

**WORKING MEMORY AND TEMPORAL
PATTERNS: THE IMPLICATIONS OF NEURAL
POPULATIONS**

by

Sergio Oscar Verduzco-Flores

Computer Engineer,

Universidad Nacional Autónoma de México,

2003

Submitted to the Graduate Faculty of
the Department of Mathematics in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2011

UNIVERSITY OF PITTSBURGH
SCHOOL OF ARTS AND SCIENCES
DEPARTMENT OF MATHEMATICS

This dissertation was presented

by

Sergio Oscar Verduzco-Flores

It was defended on

December 17th 2010

and approved by

G. Bard Ermentrout, Department of Mathematics University of Pittsburgh

Mark Bodner, Department of Mathematics University of Pittsburgh

Brent Doiron, Department of Mathematics University of Pittsburgh

Jonathan E. Rubin, Department of Mathematics University of Pittsburgh

Nathan Urban, Department of Biological Sciences Carnegie Mellon University

Dissertation Director: G. Bard Ermentrout, Department of Mathematics University of
Pittsburgh

WORKING MEMORY AND TEMPORAL PATTERNS: THE IMPLICATIONS OF NEURAL POPULATIONS

Sergio Oscar Verduzco-Flores, PhD

University of Pittsburgh, 2011

Working memory, the ability to temporarily retain information which will be used to guide subsequent behavior, is a central component of our cognitive abilities. Almost 40 years ago, electrophysiological experiments in monkeys established that persistent activity may be the neuronal substrate of working memory. Many computational models have been proposed in order to explain persistent activity, with recurrent connections playing a prominent role in many of these. No model, however, has captured all the important features in working memory networks. This work presents three related models which seek to understand some of these features. In particular, the first model explores the formation of firing rate patterns during the delay period of working memory tasks; the second model explores the dynamics of working memory networks with reduced inhibition, and their possible role in epilepsy; the third model explores the formation of temporal sequences and closed loops of activity. The three models assume the existence of densely connected neural populations (such as minicolumns), which respond similarly to the same stimuli. Another common feature is the role of dynamic synapses: the first two models rely on synaptic facilitation, whereas the third uses temporally asymmetric Hebbian plasticity.

TABLE OF CONTENTS

1.0 INTRODUCTION	1
2.0 WORKING MEMORY CELLS' BEHAVIOR MAY BE EXPLAINED BY CROSS-REGIONAL NETWORKS WITH SYNAPTIC FACILITATION	7
2.1 INTRODUCTION	7
2.2 METHODS	12
2.3 RESULTS	22
2.3.1 Single Population Model	22
2.3.2 2-Population Firing Rate Model	29
2.3.3 4-Population Firing Rate Model	32
2.3.4 Spiking Unit Network: 2-Population Model	32
2.3.5 Spiking Unit Network: 4-Population Model	36
2.3.6 Spiking Unit Network's Variability Scaling With Population Size	37
2.4 DISCUSSION	42
3.0 FROM WORKING MEMORY TO EPILEPSY: DYNAMICS OF FACILITATION AND INHIBITION IN A CORTICAL NETWORK	53
3.1 INTRODUCTION	53
3.2 METHODS	57
3.3 RESULTS	58
3.3.1 Normal Parameters	59
3.3.2 Pathology	60
3.3.2.1 All E-I disrupted	60

3.3.2.2	Partial Disinhibition	61
3.4	TWO-POPULATION MODEL	62
3.4.1	Attractors	62
3.4.2	Basins of Attraction	65
3.5	ONE POPULATION	68
3.6	SEIZURE TERMINATION	73
3.7	DISCUSSION	75
4.0	A MODEL FOR COMPLEX SEQUENCE LEARNING AND REPRO-	
	DUCTION IN NEURAL POPULATIONS	83
4.1	INTRODUCTION	83
4.2	RESULTS	85
4.2.1	Temporally Asymmetric Hebbian Learning on Asynchronous Popu-	
	lations	85
4.2.2	Overview of the Firing Rate Model	90
4.2.3	Storage and Retrieval of Sequences	93
4.2.4	Properties of the Model	99
4.3	DISCUSSION	104
4.3.1	Forming Connections in Asynchronous Populations	105
4.3.2	Which Sequences Can Be Stored and Reproduced?	105
4.3.3	Inhibition, Overlap, and Binding	107
4.3.4	Working Memory and Oscillations	107
4.4	METHODS	109
4.4.1	The Spiking Model	109
4.4.2	The Firing Rate Model	109
4.4.3	Firing Rate Model Implementation	116
4.4.4	Parameter Search	116
4.4.5	Input Patterns	117
5.0	DISCUSSION	118
	APPENDIX. SOURCE CODE	122
A.1	Source Code for Chapter 1	122

A.1.1	average9-e.ode	122
A.1.2	average10-e.ode	122
A.1.3	avnulls.ode	123
A.1.4	runner1.m	123
A.1.5	nc1.m	127
A.1.6	nc1euler.m	128
A.1.7	histAnaP.m	129
A.2	Source Code for Chapter 2	130
A.2.1	bigrunM5d.m	130
A.2.2	M5dvar.m	131
A.2.3	f.m	132
A.2.4	p.m	132
A.2.5	greatrunM5d.m	132
A.2.6	greathelper.m	133
A.3	Source Code for Chapter 3	134
A.3.1	fourPops2.sli	134
A.3.2	nestFourPops1.m	139
A.3.3	ras4NEST.m	141
A.3.4	TestSeq23gen2.m	141
A.3.5	euler23.m	145
A.3.6	euler23d.m	146
A.3.7	seqF23.m	147
A.3.8	seqF23d.m	148
A.3.9	I7.m	149
A.3.10	I8.m	149
A.3.11	I9.m	150
A.3.12	createCEI2.m	150
A.3.13	createCIE1.m	150
A.3.14	randSim1.m	150
A.3.15	randSim2.m	153

BIBLIOGRAPHY 155

LIST OF TABLES

1	Parameter Values	18
2	Connections for 2 populations	19
3	Connections for 4 populations	20
4	Average Firing Frequencies	34
5	Spiking Model Summary.	110
6	Populations in the spiking model.	110
7	Connectivity of the spiking model.	111
8	Neuron and synapse models in the spiking network.	112
9	Measurements taken for the spiking network.	112
10	Parameter values for the spiking model.	113
11	Parameter values for the firing rate model	117

LIST OF FIGURES

2.1	Network architectures in chapter 2	15
2.2	Analysis of the single population model	17
2.3	Single Population Firing Rate	27
2.4	Patterns in the single population model	28
2.5	Histograms for the 2-population Model	30
2.6	Phase Diagrams for the 2- and 4-population Models	31
2.7	The firing rate model and the average activity	33
2.8	Histograms for the 2-population Model and Real Cells	35
2.9	Distribution of Units by Pattern Type	38
2.10	Raster Plots for the 4-population Network	39
2.11	Histograms for 16 Randomly Chosen Units	39
2.12	Distribution of Patterns	41
2.13	Reliability of the Units	43
2.14	Analysis of Spiking Variability	44
3.1	20-population network with normal inhibition	59
3.2	20-population network with reduced inhibition	60
3.3	Reduced inhibition and periodic stimuli	61
3.4	Localized damage	63
3.5	Two population bifurcations	63
3.6	Steady state of the pathological network	66
3.7	Expanded view	67
3.8	Two population network	69

3.9	Bifurcations diagram for one population	69
3.10	Fast dynamics	71
3.11	Evolution of w	72
3.12	Termination of activity	74
3.13	Inhibitory inputs	76
3.14	Terminating oscillations	77
3.15	Preventing oscillations	78
4.1	Architecture of the Model Networks	86
4.2	Formation of Excitation Pathways	88
4.3	Temporally Asymmetric Hebbian Plasticity	91
4.4	Sequence Encoding and Reproduction	94
4.5	Creation of a Closed-loop Simple Sequence	95
4.6	Transition Between Sequences	97
4.7	Input with Overlap May Modify Existing Connections	98
4.8	Formation of a Complex Sequence	100

1.0 INTRODUCTION

The most quoted passage in Neuroscience, is perhaps the following one:

When an axon of cell A is near enough to excite B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.

These lines, written by Donald Hebb in 1949 [93], were a postulate in a theory which attempted to link neurophysiology and behavior. In that theory, cells firing together would associate into groups wherein neural activity could reverberate, remaining active for short periods of time. Those groups of neurons, called *cell assemblies*, would constitute the basis of mental representation, and the sequential activation of cell assemblies would be the basis of thinking [40]. These ideas continue to inspire neuroscientists. The concept of cell assemblies is still imprecise, but researchers often used the correlated activity between neurons to modulate synaptic strength, creating populations of strongly connected cells which can maintain reverberating activity. The terms *Hebbian plasticity* or *Hebbian learning* are used to refer to these types of mechanisms for synaptic modification. The first direct evidence for Hebbian plasticity surfaced in 1973 [18] in area CA1 of the hippocampus. It was significant that this mechanism—which is now known as long-term potentiation, or LTP—was found in an area shown to be closely connected to spatial learning and episodic memory, and this motivated various studies to reveal its properties [137]. Hippocampal LTP is only one of the various forms of long-term synaptic plasticity existing in the mammalian brain. Moreover, most synapses that can be strengthened by some form of LTP can also be weakened by one or more forms of an opposite mechanism, called long-term synaptic depression (LTD). A key concept that has emerged is that excitatory synapses can be bidirectionally modified by different patterns of activity [37].

The ideas of Hebbian plasticity and cell assemblies seem to agree with various anatomical [179,211] and physiological [104] studies, and they give basis to the idea of using roughly homogeneous neuronal populations in several models. In addition to the densely connected neural populations which can emerge from neural activity and Hebbian learning, the cerebral cortex of mammals has been found to present a modular organization, since columns and minicolumns form accross the cortical layers of developing animals [146]. Thus, neuronal populations with strong connections and similar response properties can have both ontogenetic and epigenetic origins. This creates a strong justification for the exploration of cortical dynamics when groups of densely interconnected populations are assumed to exist, which is the main theme of the present work.

One important type of cortical dynamics which has a direct link to behavior is the persistent activity associated with working memory. Lesions in the dorsolateral prefrontal cortex of primates impair performance in working memory tasks [16,157], and in the 1970s it was found that in tasks where the memory of a stimulus had to be maintained during a delay period, neurons in prefrontal cortex and thalamus would maintain elevated firing rates, presumably retaining the information required for future actions [70]. Further electrophysiological measurements supported this finding [64,65,139,164], and together with lesion, local inactivation [99,168], and imaging studies [43], consolidated the role of persistent activity—particularly by prefrontal cortex neurons—in studies of working memory.

Many modelling studies have focused in replicating the persistent activity observed during working memory tasks [51]. Among those studies, one dominant mechanism has been the sustained activity through recurrent connections in a cell assembly. A large part of the studies implementing this mechanism have been inspired by the Hopfield model [95], which can learn patterns through an abstract form of Hebbian learning, and can retrieve them despite noisy and incomplete cues, displaying the fault tolerance and generalization characteristic of human memory. The Hopfield model also showed an isomorphism with spin glass systems, firing up the imagination of many physicists. This might have contributed to the large amount of work poured into models similar to Hopfield’s, called *attractor networks* [6,163]. While attractor networks replicated the qualitative characteristics of prefrontal cortex neurons during delayed response tasks (namely, firing rates that were higher or lower during

the delay period), they missed several quantitative features of real neurons, including their firing rates, and the temporal variations that they would present during the delay periods. These temporal variations became important with the growing realization that they could be the indicators of interactions between prefrontal cortex and other brain areas involved in bridging perception to behavioral outputs [35, 63, 73, 171].

Chapter 2 of this work presents a model network [200] which reproduces the activation patterns observed during the delay period of working memory tasks [171]. There are three main ideas which allow this model to mimic the variability found in the activity of real neurons. The first idea is to assume strongly connected neuronal populations, as has been mentioned above. The second idea is that populations from one region (such as prefrontal or parietal cortex) send long-range connections to populations in other regions, and these connections can have either an excitatory or an inhibitory effect [35]. The third idea is to include calcium-mediated presynaptic facilitation in the long-range connections. Various forms of facilitation, such as augmentation and post-tetanic potentiation are usually observed when trains of stimulation appear at high frequencies (10-200 Hz). They often involve an increase in the probability of transmitter release due to an accumulation of intracellular calcium in the presynaptic cell [37, Pg 19]. In delayed-response working memory tasks, the presentation of the stimulus to be remembered elicits high-frequency activity in the cells which presumably encode it during the delay period; we hypothesize that this activity results in presynaptic facilitation through a buildup in intracellular calcium, resulting in a transient memory trace encoded in the synaptic state, which allows a richer set of dynamics, as shall be shown (a similar assumption is presented in [142]).

The main technical hurdle treated in chapter 2 is to find a parameter set which shows activity patterns similar to those in experimental recordings. The approach taken is to first reduce the (800-dimensional) spiking model into a mean-field model with only 3 variables. This mean-field model is further reduced into 2 dimensions by assuming a steady state in the calcium dynamics. The resultant 2-dimensional system is not only amenable to phase-plane analysis, but has fixed points whose structure follows that of the 3-dimensional mean field model, which in turn is a good description of the mean activity in a homogeneous network of spiking neurons. By analyzing the 2-dimensional model parameters are found

where its nullclines create a bottleneck, whose effect is to generate mean-field activity which decays slowly after being increased by the high firing rates when the stimulus is presented. By adjusting the strength of the connections between different populations, each individual population may or may not present a bifurcation beyond this bottleneck, which allows for a variety of activation patterns, closely resembling those found in experimental measurements.

The model in chapter 2 only requires two different populations in order to replicate the activation patterns of real neurons, and up to four populations were modeled. It is to be expected that in the brain the number of distinct populations would be much higher, and this could have a significant impact on the dynamics of the network. In chapter 3 we present an extension of the model used in chapter 2, with a larger number of populations, but simplified dynamics (calcium is no longer modeled explicitly), and populations which have coupled inhibitory and excitatory subpopulations.

The larger working memory network used in chapter 3 brings interesting consequences related to the possible source of pathological states in the brain. Epileptic seizures consist of periods with increased network excitation and variable propagation, with looping impulses being a critical factor in generating the hyperexcitability and the recruitment [107, 123, 190–192]. On the other hand, working memory networks bring selective persistent activation, presumably through reentrant connections. The main hypothesis being explored in chapter 3 is that the uncontrolled activity observed in epileptic seizures can originate from the loss of selectivity in working memory networks due to some process affecting their excitability. Furthermore, the dynamics of the working memory network in the pathological regime are such that inputs with a characteristic temporal structure can result in generalized oscillations, in a manner similar to the onset of reflex epilepsy.

In order to explain the dynamics observed in the model of chapter 3 in the pathological regime, we once more used the homogeneity in the network in order to obtain models with reduced dimensionality. We studied a model with two populations using bifurcation analysis. The attractor structure observed from this analysis provides valuable insights about the behavior of the full model, as the two population network has dynamics matching those of the 20-population one. To further understand how this attractor structure is formed we investigated the dynamics of a single population (which represents the symmetric solution

of the coupled populations), using the method of averaging and numerical analysis.

The model in chapter 3 presents isotropic all-to-all connectivity among different populations. While this is convenient for symmetry reductions, an asymmetric connectivity scheme might be required to produce activity patterns with localized spatiotemporal patterns. If cortical connections were homogeneous, they could not produce the intricate computations spanning from perception to behavior. We therefore created a new model where the connectivity between populations can be shaped by inputs external to the network. In particular, this model uses *temporally asymmetric Hebbian plasticity* in order to learn to reproduce input sequences. This model is the subject matter of chapter 4.

Temporally asymmetric Hebbian plasticity occurs when synaptic strength can be up- or down-regulated depending on the timing of pre and postsynaptic action potentials. One such mechanism was found experimentally in 1997 [134], and is known as Spike Timing Dependent Plasticity (STDP). The study of STDP and its effects is full of unsettled issues. The cellular mechanisms of STDP are not entirely understood, nor are the effects of more complex spike trains, or the modulation by other inputs [30]. Accordingly, there are various phenomenological models of STDP [145], and the particular model used can affect the distribution of synaptic strengths which is obtained [84, 167, 178].

The mechanism of STDP is suggestive for the autonomous formation of sequential activity, and some studies have suggested how feed-forward structures could be created through it in biologically plausible networks [57, 78, 103, 109, 128]. One thing to notice from these studies is that some assumption is needed in order to allow the stable autonomous formation of sequential structure. In [103], sequences can appear beyond chance level when driven by background activity, but this requires assuming a particular distribution of the axonal delays. The network in reference [109] can autonomously create synfire-chain structures, but this requires activation and silencing of synapses, as well as pruning the weak outgoing connections of neurons that have formed enough strong synapses. The model in [128] presents temporally and spatially sparse activity which circulates through the whole network in response to a stimulus, and this is achieved through a particular learning rule (presynaptic-dependent scaling) which tends to induce a target postsynaptic activity in all neurons. The study in [78] assumes that the changes in synaptic strength are slow enough so that their evolution can

be described by firing rates and their correlations, and then describes how input correlations can create a sequential structure in a plastic network. Finally, in [57] it is shown how STDP can create sequential activity of individual neurons; to overcome the lack of robustness in a sequence where only one neuron carries the activity at a time, neurons respond with bursts when driven beyond threshold, and heterosynaptic competition ensures that each neuron only has one main target after the learning phase.

The model presented in chapter 4 is not exempt from assumptions. In particular, it is assumed that there are neuronal populations present, and that there is temporally asymmetric Hebbian learning between them. Assuming the existence of neural populations is not controversial, but it is unusual to claim that the average of the connection strengths from one population to another depends on the relative timing of increments in their mean activity. The biological plausibility of the model hinges on the justification of this last assumption, and that is the reason why the first part of chapter 4 consists of a quick computational study which uses networks of integrate-and-fire neurons to illustrate how temporally asymmetric Hebbian learning could be achieved at the population level. The suggested mechanism for this takes advantage of the synchrony induced by the strong connectivity inside each population, and will be described in more detail further ahead. The second part of chapter 4 presents a firing-rate model in which populations can reproduce the order in which they were activated by an input with the right spatiotemporal characteristics.

2.0 WORKING MEMORY CELLS' BEHAVIOR MAY BE EXPLAINED BY CROSS-REGIONAL NETWORKS WITH SYNAPTIC FACILITATION

2.1 INTRODUCTION

Persistent elevation in firing rates of cortical neurons during retention of memoranda has been suggested to represent the neuronal correlate of working memory [64, 65, 70]. This activity in so called memory cells (as observed in microelectrode recordings of neurons in the cortex of primates during the performance of delay tasks) exhibits a number of different general patterns. One pattern consists of cells whose elevated firing rate persists, on average across trials of the delay task, at the same rate for the entire period during which information of the memorandum is maintained in working memory. This type of dynamics represents the canonical bistable activity which has been a major focus of theoretical and computation modeling. A second elevated firing rate pattern consists of cells whose rate either decreases or increases throughout the memory period of a delay task. In decreasing-rate memory cells, the elevated activity is attuned to the memorandum (cue) of the task, and firing rate decays as the delay progresses towards the response of the task. In increasing-rate or ramping cells, elevated activity is motor- or response-coupled, and firing rate accelerates as the response of a task approaches. These rate-changing pattern cells and their respective networks have been suggested to represent two mutually complementary and interactive representations engaged in the transfer of information of cross-temporal contingencies from memory to action in working memory. Cells exhibiting these pattern types have been found to occur anatomically intermixed in the cortex [160], with cue- and response-coupled cells appearing to be more common than fixed delay rate cells [171]. In addition to these persistently elevated firing rate patterns, neurons which presumably are constituent members of working memory

networks exhibit analogous inhibited (below baseline rate) firing patterns. Finally, many cells exhibit firing rate changes correlated with different working memory task events such as the presentation of memoranda (cue period) and/or the response of the delay task, but maintain baseline firing rates throughout the delay period during which the memorandum is retained in active short-term memory.

While the mechanism(s) by which the patterns of activity are initiated and maintained in working memory are undetermined, a number of plausible hypothesis have been proposed. With respect to persistent elevated-rate patterns, prevailing ideas which have emerged from computational and theoretical studies are that the activity arises as stable states in recurrent attractor networks [6–11, 26, 38, 50, 120, 203] and/or inherent cellular dynamics [29, 44, 45, 52, 59, 126, 133, 205]. These studies have had success in reproducing general bistable memory behavior. For example with respect to network studies, successful working memory behavior has been attained as defined by achieving persistent increased firing rates of cue-specific subpopulations of units in networks during the putative memorandum retention period of simulating delay tasks. A difficulty typically encountered however, is obtaining memory behavior, as defined within the specific range of frequency rates, statistics, and with the variability as observed in real neuronal populations of cortical working memory networks across the range of different persistent patterned behaviors. Neurons exhibiting each of the different persistent activity pattern types with some overall average frequency do so only as an average across multiple trials of a working memory task. Individual cells exhibit a significant amount of variability however, both in terms of firing frequency within and between trials of working memory tasks and may even exhibit different patterned behaviors from trial to trial [171]. Thus while cells exhibit one of the given patterns described above with some overall average firing frequency across many trials (as observed for example in an average peristimulus time histogram), they exhibit different average firing rates and/or pattern behaviors from trial to trial of the working memory task.

A potential source of these and other difficulties [170], is that they are examined within the framework of static synaptic structures. Specifically, the networks have fixed architectures, and are trained such that the strength of the connections between units (the weight matrices) produce desired memory behavior. Once memory behavior is achieved, the weight

matrix is held constant. However, the simplification of fixed synaptic strengths may not be physiologically reasonable in light of the highly dynamic structure of the cortex. Cortical networks, and their constituent neurons, receive constant input from both external and internal sources, with learning and plasticity occurring concurrently with behavior. From a functional standpoint, fixed connection strengths necessarily limits the number of activities a network can perform, which could be undesirable given the plasticity of cortical function. Further, the ubiquity of cortical working memory [66] suggests that its associated activity might not occur in fixed, dedicated networks, but rather may arise from processes present in networks performing a variety of different functions [69].

Functional architecture is a second consideration of potentially fundamental importance to the dynamics of working memory networks. Typically, efforts have focused on studying working memory within the framework of local modules or networks that exist at various specific or general locations in the cortex. However, while working memory and/or working memory-correlated neuronal activity may be maintainable within local networks (or even cellularly), considerable evidence from neurophysiological and imaging studies have shown that working memory involves widely distributed cortical networks across multiple cortical areas [69]. Such a widely distributed architecture, which, if not fundamentally necessary for producing the firing rate patterns observed in working memory network cells, is probably active in the modulation of that activity. This modulation might entail not only producing the specific range of firing rates, but also the range of pattern types.

Recent work has indicated that working memory networks incorporate dynamic synapses. One study [206] revealed that connections between pyramidal cells in the prefrontal cortex exhibit facilitation, while others have demonstrated that neocortical synapses undergo substantial synaptic plasticity following synaptic activity [33,34]. Particularly, it has been found that cells in certain cortical regions exhibit increased responses to sequences of theta burst stimulation, both from burst to burst within a given burst sequence, as well as across successive sequences. Work by Hempel et al., [94], and Galaretta and Henstrin [76] indicated that cortical synapses can exhibit augmentation (from 15 to 60 percent) that correlates with the frequency and duration of tetanic stimulation —similar to that frequently observed during the presentation of memoranda in working memory tasks.

Several computational efforts have attempted to address various aspects of the issues described above. For example, one study demonstrated that persistent activation with realistic frequencies might be achieved if working memory corresponds to attractor states other than the low firing-rate background state, and have proposed mechanisms by which such states might be stabilized [122]. Other work has emphasized the potential role of dynamic synapses in working memory processes, examining the effects of dynamic synaptic augmentation and rapid Hebbian plasticity in a recurrent network framework [170]. This work indicated that synaptic augmentation can reduce the amount of prior structure required for persistent activation to take place, while rapid Hebbian plasticity could enable persistent activity to take place within firing rate ranges observed in real cortical neurons. More recent studies have demonstrated that combinations of synaptic depression and facilitation might extend the attractor neural network framework to represent time-dependent stimuli [13]. Further efforts have indicated that calcium media synaptic facilitation could produce bistable persistent activation with firing rate increases typically observed in real cortical cells [142].

While working memory models have mostly concentrated on bistable persistent activation, some efforts have also addressed the issue of cue- or response-coupled patterns of activity that steadily increase and decrease during delay periods. For example, graded activity in recurrent networks with slow synapses has been modeled [205], while another recent study examined such activity in uniform recurrent networks with stochastic bimodal neurons without NMDA-receptor-mediated slow recurrent synapses [153]. This work has indicated that graded memory activity could be very difficult to produce within a single population or local module. Still other studies have examined the ability of networks to produce ramping behavior by maximizing the time the systems trajectory spends around a saddle in the systems phase space [49]. Other work, while not necessarily producing working memory cells with firing rate statistics of real cells, has examined networks that produce the types of general patterns observed in working memory [176,214]. A distributed network architecture may be crucial in understanding and producing those patterns of activity.

In this work we examine a cortical model of working memory incorporating dynamic synapses both within a local and a distributed cortical framework. We investigate the mechanism of dynamic synaptic facilitation in the generation of all of the different patterns of

persistent activity associated with working memory and the effect of a distributed cortical architecture on the dynamics of working memory patterns. We first examine a firing rate model incorporating dynamic synapses representing a working memory network residing locally in a given cortical area. We analyze the statistics and firing-rate-patterns of this network during simulated working memory and compare the results with that of real cortical neurons recorded from parietal and prefrontal cortex of monkeys performing working memory tasks. A reduction of this model to a 2-dimensional system enables an analysis to completely characterize the states of the system. We then examined a distributed firing rate model consisting of 2 and 4 locally interconnected networks, analyzing the possible states as a function of different long-range connectivity schemes and strengths. The expansion of the architecture to multiple networks allows the incorporation of possible heterogeneity. We compare the output of these models (local and global architecture) with the activity of the database of real cortical neurons recorded extracellularly from the prefrontal and parietal cortex of primates performing working memory tasks. The model expands on previous work examining the ability of population models with dynamic synapses to produce bistable memory states, or rate changing states (either cue dependent during the stimulus period –i.e. Barak and Tsodyks [13]– or exclusively rate changing during the delay [49]) to produce all different patterns (including inhibitory patterns) recorded during the delay period, and that these patterns can change their temporal features to accommodate a continuum of delay periods, as well as possessing relative rate changes and statistics as recorded in real cortical neurons. We also demonstrate that different patterns can occur in a distributed network concomitantly in a complimentary fashion as observed in the cortex. From the mean field firing rate model, a spiking network model is obtained whose populations mean firing rate corresponds to that of the firing rate model. This enables direct comparison of the activity with real cortical neurons. We examined the effect on unit activity with this spiking network with a distributed architecture consisting of up to four local networks connected by long range projections. The patterns and statistics of these spiking networks are analyzed and directly compared with the range of activities and firing statistics observed in the database of real cortical neurons. Finally we quantify the variability in spiking unit activity as observed in real cortical networks, and demonstrate from a nonlinear analysis how this activity arises.

The results are compared to that observed in the real cortical cell populations. The results of this work demonstrates that all of the firing patterns correlated with working memory are inherently generated in distributed networks incorporating dynamic synapses, and these exhibit variability and firing rate statistics in agreement with what is observed in the cortex.

2.2 METHODS

We start with a firing rate model of a local network (Figure 2.1A). While the population might correspond to a network anywhere in the cortex, for convenience for comparison with the real cortical data, we might associate it with a working memory network in prefrontal or parietal cortex. The network equation describing the synaptic activity of the population is given by

$$\frac{ds}{dt} = \frac{s_{min} - s}{\tau_s} + F(Cws + I(t) - \theta) \quad (2.1)$$

where s denotes synaptic activity. The second term in 2.1 corresponds to the firing rate of the population with the function $F(X)$ given by

$$F(X) = \frac{1}{\pi\tau_m} \sqrt{\frac{X}{1 - e^{-bX}}} \quad (2.2)$$

in which τ_m is the membrane time constant, and b is a parameter inversely proportional to the noise. This form of the firing rate function mimics the firing rate of a class I neuron in the presence of noise ($\sim 1/b$) [54]. The parameter C in equation 2.1 is the strength of feedback connections in the population, τ_s is the decay constant for synaptic activity, w corresponds to the synaptic facilitation, and θ is the threshold. $I(t)$ corresponds to an external current which increases during memorandum (cue) presentation in the simulated working memory task. Dynamic synaptic facilitation (w) is incorporated in the model according to

$$\frac{dw}{dt} = \frac{w_{min} - w}{\tau_w} + \gamma(w_{max} - w) \left(\frac{Ca}{Ca_0} \right)^e \quad (2.3)$$

where τ_w is the decay constant, γ is a proportionality constant controlling the amount of facilitation as a function of intra-cellular calcium, and Ca is the calcium concentration. Ca_0

is a reference parameter controlling the level of intracellular calcium at which facilitation begins to increase. The calcium concentration dynamics are given by

$$\frac{dCa}{dt} = \frac{Ca_{min} - Ca}{\tau_{Ca}} + F(Cws + I(t) - \theta) \quad (2.4)$$

where τ_{Ca} is the decay constant, and $F(x)$ is of the form given in equation 2.2. The above local architecture of the model is expanded to a distributed one, first through the addition of a second population, coupled to the first by recurrent long-range projections (Figure 2.1B). This allows the introduction of heterogeneity into the network as well as representing the first step towards investigating the effect of a distributed architecture on the dynamics and states of working memory. The system dynamics are described by the coupled network equations describing the synaptic activity

$$\frac{ds_i}{dt} = \frac{s_{i,min} - s_i}{\tau_s} + F\left(\sum_{j=1}^2 C_{ji}w_js_j + I(t) - \theta_i\right) \quad (2.5)$$

where $i = 1, 2$ corresponding to the two populations. The two populations can be considered to reside in different cortical areas (i.e. prefrontal and parietal cortex) or two populations within the same area. For convenience of description we can consider the populations to represent networks in different cortical areas, which for purposes of association with the real cortical data we take as prefrontal cortex (population 1) and parietal cortex (population 2). In these equations then, C_{12} represents the strength of the projections from prefrontal cortex to the parietal cortex population, and C_{21} is the connection strength from the parietal population to the prefrontal population, while C_{11} and C_{22} are the connections strengths within the prefrontal and parietal populations respectively. Synaptic facilitation is given by equation 2.3 and Calcium dynamics satisfy equations similar to 2.4 which are:

$$\frac{dCa_i}{dt} = \frac{Ca_{i,min} - Ca_i}{\tau_{Ca}} + F\left(\sum_{j=1}^2 C_{ji}w_js_j + I(t) - \theta_i\right) \quad (2.6)$$

The distributed architecture is further extended to one consisting of four populations (Figure 2.1C), by recurrently connecting two of the 2-population models above such that every population has projections to every other population. The system dynamics are given by:

$$\frac{ds_i}{dt} = \frac{s_{i,min} - s_i}{\tau_s} + F\left(\sum_{j=1}^4 C_{ji}w_js_j + I(t) - \theta_i\right) \quad (2.7)$$

where $i=1, 2, 3, 4$ correspond to the 4 populations, and with analogous extensions of equation 2.6 controlling the calcium dynamics. In this network each pair of populations (i.e. populations 1 and 2, and populations 3 and 4) are more strongly coupled to each other than they are to populations of the other pair. The network can be considered now to represent two local networks consisting of 2 populations each, residing within different cortical areas (i.e. prefrontal and parietal cortex). Thus the effects of heterogeneity may be examined, in addition to the effect of a distributed architecture on working memory dynamics and states. Particularly it allows the examination of the effect of heterogeneity and a distributed architecture on the occurrence of “complementary” working memory behaviors indicated by experiments to be simultaneously present in networks in the cortex.

We begin the analysis first from the single population model. The single population possesses 3-dimensional dynamics in the variables for synaptic activity (s), facilitation (w), and calcium concentration (Ca). A reduction of this model to 2 dimensions is achieved by assuming steady state calcium ($dCa/dt = 0$) allowing the system to be rigorously analyzed. While assuming steady state calcium does not have an immediate justification from a neurophysiological standpoint, it produces a system with the same attractor structure as the 3-dimensional system and thus allows the rigorous analysis. We carried out analysis of the dynamics and the stability of states of the model using XPPAUT [55]. For the 2-dimensional reduced model we examined the phase portraits (Figure 2.2), from which the fixed points of the system and their stability were determined. Through this analysis, a range of biologically plausible parameters were determined which generate persistent working memory pattern types with statistics in the range typical of real cortical neurons (Table 1). These network parameters were then used in the full 3-dimensional model with dynamic calcium. For the 3-dimensional model, the fixed points of the system were first examined to determine coincidence with the 2-D model. Simulated working memory tasks were then run, varying the magnitude of the facilitation, self connectivity, and the magnitude of the input current. The firing rate patterns and frequencies exhibited by the model were compared to the types of patterns and frequencies observed in the database of parietal and prefrontal neurons recorded from monkeys during performance of working memory tasks in other studies [19, 73, 213]. A simulated trial of a working memory task followed approximately the same generic sequence

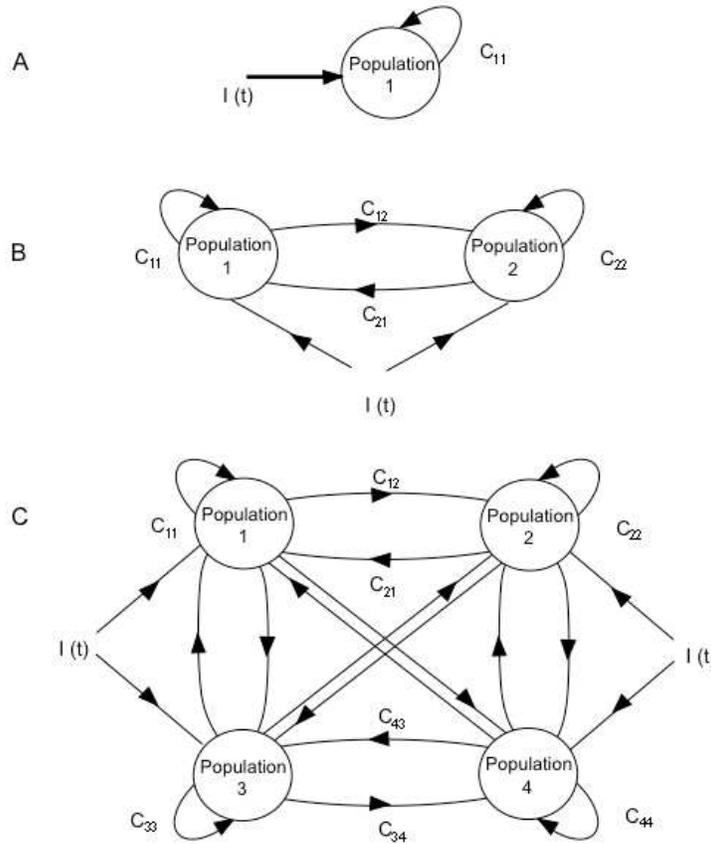


Figure 2.1: **Schematic diagrams of network architecture of the single population, two population, and four population models.** A) Singlepopulation firing rate model. In the single-population model the network receives input from an external current $I(t)$ during the cue period of a simulated working memory task. The synaptic activity $S(t)$, and consequently the firing rate increases from the input of the external current, resulting in a dynamic increase in its effective self connectivity C_{11} as a result of a concomitant dynamic change in synaptic facilitation W_1 . B) 2-population distributed model. In the distributed 2-population model, 2 local networks, 1 and 2 are connected recurrently by long-range projections C_{12} and C_{21} whose strength is weaker than the populations self connectivity C_{11} and C_{22} . Both populations receive input from an external current $I(t)$ during the cue period of simulated working memory task. This results in changes in their respective population firing frequencies along with changes in their effective long-range projections and self couplings through dynamic facilitation. In the spiking unit version of the model, populations 1 and 2 consist of networks of 100 (or 1000) spiking units each with all-to-all connectivity. C) 4-population model. In the 4-population model, four local networks (1, 2, 3, and 4) all receive input from an external current $I(t)$ during the cue period of the simulated working memory task. Each local population receives input from self connectivity, and weaker input from the long-range projections from the other populations. Long-range projections between the population pair 1 and 2, and between the pair 3 and 4 (i.e. C_{12} , C_{21} , C_{34} , and C_{43}) are stronger than the connection strength between populations 1 and 3 or 4, or between 2 and 3 or 4. In the spiking unit version of the model with 200 units, each population corresponds to a network of 50 spiking units with all to all connectivity. In the spiking unit version of the model with 2000 units, each population corresponds to a network of 500 units with all to all connectivity. The strength of the connections is scaled such that the total strength of connectivity to units is the same as the 200 unit network.

as that for which the single neuron database was acquired. The simulated task consisted of a 20-second baseline period (during which the population was in a baseline firing fixed point), followed by a 300-ms sample cue period corresponding to the period during which a memorandum was presented. After the cue period, a 12 second delay period followed. We did not consider in this study a behavior/motor response period following the delay, but rather restricted our analysis to the network behavior during these first 3 temporal aspects of the working memory task. To analyze the firing rate patterns and firing rate statistics of the model, peristimulus time (PSTH) histograms were generated and analyzed over a range of values of the self connectivity and facilitation. In addition, the phase diagram of different possible pattern states occurring over a range of values of the self connectivity and maximum synaptic facilitation were examined. For the single population model, PSTH histograms and phase diagrams were also generated for different values of dynamic synaptic depression and self connectivity, and the resulting patterns and firing rates were analyzed. Synaptic depression was incorporated by allowing the parameter for maximum facilitation (w_{max}) to range over values less than the value of the baseline facilitation (w_{min}) in equation 2.3.

For the distributed 2-population model, the parameters used were within the ranges determined and used in the single population model, and simulated working memory trials were conducted following the same course as that used for the single population model. Firing rate patterns and statistics of both populations were analyzed over a range of the inter-population connectivity values. PSTH histograms were generated to analyze the firing rate patterns and statistics. Phase diagrams of the firing rate patterns occurring in each population were generated as a function of the inter-population connectivity strength. An analysis of the behavior of the entire network was carried out through an examination of the possible pattern types occurring concomitantly in the two populations. This analysis was carried out by examining overlapping patterns in the phase diagrams of the two populations. For the distributed 4-population model, the parameters for each population were within small ranges of those determined and used in the preceding single- and 2-population models. Simulated working memory trials were conducted following the same previous course as well. Firing rate patterns and statistics occurring in all four populations were analyzed and compared to the activity of the real parietal and prefrontal neurons. Phase diagrams were generated of

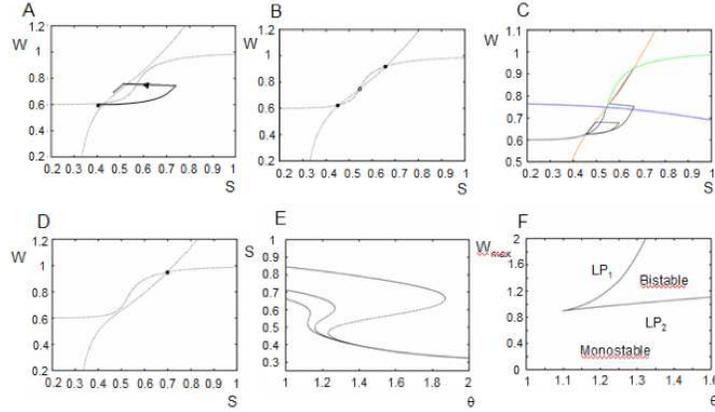


Figure 2.2: **Phase portraits and bifurcation diagrams of the single population reduced model obtained for different values of the self connectivity C_{11} and the maximum facilitation w_{max}** . The reduced 2-dimensional model is obtained at steady state calcium concentration ($dCa/dt = 0$). Shown in the phaseplane are the nullclines for Facilitation (w) and synaptic activity (s). A) Nullclines with self connectivity $C_{11} = 4.8$. The stable nodes of the system are where the nullclines cross. Here there is a single fixed point (solid black dot) corresponding to the baseline state. The trajectory of the system during a working memory task is indicated by the black line with an arrow. During the sample period, the applied current $I(t)$ raises the synaptic activity, and firing rate, with a concomitant rise in the facilitation (approximately 33% increase). After a 300 ms cue period, $I(t)$ becomes 0, and synaptic activity rapidly decreases towards the facilitation (w) nullcline, and the bottleneck. Because of the bottleneck, the trajectory then returns very slowly along the path of the facilitation nullcline towards the stable baseline state. In the present example, the synaptic activity and consequently the firing rate is still above the stable baseline rate at the end of the 10-second delay period, as indicated by the termination of the trajectory line. The system therefore maintains an increase from its baseline firing rate for the duration of the delay period. A continuum of elevated delay rates occur for different values of the parameters. B) w and s nullclines of the single population network with $C_{11} = 5$. Note there are now 3 fixed points where the nullclines cross: 2 attracting (black dots) and one saddle (white dot). C) Possible system trajectories of the network in (B) with $C_{11} = 5$. The saddle separatrix is indicated by the blue line. Two different possible trajectories of the system are shown which result from varying the magnitude of the external current $I(t)$. For $I(t) = 0.5$ the trajectory does not cross the saddle separatrix and stays in the basin of attraction of the baseline state. In this case, given sufficient time the system returns to the baseline stable state. The resulting pattern of behavior is that of cuecoupled or decaying memory cells. For $I(t) = 0.75$ the trajectory cross the saddle separatrix, entering the basin of attraction of the higher firing rate stable state. The system for this trajectory is shown to be in the higher firing rate state by the end of the delay period. The resulting pattern of behavior is that of response-coupled or ramping cells. In both cases the bottleneck can be adjusted such that the rate at which the system returns to the baseline state, or approaches the higher firing state can be arbitrarily slow, resulting in a continuum of different average firing frequencies during the delay, and apparent bistability at frequencies between the two stable fixed points. D) Phase portrait with $C_{11} = 5.2$. The system again possesses a single stable state at a higher rate than in (A). E) Bifurcation diagram of steady state synaptic activity as a function of maximum facilitation (W_{max}). Curves shown are for W_{max} equal to 0.925, 1, and 1.25 (producing synaptic facilitation in the 3060% range). Solid lines indicate stable fixed points, hashed lines are unstable fixed points. Note that 3 fixed points are present over a wide range of the facilitation. F) Bifurcation diagram for the parameters of maximum facilitation w_{max} and the threshold θ . The curves correspond to the limit points at different values of these parameters. The two branches of the limit points meet at a cusp point. For the region interior to the two curves there are three rest states and bistability, and outside them there is a single rest state and monostability.

Table 1: Model parameter values obtained from analysis of the reduced 2-dimensional network. These parameters are used in all of the local and distributed networks.

Network Parameters	
Mean w_{max}	1
Std. Dev. w_{max}	0.02
w_{min}	0.6
γ	8
τ_w	2
τ_s	0.05
τ_{Ca}	0.5
β	0.5
s_{min}	0.3
e	6
Ca_0	82
Ca_{min}	8
amp	0.2
θ	1.2
τ_m	0.03

Table 2: The mean values of connection strength between populations in the 2- population models. These values were used in the firing rate model, and also represent the mean connectivity strengths between units of different populations in the spiking model. The standard deviations for each of the mean connection strengths are given in parenthesis.

Mean connectivity, 2 populations		
	from A	from B
to A	0.05 (0.01)	-0.00134 (0.001)
to B	-0.0145 (0.04)	-0.057 (0.005)

the different firing rate patterns occurring in the populations as a function of different inter-population connectivities. An analysis of the behavior of the entire network was carried out through an examination of the possible different pattern types occurring concomitantly in the different populations. This analysis was carried out through an examination of overlapping states in the phase diagrams of the four populations. Resulting behaviors were compared with that of the 2-population model. Having determined the dynamics through the study of the firing rate models, spiking models were generated to make direct comparison with the single unit data. Spiking model versions of the 2- and 4-population firing rate models were generated by replacing the populations activities first with networks of 200 spiking units exhibiting the same overall mean firing rates. The network consisted of spiking neurons with all-to-all connectivity and random strengths. Connections between populations were both excitatory and inhibitory. Specifically, mean connectivity values were chosen from regions of the phase diagrams of the mean field model in which the range of memory cell pattern types were robustly exhibited. These connectivity values were then used as the values for setting the mean of the connectivity in the spiking model. Distributed spiking networks consisting of two populations of 100 units each, and four populations of 50 units each were generated with the average inter-area connectivity chosen to match the mean field model values within ranges of the standard deviation (Tables 2 and 3).

The spiking activity of the single units was modeled as theta neurons [54,96]. Unit firing frequency as a function of the injected current (F-I curve) can be obtained analytically in the theta model. This F-I curve is a square root function which provides a correspondence between the firing rate model and the theta model. The F-I curve for the theta model is

Table 3: The mean values of connection strength between populations in the 4- population models. These values were used in the firing rate model, and also represent the mean connectivity strengths between units of different populations in the spiking model. The standard deviations for each of the mean connection strengths are given in parenthesis.

Mean connectivity, 4 populations				
	from A	from B	from C	from D
to A	0.101 (0.004)	20.0012 (0.003)	20.0014 (0.001)	0.0005 (0.001)
to B	20.0462 (0.003)	0.1405 (0.003)	0.0011 (0.0001)	0.0002 (0.003)
to C	20.0003 (0.003)	20.0001 (0.003)	0.1025 (0.004)	20.0034 (0.004)
to D	0.0003 (0.003)	0 (0.001)	20.036 -0.036 (0.003)	0.132 (0.003)

described by

$$f(I) = \frac{1}{\pi} \sqrt{I}$$

whereas the curve for the mean field firing rate model is

$$f(I) = \frac{1}{\pi \tau_m} \sqrt{\frac{I - \theta}{1 - e^{-b(I-\theta)}}}$$

thus obtaining the correspondence between the mean field and spiking models (as the parameter b goes to infinity the above expression becomes the F-I curve for the theta model with an additional threshold). The membrane potential dynamics of a unit in the spiking model is given by the equation

$$\frac{dx_j}{\tau_m} = \left((1 - \cos(x_j)) + \left(\sum_{k=1}^{200} C_{kj} w_k s_k + I(t) + amp \sigma_j - \theta_j \right) (1 + \cos(x_j)) \right) \quad (2.8)$$

where $I(t)$ is an external current occurring during the presentation of memoranda, and amp is the amplitude of the Wiener noise. The synaptic activity of a unit s_j in equation 2.8 increases with each afferent spike according to

$$\frac{ds_j}{dt} = \frac{s_{min} - s_j}{\tau_s} + \beta \sum_m \delta(t - t_j^m) \quad (2.9)$$

where β corresponds to the increase in synaptic activity from a single afferent spike, t_m^j is the time of incidence of the m^{th} afferent spike on the j^{th} neuron, and τ_s is the decay constant. The dynamics of the synaptic facilitation w_j in equation 2.8 is given by

$$\frac{dw_j}{dt} = \frac{w_{min} - w_j}{\tau_w} + \gamma(w_{max} - w_j) \left(\frac{Ca_j}{Ca_0} \right)^e \quad (2.10)$$

where Ca corresponds to the intracellular calcium concentration which modulates the change in facilitation and increases with each spike according to

$$\frac{dCa_j}{dt} = \frac{1}{\tau_{Ca}} \left(-Ca_j + \varepsilon \sum_m \delta(t - t_j^m) \right) \quad (2.11)$$

The 2-population spiking network consisted of two “local” networks of 100 neurons each with all-to-all connectivity, and with average weaker recurrent connectivity between populations than within the populations. The activation properties of each individual network reflect that of the single populations of the firing rate models.

For the 2- and 4-population spiking networks, working memory task simulations were conducted similarly to those for the firing rate model, and the firing rate patterns and statistics were analyzed. During the baseline period of the simulated working memory task, facilitation in the models was kept low such that the firing rate of the populations was near the baseline fixed-point attractor state inherent in the model (as determined from the phaseplane analysis of the firing rate model). After 20 seconds, the baseline period ended and an external current $I(t)$ was applied for 300 ms. The external current raises the firing rate of many units in the populations, simulating the activity observed during presentation of the memorandum in working memory tasks. The current input and increased firing rate triggers dynamic facilitation through equations 2.9,2.10,2.11. After the cue period, the delay period begins. For the spiking model simulations, unit activity was analyzed over an 11-second delay period which is proportional to the delay period of the working memory tasks during which the parietal and prefrontal cells of the database were recorded. PSTH histograms of units were generated to analyze the patterns and firing rate statistics of the units. Average PSTH histograms were generated for each unit over 10 simulated working memory task trials. Pattern types appearing in the average PSTH histograms were determined and the distribution of patterns in the network were compared to the distribution of patterns observed

in the parietal and prefrontal neuron populations of the database. Variability in working memory patterns occurring across trials for each unit was analyzed and compared between the 2- and 4- population networks and the neuronal populations.

To examine the effect of network size on patterns exhibited in the networks across trials and their variability, we generated 2- and 4-population networks consisting of 2000 spiking units. For these networks the distribution of pattern types exhibited on each of 20 simulated working memory task trials was obtained and the average distribution across all 20 trials was determined. These distributions were compared to the distributions obtained with the 200 unit networks as well as that observed in the parietal and prefrontal neuron populations of the database. Variability in firing rate within trials was determined through an analysis of the coefficient of variation (CV) of the ISIs during the baseline and delay periods. Variability in working memory patterns occurring across trials for each unit was analyzed and compared to that observed in the 2- and 4-population networks of 200 units.

The database with which the different models activity is compared consists of 812 neurons recorded extracellularly from the parietal cortex (Brodmann areas 2, 3, 5, 7) and prefrontal cortex (areas 6, 8, 9 and 46) of monkeys performing working memory tasks. In parietal cortex, 521 cells were recorded from monkeys during performance of a haptic delayed matching-tosample task [213], and in prefrontal cortex, 291 neurons were recorded from monkeys during the performance of a cross-modal audiovisual delayed-response task [19, 73]. The analysis of this database and the compilation of its statistics in terms of firing rates, patterns and statistics have been presented elsewhere [171].

