On the Speed of Neuronal Populations

Dissertation
for the award of the degree
"Doctor of Philosophy" (Ph.D.)
Division of Mathematics and Natural Sciences
of the Georg-August-Universität Göttingen
submitted by

Tūreiti Keith

from Tauranga, Aotearoa.

Göttingen, 2016

Andreas Neef (Thesis Advisor)
Group Leader, Biophysics of Neural Computation,
Bernstein Center, Göttingen

Walter Stühmer (Thesis Committee Member)
Director, Dept. for the Molecular Biology of Neuronal Signals,
Max Planck Institute of Experimental Medicine, Göttingen

Fred Wolf (Thesis Committee Member)
Group Leader, Theoretical Neurophysics, Dept. of Nonlinear Dynamics,
Max Planck Institute for Dynamics & Self-Organization, Göttingen

Jörg Enderlein (Reviewer)
Biophysics / Complex Systems, Third Physical Institute,
Georg-August University, Göttingen

Ulrich Palitz (Reviewer)
Biomedical Physics Group
Max Planck Institute for Dynamics & Self-Organization, Göttingen

Florentin Wörgötter (Reviewer)
Biophysics, Third Physical Institute
Georg-August University, Göttingen

Date of the oral examination: Tuesday 7$^{\text{th}}$ March 2017

Affidavit

I declare, herewith, that the doctoral thesis "On the Speed of Neuronal Populations" is entirely my work and my work alone, prepared with neither sources nor aids other than those quoted.

Göttingen, December 31, 2016

The author, Tūreiti Keith

He waka eke noa.
A canoe upon which we embark together, without exception.

Abstract

This thesis presents novel results in the areas of closed loop electrophysiology and neuronal population coding. This work begins by presenting the first known attempt to control the spike rate of a neuron or neurons using a feedback controller that drives a stochastic photostimulus. Using this controller, *in vitro* experiments were performed that look into the response of a population of cortical neurons to subtle changes in the mean of a correlated, stochastic photostimulus. This work then presents a numerical examination of how stimuli targeting specific elements of a neuron's structure, referred to as "subcellular targeting," affect the response speed of the population, the action potential onset of its constituent neurons and their effective passive bandwidth. The results presented here demonstrate that it is indeed possible to regulate the spike rate of a neuron or neurons using a feedback controller that scales the mean and variance of a correlated, stochastic photostimulus. Furthermore, using this controller, trial-based experiments were implemented in the laboratory. These demonstrate that an *in vitro* population of cortical neurons can provide robust albeit slow responses to subtle changes in the mean of a correlated, stochastic photostimulus. Finally, in numerical experiments, mean modulated, correlated, stochastic photostimuli were applied to either the entire cell, the soma or the basal dendrites of a morphologically realistic constituent neuron. The results show that this subcellular targeting of constituents produces a faster population response when the soma is targeted, followed by the basal dendritic then the global target. Although a strong statistical argument is difficult to make, the results do suggest congruent increases in the response speed of a neuronal population under subcellular targeting, the onset rapidness of its constituent neurons and their passive bandwidths.

Acknowledgements

This work owes a great debt of gratitude to a great many people, whom I wish to thank. Firstly, I thank my supervisor, Andreas Neef, whose unwavering support, keen eye for detail and honest, approachable manner helped keep me motivated at the best and worst of times. I thank the members of my thesis committee, Andreas Neef, Walter Stühmer and Fred Wolf, as well as Jörg Enderlein, who was there at the beginning. When the going got rough, they were there to provide me with the invaluable feedback and help I needed to move forward. I thank Andreas Neef, Fred Wolf, the Göttingen Graduate School for Neurosciences, Biophysics, and Molecular Biosciences (GGNB) and the Bernstein Center for Computational Neuroscience for providing me with the financial resources to see this project through. I thank Manuel Schottdorf, Sabine Klöppner, Julian Vogel and Ahmed El Hady, for their tireless work in the laboratories at the Max Planck Inst. for Experimental Medicine (MPIEM), preparing and feeding the cortical cell cultures that I used in the course of this work. I thank Manuel Schottdorf, Ricardo Merino, Carolina León Pinzón, Ahmed El Hady and Kai Bröking for teaching this software engineer the ins and outs of the laboratory. I thank Denny Fliegner, Yorck-Fabian Beensen and Hecke Schrobsdorff at the Max Planck Inst. for Dynamics & Self-Organisation (MPIDS) for their outstanding work keeping the Grid Engine, compute cluster and network infrastructure flying. A great portion of this work would not have been possible without the tools that they provide. I thank Michael Günther in the Electronics Workshop at the MPIDS for his advice and his circuit building talents. I thank my colleagues, both past and present, at the MPIDS and MPIEM for the inspiring scientific discussions, valuable feedback and for creating an exciting and stimulating environment within which to work. I thank Ayşe Bolik, Vika Novak and Bettina Hartung at the MPIDS and Ute Rust at the MPIEM for their help in navigating the administrative side of working at the Max Planck Gesellschaft. I thank Zrinka Gattin, Kirstin Mosch, Yvonne Reimann, Susanne Kracke, Mandy Fricke, Kirsten Pöhlker, the GGNB and the office for the Program for Theoretical & Computational Neurosciences (PTCN) for their hard work administering the graduate school and the PTCN programme and for their advice and help in navigating university processes and procedures. I thank Walter Stühmer, director of the Dept. for the Molecular Biology of Neuronal Signalling at the Max Planck Inst. of Experimental Medicine, for the use of his laboratories. I thank Theo Geisel, director of the Dept. of Nonlinear Dynamics at the Max Planck Inst. for Dynamics & Self-Organisation, for the use of his offices, network and computing infrastructure. I thank Georg-August University for the use of their facilities during the course of this work. I thank Andreas Neef, Juan Daniel Florez-Weidinger and my wife Yana Breindl for putting this document under the microscope and providing me with excellent feedback. I give a very special and heartfelt thank you to my wife, Yana, for her unconditional love and support through what has been a sometimes difficult process.

Without her, this work would not have been possible. I thank my parents Wayne and Raewyn and my sister Tiraroa, who despite the 20,000 kilometres between us, are always only ever a phone call away. Last but by no means least, I thank my family: my wife Yana and my children Moanaroa and Baby who remind me on a daily basis that there is indeed life outside of work and study, who keep me firmly lodged in the here and now and who also keep me excited about what is to come.

