Irregularity in the Cortical Spike Code: Noise or Information?

Thesis by

William Russell Softky

In Partial Fulfillment of the Requirements
for the degree of

Doctor of Philosophy

California Institute of Technology
Pasadena, California

1993

(submitted February 12, 1993)

Dedication

This manuscript is dedicated to my wife Deborah, whose patience , sympathy, and ready ear made it possible.

Acknowledgements

I am most grateful to C. Koch for advice, encouragement, proofreading, and Devil's-advocacy. I am very grateful to J. Knierim, D. Van Essen, and W. Newsome for their provision of hard-won data for this analysis, and to Ö. Bernander for making his simulated neuron available to me. Special thanks are due to R. Douglas, Bartlett Mel, Ernst Niebur, Idan Segev, Wil Rall, John Rinzel, Ken Miller, Wyeth Bair, Gary Holt, and Marius Usher for many productive discussions and suggestions. I also thank the late D. H. Perkel, who pioneered these investigations and introduced my parents. The invaluable approach of approximate solutions was inspired by an "Order of Magnitude Physics" class taught by Profs. S. Phinney and P. Goldreich. The research reported here was funded by a NSF Presidential Young Investigator Award, by the Office of Naval Research, by the HPS Program in Strasbourg, and by the James S. McDonnell Foundation.

Abstract

How random is the discharge pattern of cortical neurons? We examined recordings from primary visual cortex (V1) and extrastriate cortex (MT) of awake, behaving macaque monkey, and compared them to analytical predictions. We measured two indices of firing variability: the ratio of the variance to the mean for the number of action potentials evoked by a constant stimulus, and the rate-normalized Coefficient of Variation (C_V) of the interspike interval distribution. Firing in virtually all V1 and MT neurons was nearly consistent with a completely random process (e.g., $C_V \approx 1$).

We tried to model this high variability by small, independent, and random EPSPs converging onto a leaky integrate-and-fire neuron (Knight, 1972). Both this and related models predicted very low firing variability ($C_V \ll 1$) for realistic EPSP depolarizations and membrane time constants. We also simulated a biophysically very detailed compartmental model of an anatomically reconstructed and physiologically characterized layer **V** cat pyramidal cell with passive dendrites and active soma. If independent, excitatory synaptic input fired the model cell at the high rates observed in monkey, the C_V and the variability in the number of spikes were both very low, in agreement with the integrate-and-fire models but in strong disagreement with the majority of our monkey data. The simulated cell only produced highly variable firing when Hodgkin-Huxley-like currents (I_{Na} and very strong I_{DR}) were placed on the distal basal dendrites. Now the simulated neuron acted more as a millisecond-resolution detector of dendritic spike coincidences than as a temporal integrator, thereby increasing its bandwidth by an order of magnitude above traditional estimates.

This hypothetical submillisecond coincidence detection mainly uses the cell's *capacitive* localization of very transient signals in thin dendrites. For millisecond-level events, different dendrites in the cell are electrically isolated from one another by dendritic capacitance, so that the cell can contain many independent computational units. This de-coupling occurs because charge takes time to equilibrate **inside** the cell, and can occur even in the presence of **long** membrane time constants.

Simple approximations using cellular parameters (e.g., $R_m, C_m,\ R_i, \overline{G}_{Na}$ etc) can predict many effects of dendritic spiking, as confirmed by detailed compartmental simulations of the reconstructed pyramidal cell. Such expressions allow the extension of simulated results to untested parameter regimes. Coincidence-detection can occur by two methods: (1) Fast charge-equilization inside dendritic branches creates submillisecond EPSPs in those dendrites, so that individual branches can spike in response to coincidences among those fast EPSP's, (2) strong delayed-rectifier currents in dendrites allow the soma to fire only upon the submillisecond coincidence of two or more dendritic spikes. Such fast EPSPs and dendritic spikes produce somatic voltages consistent with intracellular observations. A simple measure of coincidence-detection "effectiveness" shows that cells containing these hypothetical dendritic spikes are far more sensitive to coincident EPSPs than to temporally separated ones, and suggest a conceptual mechanism for fast, parallel, nonlinear computations inside single cells.

If a simplified model neuron acts as a coincidence-detector of single pulses, networks of such neurons can solve a simple but important perceptual problem—the "binding problem"—more easily and flexibly than traditional neurons can. In a simple toy model, different classes of coincidence-detecting neurons respond to different aspects of simple visual stimuli, for example shape and motion. The task of the population of neurons is to respond to multiple simultaneous stimuli while still identifying those neurons which respond to a *particular* stimulus. Because a coincidence-detecting neuron's output spike train retains some very precise information about the timing of its input spikes, all neurons which respond the same stimulus will produce output spikes with an above-random chance of coincidence, and hence will be easily distinguished from neurons responding to a different stimulus. This scheme uses the traditional average-rate code to represent each stimulus separately, while using precise single-spike times to multiplex information about the relation of different aspects of the stimuli to each other. In this manner the model's highly irregular spiking actually reflects information rather than noise.

Contents

List of Figures

Chapter 1

Introduction

1.1 The Basic Questions

It can be mind-boggling to think that everything we see and hear around us is almost flawlessly reflected in the activity of neurons in our brains. Perhaps the foremost scientific and philosophical question of our time is to discovering how several billion tangled and interconnected cells can re-create the world for us, using only signals entering a few small apertures in the head.

There are four major questions in understanding the brain: What does a single cell do with its inputs? How are the cells connected together? How do those connections change with time and input to "learn?" And what task is the brain accomplishing by it all? It is generally thought that only the first of those questions—how a cell works—is very close to being answered. We will challenge the conventional answer, and suggest an alternative view of what a cell might do, how it might do it, and why such a function could be useful for perceptual computation.

This challenge hinges on some simple questions with complex answers: What does a neuron do with its inputs to make an output? How reliable (or noisy) is this process? Which parts of the output signal are essential, and which parts represent irrelevant noise?

1.1.1 The Average Rate Code

It has long been observed that almost every nerve *outside* the central nervous system represents the intensity of a signal by a rate of action potentials (for example stretch-receptors in muscles; see Kuffler *et al.* 1988). Sensory neurons indicate by firing quickly that a stimulus matches their special sensitivity; the quick firing of motor neurons modifies tension in a muscle.

Only more recently has it been shown that neurons in cerebral cortex behave likewise. First came the discovery by Hubel and Wiesel (1962) that the firing rate of cells in striate cortex—the cortical area receiving the most direct input from the eyes—responds best to light/dark contours at particular orientations. Since then, cells throughout the brain have been found to have average firing rates which depend on various particular stimuli. For example, the visual system contains cells responding preferentially to an objects' motion in a particular direction, color, shape, and depth, in various combinations. In all these cases, the experimenter characterizes the cell by electrically recording the average number of spikes it produces in response to various simple stimuli.

The average spike rate evidently represents an analog signal pertaining to the stimulus. This analog code is essential to most mathematical theories of the brain, in great part because our mathematical tools (e.g., linear filter theory and differential equations) deal best with real numbers, and because no one has found a way of predicting or using the occurrence of single spikes.

1.1.2 Why Noise Isn't Investigated

But in real neurons, that analog code is contaminated with irregularity in the timing of individual spikes. No neuron anywhere can fire in a *perfectly* regular manner, if only because of the thermal nature of the chemical and electrical interactions which cause spiking. In general, the most regularly firing neurons are sensory and motor ones, while cortical cells fire in a far more irregular fashion. But while a few researchers have quantified this spiking "noise," virtually none have successfully accounted for that noise on the basis of a cell's inputs and input-output characteristics (see Chapter 2). This omission has occurred for both scientific and cultural reasons.

It is usually difficult enough to record the spikes from a neuron. But it is practically impossible to record also the numerous chemical and electrical inputs to a neuron, especially if those inputs come from thousands of disparate and inaccessible sources (as occurs in cortex). In addition, many of the most basic attributes of neurons—such as their strength and time-course of synaptic conductances and their non-linear membrane properties—remain in dispute. So attempts to explain firing irregularity usually founder on a lack of essential data.

Some cultural gaps have also contributed to this impasse. In my opinion, biologists are so accustomed to instrumental noise and to the enormous complexity of living things that neural firing irregularity may appear perfectly normal, an inevitable part of nature's unpredictability. In addition, many biologists are put off by the hard-nosed mathematical analysis necessary to understand stochastic processes, and are disinclined to trust the predictive power of mathematically formulated theories. After all, unlike in physics, biology has very little tradition of successful predictions by pure theory.

On the other hand, there *is* a thriving community of mathematicians who study neural noise (see the tome by Tuckwell 1989, and a long chapter in Jack *et al.* 1983). But that community has usually emphasized formal solutions over predictive power. Remarkably simplistic neural models are treated to exhaustive formal analysis, without relating their parameters or their results to real systems (an egregious example is a warning in one paper that "it is necessary to avoid hasty identification, for instance, between the time constant of the model and the time constant of the cell membrane," Angelini *et al.* 1982). Approximations are shunned. And the resulting exquisitely intricate formulae appear irrelevant to laboratory biologists, so that virtually no communication occurs between those who understand the living systems and those who write and solve the equations.

1.1.3 Why "Noise" Might Be Information

The analysis of the sources of firing variability has remained a backwater, in part because of its difficulty and in part because it is thought to be irrelevant to information processing.

Please note: irregularity *can be* important in computation. For example, some computations are explicitly statistical, such as the Boltzmann Machine's replicating of its input statistics (Hinton and Sejnowski 1986); some use noise to linearize otherwise pathalogical filters (Knight 1972); some use noise to explore a weight space (Mazzoni *et al.* 1991). But in these and other cases, the source of variability is usually modelled by a "random number generator" rather than by an explicit mechanism.

Unfortunately, the perceived irrelevance of the "noise" source rests on shaky assumptions about the nature of information and about the nature of perception. If one *assumes* that the average spike rate is the only important characteristic of a cell's firing, then of course any variation in it is "noise." But without that assumption, one can note that firing variability *broadens* the bandwidth of the neuron's output (the Fourier transform of a Poisson process is flat, or "white noise"; the transform of a regular spike train is a single pure frequency, along with "harmonics" at its integral multiples). So an irregularly firing neuron might at least in principle be using that extra bandwidth to carry information, with the millisecond timescale of individual action potentials and excitatory synapses giving it a kilohertz bandwidth.

Distinguishing between noise and information requires knowing where the irregularity comes from. The unpredictable nature of individual spikes might represent information about some events of no perceptual significance (such as thermal fluctuations), or it might represent a perceptual code we have not yet fathomed...we can never be sure it is not a code until we understand its source. Unpredictability by itself does not imply noise: as Carver Mead has pithily said, "One man's entropy is another man's information."

What might that information be *about*? The simple average-rate scheme correlates the neuron's output exclusively with its particular stimulus, but ignores the possible *relations* between separate stimuli, and between neurons. The simple, traditional experiments present a single isolated stimulus and record from a single, isolated neuron. But visual perception requires not only that we detect thousands of isolated contours and motions, but that we make sense of them, ignoring some and and relating others in order to represent objects and patterns. True perception involves the *inter-relationships* between primitive stimuli.

Those inter-relationships are not present in most simple experimental stimuli, and are apparently not present in the average-rate response. (There are controversial exceptions: the stimulus-induced firing-rate oscillations of Gray and Singer (1988), and the stimulus-linked variations in average-rate envelope of McClurkin *et al.* 1991). The possiblity suggested here is that the relationship of a primitive stimulus (e.g., a contour) to other stimuli is reflected in correlations between cells' individual *spike* times, while the cells' average rates carry information about the respective stimuli alone. But this multiplexing hypothesis can only be tested by presenting multiple, preceptually relevant stimuli and simultaneously recording from multiple neurons.

1.1.4 "Looking for patterns in all the wrong places..."

Researchers have searched for patterns in the "noise" of spiking irregularity, but without evident success. The late Don Perkel, a physicist who helped introduce nummerical spike analysis techniques to neurobiology (and who introduced my parents to each other), had speculated since the sixties about finding a "code" buried in the seemingly random trains of neurons, and of Geiger counters. Abeles (1980, 1990) has found and reviewed many cases of non-random signals in spike trains. In fact, there have been many other searches for recurring patterns in long trains of spikes from a single neuron (e.g., Strehler and Lestienne 1986; Legendy and Salcman 1986). As Abeles (1982) has noted, single-neuron data is relatively easy to obtain, because once a single neuron is located with a recording electrode, one can easily obtain thousands of spikes from it. With so much data, it is easy to try correlating the timing of a single neuron's spikes with either the stimulus onset or with other spikes from the same train.

But while such attempts occasionally reveal slightly-above-random spike patterns (Abeles 1990 and references therein), there are several reasons why those temporal patterns are probably *not* a significant code. First, the patterns found are faint, buried in a sea of "random" spikes. Second, no single neuron in cortex is thought to have a sufficiently long or flexible "memory" to decode a complex series of spikes. Also, the perceptual task is to relate a neuron to *other* neurons (correlation across space in short times), not to correlate a single neuron's output across a long time. Finally, the system as a whole cannot be expected to keep track of individual spikes for very long, because new input spikes are always coming in to replace the old ones.

The alternative is to compare firings of *different* neurons with each other. That is harder, because while recording from one is difficult, recording from many at once is much more so. Fortunately, the motivation and technology for such multi-neuron recordings are increasing dramatically, and have already yielded evidence of strong correlations at timescales from the submillisecond realm up to hundreds of milliseconds. The most general task is to understand the significance of those correlations for cortical information processing. The narrower task of this thesis is to explore the cellular mechanisms which may make single-spike computation possible.

1.2 Thesis Overview

Chapter 2 forms the foundation of this thesis, demonstrating a striking contradiction between the most well-accepted theory of cortical cell function and the well-known irregularity of cortical firing. This work has already appeared as a short note in *Neural Computation* (Softky and Koch 1992) and recently as a long article in the *Journal of Neuroscience* (Softky and Koch 1993).

For most cortical cells, only the mean firing frequency is reproducible under identical stimulus conditions. Because the fine time-structure of the irregularities is not reproducible, it is widely assumed that information is only carried in the average spike frequency; the fine time structure is usually assumed to be irreproducible "noise." Some electrophysiologists have focussed on the idea that the dynamics of the neuronal response may carry significant information (McClurkin *et al.* 1991; Abeles 1990; Aertsen *et al.* 1989). We here do not directly address that issue. Rather, we measure the degree of firing irregularity in cortical cells in the behaving monkey and investigate the possible neuronal sources of the high degree of observed variability. The conclusion is that such cells do not perform a temporal integration or averaging of their excitatory inputs, but rather exhibit a striking sensitivity to input fluctuations at fine timescales.

Chapter 3 takes that conclusion as an indication that cortical cells might compute with single spikes rather than with average rates. We examine in detail one hypothetical situation—the presence of spiking mechanisms on the thin, remote branches of cortical cells—in which cells might perform as high-fidelity coincidence-detectors of single input events at the submillisecond scale, without performing temporal integration of those inputs.

This approach flies in the face of much tradition. Decades of electrophysiology have produced no reliable evidence that information about a single stimulus is carried by precise spike-times in most cortical areas. Only temporal averages over much longer timescales (at least 20 ms; McClurkin *et al.* 1991) correlate with stimuli. Furthermore, it is generally believed that cortical cells are fundamentally incapable of using such millisecond-resolution information, due to their relatively longer membrane time-constants (10-30 ms) and to attenuation of high-frequency signals by their dendrites (Douglas and Martin 1991).

While some of that chapter's hypotheses are made in the absence of any experimental data to support or undermine them, the simulations and analyses are at least consistent with the relevant measurements already done on single cells—it is possible for a cell to contain those postulated spiking mechanims so that their electrical effects are nearly masked from the cell body by distributed capacitances. In fact, those same capacitive cable properties help isolate the cell's different branches from each other at fast timescales, allowing it to contain many independent, fast subunits.

An additional thrust of Chapter 3 is the introduction of some approximation techniques to supplement the brute-force numerical simulation of the model cell's differential equations. These approximations use fundamental properties of the model's geometry and electrical responses to successfully predict the simulation results **without** any free parameters ("fudge factors"), thus revealing both the primary mechanisms at work and the scaling properties of those mechanisms.

The final part, Chapter 4, constructs a Gedanken-network out of conclusions from the previous two chapters. In this network cells use single-pulse coincidence-detection to carry significant information *through* their firing irregularity. This network—while simplified to the point of silliness in a biological sense—nonetheless manages to solve an outstanding perceptual problem by multiplexing information between the the cells' average firing rates and their individual spike times. Such multiplexing would preserve the ability of single neurons to represent features by analog firing rates, but would in addition use the full kilohertz bandwidth of each neuron to link neurons together in a primitive form of perception called "binding" (Engel *et al.* 1992), while retaining the apparently random character of each neuron's output as recorded in isolation.

1.3 Cortical Physiology Oversimplified

A neuron in the cerebral cortex is a small bag of saltwater, shaped like a tree (including roots). It is surrounded by more salt water and by other neurons, some of which it is connected to. It produces as "output" electrical pulses in response to input pulses from other neurons, and those output pulses travel onward to other neurons.

Because this research focusses on the behavior of a single neuron in response to its inputs, it is necessary to review the most elementary properties of a neuron; unfortunately, the intricate and ever-changing pattern of their interconnections is well beyond this discussion.

The cell body is called the *soma*, and the branches radiating from it are *den-*

drites. The membrane forming the neuron's skin (the "bag") is about 50 $\mathring{A}$ thick. The capacitance of biological membranes depends on its thickness and dielectric constant, and is constant at about 0.7-1 $\mu F/cm^2$. But while in its purest state that membrane is impermeable to the passage of ions, the presence of small pores in the membrane *can* allow the passage of ions, so that current can flow through it (such as sodium, potassium, chlorine, and calcium, which carry the bulk of the electrical current in neurons). Herein lies the beauty and complexity of a neuron's electrical function.

The precise shape and size of those pores (called "channels") determine which ions pass through them. Because each ion species has a different concentration inside and outside the cell, the potential of an ion species differs from inside to outside, so that there is a separate voltage (or battery) associated with each ion (called E_{rev}, relative to the potential of the fluid outside the cell).

The numbers of open channels determine the amount of current passed. That current is given by the membrane conductance g_i, where i denotes the ion type (or ion mix) and its associated potential (the "reversal potential" E_{rev}, defined relative to the potential outside the cell). While a few membrane conductances (the "leak" conductances) are constant with time and independent of transmembrane voltage, the most interesting and useful ones change with voltage, time, and chemical signals.

For example, two voltage- and time-dependent conductances are activated in creating an output pulse (called an "action potential"), which are named after their discoverers Hodgkin and Huxley. The first of those Hodgkin-Huxley conductances to be engaged is the sodium conductance (g_{Na}), which opens more and more as the cell voltage increases above its usual polarized "resting" voltage of about $-75\ mV$ ($= E_{rest}$). Because there is more sodium outside the cell than inside, the sodium rushes in through the sodium channels when they begin to open and raises the cell voltage, and that depolarization opens the conductances still further, in a postive-feedback loop. (See Appendix G for the differential equations modelled.)

Three effects eventually limit the excursion of this current avalanche. A minor effect is that as the cell depolarizes, the driving potential difference ($E_{Na} - V$) decreases. A much stronger effect is that the sodium channels have a natural time-course, so that in normal circumstances they begin to close after about half a millisecond. The final effect is that the strong depolarization of the cell causes the opening of potassium channels. Because potassium is more abundant inside the cell than outside, it rushes out and repolarizes the cell. This whole process produces a voltage "spike" lasting about a millisecond, which propagates down the output fiber (the "axon") without dispersion, due to the presence of further Hodgkin-Huxley channels along the axon's length.

Another critically important type of channel is found at the "synapse," the location where the output axon of one neuron provides input to another neuron. Some synapses are direct, linear electrical connections between cells, with bi-directional current flow. But the most prominent synapse class in the brain, the chemical synapse, provides the essential function of one-way information transfer: an action potential in the axon can cause an electrical effect in the target cell, but not vice-versa . At the synapse, the arriving action potential pulse triggers the release of chemicals which (in most cases) cause the opening of channels on the target neuron. If the channels opened have a reversal potential greater than the resting potential ($E_{syn} > E_{rest}$), the synapse and synaptic current are called "excitatory," because they depolarize the target cell and make it more likely to fire. The resulting voltage change in the target neuron from such an excitatory synapse is called an EPSP (**E**xcitatory **P**ost **S**ynaptic **P**otential). If the channels have a lower voltage ($E_{syn} < E_{rest}$), the synapse is "inhibitory," leading to an IPSP. In general, excitatory synaptic currents are much briefer than inhibitory ones, and have much stronger driving potentials. Inhibitory synapses can sometimes have reversal potentials so close to resting potential that they act not by driving down the membrane voltage, but by shunting off the excitatory current ("shunting" or "silent" inhibition).

In the study of visual cortical neurons, the neurons' responses are stimulated through visual patters (e.g., bright dots or dark bars) presented in the animal's field of view; the region of the field of view over which a pattern can elicit a neural response is called that particular neuron's Receptive Field (RF). The response is usually analyzed by showing the same pattern many times, and averaging the resulting spike trains—each with the stimulus presentation time as $t = 0$—to form a Post Stimulus Time Histogram, corresponding to the average spike rate after the stimulus appears.

As a reference, a glossary of important terms follows:

action potential—sudden voltage pulse (about 1 ms duration) *adaptation*—slowing-down of output spikes during constant input

axon, axonal—the thin output cable along which a spike propagates

dendrite, dendritic—the thin (1-4 μm) "roots" branching outward

apical dendrite—a single thick (4-7 μm) dendrite, like a tree-trunk, found on pyramidal cells

distal—far away from the soma

proximal—close to the soma

soma, somatic—the cell body (typically a sphere or blob 15 μm wide)

synapse, synaptic—the one-way chemical contact from an axon to another cell

EPSP—**E**xcitatory **P**ost **S**ynaptic **P**otential

EPSC—**E**xcitatory **P**ost **S**ynaptic **C**urrent

IPSP—**I**nhibitory **P**ost **S**ynaptic **P**otential

ISI—**I**nterspike **I**nterval

PSTH—**P**ost **S**timulus **T**ime **H**istogram

RF—**R**eceptive **F**ield

Chapter 2

A Paradox: Cortical cells do not perform temporal integration of small, random EPSPs

2.1 Introduction

When a typical spiking neuron is injected with sufficient current, it fires a regular stream of action potentials. But cortical cells *in vivo* usually fire irregularly in response to a sensory stimulus. What are the cause and function of that irregularity?

The irregularity of action potential discharge has been analyzed using the mathematics of stochastic point processes and their intrinsic variability (Perkel, Gerstein, and Moore 1967; Stein 1967a,b; Lansky and Smith 1989; for a recent overview, see Tuckwell 1989 and references therein). The firing variability of thalamic and cortical spike trains has been studied experimentally (Poggio and Viernstein, 1964; Noda and Adey 1970;Burns and Webb 1976). These and similar studies measured neuronal variability— usually in the form of interspike interval distributions — and characterized that variability using various phenomenological statistical distributions (e.g., hyperbolic normal, gamma distribution etc.). They did not relate the firing variability to the quantitative biophysics of the cells.

One exception is the study by Calvin and Stevens (1968). On the basis of intracellular recordings of cat lumbrosacral motoneurons, they constructed a simple model of the spike generation mechanism. They combined the measured properties of synaptic noise with their model to account for the observed small interspike-interval variability (with an associated coefficent of variability $C_V \approx 0.05 - 0.1$). They concluded that in the majority of neurons they recorded from, synaptic noise was by itself sufficient to explain the observed variability, without invoking any additional intrinsic noise sources.

Our study uses the same starting point, measuring interspike interval histograms and their associated coefficients of variation in the case of extracellular recorded units in primary visual cortex (V1) and middle temporal visual area (area V5 or MT) of the awake behaving monkey. Unlike the lumbrosacral motoneurons, the cortical units have a very high degree of irregularity, with C_V ranging between $0.5 - 1.0$. We attempt to understand the origin of these values by two different theoretical methods: modified integrate-and-fire models, and simulations of detailed compartmental models of cortical pyramidal cells. Our analysis reveals a strong contradiction between the large observed interspike variability at high firing rates and the much smaller values predicted by well-accepted analytical and biophysical single cell models. This contradiction does not exist for high variability at *low* firing rates, which is consistent with the models of Wilbur and Rinzel (1983) and Bugmann (1990). This contradiction has been noted (but not published or resolved) in the case of retinal ganglion cells at high illumination by Barlow and Levick (H. Barlow, personal communication).

2.2 Electrophysiological Data

We used data from two different laboratories. In both cases, extracellular spike trains were recorded from cells in visual cortex of awake adult macaques. Our primary interest was neither in the nature of the stimuli used nor in the cells' selective responses to these stimuli, but only in the statistical properties of neuronal firing.

The first set of data came from primary visual cortex, or cortical area V1. In that region of cortex, neurons respond best to simple contours of a particular orientation presented in the center of small receptive fields (RFs). These recordings came from an investigation of the influence of the larger "non-classical receptive field" on single neuron activity in two alert and behaving macaque monkeys (*macaca fascicularis*; Knierim and Van Essen 1992). Data was only accepted for trials during which the monkey kept its eyes on a fixed target. The cells were stimulated by a variety of flashed bars of various orientation in the center of the classical RF, and in some cases additionally stimulated by either parallel or perpendicular oriented bars outside the classical RF (Knierim and Van Essen 1992). We used 1184 single, well-isolated spike trains of 1 *sec* duration recorded from 16 cells at a temporal resolution of 1 *msec*. Only one of these cells showed any bursting activity (as defined below), and was rejected.

The second set of data (referred to in the following simply as "MT" data) was recorded during an investigation into the relationship between motion discrimination and the behavior of single neurons in area MT (or V5), a region of extrastriate visual cortex concerned with motion processing (*M. mulatta*; Newsome, Britten and Movshon, 1989b; Britten *et al.* 1992). In brief, three monkeys were trained to report the direction of motion of a random dot display in which a fixed fraction of dots (the amount of "motion coherency") moved coherently in one direction while the remainder moved randomly in all directions (Newsome and Pare 1988). The amount of motion coherency as well as the direction of motion was varied across trials. During a single trial, the monkeys, whose heads were restrained, had to fix their gaze on a cross. If fixation was broken—as monitored by a search coil—the trial was terminated.

Standard electrophysiological procedures were used to identify and record single MT neurons in three alert and behaving monkeys (Mikami, Newsome and Wurtz 1986). The two-threshold window discriminator produced pulses corresponding to single action potentials whose time of arrival was recorded with 1 *msec* resolution. Care was taken to record only single neuron activity. Altogether, the activity from 409 neurons was recorded, each trial usually being 2 *sec* long. Figure 2.1 shows a sample spike train, a histogram of all the superposed spike trains (the Post-Stimulus Time Histogram, or PSTH), and an interspike-interval histogram from a typical MT recording.

For our analysis, we used a subset of these trials. We rejected all spike trains that contained any dominant interspike intervals (ISIs) characteristic of "bursting" behavior. A "bursting" neuron frequently fires a pair of action potentials within a short time ($< 1 - 3$ *msec*), a situation characterized by a sharp peak in that range on the ISI histogram. More specifically, we rejected any neuron whose ISI histogram (see below) contained more than twice as many counts in the 2 *msec* bin as in the 5 *msec* bin. These criteria yielded a subset of 233 non-bursty neurons.

In general, we did not find any significant difference between the degree of variability of V1 or MT neurons. Therefore, except when otherwise explicitly noted, we will lump these two sets of experimental data together.

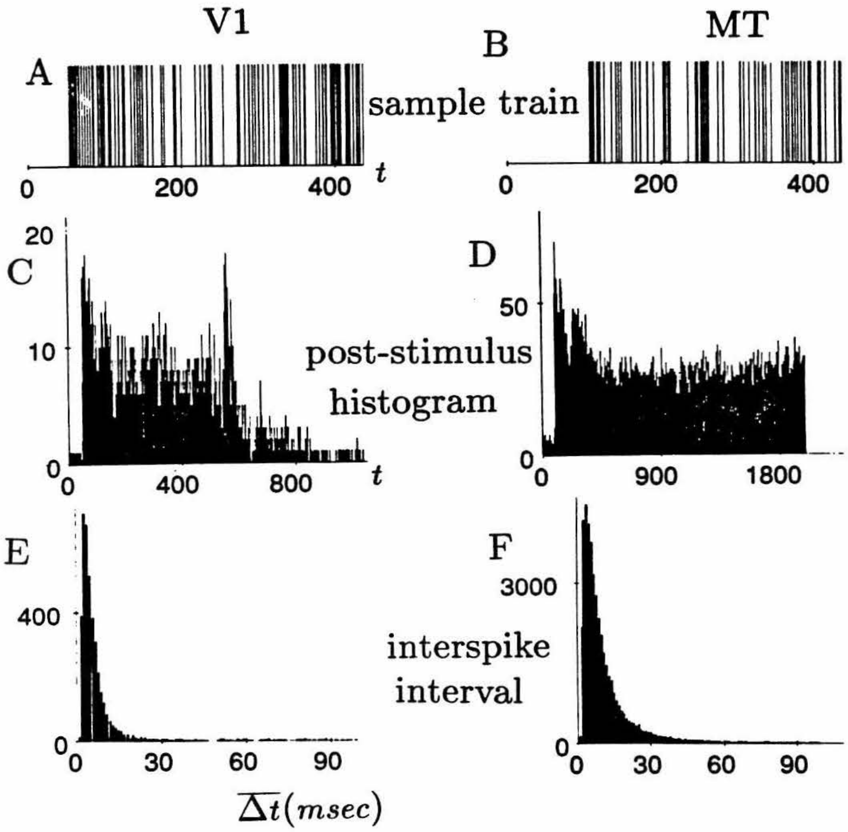

Figure 2.1: **Firing Statistics of Neurons in Areas V1 and MT** (A,B) Sample spike trains from one of the fastest-firing non-bursting neurons recorded in each area. (C,D) Post-stimulus time histograms (PSTH) from the same neuron. (E,F) Interspike-interval histograms from the same neuron. These neurons are "typical," in that their firing times seem nearly random at all observed firing rates.

2.3 Analysis Method

2.3.1 Parameters and Normalization Procedure

The spikes following the stimulus onset arrived at times $\{t_i\}$. Thus the inter-spike interval (ISI) is

$$\Delta t_i = t_{i+1} - t_i. \tag{2.1}$$

We will analyze histograms of these ISIs through two of their parameters. One is the mean of the histogram (the average interspike time $\overline{\Delta t}$):

$$\overline{\Delta t} = \frac{1}{S_j - 1} \sum_{i=1}^{S_j-1} \Delta t_i \tag{2.2}$$

where S_j is the number of spikes in the train. The other parameter is the standard deviation about that mean, which is

$$\sigma_{\Delta t} = \sqrt{\frac{1}{S_j - 1} \sum_{i=1}^{S_j-1} (\Delta t_i - \overline{\Delta t})^2}. \tag{2.3}$$

These two values together yield a measure of the variability of the spike train, the dimensionless **C**oefficient of **V**ariation, which describes the relative width of the ISI histogram:

$$C_V = \sigma_{\Delta t} / \overline{\Delta t}. \tag{2.4}$$

For a very regular spike train ("pacemaker"), the ISI histogram will have a very narrow peak and $C_V \to 0$. In the case of a random spike train (a Poisson process or shot noise), the Δt_i are exponentially distributed and $C_V = 1$. The coefficient of variation can be larger than one in the case of a multi-state neuron (Wilbur and Rinzel 1983). But such neurons would produce ISI histograms with a narrow peak on a broad base, which would be excluded by our "burstiness" measure (above, section 2.2). So this work does not measure or model multi-state neurons.

This analysis could not be applied directly to our data, because both V1 and MT neurons "adapted," in that their firing rates decreased to roughly half the initial value during the first $100 - 300$ *msec*, despite a constant visual stimulus. Moreover, because more than one stimulus was used on each cell, the number of spikes varied significantly between trains. We found that the ratios of pre-adapted and post-adapted firing rates only slightly for different stimuli, because the post-stimulus time histograms for different stimuli all had approximately the same shape.

Because such non-stationary (variable-moment) statistics are difficult to analyze, the goal of the analysis was to arrive at an *approximate* estimate of the "instantaneous" C_V, without artificially broadening the ISI histograms due to the changing mean firing rate. That is, we attempted to eliminate the artificial source of variance induced by adapting rates by separating spikes into many histograms, each representing a roughly constant firing rate. Our method was to compute the approximate instantaneous firing rate R. We then used R to segregate spikes into ten different histograms. The highest R (early times with strong stimuli) binned corresponding spikes in the "fastest" histogram, the lowest R (the tail-end of the weaker stimuli) put spikes into the "slowest" histogram, and intermediate R stored spikes in corresponding histograms in between. The predicted values of R and its resulting range were calculated separately for each cell, as follows.

The major simplifying assumption was that the cell's instantaneous rate at time t during any particular experimental trial j depended only on the total number spikes S_j in that train, and on the cell's average instantaneous response $r(t)$ averaged over all m experimental trials for that particular cell. The instantaneous response $r(t)$ was taken directly from the post-stimulus time histogram of the cell for all m stimuli, coarse-graining t to bins 20 *msec* wide (indexed by $\tilde{t} = 0, 20, 40 \ldots msec$). Thus, if $S_j(\tilde{t})$ is the number of spikes in train j falling in bin $\tilde{t}$, then

$$r(t) \quad = \quad r(\tilde{t}) = \frac{1}{20m} \sum_{j=1}^{m} S_j(\tilde{t}). \tag{2.5}$$

The true instantaneous rate $R_j(t)$ is then assumed to be the product of $r(t)$ and S_j, normalized by S_{avg}, the average number of spikes in a train for that cell (i.e., $S_{avg} = \frac{1}{m} \sum_{j=1}^{m} S_j$):

$$R_j(t) \quad = \quad \frac{S_j}{S_{avg}} \times r(t). \tag{2.6}$$

The S_j term in eq. 2.6 represents the efficiency of the stimulus, i.e., how many spikes the cell fired over the entire recording interval (e.g., 2 *sec* in the case of the MT recordings) in response to a particular visual stimulus, while $r(t)$ describes the time course of neuronal adaptation over all stimuli used for that particular cell. Note that only the parameter S_j was used in the Newsome *et al.* (1989a) study—for which the MT data analyzed here were generated—for the evaluation of neuronal sensitivity and performance.

Each ISI in any spike train for a particular cell was then placed into one of ten different histograms according to its associated R_j value. The maximum rate

R_{max} (defined over all stimulus conditions for that cell) was used to define 10 equally spaced rate intervals from 0 to R_{max} Hz, i.e., $(0 - 0.1)$ R_{max}, $(0.1 - 0.2)$ $R_{max} \ldots (0.9 - 1.0)$ R_{max}. So each cell produced ten ISI histograms, each of which had a temporal resolution of 1 $msec$ and a total range of 100 $msec$ (longer ISIs were not necessary for this analysis of high firing rates).

For each spike train j, both Δt_i and $R_j(t_i)$ were computed from the original data, following eqs. (2.1) and (2.6). Then Δt_i was assigned to the post-stimulus time t at the center of its ISI (i.e., $t = (t_i + t_{i+1})/2$). Finally, Δt_i contributed one count to the appropriate ISI bin in the particular histogram whose rate range included $R_j(t)$. Figure 2.2 illustrates this procedure for a fast-firing V1 cell. Here $R_{max} = 380$ Hz and $S_{avg} = 83$ spikes in one second (83 Hz). The average time course $r(t)$ and $R_j(t)$ (for the fastest-firing train) are shown in Figure 2.2 (a,b). Three of the associated ten histograms, into which a total of 4009 ISI values were placed, are also illustrated. Note that the instantaneous rate $R_j(t)$ was only used to determine into which histogram any particular value of Δt_i should be placed.

Because the histograms with highest R only contained the earliest spikes of the few fastest trains, they typically had far fewer spikes than the intermediate histograms (see Fig. 2.2). Nonetheless, these fast histograms usually contained enough spikes to be statistically significant, judging by the error bars in C_V as calculated below. In addition, the fastest histograms had mean rates typically twice as fast as the cell's average (adapted) response to a strong stimulus over 1-2 sec.

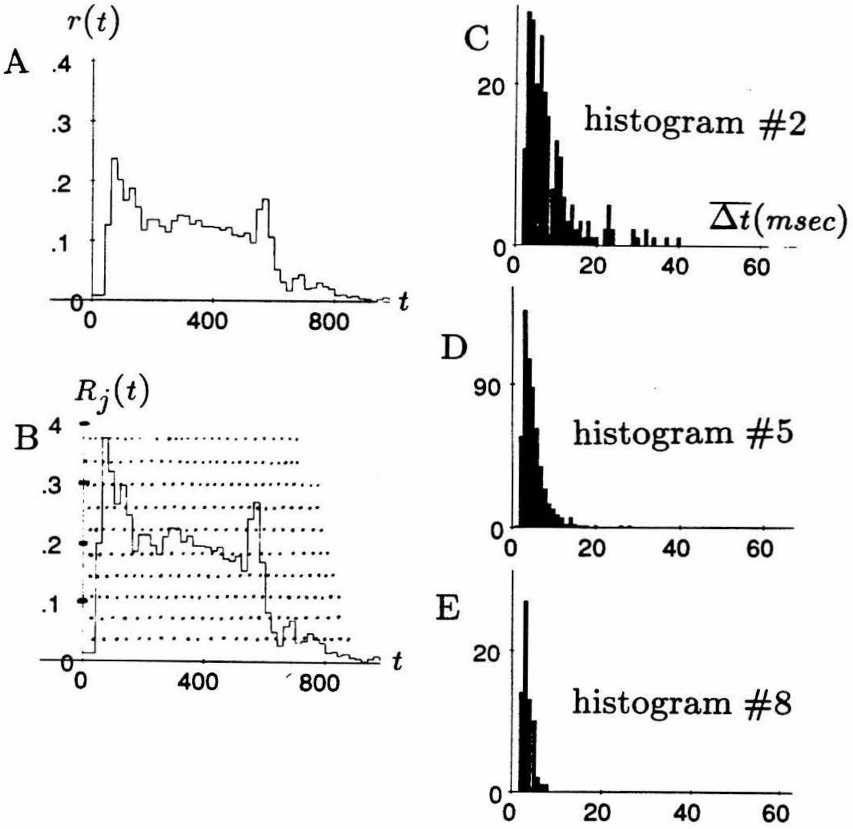

Figure 2.2: **Firing Variability Analyzed by Multiple ISI Histograms.** The firing rate of a cell depended on both stimulus efficacy and post-stimulus time. ISI histograms were made from such data by segregating ISI's according to the approximate instantaneous rate $r(t)$, according to the following steps. (A) $r(t)$ (Hz) was calculated for any particular cell from the PSTH for all the responses of that cell to stimulation. (B) The instantaneous rate $R_j(t)$ for train j was computed by multiplying $r(t)$ by the ratio of the total number of spikes S_j of that train to the cell's average number of spikes S_{avg}. Here, $S_{avg} = 83$ and $S_j = 131$. Each ISI was placed into one of ten ISI histograms, so that each histogram represented a roughly constant firing rate: histogram #0 was slowest, and histogram #9 was fastest. (C,D,E) Three of the ten histograms for the V1 cell of Fig. 1 are shown here (spikes between 50 *msec* and 100 *msec* are not shown here, but were included in our analysis.) These ISI distributions are typical of cortical cells described elsewhere: a virtual absence of ISI's below 2 *msec* indicates the refractory period and the absence of "bursting" behavior, and the distribution is very wide relative to its width. Each histogram's shape-parameter C_V contributed one point in Figs. 2.3, 2.9, and 2.13.

The parameters $\overline{\Delta t}$, $\sigma_{\Delta t}$, and C_V were then calculated from each of the ten histograms (without using $R_j(t)$). Error bars were derived from the counts in individual histogram bins by treating those bin-counts as Gaussian random variables. For example, if $M_{\Delta t}$ counts fall in a single bin Δt, then we assume the uncertainty in $M_{\Delta t}$ is $\sigma_M = \sqrt{M_{\Delta t}}$, and we propagate errors as random variables to get

$$\overline{\Delta t} = \frac{1}{100} \sum_{\Delta t=0}^{\Delta t=100} (M_{\Delta t} \pm \sqrt{M_{\Delta t}}) \Delta t \tag{2.7}$$

and

$$\sigma^2_{\Delta t} = \frac{1}{100} \sum_{\Delta t=0}^{\Delta t=100} (M_{\Delta t} \pm \sqrt{M_{\Delta t}}) (\Delta t - \overline{\Delta t})^2 \tag{2.8}$$

The resulting C_V values were plotted against $\overline{\Delta t}$ (Figure 2.3) for all but the slowest two histograms for each cell (i.e., for all histograms within which the instantaneous rate varied by no more than 33%). C_V values from histograms with less than ten counts were also excluded, so that each cell contributed eight or fewer points to a plot of C_V. This entire normalization procedure was repeated for every one of our 249 cells. Had C_V been calculated only from the total ISI histogram for a single cell (i.e., without using the multi-histogram method), equally high values would have resulted ($C_V \approx 0.7 - 1.1$); but such histograms would have confounded high and low firing rates, and would thus have been difficult to interpret.

2.3.2 Inaccuracy of Analysis Method

The statistics of spike trains are not precisely defined for non-stationary processes. But we are only concerned with the approximate variability of the

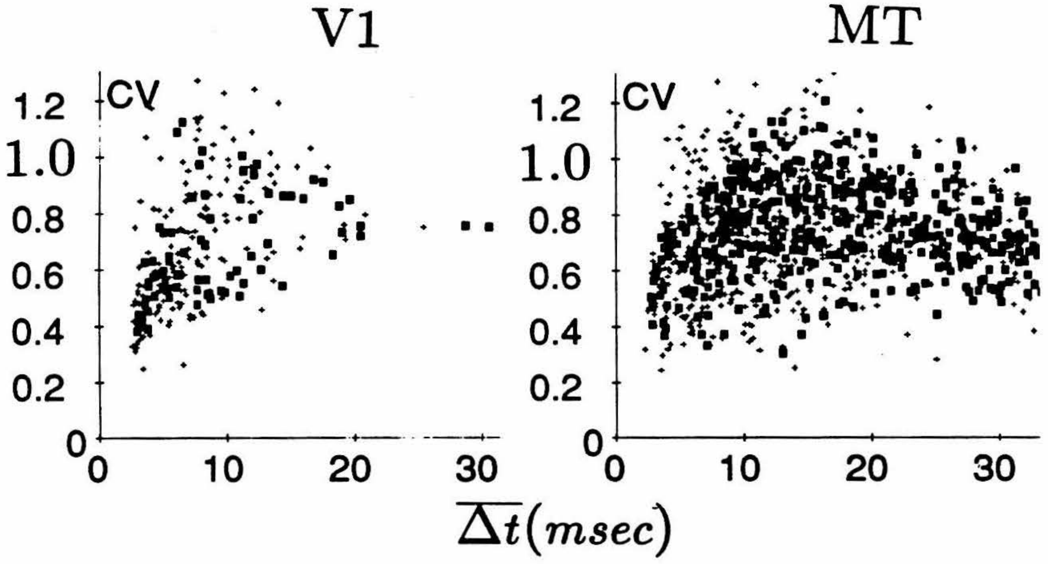

Figure 2.3: **Variability of Neurons in Areas V1 and MT.** C_V characterizes the normalized width of a histogram. The scattered points were obtained from ISI histograms like those in Fig. 2.2 (only points with $\overline{\Delta t} \leq 30$ *msec* are shown). Filled squares are reliable points ($\sigma_{C_V}/C_V \leq 0.1$), crosses are less reliable C_V values. The main systematic bias of the analysis method was to underestimate C_V for large ISI's ($\overline{\Delta t} \geq 20$ *msec*). The slightly higher firing rates of the V1 neurons resulted from the choice of such faster neurons for analysis; no other differences are apparrent between the two areas.

spike train, so let us suppose for the moment that our data represent a simplified process in which each ISI was generated randomly, according to some distribution with fixed C_V and variable rate. Would the analysis method described above reveal the true (generating) value of C_V? We will discuss some of the limitations of this multi-histogram analysis method, and then show a simulation which suggests that our method is indeed suited to our purposes.

The above method underestimated C_V for low firing rates, because some long ISIs were excluded from their proper histograms. Some spike trains were only 500 *msec* long; thus ISIs longer than that duration obviously could not be counted. A more stringent limit was the width of the ISI histogram from which C_V was calculated (100 bins of 1 *msec* each), which truncated the tails of ISI distributions with large $\overline{\Delta t}$ and high C_V (e.g., $\overline{\Delta t} \geq 25$ *msec*). In all these cases, truncating the tail of a broad ISI distribution artificially narrows the histogram, and reducing the estimated C_V below its true value.

In other cases, this analysis *overestimated* C_V. This artificial broadening of the ISI histogram can occur, for instance, when the firing rate changes during the rate-averaging period $\tilde{t}$: a smooth variation in firing rate would be misconstrued as a high random variability. Although this effect obviously occurs during the onset of spike adaptation (in the early part of the PSTH, when the average rate changes most quickly), it can also occur at the *lowest* rates measured for one cell, for which a single histogram has a higher *fractional* variability than at higher rates (e.g., a 60-90 Hz histogram contains 33% frequency variability, vs. 10% for a 270-300 Hz histogram).

A further artifact occurred at high firing rates, when the width of a single time-bin (1 *msec*) becomes comparable to the shortest ISIs observed ($\overline{\Delta t} \approx 2\ msec$). This effect is most pronounced when the true histogram is very narrow and steep-sided, so that the "rounding error" (about 0.5 *msec*) induced by shifting each ISI to a neighboring bin increases the histogram's width significantly.

In order to quantify these combined effects, we numerically simulated spike trains with the following characteristics: 1) each ISI was generated by a gamma probability distribution with constant and known C_V, variable mean rate, and a resolution of 1 *msec*; 2) each train was 500 *msec* long; 3) the mean rate dropped linearly to 0.33 of its starting value within 250 *msec* (modelling adaptation); 4) starting rates for different trains were chosen to give a range of $\overline{\Delta t} \approx 2 - 30\ msec$ (comparable to the monkey ISIs) between the very fastest and very slowest mean ISIs observed. The more variable of these artificial trains looked just like real trains from monkey. At each C_V, 500 simulated trains at different rates were analyzed together by our normalization method described above. In addition, the slowest 100 trains were separately analyzed, to better resolve the slowest ISIs.

The comparison of the C_V values yielded by this analysis with the C_V of the random processes generating the trains (Figure 2.4) confirms the two points outlined above: this method systematically *overestimates* C_V when both C_V and $\overline{\Delta t}$ are low, and systematically *underestimates* C_V when both are high. But for fast-firing, highly variable cells—like those observed in our analysis—this method introduces a systematic bias that is no greater than a few percent.

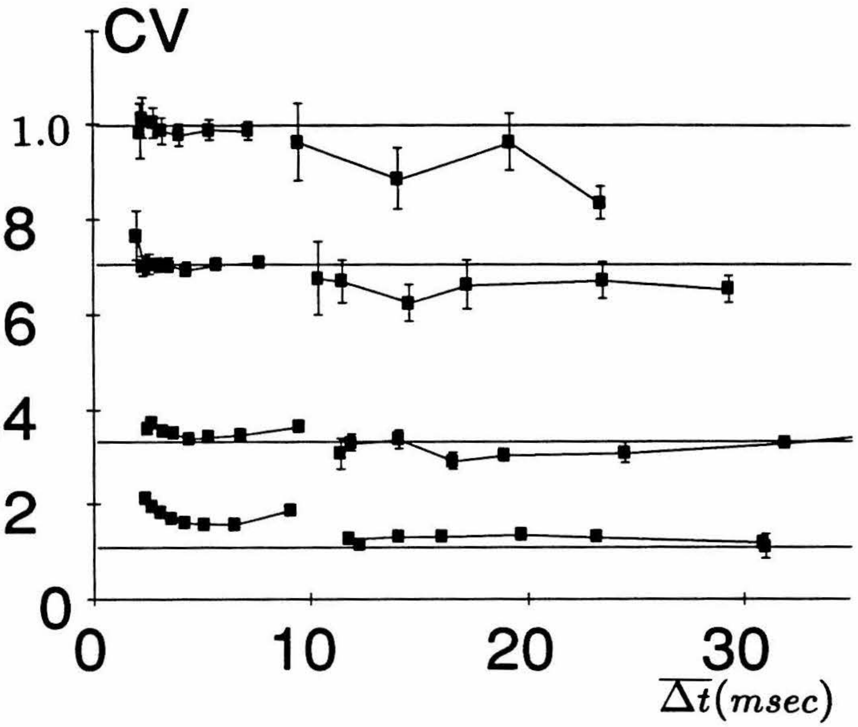

Figure 2.4: **Accuracy of the Multi-histogram Normalization Method.** Computer-generated spike trains (having roughly the same mean firing rates and adaptation course as the monkey data) were randomly generated from gamma-function ISI distributions of various C_V (1.0, 0.71, 0.33, 0.11) with 1 *msec* resolution. We analyzed these fake trains by the same method used for the monkey data, analyzing slow and fast trains separately to resolve C_V at both long and short ISI values. The resulting C_V values (connected squares) were compared with the rate-independent C_V value of the generating distribution (horizontal lines). The analysis method underestimated high C_V at long ISI values, and overestimated low C_V at short ISI values. But high C_V values at short ISI (like those observed in monkey) were not systematically biased more than a few percent. Thus, the drop in C_V at the left of Fig. 2.3 is real, but the drop at the right is an artifact of the analysis method.