2.3 RESULTS

2.3.1 Single Population Model

The reduced single-population model is 2-dimensional (in the variables s and w for synaptic activity and facilitation respectively) and thus phaseplane and numerical analysis was carried out. The nullclines of the system (Figure 2.2) correspond to the curves along which the

synaptic activity and facilitation are constant ($ds/dt = dw/dt = 0$). The steady states of the system are defined by the points at which these 2 curves intersect. For sufficiently low self connectivity strengths (or low maximum facilitation), only one such point is present, corresponding to the baseline firing rate of the population (Figure 2.2A). Stability analysis reveals this is an attracting fixed point. Thus transient perturbations from the external current during the sample period (resulting in increased synaptic activity, firing rates and facilitation) ultimately relax back to this state. As the self connection strength (or amount of facilitation for a given input current) is increased, the W-nullcline intersects the S-nullcline at 3 points (Figure 2.2B). Stability analysis reveals that 2 of these nodes are attracting fixed points, and one is a saddle node. The presence of a stable state corresponding to baseline, and a second stable state corresponding to an above baseline firing rate, enables bistable behavior, although the difference in firing rates associated with these states is much larger than typically observed in cortical data over much of the parameter space. For example, in the subpopulation of memory cells recorded from the parietal cortex, 90.1% of the cells exhibited increases from baseline to delay of less than 10 Hz, and 69.9% of frequency changes were less than 5 Hz. In the subpopulation of memory cells recorded from prefrontal cortex 100% exhibited increases of less than 10 Hz, and 95.2% were less than 4 Hz. Persistent elevated firing rates within these ranges typically observed in cortical data is inherently prevalent in the model without incorporating many of the previous mechanisms providing solutions for acquiring that behavior (see for example Latham and Nirenberg, [122]; Barak and Tsodyks [13]; Mongillo et al., [142]). An essential feature of the model allowing this behavior is the presence of a bottleneck which appears in the phaseplane near the s and w nullclines corresponding to regions of greatly diminished rates of change for the dynamic variables. Further, a bottleneck is present over a broad range of the parameter space, and its presence is not dependent on fine tuning of parameters. The bottleneck comes about because the equations of the system are a continuous map approaching zero when the nullclines are close to each other in phase space. Thus as shown in Figure 2.2, the values of ds/dt and dw/dt are reduced in those areas. There are two factors in the decay rate which include the bottleneck and the value of the time constants. While the presence of a bottleneck is not a result of the difference in time constants (but rather the shape of the nullclines), in

the present system the shape of the nullclines has a dependence on the value of the time constants, and thus the slower rate change in w than s contributes both in terms of the bottlenecks existence via nullcline shape, as well as acting to slow decay of its own accord. Thus both contribute to the slower decay. Because of the bottleneck however, the “effective time constant” or rate of decay is much slower than would be predicted from the actual time constants. When the firing rate is elevated above baseline by an external current during the cue period, facilitation also increases. The trajectory in the phaseplane is such that passing through the bottleneck, the return of the system to the baseline stable state (or procession to the higher firing rate attractor state) is “impeded”. Thus while not in a stable state, the system remains in a state of elevated (above baseline) firing frequency for an extended period of time, which can be virtually indefinite. Over a wide range of values of the parameter space, the decay to one of the stable states of the system is sufficiently slow such that no significant change in elevated firing rate is observed for the duration of the putative memory period. From the frame of reference of the memory task, this activity appears as bistable. In contrast to actual bistability of the model however, the difference in firing rates between baseline and delay periods for this apparent bistability can adopt a continuum of values within the range typically observed in real cortical cells (i.e. differences between baseline and delay rate less than 100% of baseline or typically < 5 Hz). This is also true of both the decaying memory cell and ramping cell behavior (the ramping up behaviors may occur in cells inhibited by cells ramping down, or by entering the basin of attraction of the higher fixed point). The range of facilitation values over which the different firing rate behaviors occur is affected by the value of maximum facilitation parameter. This parameter (and the threshold θ) significantly determines the fixed points of the system (where the facilitation nullcline intersects the synaptic activity nullcline). While this parameter is varied over a large percentile range (i.e. several hundred percent), the resulting change in actual facilitation realized by the network is within the range of 10% to 60%—within the range of reported increases (for example see Hempel et al. [94]).

We next analyze the stability of the attractor states of the network as a function of the threshold parameter (Figure 2.2E,F). The bifurcation diagram reveals that the 3 fixed points of the system (two stable nodes and one saddle point) are present over a wide range

of this parameter. For parameter values where the trajectory remains below the stable manifold (within the basin of attraction of the baseline node), the system ultimately returns to that attractor state. For parameter values in which the trajectory travels above the stable manifold (into the basin of attraction of the stable node corresponding to the higher firing rate), the system approaches that second stable state.

Having determined the states of the 2-dimensional system, we next use the parameters of this 2-dimensional network (Table 1) in the 3-dimensional model with dynamic calcium. An analysis of the attractor states reveals that the 3-dimensional system retains the same attractors as the 2-dimensional system. Figure 2.3 shows PSTH histograms of the model during the simulated working memory task. These histograms show that the activity patterns observed in the trajectories of the phase portraits of the 2-dimensional model are present. Particularly this analysis shows that the network inherently exhibits the range of excitatory and inhibitory patterns correlated with working memory for different values of facilitation and self-connectivity. For values of facilitation (or input strength and/or duration) such that the trajectory of the system stays below the separatrix in the basin of attraction of the baseline attractor state, the population can exhibit persistent activity which decays towards baseline throughout the delay (Figure 2.3A). The achievable increases in firing rate from baseline to delay can take on low values (i.e. <5 Hz) and can adopt any rate within a continuum. The rate of decay of the persistent activation towards baseline is also variable along a continuum. For a range of trajectories and parameter values, the decay in firing rate during the delay can become slower and slower, to the point that the population approaches for all intents and purposes bistable behavior (Figure 2.3B). Once again, the increase in firing frequency during the delay can occur along a continuum. For sufficiently low values of facilitation, the population exhibits a non-responsive pattern (Figure 2.3C). That is the trajectory returns to baseline firing rates immediately following the cue. Thus the population responds to working memory events (i.e. the cue), but exhibits baseline rates throughout the delay. As facilitation increases (or the input strength/duration increases) such that the trajectory proceeds beyond the separatrix, the pattern becomes that of a ramping increase of firing rate during the delay period (Figure 2.3D). For sufficiently large values of facilitation, the system quickly adopts the higher firing rate attractor state, thus exhibiting bistability

(with differences between baseline and delay period >10 Hz for most parameter values). Dynamic synaptic depression can be introduced into the network rather than facilitation for values of maximum facilitation (w_{max}) that are less the background value (w_{min}) in equation 2.3. For sufficiently small values of synaptic depression, as was the case for facilitation, the network exhibits the non-responsive pattern. As synaptic depression is increased the network exhibits the inhibitory pattern which is the mirror image of decaying memory cells (Figure 2.3E). As was the case for the excitatory memory cells, the decay back to the baseline state of the inhibitory pattern can be slowed to the point where the population exhibits an apparent fixed rate inhibition pattern (Figure 2.3F).

Figure 2.4 shows the phase diagram of the firing rate patterns of this model as a function of facilitation strength and self connectivity strength. It can be seen that the different excitatory patterns observed in the cortical data are produced (decaying memory, bistable memory, and ramping cells) over a wide range of parameters. As is the case in the cortical data, the non-responsive pattern behavior is the most prominent pattern type across the parameter space. For the case of synaptic depression, the decaying inhibition pattern occurs prominently over a range of parameters in addition to the non-responsive pattern. Notice that the properties of this phase diagram could change depending on other parameters, and on the properties of the stimulus.

The mechanism for the behaviors illustrated can be understood from the stability analysis and examination of the phaseplane (Figure 2.2). In all cases the firing rate begins at the lower attractor state with low values of facilitation. The input of current increases the synaptic activity s , and therefore firing rate and subsequently the facilitation w increases. If the self connectivity, facilitation or magnitude of the external current is sufficiently low such that the trajectory of the system does not cross the saddle separatrix, the trajectory is such that s quickly decreases until it approaches its nullcline. Here the trajectory proceeds such that it approaches the baseline stable attractor along the path of that nullcline. However, the bottleneck through which the trajectory proceeds slows the rate of return to the baseline state. The bottleneck can slow that rate such that the trajectory is impeded to the point that firing rate appears bistable with respect to the duration of the memory period of the task. As self-connection strength, synaptic facilitation or external current magnitude is raised

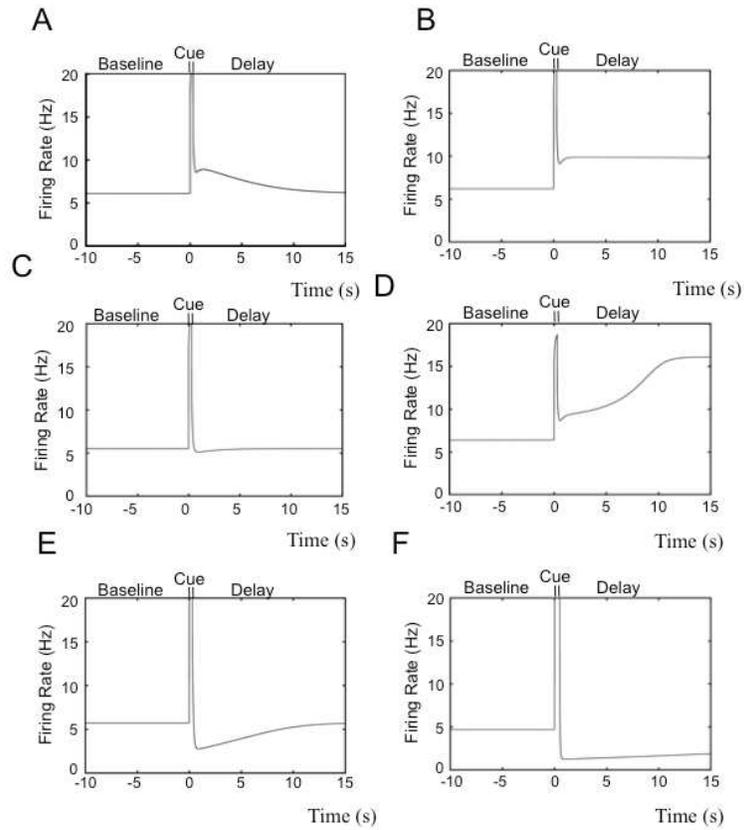


Figure 2.3: PSTH histograms of the single population firing rate model during the simulated working memory task. Different patterns correlated with working memory are obtained for different values of the facilitation (or depression) and the magnitude of the input current during the cue period. A) Decaying memory cell behavior. The baseline frequency is approximately 6 Hz, and at the beginning of the delay period the firing rate is approximately 9 Hz. The firing rate has returned to the baseline level by the end of the delay period. Average increase in firing rate from baseline to delay period is approximately 1.5 Hz. B) Apparent bistable memory cell behavior. The delay firing rate does not correspond to a stable state, but the decay towards the baseline state is sufficiently slow so that no significant decrease in firing rate occurs by the end of the period. The baseline firing rate is approximately 6 Hz, and the delay period firing rate is approximately 10 Hz. C) Nonresponsive pattern behavior. The population responds during the presentation of the cue, but immediately returns to and maintains the baseline firing rate during the delay. D) Ramping response-coupled cell behavior. The baseline frequency is approximately 6 Hz, and at the beginning of the delay period the firing rate is approximately 9 Hz. The firing rate adopts the rate of the higher fixed point attractor by the end of the delay period. Average increase in firing rate from baseline to delay period is approximately 6.5 Hz. E) Decaying inhibited pattern. In this example the value of maximum facilitation is less than the background facilitation, resulting in dynamic synaptic depression. The baseline frequency is approximately 6 Hz, and at the beginning of the delay period the firing rate has decreased to approximately 2.5 Hz. The firing rate has returned to the baseline rate by the end of the delay period, mirroring the decaying memory cell activity observed in (A). Average decrease in firing rate from baseline to delay period is approximately 1.75 Hz. F) Apparent stable inhibition. The delay firing rate does not correspond to a stable state, but the decay towards the baseline state is sufficiently slow so that no significant decrease in firing rate occurs by the end of the period. The baseline rate is approximately 4.5 Hz, and the delay period firing rate is approximately 1 Hz. Every pattern observed in the cortical database as reported in [5] is observed with the exception of delay inhibition that increases throughout the delay period.

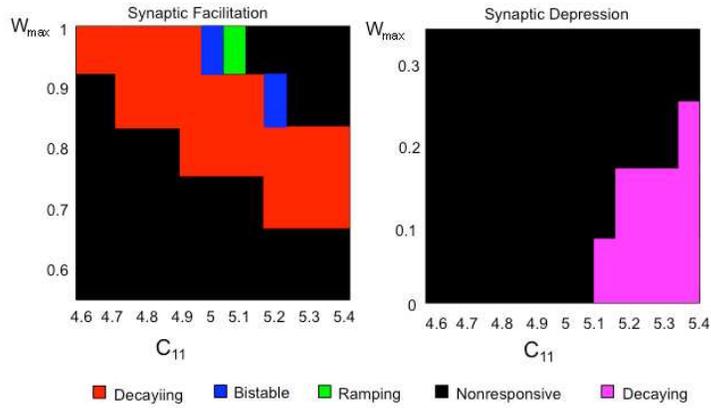


Figure 2.4: **Phase Diagram of different patterns exhibited by the single population model as a function of the maximum facilitation parameter and self-connectivity.** Left: Patterns exhibited with dynamic synaptic facilitation. Actual facilitation is less than the maximum facilitation parameter and falls within a physiological reasonable range (1060%). Right: Patterns exhibited with dynamic synaptic depression which takes place when values of W_{max} are less than the baseline level (W_{min}). Pattern activity was determined by examining the firing rate activity during 3 periods of the working memory task: the first 5 seconds of the delay period (D1) immediately following the cue, the second 5 seconds of the delay period (D2), and the last 5 seconds of the baseline period (B) immediately preceding the cue period. Significant differences were taken to be present if the absolute difference between any two periods was greater or equal to 0.5 Hz, which is approximately the lower limit of significant differences observed in the real cortical parietal and prefrontal cells. Fixed rate memory cell behavior consisted of activity in which the absolute difference in average firing rate between D1 and D2 was less the 0.5 Hz, while both those periods exhibited an average firing rate greater or equal to 0.5 Hz above the baseline rate. Ramping cue-coupled cell behavior consisted of activity in which D2 exhibited an average firing rate greater or equal to 0.5 Hz above that of D1, and the firing rate of D1 was greater or equal to that of B. Decaying memory cell behavior consisted of activity in which D1 exhibited an average firing rate greater or equal to 0.5 Hz above both B and D1. Nonresponsive cell behavior consisted of activity in which the difference between any of the periods (B, D1, and D2) was less than 0.5 Hz. Decaying inhibition cell behavior consisted of activity in which the average firing rate during D1 was at least 0.5 Hz less than that exhibited during B and D2.

beyond a critical point such that the systems trajectory goes beyond the saddle separatrix in the phaseplane, the system approaches the second stable state which corresponds to an above baseline firing rate resulting in ramping or response-coupled cell pattern behavior. Once again the rate of this increase is affected by the bottleneck, and may be arbitrarily slowed such that the firing rate appears bistable with respect to duration of the memory period of the task. This phenomenon exists for a broad range of parameter values. Thus the inherent bistability in these cases is critical in modulating patterned memory behavior, but does not in many cases in and of itself represent the memory states. Rather the activation of the network itself could represent active working memory.

2.3.2 2-Population Firing Rate Model

We next analyze the behavior of a 2-population network, recurrently connecting two of the single populations. This model is 6-dimensional and thus cannot be easily reduced and analyzed as was the case for the single population model. We analyze the patterns and statistics through the PSTH histograms (Figure 2.5) and the phase diagrams (Figure 2.6) of pattern types as a function of the strength and sign (i.e. excitatory or inhibitory) for the net effect of the inter-population projections. The phase diagrams enabled the examination of possible concomitant activities in the different networks.

Inhibition, in addition to excitation, is incorporated in the mean field 2-population model via the inter-area projections between populations. While long-range projections in the cortex are excitatory, inhibition is examined as well according to the assumption that the majority of the long-range projections may project either to inhibitory or excitatory interneurons. Thus the net effect of these projections can be excitatory or inhibitory. We analyze the behavior of the network for different possible interpopulation connectivity schemes (i.e. excitatory-inhibitory (E-I), and inhibitory-inhibitory (I-I)). Slightly different values for self feedback connections strengths within the two populations were chosen.

As was the case for the single-population model the PSTH histograms reveal that the 2-population model exhibits the excitatory patterns of memory and decaying-rate or ramping cells with a continuum of rate differences. The inclusion of inhibitory connections results

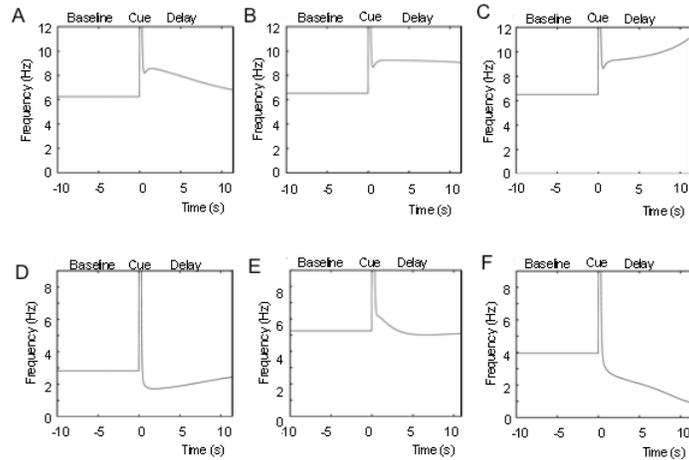


Figure 2.5: **PSTH histograms of activity exhibited by the 2-population firing rate model at different values of the inter-area connectivity.** All general patterns in the cortical data as reported in [5] are exhibited by the network, with frequencies in the range of the real data. Firing pattern behaviors shown are A) decaying memory cell, B) Stable memory cell, C) Set cell D) Decaying inhibition cell, E) Stable inhibition cell, and F) ramping inhibition cell.

in the presence of parameter ranges in which all of the inhibitory patterns (mirroring the excitatory ones) occur. These inhibitory patterns can occur purely as a function of inhibitory inter-population connectivity, without incorporating dynamic synaptic depression as was the case for the single population. In addition the inhibitory pattern of increasing inhibition throughout the delay (mirroring the excitatory ramping cells) which was absent in the single population model, now can occur (Figure 2.5F).

In the phase diagrams of the 2-populations (Figure 2.6A) it can be seen that all of the patterns of memory behavior occur over broad ranges of the parameters, and thus without fine tuning, in both populations. As in the single population model, the non-responsive type is the most prominently occurring pattern across the parameters, followed by decaying memory cells and ramping cells. Less commonly occurring types are fixed rate memory cells and the inhibitory mirror images of the excitatory patterns. While all the patterns occur over a broad range of parameters, the specific patterns present over given ranges varied considerably between populations. Thus many specific complementary patterned activities occurred simultaneously in both populations only over small ranges of the parameters, and thus some degree of fine tuning is necessary to achieve particular overall network behaviors.

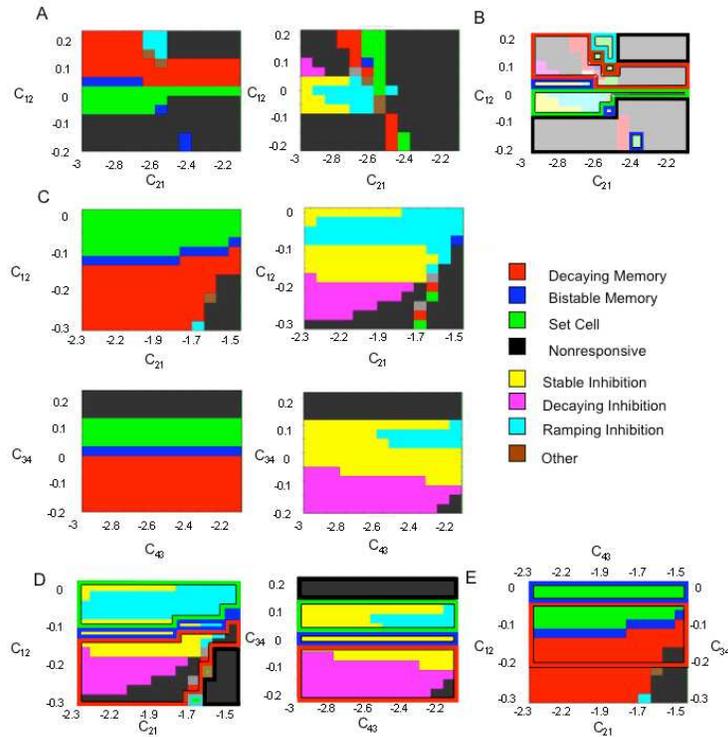


Figure 2.6: **Phase diagrams of the patterns exhibited by the 2- and 4-population firing rate models.** The phase diagrams were created using the same procedure carried out for the generation of the phase diagram of the single-population model. A) Phase diagrams of the activities of the 2-population model as a function of the inter-population connectivity. Left: Phase diagram of population 1. Right: Phase diagram of population 2. The self connectivity of population 2 (0.9) is slightly less than population 1 (1.0) so that the phase diagrams of the two populations are not identical. Excitatory-Inhibitory (E-I) and Inhibitory-Inhibitory (I-I) connectivity architectures between the populations are examined. The pattern category other corresponds to decaying or ramping cell behavior in which the average firing rate of one period of the delayD1 or D2 is significantly less than the baseline (i.e. at least 0.5 Hz less than the firing rate during the period B), while the other is significantly greater (i.e. at least 0.5 Hz greater than the firing rate during the period B). Note that all of the pattern types (excitatory, inhibitory and nonresponsive) are obtained over a wide range of the parameters. B) The overlapping simultaneously occurring patterns in the 2 population network across the range of inter-population connectivity values. The figure shows the 1 and 2 populations (left and right phase diagrams of panel A to the left) phase diagrams superimposed on each other. The 1-population states are shown as solid, and the corresponding 2-population states are shown in outline. Discrete regions of the overlapping phase diagrams reveal simultaneously occurring network behaviors observed in cortex during working memory. Included is simultaneous occurrence of decaying memory behavior in both populations (red-red overlap), and stable memory behavior in one population and ramping cell behavior in the other (blue-green overlap). Also common network behaviors observed are the simultaneous occurrence of the excitatory patterns in one population and the nonresponsive pattern in the other. Inhibitory patterns tend to occur in one population primarily with ramping cell behavior in the other. C) Phase diagrams of the activities of the 4-population model as a function of the inter-population connectivity. Upper left: Phase diagram of population 1. Upper right: Phase diagram of population 2. Lower Right: Phase diagram of population 3. Lower Right: Phase diagram of population 4. All the pattern types are exhibited by the network. However, the populations partition themselves such that they exhibit almost exclusively either excitatory or inhibitory memory patterns across the values of interpopulation connectivity. D) Simultaneously occurring patterns in populations 1 and 2, and in populations 3 and 4. E) Simultaneously occurring patterns in populations 1 and 3. The figure shows 1 and 3 (excitatory) populations phase diagrams superimposed on each other. Over a significant range of the parameters, memory cell behavior and ramping cell behavior co-occur (red and green overlap). The existence of such simultaneous population behavior has been indicated in prefrontal cortex. Memory cell behavior is also seen to simultaneously occur in multiple populations across a wide continuous range of the parameter space (red-red and redblue overlap).

For example, as can be seen in the overlapping phase diagrams (Figure 2.6B), attaining memory cell behavior simultaneously in both cortical locations, or attaining complementary cue-coupled/ response-coupled behavior, requires the connectivity of the network to be restricted to relatively small specific ranges of inter-population connectivity values.

2.3.3 4-Population Firing Rate Model

We next consider the effect on the states of the network when the model is extended to 4 populations. In the 4-population model all of the patterned activities continue to be present over a continuum range of increases and decreases in firing frequencies. However the distributed architecture results in a “specialization” of pattern activity within specific populations. As can be seen in the phase diagrams (Figures 2.6C,D) each local network (populations 1–4) exhibits the non-responsive pattern and almost exclusively either the excitatory or inhibitory memory patterns across the range of connectivity strengths. A result of this specialization or partitioning of pattern types between the local networks is that, in contrast to the 2-population model, simultaneous complementary pattern behaviors occur far more robustly across wide parameter ranges. Thus for example attaining memory cell behavior simultaneously in multiple cortical areas, or attaining complementary cue-coupled/response-coupled patterned behavior does not require fine tuning to a small restricted range of connectivity values (Figure 2.6E).

2.3.4 Spiking Unit Network: 2-Population Model

We next examine the statistics and dynamics of the spiking version of the distributed mean-field models. In the spiking network version of the 2-population mean field model we first replace the populations with two networks of 100 spiking units each, whose activity averaged across units approaches the activity of the populations in the mean field model (Figure 2.7). We first analyze the range of memory pattern types in the spiking networks during simulated working memory tasks. Average PSTH histograms over 20 simulated trials of a working memory task were generated and examined for each unit in the network (Figure 2.8). The pattern types and statistics in these units can be directly compared to those

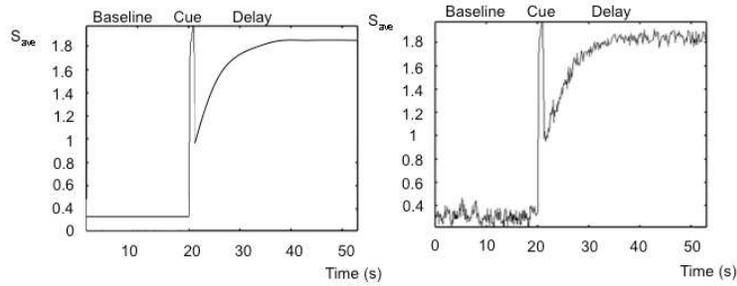


Figure 2.7: **PSTH histogram of the single-population firing rate model (left) and the corresponding average spiking unit activity in a network of 50 units obtained from the spiking model (right).** Note the general pattern is the same in both networks.

occurring in the database of real parietal and prefrontal cells. The results show that within the populations, the range of excitatory and inhibitory patterns occur, in addition to the non-responsive pattern. The average baseline frequencies, delay frequencies and deltas (changes in frequency from baseline to delay period) exhibited by the units for each pattern fall within ranges observed in the real parietal and prefrontal cells (Table 4). Figure 2.9A (left) shows the distribution of patterns types exhibited by all 200 units of the 2-population spiking network. The most commonly occurring pattern was the non-responsive pattern, followed by the excitatory patterns, and finally the inhibitory patterns. This relative distribution of pattern types is consistent with what is observed in both parietal and prefrontal cell populations (Figures 2.9B left and right). This distribution also correlates with the areas of the parameter space over which each of the patterns occurred in the phase diagrams of the 2-population model.

As is the case in real cortical cells, the specific pattern exhibited by a unit in the spiking network in any given trial can vary from the predominant pattern observed in the average PSTH histogram [171]. That is, the pattern that a unit (or real cortical cell) is classified as exhibiting, as determined from the average PSTH pattern, might not be exhibited on some subset of trials. This includes exhibiting different excitatory patterns from trial to trial in delay activated pattern cells, different inhibitory patterns from trial to trial in delay inhibited pattern cells, or even pattern types contrary to the average pattern. For example, the parietal delay activated cells exhibited delay inhibited patterns on 13.2% of the trials,

Table 4: Pattern numbers correspond to the following patterns: 1) Response-coupled ramping cell, 2) fixed rate memory cell, 3) cue-coupled (decaying) memory cell, 4) response-coupled ramping cell with no frequency difference between baseline and period D1, 5) Nonresponsive cells, 6) Decaying inhibition cells with no frequency difference between baseline and period D1, 7) decaying inhibited cells, 8) fixed rate delay inhibited cells, and 9) ramping inhibited cells. The approximate averages for each of these pattern types in the real parietal and prefrontal cell database, as determined in [5] are also indicated. Note the frequencies exhibited by the models are similar to those observed in the real cortical cells.

Average firing frequencies (Hz) and Delta values.									
Pattern	1	2	3	4	5	6	7	8	9
Baseline 2-Population Model	5.6	6.5	6.9	3.9	4.9	9.7	3.9	5.2	4.6
Baseline 4-Population Model	7.4	6.9	6.3	11.9	5.5	10.7	5.1	5.1	4.9
Baseline Parietal Neurons	14	16	14	17.5	12	11	17	15.0	10.5
Baseline Prefrontal Neurons	9.5	7.5	8	27	8.0	7.0	9.0	12.0	7.0
Delay 2-Population Model	7.0	8.9	8.7	3.8	5.0	10.2	2.8	2.7	3.1
Delay 4-Population Model	9.4	9.9	7.8	11.5	5.7	10.8	4.6	2.7	3.3
Delay Parietal Neurons	21.5	25	21	18	12.5	11	12.5	7.5	7.0
Delay Prefrontal Neurons	12.5	11.5	10.5	27.5	8.0	7.5	6.5	8.0	5.5
Delta 2-Population Model	1.4	2.4	1.8	20.1	0.1	0.5	21.1	22.5	21.5
Delta 4-Population Model	2.0	3.0	1.5	20.4	0.2	0.1	21.5	22.4	21.6
Delta Parietal Neurons	7.5	9.0	7.0	0.5	0.5	0.0	24.5	27.5	23.5
Delta Prefrontal Neurons	3.0	4.0	2.5	0.5	0.0	0.5	22.5	24.0	21.5

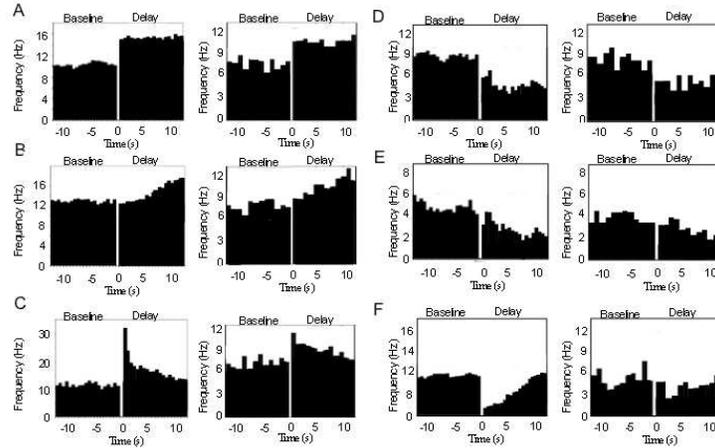


Figure 2.8: **Average histograms from example units of the 2-population spiking model showing the excitatory and inhibitory working memory patterns, and average histograms from representative single neuron recordings from the database of prefrontal and parietal neurons exhibiting the same general patterns.** PSTH histograms of spiking model units are shown in the left columns and real neuron histograms in the right columns. The gap between baseline and delay periods corresponds to the cue period (firing rate not shown). Histograms of the model units were averaged over 10 simulated trials of the working memory task. Patterns exhibited are A) Stable fixed rate memory (i.e. bistable) cell behavior, B) ramping cell behavior, C) decaying memory cell behavior, D) fixed rate inhibition cell behavior, E) ramping inhibition, and F) decaying inhibition.

and parietal delay inhibited cells exhibited delay activated patterns on 15.6% of the trials. Prefrontal delay activated cells exhibited delay inhibited firing patterns on 16.1% of the trials, and prefrontal delay inhibited cells exhibited delay activated patterns on 20.1% of the trials. Figure 2.9C (left) indicates for each unit in the 2- population network the percentage of the total number of simulated trials in which its pattern behavior differed from its dominant pattern type appearing in its average PSTH histogram (different excitatory or inhibitory pattern and/or contrary delay activity). We see that units in the network, on average, exhibit pattern activity different than their classified pattern type on approximately 52.5% of the trials, with that variability being approximately the same in both local networks. Individual units exhibited different patterns over a range from 20% to 70% of the trials. This variability is comparable to that observed in many real neurons during working memory tasks.

2.3.5 Spiking Unit Network: 4-Population Model

We next examine the range of statistic and memory pattern types occurring in the activity of units in a spiking network with four populations of 50 neurons each. In the spiking network we replace the four populations of the mean field model with four networks of 50 spiking units each whose activity averaged across units is the same as the activity of the populations of the mean field model. We first analyze the range of memory pattern types in the spiking networks during simulated working memory tasks. Average PSTH histograms over 20 simulated trials of a working memory task were generated and examined for each unit in the network. As was the case in the 2-network spiking model, the range of excitatory and inhibitory patterns, in addition to the nonresponsive pattern are exhibited by the units in the network. The specific baseline frequencies, delay frequencies and deltas (changes in frequency from baseline to delay period) exhibited by the units for each pattern fall within the ranges observed in the real parietal and prefrontal cells (Table 4). Figure 2.9A (right) shows the distribution of patterns types exhibited by all 200 units of the 4-population spiking network. The relative prominence of different pattern types is similar to that of the 2 population spiking model, with the most commonly occurring pattern being the nonresponsive pattern, followed by the excitatory patterns, and finally the inhibitory patterns. Once again this is consistent with the relative percentages of each pattern type observed for the real parietal and prefrontal neurons.

As was the case in the 2-population spiking model, the specific pattern exhibited by a unit in the 4-population spiking network in any given trial of the simulated working memory task can vary from the predominant pattern observed in the average PSTH histogram (Figure 2.10). Figure 2.9C (right) indicates for each unit in the network the percentage of the total number of trials in which its pattern behavior differed from its dominant pattern type. We see that units in the network, on average, exhibit pattern activity different than their average classified pattern type approximately 54% of the trials, with individual units exhibiting different patterns over a range of 25% to 70% of trials. This variability is very similar to that observed in the 2-population spiking model, and once again is within the range observed in the real cortical neurons. However, in contrast to the 2-population network, variability is

not uniformly distributed across populations. In the 4 population network there is a greater degree of partitioning of the activity of the networks into those primarily exhibiting non-responsive and excitatory patterns, and non-responsive and inhibitory patterns. Populations of units of primarily excitatory or the non-responsive memory pattern types exhibit their predominant pattern much more reliably than those populations of units of primarily inhibitory and non-responsive pattern types. Thus the more distributed architecture resulted in increased reliability in persistently active populations.

2.3.6 Spiking Unit Network's Variability Scaling With Population Size

We next examine the dependence of pattern type, firing rate statistics and variability as a function of population size. To do this we produced a 2- and 4-population spiking model as above consisting of 2000 units. We first analyze the range of memory pattern types in the spiking networks during simulated working memory tasks. Average PSTH histograms over 20 simulated trials of a working memory task were generated and examined for each unit in the network. As was the case in the 2- and 4 population spiking networks consisting of 200 units, the range of excitatory and inhibitory patterns, in addition to the non-responsive pattern are exhibited by the units in the network (Figure 2.11). The specific baseline frequencies, delay frequencies and deltas (changes in frequency from baseline to delay period) exhibited by the units for each pattern fall within the ranges observed in the real parietal and prefrontal cells. Figure 2.12 shows the distribution of patterns types exhibited by all 2000 units of the 2- and 4-population spiking networks. The relative percentage of excitatory and inhibitory patterns is similar to that observed in the 200 unit networks with excitatory patterns being slightly more prominent than inhibitory patterns.

The firing rate model produces a trajectory in the phase plane which corresponds to a specific pattern type. Depending on the connections and other parameters, the stimulus causes the trajectory to remain above or below the separatrix of the phase space. In terms of the spiking model the firing rate model trajectory corresponds to the mean of the trajectory of all units. Depending on how close to the separatrix that mean trajectory is after the stimulus, fluctuations about the mean from various sources of stochasticity in the spiking model will

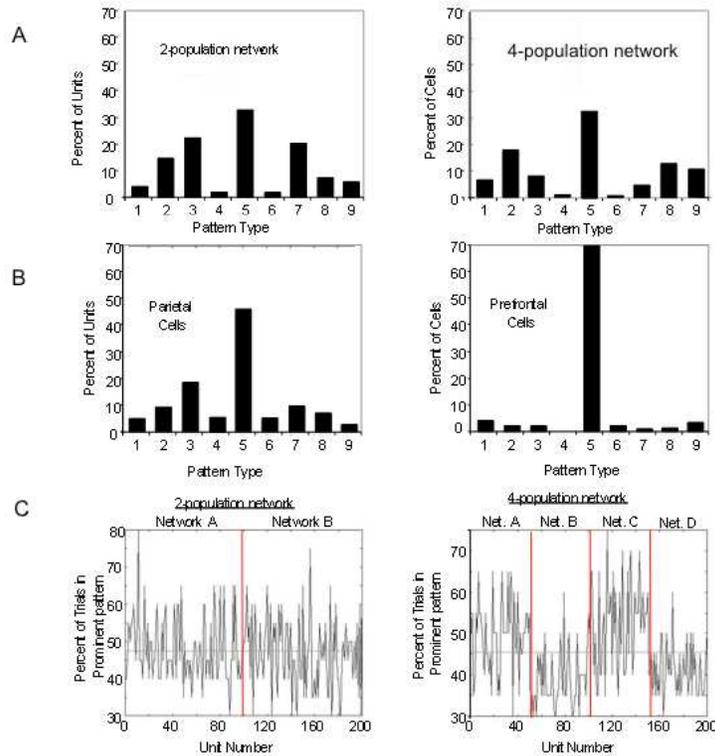


Figure 2.9: Distributions of units exhibiting the different pattern types in the spiking models and variability in pattern expressed from trial to trial in the units of those models. A) Distribution of pattern types exhibited by units in the 2- and 4-population spiking models (left and right respectively). Distributions correspond to the average number of units exhibiting each pattern type over 20 trials of the simulated working memory task. Pattern numbers correspond to the same patterns as labeled in table 4. Specifically 1) Ramping cells 2) fixed rate memory cells, 3) Decaying memory cells 4) ramping set cell with no frequency difference between D1 and baseline, 5) nonresponsive cell 6) ramping inhibited cell with no frequency difference between D1 and baseline, 7) decaying inhibition cell, 8) fixed rate delay inhibited cell, and 9) ramping delay inhibited cell B) Distribution of the percentage of parietal (left) and prefrontal (right) cells exhibiting each pattern type during the performance of unimodal and cross-modal working memory tasks. Note that the distribution of the real cells and the units in the spiking model are similar with nonresponsive cells being most common, excitatory memory cells (decaying, fixed rate, and ramping) being the next most common, and inhibitory patterns (decaying, fixed rate, and ramping) being the least common. Also patterns 4 and 6 (ramping excitatory and inhibitory cells with no frequency differences between D1 and baseline) occur infrequently in both the models and real cells. C) Consistency of pattern type expressed across trials for each of the 200 units in the 2-population spiking model (left) and the 4 population spiking model (right). Each unit was classified as displaying one of the pattern types in its average PSTH histogram (across 20 trials). The ordinate indicates the percentage of trials in which the dominant pattern (the pattern occurring most often across trials and is typically that which is observed in the averaged PSTH histogram) of the average was actually exhibited by the unit. In the 2-population graph, units 1100 are all members of one population, while units 101200 are part of the other population. In the 4-population graph, units 150, 51100, 101150, and 151200 are the members of populations 1 through 4 respectively. Note that the overall average variability in pattern expression is the same in the 2- and 4-population models, with the classified pattern type of each unit being displayed in approximately 47% of the trials, and an approximate range of 30–70% of trials showing different patterns on any given trial. However, while the variability is distributed evenly across both populations in the 2-population model, in the 4-population model the variability is significantly greater in units in those populations exhibiting primarily inhibitory patterns (populations 2 and 4).

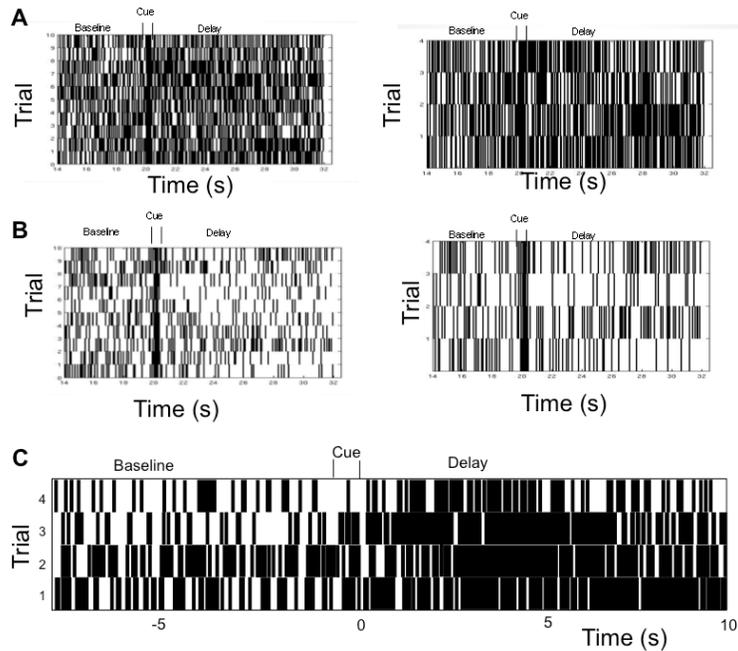


Figure 2.10: Raster plots of 10 trials (left) for an excitatory (A) and an inhibitory (B) unit from the 4-population 2000 unit spiking network. Note that different pattern types occur on different trials. Figures A and B (right) show an expanded version of 4 of the 10 trials from both units to highlight different patterns exhibited. From top to bottom of the expanded trials, the excitatory unit shows a decaying rate, nonresponsive, increasing rate, and persistent stable firing rate pattern, and the inhibitory unit shows decaying inhibition, stable inhibition, increasing rate excitation, and increasing inhibition. C) Rasters from 4 trials selected from a real neuron recorded from the prefrontal cortex during presentation of the same memorandum of the cross-modal task exhibiting stable persistent activation during the delay in the average PSTH. Although the cell exhibits stable persistent activation on average, on specific trials the cell exhibits the decaying rate activation pattern.

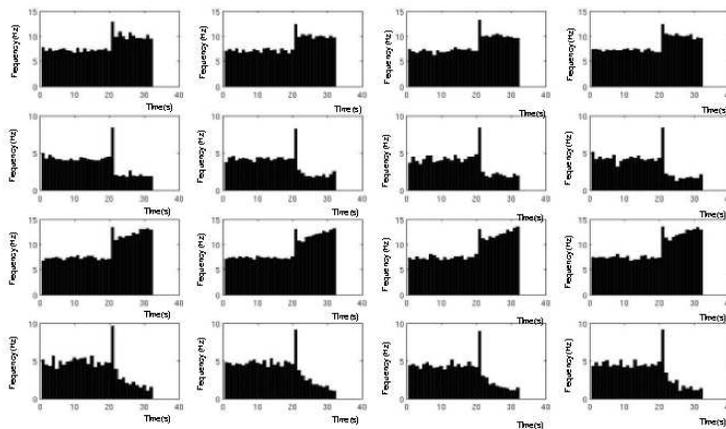


Figure 2.11: Average PSTH histograms of 16 units randomly chosen from the 2000 unit 4-population spiking model. Note that the entire range of pattern types continue to be exhibited (excitatory and inhibitory stable, decaying, and ramping delays and nonresponsive) with firing statistics similar to real neurons.

result in a probability that units will make transitions to trajectories corresponding to pattern types different than that of the mean trajectory. The resulting pattern types will have a distribution reflecting this. Conversely the closer the system is to one of the stable attractors of the system, the less probable it is for a given level of noise that the system trajectory will depart from the pattern of the mean trajectory.

There are 3 primary sources of variability in the spiking model networks, not present in the mean field model that produce fluctuations resulting in units behavior departing from the single pattern type of the mean trajectory: 1) heterogeneity in the connections between units, 2) heterogeneity in the maximum facilitation, and 3) the noise present in all the units activity. Increasing population size reduces the source of noise resulting from heterogeneous connections, and thus reduces the overall amplitude of fluctuations. Figure 2.13 shows the reliability with which units in the 2000 unit 2- and 4-population spiking models exhibit their dominant patterns. It can be seen that neurons exhibit their dominant pattern more reliably than in the 200 unit network. However, increasing population size cannot eliminate type variability across trials particularly when the system is near the separatrix. It can be seen from Figure 2.13 that the average reliability of units expressing a single pattern type across the overall network is greater for the 2000 unit networks than the 200 unit networks (75% vs. 47% respectively in the 4 population model). However, individual units still exhibit high variability in the pattern type exhibited from trial to trial ranging between exhibiting the dominant pattern type 100% of the time to approximately as low as 45% of trials. Thus while reliability in firing can be achieved by increasing population size and averaging across units of a population, this does not eliminate transitions by units from the mean pattern type or from their dominant pattern type from trial to trial.

A reduction in the source of variability due to noise present in all neurons during simulation i.e. the Wiener noise can be achieved by averaging across trials. Figure 2.12A and 2.12B (bottom) shows the distribution of the average histograms obtained for the units in the 200 and 2000 unit population models across 20 trials. It can be seen that this averaging produces a distributions which primarily consist of patterns corresponding to the canonical bistable persistent activity (activation and inhibition). While this type of averaging is not physiologically relevant in the sense that populations carry out working memory each

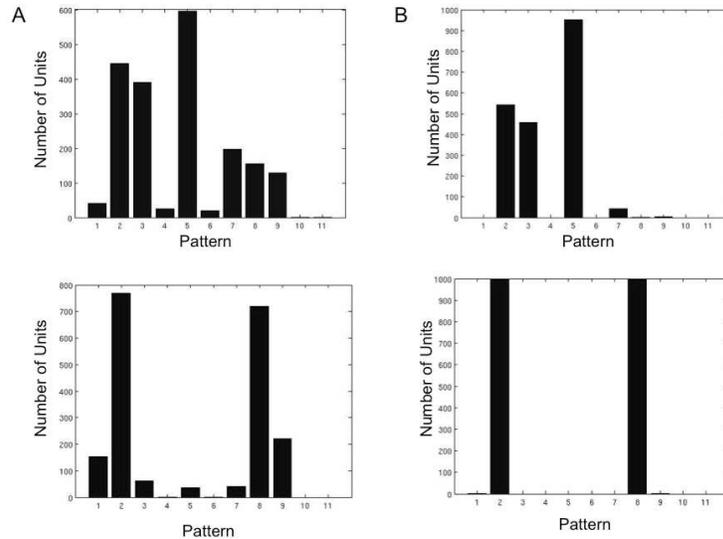


Figure 2.12: **Distributions of patterns exhibited across 20 trials of the simulated task.** A) top: Distribution of pattern occurrence for the 2- population, 2000-unit spiking network. The distribution is similar to that observed for the 200 unit networks and the parietal and prefrontal neuronal populations. Bottom: Distribution of patterns exhibited in the average (across 20 trials) PSTH histograms for the 2-population, 2000-unit spiking network. Note that the distribution of average PSTH exhibits the bistable excitatory and inhibitory patterns more prominently than actually exhibited from trial to trial by the units. B) top: Distribution of pattern occurrence for the 4-population, 2000-unit spiking network. The distribution shows similar relative occurrences of the persistent excitatory and inhibitory patterns (stable, ramping, decaying) as all previous networks and real neuronal data. Relative changes in specific pattern occurrences (i.e. decrease in the prevalence of the nonresponsive pattern) result from each populations activity approaching that of its corresponding mean firing rate model single pattern with increasing size. Thus a larger number of populations would be needed (i.e. greater than 4) to maintain the relative occurrence of all patterns and thus to maintain an invariant distribution. Bottom: Distribution of patterns exhibited in the average (across 20 trials) PSTH histograms for the 4-population, 2000-unit spiking network. Note that the average PSTHs exhibited are essentially all the bistable excitatory and inhibitory patterns. Thus analysis of average patterns across trials might falsely indicate that these are the only patterns of relevance occurring.

trial, and not as an average across trials, it does represent the typical averaging carried out to characterize cell behavior in studies of working memory (i.e. average PSTH histograms across trials determine cell pattern type).

An analysis of the intra-trial variance of firing rate in the model units revealed high variability in the distribution of ISIs during both baseline and delay periods of the model (Figure 2.14A). The CV of ISIs in the majority of units ranged in the baseline across all pattern types between 0.4 and 1 with a mode of approximately 0.6. During the delay period the distribution of intra-trial ISI CVs was bimodal with peaks at approximately 0.45 and 0.75 and most units falling within the range of 0.4 and 1 as in the baseline. These ranges of the CV overlap significantly with that observed in the real prefrontal and parietal cell populations, although their overall means are lower. Focusing on the stable excitation and inhibitory patterned activity, units exhibited decreasing average CV from baseline to the delay period in stable excitatory pattern units, and increasing average CV from baseline to the delay period in stable inhibitory pattern units (Figure 2.14B). In the real parietal and prefrontal cell populations, stable excitatory and inhibitory cells exhibit high CV in their ISIs during both baseline and delay, with the CV decreasing from baseline to delay in stable excitation cells and increasing in stable inhibitory cells. In parietal cortex, the CV of the ISIs in cells exhibiting stable persistent excitation significantly decreased ($p < 0.001$ paired t-test) from an average of 1.17 during the baseline to 1.02 in the delay. In prefrontal cortex the CV in those cells decreased insignificantly from an average of 1.03 to 1.0. In parietal cells exhibiting stable persistent inhibition in parietal cortex, the CV of the ISIs increased insignificantly from 1.19 to 1.2, while in prefrontal cortex the CV increased insignificantly from 1.02 to 1.03 on average.

2.4 DISCUSSION

The results of this study demonstrated that recurrent networks with dynamic synapses inherently produce the different persistent firing rate patterns observed in real cortical neurons during working memory. The persistent patterns produced are robust with respect to vari-

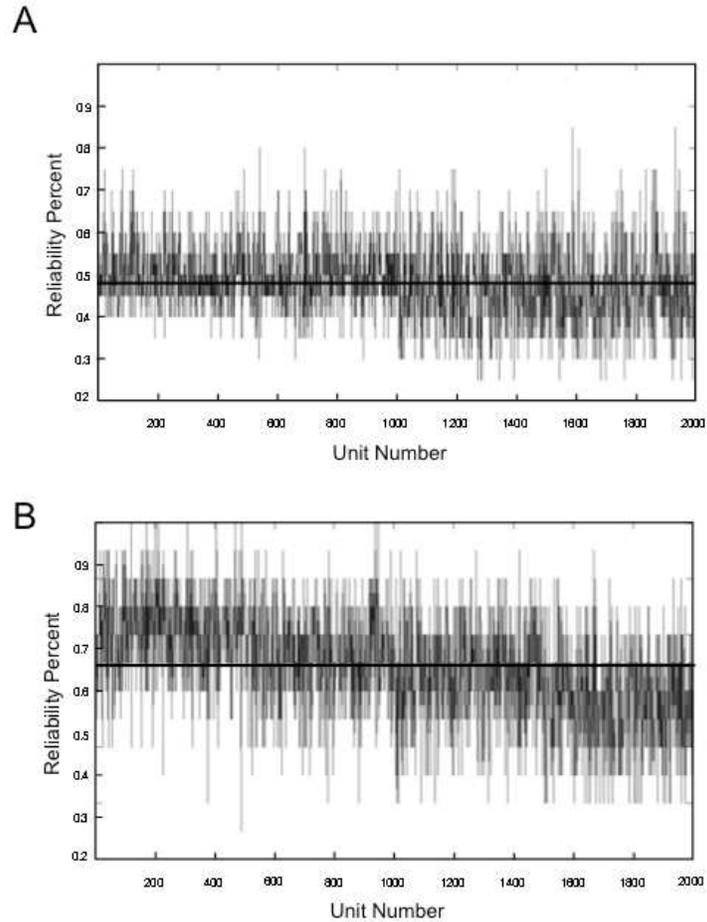


Figure 2.13: **Reliability of units in the 2000 unit networks for exhibiting from trial to trial the pattern indicated by each units average PSTH histogram.** A) Reliability of each unit in the 2-population, 2000 unit network. The average reliability (percentage of trials exhibiting the pattern observed in the average PSTH histogram) ranges from approximately 40-100% of trials, with an average of approximately 47%. This is similar to that observed for the 200 unit 2- and 4- population spiking models. B) Reliability of each unit in the 4-population, 2000 unit network. The average reliability is significantly higher across all 2000 units (approximately 75%) although the range of variability of individual units is similar to that observed in each of the previous networks. These reliability values are similar to that observed in the real cortical data of parietal and prefrontal cells.