Contents

List of Figures and Tables

Introduction

An ensemble of neurons working in concert can not only encode information, but can do it faster than a neuron working alone. Termed "population coding," this strategy has been shown to enable "ultra-fast," high bandwidth encoding of signals that modulate the mean of a correlated, stochastic process. In other words, a population of neurons can respond to a small change in the mean of its inputs within a millisecond, and the signals that it can encode can have frequencies in the hundreds of Hertz (Boucsein et al. 2009; Brunel et al. 2001; Fourcaud-Trocmé et al. 2003; Higgs and Spain 2009; Ilin et al. 2013; Köndgen et al. 2008; Naundorf et al. 2005; Ostojic et al. 2015; Tchumatchenko et al. 2011; Wei and Wolf 2011). This work looks at these results from a new perspective: Whether the speed of the population response is influenced by the application point of the stimulus, what I refer to as "subcellular targeting." This work also looks into the size of the stimulus used to generate the ultra-fast response, namely whether the size of the modulation matters. In addition to this, this thesis presents work done to extend upon current closed loop electrophysiological techniques (Bölinger and Gollisch 2012; Hodgkin et al. 1952; Neher and Sakmann 1976; Newman et al. 2015; Sharp et al. 1993; Wallach 2013; Wallach et al. 2011), to produce an experimental set up capable of controlling the rate of activity of a neuron or neurons under a correlated, stochastic stimulus.

This work extends upon emerging techniques in closed loop electrophysiology. More specifically, it examines the use of a "closed loop" technique to regulate the rate at which a cell generates action potentials (a.k.a spikes), the electrical impulses of a few milliseconds in length that a neuron uses to communicate with its downstream counterparts. Closed loop electrophysiology is, in its most abstract form, the application of a feedback loop between the measurement of potential or current on or within a cell and the application of a stimulus to said cell. It encompasses techniques that allow experimentalists to control e.g. the voltage at or current through the membrane of a cell (Hodgkin et al. 1952; Neher and Sakmann 1976; Sharp et al. 1993) or the probability of response of a neuron (Bölinger and Gollisch 2012; Wallach 2013; Wallach et al. 2011). More recently, Newman et al. (2015) have shown that it is possible to regulate the spike rate of a cell using a deteministic stimulus. In this work, I describe a novel experimental configuration, that demonstrates that this is also possible using a correlated, stochastic process.

This new closed-loop approach was crucial to the implementation of population response

experiments I performed in the laboratory. These experiments looked into how an *in vitro* neuronal population (i.e. a cell culture) responds to very small changes in the mean of a correlated, stochastic process. This is inspired by previous work performed by our group (Tchumatchenko et al. 2011), in which experiments demonstrated that a neuronal population could detect, within a millisecond, a change in the mean of a correlated, stochastic stimulus. To achieve this fast response, these experiments employed a 20 pA step in the stimulating current, which corresponds to a single strong synaptic (neuron-to-neuron) input. The experiments described in this thesis take a slightly different approach. Rather than defining an absolute step size, the size of the step is defined relative to the mean of the correlated, stochastic process required to generate a spike rate of 1 Hz. Furthermore, the chosen size of the step relative to the mean is the smallest observed that produced a response. The results of these experiments demonstrate that a small step size produces a slow neuronal population response speed. They also, however, provide testament to the extreme robustness of the population code.

Finally, this thesis discusses the idea of "subcellular targeting," i.e. stimulating a specific portion of a neuron's structure. The ubiquitous stimulation techniques of *voltage clamp* (Kenneth Cole[1]) and *current clamp* (Hamill et al. 1981) typically target the *soma* of a cell – its main body – by either manipulating the membrane potential or by injecting a current. A neuron, however, comprises not just a soma, but also *dendrites* and an *axon*. The *dendrites* are branched projections that receive inputs, which are integrated by the soma. The output is the *axon*, a longer projection whose initial segment is the point at which an action potential is initiated. Using the newer and increasingly popular optogenetic techniques (Boyden et al. 2005; Nagel et al. 2002; Nagel et al. 2003), all of these subcellular components are potential targets. With this technique, one can infect a cell such that it produces light gated ion channels in its membrane, which open and close in response to a photostimulus, allowing an ionic current to flow through the cell membrane. In this work, I look into how the response of a neuronal population is affected by the subcellular targeting of its constituents. The effect of subcellular targeting appears to be predicted, somewhat, by the literature that relates the neuronal population response to the onset of a constituent's action potential and the electrical loading of its dendritic tree. One property of a neuron's action potential, its onset rapidness, i.e. the rate at which the spike initiates, has been shown to influence the speed at which an ensemble of said neurons can respond to changes in its inputs. Namely, the faster a constituent neuron's action potential onset, the faster the response of the population (Eyal et al. 2014; Fourcaud-Trocmé et al. 2003; Ilin et al. 2013; Naundorf et al. 2005; Wei and Wolf 2011). A rapid onset is thought to be the product of fast opening sodium channels, i.e. the channels that open at the beginning of an action potential, allowing positive, ionic charge to flow into the cell (Fourcaud-Trocmé et al. 2003). However, there is theoretical evidence that points to cell morphology having an important

---

[1]The invention of voltage clamp is attributed to Kenneth Cole by Purves et al. 2004

influence, more specifically, that the size of a neuron's dendrites is proportional to the onset rapidness of its action potential (Eyal et al. 2014). They claim that, when one stimulates a cell at the soma, the dendritic tree represents a large current sink into which action potential onset currents flow. The larger this sink, the better its ability to draw spike onset currents, the faster the onset and the faster the population response. If one were to stimulate the dendritic tree such that ion channels on the dendritic membrane open, allowing current to flow into the cell through the dendrites first, then the sink effect of the dendrites should be reduced, if not eliminated. This points to the idea that, when it comes to population encoding, the point of input matters. This work demonstrates that the subcellular targeting of constituent neurons effects the speed with which a neuronal population response and also takes the first steps in understanding what lies behind this change in response speed.