While our normalization method seems to account for spurious effects introduced by a variable firing rate, there remains the fact– not modelled by the foregoing simulation– that C_V itself can vary as well. In fact, this changing C_V *is* observed in the monkey cells: lower firing rates of individual cells have a higher C_V (see sections 2.4.3 and Appendix B).

If several processes with the same rate (and thus the same $\overline{\Delta t}$) but different C_V values have their ISIs binned in the same histogram, the resulting histogram (for example, a sharp peak on a broad base) will have the same mean $\overline{\Delta t}$ as each process separately. The new value of the variance about that mean is given by the weighted mean of the two variances of the individual histograms, so that the composite C_V value will be bounded by the C_V's of those separate processes. Because our claim in this chapter is that the C_V values we observe in monkey lie outside a certain predicted range, the fact that those observed C_V's may themselves only be averages of several true values still requires that most of the true values remain outside the range.

As a further check that our high C_V values did not result from peculiarities of the normalization method, we compared rate-normalized values with those obtained from adapted, constant-rate portions of MT spike trains *without* time-dependent normalization: the two methods gave identical C_V. We also found that changing the PSTH bin-size from $\tilde{t} = 20$ *msec* to $\tilde{t} = 5$ *msec* made no difference in the computed C_V, even during strong adaptation.

2.3.3 Variability in the Interspike Interval

The approximate C_V values measured and illustrated in Fig. 2.3 are in good agreement with reports of C_V at lower firing rates of cortical cells (Burns and Webb 1976; Noda and Adey 1970): $C_V \approx 0.5 - 1$. Visual inspection of the C_V plots did not reveal any systematic differences in C_V between cells in MT and V1; we did not pursue this question further.

Both sets of data show an increase of C_V values from the shortest ISIs measured ($3\ msec$) up to longer ISIs ($10-15\ msec$). As discussed in the previous section, the possible drop in C_V at high values of the interspike interval ($30\ msec$) is most likely a measurement artifact which underestimates C_V when both C_V and Δt are large (Fig. 2.4). The drop for low values of the ISI (high firing frequencies), on the other hand, is a real effect and is in agreement with standard models (see below). While most histograms did not have sufficient counts to justify a functional fit, C_V values near unity are characteristic of the exponential ISI distributions of a Poisson process, the most random type of spike distribution possible.

2.3.4 Variability in the Number of Spikes

As a further test of the variability of these spike trains, we analyzed the number of spikes S_j occurring in a train in response to a *specific* and constant stimulus. We plotted the variance in the number of action potentials per stimulus presentation (σ_S^2) against the average number of spikes S_{avg} for the same non-bursting trains studied above. As is evident from the log-log plots in Figure 2.5, σ_S^2 is scattered widely about the mean spike number in area V1, and equal to or above the mean spike number in area MT. In the case of our large number of MT neurons, we found that the response variance in MT scales approximately as $\sigma_S^2 \propto S^{5/4}$. In the case of a pure Poisson process, the variance in the number of events is equal to the mean. Thus, their ratio should be unity, independent of firing rate. This measure can be used as an alternative to C_V when data is so sparse or average rates so variable that the multi-histogram method breaks down; the explicit relation between C_V and normalized variance is given in Appendix I.

2.4 Analytical Models

In this and the following sections, we will attempt to account for this high degree of variability using simple analytical models of the spiking process.

2.4.1 Integrate-and-Fire Neuron

A neuron is most simply modelled as a single capacitance with an associated membrane potential V, which can be stepwise increased by pulses of constant charge, each pulse incrementing V by a fixed amount. When V exceeds a certain threshold voltage, the model neuron produces an output spike and immediately resets its voltage to the resting value $V = 0$ (for references to this "integrate-and-fire" model, see Tuckwell 1989; also Knight 1972). The

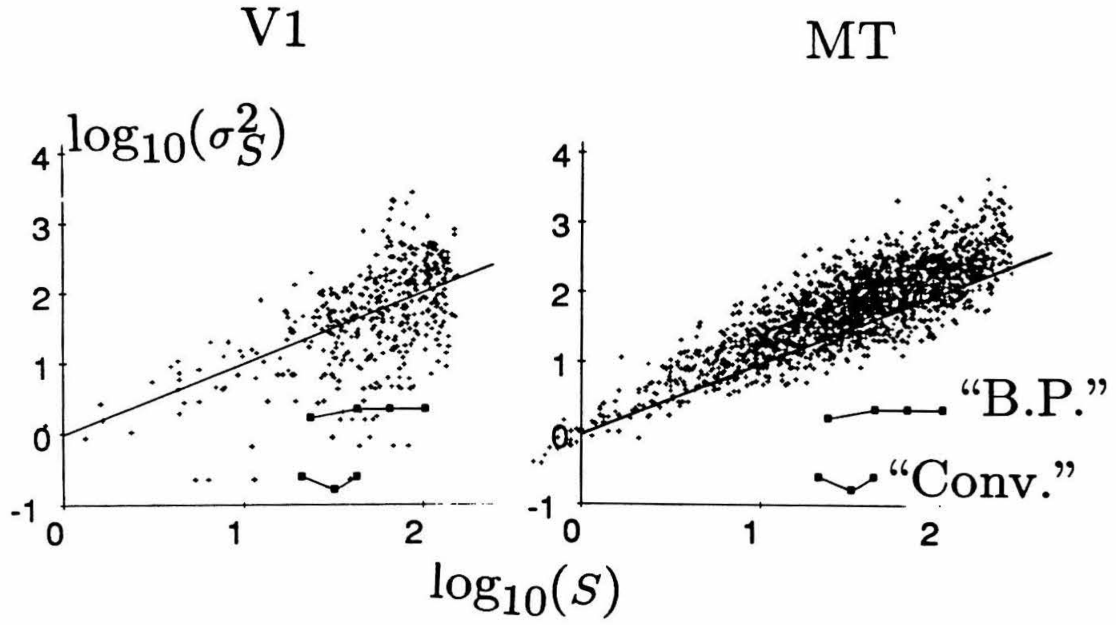

Figure 2.5: **Comparison of the Variance in Spike Count for Monkey and Simulated Pyramidal Cells.** Plots of the number of spikes S_j in a train for a continuous stimulus and the variance σ_S^2 in that number indicate the firing variability over longer times; the log-log scale contains values from a few spikes to hundreds. Values for monkey cells are crosses, in agreement with those obtained for the same areas by Snowden *et al.* (1992). The diagonal line represents the prediction for a purely random Poisson process at constant rate ($\sigma_S^2 = S$). The connected filled squares on both graphs are values given by the "barely plausible" and "conventional" simulations (section 2.5.1), and have far lower variability than that observed in real cells.

neuron's "threshold" can be expressed in terms of the number of pulses (an integer $N_{th} \geq 1$) necessary to bring the cell from rest to discharge.

We further assume that these delta-function-like pulses arrive completely randomly in time (Poisson distributed), with a mean rate of arrival R. Throughout out study, we assume that the synaptic input pulses are drawn from a Poisson distribution (we will re-examine this crucial assumption at the end of the discussion). This randomness assumption is consistent with superposing many independent but possibly non-Poisson input spike trains (Cinlar 1972). An analogy illustrates this effect: the many regular but independent handclaps from an audience can superpose to form applause which sounds like shot-noise. Large numbers of spike-trains can be superposed to produce more variability than a Poisson train only if the individual spikes from the various trains are temporally synchronized (a highly non-random but important situation, which we consider in section 2.6.5 and Appendix E).

This kind of integrating neuron gives one output pulse for every N_{th} input pulses. As a result, the ISI of the output is just the sum of the N_{th} interpulse intervals between the cell's previous spike and the final pulse which triggered the cell's response. With Poisson-distributed pulses, the probability distribution $p(\Delta t)$ of their sum— and hence the predicted shape of the output ISI histogram— is a gamma function of order $N_{th} - 1$ (Tuckwell 1989),

$$p(\Delta t) \;\; \propto \;\; (R\Delta t)^{N_{th}-1} \exp(-R\Delta t). \tag{2.9}$$

Integration of this function over Δt yields the mean and standard deviation, namely

$$\overline{\Delta t} = \frac{\int_0^{+\infty} \Delta t \, p(\Delta t) \, d\Delta t}{\int_0^{+\infty} p(\Delta t) \, d\Delta t} = \frac{N_{th}}{R} \tag{2.10}$$

and

$$\sigma^2_{\Delta t} = \frac{\int_0^{+\infty} (\Delta t - \overline{\Delta t})^2 \, p(\Delta t) \, d\Delta t}{\int_0^{+\infty} p(\Delta t) \, d\Delta t} = \frac{N_{th}}{R^2} \tag{2.11}$$

which give

$$C_V = \frac{\sigma_{\Delta t}}{\overline{\Delta t}} = \frac{1}{\sqrt{N_{th}}} \tag{2.12}$$

Thus, for this "integrate-and-fire" model of a nerve cell with independent synaptic input, C_V is independent of firing rate, since both $\overline{\Delta t}$ and $\sigma_{\Delta t}$ scale inversely with R.

To apply this model to real cells, we suppose that an approximate threshold depolarization for a pyramidal cell is 20 mV from rest to firing, and typical depolarizations for a single excitatory EPSP onto a pyramidal cell (in rat visual cortex) are in the range $0.05 - 0.5$ mV per excitatory input (reported for *detectable* monosynaptic contacts among pyramidal cells in rat cortex by Mason *et al.*, 1991). These admittedly crude values yield $N_{th} \geq 40$ EPSP's, and $C_V \leq 0.16$, i.e., the cell should spike rather regularly. The fact that eq. 2.12 predicts $C_V < 0.5$ for all threshold values $N_{th} > 3$ *spikes* (while empirically $C_V > 0.5$) constitutes the central difficulty this chapter sets out to explore.

2.4.2 Refractory Period

Real nerve cells, however, cannot fire a second action potential immediately after a first, since the sodium channels must deactivate and be repolarized before further activation. As a result, the cell undergoes a short "absolute refractory period," during which it cannot be discharged, followed by a much longer "relative refractory period" during which it is difficult to discharge.

A convenient oversimplification to this case is to modify the perfect integrate-and-fire model by the addition of an absolute refractory period t_0 ("dead time") immediately after resetting, during which the neuron is entirely inactive and after which it resumes normal function. Because the same time t_0 is added to each and every interspike interval Δt, the net effect is to shift the entire ISI histogram (eq. 2.9) rightwards by t_0:

$$\begin{aligned} p(\Delta t) &\propto [R(\Delta t - t_0)]^{N_{th}-1} \exp[-R(\Delta t - t_0)] \text{ for } \Delta t > t_0, \\ p(\Delta t) &= 0 \text{ for } \Delta t \leq t_0 \end{aligned} \tag{2.13}$$

This refractory period now gives the neuron a characteristic timescale, so we cannot expect it to have identical statistics at all firing rates. In particular, the value for $\sigma_{\Delta t}$ (eq. 2.11) now depends on $\overline{\Delta t} - t_0$ rather than on $\overline{\Delta t}$, so that the new value of C_V *does* depend on the mean ISI:

$$C_V = \frac{1}{\sqrt{N_{th}}} \left(\frac{\overline{\Delta t} - t_0}{\overline{\Delta t}} \right) \leq \frac{1}{\sqrt{N_{th}}} \tag{2.14}$$

The refractory period has little effect for $\overline{\Delta t} \gg t_0$ (since in this case $C_V \approx 1/\sqrt{N_{th}}$), but as $\overline{\Delta t} \to t_0$ the output spike train becomes extremely regular ($C_V \to 0$), regardless of N_{th} (see Figure 2.6). In general, the C_V for this simple model of a refractory period is always less than the C_V of the standard integrate-and-fire model (eq. 2.12). This result is easy enough to understand: the very fastest the cell can fire is once every t_0, when the integration period is much shorter than t_0 and contributes little variation. The sudden drop in C_V in the fastest-firing monkey cells for very small values of Δt suggests that this effect—rather than N_{th}—is the dominant influence in the regularity of those cells (see Fig. 2.9; we chose a conservative $t_0 = 1.0\ msec$ for all comparisons, because a larger t_0, leading to an even lower value of C_V, would only increase the gap between the predicted C_V and monkey data).

2.4.3 Leaky Integrate-and-Fire Neuron

It is well known that depolarizations do not persist forever, but that perturbations of membrane voltage tend to decay toward the resting potential (in this section we assume $E_{rest} = 0$ for mathematical simplicity). The simplest physical model of this current "leak" is the inclusion of a passive membrane conductance ($1/R_m$), in parallel with the capacitance of the perfect integrator. (This "leaky" or "forgetful" integrator is described in detail by Stein (1967a) and Knight (1972).) The "leaky integrator" has a decay time constant $\tau = R_m C_m$, giving a behavior between discharges of

$$\frac{dV}{dt} = -\frac{V}{\tau} + \text{input pulses} \qquad (2.15)$$

The passive decay inherent in eq. 2.15 is a simplification of the action of active, voltage-dependent conductances in the membrane of the soma and proximal dendrites. However, it does allow us to capture the essential qualitative aspects

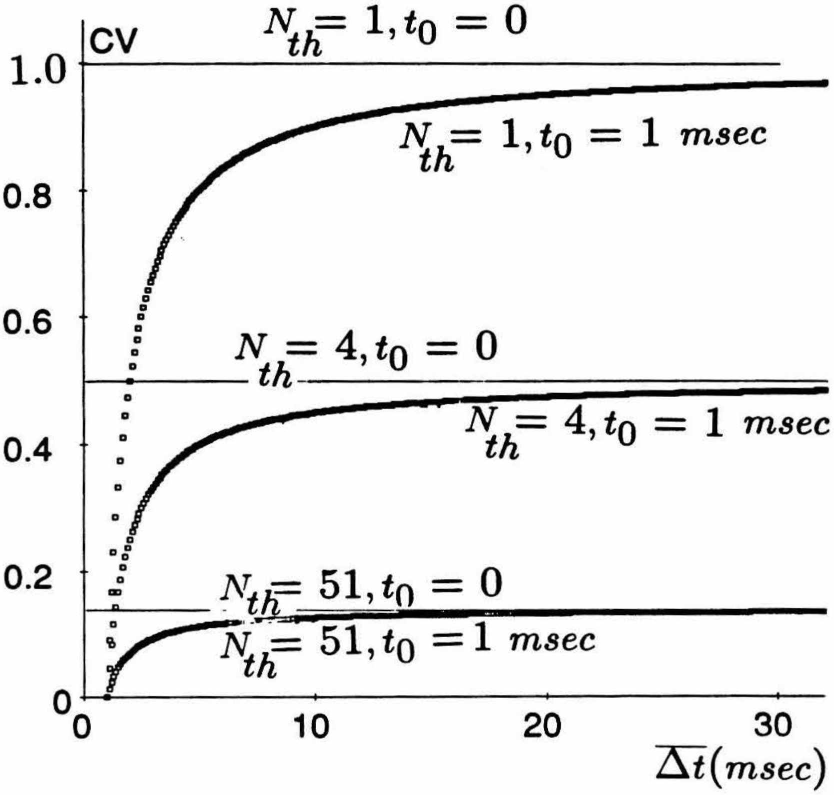

Figure 2.6: **Comparison of C_V from Integrator Models.** Straight lines represent predictions of C_V for a neuronal integrator which fires after receiving N_{th} randomly timed input impulses. The curves show C_V for such a model, modified to account for an absolute refractory period $t_0 = 1.0$ *msec* (curves computed using a different refractory period would have a similar shape, always crossing the $\overline{\Delta t}$-axis at t_0). Note that $C_V \leq 1/\sqrt{N_{th}}$ for all models, so that C_V is quite small (output spikes are regular) for large values of N_{th}.

of temporal decay. But despite decades of effort (Tuckwell 1989), the ISI histogram and C_V for even this simple model are not available in closed form. Our predictions for the C_V of this model come from numerical simulation of eq. 2.15, using a realistic value for the membrane time constant of $\tau = 13\ msec$ in the presence of random input pulses (Mason *et al.* 1991).

Qualitatively, the leak term has little effect on the C_V at high firing rates ($\overline{\Delta t} \leq \tau$), because there is not sufficient time to significantly discharge the capacitance through the leak before the threshold N_{th} is reached. But at very low firing rates ($\overline{\Delta t} \gg \tau$) the output spikes are nearly random ($C_V \approx 1$) because the neuron operates as a "coincidence detector" for occasional bursts of input pulses. In this mode, the membrane potential V "forgets" when the last firing occurred, so that the subsequent firing time is virtually independent of the previous time, i.e., the model neuron's output nearly approximates a Poisson process. Thus the neuron smoothly interpolates between a low C_V (given by eq. 2.14) and the maximum possible $C_V = 1$ as the output ISI increases.

A plot of C_V against $\overline{\Delta t}$ for this model for various levels of thresholds N_{th} illustrates the conflict between the predicted and our observed results (see Figure 2.7). These results show that $C_V > 0.5$ only occurs for $\overline{\Delta t} > 10\tau$ or $N_{th} \leq 3$ (low threshold). The case of a small τ, such that $\overline{\Delta t} \gg \tau$, corresponds to the situation where a large membrane leak exists in the cell's membrane. The conflict between theory and data is greatest for the fastest-firing cells ($\overline{\Delta t} < \tau \approx 13\ msec$); in that regime the leaky-integrator prediction is approximately given by eq. 2.14. A contour-plot for C_V as a function of N_{th} and τ, using a fixed output spike rate $R = 1/\overline{\Delta t} = 200\ Hz$ and absolute refractory period $t_0 = 1\ msec$, is shown in Figure 2.8. It is evident that in order to achieve high variability (i.e., $C_V > 0.7$) at these high rates (which are comparable to those in our faster cells), τ has to be a fraction of a msec, or N_{th} must be only 1 or 2! In fact, the model best fitting the monkey data is that for a neuron which performs no temporal integration, having $N_{th} = 1$ (Figure 2.9).

2.4.4 Realistic Parameters and Modifications

In light of the serious discrepency between the monkey data and the simple theory for random input to an integrator, we investigated several modifications to the theory. The modifications, like the foregoing analysis, are only approximate. When possible they are given as correction coefficients to the perfect integrator with refractory period (eq. 2.14). The resulting patchwork of approximations outlines the major probable influences of these various biophysical modifications on ISI variability. We reserve the detailed equations for Appendices A, B, and C, and outline here the qualitative effects.

Irregular EPSP Magnitude

Our previous model includes a crude approximation of random excitatory postysynaptic potentials (EPSPs) of *constant* amplitude and arriving randomly in time. However, the magnitude of EPSPs is expected to vary greatly, depend-

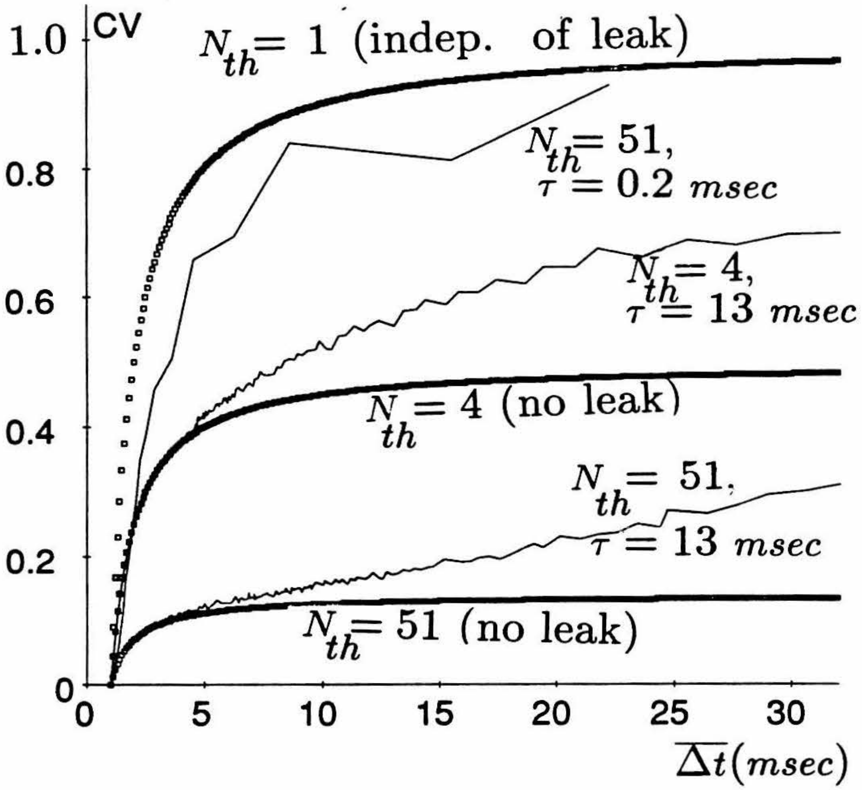

Figure 2.7: **Comparison of Leaky and Non-Leaky Integrator Models.** Squares show C_V for a the non-leaky integrator model with absolute refractory period $t_0 = 1.0\ msec$. Crooked lines show simulations of the leaky integrator for three different values of membrane time constant. The leak term has no effect on C_V for the $N_{th} = 1$ integrator, but raises C_V if $N_{th} > 1$. Only for small values of τ ($\ll \overline{\Delta t}$) does C_V approach unity.

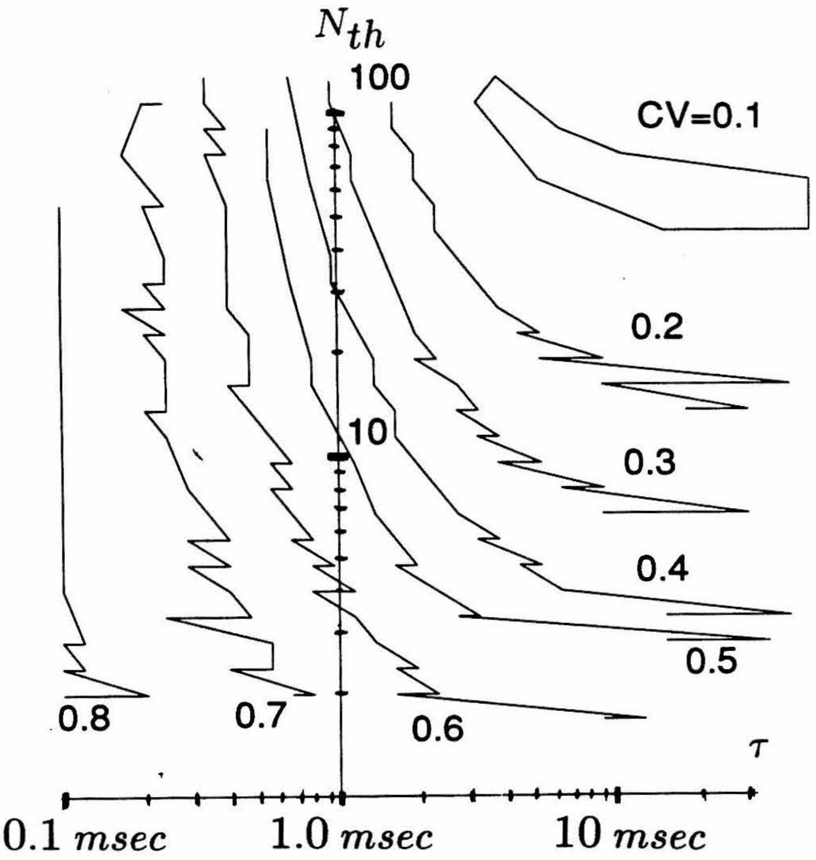

Figure 2.8: **Contour Plot of C_V for Leaky Integrator.** Simulations of the leaky integrator model for discrete values of τ and N_{th} (with refractory period $t_0 = 1.0\ msec$) give the C_V values shown when the mean output interspike interval is $\overline{\Delta t} = 5\ msec$ (corresponding to a mean firing rate of 200 Hz). The jagged contours result from simulating N_{th} and τ at discrete values. Accepted biological parameters (e.g., $N_{th} > 10, \tau > 5\ msec$) predict low C_V values (upper right region); the C_V values observed in monkey would require either $N_{th} < 3$ or $\tau < 1\ msec$ (lower left region).

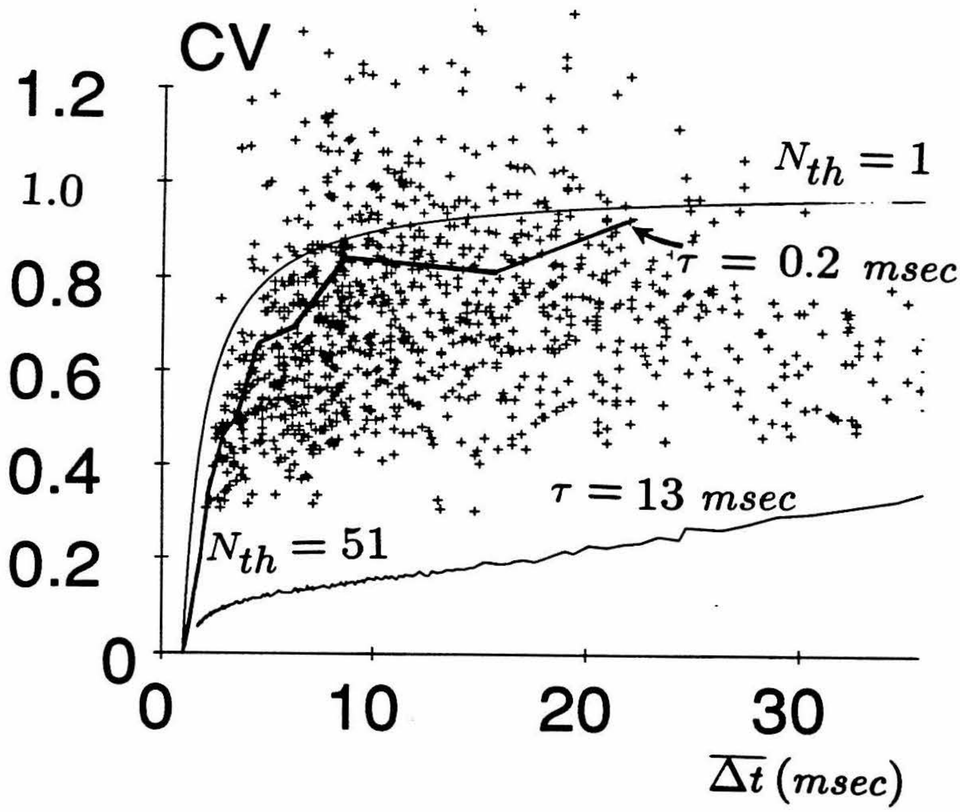

Figure 2.9: **Comparison of Macaque Cortical Neuron Variability and Leaky Integrator Model.** Scattered crosses show C_V for Macaque cortical neurons (C_V was pooled from Figure 2.3). The lower curve shows the simulated leaky integrator model with parameter values in the accepted range ($N_{th} = 51, \tau = 13$ *msec* and $t_0 = 1.0$ *msec*). The middle curve shows the same simulation, still with $t_0 = 1.0$ but $\tau = 0.2$ *msec*, a much shorter decay time than usually accepted for pyramical cells. The uppermost curve shows the theoretical upper bound on C_V for a pure Poisson spike train with "dead-time" $t_0 = 1.0$ *msec*. The observed C_V of macaque cortical cells lies much closer to the maximum possible than it does to the C_V predicted by a neuron model which performs significant temporal integration.

ing on their location on the dendritic tree, quantal fluctuations, etc. Clearly, including random EPSP amplitude as an additional source of variability will increase the variability in the cell's synaptic input and hence its firing. Recent *in vitro* two-electrode intracellular recordings in pyramidal cells in rat visual cortex have shown that the variation in amplitude of unitary EPSPs (from *different* synapses) is nearly equal to the average amplitude of these unitary EPSPs ($0.05 - 0.5\ mV$; Mason *et al.* 1991). Even after incorporating such variable-sized synaptic input into the perfect integrate-and-fire model (Appendix C), C_V only increases from its old value of $\sqrt{1/N_{th}}$ to

$$C_V = \sqrt{2/N_{th}}. \tag{2.16}$$

This factor of $\sqrt{2}$ is not sufficient to remove the discrepancy between the model and our monkey data.

Finite EPSP Width

The model above assumes that EPSCs are instantaneous current pulses, which can carry variability at arbitrarily high temporal frequencies. But even the fastest unitary synaptic currents last between $1 - 2\ msec$ (in the case of fast, non-NMDA, glutamergic synapses; Hestrin *et al.* 1990), blurring the total synaptic current and reducing its variability. Due to this smoothing effect, we conclude in Appendix A that C_V in a typical case would be *reduced* by a factor of 2 below the value predicted above, thus compounding the discrepancy between predicted and observed C_V's.

Adaptation and Hyperpolarizing Currents

The simplified model above does not take into account hyperpolarizing currents (such as the I_{AHP} potassium current which produce "spike adaptation," or the slow inhibitory $GABA_B$ synaptic inputs). The main predicted effect of such currents is to reduce the firing rate, by effectively cancelling a portion of the depolarizing current. As the rate reduces, the mean ISI increases, and allows time for more EPSPs to impinge on the cell before each firing. These extra EPSPs carry with them some added variability, so that C_V will *increase* above the predicted value as the ISI increases (Appendix B). This effect, which is very common in the monkey cells and the compartmental simulation following, is not strong enough to account for the discrepency between the model and the monkey data; in addition, it cannot change the predicted C_V for spikes during the early, non-adapted portion of the cell's response.

2.5 Compartment Models

Even with the modifications discussed above, we had to make a certain number of risky simplifications. For instance, we did not account for the complex dynamics of cellular excitability or its known dendritic morphology, nor did we include a treatment of the effect of fast synaptic inhibition. To answer these criticisms we studied the firing properties of a biophysically very detailed model of a single cortical pyramidal cell using conventional compartmental techniques.

2.5.1 Biophysical Modeling of a Cortical Pyramidal Cell

We simulated the firing properties of a layer V pyramidal cell (see inset in Figure 2.10) from primary visual cortex, whose detailed morphology was reconstructed following intracellular filling with horseradish peroxidase (HRP) during *in vivo* experiments in the anesthetized, adult cat (Douglas, Martin and Whitteridge 1991). Its dendritic tree was described as a list of 186 one-dimensional cables of specified length and diameter, all of which were assumed to be passive (see Appendix F). The cell body contained seven voltage- and calcium-dependent currents; a fast, classical sodium current I_{Na} (with peak conductance per membrane area of 200 $mScm^{-2}$), a slow, non-inactivating sodium current $I_{Na,s}$ (1 $mScm^{-2}$), a L-type calcium current I_{Ca} (0.2 $mScm^{-2}$), and four potassium currents (delayed rectifier I_{DR} (120 $mScm^{-2}$), transient I_A (1 $mScm^{-2}$), calcium-dependent $I_{K(Ca)}$ (45 $mScm^{-2}$) and a non-inactivating I_M current (0.6 $mScm^{-2}$). These currents were modelled using Hodgkin-Huxley like rate constants (Bush and Douglas 1991). This model was studied in detail by Ö. Bernander (Bernander *et al.* 1991), using the very efficient single neuron simulator NEURON, provided by Hines (1989; Appendix F). The somatic spiking threshold and response to injected current for the simulated cell matched those recorded intracellularly *in vivo* (for more details see Bernander *et al.* 1991).

We used an effective passive specific membrane resistance of 26,000 Ωcm^2 throughout the cell. Under these conditions, the somatic membrane potential stabilized at $-75\ mV$, with a spiking threshold of about $-48\ mV$, a somatic passive time constant of 30 $msec$ and an input resistance of 42 $M\Omega$, corresponding to a good and stable intracellular recording from *in vivo* cat pyramidal cells (Douglas and Martin 1991). All synaptic inputs were modelled as transient increases in the membrane conductance, $g_{syn}(t) \propto t \exp(-t/t_{peak})$ with $g(t_{peak}) = g_{peak}$, in series with the synaptic reversal battery E_{syn}.

"Conventional" and "Barely Plausible" Simulations

We then ran two distinct sets of simulations to study the temporal variability of the discharge of this pyramidal cell. In one case ("conventional" or "c." simulation), we used synaptic conductance amplitudes and distributions in rough agreement with experimental findings, while for a second set of simulations ("barely plausible" or "b.p.") we pushed these parameters to the limits of the accepted ranges in order to increase the temporal variability. The b.p. simulation therefore reflects the outer range of temporal variability compatible with a passive dendritic membrane and independent synaptic inputs.

Both sets of simulations included both excitatory as well as inhibitory synaptic input. Perhaps the single most important source of inhibition comes from "basket cell" and "chandelier cells," one of which may form up to 30 inhibitory synapses directly on the recipient neuron's cell body. In the "conventional" model, simulated input from a basket cell activated 30 synchronous somatic inhibitory $GABA_A$ synaptic events ($g_{peak} = 0.1\ nS$; $t_{peak} = 5\ msec$; $E_{rev} = -70\ mV$). All 30 synapses were randomly but jointly activated at the average rate of 450 Hz. The maximum, saturated $GABA_A$ conductance at one synapse (due to several consecutive events) was set to 0.5 nS. These 30 basket cell synapses contributed a mean somatic conductance increase of 10 nS (this value is in the range reported by Douglas and Martin (1991), for the total amount of inhibition).

For the "barely plausible" model, the number of these inhibitory synapses was kept constant, but g_{peak} was tripled to 0.3 nS and the synaptic conductance saturation was eliminated (allowing consecutive synaptic inputs to add in time), so that the net inhibitory conductance change at the soma fluctuated around 60 nS. The resting potential at the cell body stabilized (as before) at around -72 mV. By temporarily removing all the active currents at the cell body, we estimated the resulting average input resistance and passive time-constant during the inhibitory synaptic barrage as $R_N = 11\ M\Omega$ and $\tau = 7\ msec$ for the "barely plausible" cell, and $R_N = 30\ M\Omega$ and $\tau = 13\ msec$ for the "conventional" model. Since the intracellular recorded values of R_N and τ for the reconstructed cell were 23 $M\Omega$ and 20 $msec$ respectively (Bernander *et al.* 1991), the "conventional" case study represents a realistic cortical cell simulation. As mentioned above, the b.p. model will push C_V towards higher values by decreasing the effective membrane time constant.

Excitatory input was provided to the "conventional" model by placing excitatory synapses at 20 different locations throughout the basal and apical tree. The excitatory synapses were assumed to be voltage-independent of the AMPA or non-NMDA type ($E_{syn} = 0\ mV$; mean conductance $\overline{g}_{peak} = 0.5\ nS$;

$t_{peak} = 1.5\ msec$). Furthermore, since individual synaptic amplitudes may vary, the value of g_{peak} for each synaptic event was chosen from an exponential probability $P(g) \propto \exp(-g/\overline{g})$, so synaptic events had a high variability in amplitude, even events occurring at the same location. These values led to somatic EPSPs ranging from a mean of 0.4 mV peak potential and 4 $msec$ rise time for the most proximal synapse to about 0.5 μV for the most distal one. These values are within the range reported in rat visual cortex slice pyramidal cells for unitary EPSPs evoked by stimulating a single presynaptic pyramidal cell (Mason *et al.* 1991), and with values obtained with spike-triggered averaging of EPSPs in cat visual cortex (Komatsu *et al.* 1988).

In the "conventional" model, much of the variability carried by individual EPSPs was attenuated and smoothed out as their current passed through the capacitive dendritic cables toward the soma. In order to reduce dendritic attenuation and temporal smoothing in the "barely plausible" model, and thereby increase C_V, all excitatory synapses were placed on the proximal apical dendrite only 60 μm away from the soma, where the synapses' electrical coupling to the soma would be strongest, while still remaining consistent with experimental observations of synaptic position. At each synapse $t_{peak} = 0.3\ msec$ and $\overline{g}_{peak} = 10\ nS$, giving rise to a very large somatic EPSP (mean depolarization 1.6 mV) within 1 $msec$. This simulation only required the simultaneous occurrence of 19 of these "giant" EPSPs on average to bring the cell from rest to the firing threshold.

Synapse activation times were random (with a fixed probability per unit time), with average rates chosen to yield output spike rates comparable to those analyzed from the monkey ($40 - 200\ Hz$). This required total excitatory synaptic activation rates of $100 - 400\ kHz$ for the conventional model and $8.5 - 47\ kHz$ for the barely plausible simulation. The integration step size used by NEURON was $dt = 0.1\ msec$, with random synaptic activation summed over $0.02\ msec$ subintervals. The simulations generated nearly 700 spike trains of 230 ("c.") or 470 $msec$ ("b.p.") duration. In order to avoid any systematic biases, we analyzed these spike trains with the exact same normalization method outlined in section 2.3.1. Table 2.1 gives the values of the EPSP rates used as well as the number of spikes produced and Figure 2.10 gives one example of a 200 $msec$ excerpt of the somatic potential for typical c. and b.p. simulations. Figure 2.11, Figure 2.12, and Figure 2.13 show the resulting spike trains, histograms, and C_V.

In order to test whether these low C_V results depended on the details of our voltage-dependent somatic currents, we introduced two modifications to the detailed kinetic schemes of the fast sodium current responsible for the action potential in the "b.p." simulation (see Appendix G for the equations modelled). One modification lowered the firing threshold by lowering the midpoint voltage $V_{1/2}$ at which the steady-state value of the sodium activation particle (m_∞) was half its maximum (i.e., 0.5; see Appendix G for the equations modelled). When $V_{1/2}$ was lowered from $-40\ mV$ to $-50\ mV$, the firing threshold was reduced proportionately, but the cell's firing variabillity in response to random synaptic input only increased slightly (as would be expected in the integrator model, for which a lowered value of the firing threshold gives rise

SIMULATION	avg. EPSP rate	# spikes S in a train	# trains
"Barely Plausible"	47 kHz	104.6 ± 1.5	100
	26 kHz	66.5 ± 1.5	100
	16 kHz	43.7 ± 1.5	100
	8.5 kHz	23.7 ± 1.3	248
"Conventional"	400 kHz	42.6 ± 0.5	50
	232 kHz	32.1 ± 0.4	50
	103 kHz	21.0 ± 0.5	50

Table 2.1: **Predictability of Number of Spikes in Simulated Pyramidal Cell.** The compartmental-model simulation using passive dendrites produced a very predictable number of spikes for a given average EPSP rate. Shown are EPSP rates, the spike number S (averaged over all trains at that rate), and the number of trains simulated at that rate. Note that the variability in spike number is far smaller than the $1/\sqrt{S}$ variation expected for Poisson-distributed spikes.

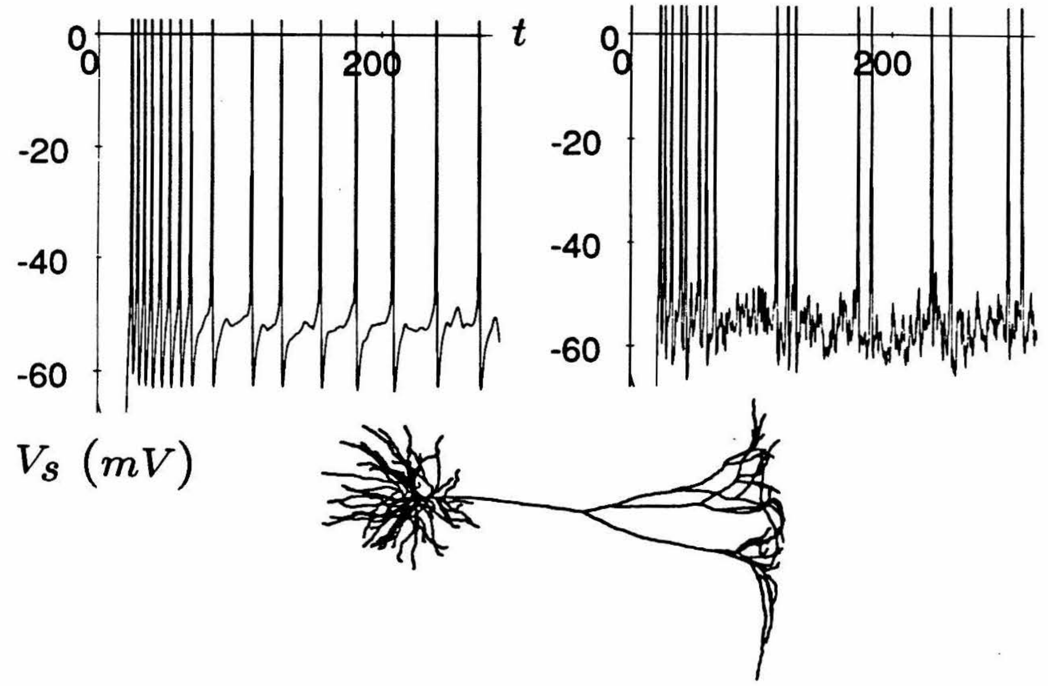

Figure 2.10: **Somatic Potential of a Simulated Pyramidal Cell.** A compartmental model (with passive dendrites) of a reconstructed striate-cortex layer-**V** pyramidal cell (inset) produced the somatic voltage traces shown upon exposure to random EPSP's and IPSP's. (left) Somatic voltage in the "conventional" simulation, with $t_{peak} = 1.5\ msec$, $g_{max} = 0.5\ nS$, with excitatory synapses distributed randomly throughout the dendritic tree. (right) Simulated somatic voltage in the "barely plausible" simulation, which used parameters at the edge of accepted ranges to create the most variability possible: a fast and strong EPSP ($t_{peak} = 0.3\ msec, g_{peak} = 10\ nS$) and all synapses located on the apical dendrite 60 μm from the soma. Note the adaptation in firing frequency following onset of the "stimulus."

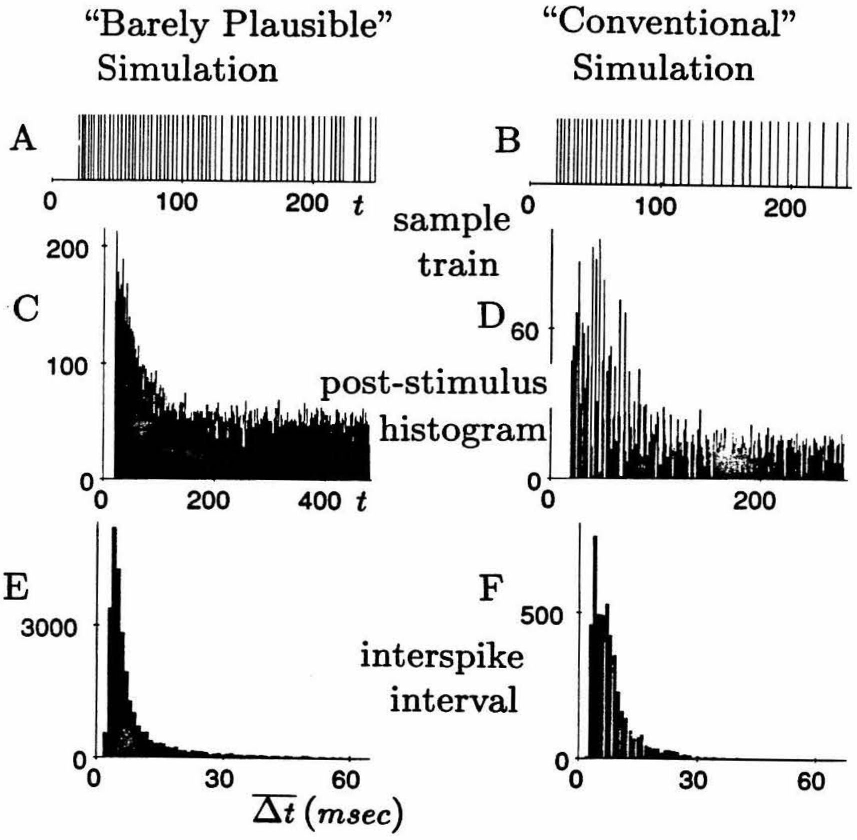

Figure 2.11: **Firing Statistics of Detailed Pyramidal Cell Simulation.** Our compartmental model produced the spike trains shown upon exposure to random EPSP's. The left column used "barely plausible" parameters to simulate a more variable output; the more regular simulation at right used "conventional" parameters. (A,B) Sample spike trains from each simulation. (C,D) PSTH from the same simulations. The prominent millisecond structure in the PSTH (especially the first 50 *msec*) result from the highly regular simulated trains, in which early spike times are well correlated with the onset of stimulation. Although the simulation parameters were fairly conventional, the highly regular spiking they produced was not observed in most cortical neurons. (E,F) Interspike-interval histograms from the same neuron. The broadness of these histograms arises from the combination of different mean firing rates in one histogram, an artifact which is eliminated in the multi-histogram analysis of C_V.

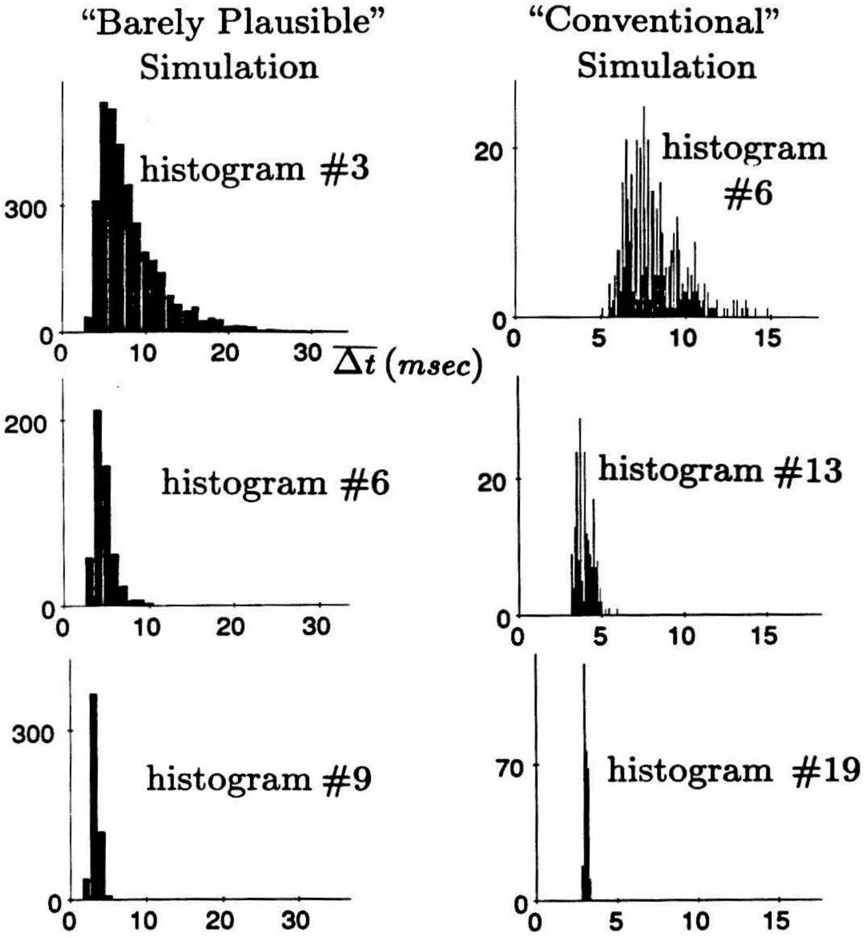

Figure 2.12: **Interspike-Interval Histograms of Simulated Spike Trains.** At left are histograms using 1.0 *msec* bins, with data from the "barely plausible" compartmental model, analyzed as described in section 2.3. At right are histograms from the "conventional" model (here shown as analyzed using 0.1 *msec* bins and 20 histograms for greater resolution). Note that both models have much narrower histograms than the Macaque data, especially at short ISI values (high firing rates), reflecting the unnaturally high regularity of these simulated trains.

