Estimation and Evaluation of Energy-Efficient Neural Communication Channels

A Dissertation

Presented to

the Faculty of the School of Engineering and Applied Science

University of Virginia

In Partial Fulfillment

of the requirements for the Degree

Doctor of Philosophy (Electrical Engineering)

by

Mustafa Sungkar

May 2018

Abstract

The brain is an energy-efficient computation device. At rest, it runs on 15 Watts of power. What can scientists and engineers learn from the brain to make computational devices more energy-efficient? This dissertation begins to address that question by studying the behavior of cortical neurons in the sensory cortex.

The task of the cortical neuron is to send information about its input to other target neurons. It performs this task by expending as little energy as possible. To quantify the performance of the neuron, Shannon's mutual information (MI) is used as a measure of neural information. The neuron is assumed to maximize MI for a fixed energy budget. Thus, an information theoretic framework can be used to analyze the energy efficiency of the neuron.

This dissertation consists of four major parts: the generalized inverse Gaussian (GIG) neuron model, assessing the energy efficiency of the model, optimizing the model, and a rate-distortion (R-D) problem inspired by the model. The GIG neuron model takes into account the fast sodium channels that allow a fast rate of increase of the postsynaptic potential (PSP). This behavior of the PSP determines the input-output behavior of the neuron, which allows the neuron to be modeled as a communication channel. Next, methods for estimating the parameters of the GIG neuron model are developed. The accuracy of the model is evaluated with simulations.

After that, the maximum MI transmitted by the GIG neuron model for a fixed energy budget is determined. Surprisingly, the input distribution that achieves the constrained capacity is discrete with a finite number of mass points for some parameter sets. This implies that the neural network (NN) should exist in discrete states to maximize the MI transmitted by the neuron. To further optimize the GIG neuron model, the parameter sets that produces the most MI for a given energy budget is discussed. An additional variance constraint is imposed to prevent MI from increasing without bound. A numerical example is used to illustrate the theory.

Finally, a R-D problem inpsired by the GIG neuron model is developed. The source distribution is given by the GIG distribution and the distortion function is related to the energy expenditure of the GIG neuron model. The result is that for some parameter sets, the reconstruction alphabet is discrete.

This dissertation is part of a nascent approach to the study of energy-efficient computing. The next steps must involve studying the network in which the neuron operates. This includes studying feedback loops that exist within the network.

APPROVAL SHEET

This Dissertation
is submitted in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Author Signature: Mustafa Sungkar

This Dissertation has been read and approved by the examining committee:

Advisor: Toby Berger

Committee Member: Stephen G. Wilson

Committee Member: William B Levy

Committee Member: Daniel Weller

Committee Member: Christian Gromoll

Committee Member:

Accepted for the School of Engineering and Applied Science:

Craig H. Benson, School of Engineering and Applied Science

May 2018

To my parents, who never gave up on me

Acknowledgements

First, I would like to express my deepest gratitude to Professor Toby Berger, whose guidance and wisdom were invaluable to my growth and development, personally and professionally. Through his parables, he has given me amusement, advice, inspiration, and encouragement. I would like to thank our esteemed colleague, Professor William B Levy, for all the thoughtful discussions, insights, and contributions he made to our endeavors. I would also like to thank Professor Stephen G. Wilson for his guidance and discussions. He is an inspiration to many of us in the ECE department. I am grateful for the members of the Ph.D. committee, Professor Daniel Weller and Professor Christian Gromoll, who have set aside time in their busy schedules to support and advise me in my endeavors. My appreciation also goes to the members of the ECE office for making all the administrative work easy and seamless.

I am grateful for my colleagues and friends for their encouragement and support. I could not ask for better friends. A special mention goes to the members of Panda House, who became my family away from home.

Last but certainly not least, I am indebted to my parents, whose support and encouragement are endless. They have stuck by me through the good and difficult times and I will never be able to repay all that they have done for me, even for a thousand lifetimes.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

One of the biggest mysteries to both engineers and scientists is the inner workings of the brain. The human brain consists of 100 billion individual units called neurons, which collectively form a neural network (NN). Figure 1.1 shows an image of stained neurons. The network can perform a wide range of tasks, e.g., prediction, estimation, discrimination, decision-making, along with keeping the organism alive and well. A single neuron may make connections to as many as 10 thousand neurons. This high level of connectivity creates an overwhelmingly complex system that requires expertise from numerous fields to analyze.

Astoundingly, NN's are highly energy efficient. The human brain functions on approximately 15 watts of power when it is at rest [1]. For comparison, a modern desktop uses power on the order of 100 watts. To further illustrate the efficiency of the brain, consider the following: when the IBM Almaden Research Center simulated the visual cortex of a cat, which contains 10 million neurons, the simulation uses power at a rate of 10 billion times that of the real thing [2]. This demonstrates that there is much to learn from the energy efficiency of the brain.

Engineers and scientists have drawn inspirations from NN's and developed artificial neural

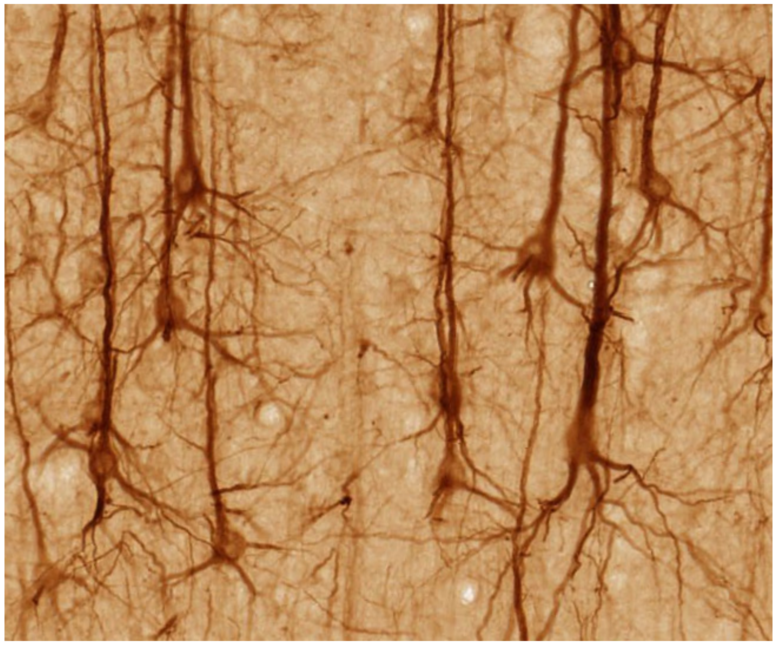

Figure 1.1: An image of stained pyramidal neurons. Image by UC Regents Davis campus is licensed under CC-BY-SA-3.0.

networks (ANN's). ANN's are heavily present in big data and artificial intelligence, two topics that dominate the field of computer science. As the push for "machine intelligence" grows, a better understanding of NN's is useful in creating not just more powerful algorithms, but energy efficient ones as well. As the size of computational devices shrink, heat density becomes a bigger concern. Hence, energy efficiency is needed so that computation generates less heat.

This dissertation addresses the energy efficiency of NN's by focusing on the energy-efficient design of a single neuron that is operating in the network. The purpose of the neuron is to transmit information efficiently. Thus, information theory is used as a tool of analysis for the neuron. Even though the single neuron is the focus, certain conclusions on NN's can be made based on the results presented in this dissertation.

1.2 Background

There are numerous types of neuron, but this dissertation focuses on pyramidal neurons in the sensory cortex. Let neuron η be a neuron of such type. Neuron η will also be referred to as just η. The reason for focusing on pyramidal neurons is because they receive input from many neurons, i.e., around 10 thousand. With a large number of input size, certain approximations need to be used to simplify the analysis.

1.2.1 Information Theory and Neural Networks

It is widely believed that neural networks process information. Sensory cells sense external stimuli and the information is passed up along the hierarchy of the NN. In each level, information is processed so that irrelevant ones are discarded and relevant details are extracted. On an individual level, it is believed that the purpose of η is to convey to its targets information about the state of its input. A quantifiable notion of information is therefore needed so that analysis on η can be performed.

A popular measure of neural information is Shannon's average mutual information, which is referred to as just mutual information (MI) in this dissertation [3]. Other notions of information exist, such as Hartley and Fisher information. However, MI is robust and can be derived from a set of intuitive axioms [4, chapter]. Mutual information has numerous interpretations, but the one adopted in this dissertation is that MI is the reduction in uncertainty of the input upon learning the output. Mutual information can be written as a difference of entropies: the entropy of the input minus the entropy of the input given the output. Entropy is a measure of uncertainty. Thus, MI is the reduction of entropy of the input upon learning the output. Entropy and MI are precisely defined in Chapter 3. By using MI, theories and developments from information theory can be used as a tool for assessing neural information.

The most widely used interpretation of MI is through the channel coding theorem. Each communication channel has a capacity given by the maximum MI that can be transmitted across the channel. The theory states that in order to send data through the channel with arbitrarily small probability of error, the transmission rate must be below the capacity of the channel. Otherwise, the error cannot be made as arbitrarily small as possible. To achieve low error, the input must be coded in a certain way. Since such coding schemes cannot exist in neurons, the channel coding theorem is not applicable. In essence, MI is not interpreted as a bound on information rate. Rather, it is simply a measure of how much information η has transmitted.

Despite the naturalness of MI, it has shortcomings. Mutual information does not identify useful information. For example, the NN may respond to a certain stimuli, e.g., Vivaldi's Four Season Suite. However, MI does not indicate from where in Vivaldi is this information coming. It just indicates that there is information extracted from Vivaldi. Regardless, MI is still a useful tool in understanding neurons and NN's as used by neuroscientists, biologists, physicists, and electrical engineers.

1.2.2 Energy-Efficient Neural Communication

Since η is energy efficient and its task is to transmit information, it is assumed that η seeks to minimize the average energy expended to send a certain value of MI. Looking at it from another perspective, η seeks to maximize the MI it transmits given an average energy budget. When η is thought of as a communication channel, neural information transmission can be related to the constrained capacity of the channel. This is what information theorists use to describe how well a power-limited channel can be made to perform.

Neuron η must also communicate using an energy-efficient scheme. This is done by communicating via all-or-none electrical pulses known as action potentials (AP's). The AP is generated when the neuron's postsynaptic potential (PSP) reaches a threshold. The time it

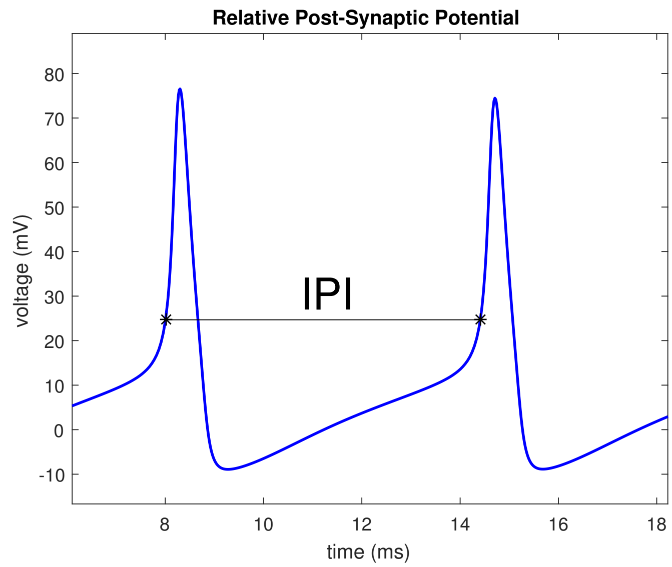

Figure 1.2: A diagram of two action potentials. The IPI is the time interval between two of the pulses.

takes the PSP to reach the threshold depends on the input intensity of the neuron. Hence, the AP's contain information about the input intensity of the neuron.

There are two major hypotheses as to how information is encoded by η: the rate coding and time coding hypotheses. In the rate coding hypothesis, information is encoded by the rate at which the AP's are produced. In the time coding hypothesis, information is encoded in the time interval between subsequent AP's. This time interval is called the interpulse interval (IPI) and is illustrated in Figure 1.2. Time coding is also considered to be more energy-efficient than rate coding. For a given energy cost, time coding provides more information than rate coding. Because η is energy efficient, it is assumed to use the time coding scheme. To communication engineers, the time coding scheme is also known as time-continuous differential pulse position modulation (TCDPPM). This coding scheme is further described in Chapter 2.

Furthermore, the time coding scheme allows for instantaneous decoding. Compare that to rate coding, where several pulses must be received before the message can be decoded. This is advantageous for when quick decisions are necessary, i.e., when danger is present to the organism. However, rate coding is less prone to error because it uses more pulses to encode a message. Regardless, it is argued in Chapter 2 that the noise in decoding the IPI duration is small. Thus, time coding is more advantageous overall for η than rate coding.

1.3 Literature Review

Information theory was first applied to neuroscience in [5], where entropy was used to analyze the information carried by neural spikes. It was not until many years later that information theory was used in experimental studies to decode the train of AP signals [6] and to measure the amount of information transmitted by a single AP [7, 8]. For cortical neurons with thousands of input lines, theoretical developments and simulations are more appealing. In [9], energy-efficient population codes for a neuron model were analyzed. In [10], information-energy tradeoff was analyzed for a neuron model with discrete inputs. The information-energy tradeoff was also analyzed for discrete poisson arrival neural models in [11, 12], diffusion models in [13, 14], and the Hodgkin-Huxley model in [15].

With regards to energy efficiency, energy minimization subject to functional constraints as a unifying principle for all neurons has been proposed [16]. Energy efficiency has also been experimentally observed in ion channels of neurons [16, 17] and action potentials [18].

1.4 Contributions

The contributions of this dissertation, along with the relevant publications, are the following:

- Connecting the generalized inverse Gaussian (GIG) neuron model to the neurobiological process of generating AP's, especially with regards to the fast sodium (Na^+) ion channels [19, 20]

- Interpreting the parameters of the GIG neuron model and estimating them from data [19, 20]

- Solving for the constrained capacity of the GIG neuron model and showing that the optimizing input distribution is discrete [21, 22]
- Formulating a MI maximization problem for the GIG neuron model over the model parameters and solving it for a special case of the GIG neuron model [23]
- Developing the conditions for the solution of the rate-distortion (R-D) problem for a general measure [24]
- Proposing a R-D problem inspired by the GIG neuron model and proving that the reconstruction alphabet is discrete for some cases [24]

1.5 Overview

In Chapter 2, the generalized inverse Gaussian neuron model is described. The role of the parameters on the behavior of the model is explored. Finally, estimation techniques of the parameter from a sample path of the model is described.

In Chapter 3, the information-energy tradeoff of the GIG neuron model is examined. In particular, the constrained capacity of the GIG neuron model is obtained. Also, it is shown that in certain cases, the input distribution that achieves the constrained capacity for the GIG neuron model is discrete. Implications of the results on the NN are discussed.

In Chapter 4, the MI maximization problem for the GIG neuron model over the model parameters is presented. Also, it is solved for a special case of the GIG neuron model.

In Chapter 5, the optimality conditions for R-D problems for a general distribution is derived. Then a R-D problem inspired by the GIG neuron model is proposed. It is proven that for certain cases, the reconstruction alphabet is discrete.

The conclusion is found in Chapter 6. Also, suggestions for future research are presented.

Chapter 2

The Generalized Inverse Gaussian Neuron Model

In order to analyze the energy efficiency of neuron η, i.e., a cortical neuron in the primary sensory cortex, the generalized inverse Gaussian (GIG) neuron model is used. It is a stochastic model based on the generalized inverse Gaussian (GIG) distribution. Associated with the GIG distribution is the generalized inverse Gaussian first hitting time (GIGHT) stochastic diffusion, which can model the postsynaptic potential (PSP) of η. The advantages of the GIG neuron model is that it gives a relatively simple mathematical description of η while also being able to capture the complexity of signaling of η.

Contributions presented in this chapter were originally published in [19] and [20]. Background on the PSP and the action potential (AP) are first discussed. Special attention is given to the fast sodium (Na^+) ion channels that give rise to the shape of the leading edge of the AP. Then the GIG neuron model is described. The GIG distribution and the GIGHT diffusion are then defined. In fact, the first hitting time (FHT) of the GIGHT diffusion is described by the GIG probability density function (PDF). The GIG PDF gives the input-output relationship for η. This leads to the development of the GIG neuron model. Then it is explained how the GIGHT diffusion can be used to model the behavior of the

PSP of η under normal circumstances. The parameters of the GIG neuron model are also described and given meanings.

Finally, techniques to estimate from data the GIG neuron model parameters are presented. The estimators are based on the Euler-Maruyama approximation of the GIGHT diffusion and based on the maximum-likelihood estimator (MLE). Simulation results of the performance of the estimators are presented. Then a possible experiment set up to test the accuracy of the GIG neuron model is discussed in the conclusion.

2.1 Background

2.1.1 Parts of Cortical Neurons

Neuron η consists of three major parts: dendrite, soma, and axon. Figure 2.1 shows a simplified diagram of η, along with the three major parts. The dendrite is essentially the input line to η. There are roughly 10 thousand neurons that connect to η's dendrite. This set of neurons is called the afferent cohort of η. The dendrite structure is tree-like to create space for connections made by the many members of η's afferent cohort. Neuron η receives signals from the members of its afferent cohort in an asynchronous fashion. That is, there is no clock that synchronizes the timing of the signals produced by each member of η's afferent cohort.

The signals received from η's afferent cohort is "integrated" into the PSP in the soma, which is the main body of the cell. The PSP is the voltage across the cellular membrane of η. To simplify the discussion, a small cell approximation for the soma is used, i.e., the voltage of the soma is assumed to be the same within the cell membrane. Thus for a "small cell", the PSP can be viewed as the internal voltage with respect to the voltage of the extracellular fluid outside the cell membrane. In this dissertation, the "excitation" of η is synonymous to its PSP.

The axon is the output line of η that carries voltage signals called action potentials (AP's). The AP's are generated at the axon initial segment (AIS), which is the junction between

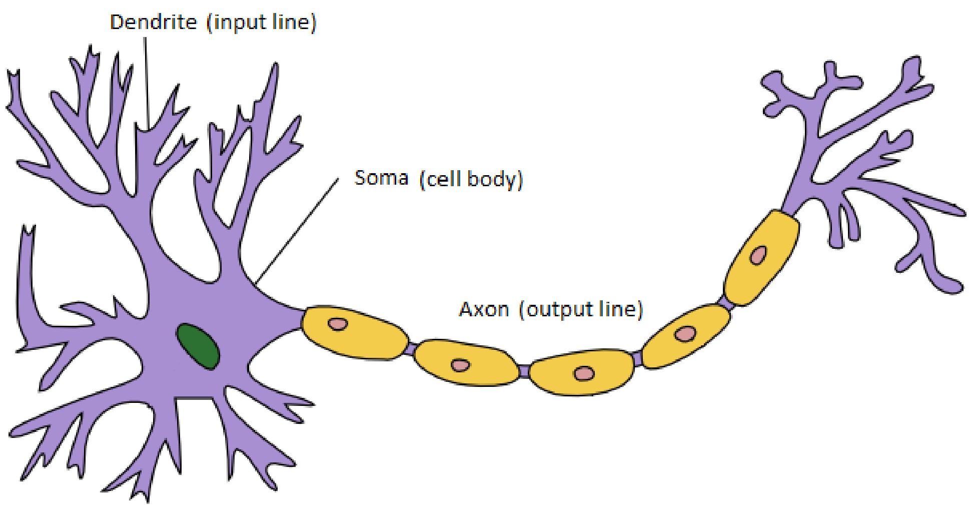

Figure 2.1: A simplified diagram of η. The dendrite is a branching structure that acts as the input line. The soma is the cell body where the vital parts, such as the nucleus, are located. The axon is an extension of the cell that branches out in the end to connect to other cells and acts as the output line.

the soma and the axon. Each AP is the same in shape and amplitude. Thus, information is encoded in the timing of the AP's. Recall from Chapter 1 that the information is encoded in the interpulse interval (IPI), which is the time interval between two adjacent AP's.

The axon is an extension of the cell that connects to η's targets. The AP traveling through the axon is regenerated either periodically or continuously, depending on the type of axon. This functions much like a repeater in the sense that it maintains the shape and strength of η's output signal, which gets attenuated and distorted as it travels along the axon. The axon terminal connects to a set of neurons called η's efferent cohort, which contains roughly 10 thousand members. The same signal generated by η is broadcasted to all members of its efferent cohort.

The connection between a neuron's axon and another's dendrite is called a synapse. A diagram of the synapse is shown in Figure 2.2 The synapse consists of a gap between the axon terminal and the dendrite called the synaptic cleft. Upon arrival of a signal at the axon terminal, chemicals called neurotransmitters are released across the synaptic cleft to receptors on the dendrite. Once the neurostransmitter binds to a receptor, a signal is generated at the

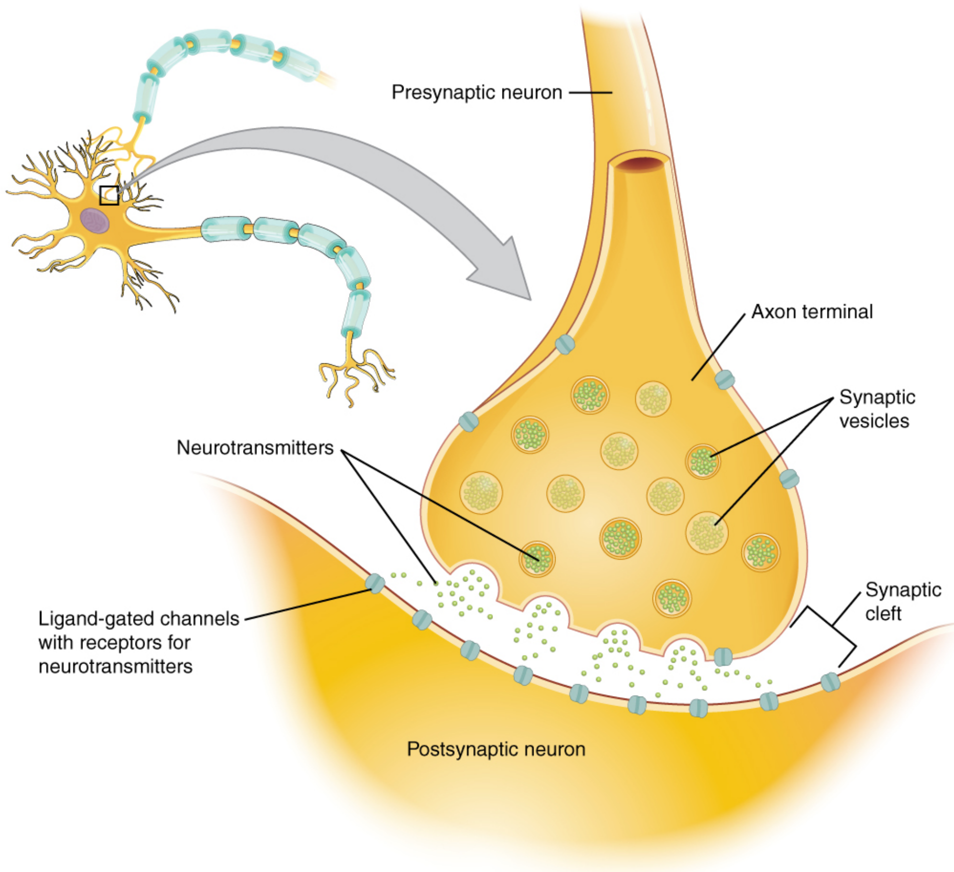

Figure 2.2: A diagram of a synapse. The synapse is the connection between the presynaptic neuron's axon to the postsynaptic neuron's dendrite. Upon arrival on an AP, the axon terminal releases neurotransmiters that diffuse across the synaptic cleft and binds to receptors at the dendrite. Image by OpenStax College is licensed under CC-BY-3.0.

dendrite and travels to the soma. The synapse can be excitatory, where the generated signals are called excitatory PSP's (EPSP's). The EPSP's add positive contributions to the PSP. In inhibitory synapses, the generated signals are called the inhibitory PSP's (IPSP's) and they negatively contribute to the PSP. Associated with each synapse is the synaptic weight, which determines the magnitude of the contributions of the signal to the PSP. Thus, a signal generated by η affects the members of its efferent cohort differently. Likewise, the signals from different members of η's afferent cohort affect η differently.

A more thorough background on neurons can be found in [25] and [26].

2.1.2 Generating an Action Potential

In η's dendrite, there are more excitatory synapses than there are inhibitory ones. Thus, the net contribution of the EPSP's and IPSP's is positive and the PSP is increasing over time. As the PSP increases, voltage-gated ion channels that are present on the cell membrane of η comes into play. There are two primary types of ion channels: sodium ion (Na^+) and potassium ion (K^+) channels. The Na^+ channels are composed of "gates" and an inactivation particle. The gates and the particle exist in an open or closed state and jump back and forth between the two states due to thermal noise. The probability of the gates being in the open state increases with voltage. The reverse is true for the inactivation particle; its probability of being in the open state decreases with voltage. When the Na^+ channels are open, positively charged Na^+ are propelled into the cell due to its concentration gradient. The K^+ channels are composed only of gates. Likewise, these gates are more likely to open if the voltage is high. However, when the gates are open, positively charged K^+ are propelled out of the cell due to its concentration gradient.

Initially, the PSP is in equilibrium where the majority of Na^+ and K^+ channels are closed. The rising PSP causes the Na^+ channels to open bringing in an influx of positively charged Na^+. An increase of positive charges in η further increases the PSP. This in turn increases the number of Na^+ channels that are open. A positive feedback loop is created and this feedback increases the voltage at a fast rate. As the voltage keeps increasing, the inactivation particles start to close and the Na^+ ion channels enter an inactivated state. This slows down the rate of increase of the PSP. The K^+ channels also begin to open, which rapidly slows the increasing voltage and brings it back down towards equilibrium. There is a delay between the rise of the PSP and the opening K^+ channels, which allows the Na^+ channels to push up the PSP before the K^+ channels can bring the PSP back down. When the voltage is low, the K^+ channels close and the PSP attempts to reach equilibrium. This surge of voltage propagates along the axon as the AP. A mathematical model for the generation of the AP was first developed in [27].

After firing an AP, there is a refractory period in which η cannot fire another AP. Not enough Na^+ channels are out of the inactivated state and cannot produce an AP. This is the absolute refractory period. Following the absolute refractory period is the relative refractory period in which η can fire an AP, but with a higher energy cost. Enough Na^+ are out of the inactivated state, but a higher voltage, thus a higher energy, is needed to open enough Na^+ to fire an AP.

As Na^+ and K^+ move in and out of η, ion pumps maintains the concentration of Na^+ and K^+ within the cell. The pump brings Na^+ into η while removing K^+. Thus, it is assumed that the ion concentrations that propel the K^+ and Na^+ remain the same at any given time for η.

The ion channels dynamics are described in greater detail in [25].

2.1.3 The Fast Sodium Channels

In η's AIS, there are multiple types of Na^+ channels. The two types of interest are the $Na_V1.6$ channels and the $Na_V1.2$ channels [28, 29]. They open at a faster rate than other types of Na^+ channels and are crucial in forming the leading edge of AP's. Both types of channels will be referred to as the fast Na^+ channels. The fast Na^+ channels causes the voltage to rise more rapidly, which forms a steep leading edge for the AP. This increases the bandwidth of the AP. In high signal-to-noise ratio (SNR) regimes, the error in estimation of the time of arrival of a signal is inversely proportional to the square of the signal's bandwidth [30]. For the AP, this error is referred to as jitter. Since the AP traveling along the axon is considered to have high SNR, the large bandwidth of the AP reduces jitter. This in turn preserves the information carried by the AP's. Since η reduces jitter, it is assumed to be negligible.

The opening of the fast Na^+ channels can be used to indicate the timing of the generation of AP's. Once a few of the $Na_V1.2$ channels begin to open, the positive feedback loop that creates the AP is virtually irreversible. Hence, the voltage where the $Na_V1.2$ channels have a high likelihood of opening can be set as a threshold level. Once the PSP reaches the threshold,

η is said to have generated an AP. This creates a precise distinction between the PSP and the AP.

2.1.4 Diffusion Models for the Postsynaptic Potential

There are numerous neurons in η's afferent cohort and each member contributes a small amount to η's PSP. It is natural to assume that each contribution is infinitesimally small and that the time interval between two input signals is also infinitesimally small. From the perspective of η, the set of input signals is random. Furthermore, it is assumed that given a height of the PSP, the dynamics of the neuron behaves independently of time. Let Y_t be the PSP at time t. With these assumptions, Y_t can be approximated by a time-homogeneous stochastic differential equation (SDE):

$$\mathrm{d}Y_t = \mu(Y_t)\mathrm{d}t + \sigma(Y_t)\mathrm{d}W_t, \quad t \geq 0, \tag{2.1}$$

where $\{W_t\}$ is the classical Wiener process (WP), μ is the drift, and σ is the square root of the infinitesimal variance σ^2. The infinitesimal contribution to the PSP is assumed to be normally distributed. The drift can be viewed as the average contribution of the afferent cohort to the PSP when the voltage is y. The infinitesimal variance can be viewed as the variance of those contributions when the voltage is y. The solution to SDE's is often called a diffusion.

All SDE's presented in this thesis are of the Itô kind, which is appropriate for modeling small and frequent events as infinitesimally small and frequent. Stochastic differential equations are rigorously defined and elaborated upon in [31].

The first diffusion model for a neuron is the classical WP with (constant) drift [32]. This model can be attained by assuming that η's cell membrane is electrically equivalent to a resistor before hitting the threshold. It can also be attained by assuming that the ion channels function to keep the drift of the diffusion constant before the threshold is hit. This linearization of the drift has been observed [33–36]. Another popular diffusion model is the Ornstein-Uhlenbeck (O-U) diffusion, which can be attained by modeling η's cell membrane with a resistor and a capacitor in parallel [37–42]. The Feller model imposes a lower bound on the diffusion model of η's PSP [43–45]. The lower bound is imposed because in a real setting, the PSP of η cannot be unbounded from below. Finally, another approach to modeling the PSP is to take existing neuron models and use a white noise input to create diffusion models. The approach has been applied to the Hodgkin-Huxley (H-H) model [46, 47] and the Fitzhugh-Nagumo (F-N) model [48].

2.2 The GIG Neuron Model

The GIG distribution and the GIGHT diffusion were first proposed as a model for a neuron's behavior in [49]. A contribution of this dissertation is to describe in detail how the GIGHT diffusion can model the behavior of the PSP of η brought about by the fast Na^+ channels. Before the GIG neuron model is explained, the input and output model for η is described.

2.2.1 Output and Input Models

Let S_k be the time of generation the k^{th} AP of η. An AP is defined to be generated when the PSP of η hits the threshold. Let S_k be modeled by a random variable (RV). The times of the generation of the AP is given by the sequence of RV's $\{S_k\}_{k=1}^{\infty}$. Let $S_0 = 0$ and define $T_k = S_k - S_{k-1}$. The random variable T_k is then the duration of the k^{th} IPI. Under the timing code hypothesis, the information trasmitted by η is encoded in the sequence of IPI's, whose durations given by $\{T_k\}_{k=1}^{\infty}$. Hence, let $\{T_k\}$ be the sequence of the output of η.

The input to η is an underlying net rate of bombardment from η's afferent cohort. Equivalently, η's afferent cohort wishes to convey some scalar to η and does so by bombarding η with AP's at some frequency. Let the underlying net rate at time t be $\Lambda(t)$. Hence, the input to η can be modeled by a random process (RP) $\{\Lambda(t)\}_{t>0}$. However, η only observes a noisy version of $\{\Lambda(t)\}$ due to the nature of the input signals. For example, the Stein-Chen approximation can be used to approximate the arrivals of signals from η's afferent cohort as a Poisson process since the number of input size is large and the signals arrive asynchronously [50]. The underlying Poisson rate is not directly observable, but can be estimated using the number of arrivals. Similarly, η can only estimate $\Lambda(t)$ from the arrivals of pulses from its afferent cohort.