The work done to examine subcellular targeting, modulation size and spike rate control is presented by first outlining the key literature in Chapter 1. This is followed by Chapter 2, which describes the experimental methods and materials employed during the course of this work. Chapters 3 and 4 present the experimental results. The former looks into the results of feedback control, the latter subcellular targeting and modulation size. Finally, Chapter 5 discusses the results, presents the final conclusions and future work.

Chapter 1

Background

The three key themes of this work are 1) closed loop electrophysiology, 2) the response speed of an *in vitro* neuronal population to subtle changes in stimulus and 3) the relationship between the response speed of a neuronal population, the onset rapidness of its constituents, their passive bandwidth and how this relates to the targeting of specific elements of a neuron's structure, what I refer to as "subcellular targeting." The background material relating to these three topics is presented in this chapter. Sections 1.1 and 1.2 present literature on closed loop electrophysiology and the proportional-integral-derivative (PID) controller respectively. This material was used in the implementation of a novel closed loop electrophysiological system, capable of regulating the spike rate of a neuron or neurons under a correlated, stochastic photostimulus. The closed loop system was used in the laboratory to measure the response of an *in vitro* neuronal population to subtle changes in the mean of a correlated, stochastic photostimulus. Section 1.3 follows. This presents the literature on onset rapidness and population coding. This is followed by Sect. 1.4, which summarises the background literature on the effects of dendritic loading on onset rapidness. The material in these last two sections is key to understanding the numerical experiments performed during the course of this work, which look into the relationship between neuronal population response, the onset rapidness of constituent neuron's, their passive bandwidth and how this relates subcellular targeting.

1.1   Closed Loop Electrophysiology

'Closed loop' electrophysiology describes, in essence, the hybridisation of neuron and silicon, whereby a neuron or neurons are stimulated and their response is recorded; this response is then used to modify the stimulus – i.e. it is fed back. The idea behind an *in silico* feedback path in a neuronal circuit forms the weft of electrophysiology's historical fabric. The arguably most famous implementations of this idea, the "voltage clamp" (Kenneth Cole; Hodgkin and Huxley 1952a) and patch clamp (Neher and Sakmann 1976), are important electrophysiological techniques, which incorporate electronic feedback circuits to clamp

membrane voltage. The techniques of voltage clamp and patch clamp led to the "dynamic clamp" (Sharp et al. 1993), where e.g. voltage is recorded from a cell or cells and currents are injected (fed back), calculated using e.g. membrane or synaptic conductances modelled *in silico*. Wagenaar (2005) extended the dynamic clamp idea by employing an *in silico* feedback path adjusting an electrical stimulus so as to reduce bursting within a neuronal culture. The "response clamp" (Bölinger and Gollisch 2012; Wallach 2013; Wallach et al. 2011) followed, generalising the idea that one can control the response probability of a neuron using feedback control of the stimulus. Eventually these closed loop techniques were employed in combination with optogenetics (Newman et al. 2015) to show the feedback control of a spike rate using a deterministic stimulus. Indeed, voltage clamp, patch clamp, dynamic clamp and response clamp form a continuous line of techniques that led to what is today described as "closed loop electrophysiology" or, in a broader context, "close loop neuroscience" (Arsiero et al. 2007; El Hady 2016; Grosenick et al. 2015; Potter et al. 2014): The insertion of a silicon-based feedback path between the output of a cell or cells and their input. In this work, I extend on the work of Newman et al. (2015), Wagenaar (2005), and Wallach et al. (2011) by demonstrating that an adaptive proportional-integral-derivative (PID) controller can be used to drive a stochastic stimulus and control the rate of neuronal activity measured at a probe point within a cortical culture.

1.2 PID Control

'Feedback control' is the term used, when a controller regulates the output of some system or 'plant' by feeding said plant's output, after some manipulation, back to its input. Such a controller performs 'closed loop' control, in contrast to 'open loop' control, where the controller drives the input of the plant without knowledge of its actual output. One exceedingly popular closed loop controller is the proportional integral derivative or PID controller. First treated academically in the 1920's and 1930's (Hazen 1934; Küpfmüller 1928; Minorsky 1922; Nyquist 1932), the PID controller has become "the standard tool for solving industrial control problems" (Åström and Hägglund 2006). The PID controller is described as

$$u(t) = K_{\mathrm{p}}e(t) + K_{\mathrm{i}}\int_0^t e(t)\,dt + K_{\mathrm{d}}\frac{de(t)}{dt}, \tag{1.1}$$

where $e(t)$ is the controller error; K_{p}, K_{i} and K_{d} are the P, I and D coefficients respectively; and $u(t)$ is the output of the controller used to drive the plant. The controller (Fig. 1.1) calculates the error, $e(t)$, between a plant's output, $y(t)$, and the target output, $r(t)$, and propagates this along three different paths, where it is amplified (proportional - P), integrated (I) or differentiated (D). The P path weights the instantaneous error, the I path the history of the error and the D path the 'future' error (i.e. the instantaneous error slope). A gain is applied to each path, and the summed result is fed back to the plant's input as the control signal. The PID controller is tuned by adjusting the gains (a.k.a. weights or
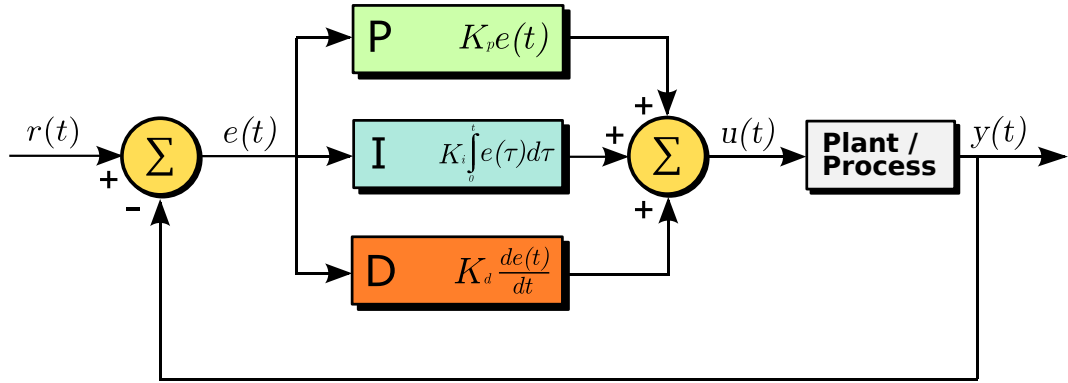
Figure 1.1: The proportional integral derivative (PID) controller (Image: Urquizo 2011), where $r(t)$ is the target output, $y(t)$ is the actual output of the plant, $e(t)$ the controller error and $u(t)$ is the output of the controller used to drive the plant. P, I and D are the proportional, integral and derivative parts respectively. K_p, K_i and K_d are the proportional integral and derivative coefficients respectively.

coefficients), K_p, K_i and K_d.