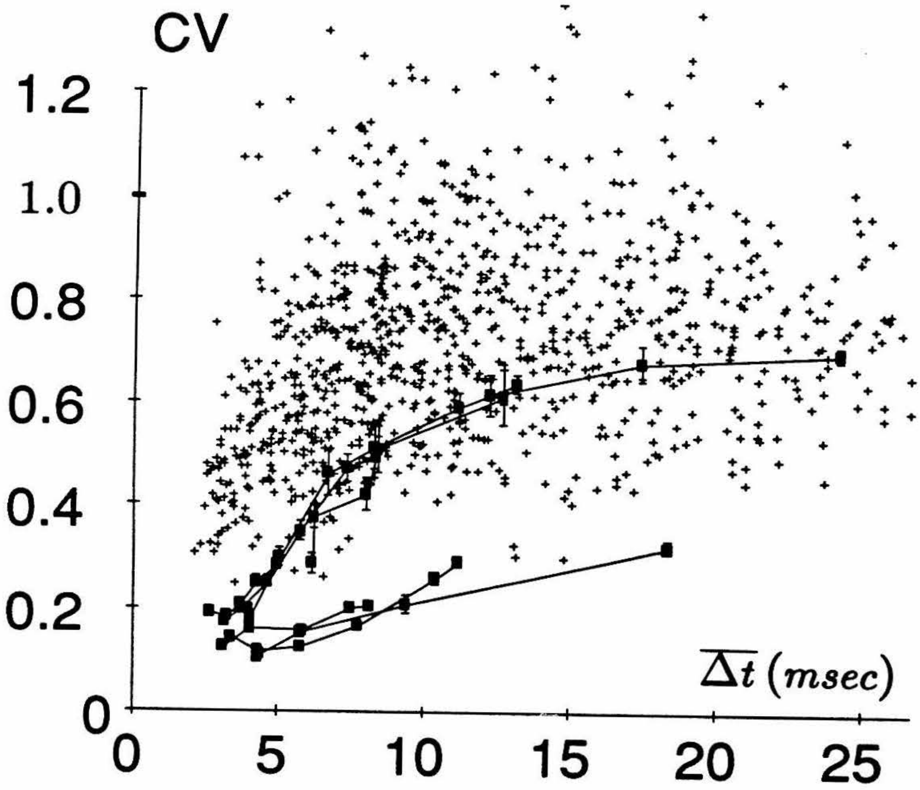

Figure 2.13: **Comparison of C_V Values from Compartmental Simulations with Macaque data.** Scattered crosses are C_V from areas V1 and MT. All sets of connected points represent simulations with random EPSP input to our detailed model of a reconstructed pyramidal cell. All data shown were analyzed with the same method (using 1.0 *msec* bins). The upper sets of filled squares are from four different EPSP rates in the "barely plausible" simulations, with fast EPSP duration, high g_{peak}, and all synapses near the soma. The lower sets of filled squares resulted from three EPSP rates in the "conventional" model, with slower and smaller synapses distributed over the dendritic tree. The C_V exhibited by this model is much lower than for the "b.p." case, because more EPSP's (130) were needed to fire the cell, and because high-frequency variation in the input is attenuated by the dendritic tree and the slow t_{peak} of the synapses. Note that for ISI's less than 10 *msec* both simulations give C_V values far less than those observed in monkey.

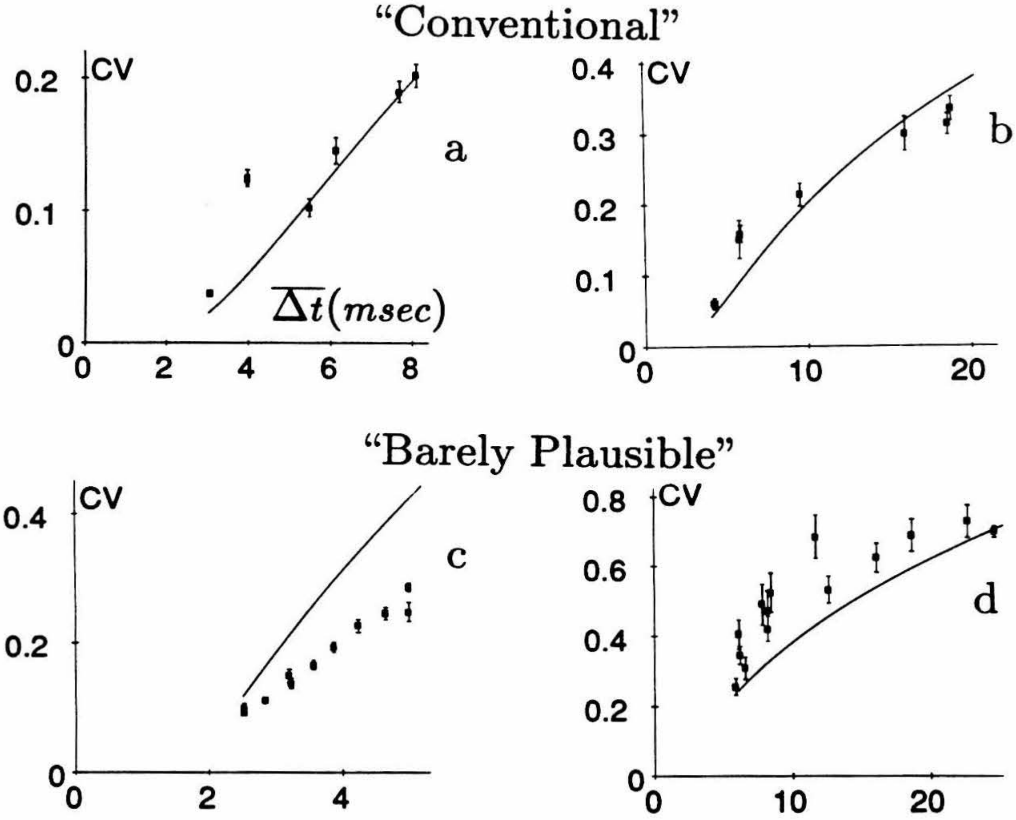

Figure 2.14: **Comparison of C_V from Compartmental Simulations with Integrator Models.** Scattered points represent the C_V's of simulated spike trains at various average EPSP rates; curves are predictions from the modified perfect integrator (eq. 2.17), using the appropriate values of N_{th}, t_{peak}, and initial $\overline{\Delta t}$, and with $t_0 = 1.5$ *msec*. C_V values were calculated with 0.1 *msec* bin-width and 20 histograms to avoid artificially broadening the histograms. The "conventional" model with average EPSP rate of 400 kHz (a) and 103 kHz (b). The "barely plausible" model with average EPSP rate of 43 kHz (c) and 21 kHz (d). These models required roughly $N_{th} = 18$ EPSP's ("b.p." model) or $N_{th} = 130$ EPSP's ("c." model) to trigger the first spike. At high firing rates the simulations produced very regular spiking, because of their refractory periods, dendritic attenuation of high-frequency signals, and non-impulse EPSP's. At lower firing rates ($\overline{\Delta t} > 5$ *msec*), those influences decreased, and the dominant effect became adaptation, as I_{AHP} increased C_V by reducing only the DC portion of the random EPSP input current. The reasonable fits in three of the four cases suggest that the modified integrator model accounts for most of the statistical properties of the *biophysical simulation* (but not the monkey data), despite the model's many drastic simplifications.

to a lower value of N_{th} and hence a slightly higher C_V).

In an alternative modification, the initial Hodgkin-Huxley-like currents (I_{Na}, I_{DR}) were kept, while all other active currents were blocked, so that the simulated cell's firing rate in response to current injection showed the steep onset of spiking typical of a Hodgkin-Huxley-like system. The response of this model to the same random synaptic events tested above only differed in an absence of adaptation; the variability remained the same (not shown).

Simulation Results

The very regular spike trains from these simulations led to low C_V values, in particular at high firing rates: for $\overline{\Delta t} < 5\ msec$ (i.e., firing rates above 200 Hz) $C_V < 0.2$, rising to 0.65 (b.p.) or 0.3 (c.) for $\overline{\Delta t} > 10\ msec$ (Figure 2.13). Thus, they fail to reproduce by a large margin our experimentally measured variability at high firing rates.

One indicator of the regularity of the generated action potential traces is that different simulated spike trains sharing a common average input EPSP rate (but with distinct time structures due to the random synaptic activation times) had virtually identical total numbers S_j of spikes. For a fixed excitatory input rate, S_j varied by only a few percent, far less than the $\sqrt{S_j}$ variation expected of a totally random point process or observed for our monkey data (see Table 2.1 and Figure 2.5).

Another indicator of the extreme spiking regularity of the "conventional" simulation was the presence of prominent peaks on the PSTH long after stimulus onset (Figure 2.11 D); the trains were so regular that a single spike's occurrence could be predicted to a few milliseconds even 150 *msec* after the first spike fired! For these simulations of a passive-dendrite pyramidal cell, there exists over one order of magnitude difference between the expected and the measured variability.

2.5.2 Comparison of Compartmental and Analytical Model

It has been argued that the leaky integrator is such a simplified model of a real neuron—especially at high firing rates—that little can be learned from it. But our simulations do not support this view.

While researchers usually believe that cortical neurons integrate synaptic inputs to produce output spikes, they often criticize the various models' simplification that the number of synchronous EPSP's required to fire (N_{th}) does not depend on firing rate or previous history. Such criticism is entirely justified. In fact, our simulation could generate values of N_{th} varying by more than a factor of four, depending on such circumstances. For instance, the "barely plausible" model required on average about 19 simultaneous EPSPs to fire from rest. After adaptation, a greater number was required.

To what extent could a simple analytical model capture the firing properties of the detailed pyramidal cell simulation? We chose as a model the perfect integrator with refractory period, modified for adaptation and random-height, non-impulse EPSPs (eq.s 2.14 and 2.16 above; A.9 and B.8 in the Appendices; we multiplied all the correction terms without considering their impact on one another). A leak term was not included, since its effect could only be studied by computer simulations; furthermore, any reasonable leak term would be overwhelmed by the adaptation term. When combined, these modifications yielded the prediction

$$C_V = \sqrt{\frac{2}{N_{th}}} \left(\frac{\overline{\Delta t} - t_0}{\overline{\Delta t}} \right) \times \frac{1}{(\pi t_{peak}/\overline{\Delta t})^2 + 1} \times \sqrt{\frac{\overline{\Delta t} - t_0}{\overline{\Delta t}_{st} - t_0}} \qquad (2.17)$$

where Δt_{st} is the mean ISI at the start (fastest part) of the spike train. We chose $t_0 = 1.5$ *msec* (the minimum possible interspike interval in our simulations), and t_{peak} directly from the corresponding simulations. For the threshold N_{th} we used the number of simultaneous EPSPs necessary to fire the cell from rest, although other definitions could have been plausibly used instead.

Equation 2.17 gave a reasonable fit to three of the four simulated C_V's (see Figure 2.14) from the "barely plausible" and "conventional" models. In the poorest match (the fastest "barely plausible" simulation), the variability predicted by eq. 2.17 was too high by a factor of two; C_V values of the other two simulations were predicted within ten to twenty percent. Such good fits are surprising, because the modified integrator model includes neither dendritic effects, shunting terms, nor relative refractory period.

While the entire modified integrator model seems to agree well with the entire compartmental simulation, it is less certain whether that agreement results from the combined agreements of each separate term, or only the fortuitous cancellation of disagreements. As a test, a baseline simulation was generated, and each separate modification added separately to it. The baseline simulation used the same synaptic conductances and membrane currents as the "conventional" model, but with each synapse location of constant strength (rather than variable) at 70 kHz over 20 locations. This simulation had higher variability than predicted at intermediate rates (open squares in Figure 2.15), probably because the rapid change in firing rate caused by spike-adaptation broadened the intermediate-rate histograms (see section 2.3.2).

The first modification to this simulation was to remove all non-linear conductances from the soma except the spiking Hodgkin-Huxley ones, so that the modified cell did not exhibit spike-adaptation; this resulted in more regular firing than the baseline simulation, while firing at a constant rate (Figure 2.15 A). This simulation matched the prediction almost exactly.

The second modification included random-amplitudes in the already randomly-timed synaptic events of the baseline simulation (re-creating the "conventional" simulation above). This resulted in a significant increase in C_V, close to that predicted (figure 2.15 B). The third modification reduced the duration of the synaptic conductances by a factor of five (from $t_{peak} = 1.5\ ms$ to $t_{peak} = 0.3\ ms$), while keeping its area constant (increasing g_{peak} from 0.5 nS to 2.5 nS). This modification was predicted to almost double C_V at the highest firing rates; the fact that it did not suggests that at those rates, the relative refractory period (not modelled) is important at limiting variability (Figure 2.15 D). This effect occurs because firing a spike during the relative

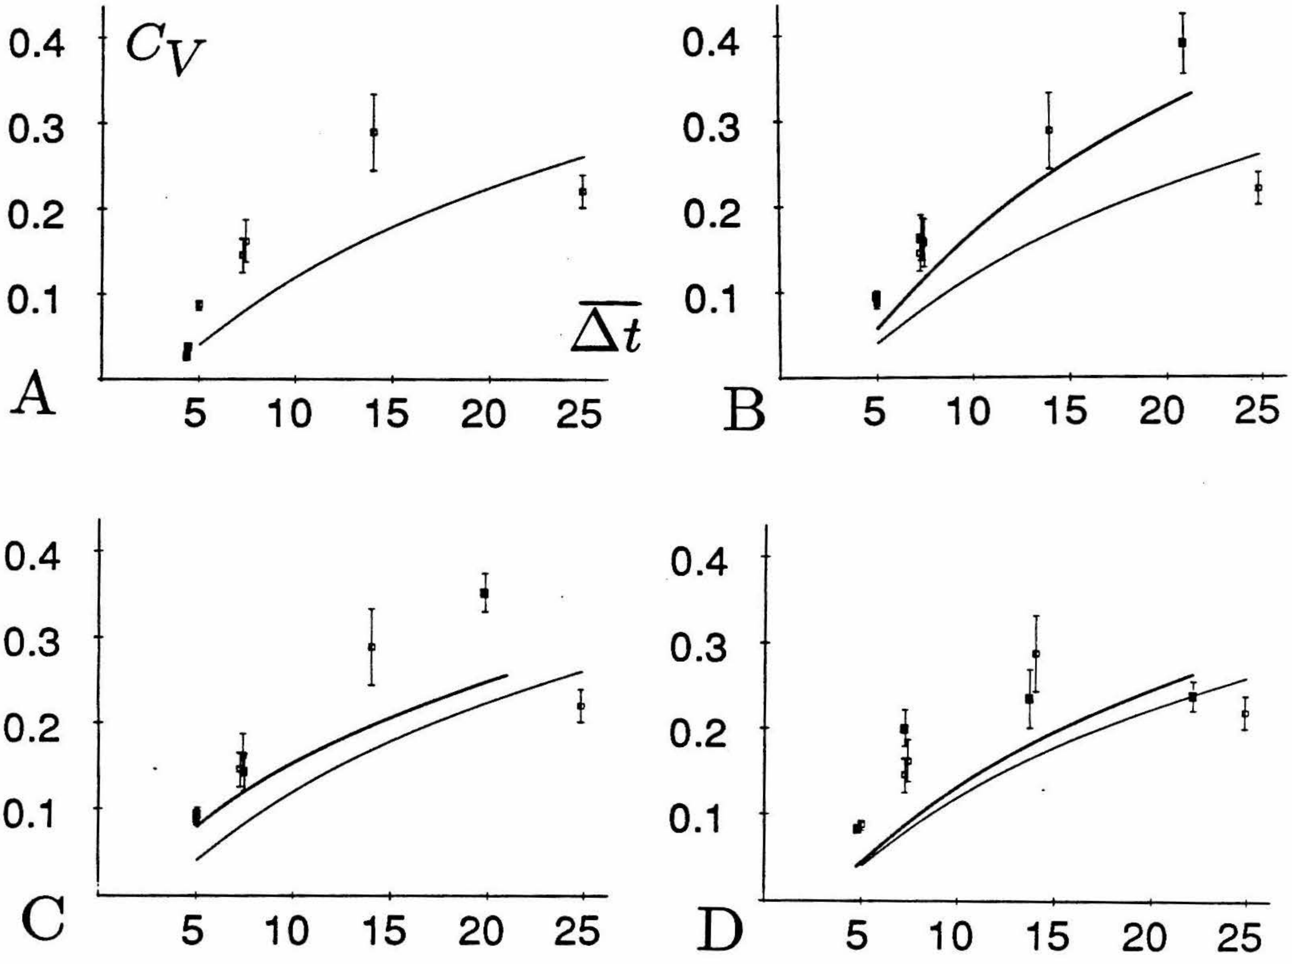

Figure 2.15: **Modifications to the Integrator Model.** All plots show simulations from the "conventional" model with fixed-amplitude synaptic currents (the "baseline"model, open squares) and predictions from the integrator model with refractory period 1.5 *ms*, $N_{th} = 80$, and spike-adapation currents (thin curves). (**A**), The removal of spike-adaptation currents causes regular firing at a fixed high rate (filled squares); the prediction of the integrator model falls just between those simulated C_V values. (**B**), When random amplitude synaptic currents are added to the baseline model, both predicted and simulated firing variability increases (thick curve, filled squares). (**C**), Synaptic conductance values fivefold faster than the baseline model's (but with similar area) lead to a predicted increase in firing variability at the highest rates (thick curve); no such increase is evident (filled squares). (**D**), correlations among inputs EPSPs at the 10% level predict only a slight increase in variability (thick curve) of simulations (filled squares).

refractory period requires many more than N_{th} EPSPs, so that the output spikes are more regular. The prediction might also have failed because additional high-frequency components of these faster EPSPs were mostly filtered out by the dendritic tree, an influence which is not included in the integrator model.

The final modification was to include moderate correlations among synaptic firing times (as described and quantified later, using the contribution coefficient C_c, in Appendix D and section 2.6.5) Here, one EPSP had a 10% above-random chance—$C_c = 0.1$—of being preceded or followed by another within 5 ms. Such coincidences lead to a slight increase in the average depolarization of each *independent* random "event" (single or paired EPSPs), which should be reflected in a slightly decreased N_{th} (eq. 2.24). The 10% correlation used here should have reduced N_{th} from about 80 to about 72, a change much smaller than the uncertainty in estimating N_{th} for this complicated cell in the first place. As expected, there was no significant change in the simulated C_V (figure 2.15 D).

2.5.3 Active Dendritic Simulation

The foregoing model included voltage-dependent conductances only at the soma, leaving the entire dendritic tree passive. In that case, the neuron will act as an integrator, with low firing variability. How could dendritic nonlinearities affect our results?

It is known that dendrites in hippocampal and neocortical pyramidal cells can generate TTX-insensitive all-or-none electrical events that most likely involve calcium conductances (Wong, Prince, and Basbaum 1979; Amitai *et al.* 1992; Jones, Kunze, and Angelides, 1989; Regehr and Tank 1991; Westenbrook, Ahlijanian, and Catterall, 1990; see also Huguenard, Hamill, and Prince 1989). But the relatively long duration of such events (20-50 *msec*) would carry little high-frequency variability in current to the soma. Because there is not much detailed data available, we conducted an explicitly unrealistic simulation of active dendritic conductances, intending only to show that they are in principle capable of producing high firing variability in response to random input.

We reasoned that since the soma spikes in response to currents from the dendrites, we must make the dendritic currents as variable as possible. Variability in dendritic current relative to its mean sustained value can arise in general from two mechanisms: 1) fast, strong depolarizing impulses (such as spikes), which add both variable and sustained components to the dendritic current, and 2) fast *repolarizing* impulses, which remove the sustained component of current contributed by the spikes. Fast repolarization increases the variable (AC) component of dendritic currents and reduces the sustained (DC) component.

In simulations of dendritic spikes, we chose strong values of the repolarizing current I_{DR}, so that the voltage at the soma returned within a few milliseconds to nearly the same voltage it had before the dendritic spike (without this very strong rectifying current, the somatic depolarization persisted, decaying slowly with the cell's passive time constant). Because the cell carried little lasting memory of a spiking event, the cell did *not* integrate dendritic spikes, but only fired upon the *coincidence* of several of them.

We therefore simulated active Hodgkin-Huxley-like conductances in the basal terminal branches, between their tip and most distal branching. We matched the mean sodium conductance to its somatic value ($\overline{G}_{Na} = 200\ mScm^{-2}$), and used a fast potassium conductance twice that value (or 3.5 times the somatic $\overline{G}_{DR}$) to accomplish the repolarization outlined above. Even with these very strong conductances, most neighboring dendritic terminal branches were decoupled, so that a spike in one would not necessarily fire its neighbor. In addition, most voltage-dependent conductances in the soma were removed, along with the basket-cell inhibition and all apical input simulated earlier. As a result this "bare" neuron contained only Hodgkin-Huxley-like mechanisms and a passive membrane with time constant of about 30 *msec*. Each basal dendritic compartment was subdivided into 20 sub-compartments (1600 sub-compartments total) to ensure that the high-frequency dendritic fluctuations were integrated faithfully. Apical dendrites were left passive and unstimulated because they consumed large computational resources while contributing little voltage to the soma.

Each of 42 active dendritic terminal branches was stimulated with single triggering pulses, each pulse instantly resetting the local membrane voltage to $-40\ mV$ and initiating a dendritic spike. Each dendritic spike caused a peak somatic depolarization between $2 - 9\ mV$. Each terminal branch was stimulated independently of the others, but the input to a single terminal branch was not random: we allowed for an absolute refractory period of 2 *msec* after each spike's firing before choosing a random time at which to fire it again. This requirement increased the regularity of the dendritic input and limited the speed at which the dendritic spikes could fire, hence keeping the soma's output spike rate lower than we desired.

As a result of this bombardment by dendritic spikes (up to 22 dendritic spikes/*msec*), the somatic voltage fluctuated strongly about a roughly constant -65 mV (Figure 2.16 A). As desired, there was a large variability in the somatic voltage without significant sustained depolarization. Each of the three dendritic spike rates used produced a constant average output spike rate, so that we could analyze with a single histogram the C_V values of each of the 20 trains simulated at that rate (Figure 2.16 B). C_V values were $0.6 - 0.8$, significantly above those for the passive-dendrite models, and in the same region as the monkey data; the variance in spike-number σ_S^2 reached nearly Poisson values ($0.7S$).

In order to isolate the relative contributions to firing variability of the dendritic sodium currents, the potassium currents, and the triggering pulses, we performed the same simulation for two related scenarios. In one scenario, we reduced I_{DR} currents by a factor of ten from their values above, leaving enough delayed rectification to reset the local sodium channels but not enough to significantly repolarize the soma after a dendritic spike. This simulation naturally required less frequent dendritic spiking to fire the soma; when

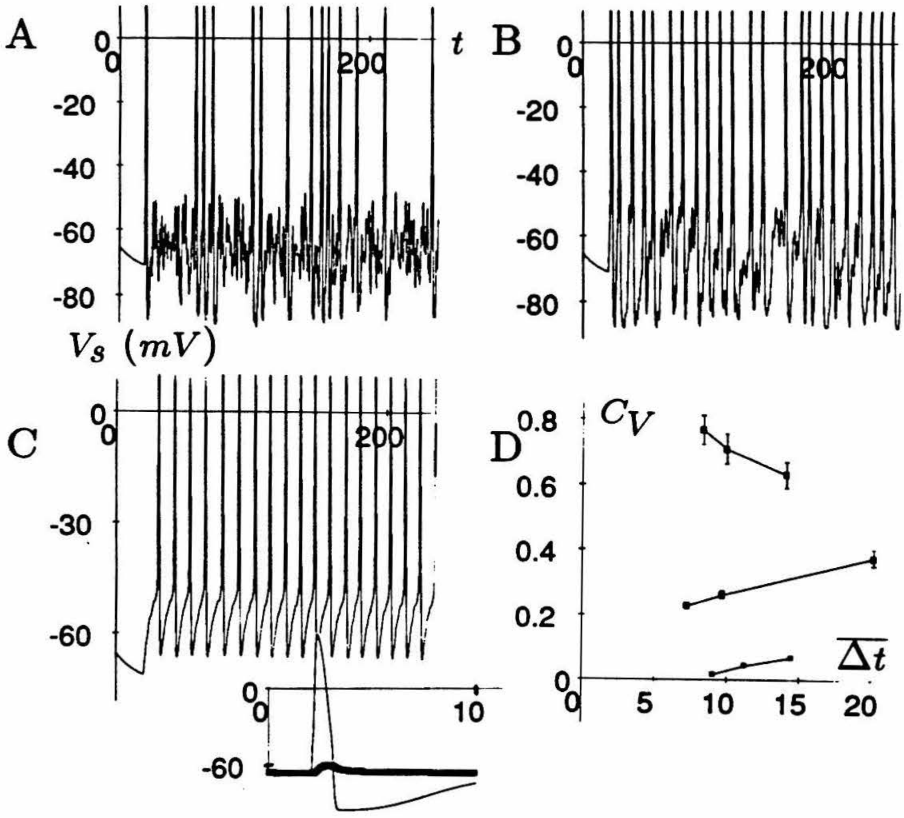

Figure 2.16: **Highly Variable Spiking Caused by Simulated Strong Dendritic Nonlinearities.** Our layer-V pyramidal cell described above was endowed with only Hodgkin-Huxley-like conductances at the soma and on the most distal branches of the basal dendrites; no other active currents or synaptic inhibition were included. Dendritic spikes (thin line in inset) were triggered in each of 42 active basal branches at a random time more than 2 *msec* after its previous firing, and independent of the other branches' triggerings. (A) The most variable somatic firing occurred for dendritic dendritic potassium conductances g_{DR} twice the strength of the sodium conductance, because the strong repolarization cut short the somatic depolarization (thick line in inset), thereby preventing temporal integration. (B) Less variable firing occurred when the dendritic g_{DR} was reduced to 1/10 of its above value, thereby allowing spikes' depolarization to accumulate in the cell body over time and permitting temporal integration of dendritic spikes. (C) The control case: highly regular spiking occurred in dendrites with no active conductances, as many triggering pulses were integrated to fire the cell. (D) C_V values for the three aforementioned simulations: strong I_{DR} (top curve), weak I_{DR} (lower curve), and passive dendrites (bottom curve). Each square represents the C_V calculated from 20 simulated trains at a constant firing rate. Only the case with strong dendritic g_{DR} yielded high output firing variability consistent with the monkey data.

its output rate was adjusted to match that of the strong-I_{DR} case above, it produced C_V values about half as large ($0.2 - 0.4$), suggesting that fast active rectifying currents in the dendrites can be an indispensable contribution to output variability (Figure 2.16 C).

A small portion of the soma's depolarization arrived not through active conductances, but merely from the triggering pulses which reset terminal branch voltages randomly to $-40\ mV$. We verified that the triggering pulses by themselves contributed virtually no variability to the output by eliminating *all* active dendritic currents, while keeping the triggering pulses; the cell's response at the same output rates above was extrememly regular ($C_V = 0.02 - 0.07$; Figure 2.16 D).

2.6 Discussion

We will now briefly review the assumptions underlying our spike train analysis, discuss the experimental data and then list possible objections to our compartmental modeling efforts. We will finish by describing some of the implications of our analysis.

2.6.1 Statistical Assumption Underlying our Data Analysis

There exists a very rich literature concerned with the statistical analysis of spike trains using the theory of stochastic point processes (Perkel *et al.* 1967; Tuckwell 1989). Almost invariably, it is assumed that the spike generating process is a *stationary* one, so that the underlying probability distribution of these point events does not change with time, or depend on a "starting" time (Burns and Webb 1976; Correia and Landolt 1977; Teich *et al.* 1977; Lansky and Radil 1987). However, the spike trains used in our study all occur following stimulation, and their response is non-stationary. The most prominent such non-stationarity is the decrease in firing rate with time (Figure 2.1 C,D), reflecting both adaption processes intrinsic to the cell as well as network effects.

Since we were not concerned with the detailed fitting of analytical distribution functions to the interspike-interval histograms of these spike trains, we tried to account for the nonstationarity inherent in the data by using the simple normalization process described in Section 2.3.1. That method calculated an approximate instantaneous firing rate from the PSTH and the stimulus efficacy, and used that rate to create several separate, near-stationary histograms for C_V analysis.

We also concluded that the firing variability arises at a fast timescale, i.e., milliseconds, rather than at the slower timescale of varying average rates, i.e., tens of milliseconds. In a different study, we had computed the autocorrelation functions associated with the single cell data derived from macaque area MT (Bair, Koch, Newsome, and Britten 1992). The absence of a broad central peak around the origin—associated with a process whose mean rate varied slowly over time (see Figure 7 in Perkel *et al.* 1967)—suggested that adjacent fast ISIs were uncorrelated, so that a fluctuating mean rate did not contribute to the high firing variability we observed.

But in some MT cells a broad autocorrelogram peak (50-200 *msec*) *did* exist in the absence of any structure in the PSTH; we interpret this to mean that the firing rate fluctuated randomly. To estimate the fast-timescale variability in that situation, we computed C_V from many tiny histograms of only ten adjacent ISIs each. Those C_V values were widely scattered about a mean 10-20% below the C_V computed by the multi-histogram method (section 2.1), suggesting that those neurons fire quite irregularly at fast as well as slow timescales.

2.6.2 The Variability of Cortical Cell Firing

We measured the degree of variability of the neuronal spike discharge in a large number of non-bursting visual cortical cells in two different but related manners. One study was of the variability in the intervals between consecutive action potentials. Its principal result is shown in Figure 2.3: for firing rates up to several hundred Hertz, the value of C_V is close to 1 (characteristic of a random Poisson process) for both V1 and MT cells.

We also measured the variability in the *number* of action potentials in a single train for both sets of data (Figure 2.5). Our finding that the variance of cell firing increases roughly linearly with the mean response rate is well known for cells in cat and monkey primary visual cortex (Heggelund and Albus 1978; Tolhurst, Movshon and Dean 1983; Parker and Hawken 1985; Vogels, Spileers and Orban, 1989; Zohary, Hillman and Hochstein 1990), and has recently also been established for cells in area MT of the alert macaque monkey (Snowden *et al.* 1992). Our results here are compatible with the known literature, and are also approximately consistent with a description of spiking as a Poisson process (see above).

2.6.3 Analytical Results

In an attempt to understand the origin of the observed variability in the neuronal discharge, we analyze the variability of the impulse activity of different integrate-and-fire models. Our primary assumption is that these models—as well as the passive compartmental models—spatially and temporally integrate synaptic input from a large number of independent processes. Our results can be qualitatively explained by the *Central Limit Theorem*, which states that as the number n of incoming independent random variables x_i goes to infinity, the random variable defined by the **mean** over x_i, i.e., $\overline{x} = (1/n)\sum_{i=1}^{n} x_i$ has an asymptotically normal (i.e., Gaussian) distribution, with mean identical to the mean of the population x_i and with standard deviation scaling as $1/\sqrt{n}$ of the population's standard deviation. In other words, if a neuron can only be brought to fire action potentials by summing over dozens or more of independent synaptic inputs, it should fire very regularly!

This intuition is born out by an analysis of the C_V of different integrator models, which are summarized in the contour plot (Figure 2.8). The high C_V value we observe experimentally can only be obtained by either assuming that N_{th} is very small, i.e., that 1 or 2 inputs are sufficent to trigger the cell, or that the time-constant τ is a fraction of a millisecond (Figure 2.9), thereby preventing any effective temporal integration from occurring. Further modifications to the leaky-integrator model, such as hyperpolarizing (adapting) currents (Appendix A) and the finite width and variable amplitude of EPSPs (Appendices B and C), do not change in any significant manner our fundamental conclusion that integrator models produce very regular output trains at high firing rates.

2.6.4 Biophysical Detailed Simulations

It can be argued that simple integrator models do not provide a realistic description of cortical pyramidal cells. In order to satisfy ourselves that a more realistic neuron which integrates many independent EPSPs is still inconsistent with the measured high variability, we numerically simulated the dynamical properties of a HRP injected and reconstructed neocortical, layer V, pyramidal cell. While the shape and electrical properties of this cell were derived from cat visual cortex, its properties are not likely to differ fundamentally from those in monkey visual cortex, the source of our variability data. We acknowledge that some of our monkey cells—especially the fastest-firing ones—may have been rapidly firing interneurons rather than pyramidal cells (Agmon and Connors 1992). But we chose a pyramidal (rather than interneuron) cell model because pyramidal cells are far more common and larger, and hence probably represent the majority of the monkey cells recorded.

Conceptually, we would like to distinguish these simulations according to whether the cell acts as an integrator or whether it acts as a high fidelity temporal coincidence detector. Accordingly, for the majority of our simulations, we assumed that the dendritic tree contained no voltage-dependent membrane conductances (integrator mode), while in a second, more exploratory set of simulations we endowed the distal part of the basal dendritic tree with strong nonlinearities (coincidence mode).

Passive Dendrites: Integrator Mode

If the simulated pyramidal cell was bombarded by massive amounts of fast, excitatory synaptic conductance inputs of the non-N-methyl-D-aspartate (non-NMDA) type, such that it fired at the high firing rates observed in our monkey data, then its output firing—even in the "barely plausible" simulation—was much more regular than the monkey cells' firing. Both the experimental and the modelling data are compared by the variability in the number of action potentials per trial, and by the rate-normalized C_V. Thus, in spite of the complex dynamics of the seven voltage- and time-dependent currents at the soma, the simulated cell essentially still acted like an integrator, and its low variability was predicted by a modified integrator model (eq. 2.17; Figure 2.14).

Similar to the integrate-and-fire model discussed above, high C_V values could only be obtained if single EPSPs were very large (greater than 10 mV, thereby reducing N_{th} to 1 or 2), or if the passive time-constant was in the submillisecond range. But these ranges are excluded by intracellular recordings. Evidence from neocortical and hippocampal slice recordings report a range of unitary EPSPs between 0.05 mV and 3 mV, with the majority of averaged EPSPs less than 0.5 mV (Thompson *et al.* 1988; Sayer *et al.* 1990; Mason *et al.* 1991; McNaughton *et al.* 1981). Those measurements **do** include the multiple connections which single input axons often make on individual cells. Occasionally, much larger EPSPs have been observed (C. Stevens, personal communication). And recordings from the cell body of cortical pyramidal cells yield values of τ on the order of 10 to 20 $msec$ in the intact animal (Creutzfeldt *et al.* 1964; Douglas *et al* 1991), and much larger values in slice neurons using the patch clamp technique (Spruston and Johnson 1992), all well beyond the necessary submillisecond range. We did not simulate any voltage-dependent (NMDA) synaptic input, since the long decay times ($20-50$ $msec$; Hestrin *et al.* 1990) of the NMDA-associated conductance would dramatically reduce variability in synaptic currents. Modification of other cellular parameters, such as lowering the threshold for initiation of action potentials or blocking all but the fast sodium and the delayed rectifier potassium current, had very little effect on the C_V values.

While an inherently random firing mechanism could in principle account for the high variability we observe in monkey cells, some research (in other neuron types) has suggested that the spike-firing mechanism is very reliable. Calvin and Stevens (1968) concluded that cat spinal motoneurons derive at least 90% of their already small firing-time variability from variability in synaptic currents. Bryant and Segundo (1976) found that various neurons in *Aplysia* gave virtually identical responses patterns to repeated injections of white-noise current. This reliability occurs because the types of statistical fluctuation expected from spike-generating mechanisms—for instance, random channel openings—have small quantal size, so their collective effects are reasonably constant over a reasonably large membrane area (Strassberg and DeFelice 1992). In general, only the largest quantal effects—such as EPSP arrivals—will contribute significantly to firing variability.

Active Dendrites: Coincidence Mode

We also simulated active dendritic conductances whose random triggering maximized the cell's firing variability. We found it very difficult to "construct" a pyramidal cell that fires as irregularly as the monkey cells. For such events to cause highly variably somatic firing, the dendritic spikes must be large, fast, and strongly repolarizing. Only under these conditions do we see high variability for spikes before adaptation sets in. Yet we do not claim that such dendritic nonlinearities exist, but only that they can, in principle, explain the observed variability.

The most important characteristic of this simulation, and the reason why it produced such high variability of output firing, was that it did *not* perform temporal integration of dendritic spikes, but only coincidence-detection among them. This property became evident in the cell's strikingly strong response to slightly synchronized inputs: when dendritic spikes were re-organized to fire in simultaneous pairs (rather than singley) at the same average rate as before, the cell's output firing rate increased by over 50%. Such coincidence-detection is analogous to the "logic operations" postulated to take place among dendritic spines (Shepherd *et al.* 1989), and has been discussed for leaky-integrator models operating at much slower rates (Bugmann 1991). But in such a scheme, the individual output spikes would represent the fundamental elements of logical computations at the millisecond scale, rather than mere "noise" in an average firing rate which is averaged out over tens to hundreds of milliseconds.

2.6.5 Network Effects

Of the many parameters which we need to reevaluate in light of this discrepency, perhaps the most intriguing is the possibility that the individual synaptic events impinging onto a cortical neuron are not independent after all. In that case, the Central Limit Theorem would not apply any more. In particular, what if the EPSPs arriving from different neurons were synchronized? What degree of synchrony could account for the firing variability of our cells, and where might it come from?

Weakly synchronized EPSPs would not be sufficient. In all of the foregoing models (except the $\tau < 1$ *msec* case), the neuron's output variability directly reflects the variability of its synaptic input current. Therefore, if many small

EPSPs are to account for the observed high output variability, they must be strongly synchronized, so that the resulting input current is *just as variable* as a current composed of individual $10 - 15\ mV$ events (i.e., the effective $N_{th} \leq 2$). Any significant number of *non*-synchronized EPSC's would create a nearly DC "background" current, which would reduce the variability of the net input current and hence of the output firing. Such strong synchrony might result from network effects such as burst synchronization (Bush and Douglas 1991; Koch and Schuster 1992).

For synchronized firing to explain the C_V results presented here, a near-majority of the EPSPs must participate in coincidences at the millisecond scale, because high output variability from an integrator-model still requires a high variability in synaptic input current, A quantitative estimate of the type and amount of synchrony necessary to create a given firing variability is given in Appendix E; the surprising result is that even moderate amounts of spiking synchrony (as measured between cell pairs by cross-correlation methods) can mask much stronger synchrony in the whole cell population, so that the highest levels measured might account for much of the variability observed. Regardless of the order of correlations assumed among a perfect integrator's EPSPs, a pairwise "contribution coefficient" of $C_c \approx 0.2$ (among the highest measured, Toyama *et al.* 1981) can yield $C_V \approx C_c^{-1} = 0.44$ (see eqs. E.19 and E.24). This represents a significant increase in variability, although not enough to solve the paradox by itself.

While those amounts of firing synchrony are still not sufficient to cause the high variability discussed here, that synchrony is too *large* to have resulted from integrator-type model neurons under any reasonable input configurations, as shown in Appendix F... but a millisecond-scale coincidence-detecting model neuron can easily account for that paradoxical synchrony of cortical firing. This second paradox results from the same root cause as the variability paradox—the fact that integrator-model firing times are better determined by their previous firing than by their present input—and may prove equally instructive.

Highly synchronized EPSPs were first proposed as the "reverberation" in a "cell assembly" by Hebb (1949), and later as "synfire chains" by Abeles (1990). As evidence, Abeles cites millisecond precision in repeated interspike intervals observed in various locations of monkey cortex in his laboratory (Abeles 1982). Similarly precise ISIs are reported by Strehler and Lestienne (1986) for monkey visual cortex, Frostig *et al.* (1985) for cat medial frontal cortex, and Legendy and Salcman (1985) for cat striate cortex. But highly synchronized inputs would raise serious questions about the "stochastic" nature of neurons (Knight 1972; Sejnowski 1981; Hinton and Sejnowski 1986), and the resulting justification for population coding and massive redundancy.

It is tempting to invoke chaotic dynamics to explain this firing irregularity, especially because these cells contain coupled nonlinear mechanisms and because chaotic behavior has been observed in other neural systems (Freeman and Van Dijk 1987). But this particular system showed no noticable chaotic behavior in all the simulations performed, and real cortical cells show virtually no firing irregularity when injected with DC current, so a chaotic explanation must await a means of producing the chaotic dynamics inside a single cell.

2.6.6 Conclusion

According to our current understanding of pyramidal cells, only a few situations could cause near-random, fast firing in these cells: a very strong inhibitory leak (leading to an effective membrane time constant $\tau \leq 0.2\ msec$); extremely strong synaptic events ($> 10\ mV$ depolarization per EPSP); strong and fast nonlinear dendritic all-or-none events, with fast repolarization; or highly synchronized, non-random synaptic input. In short: either the cell must have extremely large, fast depolarizations, or it must have a very fast mechanism for repolarizing the membrane *during* "integration." In both these cases, the high ISI variability results directly from an equally high variability in the currents arriving at the soma. Neither case corresponds to temporal integration over small, independent, excitatory synaptic events.

The traditional view of cortical firing variability has been that information is only carried in the average spike rate (frequency code); scatter about that rate represents random "noise," whose particular structure is of no use. According to this view, a neuron which fires very randomly carries uncertain information, because of the inevitable scatter in the number of counts in any time window—only a few distinct counting rates can be distinguished in a short integration time. Thus, a highly irregular neuron is the "worst possible" at carrying information in its average rate. Stein (1967b) found that such a frequency-coding neuron has a channel capacity decreasing roughly as $\log(1/C_V)$ for large integration times. While such frequency-coding is very inefficient, it is robust to perturbation of individual spike times, and it does not require complicated postsynaptic neurons to "decode" its message. Furthermore, high variability may have useful properties. It can help a neuron to "explore" its nearby synaptic vector-space during unsupervised learning (Mazzoni *et al.* 1991). And it may enable neurons to implement multiplicative (quadratic) computations (Srinivasan and Bernard 1976; Suarez and Koch 1989; Koch and Poggio, 1992;).

The alternative view is that each spike's arrival time signifies an independent message of some sort (an asynchronous binary pulse code). If each message (spike) has the same probability of arrival, independent of the other messages, then the resulting spike train is Poisson (by definition), and the spike train carries the *maximum* amount of Shannon information possible for its fixed bandwidth and firing rate (Stein 1967b; the less predictible a spike is, the more information it carries.) Thus, a highly irregular neuron would be the "best possible" for carrying information in its individual spike times, although the nature of the information encoded and the ability of subsequent neurons to use it may be unclear. Further experimental and computational studies are required to determine whether cortical computations occur at the millisecond level.

Chapter 3

A Solution: submillisecond coincidence detection in active dendritic trees

3.1 Introduction

The fundamental output of a cortical neuron is a single action potential lasting about a millisecond, which can in many cases cause an equally brief excitatory synaptic current. But can a single neuron use that temporal precision in its computation?

In search of a possible cellular basis for millisecond-scale computations, this chapter tests the upper limit of computational bandwidth in individual cortical pyramidal cells by postulating one situation—the presence of active spiking conductances in distal dendrites—in which such a cell might preferentially respond to synchronized EPSP's. Rough analytic predictions and detailed simulations of a reconstructed pyramidal cell together suggest that such a cell could in principle perform submillisecond coincidence detection. Previous related work by others has included numerical simulations which explored the computational properties of active dendritic trees (without emphasizing fast timescales; Jaslove 1992; Shepherd *et al.* 1989; Shepherd and Brayton 1987), and analytical work which yielded far more sophisitcated expressions for electrical activity in dendrites (Jack *et al.* 1983) than the simplified expressions used here.

This chapter has two sets of goals:

> 1) To explore the requirements of high-frequency coincidence discrimination; to present "proof-of-concept" simulations showing that postulated membrane properties might perform such computations in a realistic cell model; to show that such membrane properties are consistent with published intracellular recordings; and to spur debate on whether such fast computations actually occur.
>
> 2) To describe dendritic spiking by simple analytical expressions accurate at the 70-80% level, which are based only on physical principles and constants, and which use *no* free fitting parameters. These approximations should demonstrate that the effects are mostly understood, and in addition should provide simple scal-

> ing expressions which can extrapolate results from the few simulated parameter regimes into unsimulated territory. If successful, these approximations and scaling formulae will be a useful back-of-the-envelope adjunct to brute-force simulations, and may provide a conceptual link between the fundamental cable equations and numerical simulation of them.

We must beg the reader's indulgence when many of these mathematical simplifications pile so deep upon each other that the original phenomenon is obscured. We have tried to include "plain English" summaries throughout this chapter and in the Discussion, so that someone uninterested in the mathematical details can skip the equations.

This analysis invokes two unorthodox assumptions: that thin distal dendrites contain strong and fast Hodgkin-Huxley-like conductances (e.g., sodium spiking conductances), and that synaptic conductances (EPSCs) in those dendrites may have local depolarizations of tens of millivolts and durations well less than a millisecond. These assumptions are defended in the Discussion, and some experimental tests for them are proposed.

3.2 Cable Theory at Fast Timescales

The starting point for analysis of an active dendritic tree is the analysis of a passive one. We can consider a typical dendrite as a semi-infinite passive cable of fixed diameter, which contains distributed conductances and capacitances:

$$\begin{aligned}
r_a &= \frac{R_i}{\pi}\left(\frac{2}{d}\right)^2 = \text{intracellular axial resistance } (\Omega/cm) \\
g_m &= \pi d G_m = \frac{1}{r_m} = \text{membrane conductance per fiber length } (\Omega cm)^{-1} \\
c_m &= \pi d C_m = \text{ membrane capacitance per fiber length } (F/cm) \\
d &= \text{branch diameter } (cm)
\end{aligned}$$

(Jack *et al.* 1983). Synapses and active channels are usually characterized by conductances in Siemens (S, e.g., g_m), and passive properties are usually given by resistances (Ω, e.g., r_m). But both units describe the same physical mechanisms, so we will try to use whichever units are most often cited.

For a time-independent (stationary) voltage imposed at one point on the dendrite, the distance over which that voltage decays along the dendrite is the familiar electrotonic space constant

$$\lambda_{DC} = (r_a g_m)^{-1/2}, \tag{3.1}$$

which limits the physical spread of signals much as their temporal duration is limited by the membrane time-constant

$$\tau_m = r_m c_m \tag{3.2}$$

But for time-dependent voltages, the presence of membrane capacitance reduces the distance over which voltages spread below the electrotonic length. In the high-frequency regime, the capacitive (diffusive) term dominates the conductive (dissipative) leak, so that the relevant length constant is given *only* by capacitance, intracellular resistance, and time, and membrane resistance can be ignored. We will investigate this regime of timescale $t \ll \tau_m$, without discussing the case of fast events with strong membrane leaks, as simulated by Jaslove (1992) and Shepherd (1987, 1989).

For this situation we will find a time-dependent length constant λ_t which will approximate the spatial and temporal scale at which voltages spread and decay in response to a brief current pulse, such as might result from a synaptic event or a spike inside a dendrite.

When $r_m = \infty$, the cable equation reduces to the diffusion equation (Jack 1983, eq. 3.7):

$$\frac{\partial^2 V}{\partial x^2} = r_a c_m \frac{\partial V}{\partial t} \tag{3.3}$$

If an instantaneous pulse containing charge Q is injected at time $t = 0$ at the end $x = 0$ of such a semi-infinite cable, then its voltage distribution will be

$$V_\delta(x,t) = \frac{2Q}{\sigma c_m \sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{x}{\sigma}\right)^2\right) \tag{3.4}$$

where

$$\sigma = \sqrt{\frac{2t}{r_a c_m}} \tag{3.5}$$

But real membrane currents are not infinitely brief or strong, so we will adjust this solution to include two timescales: the duration t_0 of the current pulse (assumed to be rectangular), and the time t elapsed after the end of the current pulse. Let us first use the delta-function expression V_δ to find the approximate response V_{app} to a current pulse, which will have the same functional form as V_δ, but will be displaced in time by Δt (so that at $t = 0$ the charge is already somewhat spread out away from $x = 0$).

First let us find Δt. Suppose we have a steady current pulse between $t = -t_0$ and $t = 0$; we can compute the actual voltage at t_0 by the convolution of the delta function with the current's time window. We want to replace that steady pulse with a single delta-function at time $-\Delta t$, so that the approximated voltage at $t = 0$ is the same as the voltage measured at time Δt resulting from a delta-function at $t = 0$:

$$V_{app}(0,0) = V_\delta(0, \Delta t) \tag{3.6}$$

We require that these two methods produce the same voltage peak at $t = 0$, $x = 0$:

$$\begin{aligned} \frac{1}{t_0}\int_{-t_0}^{0} dt' V_\delta(0,-t') &= V_\delta(0,\Delta t) && (3.7)\\ \frac{1}{t_0}\int_{-t_0}^{0} dt' (-t')^{-1/2} &= \Delta t^{-1/2} && (3.8)\\ \frac{t_0}{4} &= \Delta t && (3.9) \end{aligned}$$

So a current pulse of duration t_0 ending at $t = 0$ is approximated by a single delta-function at $t = -t_0/4$, giving us a new σ_t to use in the expression for voltage in the dendrite:

$$V_{app}(x,t) = \frac{2Q}{\sigma_t c_m \sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{x}{\sigma_t}\right)^2\right) \tag{3.10}$$

where

$$\sigma_t = \sqrt{\frac{2(t + t_0/4)}{r_a c_m}} \tag{3.11}$$

for $t > 0$

Now we can find the length-scale over which charge is distributed, by defining a length $\lambda(t)$ over which uniformly distributed charge would have the same voltage as the peak of the approximate distribution,

$$
\begin{aligned}
\frac{Q}{\lambda(t)c_m} &= V_{app}(0,t) && (3.12)\\
\lambda(t) &= \sqrt{\frac{\pi(t+t_0/4)}{r_a c_m}} && (3.13)\\
&= \sqrt{\frac{\pi(t+t_0/4)d}{4R_iC_m}} && (3.14)\\
&\approx 100\mu m\left(\frac{d}{1.0\mu m}\right)^{1/2}\left(\frac{t+t_0/4}{0.25\ ms}\right)^{1/2}\left(\frac{200\Omega cm}{R_i}\right)^{1/2}\cdot\\
&\quad\left(\frac{10^{-6}\ \mu F/cm^2}{C_m}\right)^{1/2} && (3.15)
\end{aligned}
$$

The last expression (giving $\lambda(t) \approx 100\ \mu m$ for a dendrite of diameter 1.0 μm subject to a current pulse of duration 1.0 ms) is intended to allow easy extrapolation to other parameter values. In contrast, the value of λ_{DC} for such a fiber (with $1/G_m = 30k\Omega cm^2$) is much longer,

$$\lambda_{DC} = 610\ \mu m \qquad (3.16)$$

We can define in a similar fashion the total capacitance $C(t)$ charged by such a pulse, again assuming a constant voltage spread over a length $\lambda(t)$:

$$
\begin{aligned}
C(t) &= c_m\lambda(t) && (3.17)\\
&= \frac{(\pi d)^{3/2}}{2}\left(\frac{C_m(t+t_0/4)}{R_i}\right)^{1/2} && (3.18)\\
&\approx 3.1\times10^{-12}F\left(\frac{d}{1.0\mu m}\right)^{3/2}\left(\frac{t+t_0/4}{0.25\ ms}\right)^{1/2}\left(\frac{200\Omega cm}{R_i}\right)^{1/2}\cdot\\
&\quad\left(\frac{C_m}{10^{-6}\ \mu F/cm^2}\right)^{1/2} && (3.19)
\end{aligned}
$$

With this capacitance and the pulse charge Q we can estimate the voltage as a function of time, which we will do in the following sections. Although a real current pulse may have a complex shape, and a real cell will contain branches of various diameters and geometries, this approximation shows the dominant effect: the capacitance charged by a fast pulse will be smaller than that charged by a slow pulse, and hence the peak voltage will be greater for the same amount of charge. And that voltage decays dramatically at the timescale of the pulse itself—rather than at the much slower timescale of τ_m—allowing very strong and quickly-repolarizing EPSPs to exist inside dendrites. Such fast, localized depolarizations could in turn trigger voltage-dependent conductances in the dendrites, enabling the cell to perform (in principle) many simultaneous, near-independent computations at fast timescales.