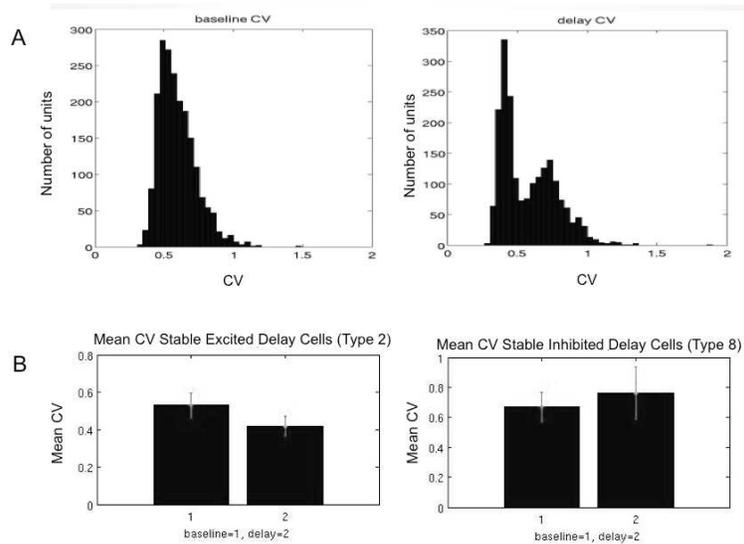


Figure 2.14: **Analysis of spiking variability.** A) Distributions of the coefficient of variation of the interspike intervals for units of the 2000 unit spiking model during the baseline and delay. Distributions correspond to cells exhibiting all patterned delay behaviors. High variability in ISI firing occurs during baseline and delay periods. Delay CVs show a bimodal distribution across all pattern types, as some modify their behavior after the stimulus, and some remain unresponsive. B) Mean baseline and delay ISI CV for stable persistent excitation delay cells (left) and stable persistent inhibition delay cells (right). The coefficient of variation decreases from baseline to delay in excited delay cells and increases in inhibited delay cells in agreement with the behavior observed in the database of prefrontal and parietal cells.

ations of the parameters in the network. That is, the different patterns occur over a wide range of values of the parameter space, and given patterns do not occur only for a very narrow set of parameter values. Further, the statistics of those patterns fall within the ranges of variation observed in firing rate pattern behavior of real cortical neurons. For example the changes in firing rate from baseline to the delay period can take values along an apparent continuum with absolute changes in firing rate of less than 100% of the baseline rate. For the majority of persistently activated cells recorded from parietal and prefrontal cortex of primates during working memory this corresponds to changes in firing rate of less than 10 Hz. The present network demonstrates a mechanism beyond previous solutions for achieving these realistic low delay firing rates [1, 13, 80, 122, 142, 165, 166]. While the expression of any particular delay frequency or rate of ramping or decay of firing rate of the units can be dependent on the particular parameters, the occurrence of any of the working memory patterns takes place across wide continuous ranges of network parameters and inputs, and thus do not involve fine tuning and are stable with respect to noise in the input.

Bistable firing rates are one of the possible activities of the model. However, the present work has focused on the range of working memory-correlated patterns of firing rate and their simultaneous, complementary occurrences in the working memory network as opposed to only fixed states that the networks or their neuronal constituents may adopt. The spiking networks exhibited all of the general patterns correlated with working memory that are observed in the database of microelectrode recordings of parietal and prefrontal cortical neurons. In addition, the statistics and firing rates of the units fall within the ranges observed in real cells, with the occurrence of the different pattern types similar in proportion to that observed in the cortical populations. In terms of the behavior of individual neurons, bistable activity is typically only observed as an average over many trials of a working memory task. Across trials, cells exhibit different average frequencies, and even within individual trials, cells exhibit significant variability in firing rather than a single stable rate [14, 161, 165, 171]. This is indicated from a high coefficient of variation in both baseline and delay periods in units exhibiting stable delay excitation and inhibition patterns. This is in agreement with database of real parietal and prefrontal stable delay units as well as previous neurophysiological studies in which CVs of within trial ISIs were around 1.0. Changes in CV

from baseline to delay period for the model units further agreed with that observed in the parietal and prefrontal database with the CV decreasing for stable excitatory pattern cells, and increasing for stable inhibitory pattern cells. High variability in ISIs has been observed in previous neurophysiological studies [14,161,165,171], although in some cases the change from baseline to delay observed has been different than that of the present cell populations. This may result from the frequencies or types of persistent patterned activity observed in those studies during the delays (e.g. bursting behavior). In addition to variability within trials of stable persistent activity cells, from trial to trial, neuron activity may adopt specific memory correlated patterns different from the most prominent one that emerges in the average across many trials. This not only includes changing between the different persistent excitatory patterns from trial to trial, but even changing between persistent activation and inhibited patterns. Thus while a population of cells may exhibit a particular pattern with consistency, individual cells of that population do not. In the 2- population spiking model with 200 units variability in firing pattern across trials was the same for both populations, with the majority of units exhibiting changes from their most prominent pattern type (including changing between persistent excitation and inhibited patterns) in 40% to 60% of the trials. In the 4-population spiking model, while the overall variability was essentially the same as in the 2-population model, the variability in pattern across trials depended on the types of patterns prominently exhibited by the particular populations. In populations exhibiting excitatory patterns the majority of units displayed a different pattern on 35% to 45% of the trials, and in populations exhibiting inhibited patterns, the majority of units displayed different patterns on 40% to 60% of the trials. Thus a more distributed architecture resulted in a more reliable occurrence of excitatory memory pattern types within units, in addition to a more stable concomitant occurrence of complimentary pattern types. Working memory therefore appeared to be more stable in the more widely distributed network. The reliability of exhibiting a given pattern across trials increases only partially with the size of the network. Looking at the 4-population network with an order of magnitude more unit results in a network that still exhibits the range of pattern types with relative proportions similar to that seen in the smaller network and in real cortical neurons. Although overall, the average percent of trials that units exhibit their most dominant pattern type increases (i.e. 75%

compared to approximately 47%), many units continue to exhibit their dominant pattern on less than a majority of the trials. The reason for this is that increasing population size decreases fluctuations due primarily to stochasticity in the connections between units, which of course does not drop to zero. In addition other sources of stochasticity remain such as noise in unit activity, and stochasticity in the facilitation. Thus given a particular stimulus, units trajectories in the phase plane will pass with some proximity to the boundary (separatrix) between the fixed bistable states of the system, and given the closeness to the separatrix and the amount of stochasticity, will have a significant chance on any given trial of crossing over to a trajectory corresponding to a different pattern than the mean trajectory of the population, or that which occurs most often for a particular unit. As a result, there is a relatively invariant distribution of pattern occurrence that changes modestly with increasing population size. It is interesting to note that artificially reducing the other sources of stochasticity by for example averaging across trials, that one produces pattern distributions exhibiting essentially only bistable patterns. That is if we look at the average PSTH activity of the networks across trials they tend to be either stable activation cells, or stable inhibition cells. While this type of reduction of stochasticity is not physiologically meaningful since working memory takes place trial to trial and cannot require averaging over many trials, data from unit recording experiments typically report unit activity as average (across trials) PSTH histograms. Thus the prevalence of bistability may be overestimated. Rather in real cortical data as well as in the model we see the variability as in the model.

While the firing pattern varies from trial to trial in cells, there are also significant variations from trial to trial in the delay frequencies for any particular patterns exhibited. The concept of a network with fixed connectivity and bistability between units is not indicated by such activity, and thus a dynamic connectivity is reasonable to consider. The idea of bistable activity corresponding to fixed attractor states may apply at the level of a population of neurons, and could be the essential neuronal correlate of working memory. However, the majority of persistent activity patterns observed consists of cells whose firing rates decay or accelerate during working memory (cue-coupled or response-coupled), and these populations should be taken into account in addition to bistability. In the present model, as stated above, bistable attractor states are present and could correspond to working memory.

Here however, a second additional role is suggested for these states in terms of modulating the firing rate activity, resulting in decaying and ramping firing patterns. That is, without necessarily representing memory states in and of themselves, the attractor states allow the network, which represents working memory and its complementary functions, to become active and behave with the necessary dynamics to mediate cross-temporal contingencies. The key component, resulting here from facilitation and observable in the phase plane, is the presence of the bottleneck through which the trajectories of the system pass. The bottleneck modulates the rate at which the trajectories approach the stable attractors, thus creating the patterned activities within the actual range of frequencies observed in the cortex. This mechanism might be present and modulate activity through other components of the network in addition to facilitation. For example Durstewitz [49] has previously demonstrated that a bottleneck or “ghost” of a line attractor could be achieved at a neuronal level through interactions between firing rate output and calcium gated ion channels, generating the climbing firing rate activity of set cells. Therefore, while the bistability (or multiple attractor states) could represent working memory, it plays the additional role of influencing the dynamics of the system such that the resulting trajectories correspond to different classes of working memory behavior (i.e. cue-coupled or response-coupled delay period patterns). The lower attractor state of the system, as is usual, is identified with baseline firing rates. Trajectories of the system not crossing the saddle separatrix remain in the basin of attraction of this attractor and ultimately return to baseline firing rates, adopting one of the classes of persistent activation associated with memory storage. In contrast, trajectories crossing the saddle separatrix approach the stable state corresponding to a higher firing frequency. Depending on the rate at which the system approaches the higher state, the firing pattern adopted is either bistable memory behavior (rapid or very slow approach), or ramping cell behavior associated with preparation for a behavioral or motor response. From these considerations it might be predicted that ramping cells and fixed rate memory cells in working memory should exhibit higher average firing rate changes from baseline than decaying-rate memory cells. An examination of the firing rate changes of prefrontal and parietal cells indicates that this is indeed the case. In parietal cortex, fixed rate memory cells exhibit an average difference between baseline and delay periods of approximately 9 Hz, set cells 7 Hz, and

decaying rate memory cells 6 Hz. In prefrontal cortex the same trend is observed with stable rate memory cells exhibiting the greatest mean change in firing rate from baseline to delay period (approximately 4.3 Hz), followed by set cells (3 Hz) and decaying rate memory cells (2 Hz). The fact that the fixed rate memory cells exhibit the largest average change in firing rate from baseline to delay is consistent with some percentage of those cells activity corresponding to the high firing rate bistable state in addition to those exhibiting only apparent bistability.

The dynamic synaptic facilitation is the component of this model which creates the bottleneck in the phase plane, and gives it its unique characteristics. Specifically it is facilitation which determines the amount of persistent activation, which, since it can adopt a continuous range of values, enables the change in firing rate from baseline to memory period to fall along a continuum. The bottleneck determines the rate at which the firing rates decay towards the baseline attractor (or increases towards the higher firing rate attractor) to adopt the continuum of firing rate values. The decay rate can be sufficiently slow such that no decay or acceleration of firing is observed for the duration of a memory period. Thus the result is an apparent or virtual bistability, which for all intents and purposes can be extended for as long as working memory is defined by the parameters of a working memory task. The fact that the rate at which persistent activation waxes or wanes is highly adjustable is consistent with the behavior of cells in the cortex during working memory. It has been observed in working memory experiments [67, 69, 110] that the rate of decay and/or the rate of acceleration of persistent activation adjust to the duration of the memory period. The dynamic synapses make this phenomenon easy to incorporate. Adjusting the maximum of facilitation or other parameters, changes the bottleneck so that the rate of decay (or ramping) can become longer or shorter along a continuum.

Another prediction from the dynamics of the model is that the rate of persistent activation correlates with baseline rate. In the majority of delay activated cells, the magnitude of firing rate increases are less than 100% of baseline, with the magnitude of the delay period firing rate change increasing nonmonotonically with baseline rate increases. The largest magnitude increases in delay period frequency are in those cells with the largest baseline firing rates, while the largest percentage changes are those with low baseline rates. This is natu-

rally incorporated in the present model. The range of rates over which the population can exhibit memory cell behavior is bounded by the saddle separatrix. Once facilitation pushes the systems trajectory beyond the separatrix, further increasing facilitation (or judiciously adjusting other parameters) does not result in further continuous increases in persistent activation delay rates, but rather a change in the activation pattern itself. The parameters of the model can be adjusted however, raising the frequency of the baseline state and incrementing the entire range of frequencies within its basin of attraction. Thus both baseline and delay rates increase in a correlated fashion, and due to the nonlinearity of the nullcline of the synaptic activity, the proportional increase in frequency is nonmonotonic.

The specific simultaneous patterns which may be exhibited in the populations are dependent on the relative strength of the interpopulation connection strength, the intrapopulation connection strength, and whether the inter-population connectivities are mutually net inhibitory, or a combination of excitatory and inhibitory. The phase diagram of the 2-population firing rate model reveals a number of behavioral trends. For an excitatory/inhibitory connectivity between populations, the networks can exhibit a range of concomitant activities which includes memory cell activity in both populations, and simultaneous cue-coupled/ response-coupled behavior. In contrast, with a mutually inhibitory connectivity between populations these particular behaviors are absent, and simultaneously occurring fixed-rate-memory/response-coupled behavior is present over only an extremely narrow range of the parameters. Thus memory being maintained simultaneously in both cortical areas occurs in the 2-population model only within the E-I connectivity scheme. During working memory, the simultaneous presence of memory cells in prefrontal cortex and another cortical area important to the sensory modality of the memorandum has been indicated by numerous studies. In addition to prefrontal cortex, memory cells have been observed for example in posterior association cortex including inferotemporal cortex [74, 140], and posterior parietal cortex [79, 114], and their simultaneous presence in multiple cortical areas have been indicated in imaging studies [35, 116, 140, 180]. The overlapping presence of cue-coupled and response-coupled cells have also been confirmed and have been implicated in working memory networks [116]. It is suggested that these populations would cooperate and be engaged in the transfer of information from a perceptual network to a motor network.

The two populations would cooperate to enable the processing of information from one network to the other with translation from perception into action. As the network becomes more distributed, increasing to four populations, simultaneous memory cell behavior and cue-coupled/ response-coupled behaviors becomes more robust with these concomitant behaviors occurring over a wide continuous range of the parameters as can be observed by the increased areas of those respective behaviors over larger continuous ranges of the parameters in the phase diagrams (Figure 2.6). While the specific connectivity between populations in different cortical areas in working memory networks is unknown, it is suggestive to consider the possible effects on each populations activity in the model when another is shut down as in reversible lesion studies. For example in the 4-population model, termination of the activity two populations changes the phase diagrams to that of the 2-population model. Depending on the specific local parameters (i.e. the connectivity between populations), the effect can be a net increase in persistent activity, a decrease, or elimination of such activity to a non-responsive pattern. Studies of reversible lesions in which one cortical area is cooled while recording cell activity in another has shown that some cells increase their firing rate during the delay, while other will show decreases or become non-responsive [35, 71, 72]. A question is whether the net effect of one cortical area on the other is excitatory or inhibitory as might as determined by more cells increasing persistent activation or decreasing it. While no definitive results exist, from these studies the data indicate that more neurons increase their activity in prefrontal cortex as a result of cooling posterior association cortices, while more neurons decrease their activity in posterior association cortex as a result of cooling prefrontal cortex. This could be indicative of an effective E-I coupling between cortical areas. Looking at the potential changes in the phase diagrams of the models going from the 4- population model (with connectivity of both the E-I and I-I type) to 2 populations, such behavior is the general trend over the majority of the parameter space. Further lesion studies or microstimulation studies may elucidate the functional connectivity of global network in light of the model.

In addition to a distributed architecture affecting the stability of memory pattern behavior and modulating activity enabling the occurrence of complimentary patterned behaviors, certain working memory pattern behaviors apparently are exclusively a function of a dis-

tributed architecture rather than the facilitation mechanism alone. Particularly the ramping delay inhibition pattern which is observed in the cortical data was present only in the distributed versions of the model. Another phenomenon is the existence of large regions of the parameter space in which one population exhibits the non-responsive pattern, while the other population exhibits memory cell behavior (fixed-rate, decaying, or ramping). In the database of real cells, the majority of neurons from parietal and prefrontal cortex exhibit the non-responsive pattern of behavior. Interspersed within these populations of nonresponsive cells are neurons that exhibit the other patterns. From the models we see that the non-responsive pattern is a common part of a working memory network coexisting with the other patterned behaviors. Studies of patterns in spike sequence of such cells [21–23] have indicated that while not exhibiting significant differences between baseline and delay firing rates, such cells can exhibit differences in the patterning of the spike sequence in these periods; indicating participation in the working memory network. The fact that the non-responsive pattern arises as a prominent one in the models, overlapping with the range of memory pattern behaviors suggests that these populations may play a role in the dynamics of working memory networks.

It should be noted that the present analysis supplements the general attractor picture rather than replacing, or invalidating it. Cells with apparent bistable activity with high firing rates above baseline, while apparently rare in the cortex [171], may still be a fundamental neuronal substrate of working memory. In the present model not only is this activity present, but also the myriad other patterns with firing statistics and variability similar to those which constitute much of the activity correlated with working memory are accounted for.

3.0 FROM WORKING MEMORY TO EPILEPSY: DYNAMICS OF FACILITATION AND INHIBITION IN A CORTICAL NETWORK

3.1 INTRODUCTION

While the mechanisms of working memory are unknown, typically in most models these states correspond to stable attractors and dynamics arising in those recurrent networks [6–11, 13, 26, 38, 50, 203]. It is assumed that derangement of these ubiquitous recurrent cortical networks play a fundamental role in various neuropathologies. Particularly, it has long been recognized that recurrent impulses are a critical factor in generating hyperexcitability and recruitment which are the essential features characterizing seizures and epilepsies [107, 123, 190–192]. Epileptic seizures represent temporary episodic periods of increased network excitation with variable propagation. It is suggestive to assume that the type of pathological activity observed in seizures/epilepsies is a function of inherent dynamics of recurrent working memory networks. Working memory has provided the archetype of persistently active states. Neuronal working memory networks remain active after the presentation of a cue (memorandum) during a delay period [64, 70]. These persistent states may be maintained through a relative balance of excitation and inhibition [87, 175], or through asynchrony, and terminated through synchronization [85]. Numerous studies have demonstrated deficits in working memory function in epileptic subjects [4, 42, 83, 115, 193].

With the exception of neurological disorders due to malnutrition, epilepsy is the most prominent disorder in the world effecting approximately 1% of the population. It is estimated that there is a 10% lifetime risk of exhibiting a single seizure, approximately one third of which will develop epilepsy. Epilepsy/seizures can arise in a number of varied forms, with potentially similar or varied underlying mechanisms. While epileptic seizures involves

paroxysmal bursting of neurons in a local circuit, the clinical manifestations of seizures result mostly from spread of activity from local circuits to involve adjacent and remote brain regions. While in working memory, widespread populations are activated in normal cognitive function, and perhaps are related to binding, in seizure activity the recruitment of cortical networks and populations occur in a nondiscriminant pathological fashion. How different brain regions or populations are recruited is not well understood and it is not known how to stop ongoing seizure propagation or prevent seizure activity. Further, little is known as to how seizures either begin or cease [188]. While it has been a long-standing belief that a connection between hyperactivity and hypersynchrony is fundamental in seizures, it has recently been shown that hypersynchrony is unnecessary to produce seizure-like bursting [151, 195]. There is much evidence suggesting that seizures described as a straightforward increase in synchronization between neurons may be too simplistic. Computational and experimental models have shown however that low levels of excitatory coupling may be a prerequisite for some types of seizure onset [56, 159, 196]. The coexistence of models involving both increases and decreases in excitation points not to a disagreement, but to the complexity of seizure causality.

Synapses between neurons are known to undergo changes in their strength and dynamics. In working memory function, dynamic synapses (i.e. through synaptic facilitation) in recurrent networks can result in the normally observed persistent activation [13, 200](Chapter 2). A reasonable postulate widely proposed is that the development of epilepsy and seizures results from a shift in the balance between excitation and inhibition towards excitation [46, 75, 150]. Further a critical role is believed to be played by recurrent synaptic excitation in epileptogenesis [107, 123, 190–192]. For example synchronized bursting is favored by strong recurrent excitation between principal neurons and by disinhibition (review in [190]). It has been recognized increasingly that epileptic seizures are a dynamic disease caused by a change in the state of the brain dynamical system [169]. Different types of seizures have been viewed as bifurcations between distinct types of non-linear dynamics [207]. Nonlinear dynamics can advance our understanding [121, 162] of the spatial and temporal behavior of seizures. Seizures may be triggered by some change in network parameters and/or inputs not evident to an observer [129]. In so-called reflex epilepsies, seizures are precipitated by some

particular influx of afferent impulses, and may be induced by a wide range of external stimuli of different modalities such as photic stimulation, geometric patterns, music or computer video games [92,189], or internal cognitive processes, such as mathematical calculation. In a normal cortex, such external or internal stimuli might cause a transient, harmless modification of cortical activity, while in a predisposed brain they can induce massive synchronous discharges leading eventually to a seizure. It has been assumed that the stimulus leads to a dynamical change of the underlying attractor that facilitates the transition to the ictal phase (bifurcation). It has been proposed for example that in neuronal networks in the brain [162] the onset of seizures occur via a transition from stable linear dynamics via linear instability to non-linear behavior.

In humans, short electrical stimulation applied during cortical mappings is able to produce repetitive or periodic excitatory discharges in the cortex. In patients with epilepsy those discharges can progress to produce clinical seizures. However, in some cases a second electrical stimulation may stop those discharges. The fact that external electrical stimulation may terminate that activity in some cases raises the possibility of a method for seizure control. Uncertainties and variability in the ability of electrical stimulation to terminate the pathological discharge activity though, implies that theoretical and model systems might be useful to understand the mechanism of action of these techniques. In contrast to the generation and termination of seizures via invasive electrical stimulation, seizures may be generated and prevented or terminated through external stimulation. In particular, while stimulation with particular music is known to induce seizures in predisposed individuals, other music has been reported to prevent or terminate epileptiform activity [100,174,194].

The potential modulation of termination seizure activity by brain stimulation is attracting considerable attention. Recently there has been growing interest in neural stimulation to reduce seizure frequency. Approaches include for example vagal nerve and thalamic stimulation, and event-driven stimulation to terminate repetitive bursting. Modeling the effects of certain characteristics in the stimulation of working memory networks, such as specific spatiotemporal patterns, could yield efficient and minimally invasive approaches for treatment of epileptic patients. A related issue to initiation and/or termination of seizures is the mechanisms and dynamics by which a seizure recruits cortical areas and spreads within

the cortex. Nonlinear dynamics can advance our understanding of the spatial and temporal behavior of seizures [121, 162]. Combining the concepts of neurophysiology of neural networks with the mathematics of nonlinear systems can help lead to an understanding of these mechanisms since neural networks are nonlinear systems with complex dynamics. This essential aspect must be accounted for in order to understand how neural network can have bistable memory states (or multi-stable states) and exhibit bifurcation between those states.

In this work we present a model of a working memory network and explore its nonlinear dynamic behavior in normal and seizure/epileptic states. Particularly we examine how the network can transition from normal working memory behavior dynamics to those characteristic of seizure activities particularly widespread recruitment of populations with varying degrees of synchronous oscillatory behavior. We propose that facilitation and inherent network parameters can bias neuronal activity to that of recruitment and seizure. We demonstrate that seizure behavior can be elicited in the model through input with specific temporal and/or spatial characteristics, simulating reflex epilepsies. Finally, we show that seizure activity may be also be terminated by input to the network with specific temporal characteristics. We start the paper by introducing a network of $N = 20$ populations, each of which consists of three variables representing the activity of the excitatory and inhibitory neurons and the degree of facilitation of the excitatory synapses. We show the “normal” behavior for this network which consists of the selection and maintenance of a salient input. We then alter the strength of lateral excitatory to inhibitory connections to mimic pathology and find a variety of disruptive states. In order to better understand these, we study a two population model using bifurcation analysis. We find the attractors and then characterize the basins of attraction for each of the stable states by varying the frequency and strength of transient stimuli. We look at a single population and use the method of averaging to clarify why there are so many attractors. Finally, we explore possible mechanisms for the termination of seizures using the results from the previous sections.

3.2 METHODS

We build on a previously defined model for working memory in a neural network based on the interactions between inhibitory and excitatory neurons as well as synaptic facilitation. The network involves coupling several modules each of which is a three-dimensional system of the form:

$$\tau_u \frac{du}{dt} = -u + f(a_{ee}(1 + kw)u - a_{ie}v - \theta_e + c_e p(t)) \quad (3.1)$$

$$\tau_v \frac{dv}{dt} = -v + f(a_{ei}u - a_{ii}v - \theta_i + c_i p(t)) \quad (3.2)$$

$$\tau_w \frac{dw}{dt} = -w + f(\gamma(u - \theta_w))[w_{max} - w]. \quad (3.3)$$

In this model, u represents the firing rate of the population of excitatory neurons and v the firing rate of the population of inhibitory neurons. The main nonlinearity is $f(x) = 1/(1 + \exp(-x))$. Coupling parameters, a_{jk} are all non-negative. w represents a slow activity-dependent facilitation of the connection strength. That is, as u fires enough above the threshold, θ_w , then w slowly moves toward w_{max} . The parameter k characterizes the importance of the facilitation. The function $p(t)$ represents external input to the system.

The simplest network involves coupling a pair of these models. Coupling is allowed only through the excitatory cell. As this pair represents a local cortical network, coupling between networks is mediated primarily through long excitatory connections which can project to either excitatory or inhibitory neurons. Thus, the u equation has a term of the form $d_{ee}(1 + kw)\hat{u}$ added to the inside of the nonlinearity f , and the v equation, a term of the form $d_{ei}\hat{u}$. In absence of the facilitation, the (u, v) -system is the classic Wilson-Cowan equation. Facilitation allows there to be a number of different coexistent stable states.

In much of the paper, we study the larger network of $N = 20$ modules with coupling from the excitatory cells to other excitatory cells and to inhibitory cells. There is only facilitation of the excitatory-excitatory cells. For the one- and two population models, we chose the following parameters: $a_{1ee} = 12.3, a_{ie} = 10.1, a_{ei} = 11, a_{ii} = 7, d_{ee} = 0.7, d_{ei} = 3.5, k = 0.7, \theta_e = 2.4, \theta_i = 2.8, \theta_w = 0.5, \gamma = 5, w_{max} = 0.7$, and $\tau_u = 0.02, \tau_v = 0.04, \tau_w = 2$. The stimulus has the form:

$$p(t) = \exp[-20(1 - \cos(2\pi t/P_0))]$$

where P_0 is the period. (All time units are in seconds.) It is applied generally for 2-3 seconds and with varying strengths to the inhibitory and excitatory populations. For the networks consisting of 20 populations, all parameters are randomly varied between 2 and 5% around the above values.

3.3 RESULTS

The idea of working memory is that a network can maintain a local area of sustained neural activity after a stimulus is removed. Almost all models of this phenomena involve selective bistability between a quiet resting state and a state of sustained activity. Essentially, if a specific stimulus arises, then the population of neurons that best responds to that stimulus will turn on and suppress other populations of neurons. Recurrent excitatory connections (and in our model, the synaptic facilitation) enable the stimulated population to stay on after the stimulus is removed so that the network can “remember” which population was stimulated. This kind of memory is also called short-term memory. In this and the ensuing parts of the paper, we want to show that (i) there can be temporal sensitivities to these networks and (ii) explore how damage to the inhibition (specifically the excitatory to inhibitory connections, d_{ei}, a_{ei}) can destroy the working memory properties in dynamically interesting and possibly clinically relevant means. To do this, we start with a network model and show a number of complex features as the network is “damaged” through the weakening of the d_{ei}, a_{ei} connections. We then turn to a two population model and use numerical bifurcation theory to understand the various attractors. Then we use a single population and the method of averaging to explain some parts of the bifurcation diagram.

Our main assumption is that damage of the inhibition is responsible for the general phenomena of reflex epilepsy, and that in this low-inhibition regime periodic stimuli can cause the onset of seizures. Thus, in general, $p(t)$ will be a periodic stimulus with some narrow range of frequencies.

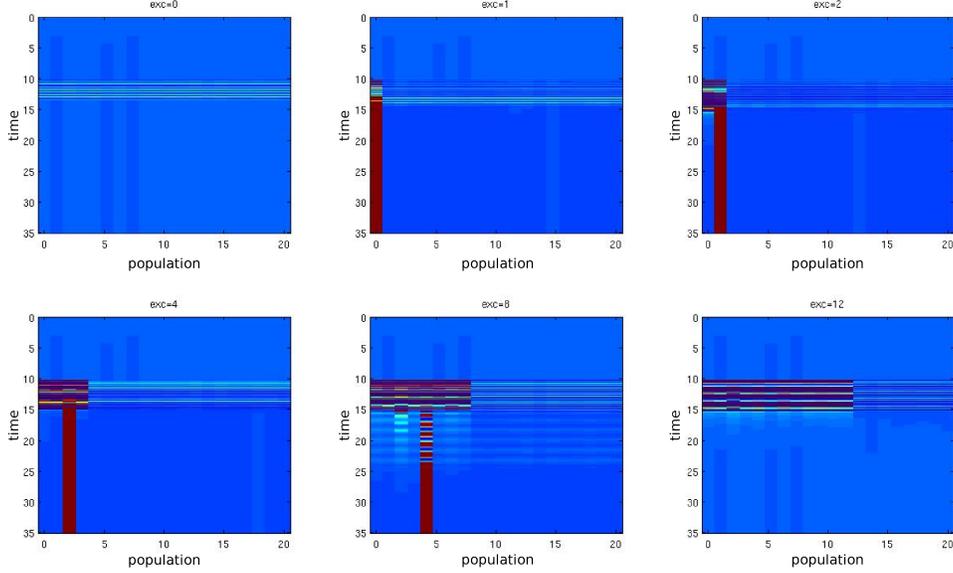


Figure 3.1: Behavior of the 20 population network with normal ($d_{ei} = 3.5$) inhibition. All cells are given a mild background stimulus and the first k cells are given a strong stimulus. Time (seconds) is indicated on the vertical axis and the population numbers are given on the horizontal axis. The excited populations are indicated above each graph. Blue indicates baseline firing rate, light blue indicates slightly higher than baseline firing rates, and red indicates activated (above baseline) firing.

3.3.1 Normal Parameters

We have tuned our model in such a way that a single population can stay excited after a stimulus. This property is implemented by assuming broad excitatory to inhibitory coupling (E-I) and strong local inhibition $I - E$ in the network. Intuitively, if the population of excitatory cells is turned on, then that will also excite the inhibitory cells of other populations which will keep these populations suppressed. Since there are also excitatory to excitatory (E-E) connections and strong local E-I connections, how is it possible to maintain activity in only one population? This is done through the slow facilitation. If a single population is strongly stimulated, then the facilitation for that population will build up and allow it to remain high once the stimulus is removed. The other populations which have not been directly stimulated, will become excited but not sufficiently to remain permanently on, especially in light of the strong inhibition.

In each of the following simulations, a periodic stimulus is given with a particular frequency to all the populations in the network and the first k populations in the network are given a larger version of the same stimulus. The stimulus lasts for 5 seconds and the

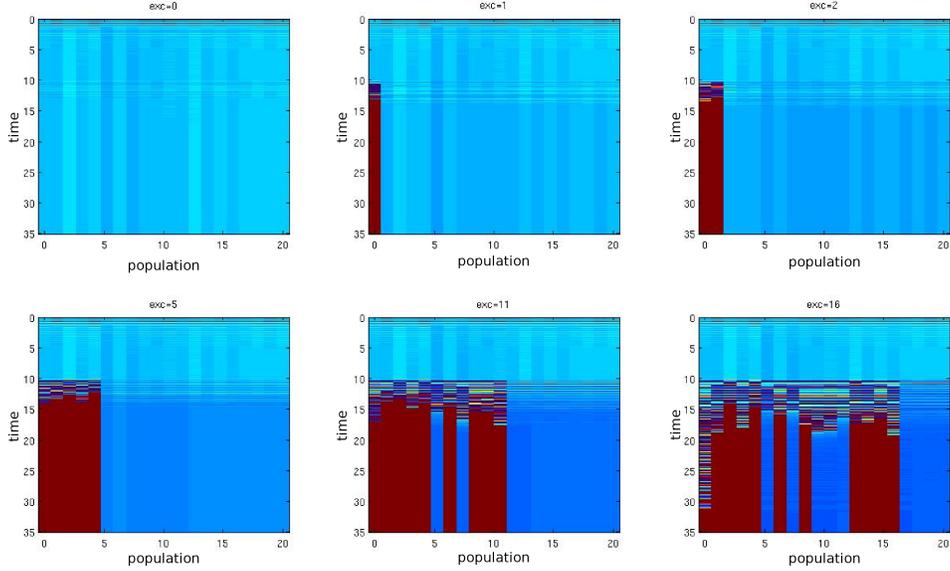


Figure 3.2: Behavior of the 20 population network with diminished ($d_{ei} = 1.59$) inhibition. All cells are given a mild background stimulus and the first k cells are given a strong stimulus.

simulation lasts for 35 seconds. Figure 3.1 shows the behavior of the normal network which undergoes winner-take-all working memory behavior as long as the number of stimulated populations is sufficiently small. The first five panels ($k = 0, 1, 2, 4, 8$) show that a single winner emerges when less than about half the network is stimulated. When only a single population is stimulated, that population will emerge as winner. However, when multiple populations are stimulated, the heterogeneity in the network (due to small random changes in the parameters) breaks the symmetry and a particular winner emerges (in this case population three). However, if more than about half the populations are stimulated (for example, 12/20), then the feedback inhibition is sufficient to prevent any population from emerging as the winner. Thus, the long-range recurrent inhibition acts to constrain the network in such a way as to prevent more than one stimulus to be selected.

3.3.2 Pathology

3.3.2.1 All E-I disrupted Our main hypothesis is that reflex epilepsy is a consequence of the breakdown of feedback inhibition. There are several ways to disrupt inhibition: change the thresholds (θ_i), the E-I connections (d_{ei}, a_{ei}) or the I-E connections (a_{ie}). In this paper, we alter E-I connections but manipulating I-E produces a similar effect (simulations not

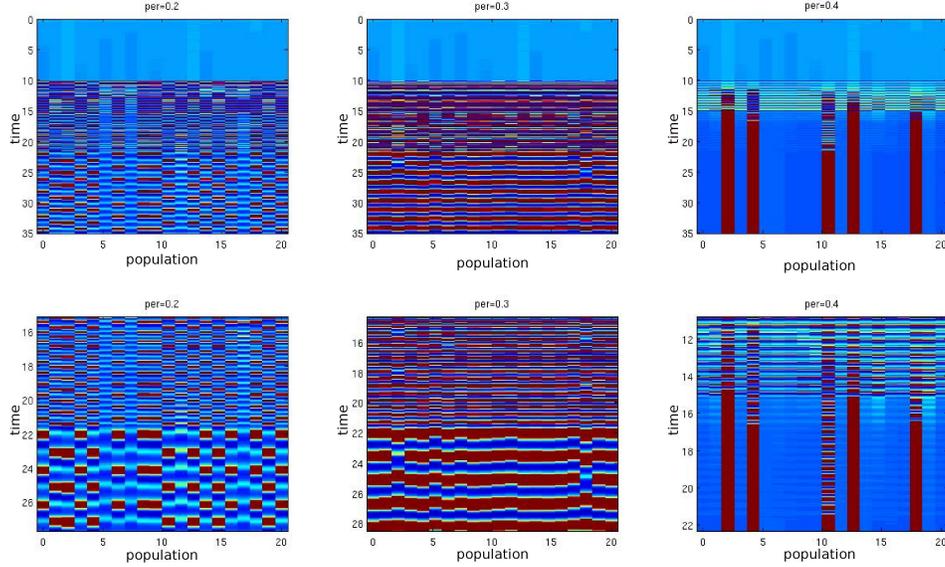


Figure 3.3: Behavior of the 20 population network with diminished ($d_{ei} = 2$) inhibition. Here, all cells are stimulated with different periodic stimuli. Lower panels are an expanded view of upper panels.

shown). Figure 3.2 shows that a strong reduction in the E-I connections (from 3.5 to 1.59) causes a loss of selectivity to the network. The resting background state remains stable to small enough perturbations and if a single population is activated ($k = 1$), then that memory can be maintained. However, if more than one population is excited, multiple populations maintain activity and selectivity is lost. An interesting transient in which populations begin to oscillate before settling to a steady state solution can be seen in several of the simulations.

Figure 3.3 shows what happens when the inhibition is not quite as reduced, (2/3.5 instead of 1.59/3.5). A distinctive frequency dependence on the stimulus emerges. Here all the populations are stimulated at three different frequencies leading to three different steady state behaviors. At 5 Hz (period is 0.2 seconds), the populations break into clusters which are separated by a half cycle. The number of clusters in each group can vary; an expanded view is shown in the lower panel. At a lower frequency of 3.3 Hz, synchronous oscillations emerge. Each population fires at the same frequency. Finally, at 2.5 Hz stimulus, there is again WTA behavior, however, the emerging patterns have nothing to do with the stimulus.

3.3.2.2 Partial Disinhibition A more biologically likely scenario would be that only some local areas are pathological. That is, the disinhibition is “broken” for a finite number

of populations. Figure 3.4 shows some simulations of the 20 population model when 1,2, or 5 populations have reduced E-I connections. A single damaged population (A) allows for selective memory and competition as long as the damaged population is not among those stimulated. If it is stimulated, then it always is selected (compare the A, 1,2 versus 4). Panel (B) shows that background stimulation is enough to keep the damaged population active at all times after the stimulus, and this is sufficient to suppress any selectivity of more strongly excited populations. Indeed the simulation shows that the constant activity of the damaged population prevents the selection of the first or second populations when they are strongly activated. Figure 3.4C shows a similar result for damage to populations 4 and 5. Interestingly in this and also to some extent in 3.4A, the selected population does not go to a fixed point but rather oscillates. Finally, with five damaged populations, 3.4D shows that it can be difficult to get an undamaged population to stay activated due to the strong surround inhibition which comes from the higher activity of the damaged population. Because the damaged population has less inhibition, it is more active and able to suppress the other populations, even though, it is not highly activated (light blue rather than red).

3.4 TWO-POPULATION MODEL

In order to gain some insight into how damage reduces selectivity and produces a variety of pathological responses, we turn to a two population model of identical groups and such the we reduce the cross inhibition. We first treat the E-I coupling as a parameter and find the attractors for the network using numerical bifurcation methods. Then we attempt to explain how stimuli effect a switch from rest to a specific attractor.

3.4.1 Attractors

Figure 3.5 shows a sketch of the bifurcation diagram for the two population model as a function of the parameter, d_{ei} which is the cross population E-I connection strength. No particular stimulus is used to produce this diagram, so we can consider that $p(t) = 0$. We

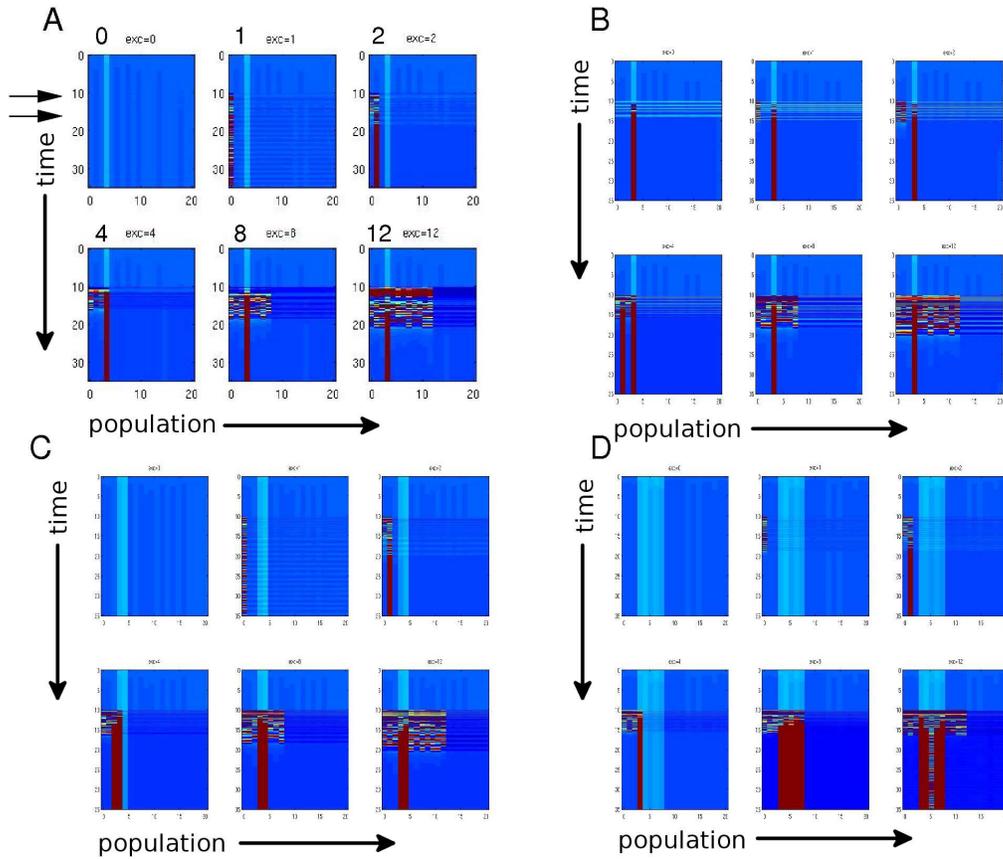


Figure 3.4: The effects of localized damage on the network. Each panel shows a simulation for thirty seconds after five seconds of stimulus (arrows) and the first k populations ($k = 0, 1, 2, 4, 8, 12$). In panel B, in addition to local stimulation, there is also background stimulation. (A) Population 4 is damaged and once stimulated will always “win”. (B) With background, the damaged population is always active even with weak (background) stimulation. (C) Similar to (A) but populations 4,5 are damaged and can both stay on when stimulated. (D) As in A,C, but 5 damaged populations.

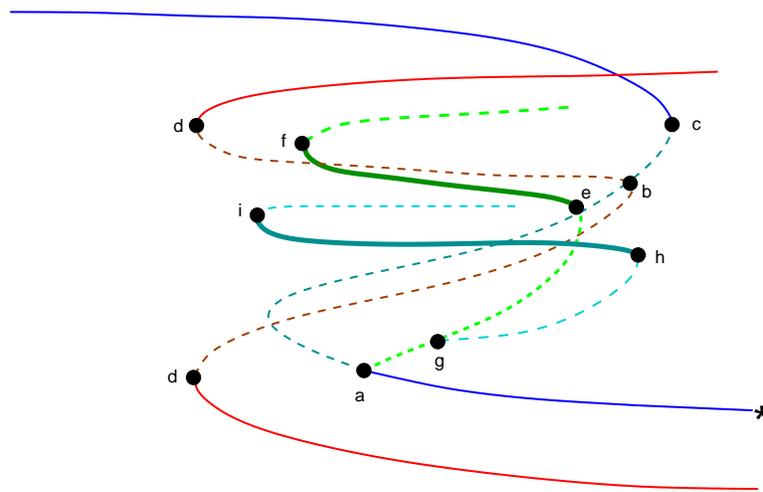


Figure 3.5: Schematic of the dynamics of a two-population network as the cross $E - I$ strength is reduced. Solid thin lines correspond to stable fixed points (red/blue) and solid thick lines correspond to stable periodic behavior (green/cyan). Black filled circles are important bifurcations. Details are in the text.

now describe the attractors and how they are connected and where they exist. We start on the dark blue curve at the asterisk and move to the left. This starting point represents a state in which both populations are at rest, in their low state. As the d_{ei} decreases, there is a Hopf bifurcation (**a**) and this state loses stability. Thus, the network cannot remain quiescent below $d_a \approx 1.713$. As we continue along this unstable symmetric branch, there is a pitchfork bifurcation (**b**) which spawns an asymmetric pair of unstable equilibria. Continuing along this branch, there is a fold bifurcation (**c**) which stabilizes the symmetric state in which both populations are turned on. This state persists for all $d_{ei} < d_c \approx 3.62$. The effect of reducing the competition is that we enable the state in which both equilibria are turned on to be stable. We now pick up the pitchfork bifurcation at **b**. This pair of unstable branches representing an asymmetric case in which one population is more active than the other, undergoes a fold bifurcation at **d** and gives rise to winner-take-all dynamics. For $d_{ei} > d_d \approx 1.162$ the network has a state in which one population is on and the other off. This corresponds to the normal state in which there is memory of the initial stimulus. A symmetric unstable branch of periodic orbits emerges from the Hopf bifurcation at **a** which undergoes a period-doubling bifurcation **g**. As we continue along the symmetric unstable branch of periodic orbits (green dashed) there is a fold of limit cycles (**e**) and the symmetric synchronous oscillation is stabilized (thick dark green curve). This solution remains until there is another fold (**f**) leading to an unstable synchronous solution (dashed green curve). Stable synchronous oscillations exist for a limited range $2.077 \approx d_f < d_{ei} < d_e \approx 2.889$. Turning our attention to the period doubling bifurcation (**g**) this branch of unstable asymmetric solutions (dashed, cyan) undergoes a fold bifurcation at **h** and another fold at **i** such that there are stable anti-phase (alternating) oscillations for $1.85 \approx d_i < d_{ei} < d_h \approx 3.17$. It is rather remarkable that all these branches are connected. In the “normal” network, we imagine that $d_{ei} > d_c$ so that the only attractors are the quiescent state and the two fixed points corresponding to WTA behavior. The most pathological state occurs for $d_{ei} < d_d$ where only the completely excited state exists. These patterns correspond to the patterns of activity seen in the 20 population model. Symmetric solutions correspond to homogenous behavior such as the synchronous oscillations seen in figure 3.3. The anti-phase oscillations are analogues of the clustered states see in figure 3.3 and the WTA behavior corresponds to the normal network

states such as seen in figure 3.1.

3.4.2 Basins of Attraction

In this section we apply a variety of stimuli to the network when d_{ei} is fixed at a value of 2.6 where all six attractors are stable. We will apply periodic stimuli at different frequencies with different amplitudes to see if it is possible to switch to the active states from the quiescent state. There are many possible stimulus parameters to vary, so we will start with the following. We stimulate for 2 or 3 seconds at a variety of different frequencies and with different amplitude ratios between the two populations. Specifically, we set the excitation to a value of 1 in population one (called the *preferred* stimulus or population) and vary the strength of the stimulus in population two between 0 and 1 (the *nonpreferred* case).

Figure 3.6 shows the phase-diagram of the steady state behavior as a function of the period (in seconds along the horizontal axis) and the magnitude of the non-preferred stimulus along the vertical axis. The behavior and transitions appear to be very complex. For example with a 2 second stimulus (A) at about 4 Hz, as the non-preferred stimulus increases, there is winner-take-all behavior, anti-phase oscillations, and synchronous oscillations. The three second stimulus (B) shows qualitatively similar behavior but there is a much larger set of initial data leading to synchronous oscillations. Whereas shorter stimuli require nearly identical preferred and non-preferred inputs, with a longer duration, the basin for synchrony is quite large. With the longer duration stimuli, synchrony takes over much of the territory of the anti-phase oscillations, while the anti-phase oscillations invade the rest state territory. Presumably, the latter effect is due to the longer stimulus allowing a greater buildup of the facilitation, w , thus making some active state more likely. In figure 3.7, we show an expanded parameter scan within the green rectangle in figure 3.6A. Based on this we suspect that the basins of these attractors are very complicated with riddled fractal structure. It seems that there is never a direct transition from synchrony to anti-phase. The WTA behavior always seems to separate these two attractors. The complex behavior shown here is a consequence of the pathology introduced in the network. In the “normal” network, we find (not shown) that for all two second stimuli (periods between 0.05 and 0.5 seconds), the network goes to

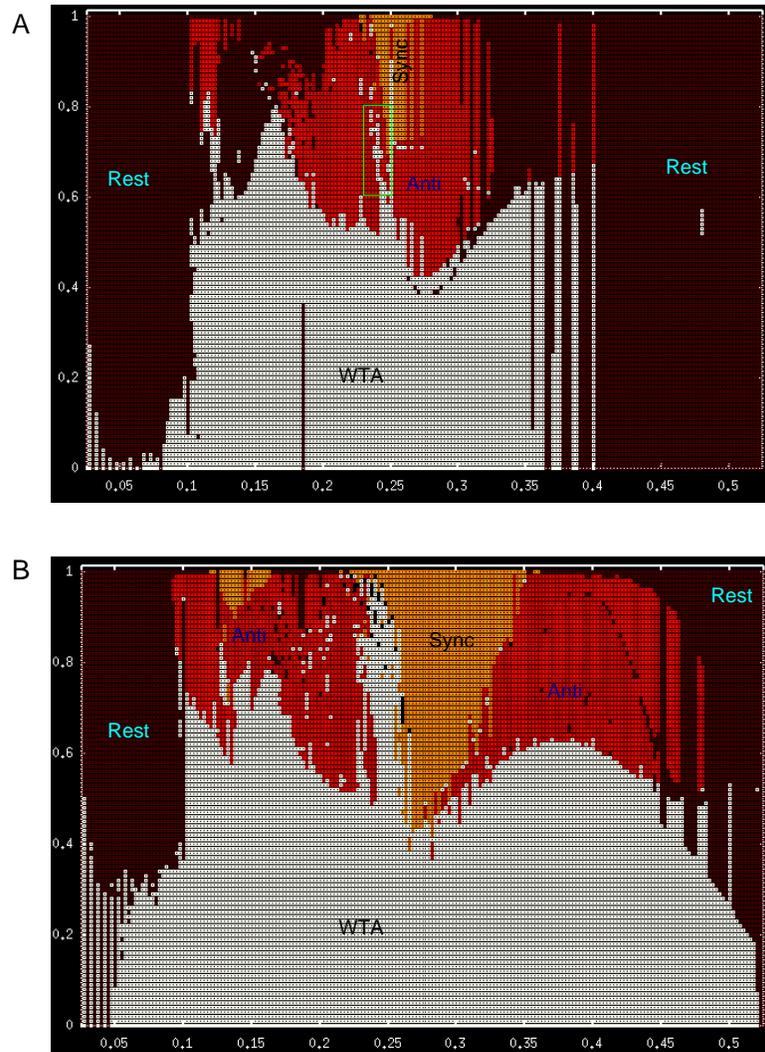


Figure 3.6: Steady state behavior of the pathological network ($E - I$ cross connections reduced to 2.6) as a function of the period of the stimulus (in seconds) and the strength of the non-preferred stimulus (preferred strength is 1). Four colors correspond to four different states: Return to rest (brown), winner-take-all (white), synchronous oscillations (orange), and anti-phase oscillations (red). The horizontal axis corresponds to the period of the stimulus, the vertical axis to the magnitude of the non-preferred stimulus. (A) two second stimulus; (B) three second stimulus.

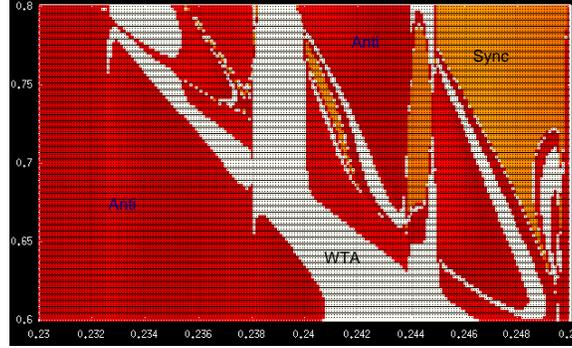


Figure 3.7: Expanded view of green rectangle in figure 3.6A. Riddled basin for input stimuli in a narrow range of periods and relative amplitudes. Parameters are the same as in figure 3.6A in the region shown with the green box.

the usual winner-take-all behavior with the preferred population always winning.

