

Noise and Information in Neural Codes

Elad Schneidman

This work was carried out under the supervision of

Prof. Idan Segev and Prof. Naftali Tishby.

Contents

1	Introduction	1
2	Background: Neuronal Biophysics: Ion Channels, Spiking, Noise and Modelling	7
2.1	Biophysical building blocks	7
2.1.1	Morphology, ions, channels, pumps, potentials etc.	8
2.1.2	Signaling and Spiking	10
2.1.3	Synapses	11
2.2	The spiking mechanism and its modelling	12
2.2.1	The Hodgkin-Huxley model	13
2.2.2	Properties of the HH model and its applicability as a general neuron model	17
2.3	Noise in Neurons	19
2.3.1	Expression of Noise in Neurons and its sources	19
2.3.2	Ion Channel Noise	21
2.4	The Computational Sketch of a Single Neuron	22
3	Background: The Neural Code – Spike Trains, Coding and Information Theory	25
3.1	Quantifying the Neural Responses	25
3.1.1	Spike train statistics	26
3.2	The Nature of the Neural Code	29
3.2.1	Rate Codes and Temporal Codes	29
3.2.2	Noise, Reliability and the neural code	30
3.2.3	Rate code and Temporal code, revisited	31
3.2.4	Experimental results of input-dependent unreliability, noise and code	31
3.2.5	Population Codes	32
3.3	Information theory and the neural code	34

3.3.1	Entropy, Relative Entropy and Mutual Information	35
3.3.2	Spike train entropy and information rates	37
3.3.3	Design principles of the code	41
4	Reliability and accuracy of spike timing in stochastic neuron models	43
4.1	Introduction	43
4.2	The stochastic HH model	45
4.3	Reliability and accuracy of spike timing in the SHH model	47
4.3.1	Encoding Reliability and Precision: Input Current Versus Channel Fluctuations	48
4.3.2	Sub-threshold oscillations, 'Spontaneous' spikes and 'missing' Spikes .	54
4.4	Spike reliability in the stochastic Wang-Buzsaki model	58
4.5	Conclusions and discussion	59
4.5.1	Summary of results	59
4.5.2	Towards a more realistic stochastic model of neurons	60
4.5.3	Other sources of noise in neurons	62
4.5.4	Implication for Neural Coding	63
5	The nature of the neural code, information optimization and biophysical design in stochastic neuron models	69
5.1	Introduction	69
5.2	Spike train properties of the SHH model	71
5.3	Information rates of the SHH model	71
5.3.1	Coding efficiency, rate coding and temporal coding in the SHH model	76
5.4	Dependence of information encoding on the parameters of the model	79
5.4.1	The effect of reducing the ion channel noise on information encoding is not unidirectional	79
5.4.2	Changing channel densities and information encoding	80
5.5	Conclusions and discussion	82
5.5.1	The nature of the neural code	82
5.5.2	Robustness and optimality of the spike generation mechanism	83
5.5.3	Is ion channel noise a bug or an unavoidable feature?	85
5.5.4	Limitations of the current work and future directions	86
6	Universality and individuality in a neural code	91
6.1	Introduction	91
6.2	An ensemble of flies and the experimental setup	92

<i>CONTENTS</i>	iii
6.3 Distinguishing among individuals	95
6.4 Spike rates and information rates of the H1 of different flies	99
6.5 Using a universal codebook	101
6.6 The nature of the 'personal' bits	103
6.7 Discussion	104
7 Concluding remarks	107

Chapter 1

Introduction

The photons reaching your retina, sound waves reaching your ears, touch on your skin, as well as taste and smell, are translated by (many) millions of sensory neurons into sequences of electrical pulses which are sent along millions of nerve fibers to your brain. Such electrical pulses, termed action potential or spikes, are the 'language' of most of the nervous system, in the blue whale as well as in a cockroach. From the different sensory modalities, to the commands to the motor system, most of the neurons use patterns of spikes to represent information and carry it from one neuron to (many) others. Spikes do not resemble the external world they represent, nor the thoughts they carry, anymore than the musical notes resemble the sound of a cello. They are the symbols of the *neural code*, which contains two basic 'letters' – spike and 'no-spike' (very much like the zeros and ones in computers). The problem of neural coding is to understand how patterns of spikes are related to sensory stimuli, motor outputs, and ultimately, thoughts and intentions.

From a mathematical or engineering standpoint, the questions to be asked about the nature of the neural code are not different from what we would ask about any other coding and communication system (or an input-output black box): What are the codewords? What do they encode and how are they decoded? Can we build a dictionary from the stimuli to the responses, and vice versa? What is the capacity of the code to convey information? How robust is the code to transmission or encoding (and decoding) errors? How efficient is the code, time-wise and cost-wise? Are the codewords fixed, or history dependent? Is it universal over the different parts of the system or between systems? A range of system identification and machine learning tools have been used in the analysis of neuronal stimulus-response paradigms, neural decoding algorithms and feature selection. Shannon's mathematical theory of communication [160], which laid the foundation of what is now known as *information theory*, gives a general mathematical framework for the analysis and quantification of coding systems and communication channels. It is therefore an appropriate tool for the study of

the neural code.

From a biological viewpoint, understanding the code is inseparable from understanding the biophysical machinery and the organism “point of view”: What is the biophysical design that enables neuronal computing and coding? How well does it perform in terms of accuracy, time and energy consumption? How does the system (or the animal) use the code and actually decode the spike trains? An implicit question is how different could the neurons and the neural code be, and can we quantify the implications? Electrophysiological studies, together with modelling and theoretical work have addressed these questions, reflecting on the nature of neuronal codewords, what they stand for (in terms of the stimulus), the accuracy of neuronal response, etc. (see Chapters 2,3) – for the animal, neuronal population and the single neuron level.

The work that I will describe deals with two fundamental issues concerning the nature of the neural code. One is the relation between the neuronal biophysical design and noise and the nature of the neural code. The other is the question of universality of the neural code.

- The nature of spike patterns as a code and the encoded content, are set by the biophysical design and properties of the single neuron and the computation it performs [88]. A neuron may receive thousands of spike trains from other neurons as input. These spike trains are transformed into voltage transients inside the neuron, and are then merged (or ‘processed’). The result is a single spike train that the neuron sends as an output to thousands of other neurons. This biophysical computation depends on the physical structure of the neuron, the ionic concentration and electric potential inside and outside the neuron membrane, and the activity of special proteins embedded in the neuronal membrane which enable ions to flow inside (or outside) of the neuron (see Chapter 2). The biophysical ‘building blocks’ of neurons are inherently noisy and unreliable [64, 88]. A clear question is then, what are the effects of biophysical design and specifically of neuronal noise on the computation and coding of neurons? (see [190]).
- Since the pioneering work of Adrian [8], we know that the common alphabet of spike and ‘no-spike’ is common to all spiking neurons, that the number of spikes is correlated with the stimulus that the neuron responds to, and that neuronal firing rate adapts in response to a fixed stimuli. Using information theory, in recent years, the capacity of the neural code has been quantified by estimating the amount of information that the spike train may convey, in different sensory and other neural systems in various animals [22]. These results show high information rates and coding efficiency, which are achieved using a rather sparse code (in terms of the number of spikes that are

used). An important, yet somewhat neglected question is what are the common coding features and what are the individual ones of different neurons, either neighboring ones or in different areas. Of special interest is the question whether the neural code of corresponding modules and neurons in different individual animals is universal.

We focus here on three questions at the core of these issues, which reflect on the design and nature of both the neural code and neuronal computation:

- **What is the source of the input-dependent spike timing reliability and accuracy, and what may be their implications?**

The main debate over the nature of the neural code has focused on whether information is encoded by the exact temporal structure of the spike train (*temporal code*), or that the fine temporal structure is noisy and information is then carried just by the average spiking rate (*rate code*). The type and resolution of code that neurons use determine the capacity of the code to convey information, and reflects on the possible decoding algorithm which neurons and the nervous system may use. However, recent experimental results *in vivo* and *in vitro* show that the same neuron may have very accurate spike timing in response to one stimulus, and highly unreliable spike timing for another. The neural code is then *input-dependent*, i.e. information may be encoded by both the average firing rate and the temporal structure of the neuronal spike train, and the division is set by the nature of the code. We study detailed biophysical models of spike generation, which incorporate ion channel noise, and investigate the nature of their spike trains. We find that the stochastic models reproduce the experimentally observed input-dependent reliability and precision of spike firing (as well as other spike train characteristics). Our results suggest that ion channel noise may have a profound effect on the nature of the neural code. We suggest that the noise inherent in neurons enable them to act as “smart” encoders (hence, it could be less of a bug and more of a feature), and thus reflects on the way information may be propagated and processed in the nervous system.

- **What are the design principles of the spiking mechanism and how may the biophysical parameters of the neuron affect the nature and content of the neural code?**

The computational “task” of the spiking mechanism is to encode selected information about the stimulus that the neuron is presented with. The biophysical design of the neuron which performs this computation, must accommodate the inherent biophysical

noise, metabolic costs and efficiency considerations. We use tools from information theory to quantify the characteristics and performance of the neuronal spiking mechanism and analyze the structure of the spike train as a code. We proceed to evaluate the effects of changing the biophysical parameters of the spiking mechanism, and learn of the 'design principles' of the 'biological hardware', its limitations and capabilities.

We calculate the information rate and efficiency that the stochastic neuron models convey about a family of stimuli, and find it to be similar to values reported experimentally. The division between the rate coding component and temporal coding component in the spike train may be highly variable for different inputs, reflecting the dynamic nature of the neural code. Study of the biophysical parameters of the models suggests that information encoding is robust to most parameters, but is sensitive to the ratio between ion channel densities. We find that there is an optimal range of parameters for information encoding, which is well within physiological range, even without taking into account any metabolic considerations. Thus, neurons may maximize their information capacity by appropriately balancing the density of the different ion channels that underlie neuronal excitability.

- **How universal and how individual is the neural code?**

An obvious question about the nature of the neural code is whether the same coding rules are used by different neurons, or by corresponding neurons in different individuals. We present a quantitative formulation of this problem using ideas from information theory, and apply this approach to the analysis of experiments in the fly visual system. We find significant individual differences in the structure of the code, particularly in the way that temporal patterns of spikes are used to convey information beyond that available from variations in spike rate. On the other hand, all the flies in our ensemble exhibit a high coding efficiency, so that every spike carries the same amount of information in all the individuals. Thus the neural code has a quantifiable mixture of individuality and universality.

We present a rather wide review of the biological and mathematical background of this work (based mostly on the books by Nicholls, Martin and Wallace [126], Koch [88], Tuckwell [182], Hille [64], Dayan and Abbott [37], Cover and Thomas [35], Rieke, Warland, de Ruyter van Steveninck and Bialek, [143]). Chapter 2 introduces the basic biophysics of neuronal design and function – morphology, ion channels, synapses and spiking (focusing on modelling). Chapter 3 presents the basic tools of spike train analysis, information theory and its application to neural coding. Chapters 4 and 5 present the study of the neural code of stochastic neuron models of spike initiation, incorporating ion channel noise. Chapter 6 present the

study of the universality and individuality in the neural code of the fly visual system. Final remarks and conclusions are given in chapter 7.

Chapter 2

Background: Neuronal Biophysics: Ion Channels, Spiking, Noise and Modelling

Although neurons may differ considerably in their shape, size, and many of their molecular components, their basic 'design principles' are similar. Sharing the same basic morphology and biophysical mechanisms enables (the vast majority of) neurons to initiate and propagate action potentials and facilitates communicating them to other neurons through synaptic connections.¹ This chapter provides an overview of the basic building blocks of neuronal biophysics, namely ion channels, ionic currents, neuronal morphology, synapses and electrical properties of the membrane. We focus on the spiking mechanism and modelling of neuronal excitability and on the noise in neuronal function.

2.1 Biophysical building blocks

Neurons are specialized for receiving information from other cells (or in the case of the periphery, the outside world), generating voltage transients (mostly action potentials) in response to these inputs, and sending them to other neurons. What enables neurons to perform their unique 'task', is their distinct morphological and biophysical design, both on the molecular level, and the macroscopic one.

¹A comparatively small class of neurons are nonspiking and use voltage graded responses as their output; e.g. the bipolar cells, horizontal cells, some of the amacrine cells in the retina and many neurons in the sensory-motor modules of invertebrates.

2.1.1 Morphology, ions, channels, pumps, potentials etc.

The typical morphology of neurons is made of three main structural components. The cell body or *soma*, where the regular cell 'machinery' resides (i.e. the nucleus, mitochondria, endoplasmic reticula, etc.), the *dendrites* which branch out of the soma – tree-like cable structures that receive inputs from other neurons, and the *axon* – a single cable leaving the soma which branches out to connect to the dendrites (or soma) of other neurons (see Figure 2.1). The complexity of the dendritic structure enables the neuron to receive inputs from many other neurons through synaptic connections (many thousands for a typical cortical neuron). Correspondingly, the axon from a single neuron may traverse very long distances, and connect to hundreds or thousands of neurons.

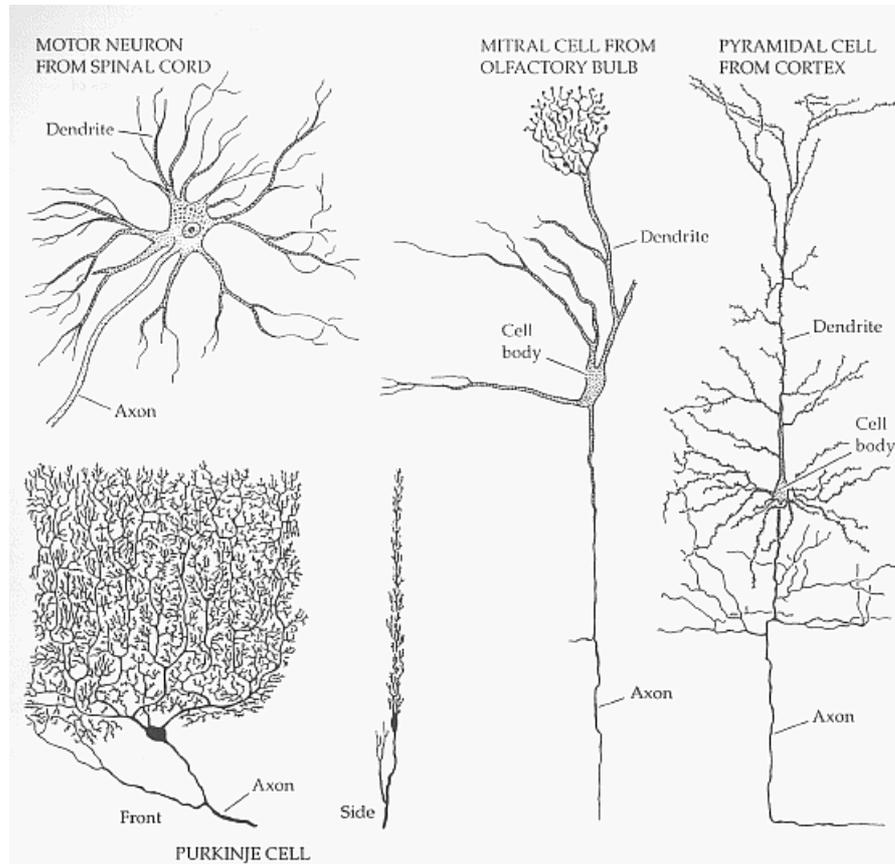


Figure 2.1: **Morphology of four types of neurons, and their common features.** Drawings of four types of neurons, shown as representative examples of neuronal morphology. While different in their size (figure is not to scale), the motor neuron (from the mammalian spinal cord), the Purkinje cell (from human cerebellum), the mitral cell (from the olfactory bulb of a rat) and the pyramidal cell (from the cortex of a mouse), share the same structure of an elaborate dendritic tree(s), cell body and an axon. Taken from [126]

Physiologically, the electrical signal of relevance to the nervous system is the difference

in electrical potential between the interior of a neuron and the surrounding extracellular medium. The ionic concentration gradients across the cell membrane and the membrane permeability to these ions, determine the membrane potential. The cell membrane is a lipid bilayer, which is impermeable to most ions. Electrically, the membrane is a capacitor separating the charges residing along its inner and outer surface, from both sides. While the resistance of the lipid bilayer by itself is quite high, the resistance of the membrane is significantly reduced by the numerous aqueous pores in the membrane, termed *ion channels*.

Ion conducting channels are proteins embedded in the membrane that allow the flow of ions, mainly Na^+ , K^+ , Ca^{2+} and Cl^- , across the membrane (see Figure 2.2). Channels may be highly selective for the ion charge or type (or have very little ion selectivity). *Voltage-gated* channels change their functional state, i.e. moving from open to close state and vice versa, as a function of the membrane potential. Other channels are activated (opened) by intracellular messengers (for example, Ca^{2+} dependent channels) or by extracellular ligands such as neurotransmitters or neuromodulators (e.g. synaptic receptor channels), and are usually called *ligand-gated* channels. Typically, a neuron would have quite a few different channel types, which may be selective to different ions. The distribution of channels in the neuronal membrane is far from uniform, and may range from a few channels up to a few hundreds of channels per square micron. The total number of channels in a cell is on the order of a few hundred thousands or millions.

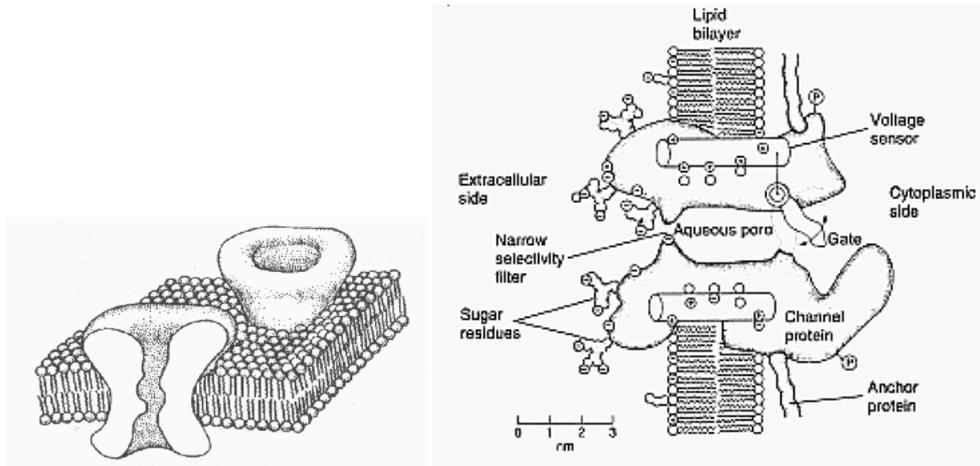


Figure 2.2: **Ion channels in the cell membrane.** Schematic representation of ion channels embedded in the cell lipid layer membrane (left), and a schematic view of a generic voltage-sensitive ion channel (both taken from [64]).

Specific ion types have distinct intracellular and extracellular concentrations (for example, potassium concentration inside the cell is higher than outside). The ions flow into and out of the cell due to both voltage and concentration gradients. However, without external

stimuli, these different forces drive the cell to an equilibrium point – the *resting potential* of a neuron, which can be explained from basic physical chemistry principles. Under these resting conditions, the electrical gradient and the ionic concentration gradient balance each other for each of the ion types. The potential inside the cell membrane of a neuron, resulting from the accumulation of charges on the membrane, is then about -70 mV relative to that of the surrounding bath, and the cell is said to be *polarized*. Selective *ion pumps* located in the cell membrane expend energy (see e.g. [64]) to maintain the potential concentration gradients that support this membrane potential difference.

The potential changes if the balance of ion flow is modified by the opening or closing of ion channels. The membrane is *hyperpolarized* when current, in the form of positively charged ions flows out of the cell (or negatively charged ions flowing into the cell) through open channels. If current flows into the cell, the membrane potential becomes less negative (or even positive), and the cell is *depolarized*. Normal neuronal membrane potentials vary over a range from about -90 mV to $+50$ mV.

2.1.2 Signaling and Spiking

The currents flowing through the ion channels are responsible for the two types of electrical signals generated by neurons – localized potentials and action potentials. Localized, graded potentials can travel only short distances before attenuation, and so they are useful in special, physically compact regions, e.g. synaptic connections (see below) or sensory nerve endings. Action potentials are voltage impulses that travel (fast) along the axon and are capable of propagating to long distances. The main reason is that the action potential is a regenerative impulse – the voltage dependent ion channels along the axon prevent its decay². Axons terminate at *synapses*, the connection points between the axon of one neuron and the dendrite or soma of another (see section 2.1.3 below).

If a neuron is depolarized sufficiently to raise the membrane potential and open enough ion channels (typically, Na^+ channels), current flows into the cell, initiating a positive feedback process, and the neuron generates an action potential. An action potential is a roughly 100 mV change in the electrical potential across the cell membrane which lasts about 1 ms. Figure 2.3 shows the membrane voltage of a layer V pyramidal neuron, recorded using an intracellular electrode.

²In many cases the axon is covered with myelin (formed by neuroglial and Schwann cells) which increases the axonal membrane resistance, resulting in a speedy and more reliable conduction of the spikes. Along myelinated axons there are usually the *nodes of Ranvier*, where the myelin cover is interrupted and the axon is exposed. In these areas, there is a high density of ion channels, which re-amplify the spike for the next axonal segment.

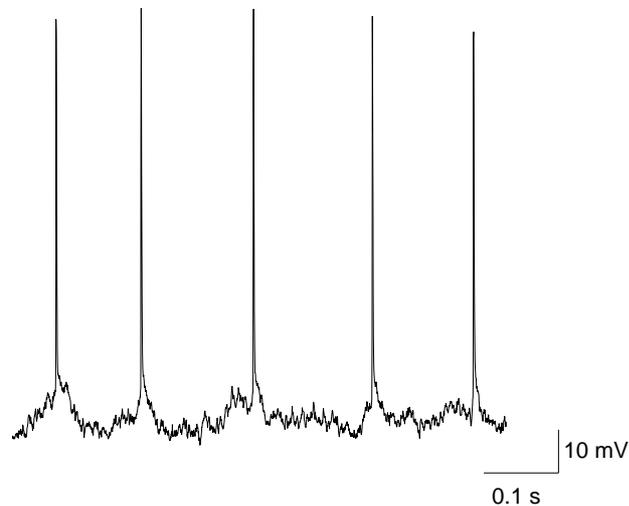


Figure 2.3: **A spike train of a cortical pyramidal neuron, *in vitro*.** The membrane voltage of a layer V cortical pyramidal neuron of a two weeks old rat, recorded *in vitro* using an intracellular patch electrode. The cell was injected with a small DC current. Experiment performed in the lab of Yosi Yarom

The generation of an action potential also depends on the recent history of cell spiking. It is virtually impossible to fire two action potential within less than a few milliseconds, a phenomena called the *absolute refractory period* of the neuron. During the *relative refractory period*, which lasts up to tens of milliseconds after a spike, it is 'difficult' to evoke an action potential. I.e., more current is needed to evoke a spike during this period.

2.1.3 Synapses

Synapses are the connection points of one neuron to another (or of a neuron with an effector cell, e.g. muscle fiber or secretory cell), where signals are passed from one cell to another. At electrical synapses (or gap junctions), current generated by an impulse in the presynaptic cell terminal spreads into the postsynaptic cell through low-resistance channels. More commonly, synapses are chemical, where there is no direct spread of current between the presynaptic and postsynaptic sides of the synapse. Synaptic transmission begins when an action potential arrives at the presynaptic terminal and causes a rise in the concentration of Ca^{2+} within the terminal (due to the activation of local voltage-dependent channels). This causes vesicles containing transmitter molecules to fuse with the cell membrane and release their contents into the gap between the presynaptic and postsynaptic membranes. Transmitter molecules diffuse across the cleft and bind to receptors on the postsynaptic membrane. Binding of transmitter molecules leads to the opening of ion channels that modify the conductance of

the postsynaptic neuron, completing the transmission of the signal from one neuron to the other.

Postsynaptic ion channels can be activated directly by binding to the transmitter (ionotropic synapses), or indirectly when the transmitter binds to a distinct receptor that affects ion channels through an intracellular second-messenger signaling pathway (metabotropic synapses). Depending on the nature of the ion flow (type and direction), the synapses can have either an excitatory, depolarizing, or an inhibitory, typically hyperpolarizing, effect on the postsynaptic neuron.

2.2 The spiking mechanism and its modelling

The pioneering experimental work of recording membrane conductances and intracellular action potentials [36, 67, 73], suggested that the source of action potentials are transient changes in the membrane conductances. Following, Hodgkin and Huxley's experimental work presented the first quantitative description of the ionic mechanisms responsible for spiking in the squid axon. Using the *voltage clamp* method³ with ionic substitutions, blockers etc., they identified the separate ion-type currents and quantified their dependency on the membrane voltage and as a function of time. They showed that changes in sodium and potassium conductances occurred during spiking, and that the timing and nature of these changes correspond to the time course and magnitude of the action potential [69, 68, 70]. They then concluded their study and series of papers by introducing a model of excitability in a single cell, which accounts qualitatively for conduction and excitability of the squid giant axon [71]. Hodgkin and Huxley suggested a physical interpretation of the changes in ionic conductances using “gating particles” which can be either in an open or closed states, and are basically a 'ball and chain' model of ion channels. However, their experimental work did not deal with single channels, (whose existence was empirically observed many years later; see below), and the formalism they introduced to describe voltage-dependent conductances (which is almost universally used to describe voltage-dependent conductances to this day) is based on **deterministic** differential equations which deal with the average behavior of ionic currents.

³A recording technique which makes it possible to set the membrane potential of the cell almost instantaneously at any desired level and keep it fixed (i.e. “clamped”), while recording the current flowing through the membrane.

2.2.1 The Hodgkin-Huxley model

Based on their experimental work, Hodgkin and Huxley [69, 68, 70] suggested that the electrical behavior of a membrane patch can be represented by the circuit presented in Figure 2.4.

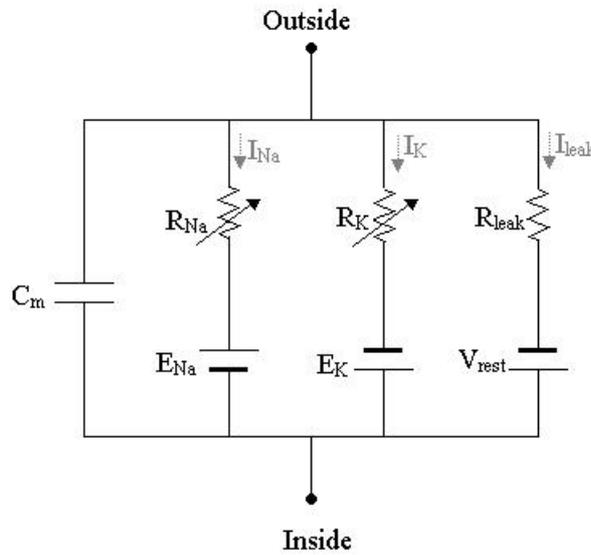


Figure 2.4: **Electrical circuit description of the Hodgkin-Huxley model of a membrane patch.** The membrane of the squid axon is modelled using four parallel branches – The voltage dependent sodium and potassium conductances ($R_K = 1/G_K$ and $R_{Na} = 1/G_{Na}$), the leak conductance ($R_{leak} = 1/G_{leak}$) and the membrane capacitance. (The nonlinearity "hides" in the voltage dependence of the sodium and potassium ionic conductances)

The Hodgkin-Huxley (HH) model is of an *isopotential* membrane patch (i.e. there are no spatial effects on the potential), or a single electrical compartment. The units of the model are per membrane unit area, and it is then straightforward to scale the model to a single compartment of any desired membrane area.

The total membrane current is the sum of the ionic currents and the capacitive current,

$$\begin{aligned} I_m(t) &= I_{ionic}(t) + C_m \frac{dV_m(t)}{dt} \\ I_{ionic}(t) &= I_{Na}(t) + I_K(t) + I_{leak}(t) \end{aligned} \quad (2.1)$$

where I_m is the membrane current density, I_{ionic} are the ionic currents densities, C_m is the membrane capacity per unit area (assumed constant) and V_m is the membrane voltage. The two main ionic conductances, sodium and potassium are independent of each other, and a third, “leak conductance” does not depend on any of the other conductances or the membranal voltage. Thus, the total ionic current is the sum of the separate ionic currents. The individual ionic currents are linearly related to the potential according to Ohm’s law,

$$\begin{aligned} I_K(t) &= G_K(V, t)(V(t) - E_K) \\ I_{Na}(t) &= G_{Na}(V, t)(V(t) - E_{Na}) \\ I_{leak}(t) &= G_{leak}(V, t)(V(t) - V_{leak}) \end{aligned} \quad (2.2)$$

where G_K , G_{Na} and G_{leak} are the potassium, sodium and leak conductances per unit area of the membrane (correspondingly) and E_K , E_{Na} and V_{leak} are the corresponding *reversal* or *equilibrium* potentials of each of the ionic species (the potential at which the ionic concentration gradient is balanced by the electrical potential gradient, and there is no net flux of the ions of this type).

The voltage-dependent conductances $G_{Na}(t)$ and $G_K(t)$ are given by

$$G_{Na}(t) = G_{Na}^{max} f_{Na}(t) \quad ; \quad G_K(t) = G_K^{max} f_K(t) \quad (2.3)$$

where G_{Na}^{max} and G_K^{max} are the maximal sodium and potassium conductances per unit membrane area and $f_{Na}(t)$ and $f_K(t)$ are each the corresponding (instantaneous) fraction of the maximal conductance which is actually open (or active). The ion channel interpretation of these equations is that each of the ion species i conductance, G_i is determined by multiplying the conductance of a single open channel by the density of channels in the membrane and by the fraction of channels that are open at time t . G_i^{max} is a membrane constant, equal to the product of the first two factors, that describes the conductance per unit area of membrane when all channels of type i are open. The fraction of channels in the open state is equivalent to the probability of finding any given channel in the open state, and it is denoted by f_i . In the original Hodgkin and Huxley interpretation, these fractions were functions of the activation and inactivation of the “gating particles” which control the ionic conductances. Hence, the potassium current is modelled as,

$$I_K = G_K^{max} n^4 (V - E_K) \quad (2.4)$$

where $G_K^{max} = 36mS/cm^2$ and $E_K = -12mV$, relative to the resting potential. The fraction of maximal conductance which is open is given by n^4 , where n is a dimensionless activation variable, ranging between 0 and 1 and whose dynamics are given by,

$$\frac{dn}{dt} = \alpha_n(V)(1 - n) - \beta_n(V)n \quad (2.5)$$

where

$$\alpha_n(V) = \frac{V - 10}{100(1 - \exp^{-(V-10)/10})} \quad (2.6)$$

and

$$\beta_n(V) = 0.125 \exp^{-V/80} \quad (2.7)$$

are voltage-dependent *rate functions*.

An alternative way to write the dynamics of n is in terms of the steady state value of $n_\infty(V)$ (the value of the activation variable n for a fixed V at $t \rightarrow \infty$) and the voltage-dependent time constant $\tau_n(V)$ (a measure of the speed of the dynamics of n),

$$\frac{dn}{dt} = \frac{n_\infty - n}{\tau_n} \quad (2.8)$$

where

$$n_\infty = \frac{\alpha_n(V)}{\alpha_n(V) + \beta_n(V)} \quad (2.9)$$

and

$$\tau_n(V) = \frac{1}{\alpha_n(V) + \beta_n(V)} \quad (2.10)$$

Similarly, the sodium current is modelled by,

$$I_{Na} = G_{Na}^{max} m^3 h (V - E_{Na}) \quad (2.11)$$

where $G_{Na}^{max} = 120 \text{ mS/cm}^2$ and $E_{Na} = 115 \text{ mV}$, relative to the resting potential. m and h are dimensionless activation and inactivation variables, with $m \geq 0$ and $h \leq 1$ and their dynamics are given by,

$$\frac{dm}{dt} = \alpha_m(V)(1 - m) - \beta_m(V)m \quad (2.12)$$

and a similar equation for h .

The corresponding rate functions are

$$\alpha_m(V) = 0.1 \frac{V - 25}{1 - \exp^{-(V-25)/10}} \quad ; \quad \beta_m(V) = 4 \exp^{-V/18} \quad (2.13)$$

and

$$\alpha_h(V) = 0.07 \exp^{-V/20} \quad ; \quad \beta_h(V) = \frac{1}{1 + \exp^{-(V-30)/10}} \quad (2.14)$$

An equivalent description of m and h using m_∞ , h_∞ , τ_m and τ_h (similar to equations 2.8- 2.10) is straightforward.

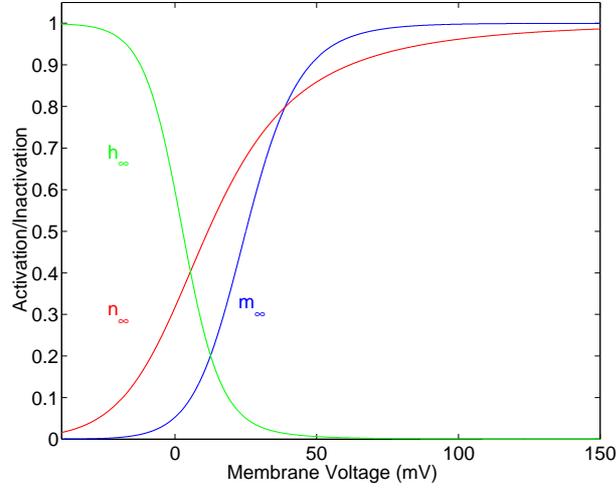


Figure 2.5: **Steady state values of activation and inactivation variables in the HH model as a function of the voltage.** The steady state sodium activation m_∞ (blue) and potassium activation n_∞ (red) increase with membrane voltage, whereas the steady state sodium inactivation, h_∞ (green), monotonically decreases with V . Activation (and inactivation) of the channel conductances is a steeper function of the voltage, due to the power law relation between the activation (inactivation) variables and the conductances.

Figure 2.5 shows the dependence of the steady state values of activation and inactivation variables on the membrane voltage.

The leak conductance density of the membrane is fixed, $G_{leak} = 0.3 \text{ mS/cm}^2$, and the leak reversal potential associated with it is $V_{leak} = 10.613 \text{ mV}$ (set so that the total membrane current at the resting potential is zero). Finally, the membrane capacity is $C_m = 1 \text{ }\mu\text{F/cm}^2$.

Thus the equation which describes the membrane potential as a function of all the currents that flow across it is,

$$\begin{aligned}
 C_m \frac{dV}{dt} &= G_{Na}^{max} m^3 h (E_{Na} - V(t)) + G_K^{max} n^4 (E_K - V(t)) \\
 &+ G_{leak} m^3 h (V_{leak} - V(t)) + I_{injected}(t)
 \end{aligned} \tag{2.15}$$

where $I_{injected}(t)$ (measured in $\mu\text{A/cm}^2$) is the current injected via an intracellular electrode. This equation, together with the sets of equations for the activation and inactivation variables, are the Hodgkin and Huxley model for an isopotential membrane patch.

The rate functions of the HH model were fit to 6.3°C , which is considered the 'standard' temperature of the HH model. To correct for other temperatures, the rate constants (i.e. α 's and β 's) are scaled by a multiplicative factor ϕ ,

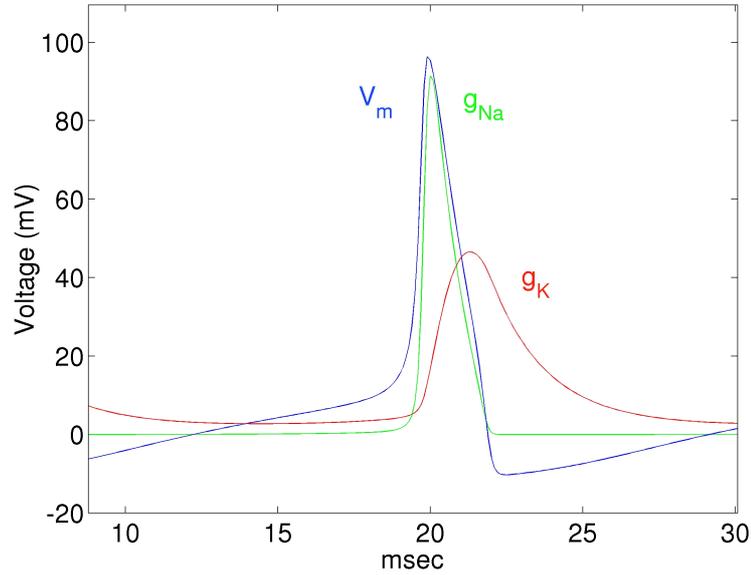


Figure 2.6: **HH spike and ion conductances.** The membrane potential of a $200 \mu m^2$ HH patch, in response to a short DC current injection is shown with the separation of ionic conductances underlying the action potential. Sodium and potassium conductances (measured in mS/cm^2) are shown in arbitrary units

$$\phi = Q_{10}^{(T-6.3)/10} \quad (2.16)$$

where the temperature coefficient Q_{10} (defined as the increase in rate when the temperature changes by $10^\circ C$) equals 3, and T is the temperature in Celsius.

Figure 2.6 shows the membrane potential during a spike (the response of the HH model of a $200 \mu m^2$ membrane patch to a short DC current injection), as well as the sodium and potassium conductances.

