech, & Mauk, 2009; Woodruff-Pak, Lavond, & Thompson, 1985). Given these results, it is predicted that for auditory CS input, other brain regions are involved in providing an input that persists though the stimulus-free interval. The ability of the cerebellum to respond to such types of inputs was tested using dual mossy fiber stimulations as the CS, where one stimulation terminated early and the other spanned the stimulus-free interval. This showed that the animal could learn to respond to the dual stimulation input (Kalmbach et al., 2011). In searching for the source of the persistent input, the medial prefrontal cortex was shown to be necessary (Kalmbach et al., 2009). Cells in the prefrontal cortex have been been observed to produce persistent activity (Siegel, Kalmbach, Chitwood, & Mauk, 2012; Siegel & Mauk, 2013) that spans the stimulus-free interval in response to the auditory input.

Animals trained using delay eyelid conditioning are capable of producing specific levels of responses; the conditioned eyelid closure responses could be partial (Kreider & Mauk, 2010). This suggests that delay eyelid conditioning can be used to study the amplitude control of the cerebellar output, which has been demonstrated previously in other behavior paradigms such as vestibulo-ocular reflex (Robinson, 1976; E. Watanabe, 1984). Delay eyelid conditioning has been used to study timing in the cerebellum because the conditioned eyelid closure responses exhibit timing (Medina et al., 2000; Ohyama & Mauk, 2001). After learning to respond to the CS, the animal does not close its eyelid at the onset of the CS, but delays the response to shortly prior to the onset of the US. The conditioned eyelid closure responses exhibit different timing with training at different intervals between the CS and US onset (White et al., 2000). For example the responses to 750ms interval are more delayed compared to 500ms interval, which suggest that the cerebellum can keep track of the amount of time that has passed since the onset of the CS (Perrett, Ruiz, & Mauk, 1993).

It is strongly suspected that the Purkinje cells are necessary to produce these well timed responses. The four primary reasons for this are as follows. 1. Purkinje cells are the sole output of the cerebellar cortex (Marr, 1969; Medina et al., 2002; Ohyama et al., 2003; Simpson et al., 1996). 2. Purkinje cell recordings during behavior show a decrease in activity immediately prior to the onset of behavior (Jirenhed & Hesslow, 2011a, 2011b; Rasmussen et al., 2008;

Svensson, Jirenhed, Bengtsson, & Hesslow, 2010). 3. Optogenetic manipulations to decrease Purkinje activity produce muscle responses (Heiney, Kim, Augustine, & Medina, 2014). Finally, 4. lesioning the cerebellar cortex disrupt well timed responses (Kalmbach et al., 2010; Perrett et al., 1993).

The granule cell activity has been proposed to produce timing information (time since the onset of the CS) for Purkinje cells. This enables Purkinje cells to decrease activity immediately prior to the onset of the US. The granule cells were proposed to produce timing information for the following reasons. 1. Purkinje cells receive a large number of inputs from granule cells. 2. Granule cells convey the input from mossy fibers (that carry the CS). 3. Plasticity at granule-Purkinje synapses have been implicated in changes in Purkinje cell activity (Ito & Kano, 1982; Ito, 2001).

In summary, delay eyelid conditioning is a behavior that directly engages the input and output of the cerebellum. The behavior output of this paradigm captures the cerebellum's ability to learn responses to new stimuli, and generate precise amplitude output. This paradigm has shown that the cerebellum is capable of producing well timed responses, and has been used to study the timing mechanisms in the cerebellum. Combined with physiological and anatomical evidence, the cerebellar granule cells have been proposed to produce the temporal information so that Purkinje cells can learn well timed responses. The theories of cerebellar computation relevant to timing and learning

are discussed below.

Theoretical models for addressing computation in the cerebellum

The detailed characterization of the cerebellar network architecture and well defined behaviors such as delay eyelid conditioning provide the tools necessary for producing experimentally testable computational models of the cerebellum. Based on the known cerebellar connectivity, many models of the cerebellum attempt to hypothesize about the computational properties of two features in the network: 1. the role of the immense population of granule cells and their connectivity to mossy fibers and Golgi cells of the input network, and 2. the role of Purkinje cells which receive a large number of granule cell inputs and are the only cells receiving climbing fibers inputs.

The first systematic model (Marr, 1969) of the cerebellum is inspired by these features of the cerebellar architecture. The theory proposes that the cerebellum is a learning system that can associate different mossy fiber input patterns (that carry different information about the world) with different output. The theory proposes that the input network is responsible for distinguishing similar mossy fiber inputs. The output network uses the granule cell activity to learn to associate specific mossy fiber inputs with specific responses.

For the input network, Marr's theory proposes that the connectivity between mossy fibers, granule cells, and Golgi cells performs pattern separation.

Specifically, The connectivity of the network enables two mossy fiber inputs that are similar to produce activity in two dissimilar groups of granule cells. This property arises from the connectivity from mossy fibers to granule cells; each granule cell has four dendrites on average, and likely receives inputs from four different mossy fibers (Ito, 1984; Palkovits et al., 1971b). Depending on the threshold of activation for a granule cell, the cell could respond to the coincident input of a combination of the mossy fiber inputs. The highest threshold for granule cell activation requires all four mossy fiber inputs to activate. Each granule cell is assumed to receive inputs from a random set of four mossy fibers. For example, with this connectivity, two mossy fiber inputs that are 90% in common could produce two groups of granule cell activity with 66% of granule cells in common. This example is assuming the highest threshold for each granule cell, which requires coincident input of all four mossy fibers for a granule cell to activate.

However, the percentage of granule cells activated given a mossy fiber input is very small when using the highest threshold. This is especially true when a mossy fiber input only activates a small percentage of all mossy fibers. For example, if a mossy fiber input only activates 5% of the mossy fibers, then the percentage of granule cells that have four mossy fiber inputs that are in the 5% of activated mossy fibers is 0.000625% of the granule cell population (six per million cells), which could be too few to support learning for the output network.

Therefore, the Golgi cells are proposed to adjust the granule cell threshold depending on the size of the mossy fiber inputs. The feedforward connectivity from mossy fibers to Golgi cells is proposed to support the ability of Golgi cells to determine the size of the mossy fiber input. With a small mossy fiber input size such as the 5% example above, the Golgi cells could decrease the granule cell threshold (by decreasing inhibition), such that fewer coincident mossy fiber inputs are necessary to activate a granule cell. This would increase the percentage of granule cells active to a mossy fiber input, which could then support learning in the output network. In contrast, if the size of the mossy fiber input is large, such as when 30% of mossy fibers are activated, then the Golgi cells can raise the granule cell threshold (by increasing inhibition) such that the granule cell population is not overly active. However, decreasing the threshold of granule cells decreases their selectivity in responding to mossy fiber inputs. For example, a threshold that requires coincident input of two mossy fibers to a granule cell will allow the cell to respond to six different combinations of two mossy fiber inputs. In comparison, a threshold that requires coincident input of all four mossy fibers will only allow the cell to be activated by a single combination of mossy fiber inputs. Therefore, there is a tradeoff between the ability to perform pattern separation and ensuring sufficient activity in the granule cell population for small mossy fiber input sizes. In summary, Marr proposes that the connectivity of the input network enables pattern separation of mossy fiber inputs in the granule cell

activity, and that Golgi cells provide a role of adjusting the granule cell threshold to enforce a relatively constant overall granule cell activity given different mossy fiber input sizes.

The Purkinje cells in the cerebellar output network are unique in receiving converging inputs from mossy fibers (by the path of granule cells) and climbing fibers. Marr proposes that the role of the climbing fibers is to convey when the output of the Purkinje cells are in error (a teaching signal). Given the connectivity, Purkinje cells are proposed to use the teaching signal to modify their activity for a given mossy fiber input by the way of modifying the granule-Purkinje synapses. As a consequence, in subsequent presentations of the same mossy fiber input, the Purkinje cell activity is modified, which could reduce or eliminate the teaching signal. The pattern separation properties in the input network support the ability of Purkinje cells to distinguish between similar mossy fiber inputs, so as to not erroneously modify output to a potentially unrelated input. The Purkinje cell activity is proposed to reflect the output of the cerebellum. Any changes to that activity should reflect the changes in the granule-Purkinje synaptic strengths in response to the teaching signal. However, it is possible for the overall granule activity to fluctuate, which could change the activity of the Purkinje cells, independent of the learned changes from the teaching signal. The basket and stellate cells are proposed to normalize the Purkinje cells' activity given fluctuating granule activity, such that any changes in the Purkinje activity only

reflect the changes in the granule-Purkinje synapses (the learned responses).

Plasticity at the granule-Purkinje synapses have been demonstrated subsequent to Marr's theory (Ito & Kano, 1982; Ito, 2001; Jörntell & Hansel, 2006; Lev-Ram, Mehta, Kleinfeld, & Tsien, 2003; Wang, Denk, & Häusser, 2000). Delay eyelid conditioning and other behaviors also strongly suggest that learning is an essential computational feature of the cerebellum (Boyden et al., 2004; Raymond, Lisberger, & Mauk, 1996). However, Marr's theory did not address the ability of the cerebellum to produce well timed responses, as suggested by delay eyelid conditioning and smooth pursuit (Li & Lisberger, 2011).

Cerebellar theories that attempt to address the temporal computation of the cerebellum focus on the role of cerebellar granule cells in producing temporal signals. The same anatomical observations which suggested that the granule cells perform pattern separation in support of learning in the Purkinje cells also suggest that granule cells are the most likely candidate to provide temporal signals to support Purkinje cells' ability to produce well-timed responses. One of the possibilities is that cerebellar granule cells could transform a tonic mossy fiber input (the CS input in eyelid conditioning) into a population activity such that different granule cells are active at different times during the input (stimulus-temporal code) (De Schutter & Bjaalie, 2001). This in turn provides the Purkinje cells the necessary timing information to generate responses at a specific time (Berthier & Moore, 1986).

There are three categories of theories that have proposed mechanisms of this possible transformation in the granule cell population. The first group of theories is synfire chains (Aviel, Mehring, Abeles, & Horn, 2003; Hosaka, Araki, & Ikeguchi, 2008; Sommer & Wennekers, 2005) and tapped delay lines (Freeman & Nicholson, 1970), both of which rely on feedforward connectivity to produce a chain of cells that are activated one after the other in response to a stimulus. This feedforward connectivity could produce cells that are active at different times. A related theory relies on the conduction velocity of granule cell axons and different conduction lengths to achieve different delayed activity at the granule-Purkinje synapses (Chapeau-Blondeau & Chauvet, 1991). Another group of theories proposes that the intrinsic physiology is different between different granule cells, such that some cells respond faster to a stimulus and others slower (Spectral timing models) (Bullock & Grossberg, 1988; Grossberg & Schmajuk, 1989; Ulloa, Bullock, & Rhodes, 2003). These differences in physiology could produce cells that are active at different times. The difference in physiology could be produced by different membrane time constants (Bullock, Fiala, & Grossberg, 1994), and/or through other means such as different synaptic strengths at mossy fiber to granule cell synapses (D'Angelo & De Zeeuw, 2009). The mossy fiber to granule cell synapses have been observed to be plastic and controlled by Golgi cell activity (Armano, Rossi, Taglietti, & D'Angelo, 2000; Mapelli & D'Angelo, 2007), which could possibly support this model for limited time scales. The final group of

theories proposes that granule cells act as oscillators and respond to a stimulus with different phased delays (M. Fujita, 1982). In summary, these theories of cerebellar timing either focus on the construction of very specific feedforward network connectivities, or physiology of the granule cells, or local synaptic interactions between mossy fiber, Golgi, and granule cells.

Using bottom-up simulations and eyelid conditioning to investigate cerebellar timing

In contrast to the models that focused on the specific physiology of certain cell types or specifically constructed connectivity, a biologically constrained simulation by Buonomano, Medina, and Mauk (Buonomano & Mauk, 1994; Medina et al., 2000) using a bottom-up approach suggests that the transformation of a tonic mossy fiber input into stimulus-temporal code can be an emergent property of the cerebellar network architecture without specifically designed elements. The simulation modeled the observed physiology and stochastic connectivity of the underlying cells of the network, and focuses on the emergent properties that arise from the interactions among these cells. The failures and successes of the simulation in reproducing eyelid conditioning behavior have been useful for generating testable predictions in experiments (Kalmbach et al., 2011; Medina et al., 2000). In this model, the interactions between the granule and Golgi cell populations provide the basis for the

emergent property of the network to keep track of time. These interactions are a consequence of recurrent connectivity between the two cell populations. The advantage of this model is its ability to partially reproduce the animal's behavior without relying on any unobserved connectivity and physiology.

However, when this model was first proposed, the available computational power constrained the number of granule cells in the simulation to 12,000. As a consequence, the connectivity of the simulation deviated by several orders of magnitude from the observed connectivity among mossy fibers and granule and Golgi cells (see (Buonomano & Mauk, 1994) for a discussion of the choice of compromise made). This leaves the possibility that the emergent computational properties of the input network (mossy fiber-granule-Golgi network) observed in the constrained simulation might not be relevant to the computation performed by these cells in the cerebellum.

In this thesis, the mechanisms of the input network for which the cerebellum can keep track of time are revisited using a simulation that expanded the number of granule cells by 100 fold from the constrained simulation. The expanded simulation contains 1 million cells, which approaches the observed connectivity in the cerebellum. Using this simulation, the mechanisms that could transform tonic mossy fiber inputs into granule stimulus-temporal code are investigated. Chapter 2 describes the implementation of the expanded simulation. The connectivity and representation of the neurons are discussed.

The algorithms that utilize graphics processing units are discussed in the context of overcoming the 100 fold increase in computation load. In chapter 3, alternate mechanisms and connectivity constraints for generating stimulus-temporal code are proposed based on the results from the expanded simulation of delayed eyelid conditioning. The interactions between granule and Golgi cells (as suggested by the constrained simulation) are found to produce stimulus-temporal code, but under connectivity constraints that are beyond what current observations would support. In searching for an alternate mechanism, a recently discovered inhibition among Golgi cells (lateral Golgi inhibition) is found to be a possibility for producing stimulus-temporal code. This connectivity allows the simulation to more closely reproduce animal behavior in delay eyelid conditioning. In chapter 4, the network mechanisms of lateral Golgi inhibition are investigated by dissecting the network in detail. The Golgi cells that are active early during the CS input (early cells) are found to be important for producing stimulus-temporal code. Further dissections found that Golgi cells that inhibit early cells and not inhibited by early cells in return (nonreciprocal inhibition) are important for producing stimulus-temporal code. The predictions from the dissections are tested in the intact simulation, demonstrating that the dissection analysis is relevant to network mechanisms in the intact simulation. Finally, in chapter 5 the limitations and significance of the results are discussed. The relevance of the methodology of using graphic processing units in high

performance computing, and further scaling of the simulation is discussed. The specific functional hypotheses of lateral Golgi inhibition in producing stimulus-temporal code in the cerebellum are discussed in the context of limitations of the model. Then, the relevance of the network computational principles discovered by the analysis to the field of recurrent neural networks is discussed. Finally, the approach of the analysis in dissecting the lateral Golgi inhibition network is discussed in the context of analyzing the mechanisms and components of complex systems in general.

CHAPTER 2:

EXPANDING THE CEREBELLAR SIMULATION TO APPROCH THE OBSERVED CONNECTIVITY

Introduction

Over 99% of neurons in the cerebellum are cerebellar granule cells (Lange, 1975). These cells receive extensive excitatory inputs from mossy fibers, which are one of the two major input pathways into the cerebellum. Recordings of the activities of these fibers suggest that they convey information about the world such as sensory stimuli (Aitkin & Boyd, 1978; K Maekawa & Takeda, 1975; Kyoji Maekawa & Takeda, 1976; Ohyama & Nores, 2003; Winfield, Hendrickson, & Kimm, 1978), and proprioception (Fuchs & Kornhuber, 1969; Murphy, MacKay, & Johnson, 1973). Anatomical observations of the cerebellum show that the mossy fibers provide extensive connections to granule cells (Eccles et al., 1967; Palkovits et al., 1971b), and a large number of granule cells converge onto Purkinje cells (Palkovits et al., 1971c) that provide the only output of the cerebellar cortex (Eccles et al., 1967; Palkovits et al., 1977). This connectivity suggests that the transformation of mossy fiber inputs by granule cells might be an important aspect of the computation performed by the cerebellum (Bullock et al., 1994; De Schutter & Bjaalie, 2001; Marr, 1969). One such transformation is suggested by the delay eyelid conditioning paradigm with mossy fiber stimulations (Kalmbach et al., 2010; McCormick & Thompson, 1984). Animals trained with this paradigm learn to close their eyelid not at the onset of the mossy fiber stimulation (conditioned stimulus, CS). Instead, the closure is timed (White et al., 2000) to anticipate the onset of the eyelid stimulation (unconditioned

stimulus, US). The timing of this learned behavior can be explained if the granule cell population responds to a tonic mossy fiber input by producing stimulus-temporal code (where different granule cells are active at different times) (Bullock et al., 1994; D'Angelo & De Zeeuw, 2009).

Buonomano and Mauk (1994) proposed that the granule cells can produce stimulus-temporal code from to the recurrent interactions between Golgi cells and granule cells, as an emergent property of the network. However, due to the computational power available at the time when the simulation was developed, the number of simulated granule cells was constrained to 12000. The resulting compromises in order to maintain a sufficient number of Golgi cells and mossy fibers to make a meaningful network for eyelid conditioning require deviating from the relevant connectivity ratios by 1-2 orders of magnitude (table 2.1). A majority of the deviation are primarily in the connectivity with granule cells.

Presyn:Postsyn	12K simulation	1M simulation	observed
Mossy fiber:granule	4:80	4:2048	4.2:400-1800
Golgi:granule	4:53.3	3:3072	3-4:5000-8000
Granule:Purkinje	500:1	32768:1	80K-200K:40

Table 2.1. Connectivity ratios of the 12000 granule cell simulation, the one million granule cell simulation, and observed connectivity. Presyn:PostSyn: convergence ratio from pre-synaptic cell and divergent ratio to post-synaptic cell.

Such large deviations from the observed connectivity leave the possibility that the emergent network properties observed in this model might not be

relevant to the actual computation performed in the cerebellum. In order to ensure that one is studying network computation that is as relevant as possible to the cerebellum, the observed network connectivity and structure must be represented faithfully in the simulated cerebellar network. Therefore, the number of granule cells must be expanded in current simulations in order to approach biologically relevant connectivity ratios across the eyelid conditioning-associated cerebellar network.

Simulating a million granule cells would allow the number of mossy fibers, Golgi cells, and Purkinje cells to be in sufficient numbers and still have biologically relevant connectivity ratios with these cells (table 1). One million granule cells represent a near 100-fold increase in the total number of cells compared to the previous simulation, and as such present significant challenges in implementation. The goal is that the simulation must run sufficiently fast to produce results overnight, in order to allow for parameter adjustments in a timely manner. The time step in the original simulation is chosen to be 1ms as a compromise between performance and the ability of the neuron models to reproduce empirically observed activity. The 12000 granule cell simulation can be executed with a speed of at least 1000 time steps a second on computers available in 2008. At this speed the simulation typically can show learned responses within an hour. With a 100-fold increase in the number of cells, the same implementation would have taken roughly 100-fold amount of time to do

the same, which would be at least 3 days—too slow to perform necessary manipulations in a timely fashion.

Considering that granule cells make up more than 99% of all the cells in the simulation, the bulk of the computation is calculating their activity and updating their inputs and outputs. The computation belongs to the easily parallel class of computations, where the calculation of a granule cell's activity and input/output during a time step does not depend on the calculation of other granule cells. Thus, all the granule cells can be processed at the same time. In addition, the instructions for calculating the activity of each granule cell is identical to every other granule cell, thus conforming to the single instruction multiple data (SIMD) computation pattern (i.e., applying the same instruction to a large array of data). Updating the input/output of the simulated granule cells also conforms to SIMD patterns, but the memory access pattern is non-sequential and more random. While the computation pattern is simple, the amount of data processed to calculate granule cell activity is immense. Each granule cell requires 128 bytes of data to describe, all of which need to be read and/or written to per time-step. Given one million granule cells the total amount of data is 128MB. Thus, to process 1000 time steps per second, (for real time performance) the amount of memory that needs to be transferred is 128GB per second.

The practical limit of the size of a silicon chip that can be produced and the

minimum size of a transistor on the chip define the transistor budget of a processor (the maximum number of transistors possible). This budget directly limits the computational power of a processor, and the CPU (central processing unit) and GPU (graphics processing unit) designs represent two optimization points that use the transistor budget differently to handle different computations. The CPUs are optimized for processing a few, complex and non-parallel instruction streams quickly, while the GPUs are optimized for processing same instructions on large amount of data, which is the type of computation that characterizes the simulation.

The design of the CPU is to use the transistor budget on a few large complex cores, with each core capable of processing a single stream of instructions quickly. This design is befitting to the programming model of the majority of existing software, where the code is designed to be executed by a single core in sequence. This type of code can be very complex because which instruction to execute next is determined by the result of the previous instruction, and therefore has no potential for execution at the same time. In practice, this type of code usually contains sections of instructions that can be executed at the same time, and CPU designs allocate the transistor budget to detect such sections and to execute them as quickly as possible. In addition, the transistor budget is used on predictively loading the instructions that are yet to be executed, so that the actual processing unit does not have to wait for the

instructions to be loaded. The ideal memory system that provide data and instructions to the CPU has low latency (i.e., minimum time between issuing a command to load instruction/data from memory and receiving said instruction/data) so to minimize the wait time while handling unpredictable access to different parts of memory. However, low latency memory requires the memory to be on the same silicon chip as the processor cores, which consumes the transistor budget and limits the amount of this memory type to typically less than 32MB, which is far less than the 128MB of memory needed for the expanded simulation. On the other hand, memory that is not on the processor chip has a high latency but can be very large (greater than 32GB). Thus, design of the memory subsystem of a CPU is a compromise between these two types of memory. The design is organized in a hierarchy, with the small but low latency memory acting as a cache for the large high latency memory. The cache stores temporary copies of data and instructions that are frequently accessed. Memory bandwidth (the amount of data that can be transferred per unit of time) is typically a secondary concern, since the computations for CPUs rarely need to access large amount of data at the same time.

In contrast, graphics processing units (GPUs) are optimized for a different kind of computation: performing identical instructions on large amount of data at the same time. The transistor budget spent on a GPU core is very small, which allows GPUs to have hundreds of processing cores. The cores are grouped

together into clusters, and each cluster contains an instruction fetcher and cache memory. Each core cannot fetch an instruction by itself, but relies on the instruction fetcher in the cluster to feed the instruction to it. Thus, a cluster of cores must perform identical instructions. To keep hundreds of cores occupied, the memory must have high bandwidth. The bandwidth of the memory is achieved by accessing large amount of data per transfer, and thus can achieve between 180-300GB/s of bandwidth, as compared to 30-50GB/s for the CPU. The performance of a GPU relies on that the computation at each core does not depend on the result from other cores, so that all the cores can process instructions at the same time. In this case, the GPU can be 1-2 orders of magnitude faster than the CPU. If the computation only engages a few cores, such as the typical computation for a CPU, the GPU would be considerably slower. For the expanded simulation, the bulk of the computation is applying identical instructions to large sets of data, which is suitable for engaging all available cores on the GPU, and thus can extract maximum performance.

The first generation of GPUs capable of general purpose computing (GPGPU) was developed in 2006 by Nvidia (Nvidia corp., www.nvidia.com) and AMD (Advanced Micro Devices, www.amd.com). With that generation, Nvidia started focusing on GPGPU in an effort to exploit the potential market in super computers. Thus, Nvidia provided the CUDA framework (Compute Unified Device Architecture) which is a combination of Nvidia's extensions to the C language, an

application programming interface, and compilers to write general purpose programs for the GPU. Prior to 2006, the GPUs could only be utilized by embedding the computation as graphics instructions, which did not cover all necessary instructions for general purpose computing, and was very complex to implement. The CUDA framework contained methods for memory management on the GPUs, data transfer between GPUs and CPU, and code execution on the GPU. These features allow the programmer to write general purpose programs without needing to work with graphics instructions. The alternative software framework is OpenCL, which at the time did not have as mature of a design or documentation. In addition, AMD did not focus on GPGPU until 2012, so while their GPUs can perform general purpose computing, these GPUs are more complex to program. Therefore, Nvidia and CUDA were chosen as the platform for developing the expanded one million cell simulation.

As described above, the cores on the Nvidia GPUs are organized into clusters(Nvidia, 2014). Within the cluster, all cores must execute the same instructions and also share the L1 cache (cache memory that is on the chip and very close to the cores). The L1 cache allows the different cores in a cluster to communicate with each other with very short latency. Different clusters of cores can be considered independent of each other. Consequently, the CUDA software architecture is designed to reflect this organization. A CUDA thread is a sequence of instructions to be executed on a core, and threads are executed in blocks on

clusters of cores. To be efficient, each thread block should contain at least 32 threads, which also reflects the minimum number of cores in each cluster. Each block can contain a maximum of 1024 threads, as defined by the CUDA engine. The programmer must decide how many threads should be in a block, which depends on the complexity of the task and the amount of registers and temporary variables required for each thread. The number of blocks scheduled by a computation task is the number of threads per block divided by the number of elements that needs to be processed for that task. The latency penalty of accessing main memory that is off the GPU chip can be compensated if multiple transfers occur for a large sequential block of data. In addition, the programmer can reserve a section of the L1 cache for specific data to allow threads within a block to perform low latency random memory access. The size of this memory is limited to 48KB for each cluster of cores (on current hardware), but is optimized for access in irregular and unpredictable patterns.

The algorithms in the expanded simulation utilized this organization in the GPU to achieve 2x real time execution with 4 GPUs. The minimum average memory bandwidth required for the expanded simulation to execute at real time (1000 time steps per second) is 128GB/s. In practice the peak memory bandwidth requirement is higher when computation and memory transfer cannot occur at the same time. Assuming that computation and memory transfer each occupies 50% of the total time, the peak memory bandwidth necessary is

256GB/s. The GPUs currently used (Nvidia GTX690) are capable of 192GB/s of bandwidth per GPU, with an aggregated bandwidth of 768 GB/s across 4 GPUs. Using the 50% assumption above, the peak memory bandwidth required is 512GB/s for 2x real time execution, which is well within the limit of the aggregated GPU bandwidth. In contrast, modern workstations only have 40 GB/s of memory bandwidth (Intel Ivy Bridge-E). The raw computation power available to the GPUs is also much higher: each GPU used currently has 1536 cores, in contrast to 8-16 cores available on the CPU. However, as discussed above, fully utilizing the cores on a GPU requires SIMD type of computation, and incurs additional software complexity. The algorithms described below are designed to utilize the GPU cores efficiently to achieve the 2x real time performance.