The steady state behavior is very difficult to predict by looking at the time series of the populations. Figure 3.8 shows the dynamics of the facilitation, w_1, w_2 and the excitatory activity, u_1, u_2 of the two populations. In figure 3.8, we show these variables for a short period of time centered around the end of the stimulus. In all cases, the non-preferred stimulus is 0.8 and the preferred is 1.0. We choose three different nearby periods for the stimulus such that there is either return to rest, anti-phase oscillations, or winner-take-all. Figure 3.8A shows the facilitation in the three cases. The red/orange curves correspond to a period of 0.1075 seconds and both populations return to rest. There simply is not enough buildup of $w_{1,2}$ to maintain them. The excitatory activity is shown in 3.8B. The green/olive curves correspond to a stimulus period (0.10875 sec) in which there is WTA behavior with the preferred (green) population winning. Note that because the non-preferred stimulus amplitude is smaller, the green curve is above the olive curve. The blue/cyan pair shows the preferred/nonpreferred when the system ends in the anti-phase state (period 0.11). Both curves are slightly higher than the green/olive combination due to the slightly higher frequency of stimulus. We first contrast the rest with the anti-phase case. The facilitation of the second nonpreferred population (cyan) is quite a bit less than that of the preferred population in the rest case (red) and yet the nonpreferred population stays active after the stimulus. The reason for this is that the preferred adaptation variable (blue) is sufficiently high to turn on and the recurrent E-E connections between the two populations give the second population a boost.

Indeed, just lowering this parameter (d_{ee}) from 0.7 to 0.6, pushes both populations to rest. The distinction between WTA and anti-phase behavior is far less clear. Figure 3.8C,D shows that the activities of the two populations (green is preferred, red is nonpreferred) are almost identical after the stimulus finishes and the first two cycles look like an anti-phase oscillation. The difference between the adaptation variables at the endpoint of the stimulus is essentially the same. Thus, there seems to be little *qualitatively* different between the transition to WTA and to anti-phase. This feature provides a potential explanation for the complex fractal nature of the basins of attraction shown in figures 3.6 and 3.7.

The two population model shows very sensitive dependence on perturbations even though each of the attracting states is very robust. It is only possible to reach the upper state in which both populations are firing at a steady state when the stimuli to both populations are very strong and nearly symmetric.

3.5 ONE POPULATION

Three of the behaviors described in the previous section can be understood by looking at the one-population model which is only a three-dimensional dynamical system. Furthermore, in fact, it is two fast variables (u, v) and one slow variable (w , the facilitation) so that we can apply standard fast-slow decomposition methods. In the two-populations (and in the N population model), one can consider the following three cases: all at rest, synchronous oscillations, and all turned on. In each of these three cases, all populations are identical, so we are left with a three-dimensional system:

$$\tau_u \frac{du}{dt} = -u + f((a_{ee} + d_{ee})(1 + kw)u - a_{ie}v - \theta_e) \quad (3.4)$$

$$\tau_v \frac{dv}{dt} = -v + f((a_{ei} + d_{ei})u - a_{ii}v - \theta_i) \quad (3.5)$$

$$\tau_w \frac{dw}{dt} = -w + f(\gamma(u - \theta_w))[w_{max} - w]. \quad (3.6)$$

Here we retain the coupling parameters d_{ee}, d_{ei} in order to emphasize that these equations represent the *symmetric solutions* of the coupled populations.

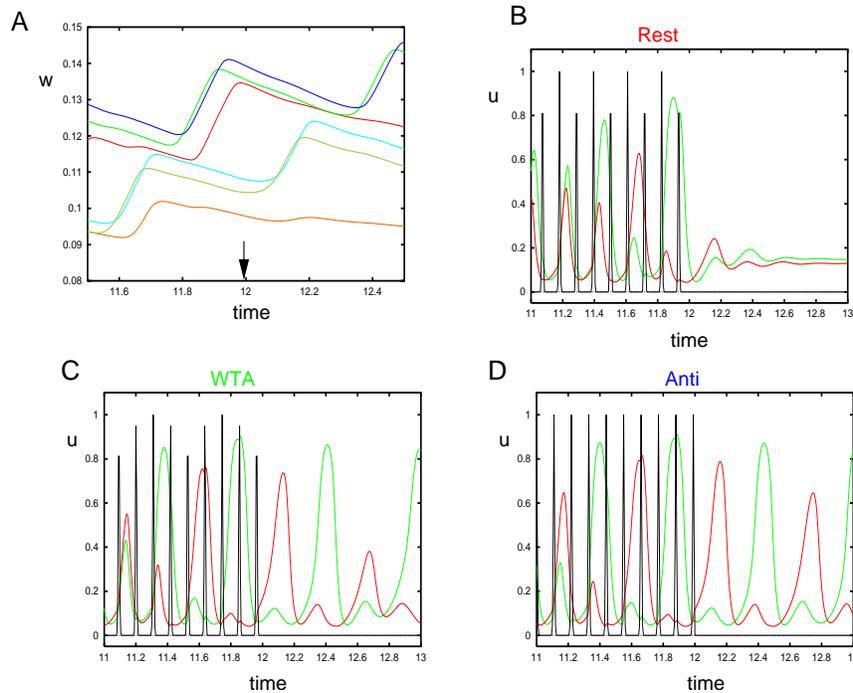


Figure 3.8: Behavior of the two-population network at the termination of a two second stimulus with different periods and with the non-preferred amplitude of 0.8. (A) the facilitation variables, w_1, w_2 at three different periods leading to three different states: green/olive (period, 0.10875) WTA with green (preferred) winning; blue/cyan (period 0.11) antiphase-oscillations; red/orange (period 0.1075) both die. (B) preferred (green), nonpreferred (red), and stimulus (black) when in the rest state basin; (C) Same as (B) with WTA (low amplitude in the oscillations of the nonpreferred population); (D) same as (B) when anti-phase oscillation occurs.

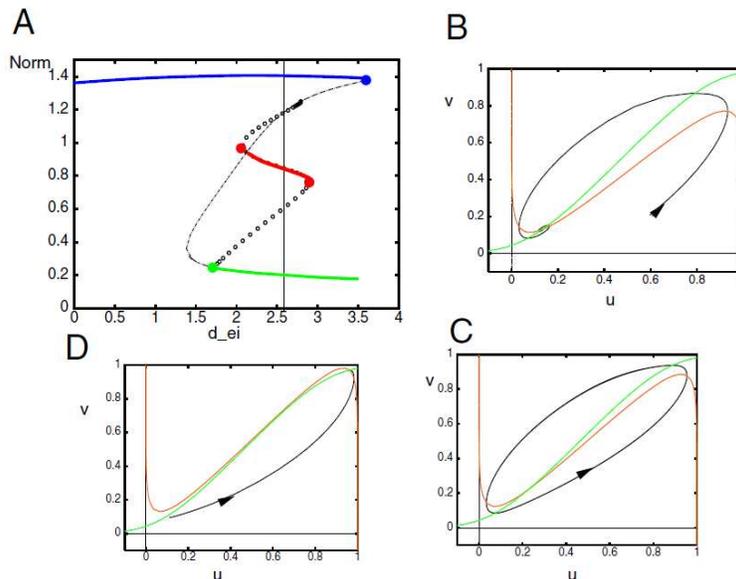


Figure 3.9: (A) Bifurcation diagram for the single population model as the cross inhibition, d_{ei} varies. Green curve is quiescent rest state, blue curve is the excited state, and red is a branch of periodic oscillations. Vertical line corresponds to $d_{ei} = 2.6$. (B, C, D) represent phase-planes of the $u-v$ system with w frozen at the steady state or average value along the line in (A). (B) w held at the value on the green curve, (C) red curve, and (D) blue curve.

Figure 3.9 shows the bifurcation diagram for the three-dimensional model when d_{ei} , the cross E-I coupling, varies. For large values, the only symmetric solution which exists is the quiescent state. As d_{ei} decreases, there is a fold and the upper symmetric state (all-on) appears and remains stable for all lower values of d_{ei} . At $d_{ei} \approx 1.713$ the lower quiescent state loses stability at a subcritical Hopf bifurcation and then loses existence at a fold. The subcritical branch of periodics turns around at $d_{ei} \approx 2.89$ and becomes a stable branch of periodic solutions. This branch again loses stability at a fold at $d_{ei} \approx 2.08$. Thus for $2.08 < d_{ei} < 2.89$ there is a stable periodic solution, two stable fixed points, two unstable periodic orbits and an unstable fixed point. w varies slowly due to its long time constant and even on periodic branches, it varies only over a small range of values. We fix $d_{ei} = 2.6$ and hold w at its steady state or average values corresponding to the three stable behaviors shown in the bifurcation diagram. Figures 3.9B,C,D show the phase-plane dynamics for the $u - v$ system with w frozen. In each case, there is a unique stable attractor corresponding to the three states in the bifurcation diagram.

Treating, w as a parameter, we can write, $u(t) = U(t; w)$ where U is the solution along the bifurcation diagram in figure 3.9A. Along the blue and green branches, U is independent of time and along the red branch, it is periodic. The slow w dynamics evolve according to equation (3.3). Since τ_w is large, we replace u by the steady state, $U(t; w)$ and obtain:

$$\tau_w \frac{dw}{dt} = -w + f(\gamma(U(t; w) - \theta_w))[w_{max} - w].$$

We note that along the equilibrium branch, $U(t; w)$ is a function of w only, so the right-hand side is only a function of w . Along the periodic branch, we can average and again obtain a function of w . Thus, we reduce the w dynamics to an equation of the form:

$$\tau_w \frac{dw}{dt} = -w + \langle f(\gamma(U(t; w) - \theta_w))[w_{max} - w] \rangle := -w + G(w). \quad (3.7)$$

Thus, we can plot the average, $G(w)$ evaluated along branches of the solutions to the fast (u, v) -dynamics with w as a parameter. Figure 3.10 shows $G(w)$ versus w for several values of d_{ei} along with the identity line, $y = w$. Intersections of $G(w)$ with w correspond to solutions to the full three-dimensional system which are either equilibria or periodic solutions. For example, with the normal value of $d_{ei} = 3.5$, there are two stable solutions, one in which

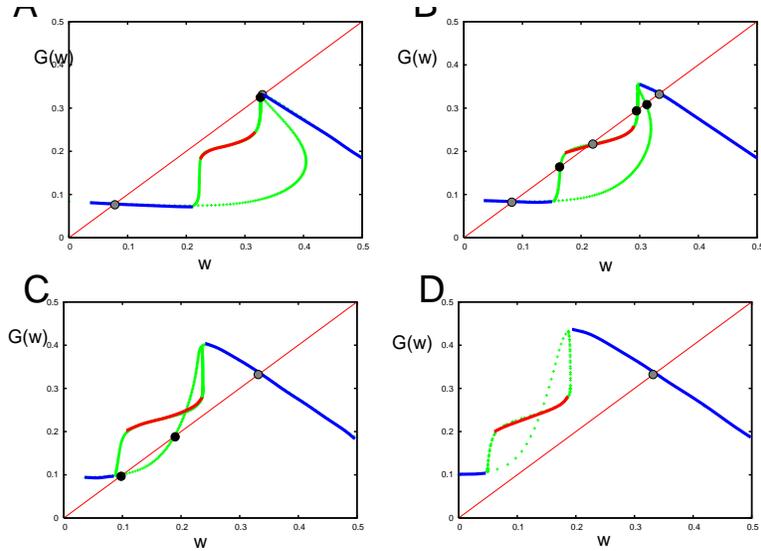


Figure 3.10: The curves depict $G(w)$ as a function of w . (See equation (3.7).) Blue curves are stable equilibria to the fast dynamics and red thick curves are stable periodic solutions. Grey circles correspond to stable solutions and black circles to unstable. Each diagram is for a different value of d_{ei} . (A) $d_{ei} = 3.5$, (B) $d_{ei} = 2.6$, (C) $d_{ei} = 2$, $d_{ei} = 1$.

the network is quiescent and one in which the population is excited (lower and upper grey circles, respectively). As d_{ei} is lowered, the curve of values, $G(w)$ rises vertically and the middle branch (red) of stable periodic orbits intersects the diagonal line. This “fixed point” represents a stable branch of periodic solutions to the full model and a synchronous oscillatory solution to the full two- (or more generally, N -) population system. As can be seen from figure 3.10B, where $d_{ei} = 2.6$, there are 6 fixed points corresponding to the two stable resting states (left- and right-most fixed points) and the synchronous orbit. Consider again figure 3.9. The vertical line corresponds to $d_{ei} = 2.6$. There are two unstable periodic orbits, one stable periodic orbit, two stable equilibria and one unstable equilibrium just as would be predicted from the slow-fast decomposition in figure 3.10B. As d_{ei} is raised further to 2, the branch of periodics is lifted above the diagonal and the stable lower equilibrium point is shifted towards and onto the unstable equilibrium of the fast dynamics (see figure 3.10C). The only stable solution to the three variable model is the upper active state. Finally, for $d_{ei} = 1$, the only equilibrium, stable or otherwise, is the upper state.

We can also use this separation of time scales to understand the dependence on frequency of the stimulation. In particular, we can see why very fast and very slow stimuli are ineffective in exciting the network. Figure 3.11A shows the evolution of the facilitation, w for 4 different

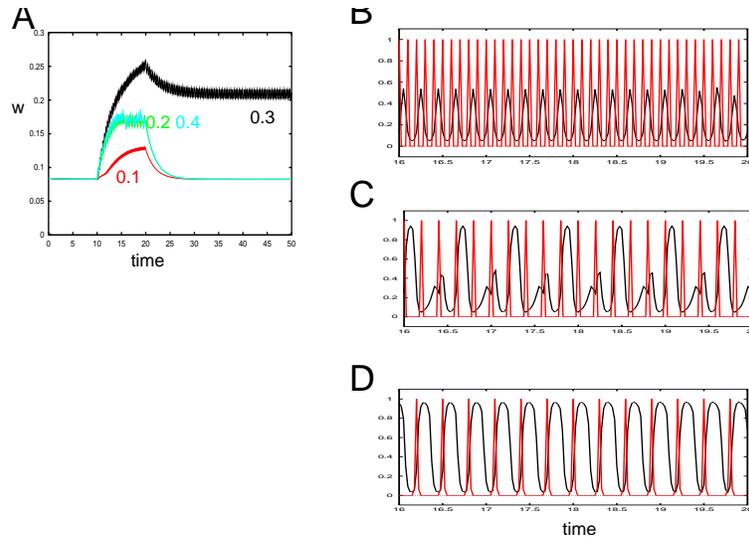


Figure 3.11: (A) The evolution of w during periodic stimuli lasting 10 seconds. Period of the stimulus is shown next to each curve. (B-C) Evolution of $u(t)$ during stimuli (red curves show the periodic stimulus) for a period of 0.1, 0.2, and 0.3 seconds.

periods of input lasting a total of 10 seconds each for somewhat reduced inhibition ($d_{ei} = 3$). Only the stimulus with period 0.3 seconds is sufficient to push the network into an excited state. In this reduced inhibition case, the fast subsystem (holding w at its resting value) is an excitable medium; there is a stable rest state but amplification before return to rest. Once the population is excited, however, it needs time for the inhibition to wear off before it is excited again. Thus, if the frequency of the stimulus is too high, the population can either never fire again or fire only on a fraction of the cycles (cf figures 3.11B,C). However at a low enough frequency, the excitatory population fires at every cycle (figure 3.11D) allowing the facilitation to build up and effect the switch into an excited state. For lower frequencies, 1:1 locking still occurs and the excitatory population fires on every cycle, but, the time between firing is such that the w can never reach a sufficient level to push the medium into an excited state. Thus, for intermediate frequencies, we can push the network into an excited state.

3.6 SEIZURE TERMINATION

A working memory network would be of no use if the persistent activity of its populations could not be terminated. It is therefore pertinent to study how the states which are reached after the stimulation of the network can be reverted back to the baseline state using a second stimulus. We study termination in three general cases: first during normal network behavior, when inhibition hasn't been disrupted and there are only one or two populations active simultaneously. The second case we study is when inhibition has been disrupted so that there are multiple populations displaying high activity, and the third case addresses the termination of oscillatory behavior. These last two could also offer some insight into possible mechanisms for the termination of ictal activity.

In the case with normal inhibition there are several ways to revert the state back to baseline. The most straightforward one consists of exciting all the populations so that lateral inhibition shuts down the active one (Figure 3.12). This method of terminating the activity is fragile, since any reduction in the inhibition will render it ineffective, and it requires the activation of nearly all populations. Moreover, there is again frequency sensitivity, with some frequencies of the stimulus being better suited to turn down the activity. Increasing the duration of the stimulus is a way to enlarge the range of frequencies that can turn down the activity. Combining selective stimulation of the inhibitory component of the active population with the stimulation of the rest of the populations largely reduces the number of populations which need to be stimulated, and allows termination with a broader set of frequencies (Figure 3.12).

In the case where we have a large number of populations active simultaneously (following a breakdown of inhibition), it is no longer possible to turn down the activity by exciting all populations (this may result in more populations becoming active). In fact, when the inhibition has been reduced to the point where exciting one population may recruit several others, the direct stimulation of the inhibitory component of the active populations is not sufficient to turn them off; a direct inhibitory stimulus to the excitatory component of the active populations is required in order to terminate the activity (Figure 3.13). Terminating the activity one population at a time requires less inhibition than the simultaneous termination

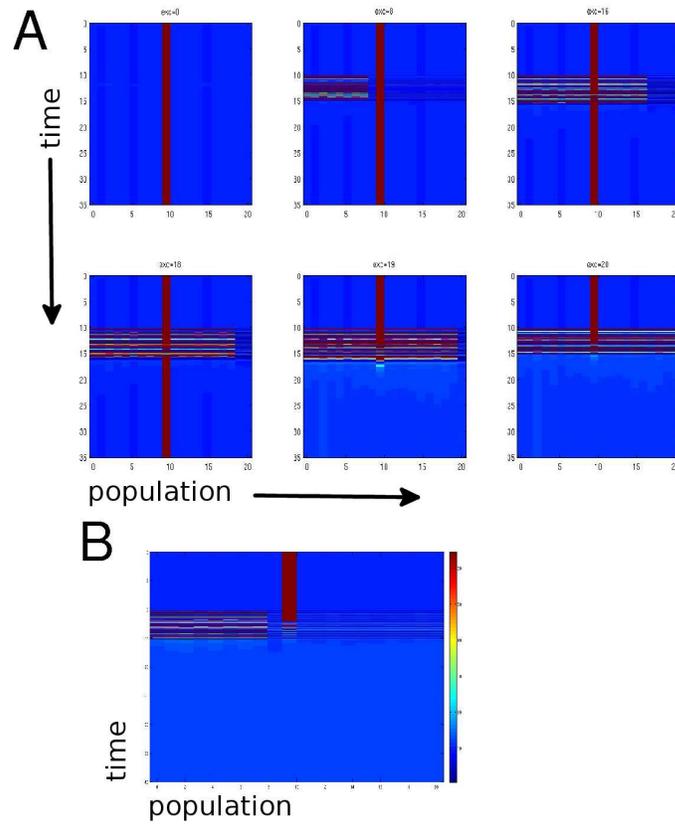


Figure 3.12: Terminating the activity of a single active population under full inhibition ($d_{ei} = 5.4$). Panel (A) shows the effect of stimulating the first k populations. Notice how the active population (number 10) only goes back to baseline when 19 or 20 populations are stimulated. Panel (B) shows the same network with the same initial state, but in this case the inhibitory component of population 10 is stimulated along with the first 8 populations. In both cases the frequency of the stimulus is 3.3 Hz.

of all activity. As in the prior case, the result of the stimulation is frequency dependent.

Unlike the case where we have a large group of active populations, both synchronous and antiphase oscillatory behavior can be turned off by a purely excitatory stimulus applied to a subset of the populations (Figure 3.14). Depending on the amount of inhibition present and on the strength of the stimulus the network state may evolve into one with many active populations. That is, the transition from oscillations may not necessarily take the system back to rest. This phenomenon can be understood using the fast-slow decomposition in the one population model of section 5. When a subset of the populations is excited, lateral inhibition to the rest of the populations is created along with the excitation. If the net effect is inhibitory, the average value of the variable u during the limit cycle will drop down (Figure 3.9B,C), and if it is excitatory the average value of u will increase. The change in the average value of u will cause the variable w to respectively decrease or increase, and that change will lead it away from the basin of attraction of the oscillatory regime (Figure 3.10B). Once a few populations stop oscillating the generalized oscillations are no longer stable, and each population falls into one of the stable attractors left, usually baseline or high activity.

In addition to terminating oscillating behavior once it has been initiated, it is interesting to observe that the onset of oscillations can be prevented by applying a strong excitatory stimulus to some of the populations along with the stimulus which would otherwise cause all the populations to oscillate, as can be observed in Figure 3.15. This phenomenon is not puzzling if we once again consider the fast-slow analysis of section 5, and notice how the appearance of stable oscillations (Figure 3.10) requires a balance in the average values of the u variables: stimuli which are too strong or too weak cannot lead the system into stable oscillations.

3.7 DISCUSSION

This work presents a physiologically-based model of working memory yielding a potential generalized description of epilepsy or seizure-like behavior. The basic premise is that seizures result from inherent states in working memory networks that come about through disinhi-

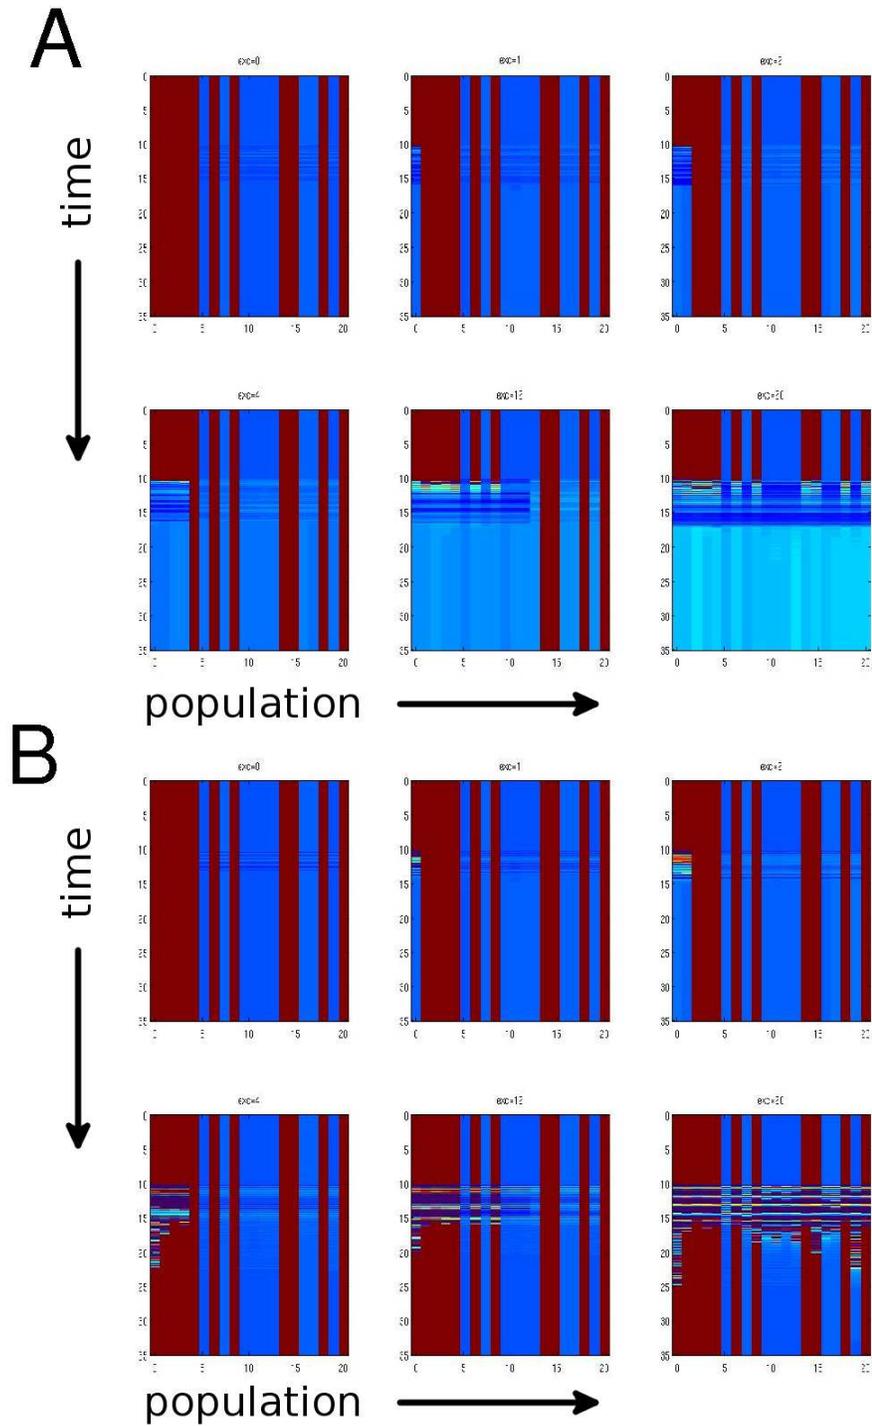


Figure 3.13: A network with low inhibition ($d_{ei} = 1.59$) and many excited populations may be reset to baseline using inhibitory inputs to the excited populations. Panel A shows a 10 Hz inhibitory stimulus being applied to 0,1,2,4,12, and 20 populations, starting with the leftmost one. Panel B is similar, but the stimulus has a frequency of 3.3 Hz. Note that in the top panel whenever a population is inhibited it goes to baseline, whereas in the bottom panel this only happens when the number of populations inhibited is small.

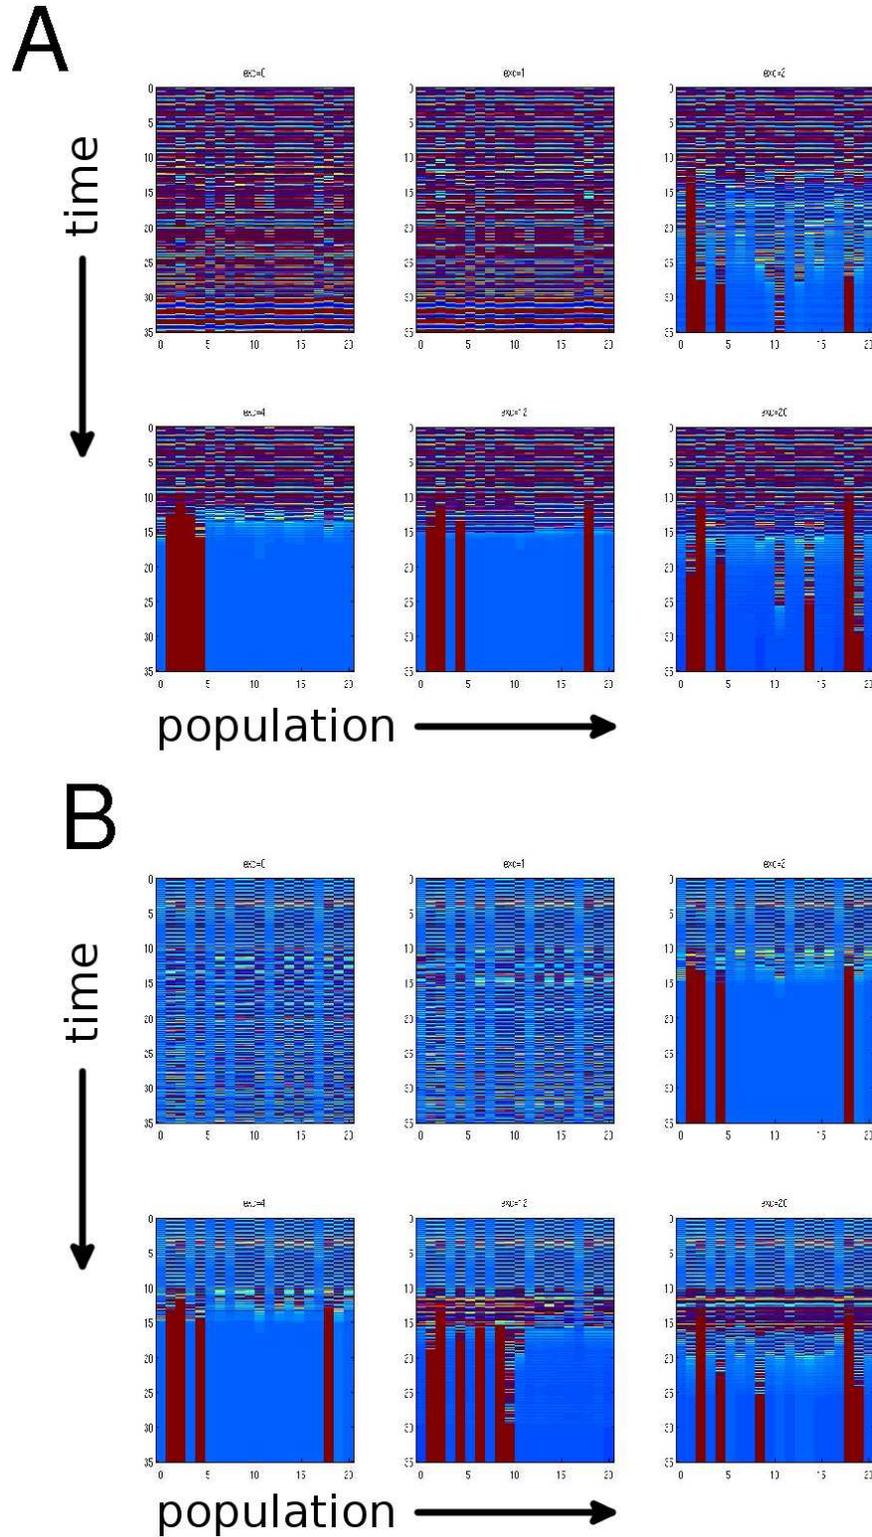


Figure 3.14: Oscillations can be terminated through the application of an excitatory stimulus to the excitatory and inhibitory components of a set of populations. The left panel shows the stimulus being applied to 0,1,2,4,12, and 20 populations in a network initially displaying synchronous oscillatory behavior. The right panel shows the same stimuli being applied to a network initially displaying antiphase oscillations. In both cases the frequency of the stimuli is 3.3 Hz and the inhibition is $d_{ei} = 2$.

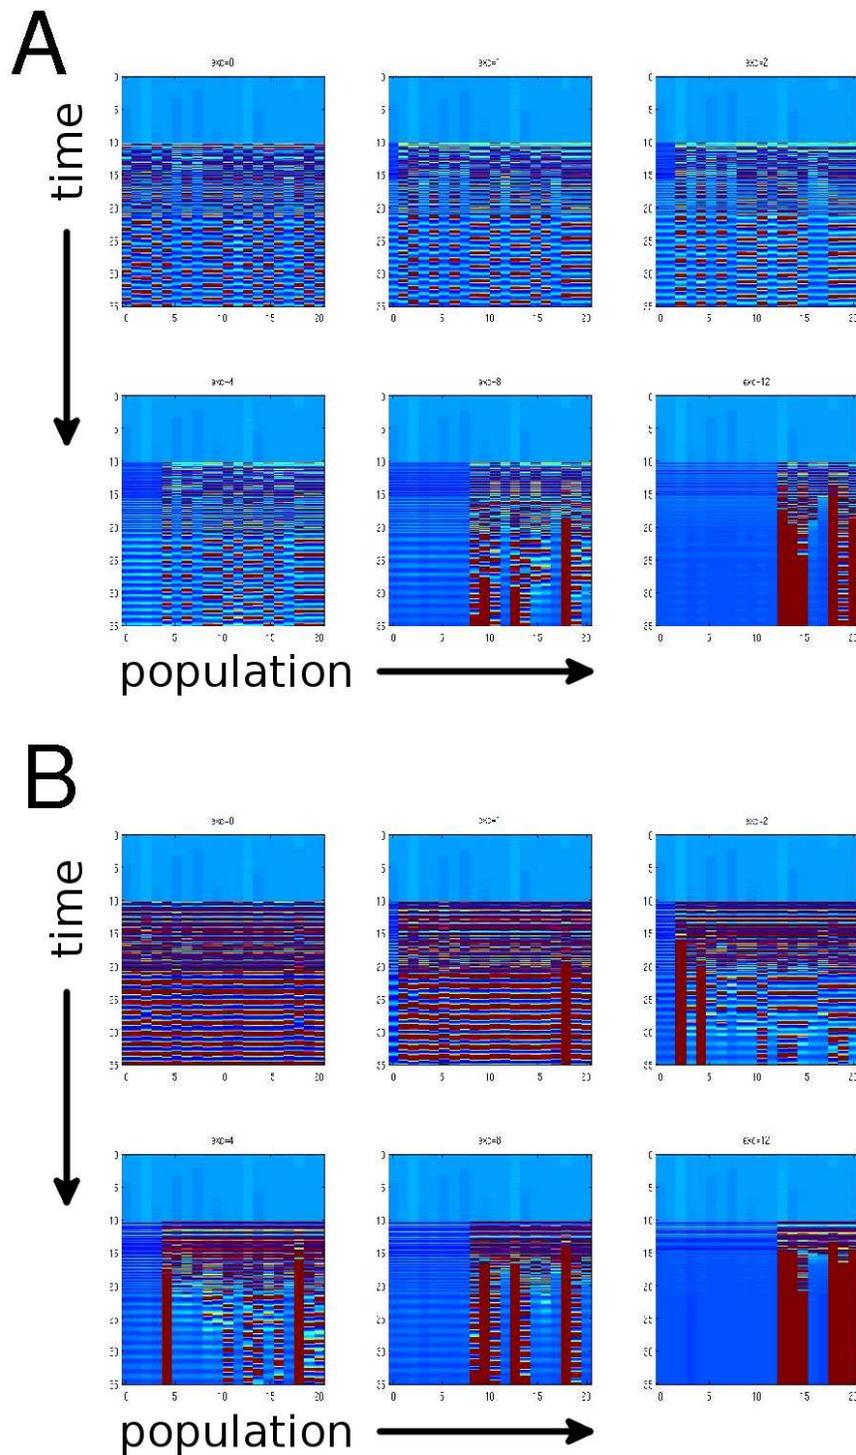


Figure 3.15: The onset of oscillations can be prevented when a certain number of populations receive a larger excitation than the rest in a case of low inhibition ($d_{ei} = 2$). In panel (A) there are 0,1,2,4,12, or 20 populations receiving the larger stimulus which tends to prevent synchronous oscillations (notice how the system goes into synchronous oscillations when no populations receive the large stimulus). In this case the stimulus frequency is 3.3 Hz. Panel (B) is analogous to panel (A), but the stimulus frequency is 5 Hz, which tends to drive the network into antiphase oscillations.

bition in the neuronal populations (either inherent imbalances between excitatory and inhibitory synapses or damage) resulting in a loss of population selectivity and potentially, concomitantly, a loss in stability of fixed point attractors. A critical component of the working memory network is dynamic synapses through facilitation which in normal working memory regions of parameter space enables the stable selective activation - to a particular input stimulus - of a given population. Changes in model parameters however, specifically reduction in inter-population inhibition, produce a series of bifurcations such that both normal and pathological states coexist for the same network parameter settings and external input (or internal perturbations) can trigger transitions from normal working memory to seizure-like behavior. Specifically here we consider a model indicative of a common type of reflex epilepsy, in which rhythmic stimulus input to hyper-excitable cortex produces seizures.

The network was not designed or fine-tuned to exhibit seizure-like or ictal activity, but rather such states and behaviors are inherent over wide ranges of the parameter space of the normal working memory network. The network exhibits working memory behavior with sufficient lateral inhibition strength between populations. Following a typical working memory paradigm, there is a baseline period during which the network populations reside in a stable attractor and exhibit resting-state firing-rate levels. During presentation of a stimulus- which is to be held active in short term memory- there is an external input (representing the stimulus) to the networks populations that subsequently causes an increase in firing frequency. After this, the external input is terminated, and a delay period ensues in which the information about the stimulus must be retained ultimately for use in some subsequent behavioral or motor response. During this delay period, a specific population (or subset of populations) representing the stimulus information being held maintains persistent activation (above-baseline elevated firing rate). The network exhibits specificity in two ways. First only a given population becomes activated (i.e. winner-take-all) as a result of the afferent memorandum stimulus. Second, whether or not the population becomes active is a function of the particular frequency of the input. Thus specific frequencies of inputs represent a memorandum, and a particular population responds to that input preferentially and becomes persistently activated, while the activity of other populations remain at baseline levels. Further this working memory activity of the network reproduces the persistent

patterned behaviors and firing statistics observed in real cortical cell populations recorded during the performance of working memory tasks. These results are presented elsewhere [200] (Chapter 2). The canonical working memory activity can be seen to be present from the schematic of the bifurcation diagram of the 2-population model for sufficiently high values of the lateral inhibition. As inhibition is reduced, while working memory behavior is still present, multiple potentially pathological states that the network may adopt as a result of specific stimuli become possible through a series of bifurcations.

Epilepsy has previously been suggested to be a dynamical disease and previous work has suggested dynamical processes leading to seizure generation including deformation of system attractors induced by changes of network parameters that lead from normal to ictal activity, bifurcations in a system possessing both normal and pathological states coexistent for the same parameter setting such that external input or internal perturbations trigger sharp transitions from normal to epileptic behavior, and a mixture of both scenarios with gradual parameter variations facilitating the transition from normal to an ictal state [129]. In the present work, we concentrate on the second of these (the coexistence of normal and pathological states). However all 3 of these routes to seizures are present in the model. Particularly facilitation in the model can create changes in the relative excitation and inhibition of the model which results in a deformation of the attractor structure.

We consider pathological activity in this work, specifically seizure activity, to be a loss of selectivity. That is, the ability of a specific subset of populations to become activated by a given stimulus input breaks down, and multiple populations are recruited by the stimulus in a nonlinear fashion. This is the general and perhaps the most common trait of all seizures and types of epilepsy. In the network we see that the recruitment of populations in pathological activity is such that different stimuli induce a loss of selectivity, with the number of populations activated (the degree of spread of the seizure) increasing in a nonmonotonic fashion. While synchronous activity of multiple populations has been implicated in normal cognitive function, it can be a double edged sword when that activation spreads. The dynamics under which normal binding and pathological recruitment and loss of selectivity occurs is as yet not understood. The elucidation of these mechanisms can lead to a better understanding of how seizures propagate and might be controlled. The specific dynamics exhibited by the

activated pathological network are such that they can exhibit a range of population activities which include fixed firing rates, synchronous oscillations, and anti-phase oscillations. Such varied states are typical of seizure in different types of epilepsies, or indeed might be observed within a given seizure [58]. In the present model, the populations of the working memory networks can transition to all of these varied behaviors.

In the model, recruitment can occur along a range of different paths exhibiting different dynamics. As can be seen from the schematic of the bifurcation diagram of two interconnected populations (Figure 3.5)—which generalizes to many populations—as inhibition is decreased the stable fixed firing rest state undergoes a Hopf bifurcation that after a pitchfork bifurcation ultimately (after a further fold bifurcation) results in winner-take-all working memory behavior. Thus normal working memory behavior is still possible in the deranged network. This is indeed the case in human epilepsy in which seizures do not occur the majority of the time. However, it is also possible from the Hopf bifurcation, for the network to proceed to a state in which all the populations of the network are active exhibiting synchronous or anti-phase oscillations. Thus the ultimate seizure state may involve hypersynchronicity, weak synchronicity, or periodic behavior depending on the specific network parameters.

A vital component of the present model is that transitions to specific states can be a function of the periodicity of the external stimulus. The dependence of transitions to a seizure-activity attractor on frequency relates to a model of reflex epilepsy. That is epilepsies involve seizures resulting from exposure to a particular external or internal stimulus (often periodic). Facilitation and dynamic synapses which have been implicated in working memory, here play a central role in which resonance with a given external stimuli causes pathological activity [142, 200]. This has been suggested to be important in working memory networks. Structural changes and particular inputs can cause the dynamic synapse mechanisms to play a fundamental role in changing the state from one attractor to another (acting as a switch), going from normal to pathological activity. The fact that such a high percentage of people exhibit a seizure in their lifetime without developing epilepsy may indicate that this is an inherent feature. More permanent parameter changes caused, for example, through learning or trauma might bias activity towards the pathological region of the state space.

An understanding that all of these behaviors can be inherent in working memory networks and how they are related might lead to potential therapeutic interventions. In the model of reflex epilepsy presented here, we see that while pathological activity is induced by specific stimuli, we also see that specific inputs are capable of terminating seizure activity once initiated, or prevent seizures from occurring depending on the specific dynamics of the seizure. In the case of termination of seizure activity once initiated, this may have relevance to the mechanisms involved in recent attempts to control seizures through electrical stimulation of the cortex [17,118,143,198,199]. In figure 3.14 we show that a general excitation of the populations, when the specific dynamics of seizure activity involves oscillations (synchronous or asynchronous), results in the termination of the seizure. Specifically from the schematic of the bifurcation diagram, we see that the general stimulation induces a transition from the stable synchronous oscillation state to the baseline state through modulation of the facilitation. Thus electrical stimulation may be most efficacious in treating seizures with that particular type of dynamics. In the case of prevention of seizure activity, recently evidence has been accumulating indicating that stimulation of the cortex with specifically patterned sensory input (i.e. particular music) can reduce or eliminate pathological interictal activity with a resulting reduction or even elimination of seizures in particular cases [100,101,119,174,194]. The mechanism for this intervention might be related to the dynamics examined in Figure 3.15. Here we see that the excitation of multiple populations by inputs of specific frequencies prevents the transition of the network to a pathological state. This models the activity of the musical stimulus which has been demonstrated to strongly excite a widely distributed population of neurons associated with working memory networks [20,147]. Recent evidence has indicated that long term exposure may result in a long-term shifting of the attractors away from pathological states.

4.0 A MODEL FOR COMPLEX SEQUENCE LEARNING AND REPRODUCTION IN NEURAL POPULATIONS

4.1 INTRODUCTION

If neuronal activity had no temporal structure, functions like movement and prediction would be severely impaired. This has motivated various studies of how the brain processes temporal information [31,32,112]. In addition, several researchers have reported a sequential structure in the activity of mammalian neocortex, both at the level of individual neurons [27,102,148,155,158,186], and local field potentials [130,131,186,187], showing that sequences can happen both at the level of single neurons, and of neuronal populations. How this sequential structure is acquired remains an open problem.

The treatment of temporal sequences of neural activity appeared early in neural network literature (e.g. [5,28]), and a variety of models have been proposed to explain sequential activity in the realm of biological modeling, as well as intelligent systems [181]. Different models have been used for the interrelated tasks of prediction, generation, and recognition of sequences, as well as sequential decision making. The first model networks which learned and reproduced sequences would usually work in discrete time, use neural elements with discrete states, impose a particular connectivity pattern, or would require some form of supervised learning ([53,86,108,113,141,173,177]), and the models with the greatest capabilities usually have one or more of these features (e.g. [48,202,210,212]). Models with increased biological plausibility have been introduced [57,89,103,109,128], but some challenges remain unresolved. Perhaps the most pervasive problem concerns robust autonomous learning, with dynamics that permit the reproduction of complex sequences. Examples like the formation of trajectories in hippocampus [148] and the learning of structure in birdsong [24] also

suggest the importance of external cues shaping the sequence.

The reproduction of a simple sequence is a phenomenon which can be regarded as the propagation of a neuronal signal. It is known that signals which propagate over individual neurons tend to be unreliable due to the stochastic regime in which neurons operate. Thus most models for propagation of activity distribute their signals over groups of neurons which operate redundantly in order to increase their robustness. Depending on whether the neurons in each group fire synchronously or asynchronously, we will have a synfire chain model [3, 36, 47], or a firing-rate propagation model [197, 201]. Several models for sequential activity have used synfire chain dynamics (e.g. [89, 109]); the precision they show in their spike timing seems promising when thinking about the autonomous formation of propagation pathways using spike-timing-dependent plasticity (STDP), and they can create associations between groups which may be germane for compositionality [48]. STDP is an attractive choice for the autonomous formation of sequences, and there have been recent results in how this may be achieved in single neurons [57, 103] and synfire chains [109]. These studies did not concern themselves with how complex sequences could be learned and reproduced or how external inputs could quickly form or modify sequences. The reproduction of a complex sequence is only possible when the sequence’s context is maintained by some form of short-term memory. The common ways to maintain the context in memory are delays (as in [86]), context units (as in [108]), or the state of a high-dimensional system (as in [183]). Persistent activity is not usually used as an option to maintain a sequence’s context, despite being regarded as the neuronal substrate of working memory [63, 70, 164]. Persistent activity would indeed be difficult to combine with STDP learning of complex sequences, and a network operating either in the synfire regime or with single neuron propagation. In this paper we provide a model which combines learning and reproduction of complex sequences through temporally asymmetric Hebbian learning, maintaining the context information through persistent activity. We achieve this by assuming the existence of densely connected populations with projections experiencing temporally asymmetric Hebbian plasticity; this simple assumption has important consequences in the creation of excitation pathways, with signals traveling from one population to another (a somewhat similar assumption is presented in [2]).

In the first part of the paper we show the feasibility of having temporally asymmetric

Hebbian learning at the population level. By combining synchronous and asynchronous dynamics in a single learning mechanism we show that this type of plasticity is biologically plausible when applied to the mean synaptic weight of the connections from one population to another. Next we present a firing-rate model capable of storing and reproducing complex input sequences of small order. We describe the dynamics of our model, and the role of its distinct elements. The transition of the activity among distinct populations is enabled by the interplay of excitation and inhibition, which causes slow frequency excitatory activity and higher frequency inhibitory activity with gamma oscillations in the mean firing rate [88]. As in synfire models, we find that correlated firing among different populations may lead to new connections forming between them; in our model, however, the level of inhibition controls whether or not different populations can associate. We also show that by forming closed-loop sequences the activity can be maintained until an exterior stimulus turns it off (by synchronizing all populations) or initiates another sequence (which by competition will terminate the activity in the first sequence). This constitutes a possible substrate for working memory, and for the generation of arbitrary rhythms, which relates our work to these biological functions.

4.2 RESULTS

4.2.1 Temporally Asymmetric Hebbian Learning on Asynchronous Populations

An important component of the model presented in this paper is a temporally asymmetric Hebbian plasticity rule which changes the connectivity among populations depending on the timing of changes in their mean firing rate. Such a rule could be justified if given a pair of small neuronal populations A and B , interconnected with synapses following an antisymmetric STDP rule, we could have that an increase in the firing rate of A followed by an increase in the firing rate of B would lead to an increase in the mean strength of connections from neurons in A towards neurons in B , and a decrease in the mean strength of connections from neurons in B towards neurons in A . Existing studies on populations

of neurons with STDP connections indicate that the distribution of synaptic weights in a spiking network with additive STDP tends to become bimodal [84, 167, 178]. We prefer to use a weight-dependent STDP rule to avoid the hard nonlinearity at saturation and to better adjust to experimental data [145]. Weight-dependent STDP rules tend to produce unimodal distributions unless the right correlations exist between presynaptic and postsynaptic activity [84, 167]. A basic characteristic of these studies was the use of initial homogeneous random connections. For the spiking model in this section we instead assume the existence of densely connected neuronal populations; neurons in different populations communicate using STDP connections. This breaks the homogeneity in the network, and helps to avoid problems such as synfire explosions and the difficulties in creating excitation pathways with synchronous inputs and STDP plasticity ([144]; Kunkel et al. 2010, submitted). In order to create connections between populations we stimulated them with Poisson spike trains, and found that by using the right timing in our inputs we could strengthen the the projections from the first stimulated population to the second stimulated population, and weaken the opposite projections (as will be described below). It should be remarked that the model with spiking neurons only intended to show the feasibility of temporally asymmetric Hebbian plasticity at the population level; it is not concerned with sequential activity and is not meant to have the same dynamics as the firing rate model in this paper.

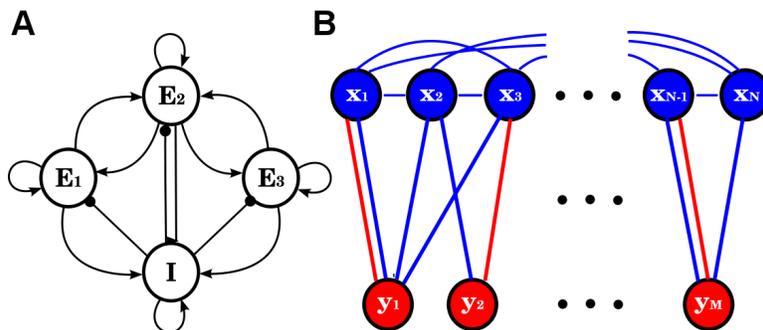


Figure 4.1: **Architecture of the model networks.** A) The spiking model consists of four excitatory populations (circles labeled E_i) and one inhibitory population. Arrowheads denote excitatory connections, and dots inhibitory connections. B) The firing rate model consists of N excitatory units (blue circles, labeled x_i), and M inhibitory units (red circles, labeled y_i). The excitatory units have weak all-to-all connectivity. The probability of x_i sending a projection to y_j is 0.5, which is also the probability of y_j sending a projection to x_i .

We used a simple model network of conductance-based integrate and fire neurons (figure 4.1A) to test our ideas. The network consisted of four populations, three excitatory (100 neurons each), and one inhibitory (80 neurons). The neurons in each population randomly sent

projections towards half of the neurons in each of the other populations. The excitatory-to-excitatory connections used STDP synapses with a power law exponent $\mu = 0.1$, as described in [84]. Each neuron received 20000 independent Poisson spike trains, 16000 excitatory with a mean firing rate of 4 Hz, and 4000 inhibitory with a mean firing rate of 5 Hz. The neurons in each population had static all-to-all connectivity between them; inhibitory populations had no connections between themselves. Since the excitatory neurons had no constant inputs and were driven by independent Poisson trains, increasing the connection strength within a population could increase its coherence. The strength of the connections within each population was not enough to cause synchronization given the normal Poisson background input, but stronger asynchronous input could increase the coherence in the population to the point where most neurons would fire together (see figure 4.2A-D). This increase in synchrony was beneficial when creating connections between populations, as it would allow a presynaptic population to emit spike volleys which could drive many postsynaptic neurons beyond threshold, in a manner similar to a synfire chain model. The intermediate inhibitory population was included in order to maintain the architecture of the firing rate model to be presented below, but it also had a beneficial function in maintaining the stability of the asynchronous firing, and reducing the positive correlation between the excitatory populations.

As a first foray into the dynamics of our spiking model, we tested whether asynchronous stimulation of population 1 could strengthen its projections towards the other two excitatory populations (2 and 3). The idea was that by increasing the firing rate in 1, the neurons in this population would cause neurons in 2 and 3 to fire, thus strengthening their STDP connections. The problem was that the neurons in 1 firing asynchronously would not provide enough excitation to drive the neurons in 2 and 3 beyond threshold. We overcame this in two ways. The first was to alter the balance between excitation and inhibition so that the distribution of voltages remained near threshold in the postsynaptic populations. Indeed, signals can propagate their firing rates when the subthreshold voltages of the underlying network have a distribution that is broad and close to threshold [197, 201]. The second solution was to have strong enough connections within the populations, so that they would fire synchronously when an additional stimulus was applied to them. The coexistence of synchronous and asynchronous dynamics, as well as the balance between excitation and

inhibition [25] ensured that upon stimulation of population 1, its outgoing projections became stronger, and its incoming excitatory projections became weaker. These results are shown in figure 4.2. An interesting additional consequence was that both targets of population 1 could strengthen their reciprocal excitatory connections, thus creating an ensemble target population (figure 4.2H).