3.2.1 Simulated Pyramidal Cell

Because the previous and following expressions are intended to augment (rather than replace) numerical simulation of the membrane equations, they will be compared to a compartmental simulation of an anatomically reconstructed and physiologically characterized layer 5 pyramidal cell from cat striate cortex. Simulations of this cell model *without* active dendrites have previously been successfully compared to the original cell's behavior at rest and under DC current clamp (Bernander *et al.* 1991). While this is still an entirely theoretical exercise—the comparison of two types of theory is in no sense an experimental test—it will help clarify the strengths and weaknesses of the approximations.

The compartmental model contained 1890 compartments (Figure 3.1). The soma and basal dendrites (where active conductances could reside) were modelled as connected cylinders with length at most 10 μm, using the NEURON program provided by Hines (1990). Apical dendrites remained passive, and had longer compartments (up to 200 μm). Passive cell properties were those described above; spines were not included, but a correction for them is made in section 3.7.

The passive time constant was about $\tau = 30\ ms$. (The actual membrane conductance was not uniform across the entire cell, but was slightly stronger within 60 μm of the soma. That conductance was equivalent to the steady-state conductance which would be produced by 500 $GABA_A$ and 500 $GABA_B$ synapses firing at 1 Hz each. See Bernander *et al.* 1991 for further details). The membrane capacitance was 1 $\mu F/cm^2$, and intracellular resistivity 200 Ωcm. Integration time-steps were 50 μsec; finer time-steps changed spiking amplitudes and time-courses by less than 5% (within the range of approximation desired here), but slowed computation significantly from the "minute-per-6 ms" usually used.

The only active properties investigated were those of the Hodgkin-Huxley-like equations (Hines 1990, and Appendix G), for which both sodium and potassium currents had fixed, voltage-independent equilibration times ($\tau(h) = 0.5\ ms, \tau(m) = 0.05\ ms, \tau(K) = 2.0\ ms$) and conductances ($\overline{G}_{Na}, \overline{G}_K$) which were adjustable as parameters. The reversal potentials of those conductances were $E_{Na} = 50\ mV$ and $E_K = -95\ mV$; the threshold potential at which each opened was $V_{1/2} = -40\ mV$. Dendritic spikes were triggered in the center of terminal branches by conductance alpha-functions with peak time 0.1 ms and peak conductance 12 nS (unless otherwise specified); these parameters

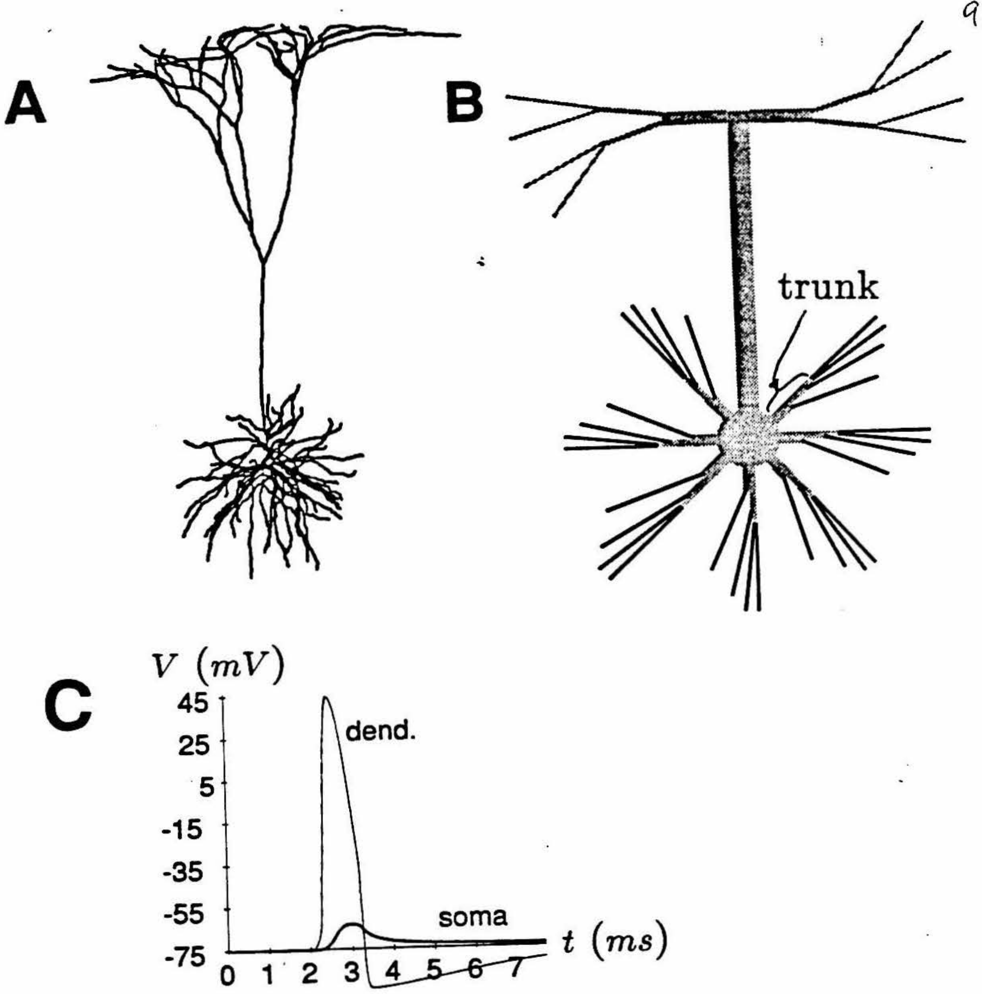

Figure 3.1: (A) A layer **V** pyramidal cell, as recorded, reconstructed, and simulated by Bernander *et al.* (1991) from cat striate cortex. (B) A caricature of the compartmental model of that cell, as modelled by Bernander *et al.* (1992) and used here. All passive-membrane areas in the simulation of dendritic spiking are shown in grey; Hodgkin-Huxley-like active membrane conductances were placed only on the 44 basal terminal branches (in black). The thicker, shorter regions between the soma and the black terminal branches are dendritic "trunks." (C) An action potential in a single terminal branch (thin curve) will create a brief, small depolarization at the soma (thick curve).

are consistent with some intracellular recordings, as discussed in section 3.9.

In most of these simulations the soma was left passive, because the simulations usually characterized only small somatic depolarizations well below firing threshold; an active soma makes virtually no difference. But the most distal basal dendrites almost always contained active conductances (I_{Na} and I_{DR}), in order to simulate the influence of dendritic spikes on each other and on the soma.

Because somatic spiking is triggered mainly by instantaneous membrane potential, the primary variable measured here was depolarization. ΔV_{soma} represents the peak increase in voltage above the initial "resting" value $E_{rest} = -75\ mV$ or $-65\ mV$,

$$\Delta V_{soma} = \max_t \left[V_{soma} - E_{rest}\right] \tag{3.20}$$

In the following expressions, E_{rest} often approximates V_{soma}, since their difference is usually much smaller than the potentials which dominate current flow.

3.3 Fast EPSPs in Thin Terminal Branches

The basal dendritic tree of this pyramidal cell has a structure which easily lends itself to a some simplification. The soma gives rise to ten thick basal dendritic trunks, each of which typically branches several times in the proximal 20-30 μm before terminating in long ($> 200\ \mu m$), thin ($< 1\ \mu m$) distal terminal branches (Figure 3.1). When a brief synaptic conductance opens in the center of such a terminal branch, only a small region is strongly depolarized, and for only a brief length of time.

The large depolarization occurs because the dendrite is thin, so that only a small, nearby capacitance can be charged by a synaptic event. The rapid repolarization occurs because the soma and other dendrites together have a much larger capacitance, onto which the synaptic charge diffuses. As this section will show, this repolarization does **not** require a fast membrane time-constant, and indeed still occurs in the limit of zero membrane conductance. Such large, quickly-repolarizing synaptic events will *not* occur at the soma, because the soma has a larger capacitance (hence a smaller peak depolarization), and because the thin dendrites around the soma do not provide much of a "sink" onto which the somatic depolarization can quickly equilibrate. Rall (1964) noted that this capacitive effect allows a faster-than-exponential EPSP decay at the soma, which could let the cell selectively fire in response to precisely timed excitatory synaptic events. Such precision in all these cases stems from his discovery that local EPSPs in branched structures always decay faster than τ_m (as reviewed in Jack *et al.* 1983).

What is the peak depolarization and time-course of a local EPSP inside a thin dendrite, if we assume only capacitive diffusion of charge (i.e., no leak terms)? We derive an approximation which gives only a very coarse estimate of the amplitude and timescale of fast EPSPs inside thin dendrites, and compare them to simulated EPSPs.

Suppose the synaptic conductance has the traditional form of an alpha-function with peak conductance g_{peak} and time-to-peak t_{peak},

$$g_{syn}(t) = \left(\frac{g_{peak}}{t_{peak}}\right) t \exp(1 - t/t_{peak}), \tag{3.21}$$

so that $g_{syn}(t_{peak}) = g_{peak}$. We can approximate the synaptic current by assuming that the local peak depolarization ΔV_{dend} is small relative to the synaptic driving potential $E_{syn} - E_{rest}$, so that the synaptic current is proportional to the synaptic conductance:

$$I_{syn}(t) \approx (E_{syn} - E_{rest}) \frac{g_{peak}}{t_{peak}} t \exp(1 - t/t_{peak}) \tag{3.22}$$

The true current only reaches zero at infinite time, so we can take the end of the current pulse as occurring at time $t = 3t_{peak}$, when over 80% of the eventual charge has passed through the synapse.

$$Q_{syn} \approx \int_0^{3t_{peak}} I_{syn}(t)dt \tag{3.23}$$

$$\approx 0.8 e g_{peak} t_{peak} (E_{syn} - E_{rest}) \tag{3.24}$$

This synapse does not pass constant current (as assumed above), but has a current peak at t_{peak} and a voltage peak (in simulations) at about $2t_{peak}$. But simplified expressions for charge distribution (eqs. 3.10 and 3.13-3.15) do not account for the structure of the current pulse; they only describe a pulse *after* it has finished. Our "best guess" is to start our approximated EPSP at time $3t_{peak}$ (when most of the synpatic current is finished), but to assume that all of that charge arrived in an impulse at the time of the actual peak current (i.e., at t_{peak}, i.e., $\Delta t = 2t_{peak}$ before we "start" the synapse, over twice as long as the value $\Delta t = 3t_{peak}/4$ which would occur in V_{app}, eq. 3.10). Because the charge distributes itself over a capacitance $C(t)$ on both sides of the synaptic site, our predicted EPSP there will be

$$\begin{aligned} V(0, t > 3t_{peak}) &\approx \frac{Q_{syn}}{2C(t)} && (3.25) \\ &= \frac{0.8 e R_i^{1/2} g_{peak} t_{peak} (E_{syn} - E_{rest})}{(\pi d)^{3/2} \left(C_m (t - 3t_{peak} + 2t_{peak})\right)^{1/2}} && (3.26) \end{aligned}$$

We cannot account for the rising part of the EPSP, during which an impulse of charge is explicitly inappropriate; such a rise is not even graphed in Figure 3.2. This approximation has many flaws: the approximation neglects current flow out the dendrite's end into the soma, the synapse saturates towards E_{rev}, and we ignore the 20% of charge contained in the alplha-function's tail. But eq. 3.26 predicts the magnitude and initial time-course of simulated EPSP's (Figure 3.2) to within 20%. The simulation here used a terminal branch of diameter 1.05 μm, with a single synaptic event located halfway along its 226 μm length ($E_{syn} = 0\ mV$; $g_{peak} = 6\ nS$; $0.05 \leq t_{peak} \leq 0.4\ ms$). The prediction worked best for short times, for which the charge was confined to the dendrite. For times $t > 3\ ms$, the simulated EPSP dropped off much faster than predicted by eq. 3.26 because the proximal end of the dendrite was effectively grounded by the soma.

This model is also accurate in predicting the EPSP decay. We can quantify the EPSP duration by the time $\tau_{1/2}$ from its peak to half its peak amplitude. The simple model decays by 1/2 when

$$\begin{aligned} \sqrt{\frac{2t_{peak}}{\tau_{1/2} + 2t_{peak}}} &= \frac{1}{2} && (3.27) \\ \tau_{1/2} &= 6t_{peak} && (3.28) \end{aligned}$$

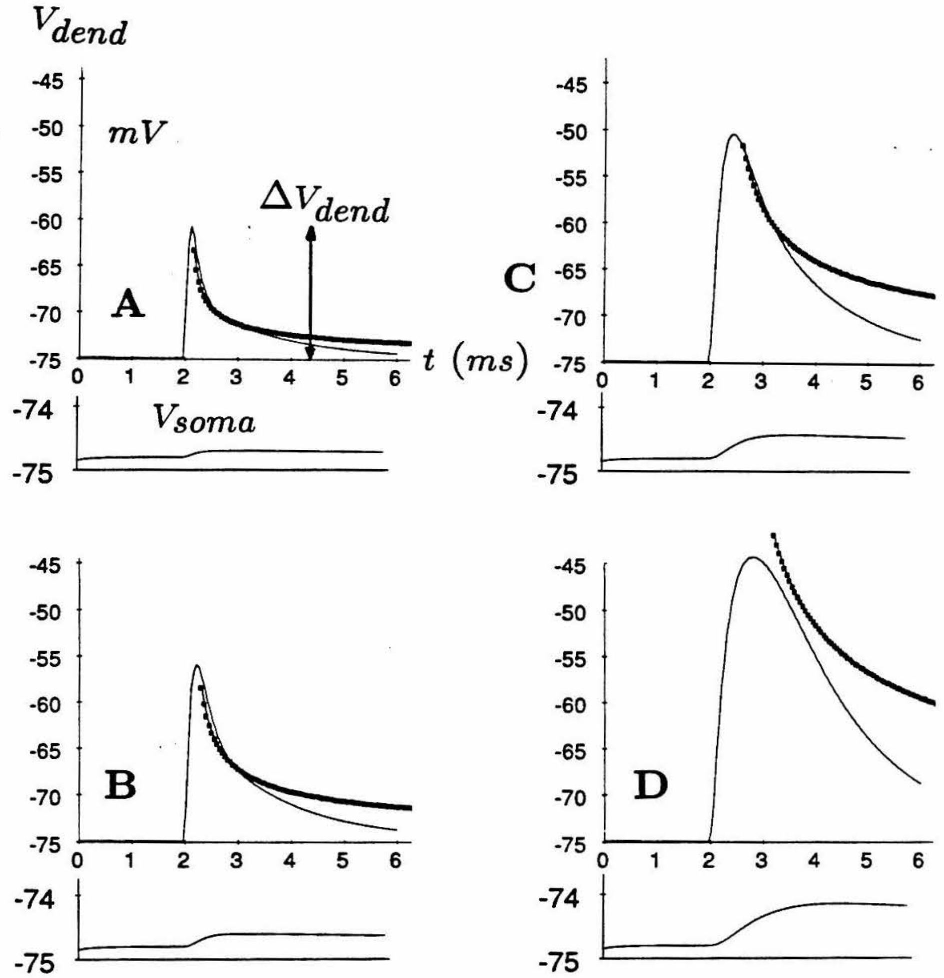

Figure 3.2: (A) A fast synaptic conductance in a terminal dendritic branch, over 100 μm from the soma, can create a strong, brief, local depolarization V_{dend}. Curves are simulated EPSP's, with $g_{peak} = 6\ nS$ and various t_{peak}; connected points are predictions by a simple model, which assumes that the peak depolarization occurs at time $4t_{peak}$, carries total charge $0.8eg_{peak}t_{peak}(E_{syn} - E_{rest})$ (eq. 3.25), and decays with time due to charge diffusion in an infinite capacitive cable. (A) $t_{peak} = 0.05\ ms$, (B) $t_{peak} = 0.1\ ms$, (C) $t_{peak} = 0.2\ ms$, (D) $t_{peak} = 0.4\ ms$. Because the simple model is based on an infinitely strong current impulse, it cannot account for the rising phase of an EPSP. But this model does explain the fast decay and the fact that the peak voltage scales roughly with $\sqrt{t_{peak}}$, rather than linearly in t_{peak} as the total charge does. Note that somatic depolarizations (shown below each curve) are typically a hundredfold weaker than dendritic EPSPs, and lack the fast repolarization.

This very narrow pulse-width matches almost exactly the $t_{1/2}$ of EPSPs with fast $t_{peak} = 0.05$ and $0.1\ ms$, and agrees within 25% for the longer t_{peak}. Such fast EPSPs, which result from capacitive charge-equilization rather than resistive decay, allow in principle submillisecond coincidence-detection among individual EPSP's. If the dendritic shaft were active, two simultaneous, co-localized synaptic events could drive the local membrane above threshold (e.g., $\Delta V_{peak} \approx 2 \times 20\ mV$ for $t_{peak} = 0.1\ ms$), but those same two events would not initiate a spike if separated by only a millisecond. In contrast, the EPSP at the soma would decay at least ten-fold more slowly ($\tau_m \approx 15 - 30\ ms$, Bernander *et al.* 1990), rather than in the submillisecond range. In addition, these capacitive effects mean that the local peak amplitude of an EPSP scales as $\sqrt{t_{peak}}$, rather than linearly as the total charge does.

While *locally* only two events might fire a dendrite, the *somatic* depolarization of the single synaptic event is a hundredfold smaller ($\Delta V_{soma} \approx 0.2\ mV$) and lasts far longer (decaying with $\tau_m \gg \tau_{1/2}$), so that the soma would need to temporally integrate many events to fire (a more rigorous distinction between temporal integration and coincidence-detection will appear in section 3.8). The small *somatic* depolarization for the $t_{peak} = 0.1\ ms$ *dendritic* synapse and its fast somatic rise-time are in the range of monosynaptic depolarizations observed in cortical somatic recordings (see section 3.9).

3.4 Active Dendritic Terminal Branches

What happens when active conductances in one of those distal terminal branches create a dendritic spike? One trustworthy solution involves solving the Hodgkin-Huxley equations (which are themselves approximations) in a one-dimensional, dissipative, dispersive membrane. Here we will only describe a very approximate view of this dendritic spiking.

Let us suppose that one entire terminal branch—but not the dendritic trunk, other terminal branches, or the soma—is homogeneously coated with Hodgkin-Huxley-like channels, which can spike in response to synaptic currents (the black areas in Figure 3.1 B). Suppose (for a moment) that this terminal branch is connected directly to the soma (as is only one terminal branch in the reconstructed cell), and that the soma remains at resting potential. How much of the terminal branch will depolarize above threshold? How much current and charge will the spiking branch deliver to the soma, and to the other terminal branches? On what parameters do these results depend?

First, we must ask which properties of the terminal branch itself will dominate: capacitive or resistive? Given a strong peak sodium conductance (during spiking) of $\overline{G}_{Na} = 0.2\ Scm^{-2}$, we calculate a temporarily fast membrane time-constant

$$\begin{aligned} \tau_m &= C_m/\overline{G}_{Na} && (3.29)\\ &= 1\ \mu F/0.2\ S && (3.30)\\ \tau_m &= .005\ ms && (3.31) \end{aligned}$$

This time-constant is far faster than the simulated time-constant of the sodium channels (here $\tau(h) = 0.5\ ms$), suggesting that the capacitance of the terminal branch membrane itself will not seriously attenuate the spike's voltage within the spiking terminal branch. Furthermore, the passive membrane conductance (about $30\ k\Omega cm^2$) is so much less than the peak active conductances that passive terms can be ignored; the only terms we will need inside the spiking branch are the axial and peak membrane conductances, r_a and $\overline{G}_{Na}$.

However, the capacitance of the soma and other terminal branches together will be substantial enough that they will *not* depolarize very much during the single dendritic spike. (We will assume this for now; this assumption will be verified by the results and simulations following). So we ask: in what manner will the terminal branch sustain a spike, given the boundry conditions that the proximal end of it is effectively kept near resting potential, and the distal end is saturated near $+50\ mV$? How much charge will it deliver to the soma during the spiking event?

The best answers to these questions come from simulations. The actual behavior of a simulated dendritic spike is very complex: a strong synaptic event in the center of the terminal branch initiates positive feedback depolarization in nonlinear sodium and potassium conductances, which propagates away from the synaptic site in both proximal and distal directions. As a result, the magnitudes and time-courses of both currents vary dramatically from one end of the terminal branch to the other.

To represent the net effect of these complicated interactions we can use several strikingly simple approximations. First, we can suppose that instead of a propagating action potential we have a single canonical Hodgkin-Huxley-type event, in which peak currents may vary along the branch, but all currents reach their peak values simultaneously. Then, we replace the complicated and temporally overlapping conductance curves $g_{Na}(t)$ and $g_K(t)$ with two non-overlapping triangular functions, whose peak value is assumed to lie at half the peak value used in the simulation $(\overline{G}_{Na}, \overline{G}_K)$ and whose duration is the conductance's simulated time-constant $(\tau(h), \tau(K))$. This approximation is a very crude one, failing to account well for the sodium conductance's time course or for the potassium conductance's amplitude (Figure 3.3), but it will prove sufficient to explain many of the spike's influences. Because the sodium and potassium conductances are simplified not to overlap in time, the peak somatic depolarization occurs before any potassium current flows, and should be independent of $\overline{G}_K$ (in simulations, changing $\overline{G}_K$ from $\overline{G}_{Na}/2$ to $2\overline{G}_{Na}$—a factor of four—changed ΔV_{soma} by less than 10%, validating this approximation). Unless otherwise noted, all simulations here used strong $\overline{G}_K = 2\overline{G}_{Na}$, as justified later in section 3.6.

Now we will estimate the sodium current's effect on the soma. The voltage profile $V(x)$ and transverse conductance $g_{Na}(x)$ of this terminal branch are clearly inhomogenous, V being near rest ($-75\ mV$) at the soma end and near E_{Na} ($= +50\ mV$) at the distal end. An axial current into the soma results from the sodium conductance. That conductance is open only *above* the threshold voltage $V_{1/2}$ (section 3.2), so that the axial current does not depend on E_{rest}, but only on $E_{Na} - V_{1/2}$. To find that current we will crudely characterize the

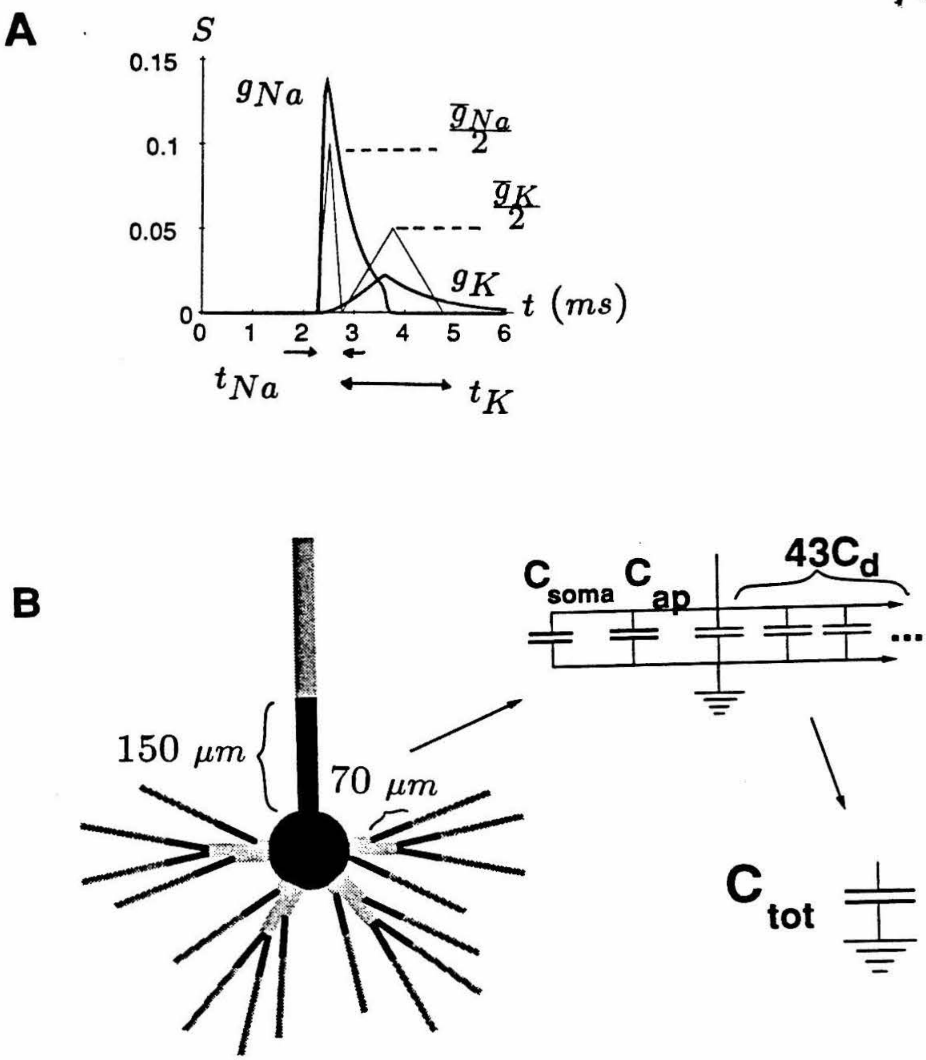

Figure 3.3: (A) Simulated active sodium and potassium currents inside a spiking basal dendrite (thick curves) are produced by Hodgkin-Huxley-like equations containing characteristic temporal responses (τ_h, τ_K) and maximum conductance ($\overline{G}_{Na}, \overline{G}_K$). These conductances overlap in time, and vary in magnitude and shape along the dendrite. They can be approximated by two non-overlapping triangle functions (thin curves) with similar duration (τ_h, τ_K) and peak conductance ($\overline{G}_{Na}/2, \overline{G}_K/2$). The obviously poor match between the simulated conductances and these crude approximations does not prevent the approximations from predicting most of the influences of dendritic spikes on the soma and other dendrites. (B) A brief pulse (≈ 1 ms) of depolarizing current from a dendritic spike only reaches the more proximal regions of the cell; the charge is deposited on a smaller region (relative to the whole cell) and produces a relatively higher peak depolarization. An approximation to the capacitance charged is the parallel combination of the capacitance of the soma and the most proximal portions (black) of the apical dendrite (150 μm) and of the 43 other terminal basal branches (70 μm each).

active portion of the dendrite by a single input conductance in series with a battery $E_{Na} - V_{1/2}$. That input conductance is roughly given by two cable properties: (1) the intracellular axial resistance r_a (Ω/cm) and (2) the mean sodium conductance $\approx 0.5\overline{G}_{Na}\pi d$ $(\Omega cm)^{-1}$ (as averaged over the voltage range E_{Na} to $V_{1/2}$). So the current out of the active region will be approximately

$$\begin{aligned} i_{cs} &\approx (E_{Na} - V_{1/2})\left(\frac{\overline{G}_{Na}\pi d}{2r_a}\right)^{1/2} && (3.32) \\ &= (E_{Na} - V_{1/2})\left(\frac{\overline{G}_{Na}\pi^2 d^3}{8R_i}\right)^{1/2} && (3.33) \end{aligned}$$

This is the "current-source" (cs) approximation (Figure 3.4). The rightmost term is the input conductance of a semi-infinite cable with "leak" $\overline{G}_{Na}/2$ (Jack *et al.* 1983). This input conductance assumes that the entire active portion of the dendrite contains uniform, half-open sodium channels, although clearly the most proximal channels are fully closed and the most distal ones fully open. Most importantly, this approximation does preserve the scaling properties, showing how input conductance would change in response to changes in dendrite diameter, peak membrane conductance, or cytoplasmic resistivity.

For example, the 1.0 μm diameter dendrite simulated in Fig 3.5 is predicted to pass a current of 3.15 nA; the simulated value, as measured from the maximum proximal slope dV/dx, is 3.0 nA. No free parameters were needed to get this agreement between approximation and simulation.

Because the voltage characteristics of spiking sodium channels are fairly constant across various cell types, the driving voltage of about 100 mV should not vary. But the *other* parameters (dendrite geometry and sodium conductance) are more germane to the scale-dependence of dendritic spiking. In particular,

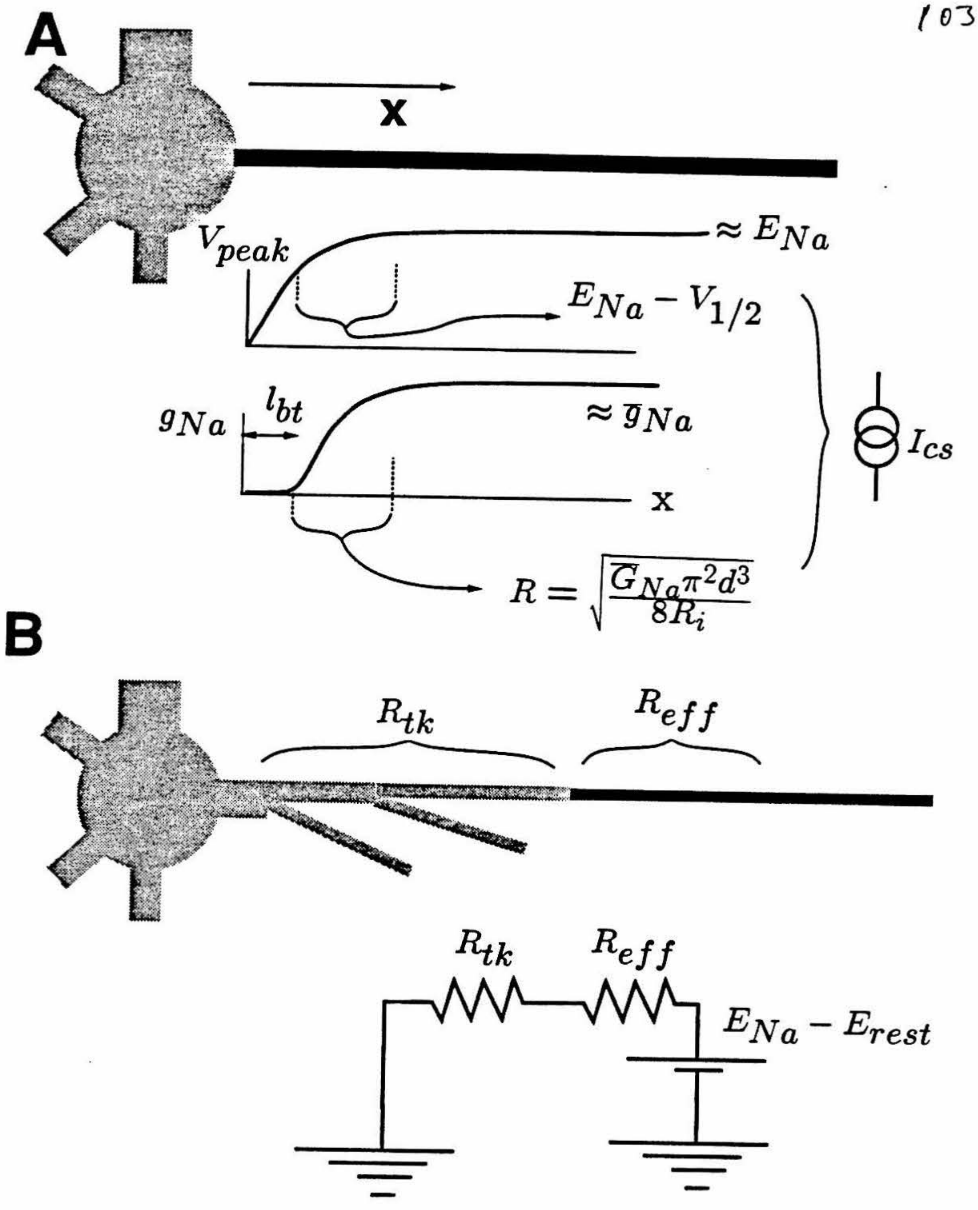

Figure 3.4: (A) A single spiking terminal branch (black) can be idealized as connected directly to the soma, with the peak voltage $V(x)$ and peak sodium conductance $g_{Na}(x)$ varying along the branch's length. Axial current to the soma is produced in the region where $V(x)$ is curved and g_{Na} is nonzero. An approximation to the peak current comes from choosing a representative input conductance and voltage for that region, and combining those into a current-source whose magnitude depends only on the membrane conductance and dendritic diameter (section 3.4). (B) When the terminal branch (black) is connected to the soma by a long dendritic trunk (grey), the branch's entire length may be driven above threshold, rendering the above current-source approximation invalid. In this case the best model for dendritic spiking is two resistors in series between E_{rest} and E_{Na}, representing the active membrane (R_{eff}) and the passive trunk (R_{tk}).

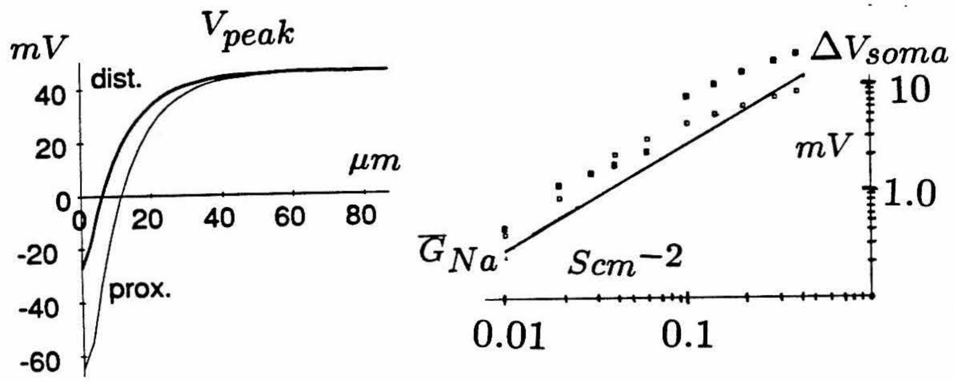

Figure 3.5: (left) Peak voltages $V(x)$ along two basal dendritic terminal branches during spiking events (thin line, a branch connected directly to soma; thick line, a moderately distal branch). Current into the soma is contributed by the curved regions; potentials below sodium threshold exist only in the left few microns of the proximal branch. (right) Somatic depolarizations for the same two spiking branches at various peak sodium conductance. The peak somatic potential change ΔV_{soma} due to spikes in the proximal branch (filled squares) or the more distal branch (open squares) both follow the prediction (thin line, eq. 3.46) that somatic depolarization increases as $\sqrt{\overline{G}_{Na}}$.

this approximation shows that the current injected by the dendrite into the soma is relatively insensitive to changes in peak sodium conductance, but is more sensitive to dendritic terminal branch diameter d, which affects both sodium conductance (per cm) and intracellular axial resistance:

$$
\begin{aligned}
i_{cs} &\propto \left(\frac{\overline{G}_{Na}\pi^2 d^3}{8r_a}\right)^{1/2} && (3.34)\\
\text{So } i_{cs} &\propto (\overline{G}_{Na})^{1/2} && (3.35)\\
i_{cs} &\propto d^{3/2} && (3.36)
\end{aligned}
$$

We can also find the length of the proximal dendrite which remains effectively "grounded" by the soma, and does not fire (called l_{bt}, for "length below threshold"). In this region ΔV_{dend} remains below $V_{1/2}$, so the axial voltage drop ΔV_{bt} due to axial current i_{cs} along that length is $\Delta V_{bt} = E_{rest} - V_{1/2} \approx 25\ mV$, so that

$$
\begin{aligned}
\Delta V_{bt} &\approx l_{bt} r_a i_{cs} && (3.37)\\
l_{bt} &\approx 5\ \mu m && (3.38)
\end{aligned}
$$

So only the few most proximal microns of the dendrite are below firing threshold. As can be seen in the simulation (Figure 3.5), the axial current in this region is constant (i.e., voltage is nearly linear in x) at a value of about $3 \times 10^{-9}\ A$. Note that Figure 3.5 shows the *peak* voltage values, which occur at slightly different times, rather than a particular "snapshot" at any one time.

Figure 5 also shows that for these parameters, almost all the current into the soma is generated in the most proximal 30 μm of the terminal branch, where dV/dx and hence I are decreasing with x (i.e., curved). Regions more distal than 30 μ have constant peak voltage, and thus contain little net axial current. Those peak voltages in the proximal 30 μm are simultaneous on the spike's timescale, so that we do not need to consider propagation of the action potential. The fact that terminal branches longer than 30 μm did not increase somatic depolarization was verified by simulation. A truncated dendrite only 30 μm long generated a ΔV_{soma} only 10% less than that generated by an otherwise identical dendrite of length 274 μm.

The above approximation was made for a dendritic terminal branch directly connected to the soma, so that its proximal end was held low, near E_{rest}. But *as long as* that end of the terminal branch was below the activation potential of the sodium currents, its potential does not enter into i_{cs}, so that the spiking terminal branch can be viewed as a current (or charge) source—independent of its dendritic location or resting potential—rather than as a voltage source. In this approximation the somatic depolarization ΔV_{soma} caused by a dendritic terminal branch should *not* depend on the intervening length of dendritic trunk, provided that the terminal branch's proximal end remains below $V_{1/2}$. But if there were a length of passive dendritic trunk with high resistance R_{tk} between the terminal branch and the soma (Figure 3.4 B), would the peak current to the soma be reduced below i_{cs}?

Yes. In the previous discussion, i_{cs} can be interpreted as resulting from a resistance R_{eff} between the location where the sodium conductances begin to open (i.e., $V_{1/2} = -40\ mV$) and E_{Na} (Figure 3.4):

$$R_{eff} \equiv \frac{E_{Na} - V_{1/2}}{i_{cs}} \tag{3.39}$$

But this second model of spiking current is of *two* resistances R_{eff} and R_{tk} in series from E_{Na} to E_{rest} (Figure 3.4), so the resulting current from the above-threshold dendrite is approximated by a "resistive-dendrite" expression,

$$i_{res} \quad = \quad \frac{E_{Na} - E_{rest}}{R_{tk} + R_{eff}} \tag{3.40}$$

We now have two separate models for the current from a spiking dendrite arriving at the soma: the current-source model (eq. 3.32) and the resistive-dendrite model (eq. 3.40). We can combine them into a single prediction by noting that each model tends to **over**predict current in the region in which it is not valid (e.g., i_{res} is too large as $R_{tk} \ll R_{eff}$, and i_{cs} is too large when $R_{tk} \approx R_{eff}$). So a simple approximation is to calculate *both* currents for a given terminal branch, and then to choose the *minimum* of the two:

$$i_{ax} \quad = \quad \min\left[i_{cs}, i_{res}\right] \tag{3.41}$$

This will be the predicted peak spiking current from an active dendrite to the soma, upon which many of this chapter's further approximations will be based.

3.4.1 Somatic Depolarization from a Spike

The net depolarization at the soma imparted by this current pulse can be calculated by knowing the total capacitance of the soma and of its proximal dendritic terminal branches. That pulse has duration $\tau(h) = 0.5\ ms$, so the lengths of dendrite charged at its conclusion will be $\lambda(t) = 70\ \mu m$ for the distal basal dendrites and $\lambda(t) = 150$ for the apical dendrite. We can then assume that the charge is evenly distributed over the soma (of area $1230\ \mu m^2$), over the proximal 70 μm of each of the 43 other (non-spiking) terminal branches (mean diameter about 0.88 μm), and over the 150 μm of the apical dendrite (diameter 4.5 μm), for a total capacitance charged of

$$C_{tot} \approx \left[17 \cdot 23\ \mu m^2 + 43(0.88 \cdot 70\ \mu m^2) + 4.5 \cdot 150 \mu m^2\right] \pi \times 1\ \mu F/cm^2 \tag{3.42}$$

$$\approx 1.1 \times 10^{-10}\ F \tag{3.43}$$

For simplicity, we ignore the variation in dendritic size and branching structure, and assume that each dendritic spike will create a voltage peak ΔV_{soma} in accordance with the definition of $C(t)$, eqs. 3.17-19.

For the highly idealized triangular sodium conductance function outlined above (which will produce an axial current with similar time-course), the net (integrated) charge Q deposited on C_{tot} will be the area under the triangle: half the product of the peak current and the current's duration. So the approximate net depolarization of the cell and dendrites due to the single dendritic spike will be given by

$$Q \approx i_{ax}\tau(h)/2 \tag{3.44}$$

$$\Delta V_{soma} = \frac{Q}{C_{tot}} \tag{3.45}$$

$$\approx \frac{\tau(h) i_{ax}}{2C_{tot}} \tag{3.46}$$

This prediction was tested for each of the 44 terminal basal branches separately, at both high and low peak sodium conductances ("strong HH," $\overline{G}_{Na} = 0.2\ S - cm^{-2}$; "weak HH," $0.033\ Scm^{-2}$). Figure 3.6 shows the peak somatic depolarization plotted against its value predicted from branch diameter and trunk resistance (eq. 3.46), so that a good match lies on the line with unit slope. Most depolarizations lie within 20% of predicted values, showing the inherent soundness of the approximations. Histograms of ΔV_{soma} show clusters around 6 mV ("strong HH") or 3 mV ("weak HH"); all of these isolated dendritic spikes are too small to fire a cell from E_{rest}.

The dependence of somatic depolarization on peak sodium conductance appears in Figure 3.5, which shows ΔV_{soma} predicted and simulated for two dendrites of nearly the same diameter (1.0 μm): one proximal, one distal. For both dendrites, ΔV_{soma} varies as $\sqrt{\overline{G}_{Na}}$, the weak dependence predicted by eq. 3.35 above. These simulations all are for active conductances which exist only on one branch at a time.

3.5 Coupling Between Active Terminal Branches

If *all* the terminal branches contain active conductances (with passive proximal trunks), can one terminal branch's firing induce another terminal branch to fire, or are the terminal branches effectively "de-coupled" from one another?

3.5.1 Predicted Depolarization of Neighboring Branches

Let us suppose terminal branch B_1 fires, sending current down its connecting trunk to the soma (Figure 3.7). We examine its effect upon terminal branch B_2, connected to the soma by another trunk. If the trunks do not overlap, then the only common link between B_1 and B_2 is the soma, and B_2 will have a depolarization no greater than ΔV_{soma} (which, as we saw above, is too small to trigger a spike in B_2). But if the trunks share some common section T, with resistance R_T, then the voltage drop across T due to B_1's firing will further raise the voltage at B_2.

Let us first use the simplest possible approximation: we assume that the peak current from B_1 (eq. 3.41) is the dominant effect, and that it reaches equilibrium in the local trunks, so that the depolarization at B_2 is approximated by the depolarization at the distal end of T (Figure 3.7 B) (this approximation neglects the dendrites' capacitive and leak effects). Then the peak voltage drop across trunk T (ΔV_T) is given by its axial resistance and the peak current through it:

$$\Delta V_T = i_{ax} R_T \tag{3.47}$$

The other component of the depolarization at B_2 is the somatic depolarization ΔV_{soma}. We have calculated its peak value above (eq. 3.46), but that peak occurs *after* the peak V_T has passed, so that adding the two peak voltages ($V_T + \Delta V_{soma}$) would give an overestimate of the true peak voltage reached at B_2. Instead we will interpolate between two limiting cases. If the somatic voltage drop dominates (i.e., the shared trunk resistance is small), then we estimate the peak voltage in the dendrite to be the somatic peak voltage (eq. 3.46; Figure 3.7C). But if the trunk voltage drop dominates ($\Delta V_T \gg \Delta V_{soma}$), then we take the peak in the dendrite to be the trunk voltage drop, added to the somatic voltage which exists at that time. Because the current pulse is a triangle, the peak spiking current occurs when only half the eventual somatic

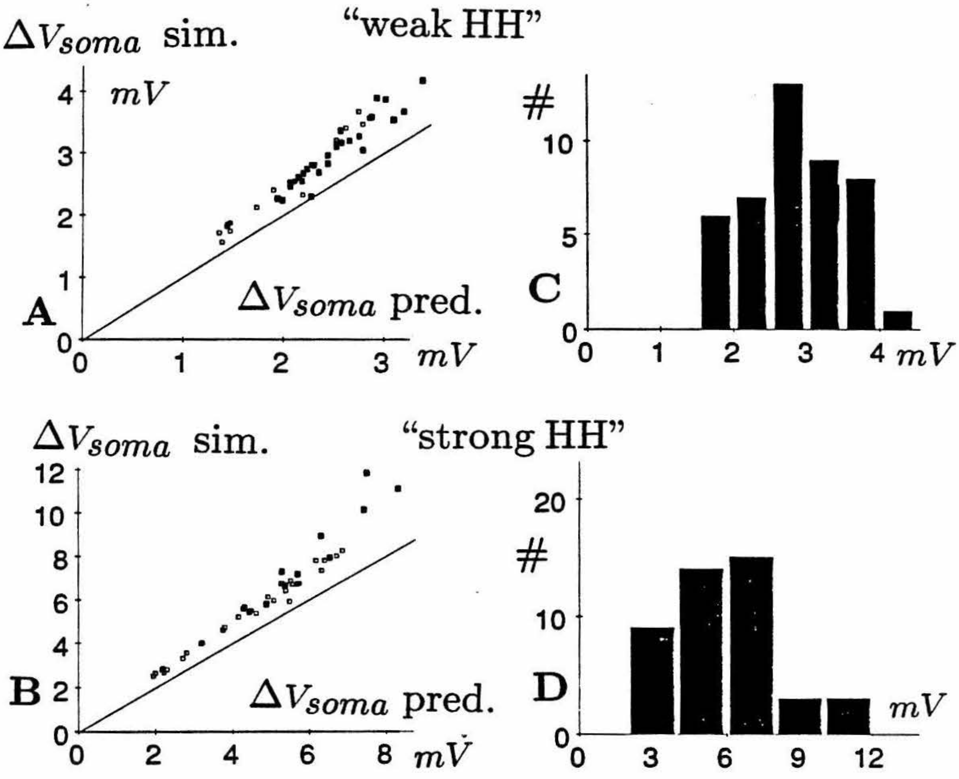

Figure 3.6: Spikes in individual terminal branches produce somatic depolarizations consistent with predictions. Each of the 44 terminal branches was separately given active Hodgkin-Huxley-like kinetics and triggered to spike, while its 43 neighbors and the remaining cell contained only passive membrane properties. Plots of ΔV_{soma} (resulting from those dendritic spikes) against the values predicted by eq. 3.46 lie near the diagonal line of unit slope, which represents a perfect prediction (A, $\overline{G}_{Na} = 0.033\ S\ cm^{-2}$; B, $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$). Histograms show that ΔV_{soma} cluster around 3 mV (C, $\overline{G}_{Na} = 0.033\ Scm^{-2}$) or 6 mV (D, $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$). In both A and B, filled points represent predictions by the current-source model (eqs. 3.32 and 3.41), while open points were predicted by the resistive model (eqs. 3.40 and 3.41).

depolarization has occurred.