Simulation connectivity

The algorithms used to connect the cells in the simulation together is intended to mimic the observed numerical and spatial relationships of the connectivity in the cerebellum. The cerebellar cortical network can be considered as two networks: the cortical input network, which contains mossy fibers, granule cells, and Golgi cells, and the output network, which receives input from the granule cells (but does not provide input to granule cells) and is composed of Purkinje, basket, and stellate cells. The Purkinje cell axons form the output of the cerebellar cortex, and connect to cerebellar deep nucleus cells that form the

output of the cerebellum. The Purkinje cells also receive inputs from the inferior olive in the form of climbing fibers.

Mossy fiber-Golgi-granule network

In order to produce a biologically relevant spatial and numerical representation of the mossy fiber-Golgi-granule network, the granule cells are represented as a grid of 512 by 2048 cells (1048576 total), which in physical dimensions represent about 1mm^2 of cerebellar cortex in the cat (estimated 1.2 million granule cells per mm^2 of the cat cerebellar cortex, (Lange, 1975; Palkovits et al., 1971b)). The Golgi cells are placed in a grid of 16 by 64 cells (1024 total), and stretched to evenly disperse in the granule cell grid. The glomeruli are important for defining the spatial relationship of connectivity between the cells. These are the synaptic structures formed by mossy fiber terminals, Golgi cell axons, Golgi cell dendrites, and granule cell dendrites, and they are placed in a 128 by 512 grid (65536 total), also evenly interspersed in the granule cell grid.

The connectivity between mossy fibers, granule cells, and Golgi cells is largely determined by their connectivity to the glomeruli, with the exception of granule cell inputs to Golgi cells. To define the connectivity between these cells, their connectivity to the glomeruli is first defined, based on which the connectivity to the other cells can be translated.

Mossy fibers have been observed to connect to 16-40 glomeruli on

average (Eccles et al., 1967; Palkovits et al., 1972), with each glomerulus containing a single mossy fiber terminal (Jakab & Hámori, 1988; Spacek et al., 1973). In the simulation, there are 2048 mossy fibers and 65536 glomeruli, and every mossy fiber is randomly assigned 32 unique glomeruli under the condition that each glomerulus can only have a single mossy fiber assigned to it.

Each dendrite of a granule cell have been observed to connect to a single glomerulus (Eccles et al., 1967; Palkovits et al., 1972). On average, a granule cell has 4 dendrites. Therefore, 4 distinct glomeruli are assigned to each granule cell in the simulation. The average reported lengths of granule dendrites are around 14um in cats (Palkovits et al., 1972), and up to 120um in turtles (Mugnaini, Atluri, & Houk, 1974). However these lengths cannot be used directly to define the span in the simulation, since the granule cell population is arranged as a purely 2 dimensional grid in the simulation and the granule cell layer in the cerebellum is arranged in 3 dimensions. The lengths of these dendrites impose a constraint of locality on the connectivity between glomeruli and granule cells, however. Therefore, the 4 dendrites for each granule cell are constrained to connect to glomeruli that are at most 64 granule cells away from the granule cell of that dendrite, which is a span of 4 by 4 glomeruli. Each dendrite randomly selects 1 of these 16 possible glomeruli to connect to.

The axon of a Golgi cell has been shown to connect to roughly 40 glomeruli within the vicinity of the cell body (Palkovits et al., 1971b), with each

glomerulus containing a single axon connection. It is unclear, however, whether Golgi axon arborization fields overlap with each other. In the simulation, each Golgi cell axon can connect to up to 48 glomeruli randomly chosen from 144 possibilities formed by a span of 12 by 12 glomeruli centered on the Golgi cell body. This produces partial overlap in the Golgi axon arborizations.

The Golgi cell descending dendrites have been observed to connect to the glomeruli (Hámori & Szentágothai, 1966). In the simulation, the Golgi descending dendrites also connect to the glomeruli with the same algorithm as the axons. However, instead of 48 glomeruli per Golgi axon, the dendrites only connect to 16 glomeruli within the same glomeruli span.

After the Glomeruli are assigned for the mossy fibers, granule cells, and Golgi cells, the identity of these assignments can then be translated to mossy fiber to granule and Golgi cell connections, and Golgi to granule cell connections. For example, the mossy fiber input to each granule cell is calculated by first identifying the glomeruli that granule cell dendrite connects to, and then finding the mossy fibers associated with these glomeruli.

The granule cell output to Golgi cells do not involve the glomeruli. Instead, their synapses involve the granule cell axons (parallel fibers) contacting the Golgi cell ascending dendrites (Palay, 1974). The parallel fibers from different granule cells all traverse parallel to each other. Therefore, relative to the position of a Golgi cell, it is possible for a granule cell that is located along the same direction

as the parallel fiber to be much further than a granule cell that is located perpendicular to provide input to the Golgi cell. This shapes the span of granule cells that provide output to a Golgi cell to be a long rectangle, with the width of the rectangle as the Golgi ascending dendrite span and the length of the rectangle as the average length of the parallel fibers. However, the Golgi ascending dendrite span has not been characterized, and in the simulation is chosen to be 40 granule cells, which represented a partial overlap with neighboring Golgi cells. The average length of parallel fibers has been characterized to be longer than 2mm in the cat (Brand et al., 1976), and in the simulation is chosen span across the entire length of the granule cell grid, at 2048. Within this span, 4096 granule cells are randomly chosen out of 81920 possible granule cells to connect to a Golgi cell. The number is chosen so that the Golgi cell population samples the granule cell population 4 times, which is consistent with the original constrained simulation.

Purkinje-basket-stellate network

The Purkinje, basket, and stellate cells all receive inputs from granule cells and are arranged to mimic the arrangement in a functional unit of the cerebellum, which is in a stripe perpendicular to the parallel fibers such that each cell population received input from the entire granule population. There are 32 Purkinje cells, 128 basket cells, and 512 stellate cells in the simulation, and

received 32768, 8192, and 2048 granule inputs respectively. The granule cells are evenly divide among the 32 Purkinje cells, such that Purkinje cell 1 receives inputs from granule cells 1-32768, and Purkinje cell 2 receives inputs from granule cells 32769-65536, and so on. The plasticity at granule-Purkinje synapses have been suggested as one of the sites of plasticity that supports learns in the cerebellum. The granule cell inputs to basket and stellate cells are arranged similarly.

Both basket and stellate cells have been observed to inhibit Purkinje cells, where stellate cells inhibit parts of the Purkinje cell dendrite, and basket cells inhibit the Purkinje cell body. In the simulation, each stellate cell inhibits a single Purkinje cell, and a Purkinje cell receives input from 16 stellate cells. The basket cell axons have been observed to inhibit 8-16 Purkinje cells along the same stripe. In the simulation, each basket cell inhibits 4 Purkinje cells (due to a lack of Purkinje cells). The Purkinje cells have also been observed to inhibit basket cells, and in the simulation each Purkinje cell inhibits 4 basket cells.

Purkinje-deep nucleus-inferior olive network

The Purkinje cell axons are the only output of the cerebellar cortex and inhibit deep cerebellar nucleus neurons. The deep cerebellar nucleus neurons also receive excitatory inputs from mossy fibers (Shinoda et al., 1992), and there is evidence to suggest that the mossy fiber to deep nucleus synapses are plastic

and controlled by Purkinje cell activity (Ohyama et al., 2006; Pugh & Raman, 2006; Zheng & Raman, 2010). Similar to granule-Purkinje synapses, this plasticity is also suggested to support learning in the cerebellum. The deep nucleus neurons have been observed provide inhibition to inferior olivary neurons (Best & Regehr, 2009; Lang et al., 1996). The axons of inferior olivary neurons in turn provide climbing fiber inputs to the Purkinje cells, and it is thought that Purkinje cells within a functional unit of the cerebellum receives inputs from a similar group of inferior olivary neurons (Ito, 2000, 2006b). There are 8 deep nucleus neurons and 4 inferior olivary neurons in the simulation. The Purkinje cell to deep nucleus connectivity have been characterized in the cat (Palkovits et al., 1977), however the number of Purkinje cells in the simulation is small and could not satisfy the observed connectivity. In the simulation, each Purkinje cell provide output to 3 deep nucleus cells, and each deep nucleus cell receives 12 Purkinje inputs. The connectivity of the mossy fibers to deep nucleus cells have been observed experimentally (Shinoda et al., 1992), and in the simulation each deep nucleus cell receives inputs from 256 mossy fibers. GABAergic neurons in the cerebellar deep nucleus have been observed to project to the inferior olivary neurons (Lang et al., 1996). In the simulation this is implemented such that each inferior olivary cell receives inputs from all deep nucleus cells. The inferior olivary neurons provide climbing fiber inputs to Purkinje cells. Each Purkinje cell have been observed to receive a single climbing fiber input (Eccles et al., 1967), and

each inferior olivary cell has been observed to provide climbing fibers to a few Purkinje cells (H. Fujita & Sugihara, 2013; Ito, 1984). In the simulation, each of the 4 inferior olivary cells provides climbing fibers to 8 Purkinje cells.

Neuron representation in the simulation

The current work is intended to explore the emergent network computational properties of the cerebellum. It is not intended to address the specific computation of inputs by individual neurons (Mauk, 2000). As such, the models of individual neurons in the simulation aim to provide phenomenological representations of empirically observed neuronal activity *in vivo*. Thus, each neuron is represented with a simplified iso-potential conductance model (Medina et al., 2000). The membrane potential V_m of this model is represented by the equation

$$\frac{dV_m}{dt} = \frac{\sum_i g_i(E_i - V_m)}{C_m}$$

where g_i is a conductance of the membrane, E_i is the reversal potential for g_i and C_m is the capacitance of the membrane. This equation is solved by simulation using discrete time steps as follows:

$$V_m(t) = V_m(t-1) + \Delta V_m(t)$$

and

$$\Delta V_m(t) = \sum_i g_i(t)(E_i - V_m(t-1))$$

where $V_m(t)$ is the membrane potential at current time step t, $t-1$ is the previous time step, $\Delta V_m(t)$ is the change in membrane potential at t, $g_i(t)$ is the magnitude of the conductance i at t. C_m is not explicitly modeled, but instead is implicitly modeled in the magnitude of conductance g_i. The conductance g_i at t is calculated as follows:

$$g_i(t) = g_i(t-1) * decay_i + s_i * input_i$$

where $decay_i$ is the decay constant for g_i, s_i is the scaling constant for g_i, and $input_i$ is the total number of action potentials (spikes) that comprise the input to that neuron at t. A free parameter s_i is included and is tuned for each conductance. The constant $decay_i$ is calculated as follows:

$$decay_i = e^{(-\Delta t/\tau_i)}$$

where Δt is the length of a time step in milliseconds, and τ_i is the decay time constant for conductance g_i (in milliseconds) which is taken from published experimental data.

The length of the time step has to be microseconds in order to model the active conductances of a spike (Lee, Neiman, & Kim, 1998). The decrease in time step size requires a proportional increase in the number of time steps for each second of simulated activity. This is computationally expensive and does not necessarily improve the fidelity of these cellular models for the purposes of

the current simulation. Thus, the representation of a spike is simplified to a number whose value is 1 at the time step when it occurs, and 0 otherwise. A spike occurs when V_m exceeds a threshold. To model the absolute and relative refractory period of spikes, the threshold TH at time step t is computed as the following:

$$TH(t)=\begin{cases} max_{TH} & \text{if} \quad AP=1 \\ TH(t-1)-(TH(t-1)-base_{TH})*decay_{TH} & \text{if} \quad AP=0 \end{cases}$$

where AP is 1 when there is a spike at t and 0 otherwise. The constants max_{TH} and $base_{TH}$ are the maximum and minimum values that threshold can reach, respectively. $decay_{TH}$ is the decay constant that determines the rate that threshold decreases toward $base_{TH}$. $decay_{TH}$ is calculated as

$$decay_{TH}=1-e^{(-\Delta t/\tau_{TH})}$$

where τ_{TH} is the decay time constant for the threshold. The threshold reaches max_{TH} immediately after a spike, and then decays toward $base_{TH}$. max_{TH} is set at or above the reversal potential of excitatory conductances and the threshold must decay below that reversal potential in order for another spike to occur (defining the absolute refractory period). Once the threshold is below that reversal potential, the amount of excitatory conductance necessary for another spike to occur decreases as the amount of time since the last spike increases (relative refractory period).

The length of the time step Δt in the simulation is one millisecond, which

provides sufficient resolution for modeling the dynamics of synaptic conductances. This is appropriate to model the spiking patterns of cerebellar neurons whose spiking frequency rarely exceeds 500 Hz.

The general procedure for calculating the activity of a neuron in the simulation for each time step is as follows: 1. calculate the number of spike inputs for each conductance, 2. update the amplitude of each conductance, 3. update the membrane potential V_m, 4. compare V_m to threshold to determine if a spike is generated, 5. update the threshold. The sources of input for each conductance of each cell is specified in a connectivity array, which contains indices that specify which cells provide input to that cell. To update the total number of spike inputs for each conductance, the corresponding activity of the input cells specified in the connectivity array are summed. The rest of the steps are straightforward implementations of the equations described above.

The focus for the next section is on the division labor of implementing the above algorithm between the GPU, which handles granule cell calculations and the CPU, which handles basket, stellate, Purkinje, deep nucleus, and inferior olive cells.

Granule cell calculations in the GPU

The activity of a million granule cells has to be calculated at each time step in the GPU, which is accomplished by scheduling 2048 blocks of 512

threads, where each thread executes the instructions for calculating the activity of a granule cell. The activity of a granule cell i is calculated by a thread with index i, where index i is calculated by

$$i = ID_{block} * \text{size}_{block} + ID_{thread}$$

where ID_{block} is the thread block index that contains the thread, size_{block} is the number of threads per block, and ID_{thread} is the index of the thread in that block. There are 5 variables that describe the state of a granule cell, each in 32 bit floating point format. There are 3 variables describing the amount of excitatory, inhibitory, and auto-inhibitory conductances, a variable for the membrane potential, and a variable for the threshold for determining if a given membrane potential should elicit a spike for the cell. For 1 million granule cells, there are 5 million variables. In the GPU, the variables are arranged into 5 arrays, so that the variables describing the amount of excitatory conductances for all cells are in a single array, the variables describing the amount of inhibitory conductances are in another array, and so on. The thread with index i accesses the element with index i in each array, with i calculated as above. This ensures that data access is sequential for each block of threads. The CUDA engine is free to schedule different blocks in the order that can maximize the utilization of all the cores. An important aspect of the algorithm is that the instructions for each thread are identical, and are not dependent on the results from other threads. Thus, this computational pattern satisfies the simplest SIMD pattern and can fully utilize the

GPU. The more difficult problem is to update the inputs and outputs of granule cells.

Updating inputs to granule cells in the GPU

Each granule cell can receive inputs from as many as 4 mossy fibers and Golgi cells. The connectivity is specified by a matrix with a million rows and 4 columns. Each row contains the mossy fiber/Golgi cell indices that provide input to a given granule cell. The connectivity is random but constrained by spatial rules. The random connectivity produces random memory access patterns as the input to granule cells are updated.

Memory latency cannot be negated when memory access pattern is random. The main memory on the GPU is unsuitable for these kinds of access patterns due its high access latency, which would result in the cores waiting for data. However, the latency for the L1 cache on the GPU is orders of magnitude lower (the specific latency varies from model to model), and so is utilized for this task.

When the mossy fiber inputs to granule cells are updated at each time step, the mossy fiber activity array is first transferred to the GPU from the CPU. Each element in the array specifies whether a mossy fiber has produced a spike at that time step. For updating granule cells, 2048 blocks of 512 threads are scheduled on the GPU. For each block of threads, the mossy fiber activity array

is first loaded into L1 cache. Then, each thread i responsible for updating the input for granule cell i reads from row i in the connectivity matrix. Each element in the connectivity row specifies a mossy fiber index k, and the thread access the k^{th} element in the mossy fiber activity array (now in L1 cache), and updates the excitatory input to the granule cell according to the general algorithm described above. Again, the instructions for each thread are identical and the result is independent from other threads. The same algorithm applies to updating the Golgi inhibitory inputs.

Updating granule output to other cells in the GPU

Each granule cell can also provide output to as many as 4 Golgi cells. Each row of the matrix that specifies this connectivity contains the Golgi cell output targets for a given granule cell, and so has 1 million rows and 4 columns. The algorithm for updating granule cell output is similar to updating granule cell input. 2048 blocks of 512 threads are scheduled for this task. Each block of threads writes to an array of integers that represent the input to Golgi cells. The length of the array is the same as the total number of Golgi cells, and is first loaded into L1 cache for each block of threads. For each granule cell that produced a spike, the corresponding thread in the block will increment the elements in the Golgi input array by 1 as specified by the connectivity matrix. Because multiple threads may need to write to the same memory location, the

writing operation has to be atomic (i.e., the thread writing the data has exclusive access to that memory location). After the threads have updated the Golgi input array, the array is saved back to the main memory. Thus, each block of threads produces an array of inputs to the population of Golgi cells. The final task is to sum the 2048 arrays into a single array by summing the k^{th} element from each array to produce the total input to the k^{th} Golgi cell.

In addition to Golgi cells, granule cells also provide input to basket cells, stellate cells, and Purkinje cells. Anatomical observations suggest that a row of functionally related Purkinje cells likely receive distinct granule cell inputs. In the expanded simulation this is also assumed to be the case for basket and stellate cells, however the anatomical observations regarding these cells are sparse.

There are 128 basket cells in the expanded simulation, each receiving 8192 granule cell inputs. The arrangement of the basket cells is such that the first basket cell receives input from granule cells 1 through 8192, the second basket cell receives input from granule cell 8192 to 16384, and so on. For each basket cell, the total number of spikes from the 8192 granule cells that provide input determines the change in its excitatory conductance. This requires that the granule activity array be summed in blocks of 8192 elements. Nvidia has a recommended algorithm for this problem, which utilizes the GPU cores to perform summation of different parts of the array and then sum the resulting partial sums until the final sum is achieved. A block of 1024 threads is used to

calculate the sum of each array of 8192 elements. Each thread j in the block is responsible for first calculating the sum of 8 elements of the array, where the index of each element is $k*1024+j$, with k between 0 and 7. The resulting 1024 sums are stored in L1 cache as a temporary array. Then, half the threads in the block (512 threads) are responsible for reducing the 1024 sums to 512 elements, with each thread summing 2 elements. 256 threads from the block then further reduces the 512 sums to 256 elements. This process repeats until there is only a single sum of the entire array of 8192 elements.

The granule cell input to stellate cells is computed using the same algorithm, except that there are 512 stellate cells, each receiving 2048 granule cell inputs. The algorithm for computing granule inputs to Purkinje cells is also similar, with 32 Purkinje cells each receiving 32768 granule cell inputs. A key difference is that the granule-Purkinje synapses have synaptic weights that are individually adjustable, according to a set of plasticity rules described below. Thus, for calculating the granule inputs to Purkinje cells, the granule cell activity array is first multiplied by their weights before summation.

Adjusting the weight of granule-Purkinje synapses

The weights of granule to Purkinje synapses are adjusted by rules based on the timing of inferior olive spikes (which carries information about the unconditioned stimulus) and the activity of the granule cells. There are 4 inferior

olive neurons, each providing output to 8 Purkinje cells. The synaptic strength of a granule to Purkinje synapse is reduced when the granule cell is active in the time period 100-200ms prior to the onset of an inferior olive spike. Granule cell activity at other times result in increased synaptic strength. As implemented here, the synaptic strength is a scaling factor that saturates between 0 and 1, with the initial strength set to 0.5 to allow for both decreases and increases to occur. The actual implementation of this rule is controlled by a timer for each Purkinje cell. This timer is reset to -200 whenever a given Purkinje cell receives an inferior olive spike, and increments by 1 per time step. At every 5 time steps, the weights of the granule-Purkinje synapses are adjusted. Each granule cell keeps a record of its activity for the previous 384 ms. The weight of granule-Purkinje synapse i for a Purkinje cell j is adjusted by the following procedure: 1. if the timer for Purkinje cell j is greater than 0, then the plasticity step is set to a positive constant (for potentiation), and if the timer is less than -100, then the plasticity step is set to a negative constant (for depression); 2. if granule cell i was active exactly 200ms ago, then the granule-Purkinje synapse i is adjusted by the plasticity step and if granule cell i was not active the synapse is not adjusted; 3. if the weight of the granule-Purkinje synapse i is less than 0 or greater than 1, then the weight is set to 0 or 1, respectively.

CPU calculations and communication to the GPU

The algorithms described utilize the GPU to compute the inputs to Granule cells and the output from granule cells to other cell types. One of the reasons to allocate the GPUs for these tasks is to minimize the data transfer between CPU and GPU, because the data bus between CPU and GPU has a maximum limit of 16GB/s of bandwidth. However, the GPU is only responsible for calculating granule cell activity, while the CPU is responsible for calculating the activity of 2048 mossy fibers, 1024 Golgi cells, 512 Stellate cells, 128 basket cells, 32 Purkinje cells, 8 deep cerebellar nuclei cells, and 4 inferior olivary cells in the expanded simulation. This is because the number of cells is relatively small for these cell types and cannot efficiently use the GPU. This division of computation between the GPU and CPU requires only a small amount of data to be transferred between the two processors.

For each time step, the activity of the mossy fibers (8192 bytes) and Golgi cells (4096 bytes) is sent to the GPU so that it can update the inputs to granule cells. The CPU in turn requests the granule cell input array to the Golgi cells (4096 bytes), to stellate cells (2048 bytes), to basket cells (512 bytes), and to Purkinje cells (128 bytes). All of these data combined requires at least 16.7MB/s of bandwidth between the CPU and the GPU to execute in real time, which is far less than the 16GB/s of bandwidth available.

The CPU is in control of which computation should be done by the GPU

and when it is performed. This gives the CPU the ability to coordinate the GPU calculations and CPU calculations to execute at the same time, reducing the amount of time either processor has to wait for the other to finish.

Summary

Buonomano, Medina, and Mauk proposed that the transformation of tonic mossy fiber inputs into a granule stimulus-temporal population code can be the result of the emergent network computation from the interactions between granule and Golgi cells. The simulation that is used to generate this hypothesis is constructed with a bottom-up approach, where the known physiology and connectivity of different types of neurons are modeled to study the emergent network properties. The computational power available at the time restricted the number of granule cells that can be represented in the simulation, such that the relevant connectivity deviated from the relevant observations by 1-2 orders of magnitude. Such deviations could mean that the emergent computational properties of the simulation is different from the computation performed in the cerebellum.

The expanded simulation described above implements one million granule cells, which allows the relevant connectivity to approach the observed connectivity. The technical performance challenges of this expanded simulation were overcome by utilizing GPUs instead of traditional CPUs. Compared to

CPUs, the GPUs are optimized for performing the same computation on a large amount of data at the same time, and are ideal for implementing the expanded simulation. The implementation described above overcame 2 additional challenges: 1. updating the input and output of granule cells, which requires random memory access pattern, and is solved by using the L1 cache memory on the GPU, and 2. overcoming the bottleneck presented by the data bus between the GPU and CPU by assigning the GPU to perform most of the calculation and only sending and receiving processed data over the data bus. The expanded simulation using this implementation is capable of executing at 2x real time, and allowed for performing manipulations in a realistic timeframe.

CHAPTER 3:

MECHANISMS OF TIMING IN THE CEREBELLUM

Introduction

The delay eyelid conditioning paradigm has been shown to directly engage the cerebellum (Garcia et al., 1999; Mauk et al., 1986; Mauk & Thompson, 1987; McCormick & Thompson, 1984; Steinmetz et al., 1989). This paradigm is used to study cerebellar computational properties (Ohyama et al., 2003). Stimulation of mossy fiber inputs to the cerebellum (Steinmetz et al., 1989) can be used as a conditioned stimulus (CS) in place of stimuli such as a tone or light (Ohyama & Nores, 2003). Air puffs to the eye, electrical stimulation of the eyelid, or stimulation of the inferior olive (Mauk et al., 1986) can be used as the unconditioned stimulus (US). For each trial of this paradigm, the CS is presented to an animal for a fixed duration, and the US is presented at a fixed delay after the CS onset. This delay is called the inter-stimulus interval, abbreviated ISI. With repeated presentations of the CS-US pairing the animal learns to close its eyelid in anticipation of the onset of the US. Importantly, after learning, the eyelid does not close at the CS onset. Instead, the response is delayed until shortly prior to the US onset (Kalmbach et al., 2010). The onset of eyelid responses is different for different ISIs (White et al., 2000). For example, if an animal is trained to 750ms ISI, the onset of the eyelid response is more delayed compared to an animal trained to 500ms ISI. Similarly, the peak of the eyelid response occurs at the US onset (Chettih, McDougle, Ruffolo, & Medina, 2011). The cerebellar cortex is hypothesized to be the responsible region for

generating this timed behavior. The supporting evidence for this hypothesis is the observation that lesions of the eyelid region of the cerebellar cortex induce animals to respond with a short fixed delay to the CS onset, as opposed delaying the response until prior to the US onset (Kalmbach et al., 2010).

These observations suggest that the cerebellar cortex can keep track of the elapsed time since the CS onset (Mauk, Medina, Nores, & Ohyama, 2000). Furthermore, a relatively small number of mossy fibers provide inputs to a large number of cerebellar granule cells (Eccles et al., 1967). These granule cells in turn provide a large number of inputs to a small number of Purkinje cells (Palkovits et al., 1971b, 1971c). Since plasticity at granule-Purkinje synapses is likely the mechanism for producing learned responses (Ito, 2005), cerebellar granule cells are a prime candidate for providing timing information to Purkinje cells.

This timing information can be in the form of stimulus-temporal code, where different granule cells respond to the tonic CS input (figure 3.1A, top) at different times (Bullock et al., 1994; De Schutter & Bjaalie, 2001) (figure 3.1A, middle, stimulus-temporal code). The stimulus-temporal code of a cell population can be measured by calculating the correlations between the population activity at one time point to another time point (Goudar & Buonomano, 2014). A matrix of correlation scores is produced when the population activity at all time points is correlated to all other time points (figure 3.1B). For a population with tonic

activity, the population activity at each time point is highly correlated with the other time points, which produces a matrix with high correlation values at all points (figure 3.1B top). In contrast, a population that produces stimulus-temporal code produces a matrix with low correlation values, especially for time points that are not close to each other (figure 3.1B, bottom, white triangle).

Granule cell stimulus-temporal code enables Purkinje cells to decrease their activity specifically prior to the US onset (figure 3.1A, bottom). This is achieved by decreasing the weight of granule-Purkinje synapses for granule cells that were active immediately prior to the US onset (Ito & Kano, 1982; Wang et al., 2000). The decrease in Purkinje cell activity then disinhibits deep cerebellar nucleus cells. This disinhibition allows nucleus cells to increase activity specifically prior to the US onset. The nucleus cells in turn are necessary to drive the eyelid response (Lavond et al., 1985).