2.2.2 Properties of the HH model and its applicability as a general neuron model

The HH model replicates many of the features of spiking of the squid giant axon: the form, duration and amplitude of a single spike (both for the membrane and the propagating spike), its sharp threshold, the conduction velocity of the spike along the axon, the refractory period of the neuron, the impedance changes during the spike, anode-break excitation, accommodation, subthreshold response and oscillations. When simulating the response to a sustained stimulus currents, it demonstrates a discontinuous onset of repetitive firing with a high spiking frequency and a limited bandwidth of the firing frequency. Figure 2.7 shows the response

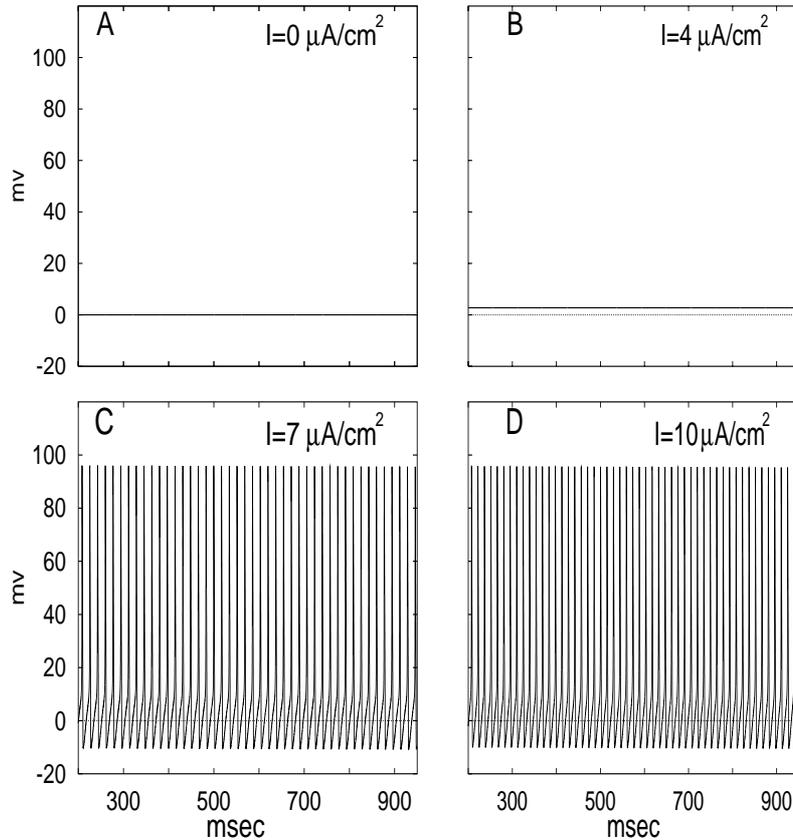


Figure 2.7: **HH spiking.** (A-D) Segments of the spiking pattern of a $200 \mu m^2$ HH patch, in response to sustained DC currents. A and B show the steady state of the membrane voltage in response to subthreshold currents; C and D show the sustained firing pattern in response to superthreshold ones

of the HH model to sustained DC current inputs reflecting the discontinuous relation between the input current amplitude and the firing frequency.

However, careful studies of the model reveal that it does not provide a good description of quite a few electrophysiological properties of the axon (see e.g. [32]), in particular the refractory behavior of the preparation in response either to sustained or periodic current pulse stimulation. Also, the model does not account for after potentials and slow changes in the squid giant axon.

As a generic model for a spiking neuron, the HH model is criticized since it has a discontinuous $f-I$ curve, exhibiting an *inverted-Hopf* current-voltage dynamics [48, 57], (also termed a Type II membrane dynamics), whereas most of the cortical neurons have a continuous and smooth *saddle-node* current-voltage relation (a Type I membrane) [182, 48, 57]. Additionally, the HH model firing pattern is non-adapting. Finally, the model is deterministic and does not exhibit the response variability most neurons demonstrate (see below). As a quantitative

neuronal model, it is clear that the HH model (like any other detailed model) needs to be modified to explain the biophysical aspects and macroscopic behavior of different kinds of neurons, as they may have a wide range of voltage-dependent currents and ligand dependent ones, different membrane areas, etc.

Still, the HH model serves as the 'golden standard' of neuronal excitability, and with minor changes - as the backbone of most neuronal spiking models. The main reason is that the HH model does capture the essence of spiking through ionic currents (Na^+ and K^+) which enter and leave the cell through voltage dependent channels. Moreover, the model is compact, and approximates well many of the features shared by different types of neurons (shape and duration of spiking, repetitive spiking in response to sustained inputs, refractoriness etc.) while incorporating biophysical aspects of the neuron. Adding the appropriate currents for other channel types (usually using similar kinetic schemes) is easily done. Accordingly, and since the model has been studied mathematically in great detail [81], it is the common choice of conductance based modelling for computational studies and theoretical ones⁴.

2.3 Noise in Neurons

2.3.1 Expression of Noise in Neurons and its sources

Like any other physical system, noise is an inherent feature of neurons resulting from their complex biophysical design and components. Unreliable and stochastic function have been observed on both the macroscopic and microscopic level of neurons, from the early days of neuronal electrophysiology. For example, neurons may spike spontaneously without stimulation, the firing of neurons exposed to near-threshold stimuli seems to be random, synaptic transmission is highly unreliable, action potentials conduction along the axon may fail, or advance with variable velocity (see e.g. [96, 97, 88, 126, 6, 113]). Figure 2.8 shows an example of the unreliable spike timing of neurons in response to repeated DC current injection *in vitro* (and interestingly – a reliable response of the same cell to repeated presentation of a fluctuating stimulus. We will return to that in the next chapters)

The unreliability and random-like function of single neurons is usually assumed to arise from the properties of the biophysical building blocks and design of the cell. The main sources of noise are associated with the membrane, synapses and ion channels: Thermal noise due to the membrane resistance (Johnson noise), quantal release of synaptic transmitters and the open-close fluctuations of the ion channels, which may originate from thermal

⁴Also common are the simplified model of Fitzhugh-Nagumo which is more mathematically tractable and the integrate and fire model (see e.g. [182]). For discussion see [154]

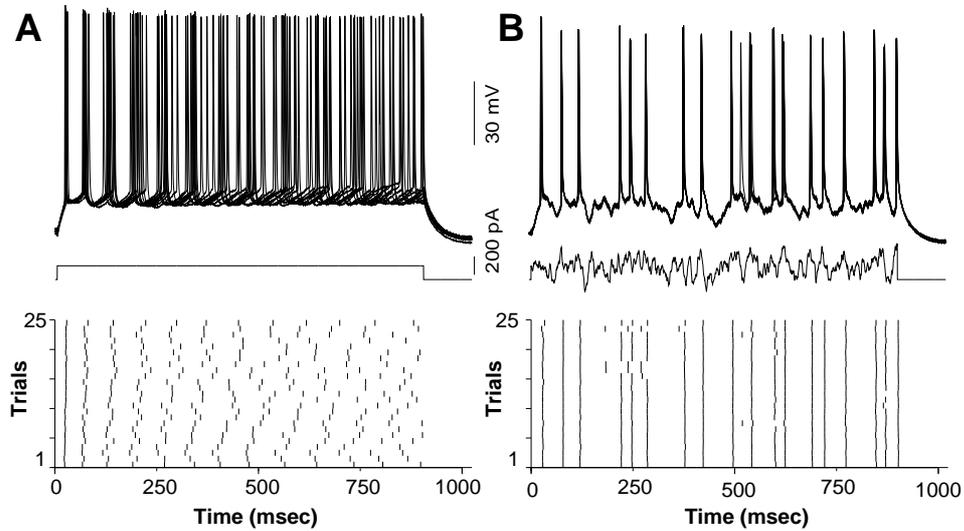


Figure 2.8: **Reliability of firing patterns of cortical neurons evoked by constant and fluctuating current.** (A) A superthreshold dc current pulse (150 pA , 900 ms , middle) evoked trains of action potentials in a layer V pyramidal neuron. Responses are shown superimposed (10 trials, top), and as a raster plot (25 trials, bottom). (B) Same cell is stimulated repeatedly, this time with a fluctuating stimulus (a Gaussian white noise, with a 150 pA mean, 100 pA std, filtered with an alpha-function with $\tau_\alpha = 3$). Taken from [110].

fluctuations, electrical ones or ion concentration fluctuations. The effects of these different noise sources have been studied experimentally and theoretically, trying to connect the biophysical properties to the operational ones [96, 97, 195, 178, 114, 115, 196].

We emphasize the difference between these internal sources of noise and the variability in the function of a neuron which stems from the input that it receives – the spike trains of many thousands of neurons. Since each of these neurons is noisy and the synapses are also stochastic, the input to the neuron is unreliable as well. Thus, if a stimulus is presented to the animal (or system) repeatedly in an identical way, the synaptic activity that the soma ‘sees’ will still be different from one presentation of the stimulus to another. The effect of the synaptic input noise on the single neuron function has been studied extensively, reflecting it may be highly significant in setting the neuronal response to its inputs [200, 169, 65] (note also the history dependent activity of synapses [118, 181, 3, 37]).

Aside from the synapses, the most prominent internal noise source of the neuron, has been argued to be the ion channel fluctuations [96, 114], that will be at the focus of large parts of our work. As we shall discuss in the following chapters, this noise may play a significant role in determining the way information is processed and transmitted by neurons.

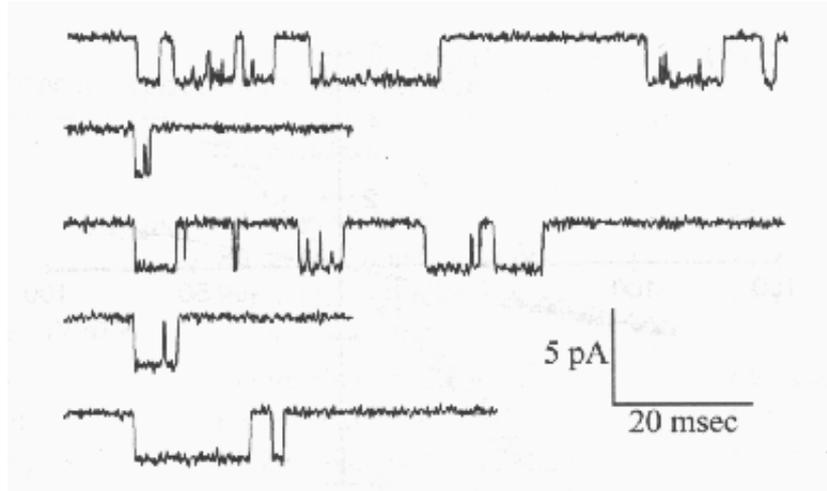


Figure 2.9: **Current flowing through a single ion channel.** Patch clamp recording of a single ion channel, held at a constant membrane potential of 0 mV . Opening of the channel (downward events), cause a 3 nA current flow. The random opening and closing of the channel is typical of all ionic channels. Baseline fluctuations are due to thermal noise. Taken from [88].

2.3.2 Ion Channel Noise

The development of the *patch clamp* recording technique [125] made it possible to observe the activity of single channels in an intact small patch of membrane, and record the currents that flow through them. Once a particular type of channel is isolated in the patch, it is possible to measure its kinetics and conductivity.

Channels open and close randomly, as one can see from the fluctuations of the membrane current. From single channels recording we learn that when activated, they may open at random intervals. Channels may also fluctuate between different open states which have different conductance levels. An example of single channel recording, reflecting the stochastic nature of channel opening and closing, and the much smaller effect of thermal noise on the membrane currents is shown in Figure 2.9.

Single channels (usually) open and close abruptly, and their kinetics are frequently described by a state diagram, as in chemical kinetics. The kinetics of single channels may be quite complex. This complexity is assumed to arise from a wide set of possible conformational and structural changes that the channel goes through as a function of membrane voltage changes, or the binding of ligands. Consequently, channel kinetics are usually modelled using a rather intricate (usually first order) Markovian kinetic model that takes into account multiple open and closed states, including 'inactivated states' (states following an open state, in which the channel is unable to open). A simple example of such a diagram would be,

$$\textit{closed} \rightleftharpoons \textit{open} \rightleftharpoons \textit{inactive} \quad (2.17)$$

with one first-order transition between the states (implying that the macroscopic permeability changes follow a single exponential relaxation after a step change, see [64]. We neglect the rate constants between the states for clarity).

Fitting such models suffers from the usual problems of model selection, as there is no clear indication of the actual number of states that the channel may have (it cannot be derived using the current knowledge of three dimensional structure of the channels in the membrane). The multiplicity of closed and inactive states is often seen as multiple kinetic time constants in the gating currents, in fluctuation measurements and in the open/close time histograms of single channel recordings. Various learning theory techniques have been utilized to find efficient and principled models (see e.g. [186, 82]). It is conceivable that the number of states is very large [177], or that channel behavior has memory of past states (i.e., the Markovian order may be larger than 1).

Although it is clear that channels are stochastic elements by nature, it has usually been assumed (and so implemented in most neuron models) that it is justified to use deterministic descriptions of the conductances arising from many channels of a given type. The reasoning is that because of the large number of channels of each type, and because the channel act independently of each other (which they do, to a good approximation), then, from the law of large numbers, the fraction of channels open at any given time is approximately equal to the probability that any one channel is in an open state. Hence, most neuron models use deterministic descriptions of the conductances arising from many channels of a given type.

Following Hodgkin and Huxley's physical interpretation [68, 71], and Fitzhugh's work [50], a stochastic version of the HH model, which is based on simulating ion channel noise has been studied in detail as a more biophysically accurate description of the spiking mechanism [163, 42, 31, 170], and is presented in chapter 6.

2.4 The Computational Sketch of a Single Neuron

We conclude this section by presenting a simplified view of the biophysical computation that the single neuron performs. As shown schematically in Figure 2.10, a single neuron may receive spike trains as inputs from many excitatory and inhibitory neurons.

These spike trains are transformed by the synapses into post synaptic potentials in the dendrites and the soma of the receiving neuron. The post-synaptic potentials are merged in a nonlinear way by the dendritic tree with its active and passive conductances, and reach

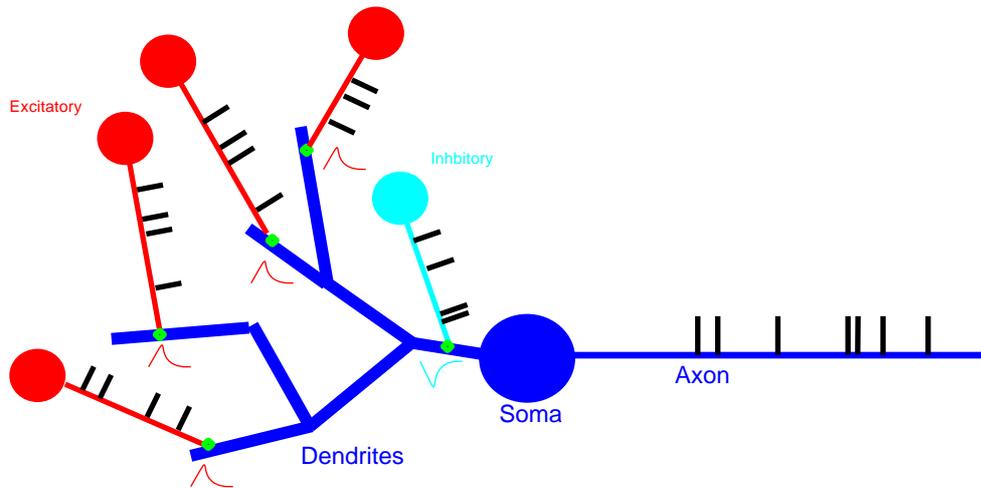


Figure 2.10: Sketch of a single neuron computation.

the soma of the neuron as a voltage trace. The many thousands of channels in the soma and the axon hillock then 'encode' this voltage trace into a sequence of action potentials that propagates down the axon to neurons 'down the road'. The output of the neuronal computation of the cell is a single spike train, but how is information encoded by this pattern of spikes?

Chapter 3

Background: The Neural Code – Spike Trains, Coding and Information Theory

Adrian's work on sensory neurons [8] formed the principles of our understanding of the *neural code*. First, he found that spiking is an all-or-none phenomena, occurring (or not) in response to incoming stimuli. As no intermediate response was found, the spiking patterns were assumed to be the sole carrier of information. Second, the number of spikes that a cell produces is related to the nature of the stimulus. Third, in response to a static stimulus, the spike rate declines (over time).

How do spike patterns encode and represent information? This is still a central question in neuroscience, usually called the problem of the neural code. The study of the neural code can be roughly divided into a few interconnected questions: What do the spikes encode? How do they encode it? how efficient is the code? How susceptible is it to noise? Is it optimal in some sense?

This chapter presents the main questions regarding the nature of the neural code and some of the answers to these questions. We present the standard tools of analysis of neuronal spike trains and the concepts of information theory, which has been a significant tool in the analysis of spike train content and nature.

3.1 Quantifying the Neural Responses

Although action potentials can vary in their amplitude, duration, and shape, these differences are small enough so that spikes are typically treated as identical stereotyped events.

Ignoring the brief duration of an action potential (which is about 1 msec), an n action potential sequence is usually characterized simply by a list of times when the spikes occurred $\{t_1, t_2, \dots, t_n\}$, or $\{t_i\}$ in an abbreviated form. Often, the sequence will be represented as a sum of idealized spikes, using the *Dirac delta function*, i.e.

$$\rho(t) = \sum_{i=1}^n \delta(t - t_i) \quad (3.1)$$

$\rho(t)$ is sometimes called the *neural response function*.

The *spike count rate* of a window of time T of the spike train, is defined as the number of spikes that occur during the time window, divided by the window length

$$\bar{r} = \frac{n}{T} = \frac{1}{T} \int_0^T \rho(t) dt \quad (3.2)$$

3.1.1 Spike train statistics

Describing the connection between the stimuli and a neuron's response, requires the construction of some form of a 'dictionary' between the two. I.e. we wish to portray the mapping that the neuron performs from any of the stimuli, $s(t)$, into the set of possible spike trains. The top panel of figure 3.1 shows the responses of a specific neuron in the fly visual system to repeated presentations of the same visual stimulus $s(t)$ (see figure caption for details). Clearly, there is no unique response and so apparently, the mapping of a stimulus to a spike train is a stochastic one (see also figure 2.8). It is then possible to define the conditional probability that a neuron will respond to a stimulus $s(t)$ with a specific spike train $\{t_i\}$, as $P(\{t_i\}|s(t))$. Using *Bayes' rule* and proper definitions of the ensemble of stimuli $P(s(t))$ and of the responses $P(\{t_i\})$,¹ it is possible to define the conditional probability of the stimulus given a response, $P(s(t)|\{t_i\})$. These conditional probability distributions, are the central tools in the analysis of spike trains and the nature of the code.

Since the spike train that the neuron responds with to a given stimulus, typically varies from trial to trial, it is customary to use the aforementioned probability densities to characterize the average behavior of the neuron. As the spike times are continuous, one needs to use probability density function for the spike timing, and so the average firing rate at time t , over a window of size Δt , is

$$r(t)\Delta t = \int_t^{t+\Delta t} \langle \rho(\tau) \rangle d\tau \quad (3.3)$$

where $\langle \rangle$ is the average over trials (input presentations). The top panel of figure 3.1, shows the different responses of the H1 neuron in the fly visual system, during an experiment in

¹We put aside the issue of defining the appropriate probability spaces, see e.g. [143].

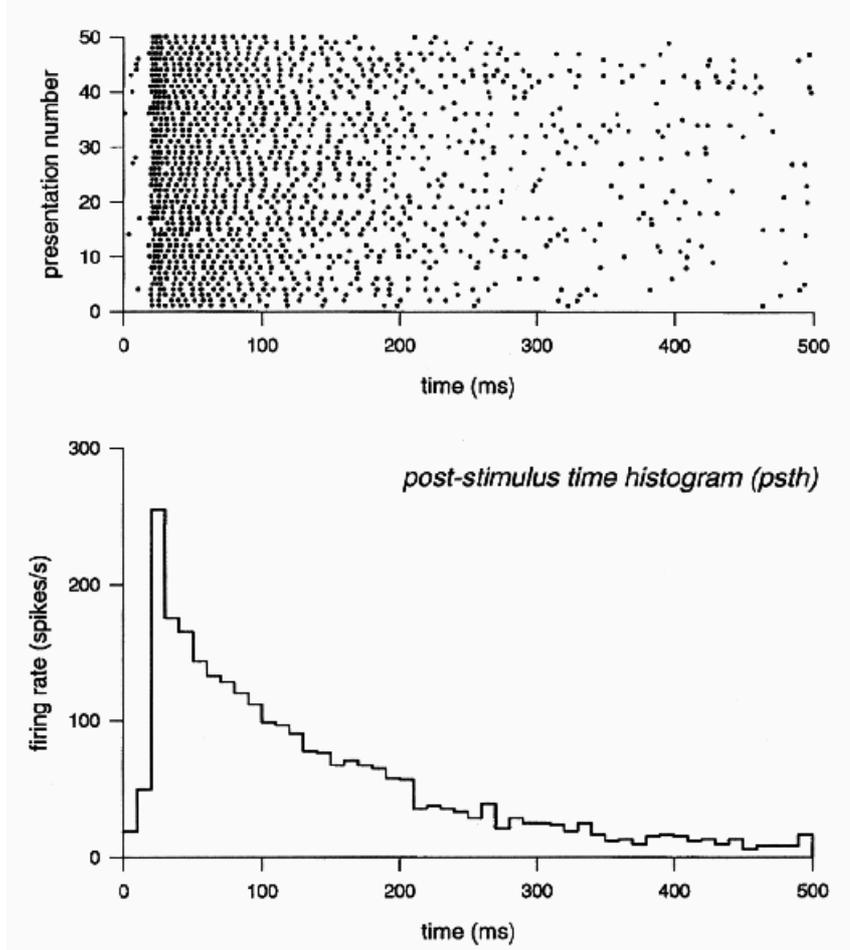


Figure 3.1: **Variability of neural responses and construction of the average response.** Top panel shows a raster plot of 50 individual spike trains of the motion sensitive H1 neuron in the fly visual system, in response to repeated presentation of the same visual stimulus at time $t = 0$. Bottom panel shows the Post Stimulus Time Histogram (PSTH) of the responses. The PSTH is calculated for consecutive time bins, by taking the average number of spikes in each bin (10 msec in this case), over stimulus presentation, and normalizing it by the bin size. Taken from [143]

which the same visual stimulus was presented repeatedly to the fly. The bottom panel shows the averaged firing rate of the H1 neuron calculated from the responses in the top panel, and is known as the *post stimulus time histogram (PSTH)* of the neuronal responses.

An alternative way to represent the spike train is to use the sequence of intervals between the spikes, i.e., $\{\delta t_1, \delta t_2, \dots, \delta t_{n-1}\} = \{t_2 - t_1, t_3 - t_2, \dots, t_n - t_{n-1}\}$. A frequently used measure of the nature of the spike train is the distribution of interspike intervals that a neuron uses, in response to a specific stimulus (*ISI distribution*). The ratio between the first two moments of the distribution, μ_{ISI} (the mean interspike interval), and Var_{ISI} (the variance of the distribution), namely $\frac{\mu_{ISI}}{Var_{ISI}}$, is termed the *coefficient of variation (CV)* of the interspike interval (*CV of ISI*).

Theoretical models of spiking often refer to spike trains as *point processes* (statistical models of a sequence of events). Due to the random-like nature of *in vivo* spike trains, it is often assumed that one spike is independent of the other spikes, in which case the spike times follow a non-homogeneous Poisson distribution. Although it is known that this assumption is incorrect, the Poisson model is a frequent approximation of spiking patterns, and considered to be a fairly reasonable one. According to the Poisson distribution, the probability to find k spikes within a time window of size Δt , is given by

$$P(k \text{ spikes}) = \frac{(\lambda \Delta t)^k \exp(-\lambda \Delta t)}{k!}$$

where λ is the average firing rate of the neuron per unit time. The probability of an interspike interval falling between t and $t + \Delta t$ is given by $r \Delta t \exp(-\lambda t)$. As can be easily verified, the *CV* of *ISI* of a Poisson spike train is 1. The value of *CV* of *ISI* of real neurons spike trains is then compared to the value of the idealized Poisson value, quantifying the randomness of the spike train (for example, the *CV* of an “ordered” spike train with fixed interspike interval equals 0).

Similar to the interval distribution, the *autocorrelation* function quantifies the probability that two spikes will occur within a certain separation, regardless of the events (spikes) in between. It is thus a useful tool to identify patterns of activity of the neuron (e.g. oscillations). Using the definitions of the neural response function and the average spike count, the autocorrelation of a spike train is given as function of the time difference between the spikes τ ,

$$C_{\rho\rho}(\tau) = \frac{1}{T} \int_0^T \langle (\rho(t) - \bar{\rho})(\rho(t + \tau) - \bar{\rho}) \rangle dt \quad (3.4)$$

where the average is over stimulus presentations. The autocorrelation is easily generalized to the correlation between the spike trains of two different neurons.

The response of a neuron may depend on many different properties of the stimulus it is presented with. Quantifying the relation between the stimulus and the response is a question of identifying the features of the stimulus that the neuron ‘cares’ about. Extracting the encoded features of a system is a difficult problem (see e.g. [21]). In neuroscience, trying to identify what neurons ‘care about’ is usually translated into mapping their *receptive field*², i.e. the set of stimuli that the neuron responds to (see e.g. [37]). Often, the receptive field properties will be summarized by the *spike-triggered average stimulus*, $A(\tau)$, which is the average value of the stimulus over a time interval τ before a spike is fired (the averaging is over spikes and stimulus presentation),

$$A(\tau) = \left\langle \frac{1}{n} \sum_{i=1}^n s(t_i - \tau) \right\rangle \quad (3.5)$$

²following to the seminal work of Hubel and Weisel [77] on the receptive field of cortical visual neurons

Figure 3.2 shows an example of the construction of a spike triggered average. It is straightforward to define similar averages of the stimulus conditioned on different response features (see e.g. [143]). We note that spike triggered averaging is a simple case of the more general

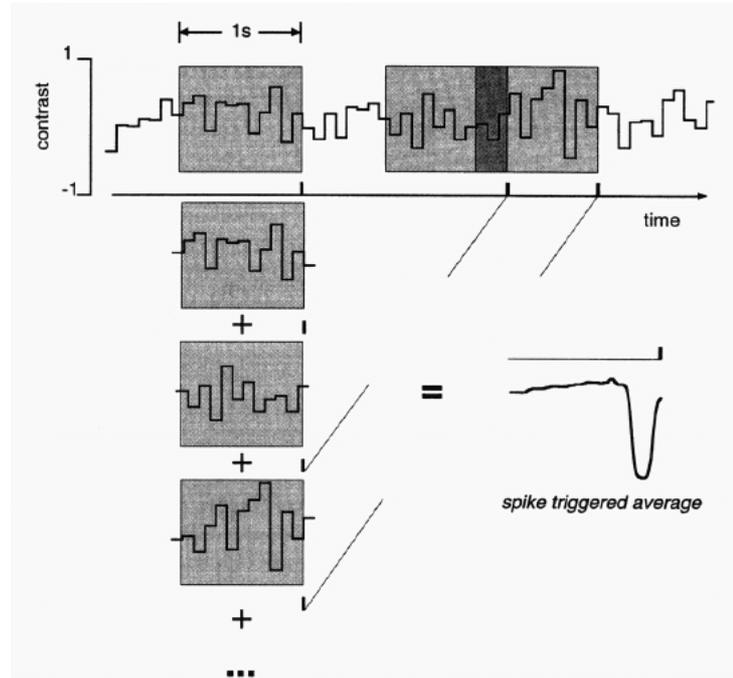


Figure 3.2: **Spike triggered average.** A sequence of light intensities which are projected on a salamander retina, is shown on top, with the resulting spike train of a retinal ganglion cell below it. The spike triggered average is constructed by averaging the stimulus waveforms preceding each of the spikes. The result is shown on the bottom right. Taken from [143].

reverse correlation method, which is a standard tool in system identification, connecting between input and output features (see [120, 143]).

3.2 The Nature of the Neural Code

3.2.1 Rate Codes and Temporal Codes

While it is clear that action potentials convey information through their timing, the nature of the neuronal codewords and their 'meaning' has been a controversial issue in neuroscience. The main debate over the nature of the neural code has focused on whether the fine temporal structure of the spike train carries information or whether the details of the fine structure are mainly noise (see e.g. [143, 158]). Obviously, the nature of spike trains as a code sets the capacity of the code to convey information, and reflects on how it could be decoded.

Much of the discussion has focused on presenting examples of neurons using a *rate code* or a *temporal code*, often without a clear definition of what these terms mean.

The ‘classical’ view of the neural code is that of *rate coding*, which stems from the empirically observed unreliability of neuronal spike trains. Since it is evident that the temporal structure of the spike train is noisy, it has often been argued that the sole carrier of information is the spike rate. The frequent version of rate coding is actually a *spike counting code*, in which the rate is defined as the number of spikes in a time window divided by the length of the window. The rate for every point time is usually estimated using the (somewhat oxymoronic) ‘instantaneous firing rate’, defined either by using a moving-window spike-counting code or by using the inter-spike interval and defining the instantaneous rate as $\frac{1}{ISI}$.

The concept of *temporal coding* arises when we consider how precisely we must measure spike times to extract most of the information from a neuronal response. This precision determines the temporal resolution of the neural code. Several studies have found that this temporal resolution is on a millisecond time scale or less, indicating that precise spike timing is a significant element in neural encoding (for example, [40, 89, 20, 143]).

3.2.2 Noise, Reliability and the neural code

Another approach in the investigation of the nature of the neural code is to connect it to the computation that the neuron performs and the way the computation is implemented. One common view of a neuron is that of an integrator of inputs. Spiking is then the result of summing of many random-like (or noisy) events. In this case, temporal information in the inputs to the neuron cannot be kept in the outgoing spike train. The alternative view suggests that the neuron may function as a coincidence detector, for coordinated spiking of a small number of input neurons (see e.g. [4]). In this case, temporal information is very likely to persist in a network, if it concerns synchronized activity of neurons (Obviously, a key issue in this respect is how many synaptic potentials are needed to create a single spike [130]. The highly irregular nature of the interspike intervals of cortical cells has been argued as an evidence for both perspectives.

Softky and Koch [164] argued that the convergence of many Poisson-like excitatory inputs onto a single integrating neuron with a membrane time constant of the order of 10 milliseconds, would result in a much more regular spike train than the usual Poisson-like spike train usually seen in cortical single cell recording. Therefore, their conclusion was that cortical cells must function as coincidence detectors, rather than integrators, and that the effective time constant is of the order of 1 – 2 milliseconds. Shadlen and Newsome [158] demonstrated that massive uncorrelated input of balanced excitatory and inhibitory contri-

butions from Poissonic-like sources (Softky and Koch neglected the inhibitory inputs) would yield a Poissonic spike train in the target cell. They argued that this result refutes the possibility of precise temporal encoding.

Balanced excitation and inhibition inputs to a neuron have been argued to result in an effective short membrane time constant [165, 159, 185], leaving the question of coincidence detection somewhat unresolved. Other possible sources for the high CV of the ISI distribution have been suggested, like the nature of repolarization after a spike, relative refractoriness, and membrane voltage instability [16, 180].

3.2.3 Rate code and Temporal code, revisited

According to the above formulation (section 3.2.1), temporal encoding might seem to be a special case of rate encoding in the limit of a small encoding time window. However, the temporal structure of a spike train or firing rate evoked by a stimulus is determined both by the dynamics of the stimulus, by the nature of the neural encoding process and the internal noise of the neuron. Stimuli that change rapidly tend to generate precisely timed spikes and rapidly changing firing rates no matter what neural coding strategy is being used (see below).

A generalized (and more accurate) view of the temporal vs. rate code problem has been given (in different ways) in [1, 175, 143, 25]. They all suggest that the question is not just a matter of the size of the encoding time window or temporal resolution, but rather a question about the symbols of the code and their content. Using very small time windows to analyze the nature of the spike train, only improves the time scale over which the rate code is considered – it is still a weighted average of the number of spikes over the encoding time window. Temporal coding should refer to the existence of information in the temporal structure of the spike train, beyond the information carried by the modulated firing rate of the neuron! (In other words, temporal coding should refer to temporal precision in the response that does not arise solely from the dynamics of the stimulus, but relates to properties of the stimulus).

3.2.4 Experimental results of input-dependent unreliability, noise and code

Clearly, the reliability and temporal accuracy of spike timing are fundamental to the nature of the neural code. Accurate spike timing can carry more information than a fuzzy one, and may imply a different decoding scheme to be used by the nervous system. The work of Bryant

and Segundo [26] and the more recent work of Mainen and Sejnowski [110] addressed this question directly, by repeatedly injecting identical current traces into the soma of neurons (while blocking synaptic transmission) and thus trying to characterize the reliability and accuracy of the neuronal spike trains. They have shown that neurons respond to fluctuating current inputs with repeatable and accurate spike trains, whereas slowly varying inputs result in lower repeatability and ‘jitter’ in the spike timing. Figure 2.8 (in the previous chapter), shows the responses of a layer V cortical pyramidal neuron from a rat to two opposite types of input currents, injected repeatedly to the neuron using an intracellular electrode. DC current injection results in very low reliability of the spike timing, reflecting a noise source dominating the timing of the spike. Since synaptic transmission was blocked, this suggests that the noise source is a cellular one. The injection of a highly fluctuating current, with the same mean value as the DC current, resulted in a similar number of spikes, which are timed very accurately across repeated trials.

Thus, the same neuron can be extremely accurate in terms of spike timing, having a precision of less than *1msec* for some inputs, and yet be very inaccurate for other inputs. This result suggests that the nature of the neural code (the outgoing spike train) is set by the nature of the stimulus and the internal noise of the neuron (which affect the neuronal computation)³. Figure 3.3 shows the similar results of [127], which also include intermediate reliability in response to different types of current stimuli.

Similar input-dependent spike train reliability and accuracy, have been reported for neurons *in vivo* when the animal was presented with a dynamic or static stimulus, and the response of a single neuron was recorded [13, 39, 139].

3.2.5 Population Codes

Until now, we have described the neural code question and analysis, in terms of the code of a single neuron. Obviously, information is typically encoded in the nervous system by populations of neurons. Population coding has several possible advantages from an information encoding viewpoint, including reduction of uncertainty due to neuronal noise or different feature selectivity, and the ability to represent a number of different attributes of a stimulus simultaneously. Individual neurons in such a population typically have different but overlapping selectiveness so that many neurons, but not necessarily all, respond to a given stimulus.

³Some arguments about experimental problems which may be the source of these result have been raised in past years. It has been suggested that instrumental noise or unreliability and electrode problems may be the reason for the unreliability of spike timing in the Mainen and Sejnowski experiment. However, the same results have been reported by different groups (see e.g. [127]) and different preparations.

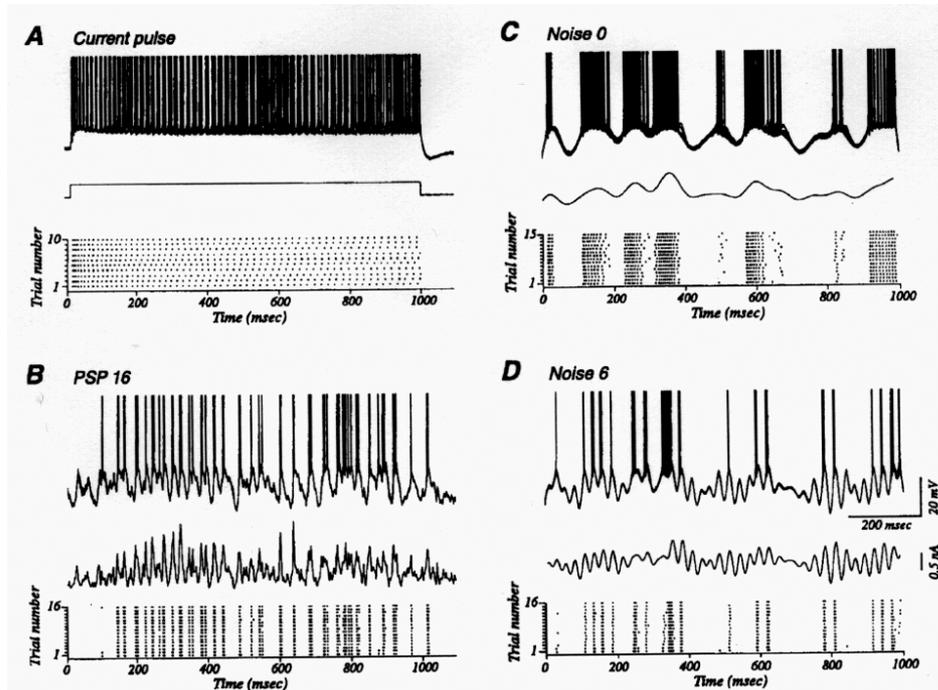


Figure 3.3: **Reliability and jitter of spike timing in a regular spiking cell.** Similar to the results of [110], shown in Fig. 2.8, the responses of a cortical neuron to repeated injection of different types of currents is shown. (A) In response to a dc current (middle), the unreliable spike patterns (superimposed, top; raster plot, bottom) is shown. (B) Response of the same cell to a current waveform derived from a visually evoked response. (C) As in B, input used is an artificial fluctuating current dominated by low frequency components. (D) As in C for a current dominated by high frequency component. Taken from [127]

Thus, in general the *population code* may have a mixture of the coding schemes. First, the cells may respond to different features in a common stimulus. Second, they may respond independently, given the stimulus, or thirdly, they may encode things in their joint response things about the common input (sometimes in a synergistic way,[53, 52, 131]).