So when R_T is large, we assume ΔV_{soma} is half the value predicted by eq. 3.46, and we add this to the trunk's voltage drop to get the peak voltage increase at B_2 due to B_1:

$$\begin{aligned}\Delta V_{B_2} &= \frac{\Delta V_{soma}}{2} + \Delta V_T && (3.48)\\ &\approx \frac{\tau(h) i_{ax}}{4C_{tot}} + i_{ax} R_T && (3.49)\end{aligned}$$

A precise interpolation between these two cases ($\Delta V_{soma} \gg \Delta V_T$ and $\Delta V_{soma} \ll \Delta V_T$) would require a precise knowledge of the time-course of spiking current. For this inexact model, we can interpolate linearly between the two cases by a variable s, which gives the fraction of somatic peak voltage ($0.5 \leq s \leq 1$) added to the peak trunk voltage drop.

$$\begin{aligned}s &\equiv 1 - \frac{1}{2}\left(\frac{\Delta V_T}{\Delta V_{soma} + \Delta V_T}\right) && (3.50)\\ &= 1 - \frac{1}{2}\left(\frac{R_T}{\frac{\tau(h)}{2C_{tot}} + R_T}\right) && (3.51)\end{aligned}$$

(i.e., $s = 0.5$ when ΔV_T dominates, and $s \approx 1.0$ when $\Delta V_T \approx 0$). The interpolated voltage change at any basal dendrite due to the spiking of any other is then about

$$\Delta V_{B_2} \approx i_{ax}\left(R_T + s\frac{\tau(h)}{2C_{tot}}\right) \qquad (3.52)$$

This approximation is very simplistic, in that it attempts to model highly transient events by interpolating between DC equations. But it works rather well. After triggering a spike in one of the most distal terminal branches on each of nine basal trunks, we recorded the peak voltages at the proximal end of several representative terminal branches sharing part of the same trunk (only about half the measurement points were on the current-carrying trunk T; some measurement points were over 90 μm from it). Those peak voltages were predicted by eq. 3.52, using R_T along the common trunk T and calculating i_{ax} for the spiking terminal branch (the other parameters $\tau(h)$ and C_{tot} were the same as before).

In this simulation the other untriggered dendritic shafts were *passive*, so that they could not spike in response to the triggered spike. In a real cell all dendrites would (presumably) contain similar active conductances, so this model of a single isolated active branch is explicitly unrealistic, intended only to test the validity of eq. 3.52.

The results are shown as the simulated ΔV_{B_2} plotted against the predicted value, so that a point on the line of unit slope would represent a prefect prediction (Figure 3.8). The fact that most points lie near that line shows that the peak voltage induced in neighboring terminal branches is determined primarily by the peak current injected and the resistance of the shared trunk, rather than by other dendritic properties.

Those peak voltage increments (caused by spikes elsewhere) range from about $0 - 80$ mV for "strong HH," and $0 - 50$ mV for "weak HH" (Figure 3.8). Because a terminal branch receiving greater than $\Delta V_{B_2} \approx 25$ mV would exceed threshold and fire (depending on E_{rest}), many of those terminal branches would be triggered by the initial distal firing, increasing the trunk's ΔV_T and probably recruiting still more terminal branches in a chain reaction. This

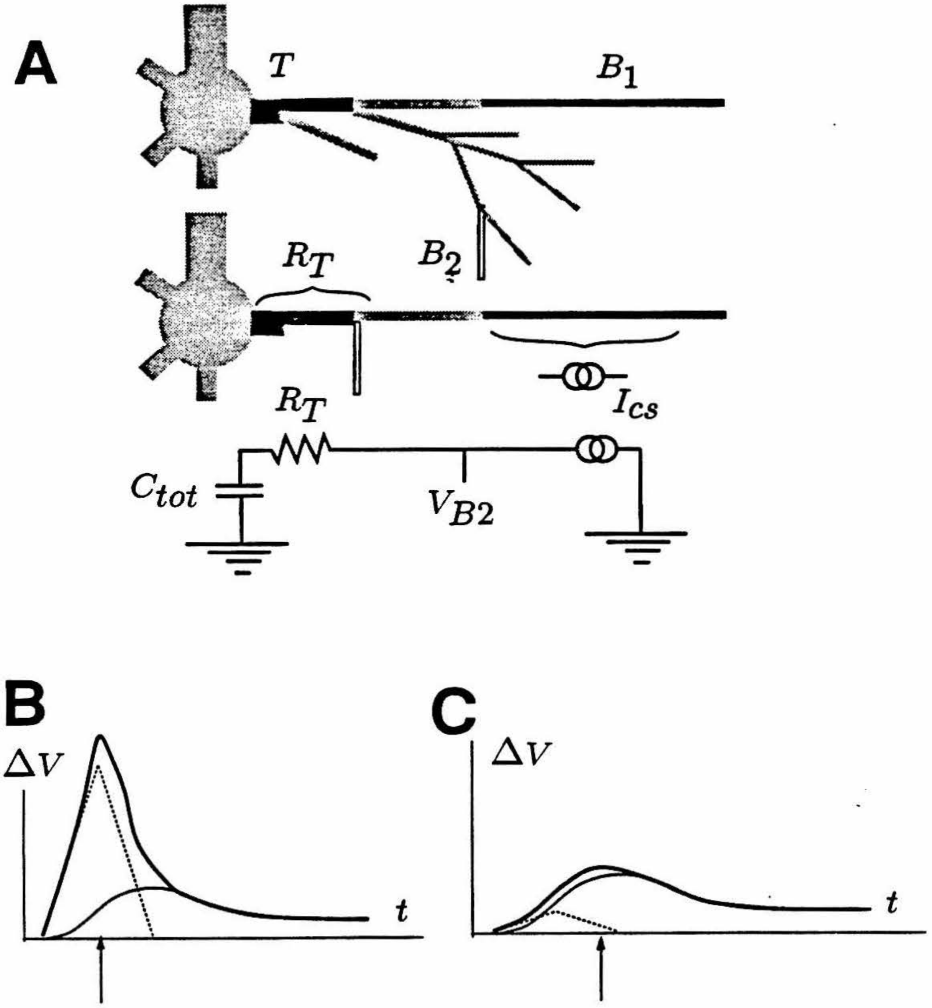

Figure 3.7: (A) When a terminal branch (B_1, black) spikes, it will cause a depolarization ΔV_{B2} in a neighboring branch (B_2, white). To calculate the magnitude of the depolarization at B_2 due to B_1, a simplified model assumes that B_2 is directly connected to the same dendritic trunk as B_1, that B_1 provides peak current in accordance with eq. 3.40, and that the two terminal branches share a common section of dendritic trunk T with resistance R_T. This model is represented by the circuit shown, where ΔV_{B2} is the sum of voltages produced across the cell's capactitance ΔV_{soma} and R_T (see section 5). (B) The contributions to ΔV_{B2} of somatic depolarization (thin curve) and voltage drop across R_T (dotted curve) depend on their relative magnitudes. When ΔV_{soma} is relatively small, the peak ΔV_{B2} (thick curve; arrow) only includes half the ΔV_{soma}. (C) When ΔV_{soma} dominates, nearly its full value contributes to ΔV_{B1} (arrow). A simple interpolation between these two cases uses a variable s ($0.5 \leq s \leq 1.0$; eq. 3.51) to determine the fractional contribution of ΔV_{soma} to ΔV_{B2}, and predicts ΔV_{B2} (eq. 3.52).

"recruitment" effect is undesirable if the spiking terminal branches are to remain decoupled, but may be useful for enhancing the depolarization caused by spikes in individual dendrites.

3.5.2 Recruitment of Neighboring Branches by Spikes

The present treatment does not examine how dendritic spikes propagate (action-potential style) with varying "safety factor" (Jaslove 1992), but whether a single dendritic spike will depolarize neighboring branches so that they also fire.

The preceding simulations suggested that when all the terminal branches contain active conductances, a strong dendritic spike in one branch would often trigger spikes in neighboring branches. Several such spikes firing in response to a single triggering event would obviously increase the event's peak somatic depolarization. Depending on the conductance strength and on the local branching geometry, that increased depolarization could either trigger still more branches, or it could saturate the local branches so that their combined somatic depolarization would be not much greater than that of an isolated branch. These competing effects make it difficult to predict the net somatic depolarization in response to a single triggering event, but simulation can resolve the question. This kind of "chain reaction" has been simulated for interactions among active dendritic *spines* by Segev and Rall (1988) and Rall and Segev (1988), whose results are qualitatively similar to the present simulated interactions among spineless branches.

In the present simulations, each terminal branch was triggered separately, and peak ΔV_{soma} plotted as a function of the independent-branch prediction (eq. 3.52). (This was identical to the method of section 3.4 and Figure 3.6, except that now all branches were active and $E_{rest} = -65\ mV$ was higher, encouraging neighboring branches to fire). The results (Figure 3.9) show that most somatic depolarizations were *above* the values predicted for isolated dendrites, indicating significant recruitment: for the "weak HH" case ($\overline{G}_{Na} = 0.033\ Scm^{-2}$) , the mean depolarization was shifted from 2.7 mV to 3.5 mV, and for the "strong HH" case ($\overline{G}_{Na} = 0.2\ Scm^{-2}$) from 6 mV to about 13 mV. The "strong" case had four dendrites whose individual firings were capable of triggering *all* the basal dendrites, and about 1/6 of the 44 dendrites in this case delivered $\Delta V_{soma} \geq 15\ mV$, which would be sufficient to bring the soma to firing threshold.

So there is a qualitative difference between the "strong" and "weak" cases: strong active dendritic conductances tend to make individual branches more likely to fire *together*, and weak conductances make them more likely to fire *independently*. In all cases, some branches are much more influential than others, and it is often (counterintuitively) the more distal ones which have the largest somatic influence, because the voltage drop across their shared R_{tk} can recruit many neighbors to fire.

In both strong and weak cases, many individual dendritic branches could trigger other branches (i.e., branches were *not* decoupled). But still most triggered shafts could not alone bring the soma to firing threshold, even when recruit-

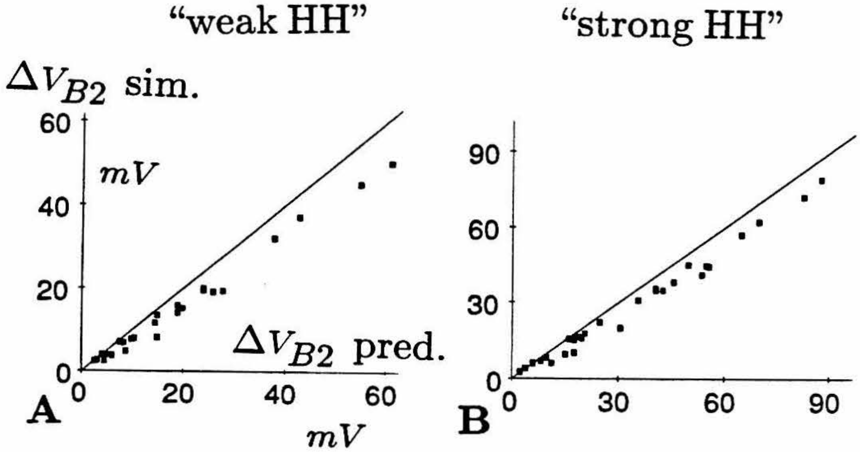

Figure 3.8: Each dendritic spike depolarizes each neighboring branch B_2 as predicted by the weighted sum of ΔV_{soma} and ΔV_T (eq. 3.52). When the most distal terminal branch on a common trunk was active and triggered, the peak depolarization ΔV_{B2} was measured at the proximal end of other terminal branches B_2 (relative to $E_{rest} = -75\ mV$). ΔV_{B2} plotted against the predicted depolarization (Figure 3.7 and section 3.5) has points lying near the diagonal line representing perfect predictions: (A) $\overline{G}_{Na} = 0.033\ S\ cm^{-2}$; (B) $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$. If the neighboring branches had also contained active conductances, most of the larger depolarizations observed (e.g., $\Delta V_{B2} > 25\ mV$) would have recruited those branches into firing in concert with the initially triggered branch.

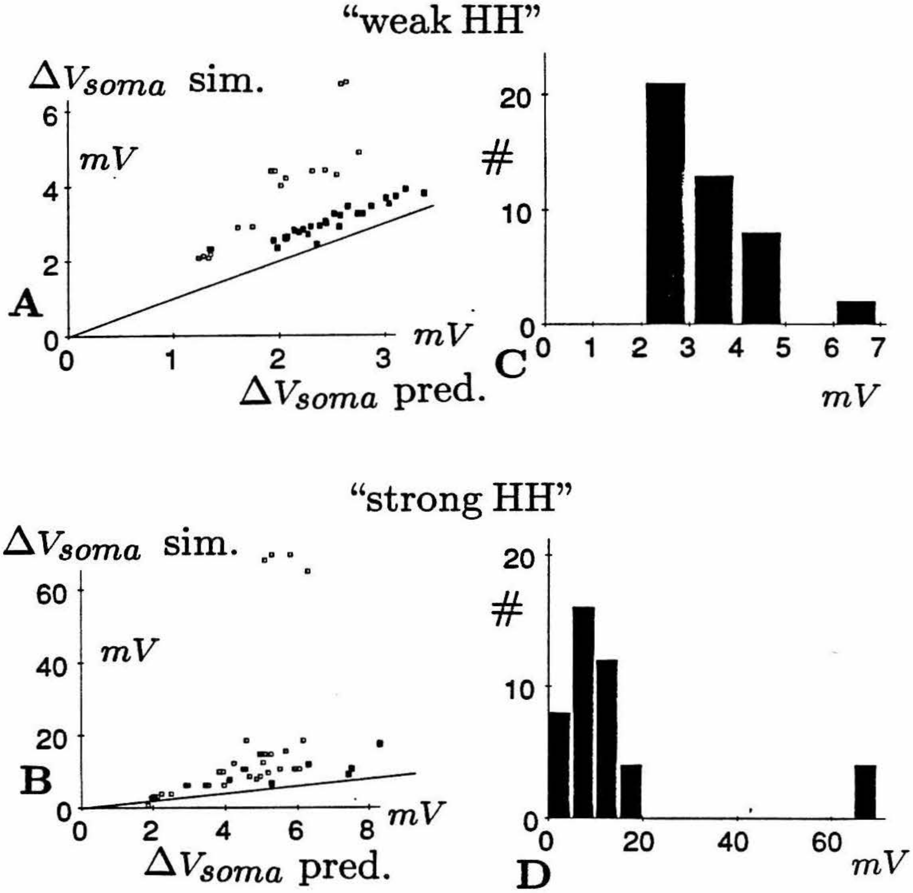

Figure 3.9: When all terminal branches contain active conductances, spikes in individual terminal branches recruit other branches to fire, and produce somatic depolarizations greater than predicted. All 44 terminal branches were given active Hodgkin-Huxley-like kinetics, but only one at a time was triggered. ΔV_{soma} resulting from those dendritic spikes (relative to $E_{rest} = -65\ mV$) is shown plotted against the values predicted by eq. 3.52 (points lying on the diagonal line represent a perfect predition): (A) $\overline{G}_{Na} = 0.033\ S\ cm^{-2}$; (B) $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$. Points well above that line represent the synchronous firing of one or more neighboring active dendrites which are triggered by the first one's firing. Histograms show that the simulated ΔV_{soma} cluster around $3-4\ mV$ when $\overline{G}_{Na}$ is small (C, $\overline{G}_{Na} = 0.033\ S\ cm^{-2}$), but that for larger $\overline{G}_{Na}$ (D, $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$) very large somatic depolarizations can sometimes result from a single synaptic event. Even in this case, most isolated events would not be sufficiently strong to bring the cell to a typical firing threshold of $\approx -50\ mV$.

ment was included. As a result, it would still take two or more events in concert to fire the soma, and the cell would act as a kind of "AND"-gate among dendritic spikes. (The next sections deal with the circumstances under which a cell could use such cooperative dendritic events as a submillisecond-resolution coincidence-detector).

If on the other hand the entire basal dendritic tree and soma (not just the terminal branches) contained strong active conductances, a single distal spike might drive the dendritic trunk above threshold. We can estimate that a typical firing *trunk*—of diameter $2 - 3.5\ \mu m$—would contribute a somatic depolarization large enought to fire the entire cell (recall that i_{ax} scales as $d^{3/2}$, eq. 3.36, leading to a peak depolarization 3-6 times greater than the typical $7\ mV$ observed from $0.8 - 1.0\ \mu m$ branches in Figure 3.6).

This situation was tested, with "strong HH" conductances simulated over the soma and entire basal tree, and $E_{rest} = -65\ mV$. Of dendritic spikes initiated in 44 of the terminal branches separately, all but two (i.e., about 95%) caused the soma to fire. Because single branches typically fired the entire cell, the cell acted as an "OR" gate among those dendritic spikes. The computation might be still more complex if dendritic spikes were generated by the coincidence of synaptic events (sections 3.3 and 3.8; the cell would compute the "OR" function of many separate "AND" functions of EPSP's, making it an "AND-OR" cell).

3.6 Somatic Repolarization by Dendritic Spiking

There are two important ways a dendritic spike can influence the soma—it can depolarize the soma (as discussed above), or its delayed-rectifier (I_{DR}) currents can *repolarize* the soma after the voltage peak has occurred, thus limiting temporal summation of sequential spikes. Both these effects strongly influence the cell's computational properties (see Discussion). Here we will qualitatively estimate the influence of I_{DR} at the soma, while saving the mathematical and simulation details for Appendix H.

A brief somatic depolarization will decay even without I_{DR}, as the charge equilibrates into the dendrites (see the upper trace of somatic voltage in Figure 3.11 A,C). What determines the additional role of I_{DR} in repolarizing the soma? Clearly, such variables as the peak conductances ($\overline{G}_K$ and $\overline{G}_{Na}$), diameter d, and resistivity R_i all can change absolute depolarizations. But if we examine the *relative* repolarization of a spike—the fraction of sodium charge removed by potassium— then we must investigate properties which differ between the two conductances, rather than properties they share in common with the dendrite.

One such difference is reversal potential. Because $E_K = -95\ mV$ is much closer to $E_{rest} =\approx -70\ mV$ than is $E_{Na} = +50\ mV$, a potassium conductance has a smaller driving voltage and will pass a far smaller current than will a sodium conductance of equal magnitude. A lower E_{rest} magnifies this effect, as the driving potential between E_K and E_{rest} diminishes. Potassium conductances' weaker current can be roughly canceled by their longer duration ($\tau(K) = 2.0\ ms$ vs. $\tau(h) = 0.5\ ms$), so the net integrated potassium and sodium charge to the soma may be very similar for equal conductances.

But when potassium conductances are much larger than sodium conductances, the peak potassium current is limited by the dendritic resistance between the active region and the soma... strong potassium conductances cannot generate correspondingly large currents. Thus, it takes very strong potassium conductances (ranging from two to four times the peak sodium value, depending on E_{rest}) to ensure that virtually no persistent somatic depolarization remains after a dendritic spike (Figure 3.10). For this reason, the ratio $\overline{G}_K = 2\overline{G}_{Na}$ was used elsewhere in the chapter, unless othewise specified.

3.6.1 Pulse Widths

A dendritic spike with little persistent somatic depolarization is very brief. In fact, we can expect that the entire duration t_p of the somatic depolarization—approximated as a triangle—will be only the sum of its rise time $\tau(h)$ and its fall time $\tau(K)$, so its full-width at half maximum $t_{1/2}$ would be about half that amount:

$$t_p \approx \tau(h) + \tau(K) \tag{3.53}$$

$$t_{1/2} \approx \frac{\tau(h) + \tau(K)}{2} \tag{3.54}$$

$$t_{1/2} \approx 1.2\ ms \tag{3.55}$$

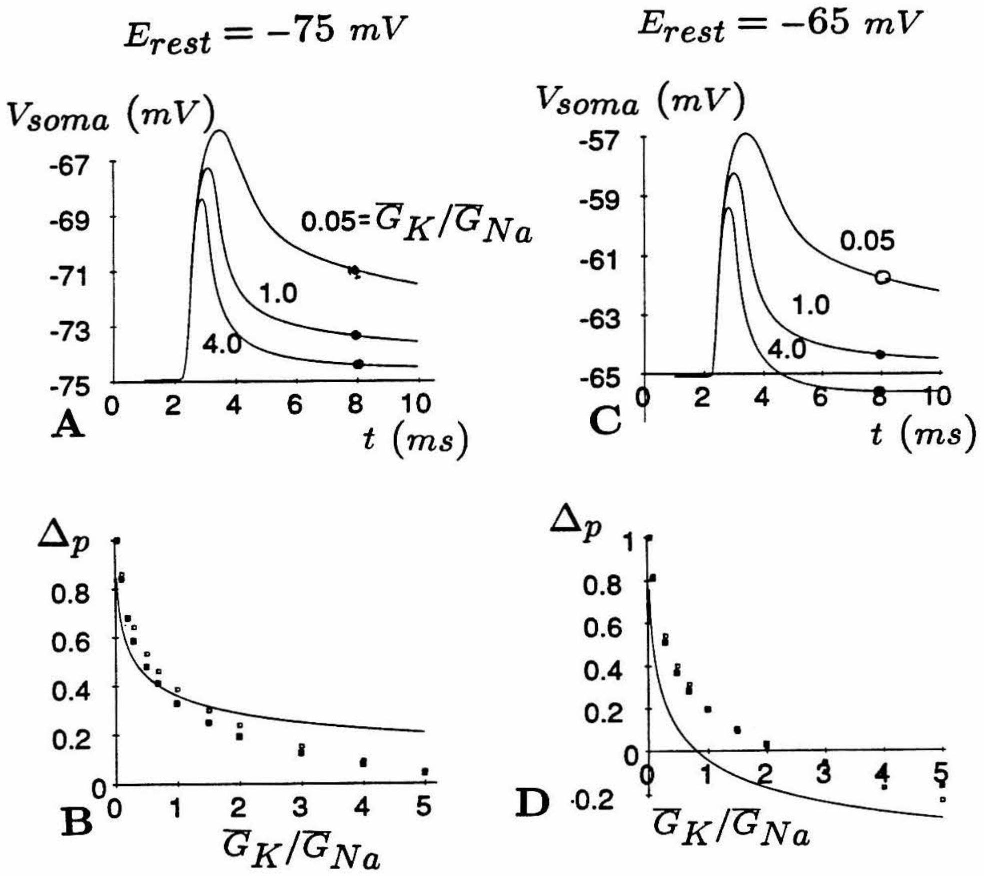

Figure 3.10: The amount of somatic depolarization (due to a dendritic spike) which persists afterwards depends on the strength of the spike's delayed-rectifier current I_{DR}. The time-course of somatic depolarization (simulated using strong active dendritic conductances $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$) is plotted for various ratios $\overline{G}_K/\overline{G}_{Na}$ (A, $E_{rest} = -75\ mV$; C, $E_{rest} = -65\ mV$). A dimensionless measure of the persistent depolarization, Δ_p, is the ratio of the somatic voltage 8 ms after triggering (filled circles) to its value in the near-absence of potassium currents (open circles); $\Delta_p = 1$ represents no attenuation due to potassium currents, and $\Delta_p < 0$ represents a spike which leaves the cell *more* polarized than before. Plots show Δ_p at various $\overline{G}_K/\overline{G}_{Na}$ for two resting potentials (B, $E_{rest} = -75\ mV$; C, $-65\ mV$) and peak sodium conductances ($\overline{G}_{Na} = 0.02\ S\ cm^{-2}$, open squares; $\overline{G}_{Na} = 0.2\ S\ cm^{-2}$, filled squares). The thin curves are predictions by the highly simplified model (eq. H.16 and Figure 10). Only at fairly strong $\overline{G}_K/\overline{G}_{Na} > 2$ and moderately high $E_{rest} = -65\ mV$ can dendritic I_{DR} currents remove persistent depolarization, a condition necessary for efficient coincidence-detection among dendritic spikes.

Measurement of $t_{1/2}$ for the dendritic spikes simulated above (section 3.5, all dendrites active) show this to be the case (Figure 3.11). Most simulated pulse widths were near 1.2 ms, and a few wider pulses ($1.3 - 1.6$ ms) occurred when one dendrite recruited others in firing, so that several sequential dendritic spikes appeared at the soma as one broader spike. This broadening was more prominent in the "strong HH" case, where stronger depolarizations caused more frequent recruitment.

3.7 Capacitance of Dendritic Spines

The primary influence of dendritic spines—which were not included in the above simulations—is to increase the membrane's effective capacitance and leak conductance by a factor of 2-3 (Jaslove 1992, Amitai *et al.* 1992). Leak conductance has proved unimportant at the fast timescales of this study, but a change in membrane capacitance can be critical.

The membrane capacitance for the entire simulated cell was thus doubled, to test the influence of spines' added capacitance and to verify that the approximations derived above can account for it. (This extra capacitance is equivalent to placing spines over the entire cell, including the soma; the simulation program could not assign separate C_m for dendrites and soma). In the case of an isolated EPSP inside a dendrite (section 4), a doubling of membrane capacitance should decrease $\lambda(t)$ by $1/\sqrt{2}$and increase $C(t)$ by $\sqrt{2}$, decreasing the peak voltage by the same amount. So the 19 mV EPSP simulated earlier (

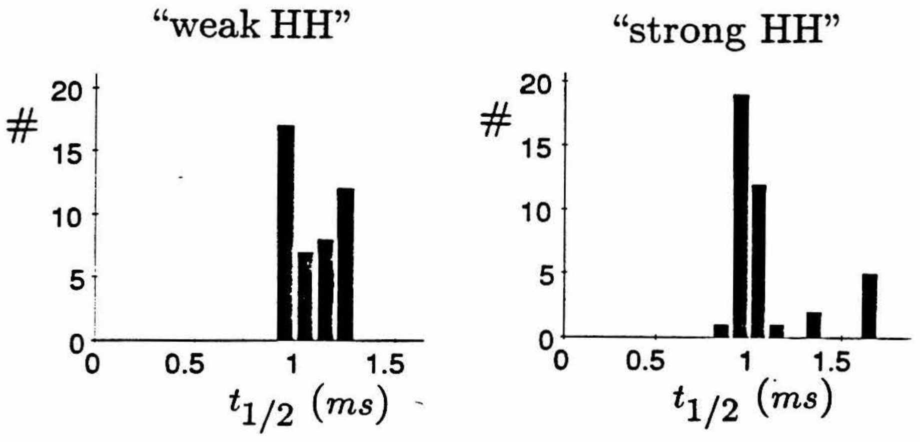

Figure 3.11: Somatic depolarizations from dendritic spikes with strong $\overline{G}_K$ have very brief durations. Histograms of somatic pulse-width $t_{1/2}$ (FWHM) for ΔV_{soma} in response to triggering synapses at each of the 44 terminal branches, with active conductances in all branches, $E_{rest} = -65\ mV$, and $\overline{G}_K = 2\overline{G}_{Na}$, as above (Figure 3.9). When sodium conductance was weak (A, $\overline{G}_{Na} = 0.033\ S\ cm^{-2}$), depolarizations were smaller and fewer neighboring dendrites were recruited to fire, so that pulse-widths clustered near the value 1.2 ms predicted in section 3.6. For larger sodium conductances ($\overline{G}_{Na} = 0.2\ S\ cm^{-2}$), larger depolarizations sometimes recruited several branches to fire sequentially, broadening the somatic pulse to values near 1.5 ms (the few cases in which *all* dendrites fired are not shown here). Depolarizations of such short duration could be used to perform sub-millisecond coincidence detection among dendritic spikes.

Figure 3.2; $t_{peak} = 0.1\ ms$) is predicted to peak at only 13.5 mV with the doubled membrane capacitance; the newly simulated value was 14.4 mV, a difference of 7% from the prediction.

A similar test was made for the somatic depolarization ΔV_{soma} due to a spike in that dendrite. The predicted effect on C_{tot} is similar, with the dendrites' effective capacitance increased by $\sqrt{2}$ (through $\lambda(t)$ and C_m) and the soma's capacitance doubled. The somatic peak earlier had been $\Delta V_{soma} = 8.3\ mV$, yielding a predicted 5.45 mV—which differed from the newly simulated peak, 5.6 mV, by only 3%. We can conclude that dendritic spines' capacitance will decrease all peak depolarizations from the values simulated in rough accordance with these predictions. These predictions do not cover another important effect: the blurring of the somatic EPSP from a fast dendritic synapse (Figure 3.2). Fast dendritic EPSPs ($T_{peak} = 0.1\ ms$) can be broadened by the spines' capacitance from a long somatic rise-time of 0.6 $msec$ (no spines, $C_m = 1\ \mu F/cm^2$) to an even longer 1.1 $msec$ (with spines, $C_m = 3 \mu m/cm^2$), which lies better within electrophysiological observations (Mason *et al.* 1991, Komatsu *et al.* 1988, Thompson *et al.* 1988; see section 3.9 for further details).

3.8 Quantifying Coincidence-Detection

A neuron's ability to "detect" coincidences requires both that it fire in response to coincident dendritic events and also that it **not** fire as frequently in response to non-coincident events. So we should investigate the cell's response to a large ensemble of inputs, in order not to ignore some detail of the interactions between synaptic events. Here we attempt to quantify a neuron's "effectiveness" at coincidence-detection by idealizing two extremes forms of pulse input: perfectly even, regular pulse trains vs. pulses which are bunched optimally to fire the neuron as fast as possible.

Consider a neuron receiving input events (either EPSPs or dendritic spikes) at fixed average rate f_{in}. If those events are *evenly* spread out in time and position (arriving at regular times on alternate dendrites, with the least chance of non-linear interaction), then the cell's output firing rate can be called f_e, representing the neuron's response to nearly "pure" temporal summation without coincidences.

But if the input events at the same rate are *optimally* arranged—usually in exactly coincident volleys barely sufficient to fire the cell, with no events between volleys—then the cell will fire at a higher rate, $f_{opt} \geq f_e$ (the "optimal" definition of f_{opt} guarantees that no more effective temporal arrangement of the same input pulses can exist). Both these arrangements of input events, the "even" and the completely synchronized, are highly unnatural in a biological context; they are meant here only to represent the two extremes over which a cell might be conceived to operate. From those rates we can produce a dimensionless number E_c, which represents the cell's "effectiveness" at distinguishing between coincidence-detection and temporal integration,

$$E_c \equiv 1 - \frac{f_e}{f_{opt}} \tag{3.56}$$

$$0 \leq E_c \leq 1, \tag{3.57}$$

where $E_c = 0$ represents a cell which cannot distinguish coincidences, and $E_c = 1$ represents a cell which performs only coincidence-detection, with no temporal integration at all.

3.8.1 Integrator-Models as Coincidence Detectors

As a benchmark, we can easily calculate E_c for three well-known integrator-models of cells (for references and descriptions of all three models, see Tuckwell 1989, or Jack *et al.* 1983, ch. 11). These integrator models assume that all input current pulses cause identical depolarization (unity) and have vanishingly brief duration, and that the cell fires and resets upon attaining a threshold depolarization of N_{th} pulses (e.g., N_{th} or more coincident input pulses will fire the cell). None of these models have spatial extent or explicit membrane conductances.

The simplest model, the "perfect integrator," is a leak-free capacitor which accumulates depolarizing pulses until it reaches threshold, at which time it instantly fires and resets. Because N_{th} inputs accumulated over a long time produce exactly the same depolarization that N_{th} coincident inputs would produce,

$$\begin{aligned} '\text{perfect}' : f_e &= f_{opt} = \frac{f_{in}}{N_{th}}, && (3.58) \\ \text{and } E_c &= 0 && (3.59) \end{aligned}$$

In contrast, the leaky integrator model loses depolarization at a rate

$$\frac{dV}{dt} = -\frac{V - E_{rest}}{\tau}, \qquad (3.60)$$

so that inputs pulses are "forgotten" with time. This model fires in response to evenly spaced inputs as

$$f_e = -\left(\tau \log(1 - \frac{N_{th}}{f_{in}\tau})\right)^{-1} \tag{3.61}$$

$$\text{for } \frac{N_{th}}{f_{in}\tau} < 1 \tag{3.62}$$

$$f_e = 0 \text{ otherwise} \tag{3.63}$$

(Stein 1967a). But its response to optimally-timed coincident inputs (in volleys of N_{th}) is the same as that of the perfect integrator model, because the cell fires before the leak term can repolarize it:

$$\text{'leaky'}: f_{opt} = \frac{f_{in}}{N_{th}} \tag{3.64}$$

$$\text{so } E_c = 1 + \frac{N_{th}}{f_{in}\tau \log(1 - \frac{N_{th}}{f_{in}\tau})} \tag{3.65}$$

$$\text{for } \frac{N_{th}}{f_{in}\tau} < 1 \tag{3.66}$$

$$E_c = 1 \text{ otherwise.} \tag{3.67}$$

This expression for E_c is near unity only when $\frac{N_{th}}{f_{in}\tau}$ is near to or greater than one, i.e., using low input rates or very strong leak terms (Figure 3.12).

The third model is the perfect integrator with refractory period, which cannot be triggered during a "dead time" t_0 after each firing. The refractory period does not change this model's response f_{opt} to coincident inputs which arrive outside the refractory period, i.e., as long as the input does not try to make the integrator spike more often than once every t_0,

$$f_{out} = \frac{f_{in}}{N_{th}} < t_0^{-1} \tag{3.68}$$

so $f_{opt} = f_{in}/N_{th}$ as before (eq. 3.58). But a certain portion of evenly spaced inputs *will* fall in the refractory period, so the output rate saturates for high

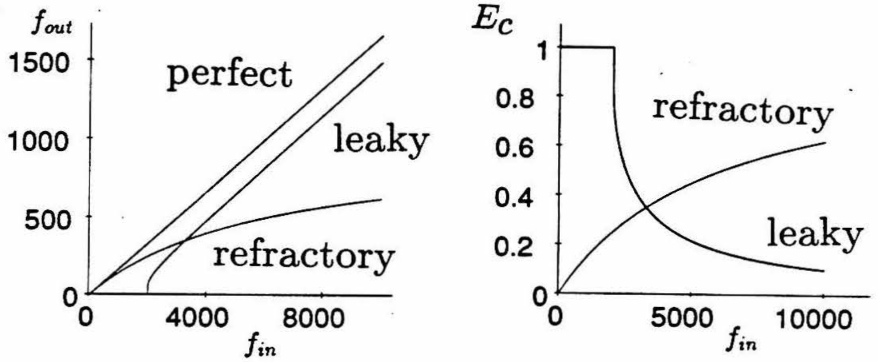

Figure 3.12: Input-output characteristics (f_{out} vs. f_{in}) of three integrator models for regular input, and their effectiveness at detecting coincident inputs pulses. (left) The perfect integrator with threshold delivers an output after every N_{th} input pulses (here $N_{th} = 6$). A leaky integrator with decay constant τ_m cannot fire at input rates below N_{th}/τ_m (here $\tau_m = 3\ ms$). A perfect integrator with absolute refractory period cannot fire again for time t_0 after firing, so f_{out} saturates at $1/t_0$ (here $t_0 = 1.0\ ms$). (right) For a fixed input rate f_{in}, both leaky and refractory models respond to coincident (or optimal) inputs as perfect integrators. Those stronger responses (relative to temporal integration of evenly timed inputs) gives an effectiveness measure E_c, where $E_c = 1$ represents a model which performs no temporal integration, and $E_c = 0$ represents a model which cannot distinguish inputs according to coincidence. Note that the leaky integrator only serves as a good coincidence-detector ($E_c \approx 1$) at low input rates.

input rates. Thus at high rates coincident inputs are more effective than evenly spaced ones, so E_c is increased:

$$\begin{aligned} (\text{“refractory''})\ f_e &= \left(\frac{N_{th}}{f_{in}} + t_0\right)^{-1} && (3.69) \\ E_c &= 1 - \left(1 + \frac{f_{in} t_0}{N_{th}}\right)^{-1}. && (3.70) \end{aligned}$$

In summary, the refractory period can detect high-resolution temporal information ($E_c > 0$) only at high input rates, while the leaky integrator does so only at low rates (Figure 3.12). These models in combination would be more realistic, but analyzing them is well beyond the benchmark-only role they serve in this chapter. In addition, a refractory period is a poor model for coincidence-detection inside cortical cells, because it operates most effectively when the cell fires *regularly*, which cortical cells do not do. (See Pratt, 1989 for computational applications of the refractory period for blocking axonal conduction.)

3.8.2 Pyramidal Cells as Coincidence Detectors

One potential method of coincidence-detection allows the coincidence of two or more submillisecond EPSPs in the same dendrite to cause a dendritic spike. In this case the dendritic branch thus acts as a kind of AND gate over synaptic events, with a dendritic spike as its output. If any one of those dendritic spikes can fire the entire cell, then the cell's output acts as an OR gate over dendritic spikes, and thus as an “AND-OR” gate over synapses. (The term “AND-OR” for this cell type is meant to characterize an approximate mode of operation;

by no means will the cell perform the exact logical function. (See Koch *et al.* (1992) for a survey of possible multiplicative interactions in neurons, and Shepherd *et al.* (1987, 1989) for a treatment of dendrites as AND-gates).

A second type of coincidence-detection might use a spike in a single terminal dendrite as the fundamental input to the soma. Such a branch can briefly influence the soma both by directly depolarizing the soma and by firing neighboring branches which help depolarize the soma. Dendritic spikes cause much stronger somatic depolarization than dendritic EPSPs alone (typically 2-15 mV vs. $100 - 200\ \mu V$), so it is possible that only a few dendritic spikes in coincidence can fire the soma. But they can only contribute to precise coincidence-detection if their persistent depolarization is immediately removed by a very strong delayed-rectifier conductance ($\overline{G}_K > \overline{G}_{Na}$), so that the cell does not temporally integrate the dendritic spikes. This mode of operation can be dubbed an "AND" cell, although the cell only approximates a logical "AND" function over dendritic spikes. Both "AND" and "AND-OR" cells have behaviors which vary, depending on peak conductance strengths and firing rates (as we will see below).

Defining or evaluating the coincidence-detecting effectiveness E_c for a realistic neural model is more difficult than for a simple integrator-model, since synaptic inputs on different dendrites have different amplitudes and interactions with one another. Thus, the "optimal" combination of coincident inputs to generate f_{opt} is far from obvious among the myriad possiblities. This chapter presents only a rough estimate of E_c, based on simulations with the same parameters as before (with the addition of Hodgkin-Huxley-like conductances at the soma, in order to create the somatic spikes which formed the output: $\overline{G}_{Na} = 0.2\ Scm^{-2}, \overline{G}_K = 0.12\ Scm^{-2}$, as used for the same cell in Bernander *et al.* 1991).

To generate the even, regular synaptic input producing f_e for both "AND" and "AND-OR" models, the 44 synaptic sites (one per distal basal dendrite) were fired in a particular listed order. In this case no synapse firing was immediately preceded or followed by another synapse sharing the same dendritic trunk (so that sequential events were electronically "far away" from each other on the dendritic tree). Only after about 5 firings would another site on the same trunk be fired, and only after all 43 other locations on the list had fired would the same synapse be fired again.

To estimate the maximum output rate f_{opt} for the "AND" cell, the same synapse order was used to generate coincident EPSPs, one EPSP per distal dendrite. These occurred in groups of M synchronous EPSPs, each EPSP typically causing a dendritic spike and all M dendritic spikes together causing a somatic spike. M should ideally just exceed the number of synchronous dendritic spikes necessary to fire the soma; in practice, the choice of M, like the choice of firing order, was estimated and led to large uncertainties in f_{opt}. M varied depending on firing rate, with $5 \leq M \leq 40$ for the "weak HH" "AND" cell and $7 \leq M \leq 10$ for the "strong HH" one. At the first firing time, the first M sites on the list were fired simultaneously; the next time fired the next M sites, and so on. Since M was chosen not to divide 44 evenly, a given synapse would participate in different groups on subsequent firings; firing times were chosen so that f_{opt} (in units of *sites/sec*) equalled f_e, as required by their definitions. This method was by no means the optimal way for coincident events to fire the cell, but it was easily understood and implemented. Each dendritic "firing" resulted from a single strong EPSP ($t_{peak} = 0.1\ msec, g_{peak} = 18\ nS$) at the terminal branch's center.

For the "AND-OR" cell, evenly-timed synaptic events occurred in the listed order (as above, with weaker synapses $g_{peak} = 6\ nS$, which could not alone fire a dendrite). The cell performed temporal integration as those events slowly depolarized the soma. Higher output rates f_{opt} resulted from *pairs* of such co-localized events (i.e., $g_{peak} = 12\ nS$ total in one dendrite), which together could typically fire the dendrite. (The ability of each single dendrite to fire the entire cell depended critically on the presence of very strong dendritic $\overline{G}_{Na} = 0.5\ nScm^{-2}$ and weaker $\overline{G}_K = 0.25\ nScm^{-2}$ throughout the dendritic tree).

Coincident inputs fired the AND-OR cell typically about three times as fast as evenly spaced inputs did (Figure 3.13A-C). This effectiveness $E_c \approx 0.5 - 0.7$ is not "perfect," but does reflect a definite preference for coincidences. This preference results entirely from the capacitive properties of the cell, which make dendritic EPSP's much stronger than somatic ones. Unlike the AND cell, this model did not use I_{DR} to prevent temporal integration, so that every EPSP in this model gave a persistent somatic depolarization and thus contributed slightly to temporal integration.

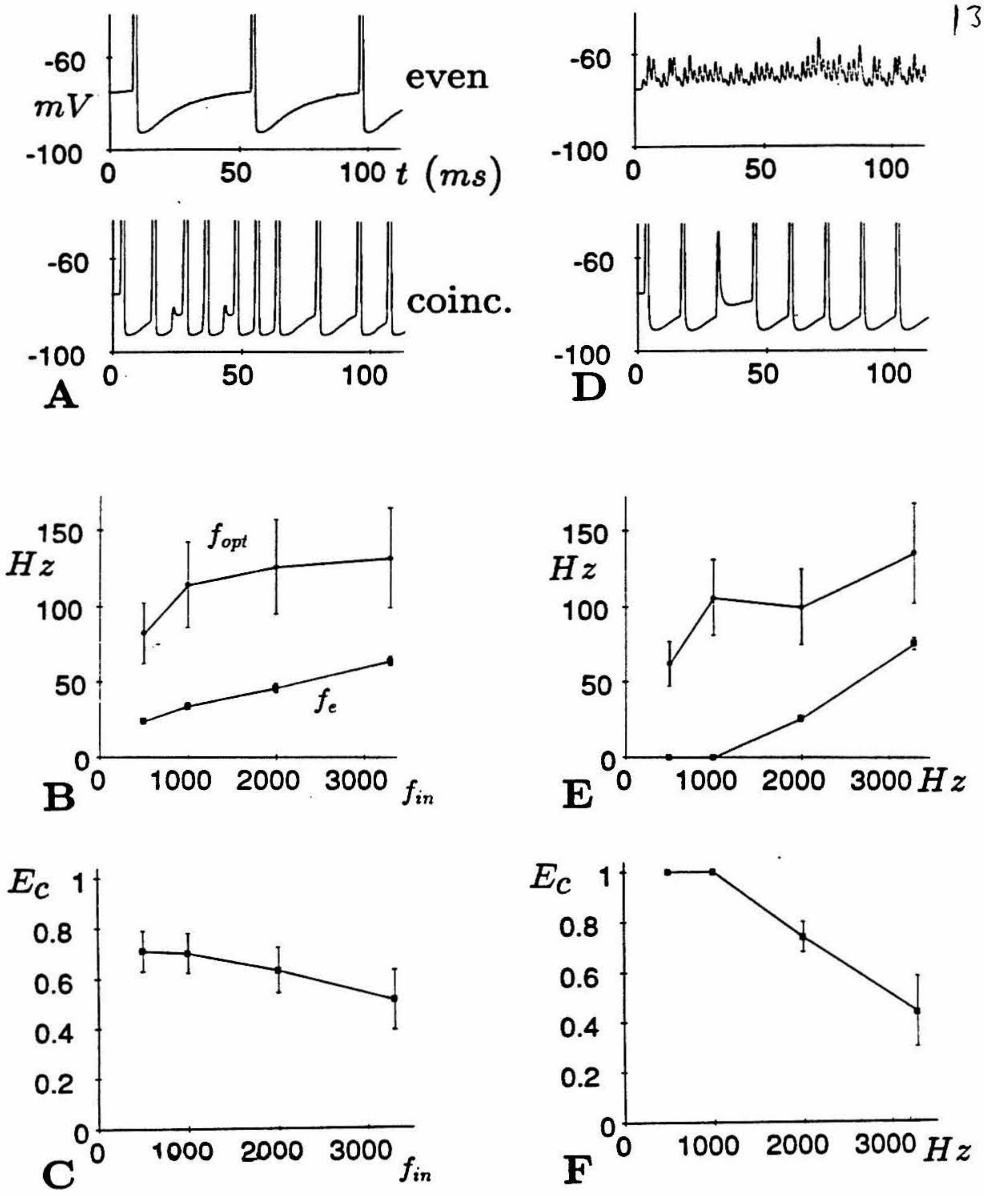

Figure 3.13: A simulated, reconstructed pyramidal cell can serve as a coincidence-detector when active conductances are present on its dendrites, as shown by these graphs of *somatic* voltage. If the entire basal tree contains strong Hodgkin-Huxley-like conductances (left column; $\overline{G}_{Na} = 0.5\ Scm^{-2}$), the cell will fire slowly (at rate f_e) due to evenly timed EPSPs (which do not initiate dendritic spikes), but fires faster (at f_{opt}) when EPSP's at the same rate occur in optimally coincident pairs inside the same dendrite and thus fire dendritic spikes (A,B). This preference for coincident inputs is quantified by values of E_c above zero (C). This cell requires simultaneous EPSP's on the same branch (an AND function), but fires the soma whenever any branch fires (OR function), so it can be dubbed an "AND-OR" cell. If instead the only dendritic conductances occur on terminal branches, but not on dendritic trunks (right column; $\overline{G}_{Na} = 0.2\ Scm^{-2}$; "strong HH"), evenly-timed spiking branches cannot usually fire the soma; only coincident dendritic spikes ($M \approx 8$ at once) can fire the cell reliably (D,E). This cell performs very little temporal integration of dendritic spikes, so that its effectiveness E_c as a coincidence-detector is high (near or equal to unity); it approximates an "AND" operation of M dendritic spikes.

The AND cell with "strong HH" dendritic conductances ($\overline{G}_{Na} = 0.2\ Scm^{-2}$) did not perform temporal integration per se, since subthreshold dendritic spikes were quickly repolarized. But a few branches were capable of firing the soma alone (as in section 3.5.2), so that the cell's response to evenly timed dendritic spikes was low but not always zero. As a result, this cell only performed as a perfect coincidence-detector ($E_c = 1$) at low firing rates, and dropped to $E_c \approx 0.4$ when firing above 100 Hz (Figure 3.13 D-F).

The AND cell with "weak HH" conductances ($\overline{G}_{Na} = 0.033\ Scm^{-2}$) acted as a perfect coincidence-detector, since it never fired under even the fastest ($10,000\ Hz$) regular stimulation by dendritic spikes (Figure 3.14 A-C; for comparison, a leaky integrator with $N_{th} = M = 10$ would need a very fast $\tau \leq 1\ ms$ to respond similarly, eqs. 3.64-3.67. As evidence that this ability resulted entirely from strong I_{DR} currents, simulations suppressing those currents (i.e., $\overline{G}_K/\overline{G}_{Na} = 0.05$) produced a dramatic increase in response to regular dendritic spikes and a resulting reduction in E_c to near zero (Figure 3.14 D-F). The difference between complete repolarization of spikes and temporal integration of them is strikingly evident in the top traces of Figure 3.14 A and D.