To simplify the problem, $\Lambda(t)$ is approximated by a sequence of RV's. In the k^{th} IPI, the average input intensity is given by

$$\Lambda_k = \frac{1}{T_k}\int_{S_{k-1}}^{S_k} \Lambda(t)\mathrm{d}t, \tag{2.2}$$

The input intensity is approximated by a step process whose height in the k^{th} IPI is Λ_k. Hence, the input of η is given by the sequence $\{\Lambda_k\}_{k=1}^{\infty}$. Upon completion of the k^{th} IPI, η received input Λ_k and sent output T_k.

To further simplify the model, a memoryless assumption is made. The PSP in a given IPI is assumed to be independent of the PSP in the previous IPI's. Thus, given the input, the output of η has the same statistics in any IPI and is independent of past inputs. In reality, the PSP of η may not reset to the same exact level every time, which affects the output of the next IPI. It is assumed that η's PSP starts at the same level during each IPI and the ion channels are reset by the start of the IPI. Partial reset of the ion channels during the relative refractory period is handled by increasing the energy required to fire an AP. This will be discussed in Chapter 3 when discussing the energy cost model of η's function. Hence, the focus can remain on a single IPI.

Let the prototypical input and output RV's for η be given by Λ and T, respectively. Note the abuse of notation for Λ, which represents a RV here. To make a distinction, the continuous RP that is the input intensity of η is denoted by $\{\Lambda(t)\}$ and the input intensity at any given time is given by $\Lambda(t)$. Thus, when one refers to simply Λ, it is the average input intensity of η for a typical IPI, and it is a RV.

2.2.2 Input-Output Relationship

For the GIG neuron model, the distribution of output $T = t$ given input $\Lambda = \lambda$ is given by

$$Q_{\text{GIG}}(t|\lambda) = M(\alpha, \beta, \gamma)^{-1} \lambda^{\alpha} t^{\alpha-1} \exp\left(-\frac{\beta}{\lambda t} - \gamma \lambda t \right), \quad \lambda, t > 0, \tag{2.3}$$

where $\alpha \le -1/2,\ \beta > 0,\ \gamma \ge 0$, and

$$M(\alpha, \beta, \gamma) = \begin{cases} 2\left(\frac{\beta}{\gamma}\right)^{\alpha/2} K_{\alpha}(2\sqrt{\beta\gamma}) & \beta > 0, \gamma > 0 \\ \beta^{\alpha}\Gamma(-\alpha) & \alpha < 0, \beta > 0, \gamma = 0. \end{cases} \tag{2.4}$$

The function Γ is the gamma function and K_{α} is the modified Bessel function of the second kind of order α, also known as the Hankel function. The PDF (2.3) is an instance of the GIG distribution, hence the name, GIG neuron model.

The GIG neuron model is appealing because (2.3) as the relationship between the input and output of η can be derived from first principles [51]. In addition, there also exists a diffusion associated with (2.3), namely the GIGHT, that can model the PSP of η. Thus, the GIG neuron model can also be considered as a diffusion model. This diffusion is described in the next subsection.

2.2.3 The GIG distribution and the GIGHT diffusion

In general, the GIG distribution is given by

$$f_{\text{GIG}}(t) = M(\alpha, \beta, \gamma)^{-1} t^{\alpha-1} \exp\left(-\frac{\beta}{t} - \gamma t \right), \quad t > 0, \tag{2.5}$$

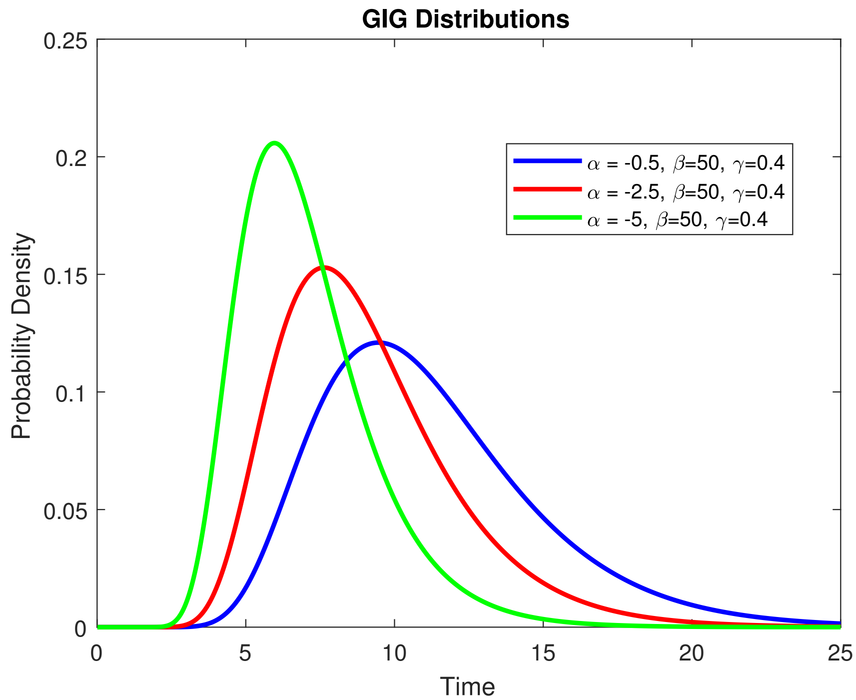

Figure 2.3: An plot of GIG distributions.

where,

$$M(\alpha, \beta, \gamma) = \begin{cases} 2\left(\frac{\beta}{\gamma}\right)^{\alpha/2} K_\alpha(2\sqrt{\beta\gamma}) & \beta > 0, \gamma > 0 \\ \gamma^{-\alpha}\Gamma(\alpha) & \alpha > 0, \beta = 0, \gamma > 0 \\ \beta^{\alpha}\Gamma(-\alpha) & \alpha < 0, \beta > 0, \gamma = 0. \end{cases} \tag{2.6}$$

The parameter space for the GIG PDF is

$$\beta \in (0, \infty),\ \gamma \in [0, \infty) \text{ if } \alpha \in (-\infty, 0),$$
$$\beta \in (0, \infty),\ \gamma \in (0, \infty) \text{ if } \alpha = 0,$$
$$\beta \in [0, \infty),\ \gamma \in (0, \infty) \text{ if } \alpha \in (0, \infty).$$

This distribution is denoted as GIG(α, β, γ). A plot of GIG distributions is shown in Figure 2.3.

When $\alpha = -1/2$, (2.5) is also known as the inverse Gaussian (IG) distribution and it is the FHT of the WP with drift [52]. When $\beta = 0$, it is also known as the gamma distribution and when $\gamma = 0$, it is also known as the inverse gamma distribution.

For $\alpha \leq 0$, the GIG distribution is the first hitting time (FHT) distribution of the GIGHT diffusion, which is the solution to the SDE (2.1) with drift given by [53]

$$\mu(y) = \sigma(y)\left(\frac{2\alpha - 1}{2\chi_\theta(y)} + \frac{\sqrt{2\gamma}K_{\alpha-1}(\chi_\theta(y)\sqrt{2\gamma})}{K_\alpha(\chi_\theta(y)\sqrt{2\gamma})}\right) + \frac{1}{4}\frac{\mathrm{d}}{\mathrm{d}y}\sigma^2(y), \quad y < \theta, \tag{2.7}$$

for $\gamma > 0$. For $\alpha < 0$ and $\gamma = 0$, the drift is given by the limit of (2.7) as $\gamma \to 0$, which is

$$\mu(y) = \sigma(y)\frac{2\alpha + 1}{2\chi_\theta(y)} + \frac{1}{4}\frac{\mathrm{d}}{\mathrm{d}y}\sigma^2(y), \quad y < \theta. \tag{2.8}$$

For both cases, $\theta > 0$ and the drift is a function of the infinitesimal variance σ^2. The function

$$\chi_\theta(y) = \int_y^\theta \frac{\mathrm{d}y'}{\sigma(y')}, \quad y < \theta. \tag{2.9}$$

The parameter β in (2.5) is given by

$$\beta = \chi_\theta(0)^2/2. \tag{2.10}$$

The diffusion begins at 0, so $Y_0 = 0$. The infinitesimal drift σ^2 satisfies

$$\chi_\theta(y) < \infty \text{ for } \theta - y < \infty$$

and

$$\chi_\theta(y) \to \infty \text{ as } \theta - y \to \infty.$$

The GIGHT diffusion takes value in $(-\infty, \theta]$ and is defined while $Y_t \leq \theta$. Once θ is hit, the diffusion is no longer defined. Hence, θ is a threshold level for the diffusion. A plot of the GIGHT diffusions that correspond to the GIG PDF's of Figure 2.3 is shown in Figure 2.4.

A property of the GIG distribution is that if T is distributed as GIG(α, β, γ), then T^{-1}

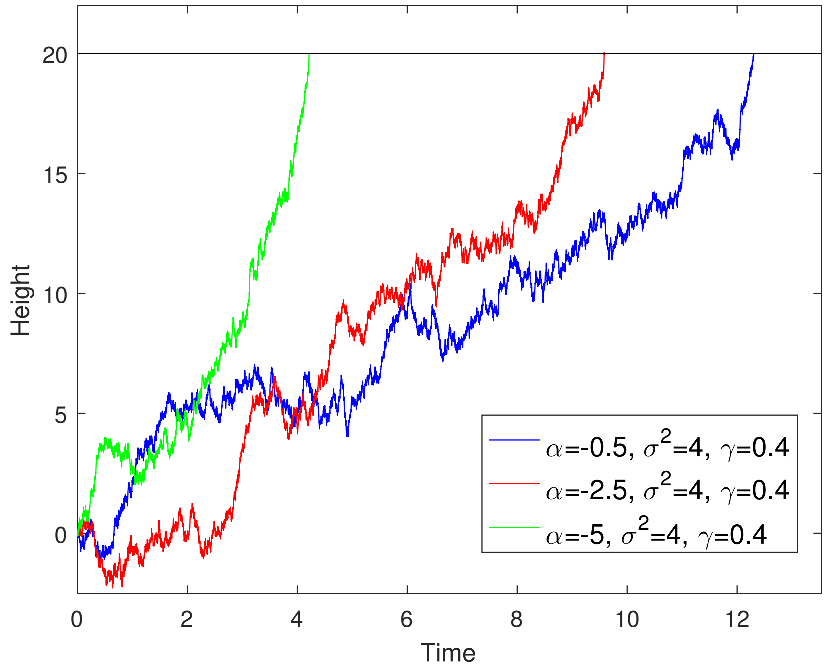

Figure 2.4: Three sample paths for the GIGHT diffusion with constant infinitesimal variance. The horizontal line is the threshold. The further the value of α gets below zero, the stronger the attraction to the threshold. For $\alpha = -1/2$, the GIGHT diffusion is also the WP with drift. The time and height units in this example are arbitrary.

is distributed as $\mathrm{GIG}(-\alpha, \gamma, \beta)$ [54]. Therefore, any GIG distribution with a positive α parameter is the reciprocal of the FHT of the corresponding GIGHT diffusion.

2.2.4 The GIGHT Diffusion with Constant Infinitesimal Variance

Like the WP with drift and O-U diffusion, take the infinitesimal variance of the GIGHT diffusion to be constant. This allows the diffusion to be described in a simple way. Using a slight abuse of notation, let this constant value be σ^2. Henceforth, σ^2 and $\sigma = \sqrt{\sigma^2}$ are constants. Then (2.7) reduces to

$$\mu(y) = \sigma^2 \frac{\alpha - \frac{1}{2}}{\theta - y} + \sigma \frac{\sqrt{2\gamma} K_{\alpha-1}(\frac{\theta-y}{\sigma}\sqrt{2\gamma})}{K_\alpha(\frac{\theta-y}{\sigma}\sqrt{2\gamma})}, \quad y < \theta. \tag{2.11}$$

In this case, The GIGHT diffusion initially increases at an approximately steady rate. For $\alpha < -1/2$, the diffusion is attracted to the threshold as it is approached. For $\alpha > -1/2$, the threshold is repulsive. For $\alpha = -1/2$, the diffusion is indifferent to the threshold. The parameter γ controls the component of the drift that is independent of the threshold. The parameters are further described in Section 2.2.7.

The behavior of GIGHT diffusions is better understood when it is either far from the threshold or near it. Define the function

$$O_\alpha(x) = \alpha - \frac{1}{2} + x\frac{K_{\alpha-1}(x)}{K_\alpha(x)}, \quad x > 0. \tag{2.12}$$

Then the drift (2.11) can be written as

$$\mu(y) = \frac{\sigma^2}{\theta - y} O_\alpha\left(\frac{\theta - y}{\sigma}\sqrt{2\gamma}\right), \quad y < \theta. \tag{2.13}$$

The function O_α can be simply described at its asymptotes. From Appendix A.3, it is clear that as $x \to \infty$,

$$O_\alpha(x) \sim x + \alpha - \frac{1}{2}, \tag{2.14}$$

where $\sim$ refers to asymptotic equality (also see Appendix A.3). The asymptote of O_α is an oblique line with slope 1. Thus, as $\theta - y \to \infty$,

$$\mu(y) \to \sigma\sqrt{2\gamma}. \tag{2.15}$$

Thus, when $y \ll \theta$, the GIGHT diffusion has approximately constant drift. Thus the GIGHT diffusion is approximately a WP with drift when $y \ll \theta$.

On the other hand, as $x \to 0$,

$$O_\alpha(x) \to -\alpha - \frac{1}{2}. \tag{2.16}$$

The function O_α approaches a constant, which may be positive, depending on the value of α. Then, as $y \to \theta$,

$$\mu(y) \sim \left(-\alpha - \frac{1}{2} \right) \frac{\sigma^2}{\theta - y}. \tag{2.17}$$

The drift approaches a negative number if $\alpha > -1/2$, which shows that the threshold is repulsive. Likewise, if $\alpha < -1/2$, the drift grows as the threshold is approached. Thus, the threshold is attractive. The drift is inversely proportional to the distance between the diffusion and the threshold. In the case that it is attractive, as the diffusion approaches the threshold, the drift approaches infinity. Hence, the threshold is hit by a rate that is infinitely big.

For $\alpha = -1/2$ it can be shown that the GIGHT diffusion is exactly a WP with drift when it is below θ. From the Hankel's expansion for K_α (see Appendix A.2),

$$K_{-1/2}(x) = \sqrt{\frac{\pi}{2x}} e^{-x}, \quad x > 0 \tag{2.18}$$

and

$$K_{-3/2}(x) = \sqrt{\frac{\pi}{2x}} e^{-x} \left(1 + \frac{1}{x} \right), \quad x > 0. \tag{2.19}$$

Hence,

$$O_{-1/2}(x) = x \tag{2.20}$$

and (2.11) reduces to

$$\mu(y) = \sigma\sqrt{2\gamma}, \quad y < \theta. \tag{2.21}$$

Thus the drift is exactly constant, which gives the WP with drift when $Y_t \leq \theta$. Hence, the GIGHT diffusion is a generalization of the WP with drift. This is consistent with the fact that for $\alpha = -1/2$, the GIG distribution is the FHT of the WP with drift.

2.2.5 The Role of the Input Distribution on the Diffusion Model

What role does the input Λ have on η's PSP? To address the question, the following assumption is made: the drift and infinitesimal variance is proportional to Λ. Consider the ideal integrate-and-fire model in [11] where the arrival of input signal to η is modeled as a Poisson process. For this model, the mean and variance of the PSP is proportional to Λ. This same property can be extended to the GIGHT diffusion by letting the drift and infinitesimal variance be $\Lambda\mu$ and $\Lambda\sigma^2$, respectively. For a given $\Lambda = \lambda$, the corresponding SDE is

$$\mathrm{d}Y_t = \mu(Y_t)\lambda \mathrm{d}t + \sigma\sqrt{\lambda}\mathrm{d}W_t, \quad t \geq 0. \tag{2.22}$$

The resulting diffusion is still a GIGHT diffusion with $\lambda\gamma$ and $\chi_\theta/\sqrt{\lambda}$ replacing γ and χ_θ, respectively. Hence, β and γ in (2.5) is replaced with β/λ and $\gamma\lambda$, respectively. Note that

$$M(\alpha, \beta/\lambda, \gamma\lambda) = M(\alpha, \beta, \gamma)\lambda^{-\alpha}, \tag{2.23}$$

so λ does not appear in the argument of M. Therefore, given $\Lambda = \lambda$, T is distributed as $\mathrm{GIG}(\alpha, \beta/\lambda, \gamma\lambda)$, whose PDF is (2.3). This gives the GIG neuron model for η. Note that for the model, $\alpha \leq -1/2$. This is explained in the next subsection.

Equation (2.3) gives the statistical relationship between the input and the output for the GIG neuron model. This can also can be viewed as a communication channel. Therefore, the GIG neuron model is also be referred to as the GIG neuron channel.

2.2.6 GIGHT Diffusion Approximates PSP of Cortical Neurons

To fit the GIGHT diffusion to the PSP of η, define the PSP level at the beginning of the IPI to be 0. Then let θ be slightly above the firing threshold of η. This is done because the GIGHT diffusion hits θ with infinite rate. Although the firing threshold of η is hit with high rate, it cannot be hit with infinite rate. Thus, the GIGHT diffusion is no longer a

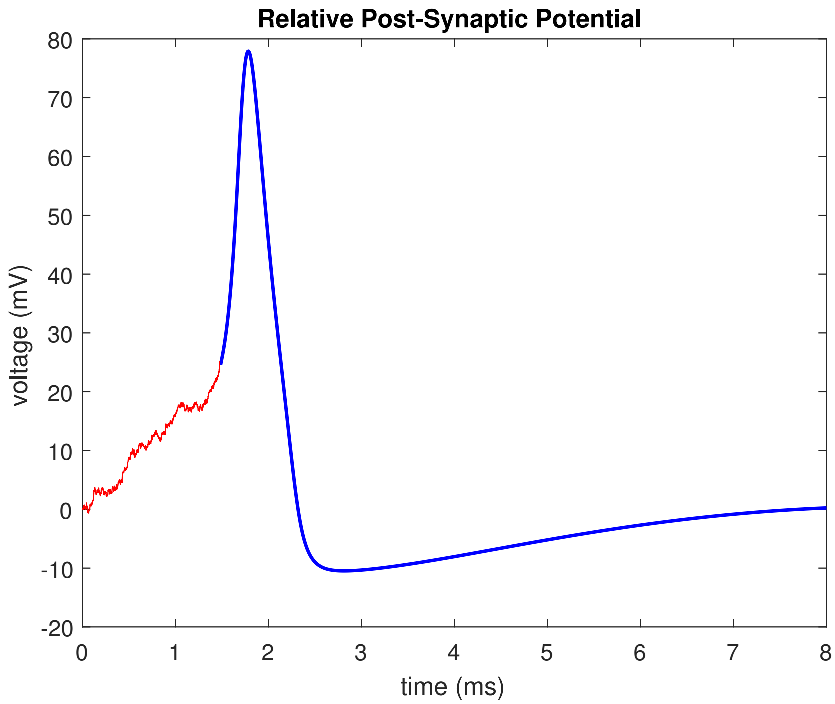

Figure 2.5: An example of a GIGHT diffusion as the PSP of a neuron leading up to an AP.

good model right before it hits θ. However, the time difference between η's PSP hitting the firing threshold and the GIGHT diffusion hitting θ is small because the GIGHT diffusion is increasing at a very fast rate. With high probability, once the GIGHT hits the firing threshold, it will hit θ in a small amount time. Thus, the first hitting time (FHT) of the GIGHT diffusion approximates the duration of the IPI of η.

The GIGHT Diffusion has certain properties that make it a good model for the behavior of the PSP of η. When the diffusion is away from the threshold, it appears to build up steadily. Likewise for neurons, it has been observed that in the beginning of an IPI, the PSP builds up approximately linearly [33–36, 55]. As described in Section 2.1, as the PSP increases, its rate of increase goes up due to the Na^+ channels, especially the $Na_V1.2$ and $Na_V1.6$ channels. For $\alpha < -1/2$, the GIGHT diffusion exhibits this behavior. This is illustrated in Figure 2.5, where the GIGHT diffusion is superimposed on the PSP leading up to an AP. Thus, for the GIG neuron model, $\alpha \leq -1/2$, where equality is allowed to include the linear case.

There several advantages of the GIGHT diffusion over other diffusion models. The GIGHT diffusion generalizes the WP with drift and allows for an increase in the drift of the diffusion as it approaches the threshold. Thus, the GIGHT is able to model the PSP more accurately than the WP with drift. Likewise, the O-U diffusion model does not exhibit an increasing drift as threshold is approached. In fact, it is assumed that the effects of the ion channels are negligible until the threshold is hit, and thus the drift is actually decreasing as the threshold is approached. As previously mentioned, the PSP should grow at a linear rate, followed by an upswing. This is not exhibited by the O-U diffusion model.

There are other diffusion models that include the upswing exhibited by PSP's, such as the quadratic and exponential integrate-and-fire models [56, 57]. More biophysically accurate models, such as the stochastic (H-H) and (F-N) models, also include this upswing. However, the FHT distribution of these models are not given explicitly, which makes analysis intractable. On the other hand, the FHT distribution of the GIGHT diffusion is the desired GIG distribution. This makes the input-output model relatively simple and makes analysis feasible.

As a note, there may be occasional cases where the GIGHT diffusion model does not fit the behavior of the PSP of η. This model only applies under normal circumstances. Long IPI's may be terminated by higher regions in the brain in order to save energy. As a result, the PSP resets to its initial value or η is forced to fire an AP. In such a case, the GIGHT diffusion is not an appropriate model.

2.2.7 The GIG Neuron Model Parameters

For a constant infinitesimal variance, the parameters of the GIGHT diffusion and GIG distribution is related to the physical entities of the diffusion in the following ways [19, 20]:

$$\alpha = -\frac{\mu_0 \theta}{\sigma^2} - \frac{1}{2}, \tag{2.24}$$

$$\beta = \frac{\theta^2}{2\sigma^2}, \tag{2.25}$$

$$\gamma = \frac{{\mu_c}^2}{2\sigma^2}, \tag{2.26}$$

where μ_0 is the initial drift solely due to the attraction of the threshold and μ_c is the constant drift sans the effect of the barrier. Note that $\mu(0) \neq \mu_0 + \mu_c$. The drift components are not additive but interact in a more complex way. This can be observed by letting $y = 0$ and substituting (2.24) and (2.26) into (2.11). Also, the product $\mu_0\theta$ is the proportionality constant between the the drift near the threshold and the reciprocal of the distance between the threshold and the drift. This can be observed from (2.17) by letting $y = 0$.

To derive the expression for γ, let $\theta - y \to \infty$. Then the drift approaches a constant as shown in (2.15). Therefore

$$\mu_c = \sigma\sqrt{2\gamma}, \tag{2.27}$$

where $\mu(y)$ from (2.15) has been replaced by μ_c to indicate that this is the drift when $\theta - y \to \infty$. Since it is not a function of y, this dependence is dropped. Solving for γ givens (2.26).

Since μ_c is constant, the drift under the $\theta - y \to \infty$ assumption is a Wiener process with constant drift, which exhibits no attraction to nor repulsion from the threshold. This same effect can be achieved by letting $\alpha = -1/2$. Since the threshold is infinitely far, it is too distant to attract or repel the diffusion. Thus μ_c can be interpreted as the drift sans the effect of the barrier.

The value of β is given in (2.10). Assuming a constant infinitesimal variance yields

$$\chi_\theta(y) = \frac{\theta - y}{\sigma}, \quad y < \theta. \tag{2.28}$$

Therefore, β reduces to (2.25).

Finally, to get an expression for α, let $\mu_c \to 0$. As a consequence, $\gamma \to 0$. Using one of the asymptotic properties of K_α (see Appendix A.3), the drift becomes

$$\tilde{\mu}(y) = \left(-\alpha - \frac{1}{2}\right)\frac{\sigma^2}{\theta - y}, \quad y < \theta, \tag{2.29}$$

where μ has been replaced by $\tilde{\mu}$, which is the drift when $\gamma = 0$. Define $\mu_0 = \tilde{\mu}(0)$. Then letting $y = 0$ and solving for α yields (2.24). Since the expression for α was attained by letting $\mu_c \to 0$, the drift $\tilde{\mu}$ is solely due to the attraction of the threshold without a constant drift component. Hence, α is linearly dependent on the initial value of the threshold-dependent drift.

Naturally, the next question is how to estimate the parameters given a sample path of the GIGHT diffusion.

2.3 Estimation of the GIG Neuron Model Parameters

The goal is to estimate σ^2, α, and γ from a realization of a GIGHT diffusion that is sampled at regular intervals. Since the PSP of η behaves like a GIGHT diffusion, the same techniques could be applied to recordings of η's PSP. Thus, a systematic way to acquire the parameters of the GIG neuron model is developed. The values of the parameters may reveal important qualities of η and can be used to test the accuracy of the GIG neuron model. Furthermore, once the parameters of the model are determined, estimators for the input intensity of the neuron are also developed. Hence, the value of Λ can be estimated from recordings. Repeated trials can be used to attain a distribution of Λ. The distribution of Λ is further discussed in Chapter 3.

Parameter estimation of diffusions has been studied in depth [58]. However, to the author's knowledge, applications to neural models have appeared only recently. Work has been done on parameter estimation of neural diffusion models, such as the O-U and Feller models [20, 45, 59–61]. Other works involve estimating parameters of stochastic models that were more biophysical, such as a stochastic version of Hodgkin-Huxley and Fitzhugh-Nagumo models [62–65].

2.3.1 The Euler-Maruyama Approximation

Two methods for estimating the parameters are developed [20]. The first is an approximation of the maximum likelihood estimators (MLE's) by using the Euler-Maruyama method [66]. These estimators are called the pseudo-maximum likelihood estimators (pMLE's). The second method involves approximating O_α via the pMLE. Then a least square fit is used to extract the estimates of α and γ. These estimators are called the pseudo-least square fit estimates (pLSFE's). Throughout the development, the value of θ is assumed to be known, whereupon β is given by (2.25) in terms of θ and σ^2.

Let T be the FHT of the GIGHT diffusion. Let time be discretized with a sampling period of Δ. Then, given that $\Lambda = \lambda$ and assuming a constant infinitesimal variance, (2.22) can be approximated by the Euler-Maruyama method:

$$\hat{Y}_k - \hat{Y}_{k-1} = \mu(\hat{Y}_{k-1})\lambda\Delta + \sigma\sqrt{\lambda}(\hat{W}_k - \hat{W}_{k-1}), \quad k = 1, 2, \ldots, K, \tag{2.30}$$

where $\hat{Y}_k = Y_{k\Delta}$, $\hat{W}_k = W_{k\Delta}$, and K is the maximum number of samples taken by the discretization scheme, excluding the initial point $Y_0 = 0$. The value of K is given by $K = \lfloor T/\Delta \rfloor$ where $\lfloor \cdot \rfloor$ is the floor function. Also, $K\Delta \to T$ as $\Delta \to 0$. It is clear that the discrete RP $\{\hat{Y}_k\}$ forms a time-homogeneous Markov process.

Given $\hat{Y}_{k-1} = y_{k-1}$, Y_k is a normally distributed RV with mean $y_{k-1} + \mu(y_{k-1})\Delta\lambda$ and variance $\sigma^2\Delta\lambda$. Let $\underline{Y} = \{\hat{Y}_k\}_{k=1}^{K}$, $\underline{y} = \{y_k\}_{k=1}^{K}$, and $y_0 = 0$. Then the joint PDF of $\underline{Y}$ is

$$f_{\underline{Y}|\Lambda}(\underline{y}|\lambda) = \frac{\exp(-\frac{1}{2\sigma^2\Delta\lambda}\sum_{k=1}^{K}(v_k - v_{k-1} - \mu(y_{k-1})\Delta\lambda)^2)}{(2\pi\sigma^2\Delta\lambda)^{K/2}}. \tag{2.31}$$

This distribution is not necessarily jointly Gaussian because μ may be a non-linear function. Therefore, $\underline{Y}$ is not necessarily a Gaussian process. The log-likelihood is attained by taking the logarithm of (2.31), which yields

$$L_{\underline{Y}|\Lambda}(\underline{y}|\lambda) = -\frac{1}{2\sigma^2\Delta\lambda}\sum_{k=1}^{K}(v_k - v_{k-1} - \mu(v_{k-1})\Delta\lambda)^2 - \frac{K}{2}\log(2\pi\sigma^2\Delta\lambda), \tag{2.32}$$

where log is the natural logarithm function.

2.3.2 The Pseudo-Maximum Likelihood Estimator

For now, assume that $\Lambda = 1$. Without a reference, the value of Λ is arbitrary. To demonstrate, suppose the estimates for σ^2, α, and γ are σ_0^2, α_0, and γ_0, respectively. Then suppose that the "real" value of Λ is λ_0. Then the estimated parameters can be scaled to form new estimates that correspond with $\Lambda = \lambda_0$: σ_0^2/λ_0, α_0, and γ_0/λ_0, respectively.

The MLE of α based on Euler-Maruyama approximation can be attained by differentiating (2.32) with respect to α and setting it equal to 0, which results in

$$\sum_{k=1}^{K}(y_k - y_{k-1} - \mu(y_{k-1})\Delta)\frac{\partial}{\partial\alpha}\mu(y_{k-1}) = 0. \tag{2.33}$$

This equation is not analytically tractable. By considering the GIGHT diffusion near the threshold, the drift can be approximated by (2.17). Since the diffusion ends by hitting the threshold, the last samples are assumed to be near the threshold. By considering the last m samples of the diffusion and using the approximation, the estimator for α is

$$\hat{\alpha}(\underline{y}) = -\frac{\sum_{k=K-m+1}^{K}\frac{y_k - y_{k-1}}{\theta - y_{k-1}}}{\hat{\sigma}^2(\underline{y})\sum_{k=K-m+1}^{K}(\theta - y_{k-1})^{-2}\Delta} - \frac{1}{2}, \tag{2.34}$$

where $\hat{\sigma}^2(\underline{y})$ is an estimator for σ^2. Potential estimators for σ^2 are considered in Section 2.3.3.

Similarly for γ, the following condition for the approximate MLE be attained:

$$\sum_{k=1}^{K}(y_k - y_{k-1} - \mu(y_{k-1})\Delta)\frac{\partial}{\partial\gamma}\mu(y_{k-1}) = 0. \tag{2.35}$$

This is also not analytically tractable. However, by considering the diffusion in the beginning and assuming that θ is large, the drift can be approximated by (2.15). By considering the first n samples of the GIGHT diffusion, the estimator for γ is

$$\hat{\gamma}(\underline{y}) = \frac{1}{2\hat{\sigma}^2(\underline{y})}\left(\frac{y_n}{n\Delta}\right)^2. \tag{2.36}$$

This is also an estimator for a constant drift [58].

2.3.3 Infinitesimal Variance Estimator

There are multiple ways to estimate σ^2. The first is based on the quadratic variation [66]:

$$\hat{\sigma}^2(\underline{y}) = \frac{\sum_{k=1}^{K}(y_k - y_{k-1})^2}{K\Delta}. \tag{2.37}$$

This estimator is referred to as the "simple" estimator of σ^2. Note that this estimator does not use any information on the drift. It is not necessarily unbiased, which can lead to inaccurate estimation as shown in a later subsection.