How the cerebellar cortex can transform tonic mossy fiber input into stimulus-temporal code has been a subject of many theories of the cerebellum. Several of these theories propose specific elements such as different time constants for different granule cells (Bullock et al., 1994) and axon conduction delays (Chapeau-Blondeau & Chauvet, 1991) as the mechanisms responsible for this transformation. In contrast, a simulation of the cerebellum by Mauk, Buonomano, and Medina suggests that the emergent network properties can be the mechanism for transforming tonic mossy fiber inputs to stimulus-temporal

code (Buonomano & Mauk, 1994; Medina et al., 2000; Ohyama, Medina, Nores, & Mauk, 2002). The emergent network properties are the result of recurrent interactions between Golgi and granule cells. However, in the simulation, the connectivity ratios among granule cells, Golgi cells, and mossy fibers are two orders of magnitude lower than anatomical observations (Buonomano & Mauk, 1994). The possibility remains that the emergent network properties discovered in this constrained simulation may not be relevant to the computations performed in the cerebellum. The expanded one million granule cell simulation is able to approach the observed connectivity. Therefore, it was used to re-examine the emergent network hypothesis from the constrained simulation as well as to explore other possible mechanisms that allow the cerebellar network to produce stimulus-temporal code.

The results from the expanded simulation suggest constraints on the connectivity for recurrent interaction between granule and Golgi cells. Specifically, for this recurrent interaction to produce stimulus-temporal code, the number of granule cell inputs per Golgi cell must be small. This limits the connectivity to be similar to that of the constrained simulation; this connectivity is unlikely to be within the ranges in the cerebellum (see discussions below). The results from the expanded simulation suggest another mechanism which could produce stimulus-temporal code. The proposed mechanism utilizes the recently discovered phenomenon of lateral Golgi inhibition (Hull & Regehr, 2012). Lateral

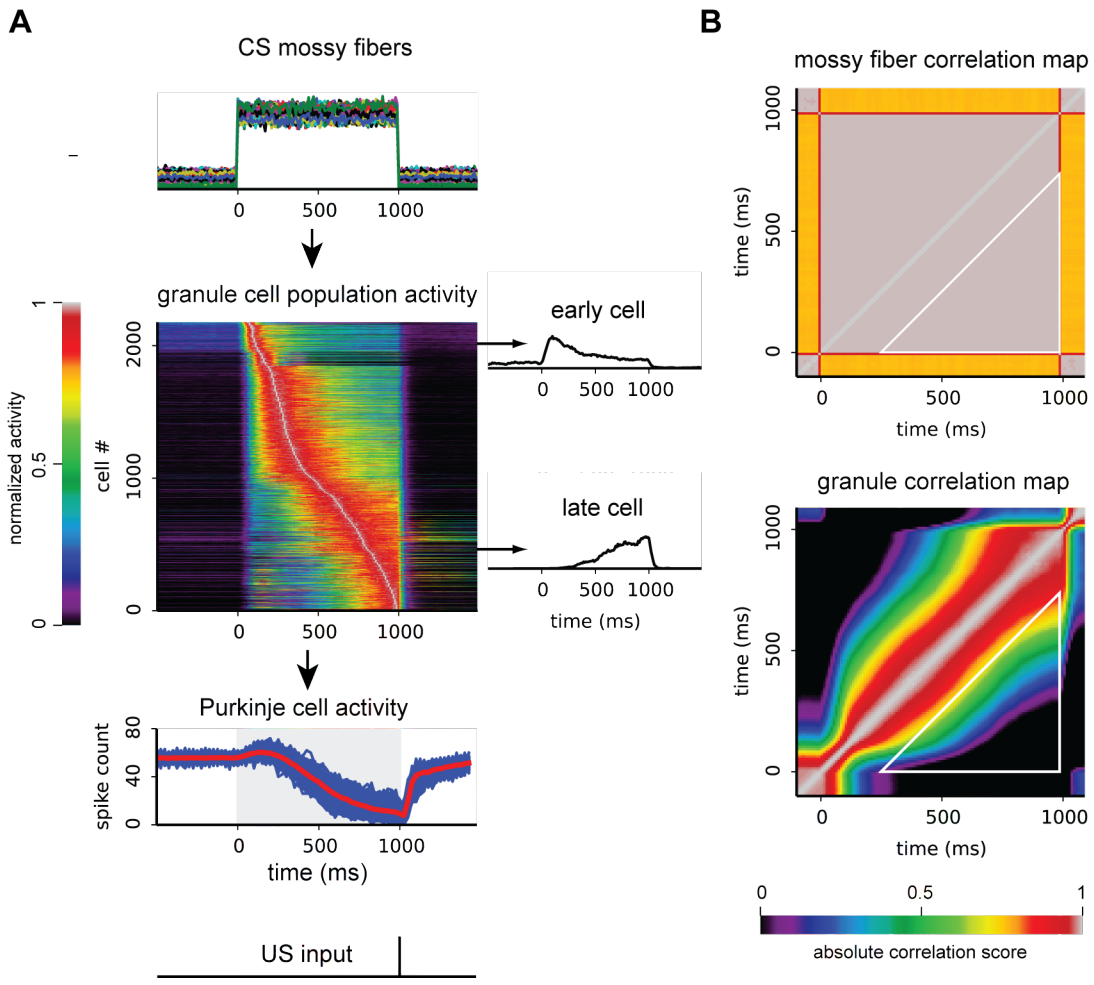

Figure 3.1. A mechanism with which the cerebellum can generate timed responses with eyelid conditioning. A. Top: conditioned stimulus (CS) mossy fiber input to the cerebellum as time-invariant step function activity. Middle: this input is transformed into stimulus-temporal code in the granule cell population, where different granule cells activate at different times. Each row is the activity of a granule cell, normalized by the maximum activity of that cell. The cells are sorted by the time of peak of activity. Examples of granule cells that are specifically active early and late are on the right. Bottom: Purkinje cells can use the granule stimulus-temporal code to learn to decrease their activity near the onset of the unconditioned stimulus (US), which drives the timing of the eyelid response. B. Correlation score matrix of the population activity at every time point vs. every other time point. The diagonal in the matrix is the correlation of the time point vs. itself. The triangle contains the points that are used to calculate a score from the matrix, which in turn reflects the quality of the stimulus-temporal population code. Top: correlation matrix of the mossy fiber activity. Bottom: population correlation matrix for the granule cell activity from A.

Golgi inhibition can produce stimulus-temporal code while maintaining biologically feasible connectivity. These results also provided constraints on the connectivity of a recurrent network that are suitable for producing stimulus-temporal code.

Methods

Simulating delay eyelid conditioning

The expanded simulation contains 1048576 granule cells, 1024 Golgi cells, 2048 mossy fibers, 32 Purkinje cells, 128 basket cells, 512 stellate cells, 8 deep cerebellar nucleus cells, and 4 inferior olivary cells. It is implemented in the C++ programming language and utilizes graphics processing units (GPUs) as described in chapter 2. The simulation uses a Linux system containing eight Nvidia GTX680 GPUs. The simulation emulates delay eyelid conditioning where the inputs are provided by mossy fiber (CS) and climbing fiber (US) stimulations. Many mossy fibers have been observed to have background activity (Aitkin & Boyd, 1978; Gould, Sears, & Steinmetz, 1993; Shinoda, Sugiuchi, & Futami, 1987), which is implemented in the mossy fiber inputs to the simulation. It is assumed that when the CS is presented to the animal through mossy fiber stimulation, only a small fraction of these fibers are activated. In the simulation, 50 (2.5% of total) mossy fibers are randomly selected to be the CS mossy fibers. These fibers have tonically elevated activity during the CS. The mossy fiber

activity is modeled by Poisson regenerative models (to model refractory periods between spikes) to provide noisy inputs to the simulation. The US is modeled as a climbing fiber spike from the inferior olivary neurons.

Data recording

The spikes that each cell generates are recorded in peri-stimulus time histograms (PSH) for 1000 trials. Each bin in the PSH is 10ms wide. The spike activity is recorded for the duration of the CS and for 500ms before and after CS. The state of the simulation is also recorded, which contains the connectivity between different cells and the internal activity state of each cell (membrane voltage, amount of conductances, and threshold).

Data analysis

The analysis code is written using the Python programming language (www.python.org). The analysis utilizes Numpy (www.numpy.org) and Matplotlib (www.matplotlib.org) for analyzing and visualizing data. PyCXX (cxx.sourceforge.net) is used to allow Python to communicate with the C++ code from the simulation in order to load and modify data from simulation recordings. The recorded Golgi and granule PSHs are sorted by time of peak activity to aid in visualizing the stimulus-temporal code (Figure 3.1A, middle). The Purkinje cell activity is recorded on a trial by trial basis, and the average population activity of

these cells over 200 trials is displayed.

In order to quantify the stimulus-temporal code of the Golgi and granule cell populations, Pearson's correlation coefficients are calculated for the population activity at different time points. A geometric interpretation of Pearson's correlation coefficient between the population activity at two time points is the cosine of the angle between the two normalized population activity vectors (Schmid, 1947). This measure is useful in quantifying the dimensionality of the population trajectory (Goudar & Buonomano, 2014). If the population activity is well correlated at all time points, the dimensionality of the activity trajectory is low and the cell population is effectively acting as one or a few cells. Conversely, if the population activity is not well correlated at different time points, the dimensionality of the activity of the population is high. High dimensional population activity is useful for producing well-timed responses by a downstream readout cell (Buonomano & Maass, 2009; Karmarkar & Buonomano, 2007). A correlation matrix is constructed (Figure 3.1B) where the population activity at each time point is correlated with all other time points. To reduce the matrix to a single measurable quantity, the average of the correlation scores between time bins that are 250ms apart (figure 3.1B, inside white triangles) is measured, and then subtracted from one so that a population with no stimulus-temporal code has a score of 0 and a population with perfect stimulus-temporal code has a score of 1.

Results

Timing in the expanded simulation with low granule-Golgi convergence ratios

The constrained simulation suggests that the interactions between granule and Golgi cells is the mechanism that transforms tonic mossy fiber input into stimulus-temporal code in the granule cell population activity. In order to test this hypothesis in the expanded simulation, there are two connections that must be defined. The first connection is the inhibitory output of Golgi cells to granule cells. The parameters for this connection were characterized by direct observations of the average number of dendrites per granule cell (Spacek et al., 1973), the number of glomeruli that a Golgi axon provides output to (Palkovits et al., 1971b), and the number of granule dendrites per glomerulus (Palkovits et al., 1972). A few observations and inferences exist to provide a basis for estimating the connectivity ratio of the second connection; the granule excitatory inputs to Golgi cells. Electron-microscopy observations of granule cell to Golgi cell synapses (Palay, 1974) suggest that a Golgi cell receives 1-6 granule synapses for every 20 microns of Golgi dendrite. Histological observations show that Golgi cell dendrites have few branches (Eccles et al., 1967; Palay, 1974), and that each Golgi cell has 3-10 of these dendrites (Eccles et al., 1967). Assuming that each dendrite pervades throughout the entire molecular layer of the cerebellar cortex

(the layer with granule cell axons and dendrites from various interneurons), the length of each dendrite should be around 300-500 microns (Eccles et al., 1967). Thus, a Golgi cell is likely to have between 1500-5000 microns of dendrite. Based on these estimates, a Golgi cell might receive hundreds to thousands of granule cell inputs. Another consideration is based on the observation that there are around 5000 granule cells for every Golgi cell (Palkovits et al., 1971b). Therefore, if the Golgi cell population receives inputs from the entire granule cell population, then on average each Golgi cell would receive at least 5000 granule cell inputs. Based on these considerations, in the expanded simulation, the connectivity ratios are set so that each granule cell provides output to up to 4 Golgi cells, and each Golgi cell receives up to 4092 granule cell inputs. However, there is considerable uncertainty in the above estimates, and the connectivity ratios in the expanded simulation are considered free parameters.

Figure 3.2 shows the results from the expanded simulation after training with the eyelid conditioning paradigm. The total Purkinje cell population responses are used to examine the ability of the simulation to produce timed anticipatory responses. This choice is appropriate because the timed decrease in activity reflects the cerebellar output that drives conditioned responses. In addition, the amplitude of the decrease reflects the robustness of the responses. The simulation performs well at short inter-stimulus-intervals (ISIs). However, the Purkinje cells in the simulation are not able respond robustly to long intervals

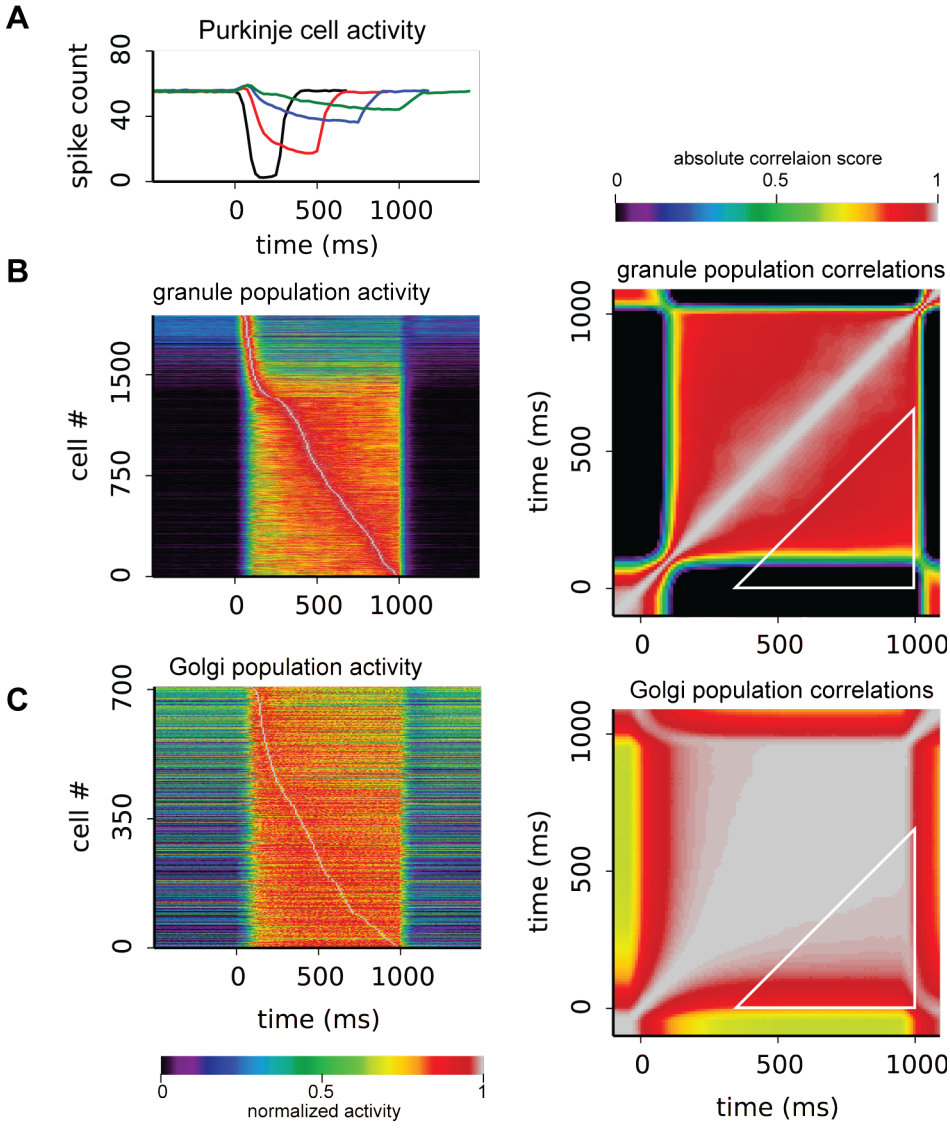

Figure 3.2. The expanded simulation with 4096 granule inputs per Golgi cell does not exhibit well timed responses. A. Average Purkinje response to mossy fiber inputs with different intervals. Intervals: Black: 250ms, red: 500ms, blue: 750ms, green: 1000ms. The Purkinje responses exhibit similar onset for all intervals and are not able to produce robust responses at longer intervals (>750ms) B. Left: Granule cell activity for 1000ms tonic mossy fiber input. The activity of each cell is normalized to its peak activity, and the cells are sorted by time of peak activity. Right: correlation matrix of the granule population activity at every time point compared to every other time point. The white triangle delineates the region of the scores used for the average of the correlation scores, (score: 0.39) see methods. C. Left: Golgi cell activity in the same simulation as the granule cells. Right: correlation matrix of the Golgi population activity. (score: 0.03)

(>750ms, figure 3.2A). At these longer intervals the onset of the responses is similar to the onset of the 500ms interval. This is contrary to the behavior observed in rabbits, where the onset of the behavior is delayed more for longer intervals (Medina & Mauk, 1999; White et al., 2000). The granule cell population activity shows that while some granule cells are phasically active at the onset of the CS, many cells produce tonic activity (figure 3.2B, left). The quality of the stimulus-temporal code is quantified by the population correlation measure (figure 3.2B, right, and see methods), which has a score of 0.39. Similarly, the Golgi cells exhibit tonic activity (Fig. 3.2C), with a score of 0.03 for the stimulus-temporal code. This activity is not surprising considering that the Golgi cells are driven by mossy fiber and granule cell inputs, and that both inputs are tonic. In summary, in the initial expanded simulation in which each Golgi cell receives 4092 granule cell inputs, the recurrent interaction between Golgi cells and granule cells does not produce stimulus-temporal code that can support timed responses for long intervals. This differs from animal behavior.

It is possible that the large number of granule cell inputs to Golgi cells is responsible for reducing the ability of the granule-Golgi recurrent interactions to produce a stimulus-temporal code. This possibility is explored because the exact convergence and divergence ratios of this connectivity remain as free parameters, as discussed previously. An error in the expanded simulation where the granule to Golgi input calculations effectively reduce the number of granule to

Golgi inputs to less than 256 provides additional evidence that this possibility exists. When this error is corrected so that each Golgi cell receives 4096 granule inputs, the ability of the expanded simulation to produce timed responses is diminished. To systematically test this possibility, the number of granule inputs to each Golgi cell is reduced from 4096 to 1024, 256, 64 and 16 in different simulations to examine the effect of the reduction on the generation of stimulus-temporal code. Figure 3.3 shows that as the number of granule inputs to each Golgi cell is scaled down, the Golgi cell populations exhibited slightly improved stimulus-temporal code (figure 3.3F). While the granule cell populations do not exhibit improved stimulus-temporal code by the correlation measures (figure 3.3E), the Purkinje cells produce more robust responses, which indicates that the stimulus-temporal code is more robust. However, the number of granule inputs per Golgi cell (64 and 16, which is similar to that in the constrained simulation) that show the best stimulus-temporal code in the expanded simulations is well beyond what the existing anatomical data supports. Therefore, an alternate mechanism where Golgi cells directly inhibit each other is explored.

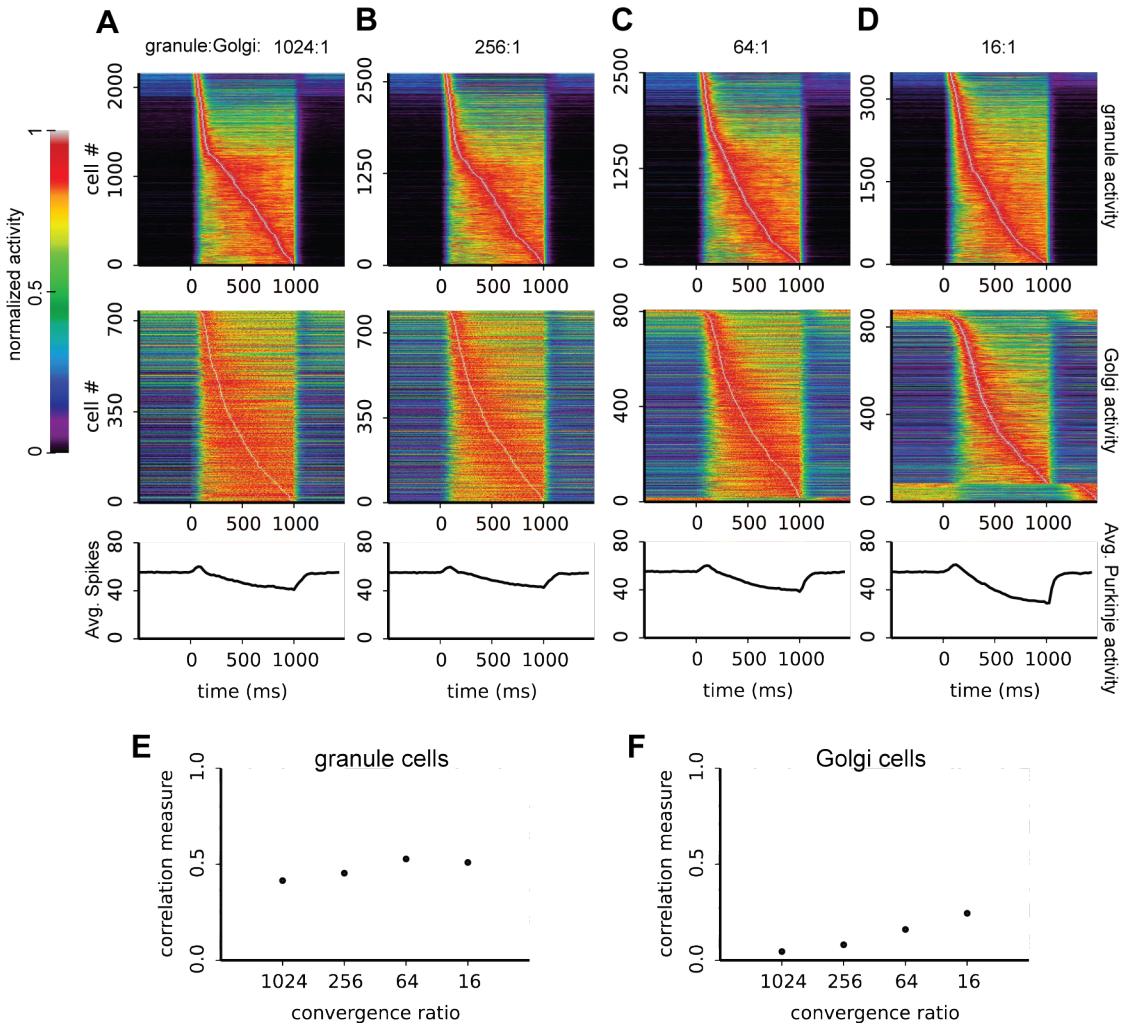

Figure 3.3. Decreasing the number of granule inputs per Golgi cell improves stimulus-temporal code for Golgi and granule cell population. A, B, C, D are 1024, 256, 64, and 16 granule inputs per Golgi cell respectively. Top: granule cell activity. The activity of each cell is normalized to its peak activity. The cells are arranged by time of peak activity. Middle: Golgi cell population activity. Bottom: averaged Purkinje cell activity. E. Correlation matrix measure of the granule activity in A-D (see methods). F. Correlation matrix measure of the Golgi activity in A-D.

Lateral Golgi inhibition is a mechanism for timing in the expanded simulation

The recurrent interactions between Golgi and granule cells are between two cell populations, without any interaction among cells within each population. However, studies in recurrent networks and temporal computation suggest that direct lateral interactions among cells in the same population can be useful in generating stimulus-temporal code (Buonomano & Maass, 2009). Such direct lateral interactions are unlikely to exist among granule cells as suggested by known anatomical observations (Ito, 1984). The Golgi cells on the other hand have been shown recently to directly inhibit each other (Hull & Regehr, 2012) (lateral Golgi inhibition). This lateral interaction might be a mechanism for producing stimulus-temporal code in the Golgi cell population activity, which can then induce stimulus-temporal code in the granule cell population activity. For example, a Golgi cell that is specifically active during the early period of the CS can inhibit the early responses of granule cells such that these cells would be active only during the late period of the CS.

There are three free parameters for implementing Golgi lateral inhibition in the expanded simulation. These are 1. the divergence ratio of the connectivity, 2. the convergence ratio of the connectivity, and 3. the spatial pattern of the connectivity. The acute slice physiology used by Hull and Regehr to report this connectivity cannot completely characterize these parameters. However, one

suggestive clue is that the findings (Hull & Regehr, 2012) are from paired recordings of nearest neighbor Golgi cells. The observed inhibition was a GABAergic synaptic conductance, likely from Golgi axons to Golgi soma or dendrites. The spread of Golgi axons in the granule layer is observed to be largely constrained to near the Golgi soma (Palkovits et al., 1971b). Therefore, assuming that the inhibitory connectivity comes from Golgi axons in the granule layer, the lateral Golgi inhibition connectivity is most likely constrained to nearest neighbors. Classical anatomical observations suggest that Golgi cells are arranged in a regular grid (Palkovits et al., 1971b), and in the simulation the Golgi cells are arranged in a grid where each cell had 8 nearest neighbors. Since many connections are cut in acute slices, it is likely that the number of connections observed (Hull & Regehr, 2012) represents an underestimate of the true connectivity in vivo. To attempt to correct for this underestimate, the proportion of connections that are observed to be reciprocal is used to estimate the probability of making a connection. Hull and Regehr report that 25 pairs of cells are recorded, each pair tested in two directions for the existence of a connection. Out of the 50 directions tested, 10 are found to connected, with 3 reciprocally connected pairs (i.e., 60% of the connections observed are reciprocal), which provides an estimate of the probability of making a connection at 0.6. This probability is implemented in the simulation as the probability of a Golgi cell making an inhibitory output to a neighbor.

Figure 3.4 shows that in the expanded simulation with a biologically feasible granule-to-Golgi convergence, the addition of Golgi lateral inhibition improves stimulus-temporal code in both granule and Golgi cell populations, as measured by the population correlation measures and Purkinje cells responses. The simulations used for this comparison are identical to the simulation in figure 3.2 (each Golgi cell in the simulations receives up to 4092 granule inputs), except with the addition of lateral Golgi inhibition. The granule cells are able to produce stimulus-temporal code that enables the Purkinje cells to learn more robust timed responses at the longer ISIs of 750, 1000, 1500, and 2000ms (figure 3.4A). These results are more consistent with delay eyelid conditioning in rabbits, for which rabbits can show timed responses beyond 1000ms (White et al., 2000). Interestingly, the ability of the expanded simulation with lateral Golgi inhibition to learn longer intervals is in contrast to the constrained simulation, which cannot produce timed responses at intervals longer than 750ms (Medina & Mauk, 1999).