Due to experimental limitations of multi-cell recordings, and since it is easier to analyze the responses of a population of neurons as independent. This has been the frequent (but not the only) approach in population decoding studies, which mainly deal with estimating an average value of the neuronal population [54, 157, 147, 148, 198, 37, 201]. But how could the correlations between the spiking of different neurons provide additional information about a stimulus that cannot be obtained by considering all of their firing patterns individually? Synchronous firing of two or more neurons is one mechanism for conveying information in a population correlation code [184, 89, 141]. Rhythmic oscillations of population activity provides another possible mechanism [90]. Both synchronous firing and oscillations are common features of the activity of neuronal populations.

One possible comprehensive view is given by the *synfire chain* model [5], which suggests that synchronized patterns of neuronal activity are the key mechanism of information encoding and propagation in the cortical networks (see also [44, 123]).

3.3 Information theory and the neural code

Information theory was formulated by Shannon [160] as a general framework for quantifying the information conveying ability of communication systems. The general form of a communication channel (be it an optic fiber line, spoken language or a neuron), is of a sender, a receiver and a physical media over which the sender conveys signals to the receiver: The sender selects a particular message it wishes to send, out of a set of possible messages. He then uses an encoder to convert the message into a set of symbols which will go over the physical channel. When arriving on the other end, the message is decoded by the receiver. In general, encoding (decoding) and transmission are stochastic and noisy processes, and so the analysis of the nature of the communication channel relies on the conditional probability densities of an “output” at the receiving end as a function of the “input” the sender wanted to send (and vice versa). The key concepts of information theory, the *entropy* and *mutual information*, depend on the probabilities with which the symbols and their combinations are used across the communication channel. Entropy measures the richness of the patterns of coding symbols and the mutual information measures how much does the output of channel, i.e. the decoded message, convey about the input to the channel, i.e. the original message of the sender.

In electrical and communication engineering, information theory evaluates the capacity of communication channels and searches coding schemes that saturate them. However, due to its mathematical rigor, and parameter-free nature, the concepts and tools of information theory have been imported to many other fields. Mutual information serves as a generalized measure of correlation between variables, an unbiased measure of the effect of changing parameters on the outcome of an experiment (used for example in psychology and biology). Maximum entropy and information maximization concepts are common in design and optimization problems (e.g. in game theory, economics and decision theory. Information theory has also been used in the study of computation and the dynamics of physical systems, since in a simplified way, every input-output “machine” is a communication channel (see e.g. [35]).

In neuroscience, information theory has been applied to questions regarding the nature of the neural code [143, 22, 37], as we discuss below. It has also been used for quantifying the nature of jitter in the propagation of action potentials along the axon [6], the analysis of the nature of a synaptic connection function and synaptic learning rules [102, 103, 116, 29],

the study of feature encoding [173]. Especially interesting are the ideas of maximization of encoded information and efficient information encoding as design principles of the nervous system [15, 11, 17, 102, 103].

3.3.1 Entropy, Relative Entropy and Mutual Information

The entropy of a set of responses of a system, measures how rich (or “surprising”) the set of responses is. If a system is presented with a variety of stimuli and responds in the same way every time, or that only a few different responses appear, we might conclude that this set of responses is uninteresting. A richer set would show a larger range of different responses, perhaps in a highly irregular and unpredictable sequence. The Shannon entropy quantifies the intuitive notion of ‘how surprising the response is’ as a function of its probability. Intuitively, an information measure should be positive, anticorrelated with the probability of a response and additive for two independent observations. The only functions that obey these requirements are logarithmic (see [7]).

Thus, if we consider the response as a random variable, the entropy of a discrete random variable X which may take K possible values in Ω was defined by Shannon as

$$H(X) = - \sum_{x \in \Omega} p(x) \log_2 p(x) \quad \text{bits} \quad (3.6)$$

where $p(x)$ is the probability of X taking the value x . The entropy is a measure of how much is unknown about the value of the random variable. The usual physical intuition is that the entropy of a system is the logarithm of the number of possible states that the system can occupy.

The entropy of a pair of discrete random variables X and Y (taking values in Ω and Υ , correspondingly), with a joint distribution $p(x, y)$ is called the joint-entropy and is defined as,

$$H(X, Y) = - \sum_{x \in \Omega} \sum_{y \in \Upsilon} p(x, y) \log_2 p(x, y) \quad \text{bits} \quad (3.7)$$

Based on the entropy measure, it is possible to quantify how much information does an output of a system (or a channel) Y , carry about the input X . Alternatively, this may apply for any two random variables, asking what does knowing the value of one variable reveals about the value of the other. If X is chosen from some probability, $P(X)$, then the entropy of this distribution is a measure of the surprise (or uncertainty) of the value of X . Once Y (the decoded received message) is observed, the range of possible inputs which were the source of it is restricted, according to the conditional distribution of X given Y . Consistent with the notion of entropy as a measure of uncertainty about the nature of X , the conditional

entropy of Y given X measures how much of the uncertainty about the nature of X has remained after seeing Y , and is then defined to be

$$\begin{aligned} H(X|Y) &= \sum_{y \in \Upsilon} p(y) H(X|Y = y) = \\ &= - \sum_{y \in \Upsilon} p(y) \sum_{x \in \Omega} p(x|y) \log_2 p(x|y) \\ &= - \sum_{y \in \Upsilon} \sum_{x \in \Omega} p(x, y) \log_2 p(x|y) \quad \text{bits} \end{aligned} \quad (3.8)$$

The similar expression of the conditional entropy of the output Y given the input X ,

$$\begin{aligned} H(Y|X) &= - \sum_{x \in \Omega} p(x) \sum_{y \in \Upsilon} p(y|x) \log_2 p(y|x) \\ &= - \sum_{x \in \Omega} \sum_{y \in \Upsilon} p(x, y) \log_2 p(y|x) \quad \text{bits} \end{aligned} \quad (3.9)$$

(In our context this would be considered as a 'noise entropy' term, as it measures what is unknown about the response when the input is known).

The mutual information between X and Y is defined as the difference between what was unknown about X prior to observing Y and the uncertainty about X left after observing Y . Measured in bits, the mutual information is symmetric in X and Y , and is given by

$$\begin{aligned} I(X; Y) &= \sum_{x \in \Omega} \sum_{y \in \Upsilon} p(x, y) \log_2 \frac{p(x, y)}{p(x)p(y)} = \\ &= H(X) - H(X|Y) = H(Y) - H(Y|X) \end{aligned} \quad (3.10)$$

Figure 3.4 gives a *Venn*-like diagram of the relation between the entropies and the mutual information of two random variables.

The relative entropy or *Kullback-Leibler divergence* between two probability distributions, $p(x)$ and $q(x)$ is defined as,

$$D_{KL}(p(x)||q(x)) = \sum_{x \in \Omega} p(x) \log_2 \frac{p(x)}{q(x)} \quad \text{bits} \quad (3.11)$$

and can serve as a measure of the dissimilarity of the two distributions. Using *Jensen's* inequality, one can see that $D_{KL}(p(x)||q(x)) \geq 0$, and is equal to zero only when $p(x)$ is identical to $q(x)$. It is an asymmetric and unbounded measure. Using the Kullback-Leibler divergence, we can rewrite the mutual information between X and Y as

$$I(X; Y) = D_{KL}(p(x, y)||p(x)p(y)) \quad (3.12)$$

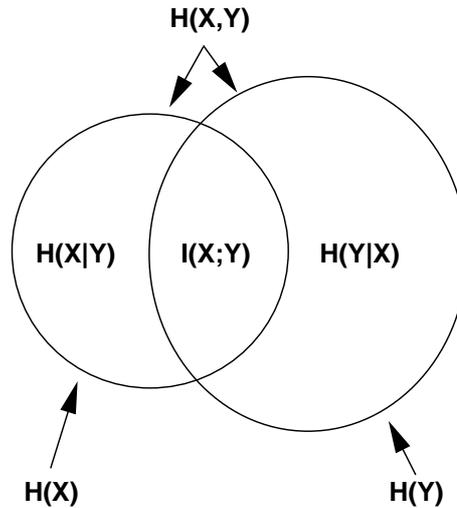


Figure 3.4: Venn-like diagram of two random variables' entropies and mutual information.

3.3.2 Spike train entropy and information rates

In the context of neural codes, information theory has mainly been used to quantify the relation between the input to the neuron (or system), and the response of the neuron. The symbols of the code are the action potentials, and so the spike timings sequence is the usual description of the response. Since the space of possible stimuli (or inputs) is usually infinite, such an analysis is limited only to the set of stimuli that are used. The choice of stimuli and the probability distribution of their presentation will determine the nature of the set of responses. The paradigm is that a set of stimuli $\{s_1(t), s_2(t), \dots, s_m(t)\}$ is presented repeatedly, and the neuronal responses are registered. Thus, for each stimuli i , we have a set of responses of the neuron, $\{r_1^i(t), r_2^i(t), \dots, r_n^i(t)\}$.⁴ The questions that can be addressed are what is the conditional entropy of the spike train of the neuron with respect to some class of stimuli? What is the mutual information between the stimuli and the spike trains? How efficient is the code or how much of the spike train is noise? How reliable is the code? Since each spike costs energy, how efficient is the code energy-wise? etc.

In practice, computing the entropy and mutual information for spike sequences can be difficult, since the probability distribution of the spike sequences and the joint (or conditional) distribution of the spike train and the stimuli must be estimated reliably. Spike train entropy calculations are typically based on the study of long-duration recordings consisting of many action potentials, which is experimentally demanding.⁵ Since the entropy of a se-

⁴As before, we take the neuron's point of view, so even when the stimulus is presented to the animal, we regard it as presented to the neuron

⁵Consequently, many information theory analyses use simplified descriptions of the response of a neuron

quence of symbols typically grows linearly with the length of the sequence, it is customary to compute entropy or information rates. These entropy and information rates are the total entropy or information divided by the duration of the spike train. Alternatively, entropy and mutual information can be divided by the total number of action potentials and reported as bits per spike rather than bits per second.

Several different methods of calculating spike train entropy (and information) have been used since the work of MacKay and McCulloch [107], who estimated the entropy of spike trains by binning them into binary sequences (and so deduced a bound on the capacity of neurons as information conveyers). Werner and Montcastle [193], Eckhorn and Popel [45, 46], and Optican and Richmond [128] have estimated the joint distribution or the conditional distribution based on a limited and simplified set of stimuli and responses of the neuron. Bialek, de Ruyter, Warland and Rieke [20, 143] have used Wiener-Volterra kernels [149] (and their causal approximations) to reconstruct the stimulus out of the spike train and thus give a lower bound on the mutual information. Heller *et al.* [63] (see also [192]) have used artificial neural networks to estimate the mutual information between the stimulus and spike trains. The context-tree-weighting algorithm [197] has also been used both for spike train entropy calculation [104]. The difficulty in estimating the entropy and information rates which leads to overestimation has been addressed by various resampling techniques (for example, [129]), approximations and bounds (e.g., [179]). We present here the technique of Strong *et al.* [171], for estimating the spike train entropy, the spike train noise entropy, and the mutual information between the stimulus and the spike train (for comparison of the methods see [28] and [192]).

Computing spike train entropy and information rates

The entropy of the spike train is a measure of the richness of the 'vocabulary' of the neuron and gives a bound on how much information the code could carry, if there was no noise. To make entropy calculations practical, a long spike train is broken into statistically independent subunits, and the total entropy is written as the sum of the entropies of the individual subunits. The spike train is discretized in Δt bins, using a sliding 'window' of size T along the discretized sequence. The train of spikes is thus transformed into a sequence of k -letter 'words' ($k = T/\Delta t$), consisting of 0's (no spike) and 1's (spike), as shown in Figure 3.5.

Clearly, any fixed choice of Δt is arbitrary. Both the spike train and noise entropy rates depend on Δt – since more information can be extracted from accurately measured spike

that reduce the number of possible symbols (i.e. responses) that need to be considered. For example, the spike timing may be replaced by the average firing rate.

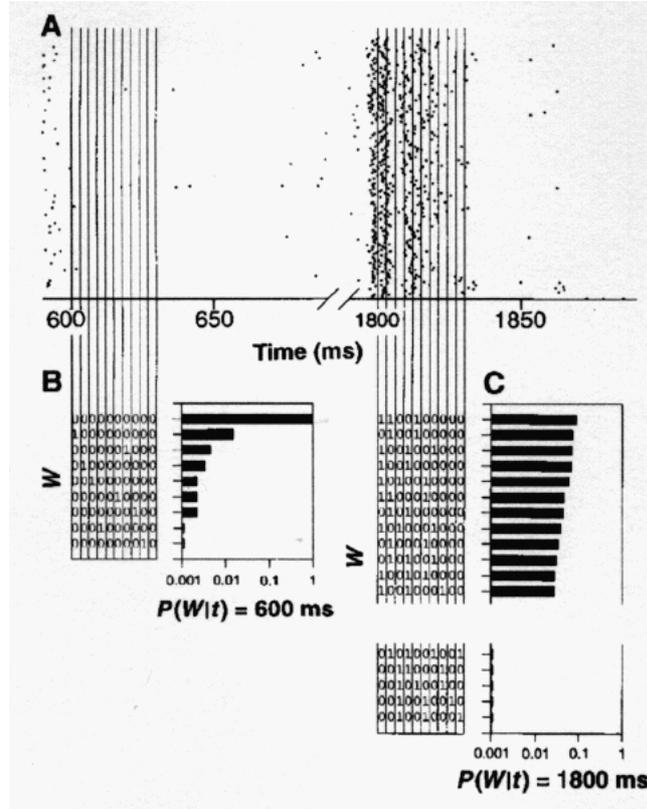


Figure 3.5: **Example of binning of spike trains and transforming them into binary word distribution.** (A) Two segments of multiple response traces of a neuron, (recorded *in vivo* from the fly H1 neuron). (B-C) Using a fixed binning time window Δt (of size 3 msec in this case), the spike trains are transformed into sequences of bins, which form 'words' (10-letter words in this case). The number of spikes in each bin (of each of the trials) is counted, forming a word W . For each point in time it is then possible to estimate the probability of observing a certain word (based on the multiple stimulus presentation). Adapted from [39].

times than from poorly measured spike times. Thus, we expect the information rate to increase with decreasing Δt , at least over some range of Δt values [25]. Often a range of Δt values is explored, usually chosen to be small enough so that the transformed spike trains are binary sequences, and large enough so that the word distribution is well sampled. Typical values will range between 0.5 msec to 5 msec .

By using a long enough spike train, we can get a reliable estimate of $P(W)$, the probability of the word W to appear in the spike train, and then compute the entropy rate of the total word distribution,

$$H_{total}^T = -\frac{1}{T} \sum_W P(W) \log_2 P(W) \quad \text{bits/sec} \quad (3.13)$$

The difference between the entropy for finite word size and the true entropy for infinite word size, should be proportional to $1/T$ for large T [35, 171]. Therefore, H_{total}^T is estimated

as a function of T (which is here assumed to be in sec), which is increased to the point where the estimation of the probability distribution becomes difficult. The entropy is then estimated, as shown in figure 3.6 for the case of the fly H1 neuron responses, by extrapolating linearly to $1/T = 0$. The resulting H_{total} gives the entropy rate of the neuronal spike train [35, 39, 171].

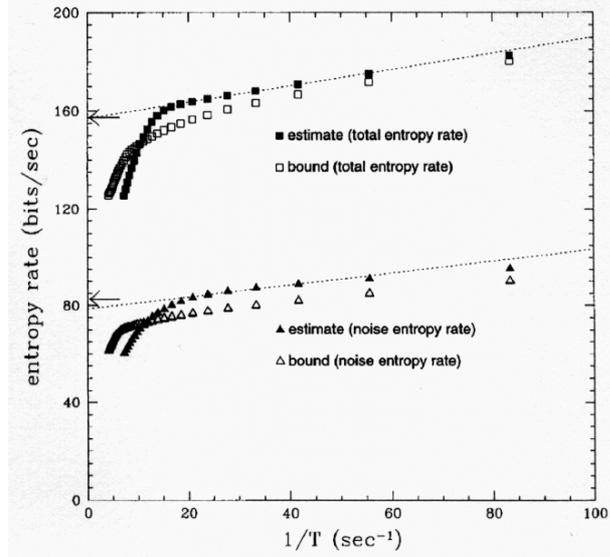


Figure 3.6: Entropy and noise entropy rates estimation for the H1 visual neuron in the fly responding to a randomly moving visual image. The activity of the motion sensitive neuron, H1, in the fly visual system was recorded in response to the repeated presentation of a randomly moving movie of a black and white bar pattern (see [171]). The spike train entropy and conditional entropies were calculated for different finite word sizes T , as explained in the text. The filled circles in the upper trace show the full spike-train entropy rate computed for different values of $1/T$. The straight line is a linear extrapolation to $1/T = 0$, which corresponds to $T \rightarrow \infty$. The lower trace shows the spike train noise entropy rate for different values of $1/T$. The straight line is again an extrapolation to $1/T = 0$. Both entropy rates increase as functions of $1/T$, and the true spike-train and noise entropy rates are overestimated at large values of $1/T$. At $1/T = .20s$, there is a sudden shift in the dependence. This occurs when there is insufficient data to compute the spike sequence probabilities. The difference between the y intercepts of the two straight lines plotted is the mutual information rate. The resolution is $t = 3$ ms. Taken from [171].

In order to calculate the mutual information we need to estimate the conditional probability of the word given the stimulus. If a long enough stimulus (which is assumed to sample “well” the possible stimulus values or features), is presented repeatedly to the neuron, we can estimate the conditional probability as a temporal average (see [171] for details). We then examine the set of words that the neuron uses at a particular time t over all the repeated presentations of the stimulus, and estimate $P(W|t)$, the time-dependent word probability distribution. At each time t we calculate the time-dependent entropy rate, and then take

the average of these entropies

$$H_{noise}^T = \left\langle -\frac{1}{T} \sum_W P(W|t) \log_2 P(W|t) \right\rangle_t \quad \text{bits/sec} \quad (3.14)$$

where $\langle \dots \rangle_t$ denotes the average over all times t ⁶. As before, taking the limit of H_{noise}^T at $T \rightarrow \infty$, gives the noise entropy rate of the spike train H_{noise} , which measures how much of the fine structure of the spike trains of the neuron is just noise (see figure 3.6).

The difference between the extrapolated total entropy rate and the noise entropy rate, gives the information rate that the spike train convey about the stimulus (with an errorbar resulting from the unreliability of the extrapolation, see [171]).

The spike train entropies and information rates of neurons have been estimated in different animals, modalities and in response to different types of stimuli [22]. The information rates conveyed about dynamic stimuli may range from a few bits per second to almost 300, whereas the average information encoded per spike, ranges between near 0 to over 3 *bits per spike*. (Since spiking is considered as one of the key energy consumption process in neurons, it is common to divide the encoded information rate by the number of spikes that were used for coding this information, and thus get energetic efficiency of the code. This assumes of course that the energy per spike is fixed.)

The range of information rates conveyed about dynamic stimuli may range from a few bits per second to almost 300, whereas the average information encoded per spike (which, assuming that every spike costs roughly the same amount of energy, measures the energetic efficiency of the code), range between near 0 to over 3 *bits per spike*.

3.3.3 Design principles of the code

Spike train entropy and information transmission have been used in 'first principles' calculations and theoretical analysis, trying to explain the nature of the neural code. Possible 'design principles' of the neuronal computation and coding, have been suggested to be the sparseness of the neural code, minimization of redundancy in the code and maximization of the amount of encoded information or information rate [15, 102, 103, 17, 11, 176]. Of special interest are studies of the possible optimization of the amount of information or information rate under metabolic constraints [95, 14].

⁶Averaging over time is equivalent to averaging over stimulus features as long as our stimulus is long enough to sample the stimulus features reliably, citeStrong-et-al-98

Chapter 4

Reliability and accuracy of spike timing in stochastic neuron models

4.1 Introduction

Following the formulation of the Hodgkin-Huxley (HH) equations for modelling spike initiation in the squid giant axon [72] (see 2.2.1), research into the electrical activity of single neurons followed two main paths: the attempt to discover further macroscopic equations governing different membrane currents (e.g., [199]), and the attempt to investigate, and mathematically describe, the behavior of the ion channels underlying these currents [64, 146]. Although part of the same general problem, mathematically these two areas of investigation are entirely different. In the Hodgkin-Huxley formulation, the ion conductances are modelled by means of deterministic differential equations, and their values range continuously from zero to a given maximum. However, because individual ion channels are discrete elements whose properties can only be given probabilistically, the electrical activity of nerve cells is most accurately described as resulting from the interaction of stochastic, discrete, units. It is commonly assumed that a large collection of such discrete units practically forms a continuous deterministic system, as is the case in numerous large physical systems. Because the number of excitable channels in the axon's spike initiation zone is estimated to be large (in the order of tens of thousands of ion channels, [64]), models for spike generation in neurons typically utilize deterministic rather than stochastic equations [111, 144, 138].

A few theoretical studies did, however, consider the effect of channel stochasticity, focusing on the question "when does the stochastic model converge to the corresponding deterministic model?". Following [71], Fitzhugh [50] suggested a kinetic model for the description of conductance change associated with the HH equations (and thus giving rise to stochas-

ticity); others used stochastic HH equations to investigate the effect of various parameters (e.g., number of channels, membrane area, etc.) on the dynamics of the membrane voltage. The main message of these studies is that the stochastic system differs considerably from the deterministic HH system when a small number (a few hundreds to a few thousands) of channels and small membrane areas are involved [42, 163, 170, 41]. Other related questions which were studied in this context are spontaneous spiking due to channel noise [163, 170, 30, 145], the effect of channel stochasticity on spike propagation in axons [145, 75, 76], and the effects on neuronal dynamics and subthreshold voltage [195] (see also the review in [196]). Correspondingly, several experimental [47, 195, 189], and theoretical [30, 96, 97, 114, 115] studies have indicated that noise from voltage-gated channels can have important effects at the cellular level (see also [18]). In a more general perspective, the effect of different kinds of noise on the firing threshold of neurons was examined by Lecar and Nossal [96, 97]. A more recent study has compared the contribution of different noise sources on synaptic signals propagating in dendrites [114, 115]. The effect of general additive noise (without specifying its biophysical origin) on the behavior of neuronal models has been studied computationally and theoretically, reflecting its effect on spiking patterns and oscillations [105, 23] and even on spiking reliability [84] and possible advantages in information encoding [33, 99].

Recent experimental studies addressed the question of reliability and accuracy of spike firing time in neocortical pyramidal cells. By blocking synaptic transmission and injecting current to the isolated neuron, Mainen and Sejnowski [110] (see also [26, 127] showed that spike timing is highly unreliable for repeated DC current inputs whereas for fluctuating current inputs the firing reliability is significantly improved. For highly fluctuating currents, they showed that the firing precision may go up to a millisecond range (see also [174]). Similar results have been presented for neurons *in vivo*, in response to static or dynamic stimuli presented to the animal [13, 139, 39].

We ask whether for realistic membrane area and number of excitable channels, a biophysically-inspired noise that is generated by channel stochasticity plays an important role in determining the reliability of spike firing times. We therefore model membrane patches of areas of a few hundred square micrometers, comprising of a total of a few thousands to tens of thousand ion channels. The model patches are injected with both DC inputs and the more biologically realistic fluctuating current inputs. We directly simulate the activity of the ion channels, thus introducing their stochasticity into the models. For a broad range of inputs, the stochastic HH model (unlike the deterministic one) replicates, at least qualitatively, the spike timing reliability and accuracy characteristics observed in cortical neurons. In addition to its significant effect on the timing of spike firing, channel noise also produces three additional experimental observations, namely voltage-dependent sub-threshold membrane

voltage oscillations for DC input, occasional "missing" spikes for supra-threshold inputs and "spontaneous" spikes for sub-threshold inputs. Similar spike timing reliability characteristics are shown to occur in a stochastic model of cortical interneurons (a stochastic version of the Wang and Buzsaki model [191]). We suggest that the noise inherent in the operation of ion channels enables neurons to act as "smart" encoders, and conclude that channel stochasticity should be considered in realistic models of neurons.

This chapter is based on work with Barry Freedman and Idan Segev, most of which was published in Refs. [152, 151].

4.2 The stochastic HH model

In both the deterministic and stochastic HH models (see section 2.2.1), the dynamics of the membrane voltage is given by

$$C_m \frac{dV}{dt} = -g_{leak}(V - V_{leak}) - g_K(V - E_K) - g_{Na}(V - E_{Na}) + I \quad (4.1)$$

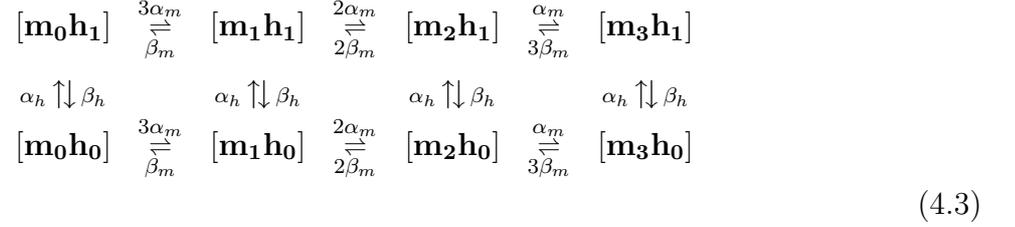
where V is the membrane potential, V_L , E_K , E_{Na} are the reversal potentials of the leakage, potassium and sodium currents, respectively, and g_L , g_K , g_{Na} , are the corresponding specific ion conductances; C_m is the specific membrane capacitance and I is the specific current injected into this membrane patch.

The stochastic model differs from the deterministic one in the description of the potassium and sodium voltage dependent conductances. The set of deterministic differential equations which describe the average response of the ion conductances (section 2.2.1), is replaced with a simulation of the thousands of single ion channels in the membrane. Following the physical interpretation of Hodgkin and Huxley [71], each ion channel is modelled as having voltage dependent 'gates' [50]. A channel is open (and allows ions to flow through it) only when all the gates are open. A potassium channel may be in one of five different states (i.e. the n^4 activation term is interpreted as having four 'gates'). The kinetics of the channel is described by a Markovian model [50, 42, 163, 31] which explicitly incorporates the internal workings of the channel,

$$[\mathbf{n}_0] \xrightleftharpoons[\beta_n]{4\alpha_n} [\mathbf{n}_1] \xrightleftharpoons[\beta_n]{3\alpha_n} [\mathbf{n}_2] \xrightleftharpoons[\beta_n]{2\alpha_n} [\mathbf{n}_3] \xrightleftharpoons[\beta_n]{\alpha_n} [\mathbf{n}_4] \quad (4.2)$$

where $[\mathbf{n}_i]$ is the state with i open gates and, hence, $[\mathbf{n}_4]$ labels the single open state of the K^+ channel. α_n, β_n are identical to the original HH rate functions.

Similarly, each Na^+ channel can exist in 8 different states (due to the m^3 activation and h inactivation variables), defined by the following scheme,



where $[\mathbf{m}_3\mathbf{h}_1]$ labels the single open state of the Na^+ channel, and $\alpha_h, \beta_h, \alpha_m$ and β_m are the rate-functions in Hodgkin-Huxley formalism. The potassium and sodium membrane conductances are given by,

$$g_K(V, t) = \gamma_K \#[\mathbf{n}_4] \quad g_{Na}(V, t) = \gamma_{Na} \#[\mathbf{m}_3\mathbf{h}_1] \tag{4.4}$$

where $\#[\mathbf{n}_4]$ is the number of K^+ channels which are open (i.e in state $[\mathbf{n}_4]$), and $\#[\mathbf{m}_3\mathbf{h}_1]$ is the number of open Na^+ channels. γ_K and γ_{Na} are the conductances of the single potassium and sodium ion channel at their open state, respectively.

By switching from the 'standard' HH model to the ion channel based model, channel stochasticity is incorporated into the voltage dynamics. Instead of keeping track of each of the channels separately, we have used a more efficient scheme to track only the total populations of channels in each of their possible states (see [163, 30], for a discussion on possible simulation methods for populations of channels). Specifically, if at time t there are n_A channels in state A and n_B channels in state B and the transfer rate of channels from state A to state B is r , then each of the channels in state A might transfer to state B between time t and $t + \Delta t$ with probability $p = r\Delta t$. Hence, for each time step we determine Δn_{AB} , the number of channels which move from A to B , by choosing a random number from a binomial distribution [135], i.e.,

$$Prob(\Delta n_{AB}) = \binom{n_A}{\Delta n_{AB}} p^{\Delta n_{AB}} (1-p)^{(n_A - \Delta n_{AB})}. \tag{4.5}$$

In the present study we used the forward Euler integration method with $\Delta t = 0.01 \text{ msec}$, as in [30].

Finally, in order to make the transition from the deterministic to the stochastic model, we need to know exactly how many channels there are in the modelled membrane patch. Once we choose the conductance of the individual channel, the number of channels can be

calculated directly from the channel densities and from the maximal conductances, \bar{g} 's, given in the HH model. We choose here $20pS$ as the value of a single channel conductance for both potassium and sodium channel types – similar to typical values of ion channel conductance reported in cortical neurons [64]. The model parameters are summarized in Table 4.1.

At the limit of infinitely large number of channels (with small channel conductance), the model converges back to the deterministic model behavior. However, when using a finite number of channels with the corresponding single channel conductance, the stochastic model will differ from the deterministic one [163, 42, 170]. It is important to note that we use here the spatially independent (space-clamped) HH equations. Clearly, this is a severe oversimplification of the realistic case and its implications will be addressed in section 4.5.

The effect of ion channel stochasticity is demonstrated in figure 4.1, by voltage clamp simulation of the model. Unlike the deterministic model, for which the current needed to clamp the voltage is fixed (different current for each voltage level), for the stochastic model the current is fluctuating. Thus, the ion channel stochasticity introduces voltage dependent noise to the neuron (compare with [195]). As a result, the input-output relation of the deterministic model is significantly changed in the stochastic model. Figure 4.2 shows the average firing rates of the stochastic model and the deterministic one, in response to DC input currents (the $f - I$ curve). Clearly, the ion channel noise transforms the discontinuous relation between the input current and the firing rate into a smooth one (which is the common behavior observed in real neuron)

4.3 Reliability and accuracy of spike timing in the SHH model

Before we proceed to the simulation results regarding the spike timing reliability and the resulting neural code, let us first try to estimate the effect of introducing stochasticity into the HH model. Suppose that the area of the membrane patch is $200 \mu m^2$. According to the model parameters (Table 4.1), this membrane patch contains 3,600 K^+ channels and 12,000 Na^+ channels. Considering the large number of modelled channels, one would naively estimate the number of fluctuating channels about the mean to be on the order of $\frac{1}{\sqrt{N}}$ channels. Hence, for $N = 3,600K^+$ channels, the size of the fluctuation is 1.7% and, hence, we would expect rather small deviations from the deterministic model. An even smaller effect would be expected for the Na^+ channels. Surprisingly, this is not the case, as shown in Figure 4.3.

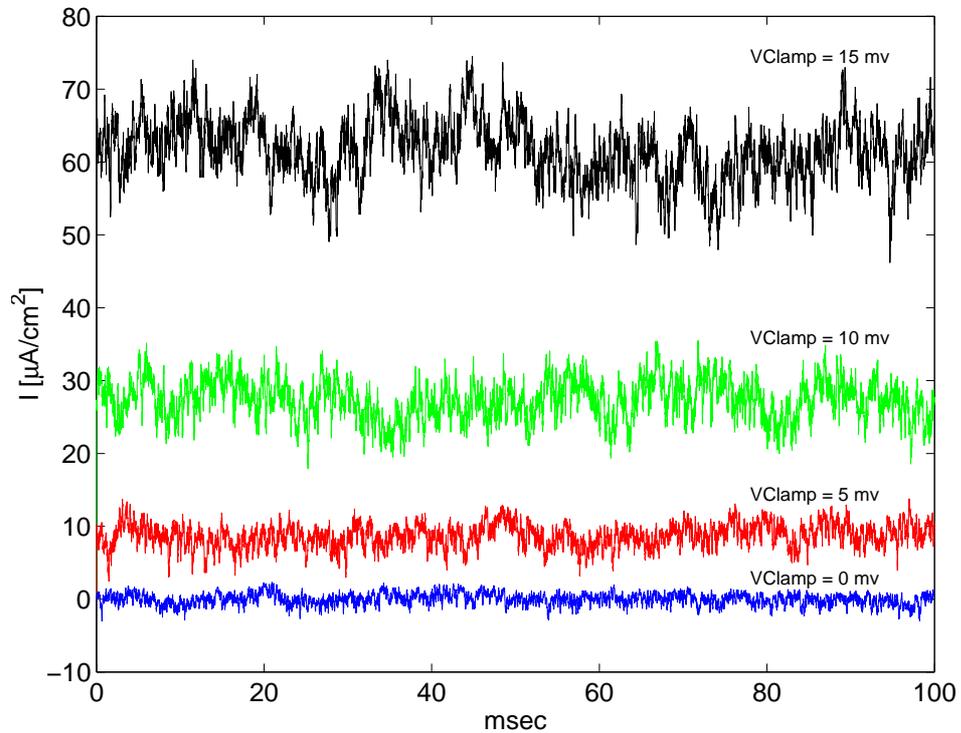


Figure 4.1: **Voltage clamp results of SHH model.** Examples of the current needed to clamp the SHH model (patch area $200 \mu\text{m}^2$) to different voltage levels (indicated near each of the waveforms) for a 100ms period.

4.3.1 Encoding Reliability and Precision: Input Current Versus Channel Fluctuations

The response of a stochastic isopotential HH compartment to repeated presentation of supra-threshold currents is shown in Figure 4.3. When the same supra-threshold DC current pulse ($10 \mu\text{A}/\text{cm}^2$, 250msec ; A, top frame) is repeatedly presented to the modelled membrane patch, the resulting spike trains vary considerably from trial to trial, i.e., the spike firing time is neither reliable nor accurate (Fig. 4.3A, bottom frame). This should be compared with the response of the corresponding deterministic model shown in the middle frames. On the other hand, when the stimulus is fluctuating (simulating the current that presumably reaches the site of spike generation following the activation of many synaptic inputs impinging on the dendritic tree, B top frame; see caption of Fig. 4.3 for details), the reliability and accuracy of the spike train in the stochastic HH model is improved compared to DC case (B, bottom frame).

As in [110], two measures of the spike timing, the *reliability* and the *precision* were calculated from the peri-stimulus time histogram (PSTH, not shown), for a wide range of

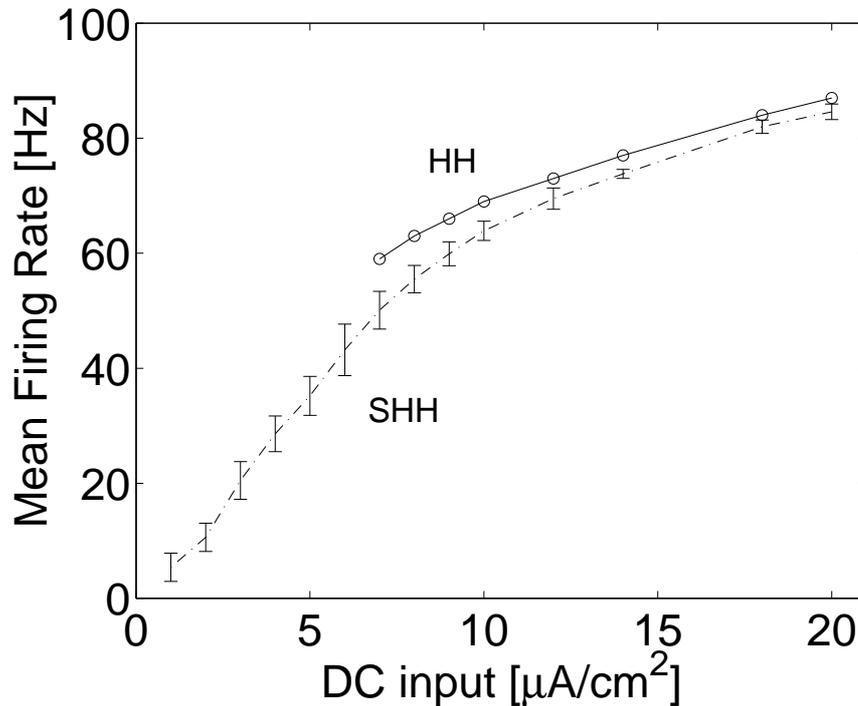


Figure 4.2: **f-I curve of the HH and SHH models.** Average firing rate in response to DC current input of both the HH and the stochastic HH model.

input patterns (see caption of Fig. 4.4 for details). The reliability and precision of the spike patterns were strongly correlated with the amplitude of the fluctuations in the input current, σ_{input} (Fig. 4.4 A and B); the reliability and precision dropped as the input was filtered with larger time constants (Fig. 4.4 C and D). In the stochastic HH model, both the reliability and the precision (which for most of the responses, was in the range of 1 to 2 *msec*), are in close agreement with the results for cortical neurons [110, 127]. It is noteworthy that there is no clear dependence of the reliability and precision on the mean value of the injected current, as was also found experimentally by Mainen and Sejnowski (personal communication).