Under coincident stimulation, both AND cells hyperpolarized to extremely low potentials ($\approx -90\ mV$) as a result of strong I_{DR} currents and no counterbalancing depolarizing currents. As a consequence, many coincident events were required to fire these models from that low potential (e.g., $M > 5$, larger than the 2 or 3 events needed to fire from $E_{rest} = -75\ mV$). Under even stimulation, V_{soma} fluctuated near $-65\ mV$, which is the potential at which each

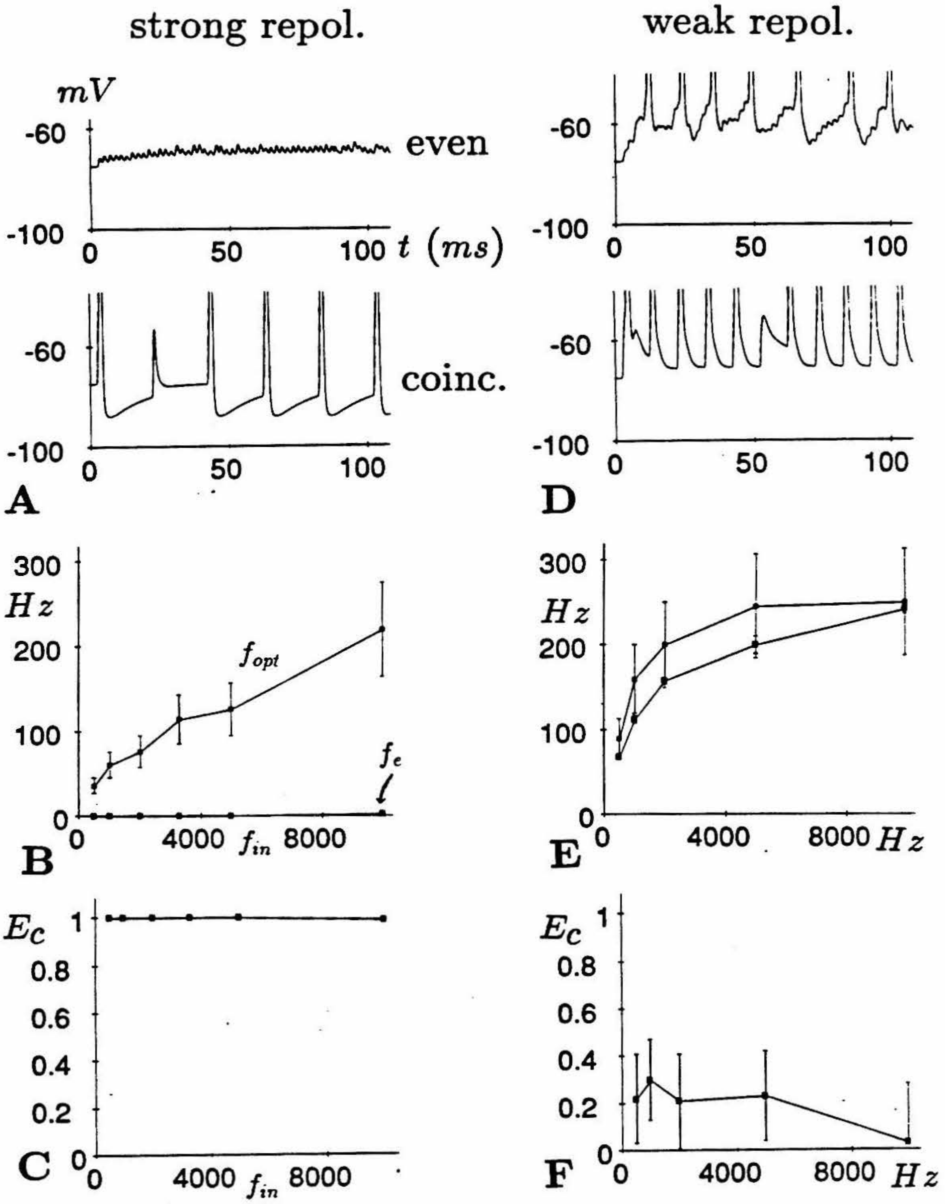

Figure 3.14: When spiking conductances on terminal dendrites are strong enough to fire (self-repolarizing) dendritic spikes but still weak enough that a single dendritic spike does not fire the soma (e.g., "weak HH", $\overline{G}_{Na} = 0.033\ Scm^{-2}$), then the cell can act as a perfect detector of coincident dendritic spikes. With strong delayed-rectifier dendritic conductances (left column; $\overline{G}_K = 0.066\ Scm^{-2}$), these graphs of somatic voltage show that individual dendritic spikes quickly repolarize the soma to about $-65\ mV$, so that no temporal integration of them occurs; the cell will not fire in response to evenly timed dendritic spikes, but only when they are coincident (A,B). Such a cell would be a perfect coincidence-detector ($E_c = 1$) at all realistic firing rates (C). This coincidence-detection ability disappears when dendritic delayed-rectifier currents are sharply reduced (right column; $\overline{G}_K = 0.0016\ Scm^{-2}$). While the cell still fires well in response to coincident dendritic spikes, it also can perform temporal integration of them (D,E; note the ramp-like voltages in the upper trace of D). The low values of E_c (F) show that this cell cannot distinguish between coincident and evenly timed dendritic spikes.

dendritic spike would be exactly repolarized for the K-conductance simulated ($\overline{G}_K/\overline{G}_{Na} = 2$; see section 3.6). This potential serves as a kind of reversal potential for the combined sodium and potassium currents which dominate the somatic voltage under regular dendritic spiking.

3.9 Discussion

Can a cortical cell perform millisecond computations, or only much slower temporal averaging? Many researchers believe that single-cell cortical computation must be inherently slow, because of the enormous attenuation of high-frequency signals in thin dendrites and because of the high spiking variability (near-Poisson "noise") associated with single-unit cortical firing (Douglas and Martin 1991). But both of these influences can instead be interpreted as *facilitating* single-spike computation.

3.9.1 Requirements for Submillisecond Computation

While the dominance of truly random "noise" in spike times would certainly preclude reliable computation at the single-spike level, the very existence of highly variable (and possibly non-random) spike trains suggests that the single neurons generating them are *not* performing much temporal integration (smoothing) of multiple EPSPs. In fact, such strong interspike-interval variability is **not** consistent with current models of temporal integration of EPSP's, and might be effective instead at fast information transmission, as described in Chapter 2. So the argument that high variability precludes high-frequency computation can be turned on its head: the high variability might indicate that the cell *does* use and transmit high-frequency information. While there is scattered indirect evidence for millisecond precision in spike patterns (Abeles 1990; Strehler and Lestienne 1986; Frostig 1985; Legendy and Salcman 1985), we so far lack evidence that such precision plays a perceptual role outside the auditory system, where submillisecond computations are routine in echolocation and binaural auditory localization. (An exception in vision is found in Burr's investigation (1979) of vernier acuity, where human observers perceived a vernier offset in a moving bar when its two halves were flashed as little as 1.5 ms apart).

The other argument against high-frequency computation—that dendritic capacitance would filter out the high frequencies—can be likewise inverted, since that very property allows different parts of the cell to be well-isolated at high frequencies, and thus to carry out nearly independent fast computations. This ability is enhanced in a cell with long, thin distal dendrites. But this proposal requires some method of delivering the results of the brief, isolated computations to the soma; passive dendrites cannot do it.

But active dendrites can. And the most effective active conductances for maintaining isolation among dendrites would be fast (e.g., sodium rather than calcium) and totally repolarized (e.g., strong $\overline{G}_K$), so that no persistent depolarization could couple the dendrites or allow temporal integration. Jaslove (1992) argues on theoretical grounds that such conductances are necessary to overcome the very short length constants of thin dendrites. Experimental evidence relating to such conductances on terminal branches is discussed in the next section.

While this chapter explores some possible functions of fast EPSPs and distal dendritic Hodgkin-Huxley conductances, it does not treat the many other conductances—voltage-dependent and otherwise—thought to operate in distal dendrites of cortical and hippocampal pyramidal neurons (e.g., Ca^{++}, Westenbroek *et al.* 1990; NMDA-mediated, Regehr and Tank 1991). This deliberate omission has two motives.

The first is that the fast sodium Hodgkin-Huxley conductances would make the best conceptual test of fast, isolated computations inside a dendritic tree. The better-known calcium conductances typically last about ten to twenty milliseconds, a timescale at which the capacitive decoupling discussed here would scarcely exist. The other motivation is simplicity—it is safer and easier first to understand and approximate the interactions between Hodgkin-Huxley-like conductances and the dendritic tree, before including other influences.

Many aspects of this subject have been addressed before. Mel (1992a) has discussed how voltage-dependent dendritic conductances (NMDA-receptor channels) can store information. Jaslove (1992) investigated coincidence-detection in similar neuron models, and noted that strong $\overline{G}_K$ can limit the propagation of dendritic spikes to the soma. But he did not discuss brief somatic depolarizations caused by dendritic spikes which do *not* propagate to the soma. Koch and Poggio (1987) reviewed several kinds of multiplicative and binary computations in dendritic cables and spines.

Shepherd *et al.* (1987, 1989) have also directly addressed the question of coincidence-detection of EPSPs in dendrites, but these studies (and Jaslove's, 1992) differ from this one in at least one important way: they used a very strong passive membrane leak conductance (Shepherd, $\tau_m = 2-4\ ms$; Jaslove, $\tau_m \approx 2\ ms$), which so effectively isolated EPSPs in time (through τ_m) and in space (through λ_{DC}) that the ability of *non*-leaky capacitive dendritic cables to isolate fast events was not emphasized. Those researchers instead discussed coincidence time-windows much longer than the submillisecond range. They also did not model their cells' net response to repeated synaptic events over the entire dendritic tree, so that the tradeoff between coincidence-detection and temporal integration was not explicit. Segev and Rall (1988) have investigated the properties of hypothetical active spines (rather than whole dendrites), pointing out that excitatory and inhibitory synapses might trigger spine-spikes with submillisecond resolution (using different mechanisms than the submillisecond excitatory-excitatory coincidences explored here).

The two mechanisms proposed here for fast membrane repolarization (which is necessary for coincidence-detection) are more energy-efficient than fast passive membrane time-constants. A strong *passive* membrane conductance draws a relatively large, constant amount of metabolic power in order to maintain the cell's resting depolarization. But the fast component of repolarization through capacitive charge-equilization requires no membane conductance at all, and repolarization by delayed-rectifier conducances only demands a large membrane conductance (and the resulting power drain) just *after* a dendritic spike, but not while the cell is unexcited.

3.9.2 Plausibility and Testability of Critical Assumptions

Submillisecond Synaptic Currents

Part of this chapter postulates that submillisecond EPSCs exist inside cortical dendrites (sections 3.3 and 3.8.2). But the existence of such EPSCs is not supported by direct recordings of EPSCs from a hippocampal dendrites. Electrical shock-induced EPSCs measured close to the soma inside cells in the hippocampal CA1 region (Hestrin *et al.* 1990; Sah *et al.* 1990) have a fast (non-NMDA) component which lasts 3-15 ms, and a slower voltage-gated NMDA component lasting 50-100 ms. Even the fast component in those recordings is far too slow to produce the submillisecond charge-equilization modelled here (section 3.3).

More importantly, measurements of EPSPs (*potentials*, not currents) in cortex **are** consistent with fast dendritic EPSCs. Spike-triggered averaging reveals subthreshold monosynaptic potentials with amplitudes in the $50 - 400\ \mu V$ range and rise-times from $0.5 - 1.0\ ms$ (10-90% rise-time in rat visual cortex *in vivo*; Mason *et al.* 1991) to $0.8 - 2.4\ ms$ (0-100% rise-time in cat visual cortex slice, Komatsu *et al.* 1988; 10-90% rise-time in hyperpolarized rat cingulate and sensorimotor cortex cells, Thompson *et al.* 1988). Although these researchers did not directly measure synaptic currents, we can assume (to first order) that the dendritic EPSC duration cannot exceed the somatic EPSP rise-time, because the potential integrates a current which has been additionally smoothed by the intervening dendritic capacitance. In that case, the submillisecond EPSP rise-times reported by those researchers are probably due to synaptic currents lasting less than a millisecond (with corresponding $t_{peak} \leq 0.3\ ms$). But the activation of NMDA-mediated conductances at moderate depolarizations could prolong the local EPSPs to tens of milliseconds (Thompson *et al.*, 1988), thereby reducing the temporal precision of single synaptic events.

All of the distal dendritic EPSCs simulated here—even the fastest ones with $t_{peak} = 0.05\ ms$—produced far longer somatic rise-times (about a millisecond), due to smoothing by realistic dendritic capacitances. Thus the simulated somatic rise-times are in the same range as the experimental recordings (see section 3.7). Direct measurements of EPSC duration in cortex at low membrane potentials ($< 75\ mV$), combined with realistic cell models, might better determine whether very fast dendritic EPSCs exist.

One provacative possibility is that the measured somatic EPSPs result from active dendritic spines. The fast time course is about the same for a sodium spiking current as for the synaptic currents simulated here. And the currents flowing though a synapse with $\overline{g}_{peak} \approx 5\ nS$ are comparable to those which would flow though an excited spine stem with diamater $0.1\ \mu m$ and length $1.0\ \mu m$ (and hence resistance $200\ M\Omega$, or $5\ nS$; see Segev and Rall 1988 for further details).

Active Dendritic Conductances

The two types of coincidence-detection proposed here—detection of coincident EPSPs inside dendrites, and of coincident dendritic spikes at the soma—both require the existence of strong Hodgkin-Huxley-like conductances inside thin terminal dendrites. Such conductances have not yet been unambiguously observed.

There is some evidence for active conductances in the *apical* dendrites of both hippocampal (Jaffe *et al.* 1992) and cortical neurons (Amitai *et al.* 1992; Huguenard *et al.* 1988; see Adams (1992) for a review). But such conductances may not be sufficiently strong to sustain spikes, or may consist of calcium conductances (Amitai *et al.* 1990), which are too slow for submillisecond coincidence detection. Because the apical dendrite is morphologically very different from the basal terminal branches, it may serve a different function and its conductances may not reflect conductances on narrow distal branches. But there are no direct recordings from distal basal dendrites of cortical cells, because those dendrites are far too thin to be impaled with a recording electrode.

If fast dendritic spikes do exist, they would be visible at the soma as potentials which repolarize much faster than the (presumed) membrane time constant. There are isolated reports of such "small spikes" in visual cortical cells which meet these criteria: they have amplitudes of $8 - 15\ mV$ (greater than single EPSP's, but weaker than full-fledged action potentials), submillisecond duration (full-width at half-maximum), and repolarization to E_{rest} within a millisecond (Ferster 1986, 1987, 1988). These small spikes result from electrical shocks to the LGN (i.e., due to many coincident synaptic events, rather than to the monosynaptic events used in spike-triggered averaging). But it is neither clear whether the spikes' high amplitude comes from sodium conductances nor whether their fast repolarization results from active dendritic potassium conductances (or instead from strong, synchronized inhibitory conductances). If other researchers reproduce such spikes, one might determine the source of the fast repolarization by blocking the action of inhibitory GABA neurotransmitters.

Such dendritic spikes might also be inferred from intracellular potentials in activated cells *in vivo*. As reported in simulations by B. Mel (1992b) and shown here, dendritic sodium spikes might only produce somatic potentials of a few millivolts, so that individual events could be nearly "invisible" as they blended together into "noise." An autocorrelation analysis of the subthreshold potential between spikes might show whether the observed fluctuations are mostly due to self-repolarizing events (e.g., dendritic spikes) or to other less structured currents.

But dendritic spikes might be sufficiently strong that a single one could fire the cell (e.g., in the "AND-OR" model). Under sensory stimulation, such a dendritic spike might be inferred from a somatic spike which rises from a voltage below the threshold found using current-injection. (A dendritic spike could reach its local "firing threshold" while the soma was still well below threshold, so that the somatic spike would appear to arise "out of nowhere" from below threshold voltage. For an example, see the difference in apparent thresholds between the two traces of Figure 3.13 A).

Distal dendritic conductances would not necessarily reveal themselves during intracellular stimulation, because a cell containing them could still behave as a temporal integrator of injected DC current (as occurs for the near-DC synaptic current in Figure 3.13 A).

Coincidence Detection in Cortical Cells

There are other types of neurons—usually auditory neurons—which operate even faster than the submillisecond regime postulated here. For instance, individual spikes in the auditory system of barn owl can phase-lock to tones of frequency up to 5-9 kHz (Sullivan 1985), although the cells responsible are morphologically very different from cortical pyramidal cells. Spikes from cells in the monaural nucleus of the echolocating big brown bat can lock to stimuli with a precision of 30 μs (Covey *et al.* 1991). And single-EPSP coincidence-detection without temporal integration can take place in the cochlear nucleus of mice (Oertel *et al.* 1989), although in this case the excitation comes from very large somatic synapses (rather than from dendritic spikes) and the repolarization comes from a fast membrane leak and inhibitory currents (rather than from active potassium conductances).

Regardless of mechanism, a role for cortical cells as coincidence-detectors is consistent with published cross-correlation data (Appendix D; the following qualitative argument is made more rigorous in Appendix F). Toyama *et al.* (1981), studying pairs of nearby cortical cells of similar response type in cat visual cortex, report that about 60% of cross-correlation histograms (CCH's) between such cells have prominent, narrow peaks at zero. Those peaks typically contain over 30% of the cells' spikes, meaning that about 30% of a pair's spikes are coincident (within 0.1 *ms* in many cases; see their figures 1,3, and 4). Other reports show wider and smaller cross-correlation peaks between cells in the same orientation column (Michalski *et al.* 1983), or as far apart as several *mm* (T'so *et al.* 1986). Nelson *et al.* (1992) find about 10% of cell pairs across cat visual areas 17 and 18 have centered CCH peaks a few milliseconds wide ("towers").

Such narrow, monosynaptic CCH peaks are traditionally interpreted as resulting from a few cells providing direct "common input" to the two cells: a substantial fraction of shared presynaptic neurons would cause a similar fraction of shared output spikes. But the narrow, centered CCH peaks are also consistent with coincidence-detection: a collection of presynaptic neurons common to the two recorded cells, and firing in occasional coincidence, might trigger a disproportionately large proportion of the output spikes recorded from two coincidence-detecting neurons. While both explanations are consistent with the data, the coincidence-detector model makes fewer demands on the network than the "common-input" model, because coincidence-detectors preferentially amplify coincident signals, and therefore require less magnitude and synchrony in their input EPSPs to generate the same magnitude of CCH peak.

The coincidence-detection model also accounts more easily for the near-absence of off-zero CCH peaks relative to centered ones, a fact which would seem to indicate that while many cell pairs share common input, the *source* of common input is somehow seldom recorded. (To solve this paradox, Nelson *et al.* (1992) postulate a tiny subset of "driving" cells which induce synchrony by means of exceptionally strong synaptic excitation to a whole cell population.) But if cortical cells are like those proposed here, responding primarily to coincidences from multiple independent sources (each source driving both recorded cells), then neither source cell would correlate well with a recorded cell, leading to the observed deficit of off-center peaks, while the recorded cells would still correlate strongly with each other. See Appendix F ("Paradoxical Cross-Correlations") for a more rigorous and complete argument.

3.9.3 Analytical and Simulation Results

The simulated "toy model" used here—a realistic pyramidal cell morphology endowed only with simplified active conductances— demonstrates two principles: 1) that many effects of dendritic spiking can be approximated by simple analytical expressions, and 2) that dendritic spikes can in principle approximate logic operations (e.g., "AND," "AND-OR") on submillisecond EPSPs with little temporal integration, even when the membrane time constant is much longer ($\approx 30\ ms$). The crucial element for coincidence-detection is the quick repolarization of the membrane after a depolarizing pulse appears. One mechanism for this repolarization is the passive voltage decay as charge from a submillisecond EPSC equilibrates across a dendrite's distributed capacitance and into the soma. Another mechanism is the active removal from the soma of charge by a dendritic spike's I_{DR}.

A quantitative measure of a model cell's effectiveness at coincidence-discrimination (E_c) compares the cell's stronger firing due to optimally coincident EPSPs to its weaker response to regularly distributed EPSPs (section 3.8). This artificial measure shows that a realistic pyramidal-cell morphology with active dendrites may discriminate fine temporal coincidences, and that the best model (the "weak HH AND" model) is a perfect coincidence-detector at even its highest firing rates (Figure 3.14).

3.9.4 Conclusion

The marriage of known cortical pyramidal cell morphology to postulated Hodgkin-Huxley-like conductances on distal basal dendrites yields a simplified model cell which can in principle discriminate EPSP arrivals at the submillisecond level. Such fast computation might be complementary to the slower and better-known coding by average spike rate, because many differently patterned pulse-trains can share a common average rate, and can thus carry independent information at both fast and slow timescales. The much higher bandwidth (kHz vs. Hz) of this hypothetical single-spike computation might prove a useful alternative to the lower-frequency oscillations proposed to solve some cognitive tasks, such as feature segmentation and the "binding problem" (Engel *et al.* 1992) and visual awareness (Crick and Koch 1990). But it remains to be seen experimentally whether single cortical cells contain the fast synaptic conductances and active distal dendrites necessary for submillisecond coincidence-detection, and whether those cells actually do perform parallel nonlinear computations with kilohertz bandwidth.

Chapter 4

An Application: Point-of-Origin Binding

> "Science is wonderful. Nowhere else can you get such a wholesale return in speculation for such a piddling investment of fact." — *attributed to Mark Twain*

4.1 Introduction

A simple neural network can only "see" one thing at a time. When a network sees just one object, many of the network's analog neurons are active at once (representing different aspects of the object...shape, texture, color, motion, etc.), and those neurons all refer to the same single object. But if multiple objects are presented, then *each feature aspect* has several active neurons, and there is no simple way to determine which active neuron refers to which object. The difficulty in associating individual neural responses with particular objects is called the "binding problem" (Figure 4.1). It is an outstanding problem not just in neural theory but in pattern recognition in general; detailed reviews of its implications are given in Engel *et al.* (1992) and Koch (1993).

There are in general two ways to solve this problem. One is give each analog neuron a very small receptive field, so that it only responds over a region so small that one would only expect to find one object at a time there. This approach requires tiling the whole visual field with identical but displaced receptive fields (thus hugely increasing the number of neurons present), with each neuron responding a much smaller fraction of the time. This is not the way the brain solves the problem; we do not yet know how it does so.

We *do* know that the brain increases visual neurons' receptive field sizes as their optimal stimuli become more specific. In a rough sense, this is like trying to make all neurons equally likely to be active, so that frequently encountered stimulus types cover small portions of the input space (small receptive fields), and rare ones cover more of the world (large receptive fields). For example, neurons in different cortical areas can respond to different features of the same object (e.g., form in V4 and info-temporal (IT), and motion in middle temporal (MT)); receptive field sizes in these areas are much larger than in V1, which responds to more common oriented contours.

In this case we still have the problem of how to label or bind together the responses of individual neurons. Because traditionally a neuron's analog re-

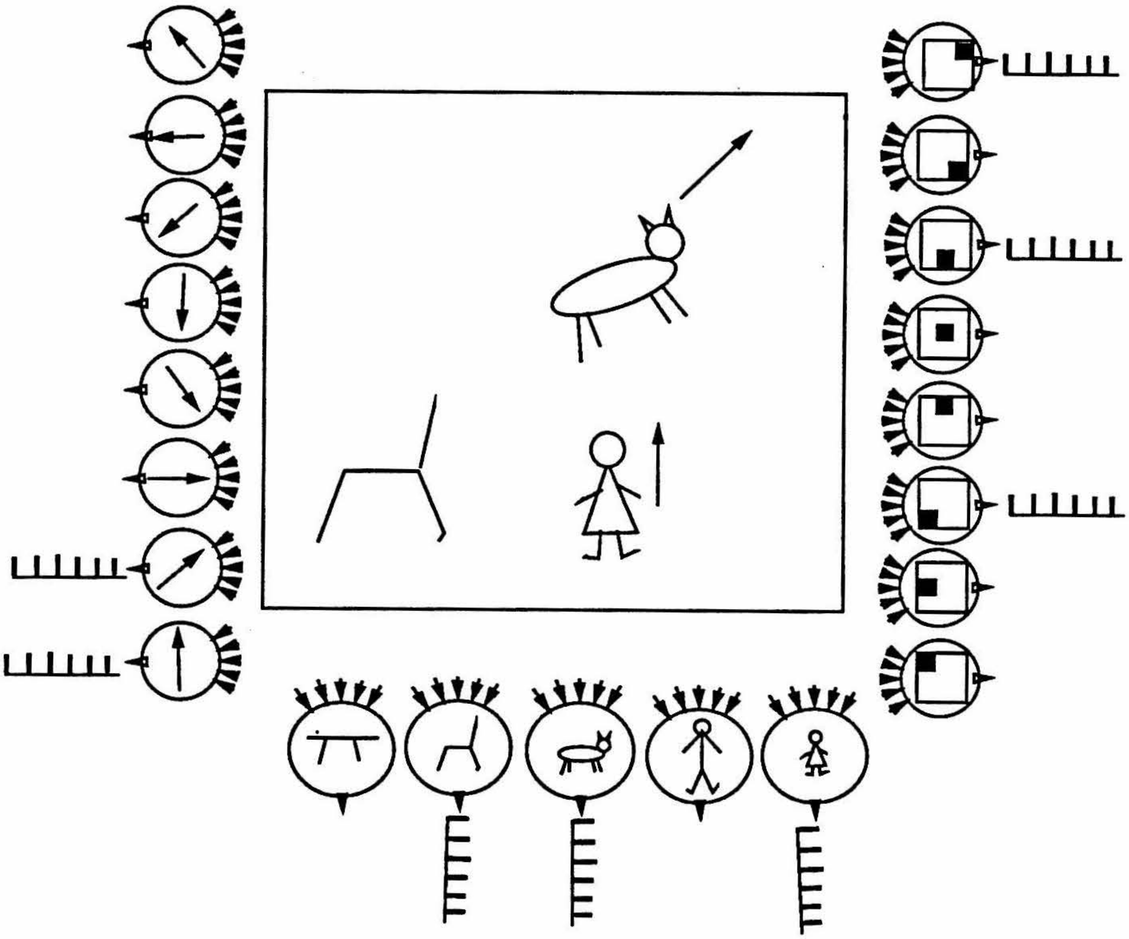

Figure 4.1: **The Binding Problem.** In this idealized example there are three classes of neurons: motion detectors (left), shape detectors (center) and position detectors (right), all with receptive fields spanning the square shown. In this example there is a chair (not moving), a moving cat, and a moving girl all stimulating the neurons. But when more than one stimulus is present in each receptive field, there is no way to use the neurons' firing rates alone to decide which shape-detector ought to be paired with which location-detector and with which motion-detector, because analog intensity signals do not carry independent information about their sources.

sponses *already* represents one stimulus property, there is no complementary information with which to distinguish two identical analog signals from one another... the only remaining way to distinguish them is through temporal modulation of the analog signal. Previous suggestions for temporal modulation to solve this binding problem have included oscillations of firing rate (Gray and Singer 1988) or synchrony of spikes within regular spike trains (von der Malsburg and Schneider 1986); a review is available in Koch (1993).

But temporal modulation demands a higher bandwidth of neural response than DC signals alone; by definition one cannot modulate a DC signal without broadening its power spectrum away from zero. But in fact cortical spike trains *do* already have a very high bandwidth (hundreds of Hz), if one interprets their white-noise power spectra (Bair *et al.* 1992) and near-Poisson firing irregularity as broadband information rather than as noise. Can this "extra" bandwidth from irregularity be the same as the "missing" bandwidth needed for binding? Can binding be performed by individual spikes in an *irregular* train?

I will speculate that the answer is "yes," outlining without hard evidence a scheme in which simple coincidence-detecting neurons produce highly irregular spike trains. Those trains' average rates will represent the individual stimulus intensity in the usual manner, but now the individual spike times will carry information relating neurons to each other.

This proposed binding mechanism—we call it "Point-of-Origin Binding"— is different from some other binding in two respects. Most proposed binding schemes assume that the neural code is an analog signal; those schemes bind or label neurons through either common values of average firing rate (Kammen et al. 1989), burst rate (von der Malsburg and Schneider 1986), or through temporal modulation of an average rate (Millner 1974; Sompolinsky *et al.* 1991; Tononi *et al.* 1992; and Engel et al. 1992, and references therein). In contrast, Point-of-Origin binding links neurons by a slightly above-average rate of *coincidences* in their individual spikes; the detection of such subtle, millisecond-scale coincidences requires a highly sensitive cell such as the one proposed above.

The second difference between this binding scheme and some others arises in deciding which neurons correspond to which objects. The binding in many conventional schemes is induced by global characteristics of features: one binds together neurons responding to similar colors, motions, positions, times of appearance, etc. Those are "top-down" approaches, requiring feedback or attractor dynamics to establish binding. The present method, like that of Horn *et al.* (1991), is purely "bottom-up." The Point-of-Origin scheme binds features together by preserving a temporal record of each neuron's spike *input* from earlier layers, so that "bound" neurons share a higher-than-random number of coincident output spikes due to some shared input spikes. This method of multiplexing binding information with average spike rates requires that the bound features share some source neurons in common, so this method works only within a single modality (e.g., vision), but not between modalities. An example follows.

4.2 Cartoon Example

Let us consider a highly idealized visual task, just to illustrate the idea (unrealistic aspects will be discussed later). The example will have only two classes of neurons: those that detect shape (independent of position), and those which detect position, independent of shape. We wish to detect separately the shape and position of an object, and to bind them together afterwards.

Visual patterns (letter shapes) will appear as black pixels on a white screen (leaving aside problems such as brightness, center-surround structure, stereo, motion, etc.). Each pixel's output is a random pulse train if the pixel is "on" and no pulses if the pixel is "off." The rate of one pixel's "on" output is R_p and the width of each spike (taken as a rectangle) is w, so that the duty cycle of a train (or the probability that its output will be 1 rather than 0) is

$$P_p \approx R_p w \tag{4.1}$$

If a pulse-width is taken to be one millisecond, then P_p is about equal to the spike rate in kHz. In fact, P_p has an upper bound of unity, while spike rate is in principle unlimited. But we will assume that spike rates are slow enough and widths narrow enough that the occasional overlap of closely-spaced pulses does not cause P_p to deviate much from spike rate. (In a real neuron, of course, refractory periods prevent this limit from being reached.)

4.2.1 Location-Detector Neurons

A location-detector cell will be a simple coincidence-detector, firing a single output pulse of duration w if two input voltage pulses overlap (i.e., $M = 2$ in Appendix F and section 3.8; the simulations of chapter 3 actually produced no values of M less than 5, but five-fold coincidences are beyond the simple treatment here). The location-detector's inputs will be from a receptive-field region of RF_L adjacent pixels (e.g., $RF_L = 9$ in Figure 4.2; the subscript "L" stands for "location," and "S" for "shape" or form). A single "on" pixel could not fire the detector, because a single input train lacks the necessary coincidences. Two "on" pixels *could* fire this detector, when spikes from those two pixels happen to coincide. In general, the rate at which this detector fires depends on the number of its input lines active ($A_L \leq RF_L$) and on the rate P_p at which each produces pulses; we will require $A_L P_p \ll 1$ to ensure that output pulses remain rare. The chance of getting two pulses (from among A_L active lines) to overlap at some point is

$$\begin{aligned} P_L &= 2P_p^2 \binom{A_L}{2} && (4.2) \\ &= A_L(A_L - 1)P_p^2 + O(P_p^3) && (4.3) \end{aligned}$$

(where the factor of two results from the non-zero width of both pulses). This firing probability P_L can be interpreted as the firing rate of the location-detector, which fires as an approximately quadratic function of the overlap A_L of the stimulus image with the N-pixel receptive field. Super-linear functions have proved useful in many computations in vision (Suarez and Koch 1989; Koch and Poggio 1987 and 1992).

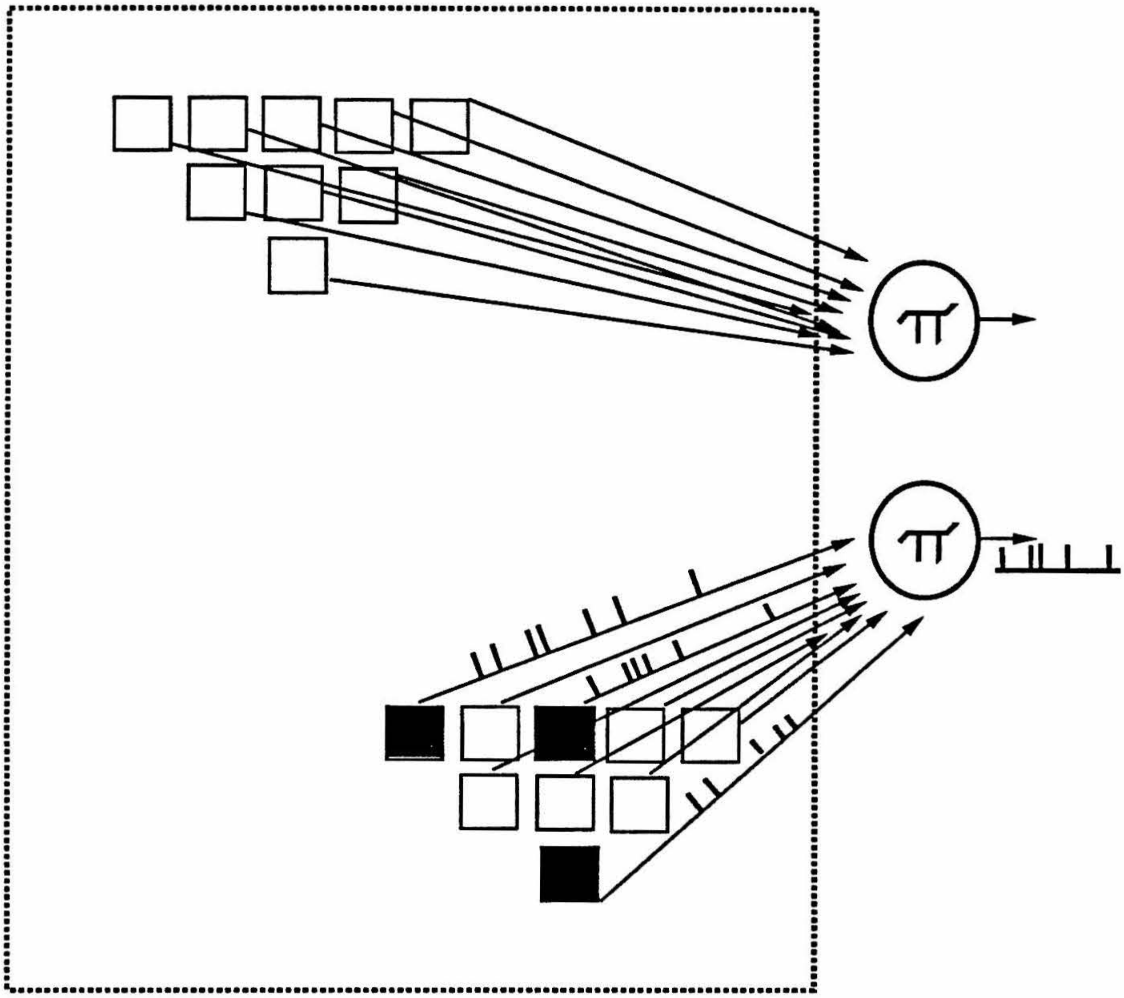

Figure 4.2: **Simple Location-Detector Neurons.** Each neuron is capable of receiving input pulses from nine pixels (arranged in the receptive field shown), and only fires upon the coincidence ("π") of two pulses. A pixel produces random pulses if it is "on," and no pulses if it is "off." The upper neuron, with no input, produces no output; the lower one has three active pixels and fires upon their occasional coincidences. In this model, such neurons tile the entire visual area with their small receptive fields.

But we are not only interested in average firing rates, but in spike-by-spike correlations. Whenever this detector fires, we know that 2 of its A_L inputs carried pulses, so that each output pulse is correlated with any active input line with probability $2/A_L$ (e.g., 2/3 in the example of Figure 4.2). The random chance of finding a pulse on an input line is $P_p \ll 1$, so that the correlation of output pulses with input pulses (to within one time window) is well above chance. This is the central principle of Point-of-Origin binding.

The entire visual field will be tiled with such detectors with overlapping receptive fields, so that several adjacent ones will be fired by any localized stimulus, such as the letter-shapes discussed below.

4.2.2 Shape-Detector Neurons

Constructing a neuron which responds to a particular shape, independent of position, is a bit more complicated. We will first construct a single localized subunit which works like the location-detector above, and then tile the visual field with many subunits to make a single position-independent neuron.

A subunit to detect a particular shape—for example 'X' or 'T' in Figure 4.3—will have a receptive field of that shape. RF_S will denote the number of pixels in a subunit's receptive field ($RF_S = 5$ in the example of Figure 4.3). The subunit's output rate will depend on the number of active input lines A_S in that receptive field as in eq. 4.2.

But an optimal response from this subunit requires that the stimulus be exactly

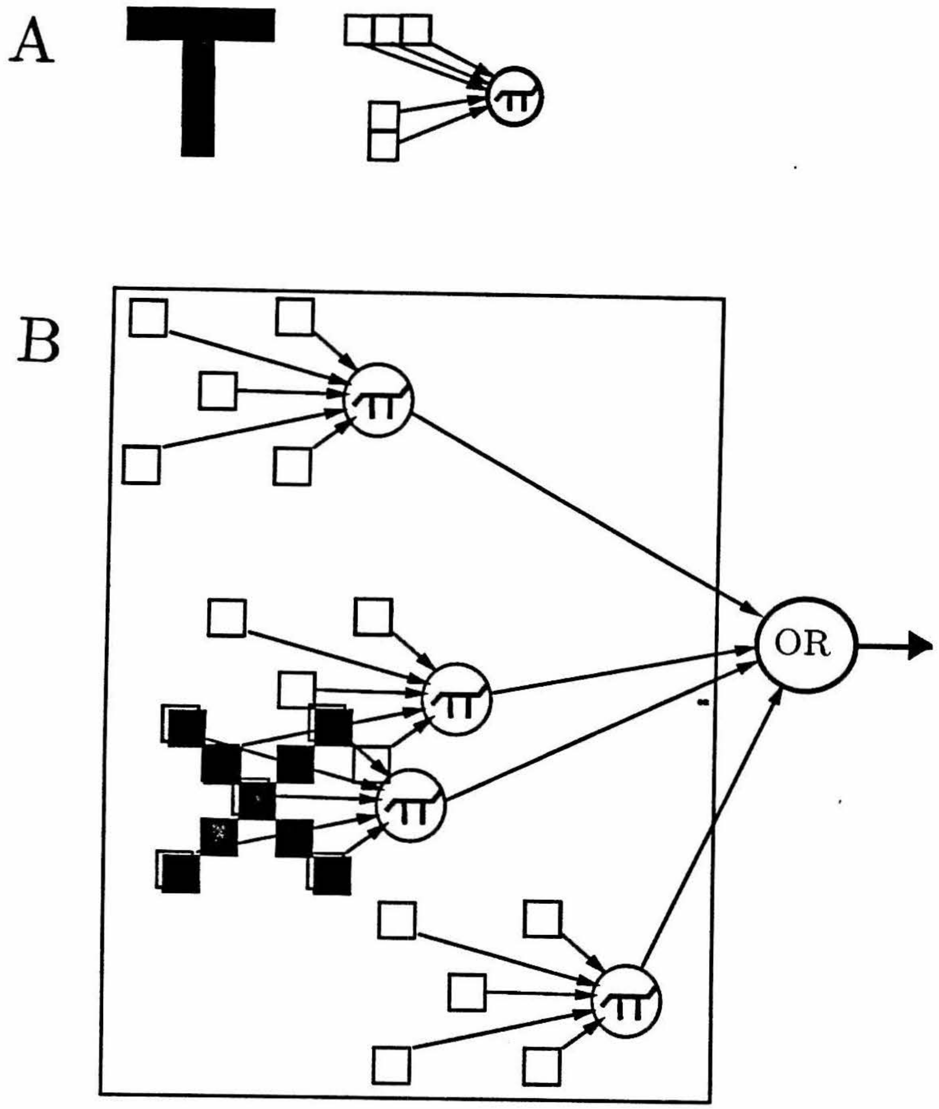

Figure 4.3: **Shape-Detector Neurons.** Each neuron has a receptive field the size of the whole visual space, and is designed to respond if its preferred shape appears on one of its subunits anywhere in that space. **A**, The 'T' shape chosen, and a single subunit for the 'T'-detector. The subunit fires upon the coincidence of any two spikes from the pixels in its input (in the same manner as the location-detector does). **B**, An entire X-detector cell collects pulses from the many subunits which tile the visual space, and fires an output when *any one* of those subunits fires (i.e., it performs an OR function on its subunits).

aligned with the subunit's receptive field. How do we make the response of the whole neuron position-*independent*? We let the shape-detector's output be the OR function of the whole set of subunits which tile the visual field: whenever *any* subunit fires, the shape-detector will fire. In the illustration given, the shape-detector's firing will be dominated by the single subunit with perfect overlap (i.e., $A_S = 5$), which will fire with probability $20P_p^2$ (eq. 4.3). For this particular pair of position and location detectors, careful overlaying of them shows that there are eight other subunits which are misaligned (e.g., with only $A_S = 2$); but outputs from any of those subunits will occur at the same time as spikes from the optimal subunit, and so will not contribute any additional output spikes. So an 'X' anywhere on the field will fire the cell at about $20P_p^2$. A different shape-detector can be made by the same principles, merely by rearranging the A_S pixels of each subunit (e.g., the 'T' detector in Figure 4.3).

The correlation of a shape-detector's output spikes with its input is higher than random, as was true for the location-detector: each output spike is coincident with (on average) $2/A_S$ of an input spike for the optimal subunit, which contributes all the output spikes. So for the 'X' example, every output spike will be caused by 2/5 of an input spike, a fraction which should be well above chance.

4.2.3 Binding the Outputs

We now have our two classes of detectors: a large set of location-detectors (one at each visual location), and two shape-detectors (one for 'X' and one for 'T') which respond independent of position. How will this array of cells respond to multiple inputs, and how might the nature of the inputs be reconstructed from the cells' activity?

Suppose the pixel screen contains exactly two objects: an 'X' in the upper left, and a 'T' in the lower right (Figure 4.4). Different neurons in both classes will fire: the 'X'-detector and the 'T'-detector (both firing at rate P_S), and several location detectors. For simplicity, let us focus only on the fastest-firing location detectors, the upper-left one (with P_{Ll} and $A_S = 5$) and the lower-right one (with P_{Lr} and $A_S = 7$; Figure 4.4).

The shape-detectors will fire primarily but not entirely due to their optimum stimuli. As we saw above, optimal stimuli create a $20P_p^2$ firing rate... but in addition, each detector sees a non-optimal stimulus (X on a T-detector, or *vice versa*) in a few of its' subunits. One can show that a T will overlap on an X-detector and an X will overlap a T detector in six ways to drive two pixels ($A_S = 2$). As an upper bound on each detector's "background rate," we can take all of those possible pairs as being independent, so that all those misaligned subunits fire together with probability $12P_p^2$ (eq. 4.3) (about half of their optimum rates), giving

$$P_S \leq 32P_p^2 \tag{4.4}$$

Now we can try to reconstruct the inputs from the outputs. If one only looks at the average spike rates (the traditional analog code), there are several ways to interpret these combined neural firings: 1) the correct way: X in upper left, T in lower right; 2) the reversed (incorrect) way: X in lower right, T in upper left; 3) a single combination of X and T in one location, and something which is neither X nor T in the other location. Only some further clues could distinguish which interpretation is correct, based on analog signals alone.

But the spike timing contains enough information to solve the problem. Consider the rate of coincidences among correctly matching neuron pairs. In this example, the location detector is driven by three of the same five pixels which drive either X or T shape-detectors, and additionally by two other pixels (X) or by four others (T) which do *not* drive the shape-detector (Figure 4.4). So any two correctly-matching detector pairs share a fraction of their inputs in common. Denoting this fraction by S, and denoting the absolute number of shared pixels by A_{com}, we have

$$S_L = A_{com}/A_L \tag{4.5}$$

$$S_S = A_{com}/A_S \tag{4.6}$$

The values of S in this example range from 3/5 to 3/7. The odd shapes of the detectors in this example were chosen to keep A_{com} the same for both X and T.

As is computed in Appendix F (eq. F.6), two coincidence-detecting neurons firing at the same rate and sharing S common input lines will have a fraction C_c of their *output* spikes coincident above chance. For this example, the X

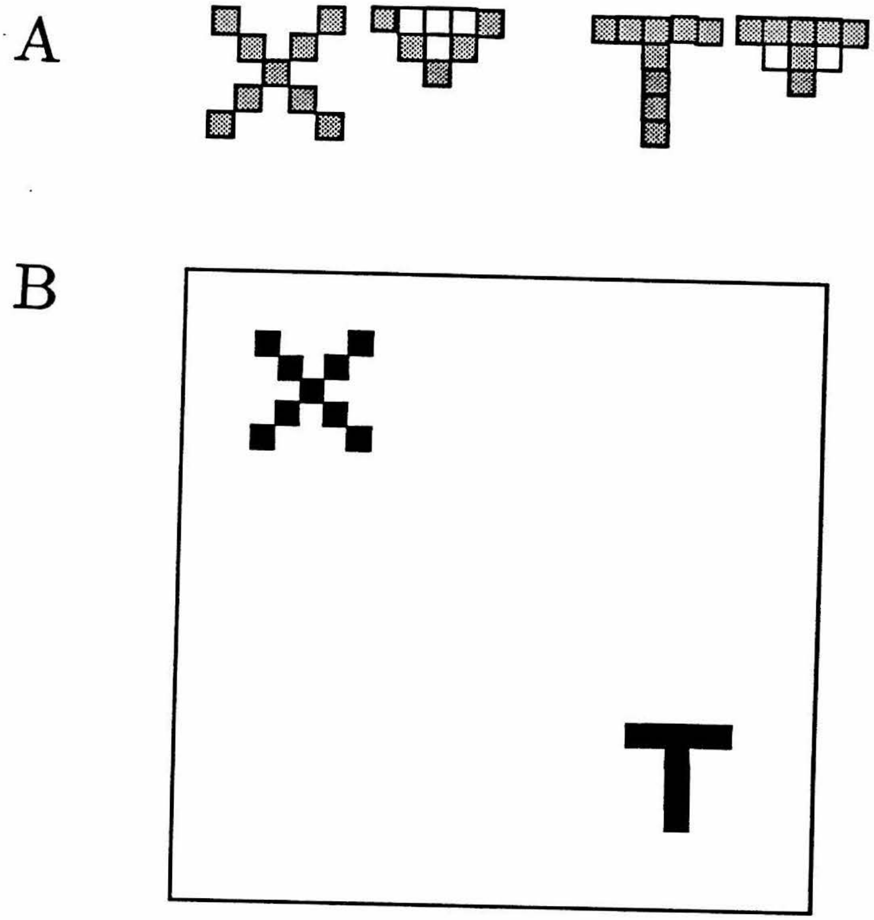

Figure 4.4: **Response to Multiple Stimuli.** **A**, The shapes X and T and their overlap with location-detectors. Note that of the pixels driving the location detectors, in both cases three of those pixels also drive the optimal shape-detector (the inverted triangle of Figure 4.3). **B**, when an X is presented in one part of the visual space and a T is presented in another part, four detectors will fire strongly: the X-detector, the T-detector, an upper-left location detector, and a lower-right location-detector. While the average firing rates of those neurons do not indicate which shape is at which location, the correctly-matched neuron pairs will share a higher than random rate of coincident output spikes as a result of their three shared input spike trains, so that the original patterns can in principle be reconstructed.

shape provides input to its location and shape detectors at the same rate ($A_L = A_S$), so that eq. F.6 gives their above-random rate of coincidences in those detectors as roughly

$$\begin{aligned} C_c &\approx S^2 + 2\binom{2}{1} S P_{non} && (4.7) \\ P_{non} &= (1 - S) P_{Ll} && (4.8) \\ C_c &= \frac{9}{25} + \frac{24}{25} P_{Ll} && (4.9) \\ &\approx \frac{1}{2} && (4.10) \end{aligned}$$

(The term P_{non} represents the chance that one of the unshared or *non*-common lines in a detector will fire).

To find the chance of coincidences from *incorrectly* matched neurons, we might naively assume they are driven by entirely different input trains, which have no overlap at all. For instance, a lower bound on the chance that a spike from the X-detector is coincident with one from the (wrong) lower-right location detector would be found from their independent, random probabilities

$$P_{x-r} \geq 2 P_{Lr} P_S \qquad (4.11)$$

These coincidences between incorrectly matching detectors will be rare, because we assume that spikes are rare ($P_L \ll 1, P_S \ll 1$). But in the example here, each shape-detector actually receives about a third of its output spikes (e.g., 12/32, eq. 4.4) from subunits which overlap the non-optimal (or incorrect) stimulus. So a better estimate is to assume that coincidences between the incorrectly matching neuron pairs will scale accordingly: the X detector will have about a sixth of its spikes coincident with those of the (wrong) location detector, while sharing about a third of its output spikes with its matching location detector.

Because the correct binding solution has a much higher coincidence rate than the incorrect solution ($1/3 > 1/6$), the necessary binding information is embedded in the spike trains. If necessary, a further coincidence-detecting neuron (in a higher layer) could distinguish bound pairs from unbound pairs.

4.3 Biological Implausibilities

There are countless drawbacks to this scheme. For instance, it becomes unworkable when the random coincidence rates between feature-detector spikes are statistically indistinguishable from coincidence rates due to common source neurons. This can obviously occur if too many features are present at once, if the non-optimal features still fire detectors efficiently, or if there are too many detector types present. But such drawbacks will probably arise in any binding scheme, not just in this one.

The disagreements with cortical biology are more glaring. At a cellular level: this scheme has ridiculously simple single cells, no explicit time-delays, no random synaptic transmission failures, no differences in synaptic strength, no inhibition, no cortical layers, no transient input components, no population coding, no columnar structure, and no feedback connections between cells or layers. At a visual-processing level: it has no realistic center-surround structure to its visual input, no scale-independence of patterns, no allowance for temporal input patterns, no orientation selectivity, no distinction between "parvo" and "magno" information streams, and no "top-down" influences.

I would like to improve upon this model in several ways. One could include input pulses which are non-rectangular, noisy, of different amounts of coincidence, and of different height (in order to limit the strength of numerous feedback connections). The calculation of firing rates could be done more flexibly using a probability distribution for membrane voltage, rather than explicitly calculating coincidence probabilities for binary inputs. One might add a feedback mechanism to allow probabilistic temporal integration without sacrificing temoral precision; another type of feedback might be "top-down," increasing the firing probability of those local units which contribute to large-scale features. Slight increases in firing threshold might serve the function of inhibition, limiting the influence of positive-feedback loops and normalizing the firing rates of the network as a whole.

4.4 Special Features

While this example uses location and shape as the two distinguishing features of an object, there is no reason to limit the features to only two, or to those two types. The arguments above can used to bind neurons responding to *any* feature-types which can share common driver neurons: texture, color, or left-eye/right-eye differences (binocular disparity). Motion-detection could be added by taking a temporal derivative of analog spike rate before coincidence-detection. There is also no reason that the driver neurons need to be individual pixels; they could equally be orientation detectors in V1, or in fact any kind of neuron which drives two or more higher-level detector types. The key point is that spikes in different feature-detectors share a common point of origin.

