

**Realistic computer modeling of the
mammalian olfactory cortex**

Thesis by

Michael Christopher Vanier

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2001

(Defended May 21, 2001)

©2001

Michael Christopher Vanier

All Rights Reserved

Acknowledgements

I have been at Caltech as a graduate student for a long time, so I have a correspondingly long list of people to thank:

Jim Bower, my supervisor, for giving me guidance about all aspects of science (the official and especially the unofficial), for being a friend, for singing in my rock band (which he mistakenly thought was his rock band), for his patience, support and encouragement even when I was veering wildly away from the official subject of my thesis work, and for demonstrating that a scientist doesn't have to be a boring person.

My Ph.D. committee members for their support, their interest, and their encouragement. I've always felt I had the best Ph.D. committee a student could ask for.

Everyone in the Bower lab for their support, friendship and encouragement.

Dave Beeman and everyone else in the GENESIS project for being a stimulating and fun group of people to be associated with.

Dave Bilitch, for teaching me an enormous amount about computer programming through a mixture of good advice and bad example.

Upi Bhalla, for interesting discussions, for supplying me with his olfactory bulb data and being willing to answer my many questions about it, for his work on GENESIS, for his sense of humor, and for being an all-around fine person. Science could use more people like you.

Lew Haberly, who I've never met, for his fine experimental work on the piriform cortex

on which much of the work described in this thesis is built. As a piriform cortex modeler, I am very much indebted to him.

Matt Wilson, for building the first version of GENESIS and for building the first version of the piriform cortex model. I can't believe how much Matt accomplished given the limited amount of data that was then available. Matt's modeling approach, especially his approach to handling scaling issues in network models, was a major influence on my model.

Alex Protopapas, for providing me with most of the intracellular data that formed the basis of this model and was an integral component of my two methodological papers. Working with Alex was a pleasure.

Karim Elaagouby, for helping me get my brain-slice electrophysiology setup working against all odds. Without Karim's help, I might never have finished my experimental work.

Dave Kewley, for sysadmin services above and beyond the call of duty. In particular, thanks to Dave for last-minute color printer fixes under desperate circumstances.

Also, on a more personal level, I'd like to thank:

My mother, Anna Dodge Vanier, for putting up with me and my somewhat rocky path towards personal growth. Mom has been and continues to be an unflinching source of love, support and inspiration even in the most difficult times of this project.

Jan Tivol, whose love and friendship did more than anyone else to keep me sane during the last few years. You're the best!

Claire Sergeant, for long-distance encouragement and for being a sounding board for my comedy routines in progress. We discussed my work, we laughed about it, and then Hemingway punched us both in the nose.

Kayla Smith, for friendship and for being my new comedy buddy.

My friends Ernesto “Boom Boom” Soares, Alfredo “Tom Jones” Fontanini, and Robert “Mr. Sensitive” Sneddon, for being fun to hang out with. Thanks to Ernesto for getting me involved in the Babe and Ricky’s blues scene (see below), to Robert for lots of politically incorrect conversations, and to Alfredo for teaching me to swear in Italian and for sharing my affinity for cheesy lounge music (don’t laugh; it’s not unusual...). I especially want to thank Alfredo for many useful conversations and for proofreading large parts of this thesis.

Shana Coates for making my second summer in Woods Hole very memorable indeed :-)
Because of Shana I will probably have a fetish for women with English accents for the rest of my life.

Melanie Walker, for her sense of humor, for being different, and for not taking herself too seriously.

The folks at Air Adventures West in Taft, California, for teaching me to skydive. I can’t possibly hope to repay that; it changed my life.

The folks at TACIT (Theater Arts at the California Institute of Technology), for teaching me to act, a very useful skill for a Caltech grad student.

Mama and the folks at Babe and Ricky’s, for letting me play the blues in a real blues

bar.

Don and Wendy Caldwell and the Caltech Chamber Singers, for teaching me to sing, for allowing me to play Dr. Harold Hill (a dream come true!) and for some great times.

Jim Boyk, for teaching me to play piano and for being an all-around amazing guy. I have met many cool people at Caltech but Jim is the coolest. If I ever make a lot of money, Jim will be set for life; that's a promise!

Linda Chappell, for helping me with housing problems while I was at Caltech. Thanks, Linda!

Richard Stallman and the GNU project, for all the great software and inspiration, and for letting me be a (very minor) part of it all.

Linus Torvalds and the Linux project, for a great operating system and for helping to rid the world of Microsoft.

Mike Hucka, hacker and sysadmin extraordinaire, for being a great friend and for being my überhacker. I'm so glad I recruited you :-)

Dedication

I dedicate this thesis to the memory of my father, Andre Guy Vanier. His love and his sense of humor will be with me always.

Abstract

A combination of experimental and computer modeling techniques were used to investigate the dynamics and computational functions of the rat olfactory (piriform) cortex. Experimental characterization of synaptic responses to afferent and associational fiber voltage shocks were performed, in the presence and absence of the neuromodulator norepinephrine. This data was used to generate computer models of synaptic transmission in piriform cortex. Models of pyramidal neurons and feedback inhibitory interneurons were constructed which accurately match intracellular experimental data in the presence and absence of norepinephrine. In order to achieve this, parameter search tools for automatically matching computer models of neurons to data were developed. Models of feedforward inhibitory interneurons were also constructed. An abstract spike generating model of the olfactory bulb was built. These components were combined to create a realistic computer model of the piriform cortex. This model can accurately replicate the response of the real system to a strong shock stimulus, as reflected in current source density plots. Two versions of the model were created to model the oscillatory response of the system to weak shocks. The first model replicates the surface field potential with considerable accuracy, but fails to replicate the current source density data. The second model replicates the current source density data and suggests a new organizing principle for the olfactory system based on non-overlapping neuronal groups. This hypothesis is experimentally testable.

Contents

I	Introduction	1
1.1	Overview of the thesis	3
1.2	The mammalian olfactory system	5
1.2.1	The olfactory epithelium and olfactory bulb	5
1.2.2	Primary olfactory (piriform) cortex	7
1.2.3	Oscillations in the olfactory system	11
1.3	Modeling the piriform cortex	16
1.3.1	Why build realistic models of piriform cortex?	16
1.3.2	Realistic modeling	19
1.3.3	The previous piriform cortex model	20
1.3.4	New model features	20
1.3.5	Tuning the model	22
1.4	Questions the model can help answer	23
1.4.1	Coding in the olfactory system	23

1.4.2	Origin and functions of oscillations	25
1.4.3	Neuromodulation in the piriform cortex	26
1.4.4	Other issues	26
1.5	Modeling methodologies	27
1.5.1	Simulation environment	27
1.5.2	Parameter searching	27
1.5.3	Bayesian methods	28
1.6	New work suggested by the model	29
1.6.1	Experimental work	29
1.6.2	Modeling work	30
1.7	Summary of thesis contributions	31
1.7.1	The new piriform cortex model	31
1.7.2	Lessons from the model	32
1.7.3	Experimental work	33
1.7.4	Development of simulation tools	33
II	Experimental studies	47
2	Synaptic Effects of Norepinephrine in Piriform Cortex	53
2.1	Abstract	53
2.2	Introduction	54

2.3	Materials and Methods	57
2.3.1	Preparation of brain slices	57
2.3.2	Preparation and application of pharmaceuticals	59
2.3.3	Electrophysiology	59
2.3.4	Data analysis	61
2.4	Results	63
2.4.1	Synaptic effects of norepinephrine	63
2.4.2	Concentration-dependence of effects	63
2.4.3	Pharmacology of effects	66
2.4.4	Effects of norepinephrine on paired-pulse facilitation	68
2.5	Discussion	70
2.5.1	Differential effects of NE on layer 1a and 1b field potentials	70
2.5.2	Pharmacological basis for the NE effects in layer 1a and 1b	72
2.5.3	Differences from previously reported results	74
2.5.4	Functional significance	75
III	Matching neural models to experimental data	91
3	A Comparative Survey of Automated Parameter Search Methods	
	for Compartmental Neural Models	95
3.1	Abstract	95

3.2	Introduction	96
3.3	Methods	99
3.3.1	Simulation environment and data analysis tools	99
3.3.2	Models	100
3.3.3	Target data sets	106
3.3.4	Computing the match between data and models	107
3.3.5	Parameter search methods	109
3.3.6	Statistical analyses	116
3.3.7	Resource requirements	117
3.4	Results	118
3.4.1	Simple spiking models	118
3.4.2	Passive cable models	121
3.4.3	Pyramidal cell model	127
3.5	Discussion	134
3.5.1	Method comparisons	134
3.5.2	Differences between the search algorithms	137
3.5.3	Matching functions	140
3.5.4	Variation in the parameter sets	141
3.5.5	Robustness of the parameter sets	143
3.5.6	Recommendations for effective parameter searching	144

3.5.7	Limitations of parameter searching techniques	145
3.6	Acknowledgements	149
3.7	Appendix: Model descriptions	157
4	On the Use of Bayesian Methods for Evaluating Compartmental	
	Neural Models	159
4.1	Abstract	159
4.2	Introduction	160
4.3	Bayesian inference	166
4.4	Bayesian compartmental modeling	171
4.4.1	Prior probabilities for compartmental models	171
4.4.2	Likelihoods for compartmental models	174
4.5	Examples	187
4.5.1	Objectives and methods	187
4.5.2	The data	189
4.5.3	The models	192
4.5.4	Parameter estimation: determining optimal parameter values	200
4.5.5	Comparing individual models from the same class	202
4.5.6	Comparing individual models from different classes	205
4.5.7	Comparing entire model classes	209

4.5.8	Estimating the optimal amount of noise for a noisy pyramidal cell model	213
4.6	Discussion	223
IV	Modeling the piriform cortex	247
5	Building Models of Single Neurons in Piriform Cortex	253
5.1	Piriform cortex neurons	253
5.2	Computer simulation environment	257
5.3	Modeling layer 2 pyramidal neurons	257
5.3.1	Data set	257
5.3.2	Simplifying the neuron morphology	258
5.3.3	Building the model	267
5.3.4	Tuning the model	267
5.3.5	Neuromodulation	269
5.4	Modeling inhibitory neurons	271
5.4.1	Data set	271
5.4.2	Building and tuning the models	273
5.4.3	Neuromodulation	273
5.5	Modeling synaptic inputs	276
5.5.1	Data sets	276

5.5.2	The basic synapse model	278
5.5.3	NMDA synapses	280
5.5.4	Synaptic facilitation and depression	281
5.5.5	Synaptic neuromodulation	287
5.6	Limitations of our approach	291
5.7	Appendix A: the neuron model simplification algorithm	293
5.8	Appendix B: model parameters	295
5.8.1	Pyramidal neuron channel model parameters	295
5.8.2	Pyramidal neuron model parameters	305
5.8.3	Feedback inhibitory interneuron model parameters	306
5.8.4	Feedforward interneuron model parameters	306
6	Building the Piriform Cortex Network Model	319
6.1	Modeling inputs from the olfactory bulb	319
6.1.1	Data set	320
6.1.2	Data analysis	322
6.1.3	Building the spike generating model	323
6.1.4	Validating the spike generating model	327
6.1.5	Conclusions	331
6.2	Modeling the piriform cortex network	332
6.2.1	Previous modeling work on piriform cortex	332

6.2.2	Construction of the model	335
6.2.3	Outputs from the cortical model	348
6.2.4	Parameterizing the model	350
6.2.5	Response to background inputs	354
6.2.6	Comparison with the Wilson/Bower model	356
6.2.7	Limitations of the model	363
7	Exploring the Network Model of Piriform Cortex	387
7.1	Response to olfactory bulb background inputs	389
7.1.1	Spontaneous activity	389
7.1.2	Spontaneous activity with NE	391
7.2	The strong shock response	392
7.2.1	Basic features of the response	392
7.2.2	Role of background excitation in the strong shock response	394
7.2.3	Role of feedforward inhibition	401
7.3	The weak shock response: model with random connectivities	404
7.3.1	Basic features of the response	406
7.3.2	Influence of feedforward/feedback interneurons	408
7.3.3	Influence of feedback inhibitory interneurons	412
7.3.4	CSD response	420
7.4	The weak shock response: model with structured connectivities	433

7.4.1	Structure of the model	433
7.4.2	Response to multiple input shocks	435
7.4.3	Discussion	446
8	Conclusions	457
8.1	Summary of simulation results	457
8.1.1	Strong shock response: influence of background synaptic inputs . . .	458
8.1.2	The strong shock response: effects of feedforward inhibition	460
8.1.3	Significance of the strong shock response	460
8.1.4	The weak shock response, oscillations, and network connectivity . .	461
8.2	Computational implications	463
8.3	Future directions	466
8.3.1	Experimental work	466
8.3.2	Modeling work	467
8.4	The value of computer modeling	469

List of Figures

1.1	The mammalian olfactory system.	6
1.2	Anatomy of the olfactory bulb.	9
1.3	Subdivisions of the piriform cortex.	10
1.4	Excitatory connections in the piriform cortex.	12
1.5	Piriform cortex neuron classes.	13
1.6	Morphology of a superficial pyramidal neuron.	14
1.7	Inhibitory connections in the piriform cortex.	15
1.8	Oscillations in the olfactory system.	17
2.1	Piriform cortex slice recording setup.	56
2.2	Basic effect of norepinephrine on synaptic transmission in piriform cortex.	64
2.3	Time-course of NE effects on synaptic transmission.	65
2.4	Dose-response curve of NE effect.	67
2.5	Pharmacology of NE effects.	69
2.6	Effects of NE on paired-pulse facilitation.	71

3.1	Comparison of parameter search methods on all models.	119
3.2	Parameter search results for model 2.	122
3.3	Search results as a function of simulation number for models 1 and 2. . . .	123
3.4	Parameter search results for model 4.	125
3.5	Search results as a function of simulation number for models 3 and 4. . . .	128
3.6	Parameter search results for model 5.	130
3.7	Search results as a function of simulation number for model 5.	131
3.8	Search results as a function of simulation number for modified model 5. . . .	135
3.9	Parameter search results for modified model 5.	136
4.1	Outputs of the target model with and without noise.	191
4.2	ISI histograms from a four-conductance model with noise.	197
4.3	Covariance matrix of a sample four-conductance model.	198
4.4	Histogram of ISI standard deviation values for a four-conductance model with noise.	199
4.5	Spike raster plots of target model versus best-matching models.	204
4.6	A plot of the sum of the negative log likelihoods versus the percent noise added to the model.	218
4.7	Raster plots of spiking times for a layer 2 pyramidal cell and spike probability profiles.	219

5.1	Piriform cortex neuron classes.	255
5.2	Pyramidal neuron response to injected currents.	259
5.3	Pyramidal neuron morphology.	261
5.4	Structure of the simplified pyramidal neuron model.	265
5.5	EPSPs in the full and the simplified pyramidal neuron model.	266
5.6	Pyramidal neuron membrane potential traces in response to current clamp inputs: data vs. simulation.	270
5.7	Pyramidal neuron membrane potential traces in response to current clamp inputs with NE: data vs. simulation.	272
5.8	Layer 3 inhibitory interneuron membrane potential traces in response to cur- rent clamp inputs: data vs. simulation.	274
5.9	f/I curves for layer 3 inhibitory interneurons.	275
5.10	f/I curves for layer 3 inhibitory interneurons: model vs. simulation.	277
5.11	Multipulse facilitation and depression in afferent and associational synapses.	283
5.12	Multipulse facilitation and depression in afferent and associational synapses for various interpulse intervals.	284
5.13	Multipulse facilitation and depression in the presence of NE.	285
5.14	Multipulse facilitation and depression in layer 1a synapses: model vs. data.	288
5.15	Multipulse facilitation and depression in layer 1b synapses: model vs. data.	289

5.16	Multipulse facilitation and depression in layer 1b synapses with $25 \mu M$ NE: model vs. data.	290
5.17	Activation (m_∞ vs. V) curves for all voltage-dependent ionic channels in the pyramidal neuron model.	298
5.18	Time constant curves for all voltage-dependent ionic channels in the pyrami- dal neuron model.	299
5.19	Activation (m_∞ vs. Ca concentration) and time constant (τ vs. Ca concen- tration) curves for calcium-dependent potassium channel in the pyramidal neuron model.	300
5.20	Activation (m_∞ vs. V) curves for voltage-dependent ionic channels in the feedback inhibitory interneuron model.	307
5.21	Time constant (τ vs. V) curves for voltage-dependent ionic channels in the feedback inhibitory interneuron model.	308
6.1	Variability in mitral cell responses to the same odor.	321
6.2	Interspike interval (ISI) distribution in mitral cells (real and simulated). . .	325
6.3	Background spiking responses of 100 simulated mitral cells.	326
6.4	Rate coding in the olfactory bulb spike generating model.	328
6.5	Synchrony coding in the olfactory bulb spike generating model.	329
6.6	Autocorrelogram of spike times in mitral cells.	330
6.7	Major subdivisions of the piriform cortex.	336

6.8	Arrangement of neuron types in the piriform cortex model.	339
6.9	Positions of neurons in the piriform cortex model in x - y coordinates.	340
6.10	Schematic diagram of connections in the piriform cortex model.	343
6.11	Strengths of afferent projections in the piriform cortex model.	344
6.12	Excitatory connections in the piriform cortex.	345
6.13	Transmembrane current loops in response to synaptic activation.	351
6.14	Voltage gradient in response to synaptic activation.	352
6.15	Background responses of pyramidal neurons in the piriform cortex model.	355
6.16	Surface CSD plots of strong shock response in piriform cortex: experiment and model.	357
6.17	Contour CSD plots of strong shock response of piriform cortex: experiment and model.	358
6.18	Pyramidal neuron soma responses to a strong shock input.	359
6.19	Surface field potential response of piriform cortex to a weak shock: experi- ment and model.	360
6.20	Pyramidal neuron soma responses to a weak shock input.	361
6.21	Interspike interval (ISI) distributions of the responses of a single mitral cell to different odors.	366
6.22	Interspike interval distributions of the responses of 18 mitral cells.	367
6.23	Pooled interspike interval distribution of all mitral cells across all odors.	368

6.24	Autocorrelograms of two mitral cells during presentation of several odors. . .	371
6.25	Autocorrelograms of different mitral cells.	372
6.26	PSTHs of mitral cells synchronized to odor onset but not to respiration onset.	374
6.27	PSTHs of mitral cells synchronized to respiration onset following odor onset.	375
6.28	PSTHs of a mitral cell showing firing rate changes after odor onset.	376
7.1	Piriform cortex model: background responses of pyramidal neurons.	390
7.2	Piriform cortex model: 2-D histograms of background firing rates in pyrami- dal neurons with NE application.	393
7.3	CSD surface profile of strong shock response in piriform cortex: experimental and model results.	395
7.4	CSD contour profile of strong shock response in piriform cortex: experimental and model results.	397
7.5	Piriform cortex model: network response to a strong shock stimulus.	399
7.6	Piriform cortex model: response of pyramidal neurons to a strong shock stimulus in superficial layer 1b dendrites.	400
7.7	Piriform cortex model: surface CSD profiles in response to a strong shock with and without background spiking inputs.	402
7.8	Piriform cortex model: 2-D histograms of pyramidal neuron spikes in re- sponse to a strong shock with and without background spiking inputs. . . .	403

7.9	Piriform cortex model without feedforward inhibition: response of pyramidal neuron somatic membrane potentials to a strong shock stimulus.	405
7.10	Surface field potential response of piriform cortex to a weak shock stimulus: experiment and model.	409
7.11	Piriform cortex model: power spectrum of weak shock response.	410
7.12	Piriform cortex model: network response to a weak shock stimulus.	412
7.13	Excitatory connections in the piriform cortex.	413
7.14	Inhibitory connections in the piriform cortex.	414
7.15	Piriform cortex model: histogram of feedforward/feedback inhibitory neuron and pyramidal neuron spiking after a weak shock stimulus.	415
7.16	Piriform cortex model: parameter dependence of weak shock-evoked oscillations.	417
7.17	Piriform cortex model: influence of connections between feedback inhibitory interneurons; intracellular traces.	418
7.18	Piriform cortex model: influence of connections between feedback inhibitory interneurons; surface field potentials and power spectra.	419
7.19	Experimental CSD contour plots of piriform cortex responses to strong and weak shocks.	421
7.20	Experimental CSD surface plots of piriform cortex responses to strong and weak shocks.	422

7.21 Piriform cortex model: CSD contour plot of the response of the random connectivity model to a weak shock stimulus.	425
7.22 Schematic diagram of how multiple inputs to layer 1a in the piriform cortex might be generated from a single weak shock.	427
7.23 Piriform cortex model: CSD contour plot of the response of the random connectivity model to a weak shock stimulus with multiple input shocks. . .	428
7.24 Piriform cortex model: intracellular response in layer 1a of a pyramidal neuron to a weak shock stimulus.	429
7.25 Piriform cortex model: CSD surface plots of the response of the random connectivity model to weak and strong shock stimuli.	431
7.26 Piriform cortex model: structured connections from the olfactory bulb to the piriform cortex.	436
7.27 Piriform cortex model: structured intracortical connections.	437
7.28 Piriform cortex model: effect of structured connections on spike latencies. .	438
7.29 Piriform cortex model: network response to a weak shock stimulus in the structured connectivity model.	441
7.30 Piriform cortex model: CSD contour plot of the response of the structured connectivity model to a weak shock stimulus with multiple input shocks with different populations of pyramidal neurons activated on successive waves of inputs.	443

7.31 Piriform cortex model: CSD surface plot of the response of the structured connectivity model to a weak shock stimulus with multiple input shocks with different populations of pyramidal neurons activated on successive waves of inputs.	444
7.32 Piriform cortex model: network response to a weak shock stimulus in the structured connectivity model with random pyramidal-to-pyramidal connections.	445
7.33 Piriform cortex model: network response to a weak shock stimulus in the structured connectivity model with the same subset of neurons activated each cycle.	447
7.34 Piriform cortex model: CSD contour plot of the response of the structured connectivity model to a weak shock stimulus with multiple input shocks with the same population of pyramidal neurons activated on successive waves of inputs.	448
7.35 Piriform cortex model: intracellular response of a pyramidal neuron in the structured connectivity model to a weak shock stimulus with multiple input shocks with the same population of pyramidal neurons activated on successive waves of inputs.	449

List of Tables

3.1	Parameter search results for model 2.	120
3.2	Parameter search results for model 4.	126
3.3	Parameter search results for model 5.	133
4.1	Parameter values for two different four-conductance models.	203
4.2	Bayesian analysis for the within-class comparison.	205
4.3	Parameter values for four- and three-conductance models.	207
4.4	Bayesian analysis for the between-class comparison.	209
4.5	Bayesian analysis for the whole-class comparison.	212
4.6	Summary of parameter estimation methods.	223
5.1	Passive parameter values for the full and simplified pyramidal neuron model.	263
5.2	Parameter values for the synaptic channel models.	280
5.3	Reversal potentials for ionic channels.	296

5.4	Conductances and gating variable exponents for pyramidal neuron ionic channels.	297
5.5	Parameters of the calcium buffer.	302
5.6	Neuromodulation of cellular parameters.	303
5.7	NMDA synapse parameters.	304
5.8	Parameters for synaptic facilitation and depression.	305
5.9	Morphology of the reduced pyramidal neuron model and synaptic channels in each compartment.	309
5.10	Conductances and gating variable exponents for feedback inhibitory neuron ionic channels.	310
5.11	Parameter values for feedback inhibitory interneurons.	310
6.1	Network scales.	337
6.2	Network connectivities. NA = not applicable.	343
6.3	Computational requirements for simulating the piriform cortex for one second of simulation time at different scales.	349

Part I

Introduction

Introduction

Listen to the technology and find out what it is telling us.

Carver Mead

1.1 Overview of the thesis

The past few years have seen a great increase in interest in olfactory neuroscience. Much of this interest is attributable to the identification of putative olfactory receptors by molecular biologists [11, 12, 52, 56], thus providing the vital link between odorants and the responses of olfactory sensory neurons which was hitherto lacking. However, the nature of the computations performed by the olfactory regions of the brain remain obscure. For the past several years, we have been using a combination of electrophysiological experimental techniques and computer modeling to help elucidate the function of the olfactory system [2, 4, 5, 25, 26, 69, 71]. This work has focused largely on the piriform cortex (primary olfactory cortex) [2, 25, 26, 69, 71] but also on the olfactory bulb [4, 5]. The construction of realistic computer simulations of the olfactory system at both the neuronal and network

levels serves as a necessary bridge between experimental data and abstract theories of olfactory computation, and also provides many insights as to which experiments to perform in order to gain the most information about the system.

This thesis describes the construction of a second-generation computer model of the piriform cortex. The first generation model was created by Matthew Wilson [71]; this was the first computer model of this brain region that attempted to simulate the entire system at a reasonable level of detail and accurately replicate experimental data. The new model was constructed in order to incorporate much more accurate models of the constituent neurons, synapses, network connectivities, and olfactory bulb inputs to the system, all based on new data acquired in the last few years. The goal of this work was to explore the oscillatory dynamics of the piriform cortex, to investigate coding strategies in the olfactory system, and to explore the effects of neuromodulation on the system dynamics.

In addition, this thesis includes experimental work which was necessary to constrain critical parameters of the model, as well as simulation tools which were constructed in the course of building the model. This introduction will provide an overview of the relevant background material, focusing on the areas pursued in the thesis.

1.2 The mammalian olfactory system

1.2.1 The olfactory epithelium and olfactory bulb

Odorants first contact the nervous system in the olfactory epithelium, where they dissolve in a thin sheet of mucus and eventually bind to olfactory receptor molecules located on cilia of olfactory sensory neurons. There appear to be about 1000 distinct olfactory receptor molecules, each of which is distributed quasi-randomly over a large group of sensory neurons in the epithelium [11, 12]. A single sensory neuron appears to express only one of the 1000 receptor molecules [52], but can nevertheless respond to a wide range of odorants [49]. Additionally, multiple olfactory receptors can respond to a given odor [49]. Olfactory sensory neurons send their axons along the olfactory nerve into the olfactory bulb.

The anatomy of the olfactory bulb is shown in figures 1.1 and 1.2 [62]. Information from the olfactory sensory neurons travels from the olfactory epithelium to the bulb along the olfactory nerve, ending in dense dendritic bundles known as glomeruli. Within the glomeruli, olfactory nerve axons make synapses with the primary output neurons of the olfactory bulb, the mitral and tufted cells. In addition, small neurons known as periglomerular cells provide intra- and interglomerular connections. Recent evidence suggests that olfactory sensory neurons projecting to a single glomerulus all express a single receptor subtype [52]. Within a glomerulus there are standard axo-dendritic connections involving sensory neurons, mitral cells and periglomerular cells as well as dendrodendritic synapses between periglomerular

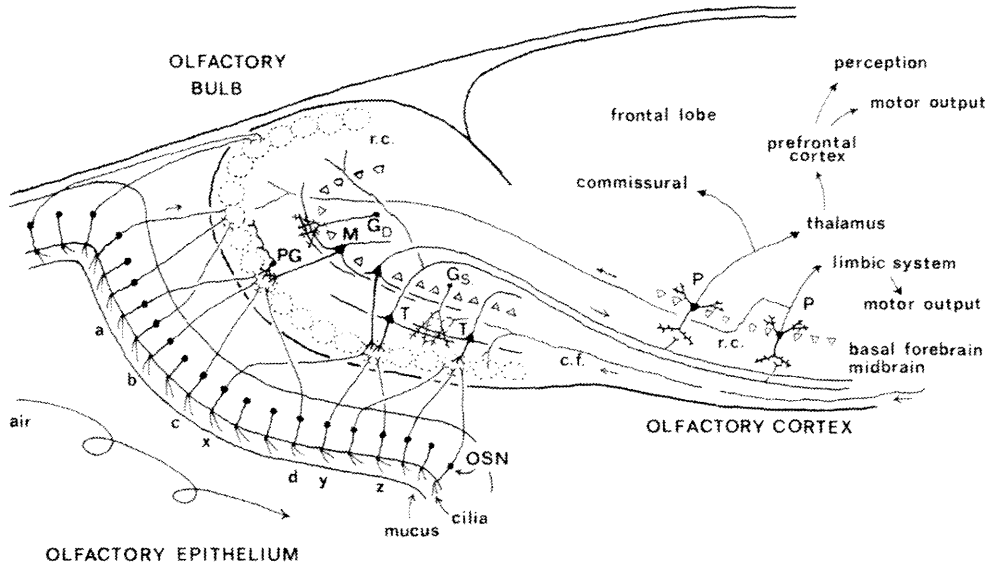


Figure 1.1: Schematic diagram of the mammalian olfactory system. Abbreviations: *OSN*, olfactory sensory neuron; *PG*, periglomerular cell; *M*, mitral cell; *T*, tufted cell; *G_S*, superficial granule cell; *G_D*, deep granule cell; *r.c.*, recurrent axon collateral; *c.f.*, centrifugal fiber; *P*, pyramidal cell. From [62].

and mitral cells.


From here, information travels down the mitral/tufted cell primary dendrite. These neurons also have secondary dendrites which form dendrodendritic synapses with the primary inhibitory neurons in the olfactory bulb, the granule cells [55, 62]. Granule cells have no axons; thus their effects appear to be mediated entirely by dendrodendritic interactions with mitral cells. Granule cells are excited by depolarization of mitral cell secondary dendrites and inhibit these same dendrites, giving rise to self-inhibition onto mitral cells.

Furthermore, nearby or distant dendrodendritic synapses of the same granule cell may also be activated, giving rise to lateral inhibitory interactions. Granule cells are also prominent targets for centrifugal input to the olfactory bulb, both from the piriform cortex and from neurons providing neuromodulatory input to the bulb [62]. Mitral cells integrate the afferent and centrifugal inputs with the inhibitory granule cell inputs and send their outputs to the piriform cortex (and other brain regions) through the lateral olfactory tract (LOT).


1.2.2 Primary olfactory (piriform) cortex

The piriform (literally, “pear-shaped”) cortex (also known as the pyriform or prepyriform cortex), the subject of this thesis, is a phylogenetically old cortical structure which receives the majority of its input from the olfactory bulb and is thus considered to be the primary olfactory cortex. The piriform cortex can be divided into three layers based on connection patterns (figures 1.3, 1.4 and 1.5) [22]. Layer 1 is mainly comprised of axons and axon collaterals and can be divided into two parts. Layer 1a consists of axons and axon collaterals from the lateral olfactory tract; this provides direct input from olfactory bulb mitral cells to the distalmost part of pyramidal neuron apical dendrites. Layer 1b consists mainly of densely packed “association fibers” which interconnect different pyramidal neurons. Layer 1 also contains cell bodies of various types of feedforward inhibitory neurons. Layer 2 is a densely packed band of cell bodies of superficial pyramidal neurons. Layer 3 consists of cell bodies of deep pyramidal neurons, excitatory and inhibitory interneurons, and fibers medi-

A.
Afferent
Fibers



B.
Principal
Cells



C.
Local
Inter-
neurons




Figure 1.2: (Previous page) Anatomy of the olfactory bulb. Abbreviations as in previous figure as well as: *ONL*, olfactory nerve layer; *GL*, glomerular layer; *EPL*, external plexiform layer; *MCL*, mitral cell layer; *IPL*, internal plexiform layer; *GRL*, granule cell layer; *SA*, short axon cells. Dashed lines represent glomeruli. From [62].

ating local connections between pyramidal neurons and interneurons. Deep to layer 3 the piriform cortex turns into the subjacent endopiriform nucleus, which, although extensively connected to piriform cortex, is generally regarded as a separate region.

Piriform cortex pyramidal neurons consist of a cell body or soma, a long apical dendrite or dendrites extending to the surface of the cortex (spanning layers 1b and 1a), and a number of basal dendrites that receive local excitatory input in layer 3 [20, 22] (figures 1.5 and 1.6). These neurons make extensive connections with other pyramidal neurons in piriform cortex, both locally (axon collaterals projecting onto synapses on basal dendrites in layer 3) and more distantly (axon collaterals projecting onto synapses on apical dendrites in layer 1b) (figure 1.4).

Piriform cortex interneurons fall into several classes [20, 22] (figures 1.5 and 1.7). In layer 1 are found superficially located horizontal neurons and small globular-soma neurons which are both believed to be inhibitory. These neurons provide synaptic input across the length of the pyramidal neuron apical dendrite and mediate feedforward inhibition onto these dendrites since they receive direct input from olfactory bulb mitral cells (figure

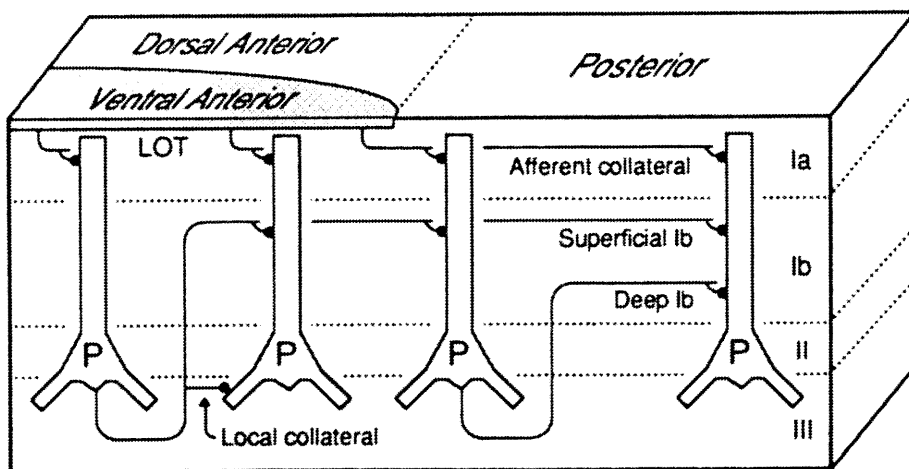


Figure 1.3: Subdivisions of the piriform cortex showing both the layered structure (layers 1a, 1b, 2 and 3) and the division into ventral anterior, dorsal anterior, and posterior regions. Abbreviation: *P*, pyramidal neuron. From [37].

1.7). In layer 3 are found multipolar neurons which either have spiny dendrites (excitatory multipolar neurons) or aspiny dendrites (inhibitory multipolar neurons). Aspiny multipolar neurons mediate feedback inhibition onto the cell bodies of pyramidal neurons (figure 1.7). Little is known about spiny multipolar neurons, or about several even rarer neuron types such as “semilunar” neurons, which resemble pyramidal neurons without basal dendrites; thus, these neuron types were not included in the present model.

The physiology of pyramidal neurons in piriform cortex is similar to that of pyramidal neurons found in other parts of the brain. These neurons display “regular spiking” behavior [3, 51], featuring fairly broad spikes and significant spike frequency adaptation to a sustained current stimulus. Although there is less physiological data for piriform cortex inhibitory neurons, they clearly fall into the “fast-spiking” category [51, 54], with narrow spikes, high excitability, and no spike frequency adaptation. A number of synaptic and voltage-dependent ionic channels are known to exist in piriform cortex pyramidal neurons including fast sodium channels, fast and slow potassium channels, calcium channels, AMPA and NMDA excitatory synapses, and GABA-A and GABA-B inhibitory synapses [22, 33, 34, 35].

1.2.3 Oscillations in the olfactory system

Oscillatory activity is a prominent feature of all olfactory systems, from invertebrates to mammals [1, 15, 18, 39, 42, 43]. Since one of the motivations of the model described in this thesis is to reproduce and explore the oscillatory dynamics of piriform cortex, these




Figure 1.4: Schematic diagram of excitatory connections between pyramidal neurons (P) in the piriform cortex. Each pyramidal neuron schematic represents a class of neurons in one of the three subdivisions of this brain region. The density of connections between regions is represented by the width of lines connecting the pyramidal neurons. Shading of pyramidal neurons and connections is just to distinguish between the three regions. Abbreviations: M/T , mitral/tufted cells of the olfactory bulb; G , granule cells of the olfactory bulb; LOT , lateral olfactory tract. From [22].




Figure 1.5: The main classes of neurons in the piriform cortex, positioned in the layers where they are most commonly found. Abbreviations: *SP*, superficial pyramidal neurons; *DP*, deep pyramidal neurons; *S*, semilunar cells; *H*, horizontal neurons; *G*, globular-soma neurons; *MS*, spiny multipolar neurons; *M*, smooth multipolar neurons. Classes *H*, *G*, and *M* are inhibitory; the rest are excitatory. From [22].




Figure 1.6: Morphology of a typical layer 2 superficial pyramidal neuron in piriform cortex. Note that the apical dendrite bifurcates into several secondary dendrites before leaving the cell body. Abbreviation: *IS*, initial segment of the axon. Bar: 100 μm . From [20].




Figure 1.7: Inhibitory connections in the piriform cortex. Abbreviations: *P*, pyramidal neuron; *FF*, feedforward inhibitory interneuron; *FB*, feedback inhibitory interneuron. From [22].

dynamics are briefly discussed here.

Extracellular field potentials recorded in both the mammalian olfactory bulb and the piriform cortex show oscillations *in vivo* in two primary frequency ranges: the theta range (4-10 Hz), which is often entrained to the sniffing rhythm, and the gamma range (35-85 Hz) (figure 1.8) [9, 10, 16, 18, 36]. In fact, when field potentials are recorded simultaneously from olfactory bulb and piriform cortex, in many cases the two field potentials exhibit a large degree of coherence [9, 10]. Coherent oscillations in the two structures may be mediated in part by the extensive back-projections from piriform cortex to the olfactory bulb. These projections synapse primarily onto inhibitory granule cells in the olfactory bulb (figure 1.4) [45], and are thus well placed to mediate phase-locking between bulbar and cortical oscillations.

1.3 Modeling the piriform cortex

1.3.1 Why build realistic models of piriform cortex?

Although mathematical and computational models are a central part of most scientific and engineering disciplines, they are considerably less prominent in neuroscience. There has been a steady growth in the number of realistic computer models of single neurons constructed in the past few years [7, 41], but realistic models of networks of neurons are still comparatively rare. This is largely because of the difficulties in acquiring the data and




Figure 1.8: Oscillatory activity in the olfactory system as represented by surface field potentials recorded in the olfactory bulb and piriform cortex. Note the high frequency gamma-band oscillations modulated by lower frequency theta-band oscillations, as well as the high degree of coherence in the theta modulation of bulb and cortex. Abbreviation: *PC*, piriform cortex. From [9].

building the simulation tools necessary to construct these models, but also because of a lack of understanding of the purpose of model building.

Whenever a system to be modeled (such as a neuron or a network of neurons) consists of a large number of interacting, nonlinear components, the behaviors of such a system can be highly unpredictable and intuition is not a reliable guide [67]. Realistic computer models can thus be very useful to an experimentalist in showing the range of possible behaviors that can be obtained when known components at one level (*e.g.*, synaptic and ionic channels, dendritic segments) are assembled into larger entities (*e.g.*, single neuron models, network models). In addition, such models act as a consistency check between data at different levels, thus highlighting which data is likely to be false or incomplete. When data is known to be incomplete, models allow investigators to perform “what-if” numerical experiments to test the plausibility of different theories, which can serve as a useful guide to further experimentation.

Realistic models of neuronal networks can also aid investigators in understanding the origins of network-level dynamical behaviors such as oscillations. Once a dynamical behavior of interest can be replicated, model components can be modified or removed to determine which aspects of the system are fundamental for achieving the correct behavior and which are not. In this way, realistic network models can be abstracted in many different ways corresponding to each type of behavior exhibited by the model. An example of this, based on the work described here, is [14]. Network models are also useful for critically evaluating

theories of neural computation and coding [24, 26, 27, 71].

1.3.2 Realistic modeling

Since this thesis describes the construction of a realistic computer model of piriform cortex, it is important to state precisely what is meant by “realistic”. The key criteria are:

1. The level of detail in the model must be a reasonable reflection of the current body of experimental knowledge given the limits of modern computers.
2. The model must reproduce relevant experimental data to a high degree of accuracy.
3. The model should provide useful suggestions for further experimental work and useful ideas about the dynamics and functions of the system being modeled.

These criteria have been met both for the previous model (at the time it was built) and the model described here; this is discussed in greater detail in the following chapters. The current model inevitably has limitations which will also be discussed in detail.

Clearly, building realistic models is critically dependent on the current state of the experimental database; in fact, it may be argued that the most important contribution of these models is to let experimentalists know which data needs to be collected to improve the model. This feedback loop between models and experiments is the primary strength of the realistic modeling approach.

1.3.3 The previous piriform cortex model

The first realistic computer model of piriform cortex was that of Matthew Wilson [71] which is the direct ancestor of the model described in this thesis. Wilson's model was able to replicate the surface field potential response of piriform cortex in response to weak and strong shocks of the LOT, and also highlighted the importance of synaptic time constants and axonal conduction velocities in generating field potentials which match experimental data. However, Wilson's model had many limitations: the simulated neurons were not parameterized to fit experimental data, the inputs were not strongly based on experimental data, the connection patterns are no longer consistent with the most recently acquired data [22, 37, 38], and many aspects of synaptic transmission were ignored (*e.g.*, fast and slow GABA-A subtypes, NMDA channels, neuromodulation, synaptic facilitation and depression). I will briefly discuss new features of the present model below and give a more detailed comparison between the two models in chapter 6.

1.3.4 New model features

A major goal of the work described in this thesis has been to incorporate accurate models of single neurons in piriform cortex into a large-scale cortical model. The two primary classes of neuron types in piriform cortex are pyramidal neurons and a number of inhibitory interneuron types. In contrast to the previous model, the piriform cortex model described here was constructed so as to accurately reproduce the input-output relations of both pyramidal

and inhibitory interneurons as well as the synaptic dynamics of the system. The pyramidal neuron model contains a variety of voltage- and calcium-dependent ionic channels, AMPA and NMDA excitatory synaptic receptors, and fast and slow GABA-A and slow GABA-B inhibitory synaptic receptors.

The synaptic connectivities of the piriform cortex model have been significantly changed to reflect new experimental data. The previous model divided piriform cortex into anterior (rostral) and posterior (caudal) subdivisions only. The modern view, reflected in the new model, is that piriform cortex is divided into three broad regions on the anterior-posterior and dorsoventral axes (figure 1.3) on the basis of external and internal synaptic connection patterns (figure 1.4) [22]. The majority of input from the olfactory bulb arrives on pyramidal neuron dendrites in the ventral anterior piriform cortex (vAPC); these neurons in turn project large numbers of long-range collaterals to the superficial layer 1b dendrites of pyramidal neurons in the other two regions. Few local projections arise from vAPC pyramidal neurons. Pyramidal neurons in the dorsal anterior piriform cortex (dAPC) and the posterior piriform cortex (PPC) project to deep layer 1b pyramidal neuron dendrites in the other regions and also give rise to significant numbers of local projections onto basal dendrites of nearby pyramidal neurons.

The realism of modeled inputs to piriform cortex from the olfactory bulb have been significantly improved with respect to the previous model. Outputs from olfactory bulb mitral cells are represented by a spike-generating object that can generate a number of

different patterns of activity across the ensemble of neurons. These patterns were derived in part from single-unit recordings done by Upinder Bhalla [5]. Artificial inputs such as weak or strong shock stimuli can also be simulated by the bulb object, as can a variety of responses of mitral cells to odors based on several possible coding schemes. This is described in more detail in chapter 6.

1.3.5 Tuning the model

The primary outputs used to constrain the piriform cortex network model were the sequence of synaptic events following brief electrical shocks to the LOT. These have recently been described experimentally in much greater detail than the data that was available for parameterizing the previous model [37, 38, 39]. Both surface field potential outputs and current-source density (CSD) responses were computed. Both strong and weak shocks were used, as they produce markedly different patterns of synaptic activation. Strong shocks evoke a single wave of excitation that spreads across piriform cortex and is then damped out, while weak shocks produce a characteristic damped oscillatory response in the field potential and CSD responses. In addition, experimental data on mitral cell firing rates and interspike interval distributions [5] were used to constrain spike inputs from the olfactory bulb.

1.4 Questions the model can help answer

1.4.1 Coding in the olfactory system

Despite decades of intensive research, there is no consensus as to how odors are encoded in the outputs of olfactory bulb mitral cells. Some authors believe that odorants are encoded by changes in mitral cell firing rate, either in a small localized group of neurons [61, 62] or in a larger distributed group of neurons [46, 57]. Others believe that odorants are encoded by synchronized firing of mitral cells (or their analogs in insects, the projection neurons of the antennal lobe [42, 43]), while still others argue for more complex spatiotemporal codes involving chaotic dynamics [19, 64].

The lack of consensus in this area causes difficulties in the construction of realistic models of piriform cortex, since the model cannot be expected to reproduce the behavior of the real system without being supplied with realistic inputs. As will be seen, this was resolved by creating simulation objects that can mimic many of the proposed olfactory coding strategies, as well as accurately replicating the first-order statistics of mitral cell responses to background odors.

The question of how odors are encoded in piriform cortex, or even what roles the piriform cortex may have in odor processing, has also not been resolved, primarily because the relevant data is so limited. Much of the functional data relating to olfactory cortex has been obtained from human studies. Lesions in piriform cortex in humans disrupt olfac-

tory perception in several ways, including odor identification and discrimination between odors [31, 32, 72], but without affecting odor intensity perception. Sobel *et al.* [65] have established through functional MRI studies that the physical act of sniffing causes marked changes in the local tissue oxygenation in piriform cortex. The presence of an odor also causes changes in oxygenation, but these decay very rapidly (within 30-40 sec [66]). In addition, studies have shown a remarkable similarity in the response patterns of neurons in piriform cortex and those in orbitofrontal cortex, a region to which piriform cortex projects [59, 60]. However, these studies still leave the question of how odors are encoded in the firing patterns of piriform cortex pyramidal neurons unanswered.

From a systems perspective, the dominant paradigm of the computational role of the piriform cortex has been that it functions as a biological autoassociative memory roughly analogous to a Hopfield network [23, 28, 29]. In contrast, the modeling work described in the last chapter suggests a completely different paradigm whereby separate nonoverlapping neuronal groups are activated in each 50 *Hz* gamma cycle. Feedback connections from posterior to anterior cortex, an essential component of all “associative memory” models of piriform cortex [23, 24, 25] are in fact not essential to replicate experimental data, although they are known to exist [21, 22]. The computational picture of the piriform cortex that emerges from the modeling work in this thesis is of a largely feedforward system driven by inputs to the olfactory bulb where feedback may have the role of assisting a pattern shift on each gamma oscillation cycle. In this case, the output code of the piriform cortex

is a complex spatiotemporal code, reminiscent in some ways of codes found in the insect olfactory system [43, 42]. One of the most intriguing aspects of this work is the fact that these computational ideas emerged as a direct consequence of the process of matching experimental data to computer models.

1.4.2 Origin and functions of oscillations

The exact functions of oscillations in the olfactory bulb and cortex are unknown as in the rest of the brain. Some investigators have postulated that the sequence of several gamma oscillations occurring within a theta oscillation may serve as a computational unit of processing [44]. The mechanism of the genesis of oscillations is also unclear in many brain structures including the piriform cortex. Previous modeling work as well as experiments involving weak afferent shocks suggested that piriform cortex oscillations can be generated endogenously [17, 71]; however, it is also possible that *in vivo* such oscillations are driven partly or even primarily by inputs from the olfactory bulb [39]. The modeling work described in this thesis strongly supports the view that oscillations in the gamma band are driven by olfactory bulb inputs. Additionally, the model also suggests a new paradigm for the organization of these inputs which in turn suggests that each gamma cycle represents the analysis of a different aspect of a single odor.

1.4.3 Neuromodulation in the piriform cortex

Piriform cortex is innervated by a number of neuromodulatory nuclei in the brainstem which supply cholinergic, noradrenergic, serotonergic, and dopaminergic inputs to the cortex [22, 63]. Some neuromodulators, predominantly acetylcholine (ACh) and norepinephrine (NE) have been strongly associated with certain aspects of learning and memory [6, 13, 53, 58, 68, 70]. Work done by Michael Hasselmo [25, 26] showed that acetylcholine, when bath-applied to a brain slice preparation in the cortex, caused a selective depression in layer 1b synaptic transmission. Work included as part of this thesis extends this to norepinephrine. However, NE also has a number of other effects, including changes in neuronal excitability in both pyramidal neurons and interneurons [47, 50]. The network-level effects of NE are not easily predictable from its cellular and synaptic effects; thus, one motivation for the network modeling efforts included in this thesis was to incorporate all these effects to be able to predict the overall effect of NE on the dynamics of the network.

1.4.4 Other issues

I used the model to explore several other aspects of the olfactory system, including the role of background spiking inputs from the olfactory bulb, the roles of feedforward and feedback inhibition onto pyramidal neurons, and the roles of various types of inhibitory synaptic receptor subtypes on the dynamical behaviors of the model.

1.5 Modeling methodologies

In the course of building models of single neurons in piriform cortex, a number of tools were developed to facilitate certain aspects of the modeling process. A description of these tools forms a major component of this thesis. These tools are briefly summarized here and discussed in more detail in the relevant chapters.

1.5.1 Simulation environment

The computer models of piriform cortex described in this thesis were all simulated using the simulation program GENESIS (General NEural SIMulation System) [7]. Single neurons were simulated by dividing them into isopotential compartments and using standard compartmental modeling techniques [41]; details of the models are to be found in the following chapters. Many extension libraries totalling approximately 60,000 lines of C code were added to GENESIS specifically in order to build the piriform cortex model. These libraries consisted of the olfactory bulb spike-generating objects, a variety of synaptic objects, objects controlling neuromodulation, and commands to establish groups of synaptic connections, weights and delays.

1.5.2 Parameter searching

I have developed a group of parameter searching methods usable within GENESIS that greatly simplify the process of assigning values to unknown parameters in single-neuron

models. Several methods have been used, including conjugate-gradient, simulated annealing, genetic algorithms, and stochastic search. The highly accurate match between the pyramidal neuron model and the experimental data on which it was based is a direct consequence of these methods. I believe that these methods will soon become an essential component of the simulation toolkit of scientists building realistic single-neuron models, since assigning parameters iteratively by hand is both much more tedious and gives much poorer results than those obtained using these methods. At the same time, a certain amount of expertise in using these methods is necessary in order to obtain the best results; this is discussed at length in chapter 3.

1.5.3 Bayesian methods

Eventually, a large enough number of realistic single-neuron models will exist that it will be possible to ask which one is the best model given some set of data to be matched. Although most modelers would currently answer this question based on a visual inspection of the results or on the basis of what aspects of the data they are most interested in, it is possible to ask this question much more rigorously if the models generate output probabilistically. I have shown that in this case one can use the Bayesian framework to compare individual models and classes of models and assign relative probabilities to the models based on how well they match the data. As models proliferate the Bayesian methodology will be essential to allow the objective evaluation of different models.

1.6 New work suggested by the model

1.6.1 Experimental work

The piriform cortex model has highlighted the importance of a more accurate understanding of the connection patterns between the olfactory bulb and the piriform cortex. As will be shown in the last chapter, simple random connectivity between the two structures results in a model which can replicate the cortical surface EEG with reasonable accuracy but which cannot replicate the CSD response to a weak shock stimulus. However, a model which has highly structured connections between bulb and cortex can replicate the CSD response. Detailed anatomical and physiological studies will be necessary to determine the true nature of the connection patterns between bulb and cortex. These studies are crucially important in that they will have a significant impact on our understanding of how computations are performed in this system.

From the perspective of improving the model, a number of experiments need to be performed. More data on pyramidal neuron responses to a variety of input stimuli will be necessary for improving that model. In addition, little experimental data exists to constrain the models of feedforward inhibitory neurons. These neurons appear to have a profound effect on responses of the network to both weak and strong shock stimuli. Therefore, experiments to better characterize these neurons are essential.

From the perspective of coding, a very important experimental study is for large-scale

multiunit recordings to be obtained from arrays of olfactory bulb mitral cells in awake behaving animals involved in odor detection tasks. This will allow us to improve the quality of the inputs delivered to the model and also to refine our understanding of coding at the level of the bulb and cortex.

1.6.2 Modeling work

The modeling work presented here has suggested a number of future paths for continued work. The pyramidal neuron model can be extended in a number of ways. One approach would be to increase the realism of the neuron morphology, which was heavily simplified in the present model for computational reasons. One question of interest concerns the possible roles of dendritic spines, which can isolate the conductance changes at synapses from the main dendritic trunk, thus effectively increasing the space constant of the neuron [40]. This may have a significant effect on synaptic integration in pyramidal neurons. In addition, the possible roles of active dendritic currents [30] and a somatic spike-initiating zone [48] remain to be established for these neurons.

There are an enormous number of components of the present model which have not been explored fully owing to time constraints. The roles of synaptic facilitation, synaptic depression, and NMDA receptors in generating network-level phenomena are not clear. Some network-level phenomena, such as the role of norepinephrine in the weak-shock response, have not been fully characterized. Some work has also been done on modeling the

surface EEG response to odors, but this work is far from complete and has therefore not been included in this thesis.

1.7 Summary of thesis contributions

1.7.1 The new piriform cortex model

A realistic computer model of piriform cortex was constructed. This model is the most accurate model of this brain structure that has been built to date. The goals motivating the construction of the model were as follows:

1. The network model incorporates models of single neurons which were required to match the input-output behavior of real neurons very accurately. In the case of pyramidal neurons, the morphology was systematically simplified from the morphology of a real piriform cortex pyramidal neuron.
2. The neuron models contain synaptic receptor types known to exist in the piriform cortex but not previously incorporated into network models of this system, including NMDA receptors, fast and slow GABA-A receptors [33, 34, 35], and synaptic facilitation and depression [8].
3. The model features more accurate inputs to the cortical model from olfactory bulb mitral cells.

4. The model includes neuromodulation with norepinephrine (NE) at both the cellular and synaptic levels.

1.7.2 Lessons from the model

The piriform cortex model has emphasized the importance of the role of background spiking input from the olfactory bulb in keeping pyramidal neurons in the ventral anterior piriform cortex close to spiking threshold. Without these background inputs, the system becomes largely unresponsive and cannot accurately replicate the strong shock response. The model predicts that norepinephrine, which increases the excitability of pyramidal neurons, must decrease the background firing rate of olfactory bulb mitral cells in order to prevent cortical pyramidal neurons from spontaneously spiking at high rates. There is some experimental support for this conclusion.

The model has also highlighted the importance of feedforward inhibition in the generation of the strong shock response, and suggested that feedback onto these neurons may be involved in the damping of the surface EEG observed in the weak shock response. The model shows that feedback inhibition alone is not sufficient to replicate the strong shock response.

Most significantly, the attempt to accurately replicate the CSD response to weak shock stimuli has suggested that mitral cells in the olfactory bulb and pyramidal neurons in piriform cortex may be divided into nonoverlapping neuronal groups, such that mitral cells

of a given group in the bulb project primarily or exclusively to pyramidal neurons of a given group in the cortex, and similarly pyramidal neurons in the cortex project primarily or exclusively to other neurons of the same group (with the possible exception of feedback projections from posterior to anterior piriform cortex). This arrangement, if true, suggests that the output code of piriform cortex will not resemble a static attractor but will be a complex spatiotemporal pattern, and that each gamma oscillation cycle may be involved in analyzing separate aspects of the same input stimulus.

1.7.3 Experimental work

In order to obtain parameters relating to the effects of norepinephrine on synaptic transmission in piriform cortex, a number of brain-slice experiments were performed which are described in chapter 2. NE was found to have profound effects on afferent synaptic transmission in layers 1a and 1b of piriform cortex, as well as effects on synaptic facilitation and depression and cellular excitability. These effects of NE were incorporated into the piriform cortex network model.

1.7.4 Development of simulation tools

Simulation tools were developed as extension libraries for the neural simulation program GENESIS to simulate synaptic, cellular, and network-level entities of the piriform cortex model. An extension library consisting of the parameter searching code was developed

and has been incorporated into recent releases of GENESIS. Simulation tools for Bayesian analysis of compartmental models were also developed.

Bibliography

- [1] E. D. Adrian. The Electrical Activity of the Mammalian Olfactory Bulb. *Journal of Physiology (London)*, 100:459–473, 1942.
- [2] P. F. Baldi, M. C. Vanier, and J. M. Bower. On the Use of Bayesian Methods for Evaluating Compartmental Neural Models. *Journal of Computational Neuroscience*, 5:285–314, 1998.
- [3] E. Barkai and M. E. Hasselmo. Modulation of the Input/Output Function of Rat Piriform Cortex Pyramidal Cells. *Journal of Neurophysiology*, 72(2):644–658, 1994.
- [4] U. S. Bhalla and J. M. Bower. Exploring Parameter Space in Detailed Single Neuron Models – Simulations of the Mitral and Granule Cells of the Olfactory-Bulb. *Journal of Neurophysiology*, 69(6):1948–1965, 1993.
- [5] U. S. Bhalla and J. M. Bower. Multiday Recordings from Olfactory Bulb Neurons in Awake Freely Moving Rats: Spatially and Temporally Organized Variability in Odorant Response Properties. *Journal of Computational Neuroscience*, 4(3):221–256, 1997.

- [6] A. Blokland. Acetylcholine: a Neurotransmitter for Learning and Memory? *Brain Research Reviews*, 21(3):285–300, 1995.
- [7] J. M. Bower and D. Beeman. *The Book of Genesis, 2nd. Ed.* Springer Verlag, New York, 1998.
- [8] J. M. Bower and L. B. Haberly. Facilitating and Nonfacilitating Synapses on Pyramidal Cells – a Correlation Between Physiology and Morphology. *Proceedings of The National Academy of Sciences of The United States of America*, 83(4):1115–1119, 1986.
- [9] S. L. Bressler. Spatial Organization of EEGs from Olfactory Bulb and Cortex. *Electroencephalography and Clinical Neurophysiology*, 57:270–276, 1984.
- [10] S. L. Bressler and W. J. Freeman. Frequency Analysis of Olfactory System EEG in Cat, Rabbit and Rat. *Electroencephalography and Clinical Neurophysiology*, 50:19–24, 1980.
- [11] L. Buck. The Molecular Architecture of Odor and Pheromone Sensing in Mammals. *Cell*, 100(6):611–618, 2000.
- [12] L. Buck and R. Axel. A Novel Multigene Family May Encode Odorant Receptors – a Molecular Basis for Odor Recognition. *Cell*, 65(1):175–187, 1991.

- [13] D. M. Compton, K. L. Dietrich, J. S. Smith, and B. K. Davis. Spatial and Nonspatial Learning in the Rat Following Lesions to the Nucleus Locus-coeruleus. *Neuroreport*, 7(1):177–182, 1995.
- [14] S. M. Crook, G. B. Ermentrout, M. C. Vanier, and J. M. Bower. The Role of Axonal Delay in the Synchronization of Networks of Coupled Cortical Oscillators. *Journal of Computational Neuroscience*, 4(2):161–172, 1997.
- [15] K. R. Delaney, A. Gelperin, M. S. Fee, J. A. Flores, R. Gervais, D. W. Tank, and D. Kleinfeld. Waves and Stimulus-Modulated Dynamics in an Oscillating Olfactory Network. *Proceedings of the National Academy of Sciences*, 91(2):669–673, 1994.
- [16] W. J. Freeman. Distribution in Time and Space of Prepyriform Electrical Activity. *Journal of Neurophysiology*, 22:644–665, 1959.
- [17] W. J. Freeman. Effects of Surgical Isolation and Tetanization of Prepyriform Cortex in Cats. *Journal of Neurophysiology*, 31:349–357, 1968.
- [18] W. J. Freeman. *Mass Action in the Nervous System*. Academic Press, New York, 1975.
- [19] W. J. Freeman. Simulation of Chaotic EEG Patterns with a Dynamic Model of the Olfactory System. *Biological Cybernetics*, 56:139–150, 1987.

- [20] L. B. Haberly. Structure of the Piriform Cortex of the Opossum. 1. Description of Neuron Types with Golgi Methods. *Journal of Comparative Neurology*, 213(2):163–187, 1983.
- [21] L. B. Haberly. Neuronal Circuitry in Olfactory Cortex – Anatomy and Functional Implications. *Chemical Senses*, 10(2):219–238, 1985.
- [22] L. B. Haberly. Olfactory Cortex. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.
- [23] L. B. Haberly and J. M.. Bower. Olfactory Cortex – Model Circuit for Study of Associative Memory? *Trends In Neurosciences*, 12(7):258–264, 1989.
- [24] M. E. Hasselmo. Acetylcholine and Learning in a Cortical Associative Memory. *Neural Computation*, 5(1):32–44, 1993.
- [25] M. E. Hasselmo, B. P. Anderson, and J. M. Bower. Cholinergic Modulation of Cortical Associative Memory Function. *Journal of Neurophysiology*, 67(5):1230–1246, 1992.
- [26] M. E. Hasselmo and J. M. Bower. Cholinergic Suppression Specific to Intrinsic not Afferent Fiber Synapses in Rat Piriform (Olfactory) Cortex. *Journal of Neurophysiology*, 67(5):1222–1229, 1992.
- [27] M. E. Hasselmo and J. M. Bower. Acetylcholine and Memory. *Trends In Neurosciences*, 16(6):218–222, 1993.

- [28] J. J. Hopfield. Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proceedings of the National Academy of Sciences – Biology*, 79(8):2554–2558, 1982.
- [29] J. J. Hopfield. Neurons with Graded Response have Collective Computational Properties Like Those of 2-state Neurons. *Proceedings of the National Academy of Sciences – Biology*, 81(10):3088–3092, 1984.
- [30] D. Johnston, J. C. Magee, C. M. Colbert, and B. R. Christie. Active Properties of Neuronal Dendrites. *Annual Review of Neuroscience*, 19:165–86, 1996.
- [31] M. Jones-Gotman and R. J. Zatorre. Olfactory Identification Deficits in Patients with Focal Cerebral Excision. *Neuropsychologia*, 26:387–400, 1988.
- [32] M. Jones-Gotman and R. J. Zatorre. Odor Recognition Memory in Humans: Role of Right Temporal and Orbitofrontal Regions. *Brain and Cognition*, 22:182–198, 1993.
- [33] E. D. Kanter and L. B. Haberly. NMDA-Dependent Induction of Long-Term Potentiation in Afferent and Association Fiber Systems of Piriform Cortex In Vitro. *Brain Research*, 525(1):175–179, 1990.
- [34] E. D. Kanter and L. B. Haberly. Associative Long-Term Potentiation in Piriform Cortex Slices Requires GABA(A) Blockade. *Journal of Neuroscience*, 13(6):2477–2482, 1993.

- [35] A. Kapur, W. W. Lytton, K. L. Ketchum, and L. B. Haberly. Regulation of the NMDA Component of EPSPs by Different Components of Postsynaptic GABAergic Inhibition: Computer Simulation Analysis in Piriform Cortex. *Journal of Neurophysiology*, 78(5):2546–2559, 1997.
- [36] L. M. Kay and W. J. Freeman. Bidirectional Processing in the Olfactory-Limbic Axis During Olfactory Behavior. *Behavioral Neuroscience*, 112(3):541–553, 1998.
- [37] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 1. Current Source-Density Analysis. *Journal of Neurophysiology*, 69(1):248–260, 1993.
- [38] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 2. Analysis with a System Model. *Journal of Neurophysiology*, 69(1):261–281, 1993.
- [39] K. L. Ketchum and L. B. Haberly. Synaptic Events that Generate Fast Oscillations in Piriform Cortex. *Journal of Neuroscience*, 13(9):3980–3985, 1993.
- [40] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, Oxford, U.K., 1999.
- [41] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.

- [42] G. Laurent and H. Davidowitz. Encoding of Olfactory Information with Oscillating Neural Assemblies. *Science*, 265:1872–1875, 1994.
- [43] G. Laurent and M. Naraghi. Odorant-Induced Oscillations in the Mushroom Bodies of the Locust. *Journal of Neuroscience*, 14(5):2993–3004, 1994.
- [44] J. E. Lisman and M. A. P. Idiart. Storage of 7 ± 2 Short-Term Memories in Oscillatory Subcycles. *Science*, 267:1512–1515, 1995.
- [45] M. B. Luskin and J. L. Price. The Topographic Organization of Associational Fibers of the Olfactory System in the Rat, Including Centrifugal Fibers to the Olfactory Bulb. *Journal of Comparative Neurology*, 216:264–291, 1983.
- [46] Kay L. M. and Laurent G. Odor- and Context-Dependent Modulation of Mitral Cell Activity in Behaving Rats. *Nature Neuroscience*, 2(11):1003–1009, 1999.
- [47] D. V. Madison and R. A. Nicoll. Noradrenaline Blocks Accommodation of Pyramidal Cell Discharge in the Hippocampus. *Nature*, 299(5884):636–638, 1982.
- [48] Z. F. Mainen, J. Joerges, J. R. Huguenard, and T. J. Sejnowski. A Model of Spike Initiation in Neocortical Pyramidal Neurons. *Neuron*, 15(6):1427–1439, 1995.
- [49] B. Malnic, J. Hirono, T. Sato, and L. B. Buck. Combinatorial Receptor Codes for Odors. *Cell*, 96:713–723, 1999.

- [50] G. J. Marek and G. K. Aghajanian. Alpha(1B)-adrenoceptor-mediated Excitation of Piriform Cortical Interneurons. *European Journal of Pharmacology*, 305(1-3):95–100, 1996.
- [51] D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative Electrophysiology of Pyramidal and Sparsely Spiny Stellate Neurons of the Neocortex. *Journal of Neurophysiology*, 54(4):782–806, 1985.
- [52] P. Mombaerts, F. Wang, C. Dulac, S. K. Chao, A. Nemes, M. Mendelsohn, J. Edmondson, and R. Axel. Visualizing an Olfactory Sensory Map. *Cell*, 87(4):675–686, 1996.
- [53] M. Ohno, M. Kobayashi, A. Kishi, and S. Watanabe. Working-Memory Failure by Combined Blockade of Muscarinic and Beta-Adrenergic Transmission in the Rat Hippocampus. *Neuroreport*, 8(7):1571–1575, 1997.
- [54] A. D. Protopapas and J. M. Bower. Physiological Characterization of Layer III Non-Pyramidal Neurons in Piriform (Olfactory) Cortex of Rat. *Brain Research*, 865(1):1–11, 2000.
- [55] W. Rall and G. R. Shepherd. Theoretical Reconstruction of Field Potentials Computed for Different Soma-Dendritic Distributions of Synaptic Input. In *The Theoretical Foundation of Dendritic Function*. MIT Press, Cambridge, Mass., 1995.

- [56] K. J. Ressler, S. L. Sullivan, and L. B. Buck. Information Coding in the Olfactory System – Evidence for a Stereotyped and Highly Organized Epitope Map in the Olfactory Bulb. *Cell*, 79(7):1245–1255, 1994.
- [57] Katz L. C. Rubin B. D. Optical Imaging of Odorant Representations in the Mammalian Olfactory Bulb. *Neuron*, 23:499–511, 1999.
- [58] S. J. Sara, C. Dyonlaurent, and A. Herve. Novelty Seeking Behavior in the Rat is Dependent Upon the Integrity of the Noradrenergic System. *Cognitive Brain Research*, 2(3):181–187, 1995.
- [59] G. Schoenbaum and H. Eichenbaum. Information Coding in the Rodent Prefrontal Cortex. 1. Single-Neuron Activity in Orbitofrontal Cortex Compared with that in Pyriform Cortex. *Journal of Neurophysiology*, 74(2):733–750, 1995.
- [60] G. Schoenbaum and H. Eichenbaum. Information Coding in the Rodent Prefrontal Cortex. 2. Ensemble Activity in Orbitofrontal Cortex. *Journal of Neurophysiology*, 74(2):751–762, 1995.
- [61] G. M. Shepherd and S. Firestein. Toward a Pharmacology of Odor Receptors and the Processing of Odor Images. *Journal of Steroid Biochemistry and Molecular Biology*, 39(4B):583–592, 1991.

- [62] G. R. Shepherd and C. A. Greer. Olfactory Bulb. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.
- [63] M. T. Shipley and M. Ennis. Functional Organization of Olfactory System. *Journal of Neurobiology*, 30(1):123–176, 1996.
- [64] C. Skarda and W. J. Freeman. How Brains Make Chaos in Order to Make Sense of the World. *Behavioral and Brain Sciences*, 10:161–195, 1987.
- [65] N. Sobel, V. Prabhakaran, J. E. Desmond, G. H. Glover, R. L. Goode, E. V. Sullivan, and J. D. E. Gabrieli. Sniffing and Smelling: Separate Subsystems in the Human Olfactory Cortex. *Nature*, 392:282–286, 1998.
- [66] N. Sobel, V. Prabhakaran, Z. Zhao, J. E. Desmond, G. H. Glover, E. V. Sullivan, and J. D. E. Gabrieli. Time Course of Odorant-Induced Activation in the Human Primary Olfactory Cortex. *Journal of Neurophysiology*, 83(1):537–551, 2000.
- [67] S. H. Strogatz. *Nonlinear Dynamics and Chaos: with Applications in Physics, Biology, Chemistry, and Engineering*. Addison-Wesley, Reading, Mass., 1994.
- [68] S. A. Thomas and R. D. Palmiter. Disruption of the Dopamine-Beta-Hydroxylase Gene in Mice Suggests Roles for Norepinephrine in Motor Function, Learning, and Memory. *Behavioral Neuroscience*, 111(3):579–589, 1997.

- [69] M. C. Vanier and J. M. Bower. A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models. *Journal of Computational Neuroscience*, 7(2):149–171, 1999.
- [70] D. A. Wilson, T. C. Pham, and R. M. Sullivan. Norepinephrine and Posttraining Memory Consolidation in Neonatal Rats. *Behavioral Neuroscience*, 108(6):1053–1058, 1994.
- [71] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.
- [72] R. J. Zatorre and M. Jones-Gotman. Human Olfactory Discrimination after Unilateral Frontal or Temporal Lobectomy. *Brain*, 114:71–84, 1991.

Part II

Experimental studies

Introduction

The work described in this section is part of an ongoing series of investigations into synaptic transmission in piriform cortex originally begun by Michael Hasselmo [1]. Hasselmo's work involved the effects of the neuromodulator acetylcholine (ACh) on synaptic transmission in piriform cortex. ACh was chosen because of its known relationship to learning and memory, and because piriform cortex has been postulated to be the site of olfactory memory storage by several groups. The neuromodulator norepinephrine (NE) also has a strong relationship with learning and memory, so the work in this section was done to extend these studies to NE. The results show that NE acts in a very similar way to ACh but has even stronger effects on afferent synapses, leading to a network which is much more heavily driven by afferent inputs than by associational (inter-network) connections.