Hence, with a realistically large number of channels, when incorporating their unavoidable stochasticity, one obtains an effect which is qualitatively similar to the behavior of real neurons, and is significant from both biophysical and computational viewpoints. Clearly, the effect of stochasticity depends on the number of ion channels and on the membrane area. It increases when decreasing the number of channels and decreases when increasing the membrane area. Still, the effect of channel stochasticity was significant even when the membrane area was increased by a factor of five (to $1,000\mu\text{m}^2$). Similar behavior was observed when, for a given membrane patch, the channel density was increased by the same factor (not shown). But why is the result of the stochastic model so different from that

C_m	Specific membrane capacitance	$1 \mu F/cm^2$
T	Temperature	$6.3^\circ C$
V_L	Leakage reversal potential	$10.6 mV$
g_L	Leakage conductance	$0.3 mS/cm^2$
V_K	Potassium reversal potential	$-12 mV$
\bar{g}_K	Maximal potassium conductance	$36 mS/cm^2$
γ_K	Potassium channel conductance	$20 pS$
D_K	Potassium ion channel density	$18 channels/\mu m^2$
V_{Na}	Sodium reversal potential	$115 mV$
\bar{g}_{Na}	Maximal sodium conductance	$120 mS/cm^2$
γ_{Na}	Sodium channel conductance	$20 pS$
D_{Na}	Sodium ion channel density	$60 channels/\mu m^2$
$\alpha_n(V)$		$\frac{0.01(10-V)}{e^{(10-V)/10}-1}$
$\alpha_m(V)$		$\frac{0.1(25-V)}{e^{(25-V)/10}-1}$
$\alpha_h(V)$		$0.07 e^{-V/20}$
$\beta_n(V)$		$0.125 e^{-V/80}$
$\beta_m(V)$		$4.0 e^{-V/18}$
$\beta_h(V)$		$\frac{1}{\exp(30-V)/10+1}$

Table 4.1: Hodgkin-Huxley parameters and rate functions used in the simulations

obtained from the corresponding deterministic HH model?

The apparent error in the previous estimation of the size of the effect of channel stochasticity lies in failing to realize that the relevant number of channels is not the total number of channels in the membrane patch, but rather the number of channels which are open near the threshold for spike firing. If this number is relatively small, the size of the fluctuations in the number of open channels in this regime is not negligible. Mathematically, the correct estimation for the size of the fluctuation should rely on the binomial statistics. For a total population of N channels and a probability p of a channel to be open, the size of the fluctuations is $\sqrt{Np(1-p)}$ and the fluctuations about the mean, Np , is $\sqrt{\frac{(1-p)}{Np}}$. If p is small, as in the case near the threshold for spike firing, the relative size of the fluctuations is rather large. In this case, the inherent stochasticity of the channels is expected to have a significant effect on the voltage dynamics, i.e., on the time of threshold crossing. When this is the case, the firing behavior of the stochastic model is expected to be considerably different from that of the corresponding deterministic model.

Figure 4.3.1 shows that this is indeed the situation. As in Figure 4.3A, ten repeated $10 \mu A/cm^2$ DC current inputs were applied, this time to a $600 \mu m^2$ membrane patch consisting of a total of 10,800 K^+ channels and 36,000 Na^+ channels. The voltage response is shown in Figure 4.3.1B, whereas the number of open K^+ and Na^+ channels near the spiking

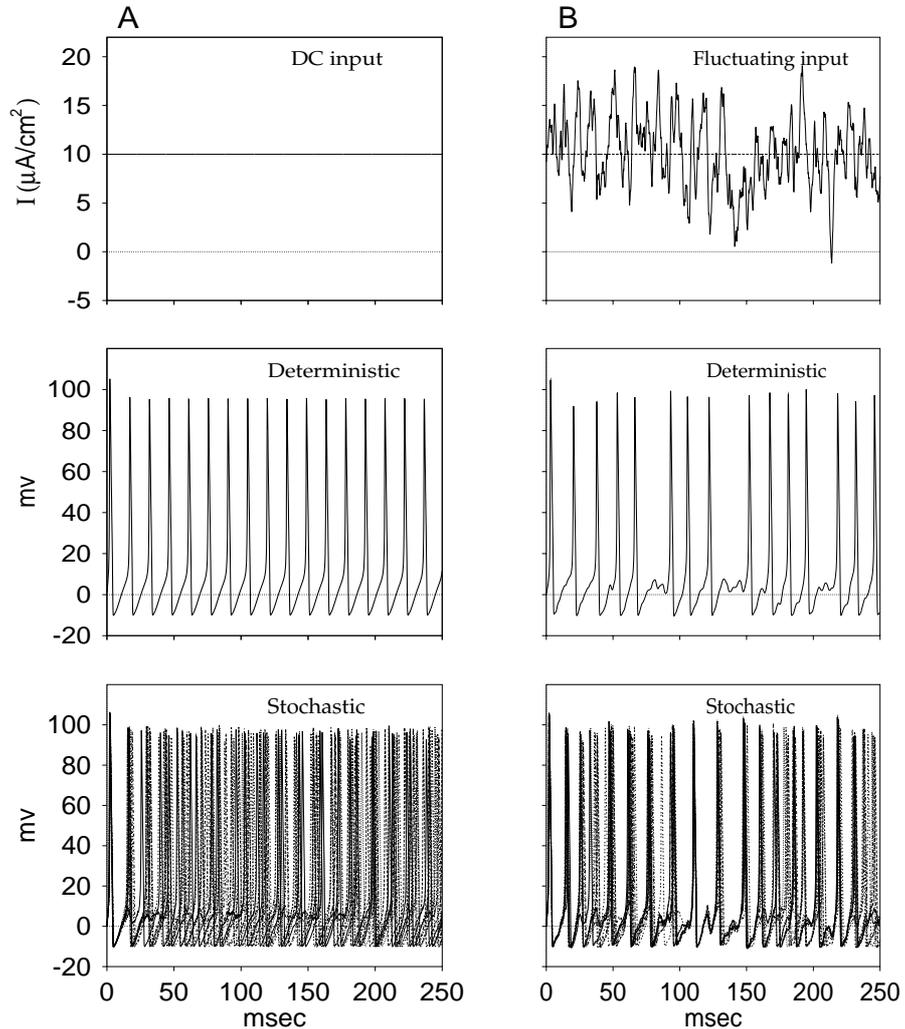


Figure 4.3:

for figure caption see footnote⁰

threshold is shown in panels C and D, respectively. As can be seen, a surprisingly small number of ion channels; approximately 300 K^+ channels and 50 Na^+ channels are opened in this voltage regime. With these small numbers, channel fluctuations become significant and critically determine the exact time in which sufficient *additional* Na^+ channels are recruited to initiate a regenerative response. When injecting the same DC current repeatedly, the fluctuations vary significantly from one trial to the other. Consequently, the time of spike firing for this input is unreliable.

⁰Figure 4.3: Reliability of firing patterns in a model of an isopotential deterministic and stochastic Hodgkin-Huxley membrane patch in response to both DC and fluctuating current input. **A** Ten superimposed responses to repeated supra-threshold DC current pulses ($10 \mu A/cm^2$, 250 msec; top frame), evoked a train of regular firing in the deterministic HH model (middle frame) and a 'jitter' in the firing in the stochastic HH model (bottom frame). **B** The same patch was again stimulated ten

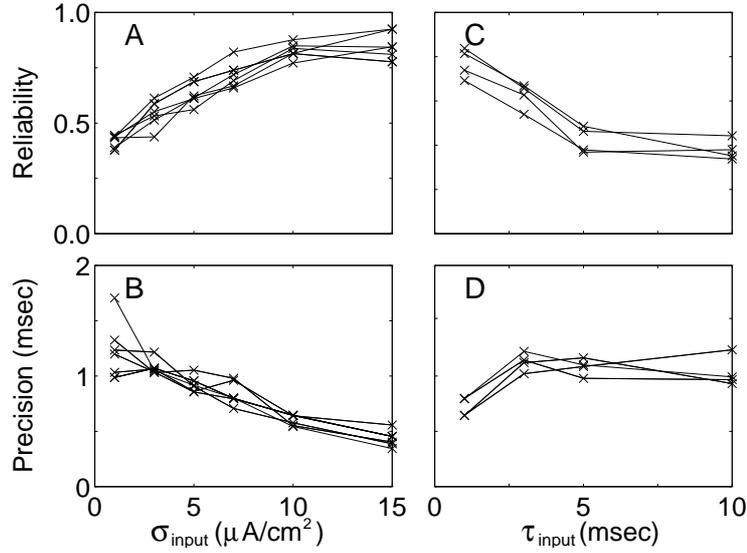


Figure 4.4: **Dependence of reliability and precision on stimulus parameters.** The reliability and accuracy of the spike train was calculated in a similar manner to [110]. The Peri-stimulus time histogram (PSTH) of 20 successive presentation of a particular stimulus was smoothed using an adaptive filter, yielding an estimate for the instantaneous firing rate. Significant elevations in the instantaneous firing rate ('events') were selected from the PSTH using a threshold of two times the mean firing rate over a given block of responses. The *reliability* of the response to a particular stimulus is defined as the average of the fraction of spikes that occurred in the events in that stimulus' PSTH. The temporal *precision* of the response is defined by the average of standard deviation of the events in that stimulus' PSTH. **A** Estimates of the reliability of the spike train in a $200\mu\text{m}^2$ stochastic HH membrane patch, for stimuli with various fluctuation amplitudes, σ_{input} . Each curve is for a different mean value of the stimulus ($\bar{I} = 7 - 20 \mu\text{A}/\text{cm}^2$, $\tau_{input} = 1 \text{msec}$). **B** The temporal precision of the same responses as in A. **C** The reliability for stimuli filtered with different time constants ($\tau_{input} = 1 - 10 \text{msec}$). Each curve is for a different mean value of the stimulus and a given σ_{input} . ($\bar{I} = 7 - 20 \mu\text{A}/\text{cm}^2$, $\sigma_{input} = 3 - 12 \mu\text{A}/\text{cm}^2$). **D** The temporal precision of the same responses as in C.

In principle, this channel noise induced unreliability can be mostly overridden by injecting a current that fluctuates significantly. If the input fluctuations are sufficiently large, the voltage dynamics will be dominated by the transients in the current input rather than by the channel noise. This effect is demonstrated in Figure 4.3.1 where the response to a fluctuating

times repeatedly, this time with a fluctuating stimulus (low-pass Gaussian white noise with a mean \bar{I} , of $10 \mu\text{A}/\text{cm}^2$, and a standard deviation σ_{input} of $7 \mu\text{A}/\text{cm}^2$ which was convolved with an 'alpha-function' with a time constant $\tau_{input} = 1 \text{msec}$, top frame (see [110])). As can be clearly seen, the 'jitter' in spike timing in the stochastic model is significantly smaller in B than in A (i.e. increased reliability for the fluctuating current input). Patch area used was $200 \mu\text{m}^2$, with 3,600 K^+ channels and 12,000 Na^+ channels. (Compare to Fig.1 in [110]).

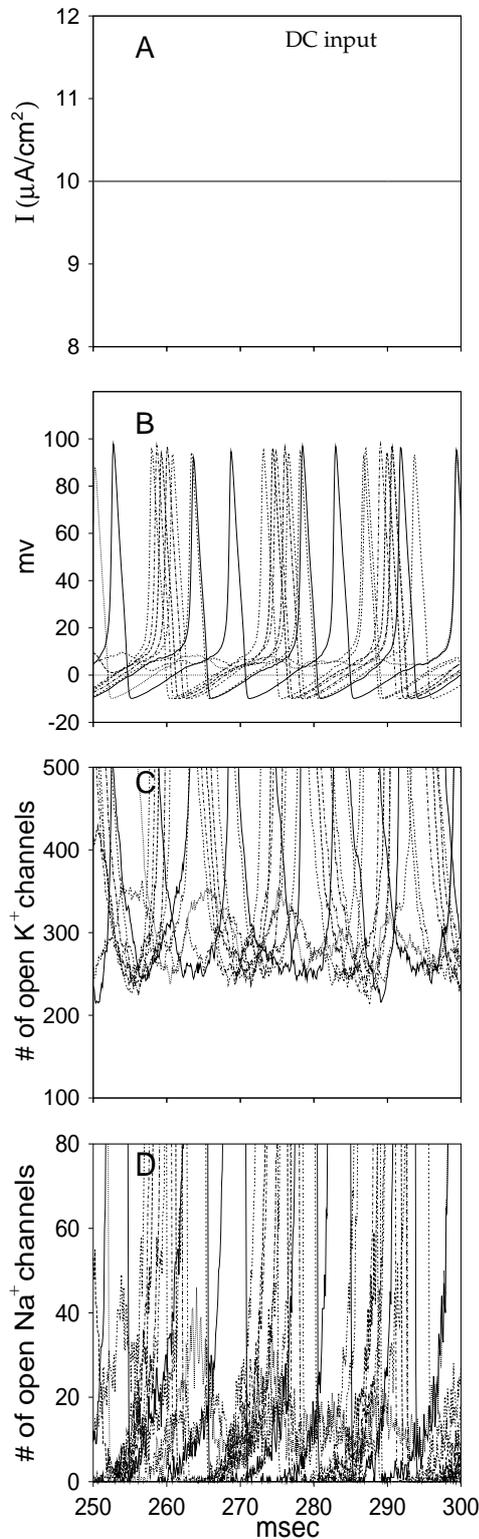


Figure 4.5: **Channel fluctuations ruins the reliability of spike timing in the case of DC current input.** A $10\mu\text{A}/\text{cm}^2$ DC current injected to a $600\mu\text{m}^2$ stochastic HH model (10,800 K^+ channels and 36,000 Na^+ channels) results with dispersed spike timings for repeated simulation (10 superimposed voltage traces) in **B**. In **C** and **D**, the number of open Na^+ and K^+ channels, respectively, corresponding to the voltage traces presented in **B** are shown.

input in a specific time window is shown. In contrast to the DC-input case (Fig. 4.3.1), here the transients in the input current partially overcome the channel fluctuations and enforce nine out of the ten spikes to occur within an approximately 1 msec time window (Fig. 4.3.1B). The reason for this relatively high reliability of spike timing becomes clear by observing panels C and D. The accuracy is determined by two parameters. The first is the variability in the time where the number of open K^+ channels reaches a sufficiently small value (note that a large outward K^+ current impedes the initiation of the spike). This variability should be small in order to obtain high accuracy. Indeed, in 9 out of 10 repetitions, this condition is satisfied (see Fig. 4.3.1C). The second parameter is the rate of the buildup of the Na^+ channel population towards threshold. For an accurate spike timing, this buildup, which is determined by the amplitude and rate of the depolarizing input current, should be sufficiently large to overcome the channel fluctuations (see Fig. 4.3.1D). We note that, for a given voltage, the size of the channel fluctuation is as large in the fluctuating-input case as in the DC-input case, but in the latter these channel fluctuations are ‘lost in the crowd’.

To examine the relative contribution of the K^+ and Na^+ channels to the reliability and precision, we simulated a hybrid system in which one of the channel populations was stochastic and the other was deterministic. Both channel types contribute to the complex reliability nature of the system. However, as expected from the larger number of K^+ channels that are open near threshold for spike firing (Fig. 4.3.1 C,D and Fig. 4.3.1 C,D), as well as from their slower kinetics, the noise introduced by the K^+ channels is more dominant in determining the reliability and accuracy of this system (not shown).

4.3.2 Sub-threshold oscillations, ‘Spontaneous’ spikes and ‘missing’ Spikes

Along with the effect of channel fluctuations on spike timing, incorporating channel stochasticity in the HH model gives rise to three additional phenomena which were observed experimentally: (i) considerable sub-threshold oscillations in the membrane voltage for DC inputs, (ii) ‘spontaneous’ spikes for ‘sub-threshold’ inputs and (iii) ‘missing’ spikes for supra-threshold inputs. These phenomena cannot be reproduced in the deterministic HH model. In the stochastic model, oscillations in the membrane voltage are already observed for zero current input (Fig. 4.7A). Occasionally, these oscillations are sufficiently large to generate ‘spontaneous’ spikes which would not have occurred in the corresponding deterministic model (current threshold for spike firing in the deterministic model is $I = 7\mu A/cm^2$). An example of ‘spontaneous’ spikes in the case of $I = 4\mu A/cm^2$ is shown in Figure 4.7B; detailed analysis of spontaneous spiking in the stochastic HH model for zero current input was recently per-

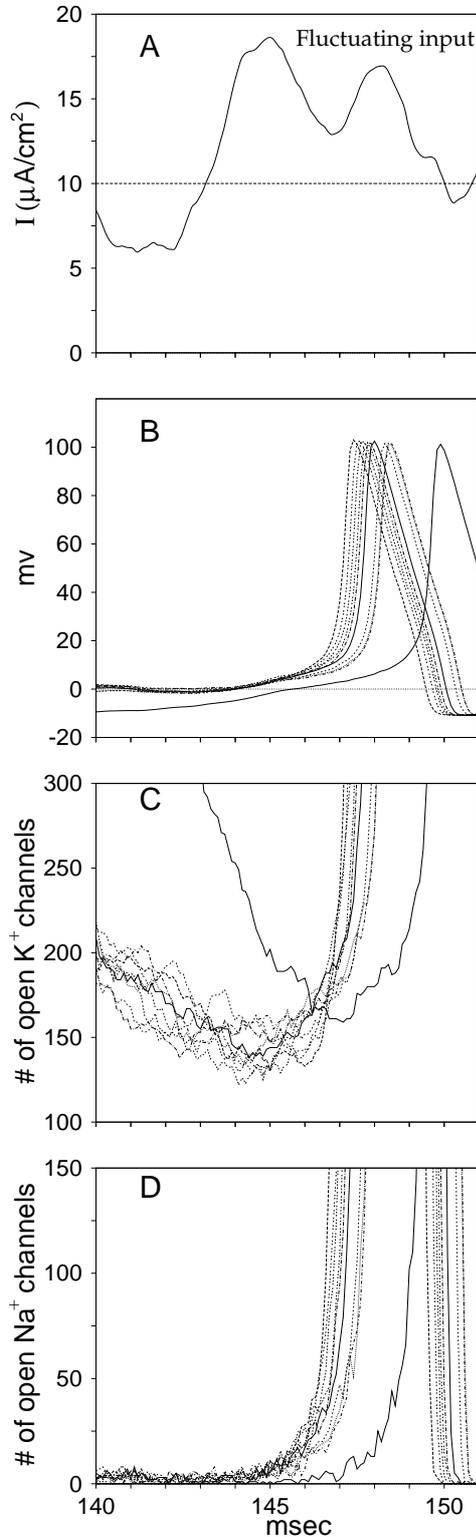


Figure 4.6: **Fluctuating input current partially overrides the channel stochasticity and increases the reliability of spike timing.** A small time window of the system behavior for the fluctuating input case is presented. **A** The input current with a mean value of $10 \mu\text{A}/\text{cm}^2$ and with $\sigma_{input} = 5 \mu\text{A}/\text{cm}^2$ and $\tau_{input} = 1 \text{ msec}$, injected to a $600 \mu\text{m}^2$ stochastic HH membrane patch is depicted. **B** Ten superimposed voltage-traces responses to repeated injection of the fluctuating current in A. In nine out of ten of the cases, a spike was fired within approximately 1 millisecond time window. **C,D** The number of open Na^+ and K^+ channels, respectively, for the voltage traces presented in B, reflecting how the fluctuations of both Na^+ and K^+ channels are overridden by the fluctuations in the input current. When a sufficient number of K^+ channels close (C), the depolarizing transient in the input current, starting at $t = 143 \text{ msec}$, results in the nearly synchronous buildup of Na^+ channels at $t = 146 \text{ msec}$ (D). The result is spike firing at $t = 147.3 - 148.4 \text{ msec}$. In the one case where insufficient number of K^+ channels was closed in time, the spike is initiated somewhat later due to the next fluctuation in the input current.

formed (see [30]). Some of the spikes that occur in repetitive firing in the deterministic model (Fig. 4.3), disappear in the stochastic model, although the input currents are supra-threshold ('missing spikes' in Fig. 4.7, C and D).

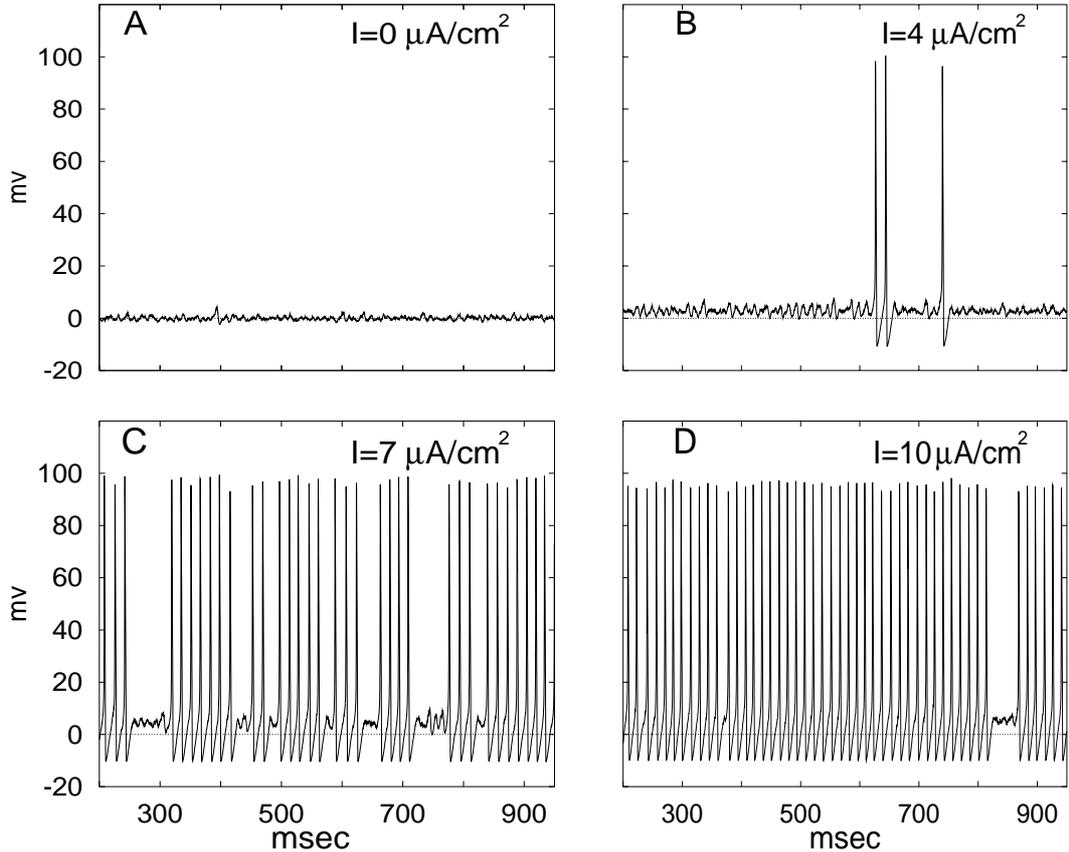


Figure 4.7: **The response of the stochastic model to injected DC input currents.** Different DC current amplitudes were injected to a stochastic HH model of an isopotential membrane patch of area $600 \mu\text{m}^2$ ($10,800 K^+$ channels and $36,000 Na^+$ Channels). **A,B** Membrane voltage oscillations are the dominant effect of the stochastic nature of the ion channels, with occasional spontaneous spiking. **C,D** Supra-Threshold DC input currents result with irregular spiking, with occasional 'missing' spikes and membrane voltage oscillations. This is not expected in the corresponding deterministic HH model where the threshold is $7 \mu\text{A}/\text{cm}^2$. Below this value, smooth sub-threshold voltage response is observed (not shown); above this value, regular firing is obtained (not shown).

It is important to note that both the amplitude and the frequency of the membrane oscillations observed in the stochastic model are voltage dependent (e.g., compare A to B in Fig. 4.7). This is also the case with the membrane voltage oscillations in neocortical neurons [56, 87], as well as in other neuron types (e.g., [79] and [91]). We suggest that in addition to the deterministic macroscopic mechanisms that were proposed to explain the generation of sub-threshold oscillations, the stochastic nature (and the limited number) of

the ion channels may have a dominant effect on the nature of these oscillations (see also [105] and [23]).

Channel stochasticity has such a dramatic effect on the voltage dynamics because it exploits a peculiar, and largely neglected, aspect of the deterministic HH equations, namely its two stable states for *supra-threshold* current input (see discussion of the bistability in the HH equations in [34] and [59]). For a DC input, one state is the well-known repetitive firing behavior (green curve in Fig. 4.8A) whereas the other state is a non-firing behavior of early damped voltage oscillations which converges to a steady voltage (Fig. 4.8A, red curve). In both cases, a $7\mu A/cm^2$ DC current was injected and the marked difference between the two curves is the result of minute perturbation in the initial conditions. These two different behaviors can be better appreciated in the phase plane diagram in panels B and C of Fig. 4.8. Translating the ion conductances to the corresponding number of open ion channels, these panels show the very different paths in phase space taken by the firing (green) and the non-firing (red) trajectories. The bottom panel shows the convergence of the non-firing behavior to a fixed point. It also shows that the distance, in terms of number of open channels, between the continuous firing cycle and the *non-firing* voltage behavior is very small. Although small, the deterministic nature of the HH equations implies that, for a DC input, the system remains in one stable state or the other. However, introduction of channel noise could, in principle, flip the system between these two states.

Figure 4.8 D-F show that, indeed, in the stochastic model, channel fluctuations do occasionally bridge the small distance in phase space between the two stable states. The stochastic opening (or closing) of a few extra K^+ and/or Na^+ channels pushes the system spontaneously from the continuous firing stable-state (blue trace) to the non-firing stable-state (red trace), where it stays for a while, and vice versa (green trace, Fig. 4.8D). This spontaneous transition between the two states is the cause for the ‘missing’ spikes and the sub-threshold membrane voltage as well as for occasional ‘spontaneous spikes’ (Fig. 4.7B). Panels E and F depict the corresponding phase-plane behavior of the system. It clearly shows that fluctuations in only a few channels are responsible for the transition between these two stable states. We conclude that the non-firing stable state in the deterministic HH model, becomes a key player in the stochastic HH model. Experimentally, the co-existence of the two stable solution in the the squid giant axon, as well as in the corresponding HH model, was demonstrated (see [59] and also [34]).

Considering the sub-threshold membrane oscillations, the role of channel fluctuations is two-fold. First, they drive the system from the firing state into the basin of attraction of the non-firing stable state. Second, the fluctuations prevent the system from converging into the fixed-point of the non-firing stable state of the corresponding deterministic model.

As a result, the system is 'cycling' around this fixed point and, thus, the sub-threshold membrane voltage oscillations emerge. The frequency of the sub-threshold oscillations is set by the period of these cycles. Based on this observation we can predict analytically the power spectrum of the oscillations with a fair degree of accuracy and we can also quantify the rate of transfer between the two states in the case of DC input (not shown). Questions regarding the effect of various parameters on the sub-threshold oscillations, such as the area of the membrane patch, the properties of the channels etc., will be briefly addressed in the Discussion but a more complete study is yet to be performed.

Finally, it is interesting to note that changing the temperature of the model, and thus accelerating the rate functions which control the dynamics and fluctuations of the ion channels, may result in a dramatic change in the nature of the spike firing. Figure 4.9 shows the effect of changing the temperature of the model to $20^{\circ}C$, which is significantly higher than the 'standard' HH temperature ($6.3^{\circ}C$) (and is not a 'normal' temperature for the squid). While the change in temperature has a small effect on the regular spiking pattern of the deterministic model (panel A), the firing of the stochastic HH model (panel B) becomes highly irregular and 'bursty' [105], and the subthreshold oscillation become even more apparent.

Increasing the temperature, thereby accelerating the rate-functions of the model, results in a dramatic change in the nature of the spike firing (Figure 4.9). The regular spiking in the deterministic model (A) turns, in the stochastic case (B) into a highly irregular and 'bursty' spike trains [[105]]. Again, this emphasizes the significant difference between the deterministic model and the more realistic stochastic one.

4.4 Spike reliability in the stochastic Wang-Buzsaki model

We turn to ask to what extent the results for the stochastic HH model may apply for other spike generation models (and other neurons). Using a similar formalism, we simulate a stochastic version of the Wang and Buzsaki model (WB) [191] of a cortical interneuron. As in the HH model, the "backbone" of the WB model is based on sodium and potassium channels, but with different channel densities and rate functions. The membrane voltage is given by equation 4.2, and the activation and inactivation variables follow the same differential equations as the HH case. $C_m = 1\mu F/cm^2$, and I_{inj} is in $\mu A/cm^2$, $G_{leak} = 0.1 mS/cm^2$ and $v_{leak} = -65 mV$. The rate functions for the sodium conductance are given

by

$$\begin{aligned}
\alpha_m(V) &= 0.1 \frac{V + 35}{1 - \exp^{-(V+35)/10}} \\
\beta_m(V) &= 4 \exp^{-(V+60)/18} \\
\alpha_h(V) &= 0.35 \exp^{-(V+58)/20} \\
\beta_h(V) &= \frac{5}{1 + \exp^{-(V+28)/10}}
\end{aligned} \tag{4.6}$$

with $G_{Na} = 35 \text{ mS/cm}^2$ and $E_{Na} = 55 \text{ mV}$. The potassium rate functions are given by

$$\alpha_n(V) = 0.05 \frac{V + 34}{1 - \exp^{-(V+34)/10}} \tag{4.7}$$

$$\beta_m(V) = 0.625 \exp^{-(V+44)/80} \tag{4.8}$$

with $G_K = 9 \text{ mS/cm}^2$ and $E_K = -90 \text{ mV}$.

Unlike the HH model, and similar to the behavior of many cortical neurons, its f - I curve is continuous, reflecting a different input-output dynamics. It is an example of the type I family of neuron models [48, 57]. Figure 4.10 shows that reliability and accuracy of the spike train of the stochastic WB model (SWB), has similar input-dependent characteristics, as the SHH model. The SWB model also shows spontaneous and 'missing' spikes, but its subthreshold voltage fluctuations do not show oscillatory behavior (not shown).

4.5 Conclusions and discussion

4.5.1 Summary of results

We have shown that with a realistically large number of ion channels, the inherent 'noise' in channel operation critically determines the timing and dynamics of spike firings for the stochastic HH model. Similar behavior was observed in a stochastic model of a cortical interneuron. The reason for this strong effect of channel stochasticity is that near the threshold for spike firing, only a very small percentage of Na^+ and a small percentage of K^+ channels is open (in the HH formalism – the activation variables, m and n , are small near threshold). Consequently, the variability in membrane voltage near threshold for excitation is large and this is reflected in the variability of spike firing time. We conclude that for a wide range of input parameters, the stochastic model captures important features of real neurons; these features are neglected in the deterministic model.

In agreement with the experimental results [110, 127] the reliability and precision of spike timing in the stochastic HH and Wang-Buzsaki models is very sensitive to the properties of

the current input. The reliability and precision of the spike timing is high for strongly fluctuating inputs and decreases for more smooth (e.g., DC) inputs. The present study shows that this effect could be explained in terms of the relation between the instantaneous shape and amplitude of the input signal and the amplitude of channel fluctuations. Strongly fluctuating inputs "override" the inherent channel fluctuations, and the spike timing is primarily dictated by the input rather than by channel stochasticity. In contrast, channel fluctuations become relatively more significant for smooth inputs and thus, spike firing time becomes less reliable.

In addition to its effect on spike timing, channel stochasticity produces three additional phenomena which do not occur in the deterministic HH model, but were all observed experimentally (e.g., see [59]). Voltage membrane oscillations are seen for sub-threshold current inputs and they also occur between spikes for supra-threshold inputs. 'Spontaneous' spikes (for sub-threshold inputs) and 'missing' spikes (for supra-threshold inputs) were also observed in this model. These three phenomena result from the "unmasking" of the non-firing stable state in the HH model by the channel fluctuations. This state, which was largely neglected in the framework of the deterministic HH model, becomes a key player in determining voltage dynamics in the stochastic model.

4.5.2 Towards a more realistic stochastic model of neurons

The present study gives only a qualitative explanation for the input dependent reliability of spiking time of neurons. First, the spatial domain of neurons was completely neglected. In particular, it is important to consider the filtering effect and the impedance load imposed by the soma and dendrites, as well as by the axon, on the excitable channels at the spike initiation zone. A multi-compartmental model (possibly composed of an axon with several highly-excitabile nodes of Ranvier, separated by passive inter-nodes and a few dendritic compartments) should be utilized to better understand the effect of channel stochasticity on the reliability and accuracy of spiking in neurons. In such a model, the input should impinge onto the dendritic compartments and be simulated by a barrage of synaptic conductance changes (rather than by current inputs). In this context, it is important to emphasize that in many neuron types, the dendritic membrane is endowed with excitable channels in low density [109] and this may imply a large variability (fluctuations) already in the receptive region of the neuron. In contrast, we expect that in the axon, most of the variability will arise in the compartment where the spikes are initiated and that, downstream along the

axon, spike timing would be encoded very reliably and with high precision ¹.

Another major issue to consider is the applicability of the results presented here to other spiking models and excitable systems, in addition to the models we have presented. Most neurons consist of a large variety of ion channel types (e.g., *A*-current, persistent and slow-inactivating Na^+ currents, low threshold Ca^{+2} current, etc.) each with different density and kinetics. Moreover, based on direct measurements of single ion channels behavior, it is possible to construct detailed kinetic diagrams of single ion channels which will better describe the channel stochasticity and voltage-dependent dynamics and open time correlations. These models usually include many more states, and result in different inherent time scales [132, 186, 51, 122, 177] (other models even consider a continuum of states, see [106]). Still we can state with confidence that the surprisingly large effect of channels stochasticity is likely to persist for other models. The important parameter that determines the size of fluctuations near threshold for spike firing is the number of open channels in this voltage regime. To the best of our knowledge, in all present models of excitability, only a small percentage of the total number of excitable channels is open near threshold. Consequently, a large variability in spike firing time is also expected in these models. Clearly, the exact nature of spike firing reliability will depend on channel properties.

What about the sub-threshold membrane oscillations, spontaneous spikes and missing spikes? The nature of the bistability of the HH model, which is set by its inverted-Hopf bifurcation, is what ‘enables’ the channel noise to spontaneously switch the system between its two stable states. For the SHH model, this is the source of the sub-threshold membrane oscillations, the ‘missing’ and the ‘spontaneous’ spikes. These phenomena may not occur in models with different types of stabilities (e.g., those with saddle-node bifurcation) and other phenomena may then arise (see [144, 194, 105]). Indeed, the results for the stochastic Wang and Buzsaki model do not show similar obvious subthreshold oscillations, but rather subthreshold fluctuations, as well as missing and spontaneous spiking. The oscillations we got from changing the temperature of the stochastic HH model suggest that the nature of subthreshold fluctuations (unlike the input-dependent spike timing reliability) may be very

¹One might wonder if the uncertainty engendered by the bistability in the HH model would make spike propagation along the axon impossible. If, at each site in the axon, there is some probability that the system will go into a stable non-firing state, the spike may fail somewhere along the axon. Also, a significant noise in the axon (see [145]) may destroy temporal correlations between the output synapses. However, except for the compartment where the spike may, or may not, be initiated as a result of the depolarizing synaptic current, all other axonal compartments downstream receive relatively sharp and large current input from the spike in the previous node. For such inputs, channel stochasticity will be masked and, consequently, the axon is expected to act as a highly reliable delay line adding only relatively small jitter [75, 76, 6, 94]. Still, complete failure may occur in axons at regions with low safety-factor for propagation (e.g., see [55]).

sensitive to the model (or neuron) parameters.

4.5.3 Other sources of noise in neurons

In the current study we only considered the effect of one source of noise in neurons – the intrinsic stochastic nature of the ion channels. A variety of other sources of noise exists, such as spontaneous synaptic release, variability in the number of transmitter molecules and in the number of available receptors. Other possible sources of neuronal noise are changes in intracellular and extracellular ion concentrations and in the concentration of neuromodulators, as well as in the activity of ion pumps. Ephaptic interactions (electric field effect) of one neuron on other neurons is yet another possible source of noise. Our study shows that the intrinsic channel stochasticity should be considered as a key source of the variability of action potential timing.