In industrial processes, the *exact* behaviour of the plant is often not only non-linear, but also unknown (Åström and Hägglund 2006; Åström and Murray 2008). Thus, many methods have been developed to adjust (or 'tune') the PID controller weights for the unknown plant. These involve either heuristic or adaptive tuning. Heuristic methods typically involve performing tests on the plant to identify response parameters such as lag or rise time, then calculating the PID coefficients using heuristics (most notably Åström and Hägglund 2006; Chien et al. 1972; Cohen and Coon 1953; Ziegler and Nichols 1942). Adaptive methods adjust the PID weights while the controller and plant are in service. Typical approaches can be categorised as 'adaptive control', where the PID coefficients are adjusted continuously, online until reaching some convergence point; and 'gain scheduling', where a set of PID coefficients are precomputed for different ranges of plant dynamics. For the purposes of this work, the former approach has been taken, more specifically, the "adaptive interaction" PID controller (Cominos and Munro 2002; Lin et al. 2000).

The "adaptive interaction" PID controller (Fig. 1.2, Lin et al. 2000), performs gradient descent, minimising the control error as a function of the PID coefficients K_p, K_i, K_d. The approach separates the closed loop control system into four separate but interacting (connected), single-input single-output, linear devices: P, I, D and plant. Using the theory of adaptive interaction (Brandt and Lin 1999), where the PID coefficients are viewed as the interaction parameters, they derive the adaptation rules

$$\dot{K}_p = -\gamma \, e \, y_p$$
$$\dot{K}_i = -\gamma \, e \, y_i \qquad\qquad (1.2)$$
$$\dot{K}_d = -\gamma \, e \, y_d,$$

where γ is the learning rate, e is the controller error and y_p, y_i and y_d are the outputs (before weighting) of the unity gain (proportional part), derivative and integrator respectively.
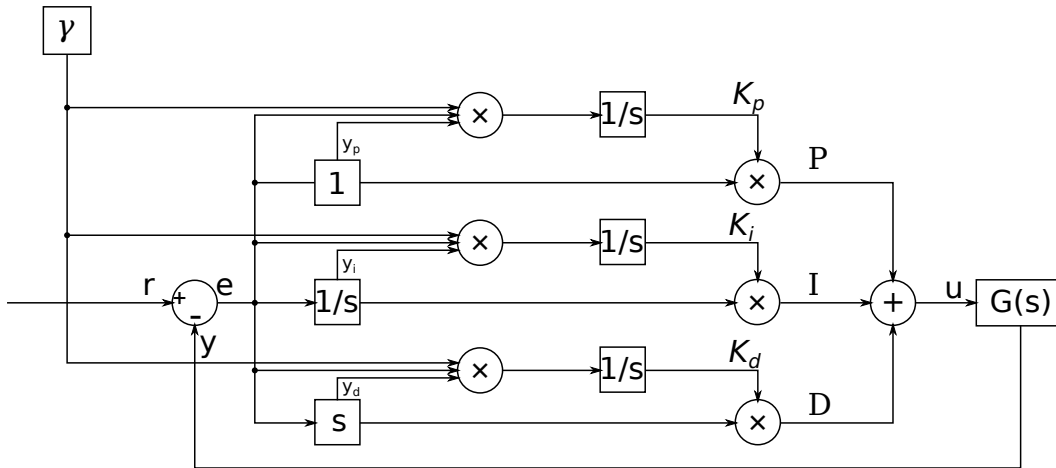
Figure 1.2: The adaptive PID controller topology. Here, the output y of the 'plant' (controlled system), $G(s)$, is subtracted from the desired output r, to produce the error signal e. This error is fed, in turn, to the controller, comprising unity gain, derivative ('s') and integrator ('$1/s$') paths. The adaptive portion of the circuit requires an adaptation rate, γ.

1.3 Population Coding & Onset Rapidness

Population coding describes a strategy by which neurons work in concert to code information. A significant body of evidence points to a relationship between the speed with which such a population can respond to a stimulus, and the onset speed of its constituents' action potentials. The following sections describe both population coding, onset rapidness and their relationship, beginning with the action potential.

1.3.1 The Action Potential

Neurons, the cellular building blocks of our nervous system and brain, can communicate with each other through their action potentials (AP). These APs, first described by Emil Heinrich Du Bois-Reymond in 1848 (Pearce 2001), can be measured electrically, at the membrane of a cell, as jumps of tens of millivolts over a few milliseconds. In the given biological context, this behaviour is somewhat brief and sharp and has earned the action potential the pseudonym 'spike'. The act of producing a spike is thus called 'spiking'. Figure. 1.3 plots an example of an action potential recorded from the giant axon of the squid.

The action potential is the composition of activity from ion channels and pumps embedded in the membrane of the neuron. These facilitate the flow of charged ions in and out of the cell, thus hyperpolarising and depolarising the cell, generating respectively the up- and down-swing observed in an action potential. A diversity of channels has been identified to date, each with different properties. These channels fall into two groups: voltage gated, which open or close depending on the membrane potential; and ligand gated, which open and close in response to chemical signals – i.e. those involved in synaptic transmission. The ionic currents involved in action potentials are comprised of Na^+, Ca^{2+}, K^+ and Cl^- ions. Typically Na^+ channels are responsible for hyperpolarisation (in some cells this role is

Figure 1.3: The first known recording of an action potential using a microelectrode inserted into a cell (Image: Hodgkin and Huxley 1939, p.711). This particular recording was taken from the giant axon of the squid, *Loligo forbesi* (Hodgkin and Huxley 1939).

filled by Ca^{2+}) and K^+ for depolarisation. Ca^{2+} channels typically modify the shape of the action potential and Cl^- channels control excitability (Bernstein 1902; Fatt and Katz 1953; Hodgkin and Huxley 1952d; Miller and White 1980; Neher and Sakmann 1976; Purves et al. 2004).