In addition to its effects on synaptic transmission, NE has a variety of other effects in piriform cortex, including modulation of membrane potential in pyramidal and inhibitory cells, modulation of synaptic facilitation and depression, and modulation of the excitability of pyramidal cells. All of these effects are included in the network model described later,

but there are undoubtedly many more effects of NE which remain to be discovered.

Bibliography

- [1] M. E. Hasselmo, B. P. Anderson, and J. M. Bower. Cholinergic Modulation of Cortical Associative Memory Function. *Journal of Neurophysiology*, 67(5):1230–1246, 1992.

Chapter 2

Synaptic Effects of Norepinephrine in Piriform Cortex

2.1 Abstract

We recorded from a brain slice preparation of rat piriform (olfactory) cortex using extracellular techniques in order to characterize the effects of norepinephrine (NE) on synaptic transmission in this brain region. $25 \mu M$ NE was shown to cause a large decrease in synaptic field potentials in the association fiber layer of piriform cortex (layer 1b), whereas the same dose of NE caused a significant increase in field potential heights in the afferent fiber layer (layer 1a). The concentration dependences of the NE effects were determined in each case. Pharmacological studies indicated that the NE effects in layer 1b are mediated primarily through a presynaptic effect dependent on α -2 adrenergic receptors, while layer 1a effects are

dependent primarily on postsynaptic α -1 adrenergic receptors. NE was also shown to cause an increase in paired-pulse facilitation in layer 1b but not in layer 1a. We consider possible reasons for differences between these results and other investigations in the literature, and discuss the possible functional significance of these modulatory effects.

2.2 Introduction

Norepinephrine (NE) is a monoaminergic neuromodulator/neurotransmitter with a wide distribution throughout the CNS [15, 17]. NE can have a wide range of effects on cortical neurons, including changing neuronal excitability [41], increasing the strength and duration of LTP in the hippocampus [32, 33] and modulating ionic channels [19]. However, its computational roles in the brain remain unclear. There is evidence for involvement of NE in a variety of brain functions including memory [2, 14, 46, 51, 50, 60, 65], arousal [3, 44], and possibly modulation of neuronal signal-to-noise ratios [31, 53].

For the last several years our laboratory has been studying the possible functional effects of neuromodulators within the circuitry of the piriform (primary olfactory) cortex using a combination of physiological [26] and modeling [29] techniques. Piriform cortex has a characteristic three-layered structure that can be investigated directly using brain slice techniques [20, 21, 22, 56] (figure 2.1). Layer 1 consists primarily of afferent fibers to pyramidal neurons coming from olfactory bulb mitral cells via the lateral olfactory tract (layer 1a) and association fibers connecting pyramidal neurons with each other (layer 1b).

Superficial pyramidal cells are located in layer 2, while deep pyramidal and excitatory and inhibitory multipolar cells are located in layer 3. This laminar arrangement makes it relatively easy to independently investigate the synaptic properties of the different fiber systems, since stimulating electrodes placed in layer 1a or 1b selectively activate afferent and association fiber synapses, respectively [26] (figure 2.1).

Both the circuitry of the olfactory cortex and the probable computational requirements of olfactory processing suggest that this structure may implement some form of associative memory [9, 20, 22, 29, 66]. We have previously reported [26] that the neuromodulator acetylcholine (ACh) can transiently decrease the strength of association fiber synapses while having no effect on afferent fiber synapses. In abstract models of piriform cortex [29], we found that this effect could reduce interference between newly stored and old memories during an associative learning process, thus preserving the fidelity of previously stored memories. Since NE is also strongly implicated in at least some memory processes (*e.g.*, [12, 46, 58, 60]) and since piriform cortex receives a significant noradrenergic projection from forebrain neurons in the locus coeruleus (LC) [56, 57], we decided to characterize the effects of NE on synaptic transmission in piriform cortex.

In these studies we demonstrate a differential effect of NE on afferent and association fiber synapses. In agreement with a previous report [31], NE transiently and substantially decreased the size of field potentials in layer 1b, presumably reflecting a decrease in synaptic transmission. In contrast to previously published results, we found that NE caused a




Figure 2.1: A schematic diagram of the piriform cortex and the setup of the recording and stimulating electrodes. Pyramidal cells are shaded gray. Note that the stimulating electrode in layer 1a is relatively far from the recording electrode in layer 1a, while the stimulating electrode in layer 1b is close to the recording electrode in layer 1b. Abbreviations: LOT, lateral olfactory tract. See text for details.

transient and significant increase in field potential amplitudes in layer 1a. Pharmacological studies indicated that the effects on association fiber synapses in layer 1b were primarily mediated by α -2 adrenergic receptors while the effects on afferent fiber synapses in layer 1a were primarily mediated by α -1 adrenergic receptors. In addition, application of NE increased the amplitude of paired-pulse facilitation (PPF) in layer 1b but had no effect on PPF in layer 1a, similar to the effects we previously observed for ACh [26]. In the discussion we suggest possible reasons for the different results obtained here and in a previous investigation [31], and suggest possible computational roles for these effects of NE. A preliminary report on these results has appeared in abstract form [63].

2.3 Materials and Methods

2.3.1 Preparation of brain slices

All experiments were performed on brain slices prepared from female albino Sprague-Dawley rats as described in previous publications [28, 26]. Slices with a thickness of 400 μ m were cut perpendicular to the laminar organization of piriform cortex in the coronal plane, using an oscillating tissue slicer (Vibratome). The location of the piriform cortex was determined visually using a rat brain atlas as a reference [47]. Slices for experiments in layer 1a were taken from the most rostral part of piriform cortex where layer 1a is the thickest. Slices for experiments in layer 1b were usually taken from rostral piriform cortex but occasionally from

more caudal regions of piriform cortex. 1b results from rostral and caudal piriform cortex were essentially identical. Slices were maintained in an artificial cerebrospinal fluid (ACSF) solution (NaHCO_3 26 *mM*; NaCl 124 *mM*; KCl 5 *mM*; KH_2PO_4 1.2 *mM*; CaCl_2 2.4 *mM*; MgSO_4 1.3 *mM*; glucose 10 *mM*) at room temperature for approximately two hours before beginning the experiments. Albumin (0.125 *g/L*) and kynurenic acid (0.66 *mM*) were added to the solution during this time but were not used in the solution bathing the slices during experiments. Albumin was added to minimize protease damage to the slices while kynurenic acid, a nonselective glutamate antagonist, was added to prevent excitotoxicity. We have found that adding albumin and kynurenic acid to the ACSF solution results in healthier and longer-lived brain slices.

Slices were placed in a submersion-type slice chamber on top of a small nylon mesh which kept both sides of the slice exposed to ACSF. ACSF was kept oxygenated with a 95% O_2 /5% CO_2 mixture bubbled through the solution. The flow rate was 4 *ml/min*. The slice chamber included a heating element which kept ACSF maintained at a temperature of 33-35° *C*. Slices were transilluminated, allowing visually guided placement of stimulating and recording electrodes. Slices were left in the chamber for at least 15 minutes before commencing recording in order to wash off all traces of kynurenic acid, which would otherwise have interfered with the recordings due to its effects as a glutamate antagonist.

2.3.2 Preparation and application of pharmaceuticals

All pharmaceuticals were obtained from Sigma Chemical Co. and were freshly prepared before each experiment. Since norepinephrine and some other pharmaceuticals used are light-sensitive, recordings were done in darkness and drug solutions were kept in bottles covered with aluminum foil to block exposure to stray light. In addition, NE oxidizes rapidly when exposed to air, so 25 μM ascorbate was added to the solutions as an antioxidant. Ascorbate was also added to the regular ACSF (without added pharmaceuticals) as a control. Ascorbate by itself had no noticeable effect on the slices except that they stayed healthy for longer periods than slices without ascorbate.

2.3.3 Electrophysiology

Figure 2.1 shows the arrangement of stimulating and recording electrodes. On any given experiment, stimulating electrodes were placed in either or both of layers 1a and 1b. Monopolar or bipolar tungsten stimulating electrodes (Micro Probe Inc.) with an impedance of 1-2 $M\Omega$ were used. Stimuli were low-amplitude (2-10 V) short-duration (0.1 $msec$) voltage shocks, with voltages greater than 5 V being used only in layer 1a. Extracellular field potential recordings were obtained with glass electrodes filled with 3 M NaCl with impedances of $\sim 5 M\Omega$. At the start of each experiment stimulating electrodes were placed either in layer 1a or layer 1b or both as shown in figure 2.1.

For recordings in layer 1a, stimulating electrodes were placed close to the lateral olfactory

tract (LOT) and high in layer 1a (among the myelinated fibers of layer 1 α [23]) to reduce the chance of inadvertently stimulating layer 1b fibers. Recording electrodes were placed 1-2 *mm* from the stimulating electrode to minimize the possibility of inadvertently recording from the layer 1b region adjacent to the stimulating electrode, which, due to the relatively high stimulation voltages necessary to elicit 1a field potentials (typically 5-10 *V*), could respond to some extent to stimulation in layer 1a. Field potentials in layer 1b decrement rapidly with distance (data not shown) so placing the recording electrode some distance from the stimulating electrode in layer 1a minimizes contamination of the layer 1a field potential with layer 1b field potentials. This arrangement is feasible since the myelinated fibers at the surface of layer 1a conduct the stimulus for relatively long distances. In addition, this arrangement reduces the size of stimulation artifacts, which is helpful since layer 1a field potentials are typically somewhat smaller than those in layer 1b.

For recordings in layer 1b, stimulating and recording electrodes were placed in the deepest part of layer 1b to reduce the possibility of inadvertently stimulating and recording from layer 1a as well as layer 1b. Stimulating electrodes in layer 1b were placed very close to the recording electrodes (< 0.5 *mm*) to give the largest signal. Typical field potential heights were 1 *mV* for layer 1b and 0.5 *mV* for layer 1a. The intertrial interval was 15 seconds.

Occasionally, simultaneous layer 1a/1b recordings were made. In order to ensure that the recordings in one layer were not contaminated by artifacts from stimulation in the other

layer, stimulations were staggered in time by 7.5 seconds with respect to each other. These recordings gave identical results to recordings done only in layers 1a or 1b.

All field potentials were allowed to stabilize for at least 15 minutes before recording began. Once field potentials had stabilized, a baseline of 10 minutes (40 trials) was recorded. NE or NE agonists were then applied for 10 or 20 minutes and washed out for 30 minutes to assess recovery from the effects of the treatment. For experiments involving NE antagonists the sequence was: baseline, antagonist only (20 minutes), antagonist + NE (20 minutes). This sequence was done to ensure that the antagonists alone had no effect on field potential amplitudes, which was the case for all antagonists used. Occasionally at the end of an experiment the chamber was perfused with a low-calcium solution (100 μM CaCl₂, 8 *mM* MgSO₄) to eliminate synaptic potentials. This verified that the field potential was in fact due to synaptic transmission.

2.3.4 Data analysis

All data analysis was done using custom-written software. Field potential amplitude was measured in terms of both peak height and initial slope. Since these gave essentially identical results, peak heights were used exclusively in the data analysis. Results of a pharmacological treatment were expressed as the ratio of the average of the final ten trials during the treatment versus the average of the baseline trials. All pharmacological and dose-response treatments were done on at least four slices from at least three different rats (usually con-

siderably more). All results are expressed as the mean \pm the standard error of the mean (SEM).

The dose-response curves were fitted to theoretical curves of the following form:

$$f_{1a}(c) = (1 + x) - \frac{x}{1 + c/K_d} \quad (2.1)$$

$$f_{1b}(c) = x + \frac{(1 - x)}{1 + c/K_d} \quad (2.2)$$

where c represents the concentration of NE and x represents either the maximal (asymptotic) effect of NE expressed as a proportion of the total response (equation 1) or the proportion of the response insensitive to NE (equation 2). Equation 2 represents a single antagonist binding equation with a Hill coefficient of 1.0 [26], and was fit to the data from layer 1b. Equation 1 is identical to equation 2 except that x has been replaced with $1 + x$ to give the proper limiting behavior; this equation was used to fit the layer 1a data. Curves were fit to the data using a nonlinear Levenburg-Marquardt algorithm [48].

2.4 Results

2.4.1 Synaptic effects of norepinephrine

Figure 2.2 shows the effects of 25 μM NE on the height of field potentials recorded in layers 1a and 1b of rat piriform cortex. NE causes a substantial decrease in the height of field potentials in layer 1b ($41.05 \pm 1.88\%$ of baseline, $n = 27$, $p < 0.001$). NE also causes a somewhat smaller increase in the height of field potentials in layer 1a ($124.52 \pm 4.89\%$, $n = 10$, $p < 0.001$).

Figure 2.3 shows the time course of the effects for each layer. The NE-induced decrease in synaptic transmission in layer 1b occurs very rapidly after NE is added to the superfusion medium (figure 2.3). The slight delay seen (less than 3 minutes) is due to the time required for the NE-containing ACSF to travel to the chamber. The NE effects in layer 1a have a somewhat longer latency (up to 4 minutes) and the field potentials rise more slowly to their maximum value. This difference may reflect a slower diffusion of NE to synapses in the myelin-rich region in layer 1a.

2.4.2 Concentration-dependence of effects

The effects of norepinephrine on field potentials in layers 1a and 1b were tested at a wide range of concentrations ranging from 0.2 μM to 500 μM . The effects of NE were tested on from 4-27 slices from at least three different rats per concentration. The results of these experiments are summarized in figure 2.4. The upper curve represents the effects of NE

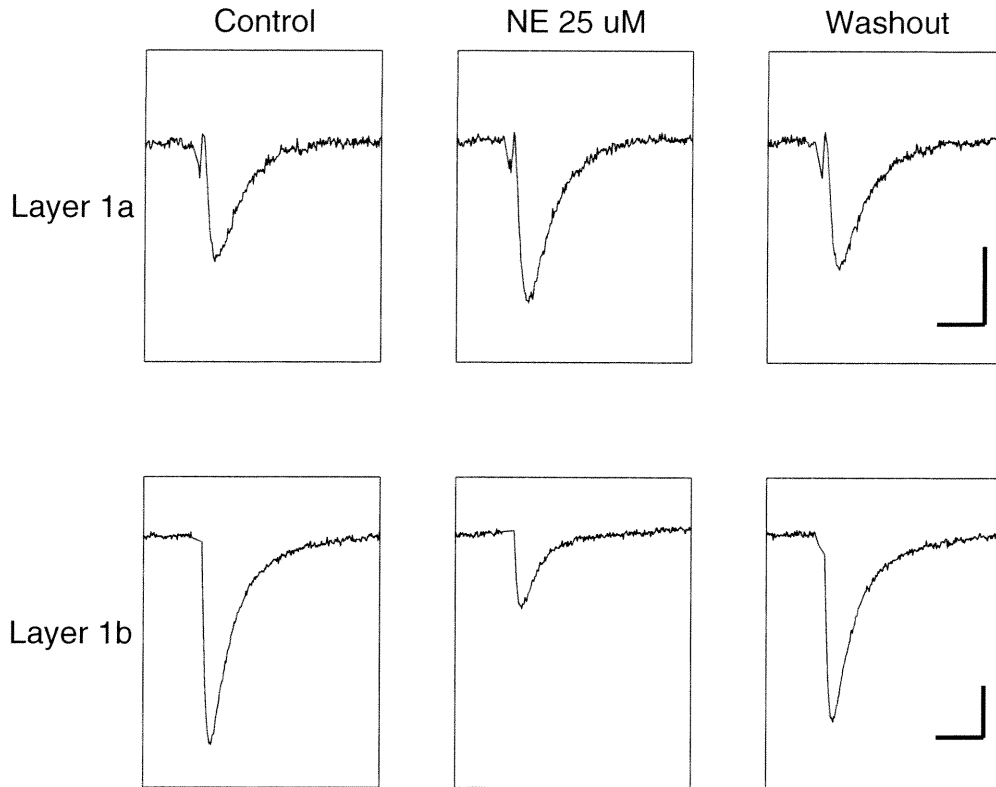


Figure 2.2: The effect of 25 μM NE on extracellular field potentials in layers 1a and 1b of piriform cortex. Stimulation artifacts have been removed. NE causes a pronounced depression in the magnitude of the field potential in layer 1b, and a smaller but significant increase in the magnitude of the field potential in layer 1a. Both effects are reversible. Horizontal bar: 5 msec. Vertical bar: 0.2 mV .




Figure 2.3: The time-course of the NE effect in layers 1a and 1b, relative to baseline. Layer 1a data are displayed with open circles while layer 1b data are displayed with filled circles. The dark bar shows the duration of NE application (10 minutes at 25 μ M).

in layer 1a and the lower curve the effects of NE in layer 1b. The curves represent the optimal fits to the binding equation described in the Methods section, as determined by a nonlinear regression procedure. The asymptotically maximum increase in field potential height in layer 1a due to NE was 23.47% with a dissociation constant K_d of 8.82 μM . The concentration giving 50% of the maximum response according to this curve was 8.8 μM . The asymptotically maximum decrease in field potential height in layer 1b due to NE was 34.02% (this is the component of the response resistant to NE) with a dissociation constant K_d of 4.66 μM . The concentration giving 50% of the maximum response was 4.37 μM .

2.4.3 Pharmacology of effects

We examined the effects of a number of noradrenergic agonists and antagonists in order to determine the likely receptor type(s) responsible for the NE effects. The results, together with the effects of NE alone, are shown in figure 2.5.

In layer 1a, the β -agonist isoproterenol at 25 μM caused a significant but very small rise in field potential heights ($105.22 \pm 1.42\%$ of baseline, $n = 6$, $p < 0.05$). The α -1 agonist phenylephrine (50 μM) reproduced the NE effect on field potentials ($122.83 \pm 5.31\%$ of baseline, $n = 7$, $p < 0.01$). The α -2 agonist clonidine (25 μM) also caused a smaller but significant increase in layer 1a field potential heights ($113.04 \pm 1.28\%$ of baseline, $n = 5$, $p < 0.001$). However, the α -1 antagonist prazosin (2 μM) completely blocked the effect of 25 μM NE in layer 1a, while the α -2 blocker yohimbine (5 μM) had virtually no effect,




Figure 2.4: Dose-response curve of NE effect in layers 1a and 1b, relative to baseline. The curves were calculated using a nonlinear regression procedure described in the Methods section. Error bars represent standard errors of the mean (SEM). 1a data are displayed with open circles while 1b data are displayed with filled circles.

suggesting that the NE effect in layer 1a is primarily mediated by α -1 and not α -2 receptors.

In layer 1b, isoproterenol at $25 \mu M$ caused a significant *increase* in field potential heights ($113.14 \pm 3.65\%$ of baseline, $n = 10$, $p < 0.01$), in sharp contrast to the effects of NE. The α -2 agonist clonidine ($25 \mu M$) caused a decrease in layer 1b field potential heights to $77.85 \pm 2.38\%$ of baseline ($n = 6$, $p < 0.001$). In addition, the α -2 antagonist yohimbine ($5 \mu M$) almost completely blocked the effects of $25 \mu M$ NE ($95.15 \pm 1.38\%$ of baseline, $n = 5$). Interestingly, the α -1 agonist phenylephrine ($50 \mu M$) also caused a slight decrease in layer 1b field potentials to $84.49 \pm 2.49\%$ of baseline ($n = 5$, $p < 0.001$). However, the α -1 receptor antagonist prazosin ($2 \mu M$) was not able to block the NE effect in layer 1b at all. This suggests that the NE effect on layer 1b field potentials is primarily mediated by α -2 receptors.

2.4.4 Effects of norepinephrine on paired-pulse facilitation

Figure 2.6 presents the effects of NE on paired-pulse facilitation (PPF) in piriform cortex. Layer 1a field potentials typically exhibit a large degree of paired-pulse facilitation, whereas layer 1b field potentials show much less PPF [11]. The effects of $25 \mu M$ NE on PPF are shown in figure 2.6 for an interpulse interval of $50 \mu sec$. NE had virtually no effect on PPF in layer 1a (PPF without NE: 1.487 ± 0.109 ; PPF with NE: 1.487 ± 0.055 ; $n = 4$), but caused a significant increase in PPF in layer 1b (PPF without NE: 1.197 ± 0.055 ; PPF with NE: 1.444 ± 0.042 ; $n = 10$, $p < 0.05$). As a result, the depression of field potential




Figure 2.5: Effects of various pharmacological agents on 1a and 1b field potential heights relative to baseline. Error bars represent standard errors of the mean (SEM). Abbreviations: NE, norepinephrine 25 μM ; clo, clonidine 25 μM ; iso, isoproterenol 25 μM ; phen, phenylephrine 50 μM ; NE/praz, norepinephrine 25 μM and prazosin 2 μM ; NE/yoh, norepinephrine 25 μM and yohimbine 5 μM .

height with NE application is considerably reduced for the second pulse. Since PPF is a presynaptic effect [24, 38] this strongly suggests that the effect of NE on association fiber synapses in layer 1b is primarily presynaptic, since if NE was acting postsynaptically both pulses should have been affected equally [24]. The lack of effect of NE on PPF in layer 1a is completely consistent with the α -1 receptor-dependence of the effect, since α -1 receptors are generally located postsynaptically [8, 17].

2.5 Discussion

2.5.1 Differential effects of NE on layer 1a and 1b field potentials

Our results show that norepinephrine causes a pronounced decrease in the height of synaptic field potentials in layer 1b of piriform cortex (figures 2.2 and 2.3). NE decreases the height of field potentials to about one-third of their original height at high doses with a half-maximal effect at $4.4 \mu M$ (figure 2.4). This effect is consistent with previous results we [63, 64] and others [31] have obtained. In contrast, norepinephrine causes a significant increase of about 25% in field potential heights in layer 1a of piriform cortex (figures 2.2 and 2.3). This effect is concentration-dependent with a half-maximal effect at $8.8 \mu M$ (figure 2.4), and has not been previously reported in the literature.

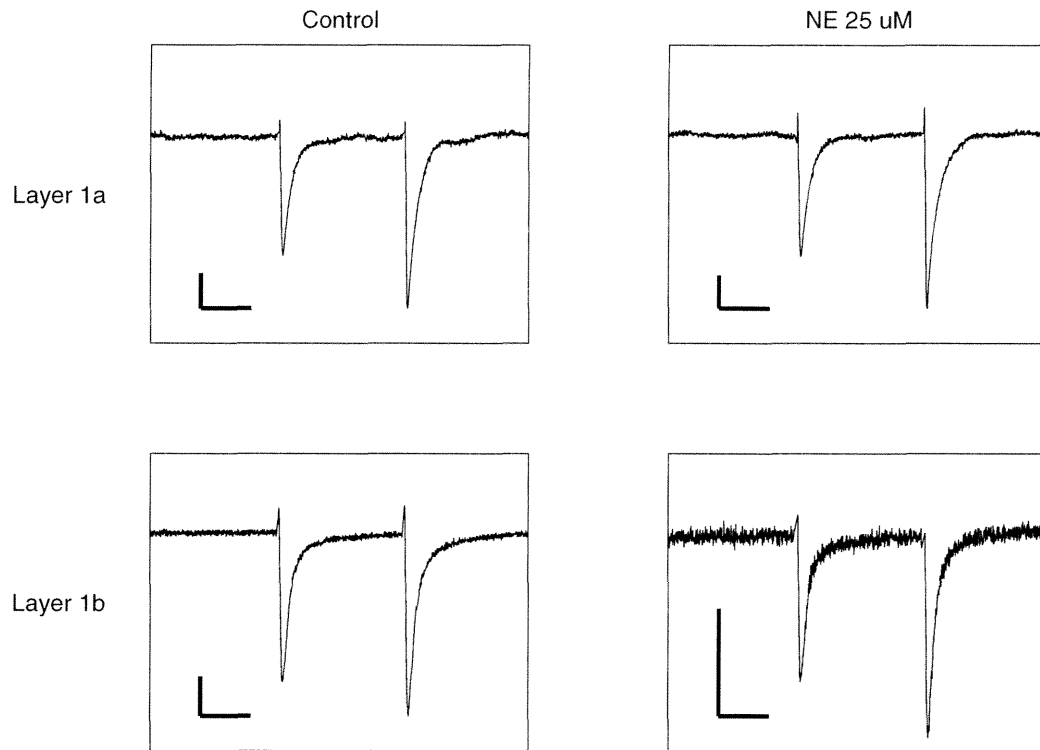


Figure 2.6: Effects of NE on paired-pulse facilitation (PPF) in layers 1a and 1b. The traces have been normalized so that the first pulse is of a constant height for comparison. The horizontal scale bar represents 20 msec while the vertical scale bar represents 200 μV . Pulses were 50 msec apart. Note that NE increases PPF in layer 1b but not in layer 1a.

2.5.2 Pharmacological basis for the NE effects in layer 1a and 1b

The results in figure 2.5 indicate that the effect of NE on layer 1b field potentials is primarily due to NE acting on an α -2 adrenergic receptor subtype. The α -2 agonist clonidine ($25 \mu M$) caused a decrease in field potential height to 77% of baseline. This is considerably less than the full NE effect; however, clonidine is very weakly water-soluble and thus it was difficult to precisely control the concentration of drug delivered to the slice chamber. Therefore, we may be overestimating the amount of clonidine that actually was in contact with the slices. In addition, clonidine is known to be a partial agonist for α -2 receptors [16, 49]. Clonidine at the concentration used ($25 \mu M$) also has some affinity for α -1 receptors [6, 34], which leaves open the possibility of some contribution to the observed effect from α -1 receptors. However, the α -1 agonist phenylephrine even at the relatively high concentration of $50 \mu M$ only caused a small decrease in 1b field potentials, to 84% of baseline. More significantly, the α -2 antagonist yohimbine ($5 \mu M$) almost completely blocked the effects of $25 \mu M$ NE (to 95% of baseline), whereas the α -1 receptor antagonist prazosin ($2 \mu M$) was not able to block the NE effect at all. Thus the contribution to this effect from α -1 receptors, if any, appears to be small.

In contrast to layer 1b, the effect of NE on layer 1a field potentials is primarily due to NE acting on an α -1 receptor subtype. The α -1 agonist phenylephrine ($50 \mu M$) fully reproduced the NE effect (122% of baseline). Clonidine ($25 \mu M$) also caused a smaller but significant increase in layer 1a field potential heights (113% of baseline), which, however,

may be due to cross-reaction with α -1 receptors. The α -1 antagonist prazosin ($2 \mu M$) completely blocked the effect of $25 \mu M$ NE in layer 1a, while the α -2 blocker yohimbine ($5 \mu M$) had virtually no effect, suggesting that the NE effect is primarily mediated by α -1 and not α -2 receptors.

There is no compelling evidence that the NE effects we observed in either layer 1a or layer 1b are mediated to any significant degree by β adrenergic receptors. The β agonist isoproterenol caused a significant *increase* in field potential heights in layer 1b, which is the exact opposite of the effects of NE. However, the size of the increase was fairly small (only about 13% of baseline). It is thus quite possible that the α -2 receptor-mediated suppression of field potentials we postulate are actually stronger than we observed, being masked to some extent by the β receptor-mediated increase in field potential heights.

From the pharmacological data and the paired-pulse facilitation data we also conclude that the NE effects on layer 1a neurons are most likely to be mediated postsynaptically, while the effects in layer 1b are most likely to be presynaptic. The evidence we present here for this conclusion is circumstantial, but plausible. First, α -1 receptors, of the sort implicated in the layer 1a responses, are generally found postsynaptically [52, 59]. α -2 receptors, which appear to mediate the layer 1b effects, are often found in the presynaptic terminal [8, 17, 35, 61, 62] (although α -2 receptors have also been found postsynaptically in some studies, *e.g.* [1, 39, 45]). Secondly, our experimental results from paired-pulse facilitation (PPF) are consistent with this hypothesis. PPF is generally believed to rely

on a presynaptic mechanism [24, 38], and NE has no effect on PPF in layer 1a (proposed postsynaptic receptors) but does affect PPF in layer 1b (proposed presynaptic receptors). Thus the pharmacology and the PPF data both suggest a second fundamental difference between the effects of NE in these two populations of synapses.

2.5.3 Differences from previously reported results

In contrast to the results reported here, Hasselmo et al. [31] reported that application of NE in layer 1a caused no effect on layer 1a field potentials for low concentrations and a slight decrease for higher concentrations. Collins et al. [13] using transverse slices also suggested a concentration-dependent effect of NE, but in their case low concentrations were reported to result in an increase in evoked potentials, while higher concentrations produced a decrease. It is difficult to directly relate the results from different slice preparations; we suspect that the recordings of Collins et al. combined both 1a and 1b field potentials to varying extents, making a direct comparison with our results impossible. The concentrations used in the experiments described here bracketed those used in both of the previous studies.

In our data it is quite clear that the effects of NE on field potentials in layer 1a are more variable than in layer 1b, ranging from a very slight increase to increases of over 200% in some slices. Nevertheless, our data show that the only effect of NE on layer 1a field potentials is an increase in field potential amplitude regardless of the concentration. We found, however, that extreme care must be taken to assure that layer 1a recording conditions

are consistent and optimal. For example, because the size of layer 1a decreases substantially in caudal piriform cortex it is almost impossible to correctly place stimulating and recording electrodes in this region. Improper placement can easily lead to confusion between layer 1a (enhanced) and layer 1b (suppressed) responses. Additionally, even with optimal placement stimulation may spread from layer 1a to layer 1b, making it very difficult to interpret the resulting field potential recordings. For this reason, in the current experiments, all the layer 1a data reported was obtained from rostral slices where layer 1a is much thicker. In addition, care was taken to place the stimulating and recording electrodes reasonably as far apart as possible within layer 1a consistent with getting an adequate signal. In practice, this meant a distance of at least 1 *mm* between these electrodes. Restriction to layer 1a was also verified using paired-pulse facilitation which has previously been shown to be strong in this layer [11]. Typically, when any of these recording conditions were not met, the effects of NE on layer 1a field potentials were greatly diminished or absent.

2.5.4 Functional significance

Piriform cortex is a popular area for computational modeling and is considered by many to be a good candidate for a biological model of associative memory [22, 26, 66]. Neuromodulators such as NE and acetylcholine have long been linked to learning and memory effects (*e.g.*, [7, 14, 46, 50, 60, 65]) although it is still far from clear how these effects are mediated at the level of single cell biophysics or network learning mechanisms. We and others have

demonstrated, however, that these modulators can have strong effects on the behavior of synapses and cells in piriform cortex [4, 5, 18, 26, 30, 25, 27, 31, 40, 42, 43, 54, 55].

Taken in the context of our efforts to build realistic models of the olfactory cortex [10, 36, 37, 66], the differential biophysical effects of both ACh and now NE on the two principle sources of excitatory inputs on the apical dendrites of pyramidal cells can be proposed to have a direct effect on the way in which this network processes incoming sensory data [9]. Specifically, we have proposed that ACh, which suppresses layer 1b association fiber synapses, may serve to make the piriform cortical network more responsive to olfactory afferent sensory inputs than to internal dynamics. We show here that NE both suppresses association fiber synaptic transmission and enhances afferent fiber synaptic transmission. Thus it would appear that NE provides a more extreme form of regulation than ACh although with effects in the same direction. Realistic modeling efforts currently underway in our laboratory will use the data presented in this paper as well as previously published results on ACh and NE to more directly contrast the consequences of these two important neuromodulators on cortical function.

Acknowledgements

We would like to thank Karim Elaagouby for expert advice on experimental technique and slice chamber construction.

Bibliography

- [1] C. Aoki, C. Venkatesan, C. G. Go, R. Forman, and H. Kurose. Cellular and Sub-cellular Sites for Noradrenergic Action in the Monkey Dorsolateral Prefrontal Cortex as Revealed by the Immunocytochemical Localization of Noradrenergic Receptors and Axons. *Cerebral Cortex*, 8(3):269–277, 1998.
- [2] A. F. T. Arnsten and J. D. Jentsch. The Alpha-1-adrenergic Agonist, Cirazoline, Impairs Spatial Working-Memory Performance in Aged Monkeys. *Pharmacology Biochemistry and Behavior*, 58(1):55–59, 1997.
- [3] G. Aston-Jones, C. Chiang, and T. Alexinsky. Discharge of Noradrenergic Locus-Ceruleus Neurons in Behaving rats and Monkeys Suggests a Role in Vigilance. *Progress in Brain Research*, 88:501–520, 1991.
- [4] E. Barkai, R. E. Bergman, G. Horwitz, and M. E. Hasselmo. Modulation of Associative Memory Function in a Biophysical Simulation of Rat Piriform Cortex. *Journal of Neurophysiology*, 72(2):659–677, 1994.

- [5] E. Barkai and M. E. Hasselmo. Modulation of the Input/Output Function of Rat Piriform Cortex Pyramidal Cells. *Journal of Neurophysiology*, 72(2):644–658, 1994.
- [6] B. Blochdaum, A. Korn, M. Wolzt, E. Schmidt, and H. G. Eichler. In vivo Studies on Alpha-adrenergic Receptor Subtypes in Human Veins. *Naunyn-Schmiedebergs Archives of Pharmacology*, 344(3):302–307, 1991.
- [7] A. Blokland. Acetylcholine: a Neurotransmitter for Learning and Memory? *Brain Research Reviews*, 21(3):285–300, 1995.
- [8] A. Borsodi and M. Wollemann. α -adrenergic Receptors. In Abel Lajtha, editor, *Handbook of Neurochemistry*, 2nd. Ed., volume 6, pages 27–38. Plenum Press, New York, 1982.
- [9] J. M. Bower. Reverse Engineering the Nervous System: an *In Vivo*, *In Vitro*, and *In Computo* Approach to Understanding the Mammalian Olfactory System. In S. F. Zorner, T. M. McKenna, and C. Lau, editors, *An Introduction to Neural and Electronic Networks*, 2nd. Ed., pages 27–38. Academic Press, San Diego, 1994.
- [10] J. M. Bower and D. Beeman. *The Book of Genesis*. Springer Verlag, New York, 1995.
- [11] J. M. Bower and L. B. Haberly. Facilitating and Nonfacilitating Synapses on Pyramidal Cells – a Correlation Between Physiology and Morphology. *Proceedings of The National Academy of Sciences of The United States of America*, 83(4):1115–1119, 1986.

- [12] F. Camacho, C. P. Smith, H. M. Vargas, and J. T. Winslow. Alpha(2)-adrenoceptor Antagonists Potentiate Acetylcholinesterase Inhibitor Effects on Passive-avoidance Learning in the Rat. *Psychopharmacology*, 124(4):347–354, 1996.
- [13] G. G. S. Collins, G. A. Probett, J. Anson, and N. J. McLaughlin. Excitatory and Inhibitory Effects of Noradrenaline on Synaptic Transmission in the Rat Olfactory Cortex Slice. *Brain Research*, 294(2):211–223, 1984.
- [14] D. M. Compton, K. L. Dietrich, J. S. Smith, and B. K. Davis. Spatial and Nonspatial Learning in the Rat Following Lesions to the Nucleus Locus-coeruleus. *Neuroreport*, 7(1):177–182, 1995.
- [15] J. R. Cooper, F. E. Bloom, and R. H. Roth. *The Biochemical Basis of Neuropharmacology*. Oxford University Press, New York, 1996.
- [16] T. J. Feuerstein, W. Sauermann, C. Allgaier, E. Agneter, and E. A. Singer. New Insights Into Receptor Theory, as Provided by an Artificial Partial Agonist Made-to-measure. *Naunyn Schmiedbergs Archives of Pharmacology*, 350(1):1–9, 1994.
- [17] M. Fillenz. *Noradrenergic Neurons*. Cambridge University Press, Cambridge, 1990.
- [18] R. L. Gellman and G. K. Aghajanian. Pyramidal Cells in Piriform Cortex Receive a Convergence of Inputs from Monoamine Activated GABAergic Interneurons. *Brain Research*, 600(1):63–73, 1993.

- [19] R. Gray and D. Johnston. Noradrenaline and Beta-Adrenoceptor Agonists Increase Activity of Voltage-Dependent Calcium Channels in Hippocampal-Neurons. *Nature*, 327(6123):620–622, 1987.
- [20] L. B. Haberly. Neuronal Circuitry in Olfactory Cortex – Anatomy and Functional Implications. *Chemical Senses*, 10(2):219–238, 1985.
- [21] L. B. Haberly and J. M. Bower. Analysis of Association Fiber System in Piriform Cortex with Intracellular-Recording and Staining Techniques. *Journal of Neurophysiology*, 51(1):90–112, 1984.
- [22] L. B. Haberly and J. M. Bower. Olfactory Cortex – Model Circuit for Study of Associative Memory? *Trends In Neurosciences*, 12(7):258–264, 1989.
- [23] L. B. Haberly and S. L. Feig. Structure of the Piriform Cortex of the Opossum. 2. Fine Structure of Cell Bodies and Neuropil. *Journal of Comparative Neurology*, 216(1):69–88, 1983.
- [24] E. W. Harris and C. W. Cotman. Effects of Synaptic Antagonists on Perforant Path Paired-pulse Plasticity – Differentiation of Presynaptic and Postsynaptic Antagonism. *Brain Research*, 334(2):348–353, 1985.
- [25] M. E. Hasselmo. Acetylcholine and Learning in a Cortical Associative Memory. *Neural Computation*, 5(1):32–44, 1993.

- [26] M. E. Hasselmo, B. P. Anderson, and J. M. Bower. Cholinergic Modulation of Cortical Associative Memory Function. *Journal of Neurophysiology*, 67(5):1230–1246, 1992.
- [27] M. E. Hasselmo and E. Barkai. Cholinergic Modulation of Activity-Dependent Synaptic Plasticity in the Piriform Cortex and Associative Memory Function in a Network Biophysical Simulation. *Journal of Neuroscience*, 15(10):6592–6604, 1995.
- [28] M. E. Hasselmo and J. M. Bower. Afferent and Association Fiber Differences in Short-term Potentiation in Piriform (Olfactory) Cortex of the Rat. *Journal of Neurophysiology*, 64(1):179–190, 1990.
- [29] M. E. Hasselmo and J. M. Bower. Cholinergic Suppression Specific to Intrinsic not Afferent Fiber Synapses in Rat Piriform (Olfactory) Cortex. *Journal of Neurophysiology*, 67(5):1222–1229, 1992.
- [30] M. E. Hasselmo and J. M. Bower. Acetylcholine and Memory. *Trends In Neurosciences*, 16(6):218–222, 1993.
- [31] M. E. Hasselmo, C. Linster, M. Patil, D. Ma, and M. Cekic. Noradrenergic Suppression of Synaptic Transmission May Influence Cortical Signal-to-Noise Ratio. *Journal of Neurophysiology*, 77(6):3326–3339, 1997.
- [32] W. F. Hopkins and D. Johnston. Frequency-Dependent Noradrenergic Modulation of Long-Term Potentiation in the Hippocampus. *Science*, 226(4672):350–352, 1984.

- [33] W. F. Hopkins and D. Johnston. Noradrenergic Enhancement of Long-Term Potentiation at Mossy Fiber Synapses in the Hippocampus. *Journal of Neurophysiology*, 59(2):667–687, 1988.
- [34] P. Jakala, M. Riekkinen, J. Sirvio, E. Koivisto, and P. Riekkinen. Clonidine, but not Guanfacine, Impairs Choice Reaction Time Performance in Young Healthy Volunteers. *Neuropsychopharmacology*, 21(4):495–502, 1999.
- [35] Y. Kamisaki, T. Hamahashi, C. M. Okada, and T. Itoh. Clonidine Inhibition of Potassium-Evoked Release of Glutamate and Aspartate from Rat Cortical Synaptosomes. *Brain Research*, 568(1-2):193–198, 1991.
- [36] C. Koch and I. Segev. *Methods in Neuronal Modeling*. MIT Press, Cambridge, Massachusetts, 1989.
- [37] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.
- [38] S. W. Kuffler, J. G. Nicholls, and A. R. Martin. *From Neuron to Brain: A Cellular Approach to the Function of the Nervous System*. Sinauer Associates, Sunderland, Massachusetts, 1984.
- [39] A. Lee, D. L. Rosin, and E. J. van Bockstaele. Ultrastructural Evidence for Prominent Postsynaptic Localization of Alpha(2C)-adrenergic Receptors in Catecholaminer-

- gic Dendrites in the Rat Nucleus Locus Coeruleus. *Journal of Comparative Neurology*, 394(2):218–229, 1998.
- [40] H. Liljenstrom and M. E. Hasselmo. Cholinergic Modulation of Cortical Oscillatory Dynamics. *Journal of Neurophysiology*, 74(1):288–297, 1995.
- [41] D. V. Madison and R. A. Nicoll. Noradrenaline Blocks Accommodation of Pyramidal Cell Discharge in the Hippocampus. *Nature*, 299(5884):636–638, 1982.
- [42] G. J. Marek and G. K. Aghajanian. Excitation of Interneurons in Piriform Cortex by 5-Hydroxytryptamine – Blockade by MDL-100,907, a Highly Selective 5-HT_{2a} Receptor Antagonist. *European Journal of Pharmacology*, 259(2):137–141, 1994.
- [43] G. J. Marek and G. K. Aghajanian. Alpha(1B)-adrenoceptor-mediated Excitation of Piriform Cortical Interneurons. *European Journal of Pharmacology*, 305(1-3):95–100, 1996.
- [44] D. A. McCormick, H. C. Pape, and A. Williamson. Actions of Norepinephrine in the Cerebral-Cortex and Thalamus – Implications for Function of the Central Noradrenergic System. *Progress in Brain Research*, 88:293–305, 1991.
- [45] T. A. Milner, A. Lee, S. A. Aicher, and D. L. Rosin. Hippocampal Alpha(2A)-Adrenergic Receptors are Located Predominantly Presynaptically but are also Found

- Postsynaptically and in Selective Astrocytes. *Journal of Comparative Neurology*, 395(3):310–327, 1998.
- [46] M. Ohno, M. Kobayashi, A. Kishi, and S. Watanabe. Working-Memory Failure by Combined Blockade of Muscarinic and Beta-Adrenergic Transmission in the Rat Hippocampus. *Neuroreport*, 8(7):1571–1575, 1997.
- [47] G. Paxinos and C. Watson. *The Rat Brain in Stereotaxic Coordinates*. Academic Press, San Diego, 1997.
- [48] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C, 2nd. Ed.* Cambridge University Press, New York, 1992.
- [49] K. Rasmussen, M. Brodsky, and C. E. Inturrisi. NMDA Antagonists and Clonidine Block C-FOS Expression During Morphine Withdrawal. *Synapse*, 20(1):68–74, 1995.
- [50] S. J. Sara, C. Dyonlaurent, and A. Herve. Novelty Seeking Behavior in the Rat is Dependent Upon the Integrity of the Noradrenergic System. *Cognitive Brain Research*, 2(3):181–187, 1995.
- [51] S. J. Sara, A. Vankov, and A. Herve. Locus Coeruleus-Evoked Responses in Behaving Rats – a Clue to the Role of Noradrenaline in Memory. *Brain Research Bulletin*, 35(5-6):457–465, 1994.

- [52] M. Segal, H. Markram, and G. Richterlevin. Actions of Norepinephrine in the Rat Hippocampus. *Progress In Brain Research*, 88:323–330, 1991.
- [53] D. Servan-Schreiber, H. Printz, and J. D. Cohen. A Network Model of Catecholamine Effects – Gain, Signal-to-Noise Ratio, and Behavior. *Science*, 249(4971):892–895, 1990.
- [54] P. W. Sheldon and G. K. Aghajanian. Serotonin (5-HT) Induces IPSPs in Pyramidal Layer Cells of Rat Piriform Cortex – Evidence for the Involvement of a 5-HT₂-activated Interneuron. *Brain Research*, 506(1):62–69, 1990.
- [55] P. W. Sheldon and G. K. Aghajanian. Excitatory Responses to Serotonin (5-HT) in Neurons of the Rat Piriform Cortex – Evidence for Mediation by 5-HT_{1c} Receptors in Pyramidal Cells and 5-HT₂ Receptors in Interneurons. *Synapse*, 9(3):208–218, 1991.
- [56] G. R. Shepherd, editor. *The Synaptic Organization of the Brain, 4th. Ed.* Oxford University Press, New York, 1997.
- [57] M. T. Shipley and M. Ennis. Functional Organization of Olfactory System. *Journal of Neurobiology*, 30(1):123–176, 1996.
- [58] L. R. Squire. *Memory and Brain.* Oxford University Press, New York, 1987.
- [59] I. Takayanagi and K. Koike. Effects of Aging on Postsynaptic Alpha-1-adrenoceptor Mechanisms in Rat Aorta. *General Pharmacology*, 22(2):211–218, 1991.

- [60] S. A. Thomas and R. D. Palmiter. Disruption of the Dopamine-Beta-Hydroxylase Gene in Mice Suggests Roles for Norepinephrine in Motor Function, Learning, and Memory. *Behavioral Neuroscience*, 111(3):579–589, 1997.
- [61] A. U. Trendelenburg, K. Starke, and N. Limberger. Presynaptic Alpha(2a)-Adrenoceptors Inhibit the Release of Endogenous Dopamine in Rabbit Caudate Nucleus Slices. *Naunyn Schmiedebergs Archives of Pharmacology*, 350(5):473–481, 1994.
- [62] P. Q. Trombley. Noradrenergic Modulation of Synaptic Transmission Between Olfactory-Bulb Neurons in Culture – Implications to Olfactory Learning. *Brain Research Bulletin*, 35(5-6):473–484, 1994.
- [63] M. C. Vanier and J. M. Bower. Differential Effects of Norepinephrine on Field Potentials in Layers 1a and 1b of Rat Olfactory Cortex. *Society for Neuroscience Abstracts*, 18:1353, 1992.
- [64] M. C. Vanier and J. M. Bower. Differential Effects of Norepinephrine on Synaptic Transmission in Layers 1a and 1b of Rat Olfactory Cortex. In *Computation and Neural Systems*, pages 267–271. Kluwer Academic Publishers, Boston, 1993.
- [65] D. A. Wilson, T. C. Pham, and R. M. Sullivan. Norepinephrine and Posttraining Memory Consolidation in Neonatal Rats. *Behavioral Neuroscience*, 108(6):1053–1058, 1994.

- [66] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.

Part III

Matching neural models to experimental data

Introduction

One problem faced by all those who build realistic models of single neurons is that of matching the outputs of the model to experimental data for the same cell type. The reason for this difficulty lies in the highly nonlinear nature of neural models and the relatively large number of poorly constrained parameter values that these models generally possess. The usual way modelers have dealt with this problem is by iteratively adjusting one parameter at a time until the model reproduces the data set. However, because of the nonlinearities inherent in the models, when one parameter is significantly changed many other parameters will usually have to be changed as well. Since this process is tedious and time-consuming, there is a need for automated methods that can do parameter-fitting with a minimum of user intervention. The first of the papers in this section describes a group of such methods implemented by the author within the framework of the GENESIS neural simulator, compares their advantages and disadvantages, and gives suggestions for the most effective use of the methods. The methods developed for this work are now available in a GENESIS library which will be included in the next version of the software. These methods make the

process of fitting single-cell models to data dramatically easier.

One question regarding single-cell models is how to compare different models which have been parameterized on the same data set. Such models might represent models from different investigators or models at different levels of complexity from the same investigator. What is needed is a rigorous quantitative framework for comparing models. The second paper in this section (done in collaboration with Pierre Baldi) uses the Bayesian probabilistic framework to compare models. Using this framework, one can compare individual models to determine which is more probable given the known data. One can also compare entire classes of models, so that one can ask more generic questions, *e.g.*, “If I remove this ion channel from this model, will the resulting class of models be more or less likely to match this data set?” Bayesian methods can also provide a rigorous underpinning for designing matching functions, which are traditionally done in an *ad-hoc* fashion. Therefore, the work in this paper is complementary to the work in the previous one; the first paper addresses the question: “How do I match a model to the data?”, while the second addresses the question “How good are the models obtained by this process?” Together, the two papers are a first step towards a more quantitative methodology of neuronal model building.

Chapter 3

A Comparative Survey of Automated Parameter Search Methods for Compartmental Neural Models

3.1 Abstract

One of the most difficult and time-consuming aspects of building compartmental models of single neurons is assigning values to free parameters in order to make models match experimental data. Automated parameter search methods potentially represent a more rapid and less labor-intensive alternative to choosing parameters manually. Here we compare the performance of four different parameter search methods on several single neuron models. The methods compared are conjugate-gradient descent, genetic algorithms, simulated anneal-

ing, and stochastic search. Each method has been tested on five different neuronal models ranging from simple models with between 3 and 15 parameters to a realistic pyramidal cell model with 23 parameters. The results demonstrate that genetic algorithms and simulated annealing are generally the most effective methods. Simulated annealing was overwhelmingly the most effective method for simple models with small numbers of parameters, but the genetic algorithm method was equally effective for more complex models with larger numbers of parameters. The discussion considers possible explanations for these results and makes several specific recommendations for the use of parameter searches on neuronal models.

3.2 Introduction

Compartmental neural models are being used increasingly in neuroscience to model structures ranging from single synapses and neurons to large networks of neurons [4, 24, 25]. Single neuron models are often required to reproduce very specific features of existing experimental data, such as transmembrane voltage trajectories and detailed neuronal spiking patterns generated experimentally in response to complex stimuli. Typically, such models include every voltage- and calcium-dependent ionic conductance known to exist in a particular neuron type, and their morphologies are based upon anatomical data collected from real neurons.

Ideally, all the constituents of the modeled neuronal structures should be derived di-

rectly from experimental data. In practice, this is usually not possible for several reasons. It is often the case that the kinetics of many of the candidate ionic conductances have not been adequately characterized experimentally. The spatial distribution and channel densities of the included ionic conductances are often unknown, because these values are relatively difficult to measure experimentally. Models may also incorporate intracellular calcium and second-messenger dynamics [25], which may not be sufficiently well characterized to completely constrain the model. For these reasons, realistic neuronal modelers are usually faced with the problem of constructing simulations containing a large number of loosely constrained parameters under expectations that the models will replicate complex experimental data with high accuracy.

Given the complexity of the models and the lack of sufficient experimental data, the traditional way in which parameters are assigned to neural models involves varying parameters iteratively by hand based on a series of educated guesses. However, the large numbers of parameters involved (which is steadily increasing as models become more sophisticated) makes manually parameterizing a model a very time-consuming and tedious exercise. In addition, ionic conductances in neural models often interact with one another in highly non-linear ways [24, 36], making it impossible to accurately adjust parameters independently of each other. Manual parameterization also does not provide any quantification of how closely the model replicates physiological data. Instead, typically several figures showing similar waveforms and spiking patterns has sufficed as proof of biological relevance. In practice,

this means that modelers are always tempted to stop when the model produces a barely adequate fit to the data, instead of trying to determine the best possible match between the model and the data.

For these reasons, we [1, 3, 40] and others [10, 11, 38] have experimented with the use of automated parameter search methods for assigning parameters to neural models. This approach requires the modeler to have a target data set that the model must match. A simulation is run, a numerical match value is calculated, and the parameter search routine automatically chooses the next set of parameters to evaluate based on the previous results and the match value. While, in theory, such a procedure could be iterated over the entire set of possible parameters (a so-called brute force search), in fact, even for simple models, limitations in computer resources make it necessary for the parameter searching software to make what are in effect its own educated guesses as to what parameters to use next. Thus the design of the searching algorithm is critical to the success of the process. Several different search algorithms exist which can in principle be applied to this problem. However, an assessment of the value of one or another automated searching technique can only come from a rigorous side-by-side comparison. In this paper, we specifically evaluate four common searching algorithms which have been applied to neural models: conjugate gradient-descent [3], stochastic search [11], genetic algorithms [1, 10, 40], and simulated annealing [1, 40]. We compare the effectiveness of these methods in obtaining good parameter sets given the same amount of computer resources. Since it is often the case that the effectiveness

of different searching algorithms depends on features of the models whose parameters are being optimized, we compare the performance of these methods on five different single-neuron models ranging from very simple to fairly complex, in order to see what general conclusions can be drawn in order to guide modelers in their choice of parameter searching methods. In the discussion we consider the benefits and limitations of these parameter searching techniques as well as how these techniques can be used most effectively by neural modelers.

3.3 Methods

3.3.1 Simulation environment and data analysis tools

The compartmental neural simulation program GENESIS [4], version 2.0, was used for all of the simulations described here. Parameter search routines and various matching functions were implemented as an extension library in GENESIS. Simulations were run using the implicit Crank-Nicolson method [4, 24] with a time step of 20 μsec . Outputs were saved to disk every 0.1 $msec$. Simulations were run on a variety of workstations running Linux or Solaris 2.3 (Sun Microsystems). Data analysis was done using Matlab (The MathWorks, Inc.) and with a variety of scripts written in the Python computer language (<http://www.python.org>).

3.3.2 Models

Five models were used to evaluate different parameter search methods. For any given model, the parameter search routines were run for a fixed number of simulations and the progress of the search was recorded as the simulation progressed. The purpose of this was to see which search methods gave the best results for a fixed number of simulations. Models 1-4 were used only to test the parameter search methods and are not realistic models of actual neurons, whereas model 5 is based on real data.

Models 1 and 2: simple spiking cell with adaptation

The first two models (known as *simple1* and *simple2* respectively) were derived from a simple one-compartment neuron model with four voltage-dependent ionic conductances: fast Na , delayed-rectifier K , inactivating K (A-current), and a slow non-inactivating K conductance (M-current). The parameters for the first three channels were derived from a model of the hippocampus by Roger Traub [39] while the last was derived from a model of a bullfrog sympathetic ganglion cell [24]. The kinetics of the channels were modified somewhat in order to produce a firing pattern similar to that of regular-spiking pyramidal cells in cerebral cortex [30]. In the first model (*simple1*) the free parameters represented the maximal conductances of the four ionic channels. The second version (*simple2*) also contained four parameters that scaled the $\tau(V)$ curves (time constant as a function of membrane potential) of the activation gating variable of each of the channels. The first

model thus had four adjustable parameters while the second had eight.

Before the searches were started, the initial values of the parameters to be varied were set to random values drawn from a log-uniform distribution centered around the correct values of the parameters. A random value from this distribution is equal to the exponential of a value drawn from a distribution uniformly distributed between the natural logarithms of the endpoints. This distribution allows us to choose random values for parameters with equal probabilities of choosing values of $start/x$ and $start * x$. This is a natural choice for nonnegative scale-dependent parameters such as maximal conductances and scaling values for $\tau(V)$ curves. The range of the distribution was from $\frac{1}{4}$ to 4 times the correct value. We generally chose the ranges for a given parameter by trial and error unless there was experimental data to constrain the range. For models 1 and 2, a range of 4 was sufficient because the initial parameter values were constrained to be no more than a factor of 4 from the correct values. For these models, the goal of the parameter searches was to match the spike timings of the target model's outputs; the shape of the interspike interval waveform was not considered. The inputs were somatic DC current injections (current-clamp) at six different current levels. The reason for using multiple current levels is that in our experience matching a model to data using only a single value of injected current is relatively easy, while matching a model using a wide range of input currents is extremely difficult because of the highly nonlinear behavior of neurons. Thus it is essential to use as much data to constrain the model's behavior as is feasible, and forcing the model to match the data over

a wide range of input currents has proven to be an easy way to strongly constrain the parameter values (see the Discussion).

Models 3 and 4: passive dendrite models

Models 3 and 4 (known as *passive1* and *passive2* respectively) were passive dendrite models, *i.e.*, they had no voltage- or calcium-dependent ionic conductances. Our approach was to choose a fixed morphology for the model and consider the passive electrical properties of the model (membrane axial resistivity (RA), specific membrane capacitance (CM) and specific membrane resistivity (RM)) to be unknown parameters. Since these models were passive and hence linear, we used a single value of current injection (0.1 nA) at one location on the cell as the input stimulus; using different values would have simply multiplied the response by a constant and so would have been redundant. We started the searches at a random point in parameter space and required the search to find the optimal match as with models 1 and 2.

Model 3 was a 100-compartment linear unbranched passive dendrite model. The three parameters were RA , RM , and CM , which were assumed to be constant for all compartments in the cable. Current was injected at one end of the dendrite. We measured the membrane potential at three points along the dendrite and used this as our output data set to be matched to the target model.

Model 4 was a branched dendrite model; in addition to the soma there were four den-

drites with 24 compartments each. We set the target values of the passive parameters RA , RM , and CM to be significantly different for each of the four dendrites and the soma (but constant within a dendrite) in order to generate waveforms at the end of each branch which differed substantially from each other. Input current was injected into the soma. The outputs of the model consisted of the membrane potentials at the soma and those at the distalmost compartments of all four branches. The 15 parameters for this model were RA , RM , and CM for each dendrite and for the soma.

It should be noted that for both passive models we measure the membrane potential at multiple points in the dendrites, which is very difficult to do reliably in real experiments. Similarly, in model 4 the passive parameters were not considered to be fixed over the entire dendrite. Therefore, finding matches for these passive models is considerably more difficult than if the membrane potential had only been measured at a single point (*e.g.*, the soma) and if the passive parameters were considered to be uniform over the entire dendrite. The latter case is much more representative of currently existing passive neuron models (*e.g.*, [19, 20, 35, 41]); however, for this type of model excellent analytical techniques exist to determine parameters [21, 24, 25, 34], and therefore parameter searches are generally unnecessary. Our purpose in using the more complex models is to demonstrate the efficacy of parameter searches on passive models for which no good analytical techniques exist. Even though such models may be beyond the current experimental state of the art, advances in multidendrite recordings (*e.g.*, [6, 28]) will make such models more relevant in the future.

Model 5: layer 2 superficial pyramidal cell in piriform cortex

This model was a simplification of a large model derived directly from an anatomical reconstruction of a layer 2 superficial pyramidal cell from piriform cortex [25]. The original model consisted of 503 compartments and included several voltage- and calcium-dependent ionic conductances. The active conductances of this model were restricted to the soma, because there is currently little evidence for active dendrites in this cell type. We simplified this model to a 5-compartment model using essentially the method of Bush and Sejnowski [5], which involves collapsing branched cables into unbranched cables while preserving total axial resistance. We simplified the calcium dynamics of the original model by keeping the calcium equilibrium (Nernst) potential constant instead of having it vary based on the calcium concentration, and by modeling the intracellular calcium concentration as being derived from a leaky integrator whose input is the current through the calcium channel (suitably scaled to change current to calcium concentration).

The ionic conductances present in the cell included fast Na , delayed-rectifier K , inactivating K (A-current), non-inactivating slow K (M-current), high-voltage-activated Ca , non-inactivating (persistent) Na , and two distinct calcium-dependent K (AHP or after-hyperpolarization) channels. One calcium-dependent K channel has a much smaller time constant ($\tau([Ca])$) than the other; we refer to this channel as the fast AHP channel. Although there is no direct evidence demonstrating this channel's presence in this cell type, we found that including this channel improved the results of the parameter search signif-

icantly (see below). Most of the kinetic parameters were taken either from existing data for piriform cortex superficial pyramidal cells [2, 7, 13, 16], or, in cases where no such data existed, from other cell types such as hippocampal pyramidal cells [12, 31, 39]. Further details on the model can be found in the appendix. The model thus contained eight voltage- or calcium-dependent ionic conductances and a simple calcium buffer, with 23 parameters chosen for fitting. In addition to the parameter types used for models 1-4, model 5 also includes parameters of the following types:

1. voltage offsets for the activation ($m_\infty(V)$ or $m_\infty([Ca])$) curve of an ionic channel gating variable
2. parameters controlling the calcium buffer model
3. voltage offsets for resting membrane potential
4. somatic and dendritic membrane resistivities

Parameters that are offsets from a starting value (*i.e.*, not scale parameters) are chosen from a uniform distribution, as opposed to the log-uniform distribution chosen for the scale parameters.

We added a separate parameter for the resting membrane potential because the values obtained from our experiments included an unspecified junction potential which could be as high as several millivolts [42]. We also considered the specific membrane resistivity of the soma and dendritic regions as separate parameters. Since the reconstructed morphology

we used did not correspond to the cell from which the data was taken, allowing membrane resistivities to vary allows us to compensate for this to some extent.

It is important to note that good parameters for the simplified pyramidal cell model cannot be obtained simply by using the corresponding parameters for the original model for at least three reasons. First, the original model was parametrized using data from sharp-electrode recordings whereas we used data from whole-cell patch-clamp recordings (see below) to parametrize the simplified model. It is well known that passive properties measured in whole-cell recordings can be significantly different from those for sharp-electrode recordings, primarily due to the relative leakiness of the soma membrane in the latter caused by impalement damage [37]. Second, the firing pattern of the cells used to parametrize the original model were quite different from the data used for this model. Finally, the passive structure of the simplified model is clearly different from the original model. However, using the parameters from a complex model as a starting point is a useful strategy to follow in parametrizing a simpler model.

3.3.3 Target data sets

For models 1 to 4, we used artificial data sets generated from the models themselves, *i.e.*, we chose values of the parameters that would give desired behaviors, generated outputs as membrane potential waveforms or lists of spike times given particular inputs, and used those outputs as our data sets. For model 5, the data was derived from brain-slice experiments

on rat piriform cortex. A layer 2 superficial pyramidal cell's responses to varying levels of current injection were recorded with standard whole-cell intracellular recording techniques. We took data from six different values of current injection (spanning the range from a level at which the cell did not spike at all to a level at which the cell spiked fairly rapidly). The data was processed to extract the spike times; this was the data set that our model had to match. The original inputs consisted of one-second current pulses; to speed up our simulations the simulated current injections were limited to 300 msec.

3.3.4 Computing the match between data and models

We used two different match functions to assess how well our models matched our data. For models 1, 2, and 5 (spiking cell models) we processed the membrane potential traces to generate lists of spike times, and then used a match function to compare the spike times of the model with those of the data set. For models 3 and 4 (nonspiking dendritic models) the match function computed the mean squared difference at each time point of the membrane potentials of the model output vs. the target data. The match functions are error functions, with 0 representing a perfect match.

Spike timing match function

This matching function computes a weighted average absolute difference in corresponding spike times between the two models, with some extra penalty terms. The spike timing

match function is as follows:

$$match = \frac{1}{N} \sum_n \left(\sum_i \left[\Theta(i, n) \cdot \frac{|t_{i,n} - o_{i,n}|}{A_{i,n}} + B_{i,n} \right] + C_n \right) ,$$

where $t_{i,n}$ is the time of the i th spike in trace n of the target data, $o_{i,n}$ is the time of the i th spike in trace n of the model data, N is the total number of spikes in all traces, the first sum is over the traces and the second sum is over the spike times in the individual traces.

$\Theta(i, n)$ equals 1 if both traces n have spike i ; otherwise it is 0. The other terms are:

$$A_{i,n} = (\max(0.010, o_{i,n}))^{pow1}$$

$$B_{i,n} = (1 - \Theta(i, n)) \cdot \frac{P}{(TC_n)^{pow2}}$$

$$C_n = Q \cdot \sum_{i=2}^{N-1} h(ISI_{i-1} - ISI_i)$$

Term A weights spike mismatches more heavily the closer they are to the start of the trace (0.010 is in seconds from the beginning of current injection). $pow1$ is a free parameter which can control how strong the weighting is. We used a value of 0.4 for $pow1$. Term B penalizes missing or extra spikes in the simulated trace relative to the target trace. P is a penalty parameter which we set to 0.5 *sec*. TC_n is the total count of spikes in target trace n . $pow2$ is a free parameter which we set to 0.6. For a $pow2$ value of 0, all extra or missing spikes are penalized the same in all cases; for a higher value of $pow2$ the penalty is reduced when there are a large number of spikes in the target trace. This prevents a single extra

spike near the end of a waveform from causing a large penalty when the trace has many spikes already. Term C is a non-monotonicity penalty that penalizes consecutive interspike intervals (ISIs) which are not monotonically increasing (as they are in the target data sets). The function $h(x)$ is zero if $x < 0$ and is x otherwise. The parameter Q is a penalty value for non-monotonically increasing ISIs which we set to 200.0 sec^{-1} .

Waveform match function

The waveform match was calculated as:

$$match = \frac{1}{M \cdot N} \sum_i^M \sum_i^N (t_i - o_i)^2 \text{ ,}$$

where o_i and t_i represent respectively the target and test waveforms sampled at 0.1 msec , N was the total number of points in a each waveform, and M is the total number of waveforms generated by the model.

3.3.5 Parameter search methods

Conjugate gradient

The conjugate-gradient (CG) method [33] is a type of gradient descent algorithm. In this method the gradient of the match function at a given point in parameter space and the direction of steepest descent with respect to this function are computed. A line minimization in the direction of steepest descent is done to find the point on the line with the best match value. Once this is done the procedure is iterated until the method converges to a local

minimum of the matching function. Successive line minimizations are done in directions which are conjugate to one another so that successive minimizations are as nearly independent of each other as possible. In theory this will guarantee that the parameter search will find a local minimum of the objective function. However, we found that when using the spike timing-based match function there were large regions of parameter space which were completely flat, *i.e.*, where the gradient was zero. This generally occurred where no spikes were produced by models whose parameters were in that region. Since no useful gradients could be calculated in these regions, the CG method could not be used directly. In these cases we began by running 100 random simulations before beginning the CG search and using the best match as a starting point. This usually resulted in a point where a gradient could be calculated.

Genetic algorithm

The genetic algorithm (GA) method [15, 18, 26, 32], treats each parameter set as an individual in a large breeding population. A new generation of the population is derived from the preceding generation by reproduction, crossing-over and mutation. This was accomplished by discretizing the parameter values into bit strings and crossing-over and mutating the different bit strings. Parameters were represented as either 8-bit or 32-bit strings (8 bits for models 1 to 4; 32 bits for model 5); we found that the bit resolution had a minimal effect on the results. We took a population of parameter sets and evaluated the fitness of each

one by computing the match as described above. Since the fitness function must increase for better matches, we calculated it by taking either the inverse of the match function (for model 5) or the square root of the inverse of the match function (for models 1-4). The square root was used to reduce the difference in fitness between models and thus has the effect of lowering the selection pressure. This was found to be useful in models 1-4 since there was often a very wide range of match values, and if selection is too strong premature convergence of the algorithm may occur. This was not found to be the case for model 5.

Once fitnesses were calculated, the next generation was determined by reproducing the current generation, with each parameter set being chosen for reproduction in proportion to its fitness (fitness-proportional reproduction). Then a fixed percentage of the resulting parameter sets were crossed over by choosing pairs of parameter sets at random, choosing a breakpoint within the bit string and exchanging the bit strings above the breakpoint (single-point recombination). We crossed over $\frac{2}{3}$ of the parameter strings per generation. Finally, each parameter set was subjected to mutation with a low probability per bit. We used a mutation probability of 0.02 per bit for models 1-4 and 0.01 for model 5. In this way, highly fit parameter sets are selected for and less fit sets are eliminated from the population over a series of generations. Furthermore, the processes of crossing-over and recombination can generate new parameter combinations whose fitness is greater than that of its predecessors.

For models 1-4, we used a population size of 100 and ran the parameter search for 100 generations, giving a total of 10,000 simulations per round. For model 5, we used

a population size of 300 and ran the search for 334 generations giving a total of about 100,000 simulations. These population sizes are quite small for GAs [15, 32], but since the total number of simulations was limited (intentionally) if a larger population size had been used the simulation would have had to have been run for fewer generations. We found that these population sizes gave reasonable results but we did not investigate the population-size dependence of the GA systematically.

Our GA method had two nonstandard features: first, the best parameter set was always preserved unaltered. This was to prevent genetic drift from eliminating good parameter sets and was necessary since the small population sizes used meant that random fluctuations could have an excessively large effect on which parameter sets were propagated into the next generation. Second, if the method went a large number of generations without improving on the best parameter set, the entire parameter table (aside from the best parameter set) was reseeded by choosing random values for all parameter sets. This provided an additional way to keep the method from getting stuck in a suboptimal region of parameter space.

Simulated annealing

We used a continuous version of simulated annealing (SA) [23] adapted from Press et al. [33]. This algorithm constructs an N -dimensional figure known as a simplex out of $N + 1$ points in the parameter space. In the deterministic form of the algorithm, the simplex is moved through the parameter space based on the calculated match values at the vertices of

the simplex. Roughly speaking, successive points in the parameter space are evaluated and accepted if their matches are better than that of the worst point on the simplex (in which case the worst point is discarded and the new point joins the simplex) or rejected. This process continues until a local minimum is reached. In the noisy version, the match values for each point are modified with noise which varies linearly with a temperature parameter which is set by an annealing schedule. The temperature parameter is a dimensionless value which starts out very high (1.0) relative to typical match values and decreases to zero; therefore, the amount of noise starts out large and gradually reduces to negligible levels. Thus, parameter sets with high match values (poor matches) can be accepted into the simplex with a nonzero probability which decreases to zero as the temperature falls, allowing the algorithm to escape from local minima. The annealing schedule was a simple proportional decrease in temperature every N simulation iterations, where N varied from 5 for the simple spiking models to 300 for the pyramidal cell model. We found the slower annealing schedule necessary to give good matches for the pyramidal cell model. With large amounts of noise, the simplex essentially performs a random walk in the parameter space, and as the noise level decreases the simplex hones in on regions of the parameter space with better match values. As noise is reduced to zero this method finds a local minimum in the parameter space, a useful property not shared by the other methods (nor by all simulated annealing algorithms).

The only difference between our procedure and that of [33] was that in the latter the

parameter space was infinite whereas we enforced wraparound boundary conditions, so that if the algorithm tried to use a parameter value that was beyond the assigned limits of the parameter space, the value was adjusted to come in from the opposite side of the parameter space by the amount of the overshoot. This was necessary since otherwise the algorithm would often set parameters to extremely small or large (*i.e.*, unphysiological) values in order to achieve very small improvements in the match. This also allowed us to enforce our constraint of keeping parameter values within specified limits without simply truncating parameter values at the limits.

If the SA algorithm converged before the parameter search run had used up all of its allocated simulations (which was always the case given the annealing schedules we used) the best results of the search were saved and the search was restarted from scratch. When plotting the results, we always give the best results of the run achieved up to that point, even when there was more than one search in the run. This was also true for the conjugate-gradient (CG) search method described above.