A similar approach to estimating γ can be used in improving the estimate of σ^2. Assume that the threshold is large and only the beginning of the diffusion is considered. Then the drift is approximately given by (2.15). Suppose only the first ℓ samples are used. Then taking the derivative of (2.32) with respect to σ^2, setting it equal to zero, and some algebra yields

$$\sum_{k=1}^{\ell}\frac{(y_k - y_{k-1})^2}{2\Delta\sigma^4} - \sqrt{\frac{\gamma}{2}}\frac{y_\ell}{\sigma^3} - \frac{\ell}{2\sigma^2} = 0. \tag{2.38}$$

Using the estimator for γ in (2.36), the estimator for σ^2 is then given by

$$\hat{\sigma}_b^2(\underline{y}) = \frac{\sum_{k=1}^{\ell}(y_k - y_{k-1})^2}{\Delta\ell} - \frac{y_\ell^2}{\ell^2\Delta}. \tag{2.39}$$

To make $\hat{\sigma}_b^2$ unbiased for $\alpha = -1/2$ or under the constant drift assumption, a scale factor of $\frac{\ell}{\ell-1}$ may be applied. Hence, the estimator for σ^2 is $\hat{\sigma}^2 = \frac{\ell}{\ell-1}\hat{\sigma}_b^2$, i.e.,

$$\hat{\sigma}^2(\underline{y}) = \frac{\sum_{k=1}^{\ell}(y_k - y_{k-1})^2}{\Delta(\ell-1)} - \frac{y_\ell^2}{\ell(\ell-1)\Delta}. \tag{2.40}$$

This is called the constant drift estimator (CDE) of σ^2.

Thus, the parameters of the GIG neuron model can be estimated via the pMLE. To reiterate, these are MLE's. These are analytical estimators based on approximations of the likelihood. There exist other estimators that are more favorable in terms of maximizing the likelihood, namely the MLE's; however, the MLE's are intractable and require a numerical solution.

2.3.4 The Pseudo-Least Square Fit Estimator

For the pLSFE, σ^2 is first estimated via the simple estimator or the CDE. Then, the sequence $\left\{O_\alpha\left(\frac{\theta-y_k}{\sigma}\sqrt{2\gamma}\right)\right\}_{k=1}^{K}$, where O_α was defined in (2.12), is estimated via the approximated MLE, i.e., by maximizing over the approximated log-likelihood function (2.32). Then a least square fit regression is used to estimate α and γ. For a sample path of the GIGHT diffusion $\{\underline{y}\}$, Let

$$o_k = O_\alpha\left(\frac{\theta - y_k}{\sigma}\sqrt{2\gamma}\right), \quad k = \{0, 1, \ldots, K-1\}. \tag{2.41}$$

Then, by using (2.13), the log-likelihood can be written as

$$L_{\underline{Y}|\Lambda}(\underline{y}|1) = -\frac{1}{2\sigma^2\Delta}\sum_{k=1}^{K}\left(y_k - y_{k-1} - \frac{\sigma^2\Delta}{\theta - y_{k-1}}o_{k-1}\right)^2 - \frac{K}{2}\log(2\pi\sigma^2\Delta), \tag{2.42}$$

where Λ is still assumed to have taken a value of 1. Using the estimate for σ^2, the approximate MLE for each o_k can be determined and is given by

$$\hat{o}_{k-1}(\underline{y}) = \frac{(y_{k-1} - y_k)(\theta - y_{k-1})}{\Delta\hat{\sigma}^2(\underline{y})}, \quad k = \{1, \ldots, K\}. \tag{2.43}$$

Then the least squares fit to O_α can be numerically determined by finding the values of α and γ that minimizes the mean square error (MSE), i.e.,

$$(\hat{\alpha}(\underline{y}), \hat{\gamma}(\underline{y})) = \arg\min_{(\alpha,\gamma)} \sum_{k=1}^{K} \left(O_\alpha \left(\frac{\theta - y_{k-1}}{\hat{\sigma}(\underline{y})} \sqrt{2\gamma} \right) - \hat{o}_{k-1}(\underline{y}) \right)^2, \tag{2.44}$$

where $\hat{\sigma} = \sqrt{\hat{\sigma}^2}$.

2.3.5 Estimator for the Input Intensity

To estimate the input intensity, the approximate log-likelihood, (2.32), is again used. Let Λ take on the value of λ. The other parameters are assumed to be known and that λ is the only unknown. Then take the partial derivative with respect to λ and set it equal to zero. Then, some algebra yields

$$\lambda^2 \Delta^2 \sum_{k=1}^{K} \mu(v_{k-1})^2 + \lambda \Delta K \sigma^2 - \sum_{k=1}^{K} (v_k - v_{k-1})^2 = 0. \tag{2.45}$$

This is a quadratic equation in λ with possibly two solutions. Since $\lambda > 0$, the only sensible solution to (2.45) is the positive one, i.e.,

$$\hat{\lambda}(\underline{y}) = \frac{-K\sigma^2 + \sqrt{K^2\sigma^4 + 4d\sum_{k=1}^{K}(y_k - y_{k-1})^2/\Delta}}{2d}, \tag{2.46}$$

where

$$d = \sum_{k=1}^{K} \mu(y_{k-1})^2 \Delta.$$

This is the the MLE based on the approximate likelihood of λ.

As $\Delta \to 0$, it can be shown that $\hat{\lambda}$ approaches something like the simple estimator for σ^2 in (2.37). Consider the diffusion in the interval $[0, t^-]$, where $t^- < T$. Note that for $t \in [0, t^-]$, $\mu(Y_t) < \infty$. Let K^- be the number of samples in the interval $[0, t^-]$. Then d is a Riemann sum that converges to a finite limit:

$$d = \sum_{k=1}^{K^-} \mu(Y_{k-1})^2 \Delta \to \int_0^{t^-} \mu(Y_t)^2 \mathrm{d}t \text{ as } \Delta \to 0. \tag{2.47}$$

Hence, $\lambda^2 \Delta d \to 0$ and (2.45) can be approximated as

$$\lambda \Delta K^- \sigma^2 - \sum_{k=1}^{K^-} (y_k - y_{k-1})^2 = 0. \tag{2.48}$$

If $K^- = K$, the estimator of λ is approximately

$$\hat{\lambda}(\underline{y}) = \frac{\sum_{k=1}^{K} (y_k - y_{k-1})^2}{\Delta K \sigma^2}, \tag{2.49}$$

for small values of Δ. This solution is almost the same the simple estimator in (2.37), except the divisor has an additional factor of σ^2 in (2.49). It is clear that (2.49) is estimating $\lambda\sigma^2/\sigma^2$, i.e., just λ.

2.4 Simulation Results

To test the accuracy of the estimators, GIGHT diffusions were simulated and the accuracy of the estimators are evaluated [20]. Two sets of parameters, called set 1 and set 2, were used and is shown in Table 2.1. Table 2.2 contains values used for the setup of the simulations. To test the estimators for σ^2, α, and γ, the value of Λ is set to $\lambda = 1$. Then to test the estimator for Λ, its value is set to $\lambda = 3$ and the values of σ^2, α, and γ are assumed to be known.

The diffusion is approximated using the Euler-Maruyama method with step size $\tau = 0.0001$. The simulation stops when the value of the diffusion hits or surpasses the threshold. Any points above the threshold are removed. Sampling was simulated by taking every other κ samples. The values of κ used are in Table 2.2. The "sampling period" is given by $\Delta = \kappa\tau$.

Parameter	Set 1	Set 2
θ	100	100
α	-50	-100
γ	10	20
σ^2	5	25

Table 2.1: The values of the parameters for set 1 and set 2.

Parameter	Value
τ	0.0001
κ	$10, 50, 100, 150, 200, 300, 400, 500, 600, 700, 800, 900, 1000$
number of trials	1000

Table 2.2: The parameter set used in the simulations.

For each value of the sampling period, the experiment was repeated a thousand times. The sample mean and mean square error (MSE) for each estimator was calculated.

For the pMLE, define $\tilde{m}$ and $\tilde{n}$ as the fraction of the samples used in estimating α and γ, respectively. For the CDE, define $\tilde{\ell}$ as the fraction of samples used in estimating σ^2.

This simulations are intended to show how well the estimators work for a GIGHT diffusion. The parameters chosen for the simulations are not tied in anyway to biological neurons. Hence, the units in the simulations are arbitrary. Before the estimators are used in experimental data, testing on data from more biophysical simulations, such as ones using stochastic H-H models, is required.

2.4.1 Variance Estimators

Figure 2.6 shows the sample mean and MSE of the variance estimators for parameter sets 1 and 2. The sample mean of the simple estimator increases roughly linearly with Δ. Excluding the case where $\tilde{\ell} = 1$, the sample means of the CDE's are near the true value of σ^2 and do not change significantly with Δ. In other words, the CDE's are unbiased, except for $\tilde{\ell} = 1$. When $\tilde{\ell} = 1$, every sample is used in the estimate, including near the end of the diffusion, where the drift increases drastically. In the CDE, a constant drift assumption is used. By including points in the non-linear regime, the drift is overestimated when using the CDE. This introduces a bias to the CDE when $\tilde{\ell} = 1$. This bias increases with Δ, as shown in Figure 2.6. For the other values of $\tilde{\ell}$, the drift is approximately constant, so the estimators are approximately unbiased. The simple estimator is the most biased among the estimators here. This is expected since the simple estimator did not use any information on the drift of the diffusion.

As for the MSE of the estimators, the rate of increase is linear in Δ when plotted on a log-log scale for all of the estimators. Therefore, the MSE data points can be fitted with a power function described by $f(\Delta) = x(\Delta)^y$. The MSE data points for the CDE's, except for $\tilde{\ell} = 1$, are collinear with a slope of approximately 1. This strongly suggests that the MSE increases linearly with Δ. On the other hand, the slope is greater than 1 for the simple estimator and CDE with $\tilde{\ell} = 1$. The MSE grows supralinearly in Δ in these two cases. The use of the simple estimator is limited to small values of Δ due to large error and bias. The best estimator in terms of the MSE is the CDE with $\hat{\xi} = 0.75$ for both parameter sets. This suggests that the GIGHT diffusion has approximately constant drift for the first 75% of it for parameter sets 1 and 2.

Figure 2.6 supports the fact that the simple estimator and CDE's are consistent because from the figure, it appears that the sample mean and MSE approach the values of σ^2 and 0, respectively, for decreasing Δ.

2.4.2 Parameter Estimators

Since estimating α requires estimating σ^2, the CDE with $\tilde{\ell} = 0.75$ was used as the estimator for σ^2 for both parameter sets 1 and 2. This estimator was used because it had the least MSE for any Δ for both parameter sets.

The sample mean and MSE for the α estimators for parameter sets 1 and 2 are plotted in Figure 2.7. For the pMLE, the sample mean of the α estimator is below the true value of

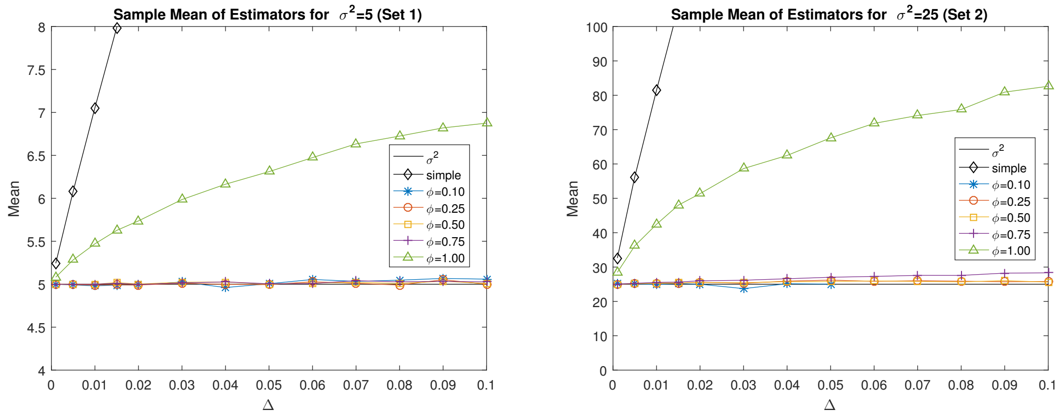

(a)

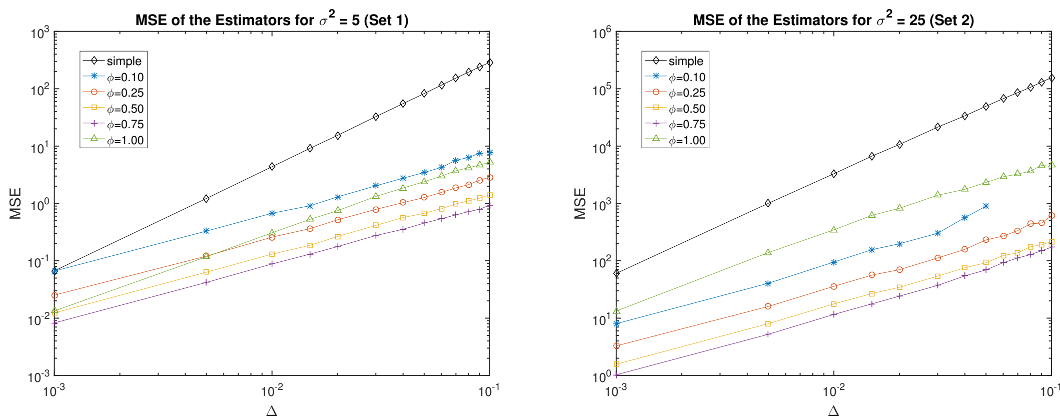

(b)

Figure 2.6: (A) The sample mean of the estimators of σ^2 is plotted against the sampling period. The sample mean of the simple estimator continues to increase beyond the plot in an approximately linear fashion. On the other hand, the CDE appears to be unbiased for $\tilde{\ell} = 0.1$, 0.25, 0.5, and 0.75. For $\tilde{\ell} = 1$, the mean of the CDE increases with Δ because the constant drift assumption is not applicable in this case. (B) The MSE of the estimators of σ^2 is plotted against the sampling period. For all of the estimators presented here, the data can be fitted with a power function, which is a straight line under log-log scaling. The higher slope of the simple estimator indicates a higher power and worse scaling with longer sampling periods. For the CDE's, with the exception of $\tilde{\ell} = 1$, the slope of the linear fit is approximately 1. This suggests that the MSE grows linearly with Δ. In parameter set 2, the pMLE for $\tilde{\ell} = 0.1$ cannot be determined past a certain value of Δ because with a low sampling period, not enough samples were taken to produce a proportion of 0.1.

α because it was assumed that the constant drift is non-existent. The estimated threshold-dependent drift must be higher to compensate for the lack of the constant drift. This corresponds to lower estimated values of α. The sample mean of the pMLE decreases with Δ. Likewise, the sample mean of the pLSFE also decreases with Δ, but at a slower rate. The sample mean of the pLSFE is closer to the true value of α, which indicates the pLSFE is less biased than the pMLE.

The MSE increases as Δ increases, which is expected. However, for the pMLE, it increases faster than that of the pLSFE. For low values of Δ, the pMLE is better than the pLSFE. As Δ increases, the pLFSE becomes the better estimator because its MSE grows more slowly than that of the other estimators.

With regards to the value of $\tilde{m}$, the MSE of the pMLE first improves by increasing $\tilde{m}$. Then as $\tilde{m}$ increases even further, the MSE worsens. There is a balance between using too few samples where there is not enough data points and using too many samples in the region where the zero constant drift assumption does not hold. This assumption is approximately true at the end of the diffusion where the barrier dependent drift swamps the constant drift. As more points farther away from the threshold are used, the estimator becomes more biased.

Note that for the pLSFE, there were a few instances where the computer program terminated early without solving for the minimum because the maximum number of iterations was reached. This was included in the results for the calculations of the mean and MSE, nonetheless. This affects the results for the pLSFE for both α and γ.

Likewise for γ, the CDE with $\tilde{\ell} = 0.75$ for the estimate of σ^2 was adopted for both parameter sets. The sample mean and MSE of the estimators of γ is plotted in Figure 2.8. For the pMLE, the sample mean is greater than the true value of γ. This is because threshold-dependent drift is assumed to be 0, so the estimated constant drift must be larger to compensate for the lack of the threshold-dependent drift. In the case of $\tilde{n} = 1$, the mean of the pMLE decreases slightly with Δ. This is because if $\tilde{n} = 1$, the last sample before the barrier is hit is used for the estimate. With larger sampling times, the last sample does not get as close to the threshold. This effect is amplified near the threshold since the effect of the attraction is more significant the closer the diffusion is to the threshold. For the other cases of the pMLE, the sample mean increases only slightly with Δ. The sample mean of the pLSFE does not change significantly with Δ for parameter set 1 and actually decreases

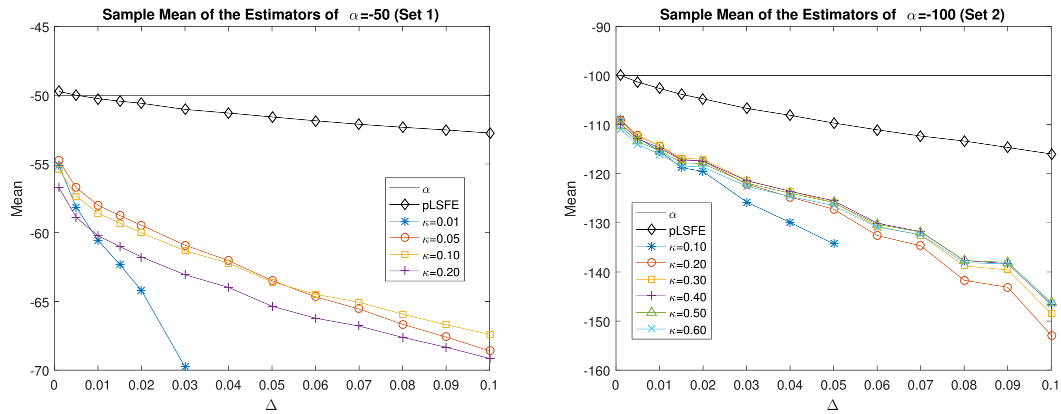

(a)

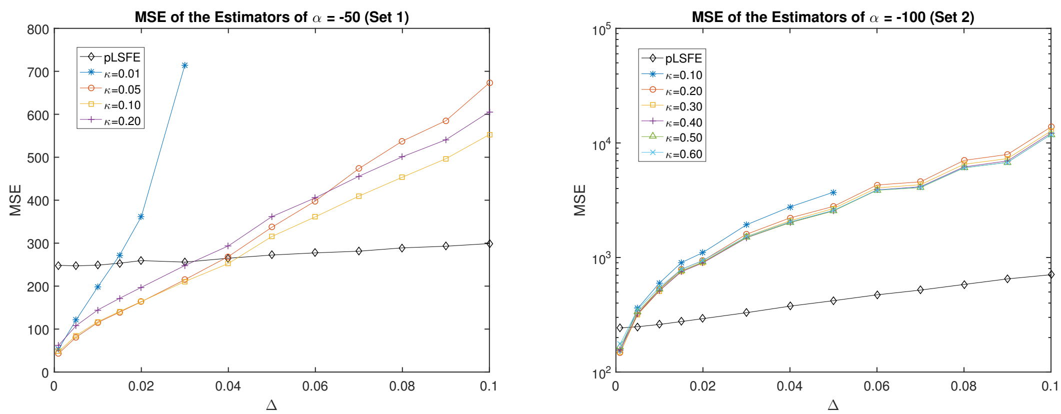

(b)

Figure 2.7: (A) The sample mean of the estimators of α is plotted against the sampling period. The sample mean decreases with Δ for all the estimators. The rate of decrease is higher for the pMLE. The pLSFE is less biased overall. (B) The MSE of the estimators for α is plotted against the sampling period. The MSE increases with Δ for all estimators. The MSE is initially lower for the pMLE, but increases and surpasses that of the pLSFE due to the higher rate of increase. Note that for parameter set 2, the scale of the ordinate is logarithmic. For a low values of $\tilde{m}$, the pMLE cannot be determined for large values of Δ. There were not enough samples to get a proportion of $\tilde{m}$ in these cases.

with Δ for parameter set 2. Its sample mean has approximately the same value as γ, which suggests that the pLSFE is approximately unbiased.

When $\tilde{n} = 1$, the MSE of the pMLE initially decreases with Δ, but then increases. This is because the estimator becomes less biased as the value of the sample mean decreases. However, the reduction of the MSE due to the unbiasedness is offset by the increase in the MSE from longer sampling periods and is eventually overcome by it. For the rest of the pMLE, the MSE increases with Δ. Likewise, the MSE of the pLSFE increases with Δ. When the bias of the pMLE is low, its MSE and that of the pLSFE are close. However, when the bias of the pMLE is high, the MSE of the pLSFE is lower. Thus, the bias of the pMLE increases the error significantly.

The value of $\tilde{n}$ that produces the least error here lies somewhere between 0.5 and 0.75, which suggests that the drift of the GIGHT diffusions for these two parameter sets is approximately constant for the first 50% to 75% of it. This agrees with the result of the CDE. Although it is constant, the threshold-dependent drift still has some contribution to the total drift, which adds to the bias in the estimator.

It is important to have large θ and μ_c for the pMLE; otherwise, the assumption that the drift is described as (2.15) in the beginning of the diffusion may not hold. This results in biased estimators such as in parameter set 2. In such a case, the pLFSE is a better choice for the estimator.

In these two parameter sets, it seems that the best choice of $\tilde{\ell}$ and $\tilde{n}$ is around 0.75. This is yet to be proven, but this suggests that the GIGHT diffusion behaves linearly for at least the first half of the diffusion. Further studies are needed but perhaps a number between 0.5 and 0.75 is a good choice for both $\tilde{\ell}$ and $\tilde{n}$. As for $\hat{m}$, the optimal choice is different for the two parameter sets.

What can be done is to anticipate the range of values the parameters σ^2, α, and γ can take and simulate many instances of the GIG diffusion. Then k-fold cross-validation can be used to find the value of $\tilde{\ell}$, $\tilde{m}$, and $\tilde{n}$. However, such anticipation is difficult at this point and the optimal selection is an open question.

2.4.3 Improved Estimators for the Parameters

The pMLE hinges on the approximation for O_α during parts of the GIGHT diffusion. However, if O_α can be approximated for the entire diffusion, a more accurate estimator can be developed. Furthermore, the hyperparameters $\tilde{\ell}$, $\tilde{m}$, and $\tilde{n}$ would not be necessary. This is especially a

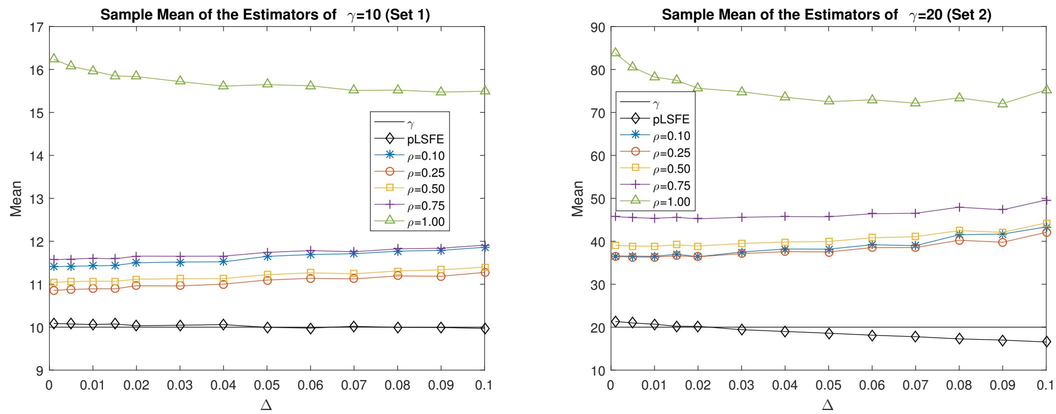

(a)

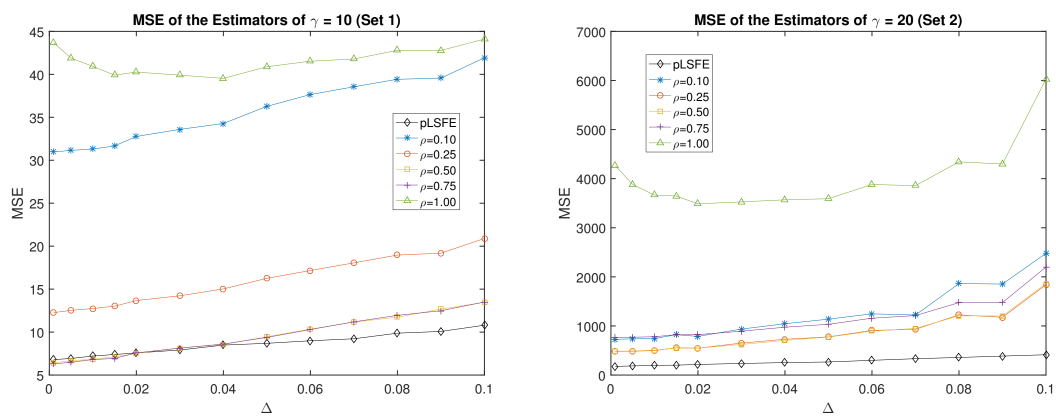

(b)

Figure 2.8: (A) The sample mean of the estimators of γ is plotted against the sampling period. For the pMLE, the mean is higher than the true value of γ. With $\tilde{n} = 1$, the sample mean decreases as a function of Δ. For the other values of $\tilde{n}$, the sample mean increases slightly with Δ. As for the pLSFE, the sample mean is constant for parameter set 1 and decreases slightly with Δ in parameter set 2. The pLSFE is less biased than the pMLE. (B) The MSE of the estimators of γ is plotted against the sampling period. For the pMLE with $\tilde{n} = 1$, the MSE initially decreases with Δ due to the decreasing mean. Then the MSE increases due to error with longer sampling periods. In all other cases, the MSE increases with Δ.

problem for γ. When the threshold is too low, the pMLE do not give an accurate estimate for γ because the drift approximation is not accurate. Since the behavior of $O_\alpha(x)$ is known for large x, it is more crucial to approximate its behavior as x approaches 0.

2.4.4 Simulations for the Input Intensity Estimator

Here, Λ was estimated with the approximate MLE. The parameters σ^2, α, and γ are assumed to be known. Let $\Lambda = 3$. The other simulation parameters are the same as in Table 2.2 and Table 2.1. Thus, the simulation was done for both parameter set 1 and set 2.

The sample mean and MSE of the estimator for Λ is plotted in Figure 2.9. The sample mean increases with Δ but with decreasing rate. The MSE is also small compared to the value of Λ, which suggests that this is a good estimator. The figure seems to indicate that MSE approaches 0 as Δ decreases. This suggests that the MLE estimator is consistent. The approximate MLE for set 1 produces less error on average than for set 2, though more research is needed to understand why.

Note that the accuracy of the estimator depends on accurate estimation of α, γ, and σ^2. In this simulation, the best case possible was assumed, i.e., the values of the parameters are known. Good parameter estimation is required in order to accurately estimate Λ.

2.5 Conclusion

Neuron η's PSP can be modeled by an SDE whose solution is the GIGHT diffusion. This diffusion can capture the upswing exhibited by η's PSP as it approaches the threshold. The GIGHT diffusion has the advantage of having an explicit form for the FHT distribution, which is given by the GIG distribution. This allows for analysis of the information-energy tradeoff of η, which is developed in Chapter 3.

In addition to developing the GIG neuron model, estimation of the parameters of such model from data was presented. The focus of this thesis was not on data collection from experiments. This requires more collaboration between information theorists and neuroscientists. Experiments are needed to see how well GIG diffusions match up to the η's PSP and to determine the values of the parameters of the model. A possible experiment to perform is to make intracellular and extracellular recordings of a neuron in V1 with the experimenters

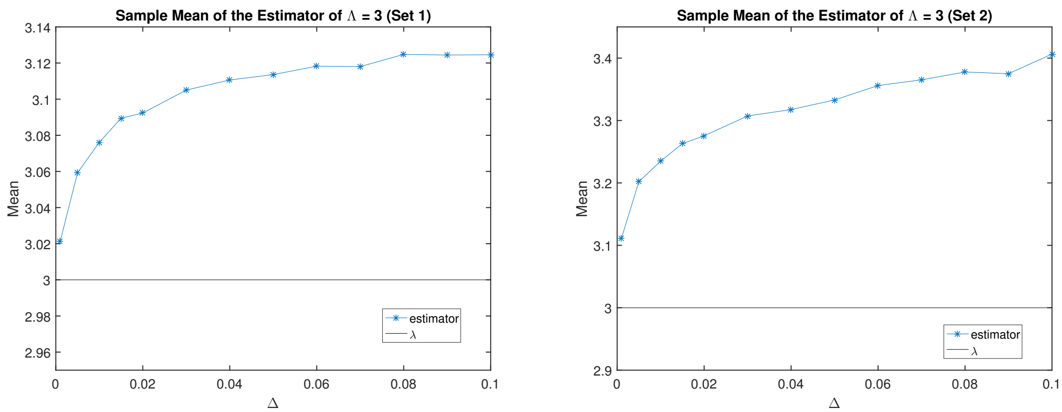

(a)

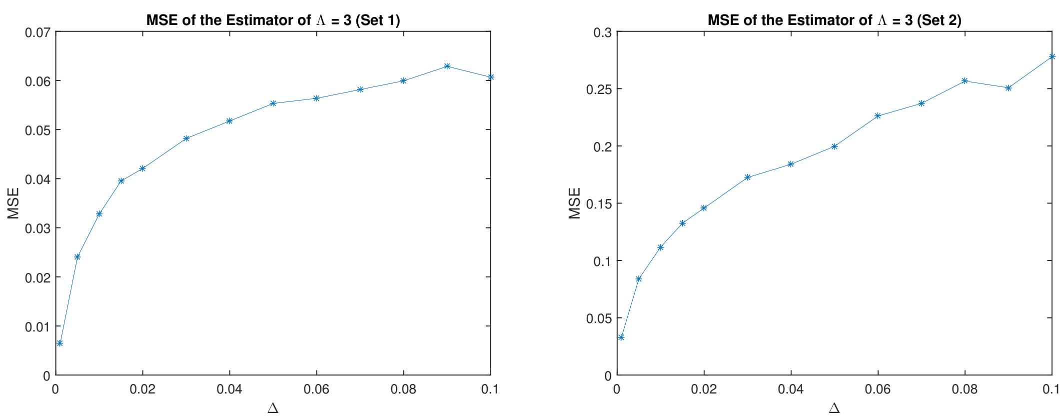

(b)

Figure 2.9: The sample mean (A) and MSE (B) of the Λ estimators is plotted against the sampling period. The black line indicates the true value of Λ. The sample mean and the MSE increases with Δ, but with decreasing rate. As $\Delta \to 0$, the sample mean approaches the true value and the MSE approaches 0, which indicates that the estimator is consistent. On average, the estimator for set 1 has less error than for set 2.

controlling its input intensity. To control the input to the neuron, find its afferent lateral geniculate nucleus (LGN) cells and transfect them with channelrhodopsin. This allows the LGN cells to be activated by light. Thus, the experimenters can use a laser scanning device to activate the transfected LGN cells to control the input intensity to the neuron. Then intracellular and extracellular recordings can be used to verify the GIG neuron model. This is a possible next step for research, which will encourage interdisciplinary approach in addressing energy-efficient neural computation.

Chapter 3

The Channel Capacity of the GIG Neuron Model

As mentioned in Chapter 1, neuron η seeks to maximize the mutual information (MI) it transmits given an average energy budget. In this chapter, η is modeled with the generalized inverse Gaussian (GIG) neuron model (see Chapter 2). This is equivalent to finding the capacity-cost (C-C) curve of the GIG neuron channel. From this curve, the tradeoff between energy and MI can be analyzed. Furthermore, the point of maximum MI per energy, i.e., bits per Joule, can be identified. It is given by the point on the curve whose tangent line passes through the origin [67]. This point of operation is the most efficient in the sense of the using the most MI per energy unit.