1.3.2 Onset Rapidness

One parameter used to characterise an action potential (AP or spike) is its onset rapidness. Onset rapidness describes how quickly an action potential rises immediately after it has been triggered. One method used to quantify onset rapidness is the slope of the AP's phase plane plot at some, specified gradient. Eyal et al. (2014) used this approach in their work. They took the example of three simulated action potentials (Fig. 1.4a) and compute their phase plane plots (Fig. 1.4b). This is simply the spike potential ($v(t)$, horizontal axis) plotted against the gradient of the spike potential ($\frac{dv}{dt}$, vertical axis). These phase plots are then used to estimate the onset rapidness by computing the gradient of the phase plane plot $\frac{d}{dv}\left(\frac{dv}{dt}\right)$ (Fig. 1.4b inset, solid lines) at a specific value of $\frac{dv}{dt}$, in this case $\frac{dv}{dt} = 10\,\mathrm{mV\,ms^{-1}}$. One sees in this example that, if the correct point of onset rapidness measurement is chosen, a sharper 'kink' (Fig. 1.4a inset) corresponds to a higher phase plane plot gradient (Fig. 1.4b inset) and therefore a higher onset rapidness value. Onset rapidness plays a key role in population coding. This role is discussed in detail in Section 1.3.4. First however, Section 1.3.3 discusses population coding.

1.3.3 Population Coding

The mammalian nervous system is fast. In $5 - 10\,\mathrm{ms}$ sensory stimuli can reach the cortex (Swadlow and Hicks 1996), in $30\,\mathrm{ms}$ a perceptual decision can be made (Stanford et al. 2010), in 150-200 ms, a natural image can be processed (Thorpe et al. 1996). However, experimental
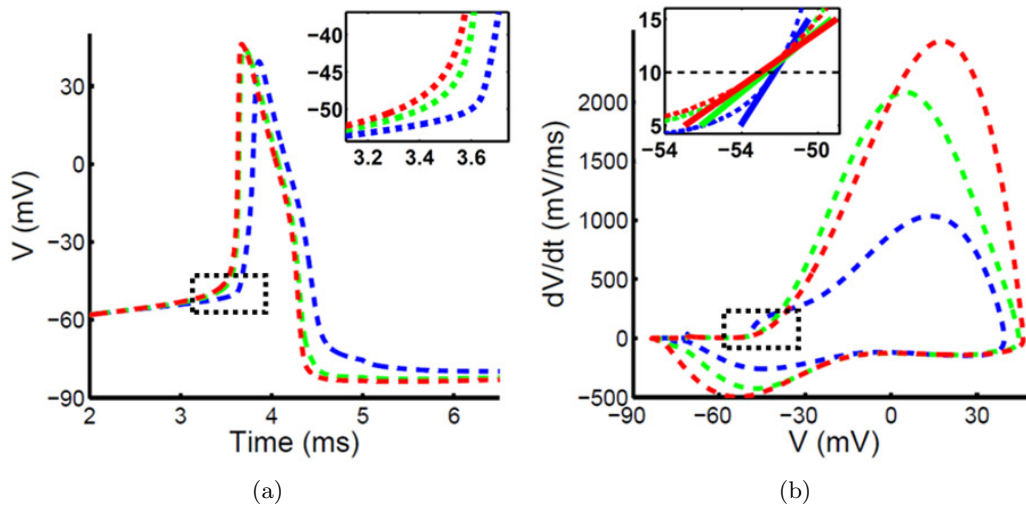
Figure 1.4: One method of measuring onset rapidness (Image: Eyal et al. 2014, p.8065). *a)* Three superimposed simulated action potentials. *a inset)* Contrasting the onset 'kinks' of each of the action potentials *b)* The corresponding phase plane plots. These show the potential of the spike on the horizontal axis against the gradient of this potential on the vertical axis. *b inset)* The reference point chosen to measure the phase plane plot's gradient, in this case, is $dv/dt = 10\,\mathrm{mV\,ms}^{-1}$ (black dashed line). The dashed coloured lines correspond to the phase plane plots. The solid coloured lines correspond to the gradients of the phase plane plots at the onset rapidness measurement point. The sharper the 'kink' the faster one expects the onset rapidness rapidness to be.

evidence shows that neurons generate action potentials at low rates and sparsely (Brecht et al. 2003; Brecht and Sakmann 2002). Thus, to achieve such high processing rates, it has been suggested that neurons act in concert as a population or ensemble, hence "population coding." Indeed, evidence of such behaviour has been observed experimentally in e.g. sensor and motor areas of the brain. For example, arm movement was shown to be encoded in a population of motor cortical neurons (Georgopoulos et al. 1986), eye movement in a neuronal population in the superior colliculus (Lee et al. 1988; Sparks et al. 1976). Maunsell and Van Essen (1983) found a population in the visual area middle temporal tuned to the direction of movement of observed objects. Given this, neuroscientists have sought to understand these populations and how they could be encoding information (Paradiso 1988; Pouget et al. 2000; Salinas and Abbott 1994; Seung and Sompolinsky 1993), how correlated noise effects the accuracy of such codes and, depending on its structure, can improve or limit the code's accuracy (Abbott and Dayan 1999; Averbeck et al. 2006; Sompolinsky et al. 2001; Vogels 1990; Wu et al. 2002; Zohary et al. 1994). This lead, of course, to work in measuring the correlation structure of noise within the brain (Cohen and Kohn 2011; Ecker et al. 2010; Zohary et al. 1994). In parallel, experimentalists have tried to measure the speed and bandwidth of populations, i.e. how quickly a population can respond to external stimuli.