Stochastic search

The stochastic-search (SS) method [11] works as follows. A random starting point in parameter space is chosen. New points are selected from a multidimensional Gaussian distribution centered on the starting point with a given (initially large) variance. These new points are evaluated, and if one is found that is a better match than the previous best match, the

Gaussian distribution is moved so that it is centered on the new (best) point. In addition, the variance of the distribution decreases linearly each simulation until it reaches some minimum value, at which time it is increased to the original value. The variance value represents a proportion of the range of the parameter. We used an initial variance of 0.5, a final variance of 0.05, and a variance-contraction of 0.95 per simulation for all models except the pyramidal cell model, where we used a variance-contraction of 0.98 in order to search parameter space more thoroughly. The variance is thus analogous to the temperature parameter in the SA method, except that it goes through successive contraction/expansion cycles instead of a single contraction to zero. In principle, choosing points from a distribution with a large variance will allow the modeler to search over large regions of parameter space, while choosing points from a distribution with a much smaller variance will allow the modeler to find locally optimal points in small regions of parameter space. Unlike our SA algorithm, however, the SS algorithm does not do a systematic search for local minima.

Random search

As a control on the performance of the other parameter search methods, we used a simple random search method as follows. For each simulation a new parameter set was chosen at random from the parameter space being searched (which always had a finite volume). The best parameter set was always saved, and the performance of the random search method equaled the best match obtained after a given number of simulations.

3.3.6 Statistical analyses

For each pair of parameter search method and model, we performed 10 separate parameter searches and recorded the best match values for each search. To compare the different methods, we computed the mean and 95% confidence intervals of the log transform of the best match values. The effect of the log transformation was to change the range of the matches from $[0, \infty]$ to $[-\infty, \infty]$. This was necessary since otherwise the confidence intervals would have in some cases included negative match values, which are inadmissible. For the time-course plots we plotted the arithmetic mean of the match values at each time point without log transformation. We separately calculated the geometric mean of the match values, which are equivalent to the exponential of the arithmetic mean of the log-transformed data. The results were qualitatively identical, with one exception: for the CG method in model 3 the geometric mean is lower than that of the GA method, while the arithmetic mean is higher; the difference is not significant in either case.

In two cases (the simulated annealing results of models 1 and 2) matches of zero were sometimes obtained, which made it impossible to log transform the matches. In these cases we separately calculated the significance of the difference between the simulated annealing results and those of the other methods using the nonparametric Mann-Whitney one-tailed U test [22]. For plotting purposes we replaced the zero match values by 10^{-8} , which gave relatively large confidence limits.

3.3.7 Resource requirements

For each of the parameter search methods, the computational overhead (in both space and time) of the search algorithm was negligible in comparison to that of the neural simulation itself. Therefore, we can compare the performance of the different methods solely by how good a match is obtained with a fixed number of simulation iterations. Thus, each run of each search method consisted of the same number of simulations for a given model. The time needed for individual simulations for all the parameter searches ranged from about five seconds per simulation for the simpler models to about 30 seconds per simulation for the more complex models (this was also highly dependent on the speed of the processor used). For models 1-4, each run of a parameter search comprised 10,000 simulations; 10 parameter search runs were performed for each model and each method in order to generate statistics. Thus we performed 4 (*models*) \ast 5 (*methods*) \ast 10 (*repeats*) \ast 10000 (*simulations per search*) = 2,000,000 simulations in all for these four models. For model 5 we found that we could not obtain reasonable results with 10,000 simulations per run, so we used 100,000 simulations per run instead. For this model we thus performed $5 \ast 10 \ast 100000$ or 5 million simulations in all.

3.4 Results

3.4.1 Simple spiking models

A typical result of a parameter search on model 2 (eight parameters) is shown in figure 3.2. This model was discovered by the genetic algorithm method. In the figure the target data is the bottom trace on each graph and has been offset to facilitate visual comparison. Note that the membrane potential trace of the model found by the search is nearly identical to that of the target model. The parameters of the model are listed in table 3.1 along with the best parameter sets found by the other search algorithms. The best parameters for the SA and GA methods are very close to the (known) parameters of the target model, while those of the other methods in some cases diverge considerably from the correct values. Even though the simulated annealing method gave better matches than the genetic algorithm method for model 2, the matches from both methods are sufficiently good that they would be difficult to distinguish by eye.

Figures 3.1 and 3.3 show how the searching methods compare for models 1 and 2. Each round consisted of 10,000 searches and each line in figure 3.3 is the average of ten rounds. For both models the simulated annealing method was superior to the other methods both in terms of how good a match was finally achieved and how quickly this match was achieved. Random search was predictably the poorest method with the other three methods ranking in between. The best (*i.e.*, lowest) match obtained for the SA method for both models was

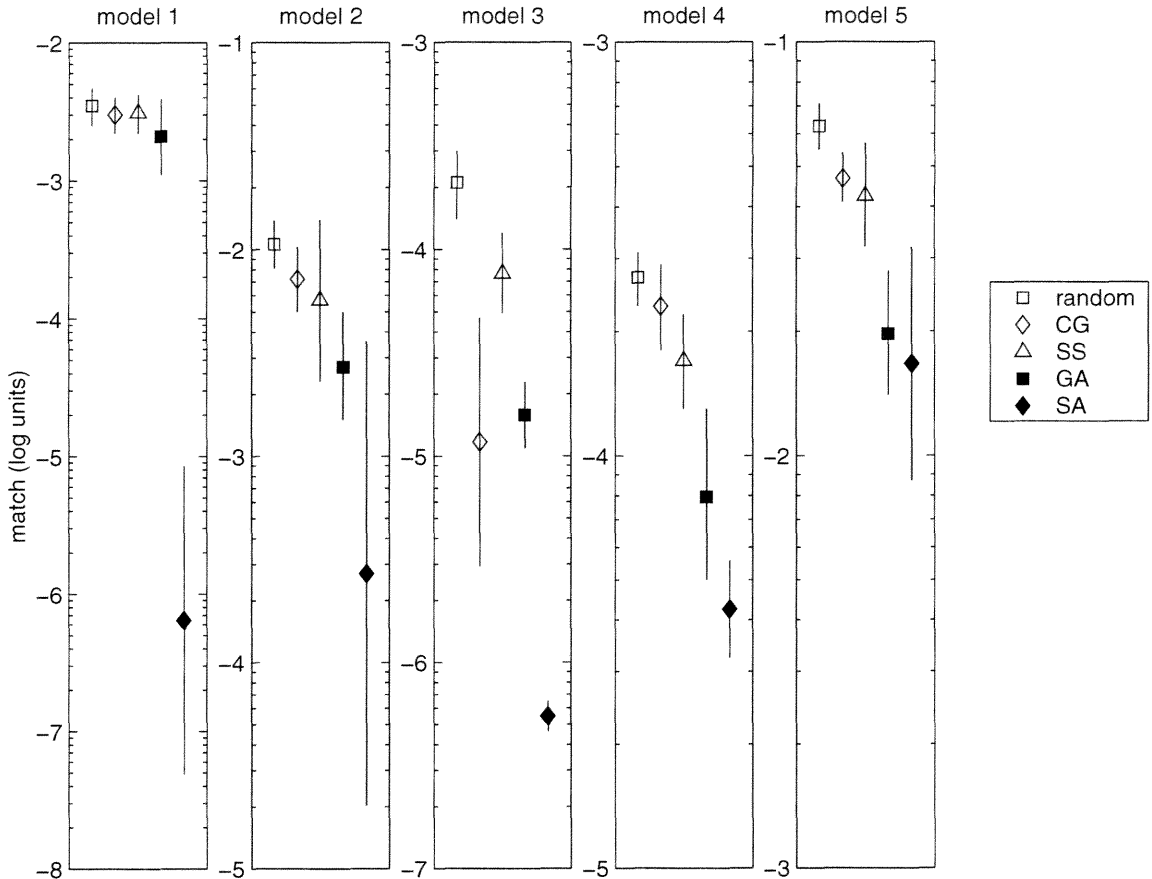


Figure 3.1: Bar graph of the results of the parameter searching methods on all five models. The y-axis represents the match values in log units (base 10). Error bars represent 95% confidence intervals. Means are of the best matches after 10,000 simulations (models 1-4) or 100,000 simulations (model 5). The large error bars on the SA results for models 1 and 2 are due to the presence of several match values of zero in the final results; however, a one-tailed Mann-Whitney U test confirmed that the SA match values are significantly smaller than those of the other methods in both cases ($p < 0.01$). Note the different y-axis scales for the different 5 graphs.

parameter	description	target	SA	GA	SS	CG	random
1	\bar{G}_{Na}	800	805	829	837	505	717
2	$\tau(V)_{Na}$	1	1.02	0.83	0.89	0.27	1.16
3	$\bar{G}_{K(DR)}$	100	100	99.1	92.8	106	67.6
4	$\tau(V)_{K(DR)}$	1	0.99	1.07	0.97	1.50	1.36
5	$\bar{G}_{K(A)}$	50	49.8	50	53.1	26.4	39.9
6	$\tau(V)_{K(A)}$	1	0.99	1.37	1.23	0.84	0.91
7	$\bar{G}_{K(M)}$	100	100	95.4	93.8	99.4	90.8
8	$\tau(V)_{K(M)}$	1	1	1.08	1.01	1.05	1.21
	match	0	0	8.88e-04	2.46e-03	3.21e-03	6.51e-03

Table 3.1: The target parameters for model 2 (simple spiking model) compared with the best parameters found by all search methods used. Maximal conductances are in S/m^2 and scaling factors for time constants (τ values) are dimensionless. All time constants refer to voltage-dependent activation gates. Na is a fast inactivating sodium channel; $K(DR)$ is a delayed-rectifier potassium current; $K(A)$ is an inactivating potassium current; $K(M)$ is a slow non-inactivating potassium current. Note that the matches found by the SA and GA methods are very close to the correct parameter values.

0.0 (a perfect match). This means that each spike in the model found by the search matched the time of the target model within the resolution of the simulation output, which was 0.1 *msec*. Searches on the eight-parameter model generally resulted in higher match values than those for the four-parameter model. This reflects the larger size of the parameter space and thus the greater difficulty of the search task. Interestingly, the GA method gave similar match values for both models; this may reflect the fact that GA methods are less sensitive to the dimensionality of the problem (see Discussion).

The statistical distribution of the best match results is displayed graphically in figure 3.1. For model 1, the difference between the non-SA methods is not significant; furthermore, all methods (even random search) were able to find reasonably good matches after 10,000 simulations. For model 2, simulated annealing is again the most effective method, but the SA results have a wide confidence interval because of a single round where a zero match was achieved. We used the Mann-Whitney one-tailed U test to confirm that the SA results are in fact significantly smaller than those of the other methods ($p < 0.01$).

3.4.2 Passive cable models

See figure 3.4 for typical outputs of model 4, with 15 parameters. In this figure the target data is plotted using dashed lines while the model found by the parameter search is plotted using solid lines. The target data has been offset by -2 *mV* for clarity; if this had not been done the models would overlap completely. The outputs from model 3 were equally

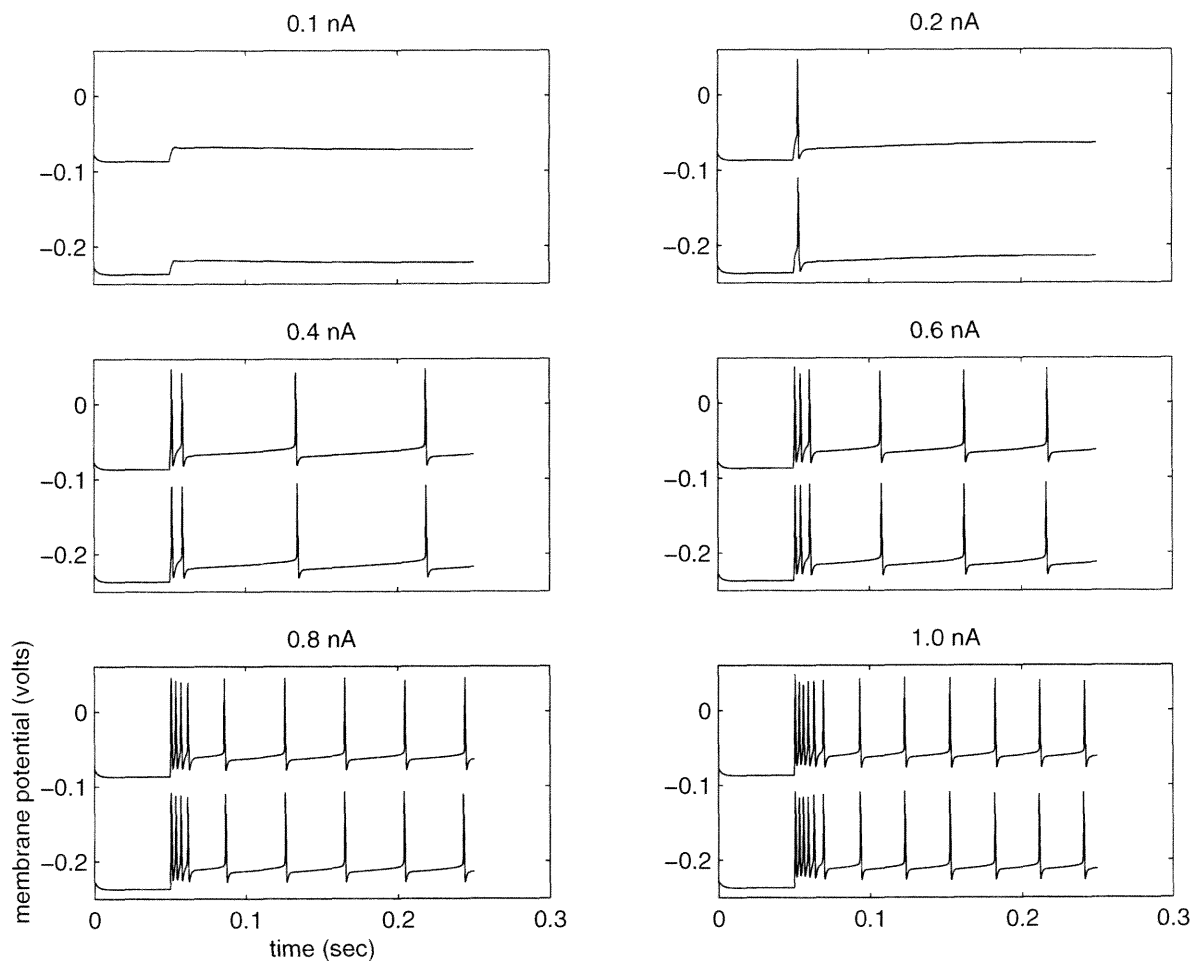


Figure 3.2: Target outputs of the simplified one-compartment spiking model. The graphs show the membrane potential changes in response to current steps of various amplitudes, from 0.1 nA to 1.0 nA . The lower traces represent the target data, while the upper traces represent the results of a genetic-algorithm parameter search using 8 parameters (model 2). The lower traces have been offset by -150 mV for clarity.




Figure 3.3: Left graph: comparison of the performance of the parameter searching methods on the simple spiking model with four parameters (model 1) as a function of simulation number. Right graph: the same comparison for the eight-parameter version of the simple spiking model (model 2). All traces are the averages of 10 runs of 10,000 simulations each.

impressive. In table 3.2 we show the parameters of the target model versus those of the model found by the parameter search for model 4. It is extremely interesting that, unlike the case with the simple spiking model, the parameter values found for this model were in many cases quite different from the correct parameters even for models which match the data very closely. This indicates that in this model there can be many optimal models based on our criteria. The capacitance values in the best parameter set (found using the SA method) were generally quite close to the correct values, whereas the resistance parameters often varied considerably from the corresponding correct values. This is not surprising since the time course of the voltage waveforms is critically dependent on the capacitance values, whereas the various resistance values can compensate for each other to some extent (*e.g.*, if one dendrite has a higher-than-average input resistance, the other can have a lower-than-average one).

Figure 3.5 demonstrates how well the different search methods work as a function of simulation number for the passive models, while figure 3.1 shows the final results of the parameter search methods along with the range of the results. Simulated annealing still outperformed the other methods, with genetic algorithms second as before. However, for model 3 the conjugate gradient (CG) method significantly out-performed the stochastic search (SS) method ($p < 0.05$) and performed as well as the genetic algorithm method. This presumably reflects the fact that for passive models with small numbers of parameters (such as model 3) the parameter space is relatively smooth, so a gradient-descent method

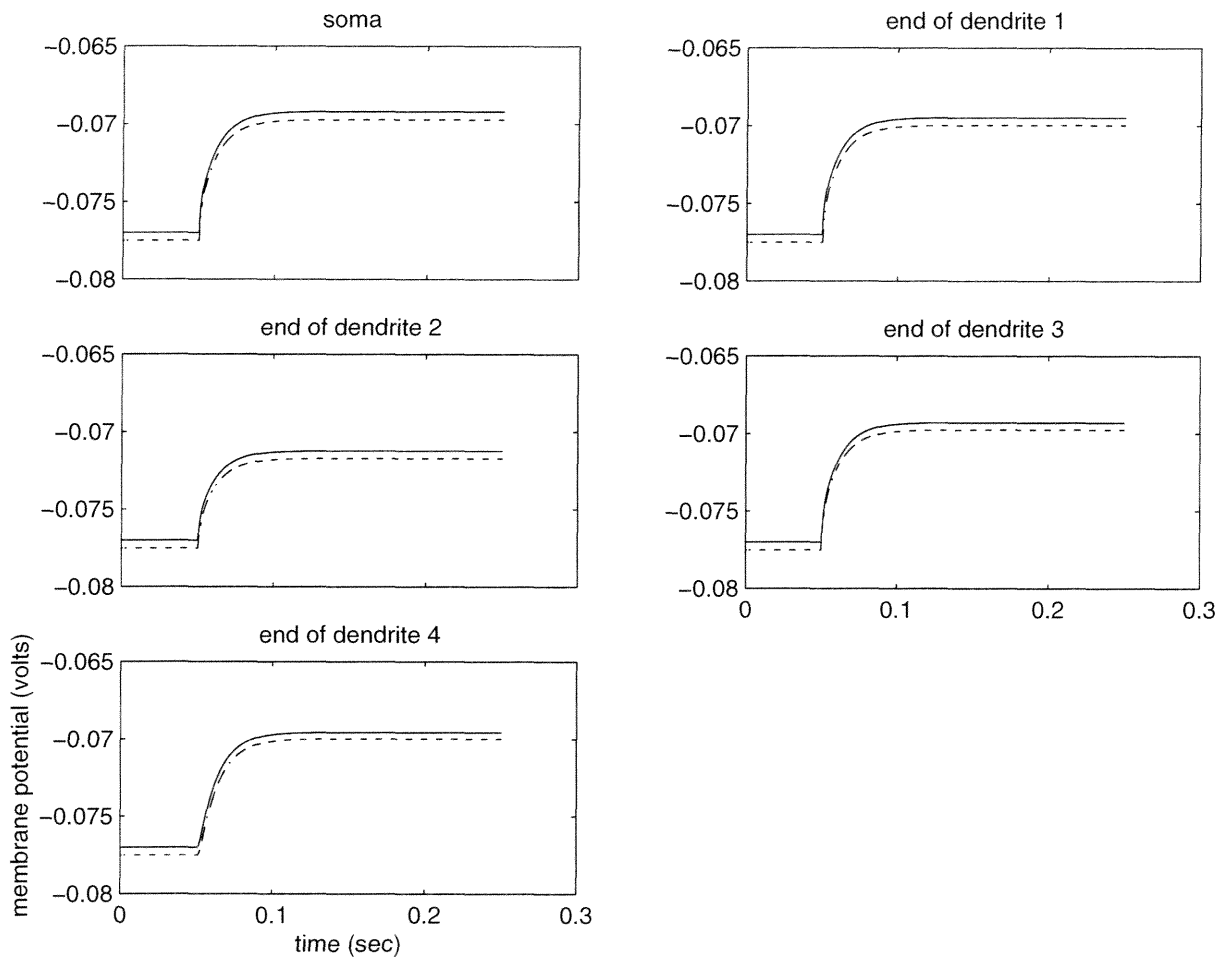


Figure 3.4: Target outputs of the passive dendritic model with 15 parameters (model 4) compared to the results of a simulated-annealing parameter search. The target data is indicated by the broken line and the results of the parameter search are indicated by the solid line. The target data is offset by -5 mV for display purposes; otherwise the traces would overlap entirely.

parameter	description	target	SA	GA	SS	CG	random
1	soma RM	0.2	0.241	0.778	0.458	0.309	1.34
2	soma RA	6.0	0.382	5.47	1.76	0.337	5.51
3	soma CM	0.008	0.0174	0.00998	0.00078	0.0137	0.0186
4	dendrite 1 RM	2.0	1.02	1.12	2.05	20.9	1.06
5	dendrite 1 RA	6.0	3.01	3.06	4.14	10.9	0.679
6	dendrite 1 CM	0.008	0.0101	0.016	0.0123	0.00634	0.0152
7	dendrite 2 RM	0.2	0.169	0.0977	0.225	0.146	0.142
8	dendrite 2 RA	6.0	5.04	2.9	6.00	4.36	3.87
9	dendrite 2 CM	0.008	0.0103	0.0185	0.0266	0.0233	0.00683
10	dendrite 3 RM	2.0	3.56	13.8	0.25	2.12	0.371
11	dendrite 3 RA	0.6	0.915	0.188	0.016	0.791	0.0952
12	dendrite 3 CM	0.008	0.0057	0.00323	0.004	0.00363	0.0255
13	dendrite 4 RM	2.0	2.54	4.73	2.47	0.729	1.08
14	dendrite 4 RA	6.0	7.08	7.98	6.23	5.46	5.35
15	dendrite 4 CM	0.08	0.0684	0.0654	0.0718	0.0757	0.0525
	match	0.0	2.14e-05	3.23e-05	7.38e-05	1.34e-04	1.80e-04

Table 3.2: The target parameters for model 4 (passive dendritic model with four branches) compared with the best parameters found by all search methods used. RM values are in units of $\Omega \cdot m^2$; RA values are in units of $\Omega \cdot m$, and CM values are in units of F/m^2 . Note the relatively large spread of parameter values around the correct values.

like the CG method can be reasonably effective. Interestingly, when we look at model 4, which is more complex, the CG method gives slightly worse results than the SS method (although this is not statistically significant). This may reflect the larger number of models with equivalent match values for this model, *i.e.*, the larger number of local minima for the CG method (but not the SS method) to get trapped in.

3.4.3 Pyramidal cell model

Figure 3.6 shows the outputs of a parameter searching method for the pyramidal cell model. This model was discovered by the SA method after about 48,000 simulations. After this the parameter search converged and subsequent searches in this run were completely new parameter searches which converged to higher match values; in fact, this represents the best SA result we achieved. The parameter values for this model are listed in table 3.3. In this case we do not know what the correct parameter values are so we cannot assess the performance of the methods on this criteria. However, figure 3.6 shows that the methods can achieve extremely good matches to the data over a wide range of input currents, at least with respect to the spike timings. The major difference between the model and the data is in the spike pre- and afterpotential, suggesting that other conductances or features need to be added to the model if greater accuracy is needed or if the spike afterpotential is of interest to the modeler. One interesting possibility would be to add a separate spike-initiation zone to the model [29]. An interesting feature of this model is that we added a fast calcium-




Figure 3.5: Left graph: comparison of the performance of the parameter searching methods on the linear dendritic model with three parameters (model 3) as a function of simulation number. Note the relatively good performance of the CG method. Right graph: comparison of the performance of the parameter searching methods on the linear dendritic model with 15 parameters (model 4) as a function of simulation number. All traces are the averages of 10 runs of 10,000 simulations each.

dependent potassium channel to the model, even though there is no experimental evidence for such a channel, because we found that adding this channel substantially improved the matches we obtained (see below and Discussion). Figure 3.7 shows a graph of the results of the different methods as a function of simulation number, while figure 3.1 shows the statistical distribution of the best match results. Note that for this model the genetic algorithm method was almost exactly as effective as the simulated annealing method. The CG and SS method are roughly equivalent, while (as expected) the random search method is the least effective.

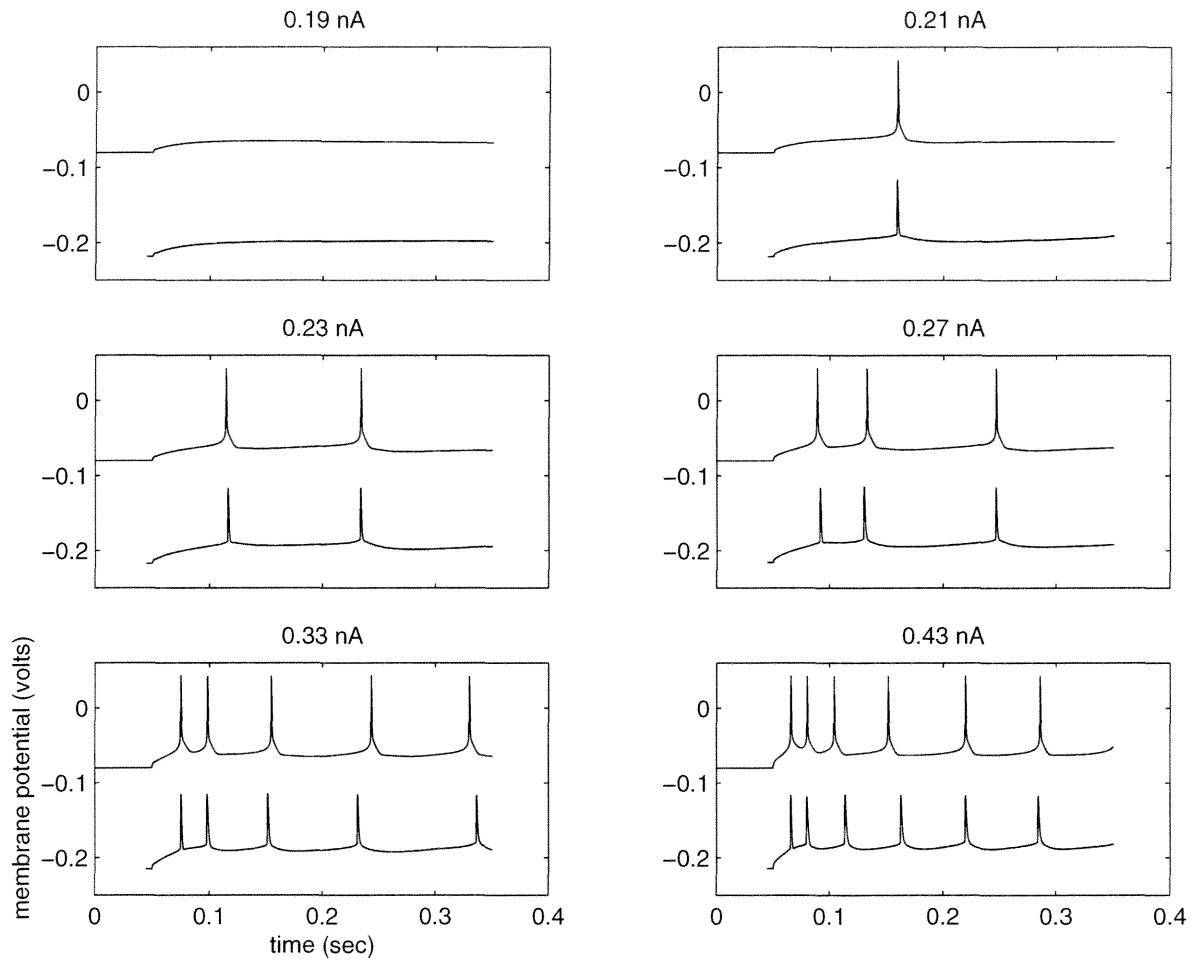


Figure 3.6: Outputs of a layer 2 pyramidal cell from piriform cortex at six different levels of current injection at the soma compared with the outputs of a simulated-annealing parameter search. The experimental data is shown below the model data in each trace and is offset by -150 mV for display purposes. Although the model did not precisely reproduce every aspect of the interspike interval waveforms the match between the spike times of the model compared to the data is extremely good.




Figure 3.7: Comparison of the performance of the parameter searching methods on a model of a layer 2 pyramidal cell from piriform cortex with 23 parameters (model 5) as a function of simulation number. All traces are the averages of 10 runs of 100,000 simulations each.

parameter	description	range	original	SA	GA	SS	CG	random
1	\overline{G}_{Na}	4	2800	3400	1400	1240	10600	718
2	$\overline{G}_{Na(pers)}$	4	65	129	160	153	125	64.4
3	$\tau(V)_{Na(pers)}$	4	1	1.09	1.29	1.65	1.74	0.421
4	$\overline{G}_{K(DR)}$	4	270	444	799	1080	917	168
5	$\overline{G}_{K(A)}$	4	200	798	107	800	52.3	54.8
6	$\tau(V)_{K(A)}$	4	1	1.61	0.310	3.86	0.540	0.597
7	$\tau_i(V)_{K(A)}$	4	1	2.06	1.58	0.324	1.96	2.20
8	$\overline{G}_{K(M)}$	10	100	112	144	22.3	59.1	106
9	$m_\infty K(M)$	10.0	0	9.0	6.0	-3.67	-6.84	7.0
10	$\tau(V)_{K(M)}$	10	1	0.414	0.575	0.422	0.750	0.739
11	$\overline{G}_{K(Ca,s)}$	10	4	20.9	17.2	22.0	7.75	9.02
12	$\tau(Ca)_{K(Ca,s)}$	10	1	2.38	0.994	4.49	3.31	5.46
13	$\overline{G}_{K(Ca,f)}$	10	20	20.9	16.4	6.55	47.2	70.2
14	$m_\infty K(Ca,f)$	3.0	0	2.27	0.774	2.98	2.25	1.44
15	$\tau(Ca)_{K(Ca,f)}$	10	1	0.892	1.00	0.681	6.150	4.18
16	\overline{G}_{Ca}	10	10	6.63	1.86	9.16	24.5	1.32
17	$\tau(V)_{Ca}$	4	1	0.285	0.334	1.43	3.11	2.21
18	$\tau_i(V)_{Ca}$	4	1	0.781	1.24	0.480	1.07	0.434
19	Ca buffer B	4	1.83	1.21	1.12	0.955	0.527	2.78
20	Ca buffer τ	4	25.0	53.3	75.8	69.5	25.3	44.0

continued on next page

<i>continued from previous page</i>								
parameter	description	range	original	SA	GA	SS	CG	random
21	soma R. P.	5.0	-73.0	-74.0	-76.3	-70.2	-74.9	-70.4
22	soma <i>RM</i>	2.0	0.50	0.411	0.707	0.676	0.462	0.259
23	dend <i>RM</i>	2.0	0.50	0.440	0.540	0.255	0.519	0.502
	match			6.15e-3	9.04e-3	0.0238	0.0281	0.0458

Table 3.3: The target parameters for model 5 (pyramidal cell model) compared with the best parameters found by the different search methods. Maximal conductances (\bar{G}) are in S/m^2 . Scaling factors for time constants τ are dimensionless. Offsets of m_∞ curves are in mV (for $K(M)$ channel) or μM (for $K(Ca)$ channels). τ_i values refer to the time constant of inactivation gates for a channel; all other τ or m_∞ values refer to activation gates. Abbreviations for channels are as follows: *Na*: fast sodium; *Na(pers)*: persistent sodium; *K(DR)*: delayed-rectifier potassium; *K(A)*: inactivating potassium; *K(M)*: slow noninactivating potassium; *K(Ca, s)*: slow *Ca*-dependent potassium; *K(Ca, f)*: fast *Ca*-dependent potassium; *Ca*: voltage-dependent calcium. *Ca* buffer *B* values are in units of $M/(A \mu sec)$. *Ca* buffer τ values are in values of *msec*. Soma R. P. = somatic resting potential (in *mV*). Soma/dend *RM*: specific membrane resistivity value for the soma or the dendrites, respectively.

As mentioned above, we added a hypothetical fast AHP channel to give better matches between the data and the model. As a control, we also ran a series of parameter searches without the fast AHP channel using the genetic algorithm method (the SA method could also have been used but was not). Figure 3.8 compares the results of the GA simulations with and without the fast AHP channel. As can be seen, the addition of the fast AHP channel improved the match value considerably. The best model we obtained using a parameter search without a fast AHP channel is shown in figure 3.8. Although this model matches the data fairly well, it is markedly inferior to the model in figure 3.6, which did include the fast AHP channel. The difference between the results of the GA searches for the model without the fast AHP channel and the model with the fast AHP channel is statistically significant ($p < 0.05$, two-tailed t -test).

3.5 Discussion

3.5.1 Method comparisons

There are two major conclusions that can be drawn from the results presented. First, parameter search methods can be extremely effective in finding very good matches between single-neuron models and a target data set in a wide variety of cases. Second, not all parameter search methods perform equally well, and the relative performance of the different methods depends on the model being optimized. In general, the simulated annealing method




Figure 3.8: Results of genetic-algorithm parameter searches on the pyramidal cell model with and without the addition of a hypothetical fast AHP channel. The addition of the fast AHP channel allowed the parameter search to find significantly better matches.

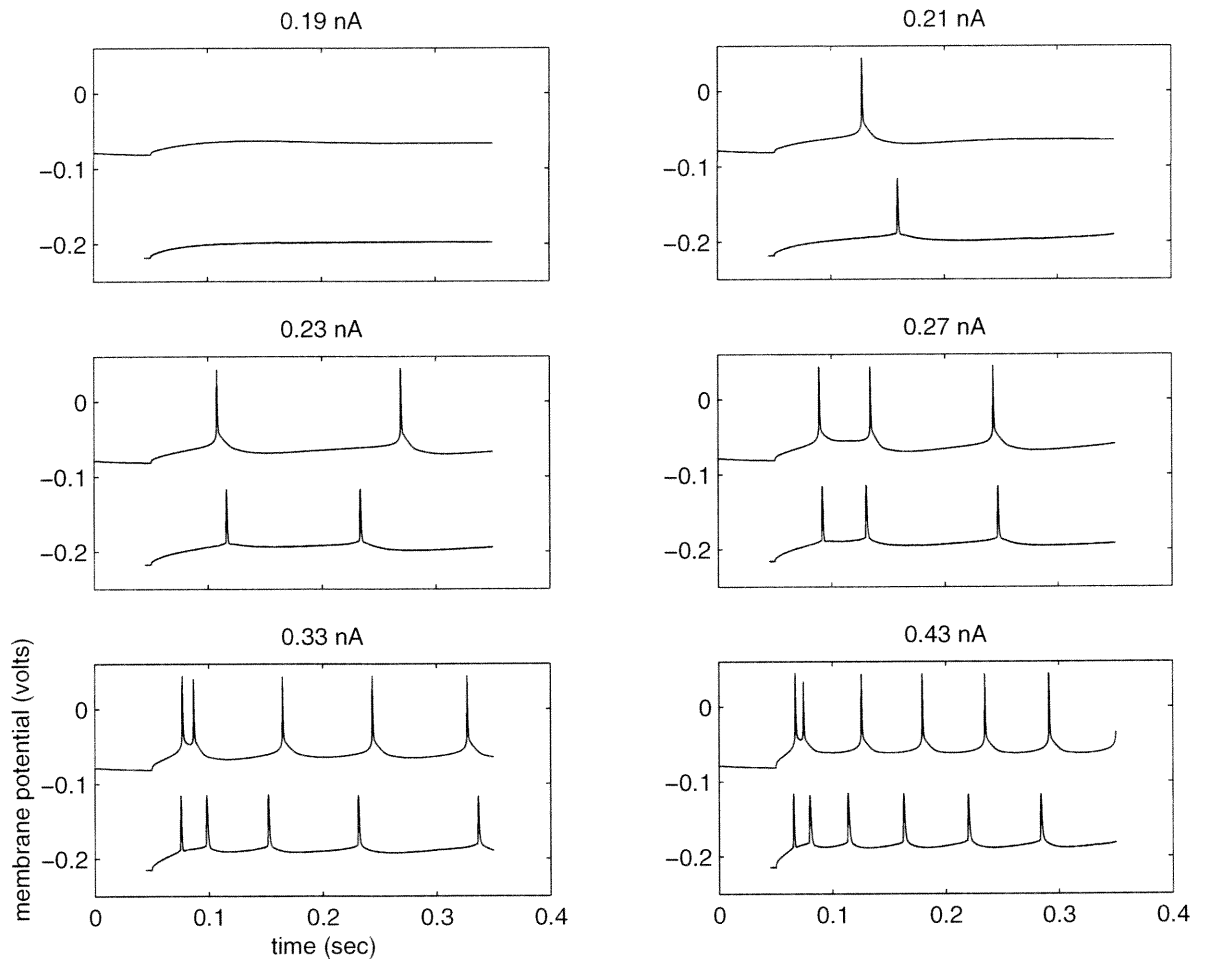


Figure 3.9: Outputs of a layer 2 pyramidal cell from piriform cortex at six different levels of current injection at the soma compared with the outputs of a genetic-algorithm parameter search. This model did not contain a fast AHP channel. Note that the model outputs do not match the spike times of the real data as well as the model outputs shown in figure 3.6, which are derived from a model containing a fast AHP channel.

and the genetic algorithm method were superior to the other methods we investigated (stochastic search and conjugate gradient) on all of our models (except for model 3, the linear passive dendrite model, for which the conjugate gradient method was as good as the genetic algorithm method). The reader should be aware, however, that all of our parameter searching methods have parameters of their own which we have determined by experimentation (see below). Therefore, it is possible that different choices of these meta-parameters could have resulted in somewhat different results. In addition, variations of each of these methods exist, especially for genetic algorithms [32]. More sophisticated algorithms in each class may very well improve on the performance of our searches. For this reason the reader should consider the data presented as a general guide providing an upper bound on the performance of these methods on the models examined.

3.5.2 Differences between the search algorithms

While it is useful to know which parameter searching algorithms are most appropriate for single-neuron compartmental models of different complexity, it would also be useful to know why one algorithm was more effective than another for a given model. Unfortunately, understanding these relationships rigorously is very difficult as it involves an interaction between the searching algorithm and the structure of the model's parameter space. However, based on previous analyses of this type [3, 11], it is possible to speculate on several possibilities for the results we obtained.

In general, it is likely that stochastic methods rather than those based on gradient-descent will be more successful for parameterizing single-neuron compartmental models. The reason for this is that simple gradient-descent methods such as conjugate gradients are greedy, *i.e.*, they move to the nearest local minimum in parameter space and stay there. This can be adequate if the space is relatively smooth and has few local minima, but it is ineffective for more ragged parameter spaces with many local minima. In fact, parameter spaces of neural models can be very ragged with sharp discontinuities separating otherwise smooth sections [3, 11]. Consistent with this analysis, the one model for which the conjugate-gradient method proved highly effective was the three-parameter passive model (model 3) which, due to its linearity and the small number of parameters, is more likely to have a smoothly varying parameter space.

Within the class of stochastic models, our results clearly demonstrate differences between genetic algorithms (GA), stochastic search (SS) and simulated annealing (SA). The SS method was markedly inferior overall to both the SA and the GA method, which is almost certainly due to the extremely unsophisticated search strategy this method employs. In general, we obtained better results with the SA method than the GA method. The SA method we used also possesses most of the advantages of gradient-descent methods without their disadvantages. At high temperatures the method can explore large expanses of parameter space, and as the temperature is lowered, the search narrows in on specific regions with better match values. The temperature-dependent noise in the method prevents the search

from getting stuck in local minima. When the temperature approaches zero the SA method converges to the closest local minimum as does a gradient-descent search. In contrast, the GA method never converges to a local minimum of the match function. However, our SA method can be expected to give poorer results with larger models, as this method can only keep a record of a small number of points at a time (the points on the simplex, which number one more than the number of parameters). For a search with a large enough number of parameters the GA method, which can simultaneously search many points in parameter space, may well overtake the SA method in performance. For our models the GA method was as effective as the SA method for the largest model (model 5). It is likely that for an even larger model GAs would have been superior. Work on GAs [15, 18, 26, 32] has suggested that this method is often optimal for problems when a large number of largely independent factors must be combined to create the best results. In contrast, in compartmental neural models optimal parameter values are typically highly dependent on each other, especially within a compartment, so that changing the value of one parameter will require changes in the values of all other parameters as well to optimize the match. This would suggest that GAs are poorly suited for compartmental models. However, while GAs are less effective for problems containing large numbers of strongly interacting parameters, they are capable of independently exploring large segments of parameter space simultaneously and in parallel, which may account for their greater relative usefulness with larger models. The intrinsic parallelism of GAs also makes them very well suited for implementation on parallel

computers [10, 14].

3.5.3 Matching functions

The choice of matching function used can have a large influence on the search process. Parameter searches can only match according to the specific criteria they are given. In our case, for our spiking models we originally tried using simple spike-matching functions which computed the RMS error between two lists of spike times. However, we found that this approach did not sufficiently penalize models whose behaviors were clearly qualitatively incorrect. For example, some models with particular parameter settings generated small bursts of spikes (*e.g.*, pairs of spikes separated by a small ISI) separated by a much larger interburst interval, in contrast to the more uniform spiking pattern we were trying to replicate. While this behavior would eventually have been selected against even using an RMS error function, we found it more efficient to put a factor into our matching function to directly penalize bursting. This allowed us to narrow down the effective search space enormously and thus sped up the searches. To be most effective, considerable thought must be applied to the selection of matching functions to achieve optimal performance. Often, initial experience with a matching function reveals that some important aspect of the data was overlooked. Therefore it is often useful to examine how well the matching function distinguishes good from bad models before committing large amounts of computer time to an extensive parameter search.

The choice of matching function can also contribute significantly to the computational cost of the parameter search, because a match must be computed for each simulation run. For this reason we chose to use spike times exclusively as the criteria for matching models 1, 2, and 5. If computational costs were not an issue, it would have been preferable to use a combination of spike timings and interspike interval (ISI) waveforms to match these models. However, since most models generate relatively few spikes, a spike timing based match function is much more efficient to calculate than a match function which also includes ISI waveforms, because for the latter the membrane potential value at each time point must be used in the computation. For some models, however, the duration of the individual simulations will still be much larger than the time needed to compute a match even if ISI waveforms are used in the match function. In such cases, one should use ISI waveforms as well as spike timings in the match function. For the passive models we had no choice but to use membrane potential waveforms, since the models do not generate spikes.

3.5.4 Variation in the parameter sets

Depending on the goals of the modeler, he or she may or may not be interested in the question of whether different runs of a parameter search result in the same set of parameter values. On one extreme, the model may be used as a phenomenological model which is a considerable abstraction of the real system, and the modeler may only wish to have a set of parameters that produces behavior similar to that seen in the real system. In this case the

uniqueness of the parameter set is of little interest. At the other extreme, a modeler may aspire to build a very detailed model whose parameters (some measured, some estimated) accurately represent the corresponding real parameters of the system. In this case the modeler is very interested in whether different parameter sets can give rise to similarly good match values, as this may indicate that the model is insufficiently constrained to accurately characterize the real system. In addition, the degree of variation of the best parameter sets gives clues to the nature of the parameter spaces being searched through.

We examined the question of non-uniqueness of the parameter sets found by our search methods by calculating statistics on a parameter-by-parameter basis for the best matches found by all the methods (data not shown). The best parameter sets found by the different methods in models 1, 2 and 3 matched the correct parameter sets very closely, with little variation in the parameter values. This suggests that the parameter spaces for these models are dominated by a single deep minimum which was reliably found by many of the searches (in particular the SA and GA searches). Models 4 and 5, however, showed a considerable degree of variation in the parameter values obtained between different successful matches, even for different parameter sets with comparable match values discovered by the same search algorithm. This suggests that parameter spaces of models with larger numbers of parameters show more degeneracy in their match values, *i.e.*, there is more than one equally good location in parameter space which can be discovered by the search algorithms. The best way to deal with this phenomenon, short of thoroughly searching parameter spaces

using brute-force methods, is to run several parameter searches with the same method on the same model and compute the variance of the parameters obtained. Typically, some parameters are highly constrained while others show a range of possible values; this in turn suggests which parameters are more or less important in determining the observed behavior.

3.5.5 Robustness of the parameter sets

Another question of interest to modelers using parameter searching techniques is the degree to which parameter sets are robust to small perturbations in the parameter values. For instance, if parameter values need to be specified with great precision in order for a good match to be achieved, the parameter set is not robust and therefore is unlikely to be physiologically relevant. In practice, estimating robustness of parameter sets is difficult because one would ideally like to vary all possible combinations of parameters. However, some insight into robustness can be achieved by varying each parameter individually (or pairs of parameters [3, 11]) around its initial value while keeping the other parameters fixed and measuring the change in the match value. Robustness estimated in this way can suggest which parameters need to be precisely set, and which can vary over a wide range. This in turn can suggest possibilities for the cellular regulation of parameter values in the real system [27]. It is also possible to build in such “robustness checks” into the matching function directly by testing robustness for each parameter set and rejecting sets that are below some threshold of robustness; however, this will dramatically increase the computational cost of

the parameter search. A better alternative is to test the robustness of the best parameter sets from a search and to reject those which are not sufficiently robust.

Another kind of robustness is robustness to small changes of the inputs. For example, one can make multiple measurements of the outputs of neurons to quantify the variability of the data source, introduce variability into the model in any of a number of different ways, and use Bayesian methods to determine the probability of the model given the data [1]. This allows modelers to test the robustness of their models to small perturbations of the inputs in a rigorous way.

3.5.6 Recommendations for effective parameter searching

While we have compared the general usefulness of different parameter searching algorithms and discussed matching functions, the process of parameter searching and the interpretation of the results can be severely hindered or improved by other factors which we discuss here.

An important issue in setting up a parameter search concerns the process of setting the parameters of the search algorithm itself. The allowable ranges of the parameter values are important parameters for search methods; ideally they should be set to physiological ranges but in many cases not enough is known to constrain these ranges. If the true parameters lie outside the initially established range, then the searching algorithm will not be able to find them. Additionally, all search algorithms have parameters of their own, such as the annealing rate for the SA method and the population size, mutation rate and recombination

rate for the GA method. Unfortunately, currently there is little analytical work which can guide the choice of these parameters, so they tend to be selected by experience and *ad hoc* rules of thumb [15, 32].

It is also necessary that the components of the model be carefully selected. A great deal of wasted time can be avoided, for example, if the initial model (before parameter searching is done) at least qualitatively matches the gross features of the target data set. Parameter searches do not remove the need for good judgment on the part of the modeler, and cannot make a bad model good. As a corollary, however, we have often found that parameter searches are often more useful when they fail than when they are successful. A failure in a well-executed parameter search suggests that a model may be incapable of reproducing the specified behavior (at least within the allowed parameter ranges). This failure naturally leads to further experimentation with either the model or the real system. For example, one can explore whether adding a new ionic channel to the model would make the model match the experimental data much more closely (*e.g.*, figures 3.6 and 3.8). If so, this can suggest new ideas for experiments.

3.5.7 Limitations of parameter searching techniques

While parameter searching techniques properly applied will add significantly to the quality and rigor of modeling efforts, several factors can limit the applicability of parameter searching methods.

Since parameter searches require a large number of evaluations of the model (tens to hundreds of thousands for the models we investigated), the time required to simulate a single iteration of a model must be relatively small. For instance, a model which requires one hour or more to go through one iteration (*e.g.*, [8, 9]) will take over a year of workstation time to run through 10,000 simulations. Obviously, unless the simulation time can be reduced significantly by running on faster computers (*e.g.*, parallel supercomputers [14]), such a model is currently not a good candidate for automated parameter searches. Under these conditions parameter searches can sometimes still be run on simplified versions of a large model with the hope that the results can then be applied to the larger model.

At the present time it does not appear that the parameter search methods we describe will be useful on models with huge numbers (*e.g.*, hundreds) of parameters. We note that moving from a model with eight parameters to one with 23 parameters required a scale-up of 10 in the number of iterations per parameter search needed to get good results. It is likely that models with very large numbers of parameters (> 100) can easily outrun the capabilities of the current generation of workstations. This is a key argument for developing parameter search methods to work on parallel supercomputers, as we have begun doing [14]. In addition, it is desirable for modelers to investigate ways of reducing the size of the parameter spaces of their models, for instance by forcing some parameters in different compartments to have the same values (*e.g.*, [8]).

Another potential difficulty with parameter searches concerns overfitting of the data.

When a model is carefully tuned to match a specific data set, it may not generalize well to similar data sets not used in the tuning. For instance, if a model is tuned to match spike times very precisely at a given value of current injection it may fail to match spike times at other current levels. For this reason, our own studies employ a wide range of inputs to provide an appropriate breadth of constraints on model behaviors. Such models are far more likely to be underfit than overfit. If overfitting is suspected, it is a simple matter to parameterize one's model on a subset of the data set (the training set) and once the model is parameterized check how well it matches the rest of the data (the test set).

It is also critical not to place too much weight on the results of a parameter search in a weakly constrained or highly simplified model. In our pyramidal cell model, many of the kinetic parameters were only approximately known, those relating to calcium buffering were only guesses, and we added a hypothetical fast AHP current. However, it is entirely possible that a combination of different calcium buffering schemes, dendritic active currents, or other channels we did not include could have given equally good or better results. For this reason, our modeling results allow us to speculate about the possible significance of the hypothetical fast AHP channel, but it would be inappropriate to make any strong statements in the absence of more definitive experimental data. Of course, one of the more useful consequences of realistic modeling is to motivate and provide context for further experimental studies.

Finally, as models proliferate it will become more important to have rigorous criteria

for comparing the quality of different models which are intended to match the same data set. We have recently begun exploring this issue, using Bayesian methods to compare neural models quantitatively [1]. As neural models become more complex, the subject of parameter searching in these models will continue to develop and will be critical to progress in computational neuroscience.

3.6 Acknowledgements

We would like to thank Alex Protopapas for helpful comments, for the experimental data on which model 5 was based, as well as for a detailed model of a layer 2 pyramidal cell in piriform cortex which formed the basis for our simplified model.

Bibliography

- [1] P. F. Baldi, M. C. Vanier, and J. M. Bower. On the Use of Bayesian Methods for Evaluating Compartmental Neural Models. *Journal of Computational Neuroscience*, 5:285–314, 1998.
- [2] M. I. Banks, L. B. Haberly, and M. B. Jackson. Layer-specific Properties of the Transient K-current (I-A) in Piriform Cortex. *Journal of Neuroscience*, 16(12):3862–3876, 1996.
- [3] U. S. Bhalla and J. M. Bower. Exploring Parameter Space in Detailed Single Neuron Models – Simulations of the Mitral and Granule Cells of the Olfactory-Bulb. *Journal of Neurophysiology*, 69(6):1948–1965, 1993.
- [4] J. M. Bower and D. Beeman. *The Book of Genesis*. Springer Verlag, New York, 1995.
- [5] P. C. Bush and T. J. Sejnowski. Reduced Compartmental Models of Neocortical Pyramidal Cells. *Journal of Neuroscience Methods*, 46:159–166, 1993.

- [6] G. Buzsaki and A. Kandel. Somadendritic Backpropagation of Action Potentials in Cortical Pyramidal Cells of the Awake Rat. *Journal of Neurophysiology*, 79(3):1587–1591, 1998.
- [7] A. Constanti, M. Galvan, P. Franz, and J. A. Sim. Calcium-dependent Inward Currents in Voltage-clamped Guinea-pig Olfactory Cortex Neurons. *Pflügers Archiv European Journal of Physiology*, 404(3):259–265, 1985.
- [8] E. DeSchutter and J. M. Bower. An Active Membrane Model of the Cerebellar Purkinje Cell I. Simulation of Current Clamps in Slice. *Journal of Neurophysiology*, 71:375–400, 1994.
- [9] E. DeSchutter and J. M. Bower. An Active Membrane Model of the Cerebellar Purkinje Cell II. Simulation of Synaptic Responses. *Journal of Neurophysiology*, 71:401–419, 1994.
- [10] R. Eichler-West and G. Wilcox. Robust Parameter Selection for Compartmental Models of Neurons Using Evolutionary Algorithms. In J. M. Bower, editor, *Computational Neuroscience: Trends in Research 1997*. Plenum Publishing, New York, 1997. In press.
- [11] W. R. Foster, L. H. Ungar, and J. S. Schwaber. Significance of Conductances in Hodgkin-Huxley Models. *Journal of Neurophysiology*, 70(6):2502–2518, 1993.

- [12] C. R. French, P. Sah, K. J. Buckett, and P. W. Gage. A Voltage-Dependent Persistent Sodium Current in Mammalian Hippocampal-Neurons. *Journal of General Physiology*, 95(6):1139–1157, 1990.
- [13] M. Galvan, A. Constanti, and P. Franz. Calcium-Dependent Action-Potentials in Guinea-pig Olfactory Cortex Neurons. *Pflugers Archiv European Journal of Physiology*, 404(3):252–258, 1985.
- [14] N. H. Goddard and G. Hood. Parallel Genesis for large-scale modeling. In J. M. Bower, editor, *Computational Neuroscience: Trends in Research 1997*, pages 911–917. Plenum Publishing, New York, 1997.
- [15] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, Mass., 1989.
- [16] J. V. Halliwell and C. N. Scholfield. Somatically Recorded Ca-Currents in Guinea-pig Hippocampal and Olfactory Cortex Neurons are Resistant to Adenosine Action. *Neuroscience Letters*, 50(1-3):13–18, 1984.
- [17] B. Hille. *Ionic Channels of Excitable Membranes, 2nd. Ed.* Sinauer Associates Inc., Sunderland, Massachusetts, 1992.
- [18] J. H. Holland. *Adaptation in Natural and Artificial Systems, 2nd. ed.* The MIT Press, Cambridge, 1992.

- [19] W. R. Holmes and W. Rall. Electrotonic Length Estimates in Neurons with Dendritic Tapering or Somatic Shunt. *Journal of Neurophysiology*, 68(4):1421–1437, 1992.
- [20] W. R. Holmes and W. Rall. Estimating the Electrotonic Structure of Neurons with Compartmental Models. *Journal of Neurophysiology*, 68(4):1438–1452, 1992.
- [21] J. J. B. Jack, D. Noble, and R. W. Tsien. *Electric Current Flow in Excitable Cells*. Clarendon Press, Oxford, 1983.
- [22] E. S. Keeping. *Introduction to Statistical Inference*. Dover Publications, Inc., New York, 1995.
- [23] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, 1983.
- [24] C. Koch and I. Segev. *Methods in Neuronal Modeling*. MIT Press, Cambridge, Massachusetts, 1989.
- [25] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.
- [26] J. R. Koza. *Genetic Programming: on the Programming of Computers by Means of Natural Selection*. The MIT Press, Cambridge, 1992.
- [27] G. Lemasson, E. Marder, and L. F. Abbott. Activity-Dependent Regulation of Conductances in Model Neurons. *Science*, 259(5103):1915–1917, 1993.

- [28] J. C. Magee and D. Johnston. Synaptic Activation of Voltage-Gated Channels in the Dendrites of Hippocampal Pyramidal Neurons. *Science*, 268(5208):301–304, 1995.
- [29] Z. F. Mainen, J. Joerges, J. R. Huguenard, and T. J. Sejnowski. A Model of Spike Initiation in Neocortical Pyramidal Neurons. *Neuron*, 15(6):1427–1439, 1995.
- [30] D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative Electrophysiology of Pyramidal and Sparsely Spiny Stellate Neurons of the Neocortex. *Journal of Neurophysiology*, 54(4):782–806, 1985.
- [31] D. A. McCormick and J. R. Huguenard. A Model of the Electrophysiological Properties of Thalamocortical Relay Neurons. *Journal of Neurophysiology*, 68(4):1384–1400, 1992.
- [32] M. Mitchell. *An Introduction to Genetic Algorithms*. The MIT Press, Cambridge, 1996.
- [33] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C, 2nd. Ed.* Cambridge University Press, New York, 1992.
- [34] W. Rall. *The Theoretical Foundation of Dendritic Function*. MIT Press, Cambridge, Mass., 1995.
- [35] W. Rall, R. E. Burke, W. R. Holmes, J. J. B. Jack, S. J. Redman, and I. Segev. Matching Dendritic Neuron Models to Experimental Data. *Physiological Reviews*, 72(4):S159–S186, 1992.

- [36] J. Rinzel. Excitation Dynamics – Insights from Simplified Membrane Models. *Federation Proceedings*, 44(15):2944–2946, 1985.
- [37] K. J. Staley, T. S. Otis, and I. Mody. Membrane-Properties of Dentate Gyrus Granule Cells – A Comparison of Sharp Microelectrode and Whole-Cell Recordings. *Journal of Neurophysiology*, 67(5):1346–1358, 1992.
- [38] B. Tawfik and D. M. Durand. Nonlinear Parameter-Estimation by Linear Association – Application to a 5-Parameter Passive Neuron Model. *IEEE Transactions on Biomedical Engineering*, 41(5):461–469, 1994.
- [39] R. D. Traub, R. K. S. Wong, R. Miles, and H. Michelson. A Model of a CA3 Hippocampal Pyramidal Neuron Incorporating Voltage-Clamp Data on Intrinsic Conductances. *Journal of Neurophysiology*, 66(2):635–650, 1991.
- [40] M. C. Vanier and J. M. Bower. A Comparison of Automated Parameter-Searching Methods for Neural Models. In J. M. Bower, editor, *Proceedings of the 1995 Computational Neuroscience Conference (CNS*95), Monterey (CA)*. Academic Press, 1996.
- [41] J. A. White, P. B. Manis, and E. D. Young. The Parameter-Identification Problem for the Somatic Shunt Model. *Biological Cybernetics*, 66(4):307–318, 1992.
- [42] L. Zhang and K. Krnjević. Whole-Cell Recording of Anoxic Effects on Hippocampal Neurons in Slices. *Journal of Neurophysiology*, 69:118–127, 1993.

3.7 Appendix: Model descriptions

GENESIS scripts containing detailed descriptions of ionic channels and cell morphologies for all models can be downloaded from

<http://www.bbb.caltech.edu/USERS/mvanier/parampaper/param.html>. Ionic channels were modeled as generalizations of the Hodgkin-Huxley model [17], in some cases adapted to use calcium concentration rather than membrane potential as the independent variable for the activation and inactivation curves. The equation we used for calcium buffering in the pyramidal cell model is the following:

$$d[Ca]/dt = B \cdot I_{Ca} - [Ca]/\tau \quad ,$$

where $[Ca]$ is the calcium ion concentration, τ is the buffer time constant, and B is a scaling factor which converts calcium current influx into changes in calcium concentration. B and τ were both parameters of the model; their initial values and ranges are in table 3.3.

Chapter 4

On the Use of Bayesian Methods for Evaluating Compartmental Neural Models

“A model should be as simple as possible, but no simpler.”

(attributed to Einstein)

4.1 Abstract

Computational modeling is being increasingly used in neuroscience. In deriving such models, inference issues such as model selection, model complexity, and model comparison must

constantly be addressed. In this paper we briefly present the Bayesian approach to inference. Under a simple set of common sense axioms, there exists essentially a unique way of reasoning under uncertainty by assigning a degree of confidence to any hypothesis or model, given the available data and prior information. Such degrees of confidence must obey all the rules governing probabilities, and can be updated accordingly as more data becomes available. While the Bayesian methodology can be applied to any type of model, as an example we outline its use for an important, and increasingly standard, class of models in computational neuroscience: compartmental models of single neurons. Inference issues are particularly relevant for these models: their parameter spaces are typically very large, neurophysiological and neuroanatomical data is still sparse, and probabilistic aspects are often ignored. As a tutorial, we demonstrate the Bayesian approach on a class of one-compartment models with varying numbers of conductances. We then apply Bayesian methods on a compartmental model of a real neuron to determine the optimal amount of noise to add to the model to give it a level of spike-time variability comparable to that found in the real cell.

4.2 Introduction

Computational modeling is being increasingly used in neuroscience, over vastly different spatio-temporal scales, from molecules to entire brain areas, and at different levels of complexity [9]. While available neurophysiological and neuroanatomical data is also rapidly increasing, what is known today constitutes only a tiny fraction of what remains to be

known. Thus in computational neuroscience in particular, and more generally in biology and all other information-rich sciences, scientists must reason in the presence of a high degree of uncertainty: many facts are missing, and some of the facts are wrong. Computational neuroscientists are thus constantly faced with inference problems related to building models from available data. For instance, it is not necessarily obvious which class of models should be used and what level of complexity is appropriate for a particular investigation. It is also not obvious which details are important and which should be ignored or discarded. Finally, it is often difficult to compare different models and to determine which model is the best, given the data available. In fact, it is relatively rare in the modeling literature to see any rigorous comparison between new and previously published models. Each of these issues is critical to determining if any particular model is a good model for a particular purpose, or constitutes an advance in our understanding.

Over the last several years there has been a dramatic growth in the number of computational neurobiologists constructing “realistic” neural models. We have previously defined these models as those that are based on the actual anatomy and physiology of the nervous system, and which generate outputs which can be directly compared to physiologically measurable outputs [6, 7]. At present the most common type of realistic neural models are compartmental models in which the structure being modeled is broken into a smaller number of interconnected compartments [26]. In the case of single neuron models, this means that the dendrites are broken up into a series of individual compartments each of which is

assumed to be isopotential [26]. The precise number of these compartments depends on the degree of anatomical detail available or desired.

As computer resources and the sophistication of both the data available and the questions being asked has increased, so has the complexity of single neuron models. In our own laboratory, for example, we have constructed and are currently studying a cerebellar Purkinje cell model containing a total of 4550 compartments and 8021 ionic channels [14, 15, 21]. This model and other complex models with large numbers of parameters are often criticized on the grounds that, with so many parameters, the model should be able to reproduce any kind of neuronal behavior. In fact, there exist many implicit constraints that emerge from the structure of the models and which render arbitrary behavior very difficult, if not impossible, to reproduce. However, it remains an open question whether a simpler model with fewer parameters could match a given data set as well as a more complex model. This issue becomes especially important when there is little experimental data to constrain the model. In these cases inference issues become fundamental.

In addition, as pointed out in [33], choosing simpler models simply because little data is available does not make much sense. This is particularly clear in the case of the nervous system, where it is obvious that the amount of data collected so far and the complexity of the nervous system itself are two completely different things. Therefore, the appropriate complexity of the model, how well it replicates the data, and whether it can be simplified without losing any essential behaviors are important questions each involving inference. In

this paper we present the Bayesian approach to inference. Our purpose is to introduce the Bayesian point of view and the advantages that may derive from it to computational neuroscientists, especially those involved in building compartmental models.

When reasoning in the presence of certainty, one uses deduction. This is how the most advanced theories in information-poor sciences, such as physics or mathematics, are presented in an axiomatic fashion. Deduction is not controversial. The overwhelming majority of people agree on how to perform deductions in a unique way: if A implies B , and A is true, then B must be true. This is the essence of Boolean algebra, and is the basis for all digital computers. When reasoning in the presence of uncertainty, one uses induction and inference: if A implies B , and B is true, then A is *more plausible*. An amazing and still poorly known fact is that there exists a simple and unique consistent set of rules for induction, model selection and comparison: it is called Bayesian induction. The Bayesian approach has been known for quite some time, but only recently has it started to infiltrate different areas of science and technology systematically with useful results [2, 16, 17, 29, 28]. To our knowledge, however, it has not previously been applied systematically to neurobiological modeling. Our purpose here is to introduce the Bayesian point of view and the advantages that may derive from it to those involved in building compartmental models. In a companion paper [1], we use a Bayesian framework to derive new algorithms for the analysis of neuronal multi-unit recordings.

The Bayesian point of view has a simple intuitive description: a degree of plausibility

is assigned to any proposition, hypothesis or model (throughout the paper, hypotheses or models are essentially synonymous: models tend to be complex hypotheses with many parameters). More precisely, in order to properly carry out the induction process, one should proceed in three steps:

1. Clearly state what the hypotheses or models are, along with all the background information and the data.
2. Use the language of probability theory to assign prior probabilities to the hypotheses.
3. Use probability calculus for the inference process, in particular to evaluate posterior probabilities (or degrees of belief) for the hypotheses in light of the available data, and derive unique answers.

Such an approach certainly seems a reasonable one ¹. But why should it be so compelling? Why use the language of probability theory, as opposed to any other method? The surprising answer is that it can be mathematically proven that this is the only consistent way of reasoning in the presence of uncertainty. Specifically, there exists a small set of very simple common-sense axioms, the Cox-Jaynes axioms, under which it can be shown that the Bayesian approach is the unique consistent approach to inference and induction. Under

¹Notice that the Bayesian approach is not directly concerned with the creative process, *i.e.*, how to generate new hypotheses or models. It is only concerned with assessing the value of models with respect to the available knowledge and data. This assessment procedure, however, may turn out to be helpful in generating new ideas.

the Cox-Jaynes axioms, degrees of plausibility satisfy exactly all the rules of probabilities. Probability calculus is then all the machinery that is required for inference, model selection and model comparison.

The axiomatic presentation of the Bayesian approach is given in appendix A for completeness. For brevity, we do not present any proofs or any historical background for the Bayesian approach, nor do we discuss any controversial issues regarding the foundations of probability theory, such as the frequentist versus Bayesian point of view ². All of these can be found in various books and articles (such as [3, 4, 12, 23, 34], and references therein).

While the Bayesian approach is universal and applies to any problem or model, our focus here is on compartmental neural models. In section 2, we briefly review Bayesian inference. In section 3, we analyze the steps compartmental modelers need to take for a proper Bayesian treatment of their models. In particular, we provide specific suggestions for the computation of priors and likelihoods for compartmental models. In section 4, we apply the methodology in detail to specific compartment models of single neurons. We

²In the frequentist approach, probabilities can be introduced only in the context of repeatable events or experiments. The probability of an event is defined as the limit of the corresponding empirical frequency (the ratio of the number of successful results N_S to the total number of experiments N) as the number of experiments goes to infinity:

$$p \equiv \lim_{N \rightarrow \infty} N_S/N$$

. Thus the notion of a probability as a number quantifying the degree of belief in a proposition has no meaning in the frequentist view.

summarize the benefits and drawbacks of the method in the discussion session, and discuss some possible extensions.

4.3 Bayesian inference

As described in appendix A, the immediate consequence of the Cox-Jaynes axioms is Bayes' theorem:

$$p(X|Y, I) = \frac{p(Y|X, I)p(X|I)}{p(Y|I)} = p(X|I) \frac{p(Y|X, I)}{p(Y|I)} \quad , \quad (4.1)$$

where X and Y are propositions and I represents background information. $p(X|Y, I)$ denotes the conditional probability of X knowing Y and I . We will apply this equation to a parametrized model $M = M(\theta)$ and a data set D . For simplicity, we will drop the background information I from the following equations, or equivalently incorporate it into the data. From Bayes' theorem we have immediately,

$$p(M|D) = \frac{p(D|M)p(M)}{p(D)} = p(M) \frac{p(D|M)}{p(D)} \quad (4.2)$$

The term $P(M)$ is referred to as the prior probability or simply the *prior*. The term $P(M|D)$ is referred to as the posterior probability or *posterior*. The term $P(D|M)$ is usually referred to as the *likelihood* (but occasionally as the *evidence*). The prior represents our estimate of the probability that model M is correct before one has obtained any data, while the

posterior represents our updated belief in the probability that model M is correct given that we have observed the data set D . For technical reasons, it is often easier to work with the corresponding logarithms, so that:

$$\log p(M|D) = \log p(D|M) + \log p(M) - \log p(D) \quad (4.3)$$

For data obtained sequentially, one has

$$p(M|D_1, \dots, D_n) = p(M|D_1, \dots, D_{n-1}) \frac{p(D_n|M, D_1, \dots, D_{n-1})}{p(D_n|D_1, \dots, D_{n-1})} \quad (4.4)$$

In other words the old posterior $p(M|D_1, \dots, D_{n-1})$ plays the role of the new prior. Bayesian modeling is entirely open to both new data and new models. Bayesian modeling is iterative by nature, as both data and models evolve.

To apply equation (4.3) to compartmental models, or to any other class of models, one must specify the prior $p(M)$ and the likelihood $p(D|M)$. Once the prior and likelihood terms are made explicit, the initial modeling effort is completed. All that remains is to apply probability theory to obtain the results. Two specific compartmental models M_1 and M_2 can be compared by comparing their probabilities $p(M_1|D)$ and $p(M_2|D)$. More generally, one objective is often to find or approximate the “best” model within a class, *i.e.*, to find the set of parameters θ maximizing the posterior $p(M|D)$ (or $\log p(M|D)$), and the corresponding error bars (see appendix B). This is called MAP (maximum a-posteriori) estimation. In order to deal with positive quantities, this is also equivalent to minimizing

$-\log p(M|D)$. Notice that the term $p(D)$ in (4.3) plays the role of a normalizing constant that does not depend on the parameters θ , and is therefore irrelevant for this optimization. If the prior $p(M)$ is uniform, then the problem reduces to finding the maximum of $p(D|M)$ (or $\log p(D|M)$). This is just maximum likelihood (ML) estimation. In most interesting models, the function being optimized is complex and its extrema cannot be solved for analytically. Thus one must resort to iterative and possibly stochastic methods such as gradient descent or simulated annealing. A comparative study of such algorithms in the case of compartmental models can be found in [42, 43].

Whereas finding the optimal model, the first level of Bayesian inference, is common practice, it is essential to notice that this is really useful only if the distribution $p(M|D)$ is sharply peaked around a unique optimum. In situations characterized by a high degree of uncertainty and relatively small amounts of available data, this is often not the case. Thus in Bayesian inference we are interested in the function $p(M|D)$ over the entire space of models, and more precisely in evaluating expectations with respect to $p(M|D)$. This is the case in higher levels of Bayesian inference, such as prediction problems, marginalization of nuisance parameters, and model class comparisons.

Consider, for instance, a prediction problem, where we are trying to predict the output value y of an unknown parametrized function f_θ , given an input x . It is easy to show that the optimal prediction is given by the expectation

$$E(y) = \int_{\theta} f_{\theta}(x)p(\theta|D)d\theta \quad (4.5)$$

This integral is just the average of the predictions made by each possible model f_{θ} , weighted by the plausibility $p(\theta|D)$ of each model. Another example is the process of marginalization, where integration of the posterior parameter distribution is carried only with respect to a subset of the parameters, the so-called nuisance parameters [16]. In a frequentist framework, the notion of distribution over the parameters is not defined and therefore nuisance parameters cannot be integrated out easily. In compartmental models, one is also often lead to the problem of comparing two model classes C_1 and C_2 . As an example, these could be compartmental models with two very different numbers of compartments, possibly representing a highly anatomically detailed model with a large number of compartments versus a simplified model with a small number of compartments. To compare C_1 and C_2 in the Bayesian framework, one must compute $p(C_1|D)$ and $p(C_2|D)$. Using Bayes' theorem again: $p(C|D) = p(D|C)p(C)/p(D)$. In addition to the prior $p(C)$, one must calculate the likelihood $p(D|C)$, and this is obtained by averaging over the entire model class:

$$p(D|C) = \int_{\theta \in C} p(D, \theta|C)d\theta = \int_{\theta \in C} p(D|\theta, C)p(\theta|C)d\theta \quad (4.6)$$

Naturally, in cases where the likelihood $p(D|\theta, C)$ is very peaked around its maximum, such expectations can be approximated using the mode (*i.e.*, the value with the highest probability). But in general, such integrals require better approximations, for instance using

Monte Carlo sampling methods [5, 32, 39]. Available computer power and its progress are of course important issues to be considered in this context.