It turns out that the result of Smith is relevant to finding the constrained capacity of the GIG neuron channel [68]. In Smith, it was shown that for an additive white Gaussian noise channel (AWGN) and an amplitude-constrained input (or equivalently, a peak power-constrained input), the capacity-achieving input distribution is discrete. In other words, the set of input alphabet that achieves maximum MI is discrete and finite. The main contribution of this chapter is to show that the discrete result applies to the GIG neuron model as well.

This chapter begins with a brief introduction on MI and the constrained capacity, which will be referred to as just capacity henceforth. Then the energy model for the GIG neuron is presented. Finally, it is shown that for some parameters, the capacity for the energy neuron model is achieved by a discrete input distribution. The approach is measure-theoretic, which is necessary for some of the proofs in this chapter. However, proofs of theorems are shown in Section 5.4 rather than immediately after the respective theorems. Then implications of the results are discussed, followed by the conclusion.

3.1 Background

In this section, relevant notations and definitions, including MI, are introduced. Then, the capacity problem is defined. This is followed by the conditions for the input distribution in order to achieve the capacity.

3.1.1 Notation

Let F, F_1, and F_2 be cumulative density functions (CDF's) of a random variable (RV) or joint CDF's of a set of RV's. The unique probability measure associated with F is denoted by μ_F. The relationship $\mu_{F_1} \ll \mu_{F_2}$ denotes that μ_{F_1} is absolutely continuous with respect to μ_{F_2}. The relationship $\mu_{F_1} \equiv \mu_{F_2}$ denotes that $\mu_{F_1} \ll \mu_{F_2}$ and $\mu_{F_2} \ll \mu_{F_1}$. If $\mu_{F_1} \ll \mu_{F_2}$, the Radon-Nikodym (R-N) derivative [69] of μ_{F_1} with respect to μ_{F_2} is given by $\frac{\mathrm{d}F_1}{\mathrm{d}F_2}$.

All integrals in this chapter are of the Lebesgue-Stieltjes kind unless otherwise stated. Lebesgue-Stieltjes integrals are described in [69]. The symbol $\mathbb{R}$ denotes the set of real numbers and $\mathbb{R}^n$ is the n-fold product of $\mathbb{R}$. If the limits of integration are absent, it is assumed that the integral is over $\mathbb{R}^n$ for the appropriate n. The symbol $\mathbb{C}$ denotes the set of complex numbers and $\mathbb{C}^n$ is the n-fold product of $\mathbb{C}$.

3.1.2 Mutual Information

Let X and Y be a pair of RV's. Let the joint CDF of (X, Y) be F. Let the marginal CDF's of X and Y be v and w, respectively. Define $F^*(x, y) = v(x)w(y)$. The function F^* is the CDF of (X, Y) if X and Y are mutually independent while retaining their respective marginal CDF's. If $\mu_F \ll \mu_{F^*}$, the MI is given by [70]

$$I(X; Y) = \int_{\mathbb{R}^2} \log \frac{\mathrm{d}F}{\mathrm{d}F^*} \mathrm{d}F. \tag{3.1}$$

Otherwise, $I(X; Y) = \infty$.

Mutual information can also be described by the CDF of X and the conditional CDF of Y given X. Let them be described by v and W, respectively. Then MI can be expressed as

$$I(v, W) = \iint \log \frac{\mathrm{d}W(y|x)}{\mathrm{d}w(y; v, W)} \mathrm{d}W(y|x)\mathrm{d}v(x), \tag{3.2}$$

where,

$$w(y; v, W) = \int W(y|x)\mathrm{d}v(x) \tag{3.3}$$

is the marginal CDF of Y. Note that the function I has been "overloaded". When the arguments are separated by a semicolon, its arguments should be random variables. When separated by a comma, its respective arguments should be distributions. It can be shown that $\frac{\mathrm{d}F}{\mathrm{d}F^*} = \frac{\mathrm{d}W}{\mathrm{d}w}$, μ_F-a.e. (almost everywhere with respect to μ_F), which implies the equivalence of (3.1) to (3.2).

If the conditional CDF can be written as a conditional probability density function (PDF), then MI can be written in terms of the conditional PDF. Let Q be the conditional PDF associated with W. Then the MI is

$$I(v, W) = \iint \log \frac{Q(y|x)}{q(y; v, Q)} Q(y|x)\mathrm{d}y\mathrm{d}v(x), \tag{3.4}$$

where

$$q(y; v, Q) = \int Q(y|x)\mathrm{d}v(x) \tag{3.5}$$

is the marginal PDF of Y. The PDF q must exist because Q exists.

If the conditional PDF of Y given X exists, MI can also be written in terms of difference of entropy. The differential entropy of Y is

$$h_Y(v, Q) = -\int q(y; v, Q) \log q(y; v, Q)\mathrm{d}y. \tag{3.6}$$

The value $h_Y(v, Q)$ is not an entropy, i.e., it does not represent the uncertainty in the RV Y. In fact, the entropy of Y is not finite because it is a continuous RV [71]. However, it can be viewed as the entropy of Y subtracted by the entropy of the uniform distribution with density 1, hence the term differential entropy.

The conditional differential entropy of Y given X is

$$h_{Y|X}(v, Q) = -\iint Q(y|x) \log Q(y|x)\mathrm{d}y\mathrm{d}v(x). \tag{3.7}$$

This gives the average differential entropy of Y given that X has taken on a particular value. Again, this can be viewed as the conditional entropy of Y given X subtracted by the entropy of the uniform distribution with density 1. Then MI can be written as

$$I(v, Q) = h_Y(v, Q) - h_{Y|X}(v, Q). \tag{3.8}$$

Since the two terms are subtracted, the reference to the entropy of the uniform distribution cancels. Thus, MI is a difference of entropies. It is the reduction from $h_Y(v, Q)$ to $h_{Y|X}(v, Q)$. If entropies are interpreted as uncertainties, then $I(v, Q)$ is, on average, the reduction of the uncertainty in Y by learning the value of X.

3.1.3 The Constrained Capacity Problem

In capacity or C-C problems, the objective is to maximize $I(\cdot, W)$ over the space of CDF's of X subject to a cost constraint. Additional constraints on the allowable values of X may be imposed. Since W is fixed and known for C-C problems, the MI $I(v, W)$ will be written as $I(v)$ without W. If W has PDF Q, the entropies will be written as $h_Y(v)$ and $h_{Y|X}(v)$. Also, the marginal CDF of Y is written as $w(y; v)$.

Let $S \in \mathcal{B}(\mathbb{R})$, where $\mathcal{B}(\mathbb{R})$ is the Borel algebra of $\mathbb{R}$. Let $\mathcal{V}(S)$ be the set of CDF's of X whose points of increase is a subset of S. In other words, if $v \in \mathcal{V}(S)$, then $\mu_v(S) = 1$. This set will be indicated by just $\mathcal{V}$ when the set S is understood. Let g be the cost function and define the average cost operator on $v \in \mathcal{V}$ as

$$G(v) = \int g(x) \mathrm{d}v(x), \quad v \in \mathcal{V}. \tag{3.9}$$

Also, define the set

$$\mathcal{V}_E(S) = \{v \in \mathcal{V}(S) : G(v) \leq E\}, \tag{3.10}$$

i.e., the subset of $\mathcal{V}$ whose average cost is no greater than E. The C-C function is given by

$$C(E) = \sup_{v \in \mathcal{V}_E} I(v), \quad E \in \mathcal{E}, \tag{3.11}$$

where $\mathcal{E} = \{E : G(v) \leq E, v \in \mathcal{V}\}$, i.e., the set of possible average costs. Define $E_{\min} = \inf_{x \in S} g(x)$, i.e., the infimum of the cost function over S. Note that $E_{\min}$ is also the infimum of $\mathcal{E}$.

3.1.4 Solution of the Capacity-Cost Problem

Equip $\mathcal{V}$ with the weak topology of probability. It should be noted that this topological space is metrizable and is therefore a metric space [69]. Then from [72] and [68], the following theorem can be stated

Theorem 3.1. *If $\mathcal{V}_E$ is compact and I is continuous, then*

$$C(E) = \max_{v \in \mathcal{V}_E} I(v), \quad E \in \mathcal{E}, \tag{3.12}$$

and is achieved by some $v^* \in \mathcal{V}_E$. If I is strictly concave, then v^* uniquely achieves the constrained capacity. Furthermore, if $E > E_{min}$ and $C(E)$ is finite, v^* maximizes I if and only if for some value of $s \geq 0$,

$$i(x; v^*) - sg(x) \leq C(E) - sE, \quad x \in S, \tag{3.13}$$

$$i(x; v^*) - sg(x) = C(E) - sE, \quad x \in S_X. \tag{3.14}$$

The set S_X is the set of points of increase of v^* and the function i is the conditional information given that $X = x$, which given by

$$i(x; v) = \int \log \frac{\mathrm{d}W(y|x)}{\mathrm{d}w(y; v)} \mathrm{d}W(y|x). \tag{3.15}$$

Proof. See [72, Sec. 5.10 Th. 2] for existence proof. See [72, Sec. 8.3 Th. 1, Sec. 8.4 Th. 1] for the optimality conditions. See [68] for its application to capacity problems and uniqueness. □

The following theorem can be stated for the C-C curve:

Theorem 3.2. *The constrained capacity $C(E)$ is an increasing concave function of E. Furthermore, if C is differentiable at E, its derivative is given by s, i.e., $s = \frac{\mathrm{d}C}{\mathrm{d}E}$.*

Proof. See [72, Sec. 8.5 Th. 1]. □

With regards to the compactness of $\mathcal{V}_E$, the following theorem can be stated,

Theorem 3.3. *Let S be closed and g be continuous. If $g(x) \to \infty$ for $x \to -\infty$ and $x \to \infty$, then $\mathcal{V}_E(S)$ is compact in the weak topology of probability.*

The proof is presented in Section 3.4.1. Note that $g(x)$ can be redefined for $x \notin S$ and not change the solution to C-C problem. Hence, if S is bounded from below, g can be redefined so that $g(x) \to \infty$ as $x \to -\infty$. The same holds if S is bounded from above. This also implies that if S is closed and bounded (and hence compact) and g is continuous over S, $\mathcal{V}_E$ is compact in the weak topology of probability.

3.2 The Constrained Capacity of the GIG Neuron Channel

Recall that the input to the GIG neuron channel is given by the input intensity Λ, which is modeled as a RV. The output is given by the duration of the interpulse interval (IPI) T, which is also a RV. Recall that the conditional PDF of T given Λ is given by the GIG distribution,

$$Q_{\mathrm{GIG}}(t|\lambda) = M(\alpha,\beta,\gamma)^{-1}\lambda^{\alpha}t^{\alpha-1}\exp\left(-\frac{\beta}{\lambda t} - \gamma\lambda t\right), \quad \lambda, t > 0.$$

Since for the GIG neuron model, it is assumed that $\alpha \leq -1/2$. Hence, $\beta > 0$ and $\gamma \geq 0$. The normalization term is given by

$$M(\alpha,\beta,\gamma) = \begin{cases} 2\left(\frac{\beta}{\gamma}\right)^{\alpha/2} K_{\alpha}(2\sqrt{\beta\gamma}) & \gamma > 0 \\ \beta^{\alpha}\Gamma(-\alpha) & \gamma = 0. \end{cases}$$

The GIG neuron channel can be viewed as a multiplicative noise channel with $1/\Lambda$ as the input and T as the output. It is described by

$$T = \frac{U}{\Lambda}, \tag{3.16}$$

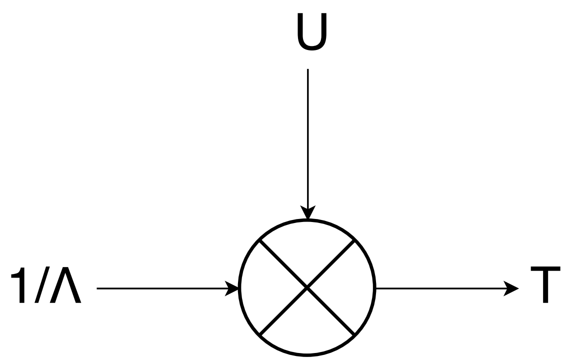

Figure 3.1: A representation of the GIG neuron channel.

where U is a RV independent of Λ. The RV U is the "noise" and is distributed as GIG(α, β,γ). So, the PDF of U is

$$f_U(u) = M(\alpha, \beta, \gamma)^{-1} u^{\alpha-1} \exp\left(-\frac{\beta}{u} - \gamma u \right). \tag{3.17}$$

A representation of this channel is shown in Figure 3.1.

3.2.1 Energy Model for Cortical Neurons

Let $g_{\text{total}}(\lambda, t)$ be the energy cost during an IPI for input intensity λ and output IPI duration t. According to the model, there are five major types of energy expenditures in a given IPI, denoted by the functions $g_i(\lambda, t),\ i = 1, \dots, 5$:

1. $g_1(\lambda, t) = z$, where $z > 0$. This term is associated with fixed energy costs such as creating and propagating an AP.

2. $g_2(\lambda, t) = ct$, where $c > 0$. This term is associated with energy costs that vary linearly in time, such as metabolic costs to maintain the neuron and keep it healthy.

3. $g_3(\lambda, t) = b/t$, where $b > 0$. This term increases the energy of two adjacent AP's that are separated by a short duration. This is because in the relative refractory period, it takes more energy to open enough channels to produce another AP (see Chapter 2).

4. $g_4(\lambda, t) = -a\log(t)$, where $a < 0$. This term is associated with maintaining neural "clocks." The clocks here are diffusive clocks where particles diffuse out of a compartment. The number of particles that left the compartment is approximately a logarithmic function and therefore should require an energy cost that is logarithmic to reset. Note that $g_4(\lambda, t) < 0$ for $t < 1$. However, this term is dominated by a positive $g_3(\lambda, t)$ as $t \to 0$.

5. $g_5(\lambda, t) = r\lambda t$, where $r > 0$. This is the cost of processing incoming spikes. For an IPI duration of t, the total intensity of input to η is given by the average intensity times the duration, i.e., λt.

The total energy is the sum of the five components,

$$g_{\text{total}}(\lambda, t) = z + b/t + ct - a\log t + r\lambda t. \tag{3.18}$$

It is useful to use the average cost given that $\Lambda = \lambda$, which is given by

$$g_\Lambda(\lambda) = \int g_{\text{total}}(\lambda, t) Q_{\text{GIG}}(t|\lambda)\mathrm{d}t, \quad \lambda > 0. \tag{3.19}$$

The mean of the cost functions g_{total} and g_Λ are equivalent. Therefore, for the purpose of analysis, these two are equivalent cost functions when Q_{GIG} is specified. From Appendix B, g_Λ is given as

$$g_\Lambda(\lambda) = z_\Lambda + b_\Lambda\lambda + c_\Lambda/\lambda + a_\Lambda\log\lambda, \quad \lambda > 0, \tag{3.20}$$

where,

$$z_\Lambda = z + k_1 r - k_g a, \tag{3.21}$$

$$a_\Lambda = a, \tag{3.22}$$

$$b_\Lambda = k_{-1}b, \tag{3.23}$$

$$c_\Lambda = k_1 c. \tag{3.24}$$

The terms k_{-1}, k_1, and k_g are constants that depend on α, β, and γ and are defined in Appendix B. The function g_Λ has a minimum at

$$\lambda^* = \frac{-a_\Lambda + \sqrt{a_\Lambda^2 + 4b_\Lambda c_\Lambda}}{2b_\Lambda}. \tag{3.25}$$

The minimum energy is then given by $E_{\min} = g_\Lambda(\lambda^*)$.

3.3 Finding the Constrained Capacity of the GIG Neuron Channel

To use Theorem 3.3, S must be closed. For the GIG neuron channel in its current form, $S = (0, \infty)$, which is not closed. However, by using a certain transformation on Λ and T, the resulting set S will be closed. Then Theorems 3.1 and 3.3 can be applied to find the constrained capacity for the GIG neuron channel.

3.3.1 Representing the Input and Output Random Variables

Define the RV's $X = -\log \Lambda$, $Y = \log T$, and $N = \log U$. The new RV's are transformations on the input, output, and noise RV's, respectively, of the GIG neuron channel. A new channel is formed by the new RV's, which is given by

$$Y = X + N, \tag{3.26}$$

where N is independent of X. The PDF of the noise is given by

$$f_N(n) = M(\alpha, \beta, \gamma)^{-1} \exp\left(\alpha n - \beta e^{-n} - \gamma e^{n}\right). \tag{3.27}$$

This type of distribution is referred to as an exponentiated GIG (EGIG) distribution, which is denoted by EGIG(α, β, γ).

Thus, the multiplicative noise channel (3.16) is converted into an additive noise channel. The conditional PDF that describes the channel is

$$Q(y|x) = M(\alpha, \beta, \gamma)^{-1} \exp\left(\alpha(y - x) - \beta e^{-(y-x)} - \gamma e^{y-x}\right). \tag{3.28}$$

Since the RV X is onto $\mathbb{R}$, $S = \mathbb{R}$, which is a closed set.

The cost function for input $X = x$ is given by $g(x) = g_\Lambda(e^{-x})$, which is equavalent to

$$g(x) = z_\Lambda - a_\Lambda x + b_\Lambda e^{-x} + c_\Lambda e^{x}. \tag{3.29}$$

The minimum of g occurs at $x^* = \exp(-\lambda^*)$ and the value of the minimum cost is $E_{\min} = g(x^*) = g_\Lambda(\lambda^*)$.

3.3.2 Existence and Uniqueness of Solution

The conditions in Theorem 3.1 must be satisfied to guarantee the existence of a solution. One of the conditions is that $\mathcal{V}_E(S)$ is compact. The function g is continuous and $g \to \infty$ for both $x \to -\infty$ and $x \to \infty$. The set $S = \mathbb{R}$ is closed. Therefore, $\mathcal{V}_E(S)$ is compact by Theorem 3.3.

The continuity of I is given by the following theorem:

Theorem 3.4. *For the GIG neuron channel* (3.28)*, the mutual information I is continuous over $\mathcal{V}_E$, $E \in \mathcal{E}$..*

The conditions of Theorem 3.1 is met. Therefore, a solution exists.

To show that the solution is unique, it must be shown that I is strictly concave. the MI I can be expressed as (3.8), i.e., a difference of two entropies, because the conditional PDF of Y given X exists. Since $h_{Y|X}$ is a linear function, it suffices to show that h_Y is strictly concave.

Define $v_\delta = \delta v_1 + (1-\delta)v_2$, where $0 < \delta < 1$, and $v_1, v_2 \in \mathcal{V}$. Since $x \log x$ is strictly convex in x, for a given value of t,

$$\begin{aligned} q(y; v_\delta) \log q(y; v_\delta) &= (\delta q(y; v_1) + (1-\delta) q(y; v_2)) \log(\delta q(y; v_1) + (1-\delta) q(y; v_2)) \\ &\leq \delta q(y; v_1) \log q(y; v_1) + (1-\delta) q(y; v_2) \log q(y; v_2), \end{aligned} \tag{3.30}$$

with equality if and only if $q(y; v_1) = q(y; v_2)$. Hence $h_Y(v_\delta) \geq \delta h_Y(v_1) + (1-\delta)h_Y(v_2)$ with equality if on only if $q(y; v_1) = q(y; v_2)$, for $y \in \mathbb{R}$ (necessity is due to continuity of $q(y; v)$ in y for any $v \in \mathcal{V}$). If it is shown that $q(\cdot; v)$ is an injective (one-to-one) transform on v, h_Y must be strictly concave because $h_Y(v_\delta) = \delta h_Y(v_1) + (1-\delta)h_Y(v_2)$ if and only if $v_1 = v_2$. Suppose that $v_1 \neq v_2$ but $q(y; v_1) = q(y; v_2)$ for all $y \in \mathbb{R}$. Then $q(y; v_1) - q(y; v_2) = 0$, i.e.,

$$\int Q(y|x) \mathrm{d}(v_1(x) - v_2(x)) = 0. \tag{3.31}$$

Let $v_- = v_1 - v_2$. Multiply both sides by $e^{-\zeta y}$ for some complex ζ, and integrate with respect to y, which yields

$$\iint \exp\left((\alpha - \zeta)(y - x) - \beta e^{-(y-x)} - \gamma e^{y-x}\right) e^{-\zeta x} \mathrm{d}v_-(x) \mathrm{d}y = 0, \tag{3.32}$$

where Q, defined in (3.28), has been substituted into the expression. By Fubini's theorem, the order of integration can be changed. From a change in variables and Appendix A.1, the inner integral can be evaluated, which yields

$$\int 2(\beta/\gamma)^{(\alpha-\zeta)/2} K_{\alpha-\zeta}(2\sqrt{\beta\gamma}) e^{-\zeta x} \mathrm{d}v_-(x) = 0. \tag{3.33}$$

Since $K_{\alpha-\zeta}(\sqrt{\beta\gamma}) \neq 0$ for real values of $\sqrt{\beta\gamma}$ [73],

$$\int e^{-\zeta x} \mathrm{d}v_-(x) = 0. \tag{3.34}$$

Equation (3.34) is a Laplace transform, which is invertible. The inverse is given by $v_-(x) = 0$ for all $x \in \mathbb{R}$. Hence, $v_1 = v_2$, which is a contradiction. As a consequence, $q(\cdot; v)$ is an injective transform on v. Thus, I is strictly concave.

Thus, the capacity achieving input distribution is unique.

3.3.3 The Constrained Capacity at the Point of Least Energy

If the energy cost $E = E_{\min}$, the second half of Theorem 3.1 is not applicable. However, the constrained capacity is still achievable by a unique input CDF. Thus, this case is handled separately in this section [21].

Since the energy cost is minimized uniquely by the input x^*, the only input CDF that achieves $E_{\min}$ is

$$v^*(x) = \Theta(x - x^*), \tag{3.35}$$

where Θ is the unit step function. Since this is the only element in $\mathcal{V}_E$, this distribution achieves the minimum of I. The output PDF is given by

$$q(y; v^*) = M(\alpha, \beta e^{x^*}, \gamma e^{-x^*})^{-1} \exp\left(\alpha(y - x^*) - \beta e^{-(y-x^*)} - \gamma e^{y-x^*}\right) \tag{3.36}$$

In fact, since $Q(y|x) = q(y; v^*)$, μ_{v^*}-a.e., $I(v^*) = 0$, which is the minimum achievable MI. Hence, $C(E_{\min}) = 0$.

Of interest is also the rate of change of C at $E_{\min}$ as E increases. Let this rate be denoted by s_0. Note that since C is concave, s_0 is also the maximum rate of change of $C(E)$. From

[74], the maximum rate of change is given as

$$s_0 = \sup_{x \neq x^*} \frac{i(x; v^*)}{g(x) - E_{\min}}. \tag{3.37}$$

Thus, the following must be satisfied for all $x \in \mathbb{R}$:

$$i(x; v^*) - s_0 g(x) \leq -s_0 E_{\min} \tag{3.38}$$

For v^*, the conditional information given that $X = x$ is

$$i(x; v^*) = \iint Q(y|x) \log \frac{Q(y|x)}{Q(y|x^*)} \mathrm{d}y \tag{3.39}$$
$$= -\alpha(x - x^*) + k_{-1}\beta e^{-(x-x^*)} + k_1\gamma e^{x-x^*} - (\beta k_{-1} + \gamma k_1). \tag{3.40}$$

This has similar terms to g and (3.38) can be expressed as

$$-a_r x + b_r e^{-x} + c_r e^{x} + z_r \leq 0, \tag{3.41}$$

where,

$$a_r = \alpha - s_0 a_\Lambda, \tag{3.42}$$
$$b_r = k_{-1}\beta e^{x^*} - s_0 b_\Lambda, \tag{3.43}$$
$$c_r = k_1\gamma e^{-x^*} - s_0 c_\Lambda, \tag{3.44}$$
$$z_r = a_r x^* - b_r e^{-x^*} - c_r e^{x^*}. \tag{3.45}$$

The values of a_r, b_r, c_r, and z_r must be so that (3.41) holds. Note that the LHS of (3.41) is 0 when evaluated at x^*. Thus, its maximum must be achieved at x^*. For $a_r > 0$, it must be that $b_r < 0$ and $c_r \leq 0$. Otherwise, the LHS of (3.41) is unbounded as $x \to \infty$ or Likewise, if $a_r < 0$, it must be that $c_r < 0$ and $b_r \leq 0$. For $a_r = 0$, it must be that $b_r < 0$ and $c_r < 0$ or

$b_r = c_r = 0$. Otherwise, the maximum will not be achieved at x^*, which means the LHS of (3.41) can become positive. This puts bounds on s_0,

$$s_0 \geq \frac{k_{-1}\beta e^{x^*}}{b_\Lambda} \text{ and } s_0 \geq \frac{k_1 \gamma e^{-x^*}}{c_\Lambda}. \tag{3.46}$$

Since s_0 is a supremum, the value of s_0 is the minimum possible value imposed by the two bounds. Hence,

$$s_0 = \max\left\{\frac{k_{-1}\beta e^{x^*}}{b_\Lambda}, \frac{k_1 \gamma e^{-x^*}}{c_\Lambda}\right\}. \tag{3.47}$$

Written in another way,

$$s_0 = \max\left\{\frac{\beta}{b\lambda^*}, \frac{\gamma\lambda^*}{c}\right\}. \tag{3.48}$$

Thus, $0 \leq s \leq s_0$.

3.3.4 The Constrained Capacity for the General Energy Cost

There are three exhaustive possibilities for S_X, i.e., the points of increase for v^*:

1. The set S_X is discrete and infinite, but with finite number of points in any finite interval.
2. The set S_X contains an infinite number of points in some finite interval. This includes the case where S_X contains an interval.
3. The set S_X is discrete and finite.

Recall that a unique solution must exist. Therefore, one of these possibilities must be true. Possibilities 1 and 2 will be proven false which leaves only possibility 3.

Disproving Possibility 1

The presented argument for disproving possibility 1 is modified from [75]. Suppose possibility 1 is true. Then S_X must be countable and its elements can be indexed by whole numbers.

Then

$$q(y; v^*) = \sum_{k=1}^{\infty} p_k Q(y|x_k), \quad k = 1, 2, \ldots, \tag{3.49}$$

where x_k, $k = 1, 2, \ldots$ is an element of S_X with probability p_k. Since for any k and y, $p_k Q(y|x_k) > 0$, taking one addend from the right-hand side (RHS) yields

$$q(y; v^*) > p_k Q(y|x_k). \tag{3.50}$$

Taking the logarithm of both sides, multiplying by $Q(y|x)$ and integrating with respect to y yield

$$\int Q(y|x) \log q(y; v^*) \mathrm{d}y > \int Q(y|x) \log(p_k Q(y|x_k)) \mathrm{d}y. \tag{3.51}$$

Then, add $-\int Q(y|x) \log Q(y|x) \mathrm{d}y + sg(x) + C(E) - sE$ to both sides. Evaluating the RHS yields

$$(C(E) - sE) - (i(x; v_0) - sg(x)) > -(sa_\Lambda - \alpha)x + (sb_\Lambda - \frac{\beta k_{-1}}{e^{-x_k}})e^{-x} + \frac{sc_\Lambda - \gamma k_1 e^{-x_k}}{e^{-x}} + k_{\text{cons}}, \tag{3.52}$$

where

$$k_{\text{cons}} = -\alpha x_k + \log p_k + \beta k_{-1} + \gamma k_1 + C(E) - sE. \tag{3.53}$$

It must be that $\{x_k\}$ is unbounded from above or below, otherwise, possibility 1 is untrue. Assume that it is unbounded from above. As x gets large, the e^x term dominates:

$$(C(E) - sE) - (i(x; v_0) - sg(x)) > (sc_\Lambda - \gamma k_1 e^{-x_k})e^x + o(e^x). \tag{3.54}$$

It must be that $s > 0$. If $s = 0$, the optimality conditions (3.13) and (3.14) in Theorem 3.1 is the same as for an unconstrained problem. However, the unconstrained problem does not have a solution. Hence, $s > 0$. Therefore, for any value of $s > 0$, there is a value of k such that $sc_\Lambda - \gamma k_1 e^{-x_k} > 0$. Since e^x is the dominating term in the RHS of (3.54), it increases without bound as $x \to \infty$. This violates the optimality condition (3.13) since it must be satisfied for any value of k. Hence, a contradiction is reached.

Now assume that $\{x_k\}$ is unbounded from below. As x gets further below zero,

$$(C(E) - sE) - (i(x; v_0) - sg(x)) > (sb_\Lambda - \beta k_{-1} e^{x_k})e^{-x} + o(e^{-x}). \tag{3.55}$$

The same argument can be applied here and another contradiction is reached. Hence, possibility 1 cannot happen.

Evaluating Possibility 2

Assume possibility 2 is true. Recall that for $v \in \mathcal{V}$, the conditional information of Y given that $X = x$ is

$$i(x; v) = \int Q(y|x) \log \frac{Q(y|x)}{q(y; v)} \mathrm{d}y. \tag{3.56}$$

Consider the optimality condition (3.14) of Theorem 3.1,

$$i(x; v^*) - sg(x) = C(E) - sE, \quad x \in S_X.$$

The left-hand side (LHS) is well defined for complex x where the principle branch of the logarithm is used. Thus, let ζ be a complex number and let Z_X be the set S_X embedded into the complex plane and the following is attained:

$$i(\zeta; v^*) - sg(\zeta) = C(E) - sE, \quad \zeta \in Z_X. \tag{3.57}$$

Note that the LHS is an analytic function of ζ.

Since S_X, and thus Z_X, has an infinite number of points in some bounded interval, Z_X has a limit point in the complex plane $\mathbb{C}$. Since (3.57) holds for a set of points with a limit point, by the identity theorem [76], (3.57) holds for all points in $\mathbb{C}$. Thus, for all $x \in \mathbb{R}$,

$$i(x; v^*) - sg(x) = C(E) - sE. \tag{3.58}$$

In [14], it was shown that the output PDF $q(\cdot; v^*)$ that satisfies (3.58) is another exponentiated GIG distribution:

$$q(y; v^*) = M(a_s, b_s, c_s)^{-1} \exp(a_s y - b_s e^{-y} - c_s e^y), \tag{3.59}$$

where $a_s = as$, $b_s = bs$, and $c_s = cs$. To show that this PDF uniquely satisfies (3.58), note that (3.58) can be rewritten as

$$\int Q(y|x) \log q(y; v^*) \mathrm{d}y = -sg(x) + k'_{\text{cons}}, \tag{3.60}$$

where $k'_{\text{cons}} = C(E) - sE + \alpha k_g - \beta k_{-1} - \gamma k_1 - \log M(\alpha, \beta, \gamma)$. It can be verified that (3.59) satisfies the equality above. The LHS is a transform on $\log q(\cdot; v^*)$. From Section 3.3.2, it was shown that this transform is invertible. Therefore, (3.59) satisfies the equation uniquely.