Clearly, the other possible sources of noise should also be considered in order to quantify the relative contribution of each of these sources, or their possible synergistic effect. Theoretical works which have analyzed possible sources of noise in neurons have reflected that synaptic noise and ion channel noise are the dominant sources of neuronal noise [96, 97, 114, 115]. Synaptic noise and background activity effects have been studied both theoretically and experimentally, and shown to have a significant effect on neuronal spiking [200, 169, 65]. However, the *in vitro* experiments on spike timing reliability have been usually performed with synaptic blockers, [110, 127], i.e. without any synaptic noise effects. We note the experimental and modelling results of White *et al* [195] who have used similar stochastic models to explain the nature of voltage noise in cortical neurons *in vitro*, and of Jensen and Gartner [84, 83], who showed that a simple additive noise could qualitatively reproduce the differences in reliability and accuracy of spike timing in response to DC versus the fluctuating input found experimentally. (However, because the neuronal noise is both voltage- and activity-dependent, it is clear that a simple additive noise is only a first-order approximation to the real case. The difference between models with simple additive noise and models with more realistic noise in terms of the fine temporal structure of spike firing requires further exploration).

Experimental studies may further clarify this issue by utilizing different manipulations; e.g., blocking synaptic receptors (as in [110]); using a dynamic clamp to "replace" the noisy channel conductance with a deterministic conductance [161]; blocking specific ion channels and observing the resultant changes in membrane noise under voltage-clamp conditions (see initial results in this direction by [189]); blocking ion pumps, etc. We also note recent experimental results reflecting different membrane voltage "states" in cortical neurons *in*

vivo, and voltage membrane fluctuations of spiking threshold [12, 10].

4.5.4 Implication for Neural Coding

The reliability and accuracy of the spike timing, in the stochastic HH and WB models and in cortical neurons ([110]; [13]) as well as in other neurons [39] range between an unreliable response to DC inputs and a very reliable response to large-amplitude, highly-fluctuating, inputs. The actual current that reaches the site of spike initiation in neurons varies between these two extreme input patterns; its exact nature is determined by the degree of correlation among the synaptic inputs that impinge onto the neuron. Highly correlated synaptic activity gives rise to sharp current transients whereas uncorrelated synaptic activity give rise to 'smooth' current traces (see also [78, 83]). Our modelling results suggest that the neuron's most basic machinery – the ion channels – enable it to act as a 'smart' encoder. Slowly varying inputs are coded with low reliability and accuracy and, hence, the information about such inputs is encoded almost exclusively by the spike *rate*. Trying to decode information about such an input, using the exact temporal structure of the spike train, would result in decoding the internal noise of the cell rather than decoding the input. On the other hand, correlated inputs are encoded with higher reliability and accuracy, giving more of a 'temporal' code. I.e. information about the input exists in the exact timing of the spikes.² In the next chapter we demonstrate and quantify the notion of temporal coding and rate coding in stochastic models which incorporate ion channel noise, using information theory.

We note that in such a system, correlated activity of a population of neurons is likely to propagate within the network with high temporal precision, as suggested by Abeles [5] in the *synfire* model as well as in recent detailed models of spike timing in a network [44, 123]. In contrast, weakly correlated activity would propagate in an imprecise temporal manner and is more likely to decay within the network. The fact that the intrinsic noise of neurons may serve as a mechanism to destroy propagation of random correlations and, at the same time, allows for an accurate chains of activity to persist within the network has no bearing on the question whether such chains do exist.

In addition to its significance for information coding, the relatively small size of the channel 'pool' in the spike initiation zone has further computational implications. One clear advantage of such a limited channel 'pool' was demonstrated in the work of [177] (see

²It seems that channel stochasticity would be very dominant in models that assume balanced excitation-inhibition (see [158, 159, 165, 16, 185]) , in which the effective 'resting' membrane voltage of the cell is near threshold. [16] and [180] suggest that complex repolarization and refractoriness schemes as another source for the high firing variability in neocortical neurons which, in many ways, coincide with the effect of channel fluctuations.

also [121]) which shows that channel inactivation and reactivation kinetics have a significant, long-lasting (minutes) effect on the ‘availability’ of channels, providing the neuron with an effective memory. Thus, the output spike train depends on both the properties of the instantaneous synaptic input as well as on the history of the presynaptic and postsynaptic activity. This memory is embedded in the distribution of channels states in the spike initiation site. The nature and resolution of this memory depends on the size of the channel pool and on the kinetics and number of states of the channels. We hypothesize that the number of channels in the spike initiation zone may be ‘optimized’ in some sense to give the reliability and accuracy discussed above, together with a short term memory of the neuron’s activity. In this context it is interesting to mention the works of [117] and [2] which demonstrates activity-dependent long-term changes in the properties of intrinsic membrane currents. In the next chapter we investigate the design principles and the effect of biophysical parameters of the stochastic model on the nature of the neural code.

Another important effect of stochasticity in a limited ‘pool’ of channels are the sub-threshold and supra-threshold membrane oscillations. Such oscillations were observed in neocortical neurons (see [56]) as well as in other neuron types ([79]; [91]) and were suggested to serve as the underlying ‘clock’ for neurons firing and even as a synchronizing and binding mechanism for neuronal activity ([74] and [189]). In the HH model, these voltage oscillation result from the channel noise; in other systems other mechanisms may be responsible for these oscillations (e.g. [56]; [194]).

We argue then that the noise inherent to the activity of ion channels must be considered in neuronal modelling and spike train analysis, if one wishes to understand what determines the firing patterns of neurons and, consequently, the nature of the neural code.

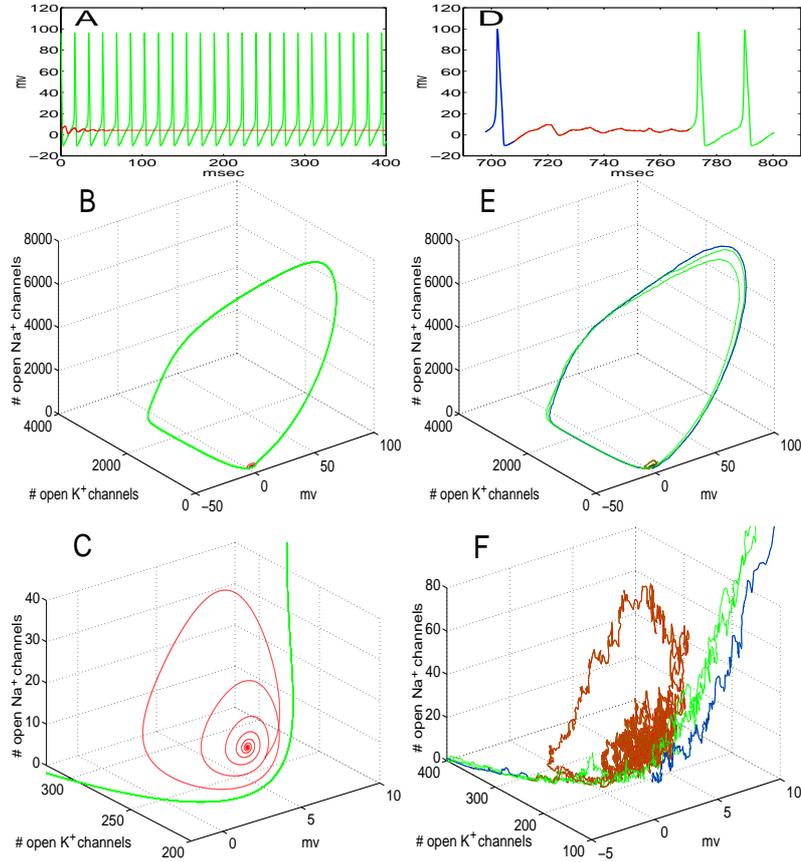


Figure 4.8: **Channel fluctuations cause flipping between firing and non-firing stable states in the stochastic HH model.** **A–C** The two stable states of the deterministic HH model are presented. In **A**, two traces of the membrane potential are shown for a $600\mu m^2$ membrane patch, injected with a $7\mu A/cm^2$ DC current. The difference between the green and red traces results from the minute difference in the initial conditions. In the continuous firing case (green curve), the initial values are: $V = 4.21 mV$; $g_K = 4121 pS$ (corresponds to 206.05 open K^+ channels) and $g_{Na} = 195.8 pS$ (corresponding to 9.79 open Na^+ channels) assuming a single-channel conductance of $20 pS$ (Table 1). In the non-firing case (red curve) the initial values are: $V = 4.23 mV$; $g_K = 4399.4 pS$ (corresponds to 219.97 open K^+ channels) and $g_{Na} = 197.4 pS$ (corresponding to 9.87 open Na^+ channels). **B** shows the 3D phase-plane of these two behaviors of the system, the green curve is for the spiking behavior, and the red curve is for the non-firing stable state. A magnification of **B** is presented in **C**, reflecting the small basin of attraction of the non-firing state and the short distance in term of number of open channels between the two states. **D–F** The corresponding behavior of **A–C** in the stochastic model. Channel fluctuations in the stochastic model spontaneously flip the system between the firing and the non-firing states. **D** A typical voltage trace of the stochastic HH patch. Different colors were used to emphasize the different segments of the trace. The corresponding phase-plane traces are shown in **E** and **F**. As can be seen in **F**, the system flips from the firing stable state (blue trace) to the non-firing stable state (red trace), where it stays for a few 'cycles'. Hence, the sub-threshold oscillations in the top trace translate to small size 'loops' in the phase-plane. The system then flips back to the firing stable state (green trace).

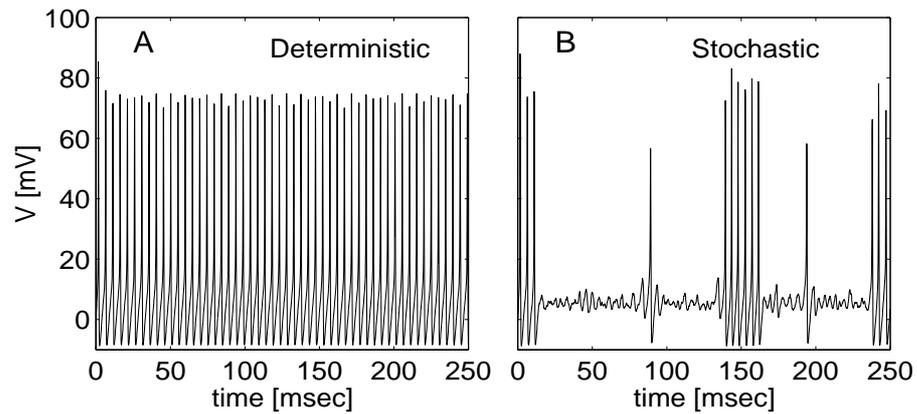


Figure 4.9: Temperature effect on the SHH model. (A) The response of a $200\mu m^2$ deterministic HH membrane patch to a $10\mu A/cm^2$ DC input with the temperature raised to $20^\circ C$. A high frequency regular spiking is seen. (B) The corresponding stochastic model responds with a highly irregular, ‘bursty’ spike train and subthreshold oscillations between the spikes.

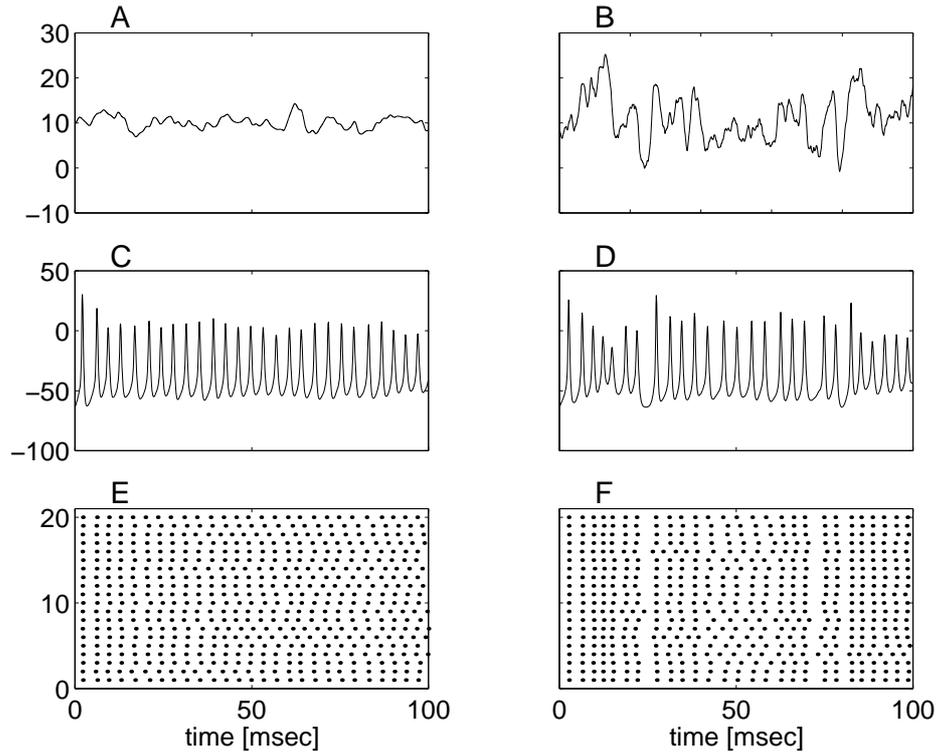


Figure 4.10: **Reliability of firing patterns in the SWB model.**

Two current traces were injected to a deterministic WB model and a stochastic version of the model. Both models are **(A)** A slowly varying input current (low-pass Gaussian white noise with a mean \bar{I} , of $10 \mu A/cm^2$, and a standard deviation σ_{input} of $3 \mu A/cm^2$ which was convolved with an 'alpha-function' with a time constant $\tau_{input} = 1 msec$). **(C)** The spike train of the deterministic model in response to the injection of the current trace in (A). **(E)** Raster plots of the spike trains of the stochastic model response to repeated injection of the current trace in (A). **(B)** A highly fluctuating input current (low-pass Gaussian white noise with a mean \bar{I} , of $10 \mu A/cm^2$, and a standard deviation σ_{input} of $15 \mu A/cm^2$ which was convolved with an 'alpha-function' with a time constant $\tau_{input} = 1 msec$). **(D)** Similar to (C), the response of the deterministic model to the injection of the current trace in (D). **(F)** Similar to (E), the responses of the stochastic model to the current trace in (B). Patch area used was $200 \mu m^2$, with 900 K^+ channels and 3,600 Na^+ channels.

Chapter 5

The nature of the neural code, information optimization and biophysical design in stochastic neuron models

5.1 Introduction

The computation that a single neuron performs, i.e. the mapping from the thousands of incoming spike trains into a single outgoing one, is implemented by a complex network of biophysical elements. The nature and result of the computation are defined by the parameters of the biological 'building blocks': the strength and dynamics of the synapses [162, 118, 3], the morphology and excitability of the dendritic tree [156, 124, 155, 172] and the types and number of ion channels in the soma and axon hillock [64, 112].

The design principles of this *biophysical computation* [88] must accommodate the computational 'task' of the neuron and the costs and limitations of using the biological machinery. First, the computation consumes significant amounts of energy, mainly due to spiking and synaptic activity [86]. Second, the synapses and ion channels are noisy and unreliable [168, 64]. More functional demands of the design may originate from the processes controlling the development of the neuron and its inherent activity-dependent plasticity [183, 43].

The effects of biological constraints on neural function, and possible optimal design strategies have been studied in various neuronal systems and modules. Metabolic costs and the efficiency of neural information coding have been studied for non-spiking and spiking neurons [100, 14, 95]. Efficient information coding has also been analyzed in the context of

computational constraints [15, 11]. The value of Na^+ conductance along the squid axon has been argued to be optimized for the conduction velocity of the action potentials [66, 9].

This chapter presents the study of the effects of the ion channel noise and other biophysical parameters of the spiking mechanism on the nature of the neural code. In chapter 4 we demonstrated that ion channel noise may control the reliability and accuracy of spike timing. One of the key conclusions was that ion channel noise may enable neurons to act as 'smart' encoders, where information may be encoded either by the spiking rate or the spiking temporal pattern, the code may be more of a rate code or a temporal code, as a function of the input. Based on these results (and claims), we proceed here to investigate and quantify the nature of the neural code of the stochastic spike generation models.

To that end, we simulate the responses of the stochastic versions of the spike generation models of Hodgkin and Huxley [71] and of Wang and Buzsaki [191] (see chapter 6), to the injection of various current stimuli. The (large) set of stimuli was chosen so that it would approximate (at least qualitatively) the currents that may reach the soma from the dendritic tree. Information theory tools [160, 35] are then used to quantify the performance of the spike generation mechanisms, measuring how much the spike trains entail about the stimulus the model was presented with. The nature of the spike trains as a code is analyzed by calculating how much information is carried by the average firing rate (rate code) and how much information is added by the temporal structure of the spike train (temporal code) [143, 171, 25], as well as various coding efficiency measures. The nature of the code is found to be variable, mixing rate and temporal components in an input-dependent manner, reflecting (again) on the decoding strategies that the nervous system may be using.

Following, we explore the effect of changing the biophysical parameters of the model on the information encoded by the model's spike trains. Reducing the level of noise which the ion channel stochasticity introduces results in a small gain in terms of the information rate encoded by the spike train about most stimuli, and in loss for other stimuli (a stochastic resonance like behavior [33, 18, 99]). The effects of changing the number of the different ion channel types (which changes the neuronal excitability) on the information encoding capabilities of the spike generation models is also studied. For most of the stimuli explored, there is an optimal combination of channels of the different types which maximizes the amount of information that the model encodes about its stimulus. These optimal channel densities lie well within physiological range, and are (for most stimuli) close to the 'standard' HH channel densities. We suggest then, that the spike generation mechanism may optimize the number of ion channels in the membrane so to maximize the average information it encodes about the stimuli it receives and discuss the possible constraints that metabolic costs may impose.

This chapter is based on work done in collaboration with Idan Segev and Naftali Tishby, parts of which have been published in [153].

5.2 Spike train properties of the SHH model

To investigate the neuronal spiking mechanism design and encoding characteristics, we simulate the responses of the stochastic HH model to a set of input stimuli. The “standard” stochastic model of reference we use in this chapter is a single $200 \mu m^2$ compartment soma which contains 3,600 K^+ and 12,000 Na^+ stochastic ion channels. The single channel conductance is taken to be $20 pS$ for both types, similar to the values found for cortical neurons [64]. As discussed in the previous chapter, this reference model replicates the input-dependent spike timing reliability and accuracy observed in real cortical neurons.

The set of stimuli we ‘inject’ into the model are current waveforms that approximate the synaptic currents that reach the soma from the dendritic tree (see [110] and chapter 2). Each stimulus is 10 – 20 seconds long, and is the result of convolving a Gaussian white noise trace (with a mean current η and standard deviation σ) with an alpha function (with a $\tau_\alpha = 3 msec$). Six different mean current values are used ($\eta = 0, 2, 4, 6, 8, 10 \mu A/cm^2$), and five different standard deviation values ($\sigma = 1, 3, 5, 7, 9 \mu A/cm^2$). Thus, the set contains 30 input current traces. We shall see that this set of stimuli spans the range of neuronal responses that real neuron exhibit. Thus, while it is by no means an exhaustive set of stimuli, it samples the space of possible stimuli in an ‘interesting’ way.

Figure 5.1 shows the average firing rates of the ‘standard’ SHH model in response to these stimuli. The firing rate range between 2 to 60 spikes per second, which is comparable to *in vivo* firing cortical rates.

A common measure of spike train irregularity, is the nature of the interspike interval distribution. Figure 5.2 shows the coefficient of variation of the interspike interval distribution (see chapter 3), as a function of the stimulus parameters. Similar to what is observed under *in vivo* conditions for cortical neurons, the coefficient of variation of the interspike interval distribution range from 0.3 to nearly 1, depending on the stimulus features (compare for example to [164, 180]).

5.3 Information rates of the SHH model

To quantify the coding properties of the SHH model, we proceed to calculate how much information is conveyed by the spike trains of the model about each of the stimuli it is

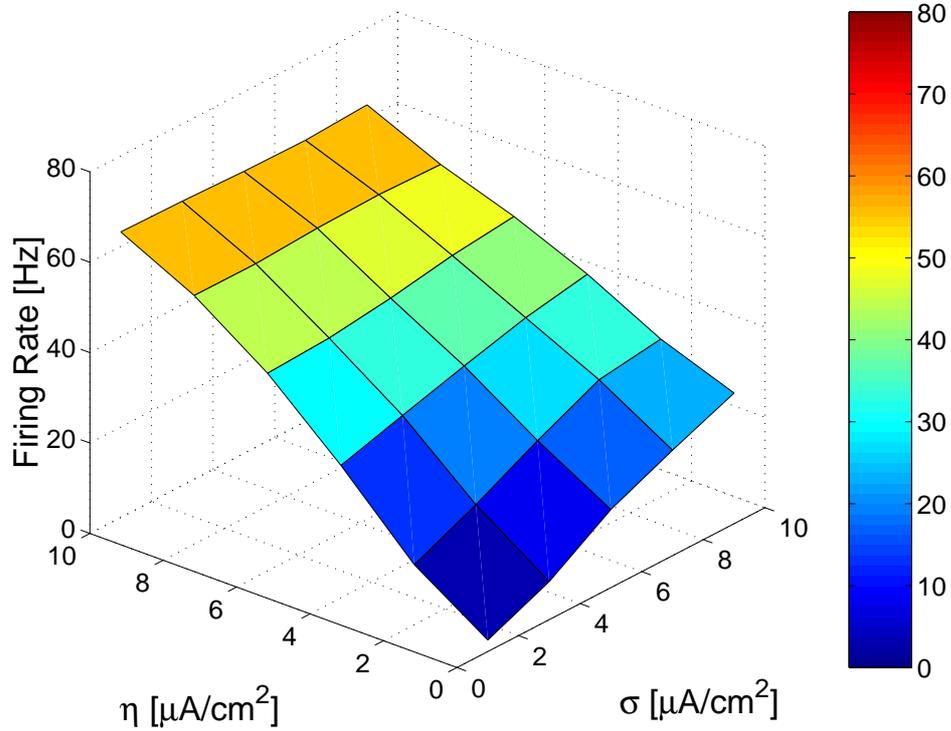


Figure 5.1: **Average firing rates of the SHH model in response to the set of stimuli.** Average spiking rate of the SHH model as a function of the current input parameters; η is the current input mean and σ is the standard deviation. The surface is an interpolation between the actual data points (which are on the surface grid).

presented with. Information theory provides a quantification of the relation between the input present to the model neuron and the spike trains that the model neuron responds with, without assuming what features of the code are important, or imposing a measure on the input-output relation (see chapter 3).

We use the measure of encoded information to complement the spike timing reliability and accuracy measures used in Chapter 4, thus quantifying the notion that accurate spiking is more informative.

Using the 'direct method' of calculating the information content of the spike train in response to a specific stimulus, following [39, 171]) (see also Chapter 3. Each of the stimuli is presented repeatedly to the model neuron and its spike trains are discretized, using bins of size $\Delta\tau$, into a binary sequence of zeros (no spike) and ones (spike). Using a sliding window of size T along the sequence, we get a sequence of K -letter binary 'words' ($K = T/\Delta\tau$). After estimating $P(W)$, the probability of the word W to appear in the spike trains, it is possible to compute the entropy rate of the total word distribution (of K -letter words),

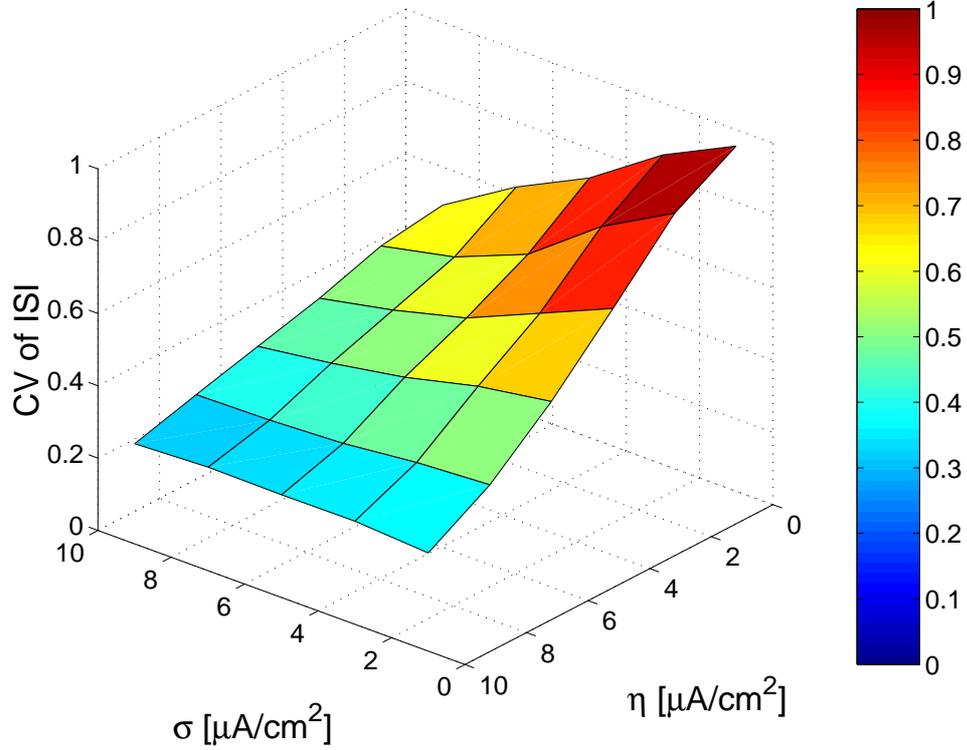


Figure 5.2: **Coefficient of variation of the interspike interval distribution of the SHH model in response to the set of stimuli.** Coefficient of variation of the inter-spike interval distribution of the SHH model in response to the current input parameters; η is the current input mean and σ is the standard deviation. The surface is an interpolation between the actual data points (which are on the surface grid).

$$H_T^{total} = - \sum_W P(W) \log_2 P(W) \quad \text{bits/word} \quad (5.1)$$

which measures the richness of the set of responses of the neuron to this stimulus. This calculation is repeated for different word sizes (different K 's). Taking the limit of infinitely long words (and normalizing by the word length), gives the spike train entropy rate, [35, 39, 171],

$$H^{total} = \lim_{T \rightarrow \infty} \frac{1}{T} H_T^{total}. \quad (5.2)$$

We then examine the set of words that the neuron model used at a particular time t over all the repeated presentations of the stimulus, and estimate $P(W|t)$, the time-dependent word probability distribution. At each time t we calculate the time-dependent entropy of the words, and then take the average (over all times) of these entropies,

$$H_T^{noise} = \langle - \sum_W P(W|t) \log_2 P(W|t) \rangle_t \quad \text{bits/word} \quad (5.3)$$

which measures how much of the fine structure of the spike trains of the neuron is just noise. $\langle \dots \rangle_t$ denotes the average over all times t . Again, this calculation is repeated for each of the inputs, using different word sizes (K values). Taking the limit of infinitely long words, gives the spike trains' noise entropy rate,

$$H^{noise} = \lim_{T \rightarrow \infty} H_T^{noise}. \quad (5.4)$$

The bin size $\Delta\tau$ was chosen to be 2 msec long, which is small enough to keep the fine temporal structure of the spike train within the word sizes used, yet large enough to avoid undersampling problems (see [171]).

Figure 5.3 shows the total entropy rate, H^{total} , of the standard SHH model in response to the set of stimuli, as a function of the stimulus parameters. The total spike train entropy range from 10 to 170 *bits/sec*.

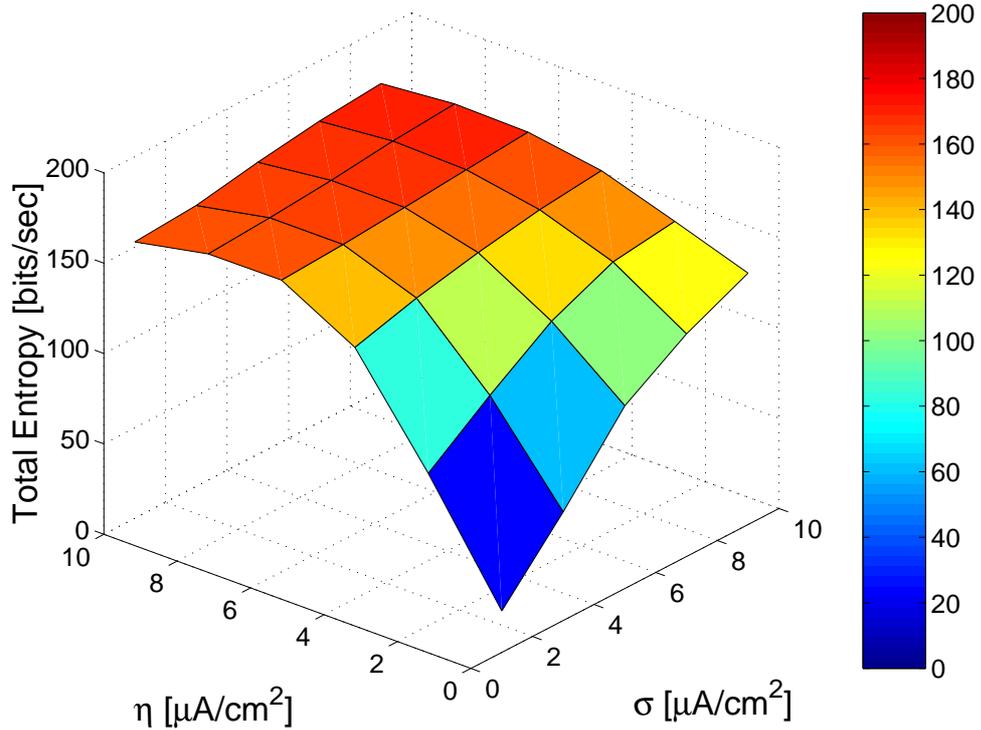


Figure 5.3: **Total entropy of the SHH spike trains** The total spike train entropy rate of the SHH model as a function of the input parameters. Error bar values range between 3 – 10% (not shown).

The noise entropy rate, H^{Noise} , depends differently on the input parameters, as shown in Fig. 5.4, and may get up to 100 *bits/sec*. Specifically, for inputs with high mean current values and low fluctuation amplitude, many of the spikes are just noise, even if the mean firing rate is high.

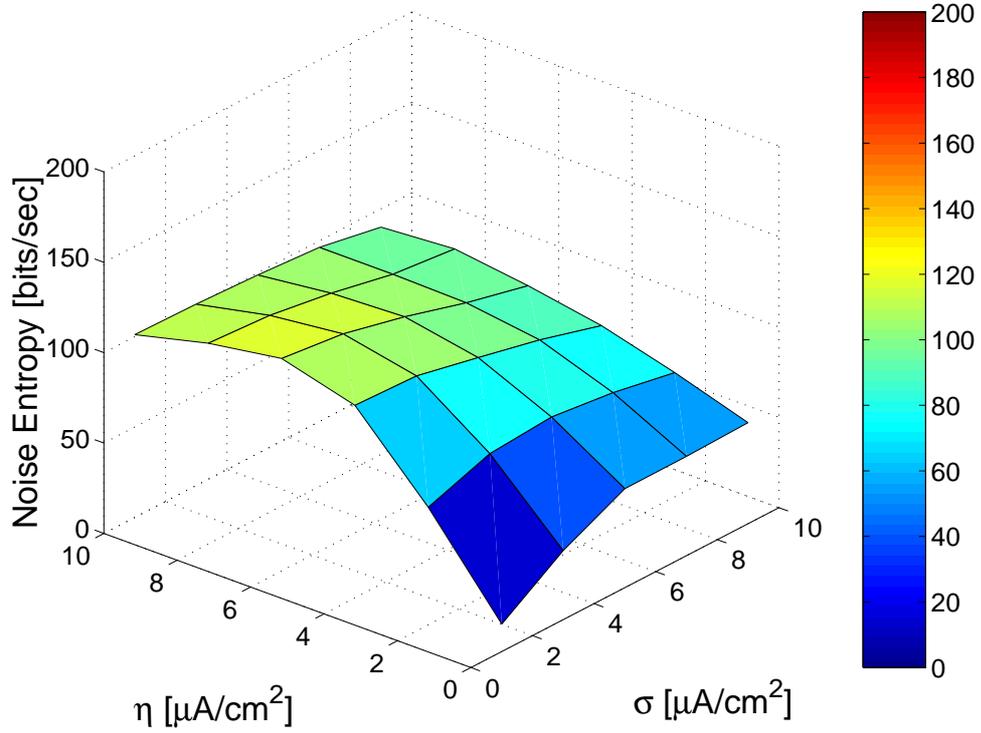


Figure 5.4: **Noise entropy of the SHH spike trains in response to the different stimuli.** Noise entropy rate as a function of the current input parameters. Error bar values of this surface range between 6 – 16% (not shown).

The difference between the neuron’s total entropy rate and the noise entropy rate, is the average information rate that the neuron’s spike trains encode about the stimulus, [35, 39],

$$I(\text{stimulus}, \text{spike train}) = H^{\text{total}} - H^{\text{noise}} \quad (5.5)$$

Figure 5.5a shows the information rate conveyed about each of the stimuli by the SHH model. The encoded information rate is more sensitive to the size of fluctuations in the input than to the inputs’ mean value, as may be expected from the spike train reliability results for cortical neurons [110]).

Dividing the information rate that the spike train of the model conveys about a certain stimulus, by the average spike rate of this spike train, gives the average information encoded per spike. Since spiking consumes considerable amounts of energy, this is a common measure of the “energetic efficiency” of the code. The resulting values, shown in Figure 5.5b, are similar to the values observed experimentally [22].

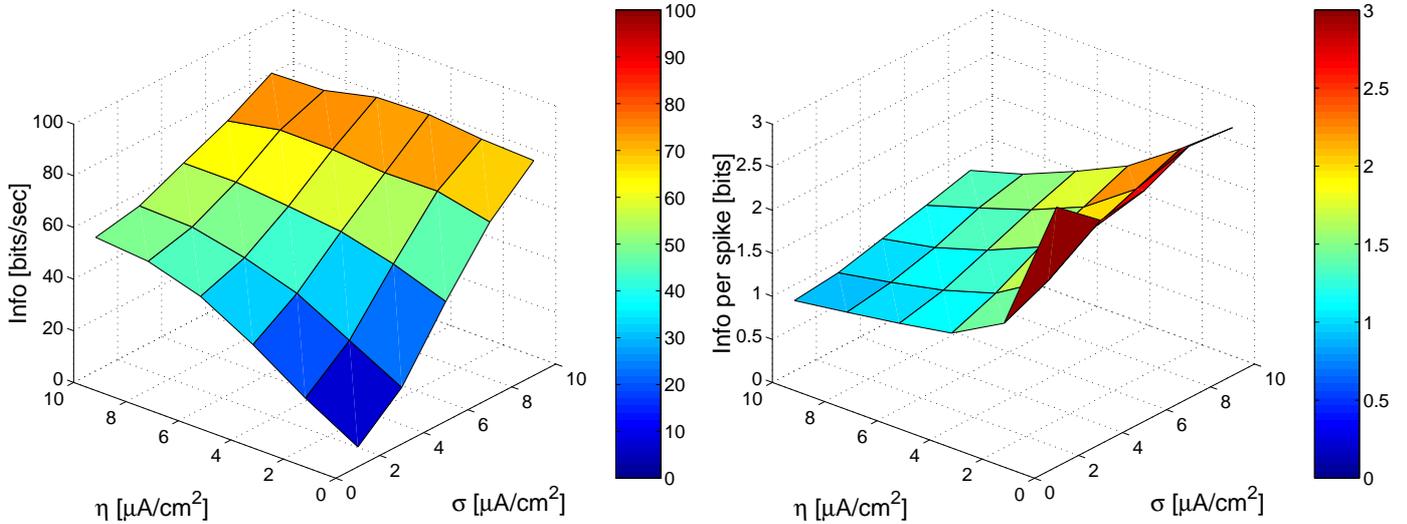


Figure 5.5: **Information rate and Information per spike encoded by the SHH model about the set of stimuli.** (a) The information rate about the stimulus in the spike trains, as a function of the input parameters, calculated by subtracting noise entropy from the total entropy. Error bars range between 6 – 14% (not shown). (b) Information per spike as a function of the input parameters, which is calculated by normalizing the results shown in a by the average firing rate of the responses to each of the inputs.

5.3.1 Coding efficiency, rate coding and temporal coding in the SHH model

Using the word probability distributions and the total word and noise entropies, we turn to characterize the nature of the neural code of the stochastic HH model, in terms of the coding efficiency and the contribution of the relative parts of the code in terms of rate and temporal components.

First, we compute a measure of coding efficiency which is the ratio between the amount of information carried by the spike train and the total spike train entropy. This measure ranges between zero and one: one extreme case is that the noise entropy is zero (i.e., the spike trains are deterministic), in which case all of the spike train structure is utilized to encode information about the stimulus, and the efficiency measure equals one. In the other extreme case, the spike train carries no information about the stimulus, and the efficiency is zero. Figure 5.6 shows the coding efficiency for each of the stimuli in the set, ranging between 0.2 to a little over 0.5 (compare to [171]).