It is possible that the improved stimulus-temporal code in the Golgi cell population activity is partially explained by the improved stimulus-temporal coding in the granule cell population activity. In other words, it is possible that an improvement in granule stimulus-temporal code enables the recurrent interactions between granule and Golgi cells to contribute to the stimulus-temporal code. To test this possibility, the total granule input to Golgi cells is examined. The results show that all Golgi cells received similar granule cell

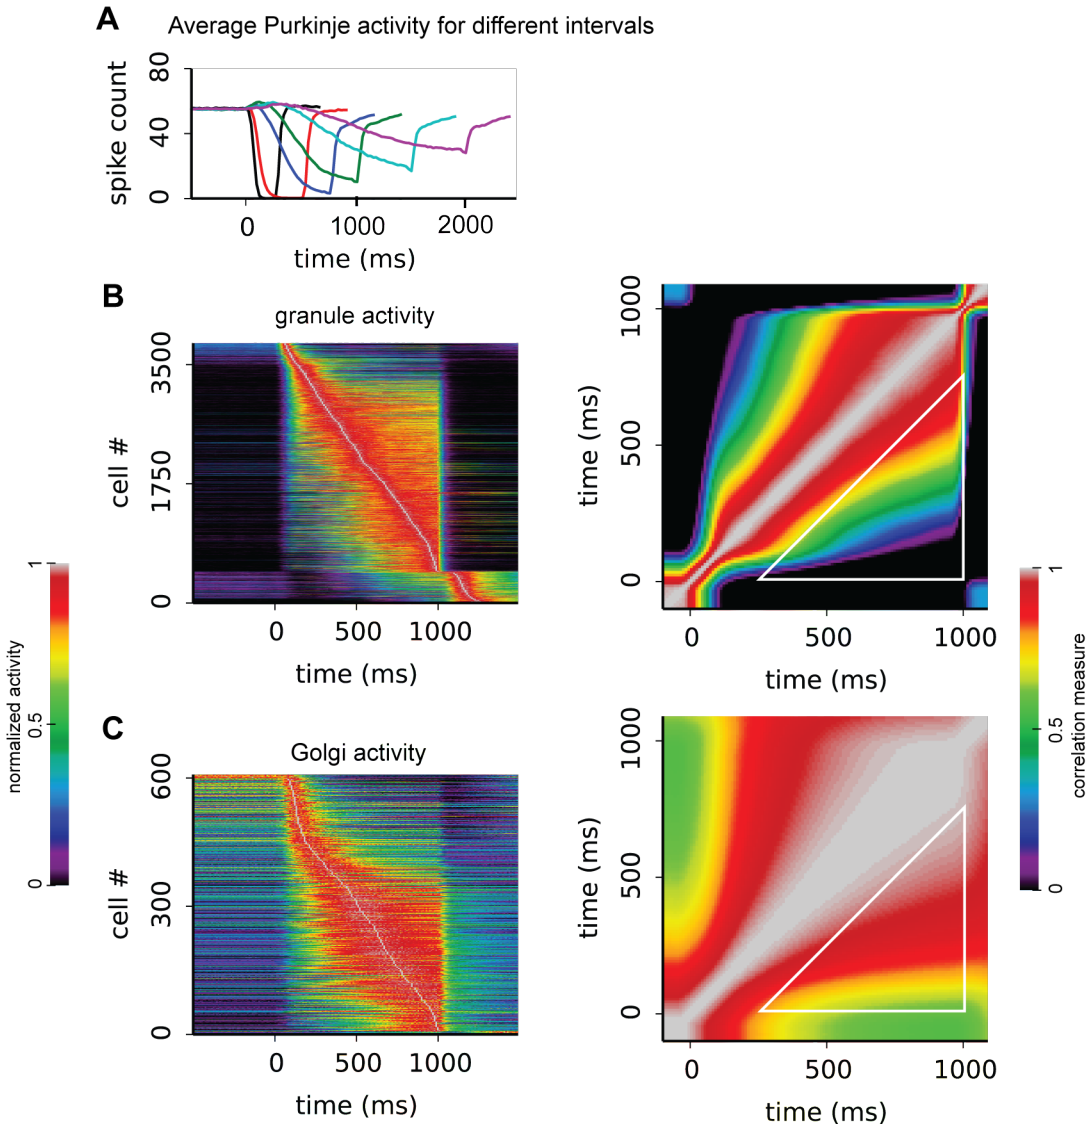

Figure 3.4. Lateral Golgi inhibition improves stimulus-temporal code in the expanded simulation and produced well timed responses. A. Purkinje responses to different ISIs. Black, red, blue, green, cyan, and magenta correspond to intervals 250ms, 500ms, 750ms, 1000ms, 1500ms, and 2000ms. B. Granule cell activity in response to 1000ms tonic mossy fiber input. Left: granule population activity. The activity of each cell is normalized to its peak activity. The cells are sorted by time of peak activity. Right: correlation matrix of the granule population activity at every time point compared to every other time point. The correlation measures are taken from the white triangle to produce a score (score: 0.66, compared to 0.39 in figure 3.2B, see methods) C. Golgi cell activity. Left: population activity. Right: Golgi correlation matrix, score: 0.14 (compared to 0.03 in figure 3.2B).

inputs regardless of the activity of the individual Golgi cell's activity (figure 3.5). This is unsurprising, since each Golgi cell receives input from 4096 granule cells, so although different granule cells are active at different times (stimulus-temporal code), the total input to each Golgi cell is tonic. This suggests that in the expanded simulation with lateral Golgi inhibition, the interaction between granule and Golgi cells does not play a role in producing stimulus-temporal code.

To further test if lateral Golgi inhibition is sufficient to generate Golgi stimulus-temporal code, the Golgi cell network from the full expanded simulation is extracted to be simulated without the granule cell population. The total excitatory input (combined mossy fiber and granule cell inputs) to each Golgi cell is recorded from the full simulation. The excitatory input is modified to be purely tonic to test if lateral Golgi inhibition is sufficient to generate stimulus-temporal code from tonic inputs. To test if lateral Golgi inhibition is necessary to generate stimulus-temporal code, this inhibition is disabled while the excitatory inputs are unmodified. Figure 3.6 shows that with lateral Golgi inhibition, the Golgi-only network is capable of generating stimulus-temporal code from tonic inputs (figure 3.6A). The excitatory inputs from granule cells to Golgi cells are insufficient to produce stimulus-temporal code without lateral Golgi inhibition (figure 3.6B).

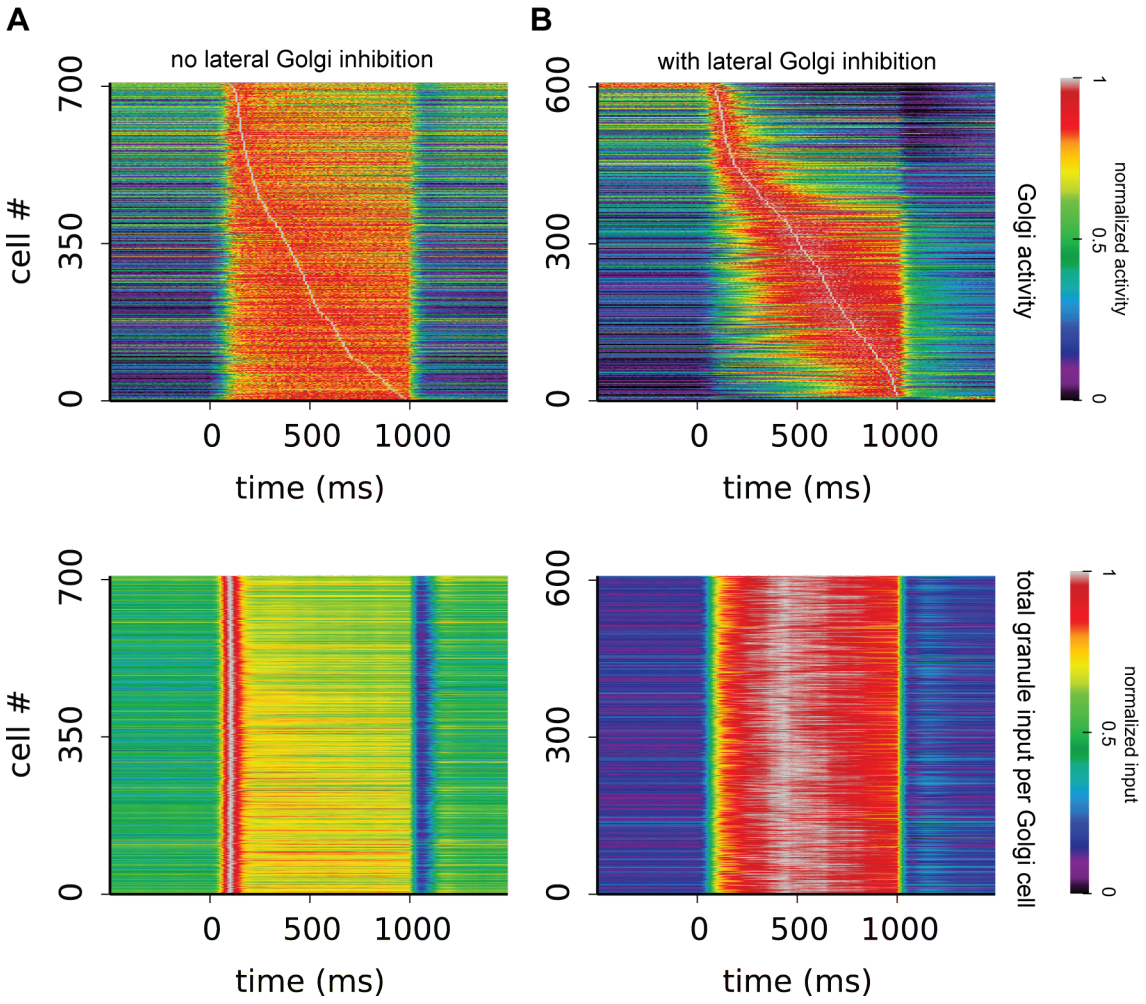

Figure 3.5. The activity of Golgi cells does not reflect the inputs from granule cells for the expanded simulation with lateral Golgi inhibition. A. Golgi cells in the simulation in figure 3.2, without lateral Golgi inhibition. Top: Golgi cell population activity. The activity of each cell is normalized to its peak activity. The cells are sorted by the time of peak activity. Bottom: total granule input per Golgi cell. The total input to each Golgi cell is normalized to the peak of input. The inputs are sorted in the same order as the Golgi cell activity. B. Input to Golgi cells In the simulation in figure 3.4, with lateral Golgi inhibition. Top: activity of Golgi cells, bottom: total granule input per Golgi cell.

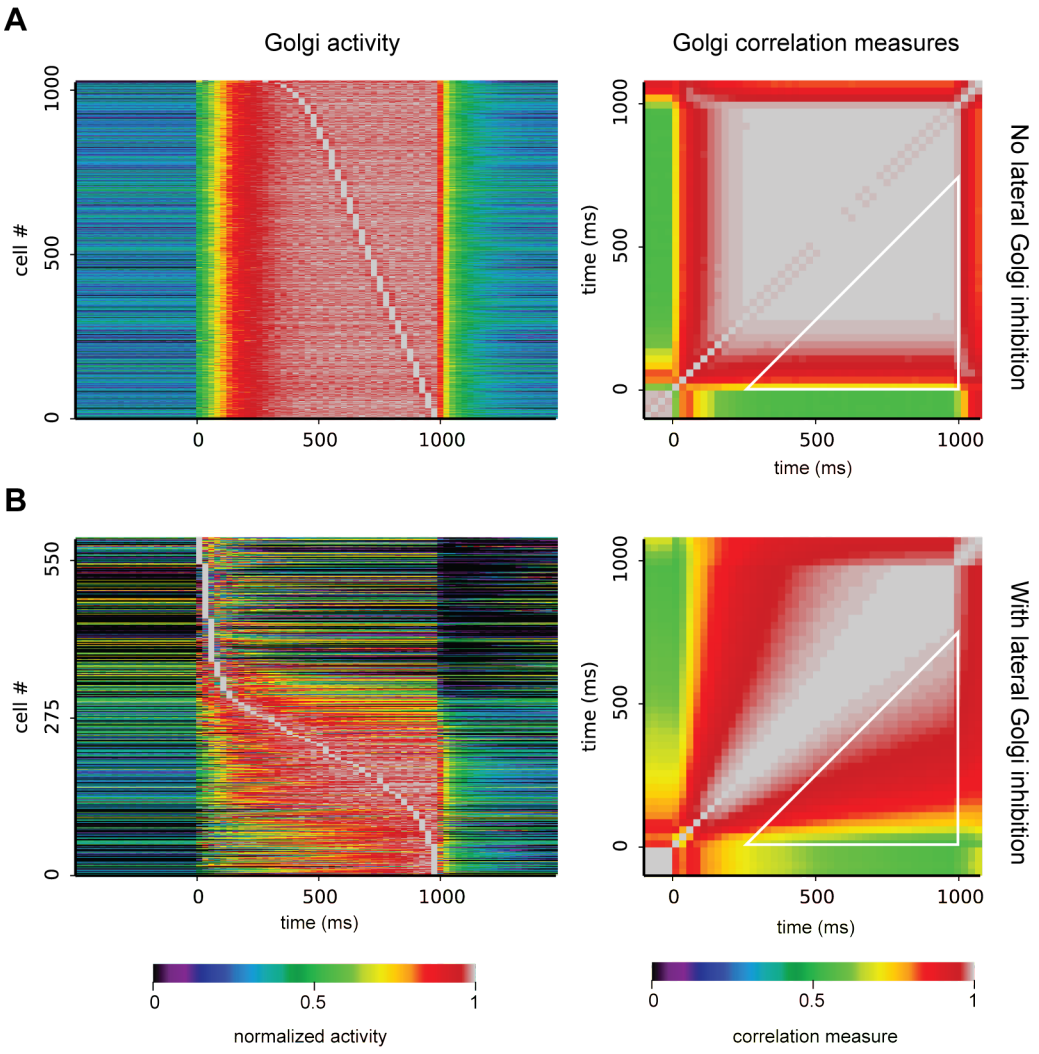

Figure 3.6. Lateral Golgi inhibition is necessary and sufficient to produce Golgi stimulus-temporal code. A. A network with only Golgi cells provided with recorded tonic excitatory inputs from the intact simulation. The network does not have lateral Golgi inhibition. Left: Golgi cell population activity. The activity of each cell is normalized to the peak of its activity. The cells are sorted by the time of peak activity. Right: correlation measure of the population activity. The measures in the white triangle is used to score stimulus-temporal code (score: 0.03, see methods). B. A network with only Golgi cells and with lateral Golgi inhibition. Left: Golgi cell activity, right: correlation measure of Golgi population activity (score: 0.13).

Summary and discussion

The results from the expanded simulation suggest that the interaction between granule and Golgi cells can transform tonic mossy fiber input into stimulus-temporal code so long as the convergence ratio of granule inputs per Golgi cell is similar to or less than the granule-Golgi convergence ratio in the constrained simulation. While there is no direct anatomical evidence regarding this ratio, a few electron-microscopy observations and the ratio between the number of granule cells to the number of Golgi cells provide the foundation for an order-of-magnitude estimation of this parameter. This estimation suggests that the conditions imposed by the connectivity in the expanded simulation are unlikely.

As a result, an alternate hypothesis for transforming tonic mossy fiber input to stimulus-temporal code is examined. The recent discovery that Golgi cells provide lateral inhibition to each other suggested the possibility that the lateral recurrent interactions within the Golgi cell population can be the primary mechanism for producing stimulus-temporal code. The results from the expanded simulation show that lateral Golgi inhibition produces stimulus-temporal code. This stimulus-temporal code yields more robust Purkinje cell responses for significantly longer intervals between CS and US onset. This is more consistent with rabbit delay eyelid conditioning behavior. Further manipulations of the network show that lateral Golgi inhibition is both necessary and sufficient to

produce stimulus-temporal code in the expanded simulation.

The results from the expanded simulation suggest constraints on the connectivity between two populations of cells that can produce stimulus-temporal code from tonic inputs. The limit on the number of granule inputs per Golgi cell that could produce stimulus-temporal code imply that only a small fraction (1/64 or less) of the granule cell population participates in the interaction with Golgi cells. More generally, the limit applies to networks where the only interactions in the network are between two populations of cells. When the size of one population is much larger than the other, these constraints limit the number of cells in the larger population that can participate to produce stimulus-temporal code.

However, these constraints only apply to situations where the interactions are strictly between two populations of cells, where there are no direct interactions within each population. The more likely situation is that one or both populations have lateral interactions within themselves. In such situations, as in the case of lateral Golgi inhibition, the above constraints no longer apply. Instead, lateral interactions within a cell population can be the primary mechanism for producing stimulus-temporal code (Buonomano & Maass, 2009).

The specific functional prediction from the expanded simulation that lateral Golgi inhibition is responsible for transforming tonic mossy fiber inputs to stimulus-temporal code has several unknowns. First, there is no direct

observation for the number of granule inputs per Golgi cell. Furthermore, the spatial connectivity pattern granule inputs to Golgi cells is unknown. This leaves open the possibility that Golgi cells receive segregated granule inputs, such that some Golgi cells only receive granule cell inputs that are active during the beginning of a mossy fiber input, while other Golgi cells receive granule cell inputs that are active near the end of a mossy fiber input. Such segregated inputs permit a relaxation of the connectivity constraints that are suggested by the simulation. If this is the case, or if future observations show that the number of granule inputs per Golgi cell is within the range suggested by the expanded simulation, then the functional role of lateral Golgi inhibition would need to be revisited.

CHAPTER 4:

COMPUTATIONAL MECHANISMS OF LATERAL GOLGI INHIBITION

Introduction

The results in chapter 3 suggest that in the expanded simulation, lateral Golgi inhibition can be the mechanism for transforming tonic mossy fiber inputs to stimulus-temporal code. This stimulus-temporal code enables the expanded simulation to more closely reproduce animal behavior in the delay eyelid conditioning paradigm. For the cerebellum, this paradigm presents the conditioned stimulus (CS) as tonic mossy fiber input, and the unconditioned stimulus (US) is presented at a fixed delay after the CS onset. The stimulus-temporal code that is produced from the tonic mossy fiber input enables the simulated network to keep track of the elapsed time since the CS onset. Purkinje cells use this stimulus-temporal code to produce responses in anticipation of the US onset.

The lateral Golgi inhibitiory network is similar to many recurrent networks. Studies of these networks emphasize that recurrent connectivity in the network can perform computation useful for keeping track of time (Buonomano & Maass, 2009; Buonomano & Merzenich, 1995, 1999; Buonomano, 2005; Laje & Buonomano, 2013; Liu & Buonomano, 2009; Lukoševičius & Jaeger, 2009; Miller, 2003; Sussillo & Abbott, 2009; Sussillo, 2014; Toyoizumi & Abbott, 2011; Wong & Wang, 2006). The computation is a result of the ongoing activity in a network of cells in response to an external stimulus. The ongoing activity is the result of the positive and negative feedback in the network (Maass, Joshi, & Sontag, 2007).

These feedback interactions produce cells with time-varying activity, such that the activity of different cells are decorrelated with each other. As a consequence, the population activity is also decorrelated between different time points (Buonomano & Maass, 2009). The pattern of population activity contains sufficient temporal information for a downstream readout cell to keep track of the elapsed time since the stimulus onset. The readout cell receives converging input from the entire network, and can learn to generate temporally specific responses by adjusting the weights of the input synapses. In contrast to networks with specifically designed neurons or circuitry (Aviel et al., 2003; Freeman & Nicholson, 1970), this class of models has stochastic connectivity and relies on the emergent network properties to perform its computation (Wiechert, Judkewitz, Riecke, & Friedrich, 2010).

However, the emergent properties of these networks present challenges in understanding their mechanisms. Understanding the properties of the components in the network generally provides very limited insight into the mechanisms of the entire network, because the network properties depend on the complex interactions among components (Funtowicz & Ravetz, 1994). These complex interactions can render manipulations of the network difficult to interpret, in part due to the potential ability of the network to to compensate for the manipulations. As a consequence of these difficulties, many of these models do not attempt to dissect the specific network mechanisms in detail. Instead, these

analyses focus on characterizing the behavior and the structure of the network, such as the connectivity (Laje & Buonomano, 2013) and the parameter spaces that produces different behaviors (Ostojic, 2014).

The analysis of the lateral Golgi inhibitory network in this chapter attempts to provide insight into the mechanisms of the network beyond describing the properties of the network. The lateral Golgi inhibitory network is dissected in detail, which shows that a small percentage of Golgi cells (~8%) that were active during early part of the CS are disproportionally important for producing stimulus-temporal code. Detailed dissections of the inhibitory sources to these early cells suggest that the nonreciprocal inhibition is important in producing early cell activity. The nonreciprocal inhibition is where cell A inhibits cell B, but cell B doesn't inhibit cell A. The results from the dissections are tested in the simulation by specifically removing the nonreciprocal inhibition to early cells (3% of the total number of connections). This manipulation disproportionally disrupts stimulus-temporal code in the simulation.

Methods

Manipulating Golgi network connectivity

Similar to the methods used in in chapter 3, Python, PyCXX, and Numpy are used for analysis. PyCXX is used to interface between the analysis code in

Python and the simulation (in C++). The simulation provides the recorded activity of Golgi cells (as peri-stimulus histograms, chapter 3 methods) and the connectivity matrix of the lateral Golgi inhibition network to the analysis code. The early cells are identified by the criteria that the time of peak activity is within the first 350 ms of the tonic mossy fiber input (CS), and that the activity at 800ms into the CS must be at most 50% of the peak. Once the early cells are identified, the cells that inhibit the early cells are also identified by using the connectivity matrix. In order to test the contributions of specific connections to producing stimulus-temporal code, the analysis instructs the simulation to add or remove specific connections. The activity of the Golgi cell activity in the manipulated network is recorded. The activity is used to observe the changes in stimulus-temporal code.

Producing isolated Golgi cell networks to test the contributions of components of the network

Isolated Golgi cell networks are used to test the contributions of specific connections for generating stimulus-temporal code in the lateral Golgi inhibitory network. These isolated networks are used to minimize possible feedback interactions from manipulating specific connections. The isolated networks are extracted from intact simulations. The excitatory inputs from the intact simulations are provided to the isolated network. These inputs are recorded from executing the intact simulation for 1000 trials of eyelid conditioning. During the execution, the input from granule cells and mossy fibers to each Golgi cell is recorded as

peri-stimulus histograms. These recorded inputs are provided to the isolated Golgi cell networks to substitute as the excitatory inputs.

Results

Early Golgi cells are important for generating a stimulus-temporal code

The Golgi cells that respond near the onset of the CS are chosen as the focus of analysis in order to understand the mechanism of the lateral Golgi inhibitory network. These neurons increase their activity during the beginning of the CS, and then decrease their activity during the late period of the CS. During the decrease, other Golgi cells (late cells) increase their activity. Therefore, it is possible that the activity patterns of the early cells delayed the responses of late cells by providing inhibition and then release of inhibition to late cells. This interaction can produce Golgi cells that respond to the CS at different times, which results in stimulus-temporal code. If this is true, disrupting the inhibition to early cells, such that their activity remain tonically elevated throughout the duration of the CS, should entirely inhibit the late cell responses. This should disrupt the stimulus-temporal code. To test this hypothesis, five simulations are constructed, each with the same connectivity parameters but with different specific connectivity. This is due to the stochastic nature of making the connections (see chapter 2, Simulation connectivity for a detailed description of the stochastic process for connecting the cells in the simulation). The reason to

use multiple simulations is to ensure that the results are reliable. These simulations are modified to remove all inhibition to early Golgi cells. The early cells are defined as cells whose time of peak activity is less than 350ms into the CS. In addition, only the early cells that decrease activity during later period of the CS (i.e., show strong temporal coding) are selected. Overall, these cells account for 9% of all Golgi cells in the five simulations. In another identical set of five simulations, a matching number of Golgi cells (that are not early cells) are randomly selected. The inhibition to these cells are removed to provide a control comparison. If the early cells are important for producing stimulus-temporal code, then disrupting the inhibitory input to these early cells should have a greater impact on the stimulus-temporal code compared to disrupting the randomly selected cells.

Prior to the manipulation, each simulation is trained for 1000 trials using the 1000ms inter-stimulus training protocol (see chapter 3 methods). At the end of 1000 trials, the Purkinje cells in the simulations are making robust learned responses in anticipation of the US. The early Golgi cells are then identified, and all inhibition to these cells is removed from the simulation. As a control, the inhibition to a matching number of randomly selected cells are removed from an identical copy of that simulation (but with intact inhibition to early cells). For both manipulations, the granule-Purkinje synaptic plasticity is frozen so that the simulations do not extinguish their acquired responses. The behavior of each

simulation is observed for 500 trials after the manipulation, during which the Golgi cell and Purkinje cell activities are recorded for analyzing the timing performance and stimulus-temporal code.

Figure 4.1 shows the representative results from one of the five simulations. In this simulation, the inhibition to 88 early Golgi cells is removed, which produced tonically elevated activity in the early cells (figure 4.1, compare A to B). Disrupting the activity of these 88 cells (out of the 1024 total number of Golgi cells in the simulation) significantly reduces the robustness and timing of the Purkinje cell responses to the CS (figure 4.1D, compare black to red lines). As a control, disrupting the inhibition to 88 randomly selected Golgi cells do not affect the Purkinje cell responses to the same degree (figure 4.1C and D, compare black to blue lines). The results from the other four simulations are similar to that in figure 4.1. These results suggest that early cells are specifically important for producing stimulus-temporal code, and that further analysis can focus on these cells.