1.3.3.1   Speed & Bandwidth

Of the experimental work to date, one area of focus has been the encoding bandwidth of neuronal populations under mean and variance modulated noise signals. Modulating the

mean could correspond to the effect of synaptic transmission on post synaptic neurons, whereas modulating the variance could correspond to the stimulation of networks where excitation and inhibition are balanced, as is implied by observed behaviour in the neocortex (Okun and Lampl 2008). In this case individual neurons observe a change in variance at their inputs, but little change to their mean. Evidence shows that *in vitro* populations are more sensitive to mean modulation of a *correlated* noise signal than to modulation of its variance (Tchumatchenko et al. 2011). Also using correlated noise, Boucsein et al. (2009) found that populations could phase lock to sinusoids modulating either mean or amplitude for frequencies of up to ≈200 Hz and ≈100 Hz respectively. Köndgen et al. (2008) found similar results for sinusoids modulating the mean of a correlated noise signal. In the *uncorrelated* noise case Silberberg et al. (2004) demonstrated that variance modulations of a noise signal could be better encoded that mean modulations.

1.3.4   Onset Rapidness Influences Population Encoding Speed

A wealth of experimental, theoretical and numerical evidence points to a relationship between the encoding bandwidth of a neuronal population and the action potential (AP) onset rapidness of its constituent neurons. From the work of Brunel et al. (2001) to more recent experimental work (Ilin et al. 2013), the evidence shows that neuronal populations are capable of encoding information at frequencies of up to hundreds of Hertz, when this information modulates a correlated noise signal. The upper bound of frequencies that can be encoded, the cutoff, is defined by the form of the action potential of the constituent neurons, more specifically, the APs onset rapidness: The faster the AP onset, the higher the frequencies that can be encoded.

One of the most important first steps along this path was the realisation that a neuron's encoding bandwidth can be greater than the cutoff frequency defined by its membrane. Brunel et al. (2001) demonstrated theoretically and numerically that the encoding capabilities of a leaky-integrate-and-fire (LIF) neuron are dependent on the qualities of the 'carrier' signal. Signals that modulate the mean of white noise are limited to the cutoff frequency of the membrane model; however, low-pass filtering this noise, i.e. introducing a time-constant, increases the bandwidth of frequencies that a LIF can encode. For the filter time-constants studied, up to $\tau = 40$ ms (a cutoff of ≈ 4 Hz), Brunel et al. (2001) showed that increasing τ increases bandwidth. E.g. where $\tau = 0$ ms gives a normalised rate response cutoff of ≈ 11 Hz, introducing $\tau = 40$ ms gives a normalised rate response of ≈ 1 in the high frequency limit. Brunel et al. (2001) tied this result to population coding by also showing the population response to a step. Visual inspection of the resulting peri-stimulus time histograms for $\tau = \{0, 2, 5, 10\}$ ms shows clear differences in the population responses. The rise time of the $\tau = 10$ ms case is ≈ 2 ms, in the $\tau = 0$ ms case, ≈ 10 ms.

Some of the first inklings that action potential initiation has an influence on population encoding bandwidth came soon after. For example, the work done by Fourcaud-Trocmé et

al. (2003) looked at the role of the sodium channel in action potential onset and population encoding. Sodium channel opening being the major physical mechanism behind the initial changes in potential observed in a neuron's spike. They used numerical simulations to demonstrate that the characteristics of fast sodium channels determine the speed at which a population of neurons can track changes to noise modulated by a sinusoidal stimulus. Fourcaud-Trocmé et al. (2003) developed a one-variable model, the "exponential integrate and fire neuron" (EIF), and showed that it is sufficient to capture the activation dynamics of near threshold sodium channels in a conductance based model. They then used this EIF model to show that above the cutoff frequency of a neuron's membrane, the decay in population coding gain follows a power law, C/f^α, where α is dependent on the non-linearity of spike initiation.

A deeper examination of the role of action potential (AP) initiation on bandwidth followed thereafter. Naundorf et al. (2005) studied the effect of AP onset dynamics on the response speed of neuronal populations. They found that a neuronal ensemble's speed, at high frequencies, is not dependent on the action potential onset speed of its constituents, but rather on the *phase* at which their action potentials are emitted. In this theoretical and numerical study, Naundorf et al. (2005) examined the response of neuronal populations to oscillations embedded in noise currents. They generalised the θ-neuron (Ermentrout and Kopell 1984; Gutkin and Ermentrout 1998), a phase oscillator model, to provide an adjustable AP onset dynamic, mimicking, in effect, the fast activation of sodium channels. They analysed the behaviour of an ensemble of generalised θ-neurons in the high frequency limit for different types of inputs (mean modulated, amplitude modulated, correlated and uncorrelated noise) and found that the response amplitude depends only on the phase of the oscillator.

In the wake of this theoretical and numerical work, the experiments presented in several key articles gave biological evidence of high frequency encoding. Köndgen et al. (2008) performed experiments in brain slices with layer V pyramidal neurons from the rat somatosensory cortex. They injected correlated noise currents, mean modulated by sinusoids. The response of cells was measured using the "modulation depth", where they generated peri-stimulus time histograms describing the population response, and fitted a sinusoid, the amplitude of this sinusoid was then normalised by the mean firing rate. Plotting modulation depth against frequency, they were able to show cutoff frequencies of approximately 200 Hz. Using a similar experimental approach and analysis, Boucsein et al. (2009) confirmed high frequency cutoffs for both mean and amplitude modulated signals, where cutoffs of approximately 100 Hz were measured in the former case, 200 Hz in the latter. Similarly, Higgs and Spain (2009) performed experiments using layer 2/3 pyramidal neurons from rat neocortical slices. They found that cells demonstrate high frequency resonances at approximately 250, 400 and 250 Hz in response to $\tau = 1$ ms exponential noise, τ=5 ms exponential noise and

'$1/f$' noise[1] respectively. Ostojic et al. (2015) also report a 200 Hz resonance in populations of Purkinje cells in rat slices.

By 2011, several experimental studies had demonstrated that neuronal populations are indeed capable of encoding information at frequencies of up to hundreds of Hertz. The biophysical models of the time predicted, however, lower cutoffs. To remedy this, Wei and Wolf (2011) introduced a dynamical model that is both capable of encoding high frequency signals and also captures action potential dynamics in an analytically tractable fashion for uncorrelated noise stimuli. For white noise modulated by a signal, they used their model to show that the cutoff frequency of a neuronal population can be modulated by, and is proportional to the action potential onset rapidness of its constituents. Although, in the white noise case they found that the encoding bandwidth was limited to the cutoff frequency of the membrane, they were able to show numerically, that the introduction of a finite correlation time lifted this restriction, as per Brunel et al. (2001).