It is also not necessary that those input or driver spikes be entirely random. But if input spikes are instead regular, then that regularity may be reflected in the feature-detectors' firing as well, so that correlations across time could appear among the detectors' output spikes. Such correlations would not invalidate the estimates of spiking and coincidence probability (as calculated above), but then a higher layer would need to average out those correlations, which would take longer than averaging over a Poisson process.

The structure of the shape-detectors is like that of a sigma-pi neuron (Mel 1990), whose output is the sum of local multiplications of inputs. But previous models of those have only considered analog computations, without temporal structure or spiking.

For the Point-of-Origin scheme to work, there must be a significant overlap S in the inputs to two different feature-detector cells. As shown above (eq. 4.7) and in Appendix F (eq. F.6), the overlap in output spikes is roughly given by S^2, so that neurons sharing a small S will have low coincidence rates and will be thus much more difficult to bind together. This effect would suggest that receptive fields be fairly coarsely tuned, to offer the maximum possible overlap between different feature types and the best chance of binding their detectors together.

4.5 Advantages

The twin assumptions in this project— coincidence-detecting neurons and Point-of-Origin binding— are too highly simplified to be biologically plausible by themselves; they are intended in the toy-model spirit of a McCulloch-Pitts neuron or a Hopfield net. Still, this general scheme has many biologically desirable properties: The computation is parallel and asynchronous. The various features are analyzed separately, but linked (through spike-synchrony) in a distributed representation. The irregularity of spiking is not a hard-to-explain nuisance which contaminates the analog firing rate, but an explicit form of information transmission. The full bandwidth of the axon is used (near the kilohertz scale), despite the fact that visual inputs arrive on a fifty-fold slower timescale. Coding explicitly requires the all-or-none nature of the action potential ("digitally"), instead of using pulses to transmit an analog firing rate. The scheme is inherently probabilistic, suggesting that a network might be robust to noise or damaged components. And coincidence-detecting neurons provide a straightforward mechanism for multiplying neural signals, and hence for the nonlinear computations which feature-detection requires.

4.6 Conclusion

Our brains can make sense of the world only because the world makes sense. Complicated as the world may be, its complexity pales in comparison to the complexity of the pattern-space it inhabits. Think of a TV screen tuned to a vacant channel: all the countless pictures which might conceivably appear, pictures of sights real and imagined, are still never produced in the "snow" which best represents the huge combinatorial possibilities of the screen's individual pixels. Compared to unrestricted probability-space, the real world is a low-entropy state.

An important aspect of the world's simplicity is the fundamental physical laws which govern the behavior of objects. Conservation of mass and charge guarantee that objects do not suddenly appear or disappear; conservation of energy and momentum govern their motion. Reproducable laws of optics determine how objects are illuminated, and how images of the world fall on our retinas and move when we do.

All these laws not only make the world simpler than it might be, but they make many aspects of it *separable*: for example, the positions of objects and the ways in which objects move can be treated *independently* of the nature of objects themselves. And the persistence of objects from one moment to the next, combined with the finite number of object types we encounter, mean together that one might categorize object types, *independently* of where they are or how they move. It is thus possible to represent many aspects of the world most efficiently using a kind of *product space*, in which various aspects of objects (e.g., motion vs. form) are computed separately and later recombined... such a scheme could represent the enormous complexity of the world by the near-infinite combinations of a finite number of computational elements.

That it *could* be done does not mean it *is* done. But the existence in the brain of separate form- and motion-sensitive pathways, with receptive fields large enough to accomodate many objects each, shows that some visual tasks are indeed divided, and must therefore be recombined. This recombination requires that a neuron's signal must not only correlate with its particular stimulus, but also with other neurons which represent complementary aspects of that stimulus. A neuron must somehow carry multiple types of information along its single output axon. And the fact that each neuron has only one stimulus to represent—while there are many other neurons to relate to, even when perceiving just two objects—suggests that the lion's share of information transmitted must correlate neurons with one another.

This can be done by multiplexing the information in time. But multiplexing requires a high bandwidth, which in turn makes three demands on cortical signals:

1) The neural signals must be capable of changing quickly. Action potentials are by nature very fast, and can in some specialized cases phase lock to the stimulus with 30 μs precision (e.g., the big brown bat, Covey *et al.* 1991). Much slower millisecond computations might in principle occur in cortical cells.

2) That fast bandwidth must be *used.* In Fourier space, this means that the signals' power-spectrum must be broad rather than narrow, which is in fact the case for the near-flat power spectra of most cortical spike trains observed (Bair *et al.* 1992), due to their near-Poisson spiking process. In probability-space, this means that individual spikes must be unpredictable, so that each

one's arrival carries new information. Such irregularity does indeed seem to be the rule for all cortical neurons firing at all rates.

3) That high-bandwidth channel must not be very noisy. When cortical cells are stimulated under controlled circumstances, their firing is quite regular and reproducible, suggesting a potentially reliable computation. But the individual synaptic inputs to such cells are known to be much less reliable; it is not yet known whether redundancy in synaptic firings can compensate for that.

But in evaluating the nature of signals and noise in the cortical spike code, it is essential that the irregularity of cortical firing **not** be considered as "noise" *a priori*. We must first understand its source, the cell-to-cell synchrony embedded in it, and the perceptual tasks it might accomplish. Until we reach that understanding, it is possible in principle that the bothersome, unreproducible crackle of cortical action potentials is actually the sound of thought.

Appendix A

EPSP Width

Here we modify the simple integrator model's prediction of C_V to account for the finite temporal duration of EPSPs. The integrate-and-fire neuron assumes that EPSPs result from instantaneous current impulses, which carry equal Fourier components at all frequencies and hence produce a frequency-independent input to the neuron. But in fact an actual Excitatory Post-Synaptic Current (EPSC) is not a δ-function, but arises from a conductance approximately of the form

$$g_{syn}(t) \quad \propto \quad t \exp(-t/t_{peak}). \tag{A.1}$$

As long as the membrane potential is well below the synapse's reversal potential E_{rev}, we can approximate the synaptic current by

$$I_{syn}(t) \quad \propto \quad t \exp(-t/t_{peak}). \tag{A.2}$$

The non-zero width of these EPSCs blurs a pulse-train's high-frequency information. We want to find the amount of that blurring in synaptic current at some test frequency f; the blurring is given by the decrease in Fourier amplitude of the blurred current train at f relative to the unblurred spike-like train.

A train of realistic current impulses is given by convolving the individual synaptic current $I(t)$ with the random "comb function" $\sum_i \delta(t-t_i)$ of the spike-like inputs,

$$I(t) = \sum_i \delta(t-t_i) * I_{syn}(t) \tag{A.3}$$

The Fourier amplitude of $I(t)$ at frequency f is just the product of the separate Fourier amplitudes $\mathcal{F}_f(I_{syn}(t))$ and $\mathcal{F}_f(\delta(t-t_i))$ (by the Fourier convolution theorem). The random comb contribution to C_V has already been determined (eq. 2.12), so we only need to examine synaptic current smoothing by the single-event term $\mathcal{F}_f(I_{syn}(t))$. That attenuation $A(f)$ of a single EPSC relative to a delta-function is the EPSCs Fourier amplitude at f, normalized by its area (found by using $f=0$):

$$\begin{aligned} A(f) &= \frac{|\mathcal{F}_f\left(I_{syn}(t)\right)|}{|\mathcal{F}_0\left(I_{syn}(t)\right)|} && \text{(A.4)} \\ &= \left|\frac{\int_{-\infty}^{\infty} I_{syn}(t)\exp(2\pi i f t)dt}{\int_{-\infty}^{\infty} I_{syn}(t)dt}\right| && \text{(A.5)} \\ &= \frac{1}{(2\pi f t_{peak})^2+1} && \text{(A.6)} \end{aligned}$$

At what frequency f do we wish to evaluate this attenuation? Suppose that some spikes occur with mean interval $\overline{\Delta t}$ and some scatter ϵ about that mean, so that sequential intervals are

$$\{\Delta t_i\} = (\Delta t+\epsilon), (\Delta t-\epsilon), (\Delta t+\epsilon), (\Delta t-\epsilon), ... \tag{A.7}$$

This simplified example, with variability present only at a single frequency, has periodicity $2\overline{\Delta t}$, so that

$$f = \frac{1}{2\overline{\Delta t}} \tag{A.8}$$

(this is identical to the result got from the Nyquist Sampling Theorem). Using this estimate of f, we conclude that

$$A(f) \approx \frac{1}{(\pi t_{peak}/\overline{\Delta t})^2 + 1} \tag{A.9}$$

$$\approx 0.5, \tag{A.10}$$

for $\overline{\Delta t} = 4\ msec$ and $t_{peak} = 1.5\ msec.$ $A(f)$ represents the attenuation of current variability at f reaching an integrate-and-fire neuron, due to synaptic blurring. If we suppose that this attenuation of current roughly corresponds to the attenuation of C_V (see eq. B.1), then $A(f)$ (eq. A.9) should be multiplied by the perfect-integrator prediction (eq. 2.12). This blurring makes it more difficult to reconcile the observed variability with theory, even at low N_{th}: how can a neuron produce output varations whose frequency is higher than that contained in a single EPSP?

Appendix B

Spike Adaptation

Here we modify the simple integrator model's prediction of C_V to account for the slower spiking which results from sustained hyperpolarizing currents. The spike rate in the monkey cells decreases by about half in the first $100-300$ *msec* of a train. Some of this decrease is likely due to a decrease in synaptic input to the cell, but a major contribution to this slowdown is probably due to the "spike adaptation" potassium currents. How will these negative currents affect the predicted C_V values if the synaptic input remains unchanged?

Because I_{AHP} (considered as a single current) has a reversal potential much lower than the resting potential, it can be modelled (to first order) as a negative sustained current, which partly cancels the inward sustained portion of the EPSC while leaving its fluctuations unchanged. As a result, the mean ISI will increase during this adaptation. We wish to approximate the influence of this increased ISI on the perfect integrator with absolute refractory period (section

2.4.2).

Let us assume a low variability in the perfect integrator model ($C_V \ll 1$), so that the proportional variation in time $C_V = \sigma_{\Delta t}/\overline{\Delta t}$ to reach a fixed threshold of $N_{th}v$ is roughly the same as the proportional variation in the relative synaptic depolarization $\delta V/V_{th}$ arriving in the mean time interval $\overline{\Delta t}$:

$$\frac{\sigma_{\Delta t}}{\overline{\Delta t}} \approx \frac{\delta V}{V_{th}} \tag{B.1}$$

$$C_{V_{AHP}} \approx \delta V_{AHP}/V_{AHP}. \tag{B.2}$$

As the negative current increases, firing frequency decreases, and the mean ISI during adaptation $\overline{\Delta t}_{AHP}$ will increase above the ISI at the start of the train $\overline{\Delta t}_{st}$,

$$\overline{\Delta t}_{AHP} > \overline{\Delta t}_{st}. \tag{B.3}$$

But by assumption the depolarization necessary to reach firing threshold will remain constant:

$$V_{AHP} = V_{st} = V_{th} \tag{B.4}$$

How much variation in depolarization (δV_{AHP}) will accumulate during that longer ISI? The rate of random EPSCs is assumed to be unchanged by the addition of negative adaptation current. So it is clear that the expected number of EPSCs arriving in $\overline{\Delta t}_{AHP}$ will increase, in proportion to $\overline{\Delta t}_{AHP}$. But the variation about that mean will not increase linearly with interval duration, but rather as its square root, as occurs in accumulating any large number of independent events in a single time period:

$$\delta V_{AHP} \approx \delta V_{st}\sqrt{\frac{\overline{\Delta t}_{AHP}}{\overline{\Delta t}_{st}}} \tag{B.5}$$

The above formula applies only during the integration time, and is thus only valid for the perfect integrator without refractory period. Recognizing that the observed ISI contains the refractory period t_0 means that the true integration time is $\overline{\Delta t} - t_0$,

$$\delta V_{AHP} \approx \delta V_{st} \sqrt{\frac{\overline{\Delta t}_{AHP} - t_0}{\overline{\Delta t}_{st} - t_0}} \tag{B.6}$$

Combining eqs. B.2, B.4, and B.6 gives us an expression for C_V as a function of ISI for an adapting spike-train:

$$\begin{aligned} C_{V_{AHP}} &\approx \left(\frac{\delta V_{st}}{V_{th}}\right) \sqrt{\frac{\overline{\Delta t}_{AHP} - t_0}{\overline{\Delta t}_{st} - t_0}} && \text{(B.7)} \\ &\approx C_{V_{st}} \sqrt{\frac{\overline{\Delta t}_{AHP} - t_0}{\overline{\Delta t}_{st} - t_0}} && \text{(B.8)} \end{aligned}$$

This rough result indicates that during spike adaptation, C_V will rise with the square root of $\overline{\Delta t}$, much faster than the leaky-integrator model suggests. Such a sharp increase in variability is indeed observed both in simulations (Figs. 2.10-2.14) and in the monkey data. But the magnitude of the rise in C_V in simulations is not sufficient to account for the larger discrepancy between theoretical and observed C_V, nor can it account for the high C_V observed for early spikes, before the onset of adaptation.

Appendix C

Irregular EPSP Magnitude

Here we modify the simple integrator model's prediction of C_V to account for variable magnitude in synaptic input pulses.

Not all EPSPs have the same magnitude. In fact, the uncertainty in EPSP depolarization σ_v may be nearly equal to the average depolarization v (as reported in Mason *et al.* 1991), so that our model's input consists of EPSPs of random depolarization ($v \pm \sigma_v$) arriving at random times. Clearly this additional source of variation will increase C_V, but by how much?

Let us assume a low variability in the perfect integrator model (as in Appendices A and B), so that the proportional variation in time $C_V = \sigma_{\Delta t}/\overline{\Delta t}$ to reach a fixed threshold of $N_{th}v$ is roughly the same as the proportional variation in the relative synaptic depolarization $\delta V/V_{th}$ arriving in the mean time interval $\overline{\Delta t}$:

$$\frac{\sigma_{\Delta t}}{\overline{\Delta t}} \approx \frac{\delta V}{V_{th}}. \tag{C.1}$$

Because on average the cell receives N_{th} EPSPs during $\overline{\Delta t}$,

$$V_{th} = N_{th} v. \tag{C.2}$$

The scatter about the average number of EPSPs is $\sqrt{N_{th}}$, so we expect that the depolarization will be composed of $N_{th} \pm \sqrt{N_{th}}$ EPSP's, each of magnitude $v \pm \sigma_v$, that is

$$V_{th} \pm \delta V \approx \sum_{1}^{N_{th} \pm \sqrt{N_{th}}} (v \pm \sigma_v) \tag{C.3}$$

Using the convenient (although inexact) assumption that σ_v is scattered about v in a Gaussian fashion, we can add the random variables in quadrature (root-sum-square), dropping second-order terms:

$$\begin{aligned} V_{th} \pm \delta V &\approx \left(\sum_{1}^{N_{th} \pm \sqrt{N_{th}}} v \right) \pm \sqrt{\sum_{1}^{N_{th}} \sigma_v^2} && \text{(C.4)} \\ &= N_{th} v \pm v\sqrt{N_{th}} \pm \sigma_v \sqrt{N_{th}} && \text{(C.5)} \\ &= N_{th} v \pm \sqrt{N_{th}(v^2 + \sigma_v^2)}. && \text{(C.6)} \end{aligned}$$

By applying our assumption that the standard deviation of EPSP amplitude is about the same as its mean ($\sigma_v = v$; Mason *et al.* 1991) and combining eqs. C.1, C.2, and C.6 we arrive at

$$\begin{aligned} \frac{\delta V}{V_{th}} &\approx \sqrt{\frac{1 + (\sigma_v/v)^2}{N_{th}}}, && \text{(C.7)} \\ C_V &\approx \sqrt{\frac{2}{N_{th}}}. && \text{(C.8)} \end{aligned}$$

(This result was derived in a different manner in Stein 1967a). Thus, the maximal realistic variation in EPSP depolarization can only multiply the perfect integrator's C_V (eq. 2.12) by $\sqrt{2}$, not a sufficient amount to account for the discrepancies observed.

Appendix D

Cross-Correlation Analysis

While the assumption that cortical cells fire completely independently has the appeal of simplicity and symmetry, it is not entirely accurate. The best-known and simplest method of describing correlations between spike trains is cross-correlation analysis, which is described in this section.

D.1 Formal analysis

Cross-correlation analysis compares spike trains from two cells which are firing at the same time (comparison of more than two cells is difficult, both because recording from more than two cells at once is difficult and because the analysis and display of higher-order correlations is cumbersome; see Abeles (1990) and Softky and Kammen (1990)). The cross-correlation histogram, or "cross-correlogram," is essentially the multiplication of one spike train by another, or the discrete form of the correlation function.

Suppose we have spike trains A and B (of equal duration T), with discrete firing times labelled by tA_i and tB_j, and total spike counts NA and NB. The cross-correlogram $C_{B-A}(\Delta t)$ is defined by

$$C_{B-A}(\Delta t) \quad = \quad \sum_{i=1}^{NA}\sum_{j=1}^{NB} \delta_{\Delta t,(tB_j-tA_i)}, \tag{D.1}$$

where $\delta_{x,y} = 0, 1$ is the Kronecker delta function. For example, if a spike in train B is more likely than chance to arrive right after a spike in train A, there will be a peak in the cross-correlogram at positive Δt. Coincidence is signalled by a peak at $\Delta t = 0$. Anti-correlations appear as dips, and the absence of correlation appears flat. If both trains have average rates which are correlated as well, there will be broad peaks in the cross-correlogram in addition to the narrower peaks of single-spike correlations. Because the Δt-axis is reversed under exchange of spike trains, the total area under the features remains unchanged.

The strength of correlations is easily measured by the "contribution coefficient" $C_c(A)$ or $C_c(B)$, or their mean C_c (Toyama *et al.* 1981). This approximate measure compares the above-chance number of spikes in a correlation peak N_{pk} to the number of spikes in train A or B, thus estimating the fraction of total spikes in one train which are causally connected with the other train. So if the peak extends from Δt_{min} to Δt_{max}, then the contribution coefficient is

$$N_{pk} \quad \approx \quad \int_{\Delta t_{min}}^{\Delta t_{max}} C_{B-A}(\Delta t) - C_{B-A}(\Delta t_{max}) d(\Delta t) \tag{D.2}$$

$$C_c(A) \quad \approx \quad \frac{N_{pk}}{NA} \tag{D.3}$$

The height of the cross-correlogram outside the peak $C_{B-A}(\Delta t_{max})d(\Delta t)$ is chosen as the baseline for determining the peak area.

In the practice of looking at cross-correlograms, on which multiple peaks may be superposed, the area N_{pk} under the peak can be subjective, depending on the width of the peak and on the time-window for "coincidence." (An example of the delicacy of this choice: in the case of a peak of half-width longer than the mean interspike time, one can in principle find more spikes under the peak than there are in the entire train, producing an uninformative and unphysical $C_c > 1$.)

We had the opportunity to examine cross-correlation data of exactly the same sort as analyzed in Chapter 2: pairs of cells in monkey area MT responding to partially coherent motion of random dots. Significant correlations between those cells—which presumably serve as input to other similar neurons nearby—might help explain the high firing variability we observed. We thus analyzed the 11 pairs of cells recorded by Zohary and Newsome. We chose $|\Delta t| \leq 7\ ms$ as the region in which to count the area under the peak (see Figure D.1 for an example). This arbitary choice—a coincidence-window even wider than many of the mean interspike intervals—still led to low values of C_c ranging from 0.12 down to 0, with mean 0.03. A shorter estimate of the coincidence-window would lead to yet lower values of C_c and smaller contributions to variability.

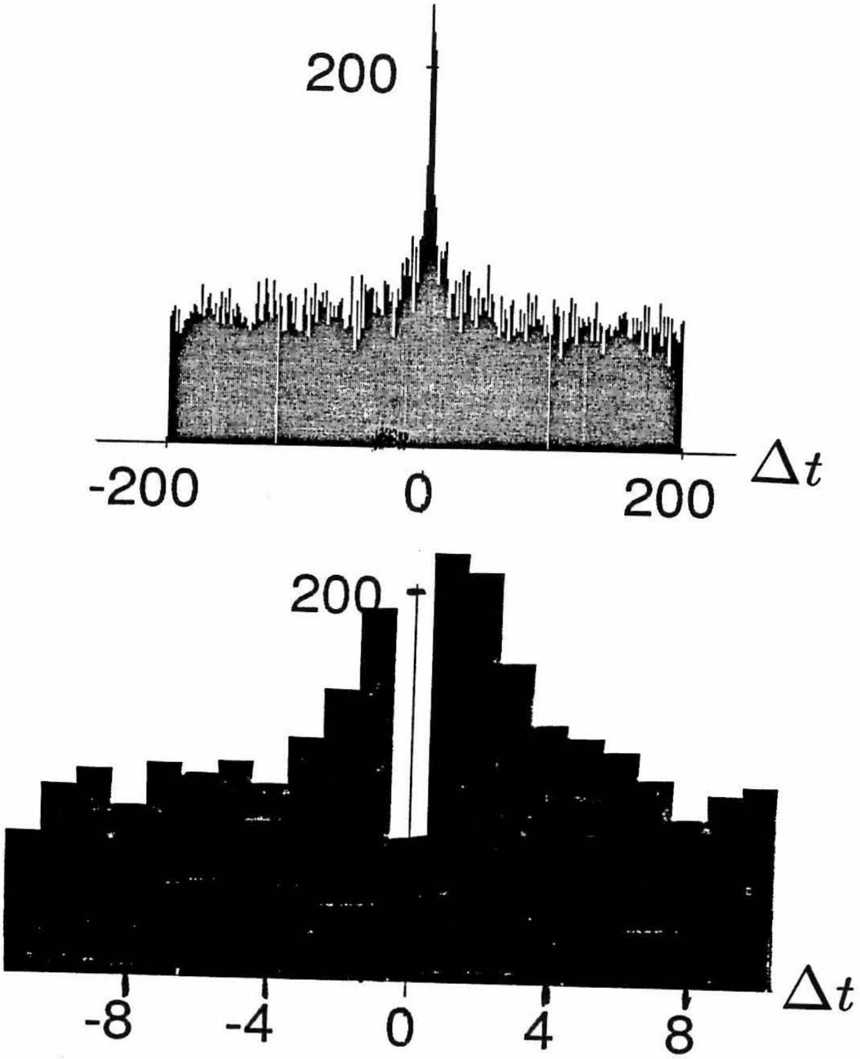

Figure D.1: **A sample Cross-Correlogram at Two Timescales.** A pairs of cells from visual area MT of monkey (Zohari and Newsome) exhibit the cross-correlations shown. A central peak (here between -4 and 6 *ms*) indicates the number of counts of a train which are coincident. Here the peak contains approximately 525 spikes, which is about 6% of the 8500 spikes in one train ($C_c = 0.06$). The dip at zero is an artifact of the recording method, which could not reliably record two simultaneous spikes on its single electrode.

D.2 Influences of Recording Method

But there are published reports of much stronger cross-correlation peaks than these, a situation which could result from the different cortical area studied or from the different recording method used.

The Zohary data above recorded two separate cells from a single electrode, distinguishing them by the different $V(t)$ shapes and peak heights of their action potentials. This is a common method, but it is usually incapable of recording two action potentials which occur within a millisecond of each other (i.e., when the separate action potentials would overlap); note the very low counts in the zero bin of Figure D.1 corresponding to few simultaneous counts from the two cells. Thus this method necessarily excludes the most tightly synchronized spikes of all... and there *is* evidence for such submillisecond synchrony in cortex.

The most quantitative cross-correlation analyses published are from studies of cell pairs in primary visual cortex (V1) of anaesthetized cat, done by Toyama *et al.* (1981). That group used two *different* electrodes (penetrating the cortex from different directions), so that completely independent recordings were possible. Their recordings of 25-70 nearby cell pairs reveal *average* values of C_c around 0.2 (about ten times Zohary's values above), and several cell pairs with central peaks at zero only 0.1 ms wide. There remains the possibility that some such data reflect the accidental recording of the same as both electrode tips converge to within 100 μm (F. Wörgötter, personal communication), although the cells' responses and reconstructed locations seem to rule this possibility out in several cases.

There are three important points about the Toyama *et al.* results. One is that they infer one cell's influence upon another (through the area of an occasional *off-center* peak) to have strength corresponding to the first cell "causing" about 1/10 of the second cell's spikes ($C_c \approx 0.1$); this connection strength is much stronger than one would infer from the measured size of EPSPs, which typically contribute about only 1% of the depolarization necessary to trigger a cortical cell, but not strong enough to account for the common-input peaks observed (see Appendix F).

A second point is that even those reports of a high average $C_c \approx 0.2$ (of *centered* peaks) would increase C_V to about 0.4 (through eq. E.19 and E.24), a significant amount but not enough to solve the C_V discrepancy by itself. A final point is that synchronization between cells at the 0.1 ms scale suggests that cortical cells are at least in principle capable of very fast information processing.

Appendix E

EPSP Synchrony Influences Firing Variability

Estimates of Synchrony in Cortex

We can estimate the synchrony of EPSP input by comparing the measured output synchrony of nearby cells in cortex, as described above. Because nearby cells in cortex are strongly connected to one another, we can further suppose that a collection of many such cells provides input to an intergrator-model, and we can see whether the observed rate of coincidences between pairs of cortical cells (thought of as "input" cells) is enough to create highly irregular firing in an "output" cell. The integrator-model here will be the simplest "perfect" integrator, with no spatial extent, leak, or refractory period.

Measurements of the cross-correlation between two spike trains can reveal the number of coincident spikes. The details of the cross-correlation method, its temporal resolution, and its application to real spike trains are discussed in Appendix D. Here we will simply assume that any coincidences are perfect (no time dispersion), and we will use experimental estimates of pair-wise coincidences (from Appendix D) to predict the rate of higher-order coincidences among a large population of cells. We then use that population estimate to produce a new estimate of the input variability to an integrator model, which will in turn give the integrator's new, higher C_V.

We assume that any one cell will have a fraction C_c of its output spikes occurring at the same time as a nearby cell's spikes ($C_c \leq 1$ is the "contribution coefficient" of the neighbor to the target cell, as discussed for cross-correlation measurements in Appendix D). Different measurements have yielded values of C_c varying from zero to 0.9 (Toyama *et al.* 1981) for individual cell pairs, and from 0.03 to 0.5 across populations, depending on measurement methods and the area of cortex investigated (for futher references and a more detailed discussion of cross-correlation analysis, see Appendix D).

Suppose that such pairwise correlations occur in a population of N neurons, each firing at the same average rate R, each pair having the same value of C_c. In this case a given spike train will have a fraction C_c of its spikes correlated with spikes from any other train, and the remaining fraction $1-C_c$ uncorrelated with spikes from that other particular train (although they may possibly be correlated with other spikes we do *not* record). Let us assume that all spikes, whether correlated or not, occur independently of others in the same train, so that the timing of events from a single neuron is still Poisson. If those N neurons provided the EPSP input to an integrator model, then we can estimate the fraction of the population's spikes which are coincident, and the mean number of cells participating in each coincidence. Although we use the pairwise correlations observed, we must make assumptions about higher-order correlations (multiple cells firing at once) in the absence of any knowledge of them. That correlation must fall between two extremes: the minimal case, in which the *only* correlations between cells are pair-wise correlations like those we measure; and the maximal case, in which coincident events involve *all* the neurons.

We will examine this simple population model in two regimes, as characterized by two parameters. One is the number of spikes (or multiplicity) $m \leq N$ in each synchronous volley; if spikes tend to fire in pairs, then $m = 2$, but if the whole population fires at once, then $m = N$. The complementary parameter is the fraction of total spikes $P_c \leq 1$ which fall those volleys (rather than being uncorrelated with anything). If every spike fired is synchronized with other spikes in a volley, then $P_c = 1$; if only a tenth of the spikes participate in volleys, and the other nine tenths are completely unsynchronized with other spikes, then $P_c = 0.1$. These parameters are constrained by the population size N and by the amount of correlation C_c as follows.

Suppose we calculate the chance that a single spike in one train is correlated (above random chance) with a spike in another train; this chance is C_c. The chance of the first spike being from some volley is P_c, as defined above. And if the volley size is fairly large ($m \gg 1$), so that the number of *other* spikes in that volley is $m - 1 \approx m$, then the net chance of pairwise synchrony is

$$C_c \approx \left(\frac{m}{N}\right) P_c \qquad \text{(E.1)}$$

We will examine two extreme cases: the smallest possible volleys (of size $m = NC_c$), with all spikes participating in a volley ($P_c = 1$); and the largest possible volleys (with $m = N$), where most spikes do *not* participate in volleys ($P_c = C_c < 1$). It will turn out that both volley sizes carry about the same variability.

First the large-volley case. If the cells are maximally correlated (within the constraint of $P_c = C_c$ measured), then all coincidences for all cells occur at exactly the same times, as the entire population fires in unison ($m = N$). In this case there are only two types of events: either all the neurons fire at once, or the various neurons fire independently of one another and of previous firings. This means that a single cell will fire a fraction C_c of its spikes as part of these giant collective events, and its remaining fraction of spikes $(1 - C_c)$ are random. In any given time interval ΔT, the whole population will fire $C_c \Delta T R$ coincident volleys of N spikes per volley, and $(1 - C_c)\Delta T R N$ single random spikes (Figure E.1 A). Note that the number of *volleys* in the population response does not scale with N, while the number of *single spikes* does... so a large population still has a correspondingly large number of independent spikes occurring in it.

In the simplest analysis, an integrator model exposed to such a population will "see" as input two superposed EPSP streams: one stream at rate $(1 - C_c)RN$ with EPSPs of unit amplitude, and one stream at rate $C_c R$ of amplitude N

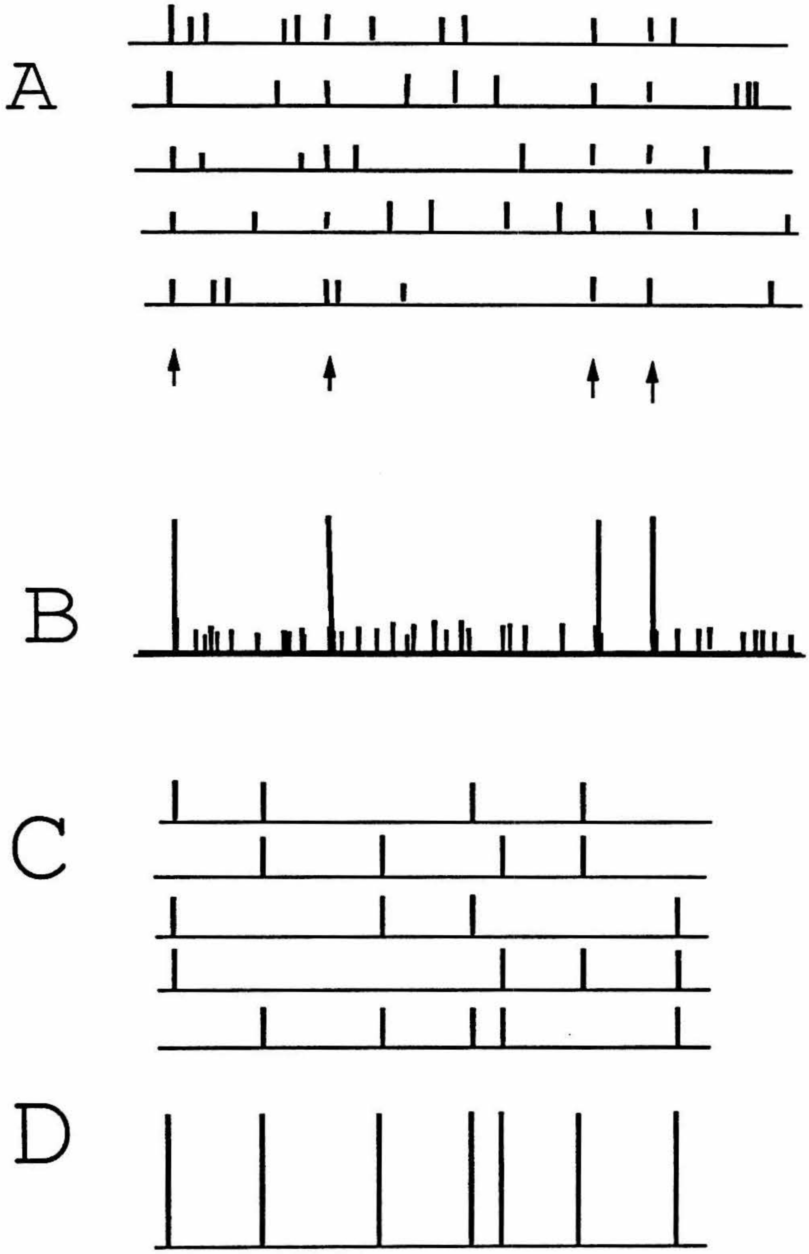

Figure E.1: (**A**), A small population of $N = 5$ neurons firing in partial synchrony. While most individual firings are random, on a few occasions the entire population fires in a synchronous volley (arrows; here the multiplicity $m = 5$ and only about $P_c \approx 1/3$ of the spikes participate in volleys). (**B**), The superposition of synaptic currents from those responses is equivalent to a single train containing many single events of unit amplitude and a few huge events of amplitude N. The standard deviation of event amplitude over that ensemble greatly increases the variability in current to an integrator model, and predicts that the model's firing irregularity will increase. (**C**), In an alternative type of synchrony, every spike participates in volleys of equal size m (i.e., $m = 3$ and $P_c = 1$). (**D**), In such a case, the number of independent events is typically far fewer, N'_{th} is larger, and hence an integrator-model will fire more irregularly.

(Figure E.1 B). Let us suppose (for simplicity) that $N = N_{th}$, so that a single volley can just barely fire the integrator. While this may seem an extreme situation for real cortical cells, it at least guarantees that the intregrator's output firing rate is the same as the firing rate of any input cell.

For this situation the number of independent EPSPs from the superposed streams necessary to fire the cell will be somewhat reduced. Of the events now appearing (as seen by the cell), most—a fraction $(1 - f)$—still have unit amplitude. The remaining fraction f have amplitude N, where

$$f = \frac{C_c R}{C_c R + (1 - C_c) R N} \tag{E.2}$$

$$\approx \frac{C_c}{(1 - C_c) N} \tag{E.3}$$

Note that a large N and a small C_c mean that $f \ll 1$. Let us examine the statistics of those events as they form a single superposed stream of input to the integrator. The average size of an event (A) has increased (slightly) above unity. Let A be the new amplitude:

$$A = f \times N + (1 - f) \times 1 \tag{E.4}$$

(Recall that we assume that the high-amplitude synchronized events have exactly N times the effect of single unsynchronized events, so that there are no amplifying or saturating nonlinearities, as there are in real cells).

But there is no single event with amplitude A, just a lot of smaller ones and a few much bigger ones. So the "average" event now has a very large variability in amplitude as well as in time,

$$\sigma_A^2 = (1 - f)(A - 1)^2 + f(A - N)^2 \tag{E.5}$$

$$\approx \quad f(N^2 - 2NA) > 1 \tag{E.6}$$

because $A - 1$ is small and N is large. Now we can estimate the variability in the output spike times (relative to $\overline{\Delta t}$) is well approximated by the relative variability in depolarization accumulated during a *fixed* time $\overline{\Delta t}$; this is the same trick used in appendices A, B, and I:

$$\frac{\sigma_{\Delta t}}{\overline{\Delta t}} \approx \frac{\delta V}{\overline{V}} \tag{E.7}$$

$$= \frac{\delta V}{\overline{N}}. \tag{E.8}$$

$$\tag{E.9}$$

We can then find δV by combining the standard deviations as Gaussian random variables, as we did in Appendix C. The number of independent events occurring in $\overline{\Delta t}$ is $(N/A) \pm \sqrt{N/A}$; the amplitude of each is $A \pm \sigma_A$:

$$V \pm \delta V = \sum_{1}^{N/A \pm \sqrt{N/A}} (A \pm \sigma_A) \tag{E.10}$$

$$= N \pm \sqrt{NA} \pm \sqrt{\frac{N}{A}}\sigma_A \tag{E.11}$$

$$= N \pm \sqrt{NA} \pm \sqrt{\left(\frac{N}{A}\right) f(N^2 - 2NA)} \tag{E.12}$$

$$= N \pm \sqrt{NA + fN(N^2 - 2NA)/A} \tag{E.13}$$

$$\approx N \pm N\sqrt{fN/A} \tag{E.14}$$

So we find the output variability when $f \ll 1$ to be

$$C_V = \frac{\delta V}{N} \tag{E.15}$$

$$= \sqrt{\frac{fN}{A}} \tag{E.16}$$

$$\approx \sqrt{\frac{fN}{1+f}} \tag{E.17}$$

$$\approx \sqrt{fN(1-f)} \tag{E.18}$$

$$= \sqrt{\frac{C_c}{1-C_c} - \left(\frac{C_c}{1-C_c}\right)^2} \tag{E.19}$$

This approximation obviously breaks down for $C_c \geq 1/3$, at which value C_V will appear to *decrease* as the correlations become stronger. But we will not need it in that regime.

What about smaller volleys? When every spike falls in some volley, and no unsynchronized spikes exist, then $P_c = 1$ and $m = C_c N$ (Figure E.1 C). This case is simpler, because all inputs are members of a volley, and all volleys—and hence all independent events—have equal size $m = C_c N$. (In assuming that every spike participates in a volley of size at least $m \geq 2$, we implicitly assume $C_c \geq 2/N$, so that this argument requires a large population or a strong cross-correlation). Because the size of each independent event has increased from unity to m, the average size has obviously increased (Figure E.1 D),

$$A = m \tag{E.20}$$

$$= C_c N, \tag{E.21}$$

so that only $N'_{th} < N_{th}$ independent events (volleys) are necessary to reach threshold:

$$N'_{th} = \frac{N_{th}}{C_c N} \tag{E.22}$$

If we again assume that $N = N_{th}$, then C_V is available from the original perfect-integrator formula (eq. 2.12):

$$C_V = \frac{1}{\sqrt{N'_{th}}} \tag{E.23}$$

$$= \sqrt{C_c} \tag{E.24}$$

This small-volley variability is slightly larger than the one calculated using large volleys (eq. E.19). Is this high enough to match the firing variability observed in cortex? Using the Zohary monkey data (which represents a probable lower bound of about $C_c \approx 0.03$, Appendix D), we would have $C_V = 0.17$; using the highest C_c values measured for populations ($\langle C_c \rangle \approx 0.2$ for cat, Toyama *et al.* 1981), we would have $C_V = 0.44$, which is remarkably high, but not yet in the range of the monkey data ($C_V \approx 0.6 - 1.0$).

There is a surprising amount of synchrony in this small-volley case, even though the correlations assumed are the lowest-order possible. The total absence of uncorrelated events, and the fact that volley size scales with population size, mean together that arbitrarily large neuron populations contain only the equivalent of $N/C_cN = C_c^{-1}$ independent trains of events. Slight degrees of pairwise correlation can still represent large amounts of synchrony—$m \gg 1$, with each spike "duplicated" many-fold—if all spikes participate in the correlations.

Why do small pairwise correlations imply large amounts of synchrony in large populations? Each neuron in a population is in principle capable of contributing a train of spikes which is independent of the other spike trains, so that in a rough sense the number of degrees of freedom is proportional to population size N. But each pairwise correlation represents an independent constraint on those degrees of freedom, and the pairwise correlations exist for *all possible* pairs of neurons in the population. So the number of constraints scales as N^2, while the number of degrees of freedom scales only as N.

Both of these models of spiking synchrony are very simplistic. They both reach the same prediction of C_V, and both suggest that spiking synchrony alone cannot account for the high irregularity of cortical firing, when appearing as input to an integrator-model. (Appendix F further suggests that integrator-models would not be able to *produce* this degree of spiking synchrony either). But the second of these models, in which all spikes participate equally in synchronous volleys, shows that even slight amounts of measured synchrony between pairs of cells may mask a huge amount of multiplicity among individual spike events, and hence may reflect a large amount of synchrony, redundancy, and degeneracy in the population firing.

Appendix F

Paradoxical Cross-Correlations

We suggested in section 3.9.2 that the abundance of narrow cross-correlation histogram (CCH) peaks centered at zero time is better explained by neurons which act as single-pulse coincidence detectors than by traditional integrator neurons. Here is the explicit argument.

While the most prominent cross-correlation features are those at 20-40 *ms* timescales (see a review in Engel *et al.* 1992), including the well-known 40 *Hz* oscillations in cat visual cortex (Gray and Singer 1988), those correlations are too wide to infer direct synaptic connections between cells. Here we will concentrate only on the very rarest and narrowest features found, dubbed "towers" by Nelson *et al.* (1992), which presumably indicate direct synaptic connections between two cells or between each of them and a common driver. Those cross-correlation studies in cat visual cortex (by and reviewed by Nelson *et al.* 1992) find that about 10% of cell pairs studied across areas of cortex share significiant spiking synchrony at zero mean time and 3 ms peak width, indicating that the cells share some common source of input; and the orientation sensitivity of the correlations suggests a cortical origin. But almost never (e.g., in only one of twenty such correlated pairs) did they find the *source* of such input, as indicated by a narrow CCH peak displaced from zero time. Why?

We will not consider in detail how one can so frequently find correlated cell pairs among millions of neurons picked essentially at random. But as one example of the problem consider that if 10% of cell pairs synchronize spikes at the 1% level (e.g., $C_c \approx 0.01$, as Nelson *et al.* suggest), then a solitary reference spike is likely to coincide with about .001 of a spike in any single other cell (above random chance). But that single other cell is drawn from a huge population, possibly 100,000— for example a couple mm^2 of cortex, with at least $40,000 - 100,000$ neurons/mm^2 (Abeles 1990). So in the population at large, the total number of spikes coincident with a reference spike—i.e., the spike's degeneracy or multiplicity m (Appendix E)—can be well over 100. While hurried experimenters are unlikely to ever find those 100 needles-in-a-haystack, that high inferred synchrony suggests two conclusions: 1) the number of independent spiking events in cortex might be a hundredfold smaller than in a random distribution, severly limiting the usefulness of averaging over populations of neurons (Britten *et al.* 1992); and 2) there exists some synchrony which might serve as input to a coincidence-detector, and which might in fact be the *output* of a coincidence-detector.

We will not consider the wider types of CCH peaks (30-200 *ms*) Nelson *et al.* examine; the focus here is on CCH peaks so narrow (3 *ms* typical) that they unambiguously indicate direct synaptic connections. We will construct predictions for such CCH peaks based on two caricature neural models sharing common sources of input. In all these cases we will assume that all cells fire at the same rate, which is in practice true only of averages over many recorded cells (although **not** of any particular pair).

First the integrator-model. Suppose we have two recorded cells, A and B, each of which temporally integrates its excitatory inputs (without significant leak). The cells each receive excitatory synapses from N_{com} "common" source cells (labelled by C_i, e.g., a pair ($N_{com} = 2$) of source cells would be C_1 and C_2). And each integrator receives an additional amount of uncorrelated, random input from other sources (the nature of which will prove to be unimportant).

In order to describe the total influence of common input, independent of the number of connections (N_{com}) it is distributed over, we can designate the relative strength S of those *combined* common connections by the total depolarization ΔV_c (out of V_{th}) they would all cause together. S is related to C_c and N_{com} as follows. S denotes the total fraction of each integrator's depolarizing input which is shared with the other integrator, $0 \leq S \leq 1$. A single one of those synaptic events (from one of the C_i) would raise a cell's potential toward firing by a smaller amount, $\Delta V/N_{com}$. Because the cell requires a depolarization of V_{th} to fire, a single event causes a relative potential change which depends on both the strength of common input and on the number of input synapses over which that input is distributed:

$$\frac{S}{N_{com}} = \frac{\Delta V_c / N_{com}}{V_{th}} \tag{F.1}$$

Since all other input is random, and independent events have independent influences, we can assume that the cell's potential is equally likely to be anywhere between zero and V_{th}, so that a single event from the common input has probability S/N_{com} of firing a recorded cell. The higher the number of cells (N_{com}) the common input S is distributed over, the smaller each individual synaptic event will be.

What fraction of output spikes will A and B have in common above chance, i.e., what will be the contribution coefficient $C_c(AB)$? (Because A and B fire at the same rate, $C_c(AB) = C_c(BA)$.) Suppose A fires a spike. The chance that its spike came from one of the common inputs is S (N_{com} common inputs with strength S/N_{com} each). In that case, the further chance that B fires in response to *the same particular event* which fired A is S/N_{com}, so the contribution coefficient is

$$C_c(AB) = \frac{S^2}{N_{com}} \tag{F.2}$$

Clearly, strong cross-correlations require strong connections from the common input source, and are more effective when the common input is concentrated in a single source rather than over many parallel, independent sources. For example, the strongest cross-correlation peaks (such as $C_c(AB) \approx 0.3$ within the same orientation column; Toyama *et al.* 1981) require that either a *single* neuron C be driving both A and B with connections effective enough to bring either one 2/3 of the way from rest to threshold, or that three common neurons

(C_1, C_2, C_3) be so strong that any one of the three can fire A or B with a single synaptic connection (neither of these cases agrees with known connectivities or synaptic strengths). Even the weaker cross-correlations (e.g., $C_c(AB) \approx 0.02$; Nelson *et al.* 1992, Figs. 1A and 6A) between cortical areas A17 and A18 of cat would require a single common cell driving both A and B with $S \approx 0.15$, corresponding to an EPSP strength of $3 - 4\ mV$ (still much larger than observed).

This difficulty was identified by Nelson *et al.* (1992), who qualitatively postulated a sub-population of very strong, very rare common-input cells which cause A and B to synchronize (they implicitly assumed the integrator model, but did not calculate the strength of connections required). The connections must be very strong because a common driver C_i is less than perfectly correlated with either A's or B's firing ($S < 1$); but A and B are linked to each other only *through* C_i, so that their correlation is even weaker than the correlation of either one with C_i.

The driver's contribution coefficient $C_c(CA)$ (equal to $C_c(CB)$) is just the chance that a firing by cell C_i causes a firing in A, i.e.,

$$C_c(CA) \quad = \quad S/N_{com} \tag{F.3}$$

From this we find that the driving cell C_i is always *better* correlated with A or B than A or B are ever correlated with each other,

$$\begin{aligned} \frac{C_c(CA)}{C_c(AB)} &= \frac{S/N_{com}}{S^2/N_{com}} && \text{(F.4)} \\ &= S^{-1} > 1 && \text{(F.5)} \end{aligned}$$

This is a firm prediction about the correlations between intergrator-type cells: the cross-correlation peaks which are offset from zero (i.e., those attributed to direct synaptic connections) should be much stronger than those centered at zero. This is not true for the research which shows quantitative CCHs (Toyama *et al.* 1981; Neslon *et al.* 1992), but unfortunately those papers have usually not calculated $C_c(AB)$, nor have they found enough offset peaks on which to base any conclusions.

(That rarity of direct connections is a mystery in itself, but is hypothetically (although implausibly) accounted for by the existence of driving cells C_i which are not only strong but also extremely rare, so that one hardly ever records from them (Nelson *et al.*, 1992)—even though they ostensibly account for a very large portion of cortical activity.)

Can a simple coincidence-detector model for cortical cells explain the cross-correlation peaks any better? Yes. Let us take a super-simplified model, in which the soma acts as a coincidence-detector among all synapses equally, regardless of their dendritic location (this model is explicitly unrealistic, intended for conceptual illustration only).

Suppose that A and B are coincidence-detecting neurons. Each receives input from a large number of independent synapses N, and fires only when M ($M \geq 2$) of them are coincident within some narrow time window (which is chosen so that A and B each fire at the same average rate as each input synapse fires). The number of shared synapses from the driving cells (the C_i) is SN, so that S denotes the fractional strength of common input relative to total input for each cell, as before (we will let $M \ll SN$ for convenience).

What is the above-random chance that A and B fire in synchrony, i.e., what is $C_c(AB)$? If A fires a spike, that means that M input pulses arrived within a time window, so we know that those M pulses were distributed somehow over the N synaptic sites of A. The chance that all M of them were among the SN sites common to B is about S^M. If the M pulses in fact *were* all among the common synapses, then neuron B must fire as well. But B could also fire if, for some integteger $i < M$, $(M - i)$ synaptic sites fired among the common input, and were coincident with i *non*-common sites firing. If we call P_{non} the probability that one of the $(1 - S)N$ *non*-common sites fires within a coincidence-widow, then the total probability that cell B fires in synchrony with cell A will be the contribution coefficient

$$C_c(AB) \;=\; \sum_{i=0}^{M-1} 2^{1-\delta_{0,i}} \binom{M}{i} S^{M-i} P_{non}^{i} \qquad \text{(F.6)}$$

The factor of $2^{1-\delta_{0,i}}$ comes from the chance of overlapping two independent pulses of non-zero width. The highest term ($i = M$), which is *not* included in this sum, is the baseline probability that two spikes will occur with no causal relationship; that term is also left out of the definition of C_c, eq. D.2, to make C_c represent an above-random probability. Note that this prediction does *not* depend on whether the common input is distributed over many input lines, and thus does not require a few very strong driver neurons, as the integrator-model does. Clearly, large M dramatically decreases the correlations between cells sharing common input (roughly as S^M), as the chance of finding enough coincidences in the common input becomes very small.