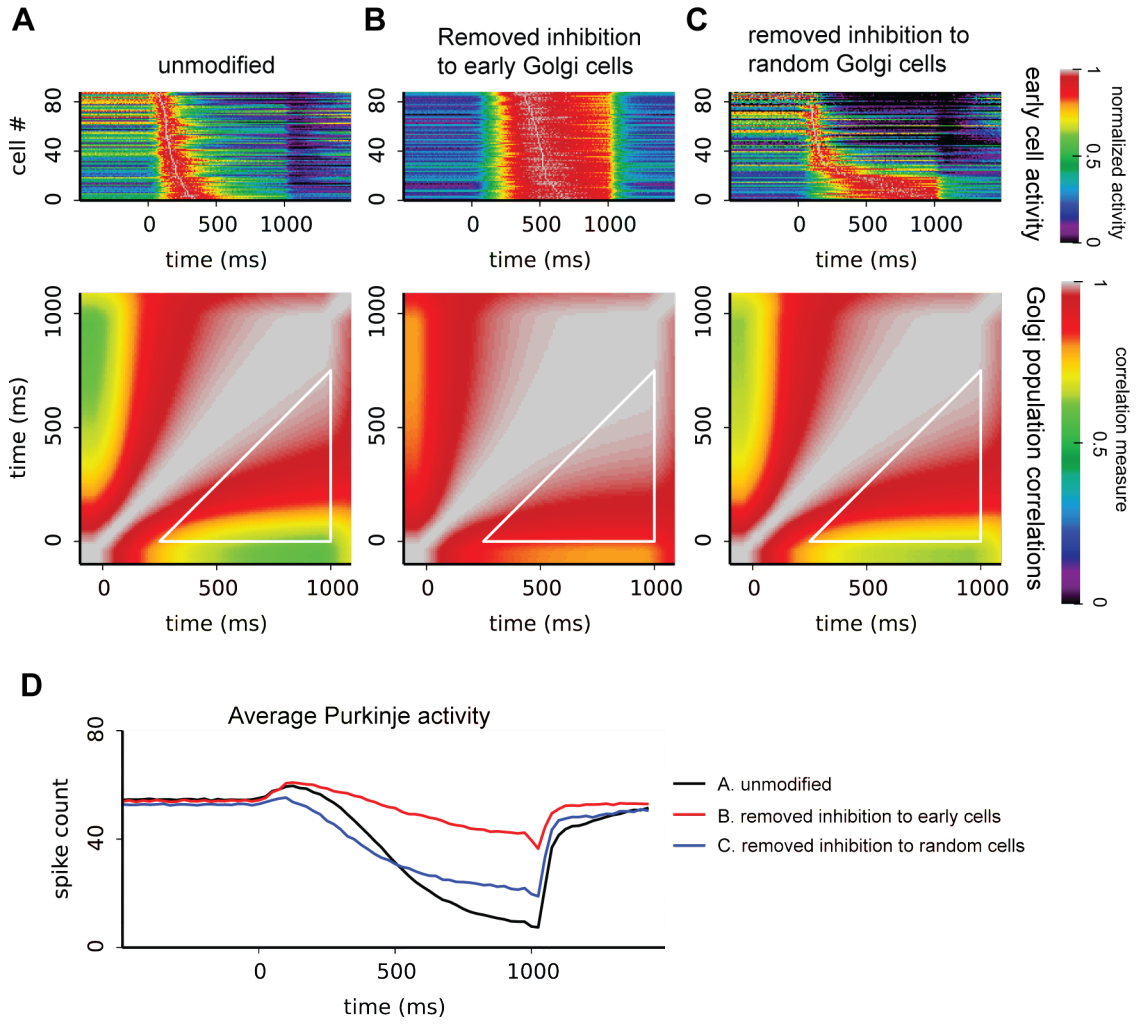

Figure 4.1. Early Golgi cells are specifically important for producing stimulus-temporal code. A. Unmodified simulation. Top: activity of early cells. The activity of each cell is normalized to its peak activity. The cells are sorted by the time of peak activity. Bottom: correlation matrix of the Golgi cell population activity at every time point compared to every other time point (see methods, chapter 3). The correlations in the white triangle are used to calculate the score (0.13, see chapter 3 methods). B. The simulation in A modified by removing inhibition to the early cells. Top: activity of the same early cells as that in A. Bottom: correlation matrix of the Golgi cell population activity, score: 0.05. C. Removing inhibition to a matching number of randomly selected Golgi cells from the simulation in A. Top: activity of the same early cells as that in A and B. Bottom: correlation matrix of the Golgi cell population activity, score: 0.10. D. Average Purkinje cell activity across 200 trials for the simulations in A-C. Black: Purkinje cell activity for the unmodified simulation (A). Red: Purkinje cell activity for the simulation that removed inhibition to early cells (B). Blue: Purkinje cel activity for the simulation that removed inhibition to randomly selected cells (C).

Component analysis of the inhibition to early cells

Figure 4.1 shows that the stimulus-temporal code is disrupted when inhibition to early cells are removed (fig. 4.1B). An important aspect of the early cell activity is that the activity decreased during later period of the CS. This is disrupted by removing the inhibition to these cells as shown in figure 4.1B. Therefore, the sources of inhibition that induce early cells to decrease activity is chosen as the focus of the analysis. To this end, a total of 274 early cells from the five simulations from the previous section are analyzed. These early cells are selected based on the criterion that each cell must show strong decrease in activity during late period of the CS. The initial observations reveal two types of inhibitory connectivity to the early cells: 2/3 of the cells that inhibited the early cells also receive inhibition from the early cells (i.e., reciprocal connections), while the other 1/3 of the cells show only nonreciprocal connectivity. To investigate which connectivity is important in inhibiting early cells, each early cell and the 1st order cells (cells that provide inhibition to the early cell) are simulated in isolation from the rest of the network. The aim of this reductionist approach is to determine which 1st order cells are sufficient to decrease the activity of an early cell during late period of the CS. However, simulating an early cell and its 1st order cells in complete isolation can result in false positives. Specifically, it is possible that a 1st order cell can decrease the early cell's activity in the isolated network, but cannot in the full simulation. In order to eliminate this possibility, all

1st order cells also receive the cumulative inhibitory inputs from 2nd order cells (cells that provided inhibition to 1st order cells). Since these 2nd order cells are not present in the isolated network, their recorded activity from the full simulation is given to the 1st order cells as a substitute. This method enables the isolated network to operate under the same inhibitory environment of the full simulation. Under this scheme, the 1st order cells receive the full inhibition as they do in the full simulation. In contrast, the early cell only receives inhibition from the different categories of 1st order cells outlined below. Under these conditions, the early cell receives less inhibition than in the full simulation, whereas the 1st order cell that is tested receives at least the same amount of inhibition as that in the full simulation. If a selected category of 1st order cells can still decrease the activity of the early cell in these conditions, then it should also be able to do so in the full simulation.

Each early cell is first tested with a single reciprocal 1st order cell. Figure 4.2A shows an example of the isolated network scheme. The trial activity of an early cell is shown. The example shows a single reciprocal 1st order cell is able to inhibit the early cell's activity. The activity of the early cell without any inhibition is shown (grey line) as a comparison for the decreased activity due to the reciprocal 1st order cell (blue line) in the isolated simulation. 15 of the 274 early cells have such a reciprocal 1st order cell. It is possible that some early cells require the full complement of reciprocal 1st order cells in order to decrease

activity. To test this possibility, the isolated simulation for each of the remaining 259 early cells contains all of the reciprocal 1st order cells. Only 10 of the 259 early cells (2.6%) decrease activity as the result of this configuration (e.g., figure 4.2B). In total, the early cell response pattern can only be partially replicated in 35 of the 274 cells (12.8%) by reciprocal inhibition from 1st order cells. The only remaining candidate category of connectivity for 1st order cells is that with nonreciprocal connectivity. For each of the 274 early cells, an isolated simulation is constructed and contained only nonreciprocal inputs from 1st order cells. In this configuration, 197 out of 274 cells (71.9%) show decreased activity (figure 4.2C). These results from dissecting the components of the inhibition to early Golgi cells suggest that nonreciprocal lateral inhibition within the Golgi cell network is the primary factor for inhibiting early cells.

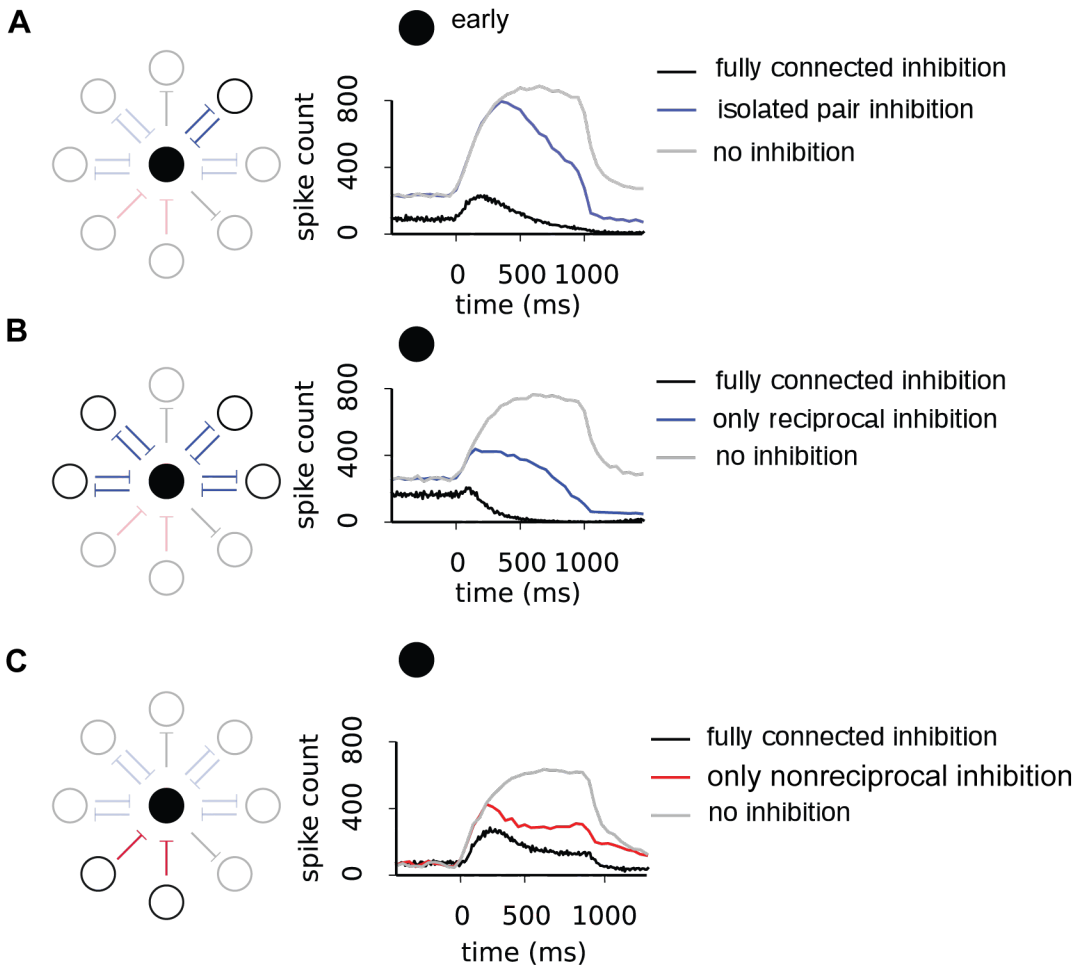

Figure 4.2. Early cell activity in isolated simulations shows that nonreciprocal 1st order cells are the primary factor for decreasing early cell activity. Each isolated simulation contains an early cell and the 1st order cell(s) of that early cell in the category that is tested. Left: schematic of each isolated simulation. Right: activity of the early cell without inhibition (grey), in the isolated network (blue and red), and in the full network (black). A. Isolated network containing a single reciprocal 1st order cell. 15 out of 274 early cells decrease activity due to a single reciprocal 1st order cell. B. Isolated network containing all reciprocal 1st order cell, 10 out of 274 early cells decrease activity due to all reciprocal cells. C. Isolated network containing all nonreciprocal 1st order cell. 192 out of 274 early cells decrease activity due to nonreciprocal cells.

Testing the necessity of the nonreciprocal inhibition to early Golgi cells for the emergence of stimulus-temporal code

The analysis of the components of the inhibition to early cells from the previous section suggests that nonreciprocal 1st order cells are important to decrease the activity of early cells during late period of the CS. The decreased early cell activity can then disinhibit the neighboring cells to allow these cells to respond during late period of the CS. If this is the case, and if the results from the component analysis are relevant to the emergent properties of the full network, then specifically removing the nonreciprocal connections to the early cells in the full simulation should disrupt the simulation's ability to produce stimulus-temporal code. To test this hypothesis, the same five simulations used above are modified so that all nonreciprocal inhibition to early cells is removed. This manipulation affects approximately 3% of all lateral connections between Golgi cells. As a control, in the same five simulations (in the unmodified state), a matching number of reciprocal inhibitory connections to the early cells are removed. If nonreciprocal inhibition is specifically necessary to decreasing the activity of the early cells, then the early cell activity should be tonically elevated when nonreciprocal inhibition is removed. The effect should be similar to removing all inhibition to these cells (shown in figure 4.1). This effect should be stronger than removing reciprocal inhibition. The stimulus-temporal code in the Golgi cell population should be more disrupted with nonreciprocal inhibition removed.

Figure 4.3B shows that removing nonreciprocal inhibition to the early cells disrupts early cell activity and Golgi population stimulus-temporal code more than removing reciprocal inhibition to the early cells (figure 4.3C). The Purkinje cells are not able to to produce robust and well-timed responses (figure 4.3D, compare blue and red lines). These results suggest that the findings from the component analysis are applicable and relevant to the emergent properties of the full network.

Nonreciprocal inhibition enhances stimulus-temporal code

If nonreciprocal lateral Golgi inhibition is important to generate stimulus-temporal code within the cerebellar network, then a simulation with only reciprocal Golgi inhibition should not be able to produce stimulus-temporal code. To test this, a simulation is constructed with only reciprocal lateral Golgi inhibition, but with a similar number of connections (6000) as that of one of the unmodified simulation (5583 connections) in figure 4.3A. Figure 4.4B shows that the simulation with only reciprocal lateral inhibition cannot produce stimulus-temporal code. The resulting Purkinje cell behavior is similar to the simulation without lateral Golgi inhibition (figure 4.4C).

An extension to the previous result is that a Golgi network with only nonreciprocal inhibition might produce better stimulus-temporal code than a

network with mixed reciprocal and nonreciprocal inhibition. To test this possibility and other network connectivity patterns and conditions that can all a network to generate stimulus-temporal code from tonic inputs, hypothetical neural networks (based on the Golgi cell network) are constructed. These networks contain a single population of neurons with lateral inhibition to each other. These networks are used to test various connectivity patterns. Each cell in the network received an external excitatory input that is independent of the network, and mimicked the excitation during the CS input in the expanded simulation. In these networks, different connectivity patterns are tested under a range of inhibitory synaptic strengths to explore the robustness of these networks in producing stimulus-temporal code. The connectivity patterns tested are as follows: 1. nearest neighbor mixed reciprocal and nonreciprocal connectivity similar to the simulation in figure 4.4A, 2. nearest neighbor fully reciprocal connectivity similar to the simulation in figure 4.4B, 3. nearest neighbor fully nonreciprocal connectivity, 4. spatially unconstrained (not constrained to nearest neighbor, but instead can connect to any cell in the network, with the constraint that the number of inputs and outputs per cell is the same as that in the nearest neighbor cases) fully reciprocal connectivity, and 5. spatially unconstrained nonreciprocal connectivity. The stimulus-temporal code produced by these networks is measured by the correlation matrix. Figure 4.5 shows that regardless of the spatial constraints of the connectivity, fully reciprocal connectivity (figure 4.5D and E) can not produce

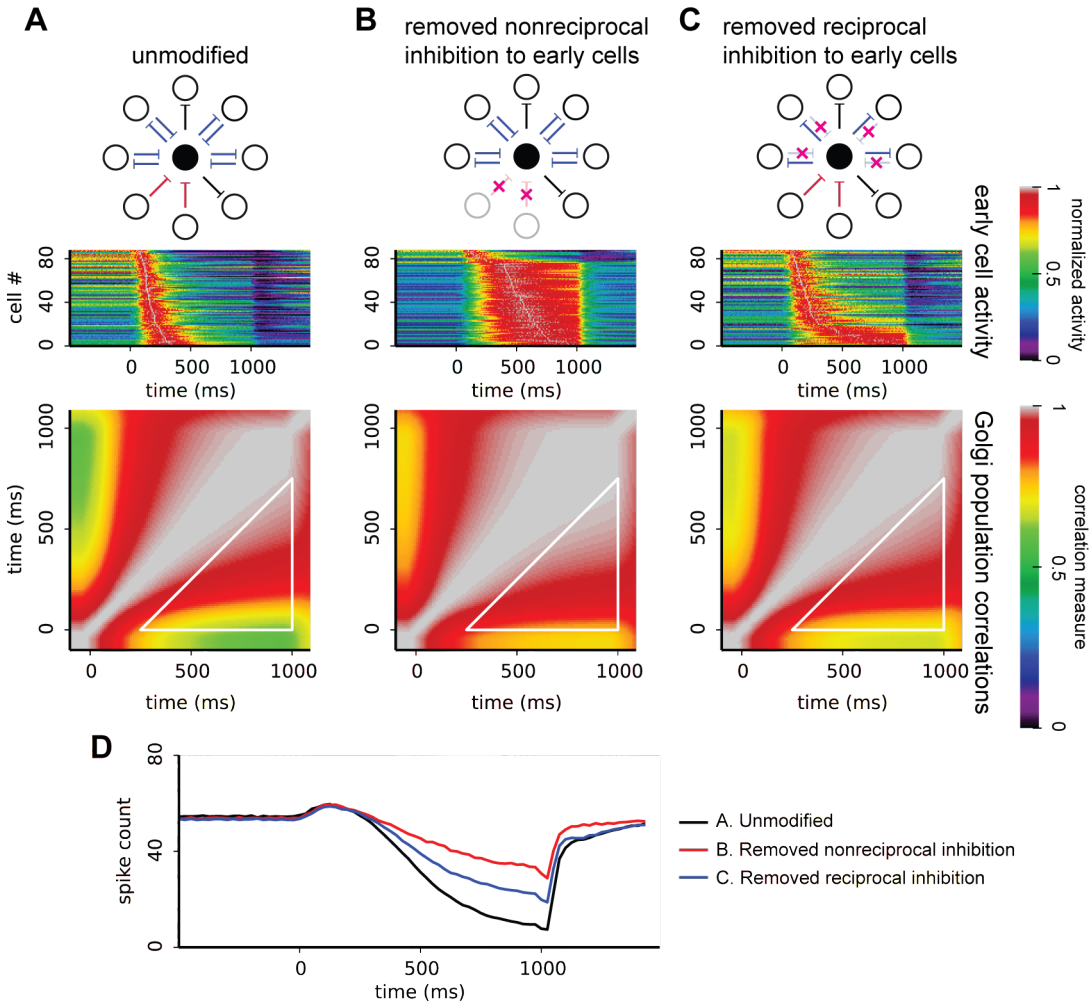

Figure 4.3. Nonreciprocal inhibition to early cells is important for producing stimulus-temporal code. A. Unmodified simulation. Top: schematic of the type of inhibitory connectivity to the early cells. Middle: activity of early cells. The activity of each cell is normalized to its peak activity. The cells are sorted by the time of peak activity. Bottom: correlation matrix of the Golgi population activity at one time point compared to another time point (see methods in chapter 3). The measures in the white triangle are used to score the stimulus-temporal code. (score: 0.13). B. The same simulation in A, with nonreciprocal inhibition to early cells removed (top). Middle: activity of the same early cells as in A. Bottom: correlation matrix of the Golgi population activity (score: 0.07). C. The same simulation in A, with reciprocal inhibition to early cells removed (matched to the same number of connections removed as in B). Middle: activity of the same early cell as in A and B. Bottom: correlation matrix of the Golgi population activity (score: 0.10). D. Averaged Purkinje cell activity for the simulations in A-C. Black: unmodified simulation, red: removed nonreciprocal inhibition to early cells, blue: removed reciprocal inhibition to early cells.

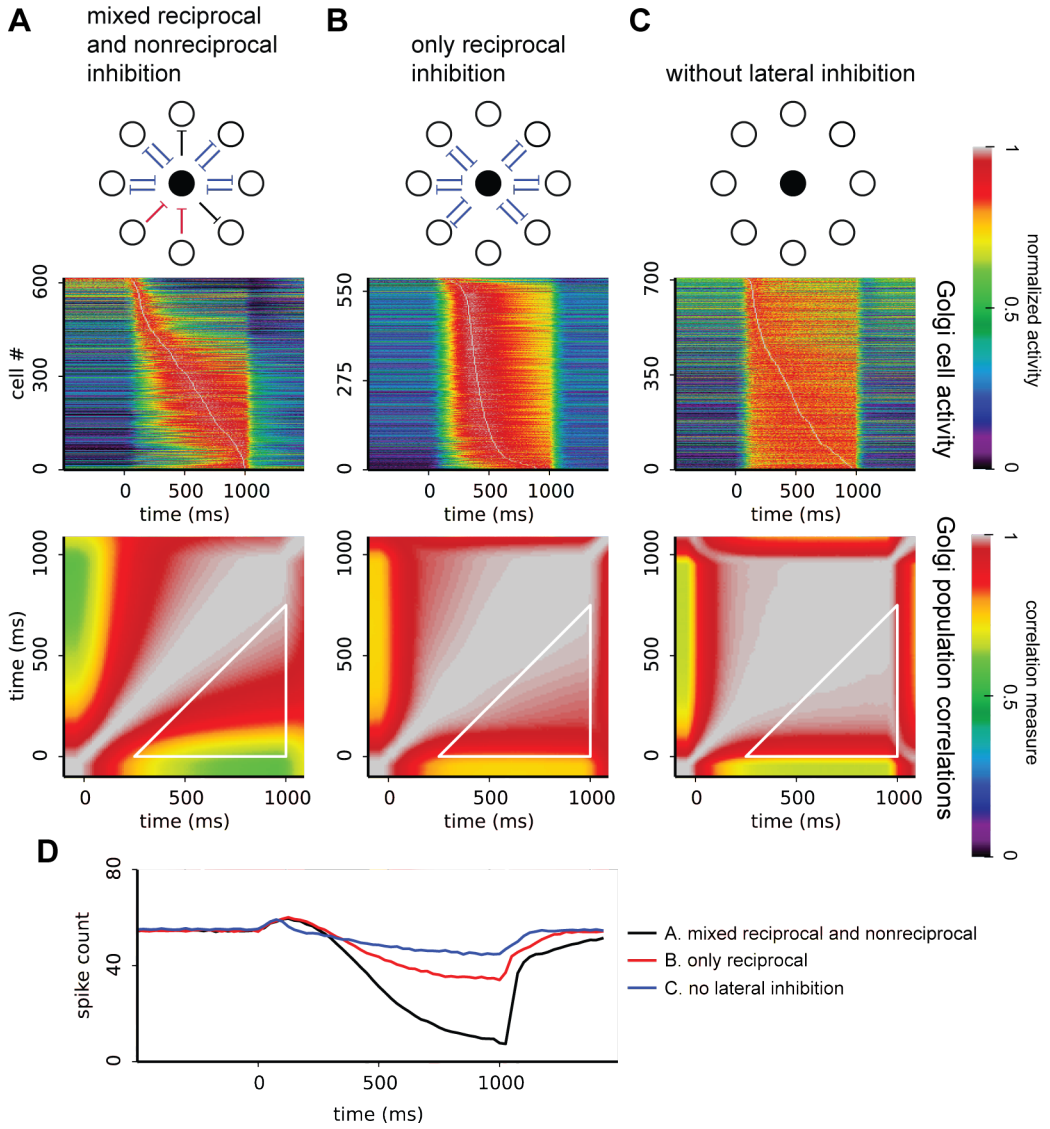

Figure 4.4. Exclusive reciprocal lateral Golgi inhibition does not produce stimulus-temporal code. A. A Simulation with mixed reciprocal and nonreciprocal lateral Golgi inhibition, similar to the simulation in figure 4.1A and 4.3A. Top: schematics of types of connectivity in the lateral Golgi inhibition. Middle: Golgi cell activity. The activity of each cell is normalized to its peak. The cells are sorted by time of peak activity. Bottom: Correlation matrix of Golgi population at every time point compared to every other time point. See methods in chapter 3. The correlations in the white triangle are used to score the stimulus-temporal code. (score: 0.13). B. Simulation with only reciprocal lateral Golgi inhibition (top). Middle: Golgi cell activity. Bottom: correlation matrix of Golgi population activity (score: 0.04). C. Simulation without lateral Golgi inhibition (top). Middle: Golgi cell activity. Bottom: correlation matrix of Golgi population activity (score: 0.03). D. Average Purkinje cell activity of the simulations in A-C.

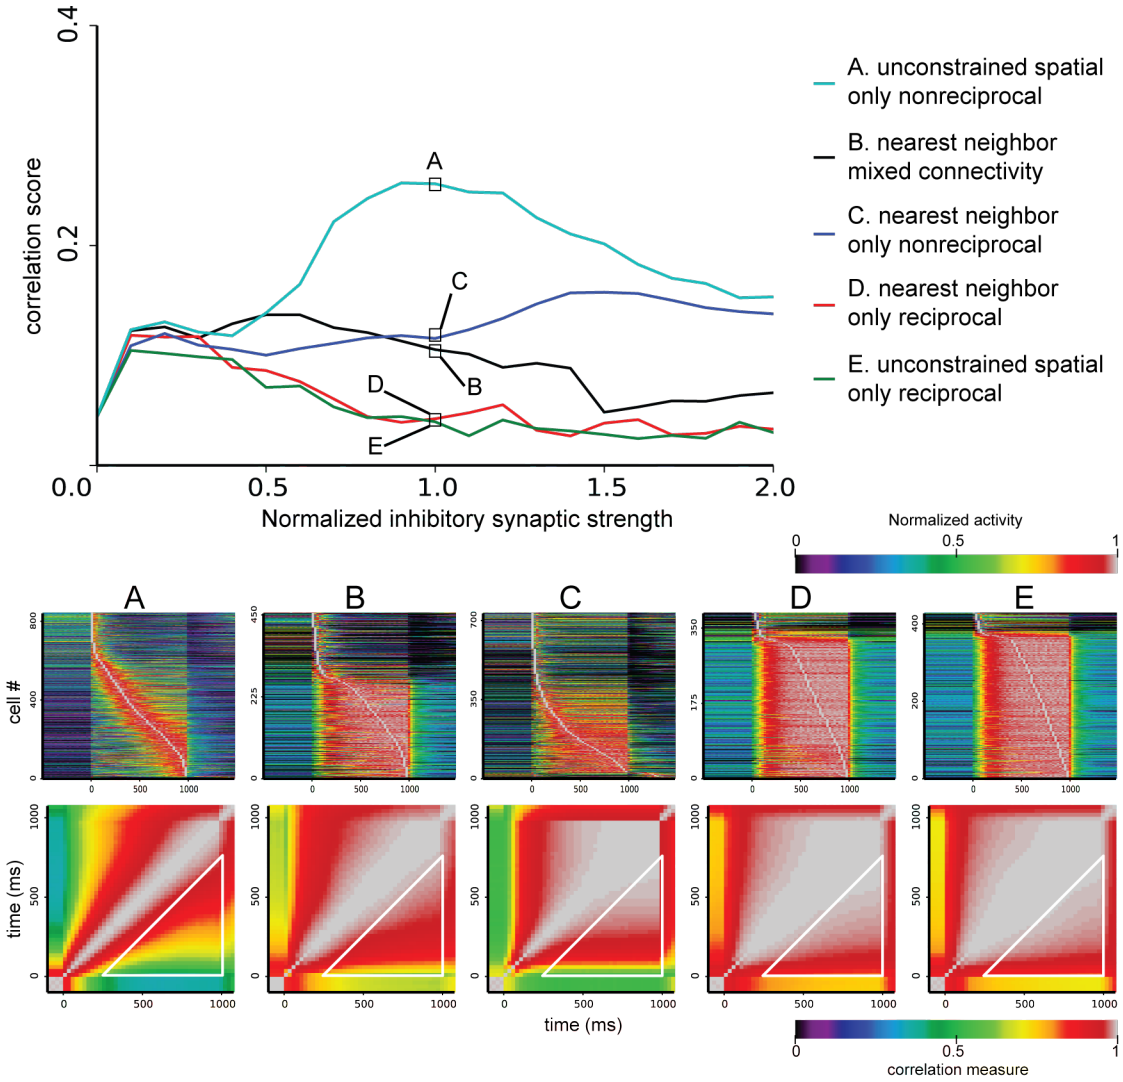

Figure 4.5. Networks containing only Golgi cells show that nonreciprocal inhibition generated stimulus-temporal code in a wide range of inhibitory synaptic strengths. Top: the quality of the stimulus-temporal code (as measured by the scoring the correlation matrix of the population activities, see chapter 3 methods) of Golgi-only networks across a range of synaptic strengths of lateral Golgi inhibition. The synaptic strengths shown are normalized to the values initially tuned for the simulations in figure 4.1A, 4.3A, and 4.4A. Bottom: the behavior of the networks at the selected points shown at top. First row: Golgi cell activity. Second row: correlation matrix of the Golgi population activity. A. Lateral inhibition not constrained to the nearest neighbor, B. constrained to the nearest neighbor, with mixed reciprocal and nonreciprocal connectivity, C. constrained to the nearest neighbor, and only nonreciprocal connectivity, D. constrained to the nearest neighbor, and only reciprocal connectivity, E. not constrained to nearest neighbor, and only reciprocal connectivity.

stimulus-temporal code within most of the inhibitory synaptic strength range. In contrast, networks that are fully (figure 4.5A and C) or partially (figure 4.5B) nonreciprocal can produce stimulus-temporal code in a wider range of synaptic strengths.