Next, we address the ongoing argument over the temporal vs. rate coding schemes of the neural code (see Chapter 3). We measure directly how much of the information conveyed by the spike trains is carried just by the modulated firing rate and how much does the

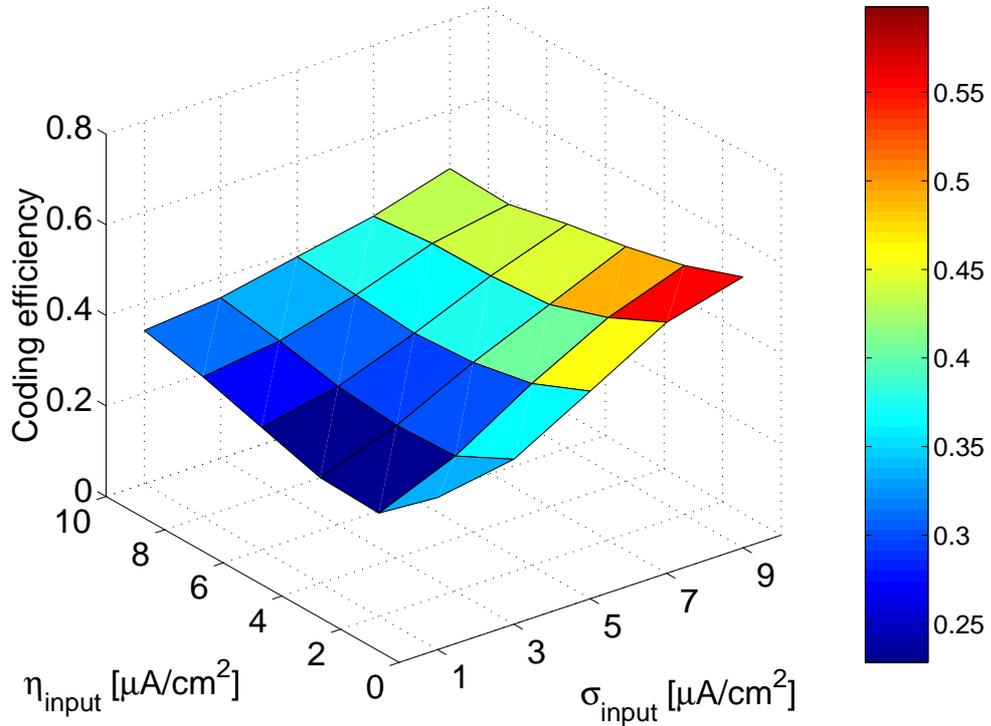


Figure 5.6: **Coding efficiency of the SHH model for the different stimuli.** The coding efficiency, measured by the ratio of the information to the total spike train entropy, as a function of the function of the stimulus parameters.

temporal structure of the spike train contributes, for each of the stimuli in our set. The information carried by the firing rate may be interpreted as the information that is carried by the individual spikes, without taking into account their relative timing (see [25]),

$$I^{\text{firingrate}} = \frac{\bar{r}}{T_{\text{stim}}} \int_0^{T_{\text{stim}}} dt \frac{r(t)}{\bar{r}} \log_2 \left[\frac{r(t)}{\bar{r}} \right] \text{ bits/sec}, \quad (5.6)$$

where \bar{r} is the average spike rate of the model and T_{stim} is the duration of the repeated stimulus. In terms of the word-based entropy calculations, this is equivalent to the information carried by words of length $K = 1$ (see also [140]). The difference between the total information and the information carried by the firing rate, is the information that the temporal structure of the spike train carries about the stimulus. Figure 5.7 shows the division between rate coding and temporal coding of the SHH model in response to each of the stimuli in the set. Each point marks the nature of the code which is used in response to a specific stimulus, showing the information rate that the temporal structure of the spike train (y-axis) adds to what is carried by the average firing rate (x-axis). It is obvious that the division between the rate coding portion and temporal coding one is highly variable, and depends on the nature of the stimulus. There is then no unique coding scheme of the SHH model, but rather an

input-dependent one. While in some cases most (or all) of the information is carried by the firing rate, for other stimuli, the temporal information may be more the twice as large as the information carried by the rate. The results for stimuli with the same mean (η) (connected by a line in Figure 5.7 for presentation purpose), reflect the non-trivial effect of the ratio between the size of the stimulus mean and std (σ) on the division between the temporal and rate coding components. As may be expected, we see that for highly fluctuating stimuli, the effect of the stimulus mean diminishes (compare to Fig. 4.4).

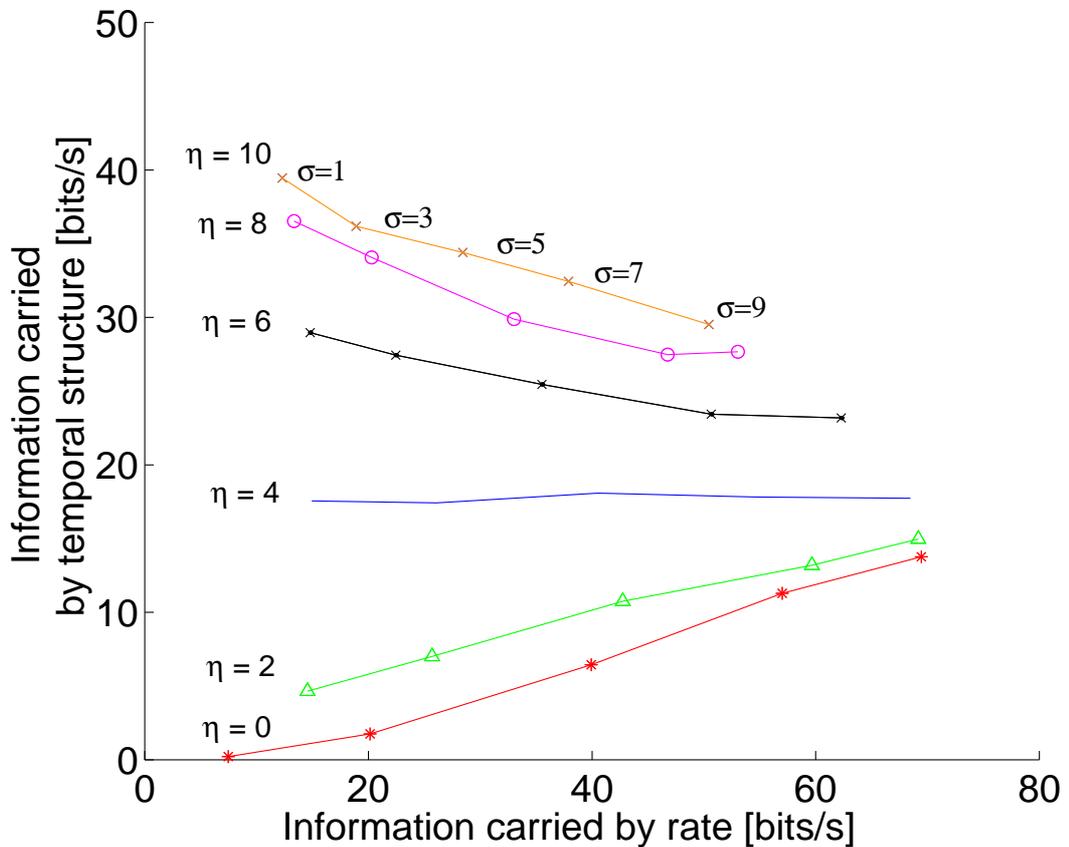


Figure 5.7: **Rate and Temporal components of the neural code of the SHH model.** For each of the stimuli, the information encoded by the temporal structure of the spike train (additional to what is coded by the rate) is shown as a function of the information encoded just by the rate modulation. The points corresponding to stimuli which had the same DC component are connected by line for presentation clarity (noted by their η value). The size of the fluctuating component of the stimuli is shown by the σ values for the top line (similar order holds for all lines).

5.4 Dependence of information encoding on the parameters of the model

Having seen that the stochastic HH model (with the standard channel density), replicates successfully the behavior of real neurons both in terms of the spike train reliability and information rates, and after quantifying the nature of the code, we turn to investigate the sensitivity of these results to changes in the (biophysical) parameters of the model.

5.4.1 The effect of reducing the ion channel noise on information encoding is not unidirectional

Given that ion channels are stochastic by nature, an obvious question regarding the design of the spiking mechanism is whether one could achieve a less noisy neuron by using the same basic building blocks, and if so, what might the gain and cost be? Since the size of the noise induced by the ion channel stochasticity is determined by the fluctuations in the number of open channels around the expected mean, one obvious way to reduce the ion channel noise is to reduce the single channel conductance and proportionally increase the number of ion channels, so that the total conductance is kept fixed. (Thus, at the limit of a very large number of channels with very small single channel conductance – keeping the total conductance fixed – one would approach the deterministic HH model [163, 42]).

Asking whether the spike initiation model will transmit more information about its inputs, if its ion channel noise were reduced in that way, we repeat the calculation of information rate encoded about the set of stimuli, for a family of stochastic HH models, with different ion channel noise levels. For the single channel conductance levels we span the full range of values reported for different types of channels and neurons ranging from a $0.1 pS$ to a few tens of pS [64]. Figure 5.8 shows the information rates of these models in response to a few representative stimuli. It turns out that for most stimuli, reducing the noise level of the model, would mean that the neuronal spike train would encode more information about the stimulus. When the single channel conductance level is reduced to $0.2 pS$ (i.e. having 100 times the number of channels of the reference stochastic model), the gain is on the order of 50%; At the limit of the deterministic HH model, the neuron model would encode a bit over that twice the information that the 'standard' stochastic model does.

Interestingly, for some of the inputs (mostly the ones with a low mean current value), reducing the noise would result in the neuron encoding less information about the stimulus. Such a *stochastic resonance* effect, where noise may improve the information encoding capabilities of a system, has been discussed theoretically (see e.g. [33]) and demonstrated

both in neuronal models and in real neurons (see e.g. [18, 99]). It is not obvious then, that reducing the level of noise that the ion channels introduce is beneficial from an information encoding point of view. Even when reducing the noise increases the encoded information rate, the improvement is not dramatic.

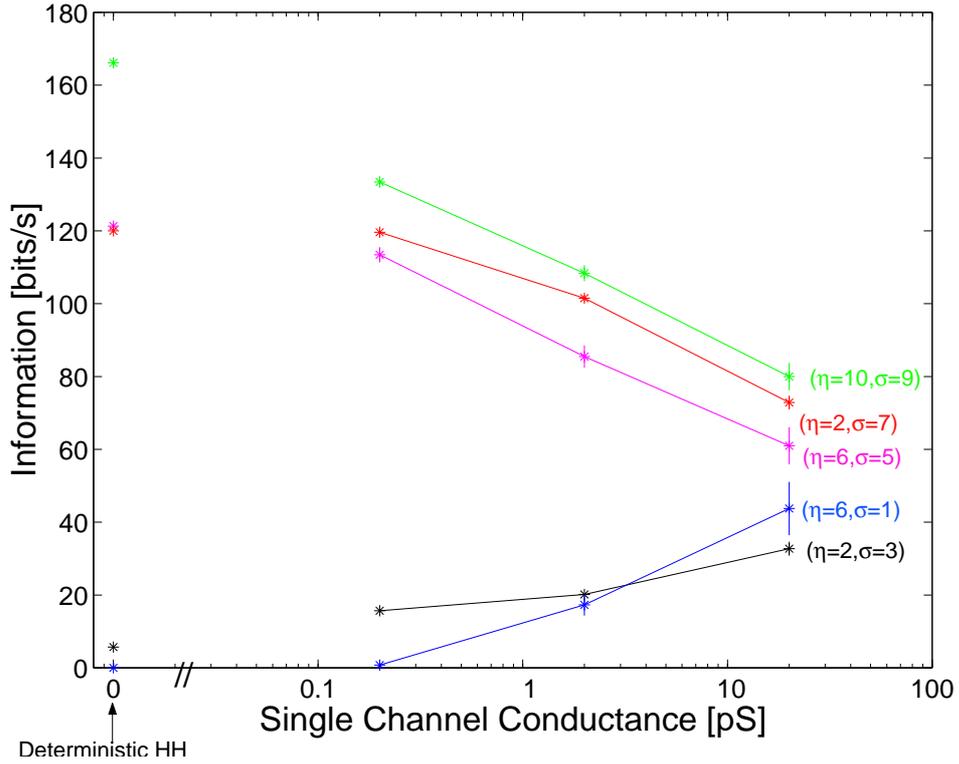


Figure 5.8: **Dependence of information rate on single channel conductance.** Information rates were calculated for SHH models with different single channel conductance level, while keeping the total ion conductance fixed (by proportionally increasing the number of channels). For relatively small values of σ and small to medium values of η (blue and black lines), information rate increases with increase in the single channel conductance, i.e. with larger channel noise. For larger values of σ (red, purple and green lines), increase in the channel noise decreases the information rate. The values for the corresponding deterministic HH model (where the single channel conductance approach zero), are marked by an arrow (on the left). Units of stimuli parameters η and σ are in $\mu A/cm^2$.

5.4.2 Changing channel densities and information encoding

The key parameters that control the nature of neuronal excitability are the numbers of ion channels of different types which reside in the neuron's membrane and are involved in the spike generation mechanism. Expressing more channels of different types and especially changing the ratio between the different types is likely to change the spiking patterns and

5.4. DEPENDENCE OF INFORMATION ENCODING ON THE PARAMETERS OF THE MODEL

information encoding properties of the neuron, as well as the nature of ion channel induced noise and the metabolic requirements of the neuron [64].

We therefore examine the effect of changing the number of sodium and potassium ion channels and the ratio between the two types of channels, on the nature of spiking and information encoding of the stochastic HH model and the stochastic Wang and Buzsaki model (SWB). We define a family of stochastic models based on the reference SHH model (and similarly for the SWB one) that differ in their Na^+ and K^+ ion channel densities and leaving the other parameters unchanged. To simplify the comparison between models, the ion channel densities of the different models will be given in units of the ion channel densities of the reference model (which result from the regular HH model). Thus in the 'normalized HH units' the ion channel densities of the reference model are both '1'.

We simulate the response of each of the models to the set of stimuli, and perform similar information rate calculations, as in the previous sections. Figure 5.9a shows the effect of changing the total Na^+ and K^+ channel densities, while keeping the ratio between them fixed. The top panel shows the average firing rate in response to the different stimuli as a function of the model channel density. Each of the curves connects the response values of 6 models to a certain stimuli (connecting the response values if for presentation purpose only). For clarity, the responses to only 9 representative stimuli are shown. The normalized channel densities of the reference SHH model (i.e. '1'), are marked by the vertical broken line. Clearly, for all the stimuli in the set, increasing the channel density of the model beyond the reference SHH one, results in a lower firing rate, and an optimum around half the standard channel density. The middle panel shows the information rate encoded by the models as a function of the channel density. Evidently, having more ion channels (of both types), results in encoding less information about the stimulus. For some of the stimuli the optimal channel densities are near the standard HH ones, and for some having even fewer channels would increase the information rate. Interestingly, the information encoded per spike (shown in the bottom panel) is robust to changes in the ion channel density, for all the stimuli used.

What would happen if we allow the ratio between the densities of the different ion channel types to change? Figure 5.9b shows firing and information rates where the Na^+ density is held fixed at the reference value, whereas the K^+ density changes. For all stimuli, the firing rate is monotonically decreasing when the K^+ density increases (top panel). For information encoding (middle panel) it turns out that for all examined stimuli, the optimal K^+ density is near the reference one (or even less). The information encoded per spike shows an opposite behavior. Thus, having more K^+ channels means that the neuron would encode less information about the stimulus it is presented with (although with higher efficiency).

Figure 5.9 shows the effect of changing the Na^+ density while keeping the K^+ density fixed at its reference value. As may be expected, having more Na^+ channels means that the model is more excitable, and the firing rate increases for all the stimuli (top frame). However, having more Na^+ channels (and a higher firing rate) does not mean a higher information rate. For most of the stimuli, the optimal Na^+ is near the reference one. For the rest of the stimuli, it is evident that the information rate is saturating. This means that the spike generation is becoming more and more noisy as the Na^+ increases. The information encoded per spike (top panel), reflects that indeed, the energetic efficiency is strongly decreasing with the Na^+ density.

To check the generic nature of these results, we also examine the dependence of information encoding on the ion channel densities for the stochastic Wang-Buszaki (WB) model [191] (as discussed in Chapters 2 and 4), the WB model of a cortical interneuron, has a different current-voltage relation, which is more similar to that of cortical neurons than the HH model.). Our reference stochastic WB (SWB) model is that of a single $200 \mu m^2$ compartment soma which contains 900 K^+ and 3,600 Na^+ stochastic ion channels. As for the SHH model, the single channel conductance is taken to be $20 pS$ for both types. (The reference model demonstrates the input-dependent spike timing reliability and accuracy, discussed in the previous chapter (see Figure 5.10).)

Similar to what was found for the SHH model, increasing the total number of ion channels of the stochastic WB model (while keeping the ratio between the different types fixed), results in a lower information rate conveyed by the outgoing spike trains. Keeping Na^+ density fixed and increasing the K^+ channel density results in a slight (saturating) increase in the information for many of the stimuli, suggesting that the optimal K^+ density is nearly twice that of the reference model (or closer to the reference model density for the rest of the stimuli used) . Increasing the Na^+ density while keeping the K^+ density fixed to its reference value, reflects a very sharp optimal Na^+ density for information encoding, lower than the reference model density.

5.5 Conclusions and discussion

5.5.1 The nature of the neural code

We have calculated the information rates and characterized the nature of the neural code utilized by the spike trains of the stochastic HH and stochastic WB spike generation models, which incorporate ion channel noise. For the stochastic HH model, the information rates encoded about a wide set of current inputs (imitating the current that may reach the soma

from the dendritic tree) depend on the nature of the stimulus, and range from nearly zero to 70 *bits/sec* (while the firing rate range between 2 to 60 *spikes/sec*). These values are similar to the information rates (and information per spike values) which were reported experimentally [22]. The coding scheme of the stochastic spiking models is found to be a combination of rate coding and temporal coding – for some stimuli, most of the information is carried by the firing rate, whereas for others the information carried by the temporal structure of the spike train may be much larger than the information carried by the rate. Complementing the results of the previous chapter, the nature of the neural code as well as the coding efficiency of the stochastic model are found to be highly input-dependent (compare to [110, 39]). Similar behavior was found for the stochastic WB model of an interneuron spike generation, although with higher firing rates and information rates.

These results suggest that the neural code is neither a "temporal code" nor a "rate code". If the same neuron uses very different coding strategies for different stimuli, combining rate and temporal components, then there is no stationary neural code – the symbols of the code stand for different 'messages' in different contexts. The experimental examples of accurate spike timing and the "contradicting" examples of noisy spiking patterns, may be just reflections of the different stimuli properties that the cells have been presented with. While the nature of spike timing, spike reliability and information content have been shown to be input-dependent [110, 127, 39], a direct measure of the coding strategy for different stimuli has not been done, to the best of our knowledge.

If real neurons use a non-stationary code *in vivo*, then efficient decoding of spike trains would require an adaptive synaptic activity or an adaptive voltage response of the decoding neuron. It would be interesting to study the nature of depressing and facilitating synapses [119]) as a reflection of such a decoding mechanism (see e.g. [118, 3, 181]). Similar questions arise in the contexts of redundancy of information encoding by population of neurons, and the individuality of the code for population of neurons (see next chapter).

5.5.2 Robustness and optimality of the spike generation mechanism

Studying the effects of changing the biophysical parameters of the models, we have found that for most stimuli, there is an optimal combination of ion channel densities that maximizes information encoding about that stimulus, within "physiological range". I.e., the optimal channel combination is well within 50 percent of the "standard" channel densities of the HH and WB models. The effect of changing other biophysical parameters, like the membrane area, the temperature or even the total ion channel densities (keeping the ratio between

sodium and potassium channels fixed) on the information encoding rate have been found to be of significantly smaller magnitude.

What is the functional significance of these optimal points? First, since the "task" of the neuron as a whole is to extract information from all the spike train it receives from other neurons (and by that to throw away information) and to propagate the result to other neurons, it is not clear that one should expect the spiking mechanism to be optimized for information encoding. Moreover, it could have been that an optimal channel combination for encoding information about these stimuli would lie well outside what one might consider as a physiological range. Second, when studying optimization principles for the neural code of neurons, it has often been suggested that neurons may optimize their information encoding properties under an energy consumption constraint [95, 14]. Finding an optimal ion channel distributions within physiological range suggests that this may not be a critical factor in setting the spiking mechanism behavior (see also below). These optimal points are not the result of metabolic constraints, but rather from the nature excitability of the spiking mechanism and its stochasticity. We find that the information encoded per spike (which is a measure of energy efficiency of the coding) may better shape the range of optimal ion channel densities found earlier.

In general, a neuron encounters a wide distribution of input stimuli, and therefore it would "make no sense" for the spike generating mechanism to be optimized for a specific stimuli, but rather to be optimized (if it was to be optimized at all) for that distribution. Accordingly, we find that the optimal ion channel densities combinations for different stimuli are not "far" from one another (in the space of ion channel densities). It is known that neurons replace their ion channels on a time scale of hours [64]. Moreover, recent experimental studies of cultured cortical neurons have shown that neurons may change their ionic conductivity in response to long periods of no spiking activity in the slice (or to excessive spiking) [43]. We therefore hypothesize that neurons may be able to adjust their ion channel expression patterns in order to maximize their information encoding, based on the distribution of stimuli they receives. We note the theoretical work of Stemmler and Koch [167] who suggested an on-line "learning" mechanism for dendritic conductances that would maximize the information that the soma would encode about the synaptic potentials. Thus, although learning is often regarded as a synaptic mechanism, the soma and spiking mechanism, may also play an active role in it.

5.5.3 Is ion channel noise a bug or an unavoidable feature?

As intuitively expected, for most of the stimuli, we have found that had it been possible for the neuron to reduce its ion channel noise by using more channels that have a lower conductance, the information rate encoded about these stimuli would increase. However, reducing the noise level to achieve a 50 % percent gain in information, requires expressing about 2 order of magnitude more ion channels. Even if packing such a large number of channels in the membrane was possible, it is not clear that the metabolic cost of expressing and maintaining so many channels in the cell membrane is practical. It is then possible that given the stochastic nature of single ion channels and the nervous' system's constraints of space allotted for single neurons, time of computation and energy consumption, alternative designs of a less noisy neuron are not "cost effective" and it is better to have several noisy and unreliable neurons devoted for the same computational task rather than a single accurate and noiseless one (see also [95]).

Other aspects of neural function and neural circuits' structure reflect on the effectiveness and value of constructing a reliable spiking mechanism (out of the known biophysical building blocks). First, the distributed nature of the nervous system (especially in mammals), which is robust to losing many neurons, suggest that it may be more efficient to use noisy neurons rather than "costly" reliable ones. Second, since the reliability of synapses is low [168, 169], it may be better not to invest in a very reliable and accurate spiking mechanism¹. Obviously, it is not clear how these different constraints and computational aspects have been weighted by evolution to drive the biophysical design, and it is currently hard to see an experimental system in which one could address these questions directly. Still, it may be interesting to compare the information encoding properties (firing patterns and content), and the biophysical parameters (noise level, neuronal size, energy consumption etc.) in different neuronal systems, and possibly compare species.

If cost effective consideration have shaped the design of the spike mechanism, it may be expected then that the nervous system will employ information encoding and processing schemes which would match the nature of the existing noise, or even take advantage of it. Corollary, we have seen that the contrary to the usual notion of noise as an obstacle, if the neuron were to be less noisy, its information encoding capabilities would diminish for some of its possible stimuli. Similar behavior, termed *stochastic resonance*, where adding noise to the input to a nonlinear system or to the system itself, results in response properties which convey more information about the nature of the input signal, has been observed in theoretical studies of models of neurons [33] as well as in real neurons [99] and in artificially

¹of course, this is somewhat of a 'chicken and egg' problem

created ion channel populations [18]. Again, in the case of spiking neurons, it is often possible to come up with "an alternative" biophysical design that would achieve the information encoding capabilities that the noise enables². However, the key point is that noise is not an added design element in neurons, but rather an inherent feature of the biological machinery. Interestingly, the "smart encoding" scheme, which we have suggested in the previous chapter, and the *synfire chain* model [5] may also be considered as giving a coding solution given a noise constraint.

5.5.4 Limitations of the current work and future directions

The current analysis of the biophysical design of the neuronal spiking mechanism and the resulting neural code, is limited in several aspects. First, we have used two specific models of neurons, and one may wonder whether our results are valid for other neuronal models, and ultimately, for real neurons. The stochastic HH and WB models include only sodium and potassium ion channels, and none of the other (less abundant) channel types which exist in real neurons [64]. Also, we have used isopotential neuron models, neglecting any spatial features of the spiking mechanism, and of the dendritic tree and the axon. We note that the models we have used belong to two different classes of spiking models [48]. It would then be surprising if it turns out that our results are due to some peculiarity of the HH and WB models, especially since these two models contain the core of most neuronal spiking models. Still, as discussed in the previous chapter, it would be interesting to add other ion channel types to the models and study the effects of the spatial and morphological properties of the neuron on the nature of the neural code. Especially interesting would be the addition of adaptation to the models, either through specific ionic currents (see for example [51]) or channel inactivation (e.g., [177]).

Second, our set of stimuli is a small sample of the full range of currents that may reach the soma from the dendritic tree. Thus, although the stimuli we used are considered as a good approximation to the currents that reach the soma from the dendritic tree, it is possible that we are missing important aspects of the correlations between synaptic activities, and filtering properties of the dendritic tree branches. Still, the responses to our set of stimuli span the range of firing rates, coefficient of variation of the interspike interval distribution and the information rates and information per spike rates reported for real neurons. Thus, while this is hardly an exhaustive set of stimuli, we argue it is interesting enough to reflect

²A common explanation of stochastic resonance behavior in neurons has suggested that adding noise to a subthreshold input could result in spiking, which would then convey more information about the stimulus. The common criticism raised in this aspect is that instead of relying on the noise, the neuron could have had a lower threshold to begin with.

on the nature of the neural code and the spiking properties of the model neurons. We expect that adding different types of stimuli would enrich our understanding of the effects of the biophysical design on the neural code, since even with the limited size of the set we were able to demonstrate the stimulus-dependent nature of information encoding in the spiking patterns and find different effects of noise reduction on the information encoding properties of the model neurons.

The main questions that stem from our results are whether (and how) real neurons indeed use different temporal and rate components in their neural code, as a function of the stimuli, and whether neurons optimize their ion channel population to improve their information encoding properties. We hope that experimental studies would address these questions by describing the response properties of neurons to rich set of stimuli, and by monitoring the changes of spiking patterns and of biophysical properties of neurons [43]. We suspect that such studies might reflect that the neural code is non-stationary by nature and may be constantly adapting on the molecular level.

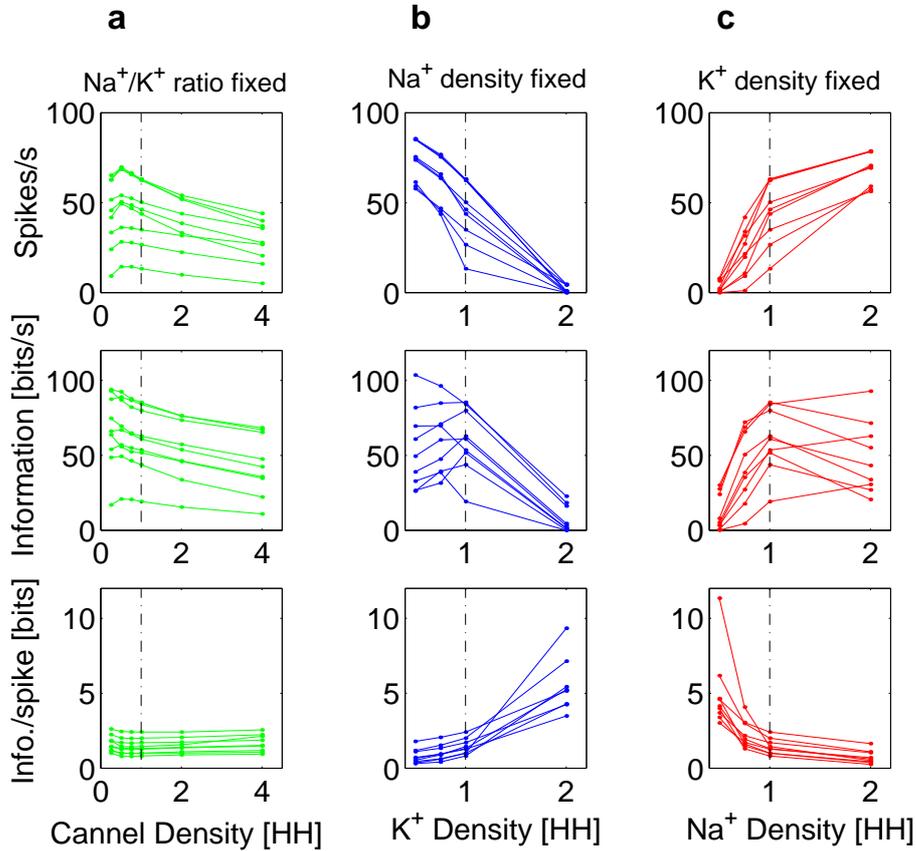


Figure 5.9: **Ion channels density effect on encoded information in the SHH model.** The firing rates and information rates that a family of SHH models (differing in their ion channel densities) encode about the set of stimuli were calculated. Denoting the initial SHH model as 'HH densities', the new models densities are measured in 'HH units'. **(a), top panel** The effect of changing both Na^+ and K^+ channel densities (keeping the ratio between the types fixed), on the firing rate of the model, shown for 9 representative stimuli out of the total stimuli set. The broken line at '1' denotes the 'regular' HH densities. The lines connect the points which correspond to the same stimuli (and are used for presentation purpose only). **(a), middle panel** Similar to a, the information rate encoded about each of the stimuli is shown as a function of the ion channel density. **(a), bottom panel** calculating the ratio between the information rate and the firing rate, the information encoded per spike is shown as a function of the model's ion channel density. **(b)** Similar to a, the firing rate, information rate and information per spike are shown as a function of the K^+ channel density used in the model (in this case the Na^+ density is kept fixed). **(c)** Similar to a and b, the firing rate, information rate and information per spike are shown as a function of the Na^+ channel density used in the model (K^+ channel density is kept fixed).

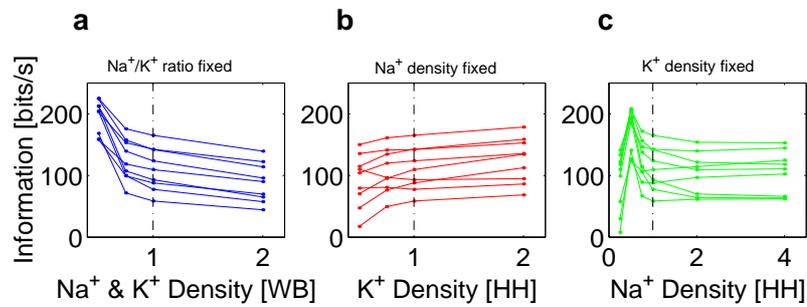


Figure 5.10: **Ion channels density effect on encoded information in the SWB model.** (a) The effect of changing both Na^+ and K^+ channel densities (keeping the ratio between the types fixed), on the information rate that the spike trains of the model convey, shown for 9 representative stimuli out of the total stimuli set. The broken line at '1' denotes the 'regular' WB densities. The lines connect the points which correspond to the same stimuli, are for presentation purpose only. rate encoded about each of (b) Similar to a, the information rate is shown as a function of the K^+ channel density used in the model (in this case the Na^+ density is kept fixed). c Similar to a and b, the information rate is shown as a function of the Na^+ channel density used in the model (K^+ channel density is kept fixed).

Chapter 6

Universality and individuality in a neural code

6.1 Introduction

When two people look at the same scene, do they see the same things? This basic question in the theory of knowledge seems to be beyond the scope of experimental investigation. An accessible version of this question is whether different observers of the same sense data have the same neural representation of these data: how much of the neural code is universal, and how much is individual? To approach this problem we must give a quantitative definition of similarity or distance among neural codes.

One way to quantify the similarities or differences among neural codes (i.e. sequences of spikes) is to imagine that each spike train is a point in an abstract space, and that there is a metric on this space [187, 188, 108]. Metric space methods have a long history of application to problems of comparing strings or sequences [98], including text. These are the standard tools for comparisons among nucleotide or amino acid sequences in molecular biology, where considerable effort has gone into the definition of metrics that are biologically plausible and computationally tractable [93]. While metric space approaches have been very useful, all such methods have several conceptual problems. First, the metric is imposed by the investigator and does not emerge from the data. Second, even within a plausible class of metrics there are arbitrary parameters, such as the relative distance cost of moving vs. deleting a spike. Finally, it is not clear that our intuitive notion of similarity among neural responses (or amino acid sequences) is captured by the mathematical concept of a metric. In contrast, we show how information theory [160, 134] can be used to quantify directly the differences among the *sources* of the sequences. This approach avoids any a priori assumption of a

metric on the sequence space and does not require a model of the process that generates the sequences.

We apply these methods to analyze experiments on an identified motion sensitive neuron in the fly’s visual system, the cell H1 [61]. Many invertebrate nervous systems have cells that can be named and numbered [27], and in many cases the total number of neurons involved in representing a portion of the sensory world is quite small, so that destruction of individual neurons can have a substantial impact on behavior (see, for example, Ref. [62]). In these cases the neural representation of sensory information is especially accessible, precisely because it is localized to a small set of identified cells. On the other hand, if a large fraction of neurons are identifiable it might seem that the question of whether different individuals share the same neural representation of the visual world would have a trivial answer.

Far from trivial, we find that the neural code even for identified neurons in flies has components which are common among flies and significant components which are individual to each fly. The existence of identified neurons thus does not preclude the expression of individuality in neural representations; we should expect that all neural circuits – both vertebrate and invertebrate – express a degree of universality and a degree of individuality. For H1 we quantify these ideas, and we hope that the methods we introduce will be applicable more generally.

This chapter is based on work with Naama Brenner, Naftali Tishby, Rob de Ruyter van Steveninck and William Bialek, parts of which has been published in [150]

6.2 An ensemble of flies and the experimental setup

We place our discussion in the context of the experiments shown in Fig. 6.2a. Nine different flies are shown precisely the same movie, which is repeated many times for each fly; as we show the movie we record the action potentials from the H1 neuron.¹ Details of the

¹Recordings were made from the H1 neuron using standard methods: the fly was immobilized in wax, a tungsten microelectrode was inserted through a small hole at the back of the fly’s head, and H1 was identified through its response properties; spikes were detected with a window discriminator. The stimulus was a rigidly moving pattern of vertical bars, randomly dark or bright, with average intensity $\bar{I} \approx 100\text{mW}/(\text{m}^2 \cdot \text{sr})$, displayed on a Tektronix 608 high brightness display; bar widths were set equal to the horizontal lattice spacing (interommatidial angle) of the compound eye. The fly viewed the display through a round diaphragm, showing approximately 30 bars. Frames of the stimulus pattern were refreshed every 2ms , and with each new frame the pattern was displayed at a new position. This resulted in an apparent horizontal motion of the bar pattern, which is suitable to excite the H1 neuron. The pattern position was defined by a pseudorandom sequence, simulating a diffusive motion or random walk. Experiments were performed by Rob de Ruyter van Steveninck and Geoff Lowen at NEC research institute, Princeton, NJ, USA.

stimulus movie should not have a qualitative impact on the results, provided that the movie is sufficiently long and rich to drive the system through a reasonable range of responses. Figure 6.2b makes clear that the responses of different flies to the same movie indeed are different: average spike rates are different, the patterns of rate modulation are different, and the degree of reliability from presentation to presentation is itself variable from fly to fly. Despite these differences, there also are some common features, both on long time scales (~ 100 ms) and in the detailed pattern of spikes on the (few) millisecond time scale. At this qualitative level, some pairs of flies (e. g., 6 and 7) seem more similar, others (e. g., 7 and 8) very different. The goal of this paper can be phrased as the problem of quantifying these observations.

Before proceeding to characterize individual variations, we must specify the “ensemble of flies” in our experiment. All of the flies are freshly caught female *Calliphora*, so that our ensemble of flies approaches the natural one and is not restricted to a highly inbred laboratory stock. We tried to minimize obvious sources of variation by doing all experiments at the same time of day, and by analyzing only stationary segments of recording that are obtained in between feedings of the animals. Recordings are rejected only if raw electrode signals are excessively noisy or unstable; in particular we do not select for flies that exhibit mean spike rates (spontaneous or driven) in a predefined range. Even with these precautions, there remain questions about whether observable differences among individuals should be ascribed to what we colloquially call “individuality” or (merely) to differences in “internal state” at the time of the experiment; one way to address this is to ask if differences are associated persistently with the individuals or if they fluctuate across experiments on the same individual at different times. In several cases we have done experiments over several days on the same flies, and the day to day variations of H1’s response within flies are significantly smaller than the variations from fly to fly. We therefore believe that the individual variations among flies in our experimental ensemble provide a fair sampling of the individual differences in nature.

The conventional measure of the neural response is the average spike rate, which varies from 22 to 63 spikes/s among our ensemble of flies. One might argue that these differences of spike rate are sufficient to establish the existence or nonexistence of individual variations, but this would be missing much of the structure in the data. First, even the largest differences of mean rate are not so large when seen in context of the rate variations across time within each fly. Second, we shall see that differences in how particular patterns of spikes are associated with visual inputs provide much more information about individual identity than that carried by the rate alone. Finally, flies with similar mean spike rates nonetheless can be distinguished based on more detailed analysis. Thus, while the mean rate provides a convenient label for each fly in our ensemble, we turn to a more quantitative approach.