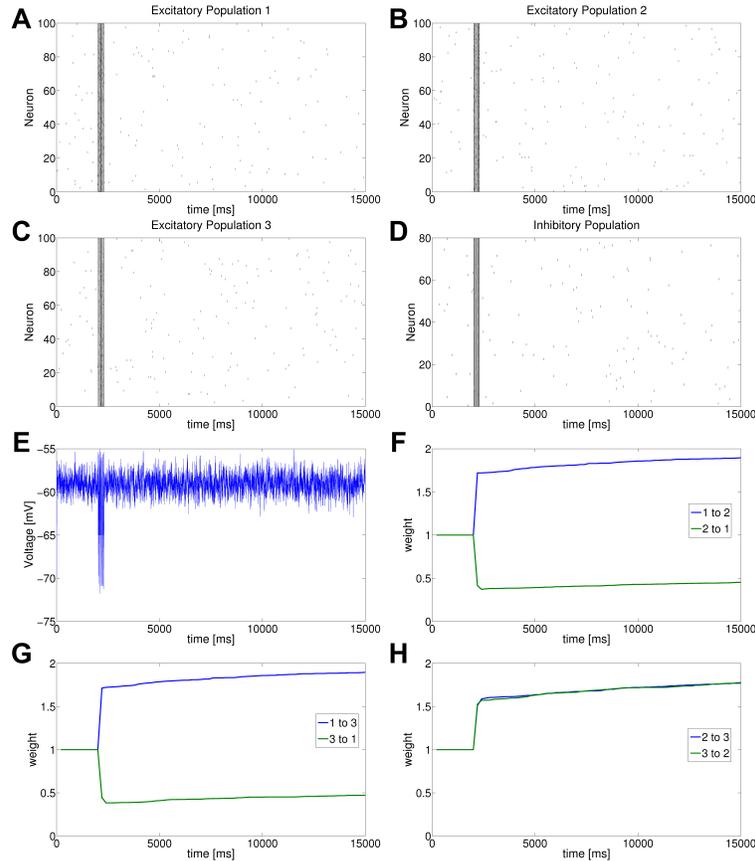


Figure 4.2: **Formation of unidirectional excitation pathways and population ensembles.** A,B,C,D) Spike rasters of the three excitatory populations and the inhibitory population. Independent Poisson stimuli are applied to population 1 from $t = 2000$ ms to $t = 2300$ ms, which causes synchronous activity. The activity in population 1 spreads to the other 3 populations, causing significant changes in the mean synaptic weight between populations. E) Voltage trace of the first neuron in population 1. F) Mean synaptic weight of the connections between populations 1 and 2. The formation of a unidirectional excitation pathway is observed here. G) Mean synaptic weight of the connections between populations 1 and 3. H) Mean synaptic weight of the connections between populations 2 and 3. These two populations experience a bidirectional strengthening of their connections.

The mechanism described so far may not be well suited for the formation of sequential activity among the populations, unless we are capable of selectively changing the balance between the excitation and inhibition of the background input. The reason for this is that the connections being formed are not specific, but arise from the stimulated population towards all the targets with mean voltages near threshold. We can specify a target by changing

the balance in its inputs so that its mean voltage is near threshold, and each transition in a sequence would be formed by a strong stimulus towards the presynaptic population combined with a shift in the excitation and inhibition of the postsynaptic population. If we use a weight-dependent STDP rule, the stability of the changes in the synaptic weights will depend on the firing rate of both populations, which also depends on the balance between excitation and inhibition.

In the last paragraph we mentioned that the stability of the changes in the synaptic weights depends on the firing rate of both populations. Indeed, when two populations are connected with weight-dependent STDP connections following a rule as in [84], and unstructured (Poisson) firing driving the activity, there is a tendency for the mean of the weights to drift towards the value where potentiation and depression are equally scaled (in our case, one half of the maximum weight). The intuition behind this is that when a weight is below this value, potentiation is stronger than depression, and vice versa. Thus, assuming no correlations between the firing of both populations, the mean of their connections will drift towards this fixed point. Having one population excite another breaks the symmetry in the connections, but if the change is small the connections will drift back towards the fixed point, at a rate that will become faster at higher firing rates. By reducing the firing rates we can have changes in connectivity that last for longer periods. On the other hand, if the symmetry in the connections is broken by a large enough alteration, this might introduce correlations in the activity of the populations, and these correlations might reinforce the weight asymmetry. This idea will be used below.

We now present another mechanism that does not require selective modulation of the background input in order to increase connections along a specific pathway, or to increase their permanence. Synchronous activity is used to form connections, and asynchronous activity is used to consolidate them. We use the same model as before, but we reduce the initial value of the STDP connections between the excitatory populations. The effect of this is that the activity in an excitatory population will not have a large effect in the activity of the other excitatory populations. As seen in figure 4.3, we apply an external Poisson stimulus to population 1 for 300 ms, and another Poisson stimulus (duration 225 ms) to population 2 25 ms before the end of the first stimulus. This resulted in a fast decrease in

the mean value of the synaptic weights in the *2 to 1* connections, and a gradual increase in the *1 to 2* connections. The fast changes in connectivity were due to synchronous activity in population *1* followed by synchronous activity in population *2*, all due to the external inputs. The gradual changes came because once the *2 to 1* connections were decreased and the *1 to 2* connections increased, the probability of neurons in population *1* activating neurons in population *2* was larger than the probability of neurons in *2* activating neurons in *1*, so that asynchronous activity increased the initial asymmetry in the connections. As can be seen in figure 4.3, the activity in *1* and *2* had little effect on population *3*, which accounts for the small size of the variations in the connections involving this population. The overall effect of the activity in *1* and *2* was to increase the connections towards population *3*, and decrease its outgoing connections, just as in the previous simulations, but on a smaller scale.

The simulation just described shows that by timing our inputs correctly we can selectively strengthen activation pathways among the populations. The possible consequences of a mechanism like this will be analyzed in the remainder of this paper using a firing rate model to describe the dynamics of a larger set of populations.

4.2.2 Overview of the Firing Rate Model

In order to analyze the behavior of multiple asynchronous neuronal populations interconnected with temporally asymmetric Hebbian plasticity, we describe the activity of each population by a single variable related to its mean firing rate. We refer to these model populations as units. Our full firing rate model contains 80 excitatory units and 16 inhibitory units (figure 4.1B). The excitatory units use an equation similar to those in the Wilson-Cowan model [208], while the inhibitory units follow a simpler, almost linear behavior. We used heterogeneous parameter values for all our equations, which are described in detail in the Methods section. A possible interpretation of our model is that it describes the activity in a single cortical column, with each unit representing the average activity in a minicolumn [146].

Excitatory units are intrinsically affected by firing-rate adaptation and noise. The inputs they receive can be divided in four categories: recurrent self-connectivity, inputs from other

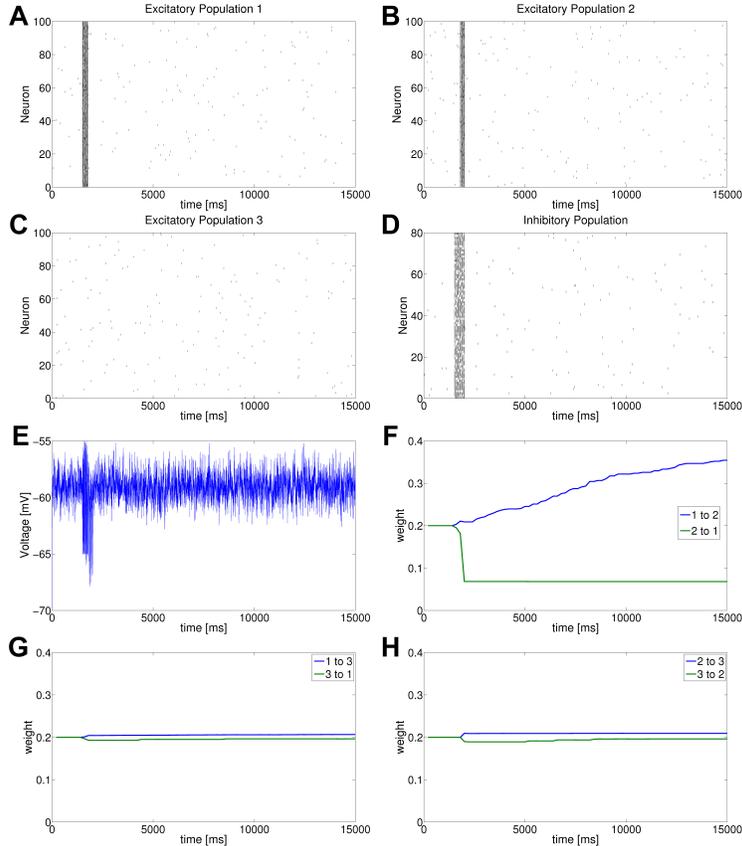


Figure 4.3: **Temporally asymmetric Hebbian plasticity at the population level.** A,B,C,D) Spike rasters for the four populations. Independent Poisson spike trains are applied to the neurons in population 1 from $t = 1500$ ms to $t = 1800$ ms, and to the neurons in population 2 from $t = 1775$ ms to $t = 2000$ ms. The activity in the first two populations had little effect in population 3, but elicited synchronous activity in the inhibitory population. E) Voltage trace of the first neuron in population 1. F) Mean synaptic weight of the connections between populations 1 and 2. A unidirectional connection starts to form during the stimulus period, and consolidates through the spontaneous activity. G) Mean synaptic weight of the connections between populations 1 and 3. H) Mean synaptic weight of the connections between populations 2 and 3.

excitatory units, inputs from inhibitory units, and external stimuli. The matrix of excitatory to excitatory (E-E) connections had all-to-all connectivity, and connections were usually initialized randomly with small values that had little effect on the dynamics. The excitatory-to-inhibitory (E-I) and inhibitory-to-excitatory (I-E) connections were randomly generated, so that given one excitatory unit e and one inhibitory unit i , the probability of e sending a projection to i was $p_{ie} = 0.5$, and the probability of i sending a projection to e was $p_{ei} = 0.5$. The excitatory units of a network whose connections are at the initial state do not have the ability to activate persistently or generate sequential behavior, and in the absence of external input there is only the spontaneous activity arising from the noise, which is not enough to

modify the connectivity. Since our goal is to store and reproduce sequences that come from an external input, it would be disadvantageous to have connections arising spontaneously from noise.

E-E connections experience synaptic plasticity. Our plasticity rule is inspired by the one found in [111], and resembles the temporally asymmetric Hebbian plasticity mechanism described for the spiking model above. The idea behind this plasticity rule is to strengthen the connection from unit A to unit B if A is active shortly before B becomes active, and to weaken that connection if B is active shortly before A becomes active. We added various modifications to the rule in [111], the most important of which is to reduce synaptic potentiation in populations where incoming activity is already high. This permits us to store complex sequences without an unstable growth of connections. Another modification is the introduction of a term which implements either heterosynaptic competition or a slow reduction of synaptic weights of intermediate strength. This term is not required by the model in order to store and retrieve sequences, but it makes it unlikely for long simulations to end with too many strong connections, a condition which can lead to instability. The implementation of this additional term is described in the Methods section.

It has been observed that temporally asymmetric Hebbian learning leads to synchronization [97, 111, 182]. This is also true in our model, which means that sequences that repeat continuously tend to synchronize their components. In order to avoid this effect we included a delay in the excitatory to excitatory connections. Such a delay could correspond to transmission delay down the axons of the activated presynaptic population, or to delay due to neurotransmitter diffusion and reception in the synaptic cleft.

Once E-E connections have been formed our system has dynamics similar to a “ k -winners-take-all” network due to the interaction between the excitatory and inhibitory populations. In our case, however, asymmetric projections and firing-rate adaptation causes one of those k winners to turn off, giving a new unit the opportunity to turn on. This periodic replacement of active units allows the network to reproduce sequences which are stored in the E-E connection matrix. More details are provided below.

4.2.3 Storage and Retrieval of Sequences

Our model has two main traits which enable it to reproduce a complex sequence of degree k : each k consecutive components in the sequence send weighted projections to the $k + 1$ -th component, and the dynamics permit the activation of k components simultaneously (figure 4.4). When the first k components of the sequence have become activated in sequential order (and remain active), the $k + 1$ -th element is unambiguously receiving the most excitation, which ensures that it becomes activated, briefly bringing the number of active units to $k + 1$. Having $k + 1$ units active simultaneously increases the excitation to inhibitory units, increasing their firing rate and the inhibition they send back to excitatory units. This increased inhibition causes one of the $k + 1$ excitatory units to return to baseline. The unit returning to baseline is the one which has been active for the longest, because it has accumulated the most synaptic adaptation. Through this process the next component becomes active and the first component returns to baseline, with a transition rate that depends on the time constants of the model and the delay in their connections. The time between the onsets of contiguous components in a sequence is about 20 milliseconds in our implementation of the model. The periodic onset of new active units in the excitatory population causes inhibitory units to oscillate due the convergent E-I projections. The frequency of these oscillations is around 50 Hz, inside the gamma range. The result is that excitatory units have low firing rates, but the population activity has oscillations in the gamma range, as is found experimentally [88]. The number k of simultaneously active units is controlled by the level of inhibition (the value of E-I and I-E connections), with lower inhibition permitting larger values of k .

The connections required to encode a sequence can be created by activating the components in the correct order a few times using an external input. The input, however, may have no effect on the connections if there is no overlap between successively activated components. In order to create a complex sequence of degree k , the input must be such that k units are active simultaneously, roughly imitating the way that the sequence would look when activated intrinsically. This requirement is a direct result of the plasticity rule (described in the Methods section), and permits the formation of simple sequences which intersect, but can nevertheless be activated unambiguously. This result is illustrated in figures 4.5, 4.6 and 4.7.

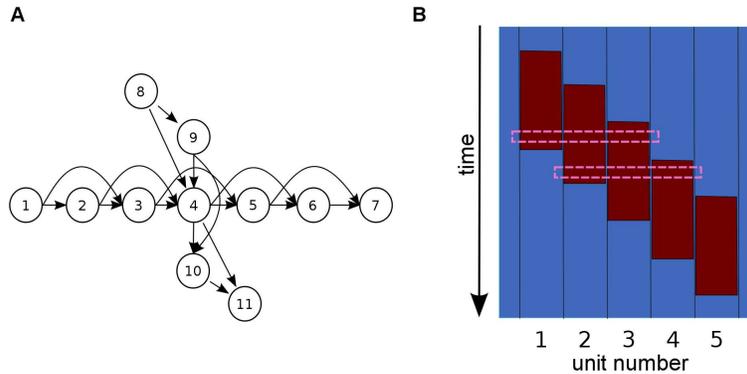


Figure 4.4: **Sequence encoding and reproduction.** A) Each circle represents an excitatory unit, and the arrows represent their projections. There are two simple sequences intersecting at unit 4, and each unit sends projections to the next two units in its corresponding sequence. Using dynamics that allow at least 2 units to be simultaneously active either sequence can be activated without ambiguity after reaching unit 4. B) The sequence 1-2-3-4-5 being generated with a maximum of 3 simultaneously active units, and a minimum of 2. Each column shows how the activity of the corresponding unit evolves through time, with blue representing a baseline value, and red representing a high frequency value. The dashed lines highlight two regions with 3 simultaneously active units.

Figure 4.5 shows the formation of a closed-loop sequence encompassing units 10 through 20, following four presentations of an input sequence. Figure 4.5A shows the activity of all excitatory and inhibitory units as the sequential input is presented after 200 milliseconds of simulated time, starting from the stable baseline state. The plasticity rule used reduction of intermediate-value weights (explained later on in this paper). It can be seen that after four repetitions the input sequence begins to be intrinsically generated by the network, with small deformations due to noise and parameter heterogeneity. The fidelity of reproduction can be altered by the input overlap and by its transition speed. The E-I and I-E connections have a magnitude which permits this intrinsically generated sequence to have at least 3 components active simultaneously, which permits the generation of complex sequences with degree 2, as well as intersections of simple sequences. Figure 4.5B partially shows the connection matrix for excitatory units after the simulation in figure 4.5A was stopped. The encoding of the sequence is clear, with each unit receiving connections from roughly 4 of the units preceding it. The connection matrix reaches stable values (slightly affected by noise fluctuations) despite the plasticity being active all the time, thanks to the delay in the connections. Although it is possible to encode sequences with no closed loops, in the rest of the paper we will present

examples using closed-loop sequences, which permit persistent activity and better illustrate the properties of the model.

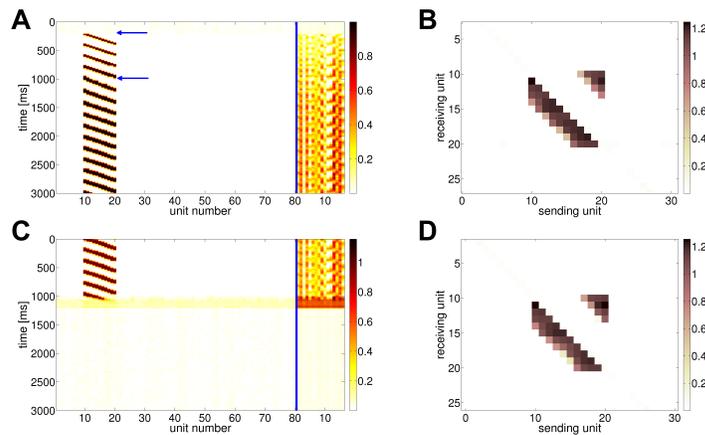


Figure 4.5: Creation of a closed-loop simple sequence. A) This panel shows the activity of all units as an input is being presented to the network. The leftmost 80 columns represent the activity of the excitatory units, and the 16 columns to the right of the black line represent the activity of the inhibitory units. A sequential input is repeated 4 times starting at 200 milliseconds, transitioning from unit 10 towards unit 20, always maintaining an overlap of two units; blue arrows indicate the start and the end of the input-driven activity. The sequence continues to be generated intrinsically after the input is over due to the excitatory connections that were formed, maintaining an overlap of 3 or 4 units. B) E-E connections for the first 30 excitatory units. Notice that each unit sends projections to (roughly) the next 4 units in the sequence. C) Continuation of the simulation shown in panel A. In this panel the activity is stopped by synchronous periodic pulses to all excitatory units, which increases the activity in the inhibitory units. D) Excitatory connections matrix after the simulation in panel C.

The activity generated in figure 4.5A persists due to the recurrent nature of the connections in a closed-loop sequence. This type of reentrant activity in a localized population is at the core of many models of working memory [51, 204]. The persistent activity needs to be terminated at some point, and to do this we use the strategy in [85], which consists of using a stimulus to synchronize the excitatory neurons. The reason why synchronization terminates the activity in [85] is because all neurons firing at once causes all recurrent inputs to arrive during the refractory periods, so that the majority of neurons remain silent after this burst of activity. Although we don't implement refractory periods, synchronization stops the activity in our model for qualitatively similar reasons: when all excitatory units fire synchronously the inhibitory units have a sharp increase in their activity, and this shuts off all the excitatory units simultaneously, leaving the system in a baseline state. In figure 4.5C the same closed-loop sequence formed in figure 4.5A is being stopped by applying a burst of synchronous pulses to all excitatory units with a period of 20 milliseconds. Figure 4.5D shows the matrix of excitatory connections after the activity was stopped by the synchronous

pulses of figure 4.5C. Notice how the synchronous pulses did not destroy the encoding of the sequence.

In order to show the ability of the model to store and reproduce multiple simple sequences that may intersect, the simulation shown in figure 4.5 was continued to store four additional sequences. The procedure to store those sequences consisted of presenting the input sequence five times, and once the sequence was being generated intrinsically all activity was stopped using synchronous input. After the activity was stopped, the next sequence would be presented. The four sequences stored were all closed-loop simple sequences involving units 20 through 30, 30 through 40, 40 through 50, and 50 through 60. Notice that units 20,30,40, and 50 were each part of 2 sequences. At this point one final sequence was added, as shown in figure 4.6. Panel A of figure 4.6 shows the sequence from unit 50 to unit 60 being generated intrinsically until a sequential input with 3 simultaneously active units begins to be applied. Due to the competitive dynamics of the network, the input extinguishes all other activity and stores a sequence going from unit 60 to unit 70, but not before some tenuous connections are formed, from the units of the sequence from 50 to 60, towards the units of the sequence from 60 to 70. These connections appear as a lighter region in the lower right section of panel B. The reduction of intermediate-value weights implemented in the plasticity rule slowly erases the connections resulting from the transitions between stored sequences. In order to illustrate how the activity may be switched from one sequence to another we continued the simulation in panel A of figure 4.6. Panel C shows the input sequence 20-21-22-23-24-25 being applied with no overlap, which results in a switch between active sequences. Panel D shows the excitatory connections after this switch, which show essentially no change.

One point that merits attention is how storing new sequences can modify previously stored ones. The sequences already in the connection matrix are not altered when the sequence being stored does not intersect them. In fact, a newly stored sequence will only affect previously stored sequences when it contains the temporal inverse of a subsequence already present in the connection matrix. We illustrate this result in figure 4.7 using an extreme case where a newly stored sequence completely replaces a previously stored sequence. The simulation in figure 4.7A starts with 6 sequences already stored. Those 6 sequences are

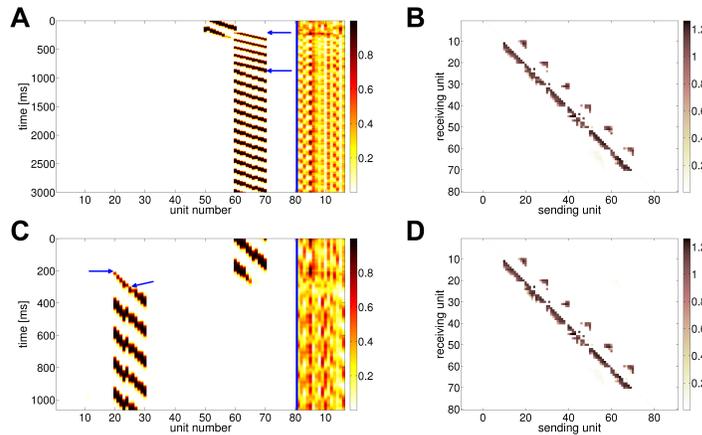


Figure 4.6: **Transition between sequences.** A) The sequence from 50 to 60 is being generated intrinsically at the top, but it stops when an input sequence is applied at 200 msec. The input sequence (from unit 60 to unit 70) is repeated five times, and then it repeats intrinsically. Blue arrows indicate the start and the end of the input. B) The matrix of excitatory connections following the simulation shown in panel A. Six simple closed-loop sequences are stored in this matrix. There are tenuous connections from the *50 through 60* sequence towards the *60 through 70* sequence; they were formed because the *50 through 60* sequence was not stopped before applying the stimulus which encoded the *60 through 70* sequence. C) Once the *60 through 70* sequence is being generated intrinsically, a sequential input with no overlap was presented from unit 20 to unit 25. This switched the activity from the *60 through 70* sequence to the *20 through 30* sequence. Blue arrows indicate the start and the end of the input. D) The connection matrix after the simulation shown in panel C. Notice that the differences in the connection matrices in panels B and D are minimal.

the ones created in the simulations of figures 4.5 and 4.6. An input sequence going from unit 30 towards unit 20 is applied 5 times. This input sequence is the inverse of the sequence going from unit 20 to unit 30, which was already stored in the connections matrix. As seen in the ensuing activity, and in the connections of figure 4.7B, the incoming sequence replaced the previously stored one (compare with figure 4.6B). No traces of the previously stored sequence were left in the connection matrix because the new sequence, being the inverse of the old sequence, weakened its connections each time it repeated. We can see that given two units, the temporally asymmetric plasticity rule can create transitions between them in only one direction, which reduces the set of sequences that can be possibly stored in this manner. This is a consequence of one aspect in the plasticity rule: “*If unit A activates before unit B, reduce the connections from B to A.*” This aspect finds its implementation in the second term of equation (4.4) of the Methods section.

Not all sequential input has to modify the matrix of excitatory connections. When the input presents no overlap it tends to create no significant connections. By having no overlap,

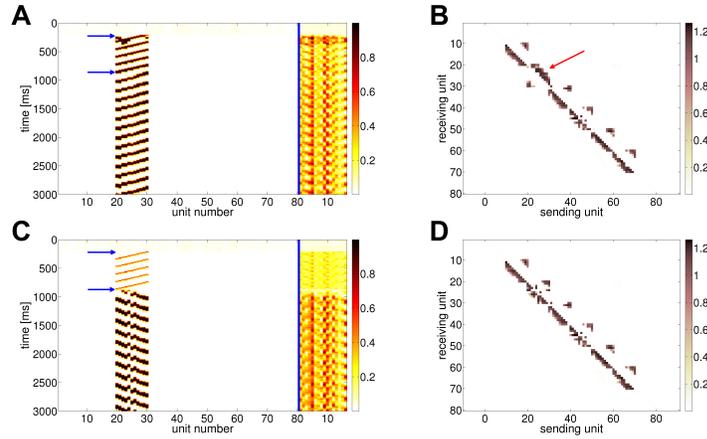


Figure 4.7: **Input with overlap may modify existing connections.** A) The input sequence going from unit 30 to unit 20 is presented 5 times with an overlap of 3 simultaneously active units. After the input stops the sequence is generated intrinsically. Blue arrows indicate the start and the end of the input. B) The matrix of excitatory connections after the simulation in panel A. The sequence from 20 to 30 has been replaced with its temporal inverse from 30 to 20 (but neighboring sequences are not modified). The red arrow indicates the region of the connection matrix where the change took place. C) The input sequence going from unit 30 to unit 20 is presented 5 times with no overlap. The time between onsets of contiguous units in the input sequence is similar to that of the input in panel A. Blue arrows indicate the start and the end of the input. D) Excitatory connections after the simulation in panel C. The connections for the 20 through 30 sequence were only minimally modified.

we mean that the input only activates one unit at a time, instead of activating several units simultaneously. Panels C and D of figure 4.7 show that input with no overlap does not modify the connections in the same manner as input with overlap. The simulation in figure 4.7C starts just like the one in figure 4.7A, but it has different input. The inputs in 4.7A and 4.7C have similar time between the onsets of contiguous units in the sequence, but in 4.7C units are activated one at a time. As can be seen in figure 4.7D, the matrix of excitatory connections was barely modified, in contrast to figure 4.7B. One instance where input with no overlap may modify the connection matrix is when there is a sequence currently active while the input is presented. In this case some projections may be formed between the active sequence and the input; reductions in the level of inhibition make this scenario more likely. Notice also, that the input sequences with the most overlap are the ones most likely to extinguish any other ongoing activity, causing them to be stored without any associations to other local sequences. Associations to sequences which do not have the same inhibitory pool are still possible.

The next result we describe is the formation of a complex sequence. It was pointed out

above that with the parameter set being used the intrinsically generated sequences had at least 3 simultaneous units. This indicates that we should be able to store and reproduce a complex sequences of degree 2. Figure 4.8 shows the formation of such a sequence, in which the subsequence 11-12 repeats in each loop. The sequence encoded is 10-**11-12**-13-...-34-**11-12**-35-36-...-54-55-10, and it is so long in order to make it stable as it loops continuously under the effect of plasticity. Consider a shorter sequence, say 10-**11-12**-13-...-20-**11-12**-21-22-23-24-25-10. In this sequence unit 12 happens before unit 17, but it also happens after it in the same loop. When unit 12 is activated, the plasticity rule can increment the connections to roughly 6 of the subsequent units (this number can be larger or shorter due to heterogeneity), so the connection from 12 to 17 would be increased and the one from 17 to 12 would be reduced. After unit 17 is active, however, unit 12 becomes active again, so the connection from 17 to 12 is increased, and the connection from 12 to 17 is reduced. The dynamics of the network ensure that only one of these transitions (from 12 to 17 or from 17 to 12) becomes dominant; thus, when there is not enough time between reappearances of a unit the sequence becomes unstable. When plasticity is turned off shorter complex sequences can be generated, as long as they are already encoded in the connection matrix. The repeated activation of unit 14 in figure 4.8A is due to heterogeneity.

4.2.4 Properties of the Model

There are several elements of the model whose purpose is not readily apparent, but nevertheless allow it to function as we have described so far. We will mention the most important ones and justify their inclusion. In particular, we will discuss the firing rate adaptation, the thresholds for potentiation, the delay in excitatory connections, and several terms in the plasticity equation. We will then finalize our presentation of results by mentioning the role played by the level of inhibition.

The function of firing rate adaptation in this model is twofold: it improves the fidelity with which the input is stored and reproduced, and it is a safeguard against unstable patterns where units remain permanently active. In order to improve the fidelity of input storage, adaptation promotes that the units which have been active for the longest times are the ones

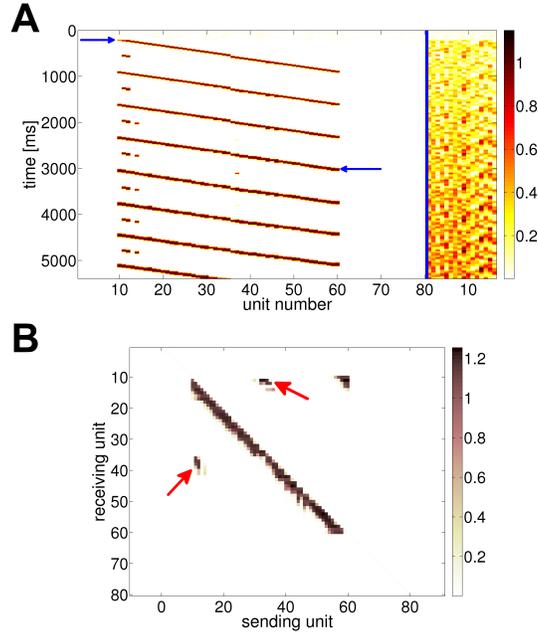


Figure 4.8: **Formation of a complex sequence.** A) Five repetitions of an input sequence with 3 simultaneously active units produce the intrinsic activation of the sequence, in a network with no previously stored patterns. Blue arrows indicate the start and end of the input. B) The resulting connection matrix. Notice the extra connections of units 11 and 12 (red arrows).

that shut off when a new unit is activated. In this way the sequence is reproduced in the right order, and the projections that form towards newly activated units come from their immediate predecessors (not from units that have been active for a long time). This function of the adaptation is particularly useful because of the heterogeneity in the parameters. When the parameters are homogeneous the system can store the order of properly timed input stimuli almost perfectly (and without error if noise is removed). As the level of heterogeneity increases, however, there will be variations in the excitability and plasticity level of individual units, which can scramble the sequential order in the input sequence (an example of this is the activity of the *20 through 30* sequence in figure 4.6C,D). The presence of adaptation reduces this effect.

The threshold θ_{inc} for synaptic potentiation is a parameter in the plasticity equation (equation (4.4) in the Methods section) which stops potentiation when the activity causing it is not strong or temporally precise enough. This threshold is required to keep the spontaneous activity from randomly increasing connections, which can lead to instability. If the threshold

were to be eliminated, the excitatory connection weights would randomly increase to the point of causing synchronous oscillations.

Another parameter we introduced is the delay in the excitatory connections. As mentioned above, we included this delay in order to avoid the synchronization of the units in a sequence, which usually leads to termination of the activity. As the delay is eliminated, continuous activation of a sequence tends to synchronize its components to the point where transitions happen so fast that some or all sequence elements stop activating. This synchronization is dependent on the plasticity, and is in agreement with other studies involving temporally asymmetric Hebbian rules [97, 111, 182].

Perhaps the most unusual part of our model is the plasticity rule, as expressed in equation (4.4) of the Methods section. The right-hand side of equation (4.4) has three terms, implementing potentiation, depression, and slow weight reduction, respectively. The first term contains four factors; the first three bound the increase in the connections, whereas the fourth ensures that increases occur when unit k activates before unit j . The factor $(\alpha + c_{jk})$ has the effect that when the connection c_{jk} has a small value it increases slowly, since α is very small. This implies that if c_{jk} is to increase from a small baseline value into a large value, several repetitions of the sequence k - j must occur. Since a connection doesn't become strong after one single activation of the units k and j with the right timing, sequences arising from spontaneous activity are less likely than sequences arising from repeated inputs. The factor $(s - c_{jk})$ ensures that no individual connection strengths become larger than the value s . This is an insufficient protection against instability, since too many connection strengths attaining their maximum value will lead to activity that cannot be stopped. This is indeed what happens in simulations where the only limit to potentiation is the factor $(s - c_{jk})$. We need to limit the growth of individual connection strengths based on a global constraint which takes into account the value other connection strengths. One common option is to limit the sum of the strengths in the connections towards the same unit, as would be done by the factor $(S - \sum_{r=1}^N c_{jr})$. This factor can be an effective guard against instability, but it has the undesirable property of not permitting the formation of sequences which intersect. In simulations using this factor, forming a sequence results in all its units developing the maximum amount of connections, all from their preceding units in the sequence. If we wanted

to form a sequence containing a unit already in another sequence, connections towards that unit could not be formed, since it already would have all its allowed connections. A robust way of overcoming this challenge is to limit the growth of connection strengths towards a unit j based on the current incoming excitation, so that new connections can always be made towards a unit as long as it is not being strongly activated. This idea is instantiated by using the factor $[S - \sum_{r=1}^N c_{jr} x_r(t - \Delta_{jr})]_+$, which both limits the unstable growth of connections, and permits connections towards units which are already part of one or more sequences. The notation $[...]_+$ indicates that the factor is zero when the quantity inside the brackets is negative; in this way we ensure that the first term of the plasticity rule only implements increases in connectivity. Notice that this factor involves the excitation arising from other excitatory units, but not from external inputs, which means that when an input induces a novel pattern of activity the factor will not limit the growth in the connections. The factor $[H_{inc} - H_{dec}]_+$ and the rest of the plasticity rule are described in the Methods section.

We will now address the term Ψ in the plasticity equation. As mentioned in the Methods section, Ψ may represent one of two expressions, corresponding to either reduction of intermediate-value weights, or heterosynaptic competition. Heterosynaptic competition can maintain stability when dozens of sequences are stored, whereas the reduction of intermediate-value weights can prevent the formation of spurious sequences arising from random inputs. One simple example to illustrate the usefulness of heterosynaptic competition is a simulation where 100 sequences (each containing between 8 and 19 randomly chosen elements) are stored by presenting each one four times in succession. When the simulation was run with $\varepsilon = 0$ (effectively removing the Ψ term), the network lost the stability of its baseline state, so that there was a repeating sequence which could not be stopped by synchronizing the excitatory units. On the other hand, when the simulation was run with $\varepsilon = 0.07$ and heterosynaptic competition, the network still retained a stable baseline state at the end. Heterosynaptic competition also has a functional effect on sequence formation, since what it does is to reduce the strength of any connection which is not taking part in stimulating an active unit. This means that when a unit u is part of two sequences A and B , and the sequence B is being repeated continuously, the projections from the units in A

towards u will be weakened. Therefore, under heterosynaptic competition, sequences that are not often activated will tend to lose the units they share with other more active sequences. Heterosynaptic competition thus tends to separate sequences so that they don't intersect. When the example of figure 4.6 is executed with heterosynaptic competition, units 20,30,40,50, and 60 tend to be removed from one of the sequences, so that we end up with 6 sequences that don't intersect. This is not observed when we use the other version of Ψ , whose effect is to reduce the connection strengths with a value between $\alpha = 0.01$ and $\mu = 0.6$. This reduction of intermediate-value weights will not affect connections whose value is larger than μ , but it may reduce connections whose value is starting to increase. This rule implies, therefore, that in order to reach a large value, a connection must increase rapidly. In order for a connection from a unit a to a unit b to reach a large value, we need several repetitions of the sequence a - b , which are unlikely to occur in close succession if the activity comes from random fluctuations. The reduction of intermediate-value weights is thus well suited to avoid the creation of spurious sequences when random fluctuations are present. If we elicit 100 random input patterns (each made by choosing 6 random units 15 times), then the reduction of intermediate-value weights prevents the formation of repeating sequences, which can be observed when the simulation is run with $\varepsilon = 0$. Reduction of intermediate-value weights can also remove the weak connections which are formed between sequences when the activity is switched directly from one to another (e.g. figure 4.6B). On the other hand, reduction of intermediate-value weights may not prevent instability when a large number of sequences are stored by repeated presentation, and the time constant of plasticity is relatively short. The reduction of intermediate-value weights does not affect sequences which have been already formed, so the stability when storing a large number of sequences depends on the balance between the potentiation and the depression terms in the plasticity rule. In the example where 100 sequences are stored, using reduction of intermediate-value weights instead of heterosynaptic competition may result in the network losing the stability of the baseline state. The goal of storing a large number of sequences for long periods, and the goal of maintaining a stable network may be at odds due to the limited storage capacity in the system. Heterosynaptic competition can maintain stability at the price of modifying existing sequences, whereas reduction of intermediate-value weights does

not modify established sequences, but storing more and more sequences will eventually make the system unstable, unless there is a bias making depression stronger than potentiation in the plasticity equation.

One last property of the model that we will address is the effect of variable inhibition. The effect of reducing the inhibition is to permit more simultaneously active excitatory units (given equal excitatory connection strengths). This does not by itself impede the storage and retrieval of sequences and may allow the reproduction of complex sequences of higher degree. Having more units active simultaneously does have several effects on the dynamics of the network. For example, two or more sequences may activate simultaneously, creating connections between themselves. Another possible effect is that inputs which extinguished ongoing activity under normal inhibition, will not do so under reduced inhibition; this can lead to new connections between input induced activity and sequences intrinsically generated. We discuss these and other effects in a separate paper (Verduzco-Flores et al., submitted), where we explore their possible pathological implications.

4.3 DISCUSSION

We have illustrated the feasibility of temporally asymmetric Hebbian plasticity at the level of asynchronous neuronal populations, and then presented a biologically plausible model for learning and retrieval of sequences at the population level. Our model uses temporally asymmetric Hebbian plasticity in order to learn input sequences, and can reproduce sequences when the plasticity is active because increments in connection strength are modulated by the incoming excitation. Different levels of input overlap can be used to activate, associate, or store sequences. An important property of the model is that it can reproduce complex sequences, and handles intersections of simple sequences without ambiguity.

4.3.1 Forming Connections in Asynchronous Populations

The simulations we presented for the spiking model are not meant to illustrate the formation of excitation pathways in random networks with asynchronous inputs. Our model assumes pre-existing populations with dense internal connectivity, such as minicolumns, which seem likely to occur when considering the local connectivity properties of cortex [179]; the possible ontogenetic and epigenetic origins of such populations is beyond the scope of this paper. The existence of neuronal populations greatly simplifies the autonomous formation of unidirectional excitation pathways, and we present two mechanisms which may achieve this. Both mechanisms take advantage of two traits which facilitate the formation of excitation pathways: background inputs which maintain the membrane voltages near threshold, and a connectivity within populations that allows them to operate asynchronously, and yet increase their synchrony in the presence of strong asynchronous inputs. This increase in synchrony is beneficial when using STDP to create excitation pathways, since it permits the emission of spike volleys in a manner similar to synfire chains, increasing the changes in the mean synaptic weights across populations. Depending on how strongly the synaptic weights are modified by the induced activity they may drift back towards a fixed point, or changes might consolidate through the spontaneous activity, with a speed that is dependent on the firing rate of all the involved populations. This is not an exhaustive description of the dynamics in populations that communicate with STDP synapses, but it suffices to establish the plausibility of the temporally asymmetric Hebbian plasticity used in the firing rate model.

4.3.2 Which Sequences Can Be Stored and Reproduced?

There are limitations on the types of sequences that can be reproduced by the firing rate model.

An obvious first limitation is in sequences where components remain active for different periods of time. In our model each unit remains active roughly the same time, and they cannot activate two or more times consecutively.

Another limitation is that complex sequences whose degree is larger than the number of simultaneously active units may be generated erroneously. Also, sequences which have closed

loops whose length is smaller than the number of simultaneously active units can become unstable.

The most prevalent limitation, however, appears when units repeat their activity too soon and their incoming connections are decreased because the plasticity rule can't decide which unit came before, and which after. This led to abnormal results described elsewhere (Verduzco-Flores et al., submitted), and is a natural consequence of bringing together temporally asymmetric Hebbian plasticity, persistent activity, and sequences with closed loops.

Having more than one sequence stored in the connection matrix can also lead to problems, as not all sequences are compatible with each other. The clearest case is when two sequences share one subsequence. If that subsequence has a length equal to, or larger than the number of simultaneously active units, then the sequences can not be reproduced unambiguously. Also, not all sequences can be stored together by the plasticity rule. As seen in figure 4.6, storing the temporal inverse of a previously stored sequence tends to erase it.

Although the model appears beset by limitations, these can be overcome in a simple manner by having two different layers of excitatory neurons uncoupled by inhibition. Our model represents an area comparable to a cortical column; excitatory layers with uncoupled inhibitory pools correspond to interconnected columns. With a second layer of excitatory neurons, several units activating consecutively in the first layer can project to a single unit in the second layer, which would maintain the unit in the second layer active for longer periods of time, overcoming the limitation on equal activation times, and permitting the creation of arbitrary rhythms. Indeed, two layers of neurons are sufficient to reproduce any type of sequence, as long as we have a long enough sequence in the first layer, and the right connections towards the second layer. The two layers require uncoupled inhibition, and might represent cortical areas which are not contiguous. Once the first layer has continuous sequential activity, the autonomous formation of connections from the first to the second layer can be achieved when the second layer is activated by an external input and a temporally asymmetric Hebbian rule is applied. Taking this idea further, the second layer could present oscillations or transitions at frequencies similar to the frequency of transitions in the first layer, and the plasticity between them could be modulated by their coherence [62]. The formation of this type of connections in order to reproduce any arbitrary sequence, and

to create a grammar (the possible combinations of sequences in both layers) is a possible direction for future research.

4.3.3 Inhibition, Overlap, and Binding

In our firing rate model, the level of inhibition controls the maximum number of simultaneously active units. Units in different sequences may form connections if they activate one after the other, but this is impossible if the level of inhibition permits only one sequence active at a time. Reducing the inhibition permits simultaneous activation of different sequences, and provides an opportunity for their binding. A modulation in the level of local inhibition could happen through interneuron targeting cells [61, 135].

Given a constant level of inhibition (and a maximum number of simultaneously active units), there may be a degree of overlap in the input that will permit the coexistence of input induced activity and currently ongoing activity. Input with no overlap creates no sequences, and input with large overlap extinguishes ongoing activity, but for sufficiently low inhibition there will be a critical value of overlap which allows to create a sequence that associates with the sequences present at the time. This constitutes a possible secondary mechanism to control the binding between populations.

Large variations in the level of inhibition can also compromise the function of the network and produce a range of pathological dynamics. These alterations are described in a separate paper (Verduzco-Flores et al., submitted).

4.3.4 Working Memory and Oscillations

Previous work has speculated on the relation between sequential activity and working memory [127, 184]. The general idea is that the persistent activity associated with working memory consists of the repeated activation of particular sequences. A mechanism to reactivate sequences during extended delay periods is thus required. The model in [184] uses background noise to activate sequences whose connections have been temporarily potentiated. On the other hand, [127] uses a combination of afterdepolarization and subthreshold oscillations. In the present model, having a closed-loop sequence permits reactivation, creating a

working memory mechanism similar to the one in [98], but with spatial structure imposed by the connection matrix. Such a mechanism entails oscillations associated with two periods: a short one for the time of transition between contiguous elements in the sequence, and a larger one for the repetition of the whole sequence. In our implementation of the model, the faster oscillations occur in the gamma range (close to 50 Hz), whereas the slower oscillations have a period equal to the period of the faster oscillations times the number of units in the looping sequence. This is reminiscent of the model in [127] in which the theta frequency serves as a reference frame for the order of memory items activated successively with the gamma frequency. Substantial evidence for such a mechanism has been found in the hippocampus, and its presence in various areas of cortex appears plausible [105, 124]. The model in [127] is an example of temporal segmentation by phase coding. Temporal segmentation refers to the ability to hold several items in short-term memory by activating them repeatedly at different points in time; in contrast, spatial segmentation assumes that individual memory items are represented by separate groups of neurons [105]. Assuming that different units encode different memory items our model is capable of temporal segmentation when a closed-loop sequence is being repeated. If that sequence contains between 5 and 12 units then we will have transitions in the gamma frequency range, and repetitions in the theta range.

Alternatively, if the memory item is represented by the full sequence then we can have a type of “combinatorial” segmentation, in which items could be combined as long as their corresponding sequences could be activated unambiguously in unison. This can happen in our model when the inhibition is reduced to allow more than one sequence simultaneously active, although plasticity would have to be reduced to avoid the sequences “fusing” together.

The main requirement to obtain persistent activity from closed-loops is that the sequence (or at least its first element) repeats during the learning phase. If each unit represented a distinct memory item in working memory, this would not be simply explained. In the case where memory items are encoded by sequences this becomes more plausible, since each representation could have formed after years of experience, wherein closed loops had the opportunity to appear (slower, long-lasting changes in synaptic strength could be involved). In either case, the activity of neuronal groups would experience fluctuations in firing rate during the delay period of working memory tasks. These fluctuations would correspond to

the activation of the group where the neuron is contained, and could be observed at the level of local field potentials. Oscillatory activity during the delay period of working memory tasks has been reported both in the gamma [156], delta-theta [149], and beta [185] ranges. The study in [39] did not find oscillatory activity in prefrontal cortex, but this was done at the level of single neurons, and was imprecise for frequencies below 5 Hz.

4.4 METHODS

4.4.1 The Spiking Model

All neurons used integrate-and-fire dynamics with conductance-based synapses. Postsynaptic potentials were modeled by alpha functions. Populations 1, 2, and 3 consisted of 100 neurons each, with homogeneous parameters. The inhibitory population consisted of 80 neurons. Each neuron had connections to all the neurons in its corresponding population, and to half of the neurons in the other two populations; the connections between populations were randomly chosen. A more detailed description of the model is provided in tables 5-10, as suggested in [152]. All simulations were implemented in NEST [77], and source code is available in the appendix.

4.4.2 The Firing Rate Model

The dynamics of the network can be described by 5 equations. Equation (4.1) determines the activity of the excitatory units.

$$\tau_e \dot{x}_j(t) = -x_j(t) + f \left(a_{ee} x_j(t) + C \sum_{k=1}^N c_{jk} x_k(t - \Delta_{jk}) - a_{ei} \sum_{k=1}^M v_{jk} y_k - dz_j + e\eta_j - \theta_e + I_j \right), \quad (4.1)$$

where $f(x) = 1/(1 + \exp(-x))$. This equation involves four variables. $x_j(t)$ is the activity of the j -th excitatory unit at time t . c_{jk} is the strength of the connection from excitatory unit k to excitatory unit j . y_k is the activity of inhibitory unit k . z_j is the adaptation of excitatory unit j . These variables correspond to 4 of the 5 equations describing the model.

Table 5: Spiking Model Summary.

Populations	Four: three excitatory, one inhibitory; external input
Topology	—
Connectivity	All-to-all within populations, random divergent across populations
Neuron Model	Conductance-based leaky integrate-and-fire
Channel Models	—
Synapse Model	Alpha functions
Plasticity	STDP on excitatory connections between different populations
Input	Independent fixed-rate Poisson spike trains to all neurons
Measurements	Spike times for all neurons, voltage trace for one neuron

Table 6: Populations in the spiking model.

Name	Elements	Size
$E_{1,2,3}$	IAF neurons	100
I	IAF neurons	80
P	Independent Poisson generators	380
O_1	Independent Poisson generators	500
O_2	Independent Poisson generators	500

Table 7: Connectivity of the spiking model.

Name	Source	Target	Pattern
EE_{int}	E_x	E_x	All-to-all for $x=1,2,3$. Weight w_{epsc} , delay wd . No autapses or multapses
EE_{ext}	E_x	$E_y (y \neq x)$	50 targets randomly selected for each source neuron, no multapses. Weight ee_{epsc} , delay eed
EI	E_x	I	40 targets randomly selected from I population for each source neuron, no multapses. Weight ei_{epsc} , delay eid
IE	I	E_x	50 targets randomly selected for each source neuron, no multapses. Weight ie_{epsc} , delay ied
P_E	P	E_x, I	16000 excitatory Poisson inputs for each target neuron, each one with static rate r_{ex}
P_I	P	E_x, I	4000 inhibitory Poisson inputs for each target neuron, each one with static rate r_{in}
$O1$	O_1	E_1	Each target neuron receives input from all generators in O_1 , each one with static rate r_{o1}
$O2$	O_2	E_2	Each target neuron receives input from all generators in O_2 , each one with static rate r_{o2}

Table 8: Neuron and synapse models in the spiking network.

Name	IAF neuron
Type	Conductance-based leaky integrate and fire, alpha-function active conductances
Subthreshold dynamics	$C_m \dot{V} = g_L(E_L - V) + g_E(E_{ex} - V) + g_I(E_{in} - V) \quad \text{if } t > t^*,$ $V_{reset} \quad \text{otherwise.}$ <p> $g_E = \kappa_e \sum_{i,j} w_i \left(e^{-\frac{t-t_i^j}{\tau_{1e}}} - e^{-\frac{t-t_i^j}{\tau_{2e}}} \right)$ where t_i^j is the j-th spike of the i-th neuron satisfying $t \leq t_i^j$; $\kappa_e, \tau_{1e}, \tau_{2e}$ are such that an event of weight 1.0 results in a peak conductance of w_i nS at $t = \tau_{synE} + t_i^j$; w_i is the weight of the connection from neuron i. </p> <p> $g_I = \kappa_i \sum_{i,j} w_i \left(e^{-\frac{t-t_i^j}{\tau_{1i}}} - e^{-\frac{t-t_i^j}{\tau_{2i}}} \right)$ where t_i^j is the j-th spike of the i-th neuron satisfying $t \leq t_i^j$; $\kappa_i, \tau_{1i}, \tau_{2i}$ are such that an event of weight 1.0 results in a peak conductance of w_i nS at $t = \tau_{synI} + t_i^j$; w_i is the weight of the connection from neuron i. </p>
Spiking	<p>If $V(t-) < \theta \wedge V(t+) \geq \theta$</p> <ol style="list-style-type: none"> 1. set $t^* = t$ 2. emit spike with time stamp t^*
Plasticity	$\Delta w = \begin{cases} \lambda f_-(w) \times K(\Delta t) & \text{if } \Delta t \leq 0, \\ \lambda f_+(w) \times K(\Delta t) & \text{if } \Delta t > 0. \end{cases}$ <p> $f_+(w) = (w_{max} - w)^\mu$ and $f_-(w) = \alpha w^\mu$. </p> <p> $K(\Delta t) = \exp(- \Delta t /\tau)$. </p>

Table 9: Measurements taken for the spiking network.

Time of each spike, and identity of the neuron emitting it. Voltage trace for first indexed neuron.

Table 10: Parameter values for the spiking model.

Name	Value	Name	Value	Name	Value
C_m	250 pF	E_L	-70 mV	E_{ex}	0 mV
E_{in}	-85 mV	V_{reset}	-60 mV	θ	-55 mV
τ_{synE}	0.2 ms	τ_{synI}	1.2 ms	τ	20 ms
μ	0.1	α	1	λ	0.1
w_{max}	2 nS	$w_{e_{epsc}}$	1.5 nS	wd	1 ms
ee_{epsc}	0.2/1 nS	eed	1 ms	ei_{epsc}	1 nS
eid	2 ms	ie_{epsc}	-2 nS	ied	2 ms
r_{ex}	4 Hz	r_{in}	5 Hz	$r_{o1/o2}$	15 Hz

The term η_j represents an Ornstein-Uhlenbeck process: $\dot{\eta}_j = -\eta_j/\tau_\eta + \xi$, where ξ is a white noise process, which is implemented at each step of the simulation by multiplying the square root of the time step by a random number drawn from a normal distribution with zero mean and unit variance.

The constants in equation (4.1) are the following. τ_e is the time constant for excitatory units. a_{ee} is the recurrent connectivity for excitatory units. C is the general excitability from excitatory inputs. N is the number of excitatory units. Δ_{jk} is the delay in the connection from unit k to unit j . a_{ei} is the general excitability from inhibitory inputs. M number of inhibitory units. v_{jk} is the connection strength from inhibitory unit k to excitatory unit j . e is the noise amplitude. θ_e is the excitatory threshold. I_j is the external input to excitatory unit j .