Summary and discussion

The analysis in this work dissects the lateral Golgi inhibitory network (a recurrent network) to understand its emergent properties in transforming tonic mossy fiber inputs (the CS) into stimulus-temporal code. The early Golgi cells are proposed to be important for this transformation, since these cells increase activity before other cells (near the onset of the CS), and subsequently decrease activity in late period of the CS. The decrease in activity is thought to be important for disinhibiting neighbor cells. These neighbor cells can then respond during late period of the CS. To test this hypothesis, all inhibition to early cells is removed in the expanded simulation, which transformed early cell activity to tonic activity, and severely disrupted the stimulus-temporal code. In contrast, removing all inhibition to a matching number of randomly selected cells has little effect. These results suggest that the inhibition to early cells is important for transforming tonic mossy fiber input into stimulus-temporal code.

The connectivity suggests two sources of inhibition to early cells: reciprocal and nonreciprocal. The early cells provide inhibition to the source cells

for the reciprocal connectivity, and not for the nonreciprocal connectivity. To determine which type of sources is important for inhibiting the early cell, the network around each early cell is dissected individually using isolated simulations. Each of these simulations contains an early cell and the neighboring cells. In these isolated networks, each neighboring cell is tested individually for its ability to decrease the early cell's activity. To eliminate possible false positives, each neighboring cell is provided with the inhibition recorded in the full simulation. This way the early cell only receives inhibition from the neighboring cell being tested, while the neighboring cell receives full inhibition as in the conditions of the full simulation. Under these conditions, if the neighboring cell can inhibit the early cell activity, then the it should be able to perform the same function in the full simulation. The results from these isolated simulations show that for the majority of early cells, the nonreciprocal inhibition is important to decrease their activity.

A critical test of these findings is to manipulate the intact full network and test the effects of this proportionately minor but very specific change in connectivity on the stimulus-temporal code generated by the network. The results from manipulations of the full simulation are consistent with the predictions from the isolated networks. These results suggest that the analysis of the isolated networks are relevant to the full network. Further tests of different connectivity patterns show that simulations with only reciprocal inhibition can not produce

stimulus-temporal code. Finally, hypothetical networks using a single population of cells with lateral inhibition show that nonreciprocal inhibition is important for the network to produce stimulus-temporal code in a wide range of inhibitory synaptic strengths.

The isolated network analysis does not explain the mechanisms for which all early Golgi cells decreased activity. Out of 274 cells, the analysis is able to explain the sources of inhibition for a total of 217 cells, which left 57 cells unexplained. This can be due to the limited fidelity of the isolated network simulation in representing the inputs from the intact simulation. This is a consequence of using the peri-stimulus histograms as the basis of the input. Each peri-stimulus histogram represents the average activity for 1000 trials and does not capture trial to trial variability. Therefore, peri-stimulus histograms are not a complete representation the network activity. It is possible that further analysis using trial to trial activity can account for the remaining 57 cells.

Despite the incompleteness of the isolated network analysis, the relevance of its results to the mechanism of the full simulation indicates that the analysis is insightful for the emergent properties of the lateral Golgi inhibitory network. Given the complexity of recurrent networks and the potential existence of many compensatory pathways, it is not expected that the predictions from the isolated networks can be directly applicable to the mechanisms in the full network. The results show that these predictions are applicable and suggest that the approach

used in dissecting the individual early cells is productive.

CHAPTER 5:

LIMITATIONS AND SIGNIFICANCE

Summary

The cerebellum is vital for precise motor control (Bastian et al., 2000; Morton & Bastian, 2004; Palliyath & Hallett, 1998) and is believed to produce the tuning signals for coordinated, smooth, and precise motor movements (Ito, 1984; Manto et al., 2012). The information carried by such tuning signals can involve precise amplitude and timing information (Ivry et al., 2002; Ulloa et al., 2003) (i.e., how much output and when to produce the output). In addition, the tuning signal may need to be adaptable to adjust to new environments and changing conditions. The delay eyelid conditioning behavior paradigm has been shown to directly engage the cerebellum and exhibits the amplitude (Kreider & Mauk, 2010), timing (Kalmbach et al., 2010; Perrett et al., 1993), and adaptability (Garcia et al., 1999) properties of the cerebellum. Therefore, this paradigm has been used as a powerful tool for understanding the computational properties of the cerebellum.

An important property of the delay eyelid conditioning behavior is that animals learn to produce well-timed responses given a tonic mossy fiber input as the (Aitkin & Boyd, 1978; Hesslow et al., 1999) conditioned stimulus (CS). After training, an animal does not respond at the CS onset, but delays its response (White et al., 2000) until shortly prior to the onset of the air puff to the eye (unconditoned stimulus, US). This aspect of the behavior suggests that the cerebellum is capable of keeping track of the elapsed time since the CS onset

(Medina et al., 2000). Previous constrained simulations by Buonomano, Medina, and Mauk (Buonomano & Mauk, 1994; Medina et al., 2000) have suggested that the emergent properties of the cerebellar network can transform tonic mossy fiber inputs into stimulus-temporal code in the granule cell population. This stimulus-temporal code can be used by the Purkinje cells to generate timed responses. The emergent properties are dependent on the stochastic connectivity between Golgi and granule cells, without specifically designed circuitry or cellular properties beyond the existing observations of the cerebellum. However, when the simulation is first constructed, the limitations of the computational power available constrained the simulation to only contain 12000 granule cells, and greatly departed from the observed connectivity. These constraints leave the possibility that the emergent properties of the constrained simulation are not relevant to the computations performed in the cerebellum. The current work expands the simulation to over a million cells to approach the observed connectivity. This expanded simulation is used to further investigate possible mechanisms for transforming tonic mossy fiber inputs into stimulus-temporal code.

The expanded simulation incorporates over a million cells which represents a nearly 100 fold increase in the number of cells compared to the constrained simulation. The amount of computation required by this expanded simulation is proportionally increased by 100 fold. For the expanded simulation to

be used practically, its speed needs to be within the same order of magnitude as the constrained simulation. This is achieved using graphics processing units (GPUs) to handle the increased computational load and allow the expanded simulation to perform at the same speed as the constrained simulation. This allows for timely implementations of manipulations.

Using this expanded simulation, the question of how the cerebellar network can keep track of time is revisited. Specifically, the network interactions that can transform tonic mossy fiber inputs into stimulus-temporal code are investigated. The results from the expanded simulation suggest that the recurrent interactions between granule and Golgi cells (as the constrained simulation suggests) are effective only when the number of granule inputs per Golgi cells is small, beyond what the existing anatomical observations support.

In searching for alternate mechanisms that can transform tonic mossy fiber inputs to stimulus-temporal code, a newly discovered inhibition among Golgi cells (lateral Golgi inhibition) is shown by the expanded simulation to be a possible mechanism. This lateral recurrent interaction provide a mechanism for producing stimulus-temporal code in the Golgi cell population, which then induce stimulus-temporal code in the granule cell population. Using this mechanism, the expanded simulation is able to produce robust responses for long CS-US intervals beyond 750ms and up to 2000ms, which better reproduced animal behavior.

To understand the mechanisms of lateral Golgi inhibition, the Golgi network is dissected. Immediately after the onset of the CS input, a set of Golgi cells (early cells) responded by increasing their activity, and then decreasing activity during late period of the CS input. It is possible that the activity patterns of the early cells can delay the responses of other cells by providing inhibition and disinhibition. Since the early cells respond first to the CS input, disrupting their activity such that these cells remain elevated throughout the CS input should not allow other cells to become active. As a consequence, the stimulus-temporal code is disrupted. This is observed in the expanded simulation, where disrupting the early cell activity is more effective than disrupting a matching number of randomly selected Golgi cells. Using isolated Golgi network simulations, it is found that the cells that provide nonreciprocal inhibition to early cells are important to decrease early cell activity during late period of the CS input. The predictions from the isolated network results are tested by specifically removing the nonreciprocal inhibition to early cells. These tests reveal that removing nonreciprocal inhibition to early cells disrupt the stimulus-temporal code more effectively than removing a matching number of reciprocal inhibition to early cells. Finally, to examine if nonreciprocal inhibition is generally important for transforming tonic input into stimulus-temporal code, hypothetical single layer networks are constructed with different connectivity patterns of lateral inhibition. These networks show that while pure reciprocal inhibition can produce stimulus-

temporal code in a very limited range of inhibitory strength parameters, nonreciprocal inhibition allow the network to produce stimulus-temporal code that is robust in a wide range of inhibitory strength parameters.

The following sections discuss the limitations and significance of the results from this work. First, the relevance of the computational methodology to further scaling the simulation and the broader field of high performance computation is discussed. Second, the specific functional hypothesis regarding lateral Golgi inhibition is discussed in the context of the limitations of the approach of the simulation. Third, the connectivity constraints identified by this work is discussed in the context of the recurrent network field. Finally, the approach in this work that dissected and analyzed the lateral Golgi inhibition network is discussed in the context of recurrent networks and more generally, complex systems.

Implications for future scaling the simulated cerebellar network

The expanded simulation contains over a million cerebellar granule cells to approach the observed connectivity in the cerebellum. The significant increase in computational load is handled by utilizing graphics processing units (GPUs) to achieve 2x real-time performance for practical use. The number of cells represents roughly 1mm^2 of cerebellar cortex and is likely only a fraction of the cells in a microzone (hypothesized to be the functional units of the cerebellum) in

larger animals such as the cat and rabbit (Ito, 2000). However, the algorithms utilized in the simulation can scale further to construct simulations with more cells. The algorithms are not limited by the number of granule cells that can be represented, but rather the number of Golgi cells and mossy fibers that must be implemented to maintain a biologically relevant connectivity ratio. This is due to the limited size of the fast on-chip memory in the GPU, which is used by the simulation for updating the mossy fiber and Golgi cell inputs to granule cells. For example, the number of mossy fibers is scaled from 2048 to 8192 when the granule cells are increased from one million to four million, which requires four times the amount of on-chip memory. The latest GPU hardware constrains the number of mossy fibers and Golgi cells to 262,144 fibers and 32,768 cells respectively. These constraints allow for 128 million granule cells to be represented. Multiple GPUs are necessary for the performance of such a simulation to be practical. The current implementation provides the algorithms to utilize multiple GPUs.

The current implementation of the expanded simulation allow near linear performance scaling when using multiple GPUs. The speed of the expanded simulation doubles when using two GPUs compared to one GPU, and doubles again when using four GPUs compared to two GPUs. However, scaling to eight GPUs only achieves the same speed as four GPUs. On the other hand, when the simulation is expanded to four million cells, it is able to utilize eight GPUs to

achieve the same speed as the one million cells simulation with two GPUs. This is consistent with linear performance scaling with size. Given these scaling results, a system with eight newest generation GPUs (GTX Titan X) should be able to execute a simulation with 128 million granule cells at between 1/16 to 1/64 realtime. The cost of such a system is between $16,000 to $20,000 US dollars.

Given the current estimates, 128 million granule cells account for half of the rat cerebellum (~260 million granule cells) (Korbo, Andersen, Ladefoged, & Møller, 1993) and 1/20 of the cat cerebellum (~2 billion granule cells) (Palkovits et al., 1971b). A simulation of this size spans a significant portion of cerebellar cortex and across multiple functional areas in many species, and can be used as a tool to investigate coordinations between these areas. These investigations can illuminate the gaps in the knowledge about the cerebellar architecture and its computation properties in more complex tasks.

Utilizing graphics processing units for high performance computing

The algorithms utilized in this simulation demonstrate the feasibility of using massively parallel hardware to simulate neural networks. The cerebellum architecture provides advantages in that the granule cells do not interact with each other directly. If the granule cells directly interact with each other, especially over long distances, the number of granule cells that can be feasibly modeled

would be much smaller. However, even in that context utilizing the fast on-chip memory of the GPU can still yield performance increases compared to the traditional processors (CPUs).

The trend in modern supercomputers is to increasingly incorporate GPUs along with CPUs (www.top500.org). This trend suggests that the future of high performance computing is in a hybrid model with a few CPU cores to handle complex tasks and the GPUs (or other parallel co-processors) to handle simple tasks that are data heavy. The concepts implemented in this simulation utilize this hybrid system to achieve maximum performance, and are likely to be relevant for the foreseeable future in high performance computing.

Simulation predictions of the timing mechanism in the cerebellum

The results from the expanded simulation provide a hypothesis for the functional role of lateral Golgi inhibition in transforming tonic mossy fiber inputs to stimulus-temporal code to support well-timed learned motor responses. The hypothesis predicts that 1. electrophysiological recordings from Golgi cells *in vivo* during eyelid conditioning with mossy fiber stimulations should exhibit temporally varying activity, and that recordings from different Golgi cells should show different temporal patterns of activity, and 2. disabling the inhibitory conductance in Golgi cells should produce Golgi cell activity that is tonic, and animals should fail to learn well-timed responses in anticipation of US onset in delay eyelid

conditioning.

If disrupting the inhibitory conductance in Golgi cells do not result in tonic Golgi activity, then it is possible that the interactions between granule cells and Golgi cells are be important in producing a stimulus-temporal code. In this case, the results from the expanded simulation suggest that the number of granule inputs per Golgi cell must be small in order for this interaction to be effective. This can be examined with more detailed anatomical observations of the average number of Golgi cells that a granule cell outputs to. More directly, this can be examined by observing the number of granule inputs each Golgi cell receives. If the number of inputs is much lower than current estimates, then the functional hypothesis for lateral Golgi inhibition needs be reevaluated.

However, when evaluating the relevance of these predictions to the computation performed by the cerebellum, there are several limitations that must be considered. The nature of these limitations is not specific to this work, but to theories in general.

First, while the current understanding of cerebellar computation suggests the possibility that the granule cell population produces a stimulus-temporal code in response to a tonic mossy fiber input, direct experimental observations of granule cell activity is sparse (Jörntell & Ekerot, 2006). The existing evidence is insufficient to support or refute the hypothesis that the granule population actually produce stimulus-temporal code during delay eyelid conditioning. This is partly

because the difficulty in recording granule cell activity *in vivo*, due to their small size, dense packing, and lack of activity. The most promising technique to observe the granule population activity is likely *in vivo* calcium or voltage imaging which can observe many cells at once.

Second, the relevance of the simulation to the biological system remains a concern. While the expanded simulation has one million granule cells and so can approximate the known connectivity ratios within an order of magnitude, key connectivity parameters are not exactly characterized, such as the number of granule inputs per Golgi cell. The connectivity of granule cell output to basket and stellate cells (cells that receive granule inputs and inhibit Purkinje cells) has been observed (Eccles et al., 1967), but without sufficiently precise connectivity parameters. In addition to these connectivity, the uni-polar brush cells (Diño et al., 2000; Dino et al., 2000; DiÑO et al., 1999; Nunzi & Birnstiel, 2001) and Lugaro cells (Lainé & Axelrad, 2002; Melik-Musyan & Fanardzhyan, 2004) are not modeled in the simulation, due to a lack of data about their connectivity parameters and spatial distribution. It is possible that these cells play an important role in the recurrent interactions in the network to produce stimulus-temporal code.

The simulation is incomplete in capturing the known physiology of the cerebellum (Armano et al., 2000; D'Angelo, De Filippi, Rossi, & Taglietti, 1995; D'Angelo & De Zeeuw, 2009; Nieus et al., 2006; D. Watanabe & Nakanishi, 2003).

It has been shown that Golgi cells have electrical synapses (gap junctions) in the distal dendrites (S. J. Mitchell & Silver, 2003; Vervaeke et al., 2010, 2012). Models using these gap junctions suggest that these junctions play an inhibitory role by relaying the after-hyperpolarization of an action potential to the neighboring cell, which hyperpolarizes the neighbor (Vervaeke et al., 2010). However, in the current simulation, the after-hyperpolarization is not explicitly modeled, and would require a significant change in the method to represent neurons. In addition, the neurons in the current simulation are modeled as iso-potential point neurons, so cannot capture the spatial distribution of these gap junctions. It is possible that the interactions between the electrical coupling and inhibition among Golgi cells are important for generating stimulus-temporal code.

In addition to the limited Golgi cell physiology, the simulation only implements two sites of synaptic plasticity, at granule-Purkinje synapses and mossy fiber-deep cerebellar nucleus synapses. However, plasticity has been observed in almost all other synapses in the cerebellum (Hansel, Linden, & D'Angelo, 2001; Kenyon, 1997; Rancillac & Crépel, 2004; Robberechts, Wijnants, Giugliano, & De Schutter, 2010), such as mossy fiber to granule cell synapses (D'Angelo & De Zeeuw, 2009). These observations lack sufficiently detailed parameters for these synapses, therefore in the simulation they are modeled as non-plastic connections. It is possible that plasticity at these synapses can play a role in shaping the stimulus-temporal code that is not

captured in the simulation.

Third, the delay eyelid conditioning behavior that the simulation attempts to model is not consistent across species. The biological predictions from the expanded simulation are in the ability of the cerebellum to generate stimulus-temporal code for time scales beyond 500ms, which is fitting for modeling eyelid conditioning in the rabbit (White et al., 2000). However, the rabbit appears unique in its ability to perform eyelid conditioning at long intervals that are not observed in mice (Chettih et al., 2011). Of interest is that the data that the simulation uses is from multiple species: cat cerebellum for connectivity (Palkovits, Magyar, & Szentágothai, 1971a; Palkovits et al., 1971c, 1972), rat and mice cerebellum for physiology (Chadderton et al., 2004; V Chan-Palay & Palay, 1972; Galliano et al., 2010; Holtzman, Rajapaksa, Mostofi, & Edgley, 2006; Wang et al., 2000), and rabbit for delay eyelid conditioning behavior (Medina et al., 2000; Medina & Mauk, 2000). The inconsistency of delay eyelid conditioning among different species, and the fact that the simulation uses data across species, suggest that the simulation might not be closely relevant to any species.

The final concern regarding the simulation's relevance to the biological system is the approach of constructing the simulation. The detailed characterization of the cerebellar network connectivity and cellular physiology allows both the constrained and expanded simulation to be constructed with a bottom-up approach. This approach models the individual components of the

network by empirical approximations in order to study the emergent property of the network itself (Mauk, 2000). A fundamental aspect of this approach to maintain biological relevance depends on reliable empirical approximations of the underlying components. In this case, these approximations are of the activity of different cell types in the cerebellum. However, beyond Purkinje cells with relatively well characterized activity in various conditions *in vivo (Bell & Grimm, 1969; Gilbert & Thach, 1977; Jirenhed & Hesslow, 2011a; Rasmussen et al., 2008)*, the other neurons in the cerebellum remain to be characterized in more detail. One group has attempted to relate the recorded activity of various cell types in the cerebellum by juxtacellular labeling (Ruigrok, Hensbroek, & Simpson, 2011; Simpson, Hulscher, Sabel-Goedknegt, & Ruigrok, 2005). However, these recordings only provide a snapshot of the activity of these cells in a very restricted context, and at best provide a single data point for approximation. Given these limitations, it is possible that the empirical approximations of these cells in the simulation are not representative of the biological system, which can limit the biological relevance of the emergent properties observed in the simulation. On the other hand, it has been suggested that at least in certain networks, some of the emergent properties are robust for a range of parameter values for the underlying cells (Prinz, Bucher, & Marder, 2004).

The nature of the limitations discussed above are not unique to this work,

since simulations and models are necessarily simplified systems that typically do not capture all of the complexity in the biological system. However, the specific biological predictions derived from this particular simulation should be considered in the context of these limitations.

Network connectivity properties illustrated by the simulation

The constrained simulation of the cerebellum by Buonomano and Mauk is one of the early models in the field of theoretical Neuroscience that focused on the computational power and properties of recurrent neural networks. This field has emphasized the computation power of recurrent interactions in a network of neurons (Buonomano, 2005; Couey et al., 2013; Laje & Buonomano, 2013; Liu & Buonomano, 2009; Lukoševičius & Jaeger, 2009; Sussillo, 2014; Toyoizumi & Abbott, 2011; Wong & Wang, 2006), especially for generating complex temporal output. With recurrent interactions, each neuron in the network can produce a pattern of activity that is decorrelated from other neurons (Wiechert et al., 2010). When perturbed by a stimulus, the network responds by propagating the perturbation throughout the entire network through the recurrent connections (Maass, Natschläger, & Markram, 2002; Yamazaki & Tanaka, 2007). The interactions among neurons in the network through these connections transform the stimulus into complex patterns of activity for each neuron. A downstream neuron that receives input from all the neurons in the network can generate any

output activity pattern by adjusting the synaptic weight of individual inputs. It should be emphasized that for this to be possible, the patterns of activity among neurons in the network must be decorrelated (Laje & Buonomano, 2013).

Following the footsteps of the constrained simulation, the results from the expanded simulation further contribute to this field by specifying the constraints on the connectivity that can contribute useful recurrent interactions for generating stimulus-temporal code. The first constraint is that when the recurrent interaction in the network is strictly between two populations of cells (no lateral recurrence within each population), the recurrent connectivity is effective in generating stimulus-temporal code when the convergence ratio is low (i.e., each cell in either population only receives a few inputs from cells of the other population). This constraint is especially relevant when the size of the two populations are very different, such as Golgi cells and granule cells. In that case, the consequence of the constraint is that only a few cells in the larger population can participate in the recurrent interactions to produce stimulus-temporal code. The second constraint is that in networks with purely inhibitory lateral recurrent connectivity, the connectivity that are nonreciprocal (i.e., cell A inhibits cell B, but not vice versa) is important for the network to produce stimulus-temporal code. While the probability of reciprocal connectivity is low when there are no spatial constraints, in spatially constrained connectivity the probability of reciprocal connectivity is much higher. In these spatially constrained networks, the necessity of

nonreciprocal connectivity is relevant for constraining the sparsity of the network connectivity.

Decomposing the Golgi network and its relevance to complex systems

In analyzing the mechanisms of Golgi lateral inhibition that transformed tonic mossy fiber input into stimulus-temporal code, the early cells are found to be important. The network is manipulated by eliminating all inhibition to these cells, which disrupted the stimulus-temporal code. The inhibition to each early cell is then dissected in detail, and the nonreciprocal inhibition is found to be important in inhibiting early cells. Finally, the results from the dissection are found to be relevant to the full network by selectively eliminating the nonreciprocal inhibition to early cells in the intact simulation. It is conceivable that the principle of nonreciprocal connectivity can be discovered by directly manipulating the intact network to eliminate nonreciprocal inhibition and entirely avoid the detailed dissections of early cells. However, the detailed dissection approach itself has relevance to understanding recurrent neural networks and complex systems in general. In analyzing the components of the network, the effort to eliminate false positives in the analysis of the isolated networks provides an exercise in analyzing the deconstructed components of a complex system. A complex system (Barrat, Barthelemy, & Vespignani, 2008), by definition is a system that is difficult to understand by evaluating its underlying components. Complex

systems tend to have many components that have non-linear complex interactions with each other (Funtowicz & Ravetz, 1994). Decomposing the system into smaller systems can result in a drastic change in the interactions due to the change in connectivity, such that the behavior of the smaller system does not provide insight to the full system (Barrat et al., 2008). Such complex systems are common in many areas of biology, and understanding their mechanisms presents significant challenges due to the interactions among the components. The approach used in this work to simulate isolated Golgi cell networks while eliminating false positives provides an insightful exercise in considering how to decompose a complex system so that the resulting components can still be relevant to the full system. The fact that the results from the isolated networks are relevant to the full network is an unexpected surprise, given the number of other recurrent connections that can potentially compensate for the removal of a few connections. The result that removing a few specific connections within a large and relatively complex network can indeed disrupt the emergent stimulus-temporal code suggests that the result from the decomposed networks is indeed relevant to the emergent properties of the intact system.

It should be noted however, that the current approach is able to identify the components of the connectivity that are specifically important, but does not provide a complete answer to the emergent mechanisms of the network. In addition, this approach to identify the network components likely will be most

insightful for network connectivity that does not have any obviously important components. For example, this approach would be unnecessary for a network whose neurons are strictly connected in a chain. In that case, disrupting any link in the chain would disrupt the behavior of the network. Finally, the approach in this work is specific to analyzing this simulation network, and benefitted from well-defined components (focusing specifically on early Golgi cells) and well defined questions (sources of inhibition that can decrease early Golgi cell's activity). The generality of this approach appears promising but remains to be tested in other complex systems.

REFERENCES

Aitkin, L. M., & Boyd, J. (1978). Acoustic input to the lateral pontine nuclei. *Hearing Research*, *1*(1), 67–77.

Albus, J. (1975). A New Approach to Manipulator Control: The Cerebellar Model Articulation Controller (CMAC). *Journal of Dynamic Systems, Measurement, and Control*. Retrieved from http://dynamicsystems.asmedigitalcollection.asme.org/article.aspx?articleid=1402197

Andersson, G., & Armstrong, D. M. (1987). Complex spikes in Purkinje cells in the lateral vermis (b zone) of the cat cerebellum during locomotion. *The Journal of Physiology*, *385*(1), 107–134. http://doi.org/10.1113/jphysiol.1987.sp016487

Arezzo, J., & Vaughan, H. (1975). Cortical potentials associated with voluntary movements in the monkey. *Brain Research*, *88*, 99–104. Retrieved from http://www.sciencedirect.com/science/article/pii/0006899375909543

Armano, S., Rossi, P., Taglietti, V., & D'Angelo, E. (2000). Long-term potentiation of intrinsic excitability at the mossy fiber-granule cell synapse of rat cerebellum. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *20*(14), 5208–16. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10884304

Asanuma, C., Thach, W. T., & Jones, E. G. (1983). Brainstem and spinal projections of the deep cerebellar nuclei in the monkey, with observations on the brainstem projections of the dorsal column nuclei. *Brain Research Reviews*, *5*(3), 299–322. http://doi.org/10.1016/0165-0173(83)90017-6

Aviel, Y., Mehring, C., Abeles, M., & Horn, D. (2003). On embedding synfire chains in a balanced network. *Neural Computation*, *1340*(1991), 1321–1340. Retrieved from http://www.mitpressjournals.org/doi/abs/10.1162/089976603321780290

Balaban, C., Schuerger, R., & Porter, J. (2000). Zonal organization of flocculo–vestibular connections in rats. *Neuroscience*, *99*(4), 669–682. Retrieved from http://www.sciencedirect.com/science/article/pii/S0306452200002323

Barrat, A., Barthelemy, M., & Vespignani, A. (2008). *Dynamical Processes on Complex Networks* (p. 361). Cambridge: Cambridge University Press.