From (3.59), the input PDF v^* can be attained. In [14], the characteristic function (CF) of X was determined. In theory, the PDF of X can then be recovered from the CF. Since $Y = X + N$ and X and N are independent, the relationship among the respective CF's is

$$\rho_Y(\omega) = \rho_X(\omega)\rho_N(\omega), \tag{3.61}$$

where ρ_Y, ρ_X, and ρ_N are the CF's for Y, X, and N, respectively. Recall that Y is the exponentiated GIG distribution. Its CF is given by

$$\rho_Y(\omega) = M(a_s, b_s, c_s)^{-1} \int \exp\left(j\omega y + a_s y - b_s e^{-y} - c_s e^y\right) \mathrm{d}y \tag{3.62}$$

$$= \frac{M(a_s + j\omega, b_s, c_s)}{M(a_s, b_s, c_s)}. \tag{3.63}$$

This integral can be evaluated by a change in variable and Appendix A.1. Since N is also distributed as an exponentiated GIG, its CF can be attained in a similar way, which yields

$$\rho_N(\omega) = \frac{M(\alpha + j\omega, \beta, \gamma)}{M(\alpha, \beta, \gamma)}. \tag{3.64}$$

The CF of X is then the ratio of the CF's of Y and N, which yields

$$\rho_X(\omega) = \frac{M(\alpha, \beta, \gamma)}{M(a_s, b_s, c_s)} \frac{M(a_s + j\omega, b_s, c_s)}{M(\alpha + j\omega, \beta, \gamma)}. \tag{3.65}$$

It turns out that this may not be a valid CF. The inverse of the CF may be a function with negative values. In this case, possibility 2 cannot happen and only possibility 3 is left, which implies that the input distribution is discrete with finite masses.

When the Input Distribution is Discrete

Recall that for the GIG neuron channel, $\alpha < 0$. For $\gamma = 0$ and $c = 0$, the input distribution has been determined in [11]. In this case, the CF of X is given by

$$\rho_X(\omega) = \frac{\Gamma(-\alpha)}{\Gamma(-a_s)} \frac{\Gamma(-a_s - j\omega)}{\Gamma(-\alpha - j\omega)} \left(\frac{b_s}{\beta}\right)^{j\omega} = \frac{B(-\alpha + a_s, -a_s - j\omega)}{B(-\alpha + a_s, -a_s)} \left(\frac{b_s}{\beta}\right)^{j\omega}. \tag{3.66}$$

The associated PDF is given by an exponentiated beta distribution:

$$p(x) = \frac{e^{a_s(x-x_0)}\left(1 - e^{-(x-x_0)}\right)^{-\alpha + a_s - 1}}{B(-a_s, -\alpha + a_s)}, \quad x > x_0, \tag{3.67}$$

where $x_0 = \log(b_s/\beta)$ is the shift. This is a valid PDF that satisfies the optimality conditions (3.13) and (3.14). So for the case $\gamma = 0$ and $c = 0$, the input distribution is continuous.

Now, let $\gamma = 0$, but $c > 0$. In this case, it is easier to work with the RV's Λ, U, and T. Since Y is an exponentiated GIG RV, T is a GIG RV whose PDF is given by

$$q_T(t) = M(a_s, b_s, c_s)^{-1} t^{a_s - 1} \exp(-b_s/t - c_s t), \quad t > 0. \tag{3.68}$$

Let p_Λ be the PDF of Λ. Then for $t > 0$,

$$q_T(t) = \int Q_{\mathrm{GIG}}(t|\lambda)p_\Lambda(\lambda)\mathrm{d}\lambda \tag{3.69}$$

$$= \int \beta^{-\alpha}\lambda^{\alpha}t^{\alpha-1}\exp\Big(-\frac{\beta}{\lambda t}\Big)p_\Lambda(\lambda)\mathrm{d}\lambda. \tag{3.70}$$

Let $\lambda' = \beta/\lambda$ and $t' = 1/t$. Then

$$\int (\lambda')^{-\alpha-2}t'^{-\alpha+1}\exp(-\lambda' t')p_\Lambda(\beta/\lambda')\mathrm{d}\lambda' = q_T(1/t';\hat{v}_0), \quad t' > 0. \tag{3.71}$$

The expression on the LHS is well defined for a complex λ'. For complex λ', the LHS is a Laplace transform. Taking the inverse transform of both sides yield [77]

$$p_\Lambda(\beta/\lambda') = \lambda'^{\alpha+2}(\lambda' - b_s)^{(a_s-\alpha-1)/2}J_{a_s-\alpha-1}\big(2\sqrt{c_s(\lambda'-b_s)}\big), \quad \lambda' > b_s, \tag{3.72}$$

where J_α is the Bessel function of the first kind of order α. However, J_α for $\alpha \in \mathbb{R}$ has some negative values, which disqualifies p_Λ, and therefore p, as a viable PDF. However, p_Λ is the only solution that satisfies (3.58). This is a contradiction, which implies that possibility 2 cannot occur in this case. Hence, possibility 3 is the only one that can occur. Therefore, the input distribution must be finite and discrete for $\gamma = 0$ and $c > 0$. The value and probability of the atoms are not known analytically, but can be determined numerically.

For $\gamma > 0$ and $c > 0$, (3.65) can be written as

$$\rho_X(\omega) = \frac{K_\alpha(\sqrt{\beta\gamma})}{K_{a_s}(\sqrt{b_s c_s})}\frac{K_{a_s+j\omega}(\sqrt{b_s c_s})}{K_{\alpha+j\omega}(\sqrt{\beta\gamma})}\left(\frac{\gamma b_s}{\beta c_s}\right)^{\frac{j\omega}{2}}. \tag{3.73}$$

The inverse of the of the CF is given by its Fourier transform:

$$p(x) = \frac{1}{2\pi}\int \rho_X(\omega)e^{-j\omega x}\mathrm{d}\omega. \tag{3.74}$$

To the best of the author's knowledge, there is no known analytical expression for the integral. Thus, such integrals must be calculated numerically. Alternatively, the discrete Fourier Transform (DFT) algorithms can be used to approximate the function p.

In numerical analysis, it turns out that p can be either non-negative or have negative values, depending on the parameters α, β, γ, a, b, c, and the energy constraint E [67]. In the case that p is non-negative, it is the capacity-achieving input distribution. However, if p has negative values, then possibility 2 cannot happen. In such a case, p is finite and discrete. The values and probabilities of the masses can be determined numerically.

An analytical way of determining whether (3.73) is a valid CF of a RV is to determine whether (3.73) is positive definite. For ρ_X to be positive definite, the matrix formed by any $\omega_i \in \mathbb{R}$, $i = 1, \ldots, K$,

$$[\rho_X(\omega_i - \omega_j)], \tag{3.75}$$

is positive semidefinite [78]. Let this matrix be M_X. The matrix M_X is positive semidefinite if for any complex vector $\underline{\zeta} \in \mathbb{C}^K$,

$$\underline{\zeta}^\dagger M_X \underline{\zeta} \geq 0, \tag{3.76}$$

where $\underline{\zeta}^\dagger$ is the conjugate transpose of $\underline{\zeta}$. Then Bochner's theorem states that p is a non-negative function. However, this is difficult to determine analytically, but can serve as a possible check of whether the inverse of (3.73) is a non-negative function.

3.3.5 Upper Bound of the Capacity-Cost Curve

Equation (3.74) can be used to obtain an upper bound on the C-C curve. It is the solution to the constrained capacity problem if probability is allowed to be negative [14]. Suppose p

is given by (3.74). For $C > 0$, the energy is given by [14]

$$E_b = z + rk_1 + \sqrt{bc}\left(\frac{K_{a_s+1}(2\sqrt{b_s c_s}) + K_{a_s-1}(2\sqrt{b_s c_s})}{K_{a_s}(2\sqrt{b_s c_s})}\right) \\ -a\left(\frac{\frac{\partial}{\partial u}[K_u(2\sqrt{b_s c_s})]_{u=a_s}}{K_{a_s}(\sqrt{b_s c_s})} + \frac{1}{2}\log\frac{b_s}{c_s}\right). \tag{3.77}$$

Using the recurrence relationship of Appendix A.4, this expression can be simplified to

$$E_b = z + rk_1 - \frac{\partial}{\partial s}\log\left[\left(\frac{b_s}{c_s}\right)^{a_s/2} K_{a_s}(2\sqrt{b_s c_s})\right]. \tag{3.78}$$

In terms of the function M, this is given by

$$E_b = z + rk_1 - \frac{\partial}{\partial s}\log M(a_s, b_s, c_s). \tag{3.79}$$

It can be shown that the expression is also valid for $c = 0$.

The upper bound of the constrained capacity when $b, c > 0$ and $\beta, \gamma > 0$ was determined in [14]. In [20], it was simplifed to

$$C_b = \log\frac{K_{a_s}(2\sqrt{b_s c_s})}{K_\alpha(2\sqrt{\beta\gamma})} + \frac{\partial}{\partial u}\left[\log\frac{K_{\alpha u}(2u\sqrt{\beta\gamma})}{K_{a_s u}(2u\sqrt{b_s c_s})}\right]_{u=1}. \tag{3.80}$$

This derivation is also presented in Section 3.4.3. In terms of the function M, this is given by

$$C_b = \log\frac{M(a_s, b_s, c_s)}{M(\alpha, \beta, \gamma)} + \frac{\partial}{\partial u}\left[\log\frac{M(\alpha u, \beta u, \gamma u)}{M(a_s u, b_s u, c_s u)}\right]_{u=1}. \tag{3.81}$$

It can also be shown that this expression is valid for $c = 0$ or $\gamma = 0$.

The graph (E_b, C_b) for all values of $s > 0$ describes the upper bound of the constrained capacity for the GIG neuron channel. This bound is tight if the input distribution to the GIG neuron channel is continuous, since p would be a non-negative function. Hence, this bound is tight for $\gamma = 0$ and $c = 0$, simultaneously.

3.3.6 Numerical Examples

First, define the RV:

$$X_c = \frac{\exp(-X)}{\exp(-X) + \exp(-x^*)}, \tag{3.82}$$

which is invertible. By using X_c as the input variable, the input space is compressed from $(-\infty, \infty)$ to $(0, 1)$. Note that $I(X;Y) = I(X_c;Y)$, so the constrained capacity is unchanged by this substitution. The value $X = x^*$ corresponds with $X_c = 1/2$, so the energy is minimized at this point. This is desirable because the distribution of X_c should center around the point of least energy, $x = 1/2$, so that minimal energy is used. This makes it relatively easy to inspect the distribution visually. The space of X_c was discretized and a numerical solver in MATLAB was used to obtain the optimal input distribution. The Gauss-Laguerre quadrature was used to approximate integrations.

Figure 3.2 shows the constrained capacity curve for a certain parameter set along with the input distribution for certain energy budgets. In this case, the actual curve is close to the upper bound. Indeed, for some points on the curve, the optimal input distribution appears to be continuous. In such a case, the upper bound and the curve coincide. For energy use above a certain level, the optimal input distribution appears continuous.

However, it cannot be said with absolute certainty that the presented continuous input is actually continuous. Using the DFT to approximate (3.74) yields a solution with a minimum negative value on the order of 10^{-10}. It is difficult to ascertain whether this is due to machine error or is part of the actual solution. Furthermore, Even though using a numerical solver seems to yield a continuous answer, at times, pushing for higher numerical accuracy yields a discrete result. This did not happen for the example in Figure 3.2. However, pushing the accuracy even further may possibly produce discrete results. There is a caveat, however. Since the space was discretized, the obtained solution is the solution to an approximation of the problem. It could be that the solution to the approximation of the problem is discrete whereas the solution to the actual problem is continuous. Thus, increasing the accuracy of

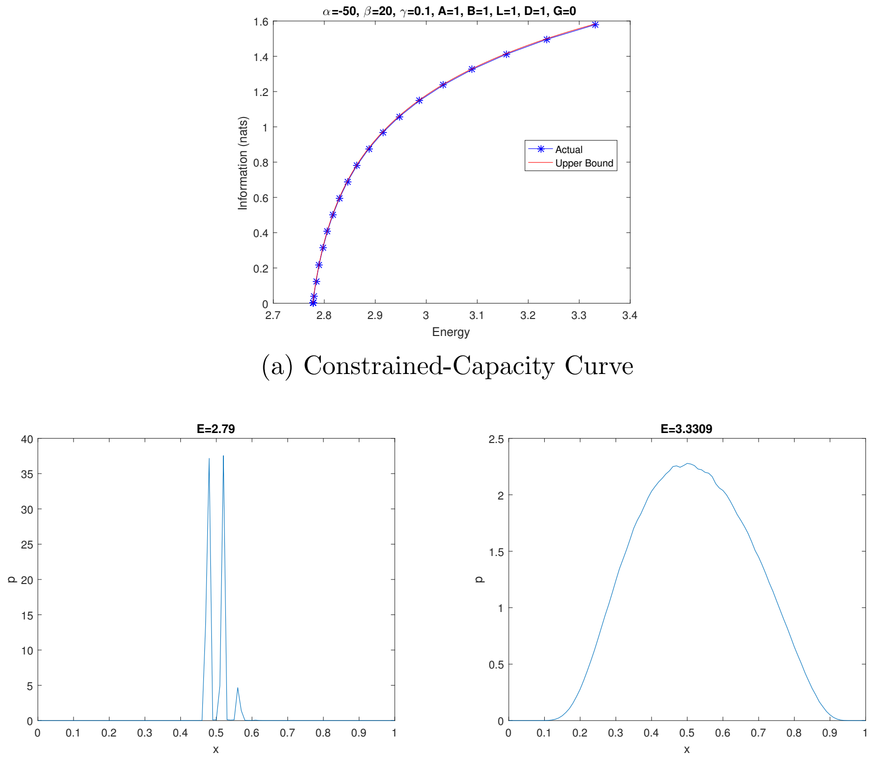

(a) Constrained-Capacity Curve

(b) Input Distribution for $\mathcal{E} = 2.79$ (c) Input Distribution for $\mathcal{E} = 3.33$

Figure 3.2: (A) The constrained capacity curve of the GIG neuron channel with parameters $(\alpha, \beta, \gamma) = (-50, 20, 0.1)$ and energy parameters $(z, a, b, c, r) = (1, 1, 1, 1, 0)$. The actual curve follows the upper bound closely. The mutual information is given in nats, while the energy unit is arbitrary. (B) The approximate input distribution when $E = 2.79$. Probability density is plotted against $X_c = x$. The input distribution appears to be discrete. (C) The approximate input distribution when $E = 3.33$. Probability density is plotted against $X_c = x$. The input distribution appears to be continuous.

the numerical solver will cause the more correct continuous solution to a discrete solution. However, if the problem is well approximated, the difference in the constrained capacity would be small, though with very different input distributions: one is discrete and the other is continuous.

Figure 3.3 shows the constrained-capacity curve for a different parameter set. The upper bound is noticeably higher than the actual curve. The input distributions are all discrete for the plotted curve. Figure 3.2 also shows the maximizing input distributions for some values for the energy budget. In this case it is easy to ascertain that the input distribution is discrete. This is because the gap between the upper bound and the actual constrained-capacity curve

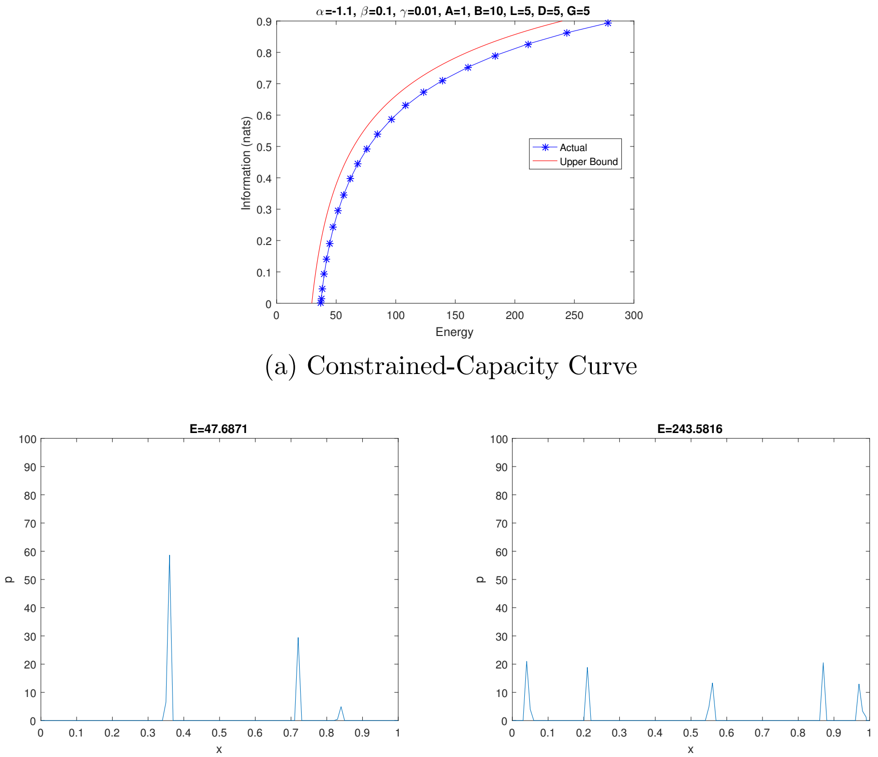

(a) Constrained-Capacity Curve

(b) Input Distribution for $\mathcal{E} = 48$

(c) Input Distribution for $\mathcal{E} = 244$

Figure 3.3: (A) The constrained capacity curve of the GIG neuron channel with parameters $(\alpha, \beta, \gamma) = (-1.1, 0.1, 0.01)$ and energy parameters $(z, a, b, c, r) = (1, 5, 5, 10, 5)$. The upper bound is noticeably higher than the actual curve. The mutual information is given in nats, while the energy unit is arbitrary. (B) The approximate input distribution when $E = 47.69$. Probability density is plotted against $X_c = x$. The input distribution appears to be discrete. (C) The approximate input distribution when $E = 243.58$. Probability density is plotted against $X_c = x$. The input distribution also appears to be discrete with more mass points.

is apparent.

For both examples, it seems that as the energy increases, the number of mass points increases. In the first example, after a certain point, it seems that a continuous input distribution is optimal, though this cannot be stated with certainty.

3.4 Proofs

3.4.1 Proof of Theorem 3.3

The following is the proof of Theorem 3.3:

Proof. First, it is shown that $\mathcal{V}_E(S)$ is tight. By Prokhorov's theorem, it is then relatively compact [69]. Then it is shown that $\mathcal{V}_E(S)$ is a closed set, which makes it sequentially compact. Sequentially compact sets are compact in a metric space. Since the weak topology of probability is metrizable [69], $\mathcal{V}_E(S)$ is a metric space therefore compact.

To show that $\mathcal{V}_E(S)$ is tight, define $g_{\min}$ as a monotonically increasing continuous function over $[0,\infty)$ where $g_{\min}(x) \leq \min\{g(x), g(-x)\}$ and $\lim_{x\to\infty} g_{\min}(x) = \infty$. Let $\mathcal{K} = [-\ell, \ell]$ for $\ell > \ell^*$, where $\ell^* = \max\{\ell > 0 : g_{\min}(\ell) = 0\}$. Note that ℓ^* must be finite. Then for $v \in \mathcal{V}_E(S)$,

$$\mu_v(\mathcal{K}^c) \overset{(a)}{\leq} \frac{\mathbb{E}[g_{\min}(|X|)]}{g_{\min}(\ell)} \overset{(b)}{\leq} \frac{\mathbb{E}[g(X)]}{g_{\min}(\ell)} \leq \frac{E}{g_{\min}(\ell)}, \tag{3.83}$$

where $\mathbb{E}$ is the expectation operator. Markov's Inequality was applied for (a) whereas (b) arises from the definition of $g_{\min}$. For any value of $\epsilon > 0$, a finite value of ℓ can be chosen such that $\frac{E}{g_{\min}(\ell)} < \epsilon$. Therefore $\mathcal{V}_E(S)$ is tight.

To show that $\mathcal{V}_E$ is closed, let $\{v_n\}$ be a sequence in $\mathcal{V}_E(S)$ that converges to v^*. Then [69],

$$\mu_{v^*}(S) \geq \limsup_n \mu_{v_n}(S) = 1, \tag{3.84}$$

since S is closed. Let $\ell > 0$ and define the operator

$$G_\ell(v) = \int_S \min(g(x), \ell)\mathrm{d}v(x). \tag{3.85}$$

Since the integrand is a continuous bounded function, G_ℓ is continuous. Note that $G_\ell(v) \leq G(v) \leq E$ for any $\ell > 0$ and $v \in \mathcal{V}_E$. By continuity of G_ℓ, $G_\ell(v_n) \to G_\ell(v^*)$ for any $\ell > 0$. Note that $G_\ell(v^*) \leq E$ and that $G_\ell(v^*)$ is increasing in ℓ. By the monotone convergence theorem, $G_\ell(v^*) \to G(v^*)$. Therefore, $G(v^*) \leq E$. Hence $\mathcal{V}_E(S)$ is closed and is sequentially compact. Therefore, $\mathcal{V}_E(S)$ compact. □

3.4.2 Proof of Theorem 3.4

Proof. Recall from 3.8,

$$I(v) = h_Y(v) - h_{Y|X}(v). \tag{3.86}$$

However,

$$h_{Y|X}(v) = \iint f_N(y-x) \log f_N(y-x) \mathrm{d}y \mathrm{d}v(x) \tag{3.87}$$
$$= \iint f_N(n) \log f_N(n) \mathrm{d}n \mathrm{d}v(x) \tag{3.88}$$
$$= h_N, \tag{3.89}$$

where $h_N = int f_N(n) \log f_N(n) \mathrm{d}n$ is the differential entropy of N, which is a constant independent of v. Thus, I is continuous if h_Y is continuous.

Let $\{v_n\}$ be a sequence in $\mathcal{V}_E$ that converges weakly to v. Note that $q(y;v) = \mathbb{E}[f_N(y - X)] \leq max_n f_N(n)$ for any $y \in \mathbb{R}$. Thus q is bounded.

Define the function $g_{\mathrm{lb}}(x) = \min\{g(-x+x^*), g(x+x^*)\}$. In words, this is the cost function that is translated so that the minimum occurs at $x = 0$ and is made symmetric by taking the minimum of the function values at equal distance away from $x = 0$. Note that the growth of g_{lb} is greater than the growth of the logarithmic function.

From Theorem 3.3, $\mathcal{V}_E$ is closed, so $v \in \mathcal{V}_E$. Also, $q(y; v_n) \to q(y; v)$ for every $y \in \mathbb{R}$, i.e., pointwise, since $Q(y|\cdot)$ is a continuous bounded function for fixed y. Thus, the conditions of Theorem 1 in [79] are met. Therefore, $h_Y(v_n) \to h_Y(v)$, which means that h_Y, and therefore I, are continuous over $\mathcal{V}_E$.

□

3.4.3 Derivation of the Compact Form of Information

The upper bound of the constrained capacity is [14]

$$\begin{aligned} C_b = \sqrt{b_s c_s}\left[\frac{K_{a_s+1}(2\sqrt{b_s c_s}) + K_{a_s-1}(2\sqrt{b_s c_s})}{K_{a_s}(2\sqrt{b_s c_s})}\right] - a_s \frac{\frac{\partial}{\partial a_s} K_{a_s}(2\sqrt{b_s c_s})}{K_{a_s}(2\sqrt{b_s c_s})} \\ -\sqrt{\beta\gamma}\left[\frac{K_{\alpha+1}(2\sqrt{\beta\gamma}) + K_{\alpha-1}(2\sqrt{\beta\gamma})}{K_{\alpha}(2\sqrt{\beta\gamma})}\right] + \alpha \frac{\frac{\partial}{\partial \alpha} K_{\alpha}(2\sqrt{\beta\gamma})}{K_{\alpha}(2\sqrt{\beta\gamma})} \\ + \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_{\alpha}(2\sqrt{\beta\gamma})}. \end{aligned} \tag{3.90}$$

To simplify the expression, first define

$$K_{\omega}^{(1,0)}(u) = \frac{\partial}{\partial \omega} K_{\omega}(u) \tag{3.91}$$

and

$$K_{\omega}^{(0,1)}(u) = \frac{\partial}{\partial u} K_{\omega}(u). \tag{3.92}$$

Next, insert a dummy variable u with value 1 into (3.90), which yields

$$\begin{aligned} C_b = \Bigg[\sqrt{b_s c_s}\left[\frac{K_{ua_s+1}(2u\sqrt{b_s c_s}) + K_{ua_s-1}(2u\sqrt{b_s c_s})}{K_{ua_s}(2u\sqrt{b_s c_s})}\right] - a_s \frac{K_{ua_s}^{(1,0)}(2u\sqrt{b_s c_s})}{K_{ua_s}(2u\sqrt{b_s c_s})} \\ -\sqrt{\beta\gamma}\left[\frac{K_{u\alpha+1}(2u\sqrt{\beta\gamma}) + K_{u\alpha-1}(2u\sqrt{\beta\gamma})}{K_{u\alpha}(2u\sqrt{\beta\gamma})}\right] + \alpha \frac{K_{u\alpha}^{(1,0)}(2u\sqrt{\beta\gamma})}{K_{u\alpha}(2u\sqrt{\beta\gamma})}\Bigg]_{u=1} \\ + \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_{\alpha}(2\sqrt{\beta\gamma})}. \end{aligned} \tag{3.93}$$

Using Appendix A.4, the recurrence relationship of K_α can be exploited. The constrained capacity can then be written as

$$\begin{aligned} C_b = \Bigg[- \frac{a_s K_{ua_s}^{(1,0)}(2u\sqrt{b_s c_s}) + \sqrt{b_s c_s} K_{ua_s}^{(0,1)}(2u\sqrt{b_s c_s})}{K_{ua_s}(2u\sqrt{b_s c_s})} \\ + \frac{\alpha K_{u\alpha}^{(1,0)}(2u\sqrt{\beta\gamma}) + \sqrt{\beta\gamma} K_{u\alpha}^{(0,1)}(2u\sqrt{\beta\gamma})}{K_{u\alpha}(2u\sqrt{\beta\gamma})}\Bigg]_{u=1} \\ + \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_{\alpha}(2\sqrt{\beta\gamma})}. \end{aligned} \tag{3.94}$$

Exploiting the differentiation yields

$$C_b = \left[-\frac{\frac{\mathrm{d}}{\mathrm{d}u} K_{ua_s}(2u\sqrt{b_s c_s})}{K_{ua_s}(2u\sqrt{b_s c_s})} + \frac{\frac{\mathrm{d}}{\mathrm{d}u} K_{u\alpha}(2u\sqrt{\beta\gamma})}{K_{u\alpha}(2u\sqrt{\beta\gamma})} \right]_{u=1} + \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_\alpha(2\sqrt{\beta\gamma})}. \tag{3.95}$$

Again, exploiting differentiation yields

$$C_b = \left[\frac{\mathrm{d}}{\mathrm{d}u} \log(K_{u\alpha}(2u\sqrt{\beta\gamma})) - \frac{\mathrm{d}}{\mathrm{d}u} \log(K_{ua_s}(2u\sqrt{b_s c_s})) \right]_{u=1} + \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_\alpha(2\sqrt{\beta\gamma})}. \tag{3.96}$$

Combining the two expressions with logarithm yields

$$C_b = \log \frac{K_{a_s}(2\sqrt{b_s c_s})}{K_\alpha(2\sqrt{\beta\gamma})} + \frac{\mathrm{d}}{\mathrm{d}u} \left[\log \left(\frac{K_{u\alpha}(2u\sqrt{\beta\gamma})}{K_{ua_s}(2u\sqrt{b_s c_s})} \right) \right]_{u=1},$$

which is the same as (3.80).

3.5 Discussions

3.5.1 Implications for Neural Networks

The result presented in this chapter suggests that for certain parameter sets, discrete and finite possibilities for the average input intensities to neuron η are desirable. Since the neural network controls the input intensity to η, it may be beneficial for the network to exist in discrete states to optimize the information transmission of η. This statement hinges on the GIG neuron model assumptions. However, there has been some evidence that neural networks have a set of possible discrete states [80–85]. If the results apply when the input intensity in each IPI forms a Markov chain, then the GIG neuron model can be related to the hidden Markov model with finite and discrete states. The input intensity Λ acts as an unobservable state and Λ can only be estimated from the output T.

However, it is also possible for the average input intensity to be continuous. Perhaps it is the case that some parts of the network exist in finite states and some in continuous

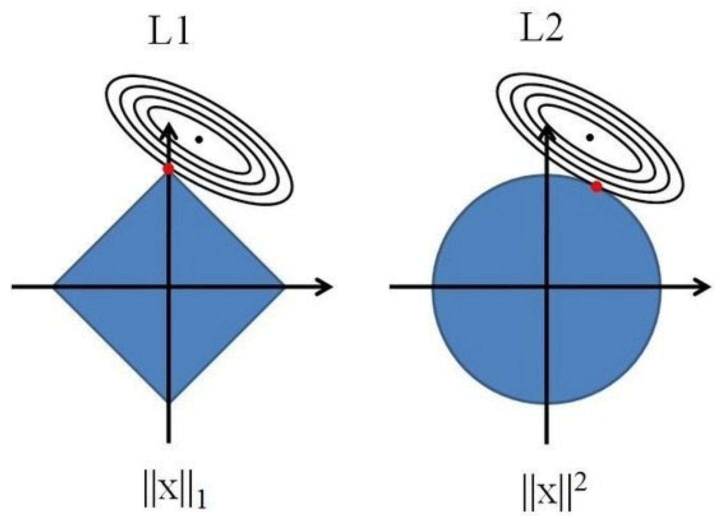

Figure 3.4: The L1-norm constraint promotes sparse solutions as opposed to the L2-norm constraint. The black ellipses are the contour lines for the objective function, i.e., the function to be optimized. The blue region is the space of possible solutions. Optimization solutions from an L1-norm constraint tends to be on an axis, which implies a sparse solution.

states. Experiments are also needed to obtain good values for the parameters of the model and determine whether the states need to be discrete to reach optimal performance.

3.5.2 Possible Connection to L1 Norm Constraint

Recent machine learning algorithms and compressed sensing techniques have used an L1-norm constraint to promote sparse solutions. The shape created by the L1-norm increases the likelihood of solutions to occur on the axes. This is illustrated in Figure 3.4. However, it is still possible for the solution to be not sparse. This depends on the shape and peak of the objective function.

The discrete result of the GIG neuron model can be viewed in some sense as a sparse solution. To explain why the solution of the GIG neuron model may be sparse, the constrained-capacity problem can be posed differently. It can be posed as searching the constrained space of generalized functions for maximizing MI. The space of generalized functions include functions and other entities like the Dirac-delta "function." Let $\mathcal{G}$ be the space of generalized functions. The constraint on the elements of $\mathcal{G}$ are the non-negativity and normalization to 1 constraints. That is, for $f \in \mathcal{G}$, the constraint space is

$$\begin{aligned}
&\int f(x)\mathrm{d}x = 1 \\
&f(x) \geq 0, \quad x \in \mathbb{R}.
\end{aligned} \tag{3.97}$$

By rewriting the normalization constraint as two inequality constraints, the following is acquired,

$$\begin{aligned}
&\int |f(x)|\mathrm{d}x \leq 1 \\
&\int |f(x)|\mathrm{d}x \geq 1 \\
&f(x) \geq 0, \quad x \in \mathbb{R}.
\end{aligned} \tag{3.98}$$

The integrand has been replaced by $|f(x)|$, i.e., its absolute value. This is is equivalent because the function is non-negative. Note that the first constraint is now the L1-norm constraint. Hence, the constraint space is a subset of the L1 unit ball. This is perhaps a starting point in analyzing the connection between the discrete input distribution for the constrained-capacity and the L1-norm constraint.

3.6 Conclusion

In this chapter, the constrained capacity was determined for the GIG neuron model. It turns out that for certain cases, the capacity-achieving input distribution may be discrete, depending on the parameter set. This was demonstrated with a numerical example, where it was shown that the optimal input distribution can change from discrete to continuous by changing the energy budget. This finding implies that for certain parameter sets, it is optimal to use a discrete number of possible average input intensity values in terms of maximizing MI. Since the network controls the input intensity to η, this implies that the network may exist in discrete states to maximize information transmission. More study is needed to determine the possible set of parameters that exists in neurons to determine which require discrete inputs to perform optimally.