The entries v_{jk} in the matrix of inhibitory to excitatory connections were generated randomly so that $v_{jk} = 1$ with probability p_{ei} , and $v_{jk} = 0$ with probability $1 - p_{ei}$. Once the entries of the matrix v were generated, its rows were normalized so that they summed to 1.

Equation (4.2) describes the synaptic adaptation of the j -th unit:

$$\dot{z}_j = \frac{-z_j}{\tau_z} + b(1-z_j)f \left(a_{ee}x_j(t) + C \sum_{k=1}^N c_{jk}x_k(t - \Delta_{jk}) - a_{ei} \sum_{k=1}^M v_{jk}y_k - dz_j + e\eta_j - \theta_e + I_j \right). \quad (4.2)$$

This equation only introduces two new constants. τ_z controls the rate of decay for the adaptation, whereas b controls its rate of growth. The function f and its argument are identical to those in equation (4.1), and are just the firing rate of that population.

Equation (4.3) describes the activity of inhibitory units:

$$\dot{y}_j = \frac{-y_j}{\tau_i} + a_{ie} \left[\sum_{k=1}^N Q_{jk} x_k - \theta_i \right]_+ + iI_j. \quad (4.3)$$

The parameter a_{ie} can adjust how strong the excitatory to inhibitory connections are. The notation $[x]_+$ is equivalent to $\max(x, 0)$, returning zero when x is negative, and x otherwise. The matrix Q of E-I connections was generated similarly to the matrix v in equation 1, but using the probability p_{ie} instead of p_{ei} . θ_i is a threshold to activate inhibitory units, and iI_j is the external stimulus to inhibitory unit j .

Equation (4.4) implements plasticity in the E-E connections, and there are two almost interchangeable versions of it. The first one incorporates reduction of weights with intermediate strength, whereas the second version uses heterosynaptic competition instead.

$$\tau_c \dot{c}_{jk} = (\alpha + c_{jk})(s - c_{jk}) \left[S - \sum_{r=1}^N c_{jr} x_r(t - \Delta_{jr}) \right]_+ \left[H_{inc} - H_{dec} \right]_+ - S c_{jk} \left[H_{dec} - H_{inc} \right]_+ - \Psi, \quad (4.4)$$

where $H_{inc} = H(x_j w_k - \theta_{inc})$ and $H_{dec} = H(x_k w_j - \theta_{dec})$, with $H(x) = 1/(1 + \exp(-qx))$ a sharp sigmoidal. The version with reduction of intermediate-value weights uses $\Psi = \varepsilon[(c_{jk} - \alpha)(\mu - c_{jk})]_+$, while the version with heterosynaptic competition uses $\Psi = \varepsilon c_{jk} x_j(t) (1 - x_k(t - \Delta_{jk}))$.

There are three terms on the right hand side of equation (4.4). The first term implements timing-dependent increments in the connection strength, while the second term creates timing-dependent decrements. The term Ψ has slower dynamics which do not depend on the timing of the inputs. Before describing the terms in equation (4.4) we introduce the last equation:

$$\tau_w \dot{w}_j = -w_j + f_w(x_j - \theta_w), \quad (4.5)$$

where $f_w(x) = 1/(1 + \exp(-\gamma x))$. w_j increases towards a fixed point when x_j is above the threshold θ_w , and otherwise decays exponentially with time constant τ_w , which makes w_j akin to a fast memory trace of the activity x_j .

Returning to equation (4.4), the first term in the equation contains four factors. The parameter α is a very small value, and the factor $(\alpha + c_{jk})$ reduces the rate at which weak connections increase their strength, which benefits stability. The parameter s is the maximum value of individual connection strengths, and the factor $(s - c_{jk})$ ensures that no connection strength increases beyond this value. The parameter S is a limit on how much incoming activity the unit j can receive before the strengths of incoming connections stop increasing. The factor $[S - \sum_{r=1}^N c_{jr} x_r(t - \Delta_{jr})]_+$ stops any increments when the excitation received by unit j reaches the value S . The value H_{inc} is near 1 when unit k is active shortly before unit j becomes active, or when both units are active; otherwise H_{inc} is near zero. Conversely, H_{dec} is near 1 when unit j is active shortly before unit k activates, or when both units are active; otherwise H_{dec} is near zero. This implies that the factor $[H_{inc} - H_{dec}]_+$ is roughly 1 when unit k is active shortly before unit j activates, and 0 otherwise. To understand how the terms H_{inc} and H_{dec} operate, it is useful to consider the values of $x_j w_k$ and $x_k w_j$. When unit k has been active for a minimum period of time, the value of w_k increases near a fixed point. If at this moment unit j becomes active, the term $x_j w_k$ will have a relatively large value, but the term $x_k w_j$ will be very small, because w_j has not increased yet. This will cause H_{inc} to be near 1, while H_{dec} is near zero, and at this point $[H_{inc} - H_{dec}]_+$ is close to 1. If both units remain active, or if unit k becomes inactive while unit j is still active, then $[H_{inc} - H_{dec}]_+$ will go back to a low value. This explains how this factor can implement timing-dependent increases in connectivity. The second term in equation (4.4) works in a similar manner to the first term. The factor S scales decreases in connection strength so that they are of similar magnitude to increases. The factor c_{jk} ensures that we create no negative connection strengths. The factor $[H_{dec} - H_{inc}]_+$ implements timing-dependent decreases in connectivity. The term Ψ slowly decreases the strength of certain connections, so that long simulations are unlikely to result in connection matrices with many large values, especially towards units which are activated continuously. When $\Psi = \varepsilon[(c_{jk} - \alpha)(\mu - c_{jk})]_+$, the connections whose values are between α and μ are reduced. When $\Psi = \varepsilon c_{jk} x_j(t) (1 - x_k(t - \Delta_{jk}))$, connections

towards unit j which take no part in activating it when x_j has an elevated value are reduced; this implements a form of heterosynaptic competition.

4.4.3 Firing Rate Model Implementation

The model was implemented in Matlab using the Euler-Maruyama approximation. Each unit in the model had its own parameter values. Equations (4.1)(4.2) and (4.5) each represent N equations, and each one of those N equations has a distinct value for its parameters which came from a base value plus a random number sampled from a normal distribution whose standard deviation was 4% of the base value. Table 11 shows the base value for all parameters. In equation (4.3) each parameter represents M distinct values. For equation (4.4) each parameter represents N values, corresponding to the unit which received the connection, except for the delay which has N^2 distinct values. All the parameters in equations (4.3) and (4.4) were also generated with heterogeneous values. The time step used for simulations corresponds to 0.5 milliseconds. Source code is available in the appendix.

4.4.4 Parameter Search

A heuristic approximation was used to find the parameters in the firing rate model. In short, for excitatory units the self-excitation and the inhibition were chosen so that a population could not reach a higher attractor by itself, but it could when receiving the input of other populations. Adaptation was strong enough to stop the activity, but slow enough to let it rise. E-I and I-E connections were adjusted to permit 2 or 3 populations active simultaneously. The time constant in the w_j variables was set to allow a plasticity window around 40 msec. The threshold for synaptic potentiation was set so that the spontaneous activity generated by the noise could not change the connectivity. The time constant of the plasticity was set so that 4 or 5 repetitions of a properly timed input sequence would be enough to create a strong change in the connections.

4.4.5 Input Patterns

Under normal conditions, simple and complex sequences were created using 5 repetitions of a stimulus with amplitude $I = 12$, overlap of 3 simultaneous units, and a 12 millisecond difference between the onset of consecutive units. Looping sequences would usually be stopped by stimulating all excitatory units simultaneously using pulses of amplitude $I = 10$ and a frequency of 50 Hz for 200 milliseconds. The demonstration where an input with no overlap caused no change in the connections used an amplitude $I = 15$ and a 12 millisecond difference between consecutive onsets.

Table 11: Standard parameter values for the simulations with 3 to 4 simultaneously active units. The units for all time constants are tens of milliseconds.

Parameter values for the firing rate model					
Parameter	Value	Parameter	Value	Parameter	Value
a_{ee}	8	α	0.01	τ_z	7
a_{ei}	15	γ	5	τ_η	0.5
a_{ie}	15	θ_e	3.8	τ_w	5
b	0.05	θ_i	0.2	μ	0.6
C	3	θ_w	0.5	Δ	2
d	11	θ_{dec}	0.3	N	80
e	1	θ_{inc}	0.3	M	16
h	0.01	ε	0.13	S	4
q	30	τ_e	1	p_{ie}	0.5
r	0.3	τ_c	7	p_{ei}	0.5
s	1.2	τ_i	0.5		

5.0 DISCUSSION

Chapters 2,3, and 4 presented 3 models describing the dynamics of working memory networks. In chapter 2 it was shown that dynamic synapses and interconnected populations are sufficient to explain the various activity patterns exhibited by real neurons during the delay period of a working memory task. Chapter 3 takes once more the ideas of dynamic synapses and interconnected populations, this time to show that a network with those characteristics, in addition to functioning as a substrate for working memory, can also respond differently to specific temporal patterns in the input. In particular, a network with reduced inhibition can exhibit synchronized oscillations (which may involve two groups with antiphase behavior), and it can present simultaneous, uncontrolled activity from several populations. The fact that these behaviors can arise in response to inputs with a particular frequency suggests that this could be one mechanism for the onset of reflex epilepsy. Chapter 4 goes beyond the model in chapter 3 by adding structure in the excitatory connections, and dynamics which ensure that the activity transitions between populations. The interaction between excitatory and inhibitory populations is crucial for this effect, bringing oscillations at the gamma frequency, with low firing rates in the excitatory populations.

The work in chapter 2 follows a common trend in models of working memory, which consists of finding a neural network that replicates some of the experimentally observed characteristics of real neurons. The first characteristic to be emphasized was the persistent activity. This turned out to fit well with the paradigm of attractor networks, which became a candidate explanation for the behavior observed in cells. It was then noticed that attractor network models had firing rates which did not agree with real neurons. Moreover, variability in attractor networks decreased when units were persistently active, but this was not the case in real neurons. The next generation of working memory models used inhibitory in-

terneurons, as well as balance between excitation and inhibition in order to adjust the firing rates and the variance of the interspike intervals [15]. In addition, mechanisms other than recurrent connections were proposed in order to generate the required behavior [44, 60, 133]. But only recently have researchers started to replicate the patterns observed in individual cortical neurons during a single trial of a delayed response task. Replicating these patterns is important, because they may be a link between modeling of detailed activation properties, and models in which the activated neurons perform a function which is relevant for behavior.

Working memory networks, and the prefrontal cortex itself, are but one part of a brain that seamlessly integrates the activity in all its components. Computational models which replicate some statistical aspects of experimental observations are a dead end if their activity is not somehow related to behavior, and that requires the integration of the model in the larger brain network. Currently, this is done mainly by pointing out the information produced by the model, and assuming that the rest of the brain is capable of using it to produce adequate behavior. In the case of chapter 2, the changes in firing rates during the delay period are assumed to encode the identity of a presented stimulus, and probably the time elapsed since its presentation (as indicated by the ramping cells). A problem with this approach is that it is unknown how much information is required, and what type of encoding is appropriate, which will determine what aspect of the activity is worth reproducing. Furthermore, different models may produce similar statistics, and it is up to the experimentalists to inquire which are the right physiological mechanisms. The process of guessing what information is represented, how it is encoded, and verifying the physiological mechanisms is a slow one. In order to improve it, neuroscientists can add further constraints on a model by making it part of a larger conceptual model with interacting specialized networks. The researchers who have worked for many years with working memory cells have created theories of how they interact with other brain areas [63, 68], and this promises to open a new level of sophistication in the next generation of models. There are thoughtful hypotheses about the main functions of dorsolateral prefrontal cortex in primates, which contains the largest proportion of working memory cells [63, 138, 154]. Likewise, conceptual models for other structures such as the hippocampus and the basal ganglia have developed over time. [81, 82, 90, 91, 125, 132, 172]. The study of working memory networks might benefit

from the type of computational studies which seek to include complementary systems [12]. The model in chapter 2 takes a modest step in this direction by showing that interactions between the prefrontal and parietal cortices may be responsible for the observed activity patterns during delayed-response tasks. We do not explicitly attach a functional meaning to this interaction, but this has been done previously [160]. Briefly, prefrontal neurons might help to integrate perceptions from parietal cortex across time; these perceptions are mostly related to spatial planning and representation. The fact that the interaction between the two areas is consistent with the firing rate patterns observed experimentally, encourages further studies about how they influence one another. Perhaps in the future the simultaneous recording of a large number of cells could provide physiological evidence on how the stimulation in one area modulates the response in the other.

The model in chapter 3 does not attempt to further probe into the functional interactions of working memory networks, but on their dysfunctional dynamics. Mnemonic activity is characterized by persistent activity, and the loss of selectivity in these networks is tantamount to ictal activity. The particular dynamics of multiple interconnected populations with dynamic synapses provide temporal sensitivities, with particular input frequencies being able to resonate with intrinsic limit cycles, giving rise to oscillations. Reflex epilepsy is often triggered by intense periodic stimuli [41, 209], such as strobing lights. Our model could be considered to represent this phenomenon, which would make decreased inhibition a culprit of reflex epilepsy. Additionally, our model suggests the type of stimulation required to stop pathological activity, and stresses the importance of activating the inhibitory populations for this effect. One way of achieving this is the activation of all the network, which, depending on the level of inhibition, may return the dynamics to baseline. This approach, however, could backfire in the case when the inhibition has been decreased beyond a certain point. Optogenetic stimulation methods which exclusively target interneurons [136] may one day be incorporated into the deep brain stimulation for the treatment of intractable epilepsy [106].

Chapter 4 takes a more ambitious approach than chapter 3, seeking to understand spatial, as well as temporal patterns. A simplistic conception of what the brain does is the mapping of spatiotemporal patterns in sensory cortex into spatiotemporal patterns in motor cortex. There is in this sense a great promise when studying how spatiotemporal patterns are

encoded in the cortex. Nevertheless, one lesson from chapter 4 is that this encoding and the subsequent reproduction are not straightforward. Shortly after the experimental discovery of STDP it was expected that the generation of sequences would be a natural consequence of this form of plasticity. The ensuing computational models which formed sequences through STDP ran into stability issues, and isolation of the excitation pathways [117,144]. Ever since, the models which achieve sequential behavior are characterized by the assumptions required in order to overcome these difficulties. In our case, our assumptions were the existence of neural populations with static internal connections, and inter-population connections experiencing temporally asymmetric Hebian plasticity. Even then, the our model required further refinement. Populations in a sequence tended to synchronize, so axonal and synaptic delays were included. The network would form unstable connections if the incoming input towards a population was not limited, but it would not encode intersecting sequences if there was a hard limit on the increase of incoming connections. Thus we limited the strengthening of connections by the current activity. Another complication was that over time too many connections would be stored, leading to instability. In response to this we included a form of heterosynaptic competition, which had a tendency to force units into having a single target for their projections. Creating a network which reproduces the order of afferent excitation has many technical difficulties, and it will be interesting to see which new mechanisms and assumptions arise in the future in order to overcome them.

APPENDIX

SOURCE CODE

A.1 SOURCE CODE FOR CHAPTER 1

A.1.1 average9-e.ode

```
# similar to average9.ode, but the exponent of Ca/Ca0 is
# now a parameter. The default parameter values are those
# from avnulls-e4-2state.ode.set.
```

```
# The parameters
par g=5, bthr=1, theta=4.7, taum=0.03
par wmin=0.2, wmax=1.5, gamma=5, tw=4
par ts=0.05, beta=0.5, smin=0.25
par tCa=0.5, eps=1.37, Ca0=50, Camin=15, e=4
par i0=2, ton=10, wid=0.5
```

```
C = 1/(pi*taum)
```

```
# The equations
s' = (smin-s)/ts + beta*v
w' = (wmin-w)/tw + gamma*(wmax-w)*(Ca/Ca0)^e
Ca' = (Camin-Ca)/tCa + eps*v
```

```
b = g*s*w + i(t) - bthr
v = C*(b/(1-exp(-b*theta)))^(1/2)
aux f=v
```

```
i(t)=i0*heav(t-ton)*heav(ton+wid-t)
```

```
# The initial conditions
init w=0.1
init s=0.1
init Ca=15
```

```
# The simulation parameters
@ total=200,dt=0.003
@ maxstor=2000000
```

```
done
```

A.1.2 average10-e.ode

```
# This file has two firing-rate networks interacting with each other.
# The parameters come from avnulls-e4-2state.ode.set
```

```

# PARAMETERS
par bthr=1, theta=4.7, taum=0.03
par wmin=0.2, wmax=1.5, gamma=5, tw=4
par ts=0.05, beta=0.5, smin=0.25
par tCa=0.5, eps=1.37, Ca0=50, Camin=15, e=4
par i0=2, ton=10, wid=0.5
par gAA=5, gBB=5.1, gAB=-0.1, gBA=0

C = 1/(pi*taum)

# EQUATIONS
sA' = (smin-sA)/ts + beta*vA
wA' = (wmin-wA)/tw + gamma*(wmax-wA)*(CaA/Ca0)^e
CaA' = (Camin-CaA)/tCa + eps*vA

sB' = (smin-sB)/ts + beta*vB
wB' = (wmin-wB)/tw + gamma*(wmax-wB)*(CaB/Ca0)^e
CaB' = (Camin-CaB)/tCa + eps*vB

bA = gAA*sA*wA + i(t) - bthr + gBA*sB*wB
vA = C*(bA/(1-exp(-bA*theta)))^(1/2)
bB = gBB*sB*wB + i(t) - bthr + gAB*sA*wA
vB = C*(bB/(1-exp(-bB*theta)))^(1/2)
aux fB=vB
aux fA=vA

i(t)=i0*heav(t-ton)*heav(ton+wid-t)

# INITIAL CONDITIONS
init wA=0.2
init sA=0.2
init CaA=15

init wB=0.2
init sB=0.2
init CaB=15

# SIMULATIONS PARAMETERS
@ total=120,dt=0.003
@ maxstor=2000000,trans=40

done

```

A.1.3 avnulls.ode

```

f(x)=sqrt(x/(1-exp(-b*x)))/(pi*tm)
par b=3
s'=(smin-s)/ts+as*nu
w'=(wmin-w)/tw+gam*(wmax-w)*(ca/car)^2
#ca'=(camin-ca)/tca+ac*nu
nu=f(kbar*s*w-thr+i(t))
ca=camin+ac*nu*tca
par kbar=7.6
par as=.95,ac=1.6
par gam=6.6,wmax=.8,car=110,tw=10,ts=.06,smin=.1,tca=.5,camin=7
par wmin=.07,tm=.03
par i0=0,ton=10,wid=1
par thr=1.4,amp=.3
i(t)=i0*heav(t-ton)*heav(ton+wid-t)
aux fr=nu
done

```

A.1.4 runner1.m

```

% runner1.m performs simulations with nc1.m and displays the results.
% It is a modified version of runner3.m in the newspike directory.

```

```

clear all;
close all;

global init AA AB AC AD BA BB BC BD CA CB CC CD DA DB DC DD      % the matrices
global wmaxA wmaxB wmaxC wmaxD
global mAa stdAA mAB stdAB mAa stdAC mAa stdAD                    % the means and SD
's

```

```

global mBA stdBA mBB stdBB mBC stdBC mBD stdBD
global mCA stdCA mCB stdCB mCC stdCC mCD stdCD
global mDA stdDA mDB stdDB mDC stdDC mDD stdDD
global mwmA stdwmA mwmB stdwmB mwmC stdwmC mwmD stdwmD
global wmin gamma tw ts as smin tCa Car aCa Camin amp th taum e % the cell's parameters
global numA numB numC numD
                                % number of cells in each region

% INITIALIZING VARIABLES

mwmA=1; stdwmA=0.01; mwmB=1; stdwmB=0.01;           % avnulls-e6-2state1/average10-ex6-2state1 parameters
mwmC=1; stdwmC=0.01; mwmD=1; stdwmD=0.01;
wmin=0.6; gamma=6; tw=2; ts=0.05; as=0.5; smin=0.3; e=6;
tCa=0.5; Car=65; aCa=6; Camin=10; amp=0.2; th=1.2; taum=.03;

mwmA=1; stdwmA=0.02; mwmB=1; stdwmB=0.02;           % avnulls-e6-2state2/average10-ex6-2state2 parameters
mwmC=1; stdwmC=0.02; mwmD=1; stdwmD=0.02;
wmin=0.6; gamma=8; tw=2; ts=0.05; as=0.5; smin=0.3; e=6;
tCa=0.5; Car=82; aCa=8; Camin=8; amp=0.2; th=1.2; taum=.03;

mwmA=1; stdwmA=0.02; mwmB=1; stdwmB=0.02;           % average10-ex6-2state2-slow parameters
mwmC=1; stdwmC=0.02; mwmD=1; stdwmD=0.02;
wmin=0.6; gamma=8; tw=4; ts=0.1; as=0.5; smin=0.3; e=6;
tCa=1; Car=93; aCa=8; Camin=8; amp=0.2; th=1.2; taum=.06;

mwmA=1; stdwmA=0.02; mwmB=1; stdwmB=0.02;           % average10-ex6-2state2-fast parameters
mwmC=1; stdwmC=0.02; mwmD=1; stdwmD=0.02;
wmin=0.6; gamma=8; tw=1; ts=0.025; as=0.5; smin=0.3; e=6;
tCa=.25; Car=74; aCa=8; Camin=8; amp=0.2; th=1.2; taum=.015;

mwmA=1; stdwmA=0.01; mwmB=1; stdwmB=0.01;           % avnulls-e6-2state3/average10-ex6-2state3 parameters
mwmC=1; stdwmC=0.01; mwmD=1; stdwmD=0.01;
wmin=0.5; gamma=8; tw=2; ts=0.05; as=0.48; smin=0.35; e=6;
tCa=0.5; Car=81; aCa=8; Camin=8; amp=0.2; th=1.2; taum=.03;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state1
mAA=0.05; stdAA=0.01; mAB=0.05; stdAB=0.01; mAC=-0.01; stdAC
=0.005; mAD=-0.01; stdAD=0.005;
mBA=0.05; stdBA=0.01; mBB=0.05; stdBB=0.01; mBC=-0.01; stdBC
=0.005; mBD=-0.01; stdBD=0.005;
mCA=-0.0015; stdCA=0.01; mCB=-0.0015; stdCB=0.01; mCC=0.055; stdCC
=0.005; mCD=0.055; stdCD=0.005;
mDA=-0.0015; stdDA=0.01; mDB=-0.0015; stdDB=0.01; mDC=0.055; stdDC
=0.005; mDD=0.055; stdDD=0.005;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- FUDGED
mAA=0.05; stdAA=0.02; mAB=0.05; stdAB=0.02; mAC=-0.016; stdAC
=0.004; mAD=-0.016; stdAD=0.005;
mBA=0.05; stdBA=0.02; mBB=0.05; stdBB=0.02; mBC=-0.016; stdBC
=0.004; mBD=-0.016; stdBD=0.005;
mCA=-0.0012; stdCA=0.002; mCB=-0.0012; stdCB=0.002; mCC=0.059; stdCC
=0.005; mCD=0.059; stdCD=0.005;
mDA=-0.0012; stdDA=0.002; mDB=-0.0012; stdDB=0.002; mDC=0.059; stdDC
=0.005; mDD=0.059; stdDD=0.005;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- for 3-7 network with i0=2.4
mAA=0.05; stdAA=0.01; mAB=0.05; stdAB=0.01; mAC=-0.0317; stdAC
=0.005; mAD=-0.0317; stdAD=0.005;
mBA=0.05; stdBA=0.01; mBB=0.05; stdBB=0.01; mBC=-0.0317; stdBC
=0.005; mBD=-0.0317; stdBD=0.005;
mCA=-0.0014; stdCA=0.001; mCB=-0.0014; stdCB=0.001; mCC=0.0802; stdCC
=0.005; mCD=0.0802; stdCD=0.005;
mDA=-0.0014; stdDA=0.001; mDB=-0.0014; stdDB=0.001; mDC=0.0802; stdDC
=0.005; mDD=0.0802; stdDD=0.005;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- for 1-9 network with i0=2.4
mAA=0.054; stdAA=0.005; mAB=0.054; stdAB=0.005; mAC=-0.017; stdAC
=0.005; mAD=-0.017; stdAD=0.005;
mBA=0.054; stdBA=0.005; mBB=0.054; stdBB=0.005; mBC=-0.017; stdBC
=0.005; mBD=-0.017; stdBD=0.005;
mCA=-0.0053; stdCA=0.001; mCB=-0.0053; stdCB=0.001; mCC=0.0654; stdCC
=0.005; mCD=0.0654; stdCD=0.005;
mDA=-0.0053; stdDA=0.001; mDB=-0.0053; stdDB=0.001; mDC=0.0654; stdDC
=0.005; mDD=0.0654; stdDD=0.005;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state3
mAA=0.055; stdAA=0.05; mAB=0.055; stdAB=0.05; mAC=-0.019; stdAC
=0.01; mAD=-0.019; stdAD=0.01;
mBA=0.055; stdBA=0.05; mBB=0.055; stdBB=0.05; mBC=-0.019; stdBC
=0.01; mBD=-0.019; stdBD=0.01;

```

```

% mCA = -0.0103; stdCA = 0.05; mCB = -0.0103; stdCB = 0.05; mCC = 0.061; stdCC
= 0.03; mCD = 0.061; stdCD = 0.03; mDC = 0.061; stdDC
% mDA = -0.0103; stdDA = 0.05; mDB = -0.0103; stdDB = 0.05; mDC = 0.061; stdDC
= 0.03; mDD = 0.061; stdDD = 0.03;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state3
% mAA = 0.061; stdAA = 0.05; mAB = 0.061; stdAB = 0.05; mAC = -0.02; stdAC
= 0.03; mAD = -0.02; stdAD = 0.03; mBC = -0.02; stdBC
% mBA = 0.061; stdBA = 0.05; mBB = 0.061; stdBB = 0.05; mBC = -0.02; stdBC
= 0.03; mBD = -0.02; stdBD = 0.03; mCC = 0.061; stdCC
% mCA = -0.018; stdCA = 0.03; mCB = -0.018; stdCB = 0.03; mDC = 0.061; stdDC
= 0.04; mCD = 0.061; stdCD = 0.04; mDB = -0.018; stdDB = 0.03;
% mDA = -0.018; stdDA = 0.03; mDB = -0.018; stdDB = 0.03; mDC = 0.061; stdDC
= 0.04; mDD = 0.061; stdDD = 0.04;

% FOUR NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- for 3-7,1-9 networks with i0=1.6
% mAA = 0.101; stdAA = 0.004; mAB = -0.0462; stdAB = 0.003; mAC = -0.0003; stdAC
= 0.003; mAD = 0.0003; stdAD = 0.001; mBC = -0.0001; stdBC
% mBA = -0.0012; stdBA = 0.003; mBB = 0.1405; stdBB = 0.003; mBC = -0.0001; stdBC
= 0.003; mBD = -0.0; stdBD = 0.001; mCC = 0.1025; stdCC
% mCA = -0.0014; stdCA = 0.001; mCB = 0.0011; stdCB = 0.0001; mDC = -0.0034; stdDC
= 0.004; mCD = -0.036; stdCD = 0.003; mDB = 0.0002; stdDB = 0.003;
% mDA = 0.0005; stdDA = 0.001; mDB = 0.0002; stdDB = 0.003; mDC = -0.0034; stdDC
= 0.004; mDD = 0.132; stdDD = 0.003;

% FOUR NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2-slow ---
% for 3-7,1-9 networks with i0=6
% mAA = 0.101; stdAA = 0.004; mAB = -0.06; stdAB = 0.003; mAC = -0.0005; stdAC
= 0.003; mAD = 0.009; stdAD = 0.001; mBC = -0.001; stdBC
% mBA = -0.00; stdBA = 0.003; mBB = 0.155; stdBB = 0.003; mBC = -0.001; stdBC
= 0.003; mBD = -0.001; stdBD = 0.001; mCC = 0.105; stdCC
% mCA = -0.005; stdCA = 0.001; mCB = 0.002; stdCB = 0.0001; mDC = -0.002; stdDC
= 0.004; mCD = -0.054; stdCD = 0.003; mDB = -0.0002; stdDB = 0.003;
% mDA = 0.005; stdDA = 0.001; mDB = -0.0002; stdDB = 0.003; mDC = -0.002; stdDC
= 0.004; mDD = 0.147; stdDD = 0.001;

% FOUR NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2-fast ---
% for 3-7,1-9 networks with i0=1.5
% mAA = 0.1; stdAA = 0.003; mAB = -0.042; stdAB = 0.003; mAC = -0.000; stdAC = 0.003;
= 0.003; mAD = 0.00; stdAD = 0.001; mBC = -0.001; stdBC
% mBA = -0.001; stdBA = 0.003; mBB = 0.13; stdBB = 0.003; mBC = -0.001; stdBC
= 0.003; mBD = -0.001; stdBD = 0.001; mCC = 0.1065; stdCC
% mCA = -0.005; stdCA = 0.001; mCB = 0.002; stdCB = 0.0001; mDC = -0.008; stdDC
= 0.006; mCD = -0.035; stdCD = 0.003; mDB = -0.0002; stdDB = 0.003;
% mDA = 0.005; stdDA = 0.001; mDB = -0.0002; stdDB = 0.003; mDC = -0.008; stdDC
= 0.004; mDD = 0.13; stdDD = 0.001;

% TWO NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- FUDGED
% mAA = 0.005; stdAA = 0.001; mAB = 0.005; stdAB = 0.001; mAC = -0.00145; stdAC = 0.0004; mAD
= -0.00145; stdAD = 0.005; mBB = 0.005; stdBB = 0.001; mBC = -0.00145; stdBC = 0.0004; mBD
% mBA = 0.005; stdBA = 0.001; mBB = 0.005; stdBB = 0.001; mBC = -0.00145; stdBC = 0.0004; mBD
= -0.00145; stdBD = 0.005; mCB = -0.000134; stdCB = 0.0001; mCC = 0.0057; stdCC = 0.0005; mCD
% mCA = -0.000134; stdCA = 0.0001; mCB = -0.000134; stdCB = 0.0001; mCC = 0.0057; stdCC = 0.0005; mCD
= 0.0057; stdCD = 0.005; mDB = -0.000134; stdDB = 0.0001; mDC = 0.0057; stdDC = 0.0005; mDD
% mDA = -0.000134; stdDA = 0.0001; mDB = -0.000134; stdDB = 0.0001; mDC = 0.0057; stdDC = 0.0005; mDD
= 0.0057; stdDD = 0.005;

% FOUR NETWORK PARAMETERS (200 NEURONS) FROM average10-ex6-2state2 --- for 3-7,1-9 networks with i0=2.4
mAA = 0.0101; stdAA = 0.0004; mAB = -0.00462; stdAB = 0.0003; mAC = -0.00003; stdAC = 0.0003; mAD
= 0.00003; stdAD = 0.001; mBC = -0.00001; stdBC = 0.0003; mBD = -0.0;
mBA = -0.00012; stdBA = 0.0003; mBB = 0.01405; stdBB = 0.0003; mBC = -0.00001; stdBC = 0.0003; mBD = -0.0;
stdBD = 0.001;
mCA = -0.00014; stdCA = 0.0001; mCB = 0.00011; stdCB = 0.00001; mCC = 0.01025; stdCC = 0.0004; mCD
= -0.0036; stdCD = 0.003;
mDA = 0.00005; stdDA = 0.0001; mDB = 0.00002; stdDB = 0.0003; mDC = -0.00034; stdDC = 0.0004; mDD
= 0.0132; stdDD = 0.003;

numA=500; numB=500; numC=500; numD=500;

i0=2.4; ton=20; wid=0.3; % current injection parameters
time=32; % running time of the simulation in seconds.
Default is 32.

% RUNNING THE SIMULATION
[X]=nc1(time,i0,ton,wid);

[cells iter vars]=size(X);
disp('nc1 done');

% PLOTTING THE CELL ARRAY

```

```

close all;
%figure;
%r=[0.2*ones(1,32) linspace(0,1,96)];
%g=[0.3*ones(1,16) linspace(0,1,48) linspace(1,0,48) 0.3*ones(1,16)];
%b=[linspace(1,0,96) 0.2*ones(1,32)];
%map=[r;g;b]';
%colormap(map);

%image(X(:,:,3)*128/max(max(X(:,:,3))));
%colorbar('vert');

% PLOTTING SOME FACILITATION GRAPHS
figure;
hold on;
plot(linspace(0,time,iter),X(1,:,3));
plot(linspace(0,time,iter),X(numA+1,:,3),'r');
plot(linspace(0,time,iter),X(numA+numB+1,:,3),'c');
plot(linspace(0,time,iter),X(numA+numB+numC+1,:,3),'k');

% CREATING A MATRIX OF FREQUENCY HISTOGRAMS
disp('Creating frequency histograms');
pause(1);
fHist=zeros(cells,floor(time));
bin = floor(iter/time); % how many iterations in each second
reftime = ceil(bin/200); % refractory time of about 5 ms
for cell=1:cells
    it=0;
    count=reftime+1;
    for sec=1:floor(time)
        for i=1:bin
            it=it+1;
            if X(cell,it,1) >= 2.9 && count > reftime
                fHist(cell,sec)=fHist(cell,sec)+1;
                count = 0;
            else
                count = count+1;
            end
        end
    end
end
end

%figure;
%map=[linspace(0,1,128) linspace(0,1,128) linspace(0,1,128)]';
%colormap(map);
%image(fHist*128/max(max(fHist)));
%colorbar('vert');

figure;
mfactor = floor((numA+numB+numC+numD)/16);
for i=1:16
    subplot(4,4,i);
    bar(fHist(mfactor*i,:));
end

% CREATING A HISTOGRAM OF NEURON TYPES FOR EACH AREA
type=histAnaP(fHist);
figure;
for i=1:11
    typeHist(i)=sum(type(1:numA)==i);
end
subplot(2,2,1);
bar(typeHist);
for i=1:11
    typeHist(i)=sum(type(numA+1:numA+numB)==i);
end
subplot(2,2,2);
bar(typeHist);
for i=1:11
    typeHist(i)=sum(type(numA+numB+1:numA+numB+numC)==i);
end
subplot(2,2,3);
bar(typeHist);
for i=1:11
    typeHist(i)=sum(type(numA+numB+numC+1:numA+numB+numC+numD)==i);
end
subplot(2,2,4);
bar(typeHist);

% CREATING A PLOT OF AVERAGE FREQUENCY
avfHist=sum(fHist)/cells;

```

```

figure;
bar(avfHist);

% CREATING A HISTOGRAM OF NEURON TYPES FOR THE WHOLE NETWORK
figure;
for i=1:11
    typeHist(i)=sum(type(1:cells)==i);
end
bar(typeHist);

```

A.1.5 ncl.m

```

function [X]=ncl(time,i0,ton,wid)
%[X]=ncl(time,i0,ton,wid) performs the simulation of newsp6-netcouple1.ode from time zero to
% time 'time', with the stimulus being of amplitude 'i0', width 'wid', and onset time 'ton'.
% The 3-D array X will contain the values of all the state variables for each one of the
% iterations. The second index of X indicates the iteration, whereas the other two indices
% indicate the variable as:

%           X(:,j,:) =      | x1 s1 w1 Ca1 |
%                           | x2 s2 w2 Ca2 |
%                           | . . . . |
%                           | xn sn wn Can |

% ncl first finds a stable initial state by calling the function ncleulerI and
% simulating for ~30 seconds; then the function ncleuler is used with the obtained initial values
% and the same parameters.

% ncl.m is a modified version of ns9.m in the newspike directory
%-----

dt=0.003;
MAX_ITER=100000;
global init AA AB AC AD BA BB BC BD CA CB CC CD DA DB DC DD      % the matrices
global wmaxA wmaxB wmaxC wmaxD
global mAa stdAA mAB stdAB mAC stdAC mAD stdAD                    % the means and SD
    's
global mBA stdBA mBB stdBB mBC stdBC mBD stdBD
global mCA stdCA mCB stdCB mCC stdCC mCD stdCD
global mDA stdDA mDB stdDB mDC stdDC mDD stdDD
global mwmA stdwmA mwmB stdwmB mwmC stdwmC mwmD stdwmD
global wmin gamma tw ts as smin tCa Car aCa Camin amp th taum    e % the cell's parameters
global numA numB numC numD                                       % number of cells in each region

% GENERATING THE INITIAL VALUES
init= repmat([-pi/2 0.2 0.2 20],numA+numB+numC+numD,1); % make sure these are below the low fixed point
numI=ceil(40/dt);
if numI > MAX_ITER
    disp('Too many iterations for the initial value run. ');
    return
end
I = zeros(1,numI);
ncleulerI(numI,dt,I);
disp('ncleulerI done');

% GENERATING THE INPUT VECTOR
clear I;
numI=ceil(time/dt);
if numI > MAX_ITER
    disp('Too many iterations. ');
    return
end
t=0;
for i=1:numI
    if t < ton
        I(i)=0;
    elseif t < ton+wid
        I(i)=i0;
    else
        I(i)=0;
    end
    t=t+dt;
end

% CALLING THE ncleuler function
X = ncleuler(numI,dt,I);

```

```
disp('ncleuler done');
```

A.1.6 ncleuler.m

```
function [X]=ncleuler(numI,dt,I)
% [X]=ncleuler(numI,dt,I) reads a global variable 'init' which is a matrix of initial values, receives a
% number of iterations 'numI', and a step size 'dt'. These arguments are used to apply the forward Euler
% method to the equations in newsp6-netcouple1 ode, creating the matrix X which contains
% the value of each of the four variables at each of the num steps for each of the neurons. The number
% of neurons in regions A,B,C,D come from parameters, and the total number of neurons must agree with
% the size of init.
% The argument I is a row vector of length numI. It specifies the current injection in each neuron
% for each iteration.

% The format of init is: | x1 s1 w1 Ca1 |
%                       | x2 s2 w2 Ca2 |
%                       | . . . . |
%                       | xn sn wn Can |

% X is a 3-d array. It's first and third indices point to the state variables as in init, at the time
% indicated by the second index.

% ncleuler.m is a modified version of ns9euler.m in the newspike directory.
%-----

global init AA AB AC AD BA BB BC BD CA CB CC CD DA DB DC DD      % the matrices
global wmaxA wmaxB wmaxC wmaxD
global mAA stdAA mAB stdAB mAC stdAC mAD stdAD                    % the means and SD
      ,s
global mBA stdBA mBB stdBB mBC stdBC mBD stdBD
global mCA stdCA mCB stdCB mCC stdCC mCD stdCD
global mDA stdDA mDB stdDB mDC stdDC mDD stdDD
global mwmA stdwmA mwmB stdwmB mwmC stdwmC mwmD stdwmD
global wmin gamma tw ts as smin tCa Car aCa Camin amp th taum e   % the cell's parameters
global numA numB numC numD
                                % number of cells in each region

ampW = amp/sqrt(dt);      % scaling the amplitude to create Wiener noise

numAB=numA+numB;
numABC=numA+numB+numC;
numABCD=numA+numB+numC+numD;

% INITIALIZING X
X=zeros(numABCD,numI+1,4);
X(:,1,:)=init;

% MAIN CYCLE
for i=2:numI+1
    xA=X(1:numA,i-1,1);                xB=X(numA+1:numAB,i-1,1);
    xC=X(numAB+1:numABC,i-1,1);        xD=X(numABC+1:numABCD,i-1,1);
    sA=X(1:numA,i-1,2);                sB=X(numA+1:numAB,i-1,2);
    sC=X(numAB+1:numABC,i-1,2);        sD=X(numABC+1:numABCD,i-1,2);
    wA=X(1:numA,i-1,3);                wB=X(numA+1:numAB,i-1,3);
    wC=X(numAB+1:numABC,i-1,3);        wD=X(numABC+1:numABCD,i-1,3);
    CaA=X(1:numA,i-1,4);                CaB=X(numA+1:numAB,i-1,4);
    CaC=X(numAB+1:numABC,i-1,4);        CaD=X(numABC+1:numABCD,i-1,4);

    pA=sA.*wA; pB=sB.*wB; pC=sC.*wC; pD=sD.*wD;

    KAA=AA*pA; KAB=AB*pA; KAC=AC*pA; KAD=AD*pA;
    KBA=BA*pB; KBB=BB*pB; KBC=BC*pB; KBD=BD*pB;
    KCA=CA*pC; KCB=CB*pC; KCC=CC*pC; KCD=CD*pC;
    KDA=DA*pD; KDB=DB*pD; KDC=DC*pD; KDD=DD*pD;

    nsA=randn(numA,1);
    nsB=randn(numB,1);
    nsC=randn(numC,1);
    nsD=randn(numD,1);

    X(1:numA,i,1)= X(1:numA,i-1,1) + dt*((1-cos(xA)) + (KAA + KBA + KCA + KDA + I(i-1) + ampW*nsA - th)
        .*(1+cos(xA)))/taum;
    X(1:numA,i,2)= X(1:numA,i-1,2) + dt*(smin-sA)/ts;
    X(1:numA,i,3)= X(1:numA,i-1,3) + dt*( (wmin-wA)/tw + gamma*(wmaxA-wA).*( (CaA/Car).^e ) );
    X(1:numA,i,4)= X(1:numA,i-1,4) + dt*(Camin-CaA)/tCa;
```

```

X(numA+1:numAB, i, 1)= X(numA+1:numAB, i-1,1) + dt*((1-cos(xB)) + (KAB+ KBB + KCB + KDB + I(i-1) + ampW*
nsB - th).*(1+cos(xB)))/taum;
X(numA+1:numAB, i, 2)= X(numA+1:numAB, i-1,2) + dt*(smin-sB)/ts;
X(numA+1:numAB, i, 3)= X(numA+1:numAB, i-1,3) + dt*( (wmin-wB)/tw + gamma*(wmaxB-wB).*( (CaB/Car).^e ) );
X(numA+1:numAB, i, 4)= X(numA+1:numAB, i-1,4) + dt*(Camin-CaB)/tCa;

X(numAB+1:numABC, i, 1)= X(numAB+1:numABC, i-1,1) + dt*((1-cos(xC)) + (KAC+ KBC + KCC + KDC + I(i-1) +
ampW*nsC - th).*(1+cos(xC)))/taum;
X(numAB+1:numABC, i, 2)= X(numAB+1:numABC, i-1,2) + dt*(smin-sC)/ts;
X(numAB+1:numABC, i, 3)= X(numAB+1:numABC, i-1,3) + dt*( (wmin-wC)/tw + gamma*(wmaxC-wC).*( (CaC/Car).^e )
);
X(numAB+1:numABC, i, 4)= X(numAB+1:numABC, i-1,4) + dt*(Camin-CaC)/tCa;

X(numABC+1:numABCD, i, 1)= X(numABC+1:numABCD, i-1,1) + dt*((1-cos(xD)) + (KAD+ KBD + KCD + KDD + I(i-1) +
ampW*nsD - th).*(1+cos(xD)))/taum;
X(numABC+1:numABCD, i, 2)= X(numABC+1:numABCD, i-1,2) + dt*(smin-sD)/ts;
X(numABC+1:numABCD, i, 3)= X(numABC+1:numABCD, i-1,3) + dt*( (wmin-wD)/tw + gamma*(wmaxD-wD).*( (CaD/Car)
.^e ) );
X(numABC+1:numABCD, i, 4)= X(numABC+1:numABCD, i-1,4) + dt*(Camin-CaD)/tCa;

for j=1:numABCD
    if X(j, i, 1) >= pi
        X(j, i, 1) = -pi;
        X(j, i, 2) = X(j, i, 2)+as;
        X(j, i, 4) = X(j, i, 4)+aCa;
    end
end
end
end

```

A.1.7 histAnaP.m

```

function [type]=histAnaP(fHist)

% [type]=histAnaP(fHist) receives as its argument a matrix of frequency histograms, as
% the one generated in runner1.m, and returns a vector 'type'. The value type(i)
% specifies the firing class of neuron i using criteria similar to Shafi et al.
% This version of histAna defines the minimum significative change as a function
% of the baseline frequency.

Bs=13; Be=19; % baseline interval start and end
D1s=21; D1e=26; % first delay interval
D2s=26; D2e=31; % second delay interval
s=0.8; % The minimum significative change lowest value
p=0.64; % min. sig. change = p*sqrt(baseline freq)

[cells secs]=size(fHist);
freqs=zeros(cells,3); % this stores the frequencies at each interval
changes=zeros(cells,3); % stores the changes: B-D1, B-D2, D1-D2

for cell=1:cells
    freqs(cell,1)=sum(fHist(cell,(Bs+1):Be))/(Be-Bs);
    freqs(cell,2)=sum(fHist(cell,(D1s+1):D1e))/(D1e-D1s);
    freqs(cell,3)=sum(fHist(cell,(D2s+1):D2e))/(D2e-D2s);
end

msc=max(s,p*sqrt(freqs(:,1)));

changes(:,1)=(freqs(:,1) < (freqs(:,2)-msc)) - (freqs(:,1) > (freqs(:,2)+msc));
changes(:,2)=(freqs(:,1) < (freqs(:,3)-msc)) - (freqs(:,1) > (freqs(:,3)+msc));
changes(:,3)=(freqs(:,2) < (freqs(:,3)-msc)) - (freqs(:,2) > (freqs(:,3)+msc));

for cell=1:cells
    switch changes(cell,1)
    case 1,
        switch changes(cell,2)
        case 1,
            switch changes(cell,3)
            case 1, type(cell)=1;
            case 0, type(cell)=2;
            case -1, type(cell)=3;
            end
        case 0,
            type(cell)=3;
        case -1,
            type(cell)=10;
    end
end

```

```

    end
case 0,
    switch changes (cell ,2)
    case 1,
        type (cell)=1;
    case 0,
        switch changes (cell ,3)
        case 1,
            type (cell)=4;
        case 0,
            type (cell)=5;
        case -1,
            type (cell)=6;
        end
    case -1,
        type (cell)=9;
    end
case -1,
    switch changes (cell ,2)
    case 1,
        type (cell)=11;
    case 0,
        type (cell)=7;
    case -1,
        switch changes (cell ,3)
        case 0,
            type (cell)=8;
        case -1,
            type (cell)=9;
        case 1,
            type (cell)=7;
        end
    end
end
end
end

```

A.2 SOURCE CODE FOR CHAPTER 2

A.2.1 bigrunM5d.m

```

% This file is used to run simulations using M5dvar.m. It produces five
% figures, each one of them with nine different subplots. The nine
% subplots correspond to different values of exc (number of excited
% neurons), and the five figures correspond to different values of per.
% With these figures I can readily characterize the behavior for a given
% value of population inhibition, when using the right input weights.

clear all;
close all;

global Cee Cei % the interpopulation connection matrices
global ce ci % the input weights
global aee aie aei aii k te ti utau vtau eps uth wmax gamma % population parameters
global Y0 % initial conditions
global beta per ton tstim; % stimulus' parameters

% 1) CREATING/LOADING THE PARAMETERS
%setM5dpars; % this creates the parameters when you don't have a set to load
load a2aI20E06h2six-ae;
POPS=length(Y0)/3;
%load Y0-asynchOscil;
%rstconn(5.4,0.6,POPS); % changing the values of I and E (i.e. Cee and Cei)
%rstLOCcei(2,[4]); % dropping local inhibiton in the populations in the vector
%rstaie(9.2,1:POPS); % originally 10.1

periods = [.2 .3]; % The values of per we'll use
exces = [0 1 2 4 8 12]; % The values of exc we'll use
np=length(periods);
ne=length(exces);

% 2) SIMULATION CYCLES
beta = 20; per=0.3; ton=10; tstim=5;
total=35;

```

```

allYs=cell(np,ne); % I use a cell array because Y will vary in size

for i=1:np
    per=periods(i)
    for j=1:ne
        exc=exces(j)

        ce=[repmat(1.0,1,exc), repmat(0.5,1,POPS-exc)];
        ci=[repmat(0.1,1,exc), repmat(0.1,1,POPS-exc)];

        [T,allYs{i,j}] = ode45('M5dvar',[0,total],Y0);
    end
end

disp('Simulations finished. Starting plots');
pause(1);

% 3) PLOTTING THE RESULTS
r=[zeros(1,24) linspace(0,1,16) ones(1,12) linspace(1,0.3,12)];
g=[zeros(1,8) linspace(0,1,16) ones(1,16) linspace(1,0,16) zeros(1,8)];
b=[linspace(0.6,1,8) ones(1,16) linspace(1,0,16) linspace(0.3,0,24)];
map=[r;g;b]';
sq=ceil(sqrt(ne));

for i=1:np
    figure('Name',[ 'per=' num2str(periods(i)) ], 'NumberTitle','off');
    title(['per=' num2str(periods(i))]);
    colormap(map);
    colorbar('vert');
    for j=1:ne
        subplot(2,3,j);
        %subplot(sq,sq,j);
        set(gca,'FontSize',20.0);
    end
end

%
% This code changes the scale of the image so that it shows time
% instead of iterations. Comment it for faster performance.
L=length(allYs{i,j}(:,1));
bucklen=2*L;
bucket=zeros(bucklen,POPS);
linearT=linspace(0,total,bucklen);
indexT=1;
for K=1:bucklen
    if T(indexT) < linearT(K)
        while T(indexT) < linearT(K) && indexT < L
            indexT=indexT+1;
        end
        bucket(K,:) = allYs{i,j}(indexT,1:3:3*POPS);
    end
end
dt=1/bucklen;
image([0,20],[0+dt,total-dt],bucket*128); %/max(max(bucket));

%
%image(allYs{i,j}(:,1:3:3*POPS)*128); %/max(max(allYs{i,j}(:,1:3:3*POPS)));
title(['k=' num2str(exces(j))]);
%title(['per=' num2str(periods(i))]);

end
end

```

A.2.2 M5dvar.m

```

function [F] = M5dvar4(T,Y)

% [F]=M5dvar4(T,Y) receives a time variable T, and a column vector Y. M5dvar4 implements equations
% like those in the M5d.ode file, but the number of populations can vary according to the size of
% the vector Y, whose length is three times the number of populations.
% The connections between populations are determined by the matrices Cei and Cee, which are global
% variables. We assume that the only connections between populations are excitatory to inhibitory,
% and excitatory to excitatory. The value Cei(i,j) is the weight of the connection from excitatory
% population j to inhibitory population i. Similarly for Cee(i,j).
% The input weights come from the global vectors ce and ci.
% All the parameters are vectors, so that each population can have unique values.

global Cee Cei % the interpopulation connection matrices
global ce ci % the input weights
global aee aie aei aii k te ti utau vtau eps uth wmax gamma % population parameters

```

```

if mod(length(Y),3) ~= 0
    error('Incorrect input function Y to M5dvar');
end

n = length(Y)/3;           % number of populations

F = zeros(3*n,1);         % So we get a column vector and not a row vector

P = p(T);
for i=1:3:(3*n)
    I=floor(i/3)+1;       % the number of the population corresponding to this cycle
    F(i) = (-Y(i) + f(aee(I)*(1+k(I)*Y(i+2))*Y(i) - aie(I)*Y(i+1) + Cee(I,:)*(Y(1:3:(3*n)))*(1+(k')*.Y
        (3:3:(3*n)))) - te(I) + ce(I)*P)/utau(I);
    F(i+1) = (-Y(i+1) + f(aei(I)*Y(i) - aii(I)*Y(i+1) + Cei(I,).*Y(1:3:(3*n)) - ti(I) + ci(I)*P))/vtau(I);
    F(i+2) = eps(I)*(-Y(i+2) + f(gamma(I)*(Y(i)-uth(I)))*(wmax(I)-Y(i+2)));
end
% COMMENTS
% f and p are functions in the same directory as this file; p's paramaters are global variables
% which should be initialized before using the function.
% Since I included aei and aee in the F equations, the main diagonal of Cee and Cei should be zero.