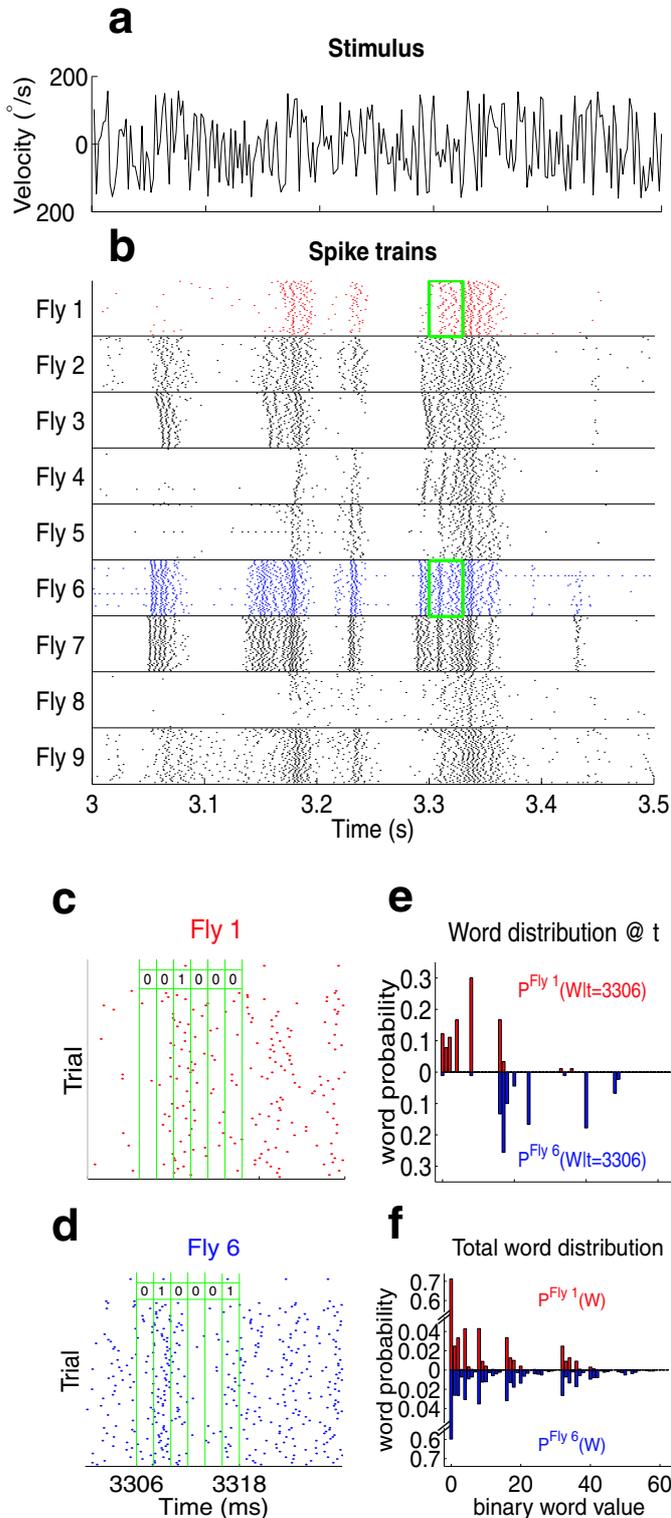


Figure 6.1: **Different flies' spike trains and word statistics.** (a) All flies view the same random vertical bar pattern moving across their visual field with a time dependent velocity, part of which is shown. In the experiment, a 40 sec waveform is presented repeatedly, 90 times. (b) A set of 45 response traces to the part of the stimulus shown in (a) from each of the 9 flies. The traces are taken from the segment of the experiment where the transient responses have decayed. Spike trains from flies 1 and 6 are colored by red and blue, respectively, which we will use as a color code for the other parts of the figure. (c) Example of construction of the local word distributions. Zooming in on a segment of the repeated responses of fly 1 to the visual stimuli (see green rectangle in (b)), the fly's spike trains are divided into contiguous 2 ms bins, and the spikes in each of the bins are counted. E.g., we get the 6 letter words that the fly used at time 3306 ms into the input trace. (d) Similar to (c) for fly 6. (e) The distributions of words that flies 1 and 6 used at time $t = 3306$ ms from the beginning of the stimulus. The time dependent distributions, $P^1(W|t = 3306 \text{ ms})$ and $P^6(W|t = 3306 \text{ ms})$ are presented as a function of the binary value of the actual 'word', e.g., binary word value '17' stands for the word '010001'. (f) Collecting the words that each of the flies used through all of the visual stimulus presentations, we get the total word distributions for flies 1 and 6, $P^1(W)$ and $P^6(W)$.

6.3 Distinguishing among individuals

Differences among neural codes may arise from various sources: First, different individuals may use different sets of coding symbols, or more subtly, the underlying symbols may be the same but the errors or noise in expression of the symbols may be individualistic. Second, they may use the same symbols to encode different stimulus features. Third, they may have different latencies, so they ‘say’ the same things at slightly different times. Finally, perhaps the most interesting possibility is that different individuals might encode different features of the stimulus, so that they ‘talk about different things.’ These different sources of variation obviously have different implications for the individuality of neural representation. We begin by quantifying the magnitude of the differences among individuals.

It is convenient to think about the responses of a neuron as being like the words in a language. One clear question then is whether different flies speak with the same vocabulary. If we can label each possible neural response as a “word” W , then the vocabulary is characterized by the distribution of words used by a particular fly’s H1 in response to the stimulus movie, $P^i(W)$ for the i^{th} fly. Evidently comparing vocabularies involves measuring the similarity among the probability distributions $P^i(W)$, and there are many possible ways of doing this. If we knew the correct family of models for sequences or words that we are trying to compare, then of course we could use specialized measures that are natural within the class of models. However, without such knowledge, and trying to avoid making apriori assumptions, we turn to information theoretic quantities such as entropy and mutual information.² In simple cases (such as χ^2 measures of difference in Gaussian distributions) it is easy to see that the “natural” measures of distance are related monotonically to the information theoretic measures. Thus the information theoretic approach automatically includes the familiar metrics, but generalizes beyond the cases where these particular metric assumptions can be justified. (The cost of this generality is that we have to sample the relevant probability distributions, which may be prohibitive in some applications.) In the present case the relevant quantity is the *information about identity*: the amount of information, in bits, that spoken words provide about the identity of the speaker. We emphasize that this is the unique measure of dissimilarity among vocabularies that is independent of any model for the structure of the underlying code.

Words are discrete objects, while spike arrival times are continuous. It is convenient, then,

²With one exception, each possible measure of (dis)similarity will produce pathological results for some set of distributions that we might observe in an experiment. As shown by Shannon [160], the only measures that work in all cases are constructed from information theoretic quantities such as entropy and mutual information

to discretize the neural response, although we shall have to make sure that our conclusions do not depend on this discretization. Specifically, we discretize the neural response into time bins of size $\Delta t = 2$ ms; at this resolution there are almost never two spikes in a single bin, so we can think of the neural response as a binary string, as in Fig. 6.2c–d. We examine the response in windows of time T , so that an individual neural response becomes a binary word W with $T/\Delta t$ letters. Any fixed choice of T and Δt is arbitrary, and so we explore a range of these parameters.

Once we have discretized the neural response into words, it is natural to compare the “vocabularies” of different flies. This can be done by looking at the distribution of words used by a particular fly’s H1 in response to the stimulus movie, $P^i(W)$ for the i^{th} fly. Figure 6.2f shows that different flies ‘speak’ with similar but distinct vocabularies. On the other hand, Fig. 6.2e shows that at the same time in the stimulus movie, different flies may choose to use very different words out of these similar vocabularies.³ As with human speech, we might imagine that individuals could be identified by their total vocabulary, or more efficiently by their choice of words in particular situations. This identifiability by words can be quantified by asking how much information, in bits, spoken words provide about the identity of the speaker.

Imagine that we record multiple speakers reading from the same text, in the same way that we record the activity of neurons from different flies responding to the same sensory inputs. There are many possible speakers, and we are shown a small sample of the speech signal: how well can we identify the speaker? If we can collect enough data to characterize the distribution of speech sounds made by each speaker then we can quantify, in bits, the average amount of information that a segment of speech gives about the identity of the speaker. Further we can decompose this information into components carried by different features of the sounds. Following this analogy, we will measure the information that a segment of the neural response provides about the identity of the fly, and we will ask how this individuality is distributed across different features of the spike train. From these distributions $P^i(W)$ we can quantify the average information that a single word of length T gives about the identity of the fly, $I_T(W; \text{identity}; T)$,

$$I_T(W; \text{identity}) = \sum_{i=1}^N P_i \sum_W P^i(W) \log_2 \left[\frac{P^i(W)}{P^{\text{ens}}(W)} \right] \text{ bits},$$

³It is important to realize that discretization allows us to exhibit the neural response as a ‘word,’ but does *not* tell us which words are similar; in particular the ordering of words along the x axis in Figs. 6.2e&f is arbitrary.

$$P^{\text{ens}}(W) = \sum_{i=1}^N P_i P^i(W). \quad (6.1)$$

where $P_i = 1/N$ is the a priori probability that we are recording from fly i and $P^{\text{ens}}(W)$ is the probability that any fly in the whole ensemble of flies would generate the word W . Thus, we measure how well we can discriminate between one individual and a mixture of all the other individuals in the ensemble, or effectively how ‘far’ each individual is from the mean of her conspecifics.

The measure $I_T(W; \text{identity})$ has been discussed by Lin [101] as the ‘Jensen–Shannon divergence’ D_{JS} among the distributions $P^i(W)$, namely $D_{\text{JS}}(P^1(W), P^2(W), \dots, P^N(W))$. We recall that the problem of finding a measure of similarity among distributions is not simple; obvious choices such as the Kullback–Leibler [35] divergence are not symmetric, and may have spurious technical requirements such as absolute continuity of one distribution with respect to the others. Lin [101] and Guttman [58] proposed D_{JS} as a way of getting around these difficulties, and showed that D_{JS} can be used to bound other measures of similarity, such as the optimal or Bayesian probability of identifying correctly the origin of a sample (as in forced choice psychophysical discrimination experiments). Here D_{JS} arises not just as an interesting possible measure of similarity (see also [137]), but as the unique answer to the question of how much information a sample provides about its source.⁴

The finite size of our data set prevents us from exploring arbitrarily long words, but happily we find that information about identity is accumulating at a more or less constant rate well before the undersampling limits of the experiment are reached (Fig. 6.2a). Thus,

$$I_T(W; \text{identity}) \approx R_T(W; \text{identity}) \cdot T \quad (6.4)$$

and $R(W; \text{identity}) \approx 5$ bits/s, with a very weak dependence on the time resolution Δt . Since the mean spike rate can be measured by counting the number of 1s in each word W , this information includes the differences in firing rate among the different flies.

Even if flies use very similar vocabularies, they may differ substantially in the way that they associate words with particular stimulus features, as is clear from the comparison of

⁴Given two empirical probability distributions (samples) $p(x)$ and $q(x)$, for every $0 \leq \lambda \leq 1$ their λ -JS divergence is defined as

$$D_{\lambda}^{\text{JS}}(p(x)||q(x)) = \lambda D^{KL}(p(x)||r(x)) + (1 - \lambda) D^{KL}(q(x)||r(x)), \quad (6.2)$$

where

$$r(x) = \lambda p(x) + (1 - \lambda) q(x) \quad (6.3)$$

can be shown to be the most likely source of both $p(x)$ and $q(x)$ [137], with λ as a prior. Without a-priori information about the relative likelihood of p and q we use $\lambda = \frac{1}{2}$. The JS divergence has various attractive properties. First, it is symmetric. But more importantly, unlike the KL-divergence, it is bounded.

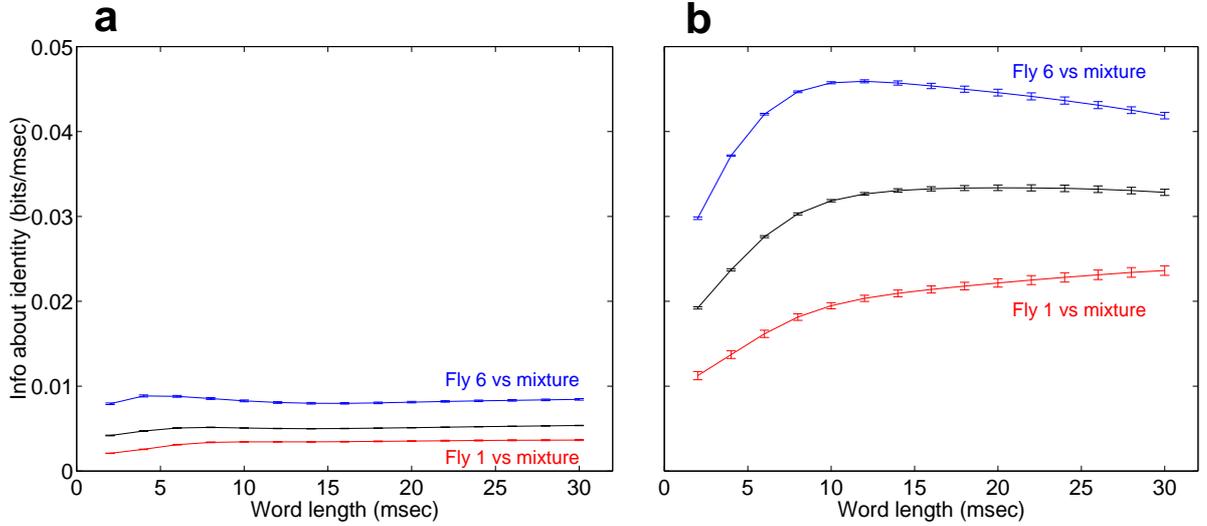


Figure 6.2: **Distinguishing one fly from others based on spike trains.** (a) The average rate of information gained about the identity of a fly, given the distribution of words that it used throughout the stimulus presentations, as a function of the word size used. The information rate is saturated even before we reach the maximal word length used. Following Figure 6.2, Red marks are the average rate of information that the word distribution of fly 1 give about its identity, compared with the word distribution mixture of all of the flies. The connecting line is used for clarification only. Blue marks the results for fly 6, and the black marks the average over all 9 flies. (b) Similar to the computation done for (a), we can compute the average amount of information that is gained about the identity of the fly, give its word distribution at a specific time, compared with the mixture of the word distribution of all of the 9 flies. Averaging over all times, we get the average amount of information gained about the identity of fly 1 based on its time dependent word distributions (red), fly 6 (blue), and the average over the 9 flies (black).

Figs. 6.2e and 6.2f. In our experiments the stimulus runs continuously in a loop, so that we can specify the stimulus precisely by giving the time relative to the start of the loop; in this way we don't need to make any assumptions about which features of the stimulus are important for the neuron, nor do we need a metric in the space of stimuli. We therefore can consider the word W that the i^{th} fly will generate at time t . This word is drawn from the distribution $P^i(W|t)$ which we can sample, as in Fig. 6.2c–e, by looking across multiple presentations of the same stimulus movie. In parallel with the discussion above, we can now ask for the average information that a word W provides about identity given that it was observed at a particular time t ,

$$I_T(W; \text{identity}|t) = \sum_{i=1}^N P_i \sum_W P^i(W|t) \log_2 \left[\frac{P^i(W|t)}{P^{\text{ens}}(W|t)} \right],$$

$$P^{\text{ens}}(W|t) = \sum_{i=1}^N P_i P^i(W|t). \quad (6.5)$$

This information depends on the time t because some moments in the stimulus are more informative than others, as is obvious from Fig. 6.2.⁵ The more natural quantity is an average over all times t , which is the average information that we can gain about the identity of the fly by observing a word W at a known time relative to the stimulus,

$$I_T(\{W, t\}; \text{identity}) = \langle I_T(W; \text{identity}|t) \rangle_t \text{ bits}, \quad (6.6)$$

where $\langle \dots \rangle_t$ denotes an average over t .

Figure 6.2b shows that observing the spike train at a known time during the stimulus movie provides 32 ± 1 bits/s about the identity of the fly. This is more than six times as much information as we can gain by observing the spike train alone, and corresponds to gaining one bit in ~ 30 ms. This one bit of information refers to identifying individuals in the full experimental ensemble; we can also ask about the discrimination between a typical pair of flies in the ensemble, and this pairwise (one bit) discrimination also becomes reliable in ~ 30 ms. This is the time scale on which flies actually use their estimates of visual motion to guide their flight [92], so that the neural codes of different individuals are distinguishable on the time scales relevant to behavior.

6.4 Spike rates and information rates of the H1 of different flies

Having seen that we can distinguish reliably among individual flies using relatively short samples of the neural response, it is natural to ask about the origins and implications of these individual differences. In particular, the vocabularies of the different flies are quite similar, but the way in which words are associated with stimulus features is much more individualistic. This is the statement that $I_T(\{W, t\}; \text{identity}) \gg I_T(W; \text{identity})$. This association of words with features is at the heart of the neural code, and it would be surprising if individuality in this association did not have implications for the representation of visual information.

In the previous discussion we measured the amount of information that the neural response provides about the identity of the fly given that we have access to the stimulus. Of course the neural response is not “designed” to represent identity, but rather the stimulus itself. Thus we can ask how much information the neural response of an individual fly provides about the stimulus. As discussed in Refs. [39, 171] (and chapter 3), this information,

⁵As above, $I_T(W; \text{identity}|t)$ is also $D_{\text{JS}}(P^1(W|t), P^2(W|t), \dots, P^N(W|t))$.

$I_T^i(W; s(t))$, is determined by the same probability distributions $P^i(W|t)$ as before:

$$I_T^i(W; s(t)) = \left\langle \sum_W P^i(W|t) \log_2 \left[\frac{P^i(W|t)}{P^i(W)} \right] \right\rangle_t. \quad (6.7)$$

As with information about identity, we expect that information about the stimulus grows with the duration of our observations, so that $I_T^i(W; s(t)) \approx R^i(W; s(t)) \cdot T$.

Figure 6.3a shows that the flies in our ensemble span a range of information rates from $R^i(W; s(t)) \approx 50$ to ≈ 150 bits/s. This threefold range of information rates is correlated almost perfectly with the range of spike rates, so that each of the cells transmits nearly a constant ($\pm 10\%$) amount of information per spike, 2.39 ± 0.24 bits/spike. Although bit rates and spike rates vary among individuals, there is a surprising universality of the ratio bits/spike. The fact that the information per spike is constant across the ensemble of flies means that cells with higher spike rates are not generating extra spikes at random, but rather each extra spike is equally informative about the visual stimulus.

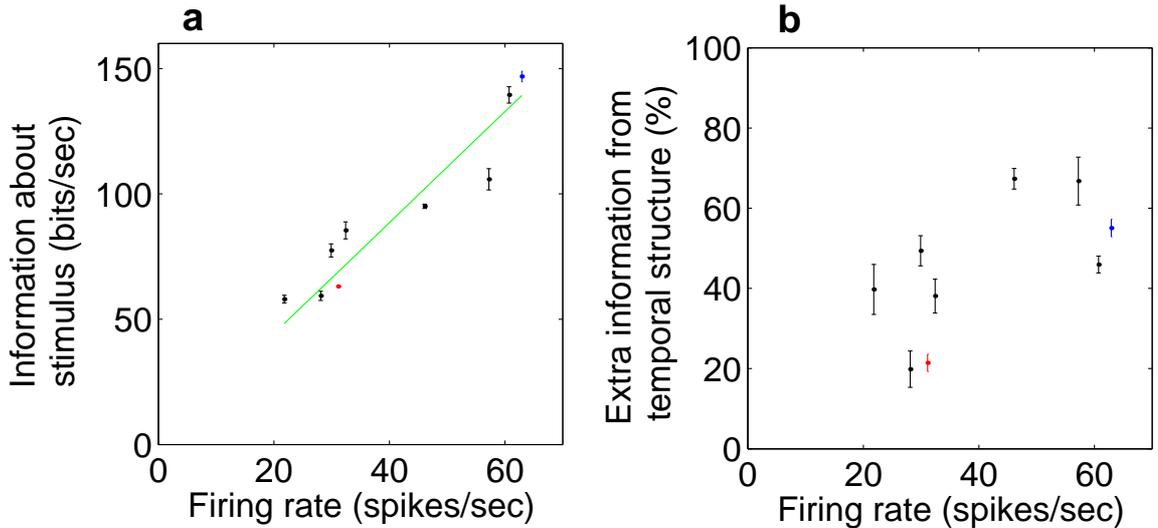


Figure 6.3: **The information about the stimulus that a fly’s spike train carries is correlated with firing rate, and yet a significant part is in the temporal structure.** (a) The rate at which the H1 spike train provides information about the visual stimulus is shown as a function of the average spike rate, with each fly providing a single data point (Fly 1 is marked by a red point and Fly 6 by a blue one). The linear fit of the data points for the 9 flies corresponds to a universal rate of 2.39 ± 0.24 bits/spike, as noted in the text. (b) The extra amount of information carried by the temporal structure of the spike train of each of the flies, as a function of the average firing rate of the fly. The average amount of additional information that is carried by the temporal structure of the spike trains, over the population is $45 \pm 17\%$.

The *capacity* of an individual code to carry information is quantified by the total entropy

of the distribution of neural responses,

$$S_{\text{total}}^i(T) = - \sum_W P^i(W) \log_2 P^i(W) \text{ bits.} \quad (6.8)$$

It is natural [143, 142] to define the efficiency of the code as the fraction of this capacity which is used to convey information about the visual stimulus, $\epsilon^i = I_T^i(W; s(t)) / S_{\text{total}}^i(T)$. Like the information per spike, this efficiency is nearly constant across the ensemble of flies, $\epsilon = 0.59 \pm 0.05$ at $\Delta t = 2$ ms, with a very weak dependence on Δt [171].

Although information rates are correlated with spike rates, this does not mean that information is carried by a “rate code” alone. Rate coding usually is distinguished from “timing codes” in which the detailed temporal structure of the spike train plays a crucial role. There are two senses in which the timing of action potentials could be important. First there is the simple question of whether marking spike arrival times to higher resolution really allows us to extract more information about the sensory inputs. The average information carried by the spike trains of H1 under these conditions is enhanced significantly when we analyze the responses with a resolution of $\Delta t = 2$ ms rather than $\Delta t = 4$ ms, and if we reduce our resolution to $\Delta t = 10$ ms we lose more than half of the information; in this sense, timing is important down to a scale of a few milliseconds [171]. A second notion of spike timing being important is that temporal patterns of spikes may carry more information than expected by adding the information carried by the individual spikes. We can address this by measuring the information carried by the arrival time of a single spike, independent of its temporal relation to the other spikes. This ‘single spike information’ can also be thought of as the information conveyed by temporal modulations in the spike rate, and can be written as an integral over the time dependent rate $r_i(t)$ for each fly [25, 140]

$$I_{\text{one spike}}^i = \frac{1}{T_{\text{loop}}} \int_0^{T_{\text{loop}}} \frac{r_i(t)}{\bar{r}_i} \log_2 \left[\frac{r_i(t)}{\bar{r}_i} \right] dt \text{ bits,} \quad (6.9)$$

where \bar{r}_i is the average spike rate in cell i and T_{loop} is the duration of the repeated stimulus movie. For all the flies in our ensemble, the total rate at which the spike train carries information is substantially larger than the ‘single spike’ information—2.39 vs. 1.64 bits/spike, on average. This extra information, shown in Fig. 6.3b, is carried in the temporal patterns of spikes.

6.5 Using a universal codebook

Although different flies encode different amounts of visual information in the spike trains of their H1 neurons, we expect that there are aspects of this information common to all the flies.

The neural code for any individual fly can be thought of as a probabilistic mapping from neural responses or words back into the space of visual stimuli [38]. The information conveyed by the spike train quantifies the specificity of this mapping: the tighter the distribution of stimuli consistent with a given response the more information is conveyed. This information [Eq. (6.7)] is available in full only to an observer who knows the mapping from responses to stimuli. If the neural codes used by different flies are different, then these conditional distributions in stimulus space also are different. The idea that there is something universal in the code used by all flies means that the neural responses are interpretable—at some cost—even without knowing the identity of the fly that generates these responses. But if we don't know the identity of the fly, all we can do is to associate each neural response with a distribution of stimuli that corresponds to an average over the individuals, and this distribution necessarily is broader than any of the individual distributions. As a result, we have less information about the visual stimulus. Once again, quantifying this information loss provides the only model independent measure of the departures from universality. Conversely the degree of universality in the code is measured by the fraction of visual information that can be captured using a “universal codebook” adapted to the ensemble of flies rather than to the individuals.⁶

The greater the differences among the neural responses of different flies, the more visual information we will lose if we don't know the identity of the individual: information gained about identity is information lost about the stimulus if we use a universal codebook. To formalize this relation between identifiability and decoding, note that if we observe the response of a neuron but don't know the identity of the individual, then we are observing responses drawn from the ensemble distributions defined above, $P^{\text{ens}}(W|t)$ and $P^{\text{ens}}(W)$. Under these conditions, the information that words provide about the stimulus is

$$I_T^{\text{ens}}(W; s(t)) = \left\langle \sum_W P^{\text{ens}}(W|t) \log_2 \left[\frac{P^{\text{ens}}(W|t)}{P^{\text{ens}}(W)} \right] \right\rangle_t \text{ bits.} \quad (6.10)$$

On the other hand, if we know the identity of the fly we gain the information $I_T^i(W; s(t))$ from Eq. (6.7). The average information loss is then

$$I_{\text{loss}}^{\text{avg}}(W; s(t)) = \sum_{i=1}^N P_i I_T^i(W; s(t)) - I_T^{\text{ens}}(W; s(t)). \quad (6.11)$$

⁶Even though flies differ in the structures of their neural responses, distinguishable responses could be functionally equivalent, as with distinct amino acid sequences that fold to the same protein structure. Thus it might be that all flies could be endowed (genetically?) with a universal or consensus codebook that allows each individual to make sense of her own spike trains, despite the differences from her conspecifics. Thus we would like to ask how much information we lose if the identity of the flies is hidden from us, or equivalently how much each fly can gain by knowing her own individual code.

This average information loss can be rewritten exactly in terms of the information about identity:

$$I_{\text{loss}}^{\text{avg}}(T) = I_T(\{W, t\}; \text{identity}) - I_T(W; \text{identity}). \quad (6.12)$$

As a practical matter, Eq. (6.12) means that the answer to our question about the efficacy of a universal codebook is contained in the results of Fig. 2. The result is that, on average, not knowing the identity of the fly limits us to extracting only 64 bits/s of information about the visual stimulus. This should be compared with the average information rate of 92.3 bits/s in our ensemble of flies: knowing her own identity allows the average fly to extract 44% more information from H1.

6.6 The nature of the 'personal' bits

Thus far we have analyzed the differences among the neural codes of different flies, and how much extra information a fly can extract by knowing its individual codebook. It is natural to ask what is being "said" by these extra bits, characterizing more explicitly the mapping from neural responses back to stimulus space for the different flies.

For each neural response W we can look back through the entire experiment and accumulate the motion trajectories that lead up to the response, and these provide samples from the distribution of stimuli conditional on the response as described above. Because the space of trajectories has many dimensions, this distribution is difficult to visualize, and so we focus here on the means of these distributions. This is a generalization of the reverse correlation or spike triggered average method [143]: rather than looking at the average stimulus that leads to a single spike, we look at the average stimulus that leads to the responses W (which consists of a pattern of spikes and empty intervals [38]).

In Fig. 6.4 we show the average waveforms of the stimulus velocity preceding a specific binary word in the spike trains of flies 1 and 6. Since fly 6 spike trains convey almost 3 times more information about the stimulus, one might have speculated that the same word was used in completely different stimulus contexts for the two flies. In fact, the differences are in the details and not in the general picture: spikes stand for pulses of positive velocity (as in Fig. 6.4b), long silent intervals stand for negative velocities (as in Figs. 6.4a&c), and the largest differences among the flies are in the widths, latencies and amplitudes of the pulses; combinations of spikes and intervals then lead to very different trajectories (as in Fig. 6.4d). For the fly which conveys less information, spikes are associated with larger positive velocities (Fig. 6.4b) and silences are associated with (slightly) larger negative velocities (Fig. 6.4a). Thus, these elementary responses come closer to exhausting the dynamic range of the inputs.

Conversely, the more informative spike train covers the dynamic range of inputs with a greater variety of composite responses.

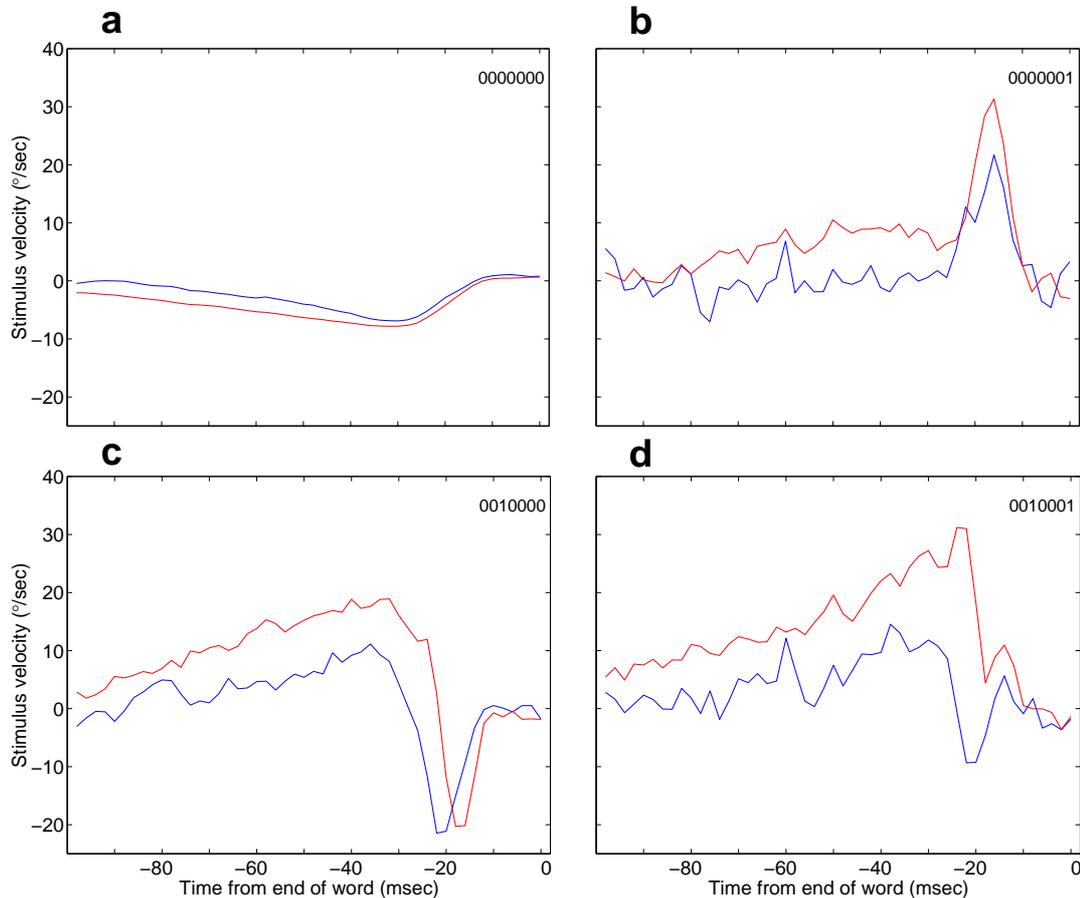


Figure 6.4: **What different flies mean by same words.** The word-triggered averages are shown for flies 1 (red) and 6 (blue) for 4 different words. Similar to the computation of spike triggered averages, we compute the average velocity profile of the movie presented to the flies, preceding 7-letter binary words. **(a)** The average stimulus waveform preceding the word '0000000' for Flies 1 (red) and 6 (blue), is shown as a function of the time relative to the end of the word (shown in actual time order on the right top side of the panel). **(b-d)** Word triggered averages for 3 other words, reflecting that the waveforms have similar rough structure, and that the difference between the flies is in the details.

6.7 Discussion

In the present work we have tried to quantify the individuality of the neural code used by a single neuron in the fly visual system. On the one hand, this individuality is sufficient to allow discrimination among individuals on time scales of relevance to behavior. Correspondingly, a significant amount ($\sim 30\%$) of the visual information carried by this neuron is accessible

only to observers who know the identity of the individual fly. We emphasize that these observations characterize the individuality of the fly's visual system up to the level of the H1 neuron, and thus provide a quantitative answer to our original question about individuality in the neural representation of sense data. A separate and obviously interesting question concerns the individuality of circuitry at subsequent stages, which might (or might not) make it possible for flies to use the individualistic components of their sensory representations in guiding motor behavior.

Although individual differences are apparent (by inspection of Fig. 6.2) and quantifiable (Fig. 6.2), we have also found significant elements of universality in the code. In the structure of the representations, the codebooks of different flies seem to differ only in matters of detail, as indicated in Fig. 6.4. More profoundly, although different flies extract very different amounts of information from the same visual inputs, all the flies achieve a high and constant efficiency in their encoding of this information (Fig. 6.3). From previous work it is known that the visual system of an individual fly exhibits substantial changes in coding strategy as it adapts to different ensembles of inputs, and at least in one case this adaptation serves to optimize information transmission [24]. Rather than converging on the same information rates in different flies, these adaptation processes seem to converge on codes of uniformly high efficiency, supporting the idea that efficiency of representation is a 'design principle' for the system [15].

On average the flies in our ensemble have neural codes in which a substantial amount of information is carried by patterns of spikes. This antiredundancy or synergy among spikes [25] is reduced substantially if we are forced to use a universal codebook. Mathematically this loss of synergy in the universal codebook is related to the fact that the rate at which we gain information about the identity of the fly (Fig. 6.2b) increases with window size to $T_c \sim 10$ ms. Discrimination among flies is enhanced by being able to see patterns of spikes in windows of size T_c , implying that the way these patterns are used to encode visual information is unique to each individual. Each individual fly thus gains nearly 50% more information through the use of a code in which patterns of spikes carry extra information, and more than half of this is lost if the fly does not have knowledge of its own identity. Not only is spike timing important for the neural code, but the way in which timing is used is specific to each individual.

One obvious difference between invertebrate and vertebrate nervous systems is the existence of identified neurons in invertebrates. The identifiability of invertebrate neurons sometimes has been interpreted to mean that these smaller nervous systems are hard wired automata; indeed the optomotor system of flies has been held up as a clear example of this extreme view. In this view, individuality plays no role, and it should even be possible to

average the results of experiments on corresponding neurons in different individuals. For vertebrates, substantial individuality arises through development and learning, and there are few if any identified neurons; at best vertebrates have identifiable modules consisting of hundreds or thousands of neurons, such as the columns in visual cortex. Against this clear dichotomy it is worth remembering that even genetically identical single celled organisms exhibit individuality in their sensory-motor behavior [166]. We hope that the techniques introduced here will help in quantifying the similarities or differences among neurons more generally, both for assessing the relatedness of function among cells in the same animal and for the comparison of corresponding cells in different individuals.

Chapter 7

Concluding remarks

We have studied three aspects of the nature of the neural code. One was the effect of ion channel noise on reliability and accuracy of spike timing. Another was the nature of the information encoding properties in such stochastic models and the effect of the biophysical parameters on the encoded information. Finally, we have compared the neural codes of an identified neuron in different animals, to quantify the individual and universal parts of their neural codes. We try to summarize what we have learned, and present the possible implications of our results as well as the future questions that arise.

Ion channel noise may determine the nature of the neural code : microscopic noise having a macroscopic effect

We have shown that ion channel noise may be the source of the stimulus-dependent reliability and timing jitter characteristics which were observed in real neurons, both *in vivo* and *in vitro* [110, 127, 13, 139, 39]. The microscopic ion channel noise can affect the macroscopic behavior of neurons, since the initiation of a spike is the result of opening of a critical number of ion channels. Because this number is relatively small, fluctuations in the number of open channels may have a significant effect on the membrane voltage, and thus on the timing and occurrence of a spike. Adding of ion channel noise to the deterministic neuron models also results in the (qualitative) replication of several other features of the spiking behavior of real neurons, namely, subthreshold oscillations of the membrane voltage, spontaneous and missing spiking. Such stochastic neuron models also better replicate the firing statistics and information encoding properties of real neurons.

These results reflect the significant connection between the detailed molecular level of the neuron and the macroscopic level of information processing and encoding. It seems that the noise inherent to the activity of ion channels must be considered in the modelling of neurons at the cellular level and also when studying neuronal networks. While this adds a significant

level of complexity to the analytical study of neuronal function, and to the simulation of population of neurons (both in terms of design and computation time), neglecting the neuronal noise may compromise the accuracy of the results and their applicability to real neurons. It remains to be seen how detailed the models of single channel activity should be in order to capture the main effects of the ion channel noise on the spiking patterns (see [154]). Our results and recent experimental results regarding the voltage-dependent nature of neuronal noise [195], and different functional 'states' of the membrane voltage [12, 10], suggest that the detailed models of the noise may (or at least a voltage dependent additive noise, which may be derived from the average fluctuations for linearized version of the full model, [114]). Adding other channel types, and using channel kinetic models that better span the history dependent nature of ion channel activity, may teach us that the microscopic level may also direct the nature of functional memory and state dependent neuronal function [177, 106].