Bastian, A. J. (2006). Learning to predict the future: the cerebellum adapts feedforward movement control. *Current Opinion in Neurobiology*, *16*(6), 645–9. http://doi.org/10.1016/j.conb.2006.08.016

Bastian, A. J., Zackowski, K. M., & Thach, W. T. (2000). Cerebellar ataxia: torque deficiency or torque mismatch between joints? *Journal of Neurophysiology*, *83*(5), 3019–30.

Bell, C. C., & Grimm, R. J. (1969). Discharge properties of Purkinje cells recorded on single and double microelectrodes. *Journal of Neurophysiology*, *32*(6), 1044–55. Retrieved from http://jn.physiology.org/content/32/6/1044.short

Berthier, N. E., & Moore, J. W. (1986). Cerebellar Purkinje cell activity related to the classically conditioned nictitating membrane response. *Experimental Brain Research*, *63*(2). http://doi.org/10.1007/BF00236851

Best, A. R., & Regehr, W. G. (2009). Inhibitory regulation of electrically coupled neurons in the inferior olive is mediated by asynchronous release of GABA. *Neuron*, *62*(4), 555–65. http://doi.org/10.1016/j.neuron.2009.04.018

Boyden, E. S., Katoh, A., & Raymond, J. L. (2004). Cerebellum-dependent learning: the role of multiple plasticity mechanisms. *Annual Review of Neuroscience*, *27*, 581–609. http://doi.org/10.1146/annurev.neuro.27.070203.144238

Brand, S., Dahl, A. L., & Mugnaini, E. (1976). The length of parallel fibers in the cat cerebellar cortex. An experimental light and electron microscopic study. *Experimental Brain Research*, *26*(1), 39–58. http://doi.org/10.1007/BF00235248

Brodal, P., Dietrichs, E., & Walberg, F. (1986). Do pontocerebellar mossy fibres give off collaterals to the cerebellar nuclei? An experimental study in the cat with implantation of crystalline HRP-WGA. *Neuroscience Research*, *4*(1), 12–24. http://doi.org/10.1016/S0921-8696(86)80058-5

Bullock, D., Fiala, J. C., & Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Networks : The Official Journal of the International Neural Network Society*, *7*(94), 1101–1114.

Bullock, D., & Grossberg, S. (1988). Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. *Psychological Review*, *95*(1), 49–90. Retrieved from

http://psycnet.apa.org/journals/rev/95/1/49/

Buonomano, D. V. (2005). A learning rule for the emergence of stable dynamics and timing in recurrent networks. *Journal of Neurophysiology*, 2275–2283. http://doi.org/10.1152/jn.01250.2004.

Buonomano, D. V, & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews. Neuroscience*, *10*(2), 113–25. http://doi.org/10.1038/nrn2558

Buonomano, D. V, & Mauk, M. D. (1994). Neural Network Model of the Cerebellum: Temporal Discrimination and the Timing of Motor Responses. *Neural Computation*, *6*(1), 38–55. http://doi.org/10.1162/neco.1994.6.1.38

Buonomano, D. V, & Merzenich, M. (1995). Temporal Information Transformed into a Spatial Code by a Neural Network with Realistic Properties. *Science*, *267*(February), 1028–1030. Retrieved from https://ftp.utdallas.edu/~kilgard/11c BuonomanoMerzenich_Science1995.pdf

Buonomano, D. V, & Merzenich, M. (1999). A neural network model of temporal code generation and position-invariant pattern recognition. *Neural Computation*, *11*(1), 103–16. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9950725

Carey, M., & Lisberger, S. (2002). Embarrassed , but Not Depressed: Eye Opening Lessons for Cerebellar Learning. *Neuron*, *35*, 223–226. Retrieved from http://www.sciencedirect.com/science/article/pii/S0896627302007717

Chadderton, P., Margrie, T. W., & Häusser, M. (2004). Integration of quanta in cerebellar granule cells during sensory processing. *Nature*, *428*(6985), 856–860. http://doi.org/10.1038/nature02467.Published

Chan-Palay, V., & Palay, S. L. (1971). The synapse en marron between Golgi II neurons and mossy fibers in the rat's cerebellar cortex. *Zeitschrift Für Anatomie Und Entwicklungsgeschichte*, *133*(3), 274–287. http://doi.org/10.1007/BF00519303

Chan-Palay, V., & Palay, S. L. (1972). The stellate cells of the rat's cerebellar cortex. *Zeitschrift Für Anatomie Und Entwicklungsgeschichte*, *136*(2), 224–48. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/15789412

Chapeau-Blondeau, F., & Chauvet, G. (1991). A neural network model of the cerebellar cortex performing dynamic associations. *Biological Cybernetics*, *65*(4), 267–279.

http://doi.org/10.1007/BF00206224

Chettih, S. N., McDougle, S. D., Ruffolo, L. I., & Medina, J. F. (2011). Adaptive timing of motor output in the mouse: the role of movement oscillations in eyelid conditioning. *Frontiers in Integrative Neuroscience*, *5*(November), 72. http://doi.org/10.3389/fnint.2011.00072

Contreras-Vidal, J. L., Grossberg, S., & Bullock, D. (1997). A neural model of cerebellar learning for arm movement control: cortico-spino-cerebellar dynamics. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *3*, 475–502. Retrieved from http://learnmem.cshlp.org/content/3/6/475.short

Couey, J. J., Witoelar, A., Zhang, S.-J., Zheng, K., Ye, J., Dunn, B., … Witter, M. P. (2013). Recurrent inhibitory circuitry as a mechanism for grid formation. *Nature Neuroscience*, *16*(3), 318–24. http://doi.org/10.1038/nn.3310

D'Angelo, E., De Filippi, G., Rossi, P., & Taglietti, V. (1995). Synaptic excitation of individual rat cerebellar granule cells in situ: evidence for the role of NMDA receptors. *The Journal of Physiology*, *484 (Pt 2*(1995), 397–413. Retrieved from http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1157902&tool=pmcentrez&rendertype=abstract

D'Angelo, E., & De Zeeuw, C. I. (2009). Timing and plasticity in the cerebellum: focus on the granular layer. *Trends in Neurosciences*, *32*(1), 30–40. http://doi.org/10.1016/j.tins.2008.09.007

Davey, N., & Romaiguere, P. (1994). Suppression of voluntary motor activity revealed using transcranial magnetic stimulation of the motor cortex in man. *Journal of Physiology*, 223–235. Retrieved from http://onlinelibrary.wiley.com/doi/10.1113/jphysiol.1994.sp020186/abstract

De Schutter, E., & Bjaalie, J. G. (2001). Coding in the granular layer of the cerebellum. *Progress in Brain Research*, *130*, 279–96. http://doi.org/10.1016/S0079-6123(01)30019-5

Deecke, L., Scheid, P., & Kornhuber, H. (1969). Distribution of readiness potential, pre-motion positivity, and motor potential of the human cerebral cortex preceding voluntary finger movements. *Experimental Brain Research*, *168*, 158–168. Retrieved from http://www.springerlink.com/index/8D1692D0EB3C6074.pdf

Desclin, J. C. (1974). Histological evidence supporting the inferior olive as the major source of cerebellar climbing fibers in the rat. *Brain Research*, *77*(3), 365–384. http://doi.org/10.1016/0006-8993(74)90628-3

Diño, M., Nunzi, M., Anelli, R., & Mugnaini, E. (2000). Unipolar brush cells of the vestibulocerebellum : afferents and targets *. *Progress in Brain Research*, *124*. Retrieved from http://www.sciencedirect.com/science/article/pii/S0079612300240132/pdf?md5=7e160d78425063472d1a7ce6c717d26a&pid=1-s2.0-S0079612300240132-main.pdf&_valck=1

Dino, M., Schuerger, R., Liu, Y., Slater, N., & Mugnaini, E. (2000). Unipolar brush cell: a potential feedforward excitatory interneuron of the cerebellum. *Neuroscience*, *98*(4), 625–636. Retrieved from http://www.sciencedirect.com/science/article/pii/S0306452200001238

DiÑO, M., Willard, F., & Mugnaini, E. (1999). Distribution of unipolar brush cells and other calretinin immunoreactive components in the mammalian cerebellar cortex. *Journal of Neurocytology*, *123*, 99–123. Retrieved from http://link.springer.com/article/10.1023/A:1007072105919

DuLac, S., Raymond, J. L., Sejnowski, T. J., & Lisberger, S. G. (1995). Learning and Memory in the Vestibulo-Ocular Reflex. *Annual Review of Neuroscience*, *18*, 409–441.

Dum, R. P., & Strick, P. L. (2003). An unfolded map of the cerebellar dentate nucleus and its projections to the cerebral cortex. *Journal of Neurophysiology*, *89*(1), 634–9. http://doi.org/10.1152/jn.00626.2002

Eccles, J. C., Ito, M., & Szentágothai, J. (1967). *The cerebellum as a Neuronal Machine* (p. 335). New York: Springer-Verlag.

Eccles, J. C., Llinas, R., & Sasaki, K. (1966). The excitatory synaptic action of climbing fibres on the Purkinje cells of the cerebellum. *The Journal of Physiology*, 268–296. Retrieved from http://jp.physoc.org/content/182/2/268.short

Eccles, J. C., Llinás, R., & Sasaki, K. (1966). The mossy fibre-granule cell relay of the cerebellum and its inhibitory control by Golgi cells. *Experimental Brain Research*, *1*(1), 82–101. http://doi.org/10.1007/BF00235211

Flament, D., Vilis, T., & Hore, J. (1984). Dependence of cerebellar tremor on proprioceptive but not visual feedback. *Experimental Neurology*, *325*, 314–325. Retrieved from http://www.sciencedirect.com/science/article/pii/0014488684902280

Flanagan, J., & Wing, A. (1993). Modulation of grip force with load force durnig point-to-point arm movements. *Experimental Brain Research*, 131–143. Retrieved from http://link.springer.com/article/10.1007/BF00229662

Flumerfelt, B. A., Otabe, S., & Courville, J. (1973). Distinct projections to the red nucleus from the dentate and interposed nuclei in the monkey. *Brain Research*, *50*(2), 408–414. http://doi.org/10.1016/0006-8993(73)90742-7

Fox, C. A., & Barnard, J. W. (1957). A quantitative study of the Purkinje cell dendritic branchlets and their relationship to afferent fibres. *Journal of Anatomy*, *91*(3), 299–313. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1244915

Freeman, J. A., & Nicholson, C. N. (1970). Space-Time Transformation in the Frog Cerebellum through an Intrinsic Tapped Delay-line. *Nature*, *226*, 640–642. Retrieved from http://www.nature.com/nature/journal/v226/n5246/abs/226640a0.html

Fuchs, A. F., & Kornhuber, H. H. (1969). Extraocular muscle afferents to the cerebellum of the cat. *The Journal of Physiology*, *200*(3), 713–722. http://doi.org/10.1113/jphysiol.1969.sp008718

Fujita, H., & Sugihara, I. (2013). Branching patterns of olivocerebellar axons in relation to the compartmental organization of the cerebellum. *Frontiers in Neural Circuits*, *7*(February), 3. http://doi.org/10.3389/fncir.2013.00003

Fujita, M. (1982). Adaptive filter model of the cerebellum. *Biological Cybernetics*, *206*, 195–206.

Funtowicz, S., & Ravetz, J. R. (1994). Emergent complex systems. *Futures*, *26*(6), 568–582. http://doi.org/10.1016/0016-3287(94)90029-9

Gabbiani, F., Midtgaard, J., & Knöpfel, T. (1994). Synaptic integration in a model of cerebellar granule cells. *Journal of Neurophysiology*, *72*(2), 999–1009. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7527078

Galliano, E., Mazzarello, P., & D'Angelo, E. (2010). Discovery and rediscoveries of Golgi cells. *The Journal of Physiology*, *588*(Pt 19), 3639–55.

http://doi.org/10.1113/jphysiol.2010.189605

Garcia, K. S., Steele, P. M., & Mauk, M. D. (1999). Cerebellar cortex lesions prevent acquisition of conditioned eyelid responses. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *19*(24), 10940–7.

Gellman, R., Gibson, A. R., & Houk, J. C. (1985). Inferior olivary neurons in the awake cat: detection of contact and passive body displacement. *Journal of Neurophysiology*, *54*(1), 40–60. Retrieved from http://jn.physiology.org/content/54/1/40.short

Gibson, A. R., Robinson, F. R., Alam, J., & Houk, J. C. (1987). Somatotopic alignment between climbing fiber input and nuclear output of the cat intermediate cerebellum. *The Journal of Comparative Neurology*, *260*(3), 362–77. http://doi.org/10.1002/cne.902600304

Gilbert, P. F. C., & Thach, W. T. (1977). Purkinje Cell Activity During Motor Learning. *Brain Research*, *128*, 309–328.

Goudar, V., & Buonomano, D. V. (2014). Useful dynamic regimes emerge in recurrent networks. *Nature Neuroscience*, *17*(4), 487–9. http://doi.org/10.1038/nn.3679

Gould, T. J., Sears, L. L., & Steinmetz, J. E. (1993). Possible CS and US pathways for rabbit classical eyelid conditioning: Electrophysiological evidence for projections from the pontine nuclei and inferior olive to cerebellar cortex and nuclei. *Behavioral and Neural Biology*, *60*(2), 172–185. http://doi.org/10.1016/0163-1047(93)90285-P

Grossberg, S., & Schmajuk, N. (1989). Neural dynamics of adaptive timing and temporal discrimination during associative learning. *Neural Networks*, *2*, 79–102. Retrieved from http://www.sciencedirect.com/science/article/pii/0893608089900269

Hámori, J., & Szentágothai, J. (1966). Participation of Golgi neuron processes in the cerebellar glomeruli: An electron microscope study. *Experimental Brain Research*, *2*(1), 35–48. http://doi.org/10.1007/BF00234359

Hansel, C., Linden, D. J., & D'Angelo, E. (2001). Beyond parallel fiber LTD: the diversity of synaptic and non-synaptic plasticity in the cerebellum. *Nature Neuroscience*, *4*(5), 467–75. http://doi.org/10.1038/87419

Harrington, D. L., Lee, R. R., Boyd, L. a, Rapcsak, S. Z., & Knight, R. T. (2004). Does the representation of time depend on the cerebellum? Effect of cerebellar stroke.

Brain : A Journal of Neurology, *127*(Pt 3), 561–74. http://doi.org/10.1093/brain/awh065

Heiney, S. a, Kim, J., Augustine, G. J., & Medina, J. F. (2014). Precise control of movement kinematics by optogenetic inhibition of Purkinje cell activity. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(6), 2321–30. http://doi.org/10.1523/JNEUROSCI.4547-13.2014

Hesslow, G., Svensson, P., & Ivarsson, M. (1999). Learned movements elicited by direct stimulation of cerebellar mossy fiber afferents. *Neuron*, *24*(1), 179–85. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10677036

Holtzman, T., Rajapaksa, T., Mostofi, A., & Edgley, S. a. (2006). Different responses of rat cerebellar Purkinje cells and Golgi cells evoked by widespread convergent sensory inputs. *The Journal of Physiology*, *574*(Pt 2), 491–507. http://doi.org/10.1113/jphysiol.2006.108282

Hosaka, R., Araki, O., & Ikeguchi, T. (2008). STDP provides the substrate for igniting synfire chains by spatiotemporal input patterns. *Neural Computation*, 1–33. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6795607

Hull, C., & Regehr, W. G. (2012). Identification of an Inhibitory Circuit that Regulates Cerebellar Golgi Cell Activity. *Neuron*, *73*(1), 149–158. http://doi.org/10.1016/j.neuron.2011.10.030

Ito, M. (1982). Cerebellar Control of the Vestibulo-Ocular Reflex-Around The Flocculus Hypothesis. *Annual Review of Neuroscience*. Retrieved from http://www.annualreviews.org/doi/pdf/10.1146/annurev.ne.05.030182.001423

Ito, M. (1984). *The Cerebellum and Neural Control* (p. 580). New York: Raven Press.

Ito, M. (2000). Mechanisms of motor learning in the cerebellum. *Brain Research*, *886*(1-2), 237–245. http://doi.org/10.1016/S0006-8993(00)03142-5

Ito, M. (2001). Cerebellar long-term depression: characterization, signal transduction, and functional roles. *Physiological Reviews*, *81*(3). Retrieved from http://physrev.physiology.org/content/81/3/1143.short

Ito, M. (2005). Bases and implications of learning in the cerebellum--adaptive control and internal model mechanism. *Progress in Brain Research*, *148*, 95–109. http://doi.org/10.1016/S0079-6123(04)48009-1

Ito, M. (2006a). Cerebellar circuitry as a neuronal machine. *Progress in Neurobiology*, *78*(3-5), 272–303. http://doi.org/10.1016/j.pneurobio.2006.02.006

Ito, M. (2006b). Cerebellar circuitry as a neuronal machine. *Progress in Neurobiology*, *78*(3-5), 272–303. http://doi.org/10.1016/j.pneurobio.2006.02.006

Ito, M., & Kano, M. (1982). Long-lasting depression of parallel fiber-Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. *Neuroscience Letters*, *33*, 253–258.

Ito, M., & Simpson, J. I. (1971). Discharges in Purkinje cell axons during climbing fiber activation. *Brain Research*, *31*(1), 215–219. http://doi.org/10.1016/0006-8993(71)90648-2

Ito, M., Yoshida, M., Obata, K., Kawai, N., & Udo, M. (1970). Inhibitory control of intracerebellar nuclei by the Purkinje cell axons. *Experimental Brain Research*, *10*(1), 64–80. http://doi.org/10.1007/BF00340519

Ivry, R., Spencer, R., Zelaznik, H. N., & Diedrichsen, J. (2002). The Cerebellum and Event Timing. *Annals of the New York Academy of Sciences*, *5*. Retrieved from http://onlinelibrary.wiley.com/doi/10.1111/j.1749-6632.2002.tb07576.x/full

Jahnsen, H. (1986). Extracellular activation and membrane conductances of neurones in the guinea-pig deep cerebellar nuclei in vitro. *The Journal of Physiology*, *372*(1), 149–168. http://doi.org/10.1113/jphysiol.1986.sp016002

Jakab, R. L., & Hámori, J. (1988). Quantitative morphology and synaptology of cerebellar glomeruli in the rat. *Anatomy and Embryology*, *179*(1), 81–88. http://doi.org/10.1007/BF00305102

Ji, Z., & Hawkes, R. (1994). TOPOGRAPHY OF PURKINJE CELL COMPARTMENTS AND MOSSY FIBER TERMINAL FIELDS IN LOBULES II AND III OF THE RAT CEREBELLAR CORTEX : SPINOCEREBELLAR AND CUNEOCEREBELLAR PROJECTIONS. *Neuroscience*, *61*(4), 935–954. Retrieved from http://www.sciencedirect.com/science/article/pii/0306452294904146

Jirenhed, D.-A., & Hesslow, G. (2011a). Learning stimulus intervals--adaptive timing of conditioned Purkinje cell responses. *Cerebellum (London, England)*, *10*(3), 523–35. http://doi.org/10.1007/s12311-011-0264-3

Jirenhed, D.-A., & Hesslow, G. (2011b). Time course of classically conditioned Purkinje

cell response is determined by initial part of conditioned stimulus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(25), 9070–4. http://doi.org/10.1523/JNEUROSCI.1653-11.2011

Jörntell, H., & Ekerot, C.-F. (2006). Properties of somatosensory synaptic integration in cerebellar granule cells in vivo. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(45), 11786–97. http://doi.org/10.1523/JNEUROSCI.2939-06.2006

Jörntell, H., & Hansel, C. (2006). Synaptic memories upside down: bidirectional plasticity at cerebellar parallel fiber-Purkinje cell synapses. *Neuron*, *52*(2), 227–38. http://doi.org/10.1016/j.neuron.2006.09.032

Kalmbach, B. E., Davis, T., Ohyama, T., Riusech, F. A., Nores, W. L., & Mauk, M. D. (2010). Cerebellar cortex contributions to the expression and timing of conditioned eyelid responses. *Journal of Neurophysiology*, *103*(4), 2039–49. http://doi.org/10.1152/jn.00033.2010

Kalmbach, B. E., Ohyama, T., Kreider, J. C., Riusech, F. A., & Mauk, M. D. (2009). Interactions between prefrontal cortex and cerebellum revealed by trace eyelid conditioning. *Learning & Memory (Cold Spring Harbor, N.Y.)*, *16*(1), 86–95. http://doi.org/10.1101/lm.1178309

Kalmbach, B. E., Voicu, H., Ohyama, T., & Mauk, M. D. (2011). A subtraction mechanism of temporal coding in cerebellar cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(6), 2025–34. http://doi.org/10.1523/JNEUROSCI.4212-10.2011

Kanichay, R. T., & Silver, R. A. (2008). Synaptic and cellular properties of the feedforward inhibitory circuit within the input layer of the cerebellar cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(36), 8955–67. http://doi.org/10.1523/JNEUROSCI.5469-07.2008

Karmarkar, U. R., & Buonomano, D. V. (2007). Timing in the absence of clocks: encoding time in neural network states. *Neuron*, *53*(3), 427–38. http://doi.org/10.1016/j.neuron.2007.01.006

Kehoe, E. J., & Holt, P. E. (1984). Transfer across CS-US intervals and sensory modalities in classical conditioning of the rabbit. *Animal Learning & Behavior*, *12*(2), 122–128. http://doi.org/10.3758/BF03213130

Kelly, R., & Strick, P. (2003). Cerebellar Loops with Motor Cortex and Prefrontal Cortex of a Nonhuman Primate. *The Journal of Neuroscience*, *23*(23), 8432–8444. Retrieved from http://www.jneurosci.org/content/23/23/8432.short

Kenyon, G. (1997). A model of long-term memory storage in the cerebellar cortex : A possible role for plasticity at parallel fiber synapses onto stellate basket interneurons. *Proceedings of the National Academy of …*, *94*(December), 14200–14205. Retrieved from http://www.pnas.org/content/94/25/14200.short

Korbo, L., Andersen, B. B., Ladefoged, O., & Møller, A. (1993). Total numbers of various cell types in rat cerebellar cortex estimated using an unbiased stereological method. *Brain Research*, *609*(1-2), 262–268. http://doi.org/10.1016/0006-8993(93)90881-M

Kreider, J. C., & Mauk, M. D. (2010). Eyelid conditioning to a target amplitude: adding how much to whether and when. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *30*(42), 14145–52. http://doi.org/10.1523/JNEUROSCI.3473-10.2010

Lainé, J., & Axelrad, H. (1996). Morphology of the Golgi-impregnated Lugaro cell in the rat cerebellar cortex: a reappraisal with a description of its axon. *The Journal of Comparative Neurology*, *375*(4), 618–40. http://doi.org/10.1002/(SICI)1096-9861(19961125)375:4<618::AID-CNE5>3.0.CO;2-4

Lainé, J., & Axelrad, H. (1998). Lugaro cells target basket and stellate cells in the cerebellar cortex. *Neuroreport*, *9*(10), 2399–403. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9694235

Lainé, J., & Axelrad, H. (2002). Extending the cerebellar Lugaro cell class. *Neuroscience*, *115*(2), 363–374. Retrieved from http://www.sciencedirect.com/science/article/pii/S0306452202004219

Laje, R., & Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nature Neuroscience*, (April). http://doi.org/10.1038/nn.3405

Lang, E. J., Sugihara, I., & Llinás, R. (1996). GABAergic modulation of complex spike activity by the cerebellar nucleoolivary pathway in rat. *Journal of Neurophysiology*, *76*(1), 255–75. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8836223

Lange, W. (1975). Cell Number and Cell Density in the Cerebellar Cortex of Man and Some Other Mammals. *Cell and Tissue Research*, *124*, 115–124. Retrieved from http://link.springer.com/article/10.1007/BF00223234

Lavond, D. G., Hembree, T. L., & Thompson, R. F. (1985). Effect of kainic acid lesions of the cerebellar interpositus nucleus on eyelid conditioning in the rabbit. *Brain Research*, *326*(1), 179–182. http://doi.org/10.1016/0006-8993(85)91400-3

Lee, S., Neiman, A., & Kim, S. (1998). Coherence resonance in a Hodgkin-Huxley neuron. *Physical Review E*, *57*(3), 3292–3297. Retrieved from http://journals.aps.org/pre/abstract/10.1103/PhysRevE.57.3292

Lev-Ram, V., Mehta, S. B., Kleinfeld, D., & Tsien, R. Y. (2003). Reversing cerebellar long-term depression. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(26), 15989–93. http://doi.org/10.1073/pnas.2636935100

Li, J. X., & Lisberger, S. G. (2011). Learned timing of motor behavior in the smooth eye movement region of the frontal eye fields. *Neuron*, *69*(1), 159–69. http://doi.org/10.1016/j.neuron.2010.11.043

Lisberger, S., & Fuchs, A. F. (1978). Role of Primate Flocculus During Rapid Behavioral . Modification of Vestibuloocular Reflex . I . Purkinje Cell Activity During Visually Guided Horizontal Smooth-Pursuit Eye Movements and Passive Head Rotation. *J Neurophysiol*, *41*(3). Retrieved from http://jn.physiology.org/content/jn/41/3/733.full.pdf

Liu, J. K., & Buonomano, D. V. (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(42), 13172–81. http://doi.org/10.1523/JNEUROSCI.2358-09.2009

Llinas, R., Lang, E. J., & Welsh, J. P. (1997). The cerebellum, LTD, and memory: alternative views. *Learning & Memory*, *3*(6), 445–455. http://doi.org/10.1101/lm.3.6.445

Llinás, R., & Mühlethaler, M. (1988). Electrophysiology of guinea-pig cerebellar nuclear cells in the in vitro brain stem-cerebellar preparation. *The Journal of Physiology*, *404*(1), 241–258. http://doi.org/10.1113/jphysiol.1988.sp017288

Llinás, R., & Welsh, J. P. (1993). On the cerebellum and motor learning. *Current Opinion in Neurobiology*, 958–965. Retrieved from http://www.sciencedirect.com/science/article/pii/095943889390168X

Lukoševičius, M., & Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, *3*(3), 127–149. http://doi.org/10.1016/j.cosrev.2009.03.005

Maass, W., Joshi, P., & Sontag, E. (2007). Computational aspects of feedback in neural circuits. *PLOS Computational Biology*, *3*(1), e165. http://doi.org/10.1371/journal.pcbi.0020165

Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, *2560*, 2531–2560. Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6789852

MacKay, W., & Murphy, J. (1979). Cerebellar modulation of reflex gain. *Progress in Neurobiology*, *13*. Retrieved from http://www.sciencedirect.com/science/article/pii/0301008279900042

Maekawa, K., & Takeda, T. (1975). Mossy fiber responses evoked in the cerebellar flocculus of rabbits by stimulation of the optic pathway. *Brain Research*, *98*, 590–595. Retrieved from http://www.sciencedirect.com/science/article/pii/0006899375903765

Maekawa, K., & Takeda, T. (1976). Electrophysiological identification of the climbing and mossy fiber pathways from the rabbit's retina to the contralateral cerebellar flocculus. *Brain Research*, *109*(1), 169–174. http://doi.org/10.1016/0006-8993(76)90388-7

Manto, M., Bower, J. M., Conforto, A. B., Delgado-García, J. M., da Guarda, S. N. F., Gerwig, M., … Timmann, D. (2012). Consensus paper: roles of the cerebellum in motor control--the diversity of ideas on cerebellar involvement in movement. *Cerebellum (London, England)*, *11*(2), 457–87. http://doi.org/10.1007/s12311-011-0331-9

Mapelli, J., & D'Angelo, E. (2007). The spatial organization of long-term synaptic plasticity at the input stage of cerebellum. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(6), 1285–96.

http://doi.org/10.1523/JNEUROSCI.4873-06.2007

Marr, D. (1969). A theory of cerebellar cortex. *The Journal of Physiology*, *202*(2), 437–70.