As mentioned previously, it does not make sense to choose a simple model simply because available data is scarce. *Everything else being equal*, however, it is true that one should prefer a simple hypothesis to a complex one. This is Ockham's razor. As pointed out by several authors, Ockham's razor is automatically embodied in the Bayesian framework [24, 28] in at least two different ways. In the first (trivial) way, one can introduce priors that penalize complex models. But even without such priors, parametrized complex models will tend to be consistent with a larger volume of data space. Since a likelihood $p(D|M)$ must sum to 1 over the space of data, if $p(D|M)$ covers a larger expanse of data space the likelihood values for given data sets will be smaller on average. Therefore, everything else being equal, complex models will tend to assign a correspondingly smaller likelihood to the observed data.

In summary, Bayesian methods provide a rigorous and provably unique framework for inference. The basic step is to compute model plausibilities with respect to the available data, and the associated expectations, using the rules of probability theory and possibly numerical approximations. We can now turn to the specific application of the general Bayesian framework to compartmental models.

4.4 Bayesian compartmental modeling in neuroscience

We now describe how to apply a Bayesian approach to compartmental modeling in neuroscience. For any type of model class, it is clear from equation (4.2) that in order to evaluate $p(M|D)$ the first step consists in making the prior $p(M)$ and the likelihood $p(D|M)$ explicit.

4.4.1 Prior probabilities for compartmental models

A typical compartmental model $M(\theta)$ has a large number of parameters θ . Different types of parameters can be used, but for single-cell models the majority of these parameters correspond to the conductances, and possibly the kinetic time constants, of ionic channels in the different compartments. Both conductance values and time constants are *scale* parameters. Thus, unless any additional information is available, a reasonable prior is a uniform logarithmic prior [29]. That is, if a conductance, or a time constant, c occupies some value in an interval $[a, b]$ and we do not have any additional specific information, it is reasonable to define the prior density

$$p(k < \log c < k + \Delta k) = \frac{\Delta k}{\log b - \log a} \quad (4.7)$$

for any $[k, k + \Delta k]$ in $[\log a, \log b]$, and 0 otherwise. So, if c is equally likely to be 1, 10 or 100, $\log_{10} c$ will have a flat distribution between 0 and 2. Starting from equation (4.3), we can calculate the prior term $\log p(M)$ by assuming that all the basic parameters are independent. That is, if a model is characterized by a set of conductances c_i satisfying

equation (4.7) over an interval $[a_i, b_i]$, then

$$\log(p(M)\Delta\vec{k}) = \sum_i \log \frac{\Delta k_i}{\log b_i - \log a_i} \quad , \quad (4.8)$$

where $\Delta\vec{k} = \Delta k_1 \Delta k_2 \dots$. In terms of probability densities, this becomes

$$\log(p(M)) = \sum_i \log \frac{1}{\log b_i - \log a_i} \quad (4.9)$$

Another useful prior for conductances, or their logarithms, is a gamma prior. The gamma density [20] with parameters α and λ , is given by

$$\Gamma(x; \alpha, \lambda) = \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\lambda x} \quad (4.10)$$

for $x > 0$, and 0 otherwise. By varying α and λ and translating x , the gamma density allows for a wide range of priors, with more mass concentrated in one specific region of parameter space.

It is also possible in Bayesian modeling to tie together different sets of parameters. For instance, one could assume that conductances are linked, or have the same values or ranges, in different regions of a large dendritic tree. An example of this in a compartmental model with four types of regions can be found in [14]. This approach is a form of hierarchical modeling, where hyperparameters are used to control the priors on the model parameters. A typical example is when there are many parameters with the same prior Gaussian distribution. In this case, the mean and standard deviation of the common Gaussian can be used

as hyperparameters. Hierarchical levels can also be iterated. Hierarchical modeling allows for the introduction of a certain amount of structure on the model parameters, and also for parameter reduction since the model prior can be calculated from a (usually smaller) number of hyperparameters. In symbolic form

$$p(\theta) = \int_{\alpha} p(\theta|\alpha)p(\alpha)d\alpha \quad (4.11)$$

where α represents hyperparameters with prior $p(\alpha)$.

The use of priors is a strength of the Bayesian approach, since it allows incorporation of prior knowledge into the modeling process. It has also been considered a weakness in the past, on the ground that priors are subjective, and that different results can be derived with different priors. There are at least four different responses to these arguments.

1. Even when priors are not mentioned explicitly, they are used implicitly. The Bayesian approach forces a clarification of one's assumptions instead of ignoring the problem of priors.
2. The effects of priors diminishes as the amount of data increases. The magnitude of the log-likelihood scales linearly with the number of data points, when these are approximately independent. Thus the more data one has, the more the likelihood will dominate the computed posterior.
3. There exist situations where objective criteria, such as maximum entropy and group

invariance, can be used to determine non-informative priors (for instance, [17]).

4. Finally, and most importantly, the effects of different priors, as well as different models and model classes, can be assessed within the Bayesian framework by comparing the corresponding probabilities.

4.4.2 Likelihoods for compartmental models

In order to define $p(D|M)$, one must consider how a model M could also give rise to a different observation set D' : in a Bayesian framework, compartmental models *must* be probabilistic. In contrast, with the exception of models that use random stimuli, such as input synapses that are activated randomly (*e.g.*, [14, 40]), most current compartmental models are deterministic. A deterministic model assigns a probability zero to all the data, except the one it can produce exactly. This is clearly inadequate for biological models. In fact, recently many experimentalists and theorists have been investigating the nature, sources, and functions of noise in real neurons (see, for instance, [30, 37, 40]). However, because of the complexity inherent to compartmental models, when constructing a model most of the effort is generally focussed on the structural and deterministic components of the model (*e.g.*, the number, morphology and conductances of the compartments) whereas the equally important probabilistic aspects are often ignored. However, a scientific discourse on compartmental models, how well they fit the data, and how they can be compared to each other, is impossible unless models are probabilistic.

It is also important to notice that results from very different sources and experiments (*e.g.*, different cells, different organisms, different experimental preparations) are routinely combined in compartmental models. For instance, channel kinetic data obtained from patch-clamp analyses performed at room temperature may be combined with intracellular current- or voltage-clamp data performed in a brain-slice preparation at a different temperature in a different lab, and with anatomical data from yet another lab. Current compartmental models in general attempt to capture the behavior of an “average” cell, or piece of circuitry. Thus mismatches between the behavior of specific cells and “average” models are bound to occur, and the probabilistic connection between compartmental models and the underlying neurobiological system being modeled is even stronger.

Since compartmental models of different types are used to study an array of very different phenomena, it is impossible to give a completely general prescription for computing likelihoods. At one extreme, we may be interested in the particular shape of an action potential or an interspike interval waveform (see [27] for a related Bayesian treatment of spike classification). Here, however, we shall concentrate on the case where we are attempting to reproduce the timing characteristics of an entire spike train while injecting constant currents of different amplitudes into the soma of the model cell. Thus in our case the data consist of a series of numbers representing spike times.

Several ad hoc criteria have been used to compare spike trains, but whatever criterion is used, it comes with an underlying probabilistic model that can be clarified and made

amenable to Bayesian analysis. Indeed, if the fit of a model $M = M(\theta)$ with parameters θ is measured by some error function $f(\theta, D) \geq 0$ to be minimized, one can always define the associated likelihoods to be

$$p(D|M(\theta)) = \frac{e^{-f(\theta, D)}}{Z} \quad , \quad (4.12)$$

where $Z = \int_{\theta} e^{-f(\theta, D)}$ is a normalizing factor (the “partition function” in statistical mechanics) which makes the probabilities integrate to 1. As a result, minimizing the error function is equivalent to maximum likelihood estimation. In particular, when the sum of squared differences in spike timings (or interspike interval timings), is used to compare two spike trains or two other data sets D and D' , a rather common practice, this implies an underlying Gaussian model (see below). Thus the Bayesian point of view clarifies the probabilistic assumptions that must underly any criteria for comparing two data sets D and D' .

We shall now consider three possible directions for building a truly probabilistic model of spike trains, using an underlying deterministic compartmental model, for the computation of likelihoods. The underlying deterministic model will be referred to as the “deterministic component” of the model.

Simple noise model

Real neurons have a number of sources of noise, such as stochastic transitions of ion channels between open and closed states, stochastic release of neurotransmitter vesicles from presynaptic terminals etc. [19]. A natural and simple way to model this noise is by incorporating Gaussian white noise in the amplitudes of injected currents, at least in the compartment(s) associated with the soma. Adding Gaussian noise to the current input to the model is appropriate since the sum of a large number of independent stochastic processes tends to a Gaussian distribution by the central limit theorem. This gives us a model capable of producing stochastic outputs which can be compared to the data. We refer to this as the “simple noise model”. While this is perhaps the most simple model, it does not by itself provide a computationally efficient way of estimating likelihoods. Likelihoods can in principle be estimated by running the simple noise model a very large number of times with different noise instantiations, collecting the spike timings of the model, and computing how often any given data set D was generated by the model. This, however, is usually impractical because any given data set has an extremely small probability of being matched exactly (where exactly means that each spike produced by the model must occur at exactly the same time as the corresponding spike in the data within the temporal resolution of the data acquisition setup and the simulation). Thus we next consider probabilistic models directly at the level of the spike trains themselves, and ways in which they can be superimposed on standard compartmental models. Such models are used extensively in the simulations

described below.

Variable-rate Poisson models

One of the most basic models for point processes, such as spike trains, is a Poisson process with rate λ . If we subdivide time into very small intervals Δt (the “bin size”), a Poisson process can be seen as a sequence of independent Bernoulli trials (*i.e.*, trials with two possible outcomes, one of which has a small constant probability λ , and the other probability $1 - \lambda$). Over a period of time T , the number of successes, or spikes, has a binomial distribution which can be approximated by a Poisson distribution with average λT . The main problem with such a model is that spike train rates are usually not stationary. Thus the natural generalization is a Poisson model with variable rate $\lambda(t)$. To generate such a process, we can imagine a large sequence of Bernoulli trials, where the probability of success in the trial indexed by time t is given by $\lambda(t)$. Clearly the rate function must satisfy $0 \leq \lambda(t) \leq 1$.

Under a variable-rate Poisson model M , the likelihood of any spike train, with spiking times (t_1, \dots, t_K) , is easily computed by

$$p(t_1, \dots, t_K | M) = \prod_{i=1}^K \lambda(t_i) \prod_{t \neq t_i} (1 - \lambda(t)) = \prod_{i=1}^K \frac{\lambda(t_i)}{1 - \lambda(t_i)} \prod_t (1 - \lambda(t)) \quad (4.13)$$

Notice that $Q = \prod_t (1 - \lambda(t))$ is the probability of having no spikes at all. The log-likelihood is given by

$$\log p(t_1, \dots, t_K | M) = \sum_{i=1}^K \log \lambda(t_i) - \sum_{i=1}^K \log(1 - \lambda(t_i)) + \log Q \quad (4.14)$$

As it stands, the variable-rate model is a generative model which has nothing to do with compartmental models. However, it is easy to couple such a model with an underlying deterministic compartmental model component by the way in which the rate function $\lambda(t)$ is computed. A first possibility is to run the simple noise model many times and collect spike histograms to estimate $\lambda(t)$, for each t . In general it would be unrealistic to set $\lambda(t)$ to 0 in a bin without spikes, since the lack of spikes may well be due to the small sample size. This issue can be easily addressed by regularizing the spike counts, which in the Bayesian framework is done by using a Dirichlet prior on $\lambda(t)$ ([2] and appendix C). A second possibility is to decompose $\lambda(t)$ as a sum of bell-shaped functions, for instance using Gaussians

$$\lambda(t) = \sum_{i=1}^K N_i e^{-(t-\mu_i)^2/2\sigma_i^2} \quad (4.15)$$

We denote such a model by $M = M(\vec{\mu}, \vec{\sigma}, \vec{N})$. When $\lambda(t)$ results from a single Gaussian curve around an isolated spike, the term $\log \lambda(t_i)$ gives rise to the usual quadratic error term. The model can be coupled to a deterministic compartmental model component by the way in which the parameters μ , σ and N are determined. In the case of $\vec{\mu}$, one approach is to run the associated deterministic component, and then superimpose a Gaussian model on each of the spike times. If the deterministic component of the model produces spikes at

times τ_1, \dots, τ_M , then we set $\mu_i = \tau_i$. This sort of technique is also known by the name of “Parzen windows”. Another approach is to run the corresponding simple noise model many times and use the mean spike times for each spike τ_i to set μ_i . In the case of σ , we can again run the corresponding simple noise model many times and use the standard deviation of the spike corresponding to τ_i to set σ_i . A faster method that may be sufficient in some cases is to use a fixed value for σ . For this, σ should not be too small (*i.e.*, below 1 ms) in order to avoid overfitting single isolated spikes, and for similar reasons it should not be too large either. A reasonable range for σ is 1 to 10 ms. In the following simulations, whenever we have used a fixed σ , we have used the value $\sigma = 2$ ms. Likewise, N cannot be too large if we do not want to replace an isolated spike with a burst. It can be shown [1] that if $\sigma > 1$ and $N \approx 1/\sqrt{2\pi}\sigma$, then the number of spikes generated by a variable-rate model associated with a single Gaussian has essentially a Poisson distribution with mean 1. For simplicity, in the following simulations, we have always used $N_i = 1/\sqrt{2\pi}\sigma_i$. Ultimately, models with variable N 's and σ 's can be expected to be more flexible and suitable for complex spike trains.

In section 4, we illustrate several variants of the variable-rate Poisson model in the following forms (see table 4.6):

- $M(\vec{\mu}, \vec{\sigma}, (N_i = 1/\sqrt{2\pi}\sigma_i))$, with $\vec{\mu}$ and $\vec{\sigma}$ estimated from running the corresponding simple noise model 1000 times (used in the individual model comparisons).
- $M(\vec{\mu}, 2, 1/2\sqrt{2\pi})$, with $\vec{\mu}$ determined by the times of the spikes in the underlying

deterministic component of the model and $\bar{\sigma}$ set to be 2 *msec* for all spikes (used in the whole-class comparisons).

- All probabilities for all bins determined directly by running the simple noise model 1000 times (used to determine the optimal noise level for the pyramidal cell model, section 4.5.8).

The variable-rate model is also just a first-order Markov model without hidden states, with the well-known advantages and limitations of such models. It has no “memory”, thus (for instance) no built-in explicit refractory period. It is clear, however, that when N is small, extremely close spikes are in any case a rare event. For instance, in the case of a single Gaussian with $N = 0.04$, the probability of observing two spikes separated by less than 5 milliseconds is less than 0.157. More complex probabilistic models of discrete time series, such as higher-order Markov models, hidden Markov models, and artificial-neural-network hybrid models could be used [1], but will not be considered here.

Gaussian ISI model

In this model, the spiking times of the neurons are treated as Gaussian random variables. For the sake of argument, assume both the data traces and the deterministic compartmental model traces have the same number L of spikes, for a given level of current injection. The first assumption is that all interspike intervals (ISIs) are independent, and each ISI I_i has a Gaussian distribution with a mean μ_i , and a standard deviation σ_i . The model then

generates spikes at times t_1, \dots, t_L , and the corresponding ISIs I_1, \dots, I_L with probability:

$$p(D|M) = p(t_1, \dots, t_L) = p(I_1, \dots, I_L) = \prod_{j=1}^L \frac{1}{\sqrt{2\pi}\sigma_j} e^{-\frac{(I_j - \mu_j)^2}{2\sigma_j^2}} \quad (4.16)$$

The only problem now is to estimate the parameters of the model (*i.e.*, the μ_j 's and the σ_j 's), which again can be done from the data, or by running the simple noise model a large number of times.

The difficulty with the above method is that in many cases real neurons (and also model neurons with noise sources) will not always generate the same numbers of spikes for a given level of stimulation. Therefore we must modify the above equation for the likelihood to take into account extra/missing spikes. Thus the second assumption is as follows. Suppose that under a fixed set of conditions, a deterministic compartmental model yields a spike train with L spikes at times τ_1, \dots, τ_L . We can imagine that with added noise the probability of observing L spikes has a probability x ($0 \leq x \leq 1$), possibly close to 1. It is reasonable to assume that the probability of the number of spikes has a maximum at L . We can further model this distribution with a second parameter $0 \leq y < 1$, by setting $p(L+k) = y^k$ for $k = 1, 2, 3, \dots$, and likewise $p(L-k) = y^k$ for $k = 1, 2, \dots, L$ (where $p(i)$ denotes the probability of observing a total of i spikes in a given trace). This simple model assumes a symmetric fall-off from the mode, according to a geometric progression with rate y . Obviously, more complex profiles for $p(i)$ could be selected if necessary. Since $\sum_i p(i) = 1$, x and y must be related by

$$x + (y + y^2 + \cdots + y^L) + (y + y^2 + \cdots) = 1 \quad (4.17)$$

where the second term on the left-hand side represents the probability of getting less than L spikes, while the third term represents the probability of getting more than L spikes. This simplifies to

$$x + \frac{2y - y^{L+1}}{1 - y} = 1 \quad (4.18)$$

This relation can be solved exactly, at least for small values of L . For instance, when $L = 0$, $y = (1 - x)/(2 - x)$. When x is close to 1, we have the approximations (these are also upper bounds)

$$y \approx 1 - x \quad \text{for } L = 0 \quad \text{and} \quad y \approx \frac{1 - x}{2} \quad \text{for } L > 0 \quad (4.19)$$

We can now adapt equation (4.16) to the case where the number of spikes K is not necessarily equal to L , by multiplying the right-hand side by a term $q(J)$ as follows:

$$p(t_1, \dots, t_J) = p(I_1, \dots, I_J) = q(J) \prod_{j=1}^J \frac{1}{\sqrt{2\pi}\sigma_j} e^{-\frac{(I_j - \mu_j)^2}{2\sigma_j^2}}, \quad (4.20)$$

where $J = \inf(K, L)$, *i.e.*, J is the smaller of K and L . If $K < L$, the data is missing spikes which are normally present in the model, and we use

$$q(J) = p(K) = y^{L-K} \quad (4.21)$$

If $K > L$ ($K = L + k, k > 0$), the data contains extra spikes not normally present in the model. We then use

$$q(J) = p(K) \frac{(K-L)!}{(T-t_L)^{K-L}} = y^{K-L} \frac{(K-L)!}{(T-t_L)^{K-L}} \quad (4.22)$$

The extra factor $(K-L)!/(T-t_L)^{K-L}$ corresponds to the assumption that the extra k spikes are selected uniformly at random over the time interval $[t_L, T]$, where T is some fixed horizon. This term is needed to ensure proper normalization, that is $p(K)$ is the probability of observing $K > L$ spikes, or

$$\int_0^T \int_{t_{L+1}}^T \dots \int_{t_{L+1}+\dots+t_{K-1}}^T dI_{L+1} \dots dI_K = \frac{(T-t_L)^{K-L}}{(K-L)!} \quad (4.23)$$

Finally, if $K = L$

$$p(J) = p(K) = x \quad , \quad (4.24)$$

which is not very different from the original equation 4.16, since x is normally close to 1.

We shall denote such type of model by $L = L(\vec{\mu}, \vec{\sigma}, x, y)$.

The corresponding log-likelihoods are given by the following equations:

If $K < L$,

$$\log p(I_1, \dots, I_K) = (L-K) \log y + \sum_{j=1}^K \log \frac{1}{\sqrt{2\pi}\sigma_j} - \sum_{j=1}^K \frac{(I_j - \mu_j)^2}{2\sigma_j^2} \quad (4.25)$$

if $K = L$,

$$\log p(I_1, \dots, I_L) = \log x + \sum_{j=1}^L \log \frac{1}{\sqrt{2\pi}\sigma_j} - \sum_{j=1}^L \frac{(I_j - \mu_j)^2}{2\sigma_j^2} \quad (4.26)$$

and if $K > L$,

$$\begin{aligned} \log p(I_1, \dots, I_L) &= (K - L) \log y + \sum_{j=1}^L \log \frac{1}{\sqrt{2\pi}\sigma_j} - \sum_{j=1}^L \frac{(I_j - \mu_j)^2}{2\sigma_j^2} \\ &\quad + \log(K - L)! - (K - L) \log(T - t_L) \end{aligned} \quad (4.27)$$

In the case of missing spikes ($K < L$), a fixed penalty equal to $\log y$ is added for each missing spike. In the case of extra spikes ($K > L$), the penalty for each extra spike is slightly more complex. It consists of a penalty term of $\log y - \log(T - t_L)$ for each additional spike, and an additional $\log(K - L)!$ for $(K - L)$ extra spikes. In the following experiments, $(K - L)$ is usually very small, typically one or two spikes at most, and the $\log(K - L)!$ term can be ignored. Likewise a time horizon T can be selected so that the terms $\log(T - t_L)$ can be neglected. As a result, for most of the following, we use a fixed penalty $\log y$ for each additional or missing spike. Alternatively, one could use a more complex parametrization of the density $p(K)$ to achieve a similar effect. One exception is with the class comparisons in the next section, where integrals over the entire space of parameters lead to parameter regions with large numbers of missing or extra spikes. In this case we use the above equations without modification.

In summary, if the compartmental model generates L spikes and a data trace contains K spikes, to compute the likelihood we let $J = \min(L, K)$. We then pair the first J spikes of each trace to each other using a linear ordering without gaps. The corresponding J ISIs are compared using quadratic penalties. Finally, a fixed penalty is used for each extra or missing spike. This is obviously not always optimal. In particular, optimal pairing of ISIs between traces could have a more complex structure. As in the case of the previous variable-rate Poisson model, such a generative model can be coupled to a compartmental model by the way in which the parameters are determined. Again these parameters can be estimated by running the corresponding simple noise model many times. Another possibility is to determine the μ_i by taking the ISI observed in the corresponding deterministic compartmental model trace (we discuss the validity of this last approximation in the next section). In the following simulations, we use the Gaussian ISI model in two forms (see table 4.6):

- $M(\vec{\mu}, \vec{\sigma}, x, y)$ where all the parameters are estimated from the corresponding simple noise model (used in the individual model comparisons; see below);
- $M(\vec{\mu}, 1, x = 0.968 \text{ or } 0.921, y = (1 - x)/2 \text{ or } y = (1 - x) \text{ for } L = 0)$ (used in the whole-class comparisons). See below for the derivation of the estimates of x .

4.5 Examples

4.5.1 Objectives and methods

We now demonstrate the Bayesian approach on actual compartmental models, using both synthetic and natural data. The first question we ask is, given a data set and two candidate models, which model is more probable? We do this in two different cases, with models in the same class (which we refer to as the “within-class” case) and in different classes (the “between-class” case). The classes are distinguished by the fact that one class of model has an extra voltage-dependent ionic channel type not present in the other class. This corresponds to the case, frequently seen in practice, where the evidence for the existence of a particular channel type is inconclusive.

The second question we ask is, which of the two *classes* is more probable given the available data? This is the “whole-class” case. It serves to demonstrate that Bayesian methods are not limited to comparing individual models, but can be used to address any inference question.

We use both variable-rate Poisson models $M = M(\vec{\mu}, \vec{\sigma}, \vec{N})$, and Gaussian ISI models $M = M(\vec{\mu}, \vec{\sigma}, x, y)$ to calculate the likelihoods for the above two questions. Whenever feasible, we have estimated these parameters from runs of the corresponding simple noise model. We have resorted to fixed values only in the most computationally intensive class comparison tasks. The resulting inhomogeneity of models is not a problem as long as

the rules are clear within a comparison task, and that likelihoods obtained with different methods are not compared directly.

For both of the previous cases, noise was added to the model by adding Gaussian noise with zero mean and a fixed standard deviation to the injected current. Although the zero mean entails no loss of generality (if the mean were nonzero this would have been equivalent to the zero-mean case with a different level of DC current injection) the standard deviation of the noise was chosen arbitrarily, which is clearly suboptimal. Also, the data sets used for the previous cases were artificial data sets; it is of interest to see how Bayesian methods perform on real biological data sets. Therefore the third question was, given a biological data set and a realistic model which has been tuned to match the “average” behaviors of this data set (*i.e.*, the average spike times of the data), how can we use Bayesian methods to assign the most probable level of noise so that the variability in spike timing in the model accurately reflects the variability present in the data? In this case we use the simplest form of the variable-rate Poisson model as a contrast with the previous cases and also to show that this method can be used directly with a minimum of assumptions (*i.e.*, without assuming that spike timings are distributed as Gaussians).

The simulations in this section were carried out on Unix workstations running Linux and Solaris 2.3, using the neural simulation program GENESIS version 2.0 [10], with an extension library to implement the parameter search methods. Data analysis was performed using Matlab (The Mathworks Inc.).

4.5.2 The data

Synthetic data sets

We first use an artificial data set generated using a simple one-compartment model. The model contains four voltage-dependent ionic conductances: fast Na , delayed rectifier K (K_{DR}), inactivating K or “A-current” (K_A), and slow non-inactivating K or “M-current” (K_M). The kinetics of the channels are adapted from models of hippocampal conductances for the first three channels [41], and from a model of a bullfrog sympathetic ganglion cell in the case of K_M [25]. These particular conductances can reproduce a number of realistic behaviors. In particular, they can produce spiking outputs reasonably similar to those of regular-spiking cortical pyramidal cells [31]. Furthermore, with the inputs we used (see below) a range of behaviors, from not spiking to rapidly spiking with significant spike frequency adaptation, can be observed, as is seen in real pyramidal cells. Nevertheless, this model is essentially a simple “toy” model used only to demonstrate the Bayesian methodology.

The mathematical details of the model are given in appendices D and E. We refer to this model as the “target” since we seek to generate a model that replicates the data produced by this model. The inputs to the target model are a series of six constant current injections, with intensities ranging from 0.1 nA to 1.0 nA . The duration of the input stimuli is 200 ms , and the stimuli start 50 ms after the beginning of the simulation. The 50 ms prestimulus period was sufficient to cause the target model without noise to settle into a steady state, where “steady state” is defined as being the state where all variables (*e.g.* membrane

potential, channel conductances) are constant. Variability in the data is simulated by adding a certain amount of independent Gaussian noise to the injected current at each time step. The mean amplitude of the injected noise is 0 nA , with standard deviation 0.1 nA . Note that this represents a fairly large amount of added noise: the standard deviation of the noise ranges from 10% to 100% of the base current. The simulation were performed using the Crank-Nicolson implicit integration method [25], with time steps of $20 \mu\text{s}$.

While the data itself is generated here by a simple noise model, from now on this fact will be entirely forgotten, and the data studied as if it had originated in a true recording session. The data consists of a sequence of interspike intervals. The first interspike interval is operationally defined as the time between current onset and the first spike. For speed purposes, we chose to use spike times or interspike intervals (*ISIs*) as the basis for our model comparisons, rather than directly comparing membrane potential (V_m) traces. The Bayesian framework could also be applied directly to V_m traces (this would be necessary, for instance, if the data consisted of voltage recordings from a non-spiking cell). Figure 4.1a shows the voltage output of the target model without noise for the six currents used, while figure 4.1b shows the output of the same model with noise for five different current injections of 0.8 nA .




Figure 4.1: (a): The output of the target model without noise. Input currents are 0.1, 0.2, 0.4, 0.6, 0.8 and 1.0 nA , from bottom trace to top respectively. Traces represent membrane potential of the cell (V_m) and are offset on the y-axis for clarity. The current injection begins at 50 $msec$ and continues to the end of the trace (at 250 $msec$). (b): The output of the same model with noise, for five different current injections of 0.8 nA . To facilitate comparison, traces were overlaid.

Neurobiological data set

For the neurobiological data set we obtained a series of current-clamp recordings from a layer 2 superficial pyramidal cell in rat olfactory (piriform) cortex in a brain-slice preparation. Details on the brain-slice preparation can be found in [18, 36]. The slices were perfused with a solution to which the compounds CNQX ($30 \mu M$), APV ($100 \mu M$) and picrotoxin ($50 \mu M$) were added to block synaptic transmission without altering the spiking characteristics of the cell. Seven levels of current injection were used, and each current was injected six times, with a 20-second delay between injections. The delay was necessary to allow slow conductances to return to their baseline values between current injections. Although the spiking patterns were quite reproducible, variability between traces did exist (see figure 4.7, lower half).

4.5.3 The models

Deterministic component

We used two different compartmental models in the following sections: a simple one-compartment “toy” model (which was a variation of the simple target model described above) and a much more complex 15-compartment pyramidal cell model which was tuned to match the average spike times of the pyramidal cell data.

For the simple one-compartment model, two candidate model classes are used to illustrate the Bayesian approach to model comparison. The first model class is identical to the

target model, except that the maximal conductances of the four ionic conductances are not known. This corresponds to the case where the electrotonic structure of a cell, and the identity and kinetics of all voltage-dependent ionic channels, has been established, but the maximal conductance values of these channels are unknown. Thus the deterministic component of each model in the first class (C_4) is characterized by four parameters. We will use subscripts to refer to the number of parameters of the deterministic component of a model or a class. Thus, for instance, $C_4(\vec{\mu}, 2, 1/(2\sqrt{2\pi}))$ denotes the sub-class of variable-rate Poisson models, with four variable conductance parameters in the underlying deterministic component, $\vec{\mu}$ determined by the deterministic component, $\sigma = 2 \text{ msec}$, and $N = 1/(2\sqrt{2\pi})$. In what follows, we do not distinguish between classes and subclasses, as things should be clear from the context.

The second model class (C_3) is identical to the first one, except that we remove the K_A conductance from the model. This represents the case where the evidence for the presence of this conductance in the real cell is inconclusive. We remove the K_A conductance specifically because of our intuition that this conductance has a very small effect on the ability of the model to generate outputs that match the target data set. The Bayesian methods allow us to test this intuition in a quantitative and precise manner. Thus each model in the second class is characterized by three parameters.

For the analysis of the pyramidal cell data a separate model with 15 compartments, seven voltage- or calcium-dependent ionic conductances and a simple model for calcium

buffering was constructed. Details on this model are included in appendix F. All pyramidal cell models used had exactly the same deterministic component.

Probabilistic component

As discussed in the previous section, we use both the variable-rate Poisson model and the Gaussian ISI model to estimate likelihoods. In most cases, we determined the parameters of these models from the corresponding simple noise model (*i.e.*, the deterministic compartmental model with a noise source). Noise was added to the compartmental models as zero-mean Gaussian noise with a fixed amplitude. The amplitude of this noise was 0.1 nA for all models except for the pyramidal cell model, where a range of noise amplitudes from 0.025 nA to 1.6 nA were used.

For the variable-rate Poisson model, we used a bin size of 1 msec throughout, on the assumption that there is a vanishingly low probability of having more than one spike in the bin. In cases where Gaussian distributions were superimposed on spike times to generate probability profiles, the height of the Gaussian centered on each spike was set to be equal to $1/(\sqrt{2\pi}\sigma)$, as described above. The only other parameters needed are the means and standard deviations of the Gaussian distributions centered at each spike. For the comparisons of individual models, we ran the models with noise 1000 times and calculated the means and standard deviations of all spikes. We used the means and standard deviation of each particular spike (*e.g.*, the second spike at current level 3) to set these parameters.

For the class comparison (section 4.5.7) this was not computationally feasible, so we ran the corresponding deterministic component and used the spike times to determine the μ 's, with a standard deviation of 2 *msec*. We used a slightly larger standard deviation than for the Gaussian ISI model (see below) because we observed that the standard deviations of the spike times were typically somewhat larger than those for the ISIs (data not shown) (see table 4.6). For the pyramidal cell model we did not use superimposed Gaussian distributions centered on the spike times; instead we merely ran simulations with noise 1000 times, collected the spike counts in each bin and used this count to directly estimate the probability of having a spike in each bin. This has the advantage of being extremely simple in practice and free from any assumptions about the distributions of spike timings; however, there remains the problem of what probability to choose when no spikes fall in a bin (see below and appendix C).

For the Gaussian ISI model, we picked several models from both the three- and four-parameter classes with conductance parameters randomly chosen within the prior range and ran each of them 1000 times with noise. For each model, we computed the covariance matrices of the ISIs across runs and across different current injections. Several sample ISI histograms are shown in figure 4.2. The line through the distributions represent the ISI of the deterministic component of the model. Notice that this line is very close to the mean. A sample (normalized) covariance matrix is shown using a grayscale colormap in figure 4.3, with white representing high covariance, black representing anticovariance, and

gray representing no covariance. The covariance is very close to zero away from the main diagonal. Some anticorrelation exists between adjacent spikes in many cases. Additionally, the last current injection (spikes 65-106) shows a positive correlation between ISI times for the earlier spikes; in this case the amplitude of the input current was sufficiently large relative to the amount of current noise that the cell's firing pattern before spike frequency adaptation set in is quite regular, leading to high covariances. Similar results were obtained for other four-conductance models with different parameters and with three-conductance models with random parameter values from within the allowed range (not shown). In figure 4.4, we show a histogram of standard deviations for the ISIs of a typical model. Most of the standard deviations (over 60%) cluster in a single bin, corresponding to standard deviations ranging from 0 to 0.55 *ms*. Thus, to a first approximation, we see that

1. The various ISIs have a Gaussian distribution and are independent (the covariance matrix terms are typically close to 0 away from the main diagonal).
2. The mean of each Gaussian can be approximated by using the corresponding value in the same model run deterministically, *i.e.*, without any noise.
3. The standard deviations of the Gaussians are typically in the [0, 5] *ms* range. Most often they are close to 1 *ms*, so that whenever necessary we can reasonably use $\sigma_i = 1 \text{ ms}$.

The procedure to obtain the parameters x and y varied depending on the comparison

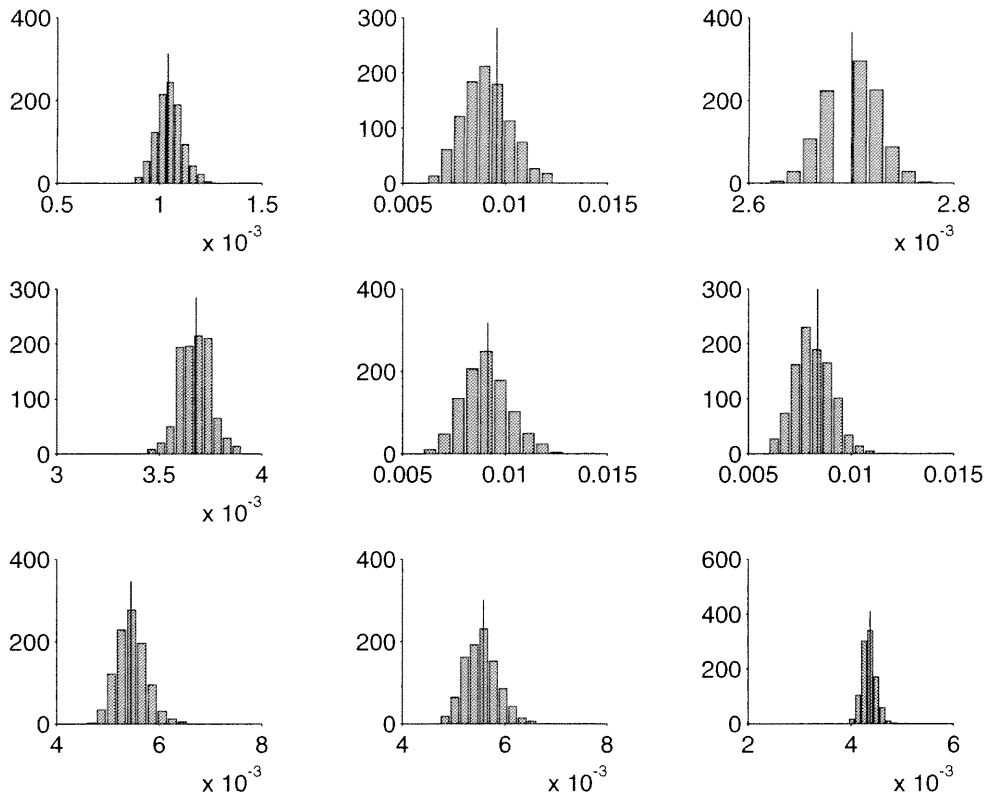


Figure 4.2: A group of ISI distribution histograms from a four-conductance model with noise and a range of input currents. The parameter settings were chosen randomly from within the allowed range. Each distribution represents a particular spike which occurred in most or all of 1000 runs of the model; nine spikes were chosen at random for display in this fashion. The x-axis is in units of seconds while the y-axis represents number per bin. The vertical line represents the position of the ISI for the model without noise. Note that the means of the distributions closely approximate the positions of the deterministic ISIs.




Figure 4.3: The normalized covariance matrix of a sample four-conductance model. The parameter settings were chosen randomly from within the allowed range. The model was run 1000 times with noise, and the standard deviations of the resulting ISI distributions were calculated for each spike. The x and y axes represent the spike number; the higher the spike number the later the spike in the trace. Spikes from different currents are given different spike numbers; thus spikes 1-2 are from current 1, spikes 3-7 are from current 2, etc. The covariance matrix has been normalized so that all diagonal entries are equal to 1.




Figure 4.4: The histogram of ISI standard deviation values for a four-conductance model with noise. The parameters of this model were chosen randomly from within the parameter ranges. Note that the majority of the values cluster in the range (0,1) msec. This analysis was repeated for five randomly chosen models in each of the two model classes (model 1 and model 2) and similar results were obtained in all cases. Absolute spike time standard deviation values were somewhat larger (not shown) but were mainly concentrated in the (0,2) msec range.

being performed. In sections 4.5.5 and 4.5.6, we ran the corresponding simple noise models 1000 times and used the incidence of missing/extra spikes to determine x and y separately for each level of current injection. In the case of the class comparison (section 4.5.7), this was not feasible, so we chose several models in each class with parameters randomly distributed in the allowable range. We used these to generate fixed estimates for x and y , which we then used for all models and all currents. The estimates were $x = 0.968$ (resp. $x = 0.921$) for the four-conductance (resp. three-conductance) class, with $y = (1 - x)/2$ ($y = (1 - x)$ when $L = 0$) (table 4.6), as described above. We also used a fixed value of $T = 400$ msec, twice the length of the current stimulus.

While the general question of the approximation of the simple noise models by the corresponding variable-rate Poisson models or Gaussian ISI models could receive a detailed Bayesian treatment, it is not the focus of the present study. From now on, the focus is on how well these models fit the given data sets. Finally, notice that another underlying assumption we use in the comparisons is that the level of the current injection in each trace is known. This is a standard assumption in real life situations, related to how the data is acquired under experimental control conditions.

4.5.4 Parameter estimation: determining optimal parameter values

For the “within-class” and “between-class” comparisons, we would like to choose reasonably optimal parameter sets for models in one or both classes. In this case, we use an automated

parameter search procedure to find a good match between the candidate models and the target model. The details of the search procedure itself are irrelevant to the Bayesian comparison, but are briefly described here for completeness. To speed up parameter searches, a modified target data set is taken by running the target model without noise. We use simulated annealing, with a modified simplex gradient-descent procedure [35], for the optimization itself. We have previously compared a number of parameter searching algorithms for neural models [42, 43], and found this one to be highly efficient. The optimization algorithm systematically varies the values of the adjustable parameters (four parameters for class 1, three parameters for class 2) until a best match is found. The matching function used is related to the likelihood of the corresponding Gaussian ISI model. M_4^{opt1} , M_4^{opt2} ... denote different four-compartment models derived by this procedure on different runs (table 4.6).

For the pyramidal cell model, a similar procedure was followed. A deterministic target data set was chosen by calculating the mean spike time for each spike in the real data. A parameter search method based on the genetic algorithm method [42, 43] was used to assign values to free parameters. Although we could have used simulated annealing as was done with the simpler models, we have found that genetic algorithms are often more suitable for models with larger numbers of parameters (in this case the pyramidal cell model had 24 parameters).

4.5.5 Comparing individual models from the same class

As the most basic test of the Bayesian method, we chose two independent models from the class of four-conductance models by running two separate parameter searches. Since the parameter search algorithm is stochastic, this produced two different candidate models M_4^{opt1} and M_4^{opt2} , both of which matched the spike timings of the data set within a fairly small error. The spike outputs of the deterministic component of the candidate models are shown as raster plots in figure 4.5(a). Each raster represents one spike. The response for different levels of current injection are offset on the y-axis, with the lowest level displayed as current 1 and the highest as current 6. The blue (middle) rasters represent the target model, the green (upper) rasters represent the first four-parameter model and the red (lower) rasters represent the second four-parameter model. The time axis begins at the beginning of the current injection, *i.e.*, at 50 msec after the beginning of the simulation. In this case, the two corresponding simple noise models almost always produced the same number of spikes per trace, so no penalty for missing/extra spikes had to be imposed. Specifically, in the Gaussian ISI model, the parameter x is always very close to 1 and is estimated by running the corresponding simple noise model 1000 times. The parameter y can be estimated using (4.19), but is not needed because $L = K$ and only (4.26) is used for the likelihoods.

The parameters of the resulting models are shown in table 4.1, along with the original parameters of the target model. It is important to notice that because of the approximations we make and the search procedure, the two models derived have good likelihoods, but do

parameter	target	model 1	model 2
$\bar{G}(Na)$	800.000	530.37	2546.44
$\bar{G}(K)$	100.000	119.74	135.07
$\bar{G}(K_A)$	50.000	25.51	31.64
$\bar{G}(K_M)$	100.000	93.40	37.38

Table 4.1: The final parameter values for two models from the same class, compared with the parameter values for the target model. The parameters were determined by a parameter search algorithm and represent maximal conductances of the ionic channel types (in S/m^2).

not necessarily correspond to a global maximum of the likelihood.

To compare two models, one needs to compare their posterior probabilities $p(M|D)$, using (4.3). Using (4.9) for the priors, and (4.14) or (4.20) for the likelihood, it is easy to compute the log-posterior probabilities of each model to within a normalizing factor. Notice that the term associated with the probability of the data ($P(D)$) is the same for both models, and therefore can be ignored in the comparison. The prior term could also have been ignored, since the prior used is uniform over the class. The results of the Bayesian analysis are shown in table 4.2, using both the variable-rate Poisson model ($M_4(\vec{\mu}, \vec{\sigma}, \vec{N})$) and the Gaussian ISI model ($M_4(\vec{\mu}, \vec{\sigma}, x, y)$) to estimate the likelihoods, with parameters determined as described above (see also table 4.6).

The analysis shows that model 1 is significantly more probable than model 2 given the




Figure 4.5: (a) A raster plot of the target model without noise versus the best-matching model for two four-compartment models (the “within-class” model comparison) without noise. (b) Raster plots for the “between-class” model comparison, also without added noise.

	variable-rate Poisson model		Gaussian ISI model	
	M_4^{opt1}	M_4^{opt2}	M_4^{opt1}	M_4^{opt2}
$\log P(D M)$	-73.83	-90.38	-8.81	-18.84
$\log P(M)$	-4.08	-4.08	-4.08	-4.08
$\log P(M D)$	-77.91	-94.46	-12.89	-22.92

Table 4.2: The likelihoods ($\log P(D|M)$), priors ($\log P(M)$), and posteriors ($\log P(M|D)$) (up to a constant) of the within-class comparison, using both the variable-rate Poisson model and the Gaussian ISI model.

data set, regardless of which statistical model is used. This is not surprising since the parameter values in model 1 are closer to the “target” parameters than those of model 2 (with the exception of parameter 3, corresponding to the K_A conductance, which is slightly closer to the correct value in model 2). Looking at the spike outputs of the corresponding deterministic components of the models in figure 4.5(a), we would agree that model 1 matches the data better than model 2, particularly at the highest current level. Bayesian methods have allowed us to determine this in a precise quantitative manner.

4.5.6 Comparing individual models from different classes

When comparing models from different classes, one may take the likelihood of the classes into account, so that Bayes’ theorem can be written as:

$$p(M, C|D) = \frac{p(D|M, C)p(M|C)p(C)}{p(D)} \quad (4.28)$$

In the present example, we will assume a uniform prior over the two model classes $p(C_3) = p(C_4) = 1/2$, so that in fact we can ignore the term $p(C)$. The term $p(M|C)$ is the prior on the model given the class and is calculated as in section 4.4.1. Thus the equation is the same as that used in the previous section, except for a multiplicative constant which is irrelevant to model comparison.

We again ran two parameter searches (one for each class) to determine optimal models in each class. In table 4.3, we show the parameters describing the target model as well as those of the two models found by the parameter search. Note that the final conductance values obtained by the parameter search method are extremely close to the “correct” values. Note also that even though the three-conductance model has no K_A conductance, the other parameter values are very close to the original values as well. In figure 4.5(b), we show the spike outputs of the two final models without noise as raster plots, compared to the original target model without noise. Here the blue (middle) rasters represent the target model, the green (upper) rasters represent the best matching four-parameter model and the red (lower) rasters represent the best matching three-parameter model.

The four-parameter model matches the spike times (and hence the ISIs) of the target data set almost perfectly, while the three-parameter model also generates an excellent match, with the exception of a spurious spike at the lowest level of current injection. In

parameter	target	four-conductance	three-conductance
$\bar{G}(Na)$	800.000	776.247	755.462
$\bar{G}(K)$	100.000	100.168	101.522
$\bar{G}(K_A)$	50.000	38.116	not applicable
$\bar{G}(K_M)$	100.000	102.119	104.974

Table 4.3: The final parameter values for a four-conductance and a three-conductance model, compared with the parameter values for the target model. The parameters were determined by a parameter search algorithm and represent maximal conductances of the ionic channel types (in S/m^2). Note that the four-parameter model is different from the two four-parameter models used in the within-class comparison.

order to calculate likelihoods, we ran the simple noise version of both models 1000 times and estimated the means and standard deviations of each spike time and ISI. The three-parameter model with noise generated a spurious spike in each of the 1000 runs, so we can only estimate an upper bound 0.001 for the parameter y in the Gaussian ISI model (*i.e.*, no spike will occur in less than one in a thousand trials). A further step could consist in putting a prior on y also, and estimating its posterior. The results of the Bayesian analysis are shown in table 4.4.

Note that the log probability for the three-conductance model is much smaller than that for the four-conductance model, both for the variable-rate and Gaussian ISI models. The magnitude of the difference is much larger for the Gaussian ISI model than for the variable-rate Poisson model, since, as discussed below, the Poisson model does not penalize missing/extra spikes as harshly as does the Gaussian ISI model (here we have a spike in the model which is missing from the data). These results are interesting in that our intuition was shown to be wrong; the K_A current appears to make a significant contribution to the spiking properties of this model. Note also that even though the prior for the M_3 model gives a smaller penalty than for the M_4 model (penalizing a more complex model), this has little effect on the final results, since the likelihood terms are much larger. The results of the Bayesian analysis are shown in table 4.4, both for the variable-rate Poisson model ($M_i(\vec{\mu}, \vec{\sigma}, \vec{N})$) and the Gaussian ISI model ($M_i(\vec{\mu}, \vec{\sigma}, x, y)$) ($i = 3$ or 4), with parameters determined as described above (see also table 4.6).

	variable-rate Poisson model		Gaussian ISI model	
	M_4^{opt}	M_3^{opt}	M_4^{opt}	M_3^{opt}
$\log P(D M)$	-69.31	-77.15	-3.24	< -81.45
$\log P(M)$	-4.08	-3.06	-4.08	-3.06
$\log P(M D)$	-73.39	-80.21	-7.32	< -84.51

Table 4.4: The likelihoods ($\log P(D|M)$), priors ($\log P(M)$), and posteriors ($\log P(M|D)$) (up to a constant) of the between-class comparison using both the variable-rate Poisson model and the Gaussian ISI model to determine the likelihoods.

4.5.7 Comparing entire model classes

Here we use equation (4.11) to calculate the likelihood values in order to compare the C_3 and C_4 model classes. To approximate the integral (4.11), we sample the entire parameter space by dividing the range of each parameter into 16 segments, running a model for each parameter set, and calculating the likelihood for each model. Since this requires $16^4 = 65536$ simulations for the four-parameter case and $16^3 = 4096$ simulations for the three-parameter case, we cannot run each one of the corresponding simple noise models 1000 times to estimate the means and standard deviations of the spike timings or the ISIs. Therefore we choose $\vec{\mu} = \vec{\mu}(\theta)$ to be the position of the spike in the corresponding deterministic component with parameters θ , both for the variable-rate Poisson and the Gaussian ISI models. The σ values used in the whole-class comparisons were fixed and estimated by running five models from

each class (with parameters randomly chosen from within the allowable range) 1000 times with noise. Although the value of σ differs between models and between currents in the same model, in general the σ for the ISIs is roughly 1 *msec* while that for the actual spike times is roughly 2 *msec*. The spike times have a larger σ since the variation in the timing of a current spike is not only due to the variation in the last interspike interval but also to the variation in all previous interspike intervals for that trace. Since we are sampling the entire parameter space, regimes where the candidate models produce extra or missing spikes are very common. This poses a problem for the Gaussian ISI model because we cannot estimate the x or y parameters directly for each model, since that would require running each of the thousands of models with noise a large number of times. Instead we used the 5 random models for each class, ran each of them 1000 times with noise, calculated an average value of x for each class and estimated y using equation (4.19). The average found was $x = 0.968$ (resp. 0.921) for the four-conductance case (resp. three-conductance). So, in the variable-rate Poisson case, we are comparing the classes $C_i(\vec{\mu}(\theta), \sigma = 2, N = 1/2\sqrt{2\pi})$ for $i = 3$ or 4. In the Gaussian ISI case, the classes are $C_4(\vec{\mu}(\theta), \sigma = 1, x = 0.968, y = (1 - x)/2)$ [(1 - x) when $L = 0$] and $C_3(\vec{\mu}(\theta), \sigma = 1, x = 0.921, y = (1 - x)/2)$ [(1 - x) when $L = 0$].

We approximate the likelihoods using our samples by

$$\int_{\theta \in C} p(D | \log(\theta), C) d \log(\theta) \approx \sum p(D | \log(\theta), C) \Delta \log(\theta) \quad (4.29)$$

and (after each parameter range has been divided into 16 increments):

$$\Delta \log(\theta) = (1/16 * 2 \log(4))^4 \quad (4.30)$$

(or $(1/16 * 2 \log(4))^3$ for the three-parameter case) which gives:

$$\begin{aligned} p(D|C) &\approx 1/(2 \log(4))^i * (1/16 * 2 \log(4))^i * \sum(p(D|\log(\theta))) \quad (4.31) \\ &\approx 1/(16^i) * \sum(p(D|\log(\theta))) \end{aligned}$$

with $i = 3$ for the three-conductance class and $i = 4$ for the four-conductance class (this is just the average of the likelihood over the class). The results are shown in table 4.5. It can be seen that the class of three-conductance models is quite inferior to the class of four-conductance models, regardless of which statistical model (variable-rate Poisson or Gaussian ISI) was used. Note also that with the Gaussian ISI model the class likelihood for the three-conductance class is of the same order as the likelihood of the “optimal” model in this class shown in table 4.4 (in fact, since the likelihood of the optimal model is an upper bound, the class likelihood is most likely greater than that of the optimal model). This is a consequence of the fact that the estimate of y for the class likelihood was much less accurate than for the optimal model (where it came from running the model with noise 1000 times) and therefore may have underestimated the penalty for missing/extra spikes for the most probable models.

Having a mapping of the space of likelihoods can be useful for other purposes as well.

	variable-rate Poisson model		Gaussian ISI model	
	C_4	C_3	C_4	C_3
$\log P(D M)$	-86.04	-136.86	-41.02	-81.45
$\log P(M)$	-0.69	-0.69	-0.69	-0.69
$\log P(M D)$	-86.73	-137.55	-41.71	-82.14

Table 4.5: The likelihoods ($\log P(D|M)$), priors ($\log P(M)$), and posteriors ($\log P(M|D)$) (up to constants) of the whole-class comparison using both the variable-rate Poisson model and the Gaussian ISI model to determine the likelihoods.

For instance, one could determine whether there is one unique maximum in the likelihood function or several, and how peaked the likelihoods are around each local maximum. One can also use this data for prediction problems, such as the problem of predicting when the first spike is most likely to occur given the class. Note also that if one has the resources to do a brute-force mapping of the parameter-space, using an optimization method for a parameter search is unnecessary. In general, this will not be possible, but Markov chain Monte Carlo methods are being developed for more efficient sampling of high-dimensional distributions [5, 38, 39].

4.5.8 Estimating the optimal amount of noise for a noisy pyramidal cell model

As a final illustration of Bayesian methods in compartmental neural modeling, we present here an application of these methods to a modeling problem involving real biological data. As we have noted, at present most compartmental models are deterministic, *i.e.*, for a given input they will produce spikes at exactly the same times. There are several reasons for this situation:

1. Most experimentalists do not systematically try to quantify the variability in their data, being more concerned with the mean responses (although this is changing; for instance see [30]).
2. It is not trivial to incorporate noise sources into a compartmental model in a way which reflects the biological processes of noise-generation. Furthermore, many possible noise sources (*e.g.*, channel noise) would be computationally expensive to model explicitly.
3. There has been no systematic procedure for comparing the outputs of a noisy compartmental model to the corresponding real data traces in order to find out how much noise needs to be added to the model in order to optimally match the variability in the data.

In this section we demonstrate how the Bayesian approach can solve problem 3. In order to do this we obtained recordings from a layer 2 superficial pyramidal cell in olfactory

(piriform) cortex in a brain-slice preparation (see [18, 36] for details on the experimental procedures). The recordings consist of seven different levels of DC current injection (total stimulation duration 1 *sec*), each repeated six times. Synaptic blockers (CNQX 30 μM , APV 100 μM and picrotoxin 50 μM) were added to the slice medium to eliminate background synaptic inputs and hence minimize externally derived noise sources. Despite this, the timing of the spikes was not identical from trial to trial, *i.e.*, there was a finite amount of variability in the spike timings (see figure 4.7, lower half for spike raster plots of the data). From the spike timing data we generated a set of mean firing times for each spike. We then built a compartmental model of this cell with 15 compartments, seven active channels and a simple model for calcium buffering (see appendix F for details). Parameters for the model were obtained using parameter search algorithms as described above. DC currents were injected into the model at the same level as in the real experiment (ranging from 0.08 to 0.2 *nA*), beginning 50 *msec* after the start of the simulation. Noise was added as zero-mean Gaussian-amplitude white noise as in previous sections. Several levels of noise differing only in the amplitude of the standard deviations were added, and for each level the model was run 1000 times for a total of 1.2 *sec* to collect spike statistics. To assess the probability of each level of noise, we used the variable-rate Poisson model as follows:

1. The spike times were collected into histograms of width 1 *msec*.
2. The probability of a spike occurring in any given bin was calculated as the total number of spikes in the bin divided by the number of runs (*nruns*; in this case 1000).

In the event that no spikes ever occurred in the bin, a probability of $0.5/nruns$ was used. This was done to prevent data traces which had spikes in these bins from being assigned a probability of zero. Similarly, bins in which spikes always occurred were assigned probabilities of $1 - 0.5/nruns$, to prevent data traces which did not have spikes in these bins from being assigned a probability of zero. In practice, the first case was fairly common but the second case never occurred with the noise amplitudes we used. Zero probabilities are avoided in this way because they would primarily be an artifact of the small sample size (see appendix C for a theoretical justification of this step).

3. Each data trace was compared to each probability profile, and the spike times in the data traces were assigned probability values. The negative log probabilities were calculated and summed for all data trace (six repeats for seven different currents).
4. This procedure was repeated for all noise levels tested.

As described above, this method gives an estimate of the negative log likelihood value for each level of noise. We ignore $P(D)$ as usual since it is irrelevant to model comparison. In this case, we also assume that all noise values have equal prior probabilities. Thus our procedure is simply a maximum-likelihood estimation of the optimal noise value.

The results of this analysis are shown in figure 4.6. The x-axis is the percent noise, which is the percentage of the energy in the noise relative to that in the DC current injection.

This was used because most of the effect of a white-noise stimulus is filtered out by the low-pass characteristics of the cell membrane, and therefore it is meaningless to compare white-noise amplitudes directly with DC current amplitudes. We performed an AC analysis of our simulated cell (not shown) and calculated that the signal attenuation with increasing frequency was such that white-noise stimulus would be only about 1.4% as effective as an equivalent DC stimulus of the same magnitude. The percent noise was calculated as follows:

$$\% \text{ noise} = \frac{(\text{noise amplitude}) * (\text{noise attenuation factor})}{\text{DC input current amplitude}} , \quad (4.32)$$

where the noise amplitude is the standard deviation of the amplitude distribution of the noise (the mean is zero) and the noise attenuation factor is about 1.4% as described above. Figure 4.6 clearly shows that there is an optimal level of noise at about 10% noise. This is shown graphically in figure 4.7, where we contrast the probability profiles of our model at three different noise levels with the same DC current level with raster plots of the spike times from the real data for one DC current level (which was the same as that used in the model). The raster plots are in the lower portion of the figure while the probability profiles are in the upper part of the figure. The lowest probability profile represents a very low level of noise, the middle profile represents a medium level, and the upper profile represents a very high level of noise. All profiles were computed using the variable-rate Poisson method. The extent of the profiles in the y direction has been scaled differently (1:2:4 ratio for low, medium and high noise, respectively) for each trace to facilitate comparison. The likelihood

of the spike data given the middle profile is significantly greater than the likelihood given either of the other profiles (it is more than 100 log units more probable than the upper profile and more than 1000 log units more probable than the lower profile). The reasons for this are clear. The low-noise profile represents a model which puts spikes very precisely at certain times; the real data has a finite amount of scatter around the mean spike times, and thus many of the spikes in the data will fall outside the distribution of the spikes in the model and will thus be assigned a very low probability value. The high-noise profile is almost a flat equiprobability distribution across the duration of the current injection. This means that while no data spike will be assigned a very low probability, neither will any of them be assigned very high probabilities. The medium-noise case is the best compromise between these extremes; the widths of the peaks for individual spikes can be seen to increase in accord with the spread of the data. Note also that models with noise are capable of producing spikes before or after the DC current was applied, since the noise current was applied for the entire duration of the simulation.

The significance of this procedure is that Bayesian methods have given us the ability to choose an important parameter for a model which is otherwise very difficult to choose in a rigorous manner. These methods can also quantify how good the choice is relative to other possible choices. This is not to say that the simple method we have used here is without its limitations. Most obviously, the variable-rate Poisson model treats all spikes as independent entities and thus misses the correlations between successive spikes in a trace.




Figure 4.6: A plot of the sum of the negative log likelihoods versus the percent noise added to the model. The negative log likelihoods consist of the sum of negative log likelihood values for all seven levels of DC current injection. The lowest value of the negative log likelihood occurs around the 10% noise level; this corresponds to the noise level with the highest likelihood (and therefore, given our assumptions, the highest probability).




Figure 4.7: Lower half: raster plots of spiking times for a layer 2 pyramidal cell with a constant DC current input of 0.1997 nA repeated six times. Current onset was at 50 msec and current duration was 1 sec . Each trial was separated by 20 sec intertrial intervals. Although the spiking pattern is fairly repeatable, the precise timing of the spikes varies from trial to trial. Upper half: three probability profiles for the pyramidal cell model using different noise levels. See text for details.

Further development of the Bayesian methodology in the context of compartmental neural models will have to address this question. Once better Bayesian methods are available, we will be able to make even stronger statements about which models are more probable than others.

Method:	Within-class		Between-class		Whole-class	
	M_4^{opt1}	M_4^{opt2}	M_4^{opt3}	M_3^{opt}	C_4	C_3
<u>Variable-rate Poisson</u>						
models run:	1000x with noise		1000x with noise		1x without noise	
<i>Parameters:</i>						
μ :	estimated from outputs of simple noise model runs		estimated from outputs of simple noise model runs		estimated as time of spikes in deterministic component	
σ :	estimated from outputs of simple noise model runs		estimated from outputs of simple noise model runs		2 msec	
C :	$1/(\sqrt{2\pi}\sigma)$		$1/(\sqrt{2\pi}\sigma)$		$1/(\sqrt{2\pi}\sigma)$	
bin size:	1 msec		1 msec		1 msec	
<i>continued on next page</i>						

continued from previous page

Method:	Within-class		Between-class		Whole-class	
	M_4^{opt1}	M_4^{opt2}	M_4^{opt3}	M_3^{opt}	C_4	C_3
<u>Gaussian ISI</u>						
models run:	1000x with noise		1000x with noise		1x without noise	
<i>Parameters:</i>						
μ :	estimated from outputs of simple noise model runs		estimated from outputs of simple noise model runs		estimated as time of spikes in deterministic component	
σ :	estimated from outputs of simple noise model runs		estimated from outputs of simple noise model runs		1 msec	
x :	estimated separately for each current from outputs of simple noise model runs		estimated separately for each current from outputs of simple noise model runs		0.968	0.921

continued on next page

<i>continued from previous page</i>						
Method:	Within-class		Between-class		Whole-class	
	M_4^{opt1}	M_4^{opt2}	M_4^{opt3}	M_3^{opt}	C_4	C_3
y :	not needed		not needed for M_4 ; see note for M_3		$(1 - x)$ for $M = 0$; else $(1 - x)/2$	
T :	400 msec		400 msec		400 msec	

Table 4.6: Parameter estimation methods used for the variable-rate Poisson model and for the Gaussian ISI model in the within-class, between-class and whole-class comparisons.

4.6 Discussion

We have briefly presented the Bayesian approach to modeling, and how it could be applied to computational neuroscience in general, and to single neuron compartmental models in particular. The main advantage of a Bayesian approach is obvious: it provides a principled approach to inference, with a strong foundation in probability theory. In fact one of the most compelling reasons in favor of Bayesian induction is its unicity, under a very small set of common sense axioms.

We have also shown that it would not be too difficult to make current compartmental

models more Bayesian, at least at the most basic level. The main advantage of Bayesian methods for compartmental modeling is the clarification of a number of modeling issues, at three different levels at least. First, a Bayesian approach forces one to clarify the prior knowledge going into the model, the data, and the hypotheses. Second, the Bayesian approach clarifies how to proceed with inference, *i.e.*, how to compare models with each other and with the data objectively, providing unambiguous unique answers to well-posed questions. In this way it establishes a formal basis for model comparisons. Finally, Bayesian methods require us to consider quantitatively issues of noise and variability in the data used to construct and test a model. This is a critical (and underconsidered) issue in computational neuroscience. Without this step, it is impossible to have a rigorous scientific discourse on models, to determine how well they fit the data, and ultimately to compare models and hypotheses. In this paper, we have demonstrated several ways to construct a probabilistic compartmental model out of its deterministic components.

The Bayesian framework is entirely open to, and actually encourages, questioning of any piece of information. It deals with the subjectivity inherent in the modeling process not by ignoring it but rather by incorporating it up front in the modeling process. It is fundamentally an iterative process in which models are progressively refined, and thus reflects and supports the nature of realistic neuronal modeling [8]. As Bayesian methods are applied to computational neuroscience, it is also likely that new computational ideas will emerge as has been the case in other fields ([29, 28, 33], and references therein). For

instance, the variable-rate Poisson model was in part suggested by the Bayesian approach, and the difficulties encountered with the Gaussian ISI model. The Bayesian approach can lead to a better understanding of the weaknesses of a model, and thereby help to generate better models. In addition, having an objective way of comparing models, and making predictions based on models, will become more important as the number, scope and complexity of compartmental models representing the same brain structures increases. It is our belief that compartmental models are still at a very early stage of development: for instance, there are not many Purkinje cell compartmental models in the world. Thus issues of model comparison and prediction are still somewhat hidden in the background, but will progressively become more central as neurobiological databases grow in size and complexity.

In the present case, the two main statistical models for the spike timings we used have complementary advantages and disadvantages. The variable-rate Poisson model handles extra or missing spikes more gracefully than the Gaussian ISI model, since it independently computes the probability of getting a spike in every time interval. The Gaussian ISI has a very ad-hoc and inelegant procedure for dealing with missing or extra spikes which requires the estimation of (at least) two new parameters. However, the Poisson model underpenalizes missing spikes (it can be shown that the penalty for a missing spike is roughly 1 log unit, in contrast to much higher values for the Gaussian ISI model on average). Thus neither model is an ideal model of the spike-generating process in this particular data set, and therefore

more such models need to be developed in the future. The variable-rate Poisson model, however, is extremely simple and flexible and can provide at least a first approximation in a wide range of situations. This was illustrated in the last section, where the variable-rate Poisson model was used to derive an estimate of the noise input needed for a compartmental model to optimally match a biological data set.

The main drawback of the Bayesian approach is that it can be computationally intensive, especially when averages need to be computed over high-dimensional distributions. For the largest compartmental models, it is unlikely that one will be able to carry out a complete Bayesian analysis on currently available computers. But continuing progress in Monte Carlo [5, 39] and other approximation techniques, as well as steady increases in raw computing power in workstations and parallel computers, is likely to make these methods feasible for compartmental neural models at a reasonable level of complexity in the near future. In addition, it should be clear that a simple method such as the variable-rate Poisson method can already be used immediately, in conjunction with the most complex existing compartmental models, to estimate likelihoods of large data sets and compare models pairwise.

Given the computational demands of the Bayesian approach, however, some modelers might wonder what these methods provide that cannot be obtained merely by comparing deterministic models according to some error criterion. First of all, one of our goals here has been to generate generic probabilistic models of spiking processes which can be overlaid on compartmental models, in contrast to error functions in parameter searches, which are

typically constructed to strongly select for particular features of the data. Secondly, error functions are not probabilities. Showing that a particular model matches a data set given a particular error criterion better than another model says nothing about how much more probable one model is relative to the other. For this a probabilistic model is clearly essential. It is not difficult to imagine an error function which overwhelmingly favors one model over another (for instance, if the error function gives an enormous weight to the first spike in the data and virtually ignores the rest) even though the two models are almost equally probable given a reasonable probabilistic spike-generation model. Thus, establishing a probabilistic model and, especially, showing that it is reasonable given the degree of variability in the data (section 4.5.8) are essential for proper model comparison.

As the scope and complexity of compartmental modeling in neuroscience continues to increase, it will become increasingly important to have an objective way of comparing models, and making predictions based on those models. Because computational neuroscience and compartmental modeling are still at an early stage of development, issues of rigorous model comparison and prediction have not yet been considered seriously by most investigators. However, as neurobiological databases grow in size and complexity, our standards for model comparison will also have to be raised in order for progress to be made [8]. We believe that Bayesian methods will provide a solid foundation for these efforts.

Appendix A: The Cox-Jaynes axioms

The objects we deal with in inference are propositions about the world. For instance, a typical proposition X is “neuron A is firing an action potential at time t ”. A proposition is either true or false, and we denote by \bar{X} the complement of a proposition X . A hypothesis H about the world is a proposition, albeit a possibly complex one composed of the conjunction of many more elementary propositions. A model M can also be viewed as an hypothesis. The difference is that models tend to be very complex hypotheses involving a large number of parameters. In discussions where parameters are important, we will consider that $M = M(\theta)$ where θ is the vector of all parameters. A complex model M can easily be reduced to a binary proposition in the form: “model M accounts for data D with an error level ϵ ” (this is a vague statement that will be made more precise in the following discussion). But for any purpose, in what follows there is no real distinction between models and hypotheses.

Whereas propositions are either true or false, we wish to reason in the presence of uncertainty. Therefore the next step is to consider that, given a certain amount of information I , we can associate with each hypothesis a degree of plausibility or confidence (also called in the literature degree or strength of belief). Let us represent it by the symbol $\pi(X|I)$. While $\pi(X|I)$ is just a symbol for now, it is clear that in order to have a scientific discourse one should be able to compare degrees of confidence. That is, for any two propositions X and Y , either we believe in X more than in Y , or we believe in Y more than in X , or we believe in both equally. Let us use the symbol “ $>$ ” to denote this relationship, so that we

write $\pi(X|I) > \pi(Y|I)$ if and only if X is more plausible than Y . It would be very hard not to agree that in order for things to be sensible, the relationship “ $>$ ” should be transitive. That is if X is more plausible than Y , and Y is more plausible than Z , then X must also be more plausible than Z . More formally this is the first axiom:

$$\pi(X|I) > \pi(Y|I) \quad \text{and} \quad \pi(Y|I) > \pi(Z|I) \Rightarrow \pi(X|I) > \pi(Z|I) \quad (4.33)$$

This axiom is trivial; it has, however, an important consequence: “ $>$ ” is an ordering relationship, and therefore degrees of belief can be expressed by real numbers. That is, from now on $\pi(X|I)$ represents a number. This does not mean that such a number is easy to calculate, but merely that such a number exists, and the ordering among hypotheses is reflected in the ordering of real numbers. To proceed any further and stand a chance of calculating degrees of beliefs we need additional axioms or rules for relating numbers representing strengths of belief. The amazing fact is that only two additional axioms are needed to constrain the theory entirely. This axiomatic presentation is usually attributed to Cox and Jaynes [13, 22]. To better understand these two remaining axioms, the reader may imagine a world of very simple neurons, where at each instant in time a given neuron can either be on or off. Thus, all the elementary hypotheses or propositions in this world, at a given time, have the simple form: “neuron A is on” or “neuron A is off”. Clearly, the more confident we are that neuron A is on, the less confident we are that neuron A is off. So, for any given proposition X , there should be a relationship between $\pi(X|I)$ and

$\pi(\bar{X}|I)$. Without assuming anything about this relationship, it is sensible to consider that, everything else being equal, the relationship should be the same for all neurons and for all types of background information, *i.e.*, for all propositions X and I . Thus, in mathematical terms, the second axiom states that there exists a function F such that

$$\pi(\bar{X}|I) = F[\pi(X|I)] \quad (4.34)$$

The third axiom is only slightly more complex. Consider this time two neurons A and B and the corresponding 4 possible joint states. Then our degree of belief that A is on and B is off, for instance, naturally depends on our degree of belief that neuron A is on, and our degree of belief that neuron B is off, knowing that A is on. Again, it is sensible that this relationship be independent of the neurons considered and the type of background information I . Thus, in mathematical terms, the third axiom states that there exists a function G such that

$$\pi(X, Y|I) = G[\pi(X|I), \pi(Y|X, I)] \quad (4.35)$$

So far, we have not said much about the information I . I is a proposition corresponding to the conjunction of all the available pieces of information. I can represent background knowledge, such as general anatomical or physiological information in neuroscience. I can also include specific experimental or other data. When it is necessary to focus on a particular corpus of data D , we can write $I = (I, D)$. In any case, I is not necessarily fixed and can

be augmented with, or replaced by, any number of symbols representing propositions, as already seen in the right hand side of equation (4.35). When data is acquired sequentially, for instance, we may write $I = (I, D_1, \dots, D_n)$. Inside a discussion where I is well defined and fixed, it can be dropped altogether from the equations.