A key work tying the experimental, theoretical and numerical findings together came from Ilin et al. (2013). In laboratory experiments they tested the theoretical prediction that action potential (AP) onset influences population encoding bandwidth (Fourcaud-Trocmé et al. 2003; Naundorf et al. 2005; Wei and Wolf 2011). Experiments with rat brain slices demonstrated that the encoding capabilities of neocortical neurons, with naturally short AP onset times, were severely reduced when either extracellular Na^+ was reduced, or with the introduction of the neurotoxin, tetrodotoxin (TTX), to the axon initial segment, both of which have the effect of reducing the onset rapidness of an action potential at the axon initial segment (AIS).

Thanks to the combined work of Boucsein et al. (2009), Brunel et al. (2001), Fourcaud-Trocmé et al. (2003), Higgs and Spain (2009), Ilin et al. (2013), Köndgen et al. (2008), Naundorf et al. (2005), Ostojic et al. (2015), and Wei and Wolf (2011) we can say, with confidence, that populations of neurons are capable of encoding signals that modulate the mean of correlated noise at frequencies in the hundreds of Hertz. Furthermore, the response speed of a neuronal population increases with the onset rapidness of its constituent's action potentials. This does, however, raise the question, what makes the onset of an action potential fast or slow? In their theoretical work, Fourcaud-Trocmé et al. (2003) found a link between fast sodium channels, onset rapidness and population response speed. More recently, scientists found a link between morphology, onset rapidness and population response speed. Sect. 1.4 describes this morphological link in more detail.

1.4 Dendritic Loading & Onset Rapidness

A neurons electrical behaviour is determined by factors typically divided into two groups, those influencing the cell's 'active' electrical response, and those influencing its 'passive'

---

[1]I.e. noise with a spectrum that decays at the rate of $1/f$.

response. In the active case, one speaks of gated ion channels, channels embedded within the membrane of a cell that can open and close allowing ionic currents to flow both into and out of the cell. A cell's passive properties, however, influence the flow of ionic currents within and their diffusion through the walls of its membrane. A cell's passive response can be understood through the analogous case of electrical transmission, the 'transmission line'. The flow of electrical currents along a transmission line or cable are influenced by the impedance properties of said cable; typically modelled as a series of connected impedances, each comprising a resistance, capacitance and inductance. In the case of a neuron, it is sufficient to consider only resistance and capacitance (Dayan and Abbott 2002).

Just as the structure, materials and dimensions of a transmission line influence its overall impedance, the passive electrical properties of a cell are a function of the cell's morphology. Work in this area has demonstrated examples of "morphology-induced resonance" (Ostojic et al. 2015) or the influence of the dendritic tree on the structure of firing patterns (Mainen and Sejnowski 1996). Specific to the work in this thesis are the results obtained by Eyal et al. (2014). They examine the role of the dendrites in action potential formation and encoding bandwidth. They were able to show that increasing dendritic load increases action potential (AP) onset rapidness and improves the encoding bandwidth of a neuronal ensemble. In their simulations of a simple ball-and-two-sticks compartment model, they increase the dendritic load by enlarging one of the two sticks, and in each case, stimulate with a fluctuating noise current, whose mean is modulated with a sinusoid. They show that, the larger the dendritic 'stick' the faster the onset rapidness and the better the cell can phase-lock to the sinusoid at higher frequencies. This work is, however, somewhat contradicted by theoretical work presented by Vacučiaková (2016). Looking at the transfer function of a passive cell, she shows that increasing dendritic diameter causes a corresponding decrease in gain, due to the increased surface which must be charged. Vacučiaková (2016) also demonstrates that increasing dendritic length has no significant influence on the response above the length constant for direct current (DC), i.e. the distance over which a DC signal has some measurable influence.

Chapter 2

Materials & Methods

This thesis presents a novel closed loop electrophysiological system, and demonstrates how it can be used to perform experiments examining the response of an *in vitro* neuronal population to subtle changes in the mean of a correlated, stochastic photostimulus. In addition to this, this work also presents the results of numerical experiments that look into the relationship between the response speed of a neuronal population, the onset rapidness of its constituents, their passive bandwidth and how this relates to the targeting of stimuli to specific elements of the neuron's structure, what I refer to as "subcellular targeting." The materials and methods required to implement these numerical and *in vitro* experiments are described in this chapter. This begins with a description of the experimental set-up used in the laboratory in Sect. 2.1. Section 2.2 presents the tools used in the numerical simulations. This is followed by Sect. 2.3 which shows the form of the stimuli applied in both the *in vitro* and in numerical experiments. Section 2.4 presents the analyses applied to experimental data. Finally, Sect. 2.5 presents the protocols that were used in the execution of both *in vitro* and numerical experiments.

2.1 Experimental Set-Up

Experiments were performed at the laboratories of the Max Plank Institute for Experimental Medicine, Dept. for the Molecular Biology of Neuronal Signals (director, Walter Stühmer). The experiments employed devices already existing in the lab. The devices were integrated using software that I both designed and implemented, to create a closed-loop electrophysiological platform.

2.1.1 Hardware

The core hardware components are the multi-electrode array (MEA) system, a personal computer (PC), a light stimulator and an inverted microscope (Fig. 2.1). The MEA system (MEA1060-Inv, Multi Channel Systems GmbH (2012a)) comprises a PC (Intel Core-2 Duo) with an on-board data acquisition card (the 'MC_Card', an analogue-to-digital converter,
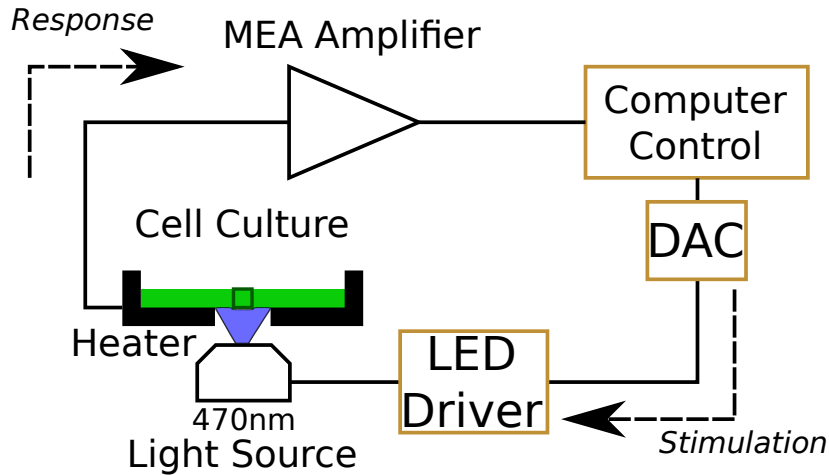
Figure 2.1: The experimental set-up. This comprises a personal computer (PC), digital to analogue converter (DAC), a blue light (470 nm) emitting diode (LED), an LED driver, and a heater above which a cell culture sits, grown on a multi-electrode array (MEA). The MEA is connected to an amplifier, which sends data to an analogue to digital converter (ADC) card sitting in the card slot of the PC.