Mauk, M. D. (2000). The potential effectiveness of simulations versus phenomenological models. *Nature Neuroscience*, *3*(7), 649–51. http://doi.org/10.1038/76606

Mauk, M. D., Medina, J. F., Nores, W. L., & Ohyama, T. (2000). Cerebellar function: Coordination, learning or timing? *Current Biology*, *10*(14), R522–R525. http://doi.org/10.1016/S0960-9822(00)00584-4

Mauk, M. D., Steinmetz, J. E., & Thompson, R. F. (1986). Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proceedings of the National Academy of Sciences of the United States of America*, *83*, 5349–5353.

Mauk, M. D., & Thompson, R. F. (1987). Retention of classically conditioned eyelid responses following acute decerebration. *Brain Research*, *403*(1), 89–95. http://doi.org/10.1016/0006-8993(87)90126-0

McCormick, D. A., & Thompson, R. F. (1984). Cerebellum: essential involvement in the classically conditioned eyelid response. *Science*, *223*, 296–299.

Medina, J. F., Garcia, K. S., & Mauk, M. D. (2001). A mechanism for savings in the cerebellum. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *21*(11), 4081–9.

Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M., & Mauk, M. D. (2000). Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *20*(14), 5516–25.

Medina, J. F., & Lisberger, S. G. (2007). Variation, signal, and noise in cerebellar sensory-motor processing for smooth-pursuit eye movements. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *27*(25), 6832–42. http://doi.org/10.1523/JNEUROSCI.1323-07.2007

Medina, J. F., & Mauk, M. D. (1999). Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber to deep nucleus synapse. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *19*(16), 7140–51.

Medina, J. F., & Mauk, M. D. (2000). Computer simulation of cerebellar information processing. *Nature Neuroscience*, *3 Suppl*(Box 1), 1205–11. http://doi.org/10.1038/81486

Medina, J. F., Nores, W. L., & Mauk, M. D. (2002). Inhibition of climbing fibres is a signal for the extinction of conditioned eyelid responses. *Nature*, *416*, 330–333.

Melik-Musyan, A., & Fanardzhyan, V. (2004). Morphological characteristics of Lugaro cells in the cerebellar cortex. *Neuroscience and Behavioral …*, *34*(6). Retrieved from http://link.springer.com/article/10.1023/B:NEAB.0000028297.30474.f9

Mertz, K., Koscheck, T., & Schilling, K. (2000). Brain-derived neurotrophic factor modulates dendritic morphology of cerebellar basket and stellate cells: an in vitro study. *Neuroscience*, *97*(2), 303–310. http://doi.org/10.1016/S0306-4522(99)00585-0

Middleton, F. A., & Strick, P. L. (1998). Cerebellar output: motor and cognitive channels. *Trends in Cognitive Sciences*, *2*(9), 348–354. http://doi.org/10.1016/S1364-6613(98)01220-0

Midtgaard, J. (1992). Stellate cell inhibition of Purkinje cells in the turtle cerebellum in vitro. *The Journal of Physiology*, 355–367. Retrieved from http://jp.physoc.org/content/457/1/355.short

Miles, F., & Lisberger, S. (1981). Plasticity in the Vestibulo-Ocular Reflex: a New Hypothesis. *Annual Review of Neuroscience*. Retrieved from http://www.annualreviews.org/doi/pdf/10.1146/annurev.ne.04.030181.001421

Miller, P. (2003). A Recurrent Network Model of Somatosensory Parametric Working Memory in the Prefrontal Cortex. *Cerebral Cortex*, *13*(11), 1208–1218. http://doi.org/10.1093/cercor/bhg101

Mitchell, S. J., & Silver, R. A. (2003). Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron*, *38*(3), 433–45.

Mitchell, S., & Silver, R. (2000). GABA spillover from single inhibitory axons suppresses low-frequency excitatory transmission at the cerebellar glomerulus. *The Journal of Neuroscience*, *20*(23), 8651–8658. Retrieved from http://www.jneurosci.org/content/20/23/8651.short

Morton, S. M., & Bastian, A. J. (2004). Cerebellar control of balance and locomotion.

The Neuroscientist : A Review Journal Bringing Neurobiology, Neurology and Psychiatry, *10*(3), 247–59. http://doi.org/10.1177/1073858404263517

Morton, S. M., & Bastian, A. J. (2006). Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(36), 9107–16. http://doi.org/10.1523/JNEUROSCI.2622-06.2006

Mugnaini, E., Atluri, R. L., & Houk, J. C. (1974). Fine structure of granular layer in turtle cerebellum with emphasis on large glomeruli. *Journal of Neurophysiology*, *37*(1), 1–29. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/4811974

Mugnaini, E., Sekerková, G., & Martina, M. (2011). The Unipolar Brush Cell: A Remarkable Neuron Finally Receiving the Deserved Attention. *Brain Research Reviews*, *66*, 220–245. http://doi.org/10.1016/j.brainresrev.2010.10.001.THE

Murphy, J. T., MacKay, W. A., & Johnson, F. (1973). Differences between cerebellar mossy and climbing fibre responses to natural stimulation of forelimb muscle proprioceptors. *Brain Research*, *55*(2), 263–289. http://doi.org/10.1016/0006-8993(73)90295-3

Nieus, T., Sola, E., Mapelli, J., Saftenku, E. E., Rossi, P., & D'Angelo, E. (2006). LTP regulates burst initiation and frequency at mossy fiber-granule cell synapses of rat cerebellum: experimental observations and theoretical predictions. *Journal of Neurophysiology*, *95*(2), 686–99. http://doi.org/10.1152/jn.00696.2005

Nowak, D. a, Topka, H., Timmann, D., Boecker, H., & Hermsdörfer, J. (2007). The role of the cerebellum for predictive control of grasping. *Cerebellum (London, England)*, *6*(1), 7–17. http://doi.org/10.1080/14734220600776379

Nunzi, M., & Birnstiel, S. (2001). Unipolar brush cells form a glutamatergic projection system within the mouse cerebellar cortex. *Journal of …*, *341*(January), 329–341. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/cne.1180/full

Nvidia. (2014). *Cuda c programming guide* (6.5 ed., p. 241). Nvidia. Retrieved from http://docs.nvidia.com/cuda/cuda-c-programming-guide/

O'Donoghue, D. (1989). Physiological and anatomical studies of the interactions between Purkinje cells and basket cells in the cat's cerebellar cortex: evidence for a unitary relationship. *The Journal of …*, (June). Retrieved from

http://www.jneurosci.org/content/9/6/2141.short

Ohyama, T., & Mauk, M. D. (2001). Latent acquisition of timed responses in cerebellar cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *21*(2), 682–690.

Ohyama, T., Medina, J. F., Nores, W. L., & Mauk, M. D. (2002). Trying to Understand the Cerebellum Well Enough to Build One. *Annals of the New York*, *978*, 425–438.

Ohyama, T., & Nores, W. L. (2003). Stimulus generalization of conditioned eyelid responses produced without cerebellar cortex: implications for plasticity in the cerebellar nuclei. *Learning & Memory*, 346–354. http://doi.org/10.1101/lm.67103.excitability

Ohyama, T., Nores, W. L., Medina, J. F., Riusech, F. A., & Mauk, M. D. (2006). Learning-induced plasticity in deep cerebellar nucleus. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(49), 12656–63. http://doi.org/10.1523/JNEUROSCI.4023-06.2006

Ohyama, T., Nores, W. L., Murphy, M., & Mauk, M. D. (2003). What the cerebellum computes. *Trends in Neurosciences*, *26*(4), 222–227. http://doi.org/10.1016/S0166-2236(03)00054-7

Oscarsson, O. (1979). Functional units of the cerebellum - sagittal zones and microzones. *Trends in Neurosciences*, (June), 143–145. Retrieved from http://www.sciencedirect.com/science/article/pii/0166223679900572

Ostojic, S. (2014). Two types of asynchronous activity in networks of excitatory and inhibitory spiking neurons. *Nature Neuroscience*, *17*(4), 594–600. http://doi.org/10.1038/nn.3658

Ozden, I., Sullivan, M. R., Lee, H. M., & Wang, S. S.-H. (2009). Reliable coding emerges from coactivation of climbing fibers in microbands of cerebellar Purkinje neurons. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *29*(34), 10463–73. http://doi.org/10.1523/JNEUROSCI.0967-09.2009

Palay, S. L. (1974). *Cerebellar cortex: cytology and organization* (p. 348). New York: Springer.

Palkovits, M., Magyar, P. A. L., & Szentágothai, J. (1971a). Quantitative Histological Analysis of the Cerebellar Cortex in the Cat. I. Number and Arrangement in Space

of the Purkinje Cells. *Brain Research*, *32*(2), 1–13.

Palkovits, M., Magyar, P. A. L., & Szentágothai, J. (1971b). Quantitative histological analysis of the cerebellar cortex in the cat. II. Cell numbers and densities in the granular layer. *Brain Research*, *32*(1), 15–30.

Palkovits, M., Magyar, P. A. L., & Szentágothai, J. (1971c). Quantitative Histological Analysis of the Cerebellar Cortex in the Cat. III. Structural Organization of the Molecular Layer. *Brain Research*, *34*, 1–18.

Palkovits, M., Magyar, P. A. L., & Szentágothai, J. (1972). Quantitative Histological Analysis of the Cerebellar Cortex in the Cat. IV. Mossy Fiber-Purkinje Cell Numerical Transfer. *Brain Research*, *45*, 15–29.

Palkovits, M., Mezey, E., Hámori, J., & Szentágothai, J. (1977). Quantitative Histological Analysis of the Cerebellar Nuclei in the Cat. I. Numerical Data on Cell and on Synapses. *Experimental Brain Research*, *28*, 189–209.

Palliyath, S., & Hallett, M. (1998). Gait in patients with cerebellar ataxia. *Movement Disorders*, *13*(6), 958–964. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/mds.870130616/abstract

Pellionisz, A. (1973). DYNAMIC SINGLE UNIT SIMULATION OF A REALISTIC CEREBELLAR NETWORK MODEL. *Brain Research*, *49*, 83–99. Retrieved from http://www.sciencedirect.com/science/article/pii/0006899373904034

Pellionisz, A., & Llinas, R. (1979). BRAIN MODELING BY TENSOR NETWORK THEORY AND COMPUTER SIMULATION . THE CEREBELLUM : DISTRIBUTED PROCESSOR FOR PREDICTIVE COORDINATION. *Neuroscience*, *4*(1959). Retrieved from http://www.sciencedirect.com/science/article/pii/0306452279900976

Pellionisz, A., Llinas, R., & Perkel, D. (1977). A computer model of the cerebellar cortex of the frog. *Neuroscience*, *2*, 19–35. Retrieved from http://www.sciencedirect.com/science/article/pii/0306452277900653

Perrett, S. P., & Mauk, M. D. (1995). Extinction of conditioned eyelid responses requires the anterior lobe of cerebellar cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *15*(3), 2074–2080.

Perrett, S. P., Ruiz, B. P., & Mauk, M. D. (1993). Cerebellar cortex lesions disrupt

learning-dependent timing of conditioned eyelid responses. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *13*(4), 1708–18.

Pijpers, A., Voogd, J., & Ruigrok, T. J. H. (2005). Topography of olivo-cortico-nuclear modules in the intermediate cerebellum of the rat. *The Journal of Comparative Neurology*, *492*(2), 193–213. http://doi.org/10.1002/cne.20707

Placantonakis, D. G., Bukovsky, A. a, Aicher, S. a, Kiem, H.-P., & Welsh, J. P. (2006). Continuous electrical oscillations emerge from a coupled network: a study of the inferior olive using lentiviral knockdown of connexin36. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(19), 5008–16. http://doi.org/10.1523/JNEUROSCI.0146-06.2006

Prinz, A. a, Bucher, D., & Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, *7*(12), 1345–52. http://doi.org/10.1038/nn1352

Pugh, J. R., & Raman, I. M. (2006). Potentiation of mossy fiber EPSCs in the cerebellar nuclei by NMDA receptor activation followed by postinhibitory rebound current. *Neuron*, *51*(1), 113–23. http://doi.org/10.1016/j.neuron.2006.05.021

Pugh, J. R., & Raman, I. M. (2008). Mechanisms of potentiation of mossy fiber EPSCs in the cerebellar nuclei by coincident synaptic excitation and inhibition. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(42), 10549–60. http://doi.org/10.1523/JNEUROSCI.2061-08.2008

Ramnani, N., Toni, I., Passingham, R. E., & Haggard, P. (2001). The cerebellum and parietal cortex play a specific role in coordination: a PET study. *NeuroImage*, *14*(4), 899–911. http://doi.org/10.1006/nimg.2001.0885

Rancillac, A., & Crépel, F. (2004). Synapses between parallel fibres and stellate cells express long-term changes in synaptic efficacy in rat cerebellum. *The Journal of Physiology*, *554*(Pt 3), 707–20. http://doi.org/10.1113/jphysiol.2003.055871

Rasmussen, A., Jirenhed, D.-A., & Hesslow, G. (2008). Simple and complex spike firing patterns in Purkinje cells during classical conditioning. *Cerebellum (London, England)*, *7*(4), 563–6. http://doi.org/10.1007/s12311-008-0068-2

Raymond, J. L., Lisberger, S. G., & Mauk, M. D. (1996). The cerebellum: a neuronal learning machine? *Science*, *272*, 1126–1131.

Robberechts, Q., Wijnants, M., Giugliano, M., & De Schutter, E. (2010). Long-term depression at parallel fiber to Golgi cell synapses. *Journal of Neurophysiology*, *104*(6), 3413–23. http://doi.org/10.1152/jn.00030.2010

Robinson, D. A. (1976). Adaptive gain control of vestibuloocular reflex by the cerebellum. *Journal of Neurophysiology*, *39*(5), 954–69. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1086347

Roland, P., & Larsen, B. (1980). Supplementary motor area and other cortical areas in organization of voluntary movements in man. *Journal of Neurophysiology*, *43*(1). Retrieved from http://jn.physiology.org/content/jn/43/1/118.full.pdf

Ruigrok, T. J. H., Hensbroek, R. a, & Simpson, J. I. (2011). Spontaneous activity signatures of morphologically identified interneurons in the vestibulocerebellum. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *31*(2), 712–24. http://doi.org/10.1523/JNEUROSCI.1959-10.2011

Schmahmann, J. D. (2004). Disorders of the cerebellum: ataxia, dysmetria of thought, and the cerebellar cognitive affective syndrome. *The Journal of Neuropsychiatry and Clinical Neurosciences*, *16*(3), 367–78. http://doi.org/10.1176/appi.neuropsych.16.3.367

Schmid, J. J. (1947). The Relationship between the Coefficient of Correlation and the Angle Included between Regression Lines. *The Journal of Educational Research*, *41*(4), 311–313. Retrieved from http://www.tandfonline.com/doi/pdf/10.1080/00220671.1947.10881608

Shinoda, Y., Sugihara, I., Wu, H. S., & Sugiuchi, Y. (2000). The entire trajectory of single climbing and mossy fibers in the cerebellar nuclei and cortex. *Progress in Brain Research*, *124*, 173–86. http://doi.org/10.1016/S0079-6123(00)24015-6

Shinoda, Y., Sugiuchi, Y., & Futami, T. (1987). Excitatory inputs to cerebellar dentate nucleus neurons from the cerebral cortex in the cat. *Experimental Brain Research*, *67*(2). http://doi.org/10.1007/BF00248551

Shinoda, Y., Sugiuchi, Y., Futami, T., & Izawa, R. (1992). Axon collaterals of mossy fibers from the pontine nucleus in the cerebellar dentate nucleus. *Journal of Neurophysiology*, *67*(3), 547–60. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1578244

Siegel, J. J., Kalmbach, B., Chitwood, R. a, & Mauk, M. D. (2012). Persistent activity in a cortical-to-subcortical circuit: bridging the temporal gap in trace eyelid conditioning. *Journal of Neurophysiology*, *107*(1), 50–64. http://doi.org/10.1152/jn.00689.2011

Siegel, J. J., & Mauk, M. D. (2013). Persistent Activity in Prefrontal Cortex during Trace Eyelid Conditioning: Dissociating Responses That Reflect Cerebellar Output from Those That Do Not. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(38), 15272–84. http://doi.org/10.1523/JNEUROSCI.1238-13.2013

Simpson, J. I., Hulscher, H. C., Sabel-Goedknegt, E., & Ruigrok, T. J. H. (2005). Between in and out: linking morphology and physiology of cerebellar cortical interneurons. *Progress in Brain Research*, *148*, 329–40. http://doi.org/10.1016/S0079-6123(04)48026-1

Simpson, J. I., Wylie, D. R., & De Zeeuw, C. I. (1996). On climbing fiber signals and their consequence(s). *Behavioral and Brain Sciences*, *19*(03), 384–398. http://doi.org/10.1017/S0140525X00081486

Sommer, F. T., & Wennekers, T. (2005). Synfire chains with conductance-based neurons: internal timing and coordination with timed input. *Neurocomputing*, *65-66*, 449–454. http://doi.org/10.1016/j.neucom.2004.10.015

Sotelo, C. (2003). Viewing the brain through the master hand of Ramon y Cajal. *Nature Reviews Neuroscience*, *4*(January), 1–7. Retrieved from http://www.nature.com/nrn/journal/v4/n1/abs/nrn1010.html

Spacek, J., Parízek, J., & Lieberman, a R. (1973). Golgi cells, granule cells and synaptic glomeruli in the molecular layer of the rabbit cerebellar cortex. *Journal of Neurocytology*, *2*(4), 407–28. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/4784778

Steinmetz, J. E. (1990). Classical nictitating membrane conditioning in rabbits with varying interstimulus intervals and direct activation of cerebellar mossy fibers as the CS. *Behavioural Brain Research*, *38*(2), 97–108. http://doi.org/10.1016/0166-4328(90)90008-3

Steinmetz, J. E., Lavond, D. G., & Thompson, R. F. (1989). Classical conditioning in rabbits using Pontine Nucleus Stimulation as a Conditioned Stimulus and Inferior

Olive Stimulation as an Unconditioned Stimulus. *Synapse*, *3*, 225–233.

Steinmetz, J. E., Logan, C. G., Rosen, D. J., Thompson, J. K., Lavond, D. G., & Thompson, R. F. (1987). Initial localization of the acoustic conditioned stimulus projection system to the cerebellum essential for classical eyelid conditioning. *Proceedings of the National Academy of Sciences of the United States of America*, *84*, 3531–3535.

Steinmetz, J. E., Logue, S. F., & Steinmetz, S. S. (1992). Rabbit classically conditioned eyelid responses do not reappear after interpositus nucleus lesion and extensive post-lesion training. *Behavioural Brain Research*, *51*(1), 103–114. http://doi.org/10.1016/S0166-4328(05)80317-1

Steinmetz, J. E., Rosen, D. J., Woodruff-Pak, D. S., Lavond, D. G., & Thompson, R. F. (1986). Rapid transfer of training occurs when direct mossy fiber stimulation is used as a conditioned stimulus for classical eyelid conditioning. *Neuroscience Research*, *3*(6), 606–616. http://doi.org/10.1016/0168-0102(86)90057-X

Stone, L. S., & Lisberger, S. G. (1990). Visual Responses of Purkinje Cells in the Cerebellar Flocculus During Smooth-Pursuit Eye Movements in Monkeys. *J Neurophysiol*, *63*(5).

Strick, P. L., Dum, R. P., & Fiez, J. a. (2009). Cerebellum and nonmotor function. *Annual Review of Neuroscience*, *32*, 413–34. http://doi.org/10.1146/annurev.neuro.31.060407.125606

Sugihara, I. (2006). Organization and remodeling of the olivocerebellar climbing fiber projection. *Cerebellum (London, England)*, *5*(1), 15–22. http://doi.org/10.1080/14734220500527385

Sussillo, D. (2014). Neural circuits as computational dynamical systems. *Current Opinion in Neurobiology*, *25*, 156–63. http://doi.org/10.1016/j.conb.2014.01.008

Sussillo, D., & Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, *63*(4), 544–57. http://doi.org/10.1016/j.neuron.2009.07.018

Svensson, P., Jirenhed, D.-A., Bengtsson, F., & Hesslow, G. (2010). Effect of conditioned stimulus parameters on timing of conditioned Purkinje cell responses. *Journal of Neurophysiology*, *103*(3), 1329–36. http://doi.org/10.1152/jn.00524.2009

Thach, W. T. (1998). A role for the cerebellum in learning movement coordination. *Neurobiology of Learning and Memory*, *70*(1-2), 177–88. http://doi.org/10.1006/nlme.1998.3846

Thach, W. T., Goodkin, H. P., & Keating, J. G. (1992). The cerebellum and the adaptive coordination of movement. *Annual Review of Neuroscience*, *15*, 403–42. http://doi.org/10.1146/annurev.ne.15.030192.002155

Topka, H., Konczak, J., & Dichgans, J. (1998). Coordination of multi-joint arm movements in cerebellar ataxia: analysis of hand and angular kinematics. *Experimental Brain Research. Experimentelle Hirnforschung. Expérimentation Cérébrale*, *119*(4), 483–92.

Toyoizumi, T., & Abbott, L. F. (2011). Beyond the edge of chaos: Amplification and temporal integration by recurrent networks in the chaotic regime. *Physical Review E*, *84*(5), 051908. http://doi.org/10.1103/PhysRevE.84.051908

Türker, K. S., & Miles, T. S. (1986). Climbing fiber lesions disrupt conditioning of the nictitating membrane response in the rabbit. *Brain Research*, *363*, 376–378.

Ulloa, A., Bullock, D., & Rhodes, B. J. (2003). Adaptive force generation for precision-grip lifting by a spectral timing model of the cerebellum. *Neural Networks : The Official Journal of the International Neural Network Society*, *16*(5-6), 521–8. http://doi.org/10.1016/S0893-6080(03)00094-7

Van Kan, P. L., Gibson, A. R., & Houk, J. C. (1993). Movement-related inputs to intermediate cerebellum of the monkey. *Journal of Neurophysiology*, *69*(1), 74–94. Retrieved from http://jn.physiology.org/content/69/1/74.short

Van Overwalle, F., Baetens, K., Mariën, P., & Vandekerckhove, M. (2013). Social Cognition and the Cerebellum: A Meta-analysis of over 350 fMRI studies. *NeuroImage*. http://doi.org/10.1016/j.neuroimage.2013.09.033

Vervaeke, K., Lorincz, A., Gleeson, P., Farinella, M., Nusser, Z., & Silver, R. A. (2010). Rapid desynchronization of an electrically coupled interneuron network with sparse excitatory synaptic input. *Neuron*, *67*(3), 435–51. http://doi.org/10.1016/j.neuron.2010.06.028

Vervaeke, K., Lorincz, A., Nusser, Z., & Silver, R. A. (2012). Gap Junctions Compensate for Sublinear Dendritic Integration in an Inhibitory Network. *Science (New York,*

N.Y.), 1624. http://doi.org/10.1126/science.1215101

Volny-Luraghi, A., Maex, R., Vos, B., & De Schutter, E. (2002). Peripheral stimuli excite coronal beams of Golgi cells in rat cerebellar cortex. *Neuroscience*, *113*(2), 363–373. http://doi.org/10.1016/S0306-4522(02)00196-3

Wang, S. S., Denk, W., & Häusser, M. (2000). Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience*, *3*(12), 1266–73. http://doi.org/10.1038/81792

Watanabe, D., & Nakanishi, S. (2003). mGluR2 postsynaptically senses granule cell inputs at Golgi cell synapses. *Neuron*, *39*(5), 821–9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12948448

Watanabe, E. (1984). Neuronal events correlated with long-term adaptation of the horizontal vestibulo-ocular reflex in the primate flocculus. *Brain Research*, *297*(1), 169–174. http://doi.org/10.1016/0006-8993(84)90555-9

Welsh, J. P., Yamaguchi, H., Zeng, X., Kojo, M., Nakada, Y., Takagi, A., … Llinás, R. (2005). Normal motor learning during pharmacological prevention of Purkinje cell long-term depression. *Proceedings of the National Academy of Sciences of the United States of America*, 1–6. Retrieved from http://www.pnas.org/content/102/47/17166.short

White, N., Kehoe, E., Choi, J., & Moore, J. (2000). Coefficients of variation in timing of the classically conditioned eyeblink in rabbits. *Psychobiology*, *28*(4), 520–524. Retrieved from http://link.springer.com/article/10.3758/BF03332010

Wiechert, M. T., Judkewitz, B., Riecke, H., & Friedrich, R. W. (2010). Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nature Neuroscience*, *13*(8), 1003–10. http://doi.org/10.1038/nn.2591

Winfield, J. A., Hendrickson, A., & Kimm, J. (1978). Anatomical evidence that the medial terminal nucleus of the accessory optic tract in mammals provides a visual mossy fiber input to the flocculus. *Brain Research*, *151*(1), 175–182. http://doi.org/10.1016/0006-8993(78)90961-7

Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(4), 1314–28.

http://doi.org/10.1523/JNEUROSCI.3733-05.2006

Woodruff-Pak, D. S., Lavond, D. G., & Thompson, R. F. (1985). Trace conditioning: Abolished by cerebellar nuclear lesions but not lateral cerebellar cortex aspirations. *Brain Research*, *348*(2), 249–260. http://doi.org/10.1016/0006-8993(85)90443-3

Xu, W., & Edgley, S. a. (2008). Climbing fibre-dependent changes in Golgi cell responses to peripheral stimulation. *The Journal of Physiology*, *586*(Pt 20), 4951–9. http://doi.org/10.1113/jphysiol.2008.160879

Yamazaki, T., & Tanaka, S. (2007). The cerebellum as a liquid state machine. *Neural Networks : The Official Journal of the International Neural Network Society*, *20*(3), 290–7. http://doi.org/10.1016/j.neunet.2007.04.004

Zheng, N., & Raman, I. M. (2010). Synaptic inhibition, excitation, and plasticity in neurons of the cerebellar nuclei. *Cerebellum (London, England)*, *9*(1), 56–66. http://doi.org/10.1007/s12311-009-0140-6