The three axioms above determine entirely, up to scaling, how to calculate degrees of belief. In particular, one can prove that there is always a rescaling w of degrees of beliefs such that $p(X|I) = w(\pi(X|I))$ is in $[0, 1]$. Furthermore, p is unique and satisfies all the rules of probability. Specifically, if degrees of belief are restricted to the $[0, 1]$ interval, then the function F and G must be given by: $F(x) = 1 - x$ and $G(x, y) = xy$. The corresponding proof will not be given here and can be found, for instance, in [13, 23]. As a result, the second axiom can be rewritten as the Sum Rule of probability:

$$p(X|I) + p(\bar{X}|I) = 1 \tag{4.36}$$

And the third axiom as the Product Rule:

$$p(X, Y|I) = p(X|I)p(Y|X, I) \tag{4.37}$$

From here on, we can replace degrees of confidence by probabilities. Notice that if uncertainties are removed, that is if $p(X|I)$ is 0 or 1, then equation (4.36) and equation (4.37) yield, as a special case, the two basic rules of deduction or Boolean algebra, for the negation and conjunction of propositions [(1) “ X or \bar{X} ” is always true; (2) “ X and Y ” is true if and

only if both X and Y are true].

The immediate consequence of equation (4.36) and equation (4.37) is the so-called Bayes' theorem:

$$p(X|Y, I) = \frac{p(Y|X, I)p(X|I)}{p(Y|I)} = p(X|I) \frac{p(Y|X, I)}{p(Y|I)} \quad (4.38)$$

Bayes' theorem is fundamental because it allows inversion: interchanging conditioning and non-conditioning propositions. In a sense, it embodies learning since it describes exactly how to update our degree of belief $p(X|I)$ in X , in light of the new piece of information provided by Y to obtain the new $p(X|Y, I)$. $p(X|I)$ is also called the prior probability (or “prior” for short) and $p(X|Y, I)$ the posterior probability with respect to Y (or “posterior” for short). This rule can also be iterated as more information becomes available.

Appendix B: Error bars

For illustration purposes, consider a modeling situation with one parameter θ , and a uniform prior. Let $f(\theta) = -\log p(D|\theta)$ be the negative log-likelihood of the data. Under mild differentiability conditions, a maximum likelihood estimator θ_0 satisfies $f'(\theta_0) = 0$. Therefore in the neighborhood of θ , we can expand $f(\theta)$ in a Taylor's series:

$$f(\theta) \approx f(\theta_0) + \frac{1}{2}f''(\theta_0)(\theta - \theta_0)^2 \quad (4.39)$$

or

$$p(D|\theta) = e^{-f(\theta)} \approx C e^{-\frac{1}{2}f''(\theta_0)(\theta-\theta_0)^2} , \quad (4.40)$$

where $C = e^{f(\theta_0)}$. So the likelihood, as well as the posterior $p(\theta|D)$, locally behave like a Gaussian, with a standard deviation $1/\sqrt{f''(\theta_0)}$, associated with the curvature of f .

Appendix C: Dirichlet distributions

Let $\lambda = \lambda(t)$ be the probability of firing in a given time bin t in a variable-rate model. A natural prior on λ is the Dirichlet distribution with parameters

$$p(\lambda) = \frac{\Gamma(\alpha)}{\Gamma(\alpha q_1)\Gamma(\alpha q_2)} \lambda^{\alpha q_1 - 1} (1 - \lambda)^{\alpha q_2 - 1} = Z \lambda^{\alpha q_1 - 1} (1 - \lambda)^{\alpha q_2 - 1} , \quad (4.41)$$

also called a Beta distribution in the two-dimensional case. The parameters $\vec{\sigma}\alpha$ and $q = (q_1, q_2)$ of the distribution must satisfy $\alpha, q_i \geq 0$ and $\sum q_i = 1$. The parameter α determines how peaked the distribution is around its mean associated with $E(\lambda) = q_1$. Assume now that the data D for the same time bin consists of a total of n observations, f of which correspond to a firing event. The likelihood associated with this observation is $p(D|\lambda) = \lambda^f (1 - \lambda)^{(n-f)}$. A maximum likelihood estimation of λ yields (not surprisingly) $\lambda = f/n$, with the problem that this estimate is 0 if $f = 0$. A simple calculation shows that the MAP estimate for λ with the Dirichlet prior is given instead by

$$\lambda = \frac{f + \alpha q_1 - 1}{n + \alpha - 2} \quad (4.42)$$

The effect of the prior amounts to adding pseudo-counts to the observed data. With a symmetric Dirichlet prior (for instance, $q_1 = q_2 = 0.5$), as long as we choose $\alpha > 2$, the MAP estimate can never be zero. The posterior distribution $p(\lambda|D)$ is also a Dirichlet distribution. For this reason, the Dirichlet distribution is said to be conjugate for the corresponding binomial/multinomial model. The Bayesian analysis of the pyramidal cell model using the variable-rate Poisson model in section 4.5.8 is essentially equivalent to using a Dirichlet prior with $q_1 = q_2 = 0.5$ and $\alpha \approx 3.001$ for $n = 1000$.

Appendix D: Model equations for the one-compartment model

The basic equation for the compartmental models is the current-balance equation:

$$C \frac{dV_m}{dt} + \frac{V_m - E_m}{R_m} + I_{ionic} = I_{stim} \quad , \quad (4.43)$$

where C is the membrane capacitance, V_m is the membrane potential, R_m is the membrane resistance, E_m is the “leak current” reversal potential (a poorly defined and yet nearly universal convention among compartmental modelers [10, 25]), I_{ionic} is the sum of currents through voltage-dependent ionic conductances, and I_{stim} is the input current. I_{ionic} can be further subdivided as follows:

$$\begin{aligned}
I_{ionic} &= I_{Na} + I_{K_{DR}} + I_{K_A} + I_{K_M} & (4.44) \\
&= G_{Na}(V_m - E_{Na}) + G_{K_{DR}}(V_m - E_K) \\
&\quad + G_{K_A}(V_m - E_K) + G_{K_M}(V_m - E_K) \ ,
\end{aligned}$$

where G_X is the conductance value for channel X , E_{Na} is the reversal potential for Na , and E_K is the reversal potential for K . These conductance values are time- and voltage-dependent according to the following equations:

$$G_{Na} = \overline{G}_{Na} m^2 h \quad (4.45)$$

$$G_{K_{DR}} = \overline{G}_{K_{DR}} n \quad (4.46)$$

$$G_{K_A} = \overline{G}_{K_A} x y \quad (4.47)$$

$$G_{K_M} = \overline{G}_{K_M} z \quad (4.48)$$

The parameters \overline{G}_C , where C represents an ionic channel species, are the variable parameters of the models. The gating variables m , h , n , x , y , and z all can be described by the Hodgkin-Huxley formalism [10, 19] as:

$$\tau_X \frac{dX}{dt} = X_\infty - X \quad (4.49)$$

$$\tau_X = \frac{1}{\alpha_X + \beta_X} \quad (4.50)$$

$$X_\infty = \frac{\alpha_X}{\alpha_X + \beta_X} \quad (4.51)$$

for a given gating variable X , where τ_X , X_∞ , α_X , and β_X are all functions of membrane potential V_m . α_X is the forward rate constant and β_X is the reverse rate constant for the channel opening reaction [19]. These functions are listed below. In some cases α and β constants were used to compute X_∞ and τ values, while in other cases the latter values were calculated directly. Note that V_m is in volts. The channel models are adapted from those in [41] except for the K_M channel which is adapted from [25]. To get the cells to spike in a manner similar to that of regular-spiking cortical pyramidal cells the kinetics of the potassium channels were altered relative to those in the references.

Na channel:

$$\alpha_m = \frac{-15008 - 32000 V_m}{-1 + \exp(-(V_m + 0.0469)/0.004)} \quad (4.52)$$

$$\beta_m = \frac{5572 + 280000 V_m}{-1 + \exp((V_m + 0.0199)/0.005)} \quad (4.53)$$

$$\alpha_h = \frac{128}{\exp((V_m - 0.043)/0.018)} \quad (4.54)$$

$$\beta_h = \frac{4000}{1 + \exp(-(V_m - 0.020)/0.005)} \quad (4.55)$$

delayed-rectifier K channel:

$$\alpha_n = \frac{-398.4 - 16000 V_m}{-1 + \exp(-(V_m + 0.0249)/0.005)} \quad (4.56)$$

$$\beta_n = \frac{250}{\exp((V_m + 0.040)/0.040)} \quad (4.57)$$

K_A channel:

$$\alpha_x = \frac{-938 - 20000 V_m}{-1 + \exp(-(V_m + 0.0469)/0.01)} \quad (4.58)$$

$$\beta_x = \frac{348.25 + 17500 V_m}{-1 + \exp((V_m + 0.0199)/0.01)} \quad (4.59)$$

$$\alpha_y = \frac{1.6}{\exp((V_m + 0.073)/0.018)} \quad (4.60)$$

$$\beta_y = \frac{50}{-1 + \exp(-(V_m + 0.0499)/0.005)} \quad (4.61)$$

K_M channel:

$$z_\infty = \frac{1.0}{1.0 + \exp(-(V_m + 0.035)/0.01)} \quad (4.62)$$

$$\tau_z = \frac{1.0}{6.6 * \exp((V_m + 0.035)/0.02) + \exp(-(V_m + 0.035)/0.02)} \quad (4.63)$$

Appendix E: Fixed parameters of the one-compartment model

$$cell\ area = 4\pi * (10\ \mu m)^2 = 1.2566 * 10^{-9}\ m^2 \quad (4.64)$$

$$C = CM * area = (0.008\ \frac{F}{m^2}) * (1.2566 * 10^{-9}\ m^2) = 1.00531 * 10^{-11}\ F \quad (4.65)$$

$$R_m = RM * area = (2.0\ \frac{ohm}{m^2}) * (1.2566 * 10^{-9}\ m^2) = 1.59155 * 10^9\ ohm \quad (4.66)$$

$$E_m = -0.077\ V \quad (4.67)$$

$$E_{Na} = 0.055\ V \quad (4.68)$$

$$E_K = -0.090\ V \quad (4.69)$$

Appendix F: The pyramidal cell model

The pyramidal cell model was based on an anatomical reconstruction of a layer 2 superficial pyramidal cell in olfactory (piriform) cortex; this reconstruction had 2670 compartments. We simplified the morphology of this model down to a 15-compartment model using essentially the method of Bush and Sejnowski [11]. Our model also included several active ionic conductances located in the soma. No active conductances were placed in the dendrites, because there is little evidence for such channels in this cell type at this time. The conductances included fast inactivating Na , fast persistent Na , delayed-rectifier K , slow non-inactivating K (“M-current”), fast and slow calcium-activated K , and high voltage-activated Ca . The model also contained a simple model of calcium buffering. Space does not permit a full description of the model here, but details on the model and full parameter sets can be obtained by contacting the authors.

Acknowledgements

We gratefully acknowledge the assistance of Alexander Protopapas, who contributed the experimental data on which section 4.5.8 is based. Alex also generated the parameters for the anatomical reconstruction of the layer 2 pyramidal cell which was a starting point for the pyramidal cell model described in that section.

The work of Pierre Baldi was supported in part by a grant from the Office of Naval Research.

Bibliography

- [1] P. F. Baldi. Probabilistic Models of Neuronal Spike Trains. In Marco Gori, editor, *Lecture Notes in Artificial Intelligence: Adaptive Processing of Temporal Information*. Springer Verlag, New York, 1997. In press.
- [2] P. F. Baldi and S. Brunak. *Bioinformatics: the Machine Learning Approach*. MIT Press, Cambridge, MA, 1998.
- [3] V. Barnett. *Comparative Statistical Inference*. John Wiley, New York, 1982.
- [4] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer Verlag, New York, 1985.
- [5] J. Besag, P. Green, D. Higdon, and K. Mengersen. Bayesian Computation and Stochastic Systems. *Statistical Science*, 10(1):3–66, 1995.
- [6] J. M. Bower. Modeling the Nervous System. *Trends in Neurosciences*, 15(11):411–412, 1992.

- [7] J. M. Bower. Reverse Engineering the Nervous System: an *In Vivo*, *In Vitro*, and *In Computo* Approach to Understanding the Mammalian Olfactory System. In S. Zornetzer, J. Davis, and C. Lau, editors, *An Introduction to Neural and Electronic Networks*, 2nd. Ed. Academic Press, San Diego, 1995.
- [8] J. M. Bower. What Will Save Neuroscience? *Neuroimage*, 4(3):S29–S33, 1996.
- [9] J. M. Bower, editor. *Computational Neuroscience: Trends in Research 1997*. Plenum Press, New York, 1997.
- [10] J. M. Bower and D. Beeman. *The Book of Genesis*. Springer Verlag, New York, 1995.
- [11] P. C. Bush and T. J. Sejnowski. Reduced Compartmental Models of Neocortical Pyramidal Cells. *Journal of Neuroscience Methods*, 46:159–166, 1993.
- [12] P. Cheeseman. An Inquiry into Computer Understanding. *Computational Intelligence*, 4:57–142, 1988. with discussion.
- [13] R. T. Cox. Probability, Frequency and Reasonable Expectation. *American Journal of Physics*, 14:1–13, 1964.
- [14] E. DeSchutter and J. M. Bower. An Active Membrane Model of the Cerebellar Purkinje Cell I. Simulation of Current Clamps in Slice. *Journal of Neurophysiology*, 71:375–400, 1994.

- [15] E. DeSchutter and J. M. Bower. An Active Membrane Model of the Cerebellar Purkinje Cell II. Simulation of Synaptic Responses. *Journal of Neurophysiology*, 71:401–419, 1994.
- [16] P. C. Gregory and T. J. Loredo. A New Method for the Detection of a Periodic Signal of Unknown Shape and Period. *The Astrophysical Journal*, 398:146–168, 1992.
- [17] S. F. Gull. Developments in Maximum Entropy Data Analysis. In J. Skilling, editor, *Maximum Entropy and Bayesian Methods*, pages 53–71. Kluwer, Dordrecht, 1989.
- [18] M. E. Hasselmo, B. P. Anderson, and J. M. Bower. Cholinergic Modulation of Cortical Associative Memory Function. *Journal of Neurophysiology*, 67(5):1230–1246, 1992.
- [19] B. Hille. *Ionic Channels of Excitable Membranes, 2nd. Ed.* Sinauer Associates Inc., Sunderland, Massachusetts, 1992.
- [20] P. G. Hoel, S. C. Port, and C. J. Stone. *Introduction to Probability Theory*. Houghton Mifflin, Boston, 1971.
- [21] D. Jaeger, E. DeSchutter, and J. M. Bower. The Role of Synaptic and Voltage-Gated Currents in the Control of Purkinje-Cell Spiking – a Modeling Study. *Journal of Neuroscience*, 17(1):91–106, 1997.

- [22] E. T. Jaynes. Bayesian Methods: General Background. In J. H. Justice, editor, *Maximum Entropy and Bayesian Methods in Statistics*, pages 1–25. Cambridge University Press, Cambridge, 1986.
- [23] E. T. Jaynes. *Probability Theory: The Logic of Science*. Unpublished, 1995. In preparation.
- [24] W. H. Jeffreys and J. O. Berger. Ockham’s Razor and Bayesian Analysis. *American Scientist*, 80:64–72, 1992.
- [25] C. Koch and I. Segev. *Methods in Neuronal Modeling*. MIT Press, Cambridge, Massachusetts, 1989.
- [26] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.
- [27] M. S. Lewicki. Bayesian Modeling and Classification of Neural Signals. *Neural Computation*, 6(5):1005–1030, 1994.
- [28] D. J. C. MacKay. A Practical Bayesian Framework for Back-Propagation Networks. *Neural Computation*, 4(3):448–472, 1992.
- [29] D. J. C. MacKay. Bayesian Interpolation. *Neural Computation*, 4(3):415–447, 1992.
- [30] Z. F. Mainen and T. J. Sejnowski. Reliability Of Spike Timing In Neocortical Neurons. *Science*, 268:1503–1506, 1995.

- [31] D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative Electrophysiology of Pyramidal and Sparsely Spiny Stellate Neurons of the Neocortex. *Journal of Neurophysiology*, 54(4):782–806, 1985.
- [32] R. M. Neal. Probabilistic Inference Using Markov Chain Monte Carlo Methods. Technical Report, Department of Computer Science, University of Toronto, Canada, 1993.
- [33] R. M. Neal. Bayesian Learning for Neural Networks. PhD Thesis, Department of Computer Science, University of Toronto, Canada, 1995.
- [34] S. J. Press. *Bayesian Statistics: Principles, Models, and Applications*. John Wiley, New York, 1989.
- [35] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C, 2nd. Ed.* Cambridge University Press, New York, 1992.
- [36] A. D. Protopapas and J. M. Bower. Noradrenaline Increases Information Without Changing Encoding in Cortical Neurons. *Unpublished*, 1997.
- [37] M. Shadlen and W. Newsome. Noise, Neural Codes And Cortical Organization. *Current Opinion in Neurobiology*, 4:569–579, 1994.
- [38] A. F. Smith and G. O. Roberts. Bayesian Computation Via the Gibbs Sampler and Related Markov Chain Monte Carlo Methods. *Journal of the Royal Statistical Society*, 55(1):3–23, 1993.

- [39] A. F. M. Smith. Bayesian Computational Methods. *Philosophical Transactions of the Royal Society of London A*, 337:369–386, 1991.
- [40] W. R. Softky and C. Koch. The Highly Irregular Firing Of Cortical-Cells Is Inconsistent With Temporal Integration Of Random EPSPs. *Journal Of Neuroscience*, 13(1):334–350, 1993.
- [41] R. D. Traub, R. K. S. Wong, R. Miles, and H. Michelson. A Model of a CA3 Hippocampal Pyramidal Neuron Incorporating Voltage-Clamp Data on Intrinsic Conductances. *Journal of Neurophysiology*, 66(2):635–650, 1991.
- [42] M. C. Vanier and J. M. Bower. A Comparison of Automated Parameter-Searching Methods for Neural Models. In J. M. Bower, editor, *Proceedings of the 1995 Computational Neuroscience Conference (CNS*95), Monterey (CA)*. Academic Press, 1996.
- [43] M. C. Vanier and J. M. Bower. A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models. *Journal of Computational Neuroscience*, 7(2):149–171, 1999.

Part IV

Modeling the piriform cortex

Introduction

The remaining chapters describe work done building a realistic model of mammalian olfactory (piriform) cortex and some of the scientific explorations I have performed on this model. This work is a continuation of Matthew Wilson's modeling work in this system [1]. However, the model described in these chapters bears almost no resemblance to the earlier model. The new model was constructed from the ground up based on new data, and incorporates major improvements in the accuracy of the component cell models, the inputs to the model from olfactory bulb mitral cells, and the network connectivities. The model also incorporates a number of entirely new neuronal features including neuromodulation at several levels, synaptic facilitation and depression, NMDA synapses and multiple types of inhibitory synapses onto apical dendrites. The resulting model is the most complex realistic network models of this brain region ever constructed and will serve as a valuable platform for exploring this system for years to come. In these chapters I will describe in some detail the process by which the model was constructed and validated, what has been learned from the model, as well as a number of areas where the model can be improved and/or extended

in the future.

The material on simplifying cell morphologies in the next chapter is adapted from chapter 12 of [2].

Bibliography

- [1] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.
- [2] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.

Chapter 5

Building Models of Single Neurons in Piriform Cortex

One feature of the piriform cortex model described in this thesis which departs significantly from the previous version of the model [48] is that the individual neuron models were built to replicate experimental data very accurately. In this chapter, I will describe the construction of these models in some detail.

5.1 Piriform cortex neurons

In the mammalian brain, piriform cortex pyramidal neurons consist of a cell body (soma), a long apical dendrite extending to the surface of the cortex (spanning layers 1b and 1a), and a number of basal dendrites that receive local excitatory input in layer 3 [16, 39]. The most

numerous type of pyramidal neurons are superficial pyramidal neurons, whose cell bodies are densely packed into layer 2 (figure 5.1, left). In addition, some pyramidal neurons are located deeper in layer 3; these deep pyramidal neurons have longer apical dendrites than superficial pyramidal neurons. There are also a small population of “semilunar” neurons in layer 2 which are similar to pyramidal neurons, but without basal dendrites.

In addition to the pyramidal neurons, which bear a close resemblance to pyramidal neurons found in other areas of the neocortex and the hippocampus, there are a variety of interneuron types found in piriform cortex [16, 39] (figure 5.1, right). In layer 1 are found superficially located inhibitory horizontal and globular-soma neurons. These neurons receive direct input from olfactory bulb mitral cells and mediate feedforward inhibition across the length of the pyramidal neuron apical dendrite. There is also speculation that they may receive feedback excitation from pyramidal neurons [18]. In layer 3 are found multipolar neurons which can be spiny or aspiny. Spiny multipolar neurons are excitatory, but little data exists on their physiology. Aspiny multipolar neurons mediate feedback inhibition onto the cell bodies of pyramidal neurons [17, 36].

The physiology of pyramidal neurons in piriform cortex is similar to that of pyramidal neurons found in other parts of the brain [39]. These neurons display “regular spiking” behavior [34], featuring fairly broad spikes and significant spike frequency adaptation to a sustained current stimulus. Although little physiology has been performed on inhibitory interneurons, they clearly fall into the “fast-spiking” category [15, 34, 38], with narrow




Figure 5.1: The main classes of neurons in the piriform cortex, positioned in the layers where they are most commonly found. Abbreviations: *SP*, superficial pyramidal neurons; *DP*, deep pyramidal neurons; *S*, semilunar cells; *H*, horizontal neurons; *G*, globular-soma neurons; *MS*, spiny multipolar neurons; *M*, smooth multipolar neurons. *I*, *II*, *III*: layers of piriform cortex. *EN*: endopiriform nucleus. Neuron classes *H*, *G*, and *M* are inhibitory; the rest are excitatory. From [17].

spikes, high excitability, and little or no spike frequency adaptation. A number of synaptic and voltage-dependent ionic channels are known to exist in piriform cortex pyramidal neurons including fast sodium channels, fast and slow potassium channels, calcium channels, AMPA and NMDA excitatory synapses, fast and slow GABA-A inhibitory synapses, and slow GABA-B inhibitory synapses [2, 9, 10, 11, 12, 17, 25, 45],

For the purposes of the network modeling described in the following chapters, I have constructed models of superficial pyramidal neurons, layer 1 feedforward inhibitory interneurons (horizontal and globular-soma neurons), and layer 3 feedback inhibitory interneurons (aspiny multipolar neurons). The present network model does not contain deep pyramidal neurons, semilunar neurons, or spiny multipolar neurons, because these neurons are much less numerous than the neurons included in the model, and because very little physiological data is available on them. Nevertheless, incorporating models of these neuron types into the framework of the network model will be an important area for future development of the model.

Another feature of piriform cortex neurons is that a number of cellular and synaptic properties can be modulated by norepinephrine (NE) (see chapter 2 and [8, 14, 21, 33]). NE transiently modulates the strength of synaptic and voltage-dependent ionic channels and alters the resting membrane potential in both pyramidal cells and interneurons. These effects of NE have been incorporated into the neuron models.

5.2 Computer simulation environment

Simulations and data analysis presented in this chapter and chapters 6 and 7 were performed on a variety of Unix workstations running Linux or Solaris. All neuronal simulations were performed using the neural simulation program GENESIS [5]. Many extension libraries were constructed as C extensions to GENESIS in order to model the variety of synaptic, neuron and network features that are specific to the model. In addition, a parameter searching library was built and used to match the parameters of the single neuron models to experimental data (chapter 3; see also [47] and below). Data analysis was performed using MATLAB (the Mathworks Inc.) and by a variety of custom-written programs in C and python (<http://www.python.org>).

5.3 Modeling layer 2 pyramidal neurons

5.3.1 Data set

Layer 2 (superficial) pyramidal neurons in piriform cortex are the primary excitatory neuron type in the piriform cortex [16, 39]. They have characteristic spiking patterns, corresponding closely to the “regular-spiking” pattern described for neocortical neurons [34]. They are characterized by moderate spiking rates with spike frequency adaptation after a current step. For the purpose of tuning this model I used a data set consisting of membrane potential traces from experiments in a brain-slice preparation [35]. A series of constant-

current pulses at a variety of current levels were used. Depolarizing current injections were used to characterize the spiking behaviors of the neurons, while weak hyperpolarizing inputs were used to compute the input resistance and primary input time constant (τ_{in}) [27] of the neurons. Synaptic blockers were added to the bathing medium (CNQX ($30 \mu M$) to block AMPA glutamate receptors, APV ($100 \mu M$) to block NMDA glutamate receptors, and picrotoxin ($50 \mu M$) to block fast GABA-A receptors) in order to prevent random synaptic inputs from altering the response to the current inputs. A typical data set is shown in figure 5.2 (bottom traces of each graph). Current-clamp experiments were also performed in the presence of $10 \mu M$ norepinephrine (NE) in order to characterize the effects of neuromodulation on the behavior of pyramidal neurons. A typical data set from the NE-treated neurons is shown in figure 5.2 (top traces of each graph). Note the large increase in neuronal excitability after NE treatment.

5.3.2 Simplifying the neuron morphology

At the time of this writing, it is unfeasible to build computer simulations of very large realistic neuronal networks (comprising *e.g.*, hundreds to thousands of neurons) where each neuron has extremely detailed morphologies (hundreds to thousands of compartments) on available computers. Thus, an essential part of building a realistic network simulation is to simplify the morphologies of the constituent neurons, while keeping as many of the physiological properties of the model intact as possible.




Figure 5.2: Response of pyramidal neuron membrane potentials to currents injected at the soma. Listed above each graph is the level of current injection. The top trace corresponds to neurons treated with $10 \mu M$ NE, while the bottom trace (offset by $-150 mV$ for clarity) represents neurons not treated with NE. Current injection starts at $100 msec$ after the start of the simulation and lasts for $1 sec$.

The morphology of the model neuron was simplified from a 2670-compartment model developed by Alex Protopapas, which in turn was an anatomical reconstruction of a piriform cortex layer 2 pyramidal neuron kindly supplied by Mark Domroese of the University of Wisconsin. The reconstruction was performed using the NeuroLucida software (<http://www.microbrightfield.com/prod-nl.htm>) (figure 5.3).

In simplifying the detailed pyramidal neuron model, we wanted to preserve as many of the physiological properties of the original full-scale single-neuron model as possible, particularly the response at the soma to synaptic inputs on distal dendrites. To accomplish this task, we chose to reduce the full model using a method that conserves axial resistance but treats unit membrane resistance and capacitance as free parameters [7]. This method reduces two cylinders located at the end of a dendritic tree to a single equivalent cylinder whose radius is the geometric mean of the radii of the two original cylinders and whose length is the average of the lengths of the two cylinders:

$$R = \sqrt{\sum_i r_i^2} \tag{5.1}$$

$$L = \frac{1}{n} \sum_i^n l_i \quad , \tag{5.2}$$

where R is the radius of the reduced equivalent cylinder, r_i is the radius of dendrite i , L is the length of the reduced equivalent cylinder, l_i is the length of dendrite i , and n is the number of cylinders being combined into one. This method can be applied iteratively to




Figure 5.3: Digitized reconstruction of the layer 2 pyramidal neuron used to generate the simplified morphology.

reduce the original model as much as desired. When any two cylinders had widely differing lengths, we removed the smaller dendrite to avoid distorting the geometry of the neuron. Because this method does not conserve total cell area, the specific membrane resistance and capacitance of the model were rescaled until the input resistance R_{in} and input time constant τ_0 are roughly the same in both models. To ensure accurate simulation results, it is also necessary that each compartment have an electrotonic length of less than 0.1λ (see chapter I and [27]). Note that the standard method for reducing complex dendritic trees (the “3/2 power-law” or equivalent cylinder method of Rall [27, 37]) was not used since the dendritic diameters at branch points on the cell reconstruction did not conform to the 3/2 power law. The methods used also assume a passive dendritic tree. This is valid in the context of our model since all the active conductances are located in the soma. In addition, there is at present no strong evidence for active conductances on the dendrites of piriform cortex pyramidal neurons (Lewis Haberly, personal communication), in contrast to neocortical and hippocampal pyramidal cells [24, 42, 43].

The neuron model simplification algorithm is described in detail in Appendix A. The resulting model is shown in figure 5.4. Table 5.1 compares the passive parameters, input resistances, and input time constants of the original and simplified neuron models. The specific resistivity and capacitance parameters were adjusted by hand to produce a good match between the full and reduced models. Figure 5.5 shows the response of both models to a weak synaptic input on the distal apical dendrite. The response of the simplified model

is very similar to that of the full model based on either set of criteria. Note that the specific values of the membrane and axial resistivities and specific membrane capacitances are quite different; as described above, this is a feature of the simplification algorithm.

	Full model	Reduced model
Number of compartments	2670	15
R_{in} ($M\Omega$)	71.1	73.9
τ_{in} ($msec$)	19.6	18.9
R_m ($k\Omega \cdot cm^2$)	30.0	5.03
R_a ($\Omega \cdot cm$)	350.0	58.3
C_m ($\mu F/cm^2$)	0.8	4.36

Table 5.1: Passive parameter values for the full and simplified pyramidal neuron model.




Figure 5.4: (Previous page) Schematic structure of the simplified 15-compartment pyramidal neuron model. Layers of the piriform cortex in which the neuron compartments are located are listed on the left side of the neuron. Each segment represents a single compartment. Voltage- and calcium-dependent ionic channels are listed on the right. Note that ionic channels are found only in the soma. Abbreviations: *sup 1b*, superficial layer 1b; *deep 1b*, deep layer 1b, *soma(II)*, somatic region in layer 2; *III*, basal dendritic region in layer 3; *Na*, fast inactivating sodium channel; *Na pers*, persistent sodium channel; *Kdr*, delayed-rectifier potassium channel, *Ka*, fast inactivating potassium channel; *KM*, slow, non-inactivating voltage-dependent potassium channel (M-current); *Kahp*, slow calcium-dependent potassium channel; *Ca*, high voltage activated calcium channel. The soma of the neuron model also contains a calcium buffer.




Figure 5.5: EPSPs recorded at the soma to a single synaptic input on the distal apical dendrite of a pyramidal neuron in layer 1a. Dashed line: full 2670-compartment model. Continuous line: simplified model. The simplified model exhibits EPSPs that are about 25% larger.

5.3.3 Building the model

In order to allow the pyramidal neuron model to respond in a realistic manner to current-clamp inputs, voltage- and calcium-dependent ionic channels were added to the soma (table 5.4). Kinetic parameters for these channels were derived either from the piriform cortex literature [2, 9, 11, 10, 12], or, where there was no such data, from models of hippocampal pyramidal neuron channels (*e.g.*, [44]). There is direct evidence for all of these channels in piriform cortex pyramidal neurons, with the exception of the persistent fast sodium channel, which has been characterized in pyramidal neurons in the hippocampus [13]. In addition, various types of excitatory and inhibitory synaptic channels were added to the model (see below and in Appendix B). Parameters for synaptic channels were also derived from experimental data in piriform cortex brain slices (chapter 2 and [25]). Specific parameter values and activation curves for all channels are listed in Appendix B.

5.3.4 Tuning the model

Once a data set for a canonical pyramidal neuron was chosen (figure 5.2), we used the automated parameter search methods described in chapter 3 to match the behavior of the model to the data. The model had 23 adjustable parameters of the following types:

1. Maximal conductance values of voltage- or calcium-dependent ionic channels.
2. Scaling factors for the activation curves of voltage- or calcium-dependent ionic chan-

nels.

3. Soma resting potential.
4. Somatic and dendritic membrane resistivities.

Since the experimental data was obtained in the presence of synaptic blockers, the parameters of the synaptic channels were not varied when tuning the neuron models. Maximal conductances of the channels were varied because these values are not available in the literature. Kinetic parameters were varied for some channels because the data sources were either not derived from experiments on piriform cortex pyramidal neurons or because the available data was not of high quality. The shape of the kinetic activation curves and time constant curves was never altered; instead, the midpoints of some of the activation curves and/or the overall time scale of the time constant curves could vary over a narrow range. The resting potential was also allowed to vary over a narrow range to compensate for unknown junction potentials, which can alter the apparent membrane potential by several millivolts [49]. Since the neuron morphology was obtained from a different cell than the physiological data, we also included the somatic and dendritic membrane resistivities as parameters, in case the morphology represented a somewhat larger or smaller cell than the physiological data. Full details on the resulting channel parameters can be found in Appendix B. This approach is reasonable given that the objective of the parameter search was to generate a phenomenologically correct model for the purposes of network modeling, rather than to

find a unique correct set of parameters.

The result of the parameter search is shown in figure 5.6. As can be seen, the resulting model replicates the intracellular current clamp data to high accuracy. The input resistance of the pyramidal neuron model was $98 M\Omega$ and the input time constant (τ_{u_0} [27]) was $20.2 msec$, which is well within the typical range for these neurons [35].

5.3.5 Neuromodulation

One of the objectives of the work described in this thesis was to examine the effect of neuromodulators on the dynamical behavior of the piriform cortex. The implementation of synaptic neuromodulation is described below. Here, we describe the modifications of the pyramidal neuron model needed to incorporate neuromodulation by norepinephrine (NE). Our experiments have shown that NE increases the resting potential and the excitability of pyramidal neurons, as well as reducing spike frequency adaptation (figure 5.2). Experiments in other systems suggest that these effects are mediated in part by suppression of slow potassium conductances [29, 30]. In order to incorporate these effects into our pyramidal neuron model, we took the model discovered by the parameter search methods and allowed several parameters to vary (see table 5.6). These parameters represented the resting membrane potential of the neuron and the maximal conductances of four voltage- and calcium-dependent potassium channels (fast delayed-rectifier channel, inactivating A-current, slow non-inactivating M-current, and calcium-dependent potassium current). All




Figure 5.6: Soma membrane potential traces in pyramidal neurons in response to current steps. The current level is listed at the top of each graph. The lower trace in each graph is experimental data derived from brain-slice experiments, while the upper trace is simulated data from the pyramidal neuron model. The lower traces have been offset in the y direction for comparison purposes, and offset in the x direction to line up the time of current onset. Current injection starts at 100 $msec$ after the start of the simulation and lasts for 1 sec .

other parameters were kept at the same values as those in the unmodulated neuron.

The response of the neuromodulated neuron to current-clamp inputs is shown in figure 5.7. The parameter changes required to turn the unmodulated neuron into a neuromodulated neuron are listed in Appendix B. As can be seen, the neuromodulated neuron matches the spiking behavior of the real neuron quite closely, although not as closely as for the unmodulated neuron. This is to be expected, since fewer parameters were allowed to vary in the neuromodulated case.

5.4 Modeling inhibitory neurons

5.4.1 Data set

Various types of inhibitory neurons are found in piriform cortex (chapters I and [16, 39]). Unfortunately, very little intracellular data is available for these cell types, with the exception of layer 3 inhibitory interneurons (putative aspiny multipolar neurons [15, 36, 38]). The spiking behavior of these inhibitory interneurons is very simple, corresponding to the “fast-spiking” class of neurons described for neocortical neurons [34]. These neurons are highly excitable, have a low spiking threshold, and show little to no spike frequency adaptation. In addition, some experiments on the effects of norepinephrine have been performed on this neuron type [14, 33]. Typical spiking behaviors of this neuron type are shown in figure 5.8, and the spiking frequencies as a function of current (f/I curve) is shown in figure

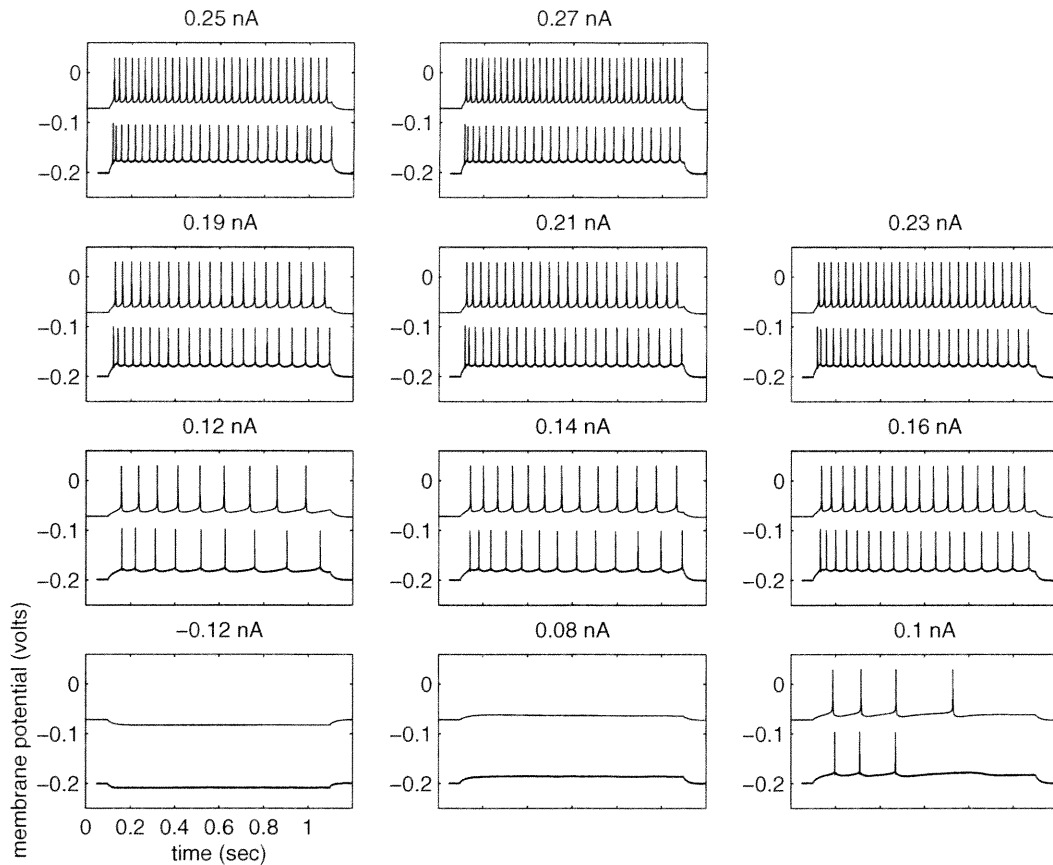


Figure 5.7: Soma membrane potential traces in pyramidal neurons in response to current steps in the presence of $10 \mu\text{M}$ NE. The current level is listed at the top of each graph. The lower trace in each graph is experimental data derived from brain-slice experiments, while the upper trace is simulated data from the pyramidal neuron model. The lower traces have been offset in the y direction for comparison purposes, and offset in the x direction to line up the time of current onset. Current injection starts at 100 msec after the start of the simulation and lasts for 1 sec .

5.9 for several neurons. Note that the excitabilities of these neurons can vary considerably from neuron to neuron.

5.4.2 Building and tuning the models

Since detailed reconstructions of the morphologies of inhibitory interneurons were not available, and since the spiking characteristics of these neurons are simple and easily modeled, we chose to use one-compartment models for these neurons. Voltage-dependent ionic channels included fast Na , delayed-rectifier K , and fast inactivating K (A-current). Parameters for these channels and for the neuron as a whole are listed in Appendix B. The A-current was added to decrease the excitability of the feedback inhibitory neurons to more physiological levels. Synaptic channels are described below. Because of the simplicity of these models, automated parameter search methods were not needed to assign parameters to the models in the cases where data was available.

No intracellular data is available for putative feedforward inhibitory interneurons in piriform cortex. Therefore, these neurons were modeled using only fast Na and delayed-rectifier K channels. Parameter sets for these neurons are listed in Appendix B.

5.4.3 Neuromodulation

NE is known to increase the excitability of layer 3 inhibitory neurons in piriform cortex, in part by increasing the resting membrane potential [14]. Although precise f/I curves for this




Figure 5.8: Intracellular recordings of soma membrane potential traces in layer 3 inhibitory interneurons in response to current steps. The current level is listed at the top of each graph. Current injection starts at 50 msec after the start of the simulation and lasts for 1 sec.




Figure 5.9: Frequency vs. injected current (f/I) curves for several layer 3 inhibitory interneurons. Each line represents a different interneuron.

neuron type under the influence of NE have not been reported in the literature, a similar strategy was used to increase the excitability of these neurons in the model. f/I curves from the model with and without NE are shown in figure 5.10. The excitability of the model interneuron is within the typical range of the data. NE causes a significant increase of the excitability of the model interneuron.

There is no data available on neuromodulation of synaptic inputs to inhibitory interneurons. Therefore, in our model these synaptic inputs are unaffected by NE.

5.5 Modeling synaptic inputs

An essential component of a realistic neuronal network model are accurate models of the various types of synaptic conductances present in neurons of the model. These are described in this section.

5.5.1 Data sets

Data on fast excitatory synaptic transmission in afferent and associational fiber pathways was derived from brain-slice experiments on piriform cortex described in detail in chapter 2. This data set also included extensive data on the effect of neuromodulators on synaptic transmission in piriform cortex, as well as information on synaptic facilitation. Additional data on synaptic facilitation and depression was derived from unpublished data from brain-slice experiments described below. Other synaptic parameters were taken from [25].




Figure 5.10: Frequency vs. injected current (f/I) curves for the feedback inhibitory interneuron model with and without $100 \mu M$ NE. The solid black curve represents the model without NE while the dashed black curve represents the model with NE. Experimental f/I curves (without NE) are shown in red for comparison.

5.5.2 The basic synapse model

The basic synaptic model used in the simulations is a generalized alpha function:

$$G_{syn}(t) = \bar{G} \cdot \frac{A}{\tau_1 - \tau_2} (e^{-t/\tau_1} - e^{-t/\tau_2}), \tau_1 > \tau_2 \quad , \quad (5.3)$$

where \bar{G} is the maximum conductance of the synaptic channel, τ_1 is the decay time constant, τ_2 is the rise time constant, and A is a normalization factor which ensures that the maximum conductance is \bar{G} . Total synaptic current for a given channel is computed using the equation:

$$I_{syn}(t) = G_{syn}(t) \cdot (V - E_{rev}) \quad , \quad (5.4)$$

where V is the membrane potential of the neuronal compartment the synaptic channel is associated with and E_{rev} is the reversal potential of the channel. Table 5.2 shows the different synaptic types used in the model along with the different parameter values used. AMPA refers to fast excitatory glutamate channels [39] which are selectively excited by AMPA (-amino-3-hydroxyl-5-methyl-4-isoxazolepropionic acid). NMDA (N-methyl-D-aspartate) channels are slow excitatory glutamate channels which are voltage-dependent as well as synaptically activated. Both types of channels are permeable to sodium and potassium ions (and calcium, in the case of NMDA channels), and have high reversal potentials. GABA-A channels are inhibitory channels permeable to chloride and have reversal potentials in the range of neuronal reversal potentials; thus, their effect on neuron behavior is predominantly

through a current-shunting mechanism [22, 25]. GABA-B channels are inhibitory channels permeable to potassium and have a hyperpolarized reversal potential.

synaptic conductance	$\tau_1(msec)$	$\tau_2(msec)$	$E_{rev}(mV)$
Afferent AMPA	1.0	1.6	0.0
Associational AMPA	0.8	2.8	0.0
NMDA	60.0	66.0	0.0
fast GABA-A	1.24	7.25	-70.0
slow GABA-A	0.74	37.0	-70.0
GABA-B	70.0	70.0	-90.0

Table 5.2: Parameter values for the synaptic channel models.

5.5.3 NMDA synapses

NMDA (N-methyl-D-aspartate-sensitive) synapses are excitatory synapses that are activated by the neurotransmitter glutamate. They are considerably slower than AMPA channels, and are also voltage-dependent [26, 39], being inactive until the postsynaptic depolarization is high enough to relieve a magnesium-mediated channel block. NMDA synapses also mediate some forms of long-term synaptic potentiation (LTP) and depression (LTD) [4, 46] and may underlie some learning processes, although this is controversial [40, 41]. Since the focus of the models described here is not to investigate learning phenomena, the NMDA synapses we modeled do not exhibit long-term synaptic potentiation (although this feature can easily be added).

We modeled NMDA synapses using the equations and parameters of [25], which were

derived from experiments in piriform cortex. The equations were as follows:

$$G_{syn}(V, t) = \bar{G} \cdot \frac{1}{N(V, t)} \cdot \frac{A}{\tau_1 - \tau_2} (e^{-t/\tau_1} - e^{-t/\tau_2}), \tau_1 > \tau_2 \quad (5.5)$$

$$N(V, t) = 1 + \eta \cdot [Mg^+] \cdot e^{-\gamma \cdot V} \quad (5.6)$$

$$I_{syn}(V, t) = G_{syn}(V, t) \cdot (V - E_{rev}) \quad , \quad (5.7)$$

where $N(V, t)$ represents the voltage dependence of the synaptic channel and the other terms are as described above. $[Mg^+]$ represents the fixed extracellular magnesium concentration (1 mM) while η and γ determine the strength of the voltage dependence. Specific parameter values are listed in Appendix B.

5.5.4 Synaptic facilitation and depression

Both afferent (layer 1a) and associational (layer 1b) excitatory synapses in piriform cortex are known to exhibit short-term synaptic facilitation (chapter 2 and [6]) and depression (unpublished data), whereby the strength of an excitatory synaptic potential (EPSP) is transiently modified based on the history of previous inputs to that synapse. The dynamics of these processes differ substantially between associational and afferent synapses. In order to characterize this phenomenon, we performed brain-slice experiments using a variety of interpulse intervals on both types of synapses in the presence and absence of 25 μM NE. Details of the experimental procedures are described in chapter 2. Figure 5.11 shows the effects

of a train of equally spaced pulses on EPSP size in afferent and associational synapses. Figures 5.12 and 5.13 show the extent of facilitation and depression as a function of interpulse interval duration in a typical experiment. In afferent synapses, considerable facilitation is seen in the short term which levels off as time goes on; little synaptic depression is seen. In contrast, synaptic facilitation is much less prominent in associational synapses than in afferent synapses, while synaptic depression is much more prominent. In addition, synaptic facilitation and depression in afferent synapses are essentially unaffected by NE (chapter 2), whereas associational synapses show a marked change in facilitation/depression behavior in response to NE; facilitation is greatly enhanced while depression is greatly reduced.

Synaptic facilitation and depression are believed to be mediated by a presynaptic mechanism involving the accumulation of calcium as a result of the depolarization of synaptic terminals [20, 28]. However, since the network simulation has many thousands of synapses, for the sake of computational efficiency we decided to model this process more abstractly as follows:

$$d(fac)/dt = \frac{-fac}{\tau_{fac}} + fps \cdot \delta_{spike}(t) \quad (5.8)$$

$$d(depr_1)/dt = \frac{-depr_1}{\tau_{depr_1}} + dps_1 \cdot \delta_{spike}(t) \quad (5.9)$$




Figure 5.11: Extracellular recordings in a brain-slice preparation of afferent (layer 1a) and associational (layer 1b) synapses in piriform cortex during stimulation with a train of voltage shocks at 50 msec intervals. The upper trace represents layer 1a synapses, the middle trace represents layer 1b synapses, and the lower trace represents layer 1b synapses with 25 μM NE added to the slice chamber. Layer 1a responses were not affected by NE.

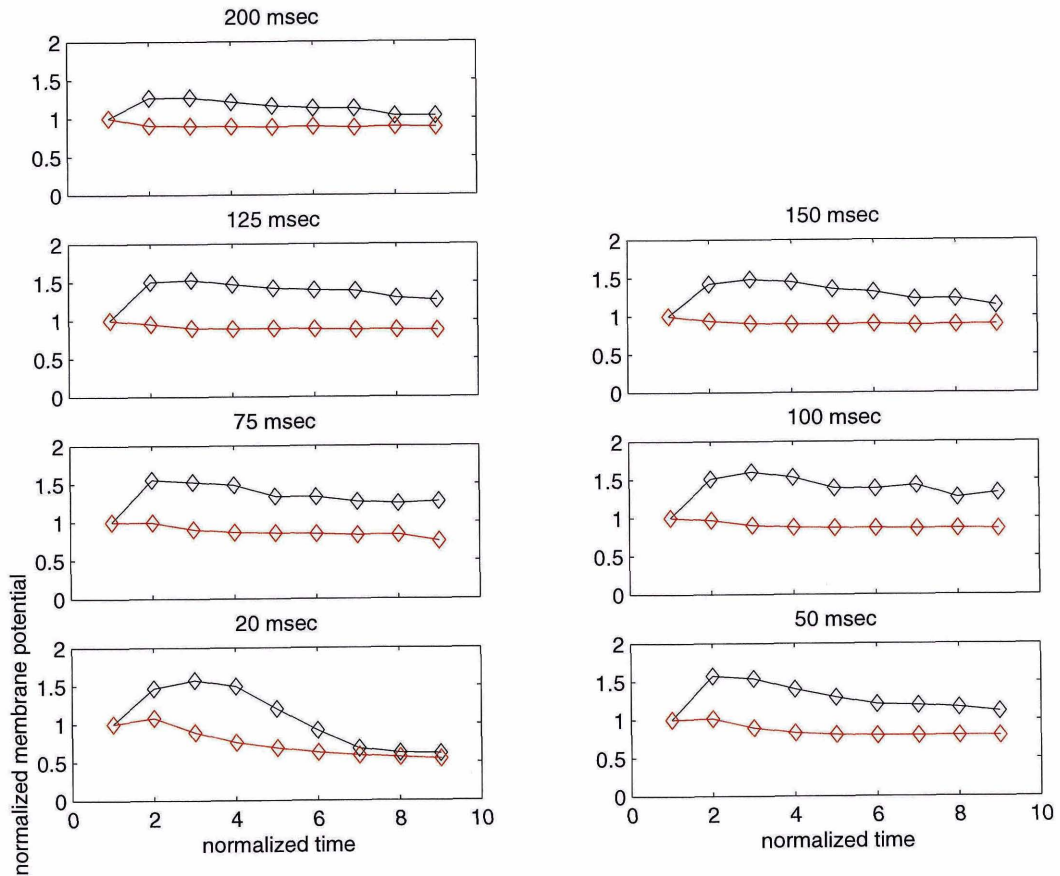


Figure 5.12: Synaptic peak heights for successive pulses in a multiple-shock paradigm. The black traces represent layer 1a synapses, while the red traces represent layer 1b synapses. The y axis represents normalized peak heights, with the first peak being 1.0 by definition. The x axis represents pulse number. The interpulse interval is shown above each graph. The time axis is normalized to multiples of the interpulse interval.

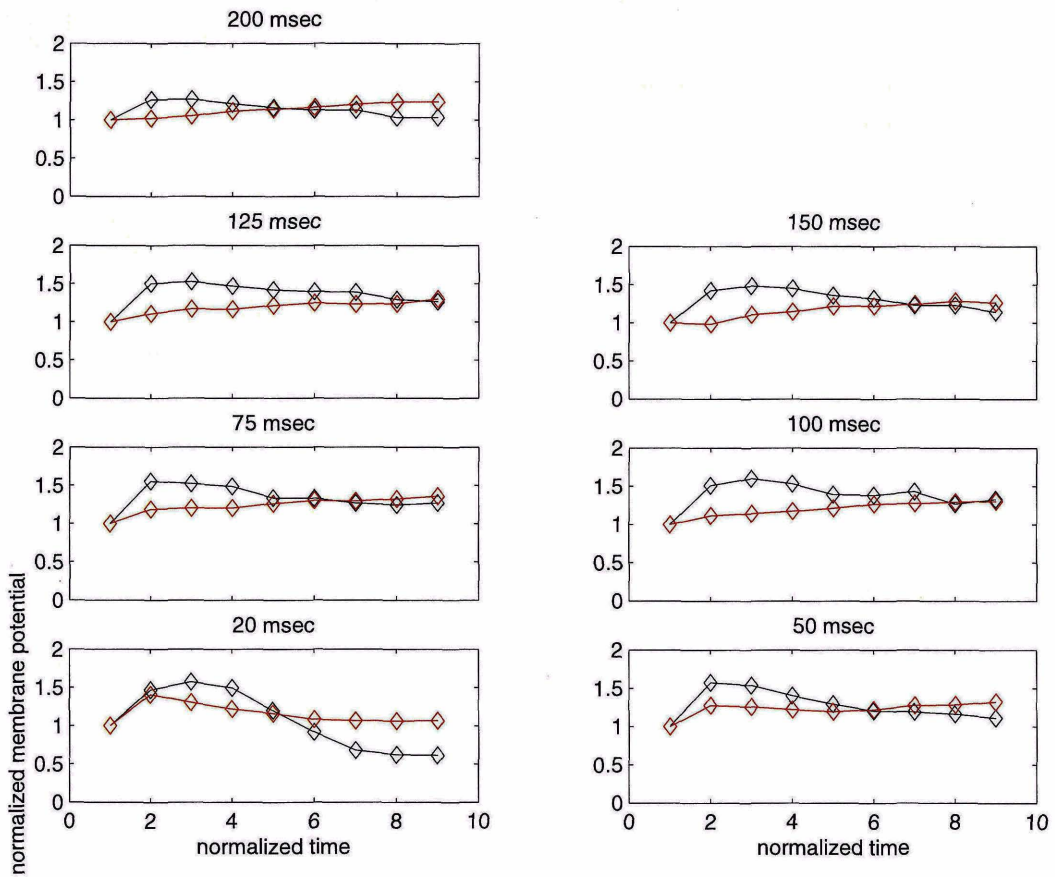


Figure 5.13: Same as figure 5.12, except that the red traces represent layer 1b responses in the presence of $25 \mu M$ NE.

$$d(depr_2)/dt = \frac{-depr_2}{\tau_{depr_2}} + dps_2 \cdot \delta_{spike}(t) \quad , \quad (5.10)$$

where fac represents synaptic facilitation, fps represents the increase in facilitation per spike, $depr_1$ and $depr_2$ represent two independent synaptic depression processes, and dps_1 and dps_2 represent the increase in synaptic depression per spike. τ values represent the time constants for each process. The combined effect of these three processes is calculated as follows:

$$wt_{dynamic} = wt_{static} \cdot \frac{(1 + fac)}{(1 + depr_1) \cdot (1 + depr_2)} \quad , \quad (5.11)$$

where wt_{static} represents the synaptic weight without facilitation or depression and $wt_{dynamic}$ is the synaptic weight after adjusting for facilitation and depression. In other words, each facilitation or depression process is modeled as an independent leaky integrator, where a facilitation/depression value of 0 indicates no facilitation or depression respectively. If all facilitation/depression state variables are 0, the weight will remain at its static value, which is equivalent to the value it would have for a spike which followed a long period of inactivity. We found that two depression processes with different time constants were necessary to adequately model synaptic depression in the real synapses. We used a parameter search to match the facilitation/depression parameters to the data; the resulting parameters are listed in Appendix B.

Figures 5.14 and 5.15 compare the resulting model with the facilitation/depression behavior of the real synapses; as can be seen, the fit is very good but not perfect. In particular,

short-term facilitation in associational (layer 1b) synapses is not captured by the model. We decided not to use a more complex model to achieve a better fit, because this would have slowed down the simulation and because synaptic facilitation is weak in layer 1b synapses (chapter 2 and [6]), and is often absent altogether (unpublished data).

We reran the parameter search for the case of associational synapses treated with norepinephrine; the results are shown in figure 5.16. The model accurately replicates the effect of NE on these synapses: synaptic depression is reduced while facilitation is augmented.

5.5.5 Synaptic neuromodulation

We have shown that norepinephrine transiently depresses synaptic transmission in associational synapses in piriform cortex, while transiently enhancing synaptic transmission in afferent synapses (chapter 2). Neuromodulation was incorporated into the synaptic activation equations described above as a multiplicative scale factor on the synaptic conductances:

$$G_{syn,mod}(t) = G_{syn,unmod}(t) \cdot N(t) \quad , \quad (5.12)$$

where $G_{syn,unmod}(t)$ and $G_{syn,mod}(t)$ are respectively the unmodulated and modulated synaptic conductances, and $N(t)$ is the time-dependent neuromodulation of that synaptic channel. $N(t)$ was channel-specific and was a function of the neuromodulator concentration.

Our experiments also showed that synaptic facilitation and depression depends on the neuromodulator level. To model this in GENESIS, we developed synaptic channel objects

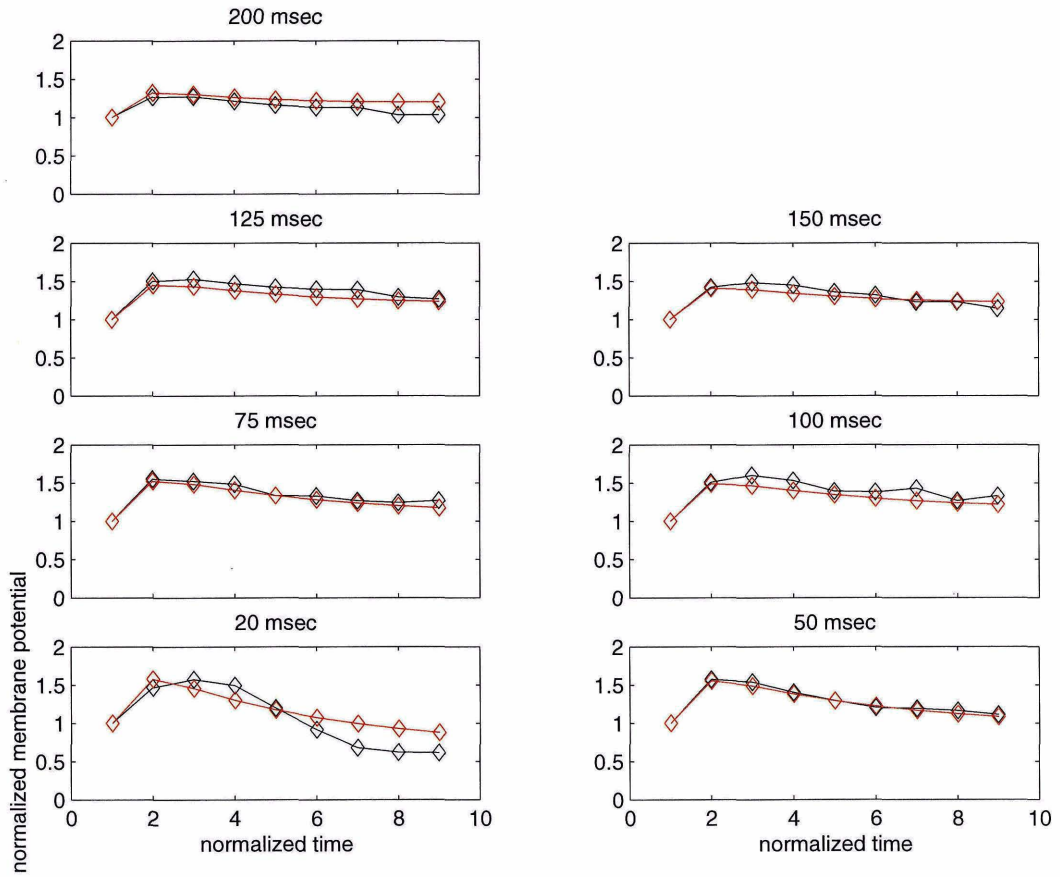


Figure 5.14: A comparison of multipulse facilitation and depression in layer 1a synapses in the data and the model. The black traces represent the experimental data while the red traces represent the model outputs. Interpulse intervals are listed above each graph.

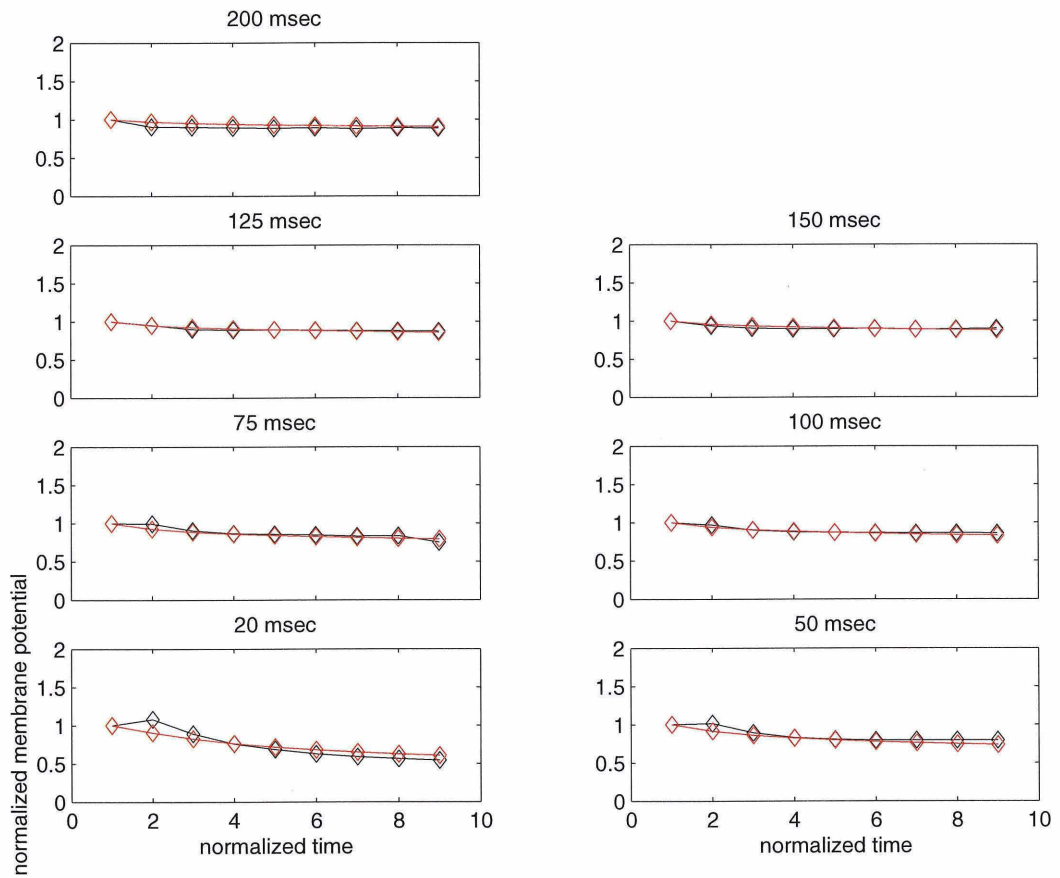


Figure 5.15: A comparison of multipulse facilitation and depression in layer 1b synapses in the data and the model. The black traces represent the experimental data while the red traces represent the model outputs. Interpulse intervals are listed above each graph.

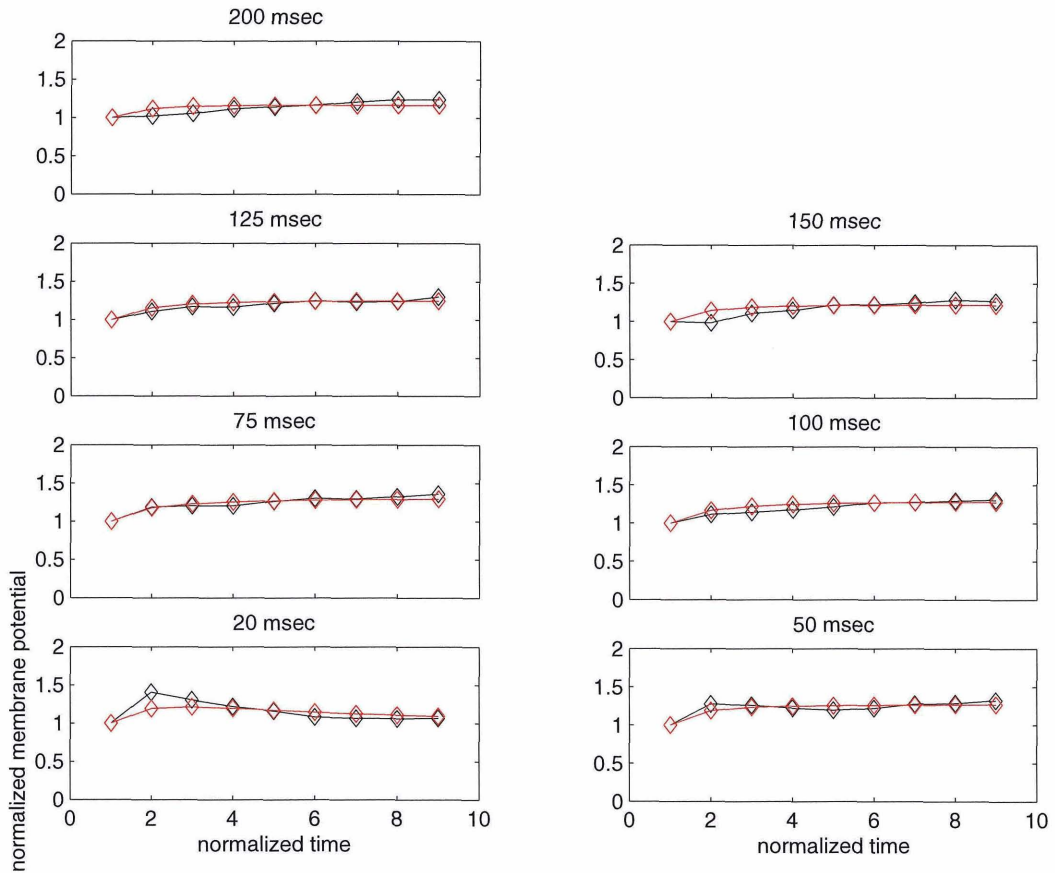


Figure 5.16: A comparison of multipulse facilitation and depression in layer 1b synapses in the data and the model. In this case, both model and experimental data were recorded in the presence of $25 \mu M$ NE (simulated in the case of the model). The black traces represent the experimental data while the red traces represent the model outputs. Interpulse intervals are listed above each graph.

which contained two sets of facilitation/depression parameters (one for unmodulated and one for neuromodulated synapses). When the neuromodulator level was above a threshold, the model would use the neuromodulated parameter set; otherwise it would use the unmodulated parameter set. This approach was necessary because neuromodulation of facilitation/depression was only systematically investigated for one concentration of NE ($25 \mu M$).

Parameters for synaptic neuromodulation are listed in Appendix B.

5.6 Limitations of our approach

Our approach to modeling single neurons is a phenomenological one, guided by the needs of the network modeling to be described in the subsequent chapters of this thesis. Many compromises had to be made in order to produce neuron models which could be simulated in large numbers on currently available workstations, and this may result in unrealistic model behavior for a number of reasons.

We used the same neuron model for every neuron of a given type in the network model. This means that we are assuming that different neurons of the same type do not differ significantly in their properties. While the evidence presently available is consistent with this view, some variation is inevitable, and may have significant effects on network behaviors. In particular, real neurons display significant anatomical variations which are not present in our models [16], and some physiological variation is seen as well [3].

The process of simplifying the neuron morphologies may obscure or distort many important phenomena [26]. The specific organization of synaptic inputs on dendrites may underlie important computational phenomena that would not be observed in our model. Synapses on fine dendritic branches will also see a larger input resistance than synapses in our model. This effect is even more pronounced for excitatory synapses, most of which are located on dendritic spines [19]. The absence of spines, which serve to partially isolate synapses from the main dendritic trunk, results in larger dendritic conductance changes in response to excitatory input than would be seen in real neurons. The localization of voltage-dependent ionic channels to the soma is also not entirely realistic; real neurons have a spike-initiation zone which has a much higher concentrations of sodium channels, with smaller concentrations of sodium channels in the soma [31]. Having a large density of sodium channels in the soma results in large transmembrane currents which can distort the field potentials measured from the model (see chapters 6 and 7). In addition, pyramidal cells in most brain regions are known to have voltage-dependent channels on their dendrites [24, 42, 43]. We did not include such channels in our pyramidal cell model because there is no evidence for their existence in piriform cortex; nevertheless, this will be an important subject for future experimental and modeling work.

The neuron models described in this chapter and used in the network model are entirely deterministic; any variability in their responses derives from variation in their inputs. In principle, this could result in artifactual behaviors if the inputs were extremely regular.

However, as will be seen in the chapters on network modeling, inputs to the network model always included a large random component. Therefore, we feel that this simplification is justified, especially since incorporating models of channel noise or other noise-generating phenomena (chapter 4 and [1, 32]) would have slowed down the simulations considerably.

All of these limitations point to areas where the present neuron models can be improved. Some changes, such as adding spike-initiation zones, are straightforward and can be done with a moderate amount of effort. Others, such as using neurons with detailed morphologies, would increase the computational complexity of the network model greatly. If the true dynamical behaviors of the network are critically dependent on these phenomena, progress in modeling brain regions at the network level will be very slow indeed. However, the hope is that at least some of the significant dynamical phenomena of the real system can be captured using reduced neuron models embedded in network models. As the following chapters demonstrate, many network-level phenomena in the piriform cortex can be reproduced faithfully using a network model built from the neuron models described in this chapter.

5.7 Appendix A: the neuron model simplification algorithm

The neuron model simplification algorithm proceeded as follows:

1. Unbranched chains of compartments were reduced to a single compartment having the same electrotonic and physical length. This was done by calculating the total electrotonic length of the branch and then adjusting the diameter of the new com-

partment to give the same total electrotonic length. The electrotonic length el for a single compartment is defined as follows:

$$el = l/\lambda \quad , \quad (5.13)$$

where l is the compartment length and λ is the compartment space constant:

$$\lambda = \sqrt{(R_m/R_a)(d/4)} \quad , \quad (5.14)$$

where R_m is the specific membrane resistivity, R_a is the specific axial resistivity, and d is the diameter of the compartment [23].