Chapter 4

Optimal Parameters for the GIG Neuron Model

This chapter addresses how the generalized inverse Gaussian (GIG) neuron model can be optimized over its parameter set. The assumption remains the same: neuron η seeks to maximize the mutual information (MI) transmitted to its targets while adhering to an average energy budget. However, additional constraints are necessary for the problem to be well-posed. The optimization problem is illustrated for the GIG neuron model when $\gamma = 0$.

4.1 Background

As discussed in Chapter 1, neuron η seeks to maximize the MI it transmits while adhering to an average energy budget. In Chapter 3, it was shown how the input to η should be configured so that it transmits the maximum MI. This is a problem for η and the neural network (NN) since the NN controls the input to η. The question that will be partially addressed here is how does η configure itself so that it transmits MI efficiently.

The structure of η determines the input-output model of the neuron. The input-output model can be viewed as a channel, whose conditional probability density function (PDF) of the output given the input is Q. Hence, η seeks a configuration such that when modeled by

Q, it transmits MI efficiently. More specifically, given a distribution for the average input intensity Λ, what is the channel Q that can transmit the most MI given an average energy budget. Compare this to Chapter 3, where the goal is to find the optimal input distribution for a given channel model. Therefore, there is a notion of double matching: the matching of the channel to the source and the matching of the source to the channel. There is a joint effort of optimizing the η by η itself and the rest of the NN.

There are restrictions on Q. Recall from Chapter 2 that η produces an output when its postsynaptic potential (PSP) hits a threshold level θ. If the PSP is modeled as a stochastic diffusion, Q must be a first hitting time (FHT) distribution. This limits the class of distributions that is possible for Q. Also, the average input rate Λ and the output interpulse interval (IPI) duration T are negatively correlated. If the input rate is high, then the rate of increase of the PSP should increase. This leads to hitting the threshold sooner, which shortens the duration of the IPI. Thus, Q should exhibit this relationship.

However, it is difficult to define the space of conditonal PDF's that are FHT's of stochastic diffusions. To simplify the problem, the space of conditional PDF's is restricted to the space of GIG distributions

4.2 Theory

4.2.1 Maximizing Over the Mutual Information

Assume that the input cumulative density function (CDF) of Λ is known and denote it by v_Λ. For the GIG conditional PDF, the MI is given by

$$I(\Lambda; T) = \iint Q_{\mathrm{GIG}}(t|\lambda) \frac{\log Q_{\mathrm{GIG}}(t|\lambda)}{q_T(t)} \mathrm{d}t \mathrm{d}v_\Lambda(\lambda), \tag{4.1}$$

where q_T is the induced output PDF of T, given by

$$q_T(t) = \int Q_{\mathrm{GIG}}(t|\lambda)\mathrm{d}v_\Lambda(\lambda), \quad t > 0. \tag{4.2}$$

Recall that the GIG conditional PDF is given by (see Chapter 2)

$$Q_{\mathrm{GIG}}(t|\lambda) = M(\alpha,\beta,\gamma)^{-1}\lambda^\alpha t^{\alpha-1}\exp\Big(-\frac{\beta}{\lambda t} - \gamma\lambda t\Big), \lambda, t > 0, \tag{4.3}$$

which is parameterized by three parameters: α, β, γ. Hence, elements of the set of GIG conditional PDF's can be identified with elements of a subset of $\mathbb{R}^3$. In fact, for GIG PDF's that are FHT's of diffusions, $\alpha \leq 0$. Hence, the parameter space of interest is given by

$$\mathcal{A} = \{(\alpha,\beta,\gamma) \in \mathbb{R}^3 : \alpha \leq 0,\ \beta > 0,\ \gamma \geq 0 \text{ and } \gamma > 0 \text{ if } \alpha = 0\}. \tag{4.4}$$

The set $\mathcal{A}$ is convex, whereas the space of GIG PDF's are not convex. Hence, it is more convenient to use $\mathcal{A}$ in place of the set of GIG PDF's, though they are essentially equivalent.

When considering the MI over the space of GIG parameters, the convexity of the MI is not guaranteed over the s. Define $I_{\mathrm{GIG}}(\alpha,\beta,\gamma) = I(\Lambda;T)$ for the GIG neuron channel with its respective parameters α, β, and γ. The function I_{GIG} is not guaranteed to be convex over $\mathcal{A}$. However, it will be argued that the supremum of I_{GIG} must occur on the boundary. Thus, the search for the supremum of I_{GIG} can be restricted to the boundary.

4.2.2 Constraints

Neuron η seeks to maximize MI given an average energy budget. As mentioned in Chapter 3, the energy expended by η for average input intensity λ and output IPI duration t is given by

$$g(\lambda, t) = z + b/t + ct - a\log t + r\lambda t, \quad \lambda, t > 0. \tag{4.5}$$

The average energy associated with the GIG model is given by

$$G_{\mathrm{GIG}}(\alpha, \beta, \gamma) = \iint g(\lambda, t) Q_{\mathrm{GIG}}(t|\lambda) \mathrm{d}t \mathrm{d}v(\lambda). \tag{4.6}$$

Evaluating the integral yields

$$G_{\mathrm{GIG}}(\alpha, \beta, \gamma) = z + bk_{-1}(\alpha, \beta, \gamma)\mathbb{E}[\Lambda] + (b\mathbb{E}[\Lambda^{-1}] + r)k_1(\alpha, \beta, \gamma) - a(k_g(\alpha, \beta, \gamma) - \mathbb{E}[\log \Lambda]), \tag{4.7}$$

where k_{-1}, k_1, and k_g are defined in Appendix B and $\mathbb{E}$ is the expectation operator, which is given by

$$\mathbb{E}[f(\Lambda)] = \int f(\lambda) \mathrm{d}v(\lambda). \tag{4.8}$$

Thus, it is assumed that $\mathbb{E}[\Lambda]$, $\mathbb{E}[\Lambda^{-1}]$, and $\mathbb{E}[\log \Lambda]$ exist. Also note that here, k_{-1}, k_1, and k_g are expressed explicitly as functions of (α, β, γ) to emphasize that their values vary as (α, β, γ) is varied.

For energy budget E, the constraint is given by

$$G_{\mathrm{GIG}}(\alpha, \beta, \gamma) \leq E. \tag{4.9}$$

The constrained space to be over which I_{GIG} is optimized is

$$\mathcal{A}_E = \{(\alpha, \beta, \gamma) \in \mathcal{A} : G_{\mathrm{GIG}}(\alpha, \beta, \gamma) \leq E\}. \tag{4.10}$$

It turns out that this constraint is not enough. The MI is unbounded over this set. This can be illustrated by the following example. The energy function with fixed input λ, $g(\lambda, \cdot)$, has a minimum. Let $t^*(\lambda)$ be that minimum. Consider a sequence of GIG conditional PDF's, each with conditional mean $t^*(\lambda)$. Suppose that the conditional variance across all $\lambda > 0$ is decreasing. The average energy in the sequence is then decreasing because the PDF's become more concentrated around $t^*(\lambda)$. Furthermore, MI is also increasing because the channel is less noisy. Hence, I_{GIG} is unbounded over $\mathcal{A}_E$.

This does not model real neurons, where the transmitted MI is expected to be finite. Thus, when considering the neuron model that maximizes MI, constraints should be considered in addition to the average energy.

Recall that the GIG neuron channel is given by $T = U/\Lambda$ where U is independent of Λ and U is distributed as $\text{GIG}(\alpha, \beta, \gamma)$. The random variable (RV) U is the noise, and thus, it is expected to have a positive variance. Otherwise, T is a deterministic function of Λ and the channel is noiseless. As described in Chapter 2, η receives a noisy version of Λ due to the random nature of incoming signals from η's afferent cohort. Thus, there should be a bound on the spread of U. The most common measure of spread is variance, thus

$$\text{Var}(U) \leq \Psi \tag{4.11}$$

where Ψ is bound on variance. The variance is defined as

$$\text{Var}(U) = \mathbb{E}\big[(U - \mathbb{E}[U])^2\big]. \tag{4.12}$$

Since $U = \Lambda T$, the bound can also be written as $\text{Var}(\Lambda T) \leq \Psi$.

Define $\nu(\alpha, \beta, \gamma) = \text{Var}(\Lambda T)$ for the GIG neuron model. Its value is given by

$$\nu(\alpha, \beta, \gamma) = k_2(\alpha, \beta, \gamma) - k_1(\alpha, \beta, \gamma)^2, \quad (\alpha, \beta, \gamma) \in \mathcal{A}, \tag{4.13}$$

where k_2 is defined in Appendix B. The set of PDF's of interest is now

$$\mathcal{A}_{E,\Psi} = \{(\alpha, \beta, \gamma) \in \mathcal{A} : G_{\text{GIG}}(\alpha, \beta, \gamma) \leq E,\ \nu(\alpha, \beta, \gamma) \geq \Psi\}. \tag{4.14}$$

Even though the variance constraint is imposed with an inequality, it must actually be satisfied by an equality for any value of $\Psi > 0$. If it were not satisfied by an equality for Ψ_0, then it would not be in that case. It was argued that if the variance is unconstrained, it will tend to 0. Thus, the variance constraint is necessary for any value of $\Psi > 0$, which implies that it must be satisfied with equality for any $\Psi > 0$. The energy constraint must also be satisfied with equality. Otherwise, the means of T given $\Lambda = \lambda$ can be spaced infinitely apart, which would yield infinite MI. Thus, the solution lies on the boundary of the constraint space.

4.2.3 Optimization

Since the solution lies on the boundary, it is assumed that the inequality constraints are satisfied with equality, i.e., $G_{\text{GIG}}(\alpha, \beta, \gamma) = E$ and $\nu(\alpha, \beta, \gamma) = \Psi$. Then the approach to use is the Lagrange method. The Lagrangian is given by

$$L(\alpha, \beta, \gamma) = I_{\text{GIG}}(\alpha, \beta, \gamma) - s_1(G_{\text{GIG}}(\alpha, \beta, \gamma) - E) - s_2(\nu(\alpha, \beta, \gamma) - \Psi), \tag{4.15}$$

where s_1 and s_2 are the Lagrange multipliers associated with the energy and variance constraints, respectively. However, the set $\mathcal{A}_{E,\Psi}$ is not compact, so the maximum is not guaranteed. Numerical trials seem to indicate that the set $\mathcal{A}_{E,\Psi}$ is bounded and the supremum of I_{GIG} is finite.

To find the maximum of L, take the partial derivative with respect to each of the parameter and the Lagrange multiplier and set it equal to zero. Then the solution can be obtained. With regards to the partial derivatives, the following theorems are presented:

Theorem 4.1. *Let*

$$F(\underline{x}; \alpha) = \frac{n(\underline{x}; \omega)}{m(\omega)} \tag{4.16}$$

be a joint PDF of a random vector $\underline{X}$ with parameter α. Then for function f,

$$\frac{\partial}{\partial \omega}\mathbb{E}[f(\underline{X})] = \int \frac{n_\omega(\underline{x}; \omega)}{d(\omega)}\{f(\underline{x}) - \mathbb{E}[f(\underline{X})]\}\mathrm{d}\underline{x} + \mathbb{E}[f_\omega(\underline{X})], \tag{4.17}$$

where n_ω and f_ω are the partials of n and f with respect to ω,respectively. The notation $\int \ldots \mathrm{d}\underline{x}$ represents the n^{th} fold integral with respect to each element of $\underline{x}$.

Proof. Take the partial derivative with respect to ω:

$$\frac{\partial}{\partial \omega}\mathbb{E}[f(\underline{X})] = \int \left(\frac{[n_\omega(\underline{x};\omega)f(\underline{x}) + n(\underline{x};\omega)f_\omega(\underline{x})]}{d(\omega)} - \frac{n(\underline{x};\omega)f(\underline{x})d_\omega(\omega)}{d(\omega)^2} \right) \mathrm{d}\underline{x}, \tag{4.18}$$

where d_ω is the partial derivative of d with respect to ω. However, since F is a PDF,

$$d(\omega) = \int n(\underline{x};\omega)\mathrm{d}\underline{x}. \tag{4.19}$$

so

$$d_\omega(\omega) = \int n_\omega(\underline{x};\omega)\mathrm{d}\underline{x}. \tag{4.20}$$

Then using the definition of expectations yields

$$\frac{\partial}{\partial \omega}\mathbb{E}[f(\underline{x})] = \int \frac{n_\omega(\underline{x},\omega)}{d(\omega)} f(\underline{x})\mathrm{d}\underline{x} + \mathbb{E}[f_\omega(\underline{X})] - \mathbb{E}[f(\underline{X})] \int \frac{n_\omega(\underline{x};\omega)}{d(\omega)}\mathrm{d}\underline{x}. \tag{4.21}$$

Rearranging the terms reveals that this is equal to (4.17). □

Theorem 4.2. *The following statement is true:*

$$\mathbb{E}[\iota_\omega(X,Y)] = 0, \tag{4.22}$$

where $\iota(x,y) = \log Q(y|x)/q(y)$ is the information density and ι_ω is its partial derivative with respect to ω. The function Q is a conditional PDF of Y given that $X = x$ and q is the PDF of Y.

Proof. Note the following:

$$\iota_\omega(x,y) = \frac{Q_\omega(y|x)}{Q(y|x)} - \frac{q_\omega(y)}{q(y)}, \tag{4.23}$$

where Q_ω and q_ω are the respective partial derivatives of Q and q with respect to ω. Integrating over the joint distribution yields

$$\mathbb{E}[\iota_\omega(X,Y)] = \iint Q(y|x)\left(\frac{Q_\omega(y|x)}{Q(y|x)} - \frac{q_\omega(y)}{q(y)}\right)\mathrm{d}y\mathrm{d}v(x) \tag{4.24}$$

$$= \iint Q_\omega(y|x)\mathrm{d}y\mathrm{d}v(x) - \int q_\omega(y)\mathrm{d}y = 0, \tag{4.25}$$

where v is the CDF of X. □

Lemma 4.3. *The following statement is true:*

$$\mathbb{E}\left[\frac{\partial}{\partial\omega}(X - \mathbb{E}[X])^2\right] = 0. \tag{4.26}$$

Proof. Note the following:

$$\frac{\partial}{\partial\omega}(x - \mathbb{E}[X])^2 = -2(X - \mathbb{E}[X])\frac{\partial}{\partial\omega}\mathbb{E}[X]. \tag{4.27}$$

It is clear that the expected value is 0. □

Hence, the partial derivatives of the Lagrangian are

$$\frac{\partial L}{\partial\alpha} = \mathbb{E}\big[\log(\Lambda T)\big(\iota(\Lambda,T) - s_1 g(\Lambda,T) - s_2(\Lambda T - \mathbb{E}[\Lambda T])^2\big)\big] = 0 \tag{4.28}$$

$$\frac{\partial L}{\partial\beta} = -\mathbb{E}\big[(\Lambda T)^{-1}\big(\iota(\Lambda,T) - s_1 g(\Lambda,T) - s_2(\Lambda T - \mathbb{E}[\Lambda T])^2\big)\big] = 0 \tag{4.29}$$

$$\frac{\partial L}{\partial\gamma} = -\mathbb{E}\big[\Lambda T\big(\iota(\Lambda,T) - s_1 g(\Lambda,T) - s_2(\Lambda T - \mathbb{E}[\Lambda T])^2\big)\big] = 0. \tag{4.30}$$

Also, the original constraints are obtained. The expectations can be approximated with the Gauss-Laguerre quadrature and the equations can be solved numerically. This is illustrated in the next section for the case $\gamma = 0$ and $c = 0$.

4.3 Numerical Example

In [23], the optimal model parameters were determined for $\beta = 0$ and $b = 0$. In this thesis, the case $\gamma = 0$ and $c = 0$ is shown.

Let the input PDF be a GIG distribution with $\gamma = 0$, i.e., an inverse Gamma distribution. This is given by

$$p_\Lambda(\lambda) = \frac{\exp(-b_i/\lambda)}{b_i^{a_i}\Gamma(-a_i)}, \quad \lambda > 0, \tag{4.31}$$

where $b_i > 0$ and $a_i < 1$. From (4.7) and Appendix B, the average energy is given by

$$G_{\mathrm{GIG}}(\alpha,\beta) = r + b\frac{b_i}{a_i+1}\frac{\alpha}{\beta} + \left(c\frac{a_i}{b_i} + r\right)\frac{\beta}{\alpha+1} - a(\log\beta - \psi(-\alpha) - \log b_i + \psi(-a_i)) \tag{4.32}$$

. From Appendix B, the variance is given by

$$\nu(\alpha,\beta) = \frac{\beta^2}{(-\alpha-1)^2(-\alpha-2)}. \tag{4.33}$$

In this case, it must be that $\alpha < 2$ so that the energy and variance exists. Hence, that additional restriction is added to $\mathcal{A}_{E,\Psi}$.

To find the mutual information, the marginal of T is first determined:

$$q_T(t) = \int \frac{t^{\alpha-1}\lambda^{\alpha+a_i-1}\exp(-(\beta/t+b_i)/\lambda)}{\beta^\alpha b_i^{a_i}\Gamma(-\alpha)\Gamma(-a_i)}\mathrm{d}\lambda \tag{4.34}$$

$$= \left(\frac{b_i}{\beta}\right)^\alpha \frac{t^{a_i-1}(t+\beta/b_i)^{\alpha+a_i}}{B(-\alpha,-a_i)}, \tag{4.35}$$

for $t > 0$ and where B is the beta function. The PDF q_T is a beta prime distribution. The information density for $\Lambda = \lambda$ and $T = t$ is given by

$$\iota_{\mathrm{GIG}}(\lambda,t) = -\frac{\beta}{\lambda t} + \alpha\log\lambda - (\alpha+a_i)\log(1/t + b_i/\beta) + \log\frac{b_i^{a_i}\Gamma(-a_i)}{\beta^{\alpha+a_i}\Gamma(-\alpha)}. \tag{4.36}$$

Note that

$$\mathbb{E}[\log(1/T + b_i/\beta)] = \frac{(\beta/b_i)^{a_i}}{B((-\alpha, -a_i)} \frac{\partial}{\partial a_i} \int t^{\alpha-1}(1/t + b_i/\beta)^{\alpha+a_i} \mathrm{d}t \tag{4.37}$$

$$= \frac{(\beta/b_i)^{a_i}\Gamma(-\alpha - a_i)}{\Gamma(-\alpha)\Gamma(-a_i)} \frac{\partial}{\partial a_i} \frac{\Gamma(-\alpha)\Gamma(-a_i)}{(\beta/b_i)^{a}{}_{i}\Gamma(-\alpha - a_i)} \tag{4.38}$$

$$= -\psi(-a_i) - a_i \log(\beta/b_i) + \psi(-\alpha - a_i). \tag{4.39}$$

Thus, the MI is given by the expectation of the information density. So,

$$I_{\mathrm{GIG}}(\alpha, \beta) = \mathbb{E}[\iota_{\mathrm{GIG}}(\Lambda, T)]$$

$$= \alpha + a_i\psi(-a_i) - (\alpha + a_i)\psi(-\alpha - a_i) + \log \frac{\Gamma(-a_i)}{\Gamma(-\alpha - a_i)}, \tag{4.40}$$

where ψ is the digamma function, given by

$$\psi(x) = \frac{\mathrm{d}}{\mathrm{d}x} \log \Gamma(x). \tag{4.41}$$

Note that I_{GIG} is not a function of β and the parameter b_i does not appear. The factor b_i acts a scaling factor for Λ, which does not affect MI. Likewise, β also acts as a scaling factor for T, so it does not have an impact on MI. However, β and b_i affects the average energy and the variance and thus cannot be ignored. Then observe that the derivative of I_{GIG} with respect to α is

$$\frac{\partial I}{\partial \alpha} = 1 - \psi(-\alpha - a) + (\alpha + a)\psi_1(-\alpha - a) + \psi(-\alpha - a) \tag{4.42}$$

$$= 1 + (\alpha + a)\psi_1(-\alpha - a), \tag{4.43}$$

where ψ_1 is the trigamma function given by

$$\psi_1(x) = \frac{\mathrm{d}}{\mathrm{d}x}\psi(x). \tag{4.44}$$

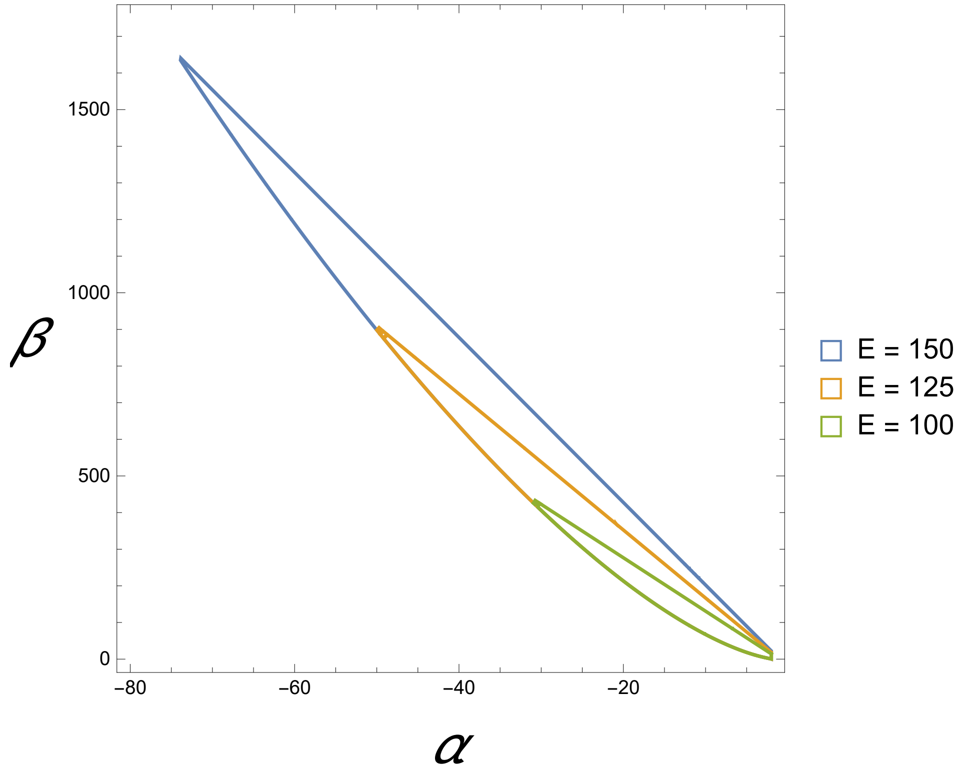

Figure 4.1: The feasibility region $\mathcal{A}_{E,V}$ for some values for the energy budget. The values of the energy parameters are $z = 10$, $a = 3$, $b = 5$, $c = 7$, and $r = 3$. The parameters of the input distribution are $a_i = -5$ and $b_i = 10$. Here $\Psi = 7$, while the the energy budget is shown for three values: $E = 100, 125$, and 150. The feasibility region is bounded, though not necessarily closed. The point that maximizes MI is the one furthest to the left.

The derivative is strictly negative. Hence, MI increases as α gets further below zero.

Some feasibility regions are plotted in Figure 4.1. The feasible point with the smallest value of α is one of the intersections of the boundaries of the constraints. Hence, as previously argued, the inequality constraints can be replaced with equality. A constrained optimization algorithm can be used to find the solution. However, since there are two equations and two unknown parameters, only the intersection of the boundaries are required to find the point of maximum MI. Note that the intersections between the two constraint boundaries are not unique. Hence, the intersection with the lowest value of α is the point that maximizes information.

The information-energy curve for different values of Ψ are shown in figure 4.2. The MI is an increasing concave function of the energy budget. As the variance constraint is decreased,

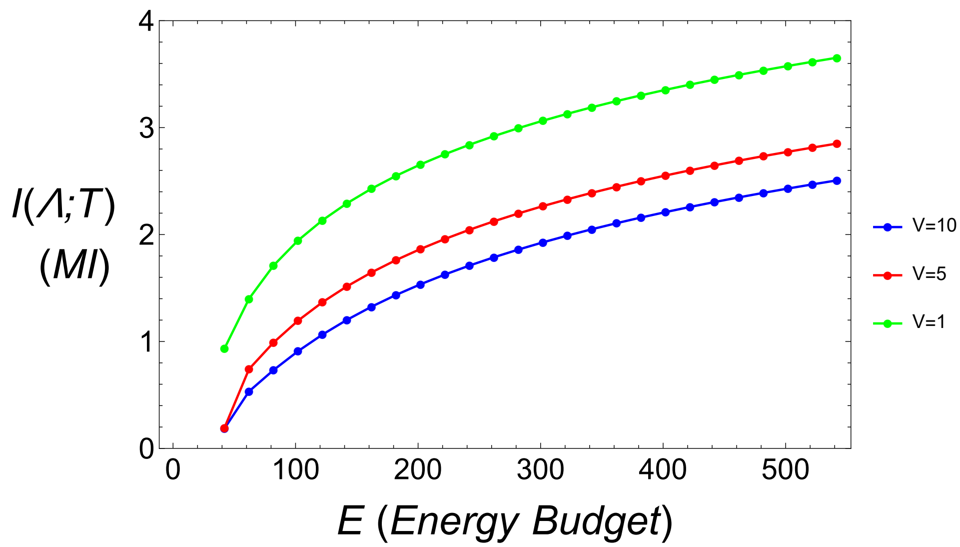

Figure 4.2: The information-energy curve for energy parameters $z = 10$, $a = 3$, $b = 5$, $c = 7$, and $r = 3$ and input distribution parameters $a_i = -5$ and $b_i = 10$. The MI is an increasing concave function of the energy budget. As the variance is decreased, MI increases. The units of MI are expressed in nats, whereas the energy units are arbitrary.

information increases. This is intuitive because as the noise decreases in variance, there is less distortion in the output, which increases information.

More investigation is needed to determine the value of the energy parameters that reflect neuron η. Likewise, more investigation is needed to determine the value for the variance constraint.

4.4 Discussion

This problem has not yet been applied to the general GIG PDF. However, as mentioned, initial studies have suggested that $\mathcal{A}_{E,\Psi}$ is bounded and I_{GIG} is finite over $\mathcal{A}_{E,\Psi}$ in the general case. Also, the MI can be reduced to a function of α and $\sqrt{\beta\gamma}$. This is similar to the case $\gamma = 0$, the MI is only a function of α. The GIG PDF can be written as

$$Q(t|\lambda) = M(\alpha,\beta,\gamma)^{-1}\lambda^{\alpha}t^{\alpha-1}\exp\left(-\sqrt{\beta\gamma}\left(\sqrt{\frac{\beta}{\gamma}}\frac{1}{\lambda t} - \sqrt{\frac{\gamma}{\beta}}\lambda t\right)\right), \quad \lambda, t > 0. \tag{4.45}$$

Hence, the factor $\sqrt{\gamma/\beta}$ is a scale factor on the output RV. Often times, the GIG PDF is parameterized with the scale parameter $\sqrt{\gamma/\beta}$, the concentration parameter $\sqrt{\beta\gamma}$, and α [54].

This process of matching the channel to the source can be followed by matching the source to the channel. Once the procedure generates the optimal channel, the theories in Chapter 3 can be used to find the optimal input distribution. Then the channel can be matched again to the new source and the process repeats. It is not yet known if MI would converge to a finite value. It is also not yet known whether the input distribution and the channel would converge. This is a possible next step for this research.

4.5 Conclusion

In this chapter, the optimal parameters for the GIG neuron model was investigated. In addition to an energy budget, a variance constraint on the model noise U was imposed so that the problem is well-posed. Optimality conditions for the problem were developed. Then, a numerical example for a simplified GIG neuron model was shown.

More research is needed to extend this to the general GIG neuron model. Also, an iterative procedure to find the input distribution and channel pair that maximizes MI given an energy budget and variance constraint was suggested. However, more research is needed to show that this procedure converges. Ultimately, it would be interesting to compare the result of the iterative procedure to real cortical neurons, i.e., will the result of the iterative procedure closely model the cortical neuron.

Chapter 5

The GIG Distribution and Rate-Distortion

A counterpart to the capacity-cost (C-C) problem is the rate distortion (R-D) problem. The goal is to reduce the rate of information with an acceptable level of distortion. It was introduced by Shannnon in 1959 [86]. Since then, R-D problems had been posed and solved for a variety of cases, including multi-terminal coding.

In R-D problems, the input distribution is assumed to be known and the mutual information (MI) is minimized over a set of "test" channels subject to a distortion constraint. The output alphabet is also referred to as the reconstruction alphabet. This is because the output is often desired to be close to the input, i.e., a reconstruction of the input.

Throughout this chapter, only the discrete memoryless source (DMS) will be studied. In a DMS, the output of a source at one time is independent of the output of the source at any other times. Also, the distribution of the source does not change over time.

This chapter shows that the reconstruction letter can be discrete even when the DMS is continuous. To the best of the author's knowledge, this result was first presented by Fix in 1978 [87]. This chapter will give a more complete background on the discreteness result. This will be demonstrated with a Gaussian input distribution with peak-limited reconstruction alphabet. Then, based on the results of Chapter 3, a discrete reconstruction result is obtained for a similar R-D problem. Namely, when the DMS is distributed as a generalized inverse Gaussian (GIG) distribution, the reconstruction alphabet is discrete for certain cases.

5.1 Rate-Distortion Problems

The same convention as Chapter 3 is used here. The reader is encouraged to refer to Sections 3.1.1 and 3.1.2 for notation and definition of MI.

Let X be designated as the input letter from a DMS. Let Y be the reconstruction, or output, letter. Let $S \in \mathcal{B}(\mathbb{R}^2)$, i.e., S is an element of the Borel algebra of $\mathbb{R}^2$. Define $\mathcal{F}_0(S)$ as the set of joint CDF's whose points of increase is a subset of S. That is, if $F \in \mathcal{F}_0(S)$, then $\mu_F(S) = 1$. This set will be indicated by just $\mathcal{F}_0$ when S is understood. The set S contains all possible input-output pairs.

Define $\mathcal{F}(S)$ as

$$\mathcal{F}(S) = \left\{F \in \mathcal{F}_0(S) : \lim_{y \to \infty} F(x, y) = v(x),\ x \in \mathbb{R}\right\}, \tag{5.1}$$

i.e., the set of CDF's where the marginal CDF of X is given by v. Also, define the set of allowable inputs:

$$\mathcal{X} = \{x \in \mathbb{R} : (x, y) \in S,\ y \in \mathbb{R}\}. \tag{5.2}$$

Likewise, define the set of allowable outputs:

$$\mathcal{Y} = \{y \in \mathbb{R} : (x, y) \in S,\ x \in \mathbb{R}\}. \tag{5.3}$$

Let S_X be the points of increase for v. For every CDF $F \in \mathcal{F}$, there exists a conditional CDF of Y given X defined by

$$W(y|x) = \frac{\mathrm{d}F(x, y)}{\mathrm{d}v(x)}, x \in S_X,\ y \in \mathbb{R}. \tag{5.4}$$

The value of $W(y|x)$ can be any real number for $x \notin S_X$, but for convention, let $W(y|x) = 0$ for $x \notin S_X$. Define $\mathcal{W}$ as the space of such conditional CDF's. Each element of $\mathcal{F}$ can be identified with an element in $\mathcal{W}$ and vice-versa. Equip $\mathcal{F}$ with the weak topology of probability. Then $\mathcal{W}$ is also equipped with the weak topology of probability. A sequence $\{W_n\}$ in $\mathcal{W}$ converges weakly in measure to W if

$$\iint f(x,y)\mathrm{d}W_n(y|x)\mathrm{d}v(x) \to \iint f(x,y)\mathrm{d}W(y|x)\mathrm{d}v(x) \tag{5.5}$$

for every continuous bounded function f. Since v is known, the MI is a function of W. Hence, the MI will be indicated by $I(W)$ and is given by

$$I(W) = \iint \log \frac{\mathrm{d}W(y|x)}{\mathrm{d}w(y;v,W)} \mathrm{d}W(y|x)\mathrm{d}v(x),$$

where,

$$w(y;v,W) = \int W(y|x)\mathrm{d}v(x)$$

is the marginal CDF of Y.