More generally, the study of the effects of fluctuations in a small number of molecules as a source of (significant) macroscopic noise, is a reoccurring theme in biology and biophysics. It has been discussed as an important component of the 'design' and reliability of biochemical switches [136, 19], synaptic modification [85, 19], genetic regulation networks and replication [133, 60] and individuality of single organisms [166, 49]. It would be interesting to compare the noise characteristics in these systems, as well as common features in terms of overcoming the noise by using a population of noisy elements and even relying on the existence of noise to improve the system performance.

Biophysical design principles of neuronal spiking

The biophysical design of the neuron must accommodate a computational 'task' with various biophysical and computational constraints, such as the physical extent of the neuron, processing time, metabolism, biophysical noise and input noise. Our study of the effect of changing the biophysical parameters on the information encoding properties of the stochastic spiking models suggests several insights to the design principles of the spiking mechanism and neuronal function.

First, even if neurons were able to reduce their ion channel noise by using many more ion channels with lower conductance, the improvement in information encoding properties will be moderate (100 times more channels would result in less than a factor of 2 in the information encoding rate for most stimuli, and much less in others). Thus, it may be that significant reduction in the ion channel induced noise is not 'cost-effective', and it is more efficient to use several noisy neurons, than building and maintaining highly reliable ones. For some stimuli, we find that noise would actually improve the information encoding. Thus, noise and information do not always play opposing roles in terms of neuronal function.

The stochastic nature of neural spiking and the input-dependent effects on information encoding imply that decoding of the ambiguous spike patterns must rely on populations of neurons (it would be interesting then to compare the noise levels of identified neurons to cortical ones) (see e.g., [202]). The nature of the noise also means that inherently, correlated patterns of neuronal activity are more likely to propagate in a network of neurons, whereas uncorrelated ones would be more likely to 'die-out' [5, 123, 44]. One obvious extension of our work is the modelling the activity patterns of networks of neurons with ion channel based noise. Other interesting questions are the relation between ion channel noise and subthreshold oscillations [80], the effect of ion channel noise as stochastic element in neuronal learning (like a finite temperature in learning), and the nature of ion channel noise in adapting neurons.

Second, we have found that information encoding properties of stochastic spiking models are robust to many of their biophysical parameters, but are relatively sensitive to changes in the ratio of densities of their different ion channel types. For most of the stimuli we checked, it turns out that the optimal ion channel combination for information encoding is within "physiological range". A possible extension of our work is to try to formulate the relation between the information encoding properties of the spiking mechanism and its metabolic costs in a rate-distortion like formulation [35, 14, 95, 176]. More generally, it is natural to ask whether the described robustness and optimality would persist when we add more ion channel types and adaptation to the model, and how might history-dependent ion channel expression serve as a learning mechanism, to optimize the neuronal information encoding properties [43, 167].

Another possible 'design principle' was found in the study of the neural codes of different flies. Even though the flies may encode very different amounts of information about the same stimulus, and use different neural codes with different noise, the average information encoded per spike is universal. An obvious question is then whether such universal coding efficiency exists in other systems, and how robust is this universality to different stimulus features.

The stochastic and non-stationary nature of the neural codes

The goal of understanding the neural code is often explained in terms of constructing a dictionary from the neuronal spike trains to the stimulus, and the opposite dictionary from the stimuli to the responses (a similar set of dictionaries would exist for motor neurons and muscles). These dictionaries must be stochastic since the response of a neuron to the repeated presentation of the same stimuli is not deterministic. Conversely, different inputs may result in the same pattern of spiking from the same neuron. Moreover, these

dictionaries must be context dependent, since neurons adapt their neural responses, either through activity dependent changes in their ion channel states, expression of ion channels, synaptic connections etc.

Suggesting a biophysical explanation to the stochastic nature of the code, which lies in the machinery of the neuron, we further argue that the code itself is input dependent. We submit that the rate and temporal components of the spike patterns depend on the nature of the stimulus (expanding the notions of input-dependent reliability and accuracy of spike timing). We hope for a direct experimental study of this prediction.

We show that not only is the code different from one stimulus to another, but rather that the neural codes of different individuals may be significantly different. Presenting a method for comparing neural codes and applying it to the H1 neuron in different flies, we find significant differences between individuals, which reflect individual and common parts of encoded information. Again, an especially interesting common feature of the codes is that the information encoded per spike is universal across the fly population.

Provocatively, we could say that the fact that all neurons use spikes, may be similar to acknowledging that human communicate using phonemes. It doesn't capture the richness of human languages, its dynamic nature and the effect on how humans may communicate. Alternatively, we easily ignore differences in pronunciation. This leads to two central questions as a key issue in the study of neural codes. First, what are the common features of the neural codes that different neurons of different types, different systems, and different animals use and how do these common features change when cells adapt their responses. The second of course is how these differences affect the nature of neuronal computation and function. The answers are likely to lie between two extreme options: One is that since a 'target' neuron must decode the different, adapting codes of population of 'input' neurons, it extracts only the common, stationary aspects of the code, thus ignoring the variability and individuality of the neural codes. The other is that neurons are constantly adapting to better decode the variety of different adapting inputs they get from other neurons.

Finally, I hope we have contributed to the understanding of the biophysical design and properties of neural codes. As the common criticism of the detailed study of the neural code goes, it may still be that none of the details of the neural code(s) matter at the behavioral level, which is an open and debatable experimental question. If nothing else, it would be a bit boring.

Bibliography

- [1] L.F. Abbott. Decoding neuronal firing and modeling neural networks. *Quarter. Rev. Biophys.*, 27:291–331, 1994.
- [2] L.F. Abbott, G. Turrigiano, G. LeMasson G., and E. Marder. Activity-dependent conductances in model and biological neurons. In D.L. Waltz, editor, *Natural and Artificial Parallel Computing*, pages 43–68. SIAM, 1996.
- [3] L.F. Abbott, J.A. Varela, K. Sen, and S.B. Nelson. Synaptic depression and cortical gain control. *Science*, 275:220–224, 1997.
- [4] M. Abeles. *Local Cortical Circuits: An Electrophysiological Study*. Springer-Verlag, 1982.
- [5] M. Abeles. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge University Press, 1991.
- [6] M. Abeles and Y. Lass. Transmission of information by the axon: II. The channels capacity. *Biol. Cybern.*, 19:121–125, 1975.
- [7] J.Z. Aczel. *On Measures of Information and Their Characterizations*. Academic Press, 1975.
- [8] E.D. Adrian. *The Basis of sensation: the action of the sense organs*. W.W. Norton: New york, 1928.
- [9] R.H. Adrian. Conduction velocity and gating current in the squid giant axon. *Proc. R. Soc. Lond. B.*, 189:81–86, 1975.
- [10] J. Anderson, I. Lampl, I. Reichova, M. Carandini, and D. Ferster. Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. *Nat. Neurosci.*, 3:617–621, 2000.

- [11] J.J. Atick. Could information theory provide an ecological theory of sensory processing? *Network: Computation in Neural Systems*, 3:213–251, 1992.
- [12] R. Azouz and C.M. Gray. Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J. Neurosci.*, 19:2209–2223, 1999.
- [13] W. Bair and C. Koch. Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comp.*, 8:1185–1202, 1996.
- [14] V. Balasubramanian, M.J. Berry 2nd, and D. Kimber. Metabolically efficient information processing. *HUTP-98/A081*, 1998.
- [15] H. B. Barlow. Possible principles underlying the transformation of sensory messages. In W. A. Rosenblith, editor, *Sensory Communication*. MIT Press, 1961.
- [16] A.J. Bell, Z.F. Mainen, M. Tsodyks, and T.J. Sejnowski. ‘Balancing’ of conductances may explain irregular cortical spiking. Technical report, # INC-9502, Feb. 1995, Institute for Neural Computation, UCSD, San Diego, CA 92093-0523, 1995.
- [17] A.J. Bell and T.J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Comput.*, 7:1129–1159, 1995.
- [18] S. Bezrukov and I. Vodyanoy. Noise-induced enhancement of signal transduction across voltage-dependent ion channels. *Nature*, 378:362–364, 1995.
- [19] W. Bialek. Stability and noise in biochemical switches. *Advances in Neural Information Processing Systems (NIPS)*, 13:101–107, 2001.
- [20] W. Bialek, F. Rieke, R. de Ruyter van Steveninck, and D. Warland. Reading a neural code. *Science*, 252:1854–1857, 1991.
- [21] C. Bishop. *Neural Networks for Pattern Recognition*. Oxford university press, 1995.
- [22] A. Borst and F.E. Theunissen. Information theory and neural coding. *Nat. Neurosci.*, 2:947–57, 1999.
- [23] H. Braun, M. Huber, M. Dewald, K. Schafer, and K. Voigt. Computer simulations of neuronal signal transduction: the role of nonlinear dynamics and noise. *Int. J. Bifurc. Chaos*, in press, 1997.
- [24] N. Brenner, W. Bialek, and R. de Ruyter van Steveninck. Adaptive rescaling maximizes information transmission. *Neuron*, 26:695–702, 2000.

- [25] N. Brenner, S.P. Strong, R. Koberle, W. Bialek, and R. de Ruyter van Steveninck. Synergy in a neural code. *Neural Comput.*, 12:1531–1552, 2000.
- [26] H.L. Bryant and Segundo J.P. Spike initiation by transmembrane current: a white-noise analysis. *J. Physiol.*, 260:279–314, 1976.
- [27] T. Bullock. *Structure and Function in the Nervous Systems of Invertebrates*. W. H. Freeman, San Francisco, 1965.
- [28] G.T. Buracas, Zador A.M., M.R. DeWeese, and T.D Albright. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, 20:959–969, 1998.
- [29] G. Chechik and N. Tishby. Temporal dependent plasticity: an information theoretic approach. *Advances in Neural Information Processing Systems (NIPS)*, 13:110–116, 2001.
- [30] C. Chow and J. White. Spontaneous action potentials due to channel fluctuations. *Biophys. J.*, 71:3013–3021, 1996.
- [31] J. Clay and L. DeFelice. Relationship between membrane excitability and single channel open-close kinetics. *Biophys. J.*, 42:151–157, 1983.
- [32] J.R. Clay. Excitability of the squid giant axon revisited. *J. Neurophysiol.*, 80:903–913, 1998.
- [33] J.J. Collins, C.C. Chow, and T.T. Imhoff. Stochastic resonance without tuning. *Nature*, 376:236–8, 1995.
- [34] J. Cooley, F. Dodge, and H. Cohen. Digital computer solutions for excitable membrane models. *J. Cell Comp. Physiol.*, 66:99–100, 1965.
- [35] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley Interscience, 1991.
- [36] H.J. Curtis and K.S. Cole. Membrane action potentials from the squid axon. *J. Cell. Comp. Physiol.*, 15:147–157, 1940.
- [37] P. Dayan. and L.F. Abbott. *Theoretical Neuroscience*. To appear, 2001.
- [38] R. de Ruyter van Steveninck and W. Bialek. Real-time performance of a movement sensitive neuron in the blowfly visual system: Coding and information transfer in short spike sequences. *Proc. R. Soc. London Ser. B*, 234:379–414, 1988.

- [39] R. de Ruyter van Steveninck, G. Lewen, S. Strong, R. Koberle, and W. Bialek. Reproducibility and variability in neural spike trains. *Science*, 275:1805–1808, 1997.
- [40] S.P. Dear, J.A. Simmons, and J. Fritz. A possible neuronal basis for representation of acoustic scenes in auditory cortex of the big brown bat. *Nature*, 364:620–623, 1993.
- [41] L. DeFelice and A. Isaac. Chaotic states in random world. *J. Stat. Phys.*, 70:339–352, 1992.
- [42] L.J. DeFelice. *Introduction to Membrane Noise*. Perseus books, 1981.
- [43] N.S. Desai, L.C. Rutherford, and G.G. Turrigiano. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nat. Neurosci.*, 2:515–20, 1999.
- [44] M. Diesmann, M.O. Gewaltig, and A. Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402:529–533, 1999.
- [45] R. Eckhorn and B. Popel. Rigorous and extended application of information theory to the afferent visual system of the cat I: basic concepts. *Kybernetik*, 16:191–200, 1974.
- [46] R. Eckhorn and B. Popel. Rigorous and extended application of information theory to the afferent visual system of the cat II: Experimental results. *Biol. Cybern.*, 17:7–17, 1975.
- [47] R.S. Eisenberg, M. Frank, and C.F. Stevens. *Membranes, channels and noise*. Plenum Press, 1984.
- [48] B. Ermentrout. Type i membranes, phase resetting curves, and synchrony. *Neural Comput.*, 8:979–1001, 1996.
- [49] J.E. Ferrell and E.M. Machleder. The biochemical basis of an all-or-none cell fate switch in xenopus oocytes. *Science*, 280:895–898, 1998.
- [50] R. Fitzhugh. A kinetic model of the conductance changes in nerve membrane. *J. Cell. Comp. Physiol.*, 66:111–118, 1965.
- [51] I.A. Fleidervish, A. Friedman, and M.J. Gutnick. Slow inactivation of Na^+ current and slow cumulative spike adaptation in mouse and guinea-pig neocortical neurons in slices. *J. Physiol.*, 493:83–97, 1996.
- [52] I. Gat and N. Tishby. Synergy and redundancy among brain cells of behaving monkeys. *Advances in Neural Information Processing Systems (NIPS)*, 11:465–471, 1999.

- [53] T.J. Gawne and B.J. Richmond. How independent are the messages carried by adjacent inferior temporal cortical neurons. *J. Neurosci.*, 13:2758–2771, 1993.
- [54] A.P. Georgopoulos, A.B. Schwartz, and R.E. Kettner. Neuronal population coding of movement direction. *Science*, 243:1416–1419, 1986.
- [55] Y. Grossman, I. Parnas, and M.E. Spira. Differential conduction block in branches of a bifurcating axon. *J. Physiol.*, 295:283–305, 1979.
- [56] Y. Gutfreund, Y. Yarom, and I. Segev. Subthreshold oscillations and resonant frequency in guinea-pig cortical neurons: physiology and modeling. *J. Physiol.*, 483:621–640, 1995.
- [57] B.S. Gutkin and G.B. Ermentrout. Dynamics of membrane excitability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics. *Neural Comput.*, 10:1047–1065, 1998.
- [58] M. Gutman. Asymptotically optimal classification for multiple tests with empirically observed statistics. *IEEE Transactions on Information Theory*, 35(2):401–408, 1989.
- [59] R. Guttman, S. Lewis, and J. Rinzel. Control of repetitive firing in squid axon membrane as a model for a neuroneoscillator. *J. Physiol.*, 305:377–395, 1980.
- [60] J. Hastay, D. McMillen, F. Isaacs, and Collins J.J. Computational studies of gene regulatory networks: in numero molecular biology. *Nat. Rev. Genet.*, 2:268–279, 2001.
- [61] K. Hausen. The lobular complex of the fly: structure, function, and significance in behavior. In M. Ali, editor, *Photoreception and vision in invertebrates*, pages 523–559. Plenum, New York, 1984.
- [62] K. Hausen and C. Wehrhahn. Microsurgical lesion of horizontal cells changes optomotor yaw responses in the blowfly calliphora erythrocephala. *Proc. R. Soc. Lond. B*, 21:211–216, 1983.
- [63] J. Heller, J.A. Hertz, T.W. Kjaer, and B.J. Richmond. Information flow and temporal coding in primate pattern vision. *J. Comput. Neurosci.*, 2:175–193, 1995.
- [64] B. Hille. *Ionic Channels of Excitable Membrane*. Sinauer Associates, 2nd ed., 1992.
- [65] N. Ho and A. Destexhe. Synaptic background activity enhances the responsiveness of neocortical pyramidal neurons. *J. Neurophysiol.*, 84:1488–1496, 2000.

- [66] A.L. Hodgkin. The optimum density of sodium channels in an unmyelinated nerve. *Phil. Trans. R. Soc. Lond. B.*, 270:297–300, 1975.
- [67] A.L. Hodgkin and A.F. Huxley. Action potentials recorded from inside a nerve fibre. *Nature*, 144:710–711, 1939.
- [68] A.L. Hodgkin and A.F. Huxley. The components of the membrane conductance in the giant axon of the *loligo*. *J. Physiol.*, 116:473–496, 1952.
- [69] A.L. Hodgkin and A.F. Huxley. Currents carried by the sodium and potassium ion through the membrane of the giant axon of *loligo*. *J. Physiol.*, 116:449–472, 1952.
- [70] A.L. Hodgkin and A.F. Huxley. The dual effect of membrane potential on sodium conductance in the giant axon of *loligo*. *J. Physiol.*, 116:497–506, 1952.
- [71] A.L. Hodgkin and A.F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.*, 117:500–544, 1952.
- [72] A.L. Hodgkin and A.F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.*, 117:500–544, 1952.
- [73] A.L. Hodgkin and B. Katz. The effect of sodium ions on the electrical activity of the giant axon of the squid. *J. Physiol.*, 108:37–77, 1949.
- [74] J. Hopfield. Pattern recognition computation using action potential timing for stimulus representation. *Nature*, 376:33–36, 1995.
- [75] H. Horikawa. Noise effects on spike propagation in the stochastic Hodgkin Huxley models. *Biol. Cybern.*, 66:19–30, 1991.
- [76] H. Horikawa. Simulation study on effects of channel noise on differential conduction at an axon branch. *Biophys. J.*, 65:680–686, 1993.
- [77] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.*, 160:106–154, 1962.
- [78] J.D. Hunter, Milton J.G., P.J. Thomas, and J.D. Cowan. Resonance effect for neural spike time reliability. *J. Neurophysiol.*, 80:1427–1438, 1998.
- [79] B. Hutcheon, R. Miura, Y. Yarom, and E. Puyil. Low threshold calcium current and resonance in thalamic neurons: a model of frequency preference. *J. Neurophysiol.*, 71:583–594, 1994.

- [80] B. Hutcheon and Y. Yarom. Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci.*, 23:216–222, 2000.
- [81] J.J.B. Jack, D. Noble, and R. Tsien. *Electric current flow in excitable cells*. Oxford, 1975.
- [82] M.B. Jackson. Inversion of Markov processes to deterministic rate constants from single-channel data. *Biophys. J.*, 73:1382–1394, 1997.
- [83] R.V. Jensen. Synchronization of randomly driven nonlinear oscillators. *Phys. Rev. E.*, 58:6907–6910, 1998.
- [84] R.V. Jensen and D.H. Gartner. Synchronization of randomly driven nonlinear oscillators and the reliable firing of cortical neurons. In *Computational Neuroscience Meeting*, 1997.
- [85] M.B. Kennedy. The biochemistry of synaptic regulation in the central nervous system. *Annu. Rev. Biochem.*, 63:571–600, 1994.
- [86] S.S. Kety. *in* metabolism of the nervous system. pages 221–37. Pergamon, 1957.
- [87] R. Klink and A. Alonso. Ionic mechanisms for the subthreshold oscillations and differential electroresponsiveness of medial enthorinal cortex layer II neurons. *J. Neurophysiol.*, 70:144–157, 1993.
- [88] C. Koch. *Biophysics of Computation: Information Processing in Single Neurons*. Oxford University Press, 1998.
- [89] P. Konig, A.K. Engel, and W. Singer. Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends Neurosci.*, 19:130–137, 1996.
- [90] I. Lampl and Y. Yarom. Subthreshold oscillations and resonant behavior: two manifestations of the same mechanism. *Neuroscience*, 78:325–341, 1997.
- [91] I. Lampl and Y. Yarom. Subthreshold oscillations and resonant behavior: two manifestations of the same mechanism. *Neuroscience*, 78(2):325–341, 1997.
- [92] M.F. Land and T.S. Collett. Chasing behavior of houseflies (*fannia canicularis*). a description and analysis. *J. Comp. Physiol.*, 89:331–357, 1974.
- [93] E.S. Lander and M.S. Waterman. *Calculating the secrets of life: Applications of the mathematical sciences in molecular biology*. National academy press, Washington, 1995.

- [94] Y. Lass and M. Abeles. Transmission of information by the axon: I. Noise and memory in the myelinated nerve fiber of the frog. *Biol. Cybern.*, 19:61–67, 1975.
- [95] S.B. Laughlin, R.R. de Ruyter van Steveninck, and J.C. Anderson. The metabolic cost of neural information. *Nat. Neurosci.*, 1:36–41, 1998.
- [96] H. Lecar and R. Nossal. Theory of threshold fluctuations in nerves. I. Relationships between electrical noise and fluctuations in axon firing. *Biophys. J.*, 11:1048–1067, 1971.
- [97] H. Lecar and R. Nossal. Theory of threshold fluctuations in nerves. II. Analysis of various sources of membrane noise. *Biophys. J.*, 11:1068–1084, 1971.
- [98] V. Levenshtein. Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Doklady*, 10:707–710, 1966.
- [99] J.E. Levin and J.P. Miller. Broadband neural encoding in the cricket cercal sensory system enhanced by stochastic resonance. *Nature*, 380:165–8, 1996.
- [100] W.B. Levy and R.A. Baxter. Energy efficient neural codes. *Neural Comput.*, 8:531–43, 1996.
- [101] J. Lin. Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, 37(1):145–151, 1991.
- [102] R. Linsker. Self organization in a perceptual network. *Computer*, 21:105–117, 1988.
- [103] R. Linsker. Local synaptic learning rules suffice to maximize mutual in a linear network. *Neural Comput.*, 4:691–702, 1992.
- [104] M. London, A. Shraibman, N. Tishby, and I. Segev. Characterization of synaptic efficacy using information theoretic methods. *NEUROSCI. LETT., suppl.*, 54, 1999.
- [105] A. Longtin and K. Hinzer. Encoding with bursting, subthreshold oscillations, and noise in mammalian cold receptors. *Neural Comp.*, 8:215–255, 1996.
- [106] S.B. Lowen, L. Liebovitch, and J.A. White. Fractal ion-channel behavior generates fractal firing patterns in neuronal models. *Phys. Rev. E.*, 59:5970–5980, 1999.
- [107] D. MacKay and W.S. McCulloch. The limiting information capacity of a neuronal link. *Bull. Math. Biophys.*, 14:127–135, 1952.

- [108] K. MacLeod, A. Backer, and G. Laurent. Who reads temporal information contained across synchronized and oscillatory spike trains? *Nature*, 395:693–698, 1998.
- [109] J. Magee, D. Hoffman, C. Colbert, and Johnston D. Electrical and calcium signaling in dendrites of hippocampal pyramidal neurons. *Annu. Rev. Physiol.*, 60:327–46, 1998.
- [110] Z. Mainen and T. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1508, 1995.
- [111] Z.F. Mainen, J. Joerges, J.R. Huguenard, and T.J. Sejnowski. A model of spike initiation in neocortical pyramidal neurons. *Neuron*, 15:1427–1439, 1995.
- [112] Z.F. Mainen, J. Joerges, J.R. Huguenard, and T.J. Sejnowski. A model of spike initiation in neocortical pyramidal neurons. *Neuron*, 15:1427–1439, 1995.
- [113] Y. Manor, J. Gonczarowski, and I. Segev. Propagation of action potentials along complex axonal trees. model and implementation. *Biophys. J.*, 60:1411–1423, 1991.
- [114] A. Manwani and C. Koch. Detecting and estimating signals in noisy cable structures, I: Neuronal noise sources. *Neural Comput.*, 11:1797–1829, 1999.
- [115] A. Manwani and C. Koch. Detecting and estimating signals in noisy cable structures, II: Information theoretical analysis. *Neural Comput.*, 11:1831–1873, 1999.
- [116] A. Manwani and C. Koch. Detecting and estimating signals over noisy and unreliable synapses: information-theoretic analysis. *Neural Comput.*, 13:1–33, 2001.
- [117] E. Marder, L. Abbott, G. Turrigiano, Z. Liu, and J. Golowasch. Memory from the dynamics of intrinsic membrane currents. *Proc. Natl. Acad. Sci.*, 93:13481–13486, 1996.
- [118] H. Markram and M. Tsodyks. Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature*, 382:807–810, 1996.
- [119] H. Markram, Y. Wang, and M. Tsodyks. Differential signaling via the same axon of neocortical pyramidal neurons. *Proc. Natl. Acad. Sci. USA*, 95:4323–5328, 1998.
- [120] P.Z. Marmarelis and V.Z. Marmarelis. *Analysis of Physiological Systems: The White-Noise Approach*. Plenum Press:New York, 1978.
- [121] S. Marom. Slow changes in the availability of voltage-gated ion channels: effects on the dynamics of excitable membranes. *J. Mem. Biol.*, 1997. in press.

- [122] S. Marom, H. Salman, V. Lyakhov, and E. Braun. Effects of density and gating delayed-rectifier potassium channels on resting membrane potential and its fluctuations. *J. Membrane Biol.*, 154:267–274, 1996.
- [123] P. Marsalek, C. Koch, and J. Maunsell. On the relationship between synaptic input and spike output jitter in individual neurons. *Proc. Natl. Acad. Sci. USA*, 94:735–740, 1997.
- [124] B.W. Mel. Information-processing in dendritic trees. *Neural Comput.*, 6:1031–1085, 1994.
- [125] E. Neher and B. Sakmann. Single-channel currents recorded from membrane of denervated frog muscle fibres. *Nature*, 260:799–802, 1976.
- [126] J.G. Nicholls, A.R. Martin, and B.G. Wallace. *From neuron to brain*, 3rd ed. Sinauer, 1992.
- [127] L. Nowak, M. Sanches-Vives, and D. McCormick. Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cerebral Cortex*, 7:487–501, 1997.
- [128] L.M. Optican and Richmond B.J. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. *J. Neurophysiol.*, 57:162–178, 1987.
- [129] L.M. Optican, T.J. Gawne, and Richmond B.J. Transmitted information and channel capacity from multivariate neuronal data. *Biol. Cybern.*, 65:305–310, 1991.
- [130] N. Otmakhov, A.M. Shirke, and Malinow R. Measuring the impact of probabilistic transmission on neuronal output. *Neuron*, 10:1101–1111, 1993.
- [131] S. Panzeri, S.R. Schultz, A. Treves, and E.T. Rolls. Correlations and the encoding of information in the nervous system. *Proc. Royal Soc. Lond. B.*, 266:1001–1012, 1999.
- [132] J. Patlak. Molecular kinetics of voltage-dependent Na^+ channels. *Physiological Rev.*, 71(4):1047–1080, 1991.
- [133] J. Paulsson and M. Ehrenberg. Random signal fluctuations can reduce random fluctuations in regulated components of chemical regulatory networks. *Phys. Rev. Lett.*, 84:5447–5450, 2000.
- [134] J.R. Pierce. *Symbols, signals and noise: The nature and process of communication*. Harper, new york, 1961.

- [135] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1992.
- [136] M. Ptashne. *Genetic Switch: Phage Lambda and Higher Organisms*. Blackwell Science, 1992.
- [137] S. Fine R. El-Yaniv and N. Tishby. Agnostic classification of markovian sequences. *Advances in Neural Information Processing Systems*, 10:465–471, 1997.
- [138] M. Rapp, Y. Yarom, and I. Segev. Modeling back propagating action potential in weakly excitable dendrites of neocortical pyramidal cells. *Proc. Natl. Acad. Sci. USA*, 93:11985–11990, 1996.
- [139] D. Reich, J. Victor, B. Knight, T. Ozaki, and E. Kaplan. Response variability and timing precision of neuronal spike trains in vivo. *J. Neurophysiol.*, 77:2836:2841, 1997.
- [140] P. Reinagel and R.C. Reid. Temporal coding of visual information in the thalamus. *J. Neurosci.*, 20:5392–5400, 2000.
- [141] A. Riehle, S. Grun, M. Diesmann, and A. Aersen. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278:1950–1953, 1997.
- [142] F. Rieke, D. Warland, and W. Bialek. Coding efficiency and information rates in sensory neurons. *Europhys. Lett.*, 22:151–156, 1993.
- [143] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spike: Exploring the Neural Code*. MIT Press, 1997.
- [144] J. Rinzel and B. Ermentrout. Analysis of neural excitability and oscillations. In C. Koch and I. Segev, editors, *Methods in Neuronal Modeling*, pages 135–169. MIT Press, 1989.
- [145] J.T. Rubinstein. Threshold fluctuations in an N sodium channel model of the node of ranvier. *Biophys. J.*, 68:779–785, 1995.
- [146] B. Sakmann and E. Neher. *Single-Channel Recording*. Plenum Publishing, 2nd edition, 1995.
- [147] E. Salinas and L.F. Abbott. Vector reconstruction from firing rates. *J. Comput. Neurosci.*, 1:89–107, 1994.

- [148] T.D. Sanger. Probability density estimation for the interpretation of neural population codes. *J. Neurophysiol.*, 76:2790–2793, 1996.
- [149] M. Schetzen. *The Volterra and Wiener theories of nonlinear systems*. Wiley and Sons, 1980.
- [150] E. Schneidman, N. Brenner, N. Tishby, R. de Ruyter van Steveninck, and W. Bialek. Universality and individuality in a neural code. *Advances in Neural Information Processing Systems (NIPS)*, 13:159–165, 2001.
- [151] E. Schneidman, B. Freedman, and I. Segev. Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural Comp.*, 10:1679–1704, 1998.
- [152] E. Schneidman, B. Freedman, and I. Segev. *Spike timing reliability in a stochastic Hodgkin-Huxley model*, chapter in *Bower J.: Computational Neuroscice: Trends in research*, pages 261–266. Plenum, NY, 1998.
- [153] E. Schneidman, I. Segev, and N. Tishby. Information capacity and robustness of stochastic neuron models. *Advances in Neural Information Processing Systems (NIPS)*, 12:178–184, 2000.
- [154] I. Segev. Single neuron models – oversimple, complex and reduced. *Trends Neurosci.*, 15:414–421, 1992.
- [155] I. Segev and M. London. Untangling dendrites with quantitative models. *Science*, 290:744–750, 2000.
- [156] I. Segev, G.M. Shepherd, and J. Rinzel. *The Theoretical Foundation of Dendritic Function : Selected Papers of Wilfrid Rall with Commentaries*. MIT press, 1996.
- [157] H.S. Seung and H. Sompolinsky. Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci.*, 90:10749–10753, 1993.
- [158] M.N. Shadlen and W.T. Newsome. Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.*, 4:569–579, 1994.
- [159] M.N. Shadlen and W.T. Newsome. Is there a signal in the noise? *Curr. Opin. Neurobiol.*, 5:569–579, 1995.
- [160] C.E. Shannon and W. Weaver. *The Mathematical Theory of Ccommunication*. Univ. of Illinois Peess, 1949.

- [161] A. Sharp, M.B. O'Neil, L.F. Abbott, and E. Marder. The dynamic clamp: Artificial conductances in biological neurons. *Trends in Neurosci.*, 16:389–394, 1993.
- [162] G.M. Shepherd. *The Synaptic Organization of the Brain*. Oxford University Press, 4 edition, 1997.
- [163] E. Skaugen and L. Walløe. Firing behavior in a stochastic nerve membrane model based upon the Hodgkin-Huxley equations. *Acta Physiol. Scand.*, 107:343–363, 1979.
- [164] W. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.*, 13:334–350, 1993.
- [165] W.R. Softky. Simple codes versus efficient codes. *Curr. Opin. Neurobiol.*, 5:239–247, 1995.
- [166] J.L. Spudich and D.E. Koshland. Non-genetic individuality: chance in the single cell. *Nature*, 262:467–471, 1976.
- [167] M. Stemmler and C. Koch. How voltage-dependent conductances can adapt to maximize the information encoded by neuronal firing rate. *Nat. Neurosci.*, 2:521–7, 1999.
- [168] C.F. Stevens. Cooperativity of unreliable neurons. *Curr. Biol.*, 4:268–269, 1994.
- [169] C.F. Stevens and A.M. Zador. Input synchrony and the irregular firing of cortical neurons. *Nat. Neurosci.*, 1:210–217, 1998.
- [170] A. Strassberg and L. DeFelice. Limits of the HH formalism: Effects of single channel kinetics on transmembrane voltage dynamics. *Neural Comp.*, 5:843–856, 1993.
- [171] S. Strong, R. Koberle, R. de Ruyter van Steveninck, and W. Bialek W. Entropy and information in neural spike trains. *Phys. Rev. Lett.*, 80:197–200, 1998.
- [172] G. Stuart, N. Spruston, and M. Hausser. *Dendrites*. UOP, 1999.
- [173] Y. Sugase, S. Yamane, S. Ueno, and K. Kawano. Global and fine information coded by single neurons in the temporal visual cortex. *Nature*, 400:869–873, 1999.
- [174] A. Tang, A.M. Bartels, and T.J. Sejnowski. Cholinergic modulation preserves spike timing under physiologically realistic fluctuating input. *Advances in Neural Information Processing System*, 9:866–872, 1997.
- [175] F. Theunissen and J. Miller. Temporal encoding in nervous systems: A rigorous definition. *J. Comp. Neurosci.*, 2:149–162, 1995.

- [176] N. Tishby, F. Pereira, and W. Bialek. The information bottleneck method. *The 37th annual Allerton Conference on Communication, Control, and Computing*, 1999.
- [177] A. Toib, V. Lyakhov, and S. Marom. Interaction between duration of activity and rate of recovery from slow inactivation in mammalian brain Na^+ channels. *J Neurosci.*, 18:1893–1903, 1998.
- [178] S.F. Traynelis and F. Jaramillo. Getting the most out of noise in the central nervous system. *Trends Neurosci.*, 21:137–145, 1998.
- [179] A. Treves and S. Panzeri. The upward bias in measures of information derived from limited data samples. *Neural Comput.*, 7:399–407, 1995.
- [180] T.W. Troyer and K.D. Miller. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Computation*, 9:733–745, 1997.
- [181] M. Tsodyks and H. Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA*, 94:719–723, 1997.
- [182] H.C. Tuckwell. *Introduction to Theoretical Neurobiology*. Cambridge University Press, 1988.
- [183] G.G. Turrigiano, K.R. Leslie, N.S. Desai, L.C. Rutherford, and S.B. Nelson. Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391:892–6, 1998.
- [184] E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Y. Prut, H. Slovin, and A. Aertsen. Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373:515–518, 1995.
- [185] C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274:1724–1726, 1996.
- [186] C.A. Vandenberg and F. Bezanilla. A sodium channel gating model based on single channel, macroscopic ionic, and gating currents in the squid axon. *Biophys. J.*, 60:1511–1533, 1991.
- [187] J.D. Victor and K. Purpura. Nature and precision of temporal coding in visual cortex: a metric-space analysis. *J. Neurophysiol.*, 76:1310–1326, 1996.

- [188] J.D. Victor and K. Purpura. Metric–space analysis of spike trains: theory, algorithms, and application. *Network*, 8:127–164, 1997.
- [189] M. Volgushev, M. Chistiakova, and W. Singer. Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience (in press)*, 1997.
- [190] J. Von Neumann. *The computer and the brain*. Yale University Press, 1958.
- [191] X.J. Wang and G. Buzsaki. Gamma oscillations by synaptic inhibition in a hippocampal interneuronal network model. *J. Neuroscience*, 16:6402–6413, 1996.
- [192] D.K. Warland, P. Reinagel, and M. Meister. Decoding visual information from a population of retinal ganglion. *J. Neurophysiol.*, 78:2336–2350, 1997.
- [193] G. Werner and V.B. Mountcastle. Neural activity in mechanoreceptive cutaneous afferents: Stimulus-response relations, Weber functions, and information transmission. *J. Neurophysiol.*, 28:359–397, 1965.
- [194] J. White, T. Budde, and A. Kay. A bifurcation analysis of neuronal subthreshold oscillations. *Biophys. J.*, 69:1203–1217, 1995.
- [195] J. White, R. Klink, A. Alonso, and A. Kay. Noise from voltage-gated channels may influence neuronal dynamics in the entorhinal cortex. *J Neurophysiol*, 80:262–9, 1998.
- [196] J.A. White, J.T. Rubinstein, and A.R. Kay. Channel noise in neurons. *Trends in Neurosci.*, 23:131–137, 2000.
- [197] F.M.J. Willems, W.M. Shtarkov, and T.J. Tjalkens. The context-tree weighting method: Basic properties. *IEEE Trans. Info. Theory*, 41:653–664, 1993.
- [198] M.A. Wilson and B.L. McNaughton. Dynamics of the hippocampal code for space. *Science*, 261:1055–1058, 1993.
- [199] W. Yamada, C. Koch, and P. Adams. Multiple channels and calcium dynamics. In C. Koch and I. Segev, editors, *Methods in Neuronal Modeling*, pages 97–134. Press, 1989.
- [200] A. Zador. Impact of synaptic unreliability on the information transmitted by spiking neurons. *J. Neurophys.*, 79:1219–1229, 1998.
- [201] R.S. Zemel, P. Dayan, and A. Pouget. Probabilistic interpretation of population codes. *Neural Comput.*, 10:403–430, 1998.

- [202] E. Zohary, Shadlen M.N., and Newsome W.T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370:140–143, 1994.