2. Branches were collapsed into single compartments, starting from the distal end and proceeding inward towards the soma, according to the algorithm of Bush and Sejnowski [7] described above. This was repeated until only five compartments remained: a basal dendritic compartment, the soma compartment, and three compartments representing the apical dendrite: the proximal (layer 1b) apical dendrite, the middle (distal layer 1b) apical dendrite, and the distal (layer 1a) apical dendrite. The apical dendritic compartments, which can be joined together, are kept separate because the connection topology of the network model (chapter 6) requires separate compartments.
3. Because the basal and apical dendritic compartments were considerably longer than the 0.1λ limit normally considered to be the maximum length compatible with nu-

merical accuracy [27], we subdivided each of the dendritic compartments into several compartments each, which resulted in fifteen compartments all less than 0.1λ in length.

5.8 Appendix B: model parameters

In this section are listed the specific parameters for the various neuron models. GENESIS simulation scripts for all neurons described here are also available online at

<http://www.bbb.caltech.edu/USERS/mvanier/thesis/thesis.html>.

5.8.1 Pyramidal neuron channel model parameters

Currents through ionic and synaptic channels are described by this equation:

$$I_{ch} = G_{ch} \cdot (V_m - E_{ch}) \quad , \quad (5.15)$$

where ch represents a channel species, V_m is the membrane potential, G_{ch} is the channel conductance and E_{ch} is the reversal potential for the channel. Reversal potentials for voltage-dependent ionic channels are listed in table 5.3. All voltage-dependent ionic channels were specific to a given ionic species and used the reversal potential for that species. In order to simplify the model, we elected not to change the Ca equilibrium potential based on the changing intracellular calcium concentration.

ion species	E_{rev} (mV)
Na	+55
K	-75
Ca	+113

Table 5.3: Reversal potentials for ionic channels.

All voltage- or calcium-dependent ionic channel conductances were computed using a generalization of the Hodgkin-Huxley model [5, 22]:

$$G_{ch} = \bar{G}_{ch} X^a Y^b Z^c \quad (5.16)$$

The parameter \bar{G}_{ch} is a variable parameter of each channel model. The gating variables X , Y , and Z represent (respectively) voltage-dependent activation (m or n in the Hodgkin-Huxley notation), voltage-dependent inactivation (h in the Hodgkin-Huxley notation), and calcium-dependent activation [5, 22]. The gating variables are computed by solving the following differential equation:

$$\tau_W(V_m) \frac{dW}{dt} = W_\infty(V_m) - W \quad (5.17)$$

for a given gating variable W , where τ_W and W_∞ are all functions of membrane potential V_m . $W_\infty(V_m)$ represents the asymptotic activation value of the gating variable when a fixed membrane potential V_m is present (*e.g.*, during a voltage clamp). $\tau_W(V_m)$ is the relaxation

time constant of the gating variable under these conditions [22].

Voltage-dependent channels

The activation and time constant kinetics curves for the voltage- and calcium-dependent channels are shown in figures 5.17, 5.18, and 5.19. Maximum conductances and gating variable exponents are listed in table 5.4. The data on which the channel models were based are described below. Note that kinetics for most of the channels were altered from the original source to fit the data by means of a parameter search (chapter 3).

channel	\bar{G}	a	b	c
Na , fast	2006.63	2	1	0
Na , persistent	64.86	1	0	0
K , delayed-rectifier	251.79	1	0	0
K , A-current	440.68	3	1	0
K , M-current	57.15	1	0	0
K , AHP	15.40	0	0	1
Ca	11.18	1	1	0

Table 5.4: Maximum conductances and gating variable exponents for the ionic channel models used in pyramidal neurons. Conductance units are S/m^3 . Gating variable exponents are dimensionless.

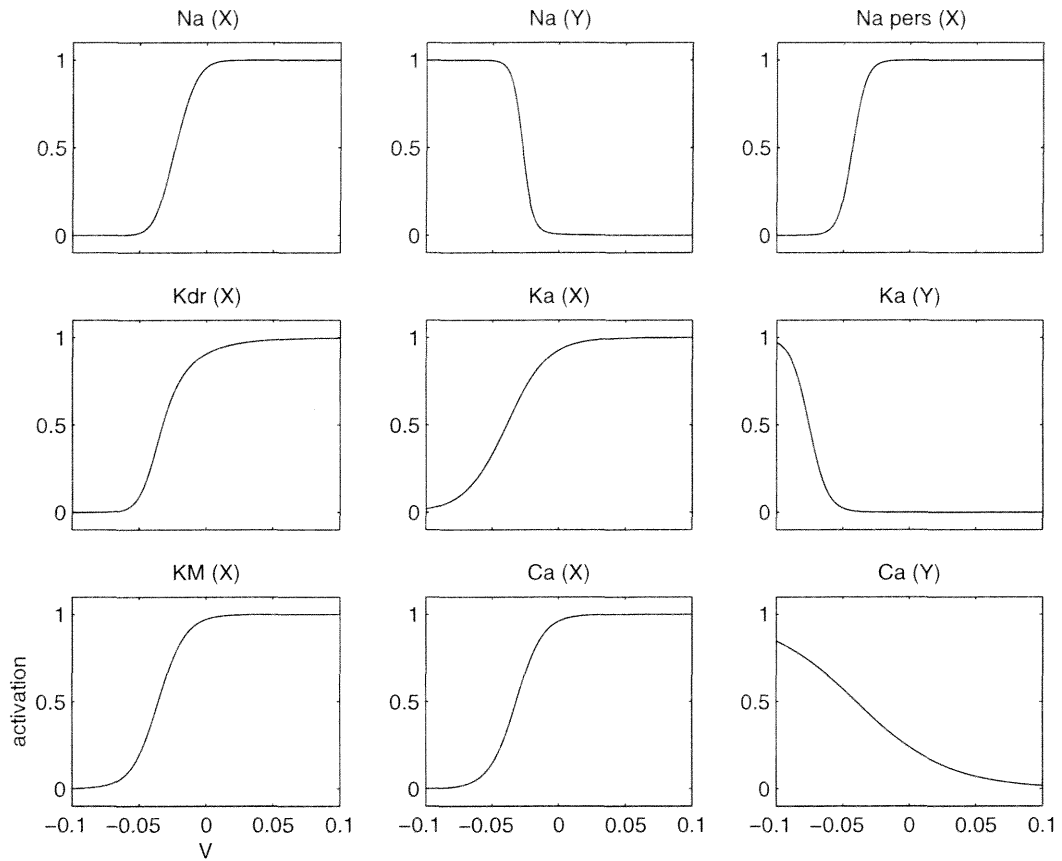


Figure 5.17: Activation (m_∞ vs. V) curves for all voltage-dependent ionic channels in the pyramidal neuron model. x -axis: membrane potential (in volts) of compartment associated with the channel. y -axis: channel activation (dimensionless). Abbreviations: X , activation gate; Y , inactivation gate; Na , fast sodium channel; $Na\ pers$, persistent sodium channel; Kdr , delayed-rectifier potassium channel; Ka , inactivating potassium channel; KM , slow non-inactivating potassium channel (M-current); Ca , high voltage activated calcium channel.

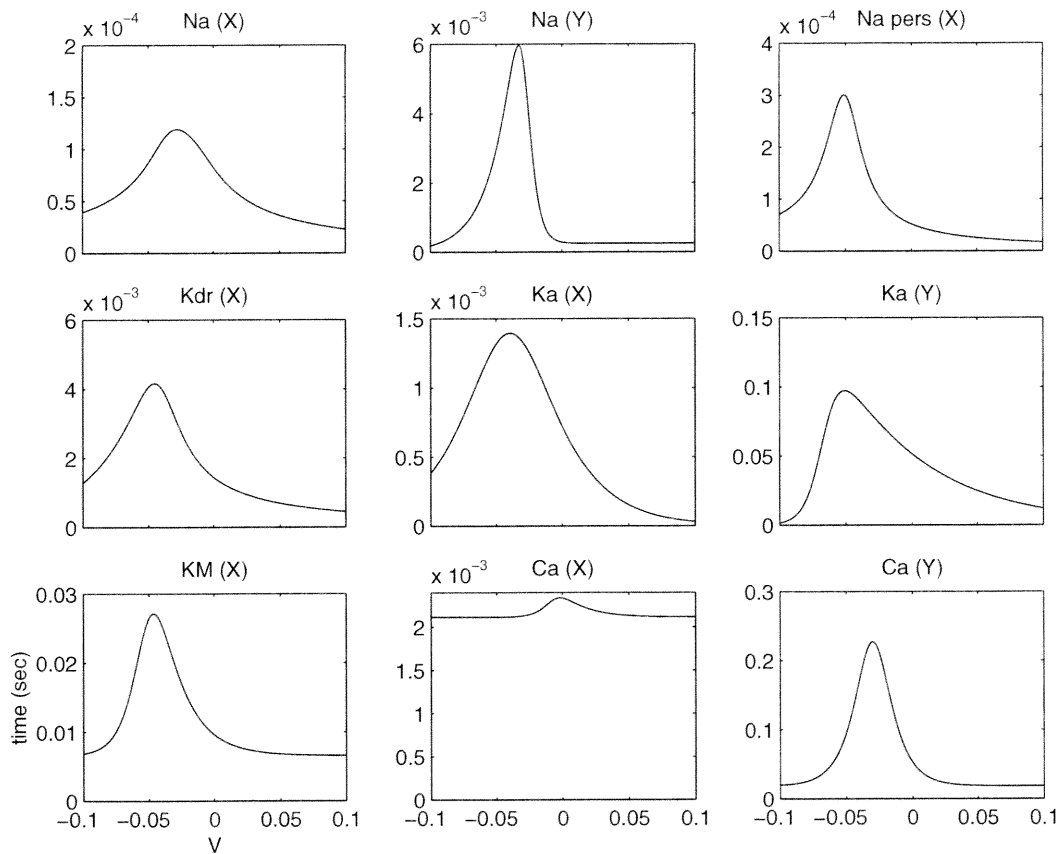


Figure 5.18: Time constant (τ vs. V) curves for all voltage-dependent ionic channels in the pyramidal neuron model. x -axis: membrane potential (in volts) of compartment associated with the channel. y -axis: time constant in seconds. Channel abbreviations as in figure 5.17.




Figure 5.19: Activation (m_{∞} vs. $[Ca]$) and time constant (τ vs. $[Ca]$) curves for calcium-dependent potassium channel in the pyramidal neuron model.

Na channels

No data is available for the fast (inactivating) *Na* channel in piriform cortex, so kinetics for this channel were taken from Roger Traub's model of hippocampal pyramidal neurons [44]. Similarly, the data for the persistent *Na* channel model was taken from the hippocampus [13].

K channels

No data is available for the delayed-rectifier *K* channel in piriform cortex, so we again used the model of Traub [44] based on data from the hippocampus. The model of the fast inactivating *K* channel ("A-current") was based on recordings in the piriform cortex by Banks and Haberly [2]. The model of the slow, non-inactivating *K* channel ("M-current") was based on recordings in the piriform cortex [9].

Ca channel

The *Ca* channel used in our model is a high-voltage-activated channel [22] based on voltage-clamp data recorded in piriform cortex pyramidal neurons [10].

Calcium-dependent channels

K_{ahp} channel

The slow, calcium-dependent K channel (“AHP” or “afterhyperpolarization” channel) in our model is based on data recorded in piriform cortex pyramidal neurons [11, 12]. This channel is not voltage-dependent.

Calcium buffer

A simple leaky-integrator model of calcium was used in our pyramidal neuron model:

$$d[Ca]/dt = B \cdot I_{Ca} - [Ca]/\tau_{Ca} \quad (5.18)$$

The B parameter converts current through calcium channels into a change in calcium concentration. The units are molarity per ampere-second. Parameters for the calcium buffer are listed in table 5.5.

parameter	value
τ (<i>msec</i>)	96.02
B ($M/(A \cdot sec)$)	2.31e+5

Table 5.5: Parameters of the calcium buffer.

Neuromodulation of voltage- and calcium-dependent channels

Neuromodulation is implemented as a sigmoidal function of concentration; for more details see chapter 2. In table 5.6 we give the effects of $100 \mu M$ NE on various cellular parameters. These parameters were discovered by a parameter search (see chapter 3). In this model NE increases the resting membrane potential of pyramidal neurons, which is consistent with experimental observations [14] and our experimental data. NE decreases the conductance of some potassium channels; this is also seen in hippocampal pyramidal neurons [29, 30]. Surprisingly, the conductance of the fast delayed-rectifier K current was increased by our parameter search; this was necessary to compensate for the loss of the slower K currents. There is currently no data available to assess whether this actually occurs *in vivo*.

parameter	value ($0 \mu M$)	value ($100 \mu M$)
soma membrane potential (mV)	-74.6	-72.1
\bar{G} of K , delayed-rectifier (S/m^3)	251.79	323.82
\bar{G} of K , A-current (S/m^3)	440.68	118.37
\bar{G} of K , M-current (S/m^3)	57.15	43.19
\bar{G} of K , AHP (S/m^3)	15.40	8.38

Table 5.6: Neuromodulation of cellular parameters.

Synaptic channels

NMDA channels

Table 5.7 gives parameter values for the NMDA synapses, taken from [25].

parameter	value
η (mM^{-1})	0.33
$[Mg]$ (mM)	1.0
$\gamma(V^{-1})$	80.0

Table 5.7: NMDA synapse parameters.

Facilitation and depression

Parameters for synaptic facilitation and depression, for afferent (layer 1a) and associational (layer 1b) synapses, are listed in table 5.8. Associational parameters are listed for both 0 and 100 μM NE. Afferent parameters are independent of the NE concentration.

Neuromodulation

Based on our experimental data (see chapter 2), 100 μM NE caused a 121% increase in the conductance of afferent (layer 1a) synapses, and a decrease in the conductance of associational (layer 1b) synapses to 23% of the original value.

parameter	afferent value	assoc value (0 μM NE)	assoc value (100 μM NE)
fps	3.64	0.13	1.25
$\tau_{fac}(msec)$	162	10	150
dps_1	0.79	0.027	0.66
$\tau_{depr_1}(msec)$	126	894	139
dps_2	0.64	0.12	0.14
$\tau_{depr_2}(msec)$	322	89	127

Table 5.8: Parameters for synaptic facilitation and depression. fps means “facilitation per spike”, and $dps_{[12]}$ mean “depression per spike” for the first or second synaptic depression process, respectively. fps and $dps_{[12]}$ are dimensionless.

5.8.2 Pyramidal neuron model parameters

The pyramidal neuron model is a linear 15-compartment model (figure 5.4); the dimensions of the compartments are listed in table 5.9. Synaptic channels are distributed across the neuron as indicated. *AMPA* and *NMDA* channels are excitatory glutamatergic synaptic channels while *GABA_A* and *GABA_B* represent fast and slow inhibitory GABAergic channels. Maximal conductance values for the synaptic channels vary according to each neuron’s position in the network and are described in chapter 6. Voltage-dependent ionic channels are found only in the soma. The somatic resting potential was $-74.7 mV$. More details on the synaptic and voltage-dependent ionic channels are given above.

5.8.3 Feedback inhibitory interneuron model parameters

Feedback inhibitory neurons were modeled as single-compartment neurons. Synaptic channels included fast Na , fast delayed-rectifier K , and inactivating K (“A-current”) (figures 5.20 and 5.21). The fast Na and delayed-rectifier K channels were similar to the corresponding channels for pyramidal neurons except that the midpoint of the activation curves was shifted slightly. The A-current was not based on experimental data but was created to reduce the excitability of these neurons to a more physiological level. Neuron parameters are listed in tables 5.10 and 5.11. Application of $100 \mu M$ NE increases the resting membrane potential of the feedback inhibitory neuron model by $10 mV$.

5.8.4 Feedforward interneuron model parameters

Feedforward inhibitory neurons, like feedback inhibitory neurons, are modeled as one-compartment neurons with fast Na and delayed-rectifier K channels. These neurons are divided into two classes: those without excitatory feedback from pyramidal cells and those with feedback. Feedforward inhibitory neurons without feedback are identical to feedback interneurons except that their resting membrane potentials are slightly lower ($-73 mV$), they do not have the inactivating K current (A-current), and NE does not alter their excitability. There is currently no experimental data on the effect of NE on these neurons. Feedforward inhibitory neurons with feedback from pyramidal cells are very similar to those without feedback, except for a higher resting membrane potential ($-65 mV$).

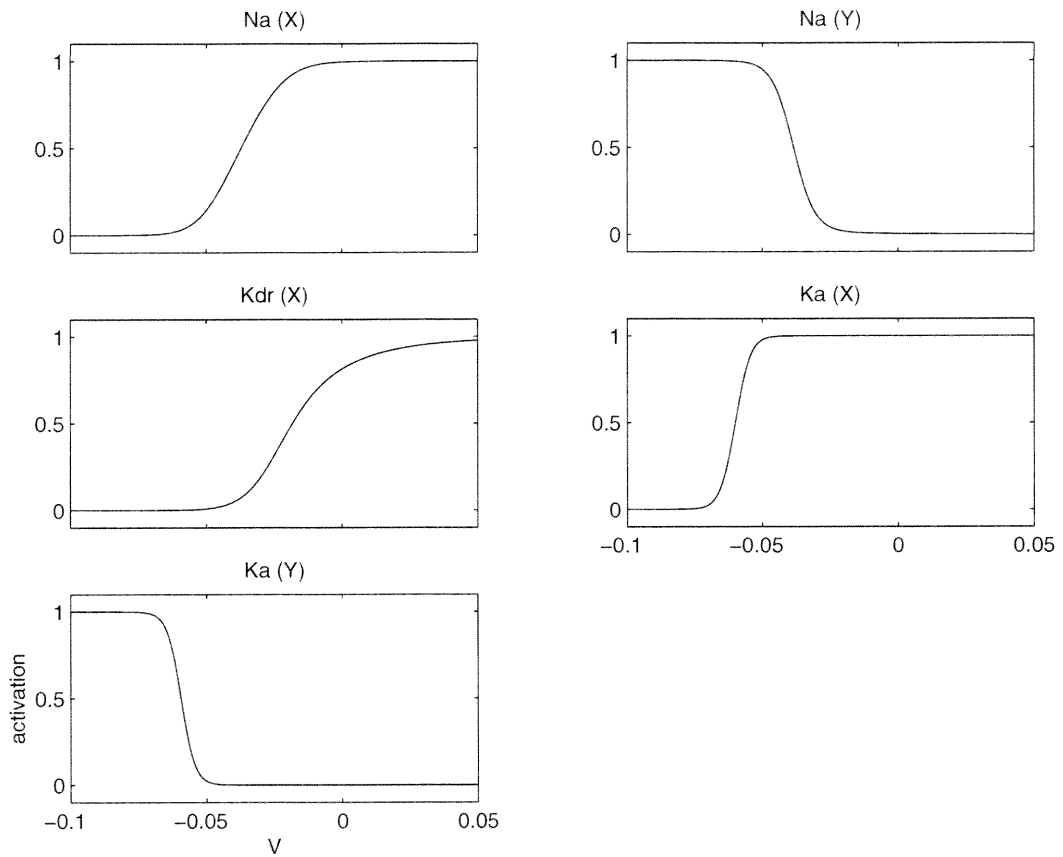


Figure 5.20: Activation (m_∞ vs. V) curves for voltage-dependent ionic channels in the feedback inhibitory interneuron model. x -axis: membrane potential (in volts) of compartment associated with the channel. y -axis: channel activation (dimensionless). Channel abbreviations as in figure 5.17.

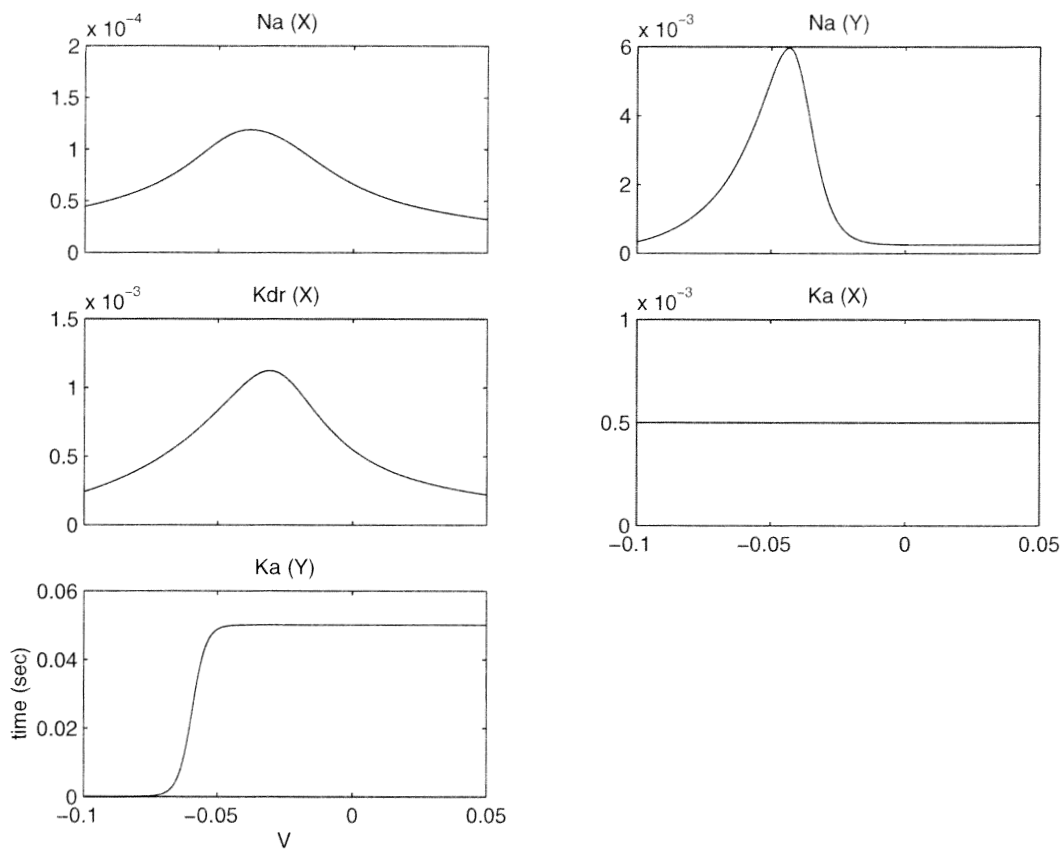


Figure 5.21: Time constant (τ vs. V) curves for voltage-dependent ionic channels in the feedback inhibitory interneuron model. x -axis: membrane potential (in volts) of compartment associated with the channel. y -axis: time constant in seconds. Channel abbreviations as in figure 5.17.

compartment	length (μm)	diameter (μm)	synaptic channels
<i>bascl6</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>bascl5</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>bascl4</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>bascl3</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>bascl2</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>bascl1</i>	79.7	3.2	AMPA, NMDA, GABA _A
<i>soma</i>	21.3	10.4	AMPA, NMDA, GABA _A
<i>deepIb1</i>	57.9	1.6	AMPA, NMDA,
<i>deepIb2</i>	57.9	1.6	AMPA, NMDA,
<i>deepIb3</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB
<i>suppIb1</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB
<i>suppIb2</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB
<i>suppIb3</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB
<i>Ia1</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB
<i>Ia2</i>	57.9	1.6	AMPA, NMDA, GABA _A , GABAB

Table 5.9: Morphology of the reduced pyramidal neuron model and synaptic channels in each compartment.

channel	\bar{G}	a	b	c
Na , fast	1700.0	2	1	0
K , delayed-rectifier	700.0	1	0	0
K , A-current	20.0	1	1	0

Table 5.10: Maximum conductances and gating variable exponents for the ionic channel models used in the feedback inhibitory interneuron model. Conductance units are S/m^3 . Gating variable exponents are dimensionless.

parameter	value
compartment length (μm)	15.0
compartment diameter (μm)	15.0
resting membrane potential (mV)	-70.0
specific membrane resistivity ($k\Omega \cdot cm^2$)	0.8
specific membrane capacitance ($\mu F/cm^2$)	10.0
input resistance ($M\Omega$)	107.0
input time constant ($msec$)	7.9

Table 5.11: Parameter values for feedback inhibitory interneurons.

Bibliography

- [1] P. F. Baldi, M. C. Vanier, and J. M. Bower. On the Use of Bayesian Methods for Evaluating Compartmental Neural Models. *Journal of Computational Neuroscience*, 5:285–314, 1998.
- [2] M. I. Banks, L. B. Haberly, and M. B. Jackson. Layer-specific Properties of the Transient K-current (I-A) in Piriform Cortex. *Journal of Neuroscience*, 16(12):3862–3876, 1996.
- [3] E. Barkai and M. E. Hasselmo. Modulation of the Input/Output Function of Rat Piriform Cortex Pyramidal Cells. *Journal of Neurophysiology*, 72(2):644–658, 1994.
- [4] M. F. Bear and W. C. Abraham. Long-term Depression in Hippocampus. *Annual Review of Neuroscience*, 19:437–462, 1996.
- [5] J. M. Bower and D. Beeman. *The Book of Genesis, 2nd. Ed.* Springer Verlag, New York, 1998.

- [6] J. M. Bower and L. B. Haberly. Facilitating and Nonfacilitating Synapses on Pyramidal Cells – a Correlation Between Physiology and Morphology. *Proceedings of The National Academy of Sciences of The United States of America*, 83(4):1115–1119, 1986.
- [7] P. C. Bush and T. J. Sejnowski. Reduced Compartmental Models of Neocortical Pyramidal Cells. *Journal of Neuroscience Methods*, 46:159–166, 1993.
- [8] G. G. S. Collins, G. A. Probett, J. Anson, and N. J. Mclaughlin. Excitatory and Inhibitory Effects of Noradrenaline on Synaptic Transmission in the Rat Olfactory Cortex Slice. *Brain Research*, 294(2):211–223, 1984.
- [9] A. Constanti and M. Galvan. M-current in Voltage-clamped Olfactory Cortex Neurons. *Neuroscience Letters*, 39(1):65–70, 1983.
- [10] A. Constanti, M. Galvan, P. Franz, and J. A. Sim. Calcium-dependent Inward Currents in Voltage-clamped Guinea-pig Olfactory Cortex Neurons. *Pflugers Archiv European Journal of Physiology*, 404(3):259–265, 1985.
- [11] A. Constanti and J. A. Sim. A Slow, Muscarine-sensitive Ca-dependent-K Current in Guinea-pig Olfactory Cortex Neurons In Vitro. *Journal of Physiology (London)*, 365:P47–P47, 1985.
- [12] A. Constanti and J. A. Sim. Calcium-dependent Potassium Conductance in Guinea-pig Olfactory Cortex Neurons In Vitro. *Journal of Physiology (London)*, 387:173–194,

1987.

- [13] C. R. French, P. Sah, K. J. Buckett, and P. W. Gage. A Voltage-Dependent Persistent Sodium Current in Mammalian Hippocampal-Neurons. *Journal of General Physiology*, 95(6):1139–1157, 1990.
- [14] R. L. Gellman and G. K. Aghajanian. Pyramidal Cells in Piriform Cortex Receive a Convergence of Inputs from Monoamine Activated GABAergic Interneurons. *Brain Research*, 600(1):63–73, 1993.
- [15] R. L. Gellman and G. K. Aghajanian. Serotonin(2) Receptor-Mediated Excitation of Interneurons in Piriform Cortex – Antagonism by Atypical Antipsychotic Drugs. *Neuroscience*, 58(3):515–525, 1994.
- [16] L. B. Haberly. Structure of the Piriform Cortex of the Opossum. 1. Description of Neuron Types with Golgi Methods. *Journal of Comparative Neurology*, 213(2):163–187, 1983.
- [17] L. B. Haberly. Olfactory Cortex. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.
- [18] L. B. Haberly and J. M. Bower. Analysis of Association Fiber System in Piriform Cortex with Intracellular-Recording and Staining Techniques. *Journal of Neurophysiology*, 51(1):90–112, 1984.

- [19] L. B. Haberly and S. L. Feig. Structure of the Piriform Cortex of the Opossum. 2. Fine Structure of Cell Bodies and Neuropil. *Journal of Comparative Neurology*, 216(1):69–88, 1983.
- [20] E. W. Harris and C. W. Cotman. Effects of Synaptic Antagonists on Perforant Path Paired-pulse Plasticity – Differentiation of Presynaptic and Postsynaptic Antagonism. *Brain Research*, 334(2):348–353, 1985.
- [21] M. E. Hasselmo, C. Linster, M. Patil, D. Ma, and M. Cekic. Noradrenergic Suppression of Synaptic Transmission May Influence Cortical Signal-to-Noise Ratio. *Journal of Neurophysiology*, 77(6):3326–3339, 1997.
- [22] B. Hille. *Ionic Channels of Excitable Membranes, 2nd. Ed.* Sinauer Associates Inc., Sunderland, Massachusetts, 1992.
- [23] J. J. B. Jack, D. Noble, and R. W. Tsien. *Electric Current Flow in Excitable Cells.* Clarendon Press, Oxford, 1983.
- [24] D. Johnston, J. C. Magee, C. M. Colbert, and B. R. Christie. Active Properties of Neuronal Dendrites. *Annual Review of Neuroscience*, 19:165–86, 1996.
- [25] A. Kapur, W. W. Lytton, K. L. Ketchum, and L. B. Haberly. Regulation of the NMDA Component of EPSPs by Different Components of Postsynaptic GABAergic Inhibi-

- tion: Computer Simulation Analysis in Piriform Cortex. *Journal of Neurophysiology*, 78(5):2546–2559, 1997.
- [26] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, Oxford, U.K., 1999.
- [27] C. Koch and I. Segev. *Methods in Neuronal Modeling, 2nd. Ed.* MIT Press, Cambridge, Massachusetts, 1997.
- [28] S. W. Kuffler, J. G. Nicholls, and A. R. Martin. *From Neuron to Brain: A Cellular Approach to the Function of the Nervous System*. Sinauer Associates, Sunderland, Massachusetts, 1984.
- [29] D. V. Madison and R. A. Nicoll. Noradrenaline Blocks Accommodation of Pyramidal Cell Discharge in the Hippocampus. *Nature*, 299(5884):636–638, 1982.
- [30] D. V. Madison and R. A. Nicoll. Cyclic Adenosine 3'-5' Monophosphate Mediates Beta-Receptor Actions of Noradrenaline in Rat Hippocampal Pyramidal Cells. *Journal of Physiology (London)*, 372:245–259, 1986.
- [31] Z. F. Mainen, J. Joerges, J. R. Huguenard, and T. J. Sejnowski. A Model of Spike Initiation in Neocortical Pyramidal Neurons. *Neuron*, 15(6):1427–1439, 1995.
- [32] Z. F. Mainen and T. J. Sejnowski. Reliability Of Spike Timing In Neocortical Neurons. *Science*, 268:1503–1506, 1995.

- [33] G. J. Marek and G. K. Aghajanian. Alpha(1B)-adrenoceptor-mediated Excitation of Piriform Cortical Interneurons. *European Journal of Pharmacology*, 305(1-3):95–100, 1996.
- [34] D. A. McCormick, B. W. Connors, J. W. Lighthall, and D. A. Prince. Comparative Electrophysiology of Pyramidal and Sparsely Spiny Stellate Neurons of the Neocortex. *Journal of Neurophysiology*, 54(4):782–806, 1985.
- [35] A. D. Protopapas. Pyramidal cell Responses to Temporally Structured Stimuli: Experiments and Computer Simulations. PhD Thesis, Department of Biology, California Institute of Technology, 1998.
- [36] A. D. Protopapas and J. M. Bower. Physiological Characterization of Layer III Non-Pyramidal Neurons in Piriform (Olfactory) Cortex of Rat. *Brain Research*, 865(1):1–11, 2000.
- [37] W. Rall. *The Theoretical Foundation of Dendritic Function*. MIT Press, Cambridge, Mass., 1995.
- [38] P. W. Sheldon and G. K. Aghajanian. Excitatory Responses to Serotonin (5-HT) in Neurons of the Rat Piriform Cortex – Evidence for Mediation by 5-HT_{1c} Receptors in Pyramidal Cells and 5-HT₂ Receptors in Interneurons. *Synapse*, 9(3):208–218, 1991.

- [39] G. R. Shepherd, editor. *The Synaptic Organization of the Brain, 4th. Ed.* Oxford University Press, New York, 1997.
- [40] T. J. Shors and L. D. Matzel. LTP: Memory, Arousal, Neither, Both. *Behavioral and Brain Sciences*, 20(4):634, 1997.
- [41] C. F. Stevens. A Million Dollar Question: Does LTP = Memory? *Neuron*, 20(1):1–2, 1998.
- [42] G. J. Stuart. Voltage-Activated Sodium Channels Amplify Inhibition in Neocortical Pyramidal Neurons. *Nature Neuroscience*, 2(2):144–150, 1999.
- [43] G. J. Stuart and B. Sakmann. Active Propagation of Somatic Action Potentials into Neocortical Pyramidal Cell Dendrites. *Nature*, 367:69–72, 1994.
- [44] R. D. Traub, R. K. S. Wong, R. Miles, and H. Michelson. A Model of a CA3 Hippocampal Pyramidal Neuron Incorporating Voltage-Clamp Data on Intrinsic Conductances. *Journal of Neurophysiology*, 66(2):635–650, 1991.
- [45] G. Tseng and L. B. Haberly. Characterization of Synaptically Mediated Fast and Slow Inhibitory Processes in Piriform Cortex in an In Vitro Slice Preparation. *Journal of Neurophysiology*, 59(5):1352–1376, 1988.

- [46] Madison D. V., Malenka R. C., and Nicoll R. A. Mechanisms Underlying Long-Term Potentiation of Synaptic Transmission. *Annual Review of Neuroscience*, 14:379–397, 1991.
- [47] M. C. Vanier and J. M. Bower. A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models. *Journal of Computational Neuroscience*, 7(2):149–171, 1999.
- [48] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.
- [49] L. Zhang and K. Krnjević. Whole-Cell Recording of Anoxic Effects on Hippocampal Neurons in Slices. *Journal of Neurophysiology*, 69:118–127, 1993.

Chapter 6

Building the Piriform Cortex Network

Model

In this chapter, we discuss the process of constructing and parameterizing the piriform cortex network model in some detail.

6.1 Modeling inputs from the olfactory bulb

A realistic network model of a brain structure is of limited usefulness without some reasonable representation of the inputs to the network. In order to provide this input, a spike generating model was built within GENESIS which can generate a variety of firing patterns either known or believed to be present in the mitral cell outputs from the olfactory bulb. This model is purely phenomenological; it produces a series of spike outputs which are sta-

tistically similar to those seen in experimental data, but it does not replicate the internal dynamics of the olfactory bulb nor does it receive feedback from the piriform cortex model. The statistical approach was taken because of the enormous variation in spike responses of individual mitral cells in response to repeated presentations of the same odor (figure 6.1). The construction of this model and the data set on which it was based are described in this section.

6.1.1 Data set

The data set which was used as the basis for the olfactory bulb spike generating object was collected by Upinder S. Bhalla as part of his thesis work in the Biology Department at the California Institute of Technology [4]. The data set consisted of extracellular recordings from mitral cells in awake behaving rats exposed to three different odors (citral, isoamyl acetate, and methyl salicylate) at nanomolar concentrations. The low concentrations were chosen to avoid problems relating to odor adaptation; however, the concentrations were still readily detectable by the rats. A total of 44 neurons from seven different rats were recorded over multiple days. Spike times were extracted from the raw data by manual clustering methods using the Brainwave software package. In addition, a thermistor was used to monitor the respiration rate of the rat. Further details of the experimental procedures are described in [4].




Figure 6.1: Raster plot of the response of a single mitral cell to the same odor upon repeated presentations. Each raster represents a single spike. Note the large variability in responses.

6.1.2 Data analysis

For the purposes of providing inputs to the piriform cortex model, we were primarily interested in the background (no odor) firing behavior of mitral cells, how this firing behavior varied with odor presentation for different odors, and how this behavior varied between different neurons. Since odors can be detected by a rat in about 150-200 *msec* [7], and since our network simulations typically range from 0.3 to 2 seconds in duration (see chapter 7), we were principally interested in the short-time response to an odor. Therefore, we restricted our analysis to three basic features of the firing patterns for single mitral cells: the interspike interval (ISI) distribution, the spike autocorrelation function, and the peristimulus time histograms (PSTHs) of averaged spike rates after odor onset. For these analyses, we found no significant differences in the responses of individual neurons on different days, so data from different days was pooled for the final analyses. Details on the data analyses are given in appendix A.

The results of the analyses show a remarkable similarity in responses of different neurons, whether responding to the same odor or to different ones. We never observed any change in the responses of any of the neurons which depended on the odor identities. In fact, we rarely observed any noticeable change from the background odor response; the only exception was in the PSTHs where a small subset of the neurons (2/44) displayed an initial rate increase followed by a decrease. Interestingly, both of these neurons were from the same rat. Most PSTHs were effectively flat (no change in response to odors). The ISI distributions (figure

6.2) and the autocorrelograms varied to some extent from neuron to neuron, but most were broadly similar, with a single peak at around 20 *msec*, corresponding to the gamma frequency band (50 *Hz*; see [5]). This suggests that different odors at low concentrations do not significantly modulate the firing rate, the interspike interval distribution or the autocorrelograms of individual mitral cells. Therefore, this data set is more consistent with coding strategies involving synchrony or spatiotemporal spike patterns than with rate modulation codes. However, the small size of the data set does not allow us to rule out rate coding altogether. Therefore, we constructed the spike generating model to be flexible enough to produce a variety of different kinds of output encodings, including both rate and synchrony coding.

6.1.3 Building the spike generating model

The spike generating model, which was constructed as a simulation object in GENESIS, was built to accurately replicate the first-order statistics of the data for background odor inputs. Since there was no real difference between the background response and the odor responses for most neurons, we used the pooled ISI distribution (figure 6.2) as the basis of the model. The model can generate arbitrary ISI distributions using the following algorithm. First, the average ISI distribution was truncated at 1 second; this is reasonable since ISIs of less than 1 second account for over 98% of the measured ISIs. This distribution is read into the model and converted into a cumulative probability distribution (*i.e.*, the probability that

the next ISI will be greater than or equal to a particular value). For each neuron modeled, the time the neuron spiked last is recorded. On each time step, a spiking probability value is determined for each neuron by looking up the cumulative probability value corresponding to the time since the last spike. Clearly, the longer the time since the last spike, the more likely it is that a given neuron will fire. Figure 6.3 shows the result of the model generating background spikes for 100 simulated mitral cells for 1 second.

In addition to modeling the background spiking activity of mitral cells, we designed the model to reproduce a variety of possible odor encoding strategies. The two main categories of coding strategies that the model can reproduce are rate coding and synchrony coding. One advantage of the algorithm used for generating background ISIs is that it can be easily adapted to generate rate-modulated output by simply modifying the effective time since the last spike. In figure 6.4, we see four different rate change patterns following a single odor onset in different subsets of modeled mitral cells. To model synchrony coding we define a subset of neurons which can respond to synchronized events; the remaining neurons fire at the background rate. A given unit can respond probabilistically to a synchronization event; if it does not respond it goes on firing at the background rate. If it does respond, it has a probability of spiking which is a uniform function of the time since the synchronization event. Once a neuron has fired in response to a synchrony event, it will not fire again until the next synchrony event occurs. Figure 6.5 shows a typical response of the mitral cell array with synchrony coding enabled. The upper 20 neurons fire in synchrony after the odor




Figure 6.2: Interspike interval (ISI) distribution in mitral cells. The black trace is the smoothed ISI histogram derived from mitral cell recordings. The red trace is an ISI histogram from the GENESIS olfactory bulb object for one simulated mitral cell run for 1000.0 seconds. The histogram counts have been normalized so that the peak of the histogram is 1.0.




Figure 6.3: Background spiking responses of 100 simulated mitral cells. Each raster represents one spike.

onset at $t = 1$ second. As can be seen, the neurons are not perfectly synchronized with each other. The model's ability to generate synchronized spikes is very useful when generating weak and strong shock inputs, which correspond to a large number of input fibers firing simultaneously due to voltage shocks to the LOT applied by the experimenter [19, 20, 21]. We have used shock stimuli extensively to parameterize the network model (see below).

6.1.4 Validating the spike generating model

Figure 6.2 shows the smoothed ISI distribution computed from the data compared to the ISI distribution computed from the spike generating model run for 1000 seconds. As can be seen, apart from some random variation in the model outputs the two curves are essentially identical.

Figure 6.6 shows the pooled and normalized autocorrelation distribution from the data compared to the autocorrelation distribution computed from the spike generating model run for 10000 seconds. The primary difference is that the model shows larger short-time correlations than the mitral cell data. This difference is modest (about 16% at its largest point), and may be a result of the strong self-inhibition in mitral cells mediated by granule cell dendrodendritic synapses [30] which is not found in the model. We decided not to modify the spike generating model to compensate for this effect since this would have greatly complicated the model.




Figure 6.4: Rate coding in the olfactory bulb spike generating model. Four different rate change patterns were applied to 100 modeled mitral cells during the odor presentation, which starts at $t = 1$ sec and ends at $t = 2$ sec. The four patterns are (from top): increase in firing rate, decrease followed by increase, increase followed by decrease, no change. Spikes occurring before and after the odor presentation are background responses of the mitral cells.




Figure 6.5: Synchrony coding in the olfactory bulb spike generating model. The upper 20 cells fire synchronously during the odor presentation, which starts at $t = 1 \text{ sec}$ and ends at $t = 2 \text{ sec}$. Spikes occurring before and after the odor presentation are background responses of the mitral cells.




Figure 6.6: Autocorrelation of spike times in experimental and modeled mitral cells. The black trace represents the experimental data while the red trace represents model output.

6.1.5 Conclusions

The spike generating model described in this section can replicate background spike input from mitral cells to piriform cortex with reasonable statistical properties, as reflected in the interspike interval distribution and the autocorrelogram. In addition, it can implement several coding strategies, including rate coding (with an arbitrary rate modulation waveform after odor onset) and synchrony coding. The model can also generate the equivalent of shock inputs to piriform cortex, where a large proportion of LOT fibers are made to fire simultaneously [19, 20, 21].

Nevertheless, the use of a spike generating model to provide inputs to the cortical model is not optimal. Strong feedback projections are known to exist between piriform cortex and olfactory bulb [13]. These projections may be partly responsible for the high degree of coherence often seen between cortical and bulbar surface field potentials (*e.g.*, [5]). The only way to adequately model these phenomena is with a combined olfactory bulb and piriform cortex model, which was outside the scope of the present work. In addition, the spike generating model only captures the first-order statistics of the mitral cell firing patterns, and can only replicate very simple collective behaviors, such as synchronized firings. If the olfactory code involves very complex spatiotemporal spike sequences in the mitral cell population [22, 31], such a code could not be replicated by the present model. This highlights an important area for future extensions of this model.

6.2 Modeling the piriform cortex network

6.2.1 Previous modeling work on piriform cortex

The earliest computational models of piriform cortex are those of Walter Freeman [9] who used lumped linear and nonlinear system models to account for the oscillatory surface field potentials in piriform cortex. Freeman has extended these models to account for large parts of the olfactory system, including piriform cortex, olfactory bulb, and the anterior olfactory nucleus [10, 11]. These models were constructed to replicate the collective behaviors of large groups of neurons (generally only the surface field potential) and do not attempt to account for individual neuron responses. As will be seen below, however, specific neuronal and synaptic phenomena can have a large effect on network dynamics in our piriform cortex model; such phenomena are outside the scope of Freeman's models.

Richard Granger and Gary Lynch produced a model of the olfactory system (olfactory bulb and piriform cortex) as a hierarchical clustering pattern recognition system [2]. Their model is interesting in that it assigns a vital role to back-projections from piriform cortex pyramidal neurons to olfactory bulb granule cells (which inhibit mitral cells). However, the model is a connectionist model, not a realistic neural model. The model does not assign a significant role for the dense associational projections in piriform cortex, and does not generate or depend on the oscillatory dynamics of the piriform cortex, both of which are important features of the work described here.

Kevin Ketchum and Lewis Haberly have built an unusual model of piriform cortex in order to account for the response of the cortex to a strong electrical shock stimulus of the LOT [20], as reflected in current-source density (CSD) data. Their approach was to build a simplified passive cable model of a pyramidal neuron and subject it to artificially generated spikes along axonal pathways with experimentally determined delay distributions. The goal of the model was to dissect out the various synaptic and passive components giving rise to the strong shock response, as well as to determine the distributions of input spikes that were most likely to give rise to this response. However, there was only one pyramidal neuron, and as this neuron was purely passive the overall model cannot be considered a realistic network model. Fortunately, the process of constructing this model led them to perform a number of experiments to estimate several classes of network parameters (notably axonal delay distributions); we have incorporated these parameters into the network model described here.

Michael Hasselmo and his collaborators have produced a series of models of piriform cortex ranging from connectionist-level [15] to moderately realistic [3]. Their models are of interest in that neuromodulation (generally mediated by acetylcholine but also occasionally by other neuromodulators including norepinephrine [16]) is an integral component of the functioning of the models. However, none of the models are at the same level of detail as the model described here. The models are theory-driven and have led to a number of interesting speculations; the most fully worked out of these is that neuromodulators may

serve to prevent previously learned olfactory memories from interfering with the learning of new memories [15, 14]. In contrast, our approach is to build as realistic a model as possible and let the model guide our understanding rather than using the theory as a starting point. Nevertheless, Hasselmo’s models have been a major influence on our decision to include neuromodulation as an integral component of our model.

The model described below is the direct descendant of a realistic neuronal network model of piriform cortex constructed by Matthew Wilson [36]. This model, like the present model, was constructed using the GENESIS simulator. The model contained pyramidal neurons, feedback inhibitory interneurons, and feedforward interneurons, and was able to replicate many of the dynamical behaviors of the piriform cortex network, including the weak and strong shock responses as reflected in the piriform cortex surface field potential (EEG). In contrast to the present model, Wilson’s model did not attempt to accurately replicate the spiking patterns of any of the constituent neurons nor did it use inputs derived from mitral cell recordings. In addition, much of the data that was used to parameterize the Wilson model has been superseded by more recent data which is incorporated in the new model. Finally, the Wilson model did not incorporate neuromodulation of cellular or synaptic properties. Our model is thus a logical development of the Wilson model, with similar goals but with much greater emphasis on accurate replication of cellular and network-level phenomena.

6.2.2 Construction of the model

Cortical structure

The piriform cortex is divided anatomically into three regions: the ventral anterior piriform cortex (vAPC), the dorsal anterior piriform cortex (dAPC) and the posterior piriform cortex (PPC) [13, 19] (figure 6.7). Afferent fibers from the olfactory bulb synapse directly onto pyramidal neurons in the vAPC through the lateral olfactory tract (LOT), and to the rest of the cortex through LOT collaterals. This arrangement is preserved in the model by designating the posterior half of the model to be the PPC, a quarter-elliptical region in the anterior half of the model to be the vAPC, and the rest to be the dAPC (figure 6.9). Both afferent and intranetwork connectivities vary greatly between different subdivisions of the network (see below).

The coordinate axes for the model were arranged as follows. The direction normal to the surface of the cortex, along which the pyramidal neurons were oriented, was assigned to be the z direction. The anterior-posterior axis was the x direction and the medial-lateral axis was the y direction. The antero-posterior extent of the cortex was 8 mm while the medio-lateral extent was 3 mm (figure 6.9). These values were obtained from a rat brain atlas [27].

Neurons

The model has four types of neurons:

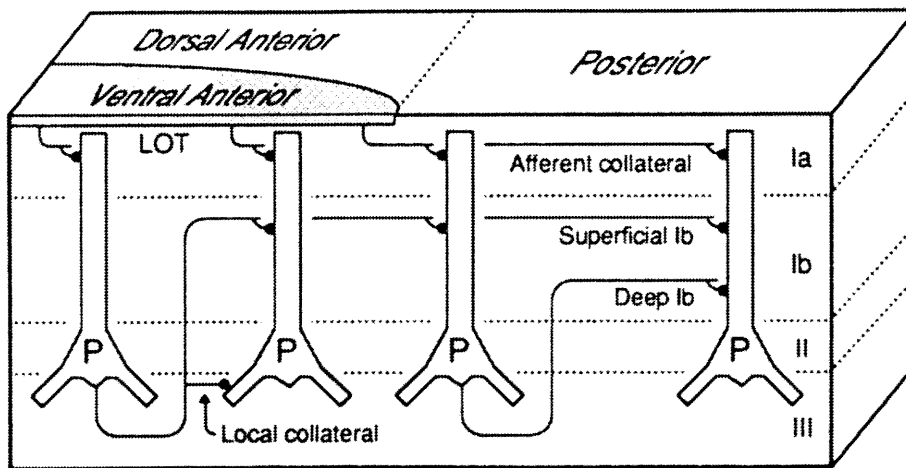


Figure 6.7: Subdivisions of the piriform cortex showing both the layered structure (layers 1a, 1b, 2 and 3) and the division into ventral anterior, dorsal anterior, and posterior regions. Abbreviation: *P*, pyramidal neuron. From [19].

1. Pyramidal neurons;
2. Feedback inhibitory neurons, which are equivalent to layer 3 aspiny multipolar neurons;
3. Feedforward inhibitory neurons, which are equivalent to layer 1 horizontal neurons;
4. Feedforward/feedback inhibitory neurons, which are equivalent to layer 1 globular-soma neurons.

The construction and parameterization of these neuron models was described in chapter 5. The number of neurons of each type in the network model are parameters of the model; the most frequently used scales are listed in table 6.1. The olfactory bulb spike generating model generated the outputs of 1000 olfactory bulb mitral cells at all network scales.

scale	number of pyramidal neurons	number of interneurons of each type
small	96 (16x6)	75 (15x5)
medium	384 (32x12)	341 (31x11)
large	1000 (50x20)	931 (49x19)

Table 6.1: Network scales. The number of neurons is listed in absolute value as well as in x-y coordinates (*e.g.*, (16x6) means an array of 16 neurons in the x direction by 6 neurons in the y direction).

The pyramidal neurons and the three different classes of interneurons are arranged at different levels in the z dimension as shown in figure 6.8, with feedforward inhibitory

interneurons located superficially and feedback inhibitory interneurons located below the pyramidal neuron cell bodies. In the x and y dimensions, the different neuronal types are equally spaced with the inhibitory interneurons located between the pyramidal neurons (figure 6.9). The number of inhibitory interneurons is comparable to the number of pyramidal neurons. Since inhibitory interneurons are much less dense than pyramidal neurons in the real cortex, the relative number of inhibitory interneurons is higher in the model than in real cortex [12, 29]. This was done to make sure that each inhibitory interneuron is locally equidistant from the same number of pyramidal neurons, which simplifies the setup of the synaptic connections. Note that the total amount of inhibitory input to pyramidal neurons is scaled to be independent of the number of inhibitory interneurons (see below).

Network connectivities

A considerable amount of data has been collected about network connectivities in the piriform cortex (summarized in [13]). In this section we briefly describe these connectivities and their implementation in the piriform cortex model. These connections are diagrammed schematically in figure 6.10.

All connections in the cortical model were set up probabilistically. Regions of source neurons and destination compartments on neurons were chosen, and then connections were made randomly at a fixed probability. Connection strengths (weights) and axonal delays were also set up probabilistically. Weights were set up with a Gaussian distribution while




Figure 6.8: Arrangement of neuron types in the piriform cortex model. Piriform cortex layers are listed on the right. Depths are in μm . Abbreviations: *FF*, feedforward inhibitory interneuron (horizontal cell); *FF/FB*, feedforward/feedback inhibitory interneuron (globular-soma cell); *FB*, feedback inhibitory interneuron (aspiny multipolar neuron); *sup 1b*, superficial layer 1b; *deep 1b*, deep layer 1b.




Figure 6.9: Positions of neurons in the piriform cortex model in x - y coordinates (“bird’s eye” view). This figure corresponds to a small-scale (96 pyramidal neuron) model. Black squares represent pyramidal neurons. Red diamonds represent inhibitory neurons (all classes). All dimensions are in mm . Abbreviations: $vAPC$, ventral anterior piriform cortex; $dAPC$, dorsal anterior piriform cortex; PPC , posterior piriform cortex.

delays had a lognormal distribution [20].

The ventral anterior region of the piriform cortex (vAPC) receives direct input from the olfactory bulb via the lateral olfactory tract (LOT). The main LOT fibers have a conduction velocity of approximately 7.0 m/sec , which results in the nearly simultaneous activation of neurons in the vAPC region when given a shock stimulus (see below) [19, 20]. LOT collaterals branch off at a 45° angle and spread over the surface of the rest of the cortex with a conduction velocity of approximately 0.8 m/sec (figure 6.11). Afferent connectivities for feedforward interneurons were similar to those for pyramidal neurons.

Pyramidal neurons make both short and long-range connections to other pyramidal neurons in the cortex. The density of connections is a function of the source region, as shown in table 6.2. Long-range projections from the vAPC synapse on the dendrites of pyramidal neurons in the superficial layer 1b region, whereas long-range projections from the PPC and dAPC synapse in the deep layer 1b region (figure 6.12). Local connections between pyramidal neurons synapse on the cell bodies and the basal dendrites of pyramidal neurons. Connections from feedforward and feedback inhibitory interneurons are local (short-range) only.

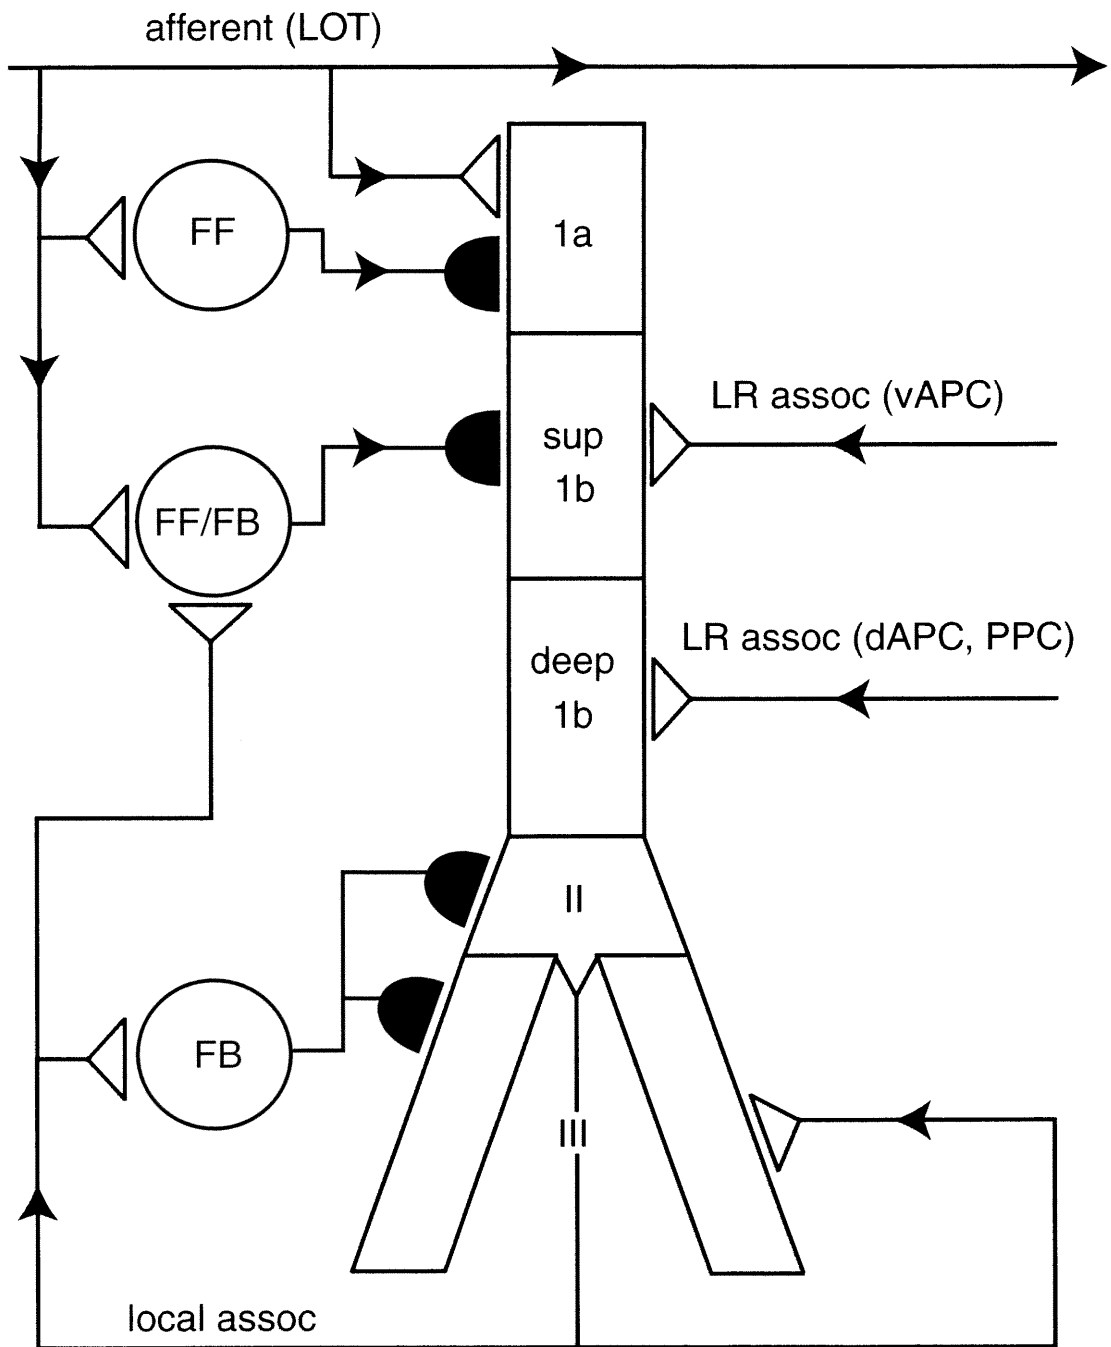


Figure 6.10: (Previous page) Schematic diagram of connections in the piriform cortex model. The large central cell is a pyramidal neuron, with cortical layers as indicated. Empty triangles represent excitatory synapses, while filled half-ovals represent inhibitory synapses. Abbreviations: *FF*, feedforward inhibitory interneuron; *FF/FB*, feedforward/feedback inhibitory interneuron; *FB*, feedback inhibitory interneuron; *LR assoc*, long-range association fibers; *local assoc*, local (short-range) association fibers; *vAPC*, ventral anterior piriform cortex; *dAPC*, dorsal anterior piriform cortex; *PPC*, posterior piriform cortex.

region	local connections	long-range connections		
		to vAPC	to dAPC	to PPC
vAPC	sparse	NA	moderate	dense
dAPC	moderate	moderate	NA	dense
PPC	dense	moderate	moderate	NA

Table 6.2: Network connectivities. NA = not applicable.

The distribution of synaptic weights from inhibitory neurons along the pyramidal neuron dendrites is not uniform. Fibers from feedback inhibitory interneurons synapse on the cell bodies and proximal basal dendrites of the pyramidal neurons (figure 6.10). Fibers from feedforward interneurons synapse fairly evenly along the apical dendrites of pyramidal neurons [18].




Figure 6.11: Relative strengths of afferent projections in the piriform cortex model. The strength of a connection is (schematically) proportional to the thickness of the line ending in a region of cortex. All dimensions are in *mm*. Abbreviations: *vAPC*, ventral anterior piriform cortex; *dAPC*, dorsal anterior piriform cortex; *PPC*, posterior piriform cortex.




Figure 6.12: Schematic diagram of excitatory connections between pyramidal neurons (*P*) in the piriform cortex. Each pyramidal neuron schematic represents a class of neurons in one of the three subdivisions of this brain region. The density of connections between regions is represented by the width of lines connecting the pyramidal neurons. Shading of pyramidal neurons and connections is just to distinguish between the three regions. Abbreviations: *M/T*, mitral/tufted cells of the olfactory bulb; *G*, granule cells of the olfactory bulb; *LOT*, lateral olfactory tract. From [13].

Scaling issues

Even at the largest scale, the network model has very few neurons relative to the real cortex, which has on the order of 10^6 neurons [13]. In addition, it is desirable for the network model to display broadly similar behaviors even when the scale is changed. We have used a number of strategies for dealing with scaling issues which we briefly describe here.

Connection probabilities in the model vary with the overall scale of the network, with smaller scale models having larger connection probabilities. This was done for two reasons. First, small connection probabilities in a small-scale model would result in too few connections in the model, each of which was relatively strong. This would have led to too much variability based on the specific locations of the connections. Conversely, high connection probabilities in a large-scale model would have been computationally very expensive since the number of synaptic connections increases as the square of the number of neurons.

Synaptic weights for afferent and associational connections decay exponentially with distance from the source (schematically illustrated in figure 6.11). This is done to compensate for the small scale of the network relative to the real cortex. The underlying probabilistic model is that of an axon collateral which has a fixed probability per unit distance of making a synapse and then stopping. This gives rise to an exponential distribution in the number of connections as a function of distance from the source. Given the small number of neurons in the model, it is more convenient (and more accurate) to modify the weights than the connection probabilities.

The model has a relatively larger number of inhibitory interneurons than the real cortex. While it would have been possible to match the relative numbers in the model, it was more convenient to set up evenly distributed lattices of neurons with the interneurons occupying the spaces between the pyramidal neurons in the xy plane. The larger number of inhibitory neurons does not affect the total amount of inhibitory input, which is set independently of synapse numbers. In fact, the larger number of inhibitory synapses probably improves the model by reducing the scale-dependent granularity of the model.

Within the cortex, synaptic connection strengths were normalized so that each destination compartment received a fixed amount of total input of each type (*e.g.*, excitatory or inhibitory). This was done to minimize edge effects whereby neurons on the edge of the simulated cortex would receive less input because there were fewer afferent neurons able to project to them over a small connection radius.

Implementing neuromodulation in the cortical model

In chapter 5, we described the effects of neuromodulation ($100 \mu M$ NE) on the synapses, ionic channels, and cellular properties of the neurons in the network. These effects were incorporated into the model, and included

1. pyramidal neurons:
 - (a) increase in layer 1a (afferent) excitatory synaptic potentials
 - (b) decrease in layer 1b (associational) excitatory synaptic potentials

- (c) increase in facilitation in layer 1b (associational) excitatory synapses
 - (d) increase in membrane potential (all parts of the neuron)
 - (e) modulation of potassium conductances in the cell body
2. feedback inhibitory neurons: increase in resting membrane potential
 3. olfactory bulb mitral cells: decrease in background firing rate

Note that feedforward and feedforward/feedback inhibitory interneurons are not neuro-modulated.

The evidence for the cellular and synaptic effects of norepinephrine was presented in chapters 2 and 5. The effects of NE on olfactory bulb mitral cells are taken from [17].

Computational requirements

The network model was simulated on pentium III workstations running Red Hat Linux 6.1 and 6.2. The computational requirements for simulating the network for one second of simulation time are listed in table 6.3 for three different network scales.

6.2.3 Outputs from the cortical model

The network model is, in principle, capable of generating output files for any set of state variables in the model. The following outputs were routinely generated for each simulation run:

scale	# pyramidal neurons	simulation duration	memory use
small	96	50 minutes	36 megabytes
medium	384	4 hours	59 megabytes
large	1000	11 hours	145 megabytes

Table 6.3: Computational requirements for simulating the piriform cortex for one second of simulation time at different scales.

1. extracellular field potentials at the surface of the cortex;
2. extracellular field potentials at various depths in the cortex;
3. membrane potentials of all neuronal compartments;
4. spike outputs of all cortical neurons;
5. spike outputs from olfactory bulb mitral cells.

In addition, occasionally other outputs such as synaptic conductances were recorded as well. Depth recordings of field potentials were used to generate current-source density (CSD) graphs as described below.

In order for extracellular field potential measurements and CSD analysis to be useful, the neurons in the network must be arranged in a laminar manner so that the current sources and sinks line up at least approximately (figure 6.13). This is the case for pyramidal neurons given shock input stimuli: since the shock inputs all arrive at roughly the same

time in a given x - y position of cortex, the source/sink pattern is roughly synchronized for all neurons in this region. Currents in the x - y plane from nearby pyramidal neurons cancel each other out whereas those along the z axis sum together (figure 6.13). The effect of the net current passing through the extracellular resistance is a voltage gradient (figure 6.14); these voltages are what is measured by field potential electrodes. The voltage gradient in a series of field potential measurements at nearby depths can be converted into current-source densities. The theoretical basis for measuring extracellular field potentials and computing CSD outputs in the model is described in appendix B. Examples of both types of measurements are shown below in section 6.2.4 and in the next chapter.

6.2.4 Parameterizing the model

In order to parameterize the piriform cortex network model, we chose two well-defined behaviors exhibited by the real cortex that are not dependent on specific inputs from the olfactory bulb. These behaviors result from the application of weak or strong electric shocks to the cortex via the lateral olfactory tract (LOT) inputs from the olfactory bulb [19, 20, 21]. The network response, measured in terms of the extracellular field potentials and/or the current-source densities, is highly stereotyped for each behavior. In addition, simulations of only 300 msec are adequate to obtain these behaviors. In order to speed up the parameterization process we used the small scale model exclusively, although we checked the results against the model at larger scales.

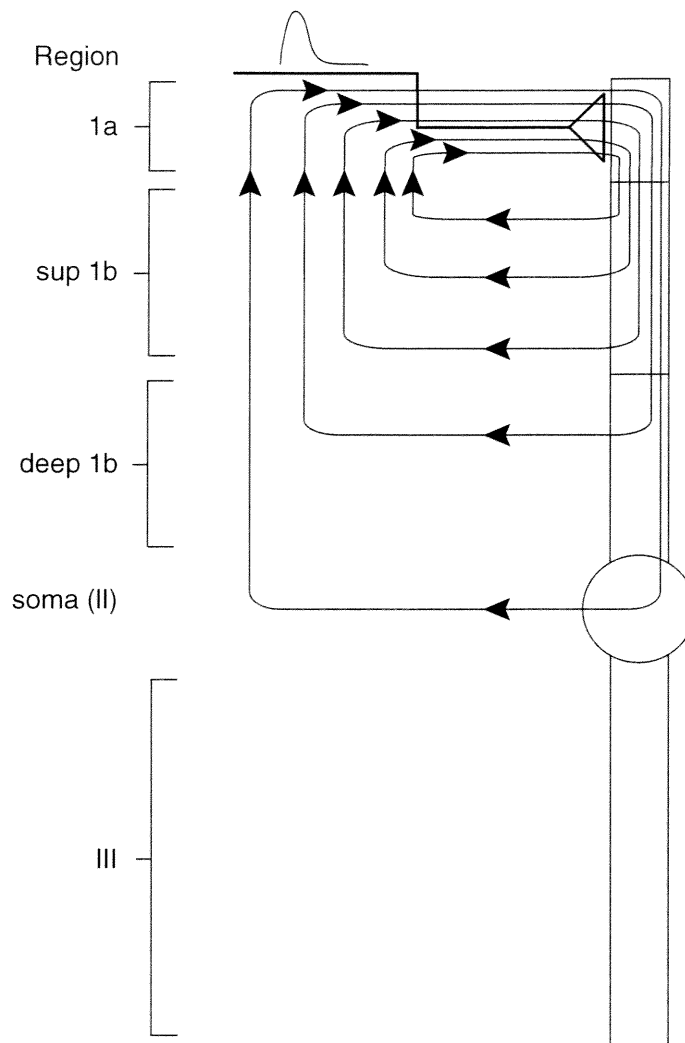


Figure 6.13: Schematic diagram of transmembrane current loops in response to synaptic activation in layer 1a of a superficial pyramidal neuron in piriform cortex. Note that the inward currents (current sinks) are concentrated near the synapse whereas the outward currents are spread out along the apical dendrite and the soma.




Figure 6.14: Schematic diagram of voltage gradient in response to synaptic activation in layer 1a of a superficial pyramidal neuron in piriform cortex. Current flowing upwards in the direction of the apical dendrite causes a voltage drop between the soma and the distal apical dendrite. Each line is an isopotential contour. Dashed lines represent negative voltages, while solid lines represent positive voltages. The thick dashed line represents a potential of zero.

The number of free or loosely constrained parameters in a realistic neuronal network model is clearly enormous. We reduced this number substantially by fixing the parameters of the constituent neuron models during this process. The remaining parameters mainly fell into the following classes:

1. synaptic connection probabilities;
2. synaptic connection topologies (connection patterns);
3. axonal conduction velocities;
4. total synaptic weight (expressed as unit synapse numbers) onto a given region of each neuron from a given source;
5. the rate of decay of synaptic weights with distance from the source.

Of these parameter classes, experimental data exists only for connection topologies and axonal conduction velocities [13, 19, 20]. In practice, the majority of parameter adjustments concerned the synaptic weights. The adjustable parameters numbered about 100 of which about 30 corresponded to synaptic weight parameters. More details are in the simulation scripts themselves, which are available online at:

<http://www.bbb.caltech.edu/USERS/mvanier/thesis/thesis.html>.

Parameterization of the network model was done manually. An important future direction of this work is to see which aspects of this process can be automated, as was done for single-neuron models in chapter 3 and [34].

In the following sections we briefly describe the network behaviors for the two types of stimuli as well as the behavior of the network when subjected to background stimulation alone. Our goal here is to show that the network is capable of accurately generating these behaviors; more details on the network dynamics underlying these behaviors are given in chapter 7.

6.2.5 Response to background inputs

As will be described in more detail in chapter 7, in order to obtain good results for the weak and strong shock responses, the strength of afferent weights into the cortex needed to be adjusted until pyramidal neurons in the ventral anterior piriform cortex (which receive the largest number of synaptic connections from the olfactory bulb) were close to the threshold for spiking. In figure 6.15, we see the distribution of membrane potentials in the pyramidal neuron cell bodies 250 msec after background inputs from the olfactory bulb were activated. Neurons in the ventral anterior piriform cortex (vAPC, lower left corner) are close to firing threshold (which is roughly at -55 mV). Further from the vAPC the membrane potentials decrease, and at the outer edge of the cortex the membrane potential is only about 3 mV above resting potential (roughly -75 mV). This gradient in membrane potential is due to the exponential falloff of the strength of the synaptic inputs from the olfactory bulb with distance. There have been very few experimental studies on the piriform cortex *in vivo*, but one study showed that pyramidal neurons in the ventral region of the piriform cortex have




Figure 6.15: Response of modeled piriform cortex pyramidal neurons to background spiking input from the olfactory bulb. Each small square represents the membrane potential of one pyramidal neuron recorded at the soma. The color bar represents the mapping between color and membrane potentials (in mV). This data is from a medium-scale model (400 pyramidal neurons) measured 250 $msec$ after the beginning of the simulation. Cortical orientation is as indicated.

resting potentials in the range of -55 to -50 mV [25] which is consistent with our model.