Suppose that distortion is measured by a real number and the distortion measure between input x and output y is given by the function ϕ. The average distortion operator on $W \in \mathcal{W}$ is

$$\Phi(W) = \iint \phi(x,y)\mathrm{d}W(y|x)\mathrm{d}v(x). \tag{5.6}$$

Then, define $\mathcal{W}_D(S)$ as the set of conditional CDF's that satisfy $\Phi(W) \leq D$, i.e.,

$$\mathcal{W}_D(S) = \{W \in \mathcal{W}(S) : \Phi(W) \leq D\}. \tag{5.7}$$

In a R-D problem, the goal is to minimize I over $\mathcal{W}_D$,

$$R(D) = \inf_{W \in \mathcal{W}_D} I(W), \quad D \in D_{\min}\mathcal{D}, \tag{5.8}$$

where $\mathcal{D} = \{D : \Phi(W) \le D,\ W \in \mathcal{W}\}$ is the set of possible average distortions. Define $D_{\min} = \int \inf_y \phi(x,y)\mathrm{d}v(x)$, i.e., the infimum of the average distortion. Note that $D_{\min}$ is also the infimum of $\mathcal{D}$.

5.2 Solution of the Rate-Distortion Problem

The following theorem states the existence of the solution to the R-D problem and gives the optimality condition. The proof is shown in Section 5.4.1.

Theorem 5.1. *If $\mathcal{W}_D(S)$ is compact, then*

$$R(D) = \min_{W \in \mathcal{W}_D} I(W), \quad D \in \mathcal{D}, \tag{5.9}$$

for some $W^* \in \mathcal{W}_D$. Furthermore, if $D > D_{min}$ and $R(D)$ is finite, W^* minimizes I if and only if

$$\iint [\iota(x,y;W^*) - s\phi(x,y)]\mathrm{d}W(y|x)\mathrm{d}v(x) \ge I(W^*) - sD, \tag{5.10}$$

for all $W \in \mathcal{W}(S)$ and for some value of $s \le 0$. The function ι is the information density and is given by

$$\iota(x,y;W) = \log \frac{\mathrm{d}W(y|x)}{\mathrm{d}w(y;W)}. \tag{5.11}$$

For the compactness of $\mathcal{W}_D$, the following theorem can be stated:

Theorem 5.2. *Suppose that $\phi(x,y) = g(f(x,y))$, where there is a one-to-one correspondence between (x,y) and $(x, f(x,y))$ and f is a continuous function. If g is continuous, and $g(x) \to \infty$ for $x \to -\infty$ and $x \to \infty$, then $\mathcal{W}_D$ is compact in the weak topology of probability.*

The proof is presented in Section 5.4.2. Like Theorem 3.3, $g(x)$ can be redefined for $x \notin S$ and not change the solution to the R-D problem.

5.2.1 Optimality Condition

The optimality condition for R-D problems are well known if W has either a conditional probability mass function (PMF) or a conditional PDF [88]. However, to the best of the author's knowledge, it was first shown for general W in [24]. It is given by the following corollary:

Corollary 5.3. *The optimality condition* (5.9) *is equivalent to*

$$\frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} \geq \lambda_0(x)e^{s\phi(x,y)}, \quad x \in S_X,\ y \in \mathcal{Y}, \tag{5.12}$$

$$\frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} = \lambda_0(x)e^{s\phi(x,y)}, \quad x \in S_X, y \in S_Y, \tag{5.13}$$

where $\lambda_0(x) = \big(\int e^{s\phi(x,y)}\mathrm{d}w(y;W^*)\big)^{-1}$ *and* S_Y *is the set of points of increase for the marginal CDF* $w(\cdot;W^*)$.

The proof is presented in Section 5.4.

In addition to the optimality condition, the following can be stated:

Corollary 5.4. *The optimality condition* (5.9) *implies*

$$\int \lambda_0(x)e^{s\phi(x,y)}\mathrm{d}v(x) = 1, \qquad y \in S_Y. \tag{5.14}$$

The proof is shown in Section 5.4.3. Based on R-D problems with discrete output letters, another condition can be postulated:

$$\int \lambda_0(x)e^{s\phi(x,y)}\mathrm{d}v(x) \leq 1, \qquad y \in \mathcal{Y}, \tag{5.15}$$

where $\mathcal{Y} = \{y \in \mathbb{R} : (x,y) \in S,\ x \in \mathbb{R}\}$. However, no satisfactory proof has been produced yet. Nevertheless, (5.14) is sufficient to prove that certain R-D problems admit a discrete output even when it is allowed to be continuous.

5.3 Discrete Reconstruction Alphabets

Here, it is shown that the resulting reconstruction alphabet is discrete for two R-D problems. The first is the classical Gaussian distributed input with the squared error distortion measure, but the amplitude of the reconstruction alphabet is restricted. The second is a GIG distributed input with a certain distortion measure based on the energy of the GIG neuron model (see Chapter 3).

5.3.1 Gaussian Input Source with Output Amplitude Constraint

Let the DMS be distributed as a Gaussian distribution with mean 0 and variance σ^2. Let the distortion function be given by the square error, $\phi(x, y) = (x - y)^2$. The change to this classic problem is this: suppose that $S = \mathcal{X} \times \mathcal{Y}$ where $\mathcal{X} = \mathbb{R}$ and $\mathcal{Y} = [-L, L]$, $L > 0$. What is $R(D)$?

First note that ϕ satisfies the conditions in Theorem 3.3, hence $\mathcal{W}_D$ is compact and $R(D)$ is attainable. Assume that $D > D_{\min}$. The points of increase of the solution, S_Y, contains either an infinite or finite number of points. Assume the former is true. By the Bolzano-Weierstrass theorem, S_Y has a limit point in $\mathcal{Y}$ [76]. Observe condition (5.14) and note that the left-hand side (LHS) is well defined for complex y where the principle branch of the logarithm is used. Hence, for a complex ζ,

$$\int \lambda_0(x) e^{s\phi(x,\zeta)} \mathrm{d}v(x) = 1, \quad \zeta \in Z_Y, \tag{5.16}$$

where Z_Y is S_Y embedded into $\mathbb{C}$. The LHS is analytic with respect to ζ. So it is satisfied over a set with a limit point, and by the identity theorem, both sides are equal over the entire complex plane [76]. Therefore, (5.14) is satisfied for $y \in \mathbb{R}$. Fix also arrived at this conclusion by the use of Liouville's theorem [87].

Since X is Gaussian distributed, its PDF and PDF conditioned on Y exist. Let those PDF's be p and P, respectively. It can be shown that $\frac{\mathrm{d}W(y|x)}{\mathrm{d}w(y;W)} = \frac{P(x|y;W)}{p(x)}$ for $x \in S_X$ and

$y \in S_Y$. So P can be obtained:

$$P(x|y;W^*) = p(x)\frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} \tag{5.17}$$

$$= \sqrt{\frac{2|s|}{\pi}}e^{s(x-y)^2}, \tag{5.18}$$

where the second equality is obtained from the fact that condition (5.14) is satisfied for $y \in \mathbb{R}$. However, there is a contradiction. Since p is a Gaussian distribution, w has to be the CDF of a Gaussian distribution. This cannot be the case as $\mathcal{Y}$ is bounded. Thus, S_Y cannot contain infinite points and must be discrete.

For the case $D = D_{\min}$, it is clear that the following CDF achieves $D_{\min}$,

$$W^*(y|x) = \begin{cases} \Theta(y+L), & x < -L \\ \Theta(y-x), & -L \le x \le L \\ \Theta(y-L), & x > L, \end{cases} \tag{5.19}$$

where Θ is the unit step function. This is the only conditional CDF that achieves $D_{\min}$. For any $x \in \mathcal{X}$, $W(y|x)$ is the CDF of a discrete RV. However, the resulting output CDF is

$$w(y;W^*) = \begin{cases} 0, & y < -L \\ v(y), & -L \le y < L, \\ 1, & y \ge L. \end{cases} \tag{5.20}$$

Its points of increase is the entire $\mathcal{Y}$. This CDF has both continuous and discrete parts. Also, the MI for W^* is given by $I(W^*) = \infty$ since for any $-L \le x \le L$, the output value is exactly the input. The RB X is continuous, so infinite bits (or nats) has been transmitted. Since infinite bits is transmitted with a non-zero probability, the average transmitted MI is infinite. Hence, $R(D_{\min}) = \infty$. This illustrates that Theorem 5.1 only applies for $D > D_{\min}$ and finite

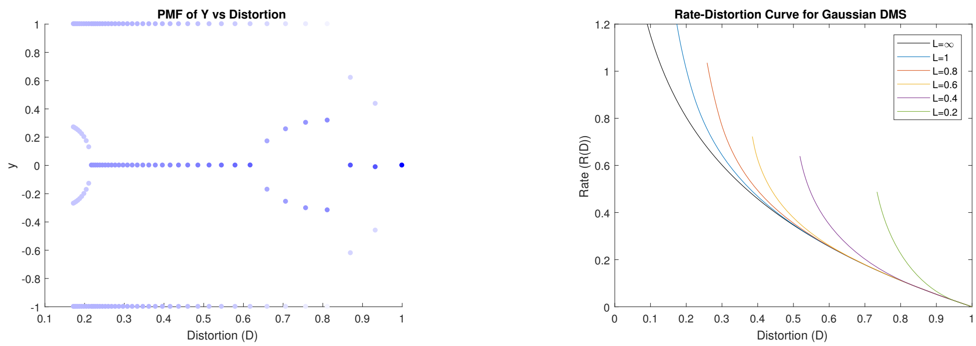

(a) Probability Mass Functions (b) Rate-Distortion Functions

Figure 5.1: (a) For $L = 1$ and $\sigma^2 = 1$, the PMF of Y is plotted with y as the ordinate an the probabilities as the color and distortion as the abscissa. The darker colors indicate higher probabilities. For $D > D_{\min}$, four atoms first appear with two at the edges of $\mathcal{Y}$. As D increases, the middle two atoms converge and eventually split again. As $D = \sigma^2$ is approached, all of the atoms converge to $y = 0$. (b) The R-D curves are plotted for $\sigma^2 = 1$ and various values of L. Although only a part of the curve is shown, each curve actually begins at $(D_{\min}, \infty)$ and descend in a convex manner to $(\sigma^2, 0)$. As L decreases, $D_{\min}$ decreases and for a fixed value of R, the distortion increases.

$R(D)$.

The infimum of the average distortion value where $R(D) = 0$ is given by $D_{\max} = \inf_y \int \phi(x, y) \mathrm{d}v(x)$. In this case, it is the same for classical problem, which is given by $D_{\max} = \sigma^2$.

The location and weight of the discrete points were obtained numerically and is plotted in Figure 5.1. The R-D curve is also given in Figure 5.1. As a comparison, the solution to the classic R-D problem, where $L = \infty$, is also plotted.

The value of $D_{\min}$ depends on L. As L increases, $D_{\min}$ approaches 0. The R-D curve begins at infinity at $D_{\min}$ and decreases in a convex fashion to 0 at $D_{\max} = \sigma^2$. For smaller values of L, the distortion is larger as the range for the reconstruction letter is decreased. For the evolution of the PMF as D increases, there are three points of interest. The first is that the number of masses changed from four to three as the middle two converge. Then the middle mass splits to become four masses. The final point of interest is when the two middle mass combined again and the masses at the end of $\mathcal{Y}$ start moving in towards the middle. Eventually, all the masses combine at $y = 0$ at Dmax.

5.3.2 GIG Distributed Memoryless Input Source

Let Λ be distributed by the generalized inverse Gaussian (GIG) distribution, given by

$$p_{\mathrm{GIG}}(\lambda) = M(\alpha, \beta, \gamma)^{-1}\lambda^{\alpha-1}\exp(-\beta/\lambda - \gamma\lambda), \quad \lambda > 0, \tag{5.21}$$

where M was defined in (2.6). Let the distortion measure be

$$\phi_{\mathrm{GIG}}(\lambda, t) = z - a\log(\lambda t) + b/\lambda t + c\lambda t, \tag{5.22}$$

where,

$$b \geq 0, c > 0 \text{ if } a > 0 \tag{5.23}$$

$$b > 0, c \geq 0 \text{ if } a < 0 \tag{5.24}$$

$$b > 0, c > 0 \text{ if } a = 0. \tag{5.25}$$

To see how this can be a distortion measure, consider the following channel

$$T = \frac{U}{\Lambda}, \tag{5.26}$$

where Λ is the input, T is the output, and U is the multiplicative noise and is independent of Λ. This is exactly the GIG neuron channel of Chapter 3. The noise U is responsible for the distortion. Since $U = \Lambda T$, it is reasonable for the distortion function to be a function of λt, which is the noise for having input λ and output t.

Now consider the function

$$g_{\mathrm{RD}}(u) = z - a\log(u) + b/u + cu, \quad u > 0, \tag{5.27}$$

where a, b, and c obeys (5.25). This function is similar to the conditional energy function g of Chapter 3. As shown previously, this function has a minimum given by

$$u^* = \frac{a + \sqrt{a^2 + 4bc}}{2c}. \tag{5.28}$$

The value of z can be chosen so that the minimum of g_{RD} is 0, hence

$$z = -cu^* - b/u^* + a\log(u^*). \tag{5.29}$$

Also note that with this choice for z, the function g_{RD} is non-negative. The value $g_{\text{RD}}(U)$ is a viable measure of deviance of U from the value u^*. Hence, $g_{\text{RD}}(\lambda t)$ is a viable distortion measure between λ and t and is given by (5.22). No distortion occurs when $t = u^*/\lambda$.

Like the C-C problem in Chapter 3, it is more convenient to work with the following RV's: $X = \log \Lambda$ and $Y = -\log T$. The input distribution is then given by

$$p(x) = M(\alpha, \beta, \gamma)^{-1} \exp(-\alpha x - \beta e^{-x} - \gamma e^{x}), \tag{5.30}$$

which is an exponentiated GIG distribution. The distortion function is then

$$\phi(x, y) = z - a(x - y) + be^{-(x-y)} + ce^{x-y}. \tag{5.31}$$

Note that in order to make the minimum distortion 0, the value of z remains the same as in (5.29). Now the distortion is 0 when $x - y = n^*$, where $n^* = \log u^*$. Let the set of allowable input-output pair be $S = \mathbb{R}^2$, which is a closed set.

The distortion ϕ satisfies the conditions of Theorem 5.2. Hence, $\mathcal{W}_D$ is compact and $R(D)$ is achievable. For $D = D_{\min} = 0$, the only element of $\mathcal{W}_D$ is

$$W^*(y|x) = \Theta(y - x - n^*). \tag{5.32}$$

Since in this case the channel is deterministic, the transmitted MI is infinite. Thus, $R(D_{\min}) = \infty$. The output PDF is $q(y; W^*) = p(y - n^*)$. The output is continuous in this case.

For $D > D_{\min}$, S_Y either has an infinite number of points in some finite interval or a finite number in every interval. Assume that the former is true. The LHS of condition (5.15) is well defined for complex y and is analytic with respect to the complex y. Then using the same identity theorem argument as in the previous example (Section 5.3.1), the condition (5.15) is satisfied for all $y \in \mathbb{R}$. This implies that the PDF of X conditioned on Y is

$$P(x|y; W) = M(a_s, b_s, c_s)^{-1} \exp(a_s(x - y) - b_s e^{-(x-y)} - c_s e^{x-y}), \tag{5.33}$$

where $a_s = -sa$, $b_s = -sb$, and $c_s = -sc$. In order to get w, the following relationship can be stated

$$p(x) = \int P(x|y; W) \mathrm{d}w(y; W). \tag{5.34}$$

As shown in Chapter 3, the only value of w that satisfies this equation has as its characteristic function (CF)

$$\rho_Y(\omega) = \frac{M(\alpha, \beta, \gamma)}{M(a_s, b_s, c_s)} \frac{M(a_s + j\omega, b_s, c_s)}{M(\alpha + j\omega, \beta, \gamma)}. \tag{5.35}$$

As discussed in Chapter 3, this may not be a valid CF because its inverse may not be a non-negative function. In particular, if $\gamma = 0$ and $c > 0$ (or similarly $\beta = 0$ and $b > 0$), then the inverse of ρ_Y contains a Bessel function of the first kind, which is an oscillatory function with negative values. This is a contradiction and S_Y cannot contain infinite points in some interval. In other words, S_Y must contain at most countable discrete points. The location and probabilities of the points were calculated numerically and is plotted in Figure 5.2. Numerical integrations were approximated with the Gauss-Laguerre quadrature. The R-D curve is also plotted in Figure 5.2.

The R-D curve starts at infinity for $D = 0$ and decreases in a convex fashion towards to 0 at about $D = 0.6$. As for the PMF, the number of masses decrease as D increases. The masses seem to to just die out rather than combine like the Gaussian example in Section

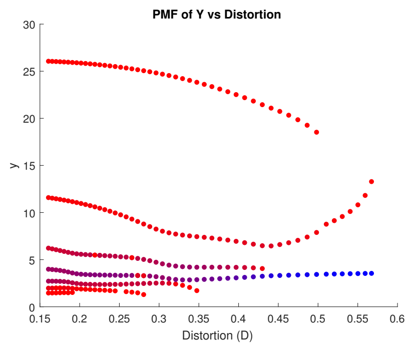

(a) Probability Mass Functions

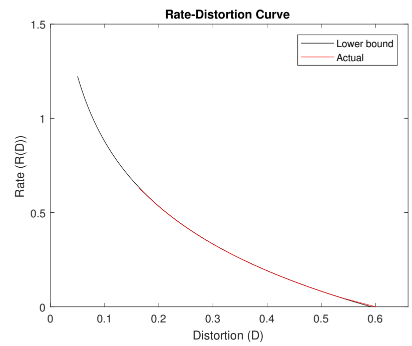

(b) Rate-Distortion Function

Figure 5.2: (a) For each value of the distortion, the PMF of Y is plotted with y as the ordinate an the probabilities as color. The more blue content indicates a higher probability and for visibility, low probabilities are indicated by the color red. Here, $\alpha = 5, \gamma = 20, a = b = c = 1$ and z is chosen so that $D_{\min} = 0$. For $D > D_{\min}$, there are four mass points that initially move closer together as D is increased. However, some of the atoms disappear as D keeps increasing, until eventually, only one atom remain. (b) Only part of the R-D curve is plotted, but it would begin at $(D_{\min}, \infty)$ and decrease in a convex fashion towards the abscissa. There is not much difference between the actual R-D curve and the lower bound.

5.3.1. As D increases, one mass point becomes dominant and remains that way until at $D_{\min}$, where it is the only mass point that survives. The path of the masses as D increases do not follow a predictable trajectory, but seems to follow a smooth path. Though it has not yet been proven that the number of mass points are finite, numerical results seem to suggest that. This would make it consistent with the C-C result for the GIG neuron channel in Chapter 3.

Lower Bound on the Rate-Distortion Curve

Similar to the C-C problem, the lower bound of the R-D problem with a GIG input source can be obtained. Based on (5.33), the conditional PDF of Λ given T is

$$P_{\text{GIG}}(\lambda|t; W) = M(a_s, b_s, c_s)^{-1} t^{a_s} \lambda^{a_s - 1} \exp(-b_s/t\lambda - c_s t\lambda). \tag{5.36}$$

Since this conditional PDF allows the output PDF q with negative values, using this conditional PDF gives a lower bound on the MI. The average distortion is then given by

$$D_b = \iint q_T(t) P_{\mathrm{GIG}}(\lambda|t)\phi_{\mathrm{GIG}}(\lambda, t)\mathrm{d}\lambda\mathrm{d}t \tag{5.37}$$

$$= z - a\frac{\frac{\partial}{\partial a_s}M(a_s, b_s, c_s)}{M(a_s, b_s, c_s)} + b\frac{M(a_s - 1, b_s, c_s)}{M(a_s, b_s, c_s)} + c\frac{M(a_s + 1, b_s, c_s)}{M(a_s, b_s, c_s)}, \tag{5.38}$$

where q_T is the inverse CF of (5.35). Using the identities in Appendix B, this can be simplified to

$$D_b = z - \frac{\partial}{\partial s}\left[\log\left(\frac{b_s}{c_s}\right)^{a_s/2} K_{a_s}(2\sqrt{b_s c_s})\right]. \tag{5.39}$$

A similar derivation was done for the the average energy of the GIG neuron channel in Chapter 3.

Using this assumption, it can be shown that the lower bound on the rate is given by

$$R_b = \log\frac{M(\alpha, \beta, \gamma)}{M(a_s, b_s, c_s)} + \frac{\partial}{\partial u}\left[\log\frac{M(a_s, b_s, c_s)}{M(\alpha, \beta, \gamma)}\right]_{u=1}. \tag{5.40}$$

This derivation is similar to the one for the upper bound of the constrained capacity of the GIG neuron channel in Section 3.4.3 in Chapter 3. In fact, the lower bound for the rate is the same expression as the upper bound of the C-C problem with the roles of (α, β, γ) and (a, b, c) reversed. The graph (D_b, R_b) describes the lower bound of the R-D curve. The bound is tight for $\gamma = 0$ and $C = 0$, or $\beta = 0$ and $B = 0$ because the reconstruction alphabet is continuous in this case. The lower bound is plotted for the numerical example in Figure 5.2. It seems that the actual R-D curve follows the lower bound closely.

5.4 Proofs

5.4.1 Proof of Theorem 5.1

Proof. A minimum exists if $\mathcal{W}_D$ is convex and compact and I is lower semicontinuous [72]. It is clear that $\mathcal{W}_D$ is convex and by assumption, it is compact. For Borel measures, I is lower semicontinuous (LSC) [89]. Since X and Y are real, $F \in \mathcal{F}_D$ is a Borel measure. Hence I is LSC and the minimum exists.

The mutual information I is a convex function [88]. Also, Φ is a convex function. For any $D > D_{\min}$, there exists $W \in \mathcal{W}_D$ such that $\Phi(W) < D$ because $D_{\min}$ is the infimum of the possible average distortion. Then by [72], if $D > D_{\min}$ and $R(D)$ is finite,

$$R(D) = \min_{W \in \mathcal{W}} J(W), \tag{5.41}$$

where $J(W) = I(W) - s\Phi(W)$ for some $s \leq 0$ [72]. Also $s\Phi(W^*) = sD$. It is easy to show that J is Gateaux differentiable everywhere on $\mathcal{W}$. The Gateaux differential of J at W with increment f is

$$\delta J(W; f) = \lim_{\epsilon \to 0} \frac{1}{\delta}(J(W + \epsilon f) - J(W)). \tag{5.42}$$

If W^* minimizes J, δJ satisfies [72, Sec. 7.4 Th. 2],

$$\delta J(W^*; W - W^*) \geq 0, \tag{5.43}$$

for all $W \in \mathcal{W}$. Then using $s\Phi(W^*) = sD$ yields (5.9). $\square$

5.4.2 Proof of Theorem 5.2

This proof is similar to the proof of Theorem 3.3, but with more nuance.

Proof. It is first shown that $\mathcal{W}_D$ is tight, which makes it relatively compact by Prokhorov's theorem [69]. Then it is shown that $\mathcal{W}_D$ is a closed set, which makes it sequentially compact.

Finally, since sequentially compact sets are compact in a metric space, it can be concluded that $\mathcal{W}_D$ is compact.

First, define the random variable $Y' = \xi(X, Y)$. Let $g_{\min}$ be a monotically increasing function over $[0, \infty)$ such that $g_{\min}(u) \leq \min\{g(u), g(-u)\}$ and $\lim_{u\to\infty} g(u) = \infty$. Let S' be the corresponding set of allowable input-output pairs where X is the input and Y' is the output. Let the CDF of Y' given X be W' and let the joint CDF of X and Y' be F'.

Let $\mathcal{K} = [-k, k] \times [-\ell, \ell]$ for $k > 0$ and $\ell > \ell^*$, where $\ell^* = \max\{\ell > 0 : g_{\min}(\ell) = 0\}$. Note that ℓ^* must be finite. Then,

$$\mu_{W'(\cdot|x)}([-\ell, \ell]^c) \stackrel{(d)}{\leq} \frac{\mathbb{E}[g_{\min}(|U|)|X = x]}{g_{\min}(\ell)} \stackrel{(e)}{\leq} \frac{\mathbb{E}[g(U)|X = x]}{g_{\min}(\ell)}. \tag{5.44}$$

Markov's Inequality was applied for (d), and (e) arises from the definition of $g_{\min}$. Then

$$\int \mu_{W'(\cdot|x)}([-\ell, \ell]^c)\mathrm{d}v(x) \leq \frac{D}{g_{\min}(\ell)}. \tag{5.45}$$

Let δ be a value such that $\mu_v([-k, k]) \geq 1 - \delta$. Then,

$$\begin{aligned}
1 - \frac{D}{g_{\min}(\ell)} &< \int \mu_{W'(\cdot|x)}([-\ell, \ell])\mathrm{d}v(x) && (5.46)\\
&= \int_{[-k,k]} \mu_{W'(\cdot|x)}([-\ell, \ell])\mathrm{d}v(x) + \int_{[-k,k]^c} \mu_{W'(\cdot|x)}([-\ell, \ell])\mathrm{d}v(x) && (5.47)\\
&< \mu_{F'}(\mathcal{K}) + \delta. && (5.48)
\end{aligned}$$

Hence,

$$\mu_{F'}(\mathcal{K}^c) \leq \frac{D}{g_{\min}(\ell)} + \delta. \tag{5.49}$$

For any value of $\epsilon > 0$, we can pick a finite value of k and ℓ such that the right-hand side (RHS) above is smaller than ϵ. Since $\mathcal{K}$ is compact in $\mathcal{X} \times \mathcal{Y}'$, $\mathcal{K}$ is also compact in $\mathcal{X} \times \mathcal{Y}$ because ξ is continuous. Therefore $\mathcal{W}_D$ is tight.

To show that $\mathcal{W}_D$ is closed, let $\{W^{(n)}\}$ be a sequence in $\mathcal{W}_D$ whose limit is $W^{(0)}$. The conditional CDF $W^{(0)}$ is an element of $\mathcal{W}$ by relative compactness of $\mathcal{W}_D$. For $\ell > 0$, define the operator

$$\Phi_\ell(W) = \iint \min(\phi(x,y), \ell) \mathrm{d}W(y|x) \mathrm{d}v(x). \tag{5.50}$$

Note that $\min(\phi(x,y), \ell)$ is a continuous bounded function. Also note that $\Phi_\ell(W) \leq \Phi(W) \leq D$ for any $\ell > 0$ and $W \in \mathcal{W}_D$. By weak convergence, $\Phi_\ell(W^{(n)}) \to \Phi_\ell(W^{(0)})$ for any $\ell > 0$. Note that $\Phi_\ell(W^{(0)})$ is increasing in ℓ. By the monotone convergence theorem, $\Phi_\ell(W^{(0)}) \to \Phi(W^{(0)})$. But since $\Phi_\ell(W^{(0)}) \leq D$, $\Phi(W^{(0)}) \leq D$. Hence $\mathcal{W}_D$ is closed and is sequentially compact.

Finally, since the weak topology is metrizable, it is a metric space [69]. Hence $\mathcal{W}_D$ is compact. □

5.4.3 Proof of Corollary 5.3

Proof. To simplify, let $w^* = w(\cdot; W^*)$. To prove that (5.12) and (5.13) implies (5.9), it suffices to show that

$$\int \log \lambda_0(x) \mathrm{d}v(x) = I(W^*) - sD. \tag{5.51}$$

Recall that $R(D)$ is finite by assumption, hence $\mu_{W^*(\cdot|x)} \ll \mu_{w^*}$, μ_v-a.e. Also, note that the right-hand sides (RHS's) of (5.12) and (5.13) are strictly greater than 0. Hence, $\frac{\mathrm{d}w^*}{\mathrm{d}W^*(\cdot|x)} = \left(\frac{\mathrm{d}W^*(\cdot|x)}{\mathrm{d}w^*}\right)^{-1}$ is finite for $x \in S_X$, which implies $\mu_{w^*} \ll \mu_{W^*(\cdot|x)}$, μ_v-a.e. Consequently,

$$\begin{aligned}
\log \lambda_0(x) &= -\log \int e^{s\phi(x,y)} \mathrm{d}w^*(y) \\
&= -\log \int \frac{\mathrm{d}w^*(y)}{\mathrm{d}W^*(y|x)} e^{s\phi(x,y)} \mathrm{d}W^*(y|x) \\
&\overset{(a)}{\geq} -\int \log \left(\frac{\mathrm{d}w^*(y)}{\mathrm{d}W^*(y|x)} e^{s\phi(x,y)} \right) \mathrm{d}W^*(y|x) \\
&= \int \log \left(\frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w^*(y)} e^{-s\phi(x,y)} \right) \mathrm{d}W^*(y|x), \quad \mu_v\text{-a.e.},
\end{aligned} \tag{5.52}$$

where Jensen's inequality was used for (a). However, (a) is satisfied by equality if and only if (5.13) is true. Hence, (a) is satisfied with an equality and integrating both sides with respect to v yields (5.51).

Now the converse is proven. Suppose that (5.12) is false and (5.13) is true. There must exist $y^* \in \mathcal{Y}$, such that the set

$$\mathcal{G} = \left\{ x \in S_X : \frac{\mathrm{d}W^*(y^*|x)}{\mathrm{d}w(y^*; W^*)} < \lambda_0(x) e^{s\phi(x,y^*)} \right\}. \tag{5.53}$$

has positive measure, i.e., $\int_{\mathcal{G}} \mathrm{d}v(x) > 0$. Let $y^+ \in S_Y$. Then let

$$W(y|x) = \begin{cases} \Theta(y - y^*), & x \in \mathcal{G} \\ \Theta(y - y^+), & x \in S_X \backslash \mathcal{G}, \end{cases} \tag{5.54}$$

where Θ is the unit step function and $\backslash$ is the set difference operation. Consequently,

$$\begin{aligned} &\iint_S [i(x, y; W^*) - s\phi(x, y)] \mathrm{d}W(y|x) \mathrm{d}v(x) \\ &\quad = \int_{\mathcal{G}} \int_{\mathcal{Y}} [i(x, y; W^*) - s\phi(x, y)] \mathrm{d}\Theta(y - y^*) \mathrm{d}v(x) \\ &\qquad + \int_{\mathcal{X} \backslash \mathcal{G}} \int_{\mathcal{Y}} [i(x, y; W^*) - s\phi(x, y)] \mathrm{d}\Theta(y - y^+) \mathrm{d}v(x) \\ &\quad \overset{(b)}{<} \int_{\mathcal{G}} \log \lambda_0(x) \mathrm{d}v(x) + \int_{\mathcal{X} \backslash \mathcal{G}} \log \lambda_0(x) \mathrm{d}v(x) \\ &\quad \overset{(c)}{=} I(W^*) - s\Phi(W^*). \end{aligned} \tag{5.55}$$

The inequality (b) is due to the definition of $\mathcal{G}$ and (5.13). For (c), the argument is similar to (5.52) in the direct proof. Equation (5.55) contradicts (5.9). Hence, (5.12) must be true.