```

A.2.3 f.m

```

function [y]=f(x)

% [y]=f(x) implements the standard sigmoidal nonlinearity y=1/(1+exp(-x))
% in a vector fashion.

y = 1./(1+exp(-x));

```

A.2.4 p.m

```

function [y]=p(t)

% [y]=p(t) implements the scalar input function in M5d.ode. It requires for all its parameters
% to be global variables already initialized.

global beta per ton tstim;

%y = exp(-beta*(1-cos(2*pi*t/per)))*(1+sign(t-ton))*(1+sign(ton+tstim-t))/4;
y = exp(-beta*(1-cos(2*pi*t/per)))*heaviside(t-ton)*heaviside(ton+tstim-t);

```

A.2.5 greatrunM5d.m

```

% greatrunM5d.m performs several bigrunM5d.m simulations, and keeps all the
% resulting figures open. The simulations differ in their values of I and E
% (mean inhibition and excitation), as well as values of aie. Since
% bigrunM5d.m doesn't take any arguments, and since it closes every window
% on each run, greatrunM5d.m calls instead the function
% greathelper(I,E,ie), which is a modified version of bigrunM5d.m.

% The number of simulations performed by greatrunM5d will depend on the
% length of the vectors I, E, ie, periods, and exces. The latter two are in
% greathelper.m. Usually, this will result in periods*I*E*ie windows, with
% 'periods' simulations in each one. These simulations almost fully
% characterize a parameter set. The only unconsidered aspects are the
% effect of different background and excitation amplitudes, and the effect of
% resetting the connections' heterogeneity in each run of greathelper.m.

clear all;
close all;

I=[3.4 3.2 3 2.8 2.5];
E=[0.6];
ie=[10.2 10.1 9.6 9.4 9.2];

for i=1:length(I)
    for e=1:length(E)
        for a=1:length(ie)

```

```

        disp(['I=' num2str(I(i)) ' E=' num2str(E(e)) ' aie=' num2str(ie(a))]);
        greathelper(I(i),E(e),ie(a));
    end
end
end

```

A.2.6 greathelper.m

```

function greathelper(I,E,ie)

% greathelper(I,E,ie) is just like bigrunM5d, but it doesn't close any windows
% before it starts, and it now receives the mean inhibition, excitation,
% and aie arguments. This function is used by greatrunM5d.m.

%global Cee Cei % the interpopulation connection matrices
global ce ci % the input weights
%global aee aie aei aii k te ti utau vtau eps uth wmax gamma % population parameters
global Y0 % initial conditions
global beta per ton tstim; % stimulus' parameters

% 1) CREATING/LOADING THE PARAMETERS
%setM5dparams; % this creates the parameters when you don't have a set to load
load a2a135E06h2six-ae;
POPS=length(Y0)/3;
%rstconn(I,E,POPS); % changing the values of I and E (i.e. Cee and Cei)
rstaie(ie,1:POPS); % originally 10.1

periods = [.15 .2 .25 .3 .4]; % The values of per we'll use
exces = [0 1 6 8 12 16];
np=length(periods);
ne=length(exces);

ceb = 0.4 - (E-0.7) + (I-2)/10 + (ie-10.1)/2; % heuristic values of excitation
cef = 1.5 - (2/5)*(3-I) - (3/5)*(E-.7) + (ie-10.1);
%ceb=0.08;
%cef=0.5;

% 2) SIMULATION CYCLES
beta = 20; per=0.3; ton=10; tstim=5;
total=35;
allYs=cell(np,ne); % I use a cell array because Y will vary in size

for i=1:np
    per=periods(i);
    for j=1:ne
        exc=exces(j);

        ce=[repmat(cef,1,exc), repmat(ceb,1,POPS-exc)];
        ci=[repmat(0.1,1,exc), repmat(0.1,1,POPS-exc)];

        [T, allYs{i,j}] = ode45('M5dvar',[0,total],Y0);
    end
end

%disp('Simulations finished. Starting plots');
%pause(1);

% 3) PLOTTING THE RESULTS
%r=[linspace(0,1,64) ones(1,64)];
%g=[linspace(1,0.2,64) linspace(0.2,1,64)];
%b=[ones(1,64) linspace(1,0.3,32) linspace(0.3,0.8,32)];
%map=[r;g;b]';
sq=ceil(sqrt(ne));

for i=1:np
    figure('Name',[ 'per=' num2str(periods(i)) ' I=' num2str(I) ' E=' num2str(E), ' aie=' num2str(ie)], '
        NumberTitle','off');
    %colormap(map);
    for j=1:ne
        subplot(sq,sq,j);
        image(allYs{i,j}(:,1:3:3*POPS)*128); %/max(max(allYs{i,j}(:,1:3:3*POPS)));
        title(['exc=' num2str(exces(j))]);
    end
    colorbar('vert');
end
end

```

A.3 SOURCE CODE FOR CHAPTER 3

A.3.1 fourPops2.sli

```
/*      fourPops2.sli
      Simulates four populations of IAF neurons
      to investigate population-level
      temporally asymmetric Hebbian learning

      The only difference with fourPops1 is that in
      this version the connections are not retrieved
      after each chunk of simulation
*/

%-----
% Setting parameters
%-----
tic
200. /chunk Set % time bin to write synaptic weights to file
15000.0 ms /t_sim Set % simulation time (multiple of chunk)
-65.0 mV /V_res Set % reset voltage for all neurons
100 /N1 Set % Size of population 1
100 /N2 Set % Size of population 2
100 /N3 Set % Size of population 3
80 /N4 Set % Size of population 4 (the inhibitory population)
% if you change the number of neurons, you must change the name
% of the files in nestFourPops1.m

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% input rates %%%%%%%%%%%%%
16000 /n_ex Set % size of the excitatory populations
4000 /n_in Set % size of the inhibitory populations
500 /n_ex1 Set % size of extra Exc. Pop. to population 1
500 /n_ex2 Set % size of extra Exc. Pop. to population 2
4.0 Hz /r_ex Set % mean rate of the excitatory populations
5.0 Hz /r_in Set % rate of the inhibitory populations
15.0 Hz /r_ex1 Set % rate of extra Exc. Pop. to population 1
15.0 Hz /r_ex2 Set % rate of extra Exc. Pop. to population 2

n_ex r_ex mul /ple_rate Set % rate of Exc. input for population 1
n_in r_in mul /pli_rate Set % rate of Inh. input for population 1
n_ex r_ex mul /p2e_rate Set % rate of Exc. input for population 2
n_in r_in mul /p2i_rate Set % rate of Inh. input for population 2
n_ex r_ex mul /p3e_rate Set % rate of Exc. input for population 3
n_in r_in mul /p3i_rate Set % rate of Inh. input for population 3
n_ex r_ex mul /p4e_rate Set % rate of Exc. input for population 4
n_in r_in mul /p4i_rate Set % rate of Inh. input for population 4
n_ex1 r_ex1 mul /p1o_rate Set % rate of extra input to population 1
n_ex2 r_ex2 mul /p2o_rate Set % rate of extra input to population 2
n_ex r_ex 0.8 sub mul /pxe_rate Set % post-stim Exc. rate for all
n_in r_in 0.8 sub mul /pxi_rate Set % post-stim Inh. rate for all

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% input timing %%%%%%%%%%%%%
1.0 /start1 Set % start time for Poisson input at population 1
1.0 /start2 Set % start time for Poisson input at population 2
1.0 /start3 Set % start time for Poisson input at population 3
1.0 /start4 Set % start time for Poisson input at population 4
2000. /start1o Set % start time for 2nd Poisson input at Pop 1
t_sim /start2o Set % start time for 2nd Poisson input at Pop 2
t_sim /end1 Set % end time for Poisson input at population 1
t_sim /end2 Set % end time for Poisson input at population 2
t_sim /end3 Set % end time for Poisson input at population 3
t_sim /end4 Set % end time for Poisson input at population 4
2300. /end1o Set % end time for 2nd Poisson input at Pop 1
2000. /end2o Set % end time for 2nd Poisson input at Pop 2
t_sim /startx Set % start time for post-stim Poisson input
t_sim /endx Set % end time for post-stim Poisson input

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% synaptic weights and delays %%%%%%%%%%%%%
1. pA /ee_epsc Set % weight for Exc. Conn. between populations
1.0 pA /ei_epsc Set % weight for E-I connections
-2.0 pA /ie_epsc Set % weight for I-E connections
1.5 pA /we_epsc Set % weight for connections within Exc. populations
0.0 pA /wi_epsc Set % weight for connections within Inh. populations
```

```

0.5 pA /p1e.epsc Set % peak amp for Exc. synapses at population 1
-0.5 pA /p1i.epsc Set % peak amp for Inh. synapses at population 1
0.5 pA /p2e.epsc Set % peak amp for Exc. synapses at population 2
-0.5 pA /p2i.epsc Set % peak amp for Inh. synapses at population 2
0.5 pA /p3e.epsc Set % peak amp for Exc. synapses at population 3
-0.5 pA /p3i.epsc Set % peak amp for Inh. synapses at population 3
0.5 pA /p4e.epsc Set % peak amp for Exc. synapses at population 4
-0.5 pA /p4i.epsc Set % peak amp for Inh. synapses at population 4
0.5 pA /p1o.epsc Set % peak amp for 2nd. Exc. synapses at population 1
0.5 pA /p2o.epsc Set % peak amp for 2nd. Exc. synapses at population 2
1.0 ms /d Set % default synaptic delay
1.0 ms /wd Set % synaptic delay within populations
1.0 ms /eed Set % synaptic delay in E-E connections
2.0 ms /eid Set % synaptic delay in E-I connections
2.0 ms /ied Set % synaptic delay in I-E connections

##### other synaptic parameters #####
0.1 /mu_plu Set % potentiation exponent
0.1 /mu_min Set % depression exponent
0.5 /w_init Set % initial value of connections across neurons
2.0 /Wmax Set % maximum connection weight
0.1 /lambda Set % change per spike
1.0 /alpha Set % ratio of inhibitory to excitatory changes
1.2 /tau_in Set % time constant for inhibitory alpha function
0.2 /tau_ex Set % time constant for excitatory alpha function

%-----
% Create 4 neuronal populations
%-----

/subnet Create /sub1 Set % population 1
/subnet Create /sub2 Set % population 2
/subnet Create /sub3 Set % population 3
/subnet Create /sub4 Set % population 4

/iaf_cond_alpha << /V_reset V_res % calibrating
/tau_syn_in tau_in /tau_syn_ex tau_ex >> SetDefaults % alpha functions

sub1 ChangeSubnet % Pop. 1 is current working subnet
/iaf_cond_alpha N1 Create ; % creating N1 IAF neurons in Pop. 1
/pop1 sub1 GetNodes def % pop1 is an array with all pop1 neurons' IDs

sub2 ChangeSubnet % Pop. 2 is current working subnet
/iaf_cond_alpha N2 Create ; % creating N2 IAF neurons in Pop. 2
/pop2 sub2 GetNodes def % pop2 is an array with all pop2 neurons' IDs

sub3 ChangeSubnet % Pop. 3 is current working subnet
/iaf_cond_alpha N3 Create ; % creating N3 IAF neurons in Pop. 3
/pop3 sub3 GetNodes def % pop3 is an array with all pop3 neurons' IDs

sub4 ChangeSubnet % Pop. 4 is current working subnet
/iaf_cond_alpha N4 Create ; % creating N4 IAF neurons in Pop. 4
/pop4 sub4 GetNodes def % pop4 is an array with all pop4 neurons' IDs

%-----
% Connect the populations among themselves
%-----

##### all-to-all connections within populations #####
sub1 pop1 [we.epsc] [wd] DivergentConnect % connecting pop1
sub2 pop2 [we.epsc] [wd] DivergentConnect % connecting pop2
sub3 pop3 [we.epsc] [wd] DivergentConnect % connecting pop3
sub4 pop4 [wi.epsc] [wd] DivergentConnect % connecting pop4

##### random connections across populations #####
%% setting the potentiation exponents for stdp connections
/stdp_synapse << /mu_minus mu_min /mu_plus mu_plu
/Wmax Wmax /weight w_init /lambda lambda /alpha alpha >> SetDefaults
%% forbidding autapses and multapses
/RandomDivergentConnect
<< /allow_autapses false /allow_multapses false >> SetOptions

/*
For some reason RandomDivergentConnect won't accept arrays with
a single element when specifying weights and delays. Thus I'm
creating the next arrays.
*/
N1 2 div /CX1 Set % number of connections towards Pop. 1
N2 2 div /CX2 Set % number of connections towards Pop. 2

```

```

N3 2 div /CX3 Set          % number of connections towards Pop. 3
N4 2 div /CX4 Set          % number of connections towards Pop. 4
CX1 {ee_epsc} repeat CX1 arraystore /dubyasX1 Set
CX2 {ee_epsc} repeat CX2 arraystore /dubyasX2 Set
CX3 {ee_epsc} repeat CX3 arraystore /dubyasX3 Set
CX4 {ei_epsc} repeat CX4 arraystore /dubyasEI Set
CX1 {ie_epsc} repeat CX1 arraystore /dubyasIE1 Set
CX2 {ie_epsc} repeat CX2 arraystore /dubyasIE2 Set
CX3 {ie_epsc} repeat CX3 arraystore /dubyasIE3 Set
CX1 {eed} repeat CX1 arraystore /DlaysX1 Set
CX2 {eed} repeat CX2 arraystore /DlaysX2 Set
CX3 {eed} repeat CX3 arraystore /DlaysX3 Set
CX4 {eid} repeat CX4 arraystore /DlaysEI Set
CX1 {ied} repeat CX1 arraystore /DlaysIE1 Set
CX2 {ied} repeat CX2 arraystore /DlaysIE2 Set
CX3 {ied} repeat CX3 arraystore /DlaysIE3 Set

%% Here 's the actual connecting across populations
sub1 CX2 pop2 dubyasX2 DlaysX2 /stdp_synapse RandomDivergentConnect % 1->2
sub1 CX3 pop3 dubyasX3 DlaysX3 /stdp_synapse RandomDivergentConnect % 1->3
sub1 CX4 pop4 dubyasEI DlaysEI /static_synapse RandomDivergentConnect % 1->4
sub2 CX1 pop1 dubyasX1 DlaysX1 /stdp_synapse RandomDivergentConnect % 2->1
sub2 CX3 pop3 dubyasX3 DlaysX3 /stdp_synapse RandomDivergentConnect % 2->3
sub2 CX4 pop4 dubyasEI DlaysEI /static_synapse RandomDivergentConnect % 2->4
sub3 CX1 pop1 dubyasX1 DlaysX1 /stdp_synapse RandomDivergentConnect % 3->1
sub3 CX2 pop2 dubyasX2 DlaysX2 /stdp_synapse RandomDivergentConnect % 3->2
sub3 CX4 pop4 dubyasEI DlaysEI /static_synapse RandomDivergentConnect % 3->4
sub4 CX1 pop1 dubyasIE1 DlaysIE1 /static_synapse RandomDivergentConnect % 4->1
sub4 CX2 pop2 dubyasIE2 DlaysIE2 /static_synapse RandomDivergentConnect % 4->2
sub4 CX3 pop3 dubyasIE3 DlaysIE3 /static_synapse RandomDivergentConnect % 4->3

%-----
% Create inputs and detectors
%-----
/poisson-generator Create /poisson1e Set
/poisson-generator Create /poisson1i Set
/poisson-generator Create /poisson2e Set
/poisson-generator Create /poisson2i Set
/poisson-generator Create /poisson3e Set
/poisson-generator Create /poisson3i Set
/poisson-generator Create /poisson4e Set
/poisson-generator Create /poisson4i Set
/poisson-generator Create /poisson1o Set
/poisson-generator Create /poisson2o Set
/poisson-generator Create /poissonxe Set
/poisson-generator Create /poissonxi Set

poisson1e << /rate p1e_rate /start start1 /stop end1 >> SetStatus
poisson1i << /rate p1i_rate /start start1 /stop end1 >> SetStatus
poisson2e << /rate p2e_rate /start start2 /stop end2 >> SetStatus
poisson2i << /rate p2i_rate /start start2 /stop end2 >> SetStatus
poisson3e << /rate p3e_rate /start start3 /stop end3 >> SetStatus
poisson3i << /rate p3i_rate /start start3 /stop end3 >> SetStatus
poisson4e << /rate p4e_rate /start start4 /stop end4 >> SetStatus
poisson4i << /rate p4i_rate /start start4 /stop end4 >> SetStatus
poisson1o << /rate p1o_rate /start start1o /stop end1o >> SetStatus
poisson2o << /rate p2o_rate /start start2o /stop end2o >> SetStatus
poissonxe << /rate pxe_rate /start startx /stop endx >> SetStatus
poissonxi << /rate pxi_rate /start startx /stop endx >> SetStatus

/spike_detector Create /spikes1 Set
/spike_detector Create /spikes2 Set
/spike_detector Create /spikes3 Set
/spike_detector Create /spikes4 Set
/voltmeter Create /volt1 Set
/voltmeter Create /volt2 Set
/voltmeter Create /volt3 Set
/voltmeter Create /volt4 Set

%%%%%%%%% sending records to output files %%%%%%%%%%
spikes1 << /to_file true /to_memory false /label (spikes1)
/close_on_reset false >> SetStatus % Cuz the Sim. stops and starts again
spikes2 << /to_file true /to_memory false /label (spikes2)
/close_on_reset false >> SetStatus
spikes3 << /to_file true /to_memory false /label (spikes3)
/close_on_reset false >> SetStatus
spikes4 << /to_file true /to_memory false /label (spikes4)
/close_on_reset false >> SetStatus
volt1 << /to_file true /to_memory false /label (volt1)

```

```

/close_on_reset false >> SetStatus
volt2 << /to_file true /to_memory false /label (volt2)
/close_on_reset false >> SetStatus
volt3 << /to_file true /to_memory false /label (volt3)
/close_on_reset false >> SetStatus
volt4 << /to_file true /to_memory false /label (volt4)
/close_on_reset false >> SetStatus

% next line ensures that existing output files are overwritten
0 << /overwrite_files true >> SetStatus

%-----
% Connect inputs and detectors to populations
%-----
pop1 spikes1 ConvergentConnect % connecting to spike recorders
pop2 spikes2 ConvergentConnect
pop3 spikes3 ConvergentConnect
pop4 spikes4 ConvergentConnect

poisson1e pop1 [p1e.epsc] [d] DivergentConnect
poisson1i pop1 [p1i.epsc] [d] DivergentConnect
poisson2e pop2 [p2e.epsc] [d] DivergentConnect
poisson2i pop2 [p2i.epsc] [d] DivergentConnect
poisson3e pop3 [p3e.epsc] [d] DivergentConnect
poisson3i pop3 [p3i.epsc] [d] DivergentConnect
poisson4e pop4 [p4e.epsc] [d] DivergentConnect
poisson4i pop4 [p4i.epsc] [d] DivergentConnect
poisson1o pop1 [p1o.epsc] [d] DivergentConnect
poisson2o pop2 [p2o.epsc] [d] DivergentConnect
poissonxe pop1 [p1e.epsc] [d] DivergentConnect
poissonxi pop1 [p1i.epsc] [d] DivergentConnect
poissonxe pop2 [p2e.epsc] [d] DivergentConnect
poissonxi pop2 [p2i.epsc] [d] DivergentConnect
poissonxe pop3 [p3e.epsc] [d] DivergentConnect
poissonxi pop3 [p3i.epsc] [d] DivergentConnect
poissonxe pop4 [p4e.epsc] [d] DivergentConnect
poissonxi pop4 [p4i.epsc] [d] DivergentConnect

volt1 pop1 0 get Connect % the voltmeters are connected
volt2 pop2 0 get Connect % to the first neuron in Pops.
volt3 pop3 0 get Connect
volt4 pop4 0 get Connect

%-----
% Functions to obtain synaptic weights
%-----

[] /Wendy12 Set % All the connection weights from pop1 to pop2
[] /Wendy13 Set % All the connection weights from pop1 to pop3
[] /Wendy21 Set % All the connection weights from pop2 to pop1
[] /Wendy23 Set % All the connection weights from pop2 to pop3
[] /Wendy31 Set % All the connection weights from pop3 to pop1
[] /Wendy32 Set % All the connection weights from pop3 to pop2
pop1 GetMax /max1 Set % last element in pop1
pop2 GetMax /max2 Set % last element in pop2

[] /PIPX Set % All STDP connections arising from pop1
pop1 % The next procedure initializes PIPX
{
    /top Set << /source top /synapse_type /stdp_synapse >>
    FindConnections
    PIPX join /PIPX Set
} forall

[] /P2PX Set % All STDP connections arising from pop2
pop2
{
    /top Set << /source top /synapse_type /stdp_synapse >>
    FindConnections
    P2PX join /P2PX Set
} forall

[] /P3PX Set % All STDP connections arising from pop3
pop3
{
    /top Set << /source top /synapse_type /stdp_synapse >>
    FindConnections
    P3PX join /P3PX Set
} forall

```

```

%% Updates the Wendy12 and Wendy13 arrays
/UpdateW1x
{
    [] /Wendy12 Set
    [] /Wendy13 Set

    P1PX
    {
        % next line puts the weight and target on the stack
        GetStatus [[/weight /target]] get arrayload ;

        max2 leq      % if target is smaller or equal to max2
        { Wendy12 exch append /Wendy12 Set }
        { Wendy13 exch append /Wendy13 Set } % else case
        ifelse

    } forall
} bind def

/UpdateW2x
{
    [] /Wendy21 Set
    [] /Wendy23 Set

    P2PX
    {
        % next line puts the weight and target on the stack
        GetStatus [[/weight /target]] get arrayload ;

        max1 leq      % if target is smaller or equal to max1
        { Wendy21 exch append /Wendy21 Set }
        { Wendy23 exch append /Wendy23 Set } % else case
        ifelse

    } forall
} bind def

/UpdateW3x
{
    [] /Wendy31 Set
    [] /Wendy32 Set

    P3PX
    {
        % next line puts the weight and target on the stack
        GetStatus [[/weight /target]] get arrayload ;

        max1 leq      % if target is smaller or equal to max1
        { Wendy31 exch append /Wendy31 Set }
        { Wendy32 exch append /Wendy32 Set } % else case
        ifelse

    } forall
} bind def

toc
(Build time = ) =only =
%-----
% Simulate and record synaptic weights
%-----
tic
MLWARNING setverbosity % so it doesn't display stuff each Sim.
pclocks First 1 arraystore /seed Set % seed comes from clock
0 << /rng-seeds seed >> SetStatus % seeding the RNG

(P3P2.dat) ofstream ; % stream atop stack (unless can't open)
(P3P1.dat) ofstream ; % stream atop stack (unless can't open)
(P2P3.dat) ofstream ; % stream atop stack (unless can't open)
(P2P1.dat) ofstream ; % stream atop stack (unless can't open)
(P1P3.dat) ofstream ; % stream atop stack (unless can't open)
(P1P2.dat) ofstream ; % stream atop stack (unless can't open)
0. /sim-t Set % initializing a time counter

t_sim chunk div cvi % number of times to repeat the simulation
{
    chunk Simulate
    sim_t chunk add /sim_t Set % sim_t = sim_t + chunk
    UpdateW1x % updating weight arrays
    UpdateW2x
    UpdateW3x
    sim_t <- ( ) <- Wendy12 Mean <- endl % writing mean weight
}

```

```

        6 1 roll          % rolling streams in stack
        sim_t <- ( ) <- Wendy13 Mean <- endl % writing mean weight
        6 1 roll          % rolling streams in stack
        sim_t <- ( ) <- Wendy21 Mean <- endl % writing mean weight
        6 1 roll          % rolling streams in stack
        sim_t <- ( ) <- Wendy23 Mean <- endl % writing mean weight
        6 1 roll          % rolling streams in stack
        sim_t <- ( ) <- Wendy31 Mean <- endl % writing mean weight
        6 1 roll          % rolling streams in stack
        sim_t <- ( ) <- Wendy32 Mean <- endl % writing mean weight
        6 1 roll          % rolling streams in stack
    } repeat

    close close close      % closing weights' file streams
    close close close

    toc
    (Simulation Time = ) =only =

    %-----
    % Display the average firing rates during stimulation
    % (this code may fail depending on the relevant intervals)
    %-----
    cout          % putting the standard output in stack
    (firing rate 1:) <- endl      % printing text
    spikes1 GetStatus
    /n.events get % putting the number of spikes in stack
    1000.0 N1 div mul t_sim div % 1000*n/(t_sim*N1)
    %1000.0 N1 div mul endl start1 sub div % 1000*n/((end-start)*N1)
    %1000.0 N1 div mul endx startx sub div % 1000*n/((end-start)*N1)
    <- endl      % send result to output stream

    (firing rate 2:) <- endl      % printing text
    spikes2 GetStatus
    /n.events get % putting the number of spikes in stack
    1000.0 N2 div mul t_sim div % 1000*n/(t_sim*N2)
    %1000.0 N2 div mul end2 start2 sub div % 1000*n/((end-start)*N2)
    %1000.0 N2 div mul endx startx sub div % 1000*n/((end-start)*N2)
    <- endl      % send result to output stream

    (firing rate 3:) <- endl      % printing text
    spikes3 GetStatus
    /n.events get % putting the number of spikes in stack
    1000.0 N3 div mul t_sim div % 1000*n/(t_sim*N3)
    <- endl      % send result to output stream

    (firing rate 4:) <- endl      % printing text
    spikes4 GetStatus
    /n.events get % putting the number of spikes in stack
    1000.0 N4 div mul t_sim div % 1000*n/(t_sim*N4)
    <- endl ;      % send result to output stream and pop it

    %-----
    % Call a Matlab process to visualize results
    %-----

    (matlab < nestFourPops1.m) 0 system ; ;

```

A.3.2 nestFourPops1.m

```

% function nestFourPops1.m

% nestFourPops1 loads the results from simulations of the fourPops1.sli
% program and displays the results. Make sure this is run in the same
% directory where fourPops1.sli writes its recordings.

close all;
clear all;

sp1 = load('spikes1-397-0.gdf');
sp2 = load('spikes2-398-0.gdf');
sp3 = load('spikes3-399-0.gdf');
sp4 = load('spikes4-400-0.gdf');
volt1 = load('volt1-401-0.dat');
volt2 = load('volt2-402-0.dat');
volt3 = load('volt3-403-0.dat');
volt4 = load('volt4-404-0.dat');

```

```

P1P2 = load('P1P2.dat');
P1P3 = load('P1P3.dat');
P2P1 = load('P2P1.dat');
P2P3 = load('P2P3.dat');
P3P1 = load('P3P1.dat');
P3P2 = load('P3P2.dat');

figure;
set(gca,'FontSize',32);
%title('Voltage trace for a Pop.1 neuron','FontSize',34,'FontName','Arial');
plot(volt1);
xlabel('time[ms]','FontSize',34,'FontName','Arial');
ylabel('Voltage[mV]','FontSize',34,'FontName','Arial');
hold on
%plot(volt2,'r');
figure;
plot(volt4,'c');
title('Inhibitory neuron');

[r1 c1] = size(sp1);
[r2 c2] = size(sp2);
[r3 c3] = size(sp3);
[r4 c4] = size(sp4);

if c1 > 1
    figure;
    sp1(:,1) = sp1(:,1)-4;
    ras4NEST(sp1);
    set(gca,'FontSize',32);
    xlabel('time[ms]','FontSize',34,'FontName','Arial');
    ylabel('Neuron','FontSize',34,'FontName','Arial');
    title('Excitatory Population 1','FontSize',34,'FontName','Arial');
else
    disp('No spikes in first raster');
end
if c2 > 1
    figure;
    sp2(:,1) = sp2(:,1)-104;
    ras4NEST(sp2);
    set(gca,'FontSize',32);
    xlabel('time[ms]','FontSize',34,'FontName','Arial');
    ylabel('Neuron','FontSize',34,'FontName','Arial');
    title('Excitatory Population 2','FontSize',34,'FontName','Arial');
else
    disp('No spikes in second raster');
end
if c3 > 1
    figure;
    sp3(:,1) = sp3(:,1)-204;
    ras4NEST(sp3);
    set(gca,'FontSize',32);
    xlabel('time[ms]','FontSize',34,'FontName','Arial');
    ylabel('Neuron','FontSize',34,'FontName','Arial');
    title('Excitatory Population 3','FontSize',34,'FontName','Arial');
else
    disp('No spikes in third raster');
end
if c4 > 1
    figure;
    sp4(:,1) = sp4(:,1)-304;
    ras4NEST(sp4);
    set(gca,'FontSize',32);
    xlabel('time[ms]','FontSize',34,'FontName','Arial');
    ylabel('Neuron','FontSize',34,'FontName','Arial');
    title('Inhibitory Population','FontSize',34,'FontName','Arial');
else
    disp('No spikes in fourth raster');
end

figure;
H=plot(P1P2(:,1),P1P2(:,2),P2P1(:,1),P2P1(:,2),'LineWidth',3);
set(gca,'FontSize',32);
ylim=get(gca,'YLim');
xlabel('time[ms]','FontSize',34,'FontName','Arial');
ylabel('weight','FontSize',34,'FontName','Arial');
legend(H,'1_to_2','2_to_1');
%title('Mean synaptic weight from Pop.1 to Pop.2','FontSize',34,'FontName','Arial');

```

```

figure;
H=plot(P1P3(:,1),P1P3(:,2),P3P1(:,1),P3P1(:,2),'LineWidth',3);
set(gca,'FontSize',32);
xlabel('time[ms]','FontSize',34,'FontName','Arial');
ylabel('weight','FontSize',34,'FontName','Arial');
legend(H,'1to3','3to1');
ylim(ylima);
%title('Mean synaptic weight from Pop.1 to Pop.3','FontSize',34,'FontName','Arial');

```

```

figure;
H=plot(P2P3(:,1),P2P3(:,2),P3P2(:,1),P3P2(:,2),'LineWidth',3);
set(gca,'FontSize',32);
legend(H,'2to3','3to2');
xlabel('time[ms]','FontSize',34,'FontName','Arial');
ylabel('weight','FontSize',34,'FontName','Arial');
ylim(ylima);
%title('Mean synaptic weight from Pop.2 to Pop.3','FontSize',34,'FontName','Arial');

```

```

menu('Windows will die when you press something','OK','ok','Ok','oK');

```

A.3.3 ras4NEST.m

```

function ras4NEST(T)

% ras4NEST(T) receives a matrix T with n rows and 2 columns. For each row,
% the first column is the number of a neuron, and the second column is the
% time at which it spiked. ras4NEST(T) produces the corresponding raster
% plot.
% The input to this plotting function are the output files produced by the
% spike_detector devices in NEST.

X= repmat(T(:,2) ',2,1);

Y=[(T(:,1)-1)';T(:,1)'];

plot(X,Y,'k')

```

A.3.4 TestSeq23gen2.m

```

% TestSeq23gen2.m
% This version of TestSeq23gen has different, more generalized menu
% options, which are better suited to study pathological regimes.

clear all;
close all;
home;

global aee aei aie alpha eps b c d e h s S r q mid del sdel gamma the thi thw thdec thinc tc te ti tv tw
tz tn N M tmax I Tin Tout V freq CIE lEI dt rsqdt inpFun

%*****
%***** INITIALIZE PARAMETERS *****
%*****

sim = @euler23; % use euler23d for heterogeneous delays
% PARAMETER VALUES
%-----
aee=8; aei=15; aie=15; b=0.05; c=3; d=11; e=1; h=0.01; S=4; gamma=5;
r=0.3; s=1.2; q=30; mid=0.6; the=3.8; thi=0.2; thw=0.5; thdec=0.3; thinc=0.3;
sdel=2; alpha=0.01; eps=0.1; tc=7; te=1; ti=0.5; tv=40e9; tz=7; tn=0.5; tw=5;
N=80; M=16; % number of excitatory and inhibitory units respectively
PEI=0.5; PIE=0.5; % probability of EI and IE connections
dt=0.05; % step size for Euler method
rsqdt=1/sqrt(dt); % used to avoid extra computations

% MAKING PARAMETERS HETEROGENEOUS
std=0.04; % 0.04
aee=aee+std*aee*randn(N,1); % for aie and aei the heterogeneity comes from the connection matrix
b=b+std*b*randn(N,1); d=d+std*d*randn(N,1); e=e+std*e*randn(N,1);
gamma=gamma+std*gamma*randn(N,1); h=h+std*h*randn(N,1); s=s+std*s*randn(N,1);
S=S+std*S*randn(N,1); r=r+std*r*randn(N,1); q=q+std*q*randn(N,1); del=sdel+std*sdel*randn(N*N,1);

```



```

end
disp('')
switch cas
    case 1
        % apply a stimulus

        %+++++ SETTING THE STIMULUS +++++
        syn = input('Press 1 for synchronous stimulus, anything else for sequential:');
        if isempty(syn) || syn ~= 1 % sequential stimulus
            inpSeq = input('Enter the input vector (column!):');
            reps = input('Enter the number of repetitions:');
            dur = input('Enter the duration of each input step [Default 3.6]:');
            if isempty(dur)
                dur = 3.6;
            end
            end
            I = input('Enter input amplitude [Default 12]:');
            if isempty(I)
                I = 12;
            end
            end
            ovr = input('Enter the degree of overlap [Default 3]:');
            if isempty(ovr)
                ovr = 3;
            end
            end
            switch ovr
                case 1
                    V= repmat(inpSeq, reps, 1);
                    inpFun = @I8;
                    Tin=20; Tout=Tin + ceil(dur*length(V));
                case 2
                    len = length(inpSeq);
                    V = zeros(len, 2);
                    V(:,1) = inpSeq;
                    V(1:end-1,2) = inpSeq(2:end);
                    V(end,2) = V(1,1);
                    V = repmat(V, reps, 1);
                    inpFun = @I9;
                    Tin=20; Tout=Tin + ceil((dur/2)*length(V));
                otherwise % degree of overlap 3 is the default case
                    len = length(inpSeq);
                    V = zeros(len, 3);
                    V(:,1) = inpSeq;
                    V(1:end-1,2) = inpSeq(2:end);
                    V(1:end-2,3) = inpSeq(3:end);
                    V(end,2) = V(1,1);
                    V(end,3) = V(2,1); V(end-1,3) = V(1,1);
                    V = repmat(V, reps, 1);
                    inpFun = @I9;
                    Tin =20; Tout=Tin + ceil((dur/3)*length(V));
            end
            end
            else % synchronous stimulus
                I = input('Enter input amplitude [Default 12]:');
                if isempty(I)
                    I = 15;
                end
                end
                per = input('Enter the period of the stimulus:');
                Tstim = input('Enter the duration of the stimulus:');
                V = input(['Enter the input vector [default (1: num2str(N) )]:']);
                if isempty(V)
                    V = (1:N)';
                end
                end
                freq=2*pi/per;
                Tin=100; Tout=Tin + Tstim;
                inpFun = @I7; % I7 implements simultaneous pulses
            end
            %+++++ SIMULATING +++++
            %=====
            tmax = 2*Tout; tmax=540;
            [Y] = sim([0 tmax], X0, X00);
            %=====
            %***** DISPLAYING RESULTS *****
            CF=Y(4*N+M+1:N*(N+4)+M, end); CF=reshape(CF,N,N)';
            figure;
            colormap(flipud(hot));
            %h=imagesc([1 2*N+M], [0 tmax], Y([1:2*N, 4*N+1:4*N+M], :));
            h=imagesc([1 N+M], [0 10*tmax], Y([1:N, 4*N+1:4*N+M], :));
            set(gca, 'FontSize', 32);
            set(gca, 'xTick', (10:10:90));
            set(gca, 'xTickLabel', {'10'; '20'; '30'; '40'; '50'; '60'; '70'; '80'; '10'});
            xlabel('unit number', 'FontSize', 34, 'FontName', 'Arial');
            ylabel('time [ms]', 'FontSize', 34, 'FontName', 'Arial');
            %title('Excitatory and inhibitory activity', 'FontSize', 32);

```


A.3.6 euler23d.m

```

function [Y]=euler23d(tspan,X0,X00)

% [Y]=euler23d(tspan,X0,X00). This function will simulate the model in
% seq23.m using the forward Euler method starting from the initial
% conditions X0, and assuming that at 'del' time units before X0 the state
% of the system was X00.
%
% Notice that each connection has its own delay, so del is a N*N vector.
% Since we have N*N possibly different delays, to begin simulating we would
% need N*N activity values, each one corresponding to a different delay.
% The argument X00 we receive, however, is a single initial initial
% state of size N*(N+M+4)+M, so what we do is to create a matrix Y00
% with N*(N+M+4)+M rows and ceil(max(del/dt)) columns. Y00(:,end)=X0, and
% Y00(:,1) is the state of the system at max(del) time units before the
% start of the simulation. The states between Y00(:,1) and Y00(:,end) will
% be obtained from linear interpolation. Since most delays are almost
% identical, we simplify things and make Y00(:,1)=X00.
%
% The simulation will range in time from tspan(1) until tspan(2).

global del dt N

NSTEPS = ceil((tspan(2)-tspan(1))/dt); % total number of simulation steps
STEPS = max(1,ceil(del/dt)); % number of steps for each one of the delays
maxSTEPS = max(STEPS); % maximum number of steps between delayed and current values
minSTEPS = min(STEPS); % minimum number of steps between delayed and current values

Y=zeros(length(X0),NSTEPS); % This matrix will contain the simulation data
Y(:,1)=X0;

% the next line create maxSTEPS linear interpolations from X00 to X0.
Y00 = repmat(X00,1,maxSTEPS) + repmat(X0-X00,1,maxSTEPS).*repmat(linspace(0,1,maxSTEPS),length(X0),1);

% In order to calculate the next step, you only need N*N prior activity
% values. The value which multiplies connection C(j,k) in the dynamic
% equations is the k-th activity with the (j,k) delay. The (j,k) delay is
% del(N*(j-1)+k), and the corresponding activity is
% Y00(k,maxSTEPS+i-STEPS(N*(j-1)+k)). What we are sending into seqF23d is
% a vector Xdel with N*N elements:
% Xdel(1:N) = Xk(t - del(k)), for k=1:N; <-- inputs to first unit
% Xdel(N+1:2*N) = Xk(t - del(N+k)), k=1:N; <-- inputs to second unit
% Xdel((j-1)*N:j*N) = Xk(t - del(N*(j-1)+k)), k=1:N; <-- inputs to unit j
% Notice that Xdel((j-1)*N:j*N) =
% diag(Y00(1:N,maxSTEPS+i-STEPS(N*(j-1)+1:N)))
% Xdel consists of these diagonals concatenated into a column vector.

% In order to obtain Xdel in a single line we use linear indexing of the
% EY00 matrix, which contains the N activity values at the N*N delays.
% The index vector is complicated, but you only have to obtain it once.
% Notice that the index vector doesn't work if we use Y00 instead of EY00.

base = 0:N*N:(N-1)*N*N; % the index will jump through square matrices
base = reshape(repmat(base,N,1),N*N,1);
prog = 1:N+1:N*N; % the index will get the main diagonal
prog = repmat(prog,1,N)';
index = base+prog; % the index is ready

Xdel = zeros(N*N,1);
t=tspan(1);

for i=2:minSTEPS
    EY00 = Y00(1:N,maxSTEPS+i-STEPS);
    Xdel = EY00(index);
    Y(:,i) = Y(:,i-1) + dt*seqF23d(t,Y(:,i-1),Xdel);
    t=t+dt;
end

for i=minSTEPS+1:maxSTEPS % This segment is short, so I just use for loops
    for j=1:N
        for k=1:N
            unit = (j-1)*N+k;
            if STEPS(unit) >= i % delay sends you back to Y00
                Xdel(unit) = Y00(j,maxSTEPS+i-STEPS(unit));
            else
                Xdel(unit) = Y(i-STEPS(unit));
            end
        end
    end
end

```

```

        end
    end
    Y(:,i) = Y(:,i-1) + dt*seqF23d(t,Y(:,i-1),Xdel);
    t=t+dt;
end
clear Y00 % releasing some memory

for i=maxSTEPS+1:NSTEPS
% A for loop is my insurance if the linear indexing doesn't work
%   for j=1:N
%       Xdel((j-1)*N:j*N) = Y(j,i-STEPS((j-1)*N+(1:N)));
%   end
    EY00 = Y(1:N,i-STEPS);
    Xdel = EY00(index);
    Y(:,i) = Y(:,i-1) + dt*seqF23d(t,Y(:,i-1),Xdel);
    t=t+dt;
end

```

A.3.7 seqF23.m

```

function [Y] = seqF23(T,X,Xdel)

% [Y] = seqF23(T,X,Xdel) implements the dynamics of the model in seq23.m.
% All the parameters are passed as global variables.
% The argument T is time (a scalar); X is the current state vector, and
% Xdel is the state vector at time T-sdel.
% The state vectors which are received have N*(N+M+4)+M variables:
% [ x(1:N); w(1:N); z(1:N); n(1:N); y(1:M); C(1:N*N); V(1:N*M) ]
% where x(i) = firing rate of excitatory unit i
%       w(i) = memory of the activity in excitatory unit i
%       z(i) = adaptation of excitatory unit i
%       n(i) = noise in excitatory unit i
%       y(i) = firing rate of inhibitory unit i
%       C(i) = excitatory connection matrix reshaped as a column.
%           C(1:N)=CEE(1,1:N); C(N+1:2*N)=CEE(2,1:N); ...
%       V(i) = inhibitory to excitatory connections reshaped as a column.
%           V(1:M)=CEI(1,1:M); V(M+1:2*M)=CEI(2,1:M); ...

global aee aei aie alpha eps b c d e h s S r q mid gamma the thi thw thdec thinc tc te ti tv tz tn tw NM
      IEI CIE rsqdt inpFun

%----- EXTRACTING THE VARIABLES FROM THE INPUT VECTORS -----
Y = zeros(N*(N+M+4)+M,1);
E = X(1:N); % excitatory units
W = X(N+1:2*N); % activity memory
Z = X(2*N+1:3*N); % adaptation
H = X(3*N+1:4*N); % noise
y = X(4*N+1:4*N+M); % inhibitory units
Edel = Xdel(1:N); % delayed excitation
C = reshape(X(4*N+M+1:N*(N+4)+M),N,N)'; % excitatory connections
V = reshape(X(N*(N+4)+M+1:end),M,N)'; % I-E connections
I = inpFun(T); % inpFun is a handle to the input function (I7-I9)
%-----

% ++++++ IMPLEMENTING THE DYNAMICS ++++++
Cin = C*Edel;
argx = aee.*E + c*Cin - d.*Z - aei.*V.*y - the + e.*H + I(1:N);
argy = max(aie.*CIE.*E - thi + I(N+1),zeros(M,1));
fx = ones(N,1)./(1+exp(-argx)); fw = ones(N,1)./(1+exp(-gamma.*(E-thw)));

Y(1:N) = (-E + fx)./te; % excitatory units
Y(N+1:2*N) = (-W + fw)./tw; % activity memory
Y(2*N+1:3*N) = -Z./tz + b.*(1-Z).*fx; % adaptation
Y(3*N+1:4*N) = -H./tn + rsqdt.*randn(N,1); % noise
Y(4*N+1:4*N+M) = -y./ti + argy; % inhibitory units

% next comes the E-E synaptic plasticity
Erep=repmat(E',N,1); Erep=reshape(Erep,N*N,1);
Eprog=repmat(E,N,1);
Wrep=repmat(W',N,1); Wrep=reshape(Wrep,N*N,1);
Wprog=repmat(W,N,1);
repCin=repmat(Cin',N,1); repCin=reshape(repCin,N*N,1);
repEdel=repmat(Edel',N,1); repEdel=reshape(repEdel,N*N,1);
base1=4*N+M+1; base2=N*(N+4)+M;
Hinc=ones(N*N,1)./(ones(N*N,1)+exp(-q.*(Erep.*Wprog-thinc)));
Hdec=ones(N*N,1)./(ones(N*N,1)+exp(-q.*(Eprog.*Wrep-thdec)));

```

```

poten=(alpha+X(base1:base2)).*(s-X(base1:base2)).*max(zeros(N*N,1),(S - repCin)).*max(zeros(N*N,1),(Hinc-
Hdec));
depl=S.*X(base1:base2).*max(zeros(N*N,1),(Hdec-Hinc));
%hetero=eps.*X(base1:base2).*Erep.*(1-repEdel); % heterosynaptic competition
inter=eps.*max(zeros(N*N,1),(X(base1:base2)-alpha).*(mid-X(base1:base2))); %reduction of intermediate
weights

%Y(base1:base2) = (poten - dep1 - hetero)./tc; % excitatory plasticity
Y(base1:base2) = (poten - dep1 - inter)./tc; % excitatory plasticity

Y(base1 + (N+1)*(0:N-1)) = zeros(N,1); % making C(i,i) = 0

% and finally the I-E plasticity
ErepM=repmat(E,1,M); ErepM=reshape(ErepM',N*M,1);
Y(base2+1:end) = (X(base2+1:end).*(1EI + r.*ErepM - X(base2+1:end)).*max((ErepM - h),zeros(N*M,1)))./tv;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

A.3.8 seqF23d.m

```

function [Y] = seqF23d(T,X,Xdel)

% [Y] = seqF23d(T,X,Xdel) implements the dynamics of the model in seq23.m.
% All the parameters are passed as global variables.
% The argument T is time (a scalar); X is the current state vector, and
% Xdel has the N*N delayed activities required to compute this function. A
% description of Xdel can be found in euler23d.m.