This scheme more easily accounts for strong common-input correlations. As a lower bound, we can assume that $P_{non} \approx 0$, so that only the S_M term contributes. In that case, for two neurons sharing fiteeen percent of their inputs ($S = 0.15$) and firing upon the coincidence of two, this yields $C_c(AB) > 0.02$, in the range found for correlations between cortical areas. (Even higher values, such as $C_c(AB) \approx 0.3$, are consistent with the known high connection overlap among neurons in the same orientation column).

A slightly more accurate approach would coarsely estimate P_{non} from some typical assumptions. For example, suppose the chance that any one (not each one) of the cell's inputs fires within a coincidence-window is P_1, and the chance that the cell itself fires is P_{out}. Then for a one-millisecond coincidence-window and an output firing rate of 50 Hz,

$$\begin{aligned} P_1^M &\approx P_{out} && \text{(F.7)} \\ P_1 &\approx P_{out}^{1/M} && \text{(F.8)} \\ &\approx (50Hz \times 1ms)^{1/M} && \text{(F.9)} \\ &= 0.02^{1/M} && \text{(F.10)} \end{aligned}$$

The chance that one of the non-common inputs fires is scaled down from this estimate by $(1 - S)$, so that we have

$$C_c(AB) = \sum_{i=0}^{M-1} 2^{1-\delta_{0,i}} \binom{M}{i} S^{M-i} (0.02)^{i/M} (1-S)^i \qquad \text{(F.11)}$$

For $S = 0.15$ and $M = 2$ (as above), this new estimate for C_c is about triple the previous estimate, being 0.058 instead of 0.02. This degree of correlation can occur because every input C_i (out of possibly hundreds or thousands) is capable of gating A or B's firing, so that C_i's instantaneous influence is high (while its average influence is low, due to an absence of temporal integration).

But what about the problematic correlation of A (or B) with its *driving* cells? Each driving cell has an equal chance of firing A, and each C_i and A fire at the same rate. So every time A fires, it does so in response to the firing of M of the N drivers. That means that each C_i has an average chance M/N of firing A, so that its contribution coefficient is

$$C_c(CA) = \frac{M}{N} \tag{F.12}$$

If M is typically between 2 and 5 (a reasonable range for the number of coincidences required to fire a cell, as in Chapter 3), and N is at least several hundred, then the connection between a driver C_i and a cell A would be very weak ($C_c(CA) \leq .01$). The crucial point is that in this coincidence-detector model with many inputs, the inevitable correlation between driver and cell is expected to be much *weaker* than the correlation between two cells sharing common input, and might be easily lost in the noise, or in the tail of a stronger central peak (as might occur if the driver and reciever cells happen to share common input). On the other hand, the integrator model predicts that direct connections should appear *stronger* on a CCH; there is no way such strong peaks could be lost in the noise.

In summary: the coincidence-detector model can easily produce strong common-input peaks when common input is shared among many input lines; the integrator-model can only produce strong central peaks if a few super-strong neurons provide common input. Also, two coincidence-detectors sharing common input will typically exhibit much stronger correlations with each other than with their drivers (matching observations in cortex); in the integrator's case the opposite is true.

The fundamental difference between the two models is that a spike out of an integrator depends primarily on the timing of its *previous* output spike, and only secondarily on its instantaneous input (common or otherwise). So an integrator is difficult to synchronize for exactly the same reason that it is difficult to make fire irregularly: it has a strong individual memory of its previous firing, and other cells do not necessarily share that memory. Coincidence-detectors, on the other hand, have no memory at all, and respond only to their instantaneous input. When that input is shared among cells, their spikes reflect it directly, unaffected by history.

So narrow CCH peaks at zero time are much easier to explain using coincidence-detecting neurons. But if cortical cells are indeed coincidence-detectors, we may have to reexamine the traditional interpretation of how we infer cortical connectivity from cross-correlation histograms.

Appendix G

Compartmental Modelling

The simulation of the electrical behavior of single neurons has become so common and standardized that programs for it are available virtually "off-the-shelf." We used one of the most popular such programs, NEURON (Hines 1990), which is described below.

Compartmental modelling takes advantage of the fact that most neurons—including cortical cells—are composed of long, thin branches, filled with a fairly uniformly conductive saline solution and surrounded by a conductive, capacitive membrane. Each branch can be locally modelled first as a uniform cylinder, thus approximating the entire cell as a collection of linked cylinders of various dimensions and membrane properties. Furthermore, at the timescales of known neural operation (kHz and below), the cylinders' electrical properties have no radial dependence (only axial), so that each cylinder can be approximated by a single capacitor C_i in parallel with a linear resistor

$R_i = 1/G_i$ (and possibly in parallel also with one or more nonlinear resistors, whose conductances may depend on local voltage and voltage history). (See Figure F.2). These parallel combinations (representing the cell membrane) are connected together by the resistors representing the axial resistance through the intracellular saline fluid. Book-keeping routines in the simulation program ensure that the axial resistance is calculated directly from the fluid's bulk resistivity and the cylinder's dimensions, just as the parallel capacitor-resistor values are calculated from the cylinder's membrane area.

In the case of a passive branch (which contains no nonlinear resistors), the simulation merely adjusts the voltage on each cylinder section according to the local differential equations. For example, a new voltage on a section of a single cylinder i (with no branches) would be computed from the voltages on it and on its neighbors $i+1$ and $i-1$ at the previous time-step as

$$\begin{aligned} V_i(t+\Delta t) &= V_i(t) + \frac{\Delta t((V_{i+1}(t) - V_i(t))}{C_i R_{i,i+1}} + \frac{\Delta t((V_{i-1}(t) - V_i(t))}{C_i R_{i,i-1}} \\ & - \frac{(V_i - V_{ext})G_i}{C_i} \end{aligned} \tag{G.1}$$

The nonlinear membrane properties are numerically integrated in the same way. But these equations are more complicated, depending on the voltage history through other intermediate variables. For example, the most important nonlinear membrane conductances, the Hodgkin-Huxley-like spiking conductances used here, are modelled as depending at any point in time on the values of parameters m and h (thought to represent the concentrations of ionic species which instantaneously adjust pores in the membrane). The two separate conductances (for sodium and potassium) each have a parameter denoting the

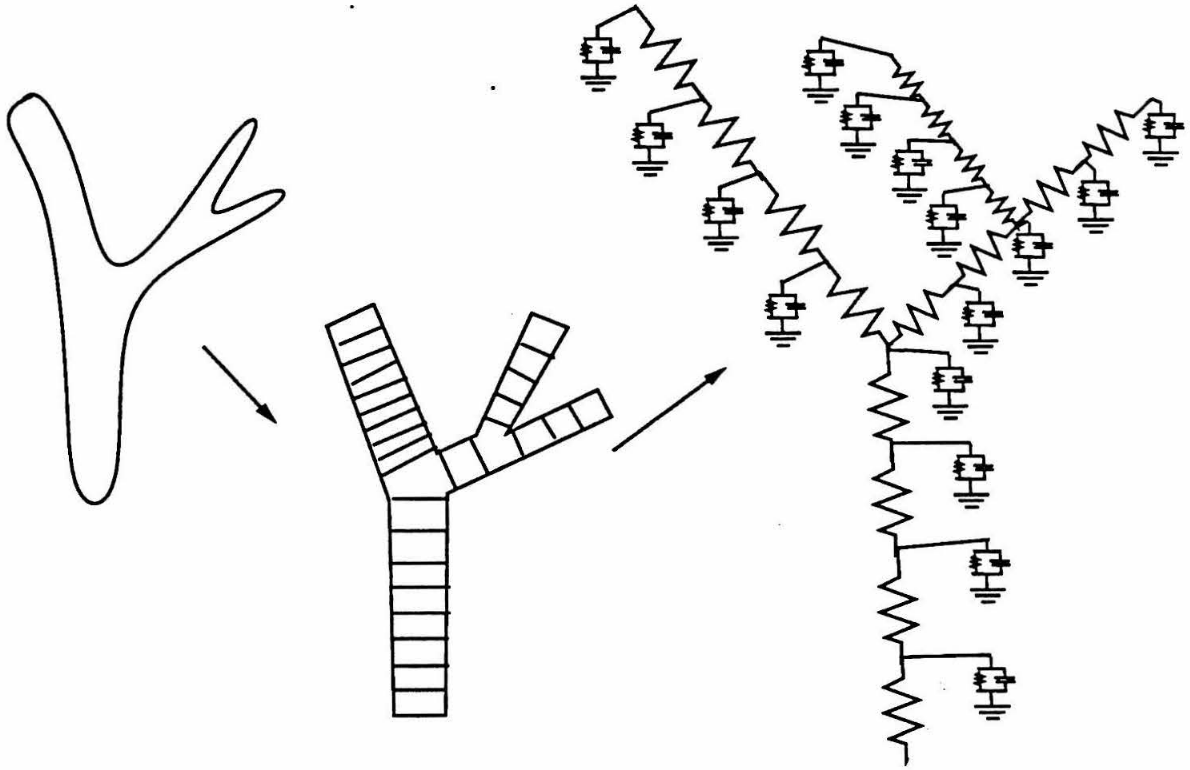

Figure G.1: **Compartmental Simulation.** A cell's dendrite is a branched tube of salt water (left), which can be electrically approximated as a collection of connected cylinders (center), and further approximated as a collection of one-dimensional circuit elements (right), whose interactions are then numerically simulated.

maximum possible conductance (g_{peak}), a driving voltage (E_{rev}), and the dependence on history through m and h:

$$g_{Na} = g_{Na,peak}(V - E_{Na})m_{Na}^2 h \tag{G.2}$$

$$g_K = g_{K,peak}(V - E_K)m_K^2 \tag{G.3}$$

The m and h values reflect the voltage history as if they were physical particles, flowing into and out of the cell according to the dynamical equations

$$\frac{dm_{Na}}{dt} = \frac{m_{Na,ss} - m_{Na}}{\tau_{Na,m}} \tag{G.4}$$

$$m_{Na,ss} = \frac{1}{1 + \exp\left(\frac{V - V_{1/2,Na,m}}{\Theta_{Na,m}}\right)} \tag{G.5}$$

$$\tau_{Na,m} = 0.05\ ms \tag{G.6}$$

$$\Theta_{Na,m} = -3\ mV \tag{G.7}$$

$$V_{1/2,Na,m} = -40\ mV \tag{G.8}$$

$$\frac{dm_K}{dt} = \frac{m_{K,ss} - m_K}{\tau_{K,m}} \tag{G.9}$$

$$m_{K,ss} = \frac{1}{1 + \exp\left(\frac{V - V_{1/2,K,m}}{\Theta_{K,m}}\right)} \tag{G.10}$$

$$\tau_{K,m} = 2\ ms \tag{G.11}$$

$$\Theta_{K,m} = -3\ mV \tag{G.12}$$

$$V_{1/2,K,m} = -40\ mV \tag{G.13}$$

$$\frac{dh}{dt} = \frac{h_{ss} - h}{\tau} \tag{G.14}$$

$$h_{ss} = \frac{1}{1 + \exp\left(\frac{V - V_{1/2,h}}{\Theta_h}\right)} \tag{G.15}$$

$$\tau_h = 0.5\ ms \tag{G.16}$$

$$\Theta_m = 3\ mV \tag{G.17}$$

$$V_{1/2,h} = -45\ mV \tag{G.18}$$

These equations are similar to the original Hodgkin-Huxley equations, differing in the order of m used (m^3 in the original H-H model, but here m^2), and in the absence of any voltage-dependence in τ. But both sets of equations are phenomenological models which can be adjusted to fit recordings from real neurons, and the differences between these two models pale in comparison to our uncertain knowledge of their parameters for use in cortical cells.

Appendix H

Somatic Repolarization by Dendritic Spiking

Here we will estimate the ability of I_{DR} to repolarize the soma after a dendritic spike. To simplify this task let us only consider the current-source model (eq. 3.33) of dendritic spiking.

As we saw in the current-source approximation above (section 3.4), the sodium current localizes its activity in the dendrite so that it provides a current to the soma which is roughly independent of somatic potential or intervening trunk resistance. The potassium current has no such choice of position, being activated in approximately the same physical place as the sodium currents (because the two conductances have similar threshold voltages). So let us think of the sodium current as arising from a resistance R_+, composed of both dendritic trunk and part of the terminal branch, which (briefly) connects the soma to the sodium reversal potential (Figure H.1) and delivers current i_{cs} (eq. 3.33):

$$R_+ = \frac{E_{Na} - E_{rest}}{i_{cs}} \tag{H.1}$$

Because we pretend that the sodium and potassium currents do not overlap in time and are small, we can sum their contributions separately and ignore their interactions (Figure 3.3, Figure 3.10). If the sodium and potassium conductances are exactly equal and identically distributed on the terminal branch, then the axial potassium current to the soma—i_K—uses the same R_+ as sodium to connect the soma with E_K, and

$$i_K = \frac{E_K - E_{rest}}{R_+} \tag{H.2}$$

But what if the peak potassium conductance is adjusted to a value different from the sodium conductance? We can add a correction term ΔR to R_+, so that $\Delta R = 0$ when the two conductances are equal, but ΔR deviates from zero as the potassium conductance deviates from the sodium conductance. The magnitude of that deviation can be seen from eq. 3.33 (containing the input resistance of an infinite cable) to scale with the square root of the potassium conductance, and is zeroed by the sodium conductance:

$$\Delta R = \sqrt{\frac{r_a}{\overline{G}_K \pi d}} - \sqrt{\frac{r_a}{\overline{G}_{Na} \pi d}} \tag{H.3}$$

So that

$$i_K = \frac{E_K - E_{rest}}{R_+ + \Delta R} \tag{H.4}$$

The current-source approximation allows us to compare the potassium current to the sodium current by a ratio which does not depend on the absolute con-

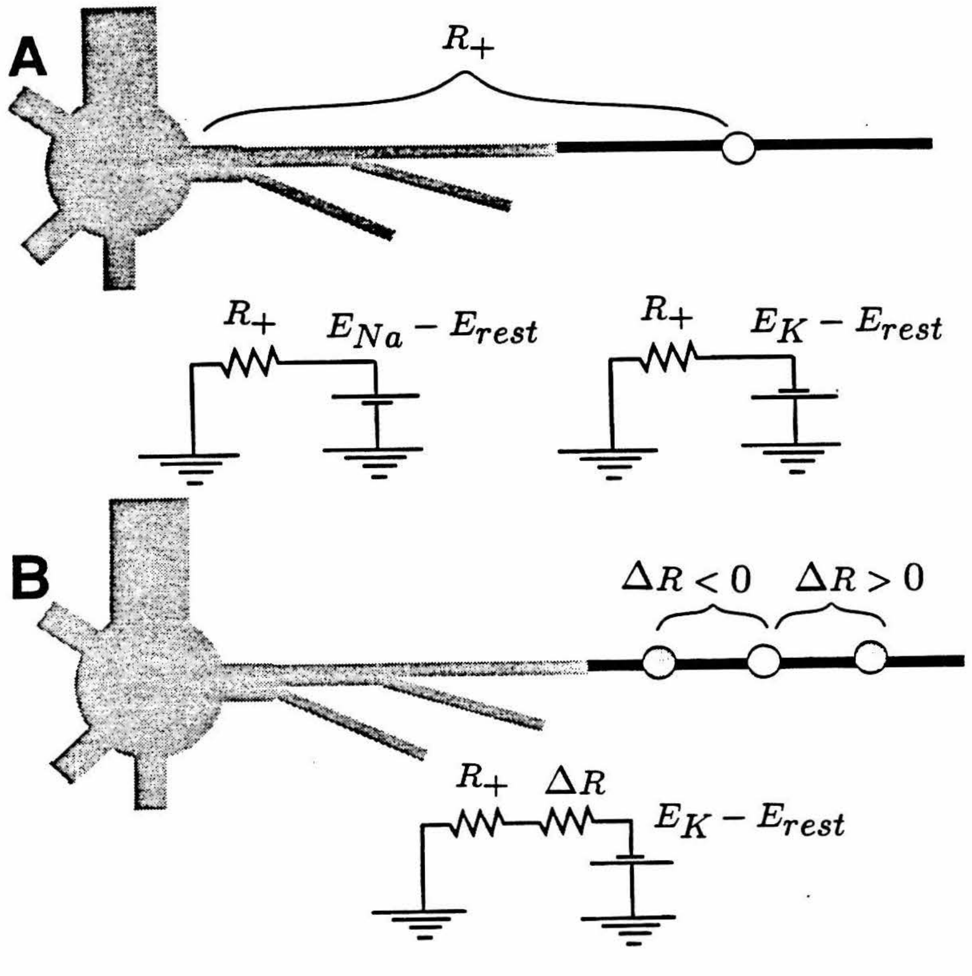

Figure H.1: (A) The sodium conductance in a spiking terminal branch (black) can be thought of as briefly creating an effective resistance R_+ between E_{rest} and E_{Na}. If the subsequent potassium conductance has the same magnitude and spatial distribution, then it has an identical circuit, using E_K in place of E_{Na}. (B) If the potassium conductance is stronger or weaker than the sodium conductance, then the above circuit can be modified by a correction term ΔR.

ductance values or on the diameter of the terminal branch, using eqs. H.1-H.4:

$$\begin{aligned} \frac{i_K}{i_{cs}} &= \frac{1}{i_{cs}}\left(\frac{E_K - E_{rest}}{R_+ + \Delta R}\right) && \text{(H.5)} \\ &= \frac{1}{i_{cs}}\left(\frac{(E_K - E_{rest})}{\frac{E_{Na}-E_{rest}}{i_{cs}} + \Delta R}\right) && \text{(H.6)} \\ &= \frac{E_K - E_{rest}}{E_{Na} - E_{rest} + i_{cs}\Delta R} && \text{(H.7)} \end{aligned}$$

where

$$\begin{aligned} i_{cs}\Delta R &= \left(\sqrt{\frac{8R_i}{\overline{G}_K\pi^2 d^3}} - \sqrt{\frac{8R_i}{\overline{G}_{Na}\pi^2 d^3}}\right)(E_{Na} - V_{1/2})\sqrt{\frac{\overline{G}_{Na}\pi^2 d^3}{8R_i}} && \text{(H.8)} \\ &= (E_{Na} - V_{1/2})\left(\sqrt{\frac{\overline{G}_{Na}}{\overline{G}_K}} - 1\right) && \text{(H.9)} \end{aligned}$$

The most important variable for somatic spike triggering is not current but depolarization. In particular, the repolarization by I_K relative to the depolarization by I_{Na} will be given by the ratio of charges deposited at the soma:

$$\begin{aligned} Q_K &= \int i_K(t)dt \approx \frac{i_K\tau(K)}{2} && \text{(H.10)} \\ \text{So } \frac{Q_K}{Q_{Na}} &= \frac{i_K\tau(K)}{i_{cs}\tau(h)} && \text{(H.11)} \end{aligned}$$

When this ratio is unity, we can say that the potassium current has removed the depolarization caused by the sodium currents. Because the persistent somatic depolarization is the integral of these two opposing currents, we want to compare the voltage after the spiking event to the voltage which would persist in the absence of repolarizing currents (Figure 3.10). A convenient measure of that quantity is the persisitent somatic depolarization at time t_a *after* the dendritic spike has occurred,

$$\Delta V(t_a) = V_{soma}(t_a) - E_{rest}. \qquad \text{(H.12)}$$

In the near-absence of any potassium currents, $\Delta V(t_a)$ would have a maximum value of $\Delta V_{Na}(t_a)$ (which is less than the peak ΔV_{soma} because charge equilibrates in the cell after the peak depolarization occurs). We can use these measures of persistent depolarization to define a dimensionless ratio Δ_p:

$$\begin{aligned}
\Delta_p &\equiv \frac{\Delta V(t_a)}{\Delta V_{Na}(t_a)} && \text{(H.13)}\\
&= \frac{Q_K + Q_{Na}}{Q_{Na}} && \text{(H.14)}\\
&= 1 + \frac{Q_K}{Q_{Na}} && \text{(H.15)}\\
\Delta_p &= 1 + \left(\frac{\tau(K)}{\tau(h)}\right) \frac{E_K - E_{rest}}{E_{Na} - E_{rest} + \left(\sqrt{\frac{\overline{G}_{Na}}{\overline{G}_K}} - 1\right)(E_{Na} - V_{1/2})} && \text{(H.16)}
\end{aligned}$$

This peristent depolarization Δ_p nears unity in the near-absence of potassium currents, crosses zero when the dendritic spike has no lasting somatic polarization, and dips below zero if the soma is left at a *lower* potential than before the dendritic spike. The soma is more effectively repolarized when the resting potential is higher (further from E_K), and when the potassium conductance is stronger. (All other terms, such as reversal potentials and time-constants, are presumably fixed by properties of the individual channels or extracellular fluid, so they were not varied.)

This coarse estimate (eq. H.16) compares well with simulations of a representative terminal branch for values of $\overline{G}_K$ spanning a factor of one hundred ($.05\ Scm^{-2}$–$5.0\ Scm^{-2}$) at two different reversal potentials (-75 and $-65\ mV$, Figure 3.10); simulations on other branches gave almost identical results. The only problem is that when $\overline{G}_K = 0$, the sodium channels in the dendrite would "latch on" and fire repeatedly—an undesirable situation. The lowest potassium conductance for which the sodium channels did not "latch on" was $\overline{G}_K = .05\ Scm^{-2}$, so this value simulated the "no-potassium" depolarization $\Delta V_{Na}(t_a)$. In simulations, the somatic potential dropped to about half its peak value during the $3 - 4\ ms$ it took for the charge to equilibrate in the cell (Figure 3.10 A,B), after which time the somatic depolarization decayed with the membrane time constant. The particular choice of measurement time ($t_a = 8\ ms$) did not significantly affect Δ_p.

The predicted values of persistent depolarization Δ_p for various values of $\overline{G}_K$ and E_{rest} agree qualititively with the simulations, showing a sharp drop in persistent depolarization for small $\overline{G}_K$ and a saturated minimum value of $\Delta_p \approx 0$ for $\overline{G}_K \gg \overline{G}_{Na}$ (Figure 3.10). The increasing deviation of the predicted curve from the simulated ones at these high-$\overline{G}_K$ values occurs in part because of the poor model for $g_K(t)$ (Figure 3.3), and in part because the prediction subtracts two opposing approximations (i_{cs} and i_K), whose inaccuracies still add. Because there are *no* free parameters, it is a bit surprising that this highly simplified model worked even as well as this at accounting for these highly nonlinear repolarizations over different dendrites and parameter values.

These expressions and simulations suggest that a peak potassium conductance about twice the peak sodium conductance will leave virtually no persistent somatic depolarization after a dendritic spike for $E_{rest} = -65\ mV$ (Figure 3.10). This ratio of conductances was used throughout Chapter 3, unless otherwise specified.

Appendix I

Relation Between Spike-number Variance and C_V

In Chapter 2 we tried to measure the variability of a spike-generating process. There are two easy methods of doing that; here we will find the relation between those two measures.

The measure we concentrated upon was looking at the variability in *time* between successive intervals, through the normalized width (C_V) of an interspike-interval histogram. This method can work for trains of arbitrary length, but all trains must have the same stationary firing rate.

The other method is the variability in *number* of spikes occurring in a given time period, through the normalized variance in the number of spikes S per trial (σ_S^2/S). This method requires that all trials have the same duration, but can accomodate firing rates which change with time (as long as that change is the same within each trial). Both measures yield values of unity for Poisson spike trains and zero for perfectly regular trains. But how are the two measures related?

Let us evaluate the simplest possible case: we generate spikes with stationary rate and C_V. From this process we will construct two distinct batches of spike trains: one batch of trains with a fixed duration but a variable number of spikes (the usual experimental protocol), and a second batch of trains with a fixed number of spikes and variable duration. By calculating each sort of variability separately, we can find their relationship.

First the batch with a large but fixed number of spikes S per train. We can treat each separate ISI Δt_i in each train as an independent random variable with standard deviation $\sigma_{\Delta t}$ and mean $\overline{\Delta t}$. The duration of each whole train is different, fluctuating randomly about the mean T with standard deviation σ_T. If we assume that all random variables are Gaussian-distributed, then

$$\begin{aligned} T \pm \sigma_T &= \sum_{i=1}^{S-1} \Delta t_i && \text{(I.1)} \\ &\approx \sum_{i=1}^{S} (\overline{\Delta t} \pm \sigma_{\Delta t}) && \text{(I.2)} \\ &= S\overline{\Delta t} \pm \sigma_{\Delta t}\sqrt{S} && \text{(I.3)} \end{aligned}$$

We can also interpret this variability in time as being a variability in firing rate about a mean value R, so that

$$\begin{aligned} R &= \frac{S}{T} && \text{(I.4)} \\ R \pm \sigma_R &= \frac{S}{T \pm \sigma_T} && \text{(I.5)} \end{aligned}$$

$$\approx \frac{S}{T}\left(1 \pm \frac{\sigma_T}{T}\right) \tag{I.6}$$

The other batch of spike trains has fixed duration T (the same as the mean duration of the first batch), but the number of spikes per train varies as $S \pm \sigma_S$. This variability can be gotten directly from the variability in firing rate,

$$\begin{aligned} R \pm \sigma_R &= T^{-1}(S \pm \sigma_S) && \text{(I.7)} \\ &= \frac{S}{T}\left(1 \pm \frac{\sigma_S}{S}\right) && \text{(I.8)} \end{aligned}$$

Equations I.6 and I.8 together show that the relative variability in spike number is equal to the relative variability in train duration (as in eq. B.1),

$$\begin{aligned} \frac{\sigma_S}{S} &= \frac{\sigma_T}{T} && \text{(I.9)} \\ &= \frac{\sqrt{S}\sigma_{\Delta t}}{S\overline{\Delta t}} && \text{(I.10)} \end{aligned}$$

So the standard deviation in spike number will be

$$\begin{aligned} \sigma_S &\approx \frac{\sqrt{S}\sigma_{\Delta t}}{\overline{\Delta t}} && \text{(I.11)} \\ &= C_V\sqrt{S}, && \text{(I.12)} \end{aligned}$$

using the definition of C_V (eq. 2.4). So the normalized variance will be given by the square of C_V:

$$\frac{\sigma_S^2}{S} = C_V^2 \tag{I.13}$$

This formula only applies for a train at constant rate, for which C_V can be calculated directly from the ISI histogram. But if we apply it to a train with variable rate (such as real neural data), we find that the C_V values this formula yields are in the range of multiple C_V values found by the multi-histogram method. For instance, the normalized variance is about 0.02 for the fastest "barely plausible" simulation, and about 0.005 for the fastest "conventional" simulation. The above formula yields $C_V = 0.14$ and $C_V = 0.07$ respectivey, while the multi-histogram analysis gives $C_V = 0.1 - 0.3$ and $C_V = 0.03 - 0.2$ (Figure 2.14). So even when significant temporal structure in the average spike rate (or PSTH) makes multi-histogram analysis unreliable, the variance in spike number can still indicate the intrinsic variability of the neuron's firing.

Appendix J

References

Abeles, M. (1980) "Local Cortical Circuits," Springer-Verlag, New York.

Abeles, M. (1982) "Role of Cortical Neuron: Integrator or Coincidence Detector?" *Israeli Journal of Medical Sciences* **18**, 83-92.

Abeles, M. (1990) "Corticonics," Cambridge University Press, Cambridge.

Adams, P. (1992) "The Platonic Neuron Gets the Hots," *Current Biology* **7**, 625-627.

Aertsen, A., Gerstein, G., Habib, M., & Palm, G. (1989) "Dynamics of Neuronal Firing Correlation: Modulation of 'Effective Connectivity,' " J. Neurophysiol. **61**, 900-917.

Agmon, A. & Connors, B. (1992) "Correlation between Intrinsic Firing Patters and Thalamocortical Responses of Neurons in Mouse Barrel Cortex," *J.*

Neurosci. **12**, 319-329.

Amitai, Y., Firedman, A., Connors, B., & Gutnick, M. (1992) "Regenerative Electrical Activity in Apical Dendrites of Pyramidal Cells in Neocortex," submitted to *J. Neurophysiol.*

Angelini, F., Barbi, M., Chillemi, S., & Petracchi, D. (1982) "Operational Models of Neural Encoding," in *Biomathematics in 1980*, ed. L. M. Ricciardi and A. C. Scott, North-Holland.

Bair, W., Koch, C., Newsome, W. & Britten, K. (1992) "Power Spectrum Analysis of MT Neurons in the Behaving Monkey," in preparation.

Bernander, Ö., Douglas, R., Martin, K., & Koch, C. (1991) "Synaptic Background Activity Determines Spatio-Temporal Integration in Single Pyramidal Cells," Proc. Natl. Acad. Sci., USA **88**, 1569-1573

Bernander. Ö., Douglas, R., Koch, C.,& Niebur, E. (1991) "Network Activation Level Influences Single-Cell Properties," submitted to *Neuroscience Abstracts.*

Britten, K., Shadlen, M., Newsome, W., & Movshon, A. (1992) "The Analysis of Visual Motion: A Comparison of Neuronal and Psychophysical Performance," *J. Neurosci.* **12**, 4745-4765.

Bryant, H. & Segundo, J. (1976) "Spike Initiation by Transmembrane Current: A White-Noise Analysis," *J. Physiol.* **260**, 279-314.

Bugmann, G. (1990) "Irregularity of Natural Spike Trains Simulated by an Integrate-and-Fire Neuron," Extended Abstracts, 3rd Int'l Symposium on Biolelectronic and Molecular Electronic Devices, Kobe, 105-105.

Bugmann, G. (1991) "Neural Information Carried by One Spike," proc. 2nd Australian Conf. on Neural Networks (ACNN'91), Sydney, 235-238.

Burns, B.D. & Webb, A.C. (1976) "The Spontaneous Activity of Neurones in the Cat's Visual Cortex," *Proc. R. Soc. Lond. B* **194**, 211-223.

Burr, D.C. (1979) "Acuity for Apparent Vernier Offset," *Vision Res.* **19**, 835-837.

Bush, R. & Douglas, P. (1991) "Synchronization of Bursting Action Potential Discharge in a Model Network of Neocortical Neurons," *Neural Compuatation* **3**, 19-30.

Calvin, W. & Stevens, C. (1968) "Synaptic Noise and Other Sources of Randomness in Motoneuron Interspike Intervals," *J. Neurophysiol.* **31**, 574-587.

Cinlar, (1972) "Superposition of Point Processes," in *Stochastic Point Processes*, ed. P.A.W. Lewis, 549-606, John Wiley, New York.

Correia, M. J., & Landolt, J. P. (1977) "A Point Process Analysis of the Spontaneous Activity of Anterior Semicircular Canal Units in the Anesthetized Pigeon," *Biol. Cybernetics* **27**, 199-213.

Covey, E. & Casseday, J.H. (1991) "The Monaural Necleu of the Lateral Lemniscus in an Echolocating Bat—Parallel Pathways for Analyzing Temporal Features of Sound," *J. Neurosci.* **11**, 3456-3470.

Creutzfeld, O., Lux, H., & Nacimiento, A. (1964) "Intrazelluläre Reizung Corticalen Nervenzellen," *Pflügers Archiv* **281**, 129-151.

Crick, F. & Koch, C. (1990) "Toward a Neurobiological Theory of Consciousness," *Seminars in the Neurosciences* **2**, 263-275.

Douglas, R., Martin, K., & Whitteridge D. (1991) "An Intracelluar Analysis of the Visual Responses of Neurones in Cat Visual Cortex," *J. Physiol.* **440**, 659-696.

Douglas, R. & Martin, K. (1991) "A Functional Microcircuit for Cat Visual Cortex," *J. Physiol.* **440**, 735-769.

Douglas, R. & Martin, K. (1991) "Opening the Grey Box," *TINS* **14**, 286-293.

Engel, A., Koenig, P. Kreiter, A., Schillen, T. and Singer, W. (1992) "Temporal Coding in the Visual Cortex: New Vistas on Integration in the Nervous System," *TINS* **15**.

Ferster, D. (1986) "Orientation Selectivity of Synaptic Potentials in Neurons of Cat Primary Visual Cortex", *J. Neurosci.* **6**, 1284-1301.

Ferster, D. (1987) "Origin of Orientation-Selective EPSPs in Simple Cells of Cat Visual Cortex," *J. Neurosci.* **7**, 1780-1791.

Ferster, D. (1988) "Spatially Opponent Excitation and Inhibition in Simple

Cells of the Cat Visual Cortex," *J. Neurosci.* **8** 1172-1180.

Fetz, E., Toyama, K., & Smith, W. (1991) "Synaptic Interactions between Cortical Neurons," in *Cerebral Cortex* **9**, ed. A. Peters, Plenum Publishing.

Freeman, W. & van Dijk, B. (1987) "Spatial Patterns of Visual Cortex Fast EEG Response During Conditioned Reflex in a Rhesus Monkey," *Brain Res.* **422**, 267-276.

Frostig, R. D., Frostig, Z., Frysinger, R.C., & Schechtman, V.L. (1985) "Multineuron Analysis Reveals Complex Patterns of Interaction Among Neurons in Forebrain Networks and Cardiorespiratory Parameters During Sleep-Waking States," *Soc. Neurosci. Abstr.* **11**, 1020.

Gray, C. M., and Singer, W. (1988) "Stimulus-Specific Neuronal Oscillations in Orientation Columns of Cat Visual Cortex," *Neuroscience* **22**

Hebb, D. O. (1949) "The Organization of Behavior," Wiley, New York.

Heggelund, P. & Albus, K. (1978) "Response Variability and Orientation in Single Cells in Striate Cortex of Cat," *Exp. Brain Res.* **32**, 197-211.

Hestrin, S., Sah, P., & Nicoll, R. (1990) "Mechanisms Generating the Time Course of Dual Component Excitatory Synaptic Currents Recorded in Hippocampal Slices," *Neuron* **5**, 247-253.

Hines, M. (1989) "A Program for Simulation of Nerve Equations with Branching Geometries,"*Int. J. Biomed. Comput.* **24**, 55-68.

Hinton, G. & Sejnowski, T. (1986) "Learning and Relearning in Boltzmann Machines," in *Parallel Distributed Processing* **1**, ed. Rumelhart & McClelland, MIT Press, Cambridge, Massachussetts.

Horn, D., Sagi, D. & Usher, M. (1991) "Segmentation Binding and Illusory Conjunctions," *Neural Comp.* **3**, 510-525.

Hubel, D. H. & Wiesel, T. N. (1962) "Receptive Fields, Binocular Interaction, and Functional Architecture in the Cat's Visual Cortex," *J. Physiol.* **160**, 106-154.

Huguenard, J., Hamill, O., & Prince, D. (1988) "Sodium Channels in Dendrites of Rat Cortical Pyramidal Neurons," *Proc. Natl. Acad. Sci. USA*, **86**, 2473-2477.

Jack, J., Noble, D. & Tsien, R. (1983) "Electric Current Flow in Excitable Cells," Oxford University Press, Oxford.

Jaffe, D., Johnston, D., Lasserross, N., & Lisman, J. (1992) "The Spread of Na^+ Spikes Determines the Pattern of Dendritic Ca^{2+} Entry into Hippocampal Neurons," *Nature* **357**, 244-246.

Jaslove, S. (1992) "The Integrative Properties of Spiny Distal Dendrites," *Neuroscience* **47**, 495-519.

Jones, O.W., Kunze, D.J., & Angelides, K.J. (1989) "Localization and Mobility of Ω-cytotoxin Sensitive Ca^{++} Channels in Hippocampal CA1 Neurons," *Science* **244**, 1189-1191.

Kammen, D., Holmes, P., and Koch, C. (1989) "Cortical Oscillations: Feedback vs. Local Coupling," in *Models of Brain Function*, Cambridge Univ. Press.

Knierim, J. & Van Essen, D. (1992), "Neuronal Responses to Static Textural Patterns in Area V1 of the Alert Macaque Monkey," *J. Neurophysiol.* **67**, 961-980.

Knight, B. (1972), "Dynamics of Encoding in a Population of Neurons," *J. Gen. Phys.*, **59**, 734-766.

Koch, C. & Poggio, T. (1987) "Biophysics of Computation: Neurons, Synapses, and Membranes," in *Synaptic Function*, ed. G.M. Edelman, W.E. Gall, & W.M. Cohen, 637-694, John Wiley and Sons, New York.

Koch, C. and Poggio, T. (1992) "Multiplying with Synapses," in *Single Neuron Computation*, ed. McKenna T., Davis J., and Zornetzer S., Academic press, New York.

Koch, C. & Schuster, H. (1992) "A Simple Network Showing Burst Synchronization without Frequency-Locking," *Neural Computation* **4**, 211-223.

Koch, C. (1993) "Computational Approaches to Cognition: The Bottom-Up View," *Current Opinion in Biology*, in press.

Komatsu, Y., Nakajima, S., Toyama, K., & Fetz, E. (1988) "Intracortical Connectivity Revealed by Spike-Triggered Averaging in Slice Preparations of Cat Visual Cortex," *Brain Res.* **442**, 359-362.

Kuffler, S., Nicholls, J., & Martin, A. (1984) "From Neuron to Brain," 2nd ed., Sinauer Assoc., Sunderland, Massachussetts.

Lansky, P. & Smith, C. (1989) "The Effect of a Random Initial Value in Neural First-Passage-Time Models," *Mathematical Biosciences* **93**, 191-215.

Lansky, P. & Radil, T. (1987) "Statistical Inference on Spontaneous Neuronal Discharge Patterns," *Biol. Cybern.* **55**, 299-311.

Legendy, C. & Salcman, M. (1985) "Bursts and Recurrences of Bursts in the Spike Trains of Spontaneously Active Striate Cortex Neurons," J. Neurophysiol. **53**, 926-939.

Mason, A., Nicoli, A., & Stratford, K. (1991) "Synaptic Transmission between Individual Pyramidal Neurons of the Rat Visual Cortex *in vitro*," *J. Neurosci.* **11**, 72-84.

Mazzoni, P. Andersen, R., & Jordan, M. (1991) "A More Biologically Plausible Learning Rule for Neural Networks," *PNAS* **88**, 4433-4437.

McClurkin, J., Optican, L., Richmond, B., & Gawne, T. (1991) "Concurrent Processing and Complexity of Temporally Encoded Neuronal Messages in Visual Perception," *Science,* **253**, 675-677.

McNaughton, B., Barnes, C., & Andersen, P. (1981) "Synaptic Efficacy and EPSP Summation in Granule Cells of Rat Fascia Denta Studied in Vitro," *J. Neurophysiol.* **46**, 952-966.

Mel, B. & Koch, C. (1990) "Sigma-Pi Learning: On Radial Basis Functions and cortical Associative Learning," in *Advances in Neural Information and Processing Systems 2*, ed. D. Touretzky, Morgan Kaufmann, San Mateo, California.

Mel, B. (1992a) "NMDA-Based Pattern Discrimination in a Model Cortical Neuron," *Neural Computation* **4**, 502-517.

Mel, B. (1992b) "Information Processing in an Excitable Dendritic Tree," *CNS Memo* **17**, Computation and Neural Systems Program, Caliornia Institute of Technology, Pasadena, CA.

Michalski, A., Gerstein, G., & Tarnecki, R. (1983) "Interactions Between Cat Striate Cortex Neurons," *Exp. Brain Res.* **51**, 97-107.

Mikami, A., Newsome, W., & Wurtz, R. (1986) "Motion Selectivity in Macaque Visual Cortex: **I**. Mechanisms of Direction and Speed Selectivity in Extrastriate Area MT," *J. Neurophyiol.* **55**, 1308-1327.

Millner, P.M (1974) "A Model for Visual Shape Recognition," *Psychol. Review* **81**, 521-535.

Nelson, J., Salin, P. Munk, H.-J., Arzi, M., & Bullier, J. (1992) "Spatial and Temporal Coherence in Cortico-Cortical Connections: A Cross-Correlation Study in Areas 17 and 18 in the Cat," *Vis. Neurosci.* **9**, 21-37.

Newsome, W. & Pare, E. (1988) "A Selective Impairment of Motion Perception Following Lesions of the Middle Temporal Visual Area (MT)," *J. Neurosci.* **8**, 2201-2211.

Newsome, W., Britten, K., Movshon, J. A., & Shadlen, M. (1989a) "Single Neurons and the Perception of Motion," in *Neural Mechnanisms of Visual Perception*, ed. D. Man-Kit Lam and C. Gilbert, 171-198, Portfolio Publishing Co.

Newsome, W., Britten, K., & Movshon, J. A. (1989b) "Neural Correlates of a Perceptual Decision," *Nature*, **341**, 52-54.

Noda, H. & Adey, R. (1970) "Firing Variability in Cat Association Cortex During Sleep and Wakefulness," *Brain Res.* **18**, 513-526.

Oertel, D., Shu, H., and Hirsch, J. (1989) "Electrical Characteristics of Cells and Neuronal Circuitry in the Cochlear Nuclei Studied with Intracellular Recordings from Brain Slices," in *Auditory Function: Neurobiological Bases of Hearing*, ed. G. Edelman, W. GAll, and W. Cowan, New York (Wiley Interscience).

Parker, A. & Hawken, M. (1985) "Capabilities of Monkey Cortical Cells in Spatial-resolution Tasks," *J. Opt. Soc. Am.* **2**, 1101-1114.

Perkel, D.H., Gerstein, G., & Moore, G. (1967) "Neuronal Spike Trains and Stochastic Point Processes," *Biophysical J.* **7**, 391-418.

Poggio, T. & Viernstein, L. (1964) "Time Series Analysis of Impulse Sequences of Thalamic Somatic Sensory Neurons," *J. Neurophysiol.* **27**, 517-545.

Pratt, G. A. (1989) "Pulse Computation," Ph.D. Thesis, Massachussetts Institute of Technology, Cambridge, MA.

Rall, W. & Segev, I. (1988) "Synaptic Integration and Excitable Dendritic Spine Clusters: Structure/Function," in *Intrinsic Determinants of Neuronal Form and Function*, Alan Liss Inc., 263-282.

Rall, W. (1964) "Theoretical Significance of Dendritic Trees for Neuronal Input-Output Relations," in *Neural Theory and Modelling*, ed. R. Reiss, Stanford Univ. Press, 73-97.

Regehr, W.G. & Tank, D.W. (1991) "Postsynaptic NMDA Receptor-Mediated Calcium Accumulation in Hippocampal CA1 Pyramidal Cell Dendrites," *Nature* **345**, 807-810.

Sah, P. Hestrin, S., and Nicoll, R. (1990) "Properties of Excitatory Postsynaptic Currents Recorded *in vitro* from Rat Hippocampal Interneurones," *J. Physiol.* **430**, 605-616.

Sayer, R., Friedlander, M., & Redman, S. (1990) "The Time Course and Amplitude of EPSPs Evoked at Synapses Between Pairs of CA3/CA1 Neurons in the Hippocampal Slice," *J. Neurosci.* **10**, 826-836.

Segev, I. & Rall, W. (1988) "Computational Study of an Excitable Dendritic Spine," *J. Neurophysiol.* **60**, 499-523.

Sejnowski, T. (1981) "Skeleton Filters in the Brain," in *Parallel Models of Associative Memory* (ed. Hinton, G. and Anderson, J.), Lawrence Erlbaum, Hillsdale, New Jersey.

Shepherd, G, Woolf, T., & Carnevale, N. (1989) "Comparisons Between Active Properties of Distal Dendritic Branches and Spines: Implications for Neuronal Computations," *J. Cog. Neurosci.* **1**, 273-286.

Shepherd, G. & Brayton, R. (1987) "Logic Operations are Properties of Computer-Simulated Interactions Between Excitable Dendritic Spines," *Neuroscience* **21**, 151-165.

Snowden, R.J., Treue, S., & Andersen, R. A. (1992) "The Response of Neurons in Areas V1 and MT of the Alert Rhesus Monkey to Moving Random Dot Patterns," *Exp. Brain Res.* **88**, 389-400.

Softky, W. & Kammen, D. (1990) "Correlations in High Dimensional or Asymmetric Data Sets: Hebbian Neuronal Processing," *Neural Networks* **4**, 337-347.

Softky, W. & Koch, C. (1992) "Cortical Cells Should Fire Regularly, But Do Not," *Neural Computation* **4**, 643-646.

Softky, W,, and Koch, C. (1993) "The Highly Irregular Firing of Cortical Cells is Inconsistent with Temporal Integration of Random EPSP's," *J. Neurosci.* **13**, 334-350.

Softky, W. "Submillisecond Coincidence Detection in Active Dendritic Trees," submitted to *Neuroscience.*

Sompolinsky, H., Golomb, D., & Kleinfeld, D. (1991) "Cooperative Dynamics in Visual Processing," *Phys. Rev. A* **43**, 6990-7011.

Spruston, N., & Johnson, D. (1991) "Perforated Patch-clamp Analysis of the Passive Membrane Properties of Three Classes of Hippocampal Neurons," *J. Neurophysiol.* **67**, 508-529.

Srinivasan, M. and Bernard, G. (1976) "A Proposed Mechanism for Multiplication of Neural Signals," *Biol. Cybern.* **21**, 227-236.

Stein, R. (1967a) "Some Models of Neuronal Variability," *Biophys. J.* **7**, 37-68.

Stein, R. (1967b) "The Information Capacity of Nerve Cells Using a Frequency Code," *Biophys. J.* **7**, 797-826.

Strassberg, A. & DeFelice, L. (1992) "Limitations of the Hodgkin-Huxley Formalism," submitted to *Neural Computation*

Strehler & Lestienne (1986), "Evidence on Precise Time-Coded Symbols and Memory of Patterns in Monkey Cortical Neuronal Spike Trains," *Proc. Nat. Acad. Sci. USA*, **83**, 9812-9816.

Suarez, H. and Koch, C. (1989) "Linking Linear Threshold Units with Quadratic Models of Motion Perception," *Neural Comput.* **1**, 318-320.

Sullivan, W. E. (1985) "Classification of Response Patterns in Cochlear Nucleus of Barn Owl: Correlation with Functional Response Properties," *J. Neurophys.* **53**, 201-216/

Teich, M., Matin, L., & Cantor, B. (1977) "Refractoriness in the Maintained

Discharge of the Cat's Retinal Ganglion Cell," *J. Opt. Soc. Am.* **68**, 386-401.

Thompson, A., Girdlestone, D., & West, D. (1988) "Voltage-Dependent Currents Prolong Single-Axon Postsynaptic Potentials in Layer **III** Pyramidal Neurons in Rat Neocortical Slices," *J. Neurophysiol.* 60, 1896-1907.

Tolhurst, D., Movshon, J. , & Dean, A. (1983) "The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex," *Vis. Res.***23**, 775-785.

Tononi, G., Sporns, O., & Edelman, G. (1992) "Modeling Integration in the Visual Cortex," *Society for Neuroscience Abstracts* **18** , 313.8, 741.

Toyama, K., Kimura, M., & Tanaka, K. (1981) "Organization of Cat Visual Cortex as Investigated by Cross-Correlation Technique," *J. Neurophysiol.* **46**, 202-214.

T'so, D., Gilbert, C., & Wiesel, T. (1986) "Relationships Between Horizontal Interactions and Functional Architecture in Cat Striate Cortex as Revealed by Cross-Correlation Analysis," *J. Neurosci.* **6**, 1160-1170.

Tuckwell, H.C. (1989). "Stochastic Processes in the Neurosciences," Soc. for Indust. and App. Math., Philadelphia.

Vogels, Spileers, & Orban (1989) "The Response Variability of Striate Cortical Neurons in the Behaving Monkey," *Exp. Brain Res.* **77**, 432-436.

von der Malsburg, C. & Schneider, W. (1986) "A Neural Cocktail-Party Processor," *Biol. Cybern.* **54**, 29-40.

Westenbrook, R., Ahlijinian, M.K., & Catterall, W.A. (1990) "Clustering of L-type Ca^{2+} Channels at the Base of Major Dendrites in Hippocampal Pyramidal Neurons," *Nature* **347**, 281-284.

Wilbur, W.J., & Rinzel, J. (1983) "A Theoretical Basis for Large Coefficient of Variation and Bimodality in Neuronal Interspike Interval Distributions," *J. Theor. Biol.*, **105**, 345-368.

Wong, R., Prince, D., and Basbaum, A. (1979) "Intradendritic Recordings from Hippocampal Pyramidal Neurons," *Proc. Natl. Acad. Sci. USA* 76, 986-990.

Zohary, E., Hillman, P., & Hochstein S. (1990) "Time Course of Perceptual Discrimination and Single Neuron Reliability," *Biol. Cybern.* **62**, 475-486.

Zohary, E., Shadlen, M., & Newsome, W. (1992) "Correlated Activity of Neurons in Area MT," *Soc. for Neuroscience Abstracts* **18**, 464.4, 1101.