Now assume that (5.12) is true and (5.13) is false. Condition (5.12) implies that

$$\int \log \lambda_0(x) \mathrm{d}v(x) \leq I(W^*) - sD. \tag{5.56}$$

However, repeating (5.52) implies that

$$\int \log \lambda_0(x) \mathrm{d}v(x) \geq I(W^*) - sD. \tag{5.57}$$

As a consequence, both sides must be equal. However (a) in (5.52) cannot be satisfied with equality because this requires (5.13) to be true, which contradicts our premise. Both (5.12) and (5.13) must be true. □

5.4.4 Proof of Corollary 5.4

Proof. It suffices to show that

$$\int \frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} \mathrm{d}v(x) = 1, \quad \mu_w\text{-a.e.}, \tag{5.58}$$

where μ_w-a.e. means almost everywhere with respect to μ_w. Then, integrate (5.13) with respect to $\mathrm{d}w$ to get (5.14).

By R-N theory, for any set $A \in \mathcal{B}(\mathbb{R})$, i.e., the Borel set of $\mathbb{R}$,

$$\mu_{W^*(\cdot|x)}(A) = \int_A \frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} \mu_{w(\cdot;W^*)}(\mathrm{d}y). \tag{5.59}$$

Integrate both sides with respect to v and use Tonelli's theorem to switch the order of integration of the RHS,

$$\mu_{w(\cdot;W^*)}(A) = \int_A \int_{\mathcal{X}} \frac{\mathrm{d}W^*(y|x)}{\mathrm{d}w(y;W^*)} \mathrm{d}v(x) \mu_{w(\cdot;W^*)}(\mathrm{d}y). \tag{5.60}$$

On the RHS, integral is with respect to the same measure as the left-hand side (LHS). Hence, by the R-N theory, the integrand must be equal to 1, μ_w-a.e., which is the desired result. □

5.5 Conclusion

In this chapter, it was proven that the solution to R-D problems exist when $\mathcal{W}_D$ is compact. Furthermore, the well-known optimality conditions for R-D problems were extended to include output RV's that are neither discrete nor continuous. Then it was shown that it is possible to have a discrete reconstruction alphabet for the R-D problem. This was first demonstrated on a modification of the classic R-D problem where the DMS is Gaussian distributed and the distortion is the mean square. An additional amplitude constraint for the reconstruction alphabet was imposed. This produced a discrete reconstruction alphabet.

Then a R-D problem based on the GIG neuron channel was also proposed. A GIG distributed DMS with a certain distortion function was used. It was shown that for certain parameters, the reconstruction alphabet is discrete. This was a result similar to the GIG neuron channel in Chapter 3, where the optimal input distribution is discrete for some parameter sets. This shows the interconnection of C-C and R-D problems.

Chapter 6

Conclusion

A model for energy efficient cortical neurons was proposed in this dissertation. The GIG neuron model was developed for cortical neurons in the sensory cortex. Let η be one of such neuron type. The model describes the PSP of η with the GIGHT stochastic diffusion. Such a diffusion has the advantage of taking into account the upswing exhibited by PSP as it nears the threshold. Another advantage of the GIGHT diffusion is it has a closed form equation for the FHT distribution of the diffusion. Thus, a neuron "channel" can be characterized. This channel is a multiplicative noise channel where the noise is distributed as a GIG distribution.

The GIG neuron model has three parameters: α, β, and γ. Assume that the threshold and infinitesimal variance are fixed. The α parameter controls the attraction of the GIGHT diffusion to the threshold. The lower the value of α, the more attraction there is to the threshold. The β parameter is determined from the threshold and the infinitesimal variance. The γ parameter controls a constant drift component of the GIGHT diffusion.

Two estimation techniques were developed for estimating the parameters of the GIG neuron model from a realization of the GIGHT diffusion or an intracellular recording of the neuron: the pMLE and the pLFSE. Simulations show that the pLFSE gives a more unbiased estimate for the α and γ parameters. Also, when the constant drift of the GIGHT is dominant, the pMLE performs as well as the pLFSE. Otherise, the pLFSE performs better than the pMLE because the pLFSE is less biased. Finally, a technique for estimating the input intensity from the intracellular recording was also developed. The estimator may be biased, but it appears to be consistent.

The next step is to evaluate the energy efficiency of the GIG neuron model. Given a fixed energy budget, how much information can η deliver to its target? Thus, two things are needed: an energy model and a measure of information.

The energy model was determined from assumptions about η's energy costs, such as a fixed cost of propagating an AP, linear metabolic cost, and logarithmic biological clock cost. As for information, Shannon's MI was used as a measure of information. Since η is energy efficient, it is assumed that it maximizes MI given an average energy constraint. In information theoretic terms, the constrained capacity of the GIG neuron channel is what is sought after.

It turns out that for some parameter sets, the input distribution that achieves the constrained capacity is a discrete distribution with a finite number of mass points. This surprising result has implications on how the optimal network should be designed. Since the network has influence over the input intensity, this implies that in order for η to perform most efficiently, the network must exist in discrete states.

Next, optimization of the parameters of the GIG neuron model was addressed. The input distribution is assumed to be known. For a fixed energy budget, the parameter set that yields the highest MI is sought after. However, it turns out that this would produce infinite MI. Hence, a variance constraint was added. The variance constraint prevents the GIG neuron channel from approaching the noiseless case, i.e., approaching a perfect decoding of the output IPI's to the input rate.

The result was demonstrated numerically for a special case of the GIG neuron model. For a fixed variance constraint, the MI is a concave function of energy. For fixed energy, the MI has an inverse relationship with the variance constraint.

With this result, it is possible to envision a double matching problem for η: the matching of the channel to the source and the matching of the source to the channel. This could be iteratively determined. Assume a fixed energy budget. First, the source is chosen to maximize the MI. Then, with the source fixed, the channel parameters can be selected to maximize MI. The process can be repeated. It is not yet known that such a process will converge to a maximum MI, but it is an idea worth exploring.

Finally, the discrete result from the GIG neuron channel was extended to a parallel R-D problem. First, the existence of the solution to R-D problems for certain distortions are proven. Then, the optimal conditions for $R(D)$ and the achieving test channel was proven. Let the source be distributed as a GIG distribution. Then for a certain distortion function, it was proven that the resulting reconstruction alphabet can be discrete. This was then demonstrated by a numerical solution.

6.1 Future Research

Once the individual neurons are better understood, the natural next step of the research is to understand how the neurons work together in the network. Different parts of the nervous system have different purposes, but a possible purpose of the sensory cortex is to reduce the statistical dependence of each input line to the NN [90–92]. The purpose of this is to reduce complexity in the input signal and make information processing easier for the higher regions of the brain. With energy efficient neurons, how does energy play into the process of reducing information complexity? That is one possible path of future research.

Another question to be addressed is how should the neurons be connected within one region. Many algorithms exist for updating synaptic weights in ANN's [93]. However, what can be learned from how NN's actually create their synaptic connections and weights. Other than the sheer number of number of neurons in ANN's versus NN's, the neurons in NN's have relatively fewer connections per neurons than ANN's. One possibility is that for NN's, a connection between neurons needs energy to maintain. Thus, as an energy saving strategy, unimportant connections are dropped. In current implementation of ANN's, energy is not needed to maintain connections, but the synaptic weights require memory. What is the memory-performance tradeoff of having fewer connections? Furthermore, if a computer based on neural elements are built, such a connection may require energy to maintain. Thus, energy can be saved in such a case.

Another high level approach is to understand how interactions between different regions can promote energy efficiency. Consider a classification problem with two sensors, e.g. sight and sound. Suppose the object to be identified is the type of bird. By using sound and vision, the type of bird can be identified. However, if the bird's view is blocked, the visual signal is not as valuable as the sound signal. Hence, less attention can be paid to the visual cortex. To save energy, a lower energy budget can be assigned to this region since it yields low information. The higher up region of the brain can assign a lower energy budget to the visual cortex and increase the energy budget of the auditory cortex. Thus, based on the information the higher regions receive, they send a feedback signal to the lower level. This type of feedback can be further studied in order to understand how the network can be designed to be energy efficient.

The challenge of energy efficient neural computation is daunting. Hopefully, this dissertation will begin to answer some important questions about energy efficient neural computation and raise important questions and considerations.

Appendix A

Properties of the Modified Bessel Function of the Second Kind

A.1 Integral Representation

The function $K_\omega(\zeta)$ for complex ω and ζ can be represented by the following integral [73]:

$$K_\omega(\zeta) = \frac{1}{2}\left(\frac{1}{2}\right)^\omega \int_0^\infty \exp(-u - \zeta^2/4u)\frac{\mathrm{d}u}{u^{\omega+1}}, \quad |\angle\zeta| < \frac{\pi}{4}, \tag{A.1}$$

where $\angle\zeta$ is the phase of ζ.

A.2 Hankel Expansion

For complex ω and ζ, $K_\omega(\zeta)$ has the following Hankel's expansion for $\zeta \to \infty$ [73]:

$$K_\omega(\zeta) \sim \sqrt{\frac{\pi}{2\zeta}} e^{-\zeta} \sum_{k=0}^{\infty} \frac{\xi_k(\omega)}{\zeta^k}, \quad |\angle\zeta| < \frac{3\pi}{2} \tag{A.2}$$

where $\xi_0(\omega) = 1$ and

$$\xi_k(\omega) = \frac{1}{k!8^k} \prod_{\ell=1}^{i} (4\omega^2 - (2\ell - 1)^2), \quad k \in \{1, 2, \ldots\}. \tag{A.3}$$

The notation $\sim$ refers to asymptotic equality. Therefore for $\omega = -1/2$ and $\omega = -3/2$ and equivalently for $\omega = 1/2$ and $\omega = 3/2$,

$$K_{-1/2}(\zeta) = K_{1/2}(\zeta) = \sqrt{\frac{\pi}{2\zeta}} e^{-\zeta} \tag{A.4}$$

and

$$K_{-3/2}(\zeta) = K_{3/2}(\zeta) = \sqrt{\frac{\pi}{2\zeta}} e^{-\zeta} \left(1 + \frac{1}{\zeta}\right). \tag{A.5}$$

However, equality holds for all complex ζ in this case [73].

A.3 Asymptotic Forms

There are two asymptotes of interest. The first asymptote is for $\omega \in \mathbb{C}$, as $\zeta \to 0$,

$$K_\omega(\zeta) \sim \begin{cases} \frac{\Gamma(\omega)}{2} \left(\frac{2}{\zeta}\right)^\omega & \mathrm{Re}\{\omega\} > 0, \\ -\log\left(\frac{\zeta}{2}\right) - \xi & \omega = 0, \\ \frac{\Gamma(-\omega)}{2} \left(\frac{2}{\zeta}\right)^{-\omega} & \mathrm{Re}\{\omega\} < 0, \end{cases} \tag{A.6}$$

where Γ is the gamma function, $\mathrm{Re}\{\omega\}$ is the real part of ω and ξ is the Euler-Mascheroni constant [73].

The second asymptote is for a complex ω, as $\zeta \to \infty$ [73],

$$K_\omega(\zeta) \sim K_{1/2}(\zeta) = \sqrt{\frac{\pi}{2\zeta}} e^{-\zeta}. \tag{A.7}$$

A.4 Recurrence Relations

The following recurrence relations can be stated [73],

$$K_{\omega}^{(0,1)}(\zeta) = -\frac{1}{2}(K_{\omega-1}(\zeta) + K_{\omega+1}(\zeta)), \tag{A.8}$$

where $K_{\omega}^{(0,1)}(\zeta) = \frac{\partial}{\partial\zeta}K_{\omega}(\zeta)$

Appendix B

Moments of the Generalized Inverse Gaussian Distribution

Let $U \sim \mathrm{GIG}(\alpha,\beta,\gamma)$, i.e., U is distributed as a GIG distribution. Let $\mathbb{E}$ denote the expectation operator. Recall that

$$M(\alpha,\beta,\gamma) = \begin{cases} 2\left(\frac{\beta}{\gamma}\right)^{\alpha/2} K_\alpha(2\sqrt{\beta\gamma}) & \beta > 0, \gamma > 0 \\ \gamma^{-\alpha}\Gamma(\alpha) & \alpha > 0, \beta = 0, \gamma > 0 \\ \beta^{\alpha}\Gamma(-\alpha) & \alpha < 0, \beta > 0, \gamma = 0. \end{cases} \tag{B.1}$$

The n^{th} moment of U is given by

$$\begin{aligned} \mathbb{E}[U^n] &= M(\alpha,\beta,\gamma)^{-1} \int_0^\infty u^{\alpha+n-1} \exp(-\beta/u - \gamma u) \mathrm{d}u \\ &= \frac{M(\alpha+n,\beta,\gamma)}{M(\alpha,\beta,\gamma)} \end{aligned}$$

To simplify, let $k_n = \mathrm{E}[U^n]$. Evaluating for each case yields

$$k_n = \begin{cases} \left(\frac{\beta}{\gamma}\right)^{n/2} \frac{K_{\alpha+n}(2\sqrt{\beta\gamma})}{K_\alpha(2\sqrt{\beta\gamma})}, & \beta > 0, \gamma > 0 \\ \frac{\Gamma(\alpha+n)}{\Gamma(\alpha)\gamma^n}, & \alpha > 0, \beta = 0, \gamma > 0, \\ \frac{\Gamma(-\alpha-n)\beta^n}{\Gamma(-\alpha)}, & \alpha < 0, \beta > 0, \gamma = 0. \end{cases} \tag{B.2}$$

The log-moment is given by

$$\begin{aligned} \mathbb{E}[\log U] &= \int_0^\infty M(\alpha,\beta,\gamma)^{-1} \log(u) u^{\alpha-1} \exp(-\beta/u - \gamma u) \\ &= M(\alpha,\beta,\gamma)^{-1} \frac{\partial}{\partial\alpha} \int_0^\infty u^{\alpha-1} \exp(-\beta/u - \gamma u) \\ &= \frac{\frac{\partial}{\partial\alpha} M(\alpha,\beta,\gamma)}{M(\alpha,\beta,\gamma)} \\ &= \frac{\partial}{\partial\alpha} \log M(\alpha,\beta,\gamma) \end{aligned} \tag{B.3}$$

Let $k_g = \mathbb{E}[\log U]$. Evaluating for each case yields

$$k_g = \begin{cases} \frac{\alpha}{2} \log \frac{\beta}{\gamma} + \frac{\frac{\partial}{\partial\alpha} K_\alpha(2\sqrt{\beta\gamma})}{K_\alpha(2\sqrt{\beta\gamma})}, & \beta > 0, \gamma > 0 \\ \psi(\alpha) - \log\gamma, & \alpha > 0, \beta = 0, \gamma > 0, \\ \log\beta - \psi(-\alpha), & \alpha < 0, \beta > 0, \gamma = 0, \end{cases} \tag{B.4}$$

where ψ is the digamma function.

The variance of the GIG distribution is given by

$$\begin{aligned} \mathrm{Var}(U) &= \mathbb{E}[U^2] - \mathbb{E}[U]^2 \\ &= k_2 - k_1^2. \end{aligned} \tag{B.5}$$

For $\alpha > 0$ and $\beta = 0$, it can be written as

$$\mathrm{Var}(U) = \frac{\alpha}{\gamma^2}. \tag{B.6}$$

For $\alpha < 0$ and $\gamma = 0$,

$$\mathrm{Var}(U) = \frac{\beta^2}{(\alpha - 1)^2(\alpha - 2)}, \quad \alpha < -2. \tag{B.7}$$

For the other cases,

$$\mathrm{Var}(U) = \frac{\beta}{\gamma}\left[\frac{K_{\alpha+2}(2\sqrt{\beta\gamma})}{K_{\alpha}(2\sqrt{\beta\gamma})} - \left(\frac{K_{\alpha+1}(2\sqrt{\beta\gamma})}{K_{\alpha}(2\sqrt{\beta\gamma})}\right)^2\right] \tag{B.8}$$

Bibliography

[1] L. C. Aiello and P. Wheeler. The expensive-tissue hypothesis: The brain and the digestive system in human and primate evolution. *Current Anthropology*, 36(2):pp. 199–221, 1995.

[2] D. Modha. Private communications, 2009.

[3] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, 1948.

[4] R. G. Gallager. *Information theory and reliable communication*, volume 2. Springer, 1968.

[5] D. M. MacKay and W. S. McCulloch. The limiting information capacity of a neuronal link. *The Bulletin of Mathematical Biophysics*, 14(2):127–135, 1952.

[6] W. Bialek, F. Rieke, R. R. de Ruyter van Steveninck, and D. Warland. Reading a neural code. *Science*, 252(5014):pp. 1854–1857, 1991.

[7] J. E. Niven, J. C. Anderson, and S. B. Laughlin. Fly photoreceptors demonstrate energy-information trade-offs in neural coding. *PLoS Biology*, 5:e116, 03 2007.

[8] I. Nemenman, G. D. Lewen, W. Bialek, and R. R. de Ruyter van Steveninck. Neural coding of natural stimuli: Information at sub-millisecond resolution. *PLoS Computational Biology*, 4(3):e1000025, 03 2008.

[9] William B Levy and Robert A. Baxter. Energy efficient neural codes. *Neural Computation*, 8(3):531–543, April 1996.

[10] V. Balasubramanian, D. Kimber, and M. J. Berry. Metabolically efficient information processing. *Neural Computation*, 13(4):799–815, April 2001.

[11] T. Berger and W. B Levy. A mathematical theory of energy efficient neural computation and communication. *IEEE Transactions on Information Theory*, 56(2):852–874, 2010.

[12] Jie Xing, T. Berger, and T. J. Sejnowski. A berger-levy energy efficient neuron model with unequal synaptic weights. In *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, pages 2964–2968, July 2012.

[13] J. Xing and T. Berger. Energy efficient neurons with generalized inverse gaussian conditional and marginal hitting times. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 1824–1828, July 2013.

[14] J. Xing, T. Berger, M. Sungkar, and William B Levy. Energy efficient neurons with generalized inverse Gaussian conditional and marginal hitting times. *IEEE Transactions on Information Theory*, 61(8):4390–4398, 2015.

[15] L. Kostal and R. Kobayashi. Optimal decoding and information transmission in Hodgkin-Huxley neurons under metabolic cost constraints. *Biosystems*, 136:3 – 10, 2015. Selected papers presented at the Eleventh International Workshop on Neural Coding, Versailles, France, 2014.

[16] A. Hasenstaub, S. Otte, E. Callaway, and T. J. Sejnowski. Metabolic cost as a unifying principle governing neuronal biophysics. *Proceedings of the National Academy of Sciences*, 107(27):12329–12334, jun 2010.

[17] J. E. Niven, M. Vähäsöyrinki, and M. Juusola. Shaker K^+-channels are predicted to reduce the metabolic cost of neural information in drosophila photoreceptors. *Proceedings of the Royal Society of London B: Biological Sciences*, 270(Suppl 1):S58–S61, 2003.

[18] B. Sengupta, M. Stemmler, S. B. Laughlin, and Jeremy E. Niven. Action potential energy efficiency varies among neuron types in vertebrates and invertebrates. *PLoS Computational Biology*, 6(7):e1000840, 07 2010.

[19] T. Berger and M. Sungkar. Barndorff-Nielsen diffusions model neurons in primary sensory cortex. In *Information Theory and Applications Workshop*, February 2016. Manuscript available at: www.ece.virginia.edu/~nitlab/.

[20] Mustafa Sungkar, Toby Berger, and William B Levy. Mutual information and parameter estimation in the generalized inverse Gaussian diffusion model of cortical neurons. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(2):166–182, 2016.

[21] M. Sungkar, T. Berger, and W. B Levy. The Karush-Kuhn-Tucker conditions for a generalized inverse gaussian neural model. In *Information Theory Applications and Workshop*, Feb 2017.

[22] M. Sungkar, T. Berger, and W. B Levy. Capacity achieving input distribution to the generalized inverse gaussian neuron model. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 860–869, Oct 2017.

[23] T. Berger and M. Sungkar. Design of efficient channels with given input statistics. In *IEEE International Symposium on Information Theory*, June 2015.

[24] M. Sungkar and T. Berger. Discrete reconstruction alphabets in discrete memoryless source rate-distortion problems. In *IEEE International Symposium on Information Theory*, 2018. To be presented.

[25] T. F. Weiss. *Cellular Biophysics: Electrical Properties*, volume 2. The MIT Press, 1996.

[26] E. R. Kandel, J. H. Schwartz, and T. M. Jessel. *Principles of Neural Science*. McGraw-Hill, 4 edition, 2000.

[27] A. L. Hodgkin and A. F. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544, 1952.

[28] M. Royeck, M.-T. Horstmann, S. Remy, M. Reitze, Y. Yaari, and H. Beck. Role of axonal $Na_V 1.6$ sodium channels in action potential initiation of CA1 pyramidal neurons. *Journal of Neurophysiology*, 100(4):2361–2380, 2008.

[29] M. H. P. Kole, S. U. Ilschner, B. M. Kampa, S. R. Williams, P. C. Ruben, and G. J. Stuart. Action potential generation requires a high sodium channel density in the axon initial segment. *Nature neuroscience*, 11(2):178–186, 2008.

[30] C. W. Helstrom. *Statistical Theory of Signal Detection*, volume 9. Pergamon Press, 1960.

[31] P. E. Kloeden and E. Platen. *Numerical solution of stochastic differential equations*, volume 23. Springer Science & Business Media, 1992.

[32] G. L. Gerstein and B. Mandelbrot. Random walk models for the spike activity of a single neuron. *Biophysical journal*, 4(1 Pt 1):41, 1964.

[33] S. Cash and R. Yuste. Input summation by cultured pyramidal neurons is linear and position-independent. *Journal of Neuroscience*, 18(1):10–15, 1998.

[34] S. Cash and R. Yuste. Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron*, 22(2):383 – 394, 1999.

[35] R. K. Powers and M. D. Binder. Summation of effective synaptic currents and firing rate modulation in cat spinal motoneurons. *Journal of Neurophysiology*, 83(1):483–500, 2000.

[36] R. Araya, K. B. Eisenthal, and R. Yuste. Dendritic spines linearize the summation of excitatory potentials. *Proceedings of the National Academy of Sciences*, 103(49):18799–18804, 2006.

[37] P. I. M. Johannesma. *Neural Networks: Proceedings of the School on Neural Networks - June 1967 in Ravello*, chapter Diffusion Models for the Stochastic Activity of Neurons, pages 116–144. Springer Berlin Heidelberg, Berlin, Heidelberg, 1968.

[38] R. M. Capocelli and L. M. Ricciardi. Diffusion approximation and first passage time problem for a model neuron. *Kybernetik*, 8(6):214–223, 1971.

[39] L. M. Ricciardi and L. Sacerdote. The Ornstein-Uhlenbeck process as a model for neuronal activity. *Biological Cybernetics*, 35(1):1–9, 1979.

[40] H. C. Tuckwell and D. K. Cope. Accuracy of neuronal inerspike times calculated from a diffusion approximation. *Journal of Theoretical Biology*, 83(3):377–387, 1980.

[41] B. Lindner and A. Longtin. Effect of an exponentially decaying threshold on the firing statistics of a stochastic integrate-and-fire neuron. *Journal of Theoretical Biology*, 232(4):505 – 521, 2005.

[42] P. Lansky and J. P. Rospars. Ornstein-Uhlenbeck model neuron revisited. *Biological Cybernetics*, 72(5):397–406, 1995.

[43] Floyd B. Hanson and H. C. Tuckwell. Diffusion approximations for neuronal activity including synaptic reversal potentials. *Journal of Theoretical Neurobiology*, 2:127–153, 1983.

[44] P. Lansky and V. Lanska. Diffusion approximation of the neuronal model with synaptic reversal potentials. *Biological Cybernetics*, 56(1):19–26, 1987.

[45] E. Bibbona, P. Lansky, and R. Sirovich. Estimating input parameters from intracellular recordings in the feller neuronal model. *Physical Review E*, 81:031916, 2010.

[46] R. Rodriguez and H. C. Tuckwell. A dynamical system for the approximate moments of nonlinear stochastic models of spiking neurons and networks. *Mathematical and computer modelling*, 31(4):175–180, 2000.

[47] H. C. Tuckwell and F. Y. M. Wan. Time to first spike in stochastic hodgkinhuxley systems. *Physica A: Statistical Mechanics and its Applications*, 351(24):427 – 438, 2005.

[48] H. C. Tuckwell, R. Rodriguez, and F. Y. M. Wan. Determination of firing times for the stochastic Fitzhugh-Nagumo neuronal model. *Neural computation*, 15(1):143–159, 2003.

[49] S. Iyengar and Q. Liao. Modeling neural activity using the generalized inverse Gaussian distribution. *Biological Cybernetics*, 77(4):289–295, 1997.

[50] A. D. Barbour, L. Holst, and S. Janson. *Poisson approximation.* Clarendon Press Oxford, 1992.

[51] William B Levy, Toby Berger, and Mustafa Sungkar. Neural computation from first principles: Using the maximum entropy method to obtain an optimal bits-per-joule neuron. *IEEE Transactions on Molecular, Biological and Multi-Scale Communications*, 2(2):154–165, Dec 2016.

[52] E. Schrödinger. Zur theorie der fall-und steigversuche an teilchen mit brownscher bewegung. *Physikalische Zeitschrift*, 16:289–295, 1915.

[53] O. Barndorff-Nielsen, P. Blæsild, and C. Halgreen. First hitting time models for the generalized inverse Gaussian distribution. *Stochastic Processes and their Applications*, 7(1):49 – 54, 1978.

[54] B. Jørgensen. *Statistical properties of the generalized inverse Gaussian distribution*, volume 9. Springer-Verlag, 1982.

[55] Danielle Morel, Chandan Singh, and William B. Levy. Linearization of excitatory synaptic integration at no extra cost. *Journal of Computational Neuroscience*, 44(2):173–188, Apr 2018.

[56] N. Brunel and P. E. Latham. Firing rate of the noisy quadratic integrate-and-fire neuron. *Neural Computation*, 15(10):2281–2306, 2003.

[57] R. Brette and W. Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, 94(5):3637–3642, 2005.

[58] B. L. S. Prakasa Rao. *Statistical inference for diffusion type processes.* Arnold, 1999.

[59] S. Ditlevsen and P. Lansky. Estimation of the input parameters in the Ornstein-Uhlenbeck neuronal model. *Physical Review E*, 71:011907, Jan 2005.

[60] P. Lansky and S. Ditlevsen. A review of the methods for signal estimation in stochastic diffusion leaky integrate-and-fire neuronal models. *Biological Cybernetics*, 99(4):253–262, 2008.

[61] P. Jahn, R. W. Berg, J. Hounsgaard, and S. Ditlevsen. Motoneuron membrane potentials follow a time inhomogeneous jump diffusion process. *Journal of Computational Neuroscience*, 31(3):563–579, 2011.

[62] Q. J. M. Huys, M. B. Ahrens, and L. Paninski. Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, 96(2):872–890, 2006.

[63] Q. J. M. Huys and L. Paninski. Smoothing of, and parameter estimation from, noisy biophysical recordings. *PLoS Comput Biol*, 5(5):1–16, 05 2009.

[64] A. C. Jensen, S. Ditlevsen, M. Kessler, and O. Papaspiliopoulos. Markov chain monte carlo approach to parameter estimation in the Fitzhugh-Nagumo model. *Phys. Rev. E*, 86:041114, Oct 2012.

[65] S. Ditlevsen and A. Samson. Estimation in the partially observed stochastic morris-lecar neuronal model with particle filter and stochastic approximation methods. *The Annals of Applied Statistics*, 8(2):674–702, 2014.

[66] I. Shoji and T. Ozaki. Comparative study of estimation methods for continuous time stochastic process. *Journal of Time Series Analysis*, 18(5):485–506, 1997.

[67] M. Sungkar. Energy-efficient cortical neuron models with gig likelihood functions. Master's thesis, University of Virginia, 2015.

[68] J. G. Smith. The information capacity of amplitude- and variance-constrained sclar gaussian channels. *Information and Control*, 18(3):203 – 219, 1971.

[69] A. N. Shiriaev. *Probability.* Springer, 2 edition, 1996.

[70] M. S. Pinsker. *Information and Information Stability of Random Variables and Processes*. Holden-Day, 1964. Translated to English.

[71] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, Inc., Hoboken, N.J., 2 edition, 2006.

[72] D. G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, 1969.

[73] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, editors. *NIST handbook of mathematical functions*. Cambridge University Press, 2010.

[74] S. Verdu. On channel capacity per unit cost. *IEEE Transactions on Information Theory*, 36(5):1019–1030, Sep 1990.

[75] I. C. Abou-Faycal, M. D. Trott, and S. Shamai (Shitz). The capacity of discrete-time memoryless rayleigh-fading channels. *IEEE Transactions on Information Theory*, 47(4):1290–1301, May 2001.

[76] E. Hille. Analytic function theory. vol. ii, introductions to higher mathematics, ginn and co., boston, mass. *New York-Toronto, Ont*, 1962.

[77] F. Oberhettinger and L. Badii. *Tables of Laplace Transforms*. Springer-Verlag, New York, 1973.

[78] E. Lukacs. Characteristic functions, griffins statistical monographs& courses, no. 5, 1960.

[79] J. Fahs and I. Abou-Faycal. On the finiteness of the capacity of continuous channels. *IEEE Transactions on Communications*, 64(1):166–173, 2016.

[80] M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, and E. Vaadia. Cortical activity flips among quasi-stationary states. *Proceedings of the National Academy of Sciences*, 92(19):8616–8620, 1995.

[81] E. Seidemann, I. Meilijson, M. Abeles, H. Bergman, and E. Vaadia. Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *Journal of Neuroscience*, 16(2):752–768, 1996.

[82] L. M. Jones, A. Fontanini, B. F. Sadacca, P. Miller, and D. B. Katz. Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles. *Proceedings of the National Academy of Sciences*, 104(47):18772–18777, 2007.

[83] A. P. Nikitin, N. G. Stocks, R. P. Morse, and M. D. McDonnell. Neural population coding is optimized by discrete tuning curves. *Phys. Rev. Lett.*, 103:138101, Sep 2009.

[84] J. Niessing and R. W. Friedrich. Olfactory pattern classification by discrete neuronal network states. *Nature*, 465(7294):47–52, May 2010.

[85] Y. Mochizuki and S. Shinomoto. Analog and digital codes in the brain. *Phys. Rev. E*, 89:022705, Feb 2014.

[86] C. E. Shannon. Coding theorems for a discrete source with a fidelity criterion. In *IRE National Convention Record, part 4*, pages 142–163, 1959.

[87] S. L. Fix. Rate distortion functions for squared error distortion measures. In *Annual Allerton Conference on Communication, Control and Computing, 16th*, pages 704–711, Monticello, Ill., 1978.

[88] T. Berger. *Rate distortion theory: A mathematical basis for data compression.* Prentice-Hall, 1971.

[89] E. Posner. Random coding strategies for minimum entropy. *IEEE Trans. Inf. Theory*, 21(4):388–391, 1975.

[90] D. M. Adelsberger-Mangan and W. B Levy. Information maintenance and statistical dependence reduction in simple neural networks. *Biological Cybernetics*, 67(5):469–477, September 1992.

[91] D. M. Adelsberger-Mangan and W. B Levy. Character recognition using adaptively constructed feed-forward networks. In *INNS World Congress on Neural Networks*, volume 11, pages 423–426, 1993.

[92] D. M. Adelsberger-Mangan and W. B Levy. Adaptive synaptogenesis constructs networks which allocate network resources by category frequency. In *Neural Networks, 1994. IEEE World Congress on Computational Intelligence., 1994 IEEE International Conference on*, volume 4, pages 2245–2249. IEEE, June 1994.

[93] R. Rojas. *Neural networks: a systematic introduction.* Springer Science & Business Media, 2013.