% The state vector which is received has N*(N+M+4)+M variables:
% [ x(1:N); w(1:N); z(1:N); n(1:N); y(1:M); C(1:N*N); V(1:N*M) ]
% where x(i) = firing rate of excitatory unit i
% w(i) = memory of the activity in excitatory unit i
% z(i) = adaptation of excitatory unit i
% n(i) = noise in excitatory unit i
% y(i) = firing rate of inhibitory unit i
% C(i) = excitatory connection matrix reshaped as a column.
% C(1:N)=CEE(1,1:N); C(N+1:2*N)=CEE(2,1:N); ...
% V(i) = inhibitory to excitatory connections reshaped as a column.
% V(1:M)=CEI(1,1:M); V(M+1:2*M)=CEI(2,1:M); ...
%
% Unlike seqF23.m, this version deals with heterogeneous delays. Since we
% have N*N excitatory connections, there are N*N different delays, and
% therefore Xdel contains the N*N previous states required to compute the
% next one. Xdel(1,1:N) are the states required by the delays of
% connections C(1:N), namely, the connections incoming to neuron 1.
% Xdel(2,1:N) are the activities whose corresponding delays are those of
% the connections CEE(2,1:N), etc.

global aee aei aie alpha eps b c d e h s S r q mid gamma the thi thw thdec thinc tc te ti tv tz tn tw NM
1EI CIE rsqdt inpFun

%----- EXTRACTING THE VARIABLES FROM THE INPUT VECTORS -----
Y = zeros(N*(N+M+4)+M,1);
E = X(1:N); % excitatory units
W = X(N+1:2*N); % activity memory
Z = X(2*N+1:3*N); % adaptation
H = X(3*N+1:4*N); % noise
y = X(4*N+1:4*N+M); % inhibitory units
C = reshape(X(4*N+M+1:N*(N+4)+M),N,N); % transpose of the excitatory connection matrix
V = reshape(X(N*(N+4)+M+1:end),M,N); % I-E connections
D = reshape(Xdel,N,N); % delay matrix, ready for dot multiplication with C
I = inpFun(T); % inpFun is a handle to the input function (I7-I9)
%
%
% ++++++ IMPLEMENTING THE DYNAMICS ++++++
Cin = sum(C.*D)';
argx = aee.*E + c.*Cin - d.*Z - aei.*V.*y - the + e.*H + I(1:N);
argy = max(aie.*CIE.*E - thi + I(N+1),zeros(M,1));
fx = ones(N,1)./(1+exp(-argx)); fw = ones(N,1)./(1+exp(-gamma.*(E-thw)));

Y(1:N) = (-E + fx)./te; % excitatory units
Y(N+1:2*N) = (-W + fw)./tw; % activity memory
Y(2*N+1:3*N) = -Z./tz + b.*(1-Z).*fx; % adaptation
Y(3*N+1:4*N) = -H./tn + rsqdt.*randn(N,1); % noise
Y(4*N+1:4*N+M) = -y./ti + argy; % inhibitory units

% next comes the E-E synaptic plasticity
Erep=repmat(E',N,1); Erep=reshape(Erep,N*N,1);

```

```

Eprog=repmat(E,N,1);
Wrep=repmat(W,N,1); Wrep=reshape(Wrep,N*N,1);
Wprog=repmat(W,N,1);
repCin = repmat(Cin',N,1); repCin = reshape(repCin,N*N,1);
base1=4*N+M+1; base2=N*(N+4)+M;
Hinc=ones(N*N,1)/(ones(N*N,1)+exp(-q.*(Erep.*Wprog-thinc)));
Hdec=ones(N*N,1)/(ones(N*N,1)+exp(-q.*(Eprog.*Wrep-thdec)));
poten=(alpha+X(base1:base2)).*(s-X(base1:base2)).*max(zeros(N*N,1),(S-repCin)).*max(zeros(N*N,1),(Hinc-Hdec));
depl=S.*X(base1:base2).*max(zeros(N*N,1),(Hdec-Hinc));
%hetero=eps.*X(base1:base2).*Erep.*(1-Xdel);
inter=eps.*max(zeros(N*N,1),(X(base1:base2)-alpha).*(mid-X(base1:base2))); %reduction of intermediate weights

%Y(base1:base2) = (poten - depl - hetero)./tc; % excitatory plasticity
Y(base1:base2) = (poten - depl - inter)./tc; % excitatory plasticity

Y(base1 + (N+1)*(0:N-1)) = zeros(N,1); % making C(i,i) = 0

% and finally the I-E plasticity
ErepM=repmat(E,1,M); ErepM=reshape(ErepM',N*M,1);
Y(base2+1:end) = (X(base2+1:end)).*(1EI + r.*ErepM - X(base2+1:end)).*max((ErepM - h),zeros(N*M,1))./tv;
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```

A.3.9 I7.m

```

function [In] = I7(t)

% [In] = I7(t). This function is used to provide structured input to
% the seqI7.m model. I7(t) is called in the seqF17.m file.
% The input provided consists of periodic pulses to all the units indicated
% in the vector V. The frequency is in the global variable freq, and will
% start and stop according to the value of the variables Tin and Tout.
% The first N elements of In correspond to the excitatory units. The last
% element corresponds to the inhibitory unit.

global Tin Tout N I V freq

In = zeros(N+1,1);

if Tin <= t && t < Tout
    t = t - Tin;
    In(V) = I*exp(-3*(1-cos(freq*t)));
end

```

A.3.10 I8.m

```

function [In] = I8(t)

% [In] = I8(t). This function is used to provide structured input to
% the seqI9.m model. I8(t) is called in the seqF19.m file.
% The input provided activates the units in a sequence. The sequence will
% follow the order of the vector V (a global variable), and will last from
% Tin until Tout (also global variables). Unlike I6.m, the sequence may be
% have repeated elements.
% The first n elements of In correspond to the excitatory units. The last
% element corresponds to the inhibitory unit. There is an additional line
% which always adds input to the inhibitory unit. It seems to be useful
% when modifying existing patterns.

global Tin Tout N I V

In = zeros(N+1,1);
M = length(V);
%steps = (Tout - Tin)/M;

if Tin < t && t <= Tout
    t = t - Tin;
    In(V(ceil(M*t/(Tout-Tin)))) = I;
    In(N+1) = I/40; % additional inhibition
end

```

A.3.11 I9.m

```
function [In] = I9(t)

% [In] = I9(t). This function is used to provide structured input to
% the seq20.m model. I9(t) is called in the seqF20.m file.
% The input provided activates vectors of units in a sequence. The
% sequence of vectors consists of the rows in matrix V (a global variable),
% and it will start at Tin, finish at Tout, and give each vector an equal
% amount of time as the input.
% The first n elements of In correspond to the excitatory units. The last
% element corresponds to the inhibitory unit. There is an additional line
% which always adds input to the inhibitory unit. It seems to be useful
% when modifying existing patterns.

global Tin Tout N I V

In = zeros(N+1,1);
[r c] = size(V);
%steps = (Tout - Tin)/M;

if Tin < t && t <= Tout
    t = t - Tin;
    In(V(ceil(r*t/(Tout-Tin)),:)) = I;
    %In(N+1) = I/35; % additional inhibition
end
```

A.3.12 createCEI2.m

```
function [CEI]=createCEI2(n,m,pei)

% [CEI]=createCEI2(n,m,pei) creates the matrix of inhibitory to excitatory
% projections in the seq20 model. n is the number of excitatory units, m is
% the number of inhibitory units, and pei is the probability that a given
% inhibitory unit will project to any given excitatory unit.

CEI = rand(n,m) < pei;

% normalizing row sums to 1
CEI = CEI./(repmat(sum(CEI,2),1,m)+1e-5);

% making heterogeneous connections - not used
%CEI = abs(CEI + (0.05/m)*randn(n,m));
```

A.3.13 createCIE1.m

```
function [CIE]=createCIE1(n,m,pie)

% [CIE]=createCIE1(n,m,pie) creates the matrix of inhibitory to excitatory
% projections in the seq18 model. n is the number of excitatory units, m is
% the number of inhibitory units, and pie is the probability that a given
% excitatory unit will project to any given inhibitory unit.

CIE = rand(m,n) < pie;

% normalizing row sums to 1
CIE = CIE./(repmat(sum(CIE,2),1,n)+1e-5);

% making heterogeneous connections
CIE = abs(CIE + (0.05/n)*randn(m,n));
```

A.3.14 randSim1.m

```
% randSim1.m
% This program repeatedly runs simulations where the input consists of
% random sequences. See the 6/13/10 entry in log.txt.
```

```

clear all;
close all;
home;

global aee aei aie alpha eps b c d e h s R r q mid del sdel gamma the thi thw thdec thinc tc te ti tv tw
tz tn N M tmax I Tin Tout V freq CIE lEI dt rsqdt inpFun

%*****
%***** INITIALIZE PARAMETERS *****
%*****

% PARAMETER VALUES
%-----
aee=8; aei=15; aie=15; b=0.05; c=3; d=11; e=1; h=0.01; S=4; gamma=5;
r=0.3; s=1.2; q=30; mid=0.7; the=3.8; thi=0.2; thw=0.5; thdec=0.3; thinc=0.3;
sdel=2; alpha=0.01; eps=0.15; tc=7; te=1; ti=0.5; tv=40e9; tz=7; tn=0.5; tw=5;
N=80; M=16; % number of excitatory and inhibitory units respectively
PEI=0.5; PIE=0.5; % probability of EI and IE connections
dt=0.05; % step size for Euler method
rsqdt=1/sqrt(dt); % used to avoid extra computations

% MAKING PARAMETERS HETEROGENEOUS
std=0.04; % 0.04
aee=std*aee*randn(N,1); % for aie and aei the heterogeneity comes from the connection matrix
b=std*b*randn(N,1); d=std*d*randn(N,1); e=std*e*randn(N,1);
gamma=std*gamma*randn(N,1); h=std*h*randn(N,1); s=std*s*randn(N,1);
S=std*S*randn(N,1); r=std*r*randn(N,1); q=std*q*randn(N,1); del=std*del*randn(N*N,1);
tn=std*tn*randn(N,1); tz=std*tz*randn(N,1); te=std*te*randn(N,1);
tw=std*tw*randn(N,1); tc=std*tc*randn(N,1); tv=std*tv*randn(N,1);
the=std*the*randn(N,1); thi=std*thi*randn(N,1); thw=std*thw*randn(N,1);
thdec=std*thdec*randn(N,1); thinc=std*thinc*randn(N,1);
alpha=std*alpha*randn(N,1); eps=std*eps*rand(N,1); mid=std*mid*randn(N,1);

thdec=repmat(thdec',N,1); thdec=reshape(thdec,N*N,1); % STDP seems to be a postsynaptic affair
thinc=repmat(thinc',N,1); thinc=reshape(thinc,N*N,1);
tc=repmat(tc',N,1); tc=reshape(tc,N*N,1);
tv=repmat(tv',M,1); tv=reshape(tv,N*M,1);
h=repmat(h,1,M); h=reshape(h',N*M,1);
S=repmat(S',N,1); S=reshape(S,N*N,1);
s=repmat(s',N,1); s=reshape(s,N*N,1);
q=repmat(q',N,1); q=reshape(q,N*N,1);
r=repmat(r,1,M); r=reshape(r',N*M,1);
eps=repmat(eps',N,1); eps=reshape(eps,N*N,1);
alpha=repmat(alpha',N,1); alpha=reshape(alpha,N*N,1);
mid=repmat(mid',N,1); mid=reshape(mid,N*N,1);
%-----

% INITIAL STATE, INITIAL CONNECTIONS
%+++++
X0=zeros(N*(N+M+4)+M,1);
X0(1:N)=0.2*rand(N,1);
X00=zeros(N*(N+M+4)+M,1);
X00(1:N)=0.2*rand(N,1);

CEE = abs(0.05*rand(N));
%CEE = createCEE1(N,[20:30,20])+createCEE2(N,[40:50,40]);
CEI = createCEI2(N,M,PEI); % creates a matrix with no heterogeneity
lEI = reshape(CEI',M*N,1); % form of CEI used in seqF20
CIE = createCIE1(N,M,PIE);
CF = CEE; % so we can select case 5 in the first loop repetition

X0(4*N+M+1:N*(N+4)+M) = reshape(CEE',N*N,1); % initial E-E connections
X0(N*(N+4)+M+1:N*(N+M+4)+M) = reshape(CEI',N*M,1); % initial I-E connections
X00(4*N+M+1:N*(N+4)+M) = X0(4*N+M+1:N*(N+4)+M); % same E-E connections
X00(N*(N+4)+M+1:N*(N+M+4)+M) = X0(N*(N+4)+M+1:N*(N+M+4)+M); % same I-E connections
%+++++

% MODIFICATIONS TO STANDARD PARAMETERS
%-----
%load T23A1 % default parameters, 6 intersecting simple loops
%load T23A2CS2 % default pars, complex sequence order 2
%load T23B1MS % 4 intersecting simple loops, thdec=0.6.
%load /home/z/Documents/athena/paper3/figures/A1
%load randSim1pars;

%eps=0;
%thinc=0.1;
%aie=10; aei=10;
%r=0; % eliminating homeostatic inhibition
%tv=40e9; %tc=10e10;

```


A.3.15 randSim2.m

```

% randSim2.m
% This program repeatedly runs simulations where the input consists of
% random sequences. See the 7/19/10 entry in log.txt.

clear all;
close all;
home;

global aee aei aie alpha eps b c d e h s S r q mid del sdel gamma the thi thw thdec thinc tc te ti tv tw
      tz tn N M tmax I Tin Tout V freq CIE lEI dt rsqdt inpFun

%*****
%***** INITIALIZE PARAMETERS *****
%*****

% PARAMETER VALUES
%-----
aee=8; aei=15; aie=15; b=0.05; c=3; d=11; e=1; h=0.01; S=4; gamma=5;
r=0.3; s=1.2; q=30; mid=0.6; the=3.8; thi=0.2; thw=0.5; thdec=0.3; thinc=0.3;
sdel=2; alpha=0.01; eps=0.07; tc=7; te=1; ti=0.5; tv=40e9; tz=7; tn=0.5; tw=5;
N=80; M=16; % number of excitatory and inhibitory units respectively
PEI=0.5; PIE=0.5; % probability of EI and IE connections
dt=0.05; % step size for Euler method
rsqdt=1/sqrt(dt); % used to avoid extra computations

% MAKING PARAMETERS HETEROGENEOUS
std=0.04; % 0.04
aee=aee+std*aee*randn(N,1); % for aie and aei the heterogeneity comes from the connection matrix
b=b+std*b*randn(N,1); d=d+std*d*randn(N,1); e=e+std*e*randn(N,1);
gamma=gamma+std*gamma*randn(N,1); h=h+std*h*randn(N,1); s=s+std*s*randn(N,1);
S=S+std*S*randn(N,1); r=r+std*r*randn(N,1); q=q+std*q*randn(N,1); del=sdel+std*sdel*randn(N*N,1);
tn=tn+std*tn*randn(N,1); tz=tz+std*tz*randn(N,1); te=te+std*te*randn(N,1);
tw=tw+std*tw*randn(N,1); tc=tc+std*tc*randn(N,1); tv=tv+std*tv*randn(N,1);
the=the+std*the*randn(N,1); thi=thi+std*thi*randn(M,1); thw=thw+std*thw*randn(N,1);
thdec=thdec+std*thdec*randn(N,1); thinc=thinc+std*thinc*randn(N,1);
alpha=alpha+std*alpha*randn(N,1); eps=eps+std*eps*rand(N,1); mid=mid+std*mid*randn(N,1);

thdec=repmat(thdec',N,1); thdec=reshape(thdec,N*N,1); % STDP seems to be a postsynaptic affair
thinc=repmat(thinc',N,1); thinc=reshape(thinc,N*N,1);
tc=repmat(tc',N,1); tc=reshape(tc,N*N,1);
tv=repmat(tv',M,1); tv=reshape(tv,N*M,1);
h=repmat(h,1,M); h=reshape(h',N*M,1);
S=repmat(S',N,1); S=reshape(S,N*N,1);
s=repmat(s',N,1); s=reshape(s,N*N,1);
q=repmat(q',N,1); q=reshape(q,N*N,1);
r=repmat(r,1,M); r=reshape(r',N*M,1);
eps=repmat(eps',N,1); eps=reshape(eps,N*N,1);
alpha=repmat(alpha',N,1); alpha=reshape(alpha,N*N,1);
mid=repmat(mid',N,1); mid=reshape(mid,N*N,1);
%-----

% INITIAL STATE, INITIAL CONNECTIONS
%+++++
X0=zeros(N*(N+M+4)+M,1);
X0(1:N)=0.2*rand(N,1);
X00=zeros(N*(N+M+4)+M,1);
X00(1:N)=0.2*rand(N,1);

CEE = abs(0.05*rand(N));
%CEE = createCEE1(N,[20:30,20])+createCEE2(N,[40:50,40]);
CEI = createCEI2(N,M,PEI); % creates a matrix with no heterogeneity
lEI = reshape(CEI',M*N,1); % form of CEI used in seqF20
CIE = createCIE1(N,M,PIE);
CF = CEE; % so we can select case 5 in the first loop repetition

X0(4*N+M+1:N*(N+4)+M) = reshape(CEE',N*N,1); % initial E-E connections
X0(N*(N+4)+M+1:N*(N+M+4)+M) = reshape(CEI',N*M,1); % initial I-E connections
X00(4*N+M+1:N*(N+4)+M) = X0(4*N+M+1:N*(N+4)+M); % same E-E connections
X00(N*(N+4)+M+1:N*(N+M+4)+M) = X0(N*(N+4)+M+1:N*(N+M+4)+M); % same I-E connections
%+++++

% MODIFICATIONS TO STANDARD PARAMETERS
%-----
%load T23A1 % default parameters, 6 intersecting simple loops
%load T23A2CS2 % default pars, complex sequence order 2
%%load T23B1MS % 4 intersecting simple loops, thdec=0.6.

```


BIBLIOGRAPHY

- [1] A Treves and D J Amit. Low firing rates: an effective Hamiltonian for excitatory neurons. *Journal of Physics A: Mathematical and General*, 22(12):2205, 1989.
- [2] L. F. Abbott and K. I. Blum. Functional significance of long-term potentiation for sequence learning and prediction. *Cerebral cortex (New York, N.Y. : 1991)*, 6(3):406–16, 1996.
- [3] M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, New York, 1991.
- [4] S. Abrahams, R. G. Morris, C. E. Polkey, J. M. Jarosz, T. C. Cox, M. Graves, and A. Pickering. Hippocampal involvement in spatial and working memory: a structural MRI analysis of patients with unilateral mesial temporal lobe sclerosis. *Brain and cognition*, 41(1):39–65, Oct. 1999.
- [5] S.-I. Amari. Learning Patterns and Pattern Sequences by Self-Organizing Nets of Threshold Elements. *IEEE Transactions on Computers*, c-21(11):1197–1206, 1972.
- [6] D. J. Amit. *Modelling brain function: the world of attractor neural networks*. Cambridge University Press, Cambridge, 1989.
- [7] D. J. Amit. The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behavioral and Brain Sciences*, 13(4):617–657, Dec. 1995.
- [8] D. J. Amit and N. Brunel. Dynamics of a recurrent network of spiking neurons before and following learning. *Network: Computation in Neural Systems*, 8:373–404, Nov. 1997.
- [9] D. J. Amit and N. Brunel. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cerebral cortex (New York, N.Y. : 1991)*, 7(3):237–52, 1997.
- [10] D. J. Amit and M. V. Tsodyks. Quantitative study of attractor neural network retrieving at low spike rates: I. substratespikes, rates and neuronal gain. *Network: Computation in Neural Systems*, 2(3):259–273, 1991.

- [11] D. J. Amit and M. V. Tsodyks. Quantitative study of attractor neural networks retrieving at low spike rates: II. Low-rate retrieval in symmetric networks. *Network: Computation in Neural Systems*, 2(3):275–294, 1991.
- [12] H. E. Atallah, M. J. Frank, and R. C. O’Reilly. Hippocampus, cortex, and basal ganglia: insights from computational models of complementary learning systems. *Neurobiology of learning and memory*, 82(3):253–67, Nov. 2004.
- [13] O. Barak and M. Tsodyks. Persistent activity in neural networks with dynamic synapses. *PLoS computational biology*, 3(2):e35, Feb. 2007.
- [14] F. Barbieri and N. Brunel. Irregular persistent activity induced by synaptic excitatory feedback. *Frontiers in Computational Neuroscience*, 1(5):1–12, 2007.
- [15] F. Barbieri and N. Brunel. Can attractor network models account for the statistics of firing during persistent activity in prefrontal cortex ? *Frontiers in Neuroscience*, 2(1):114–122, 2008.
- [16] A. Bechara, H. Damasio, D. Tranel, and S. W. Anderson. Dissociation Of working memory from decision making within the human prefrontal cortex. *The Journal of Neuroscience*, 18(1):428–37, Jan. 1998.
- [17] E. Ben-Menachem. Modern management of epilepsy: Vagus nerve stimulation. *Baillière’s clinical neurology*, 5(4):841–8, Dec. 1996.
- [18] T. Bliss and T. Lømo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of physiology*, 232(2):331–356, 1973.
- [19] M. Bodner, J. Kroger, and J. M. Fuster. Auditory memory cells in dorsolateral prefrontal cortex. *Neuroreport*, 7(12):1905–8, Aug. 1996.
- [20] M. Bodner, L. T. Muftuler, O. Nalcioglu, and G. L. Shaw. FMRI study relevant to the Mozart effect: brain areas involved in spatial-temporal reasoning. *Neurological research*, 23(7):683–90, Oct. 2001.
- [21] M. Bodner, M. Shafi, Y.-D. Zhou, and J. M. Fuster. Patterned firing of parietal cells in a haptic working memory task. *The European journal of neuroscience*, 21(9):2538–46, May 2005.
- [22] M. Bodner, Y. D. Zhou, and J. M. Fuster. Binary mapping of cortical spike trains in short-term memory. *Journal of neurophysiology*, 77(4):2219–22, Apr. 1997.
- [23] M. Bodner, Y. D. Zhou, and J. M. Fuster. High-frequency transitions in cortical spike trains related to short-term memory. *Neuroscience*, 86(4):1083–7, Oct. 1998.
- [24] E. A. Brenowitz, D. Margoliash, and K. W. Nordeen. An introduction to birdsong and the avian song system. *Journal of neurobiology*, 33(5):495–500, Nov. 1997.

- [25] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of computational neuroscience*, 8(3):183–208, 2000.
- [26] N. Brunel and X. J. Wang. Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of computational neuroscience*, 11(1):63–85, 2001.
- [27] D. V. Buonomano. Timing of neural responses in cortical organotypic slices. *Proceedings of the National Academy of Sciences of the United States of America*, 100(8):4897–902, Apr. 2003.
- [28] E. Caianiello, A. de Luca, and L. Ricciardi. Reverberations and Control of Neural Networks. *Kybernetik*, 4:10–18, 1967.
- [29] M. Camperi and X. J. Wang. A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of computational neuroscience*, 5(4):383–405, Dec. 1998.
- [30] N. Caporale and Y. Dan. Spike timing-dependent plasticity: a Hebbian learning rule. *Annual review of neuroscience*, 31:25–46, Jan. 2008.
- [31] P. A. Cariani. Temporal codes and computations for sensory representation and scene analysis. *IEEE transactions on neural networks / a publication of the IEEE Neural Networks Council*, 15(5):1100–11, Sept. 2004.
- [32] C. E. Carr. Processing of temporal information in the brain. *Annu. Rev. Neurosci.*, 16:226–243, 1993.
- [33] M. A. Castro-Alamancos and B. W. Connors. Short-term synaptic enhancement and long-term potentiation in neocortex. *Proceedings of the National Academy of Sciences of the United States of America*, 93(3):1335–9, Mar. 1996.
- [34] M. A. Castro-Alamancos, J. P. Donoghue, and B. W. Connors. Different forms of synaptic plasticity in somatosensory and motor areas of the neocortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 15(7 Pt 2):5324–33, July 1995.
- [35] M. V. Chafee and P. S. Goldman-Rakic. Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *Journal of neurophysiology*, 83(3):1550–66, Mar. 2000.
- [36] W. Chang and D. Jin. Spike propagation in driven chain networks with dominant global inhibition. *Physical Review E*, 79(5):1–5, May 2009.
- [37] A. Citri and R. C. Malenka. Synaptic plasticity: multiple forms, functions, and mechanisms. *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, 33(1):18–41, Jan. 2008.

- [38] A. Compte, N. Brunel, P. S. Goldman-rakic, and X.-j. Wang. Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex*, 10:910–923, 2000.
- [39] A. Compte, C. Constantinidis, J. Tegner, S. Raghavachari, M. V. Chafee, P. S. Goldman-Rakic, and X.-J. Wang. Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. *Journal of neurophysiology*, 90(5):3441–54, Nov. 2003.
- [40] S. J. Cooper. Donald O. Hebb’s synapse and learning rule: a history and commentary. *Neuroscience and biobehavioral reviews*, 28(8):851–74, Jan. 2005.
- [41] A. Covanis. Photosensitivity in idiopathic generalized epilepsies. *Epilepsia*, 46 Suppl 9:67–72, Jan. 2005.
- [42] C. M. Cowey and S. Green. The hippocampus: a ”working memory” structure? The effect of hippocampal sclerosis on working memory. *Memory (Hove, England)*, 4(1):19–30, Jan. 1996.
- [43] C. Curtis and M. D’Esposito. Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, 7(9):415–423, Sept. 2003.
- [44] B. Delord, P. Baraduc, R. Costalat, Y. Burnod, and E. Guigon. A model study of cellular short-term memory produced by slowly inactivating potassium conductances. *Journal of computational neuroscience*, 8(3):251–73, 2000.
- [45] B. Delord, A. J. Klaassen, Y. Burnod, R. Costalat, and E. Guigon. Bistable behaviour in a neocortical neurone model. *Neuroreport*, 8(4):1019–23, Mar. 1997.
- [46] M. A. Dichter and G. F. Ayala. Cellular mechanisms of epilepsy: a status report. *Science (New York, N.Y.)*, 237(4811):157–64, July 1987.
- [47] M. Diesmann, M. O. Gewaltig, and a. Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–33, Dec. 1999.
- [48] R. Doursat and E. Bienenstock. Neocortical Self-Structuration as a Basis for Learning. In *5th International Conference on Development and Learning (ICDL 2006)*, number Icdl, pages 1–6, 2006.
- [49] D. Durstewitz. Self-organizing neural integrator predicts interval times through climbing activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(12):5342–53, June 2003.
- [50] D. Durstewitz, J. K. Seamans, and T. J. Sejnowski. Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex. *Journal of neurophysiology*, 83(3):1733–1750, Mar. 2000.

- [51] D. Durstewitz, J. K. Seamans, and T. J. Sejnowski. Neurocomputational models of working memory. *Nature neuroscience*, 3 Suppl(november):1184–91, Nov. 2000.
- [52] A. Egorov, B. Hamam, E. Fransén, M. Hasselmo, and A. Alonso. Graded persistent activity in entorhinal cortex neurons. *Nature*, 420(6912):173–178, 2002.
- [53] J. L. Elman. Finding Structure in Time. *Cognitive Science*, 14(2):179–211, Mar. 1990.
- [54] B. Ermentrout. Type I membranes, phase resetting curves, and synchrony. *Neural computation*, 8(5):979–1001, July 1996.
- [55] G. B. Ermentrout. *Simulating, analyzing and animating dynamical systems: A guide to XPPAUT for researchers and students*. SIAM, New York, 1st edition, 2002.
- [56] Z. Feng and D. M. Durand. Suppression of excitatory synaptic transmission can facilitate low-calcium epileptiform activity in the hippocampus in vivo. *Brain research*, 1030(1):57–65, Dec. 2004.
- [57] I. R. Fiete, W. Senn, C. Z. H. Wang, and R. H. R. Hahnloser. Spike-time-dependent plasticity and heterosynaptic competition organize networks to produce long scale-free sequences of neural activity. *Neuron*, 65(4):563–76, Feb. 2010.
- [58] P. J. Franaszczuk, G. K. Bergey, P. J. Durka, and H. M. Eisenberg. Time-frequency analysis using the matching pursuit algorithm applied to seizures originating from the mesial temporal lobe. *Electroencephalography and clinical neurophysiology*, 106(6):513–21, June 1998.
- [59] E. Fransén, A. A. Alonso, and M. E. Hasselmo. Simulations of the role of the muscarinic-activated calcium-sensitive nonspecific cation current INCM in entorhinal neuronal activity during delayed matching tasks. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(3):1081–97, Feb. 2002.
- [60] E. Fransén, B. Tahvildari, A. V. Egorov, M. E. Hasselmo, and A. a. Alonso. Mechanism of graded persistent cellular activity of entorhinal cortex layer v neurons. *Neuron*, 49(5):735–46, Mar. 2006.
- [61] T. Freund and G. Buzsáki. Interneurons of the hippocampus. *Hippocampus*, 6(4):347–470, Dec. 1998.
- [62] P. Fries. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, 9(10):474–80, Oct. 2005.
- [63] S. Funahashi. Prefrontal cortex and working memory processes. *Neuroscience*, 139(1):251–61, Apr. 2006.
- [64] S. Funahashi, C. J. Bruce, and P. S. Goldman-Rakic. Mnemonic Coding of Visual Space in the Monkeys Dorsolateral Prefrontal Cortex. *Journal of Neurophysiology*, 61(2):331–349, 1989.

- [65] J. Fuster. Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *Journal of Neurophysiology*, 36(1):61, 1973.
- [66] J. M. Fuster. *Memory in the cerebral cortex*. MIT Press, Cambridge, 1995.
- [67] J. M. Fuster. *The prefrontal cortex*. Lippincott-Raven, Philadelphia, 3rd edition, 2001.
- [68] J. M. Fuster. The prefrontal cortex—an update: time is of the essence. *Neuron*, 30(2):319–33, May 2001.
- [69] J. M. Fuster. *Cortex and Mind: Unifying Cognition*. Oxford University Press, 2003.
- [70] J. M. Fuster and G. E. Alexander. Neuron Activity Related to Short-Term Memory. *Science*, 173(3997):652–654, 1971.
- [71] J. M. Fuster and R. H. Bauer. Visual short-term memory deficit from hypothermia of frontal cortex. *Brain research*, 81(3):393–400, Dec. 1974.
- [72] J. M. Fuster, R. H. Bauer, and J. P. Jervey. Functional interactions between inferotemporal and prefrontal cortex in a cognitive task. *Brain research*, 330(2):299–307, Mar. 1985.
- [73] J. M. Fuster, M. Bodner, and J. K. Kroger. Cross-modal and cross-temporal association in neurons of frontal cortex. *Nature*, 405(May):347–351, 2000.
- [74] J. M. Fuster and J. P. Jervey. Neuronal firing in the inferotemporal cortex of the monkey in a visual memory task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 2(3):361–75, Mar. 1982.
- [75] M. Galarreta and S. Hestrin. Frequency-dependent synaptic depression and the balance of excitation and inhibition in the neocortex. *Nature neuroscience*, 1(7):587–94, Nov. 1998.
- [76] M. Galarreta and S. Hestrin. Burst firing induces a rebound of synaptic strength at unitary neocortical synapses. *Journal of neurophysiology*, 83(1):621–4, Jan. 2000.
- [77] M.-O. Gewaltig and M. Diesmann. NEST (Neural Simulation Tool). *Scholarpedia*, 2(4):1430, 2007.
- [78] M. Gilson, A. N. Burkitt, D. B. Grayden, D. a. Thomas, and J. L. van Hemmen. Emergence of network structure due to spike-timing-dependent plasticity in recurrent neuronal networks IV: structuring synaptic pathways among recurrent connections. *Biological cybernetics*, 101(5-6):427–44, Dec. 2009.
- [79] J. W. Gnadt and R. A. Andersen. Memory related motor planning activity in posterior parietal cortex of macaque. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 70(1):216–20, Jan. 1988.

- [80] D. Golomb, N. Rubin, and H. Sompolinsky. Willshaw model: Associative memory with sparse coding and low firing rates. *Physical review. A*, 41(4):1843–1854, Feb. 1990.
- [81] A. M. Graybiel. The basal ganglia: learning new tricks and loving it. *Current opinion in neurobiology*, 15(6):638–44, Dec. 2005.
- [82] A. M. Graybiel. Habits, rituals, and the evaluative brain. *Annual review of neuroscience*, 31:359–87, Jan. 2008.
- [83] A. Grippo, L. Pelosi, V. Mehta, and L. D. Blumhardt. Working memory in temporal lobe epilepsy: an event-related potential study. *Electroencephalography and clinical neurophysiology*, 99(3):200–13, Sept. 1996.
- [84] R. Gütiğ, R. Aharonov, S. Rotter, and H. Sompolinsky. Learning input correlations through nonlinear temporally asymmetric Hebbian plasticity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 23(9):3697–714, May 2003.
- [85] B. S. Gutkin, C. R. Laing, C. L. Colby, C. C. Chow, and G. B. Ermentrout. Turning On and Off with Excitation : The Role of Spike-Timing Asynchrony and Synchrony in Sustained Neural Activity. *Journal of Computational Neuroscience*, 11:121–134, 2001.
- [86] I. Guyon, L. Personnaz, J. Nadal, and G. Dreyfus. Storage and retrieval of complex sequences in neural networks. *Physical Review A*, 38(12):6365–72, 1988.
- [87] B. Haider, A. Duque, A. R. Hasenstaub, and D. A. McCormick. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 26(17):4535–45, Apr. 2006.
- [88] N. Hájos, J. Pálhalmi, E. O. Mann, B. Németh, O. Paulsen, and T. F. Freund. Spike timing of distinct types of GABAergic interneuron during hippocampal gamma oscillations in vitro. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 24(41):9127–37, Oct. 2004.
- [89] A. Hanuschkin, M. Diesmann, and A. Morrison. A reafferent model of song syntax generation in the Bengalese finch. *BMC Neuroscience*, 11(Suppl 1):P33, 2010.
- [90] M. E. Hasselmo. Temporally structured replay of neural activity in a model of entorhinal cortex, hippocampus and postsubiculum. *The European journal of neuroscience*, 28(7):1301–15, Oct. 2008.
- [91] M. E. Hasselmo. A model of episodic memory: mental time travel along encoded trajectories using grid cells. *Neurobiology of learning and memory*, 92(4):559–73, Nov. 2009.

- [92] T. Hayashi, T. Ichiyama, M. Nishikawa, H. Isumi, and S. Furukawa. Pocket Monsters, a popular television cartoon, attacks Japanese children. *Annals of neurology*, 44(3):427–8, Sept. 1998.
- [93] D. O. Hebb. *The Organization of Behavior: a neuropsychological theory*. Wiley, New York, 1949.
- [94] C. M. Hempel, K. H. Hartman, X. J. Wang, G. G. Turrigiano, and S. B. Nelson. Multiple forms of short-term plasticity at excitatory synapses in rat medial prefrontal cortex. *Journal of neurophysiology*, 83(5):3031–41, May 2000.
- [95] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA*, 79(April):2554–2558, 1982.
- [96] F. Hoppensteadt and E. M. Izhikevich. *Canonical neural models*. MIT Press, Cambridge, 2002.
- [97] D. Horn, N. Levy, and E. Ruppin. Distributed synchrony in an attractor of spiking neurons. *Neurocomputing*, 32-33:409–414, June 2000.
- [98] D. Horn and M. Usher. Oscillatory Model of Short Term Memory. In J. Moody, editor, *Advances in Neural Information Processing Systems, Vol. 4*. Morgan Kaufmann, 1992.
- [99] N. K. Horst and M. Laubach. The role of rat dorsomedial prefrontal cortex in spatial working memory. *Neuroscience*, 164(2):444–56, Dec. 2009.
- [100] J. R. Hughes, Y. Daaboul, J. J. Fino, and G. L. Shaw. The "Mozart effect" on epileptiform activity. *Clinical EEG (electroencephalography)*, 29(3):109–19, July 1998.
- [101] J. R. Hughes and J. J. Fino. The Mozart effect: distinctive aspects of the music—a clue to brain coding? *Clinical EEG (electroencephalography)*, 31(2):94–103, Apr. 2000.
- [102] Y. Ikegaya, G. Aaron, R. Cossart, D. Aronov, I. Lampl, D. Ferster, and R. Yuste. Synfire chains and cortical songs: temporal modules of cortical activity. *Science (New York, N.Y.)*, 304(5670):559–64, Apr. 2004.
- [103] E. M. Izhikevich. Polychronization: computation with spikes. *Neural computation*, 18(2):245–82, Feb. 2006.
- [104] A. Jackson, J. Mavoori, and E. E. Fetz. Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, 444(7115):56–60, Nov. 2006.
- [105] O. Jensen. Maintenance of multiple working memory items by temporal segmentation. *Neuroscience*, 139(1):237–49, Apr. 2006.
- [106] B. C. Jobst. Electrical stimulation in epilepsy: vagus nerve and brain stimulation. *Current treatment options in neurology*, 12(5):443–53, Sept. 2010.

- [107] D. Johnston and T. H. Brown. Giant synaptic potential hypothesis for epileptiform activity. *Science (New York, N.Y.)*, 211(4479):294–7, Jan. 1981.
- [108] M. Jordan. Attractor dynamics and parallelism in a connectionist sequential machine. In *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, pages 531–546, Hillsdale, NJ, 1986. Lawrence Erlbaum Associates, Inc.
- [109] J. K. Jun and D. Z. Jin. Development of neural circuitry for precise temporal sequences through spontaneous activity, axon remodeling, and synaptic plasticity. *PLoS ONE*, 2(1):e723, Jan. 2007.
- [110] A. L. Kaas, H. van Mier, and R. Goebel. The neural correlates of human working memory for haptically explored object orientations. *Cerebral cortex (New York, N.Y. : 1991)*, 17(7):1637–49, July 2007.
- [111] J. Karbowski and G. Ermentrout. Synchrony arising from a balanced synaptic plasticity in a network of heterogeneous neural oscillators. *Physical Review E*, 65(3):1–5, Feb. 2002.
- [112] U. R. Karmarkar and D. V. Buonomano. Timing in the absence of clocks: encoding time in neural network states. *Neuron*, 53(3):427–38, Feb. 2007.
- [113] D. Kleinfeld and H. Sompolinsky. Associative Neural Network Model for the Generation of Temporal Patterns: Theory and Application to Central Pattern Generators. *Biophysical Journal*, 54:1039–1051, 1988.
- [114] K. W. Koch and J. M. Fuster. Unit activity in monkey parietal cortex related to haptic perception and temporary memory. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale*, 76(2):292–306, Jan. 1989.
- [115] M. J. Koepp. Juvenile myoclonic epilepsy—a generalized epilepsy syndrome? *Acta neurologica Scandinavica. Supplementum*, 181:57–62, Jan. 2005.
- [116] Y. Ku, S. Ohara, L. Wang, F. A. Lenz, S. S. Hsiao, M. Bodner, B. Hong, and Y.-D. Zhou. Prefrontal cortex and somatosensory cortex in tactile crossmodal association: an independent component analysis of ERP recordings. *PloS one*, 2(1):e771, Jan. 2007.
- [117] S. Kunkel, M. Diesmann, and A. Morrison. Limits to the development of feed-forward structures in large recurrent neuronal networks. *Submitted*, 2010.
- [118] D. Labar, J. Murphy, and E. Tecoma. Vagus nerve stimulation for medication-resistant generalized epilepsy. E04 VNS Study Group. *Neurology*, 52(7):1510–2, Apr. 1999.
- [119] N. Lahiri and J. S. Duncan. The Mozart effect: encore. *Epilepsy & behavior : E&B*, 11(1):152–3, Aug. 2007.
- [120] C. R. Laing and C. C. Chow. Stationary bumps in networks of spiking neurons. *Neural computation*, 13(7):1473–94, July 2001.

- [121] R. Larter, B. Speelman, and R. M. Worth. A coupled ordinary differential equation lattice model for the simulation of epileptic seizures. *Chaos (Woodbury, N.Y.)*, 9(3):795–804, Sept. 1999.
- [122] P. E. Latham and S. Nirenberg. Computing and stability in cortical networks. *Neural computation*, 16(7):1385–412, July 2004.
- [123] W.-L. Lee and J. J. Hablitz. Involvement of non-NMDA receptors in picrotoxin-induced epileptiform activity in the hippocampus. *Neuroscience Letters*, 107(1-3):129–134, 1989.
- [124] J. Lisman and G. Buzsáki. A neural coding scheme formed by the combined function of gamma and theta oscillations. *Schizophrenia bulletin*, 34(5):974–80, Sept. 2008.
- [125] J. Lisman and a. D. Redish. Prediction, sequences and the hippocampus. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364(1521):1193–201, May 2009.
- [126] J. E. Lisman, J. M. Fellous, and X. J. Wang. A role for NMDA-receptor channels in working memory. *Nature neuroscience*, 1(4):273–5, Aug. 1998.
- [127] J. E. Lisman and M. a. Idiart. Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science (New York, N.Y.)*, 267(5203):1512–5, Mar. 1995.
- [128] J. K. Liu and D. V. Buonomano. 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, Oct. 2009.
- [129] F. H. Lopes da Silva, W. Blanes, S. N. Kalitzin, J. Parra, P. Suffczynski, and D. N. Velis. Dynamical diseases of brain systems: different routes to epileptic seizures. *IEEE transactions on bio-medical engineering*, 50(5):540–8, May 2003.
- [130] A. Luczak, P. Barthó, and K. D. Harris. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*, 62(3):413–25, May 2009.
- [131] A. Luczak, P. Barthó, S. L. Marguet, G. Buzsáki, and K. D. Harris. Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences of the United States of America*, 104(1):347–52, Jan. 2007.
- [132] J. R. Manns and H. Eichenbaum. A cognitive map for object memory in the hippocampus. *Learning & memory (Cold Spring Harbor, N.Y.)*, 16(10):616–24, Jan. 2009.
- [133] E. Marder, L. F. Abbott, G. G. Turrigiano, Z. Liu, and J. Golowasch. Memory from the dynamics of intrinsic membrane currents. *Proceedings of the National Academy of Sciences of the United States of America*, 93(24):13481–6, Nov. 1996.
- [134] H. Markram, J. Lubke, M. Frotscher, and B. Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275(5297):213–215, 1997.

- [135] V. Meskenaite. Calretinin-immunoreactive local circuit neurons in area 17 of the cynomolgus monkey, *Macaca fascicularis*. *The Journal of comparative neurology*, 379(1):113–32, Mar. 1997.
- [136] G. Miesenböck. The optogenetic catechism. *Science (New York, N.Y.)*, 326(5951):395–9, Oct. 2009.
- [137] R. Miles, J.-C. Ponce, D. Fricker, and X. Leinekugel. The birth (and adolescence) of LTP. *The Journal of physiology*, 568(Pt 1):1–2, Oct. 2005.
- [138] E. K. Miller and J. D. Cohen. An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, 24:167–202, Jan. 2001.
- [139] E. K. Miller, C. A. Erickson, and R. Desimone. Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 16(16):5154–5167, Aug. 1996.
- [140] E. K. Miller, L. Li, and R. Desimone. Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(4):1460–78, Apr. 1993.
- [141] A. Minai and W. Levy. Sequence learning in a single trial. In *INNS World Congress Neural Networks II*, volume 2, pages 505–508, 1993.
- [142] G. Mongillo, O. Barak, and M. Tsodyks. Synaptic theory of working memory. *Science (New York, N.Y.)*, 319(5869):1543–6, Mar. 2008.
- [143] G. L. Morris and W. M. Mueller. Long-term treatment with vagus nerve stimulation in patients with refractory epilepsy. The Vagus Nerve Stimulation Study Group E01-E05. *Neurology*, 53(8):1731–5, Nov. 1999.
- [144] A. Morrison, A. Aertsen, and M. Diesmann. Spike-timing-dependent plasticity in balanced random networks. *Neural computation*, 19(6):1437–67, June 2007.
- [145] A. Morrison, M. Diesmann, and W. Gerstner. Phenomenological models of synaptic plasticity based on spike timing. *Biological cybernetics*, 98(6):459–78, June 2008.
- [146] V. B. Mountcastle. The columnar organization of the neocortex. *Brain : a journal of neurology*, 120:701–22, Apr. 1997.
- [147] T. Muftuler, M. Bodner, G. L. Shaw, and O. Nalcioglu. fMRI study to investigate spatial correlates of music listening and spatial-temporal reasoning. In *12th Annual Meeting of the International Society of Magnetic Resonance in Medicine*, 2004.
- [148] Z. Nádasdy, H. Hirase, a. Czurkó, J. Csicsvari, and G. Buzsáki. Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 19(21):9497–507, Nov. 1999.

- [149] K. Nakamura, A. Mikami, and K. Kubota. Oscillatory neuronal activity related to visual short-term memory in monkey temporal pole. *Neuroreport*, 3(1):117–120, 1992.
- [150] S. B. Nelson and G. G. Turrigiano. Synaptic depression: a key player in the cortical balancing act. *Nature neuroscience*, 1(7):539–41, Nov. 1998.
- [151] T. I. Netoff and S. J. Schiff. Decreased neuronal synchronization during experimental seizures. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(16):7297–7307, 2002.
- [152] E. Nordlie, M.-o. Gewaltig, and H. E. Plesser. Towards Reproducible Descriptions of Neuronal Network Models. *PLoS Computational Biology*, 5(8), 2009.
- [153] H. Okamoto, Y. Isomura, M. Takada, and T. Fukai. Temporal integration by stochastic recurrent network dynamics with bimodal neurons. *Journal of neurophysiology*, 97(6):3859–67, June 2007.
- [154] R. C. O’Reilly, S. a. Herd, and W. M. Pauli. Computational models of cognitive control. *Current opinion in neurobiology*, 20(2):257–61, Apr. 2010.
- [155] E. Pastalkova, V. Itskov, A. Amarasingham, and G. Buzsáki. Internally generated cell assembly sequences in the rat hippocampus. *Science (New York, N.Y.)*, 321(5894):1322–7, Sept. 2008.
- [156] B. Pesaran, J. S. Pezaris, M. Sahani, P. P. Mitra, and R. a. Andersen. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature neuroscience*, 5(8):805–11, Aug. 2002.
- [157] M. Petrides. The role of the mid-dorsolateral prefrontal cortex in working memory. *Experimental brain research*, 133(1):44–54, July 2000.
- [158] Y. Prut, E. Vaadia, H. Bergman, I. Haalman, H. Slovin, and M. Abeles. Spatiotemporal structure of cortical activity: properties and behavioral relevance. *Journal of neurophysiology*, 79(6):2857–74, June 1998.
- [159] R. Pumain, C. Menini, U. Heinemann, J. Louvel, and C. Silva-Barrat. Chemical synaptic transmission is not necessary for epileptic seizures to persist in the baboon *Papio papio*. *Experimental neurology*, 89(1):250–8, July 1985.
- [160] J. Quintana and J. M. Fuster. From perception to action: temporal integrative functions of prefrontal and parietal neurons. *Cerebral cortex (New York, N.Y. : 1991)*, 9(3):213–21, 1999.
- [161] A. Renart, R. Moreno-Bote, X.-J. Wang, and N. Parga. Mean-driven and fluctuation-driven persistent activity in recurrent networks. *Neural computation*, 19(1):1–46, Jan. 2007.

- [162] P. A. Robinson, C. J. Rennie, and D. L. Rowe. Dynamics of large-scale brain activity in normal arousal states and epileptic seizures. *Physical review. E, Statistical, nonlinear, and soft matter physics*, 65(4 Pt 1):041924, Apr. 2002.
- [163] E. T. Rolls. Attractor networks. *Wiley Interdisciplinary Reviews: Cognitive Science*, 1(1):119–134, Jan. 2010.
- [164] R. Romo, C. D. Brody, A. Hernández, and L. Lemus. Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature*, 399(6735):470–3, June 1999.
- [165] Y. Roudi and P. E. Latham. A balanced memory network. *PLoS computational biology*, 3(9):1679–700, Sept. 2007.
- [166] N. Ruben and H. Sompolinsky. Neural networks with low local firing rates. *Europhysics Letters*, 10:465–470, 1989.
- [167] J. Rubin, D. Lee, and H. Sompolinsky. Equilibrium Properties of Temporally Asymmetric Hebbian Plasticity. *Physical Review Letters*, 86(2):364–367, Jan. 2001.
- [168] T. Sawaguchi and M. Iba. Prefrontal cortical representation of visuospatial working memory in monkeys examined by local inactivation with muscimol. *Journal of neurophysiology*, 86(4):2041–2053, 2001.
- [169] S. J. Schiff. Forecasting brain storms. *Nature medicine*, 4(10):1117–8, Oct. 1998.
- [170] M. Shafi. *Working memory: patterns variability and computational modelling*. Ph.d. dissertation, UCLA, 2005.
- [171] M. Shafi, Y. Zhou, J. Quintana, C. Chow, J. Fuster, and M. Bodner. Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*, 146(3):1082–108, May 2007.
- [172] M. Shapiro. Plasticity, hippocampal place cells, and cognitive maps. *Archives of neurology*, 58(6):874–81, June 2001.
- [173] G. Shaw, D. Silverman, and Pearson JC. Model of cortical organization embodying a basis for a theory of information processing and memory recall. *Proc. Natl. Acad. Sci. USA*, 82:2364–2368, 1985.
- [174] G. L. Shaw and M. Bodner. *Mountcastle principle of columnar cortex as basis for theory of higher brain function: Clinical relevance*, pages 187–203. Nova, New York, 2005.
- [175] Y. Shu, A. Hasenstaub, and D. A. McCormick. Turning on and off recurrent balanced cortical activity. *Nature*, 423(6937):288–93, May 2003.

- [176] R. Singh and C. Eliasmith. Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 26(14):3667–78, Apr. 2006.
- [177] H. Sompolinsky and I. Kanter. Temporal Association in Asymmetric Neural Networks. *Physical Review Letters*, 57(22):2861–2864, 1986.
- [178] S. Song, K. D. Miller, and L. F. Abbott. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–26, Sept. 2000.
- [179] S. Song, P. J. Sjöström, M. Reigl, S. Nelson, and D. B. Chklovskii. Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, 3(3):e68, Mar. 2005.
- [180] M. C. Stoeckel, B. Weder, F. Binkofski, G. Buccino, N. J. Shah, and R. J. Seitz. A fronto-parietal circuit for tactile object discrimination: an event-related fMRI study. *NeuroImage*, 19(3):1103–14, July 2003.
- [181] R. Sun and C. Giles. Sequence learning: from recognition and prediction to sequential decision making. *IEEE Intelligent Systems*, 16(4):67–70, July 2001.
- [182] R. E. Suri and T. J. Sejnowski. Spike propagation synchronized by temporally asymmetric Hebbian learning. *Biological cybernetics*, 87(5-6):440–5, Dec. 2002.
- [183] D. Sussillo and L. F. Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–57, Aug. 2009.
- [184] B. Szatmáry and E. M. Izhikevich. Spike-Timing Theory of Working Memory. *PLoS Computational Biology*, 6(8):e1000879, Aug. 2010.
- [185] C. Tallon-Baudry, O. Bertrand, and C. Fischer. Oscillatory synchrony between human extrastriate areas during visual short-term memory maintenance. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 21(20):RC177, Oct. 2001.
- [186] A. Tang, D. Jackson, J. Hobbs, W. Chen, J. L. Smith, H. Patel, A. Prieto, D. Petrusca, M. I. Grivich, A. Sher, P. Hottowy, W. Dabrowski, A. M. Litke, and J. M. Beggs. A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *The Journal of Neuroscience: the official journal of the Society for Neuroscience*, 28(2):505–18, Jan. 2008.
- [187] T. C. Thiagarajan, M. A. Lebedev, M. A. Nicolelis, and D. Plenz. Coherence potentials: loss-less, all-or-none network events in the cortex. *PLoS biology*, 8(1):e1000278, Jan. 2010.
- [188] I. Timofeev and M. Steriade. Neocortical seizures: initiation, development and cessation. *Neuroscience*, 123(2):299–336, Jan. 2004.

- [189] S. Tobimatsu, Y. M. Zhang, Y. Tomoda, A. Mitsudome, and M. Kato. Chromatic sensitive epilepsy: a variant of photosensitive epilepsy. *Annals of neurology*, 45(6):790–3, June 1999.
- [190] R. D. Traub and R. Miles. *Neuronal Networks of the Hippocampus*. Cambridge University Press, New York, 1991.
- [191] R. D. Traub, R. Miles, and J. G. R. Jefferys. Synaptic and intrinsic conductances shape picrotoxin-induced synchronized after-discharges in the guinea-pig hippocampal slice. *Journal of Physiology*, 461:525–547, 1993.
- [192] R. D. Traub and R. K. S. Wong. Cellular mechanism of neuronal synchronization in epilepsy. *Science*, 216(4547):745–747, 1982.
- [193] F. H. Treitz, I. Daum, P. M. Faustmann, and C. G. Haase. Executive deficits in generalized and extrafrontal partial epilepsy: long versus short seizure-free periods. *Epilepsy & behavior : E&B*, 14(1):66–70, Jan. 2009.
- [194] R. P. Turner. The acute effect of music on interictal epileptiform discharges. *Epilepsy & behavior : E&B*, 5(5):662–8, Oct. 2004.
- [195] W. Van Drongelen, H. Koch, C. Marcuccilli, F. Pena, and J.-M. Ramirez. Synchrony levels during evoked seizure-like bursts in mouse neocortical slices. *Journal of neurophysiology*, 90(3):1571–80, Sept. 2003.
- [196] W. van Drongelen, H. C. Lee, M. Hereld, Z. Chen, F. P. Elsen, and R. L. Stevens. Emergent epileptiform activity in neural networks with weak excitatory synapses. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 13(2):236–41, June 2005.
- [197] M. C. W. van Rossum, G. G. Turrigiano, and S. B. Nelson. Fast propagation of firing rates through layered networks of noisy neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(5):1956–66, Mar. 2002.
- [198] F. Velasco, M. Velasco, F. Jiménez, A. L. Velasco, F. Brito, M. Rise, and J. D. Carrillo-Ruiz. Predictors in the treatment of difficult-to-control seizures by electrical stimulation of the centromedian thalamic nucleus. *Neurosurgery*, 47(2):295–304; discussion 304–5, Aug. 2000.
- [199] F. Velasco, M. Velasco, A. L. Velasco, F. Jimenez, I. Marquez, and M. Rise. Electrical stimulation of the centromedian thalamic nucleus in control of seizures: long-term studies. *Epilepsia*, 36(1):63–71, Jan. 1995.
- [200] S. Verduzco-Flores, M. Bodner, B. Ermentrout, J. M. Fuster, and Y. Zhou. Working memory cells’ behavior may be explained by cross-regional networks with synaptic facilitation. *PLoS one*, 4(8):e6399, Jan. 2009.

- [201] T. P. Vogels and L. F. Abbott. Signal propagation and logic gating in networks of integrate-and-fire neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 25(46):10786–95, Nov. 2005.
- [202] L. Wang. Multi-associative neural networks and their applications to learning and retrieving complex spatio-temporal sequences. *IEEE transactions on systems, man, and cybernetics. Part B, Cybernetics : a publication of the IEEE Systems, Man, and Cybernetics Society*, 29(1):73–82, Jan. 1999.
- [203] X. J. Wang. Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 19(21):9587–9603, Nov. 1999.
- [204] X. J. Wang. Synaptic reverberation underlying mnemonic persistent activity. *Trends in neurosciences*, 24(8):455–63, Aug. 2001.
- [205] X.-J. Wang. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5):955–68, Dec. 2002.
- [206] Y. Wang, H. Markram, P. H. Goodman, T. K. Berger, J. Ma, and P. S. Goldman-Rakic. Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature neuroscience*, 9(4):534–42, Apr. 2006.
- [207] F. Wendling, F. Bartolomei, J. J. Bellanger, and P. Chauvel. Epileptic fast activity can be explained by a model of impaired GABAergic dendritic inhibition. *European Journal of Neuroscience*, 15(9):1499–1508, May 2002.
- [208] H. R. Wilson and J. D. Cowan. Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical journal*, 12(1):1–24, Jan. 1972.
- [209] L. Y. Xue and A. L. Ritaccio. Reflex seizures and reflex epilepsy. *American journal of electroneurodiagnostic technology*, 46(1):39–48, Mar. 2006.
- [210] Y. Yamashita, M. Takahashi, T. Okumura, M. Ikebuchi, H. Yamada, M. Suzuki, K. Okanoya, and J. Tani. Developmental learning of complex syntactical song in the Bengalese finch: a neural network model. *Neural networks : the official journal of the International Neural Network Society*, 21(9):1224–31, Nov. 2008.
- [211] Y. Yoshimura, J. L. M. Dantzker, and E. M. Callaway. Excitatory cortical neurons form fine-scale functional networks. *Nature*, 433(February), 2005.
- [212] M. Yoshioka, S. Scarpetta, and M. Marinaro. Spike-Timing-Dependent Synaptic Plasticity to Learn Spatiotemporal Patterns in Recurrent Neural Networks. In M. D. S. E. A. J, editor, *ICANN 2007, Part I, LNCS 4668*, volume 1, pages 757–766, Berlin Heidelberg, 2007. Springer-Verlag.

- [213] Y. D. Zhou and J. M. Fuster. Mnemonic neuronal activity in somatosensory cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 93(19):10533–7, Sept. 1996.
- [214] D. Zipser, B. Kehoe, G. Littlewort, and J. Fuster. A spiking network model of short-term active memory. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(8):3406–20, Aug. 1993.