Strong shock response

Strong electric shocks of brief duration (< 1 $msec$) applied to the LOT give rise to a characteristic response in the piriform cortex as measured by the current-source density (CSD) response. Figures 6.16 and 6.17 show these for the model and the real system in posterior piriform cortex. As can be seen, in both cases an initial current source is seen

in layer 1a in response to afferent inputs, with a second (slightly larger) source in layer 1b in response to associational inputs following about 10 *msec* later. The current sources are balanced by current sinks in layers 2 and 3. Figure 6.18 shows a sequence of activations of the pyramidal neuron cell bodies during the course of the strong shock response. As in the real system [19], a large proportion of neurons in the ventral anterior piriform cortex spike, with little spiking in the rest of the cortex.

Weak shock response

In contrast to the strong shock response, weak electrical shocks elicit oscillatory patterns of activity as seen in surface field potentials (figure 6.19) in both the model and in experimental recordings. The field potential oscillates at roughly 50 *Hz* and damps out after several cycles. The exact sequence of cellular events during the weak shock response is not known experimentally, but in the model it is due to a succession of cycles of activity which sweep across the network, activating pyramidal neurons and interneurons in succession (figure 6.20). Current-source density plots have also been recorded from the model; these are discussed extensively in the next chapter, as they gave rise to numerous unexpected insights.

6.2.6 Comparison with the Wilson/Bower model

The model described here is a second-generation version of the Wilson/Bower (WB) model of piriform cortex [36]. The new model improves on the realism of the WB model in the

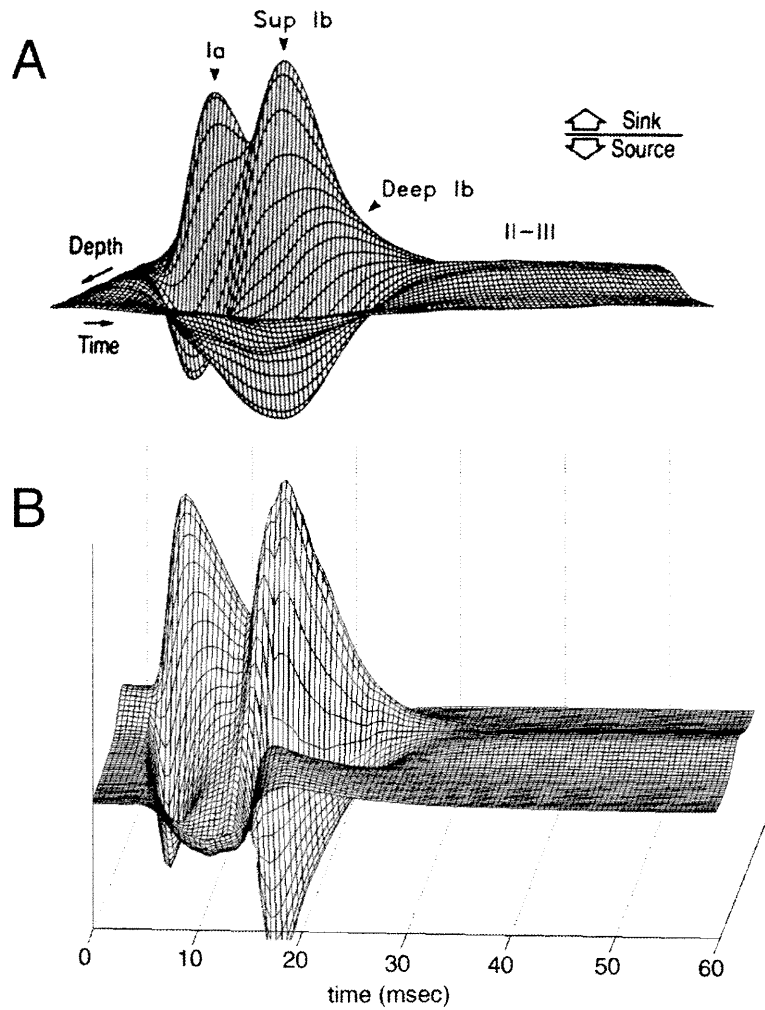


Figure 6.16: Surface CSD plots of the strong shock response in piriform cortex recorded experimentally (A) (from [19]) and in the piriform cortex model (B). The time axis applies to both figures. The shock occurred at time = 0. Direction of current flow (source/sink) is the same for both figures.

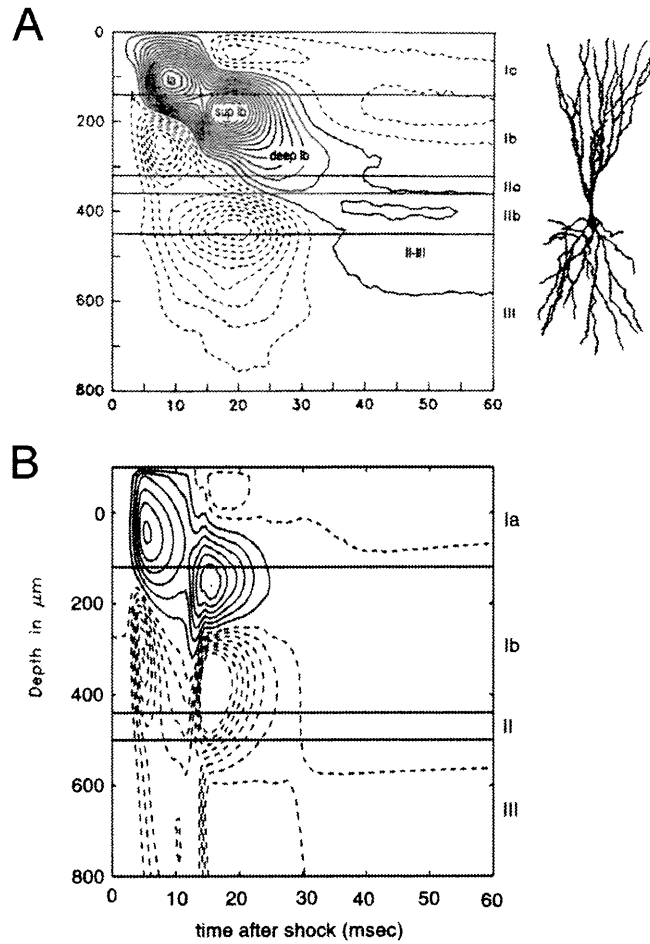


Figure 6.17: Contour CSD plots of the strong shock response in piriform cortex recorded experimentally (A) (from [19]) and in the piriform cortex model (B). The shock occurred at time = 0. Contour lines are iso-CSD contours. Continuous lines represent current sinks while dashed lines represent current sources. The pyramidal neuron figure to the right of the upper plot represents the orientation of a typical superficial pyramidal neuron at the scale of the plot. Piriform cortex layers are listed to the right of each plot.




Figure 6.18: Pyramidal neuron soma responses to a strong shock input in the piriform cortex model. Each colored square represents the soma of a single pyramidal neuron during a strong shock stimulus. The color bar represents the mapping between color and membrane potentials (in mV). This data is from a medium-scale model (400 pyramidal neurons). Each frame represents a specific time (in $msec$) after the shock (indicated at the left of each frame). Cortical orientation is as indicated.




Figure 6.19: Surface field potential response of piriform cortex to a weak shock stimulus. Trace (A) represents experimental data from cat, from [8]. Weak shock stimuli applied to rat piriform cortex gives rise to similar field potentials [13]. Trace (B) represents output from the piriform cortex model. Both horizontal bars represent time intervals of 40 *msec*.




Figure 6.20: Pyramidal neuron soma responses to a weak shock input in the piriform cortex model. Figure layout is the same as in figure 6.18. This data is from a medium-scale model (400 pyramidal neurons). Each frame represents a specific time (in *msec*) after the weak shock (indicated at the left of each frame). The figure represents one wave of activation spreading across the cortex. Cortical orientation is as indicated.

following ways:

1. The pyramidal and feedback inhibitory interneurons in the new model were parameterized to accurately replicate intracellular experimental data.
2. The inputs to the new model were based on the statistical properties of mitral cell recordings from the olfactory bulb.
3. The new model incorporates neuromodulation at many different levels (see above).
4. The WB model divided piriform cortex into anterior and posterior halves. The new model, based on more recent data, divides the model into three regions (ventral anterior, dorsal anterior, and posterior).
5. The pyramidal neurons in the WB model had five compartments and were not electrotonically correct (compartments less than 0.1λ (space constants) in length; see chapter 5), making it difficult to accurately replicate the electrotonic properties of the neuron. The size of compartments in the new model is less than 0.1λ .
6. The WB model used explicit integration methods and a large time step (0.1 msec). For the new model, we used implicit integration methods and a $20 \mu\text{sec}$ time step. We believe that neurons in the present model are consequently much more accurately simulated.

The net effect of these changes is that the new model is a considerably more accurate representation of the piriform cortex than the WB model. In chapter 7, we discuss the functional implications of these changes.

6.2.7 Limitations of the model

Although the model of piriform cortex described here is the most accurate model of this brain structure ever constructed, there are several limitations of the model which may be computationally significant, and which will need to be removed in future models of this system.

The most significant omission of the current model is that the modeled inputs from olfactory bulb mitral cells are based only on the first-order statistics of spike trains from these neurons. A more accurate model would account for both higher-order statistical properties of the input spike trains and also for the statistical relationship between spike trains of different mitral cells. In addition, the model should incorporate data on a large number of odorants and a large number of odor concentrations. Clearly, this will require a considerable amount of experimental data collection and data analysis, but it will be essential in order to provide the cortical model with a reasonable representation of natural inputs. Similarly, since the piriform cortex and the olfactory bulb are in fact tightly coupled in the brain [10, 11, 13], it will be necessary to build combined bulb/cortex models to understand many dynamical phenomena, such as how feedback from the cortex affects

inputs from the bulb.

Several neuron types were omitted from the present model; these principally include deep pyramidal neurons and spiny multipolar neurons. This omission was primarily due to a lack of experimental data (both intracellular data and data on network connectivities), and also because these neuronal types are considerably less numerous than the neurons modeled. There is some data on deep pyramidal neurons [32, 33], but we made the approximation that these neurons are not significantly different functionally from the superficial pyramidal neurons in the model.

The models for feedforward neurons were essentially unconstrained by experimental data. This is unfortunate because, as will be seen in the next chapter, these neurons have a large effect on the dynamics of the system in the model. The model has therefore informed us that these neurons must be characterized experimentally in order to take the model to the next level of realism, and this experimental characterization is in progress in our lab.

Many phenomena observed in the model may differ from those in the real system due to network scaling effects. Ideally, we would like to build much larger models, especially in terms of the number of pyramidal neurons simulated. As computers increase in speed, this will become increasingly feasible.

Appendix A: Data analysis for mitral cell recordings

In order to characterize the firing behaviors of olfactory bulb mitral cells in response to odors, we performed a number of statistical analyses which are described in this section.

ISI analysis

When the interspike interval (ISI) distributions from individual mitral cells was examined for changes in response to different odors, no significant changes were seen in any of the neurons (figure 6.21). When data from all odors was pooled for individual neurons, different neurons generally displayed broadly similar ISI distributions, with the mode of the distribution located at about 20 *msec*, corresponding to the gamma frequency range (50 *Hz*) [6] (figure 6.22). A small subset of neurons (3) had the mode located at a lower ISI, corresponding to a higher overall spike rate, while one neuron had the mode located at a higher ISI (not shown). The average firing rate of the mitral cells was approximately 7 *Hz*; this also did not vary appreciably between odors. The pooled ISI data for all neurons and all odors (except for the few neurons with skewed ISI distributions) is shown in figure 6.23. This was the ISI distribution that was used in the olfactory bulb spike generating model.

Autocorrelation analysis

Given that the olfactory bulb exhibits field potential oscillations in the gamma range (35-85 *Hz*) [6], an important question is whether individual mitral cells also spike at this rate.




Figure 6.21: Interspike interval (ISI) distributions of the responses of a single mitral cell to different odors. The black line represents background responses, the red line represents citral, the blue line represents isoamyl acetate, and the green line represents methyl salicylate. ISI distributions have been offset by 0.05 on the *y*-axis for clarity. All ISI distributions are essentially identical.




Figure 6.22: Interspike interval distributions of the responses of 18 mitral cells.

All responses (background and the three odors) have been pooled for each cell.




Figure 6.23: Pooled interspike interval distribution of all mitral cells across all odors.

Given the low average spiking rate, it is clear that most neurons do not; however, since the ISI distribution shows a peak at about 50 *Hz* it remains possible that a subset of neurons fired at this rate for some period of time. To ascertain this we computed the autocorrelation function for the data sets [1]. If a mitral cell fired repeatedly every 20 *msec* (50 *Hz*), then peaks should be found in the autocorrelogram at 20 *msec*, 40 *msec*, etc. This would not show up in the ISI distribution because each interspike interval would be the same.

The autocorrelation algorithm we used was optimized for spike data as follows. For each neuron, each spike from a given trial was taken and compared to all subsequent spikes on that trial. For each such pair, the time difference was calculated and recorded, and the data was binned to produce a (one-sided) autocorrelogram. Our data sets were generally 5 seconds in duration [4] but we were only interested in the autocorrelogram for the first second after odor onset (or after odor purge for background odors [4]) because this is the time in which odor identification is believed to occur [7, 23].

The results show that different odors did not affect the autocorrelogram of a given neuron to any significant degree at the concentrations used (figure 6.24). In addition, although nearly all autocorrelograms showed a peak at around 20 *msec*, there was no evidence whatsoever for peaks at larger multiples of 20 *msec*. This result is interesting in that it suggests that different neurons are involved in different 50 *Hz* cycles, and that a given neuron rarely fires on more than two consecutive cycles. Different neurons generally exhibited broadly similar autocorrelograms, with a sharp peak at 20 *msec* and a slow

falloff (figure 6.25, black traces), although some neurons showed higher short-term and/or lower longer-term autocorrelations (figure 6.25, red traces). The autocorrelogram for all the combined data is shown in figure 6.6.

PSTH analysis

An important question when analyzing the response of any neuron to a stimulus is whether its average firing rate changes in response to the input stimulus. This is measured by peri-stimulus time histograms (PSTH) which bin spikes on repeated trials based on the time after the stimulus onset that the spikes occurred. For mitral cells the situation is complicated by the fact that rats generally sniff rapidly in the presence of an odor [4] and the firing rate may be modulated by the sniffing rate [4]. Thus, we computed the PSTHs of the mitral cells in response to all three odors under two conditions: locked to odor onset, and locked to the first inspiration following odor onset.

Although the small data set makes it difficult to draw firm conclusions, we could not detect any consistent differences in the PSTH response to individual odors when not synchronized to respiration (figure 6.26) or when synchronized to respiration (figure 6.27). In fact, most neurons appeared to have approximately a flat PSTH (no change in rate in response to the odor) for all odors. This was surprising to us since other authors have described a variety of rate changes in mitral cells in response to odors [23, 35]. A typical example of this is an initial increase in rate followed by a decrease; we observed this in only




Figure 6.24: Autocorrelograms of two mitral cells during presentation of several odors. Black traces represent autocorrelograms from one mitral cell while red traces represent autocorrelograms from a different cell. For each cell, four different responses are plotted: background response and response to each of three odors. Note that autocorrelograms from different odors are indistinguishable except for slight random fluctuations.




Figure 6.25: Autocorrelograms of different mitral cells. Results for background activity and for responses to different odors were pooled for each cell. The black traces represent typical autocorrelograms (19 cells). Red traces represent atypical autocorrelograms (6 cells).

two neurons out of 44 (*e.g.*, figure 6.28). Interestingly, this change was only observed in the PSTHs synchronized to respiration, and was observed for all three odors. The rate change was slightly more than twice the background firing rate. We believe that the difference between these results and those of other groups that have found rate increases in response to odors is largely due to the low concentrations used; when high concentrations are used [35], significant rate changes are routinely seen.

Appendix B: Extracellular field potential measurement and current-source density (CSD) calculation

The network model was designed to output extracellular field potentials from any set of arbitrary locations in the cortex. These field potential measurements can be processed offline to generate current-source density (CSD) plots, which give a graphical representation of the density of current sources and sinks at different depths of the cortex (in a particular x - y location) as a function of time. These outputs were used extensively in parameterizing the behaviors of the cortical model. In this section we briefly describe the theory behind these measurements and their implementation in the cortical model.

Extracellular field potentials

When an electrode is placed into the extracellular medium in a cortical structure, the voltage measured at the electrode tip (which is referred to as the extracellular field potential) is




Figure 6.26: PSTHs of several mitral cells synchronized to odor onset but not to respiration onset. Each PSTH is a single connected line on the graph representing the pooled response of a single mitral cell to a single odor. Histogram counts have been normalized to an average of 1.0. The x axis represents time after odor onset. Note the absence of any obvious structure in the PSTHs.




Figure 6.27: PSTHs of mitral cells synchronized to respiration onset following odor onset. Each PSTH is a single connected line on the graph representing the pooled response of a single mitral cell to a single odor. Histogram counts have been normalized to an average of 1.0. The x axis represents time after odor onset. Again, there is no obvious structure in the PSTHs.




Figure 6.28: PSTHs of a mitral cell showing firing rate changes after odor onset. PSTH data from each of three odors was pooled. Each trace represents a single odor. The x axis represents time after odor onset.

the result of the superposition of the effects of a large number of current sources and sinks propagating through a resistive medium (figures 6.13 and 6.14). Current sources represent current flowing from the intracellular to the extracellular medium, while current sinks represent current flowing in the reverse direction. These current sources/sinks are due to synaptic channels, voltage-dependent and calcium-dependent ionic channels, and capacitive and resistive membrane currents [19, 20, 24, 28]. The value of the field potential depends on the number, strength, location and polarity of current sources near the recording electrode, as well as the extracellular resistivity.

By Kirchoff's law, the total amount of current flowing into a neuron (current sinks) equals the total current flowing out (current sources) (figure 6.13). If sources and sinks are distributed evenly over the neuron, the measured field potential will be very small. Therefore, field potential measurements are only useful when there is a spatial separation between sources and sinks for a large number of neurons and a given input stimulus. This is the case for piriform cortex pyramidal neurons when voltage shock stimuli are applied to the LOT, because the temporal pattern of activation of synapses on the neuronal dendrites are very different in different layers of cortex [19, 28]. Synapses in particular layers are stimulated nearly synchronously, giving rise simultaneously to large numbers of current sources in some layers and large numbers of current sinks in other layers. Typically, the sinks are localized in a narrow region of dendrite where synaptic input arrives, while the sources are distributed over a broader expanse of dendrite (representing passive resistive

and capacitive equalizing currents) (figure 6.13).

The field potential in a non-capacitive extracellular medium of uniform resistivity driven by point current sources is given by [26]:

$$V_{fp}(x, y, z, t) = \frac{R_e}{4\pi} \sum_i^{n_{sources}} \frac{I_m(x_i, y_i, z_i, t)}{d_i}, \quad (6.1)$$

where V_{fp} is the field potential, R_e is the extracellular resistivity (in units of $\Omega \cdot m$), $n_{sources}$ is the total number of current sources, I_m are the current sources, and

$$d_i(x, y, z, x_i, y_i, z_i) = \sqrt{(x - x_i)^2 + (y - y_i)^2 + (z - z_i)^2} \quad (6.2)$$

represents the Euclidean distance between the source (at (x_i, y_i, z_i)) and the electrode tip (at (x, y, z)). Compartment objects in GENESIS compute the value of I_m due to all current sources associated with the compartment. Virtual electrodes can be placed at arbitrary points in the simulated cortex, and the field potentials are calculated at each time step according to equation 6.1. In practice, two types of field potentials were computed: surface field potentials at the top of the simulated cortex, and field potentials at varying depths in cortex. The latter were used to compute current-source densities as we describe next.

Current-source densities

Although some insights can be gained directly from examining extracellular field potentials, in certain circumstances, the data can be processed to give what are known as current-

source density (CSD) profiles. These profiles give a graphical representation of the local density of current sources and sinks in a cortical network, and allow a much more detailed understanding of the synaptic interactions in the network. CSD profiles are also much easier to interpret than raw field potentials.

Measuring CSDs experimentally consists of recording extracellular field potentials at a series of depths at a particular x - y position of cortex. Typically, a series of recordings are made at a single depth in response to a single shock stimulus, then the electrode is advanced a short distance deeper and another series of recordings are made, and so on [19, 28]. Field potential responses from each depth are averaged before computing CSDs. This is necessary because second derivatives of the field potentials have to be computed to obtain the CSD (see below), which increases the effect of any noise in the recordings.

Field potential recordings are converted to CSDs using Poisson's equation [24]:

$$\frac{1}{R_e} \nabla^2 \phi = -I_m \quad (6.3)$$

where R_e is the extracellular resistivity as before, $\nabla^2 \phi$ is the spatial Laplacian of the scalar field potential (in units of V/m^2), and I_m is the amplitude of the CSD (in units of A/m^3). A full derivation is included in [24]. Note that equation 6.3 assumes a uniform extracellular resistivity; this assumption has found to be approximately true in piriform cortex [28]. The equation also assumes that Ohm's law is valid in the extracellular space and that the electric field generated by the transmembrane currents is quasistatic. These assumptions have also

been verified in piriform cortex [28].

Because of the laminated structure of piriform cortex and the fact that the cortex is activated homogeneously with respect to the laminar planes (layers 1a, 1b, 2 *etc.*), the field potential is locally translationally invariant in the two directions parallel to the laminar planes (the x and y dimensions in our simulations) [24]. Intuitively, this is because the superposition of currents in these directions from many nearby pyramidal neurons (receiving similar input stimuli) cancel each other out to a large extent, leaving little net current in these directions and thus little or no potential gradient. Under these circumstances, equation 6.3 can be reduced to the one-dimensional Poisson equation:

$$\frac{1}{R_e} \frac{\partial^2 \phi}{\partial z^2} = -I_m \quad (6.4)$$

We used this equation to calculate CSD profiles from field potential measurements taken at depth increments of $20 \mu m$ at a number of different x - y locations evenly spaced in the piriform cortex model. Resistivities are not available for piriform cortex, so the scale of the CSDs is arbitrary. We found that the response of the network model was sufficiently regular that repeated trials and smoothing of the response were not necessary in order to generate CSD profiles.

Bibliography

- [1] M. Abeles. *Corticonics*. Cambridge University Press, Cambridge, 1991.
- [2] J. Ambros-Ingerson, R. Granger, and G. Lynch. Simulation of Paleocortex Performs Hierarchical Clustering. *Science*, 247(4948):1344–1348, 1990.
- [3] E. Barkai, R. E. Bergman, G. Horwitz, and M. E. Hasselmo. Modulation of Associative Memory Function in a Biophysical Simulation of Rat Piriform Cortex. *Journal of Neurophysiology*, 72(2):659–677, 1994.
- [4] U. S. Bhalla and J. M. Bower. Multiday Recordings from Olfactory Bulb Neurons in Awake Freely Moving Rats: Spatially and Temporally Organized Variability in Odorant Response Properties. *Journal of Computational Neuroscience*, 4(3):221–256, 1997.
- [5] S. L. Bressler. Spatial Organization of EEGs from Olfactory Bulb and Cortex. *Electroencephalography and Clinical Neurophysiology*, 57:270–276, 1984.

- [6] S. L. Bressler and W. J. Freeman. Frequency Analysis of Olfactory System EEG in Cat, Rabbit and Rat. *Electroencephalography and Clinical Neurophysiology*, 50:19–24, 1980.
- [7] W. J. Freeman. Personal communication.
- [8] W. J. Freeman. The Electrical Activity of a Primary Sensory Cortex: Analysis of EEG Waves. *International Review of Neurobiology*, 5:53–119, 1963.
- [9] W. J. Freeman. A Linear Distributed Feedback Model for Prepyriform Cortex. *Experimental Neurology*, 10:525–547, 1964.
- [10] W. J. Freeman. *Mass Action in the Nervous System*. Academic Press, New York, 1975.
- [11] W. J. Freeman. Simulation of Chaotic EEG Patterns with a Dynamic Model of the Olfactory System. *Biological Cybernetics*, 56:139–150, 1987.
- [12] L. B. Haberly. Structure of the Piriform Cortex of the Opossum. 1. Description of Neuron Types with Golgi Methods. *Journal of Comparative Neurology*, 213(2):163–187, 1983.
- [13] L. B. Haberly. Olfactory Cortex. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.
- [14] M. E. Hasselmo. Acetylcholine and Learning in a Cortical Associative Memory. *Neural Computation*, 5(1):32–44, 1993.

- [15] M. E. Hasselmo, B. P. Anderson, and J. M. Bower. Cholinergic Modulation of Cortical Associative Memory Function. *Journal of Neurophysiology*, 67(5):1230–1246, 1992.
- [16] M. E. Hasselmo, C. Linster, M. Patil, D. Ma, and M. Cekic. Noradrenergic Suppression of Synaptic Transmission May Influence Cortical Signal-to-Noise Ratio. *Journal of Neurophysiology*, 77(6):3326–3339, 1997.
- [17] M. Jiang, E. R. Griff, M. Ennis, L. A. Zimmer, and M. T. Shipley. Activation of Locus Coeruleus Enhances the Responses of Olfactory Bulb Mitral Cells to Weak Olfactory Nerve Input. *Journal of Neuroscience*, 16(19):6319–6329, 1996.
- [18] A. Kapur, W. W. Lytton, K. L. Ketchum, and L. B. Haberly. Regulation of the NMDA Component of EPSPs by Different Components of Postsynaptic GABAergic Inhibition: Computer Simulation Analysis in Piriform Cortex. *Journal of Neurophysiology*, 78(5):2546–2559, 1997.
- [19] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 1. Current Source-Density Analysis. *Journal of Neurophysiology*, 69(1):248–260, 1993.
- [20] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 2. Analysis with a System Model. *Journal of Neurophysiology*, 69(1):261–281, 1993.

- [21] K. L. Ketchum and L. B. Haberly. Synaptic Events that Generate Fast Oscillations in Piriform Cortex. *Journal of Neuroscience*, 13(9):3980–3985, 1993.
- [22] G. Laurent and H. Davidowitz. Encoding of Olfactory Information with Oscillating Neural Assemblies. *Science*, 265:1872–1875, 1994.
- [23] Kay L. M. and Laurent G. Odor- and Context-Dependent Modulation of Mitral Cell Activity in Behaving Rats. *Nature Neuroscience*, 2(11):1003–1009, 1999.
- [24] U. Mitzdorf. Current Source-Density Method and Application in Cat Cerebral-Cortex – Investigation of Evoked-Potentials and EEG Phenomena. *Physiological Reviews*, 65(1):37–100, 1985.
- [25] J. W. Nemitz and S. J. Goldberg. Neuronal Responses of Rat Piriform Cortex to Odor Stimulation – An Extracellular and Intracellular Study. *Journal of Neurophysiology*, 49(1):188–203, 1983.
- [26] P. L. Nunez. *Electric Fields of the Brain: the Neurophysics of EEG*. Oxford University Press, New York, 1981.
- [27] G. Paxinos and C. Watson. *The Rat Brain in Stereotaxic Coordinates*. Academic Press, San Diego, 1997.

- [28] R. Rodriguez and L. B. Haberly. Analysis of Synaptic Events in the Opossum Piriform Cortex with Improved Current Source Density Techniques. *Journal of Neurophysiology*, 61(4):702–718, 1989.
- [29] G. R. Shepherd, editor. *The Synaptic Organization of the Brain*, 4th. Ed. Oxford University Press, New York, 1997.
- [30] G. R. Shepherd and C. A. Greer. Olfactory Bulb. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.
- [31] C. Skarda and W. J. Freeman. How Brains Make Chaos in Order to Make Sense of the World. *Behavioral and Brain Sciences*, 10:161–195, 1987.
- [32] G. Tseng and L. B. Haberly. Deep Neurons in Piriform Cortex. 1. Morphology and Synaptically Evoked-Responses Including a Unique High-Amplitude Paired Shock Facilitation. *Journal of Neurophysiology*, 62(2):369–385, 1989.
- [33] G. Tseng and L. B. Haberly. Deep Neurons in Piriform Cortex. 2. Membrane Properties that Underlie Unusual Synaptic Responses. *Journal of Neurophysiology*, 62(2):386–400, 1989.
- [34] M. C. Vanier and J. M. Bower. A Comparative Survey of Automated Parameter-Search Methods for Compartmental Neural Models. *Journal of Computational Neuroscience*, 7(2):149–171, 1999.

- [35] D. A. Wilson. Habituation of Odor Responses in the Rat Anterior Piriform Cortex. *Journal of Neurophysiology*, 79(3):1425–1440, 1998.
- [36] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.

Chapter 7

Exploring the Network Model of Piriform Cortex

The network model of piriform cortex described in the previous chapter has been used to explore the dynamics of this system in a variety of ways. Some of these explorations will be described in this chapter. These explorations form only a tiny fraction of all the possible investigations that the model makes possible; nevertheless, they have already substantially increased our understanding of several aspects of the system dynamics, suggested numerous new experiments, and given rise to new ideas regarding how the piriform cortex may process information.

First, the response of the piriform cortex model to background inputs from the olfactory bulb and to strong electrical shock stimuli of the lateral olfactory tract (LOT) will be

described. This has suggested that background inputs are essential for depolarizing neurons in the ventral anterior region of cortex to the point where they fire readily in response to input stimuli. I will also show that background inputs need to be substantially reduced in the presence of norepinephrine, or else the cortical model will spontaneously generate large numbers of spikes. The strong shock response in the model was found to depend critically on both background inputs from the olfactory bulb and the presence of strong feedforward inhibition to prevent neural activity from propagating out of the ventral anterior piriform cortex.

I will then contrast two significantly different versions of the piriform cortex model from the standpoint of their ability to replicate the experimental response to weak electrical shocks to the LOT (the weak shock response). The first version of the model features random connectivity between the olfactory bulb and cortex, as well as strong feedback connections between posterior and anterior cortex. The connectivity pattern of this model is similar to that of the Wilson/Bower piriform cortex model [21]. This version of the new model replicates the oscillatory surface field potential, but fails to replicate crucial features of the current-source density (CSD) response. The second version of the model, in contrast, features highly structured connections between the olfactory bulb and piriform cortex and within the cortex; this model replicates the CSD response much more accurately. While highly speculative in some ways, the latter model casts a completely new light on how piriform cortex may encode odors and may function computationally, and is experimentally

testable.

7.1 Response to olfactory bulb background inputs

7.1.1 Spontaneous activity

When the piriform cortex model is subjected to background spiking activity with a statistical distribution similar to that of olfactory bulb mitral cells (see chapter 6 for details), the membrane potentials of pyramidal neurons in the ventral anterior piriform cortex (vAPC) rise close to spiking threshold, whereas the membrane potentials of pyramidal neurons in the posterior piriform cortex (PPC) remain close to resting potential (figure 7.1). The reason for the difference is the much denser projection from mitral cells to pyramidal neurons in vAPC as compared to the rest of piriform cortex [5]. Experimentally, piriform cortex pyramidal neurons *in vivo* have been recorded with membrane potentials in the range of the vAPC neurons in the model ($-55mV$ to $-50mV$) [14]; as in the model, these neurons were located in the ventral half of the piriform cortex. Pyramidal neurons *in vivo* also often spike spontaneously at roughly 1-5 *Hz* [14, 16, 17]; however, in the model background input is not sufficient to cause spiking in pyramidal neurons. This is because in the model the pyramidal neurons are sufficiently close to firing threshold that if one neuron fires, typically all of the neurons fire in a wave-like pattern which spreads across the cortex (see below). The discrepancy between the model and the experimental data is due to scaling effects; each




Figure 7.1: Response of pyramidal neurons to background spiking input from the olfactory bulb in the piriform cortex model. Each small colored square represents the membrane potential of one pyramidal neuron recorded at the soma. The color bar represents the mapping between color and membrane potentials (in mV). This data is from a medium-scale model (400 pyramidal neurons) measured 250 $msec$ after the beginning of the simulation. Cortical orientation is as indicated.

neuron in the model takes the place of over 1000 real neurons (piriform cortex consists of over 10^6 pyramidal neurons [5]), so in effect one model neuron firing is equivalent to over 1000 neurons firing synchronously in the real cortex. A very large scale model (10,000 to 100,000 pyramidal neurons) with very sparse connections might replicate the low background spiking activity of pyramidal neurons; however, such a model is beyond the computational resources presently at our disposal.

7.1.2 Spontaneous activity with NE

Experimentally, addition of $100 \mu M$ NE to a brain-slice preparation of piriform cortex causes a number of changes in synaptic transmission and neuronal excitability (described in detail in chapters 2, 5, and 6). Briefly, the most important of these effects include

1. increasing the strength of synaptic transmission onto piriform cortex pyramidal neurons along the afferent (layer 1a) pathway;
2. decreasing the strength of synaptic transmission between pyramidal neurons along the associational (layer 1b) pathway);
3. increasing the excitability of pyramidal neurons and feedback inhibitory interneurons.

These effects were incorporated into the piriform cortex network model. When a concentration of $100 \mu M$ NE was imposed on the model in the presence of random background inputs, the pyramidal neurons began to spike spontaneously at roughly $10 Hz$ (figure 7.2, top). Although the spontaneous spiking rate of pyramidal neurons in piriform cortex under the influence of NE is not known, it is extremely unlikely that these neurons fire at such a high spontaneous rate *in vivo* since most of the existing experimental data indicates that piriform cortical neurons generally have low spiking rates ($< 20 Hz$) even during odor presentation [13, 14]. In order to prevent NE application from causing pyramidal neurons to fire spontaneously in response to background inputs, it was necessary to reduce the background input firing rate (*i.e.*, the background firing rate of mitral cells) to 40% of the rate

in the absence of NE (figure 7.2, bottom). It has been determined experimentally that NE reduces background firing of mitral cells by roughly 50% *in vivo* [8], which agrees well with the model.

7.2 The strong shock response

7.2.1 Basic features of the response

Experimentally, stimulation of the lateral olfactory tract (LOT) with very brief (0.1 *msec*) electrical shocks causes characteristic responses in piriform cortex as seen in current-source density (CSD) profiles. These shocks have been used extensively by Lewis Haberly and his collaborators [9, 10, 11] and Walter Freeman [4] to characterize the dynamical behaviors of the piriform cortex network. These behaviors are qualitatively different for strong (high voltage) and weak (low voltage) electrical shocks. A typical experimental strong shock response is shown in the top half of figures 7.3 and 7.4 [9, 10]. The response consists of two separate peaks occurring after the shock stimulus. The first peak is due to direct afferent stimulation in layer 1a of the piriform cortex, while the second is believed to be due to associational fiber activity in layer 1b, which in turn is a consequence of pyramidal neuron firing in ventral anterior piriform cortex [5, 10]. Few neurons fire in response to the strong shock in dorsal anterior or posterior piriform cortex [9].

In the piriform cortex model, the strong shock is simulated by making the entire pop-

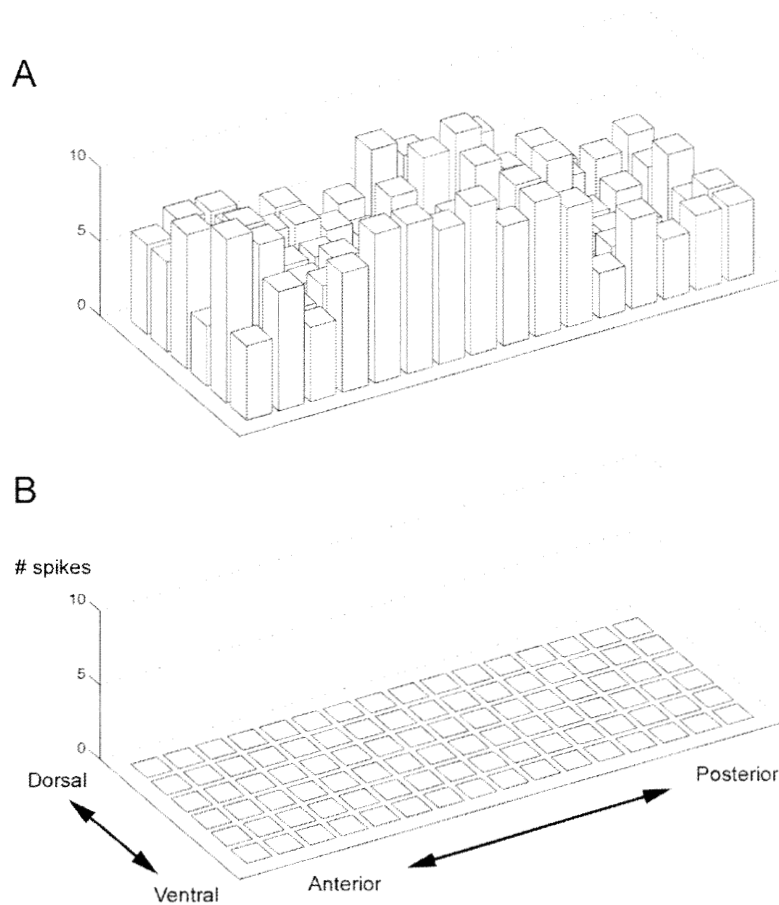


Figure 7.2: 2-D histograms of pyramidal neuron background firing rates in the piriform cortex model under the influence of $100 \mu M$ NE. Each bin represents the total number of spikes in one pyramidal neuron during a 500 msec simulation. **A.** Without reduction in the mitral cell background firing rate. **B.** With a 40% reduction in mitral cell background firing rate. This data is from a small scale (96 pyramidal neuron) model.

ulation of mitral cell inputs to the cortex spike simultaneously. The network response to this input is shown in figure 7.5. The model faithfully reproduces the experimental results (figures 7.3 and 7.4) and gives insight into possible mechanisms underlying them. The shock causes a large EPSP in the distal dendrites of pyramidal neurons, which can be seen in the CSD response centered at the level of layer 1a. This in turn leads to the firing of pyramidal neurons primarily in the ventral anterior piriform cortex (whose membrane potential is higher than in other regions of cortex, due to the influence of mitral cell background firing as discussed previously) (figure 7.5). This leads to a second peak in the CSD response centered at the level of layer 1b. In the model, the two peaks can also be seen in intracellular traces in pyramidal cell dendrites in layer 1b (figure 7.6); the blue trace (from a pyramidal neuron in ventral anterior piriform cortex) shows one large peak occurring shortly after the shock input, while the red trace (from a pyramidal neuron in posterior piriform cortex) also shows a second peak due to associational activity. These behaviors of the model are consistent with available experimental data [9].

7.2.2 Role of background excitation in the strong shock response

An unexpected insight derived from these simulations concerns the significance of the background spiking inputs from olfactory bulb mitral cells in the genesis of the strong shock response. As we have seen, background inputs in the network model cause the cell bodies of neurons in the ventral anterior piriform cortex to remain very close to spiking threshold.

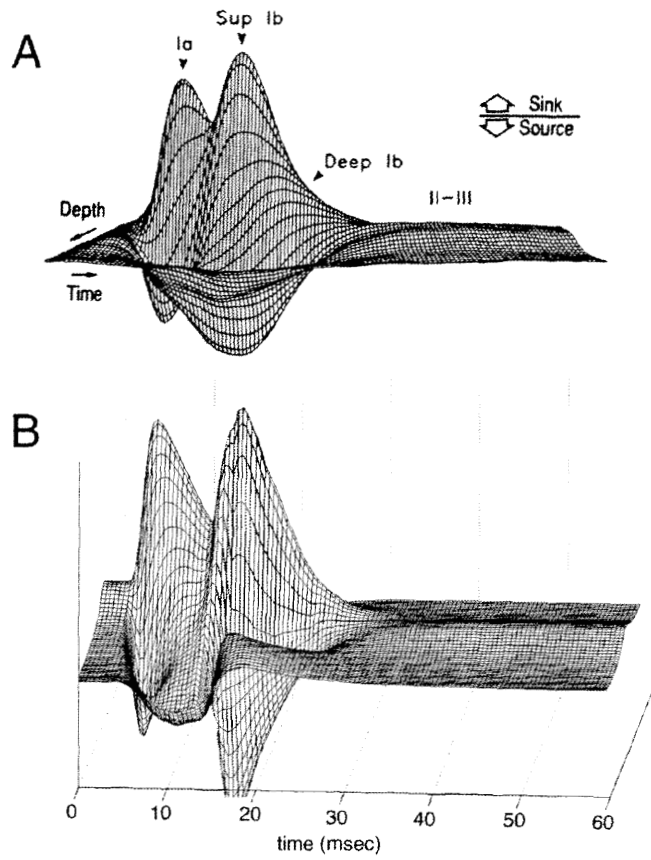


Figure 7.3: CSD surface profile of the response of the piriform cortex to a strong shock stimulus. **A.** Experimental results from [9]. **B.** Results from the piriform cortex model. The shock occurred at time = 0. The time base and direction of current flow are the same for both figures.





Figure 7.4: (Previous page) CSD contour profile of the response of the piriform cortex to a strong shock stimulus. **A.** Experimental results from [9]. **B.** Results from the piriform cortex model. Contour lines are iso-CSD contours. Continuous lines represent current sinks while dashed lines represent current sources. The shock occurred at time = 0. The pyramidal neuron figure to the right of the upper plot represents the orientation of a typical superficial pyramidal neuron at the scale of the plot. Piriform cortex layers are listed to the right of each figure. In the model, the layer 2a region is considered to be part of superficial layer 1b; layer 2 only contains the pyramidal neuron cell bodies.

The effect of this is to minimize the delay between the shock and the spiking of pyramidal neurons in the ventral anterior piriform cortex. This delay (caused by passive propagation of the afferent input down the pyramidal neuron apical dendrite) is still substantial; in the model the shock occurs at 75 *msec* after the start of the simulation, and the first pyramidal neuron spike typically occurs at roughly 84 *msec*. Since the magnitude of the afferent weights from mitral cells to pyramidal neurons is a free parameter of the model, the possibility existed that increased afferent weights could substitute for the lack of background synaptic input. This was not found to be the case (figure 7.7). When background input is removed from the model, the afferent weights needed to be at least twice as large in order to cause a comparable amount of spiking in the pyramidal neurons in response to a strong shock. Even given this large increase in afferent weights, the latency to first spike is still much larger when compared to the shock stimulus with background input (18 *msec*

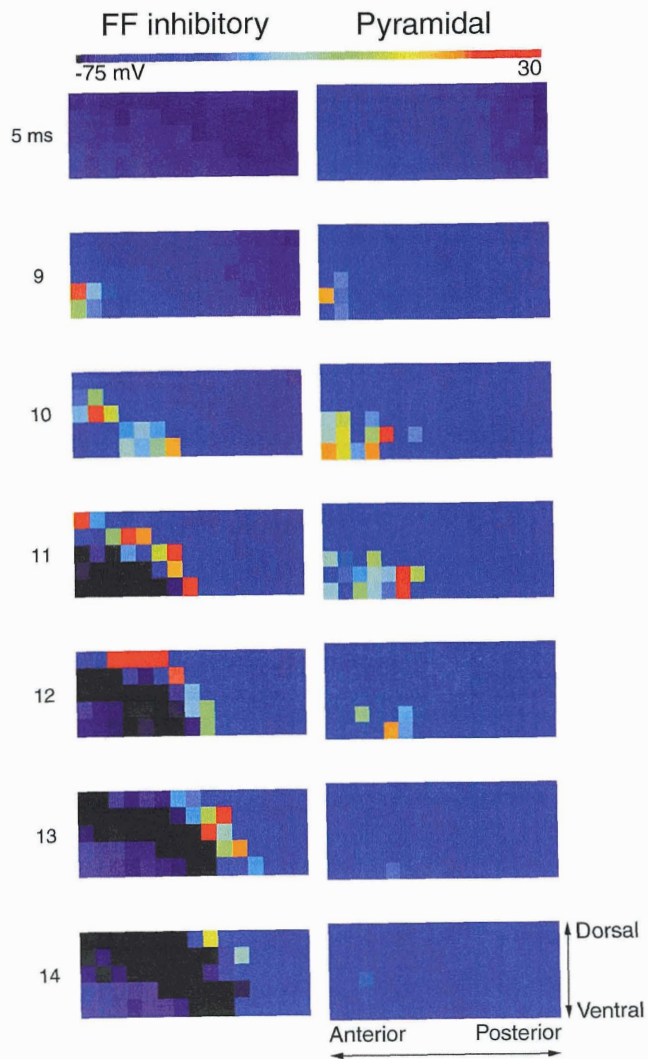



Figure 7.5: (Previous page) Response of pyramidal neurons (measured at the soma) and feedforward (FF) inhibitory interneurons in the piriform cortex model to a strong shock stimulus. This data is from a small-scale model (96 pyramidal neurons). The color bar represents the membrane potential in mV . Each horizontal pair of frames represents a specific time (in $msec$) after the shock stimulus (indicated at the left of each pair of frames). Note the spreading of activation among the feedforward inhibitory interneurons; spikes in these neurons prevent spiking in pyramidal neurons outside of the ventral anterior piriform cortex. The schematic diagram at the top of the page represents afferent connections to feedforward interneurons (left) and pyramidal neurons (right), as well as the inhibitory connections between feedforward interneurons and pyramidal neurons.

vs. 9 $msec$). In figure 7.7, we see the CSD profile for the strong shock with and without background input; without background inputs the two peaks are much farther apart than in the experimental data (or in the model data with background inputs). If the afferent weights in the model are increased still further in order to decrease the latency, pyramidal neurons outside the ventral anterior region begin to spike, particularly in the posterior piriform cortex where feedforward inhibition is relatively weak (figure 7.8); this contradicts the experimental data [9].

Thus, the model suggests that background spiking inputs from the olfactory bulb are necessary in order to give rise to a strong shock response that accurately reflects the experimental data. Experimentally, it was shown many years ago by Walter Freeman that




Figure 7.6: Response of pyramidal neurons in the piriform cortex model to a strong shock stimulus measured in superficial layer 1b dendrites. Blue trace: recording from a pyramidal neuron in ventral anterior piriform cortex (vapo). Red trace: recording from a pyramidal neuron in posterior piriform cortex (ppc). The shock occurred at time = 0.

cutting the LOT before giving shock stimuli (effectively removing background input) caused the entire piriform cortex to become quiescent [4]; this is consistent with our model. The predecessor of the current model, the Wilson/Bower model [21], did not feature background spiking inputs to piriform cortex. In that model, such inputs were not needed to replicate the strong shock response. However, neuron models in the Wilson/Bower model were not based on experimental data, allowing the use of artificially high resting potentials (-55 mV) which kept the neurons closer to their spiking thresholds than is the case with the current model. In addition, the dendritic structure of the neurons was not constrained to match the cable properties of real pyramidal neurons. In a more realistic setting, background inputs appear to be essential in order to replicate the system dynamics accurately.

7.2.3 Role of feedforward inhibition

The Wilson/Bower model of piriform cortex [21] responded to strong shock inputs with pyramidal neurons firing across most of the extent of piriform cortex. Experimental data unavailable when that model was constructed now suggests that only pyramidal neurons in the ventral anterior cortex spike to any significant extent [9]. Paradoxically, a weak shock gives rise to a much longer-lived oscillatory response which presumably involves pyramidal neuron spiking across the full extent of cortex ([4, 11] and sections 7.3 and 7.4). There is currently no experimental data to explain this difference in response with shock strength. However, in the piriform cortex model, in response to a strong shock stimulus, pyramidal

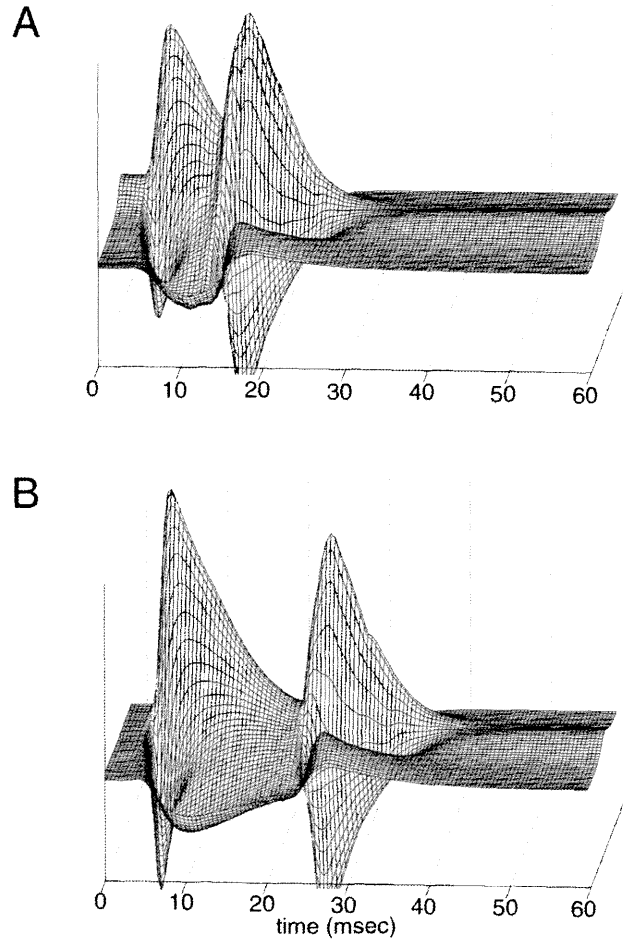


Figure 7.7: Surface CSD profiles of the response to a strong shock in the piriform cortex model in the presence and absence of background spiking inputs from the olfactory bulb. **A.** With background spike inputs from olfactory bulb mitral cells. **B.** Without background inputs. Afferent synaptic strengths in **B** were increased to compensate for the lack of background synaptic activity.

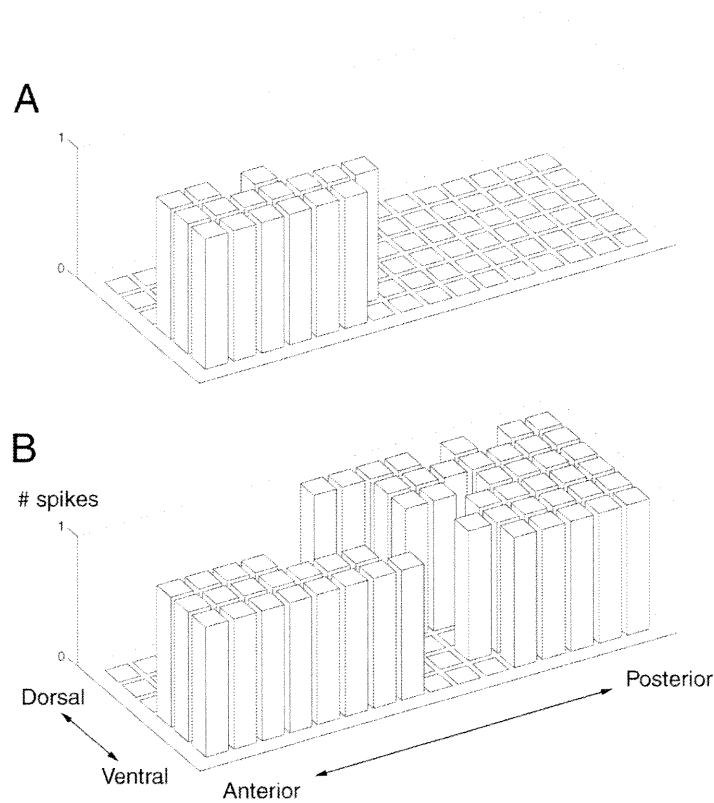


Figure 7.8: 2-D histograms of piriform cortex model pyramidal neuron spikes in response to a strong shock with and without background spiking inputs. Each bin represents the total number of spikes in one pyramidal neuron during a 250 msec simulation. **A.** With background spike inputs from olfactory bulb mitral cells. **B.** Without background inputs. Strength of afferent inputs have been increased to reduce latency to 1b peak, causing neurons outside the ventral anterior piriform cortex to spike. This data is from a small scale (96 pyramidal neuron) model.

neuron spiking in regions other than the ventral anterior piriform cortex (vAPC) is blocked by feedforward inhibition. Feedforward inhibitory neurons across most of the extent of piriform cortex spike in response to the strong shock stimulus (figure 7.5). This spiking occurs too late to prevent pyramidal neuron spiking in the ventral anterior region, but not too late to prevent pyramidal neuron spiking in the rest of cortex. If feedforward inhibition is disabled, pyramidal neurons across the entire cortex spike, and the resulting behavior is similar to the weak shock response (see below). Figures 7.5 and 7.9 show the somatic membrane potentials of all pyramidal neurons in the model during a strong shock with and without feedforward inhibition. With feedforward inhibition, only neurons in the ventral anterior region spike, whereas without feedforward inhibition neurons are activated in a wavelike pattern spreading out from the vAPC to all regions of piriform cortex. Purely feedforward interneurons in the model have a high spiking threshold; the effect of this is that they spike readily in response to a strong shock (figure 7.5) but do not spike in response to a weak shock. This is a prediction of the model.

7.3 The weak shock response: model with random connectivities

It has been assumed for many years that both afferent connectivities to the piriform cortex and associational connectivities within the cortex are essentially random based on anatom-




Figure 7.9: Pyramidal neuron soma responses to a strong shock stimulus in the piriform cortex model in the absence of feedforward inhibition. This data is from a small-scale model (96 pyramidal neurons). Each frame represents a specific time (in *msec*) after the shock. The color bar represents the membrane potential in *mV*. Note that pyramidal neuron activation spreads outside the ventral anterior piriform cortex.

ical data showing an absence of topographical organization in the two pathways [7, 12, 18]. A random connectivity pattern was also used in the Wilson/Bower piriform cortex model; however, CSD responses to weak shocks were not measured in that model [21], because the relevant data [11] was not available at the time that model was built.

This section describes the weak shock response in a version of the piriform cortex model which features random connectivity between the olfactory bulb and cortex and strong feedback connections from posterior to anterior piriform cortex. In addition to replicating the strong shock response, this model also reproduces the piriform cortex oscillatory field potential in response to weak shock stimuli. The dependence of this response on several parameter settings will be explored. The CSD response to weak shocks will then be examined. This version of the model does not accurately reproduce this response; the reasons for this will be discussed in some detail, and an alternative model which does match the CSD response will be presented in the next section.

7.3.1 Basic features of the response

Experimentally, stimulation of the lateral olfactory tract (LOT) with a weak electrical shock causes a damped oscillatory response in piriform cortex as seen in the surface field potential (figure 7.10, top) [4, 11]. The oscillatory frequency is in the gamma range (40 to 70 Hz); the precise frequency varies between species [1], but in rats is usually roughly 50 Hz . Several oscillatory cycles occur before the response damps out completely.

In the version of the network model with random connectivity, the weak shock was simulated by making 20% of the afferent fibers from the olfactory bulb spike simultaneously. As in the data, the model produces a damped oscillatory pattern (figure 7.10, bottom) with oscillations in the gamma range (figure 7.11). The network response is shown in figure 7.12 for one cycle of this oscillation. The mechanism of the weak shock response in this model will now be described in some detail.

The synchronized spike inputs from the weak shock stimulus cause a rise in membrane potentials in distal dendrites of pyramidal neurons (layer 1a). These inputs propagate down the dendrite and bring a substantial fraction of the neurons in the ventral anterior piriform cortex above spiking threshold. These neurons have extensive long-range excitatory connections with the rest of piriform cortex which are principally located on superficial layer 1b dendrites [5, 6] (figure 7.13). In addition, short-range excitatory connections onto basal dendrites of pyramidal neurons are also activated. The combined effect of these connections in the model is to cause a wave of spiking which spreads across pyramidal neurons in the cortex (figure 7.12). Pyramidal neurons also activate interneurons mediating feedback inhibition onto other pyramidal neurons (figure 7.14); spiking in inhibitory interneurons in the model lags spiking in pyramidal neurons by no more than a few milliseconds (figure 7.12). When the wave of pyramidal neuron spiking reaches the posterior end of piriform cortex, neurons in the anterior end are re-excited by long-range connections from pyramidal neurons in the posterior piriform cortex which synapse on the apical dendrites of pyramidal

neurons in the anterior piriform cortex in the deep layer 1b region (figure 7.13). This gives rise to the next wave of spiking. Thus, in this model the response of piriform cortex to a weak shock consists of a series of traveling-wave oscillations. These traveling-wave dynamics have been predicted by other models (notably the Wilson/Bower piriform cortex model [21]) and have also been suggested by some experimental studies [11]. In contrast to the strong shock response in this model, the weak shock does not cause activation of purely feedforward inhibitory interneurons. This is due to the more hyperpolarized resting membrane potentials in these neurons in the model (section 7.2.3).

7.3.2 Influence of feedforward/feedback interneurons

The damping out of the surface field potential in this model, as in the Wilson/Bower model [21] is primarily due to feedback projections from pyramidal neurons onto one of the two groups of feedforward interneurons (the feedforward/feedback interneurons described in chapter 6). The presence of this feedback pathway has not been definitively established, but there is indirect evidence for it [5, 6]. In the piriform cortex model, a combination of afferent (feedforward) activation and feedback activation from pyramidal neurons causes these neurons to spike starting at about 50 *msec* after the weak shock stimulus. This spiking continues until the pyramidal neurons cease spiking and the shock response terminates (figure 7.15).

If the synaptic weights of the pyramidal neuron to feedforward/feedback interneuron




Figure 7.10: Surface field potential response of piriform cortex to a weak shock stimulus. Trace (A) represents experimental data from [3]. Trace (B) represents output from the piriform cortex model. Both horizontal bars represent time intervals of 40 *msec*.




Figure 7.11: Power spectrum of the response of the piriform cortex model to a weak shock stimulus. Note the prominent peaks in the theta (5 Hz) and gamma (40 – 50 Hz) ranges.

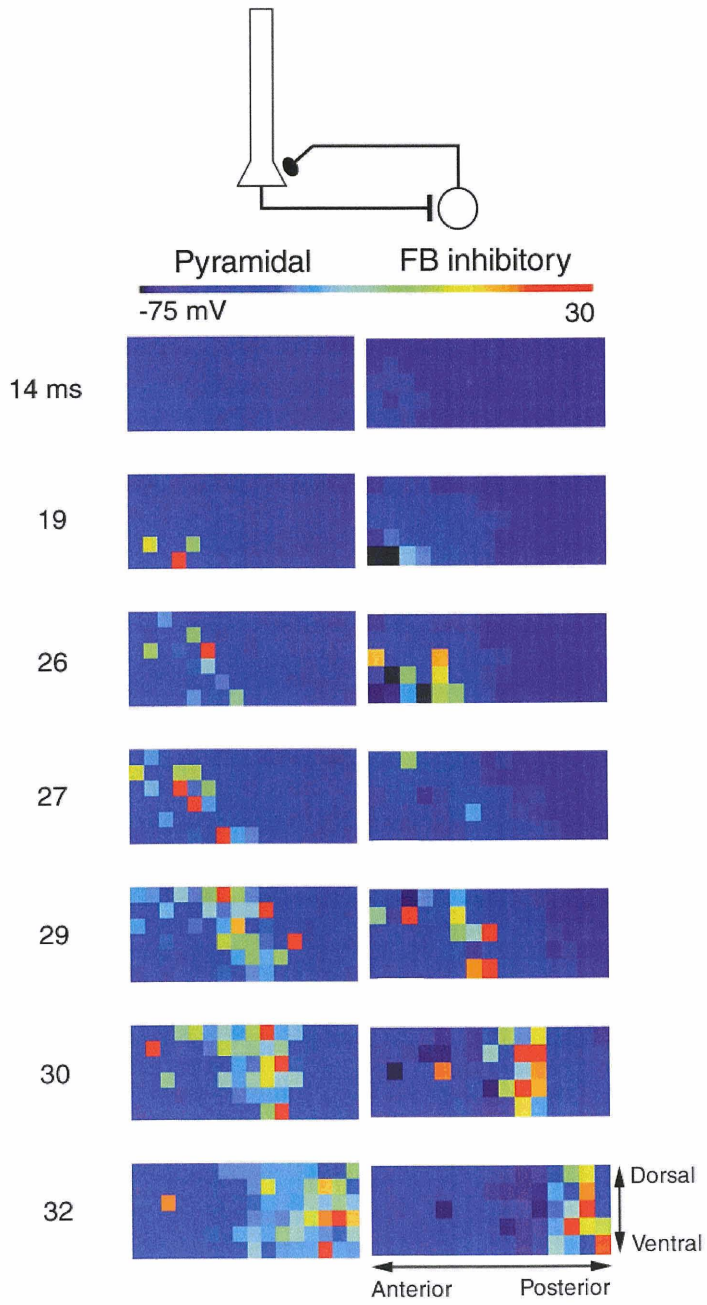


Figure 7.12: (Previous page) Response of pyramidal neurons (measured at the soma) and feedback (FB) inhibitory interneurons in the piriform cortex model to a weak shock stimulus. This data is from a small-scale model (96 pyramidal neurons). The color bar represents the membrane potential in mV . Each horizontal pair of frames represents a specific time (in $msec$) after the shock stimulus (indicated at the left of each pair of frames). The schematic diagram at the top of the page represents connections between pyramidal neurons (left) and feedback inhibitory neurons (right).

pathway are doubled, damping of the oscillatory field potential occurs much more rapidly (figure 7.16, B); only three oscillatory cycles are seen instead of five. Conversely, if this pathway is abolished, damping of the response does not occur; instead, persistent oscillations are seen. Depending on the specific parameter values, periodically modulated oscillations or constant amplitude oscillatory patterns can be observed in the field potentials (figure 7.16, C and D).

7.3.3 Influence of feedback inhibitory interneurons

The gamma oscillation frequency of this model is critically dependent on parameters relating to feedback inhibitory interneurons. There is some experimental evidence for inhibitory interactions between these interneurons [5]. In the model, these connections prevent the highly excitable feedback interneurons from firing more than once per gamma cycle. If these mutually inhibitory interactions are removed, each feedback interneuron spikes twice




Figure 7.13: Schematic diagram of excitatory connections between pyramidal neurons (*P*) in the piriform cortex. Each pyramidal neuron schematic represents a class of neurons in one of the three subdivisions of this brain region. The density of connections between regions is represented by the width of lines connecting the pyramidal neurons. Shading of pyramidal neurons and connections is just to distinguish between the three regions. Abbreviations: *M/T*, mitral/tufted cells of the olfactory bulb; *G*, granule cells of the olfactory bulb; *LOT*, lateral olfactory tract. From [5].




Figure 7.14: Inhibitory connections in the piriform cortex. Abbreviations: *P*, pyramidal neuron; *FF*, feedforward inhibitory interneuron; *FB*, feedback inhibitory interneuron. From [5].

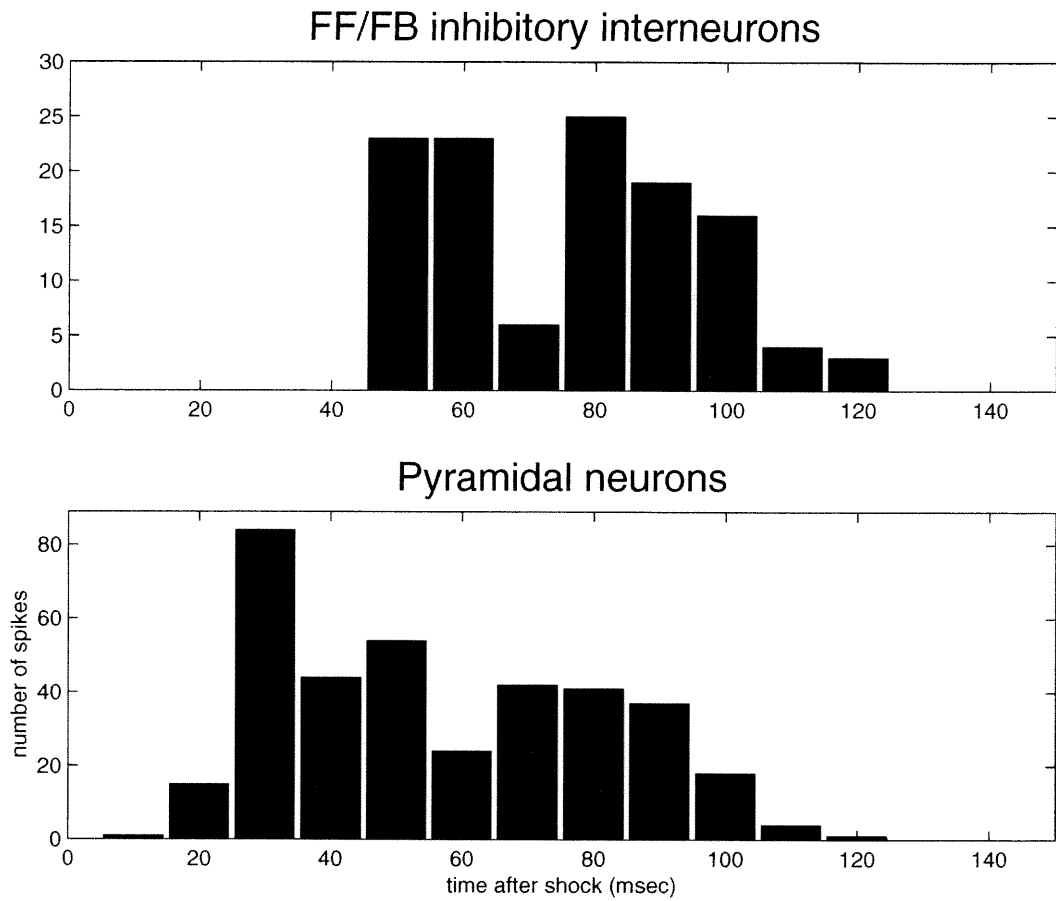


Figure 7.15: Histogram of spikes in feedforward/feedback inhibitory neurons and pyramidal neurons after a typical simulation of the piriform cortex model in response to a weak shock stimulus.

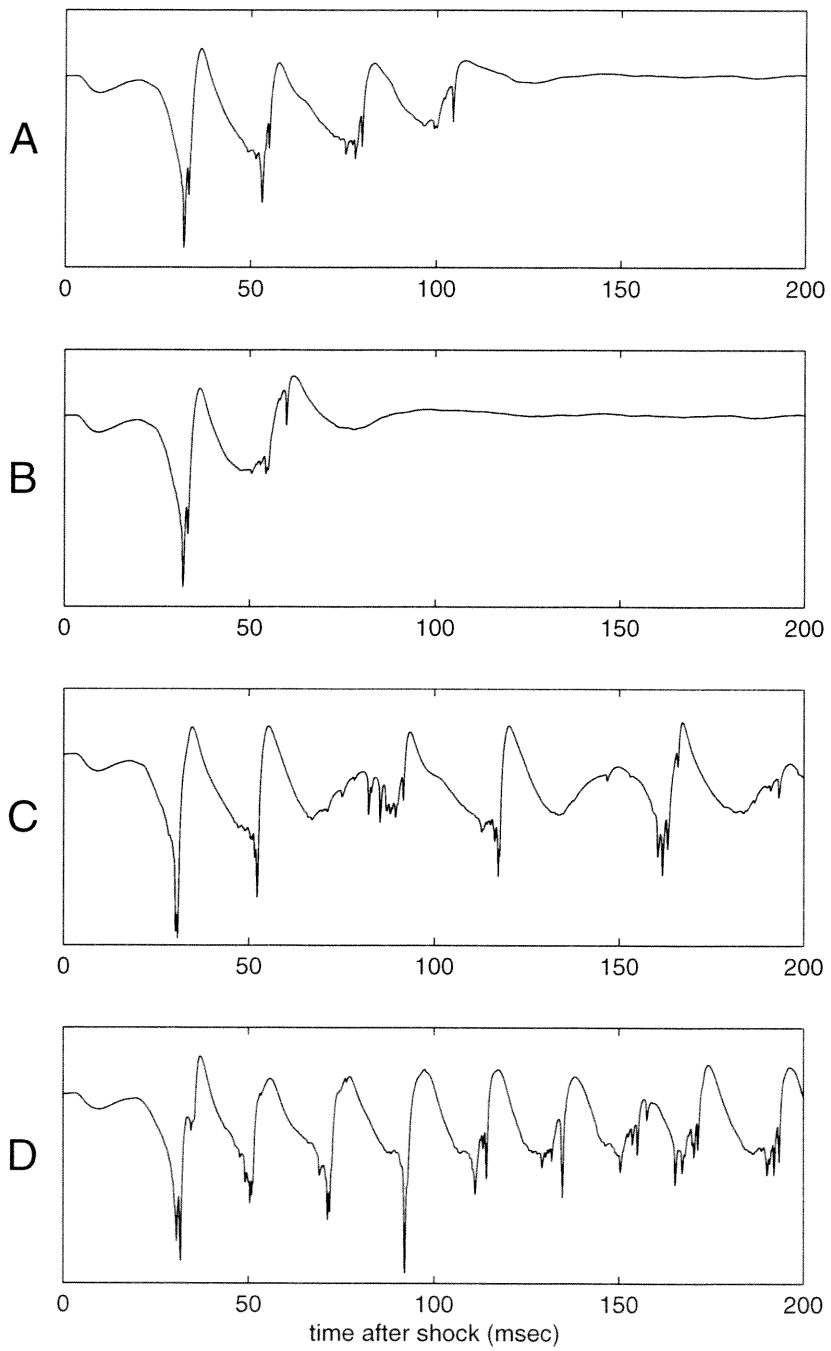


Figure 7.16: (Previous page) Dependence of weak shock-evoked oscillations on parameter settings in the piriform cortex model. Each figure represents the surface field potential recorded in posterior piriform cortex. **A.** Standard parameter set. **B.** Synaptic weights from feedforward/feedback interneurons are doubled. **C.** Synaptic weights from feedforward/feedback interneurons are set to zero. **D.** Same as **C**, but connections between feedback interneurons were also set to zero. Traces are unfiltered.

per oscillatory cycle (figure 7.17), which greatly increases the effective time constant of inhibition. This in turn reduces the oscillatory response to a weak shock stimulus from four to five successive traveling waves of spiking activity to three waves, and also changes the shape and the power spectrum of the surface field potential substantially (figure 7.18). Specifically, whereas the power spectrum in the original model has most of the power in the theta (4-10 Hz) and gamma (40-70 Hz) band, as in piriform cortex [1], the model without interconnections between feedback inhibitory interneurons has a much broader distribution with significant power in the beta (20-35 Hz) band, reflecting the slower oscillation rate. Some versions of the Wilson/Bower piriform cortex model [20, 21] also contained synaptic connections between feedback inhibitory interneurons, but the possible functions of these connections were not explored.




Figure 7.17: Intracellular traces of feedback inhibitory interneurons in the piriform cortex model in the presence and absence of inhibitory connections between feedback inhibitory interneurons. Top: with inhibitory to inhibitory connections. Bottom: without inhibitory to inhibitory connections.

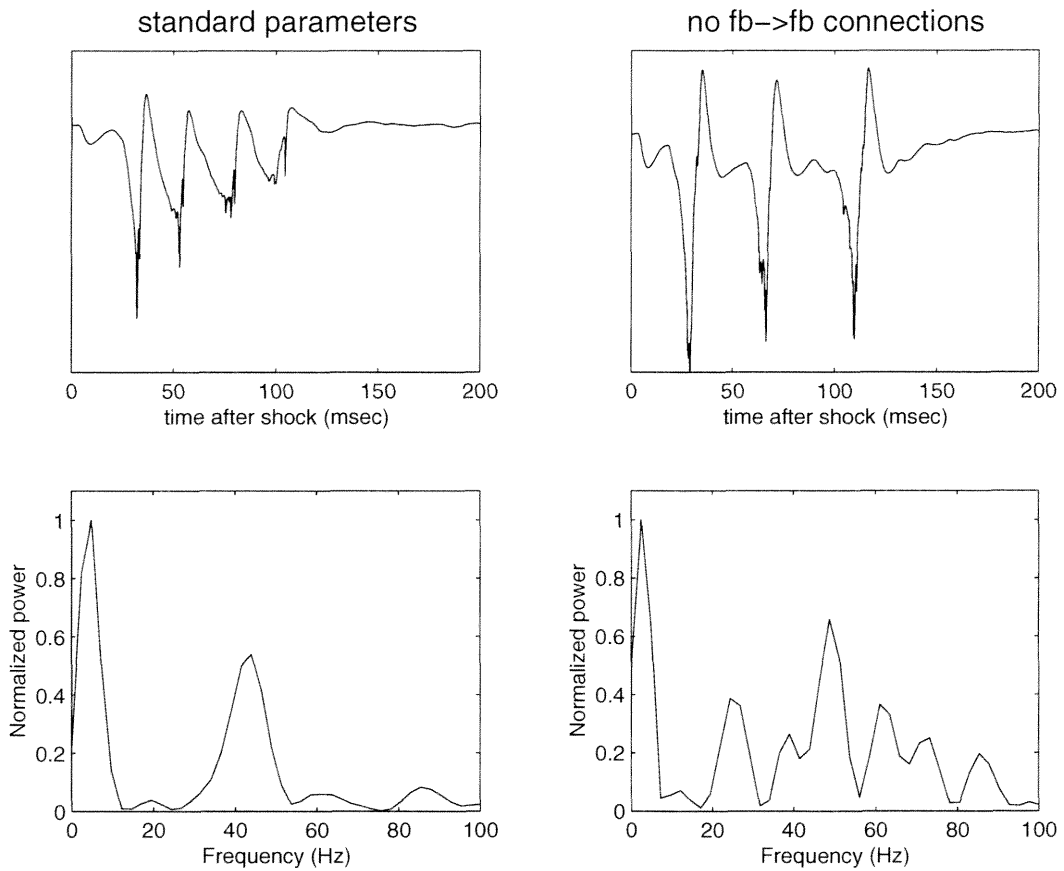


Figure 7.18: Surface field potentials and power spectra in the piriform cortex model in the presence and absence of inhibitory connections between feedback inhibitory (fb) interneurons. Top left: surface field potential with inhibitory-to-inhibitory connections. Bottom left: power spectrum of surface field potential with inhibitory-to-inhibitory connections. Top right: surface field potential without inhibitory-to-inhibitory connections. Bottom right: power spectrum of surface field potential without inhibitory-to-inhibitory connections.

7.3.4 CSD response

The Wilson/Bower model of piriform cortex [20, 21] did not include CSD results for weak shock stimuli, because at the time that model was constructed there was no relevant experimental data available. Recently, experimental CSD analysis has been performed in the piriform cortex of rats receiving weak shock stimuli [11] and the results are very revealing. In figure 7.19 [9] we see contour plots of the CSD responses to strong and weak shock stimuli. In figure 7.20 we see the same data presented as CSD surface plots. There are a number of interesting features of these results which point out serious flaws in the piriform cortex model with random connectivities described above. Here, I will describe both the essential features of the experimental weak shock response and why the random connectivity model cannot replicate them.

A technical detail needs to be mentioned at this point. CSD responses as measured in our piriform cortex models feature very large ionic currents in the vicinity of the soma and extending well into the apical dendrite. These currents are a consequence of the pyramidal neuron model used and are not observed in experimentally measured CSDs. The reason for this discrepancy is unclear, but may reflect differences in cable properties between the morphologies of the real neuron and the simplified model neurons used in the network model. These large ionic currents make interpretation of the model CSDs extremely difficult, as they overlap with synaptically generated current sinks in layer 1b. In order to circumvent this problem, synaptic CSDs were measured from the model by ignoring all but the synaptic

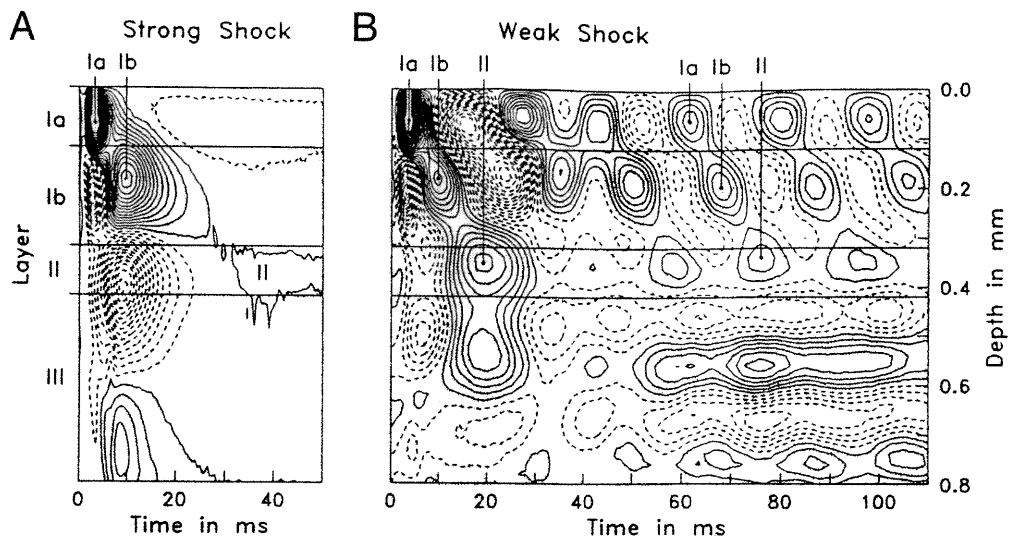


Figure 7.19: Experimental CSD contour plots of piriform cortex responses to strong and weak shocks. **A.** Strong shock response. **B.** Weak shock response. Time is in *msec* after the shock. From [11]

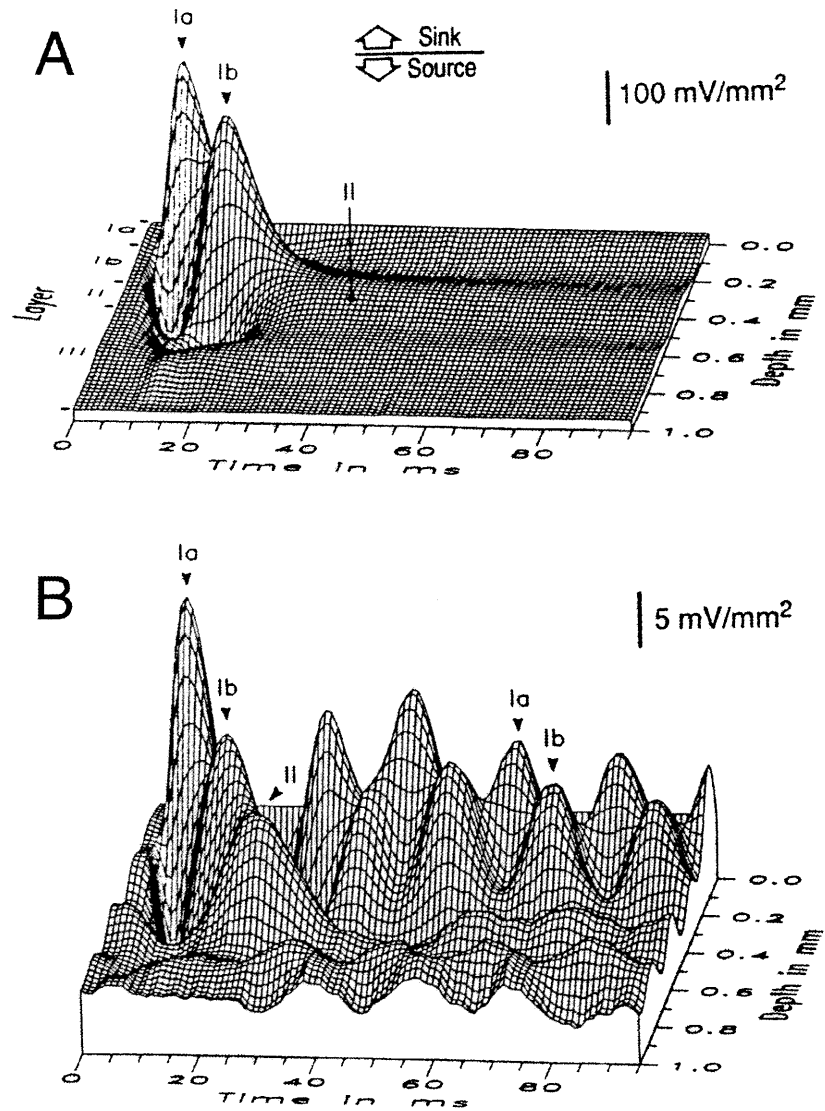


Figure 7.20: Experimental CSD surface plots of piriform cortex responses to strong and weak shocks. **A.** Strong shock response. **B.** Weak shock response. Time is in msec after the shock. From [11]

currents in the computation of the CSD profiles. Since synaptically derived currents are believed to comprise a large majority of the experimentally measured CSD currents as well [10], this approximation is reasonable and facilitates comparison between the model and the experimental data.

A contour plot of the weak shock CSD response from the random connectivity model is shown in figure 7.21. The most obvious difference between this response and the experimental weak shock CSD response (figure 7.19) is that the model gives rise to only one peak in layer 1a, while the experimental data exhibits multiple 1a peaks. These peaks are most likely due to direct inputs from the olfactory bulb, as these inputs comprise a large majority of the afferent inputs to the piriform cortex [5]. It is not clear, however, how a single weak shock can elicit multiple waves of activation on the bulbar inputs to the piriform cortex. One possibility is that the shock causes back-propagating action potentials along the LOT axons to the olfactory bulb mitral cells. This in turn might induce an oscillatory “ringing” in the bulbar network, which then could generate successive waves of inputs at a frequency of approximately 50 *Hz*. In this case, back-propagating action potentials only occur on the first cycle, in direct response to the shock (figure 7.22, A) while subsequent afferent inputs are due to the internal circuitry of the olfactory bulb causing mitral cells to spike on each gamma cycle. Another possibility is that back-projections from piriform cortex to the olfactory bulb [5, 19] might induce multiple waves of activity (figure 7.22, B). In the random connectivity model, however, a single weak shock input results in only a single

peak in layer 1a (figure 7.21). In order to model the phenomenon of multiple 1a peaks, simulations were run whereby the weak shock was repeated multiple times at 50 *Hz* (figure 7.23). Interestingly, not all of the multiple input shocks (at 20, 40, 60, 80, and 100 *msec* after the first shock) can be seen in the CSD profile. This is because activation of layer 1b afferents raise the membrane potential in distal apical dendrites close to the synaptic reversal potential, which dramatically reduces the input currents from subsequent shocks (figure 7.24) because the synaptic driving force is much smaller. Other than this, the resulting CSD profiles were not substantially different from those elicited by a single weak shock. In both cases, large current sinks were generated in layer 1b as a result of feedforward projections to superficial layer 1b dendrites on pyramidal neurons and feedback projections to deep layer 1b dendrites.

Another significant discrepancy between the experimental CSD results and those of the random connectivity model involves the latency between the CSD peaks in layer 1a (caused by afferent input) and layer 1b (caused by associational input). The experimental data shows that the latency between these peaks is roughly the same for both the weak and the strong shock response (about 6 to 8 *msec*) (figures 7.4 and 7.19). In the random connectivity model, the latency between these peaks is much higher for the weak shock than the strong shock; for the strong shock the latency is 8-9 *msec* (figure 7.4) while for the weak shock it is over 20 *msec* (figures 7.21 and 7.25), effectively skipping an entire gamma cycle (at 50 *Hz*, *i.e.*, 20 *msec*).




Figure 7.21: CSD contour plot of the response of the piriform cortex model (random connectivity model) to a weak shock stimulus. Contour values are shown at right on the color bar, with units of mV/mm^2 . Piriform cortex layers are separated with black lines and are labeled on the right. The arrowhead denotes the time of the shock.

Figure 7.22: (Previous page) Schematic diagram of how multiple inputs to layer 1a in the piriform cortex might be generated from a single weak shock. Numbers represent the order in which events occur. **A.** Multiple inputs caused by back-propagation of the weak shock to olfactory bulb mitral cells, which induces a series of inputs along the LOT. 1. Shock occurs. 2. Spikes propagate anterogradely towards pyramidal neurons in piriform cortex and retrogradely towards olfactory bulb mitral cells. 3. Mitral cells spike, activating olfactory bulb granule cells via dendrodendritic synapses. 4. Granule cells inhibit mitral cells via dendrodendritic synapses. 5. Mitral cells generate a series of inputs to the piriform cortex because of the activation of the bulbar circuitry. **B.** Multiple inputs caused by projections from piriform cortex pyramidal neurons to olfactory bulb granule cells, which set up an oscillatory pattern in the bulb. 1. and 2. as above. 3. Pyramidal neurons in the piriform cortex spike and send inputs to olfactory bulb granule cells. Subsequent activation of the bulbar circuitry and generation of multiple inputs to piriform cortex is as above.




Figure 7.23: CSD contour plot of the response of the piriform cortex model (random connectivity model) to a weak shock stimulus with multiple input shocks occurring at a rate of 50 Hz . Contour values are shown at right on the color bar, with units of mV/mm^2 . Piriform cortex layers are separated with black lines and are labeled on the right. Arrowheads denote the times of the shocks.




Figure 7.24: Intracellular response of a typical pyramidal neuron (layer 1a dendrite) in the piriform cortex model to a weak shock stimulus consisting of multiple shocks. This neuron is located in the dorsal anterior region of piriform cortex. Arrowheads denote the times of the shocks. Time is in *msec* after the first shock.

The CSD peak amplitudes of the weak shock response in the experimental data are also much smaller than those of the strong shock response; in figure 7.20, the amplitude of the strong shock response is roughly twenty times greater than that of the weak shock response. In the random connectivity model, however, the CSD responses are of comparable amplitudes for both weak and strong shocks (figure 7.25). In addition, the ratio of layer 1b to layer 1a peak heights, which is close to unity for weak and strong shocks experimentally (figure 7.20) and for strong shocks in the model (figure 7.3), is much higher for weak shocks in the model.

These results suggest several points about piriform cortex dynamics. First, the amplitude of the layer 1b peak in the CSD plots is roughly proportional to the number of pyramidal neurons firing in response to the input stimulus, since this is what determines the amount of synaptic activity which is directly measured in the CSDs. Thus, the fact that this peak is twenty times smaller in the weak shock response than in the strong shock response implies that roughly $1/20$ as many neurons fire in response to the weak shock as compared to the strong shock. We have assumed that a strong shock causes all pyramidal neurons in the ventral anterior region of piriform cortex to fire. Thus, we would naively expect a much smaller number of neurons to fire in response to a weak shock per gamma cycle. In contrast, in the cortical model the random connectivity between olfactory bulb mitral cells and piriform cortex pyramidal neurons has the effect that activating (in this case) $1/5$ of the bulbar inputs causes all of the pyramidal neurons to receive on average

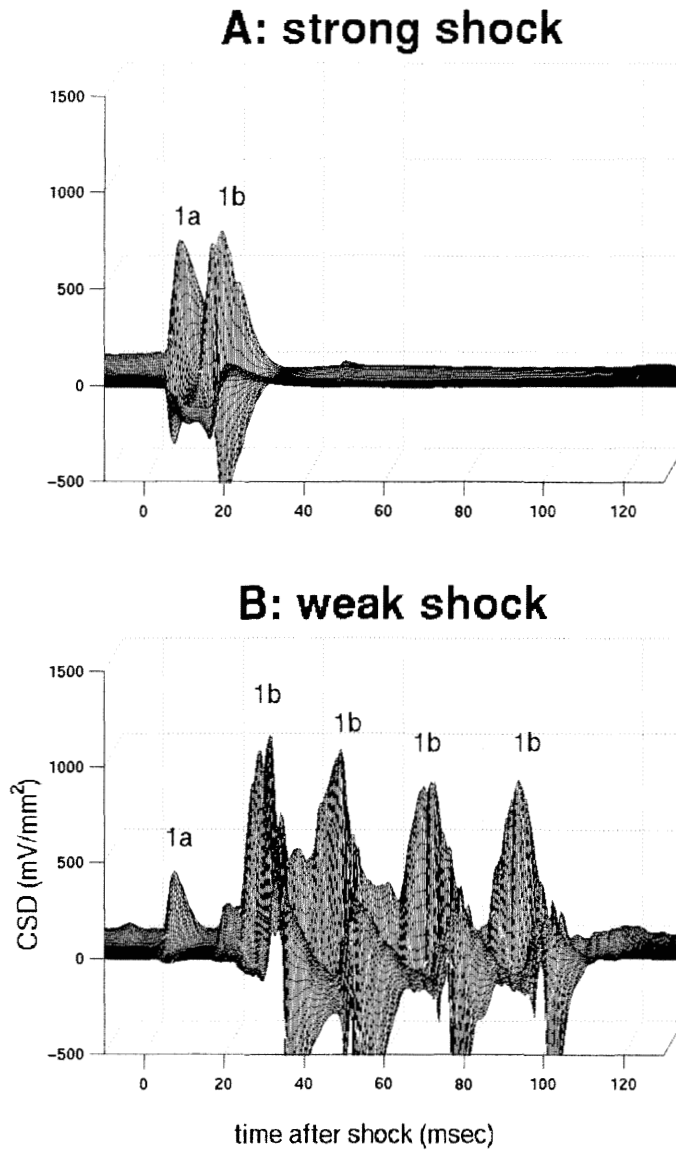


Figure 7.25: CSD surface plot of the response of the piriform cortex model (random connectivity model) to weak and strong shock stimuli. Current sinks point upwards. Peaks are labeled with the layer of piriform cortex in which they occur. **A.** Strong shock response. **B.** Weak shock response.

about $1/5$ as much excitatory input, which causes all the neurons to spike with a very long latency as opposed to causing $1/5$ of the neurons firing with a short latency (figures 7.12, 7.21, and 7.25). This accounts for the much larger CSD amplitudes in response to the weak shock seen in the model as compared to the data, and also accounts for the much larger size of the 1b peaks compared to the 1a peak(s).

These results suggest that there is something fundamentally wrong with the random connectivity model despite the fact that it replicates the oscillatory surface field potential fairly accurately. Interestingly, this model nevertheless reflects the view of many neuroscientists on the genesis of the oscillatory dynamics in piriform cortex [4, 21]. In particular, the Wilson/Bower piriform cortex model, which was the immediate predecessor of the random connectivity model described above, exhibited the same kinds of dynamical behaviors in response to weak shock stimuli (although CSDs for weak shock stimuli were not measured in that model). Below I describe a different model which replicates the CSD response to weak shocks much more accurately. As will be seen, the difference between the two models has major consequences for theories of neural coding and information processing in the piriform cortex (discussed in chapter 8).

7.4 The weak shock response: model with structured connectivities

This section describes the weak shock response in the second version of the piriform cortex model. This model features highly structured connectivity between the olfactory bulb and piriform cortex and within cortex. Like the random connectivity model, this model can accurately reproduce the strong shock response. Unlike the previous model, however, this model also replicates the weak shock CSD response of the system with considerable accuracy.

7.4.1 Structure of the model

The structured connectivity model has several significant differences from the random connectivity model, shown schematically in figures 7.26 and 7.27. They are:

1. Mitral cells in the olfactory bulb and pyramidal neurons in piriform cortex are divided into non-overlapping neuronal groups on the basis of their connection patterns. For convenience in recording simulated CSDs, I used four groups in both olfactory bulb and cortex (figure 7.26). However, it is possible that a much larger number of groups may exist (see chapter 8).
2. Mitral cells in the olfactory bulb in a particular group project exclusively to pyramidal neurons in piriform cortex of a particular group. Thus, mitral cells in bulbar group 1 (red color in the figure) project to pyramidal neurons in cortical group 1, etc. (figure

7.26).

3. Excitatory associational connections from pyramidal neurons in piriform cortex project exclusively to pyramidal neurons of the same group (figure 7.27).
4. Feedback excitatory connections from pyramidal neurons in posterior piriform cortex to pyramidal neurons in anterior piriform cortex were removed. The reason for this is discussed below.

The motivation for using the structured connections is as follows. The latency between the time a shock input stimulus hits a pyramidal neuron in layer 1a and the time the neuron spikes is roughly equal to the time delay between the layer 1a and layer 1b peaks in the CSD. The CSD data indicates that this delay is not substantially different for weak and strong shock stimuli. The CSD data also suggests that far fewer neurons are spiking in response to a weak shock than to a strong shock. From this it can be inferred that neurons that spike in response to a weak shock are receiving roughly as much input as they receive during a strong shock, while those which are not are receiving very little or no input. Let us also assume that a weak shock activates a much smaller proportion of LOT fibers than a strong shock (say 20%). If all bulbar connections to all pyramidal neurons are set up randomly, then on average all the neurons would be receiving only 20% as much input. This results in a much greater latency to firing, which contradicts the data (figure 7.28). What is more likely is that the active 20% of the inputs are mainly projecting to a very

well-defined subset of pyramidal neurons, and that these neurons receive essentially all their afferent input from that subset of input fibers. In this case, these neurons will then fire with the same latency as with a strong shock, since they are receiving the same amount of input in both cases, even though only 20% of all afferent inputs are active. This then suggests that the inputs are segregated into groups which project exclusively to particular subsets of neurons, which is the hypothesis explored here.

One important question is: how can a weak shock stimulus created by the experimentalist activate neurons of only one group? In fact, the data (figure 7.20, B) suggests that it does not. The first layer 1a peak in the experimental CSD plot is about twice as large as subsequent peaks. We suggest that this peak is due to spiking in LOT fibers projecting from mitral cells belonging to more than one neuronal group, but that subsequent peaks (generated by the internal circuitry of the olfactory bulb) are mitral cells belonging to single neuronal groups. This will be discussed further below.

7.4.2 Response to multiple input shocks

As mentioned above, in order to simulate multiple CSD peaks in layer 1a it is necessary to subject the network model to multiple weak shocks at a frequency of approximately 50 Hz (corresponding to the gamma rhythm). For the structured connectivity model we activated a single group of neurons each cycle. The two alternatives explored were: activating a different group of neurons each gamma cycle or activating the same group of neurons each

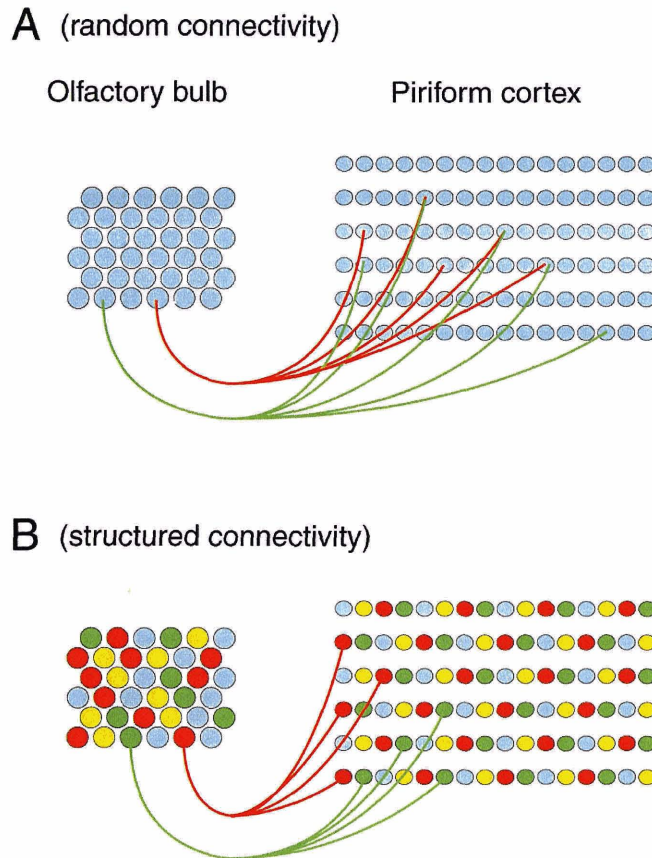



Figure 7.26: Schematic diagram of olfactory bulb mitral cells and piriform cortex pyramidal neurons in the two piriform cortex models. **A.** Model with random connectivities. Individual mitral cells project to broad and overlapping subsets of pyramidal neurons in the piriform cortex. **B.** Model with structured connectivities. Different neuronal groups are identified with one of four colors. Mitral cells from a given neuronal group only project to pyramidal neurons of the corresponding group. Neuronal groups are non-overlapping at the level of both the bulb and cortex.

Piriform cortex

A (random connectivity)



B (structured connectivity)





Figure 7.27: Schematic diagram of associative connections between pyramidal neurons in the two piriform cortex models. **A.** Model with random connectivities. Individual pyramidal neurons project to broad and overlapping subsets of pyramidal neurons in the piriform cortex. **B.** Model with structured connectivities. Different neuronal groups are identified with one of four colors. Pyramidal neurons from a given neuronal group only project to pyramidal neurons of the corresponding group.

A random connectivity



B structured connectivity




Figure 7.28: Schematic diagram of how connection topologies affect spiking latencies in response to a weak shock in the piriform cortex model. **A.** Model with random connectivities. All neurons receive 20% of their maximum afferent inputs, leading to a large spike latency. **B.** Model with structured connectivities. Some neurons receive near-maximal afferent inputs while others receive very little input, leading to a small spike latency in neurons receiving large inputs. Circles at the right represent pyramidal neurons; the shading represents the level of input (black = maximal; white = none; gray = partial).

cycle.

Different neuronal groups activated each cycle

Activating a different neuronal group each cycle gives rise to an ordered series of waves of excitation in the cortical model alternating between different neuronal groups (figure 7.29). The neuronal groups in the cortex were set up with a regular spatial spacing of period 4 (figure 7.27). This arrangement was used in order to produce more consistent CSD plots, since each simulated CSD recording electrode is equidistant from neurons of each group. This is necessary in a small scale model; in the real system the number of neurons is so much larger that activation of any one group would have a roughly equivalent effect on the CSD profile as long as neurons from all groups are randomly distributed throughout the cortex. As can be seen in figure 7.29, the first wave activates a particular group of neurons (group 1), the next wave activates a different (non-overlapping) group (group 2), etc.

The CSD profile for the structured connectivity model is shown in figure 7.30. If we compare this to the experimental data (figure 7.19, B), we see that the model replicates the main features of the data. In particular, the model displays an ordered sequence consisting of a layer 1a peak followed by a layer 1b peak, then another layer 1a peak, etc. The 1a/1b pattern is repeated every 20 *msec*. A closeup of the first two peaks shown as a surface plot (figure 7.31) shows that the two peaks are of nearly the same height, with the layer 1a peak somewhat larger. The latency between the 1a and 1b peaks is roughly 9 *msec*,

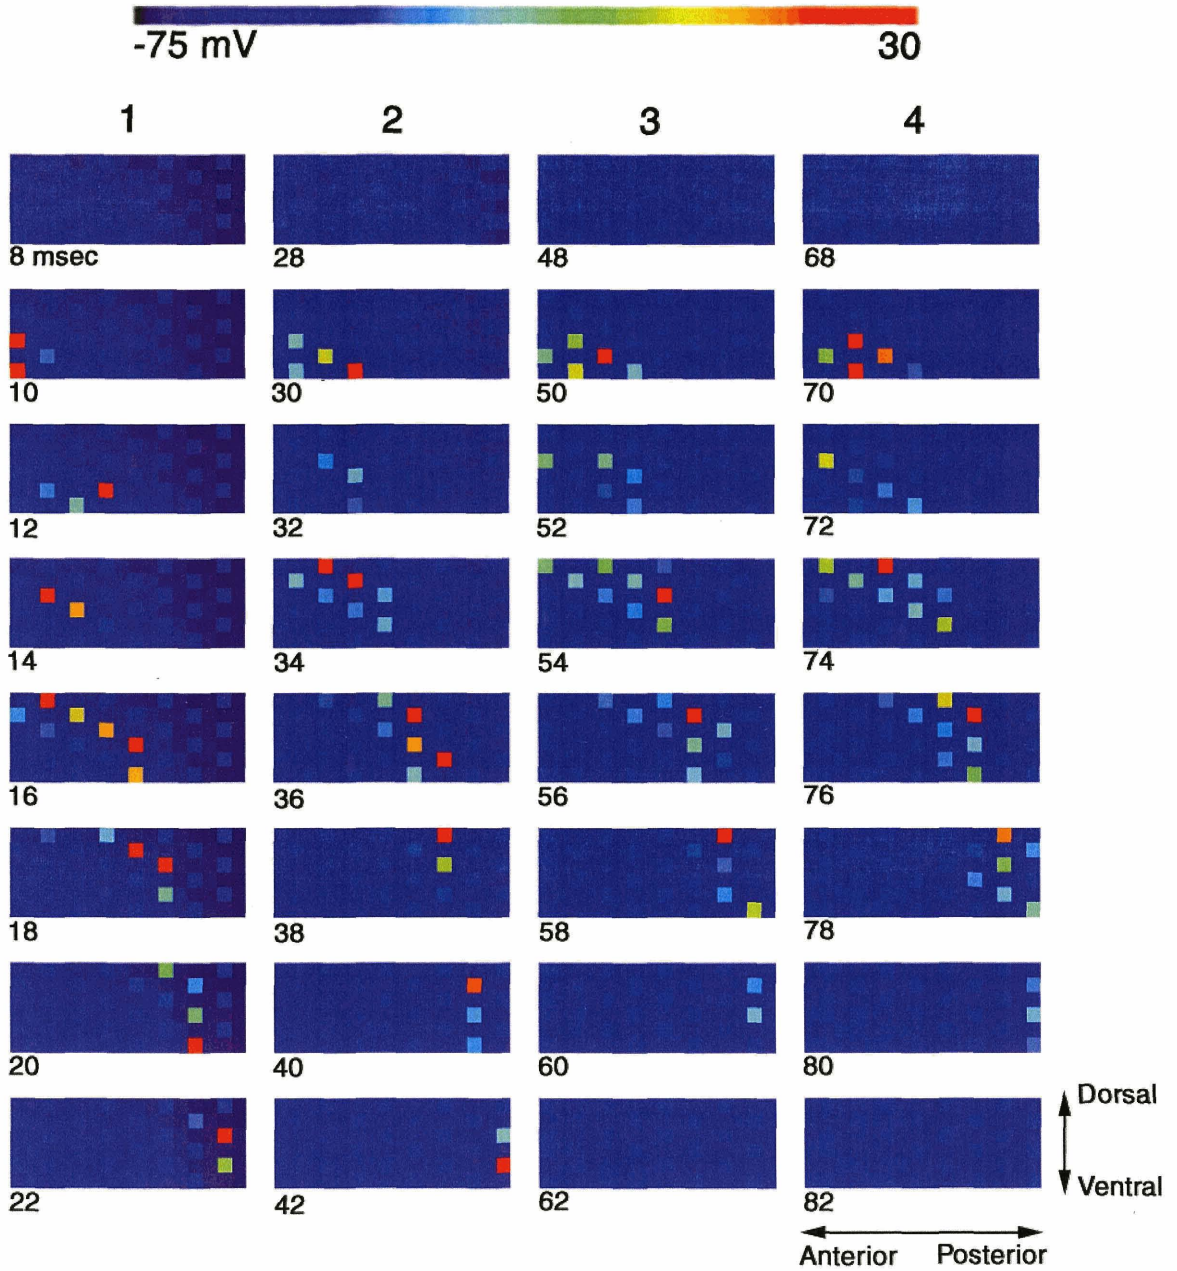


Figure 7.29: (Previous page) Response of the piriform cortex model with structured connectivities to a weak shock stimulus. This data is from a small-scale model (96 pyramidal neurons). Each frame represents the state of the ensemble of pyramidal neuron somata at a particular time (measured in *msec* after a weak shock). Each colored square in a frame represents a single pyramidal neuron soma. Time after the shock is listed at the lower left corner of each frame. The color scale is shown above the frames. Cortical orientation is as indicated. Each vertical strip represents a single wave of activation of inputs. In this version of the model, inputs to a different subset of pyramidal neurons are activated in each wave. Only the first four waves of activation (out of six) are shown. Note that each wave consists of the activation of non-overlapping groups of neurons.

which is comparable to that seen in the experimental data. Similar results to these hold for larger-scale cortical models; I have tested models using up to 400 pyramidal neurons and seen the same phenomena (data not shown).

One area where the model differs from the data is that the CSD peaks in the model cover a broader expanse of dendrite than in the data; this presumably means that synaptic activation is more concentrated on particular regions of dendrites in the real system than in the model. Interestingly, the location of synaptic connections and their relative densities on pyramidal neuron dendrites were taken from a modeling paper by Ketchum and Haberly ([10], fig. 2), who also supplied the CSD data in response to a weak shock (figure 7.19). Our model therefore suggests that the synaptic connectivity values taken from [10] are

not completely consistent with experimental CSD data obtained in the same laboratory. However, our model can easily be modified to accurately match this aspect of the data.

Another minor difference between the model CSD and the data is that the afferent peaks in the data after the first are more dispersed in time than in the model; this suggests that the inputs following the initial weak shock (which are generated from the olfactory bulb circuitry) are not as precisely synchronized as they are in the model inputs. Additionally, the first afferent (1a) peak in the data is about twice as large as the first 1b peak, whereas the difference is considerably smaller in the model. This will be discussed further below.

If only the afferent connections are structured, and the intracortical pyramidal to pyramidal excitatory connections are randomly distributed as in the random connectivity model, then the ordered firing pattern very quickly breaks down into an irregular pattern (figures 7.32). This is because spiking of neurons of a single group gives rise to inputs onto all pyramidal neurons due to the randomly distributed intracortical connections. In particular, some neurons which are not in the same neuronal group as the neurons which fired first also fire, and these activate other neurons through local and long-range connections. The net result is a rapid breakdown of ordered firing, which strongly suggests that if afferent connections are ordered into groups, associational connections between pyramidal neurons must also be ordered into groups.




Figure 7.30: CSD contour plot of the response of the piriform cortex model with structured connectivities to a weak shock stimulus with multiple input shocks occurring at a rate of 50 Hz . A different subset of pyramidal neurons was activated on successive waves of inputs. Contour values are shown at right on the color bar, with units of mV/mm^2 . Piriform cortex layers are separated with black lines and are labeled on the right. Arrowheads denote the times of the shocks.




Figure 7.31: CSD surface plot of the response of the piriform cortex model with structured connectivities to a weak shock stimulus with multiple input shocks occurring at a rate of 50 Hz . A different subset of pyramidal neurons was activated on successive waves of inputs. Only the first two peaks are shown. Layer in which each peak occurs is as indicated.

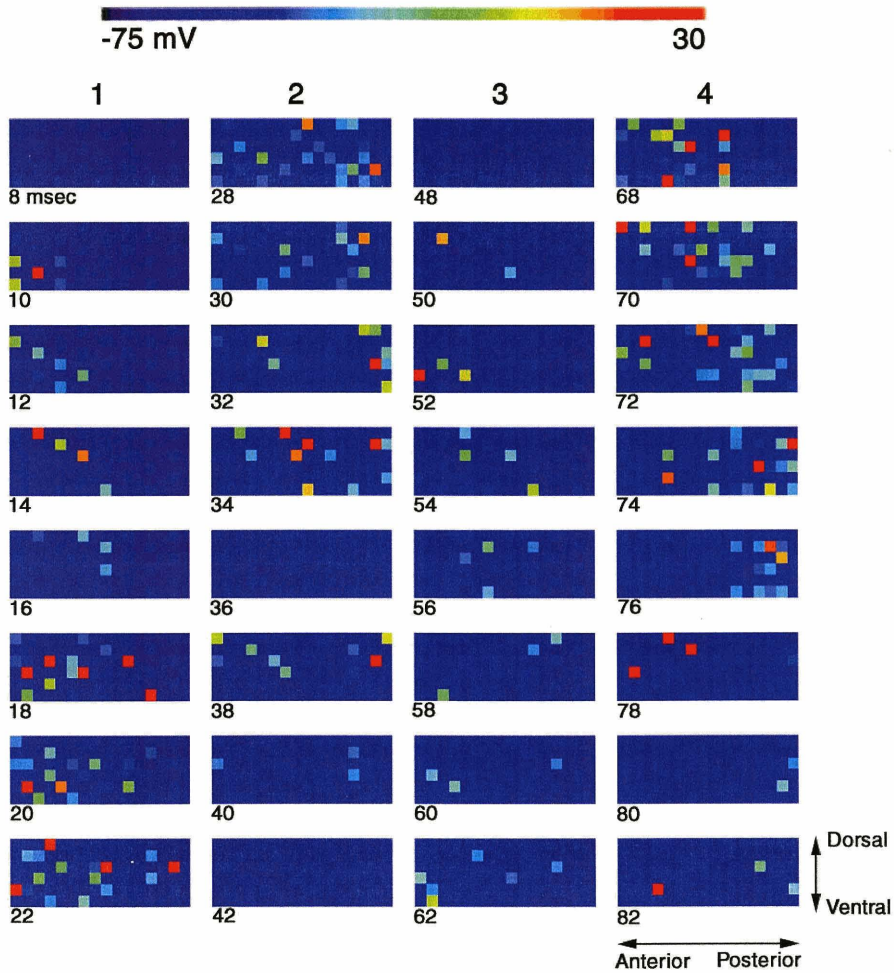


Figure 7.32: Response of the piriform cortex model with partially structured connectivities to a weak shock stimulus. This version of the model contains structured connectivities only in the mitral cell to pyramidal neuron connections; connections between pyramidal neurons are set up randomly. Figure layout is as in figure 7.29.

Same neuronal group activated each cycle

If the same neuronal group is activated on each gamma cycle, the ordered series of neuronal activations seen previously again breaks down to give rise to a more disordered firing pattern (figure 7.33). In particular, note that the second wave of activation does not cause pyramidal neurons in the ventral anterior piriform cortex to spike at all. The CSD contour plot (figure 7.34) also shows that most of the layer 1b peaks after the first are substantially attenuated, which is not surprising given that the pyramidal neurons are firing quasi-randomly after the first wave of excitation. These results depend on the presence of slow potassium channels (predominantly K_M) which are activated when pyramidal neurons spike (figure 7.35). These channels have slow enough time constants that they have not returned to their baseline state when the next wave of inputs arrives, which prevents the pyramidal neurons from firing in response to these inputs.

7.4.3 Discussion

I have shown that a piriform cortex model can reproduce experimental CSD data for a weak shock stimulus as long as the mitral cells in the olfactory bulb and the pyramidal neurons in the olfactory cortex are divided into non-overlapping groups which are predominantly connected to other neurons of the same group. This is a radical hypothesis that has many consequences from the standpoint of computations that the olfactory system may be performing. I will summarize the evidence here and in the next chapter I will present my view

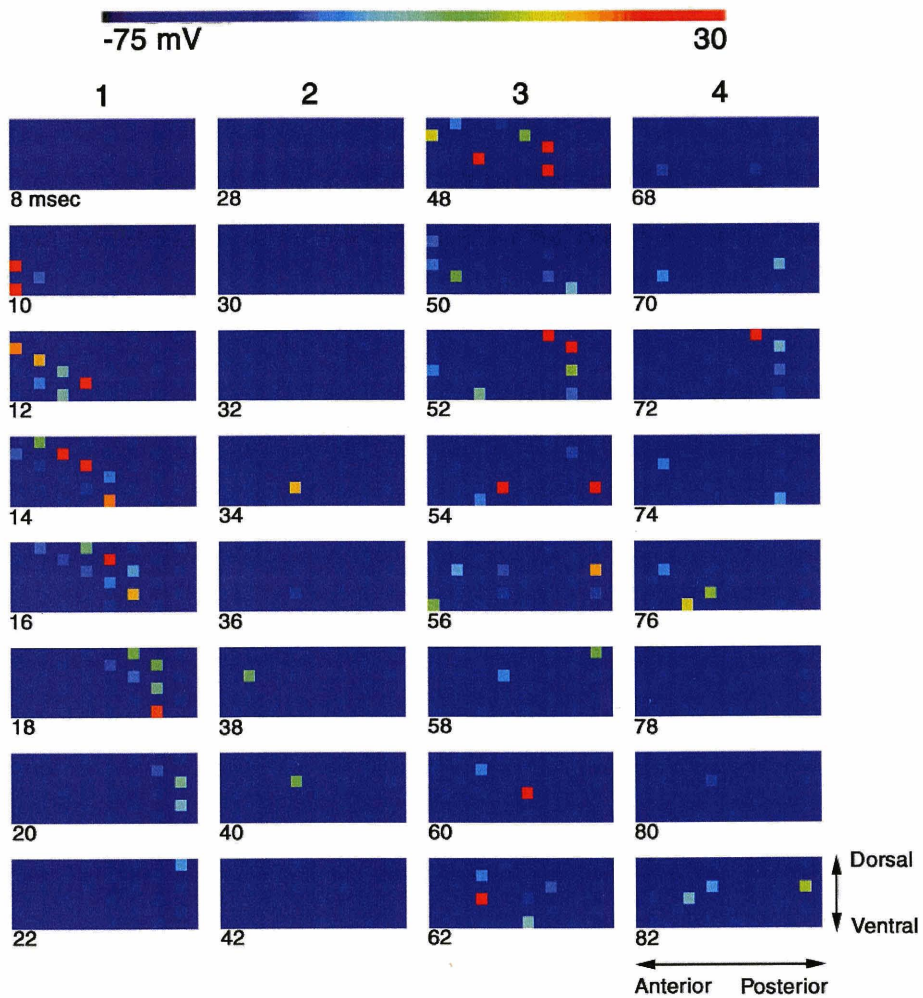


Figure 7.33: Response of the piriform cortex model with structured connectivities to a weak shock stimulus. In this version of the model, the same subset of pyramidal neurons are activated on each wave of inputs. Figure layout is as in figure 7.29.




Figure 7.34: CSD contour plot of the response of the piriform cortex model with structured connectivities to a weak shock stimulus with multiple input shocks occurring at a rate of 50 *Hz*. The same subset of pyramidal neurons was activated on successive waves of inputs. Contour values are shown at right on the color bar, with units of mV/mm^2 . Piriform cortex layers are separated with black lines and are labeled on the right. Arrowheads denote the times of the shocks. Note the attenuation of the layer 1b peaks after the second shock.




Figure 7.35: Intracellular plots of the response of a single pyramidal neuron (recorded at the soma) in the piriform cortex model with structured connectivities to a weak shock stimulus with multiple input shocks occurring at a rate of 50 Hz . The same subset of pyramidal neurons was activated on successive waves of inputs. Top: membrane potential. Middle: currents through potassium channels. Bottom: closeup of currents through potassium channels. K channel abbreviations: Kdr : delayed-rectifier; Ka : fast inactivating; KM , slow voltage-dependent; $Kahp$: calcium-dependent. Arrowheads denote the times of the shocks. Note that current through the KM channel does not return to its baseline value between shocks, which prevents the neuron from spiking on successive waves of input.

of the computational implications of these results.

As shown above, a cortical model with random connectivities between the olfactory bulb and the olfactory cortex cannot replicate the CSD response to a weak shock stimulus. If a weak shock activates only a small subset of input fibers, as seems likely, then each pyramidal neuron in ventral anterior piriform cortex of a randomly connected model will receive a correspondingly smaller input, which will result in a long latency to firing (figure 7.28). In contrast, in the data we see a short latency to firing in response to both strong and weak shocks (figure 7.19). This suggests that only a subset of neurons in the vAPC are firing in response to a weak shock, but that these neurons are receiving a similar amount of input when either weak or strong shock stimuli are imposed on the system. If we divide neurons into neuronal groups in both the bulb and the cortex (the structured connectivity model; figure 7.26), the latencies are similar for both weak and strong shock stimuli. Note that the division into groups is irrelevant for the strong shock response, since all input fibers are activated. The structured connectivity model also replicates other features of the CSD response, such as the relative peak heights in layer 1a and 1b.

In addition, I have shown that activating the same neuronal group each cycle results in disorganized firing, due to the activity of slow potassium currents in pyramidal neurons which cannot return to their baseline levels by the time the next wave of inputs arrives. This suggests that successive cycles of activity in piriform cortex should consist of different neuronal groups. This was why the feedback connections from posterior to anterior piriform

cortex were disabled in model 2: it was not clear which groups in posterior piriform cortex should project to which groups in anterior piriform cortex. For instance, if the groups are activated in the sequence 1/2/3/4 etc., then neurons in group 1 in the posterior piriform cortex should presumably have feedback projections to neurons in group 2 in anterior piriform cortex in order to reinforce the pattern. Alternatively, if the groups are activated in the sequence 1/1/1/1 etc. then neurons in posterior piriform cortex should have feedback projections to neurons of the same group. In either case, these feedback connections in the model are clearly not needed in order to replicate the experimental weak shock CSD data.

As mentioned above, the first afferent peak in the CSD data is disproportionately large in comparison with subsequent peaks, a feature which is not replicated by our model. This is probably due to the fact that the first peak is the result of the experimental stimulus, whereas the subsequent peaks are most likely to be generated by the bulbar circuitry on its own (figure 7.22). If we assume that neurons in the olfactory bulb and piriform cortex are in fact organized into non-overlapping groups as proposed above, then it is very unlikely that an experimentally applied weak shock stimulus would activate most or all neurons from a single group and no others; rather, the shock would activate a subset of neurons from many neuronal groups. However, subsequent waves of excitation, generated by the olfactory bulb, might well consist only of specific neuronal groups, and consist of most of the neurons of each activated group. The main role of the weak shock may thus be to activate a pattern corresponding to a particular odor in the olfactory bulb, which then plays back the rest

of the odor pattern to the cortex. In other words, the weak shock may effectively excite neurons from several patterns, of which the most strongly excited pattern will be played back in subsequent cycles. These subsequent cycles would consist of inputs from a smaller number of neurons comprising most neurons of a few neuronal groups.

In this context it is important to note that several experimentalists have recently found that individual mitral cells do not project randomly onto pyramidal neurons in piriform cortex, but rather appear to have a small number of dense patches of connections at widely separated locations throughout the piriform cortex [2, 15]. In addition, when mitral cells that are sufficiently close to each other (so that they might be assumed to receive input from the same glomerulus) were stained, their projection patterns were found to be very similar [2]. This is consistent with the hypothesis that neurons in the olfactory bulb may be organized into groups as long as the groups are organized along glomerular lines. Similar investigations have not been reported for projections within piriform cortex. However, such experiments will be necessary to test the hypothesis described above. Computational implications of this hypothesis are discussed in the next chapter.

Bibliography

- [1] S. L. Bressler. Spatial Organization of EEGs from Olfactory Bulb and Cortex. *Electroencephalography and Clinical Neurophysiology*, 57:270–276, 1984.
- [2] N. Buonviso, M. F. Revial, and F. Jourdan. The Projections of Mitral Cells from Small Local Regions of the Olfactory-bulb – An Anterograde Tracing Study Using PHA-L (Phaseolus-Vulgaris Leukoagglutinin). *European Journal of Neuroscience*, 3(6):493–500, 1991.
- [3] W. J. Freeman. The Electrical Activity of a Primary Sensory Cortex: Analysis of EEG Waves. *International Review of Neurobiology*, 5:53–119, 1963.
- [4] W. J. Freeman. Effects of Surgical Isolation and Tetanization of Prepyriform Cortex in Cats. *Journal of Neurophysiology*, 31:349–357, 1968.
- [5] L. B. Haberly. Olfactory Cortex. In G. M. Shepherd, editor, *The Synaptic Organization of the Brain*. Oxford University Press, Oxford, England, 1998.

- [6] L. B. Haberly and J. M. Bower. Analysis of Association Fiber System in Piriform Cortex with Intracellular-Recording and Staining Techniques. *Journal of Neurophysiology*, 51(1):90–112, 1984.
- [7] L. B. Haberly and J. L. Price. The Axonal Projection Patterns of the Mitral and Tufted Cells of the Olfactory Bulb in the Rat. *Brain Research*, 129:152–157, 1977.
- [8] M. Jiang, E. R. Griff, M. Ennis, L. A. Zimmer, and M. T. Shipley. Activation of Locus Coeruleus Enhances the Responses of Olfactory Bulb Mitral Cells to Weak Olfactory Nerve Input. *Journal of Neuroscience*, 16(19):6319–6329, 1996.
- [9] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 1. Current Source-Density Analysis. *Journal of Neurophysiology*, 69(1):248–260, 1993.
- [10] K. L. Ketchum and L. B. Haberly. Membrane Currents Evoked by Afferent Fiber Stimulation in Rat Piriform Cortex. 2. Analysis with a System Model. *Journal of Neurophysiology*, 69(1):261–281, 1993.
- [11] K. L. Ketchum and L. B. Haberly. Synaptic Events that Generate Fast Oscillations in Piriform Cortex. *Journal of Neuroscience*, 13(9):3980–3985, 1993.
- [12] M. B. Luskin and J. L. Price. The Topographic Organization of Associational Fibers of the Olfactory System in the Rat, Including Centrifugal Fibers to the Olfactory Bulb.

- Journal of Comparative Neurology*, 216:264–291, 1983.
- [13] J. McCollum, J. Larson, T. Otto, F. Schottler, R. Granger, and G Lynch. Short-Latency Single Unit Processing in Olfactory Cortex. *Journal of Cognitive Neuroscience*, 3(3):293–299, 1991.
- [14] J. W. Nemitz and S. J. Goldberg. Neuronal Responses of Rat Piriform Cortex to Odor Stimulation – An Extracellular and Intracellular Study. *Journal of Neurophysiology*, 49(1):188–203, 1983.
- [15] H. Ojima, K. Mori, and K. Kishi. The Trajectory of Mitral Cell Axons in the Rabbit Olfactory Cortex Revealed by Intracellular HRP Injection. *Journal of Comparative Neurology*, 230(1):77–87, 1984.
- [16] G. Schoenbaum and H. Eichenbaum. Information Coding in the Rodent Prefrontal Cortex. 1. Single-Neuron Activity in Orbitofrontal Cortex Compared with that in Piriform Cortex. *Journal of Neurophysiology*, 74(2):733–750, 1995.
- [17] G. Schoenbaum and H. Eichenbaum. Information Coding in the Rodent Prefrontal Cortex. 2. Ensemble Activity in Orbitofrontal Cortex. *Journal of Neurophysiology*, 74(2):751–762, 1995.
- [18] J. W. Scott, R. L. McBride, and S. P. Schneider. The Organization of Projections from the Olfactory Bulb to the Piriform Cortex and Olfactory Tubercle in the Rat. *Journal*

of Comparative Neurology, 194:519–534, 1980.

- [19] G. R. Shepherd, editor. *The Synaptic Organization of the Brain*, 4th. Ed. Oxford University Press, New York, 1997.
- [20] M. A. Wilson. An Analysis of Olfactory Cortical Behavior and Function Using Computer Simulation Techniques. PhD Thesis, Department of Computation and Neural Systems, California Institute of Technology, 1990.
- [21] M. A. Wilson and J. M. Bower. Cortical Oscillations and Temporal Interactions in a Computer Simulation of Piriform Cortex. *Journal of Neurophysiology*, 67(4):981–995, 1992.

Chapter 8

Conclusions

In this chapter I will summarize the main results that have been obtained from the piriform cortex model, discuss the dynamical and computational implications of these results, and point out the many areas of continuing exploration, both experimental and computational.

8.1 Summary of simulation results

A realistic computer simulation model of mammalian olfactory (piriform) cortex was constructed. This model extends the previous simulation work of Matthew Wilson and James Bower [31, 32]. The new model incorporates much more anatomical and physiological detail than the Wilson/Bower piriform cortex model, and is the most accurate model of this brain structure that has thus far been constructed. The new model features simulated pyramidal neurons and inhibitory interneurons which accurately replicate experimental intracellular

responses to current injection. The model also incorporates patterns of synaptic connections based on the most current anatomical and physiological data. The model replicates large scale features of the system dynamics, including the response of the system to brief weak and strong electrical shocks of the lateral olfactory tract (LOT). The next sections summarize the modeling results.

8.1.1 Strong shock response: influence of background synaptic inputs

The notion that background synaptic activity can have a significant effect on the integrative properties of neurons has been explored by many researchers in both experiments and simulations, including Bernander and Koch [3], DeSchutter and Bower [6], Pare and Destexhe [7, 27] and Walter Freeman [10]. Background spiking activity can cause significant membrane depolarizations, reduce input resistances, and increase fluctuations in somatic membrane potentials. Walter Freeman looked at background inputs in the context of experiments on piriform cortex and demonstrated the need for these inputs in the production of weak shock evoked cortical oscillations [10].

The piriform cortex model is able to accurately replicate the response of the piriform cortex to a strong electrical shock stimulus to the LOT, both from the standpoint of the surface field potential and current-source density (CSD) responses. The model described in the previous chapters has provided insights into the possible role of background spiking activity in olfactory bulb mitral cells (which project directly to piriform cortex pyramidal

neurons) in terms of its effects on piriform cortex, and specifically on the strong shock response. Background spiking in these inputs to the model induces a gradient of excitability in piriform cortex, with pyramidal neurons in ventral anterior piriform cortex (vAPC) being most strongly excited. Without background excitation, neurons in vAPC are too unresponsive to fire even in response to a strong shock stimulus. Increasing the strength of shock-driven afferent inputs cannot compensate for the loss of background inputs because the absence of background inputs results in a prolonged delay before pyramidal neurons can fire in response to shock inputs. Background spiking was not needed to replicate the strong shock response in the Wilson/Bower model, because the neurons used in that model were not constrained to have realistic behaviors in terms of spiking and dendritic integration. Thus, adding a new level of detail to the simulated neurons puts a significant constraint on the network model which was not previously apparent. Application of the neuromodulator norepinephrine leads to runaway spiking in the cortical model unless the mitral cell background spiking rate is decreased by approximately 50%. Experimental data supports this mechanism.

The most detailed experimental study to date on the strong shock response, that of Ketchum and Haberly [20], did not investigate the role of background inputs in this response. Ketchum and Haberly's modeling study of the strong shock response [21] also did not address this question. However, in their model intracortical spikes were imposed upon a passive cable model meant to represent a typical pyramidal neuron. When the model neurons must

themselves generate spikes, as is the case in the network model described here, background inputs are necessary to bias neurons in the ventral anterior piriform cortex into a regime where they can spike readily in response to inputs.

8.1.2 The strong shock response: effects of feedforward inhibition

Feedforward inhibition was found to be necessary in reproducing the strong shock response. Without feedforward inhibition, excitation spreads across the full extent of the network, whereas experimental data suggests that only neurons in the vAPC fire in response to strong shocks [20]. In contrast to the present work, feedforward inhibition was not needed to damp out the strong shock stimulus in the Wilson/Bower model; instead, a strong shock led to activation of neurons across the entire cortex. This was reasonable given the data available at that time but is no longer consistent with more recent data.

8.1.3 Significance of the strong shock response

The question of the significance of the strong shock response is interesting. Given that most or all of the LOT input fibers in the model (and presumably *in vivo*) generate spikes in response to a strong shock, it is highly unlikely that the response of the system to these shocks has any computational significance. Instead, the strong shock response is more likely to be a consequence of mechanisms for protecting the piriform cortex from injury. Piriform cortex, and especially the deep layer of piriform cortex, is known to be highly susceptible

to epileptic activity [5, 8, 9]. In the model, feedforward inhibitory neurons fire in response to a strong shock stimulus but not in response to a weak shock. The model suggests that the activation of high-threshold feedforward inhibitory neurons prevents the majority of the cortical pyramidal neurons from spiking in response to a strong shock, which in turn would prevent the initiation of epileptic activity in response to prolonged and excessive afferent activity.

8.1.4 The weak shock response, oscillations, and network connectivity

A number of mechanisms have been proposed for the genesis of oscillatory activity in various brain regions by Walter Freeman [11], Wilson and Bower [32], Roger Traub and his collaborators [19, 30], Mircea Steriade and his collaborators [25], and others. Most of these mechanisms involve oscillations produced by the intrinsic circuitry of the region. Several such mechanisms have been described, including recurrent connections between excitatory neurons [32], recurrent connections between excitatory and inhibitory neurons [11], mutual inhibition between populations of inhibitory neurons [30], and/or intrinsic neuronal oscillators driving the rhythm [25].

In piriform cortex, oscillatory activity, as reflected in surface field potentials and the CSD response [10, 22], is elicited by weak electrical shock stimulation of the LOT. Until recently, the generally held view of oscillations in the piriform cortex has been that they are due either to recurrent connections between excitatory and inhibitory neurons [11] or

to a combination of recurrent excitatory connections and recurrent excitatory/inhibitory connections [32]. However, more recent experimental data from Ketchum and Haberly [20] suggest that oscillations in piriform cortex are driven largely by oscillations in the olfactory bulb.

Two versions of the piriform cortex model were built to explore this phenomenon. One version of the model (the random connectivity model) features random connectivity between the olfactory bulb and cortex and within piriform cortex, and also features strong feedback connections between posterior and anterior piriform cortex. This model, like the Wilson/Bower model, is able to reproduce the weak shock evoked oscillatory surface potential, and exhibits a natural frequency of oscillation at 50 Hz . However, this model was unable to replicate the CSD response to weak shocks. The second version of the model (the structured connectivity model) features highly structured afferent connections between olfactory bulb mitral cells and piriform cortex pyramidal neurons, as well as highly structured associational connections between pyramidal neurons within piriform cortex (figures 7.26 and 7.27). Specifically, neurons in both the bulb and cortex were divided up into non-overlapping neuronal groups, and connections were allowed only between neurons of corresponding groups. This model replicated the main features of the weak shock CSD response. The model also predicted that successive gamma (50 Hz) cycles of activity in piriform cortex are mediated by distinct sets of neurons. The modeling work does not rule out the possibility that some component of the oscillatory activity of piriform cortex

is internally generated; in fact, the work of Freeman [10] supports this view. The possible computational consequences of highly structured connectivities in piriform cortex are discussed in the next section.

8.2 Computational implications

The two models of piriform cortex used to explore the weak shock response suggest different computational strategies. The connection pattern in the random connectivity model is reminiscent of an autoassociative memory or attractor neural network (ANN) [2, 17, 18]. ANN models typically feature broad, diffuse, and unstructured connectivity between large populations of neurons in the network as well as broadly distributed inputs to the network. The computational power of such networks reside not in the presence or absence of connections between neurons but in the specific weights assigned to each connection (although these concepts are identical in the limit of zero weights). This idea has been postulated several times as a possible computational paradigm for piriform cortex [13, 14, 15]. However, our simulation results show that a model with highly structured connectivities provides a much better match to the weak shock CSD data. This in turn suggests that the paradigm of the piriform cortex as a biological analog of an ANN needs to be revised.

The computational paradigms suggested by the structured connectivity model fall into two categories depending on the extent of feedback from posterior to anterior piriform cortex. Without this feedback, the system would function as a primarily feedforward sys-

tem, whereas with feedback the system might resemble an ANN in which the attractor is a sequence of patterns rather than a single pattern [16]. In either case, the system is better described as a collection of parallel semi-independent subnetworks (one subnetwork per group) which interact with each other primarily through inhibitory connections rather than as a single homogeneous network. Ketchum and Haberly, who performed the CSD experiments on piriform cortex involving weak shock stimuli [22], recognized that their data suggested that oscillations in piriform cortex are primarily driven by oscillations in the olfactory bulb as opposed to endogenously generated oscillations. However, they apparently did not realize the implications of their data on the connection structure of piriform cortex.

Other work, both experimental and modeling, has suggested that olfactory codes may have multiple temporal as well as spatial components. Evidence from the locust [23] suggests that an encoding strategy similar to that suggested by the cortical model described here may exist in insect olfactory systems. In that system, odor information appears to be encoded as an evolving sequence of synchronized oscillatory assemblies. The neuronal activation sequence shown in the current piriform cortex model is also somewhat reminiscent of the Granger/Lynch model of the olfactory system [1], which described the olfactory system as a hierarchical clustering network. In their model, the activation of different neuronal groups each gamma cycle is a function of feedback from piriform cortex to olfactory bulb. The models described in this thesis cannot assess the plausibility of this idea, since they do not include cortico-bulbar feedback. This remains an interesting avenue for future work.

The notion that neurons in the olfactory bulb and cortex are divided into non-overlapping groups raises a large number of questions. If such groups exist, a fundamental question is how many neuronal groups are present. In the modeling work described above only four groups were used; however, this was mainly a reflection of the small scale of the model, and a much larger number of neuronal groups might exist in the real system. If so, a particular odor might be encoded not by which neuronal group is activated, but by which set of neuronal groups are co-activated. One intriguing possibility is that there may be one group for each glomerulus in the olfactory bulb, giving a total of roughly 1000 groups. This could provide an enormous computational substrate for odor encoding; for example, if for a given odor ten distinct glomeruli are activated on each gamma cycle, the total number of possible odors encodable by this strategy would be (assuming five gamma cycles of activity and no overlap in group identities between cycles)

$$n_{odors} = \prod_{i \in (1000, 990, 980, 970, 960)} \binom{i}{10}, \quad (8.1)$$

or 4.6×10^{116} different odors. In actual fact, the real number is almost certainly substantially lower, since other variables besides odor identity (*e.g.*, odor concentration) need to be encoded as well. In this regard, it is of interest that different odors are known to activate different subsets of glomeruli [28, 29], and that increasing the concentration of an odor recruits additional sets of glomeruli [28]. It is also of interest that removing the ol-

factory bulb from neonatal mice causes the olfactory receptor neurons to project directly into piriform cortex, where they induce the formation of glomerular-like structures [12]. This implies that the substrate for glomerular organization may exist in the piriform cortex as well as in the olfactory bulb, which is consistent with the arguments presented above. However, our current models do not address the issue of whether neuronal groups are set up developmentally, are adaptively set up in response to inputs, or both. Yet another important question is the degree of overlap between the responses of neurons of different groups. In the extreme, there would be no connections whatsoever between neurons of different groups. Alternatively, some interconnectivity may exist, especially between groups that are frequently co-activated.

8.3 Future directions

8.3.1 Experimental work

A crucial aspect of the future development of the piriform cortex modeling effort is the continued acquisition of new experimental data. Here we will emphasize the data which is most necessary for improving the model in the short term.

If the piriform cortex is organized along glomerular lines, which is suggested by the modeling work described above, it would be desirable to do experiments wherein a single glomerular region of the olfactory bulb is stimulated while recording intracellularly from

pyramidal neurons in piriform cortex. The theory would predict that a given pyramidal neuron would respond very strongly to the stimulation of a small number of glomeruli, and would respond very little to stimulation of other glomeruli. In addition, the theory suggests that a given pyramidal neuron will be unlikely to spike on consecutive gamma cycles during a weak shock stimulus, which can also be tested experimentally. These experiments are currently under way in our laboratory.

Although there is evidence for nonrandom connectivity between olfactory bulb and cortex [4, 26], no such evidence exists for connections within the piriform cortex. One prediction of the model is that individual pyramidal neurons project to a comparatively small number of other pyramidal neurons which also receive input primarily from the same subset of mitral cells. This could be tested using standard anatomical methods.

In order to improve the models of single neurons, more experimental data will need to be obtained. This is particularly important for feedforward inhibitory interneurons, for which no intracellular data currently exists. In addition, detailed information about dendritic structure and active dendritic currents in pyramidal neurons in piriform cortex is needed in order to improve the pyramidal neuron model.

8.3.2 Modeling work

There are many ways in which the piriform cortex model can be extended. These divide naturally into improvements in modeled neurons and network-level improvements.

The feedforward interneuron models in the network need to be re-parameterized once intracellular data becomes available. The pyramidal neuron model needs to be re-examined in order to understand why the model generates large extracellular ionic currents which are not seen in experimental CSD results. We are also interested in several neuron features which were left out of the current pyramidal neuron model, including active dendritic currents, dendritic spines, and spike-initiating zones. These may affect not only the CSD results but also the neurons' sensitivity to afferent inputs. At the network level, the possible roles of feedback from posterior to anterior piriform cortex needs to be investigated in the context of the neuronal group hypothesis presented earlier.

Some work on modeling the background surface EEG response of the system has already been done. In experiments, this response shows a sequence of oscillations at the gamma frequency modulated by a much slower theta rhythm (4-7 Hz). A key question is whether this modulation is generated inside the piriform cortex, as suggested by previous modeling work [32] or imposed externally. Preliminary results suggest that this modulation is likely to be external, in contrast to previously accepted views; however, much work remains to be done to establish this conclusively.

There are numerous synaptic phenomena which have not yet been explored systematically in the current model. These include synaptic facilitation, synaptic depression, and synaptic long-term potentiation (LTP). The first two phenomena have been included in the model, but no attempt was made to investigate their computational or dynamical effects

on the system. LTP has not yet been incorporated into the model; however, it will almost certainly need to be included in future versions of the model in order to create a complete theory of how the system processes information and in order to understand how neuronal groups form and/or change with experience. Neuromodulation by norepinephrine and its effect on weak and strong shock responses has been explored to some extent, but more work remains to be done on this subject as well.

Finally, the neuronal group hypothesis could be investigated in the context of more abstract models of cortical function, in order to better understand the possible computational advantages that may result from such an arrangement.

8.4 The value of computer modeling

New experimental techniques in neurobiology have been likened to new kinds of microscopes; they allow scientists to probe more deeply into the structure of the brain and collect information on new details of the system [24]. In contrast, realistic computer modeling of neural systems is a kind of “macroscope” which allows scientists to test the consistency of facts obtained at different levels. This process often gives rise to new and unexpected insights and suggests new experiments which would not have been considered worth doing in the absence of the model. Although realistic modeling of large neuronal networks is extremely difficult and time-consuming, I believe that no other technique has as much potential for integrating the diverse aspects of neurobiology into a coherent framework. The continued

development of realistic models of neuronal networks and tools for network modeling, such as the work described in this thesis, will be essential in order to understand how brains process information.

Bibliography

- [1] J. Ambros-Ingerson, R. Granger, and G. Lynch. Simulation of Paleocortex Performs Hierarchical Clustering. *Science*, 247(4948):1344–1348, 1990.
- [2] D. J. Amit. *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press, Cambridge, 1989.
- [3] O. Bernander, R.J. Douglas, K.A.C. Martin, and C. Koch. Synaptic Background Activity Influences Spatiotemporal Integration in Single Pyramidal Cells. *Proceedings of The National Academy of Sciences of The United States of America*, 88(24):11569–11573, 1991.
- [4] N. Buonviso, M. F. Revial, and F. Jourdan. The Projections of Mitral Cells from Small Local Regions of the Olfactory-bulb – An Anterograde Tracing Study Using PHA-L (Phaseolus-Vulgaris Leukoagglutinin). *European Journal of Neuroscience*, 3(6):493–500, 1991.

- [5] R. Demir, L.B. Haberly, and M.B. Jackson. Characteristics of Plateau Activity During the Latent Period Prior to Epileptiform Discharges in Slices from Rat Piriform Cortex. *Journal of Neurophysiology*, 83(2):1088–1098, 2000.
- [6] E. DeSchutter and J. M. Bower. An Active Membrane Model of the Cerebellar Purkinje Cell II. Simulation of Synaptic Responses. *Journal of Neurophysiology*, 71:401–419, 1994.
- [7] A. Destexhe and D. Pare. Impact of Network Activity on the Integrative Properties of Neocortical Pyramidal Neurons in Vivo. *Journal of Neurophysiology*, 81(4):1531–1547, 1999.
- [8] J. Doherty and K. Gale. Evoked Epileptiform Discharges in the Rat Anterior Piriform Cortex: Generation and Local Propagation. *Brain Research*, 861(1):77–87, 2000.
- [9] U. Ebert, P. Wlaz, and W.G. Loscher. High Susceptibility of the Anterior and Posterior Piriform Cortex to Induction of Convulsions by Bicuculline. *European Journal of Neuroscience*, 12(12):4195–4205, 2000.
- [10] W. J. Freeman. Effects of Surgical Isolation and Tetanization of Prepyriform Cortex in Cats. *Journal of Neurophysiology*, 31:349–357, 1968.
- [11] W. J. Freeman. *Mass Action in the Nervous System*. Academic Press, New York, 1975.

- [12] P.P.C. Graziadei and G.A.M. Graziadei. Neuronal changes in the forebrain of mice following penetration by regenerating olfactory axons. *Journal of Comparative Neurology*, 247(3):344–356, 1986.
- [13] L. B. Haberly. Neuronal Circuitry in Olfactory Cortex – Anatomy and Functional Implications. *Chemical Senses*, 10(2):219–238, 1985.
- [14] L. B. Haberly and J. M.. Bower. Olfactory Cortex – Model Circuit for Study of Associative Memory? *Trends In Neurosciences*, 12(7):258–264, 1989.
- [15] M. E. Hasselmo. Acetylcholine and Learning in a Cortical Associative Memory. *Neural Computation*, 5(1):32–44, 1993.
- [16] J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Redwood City, CA, 1991.
- [17] J. J. Hopfield. Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proceedings of the National Academy of Sciences – Biology*, 79(8):2554–2558, 1982.
- [18] J. J. Hopfield. Neurons with Graded Response have Collective Computational Properties Like Those of 2-state Neurons. *Proceedings of the National Academy of Sciences – Biology*, 81(10):3088–3092, 1984.