(a)	(b)
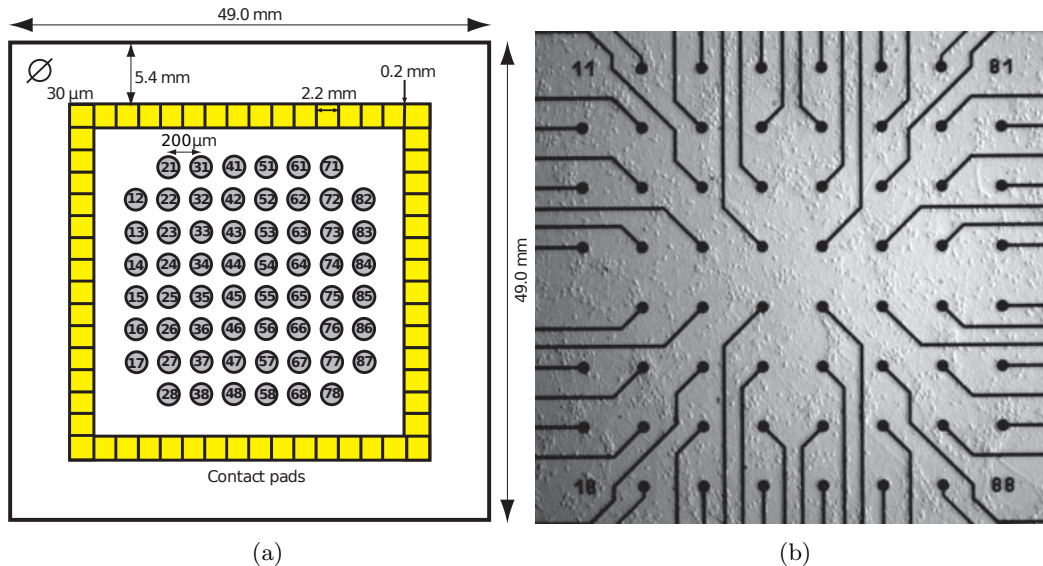
Figure 2.2: The multi-electrode array (Image: Multi Channel Systems GmbH 2012b). *a)* The key measurements for the 60MEA multi-electrode array used in this work. Electrodes are circular with a diameter of $30\,\mu()$m and a centre-to-centre spacing of $200\,\mu$m. Electrodes are laid out in an 8 by 8 grid without electrodes in the corners, making 60 electrodes, one of which (electrode 15) acts as a reference, thus leaving 59 from which to record signals. *a)* An image of the electrode array showing tracks and electrodes. The tracks and electrodes (indium tin oxide) are embedded in a glass substrate. Tracks are insulated from the culture with a thin layer of glass.

ADC), which accepts signals from an amplifier, connected, in turn, to the multi-electrode array. The PC is also responsible for driving stimuli via a digital-to-analogue converter (DAC, the STG2008, Multi Channel Systems GmbH 2008), which is connected to an LED driver, built by Michael Günther in the electronics workshop of the Max Planck Institute for Dynamics & Self-Organization. The LED driver powers a Philips Luxeon Rebel Blue LED (Koninklijke Philips N.V. 2016), with a wavelength of approximately 470 nm. This

LED performs 'whole-field' illumination, i.e. it illuminates the entire recording-electrode field of the multi-electrode array. During experiments, a multi-electrode array is placed in its amplifier which sits on the stage of an inverted microscope (Axiovert 200, Carl Zeiss 2001), the stimulating LED is mounted in a brass enclosure which sits in the microscope's turret. For an external reference of the stimulus, a photodiode (OPT101, Texas Instruments Inc. 1994) is mounted in the microscope's condenser. Its output is also connected to the data-acquisition card. The multi-electrode array used in this work (60MEA200/30iR-Ti, Fig. 2.2, Multi Channel Systems GmbH 2012b) comprises a grid of $60 \times 30\,\mu$m diameter titanium nitride electrodes, with a centre-to-centre spacing of $200\,\mu$m, one electrode is reserved as a reference electrode.

2.1.2 Software

The software that I built for this project interfaces with the aforementioned analogue-to-digital and digital-to-analogue converters (Section 2.1.1) via an application programming interface (API) called MscUsbNet.dll (Multi Channel Systems GmbH 2016) written for Microsoft's .NET Framework (Microsoft Corporation 2016). Using these interfaces, the software implements feedback control of the spike-rate measured at a user-selected electrode, and can also embed a protocol in the stimulus that it generates. The software does this by recording voltage traces captured by the ADC from user-selected electrodes. It then filters these traces, removing low frequency oscillations, and performs simple threshold spike detection. The software uses an exponential window to calculate a spike rate which is fed to an adaptive PID controller, the output of which is used to scale the mean (μ) of an Ornstein-Uhlenbeck (OU) process. The standard deviation (σ) of the OU process is scaled relative to the mean, such that $\sigma = \mu/2$. As the signal is eventually fed to the LED driver, which is active only for positive input voltages, the OU process is clipped for values below 0. If a user-provided protocol is present, this is then added to the OU process. Finally, the software sends the stimulus to the DAC, which is ultimately used to illuminate the multi-electrode array.

2.1.3 Optogenetic Photostimulation

Light stimulation has grown in popularity in the last decade, becoming a standard tool in the electrophysiologist's tool kit. Although earlier examples of photostimulation were proven successful (Callaway and Katz 1993; Fork 1971), it is the introduction of genetic techniques that brought about a paradigm shift in the approach not only to light stimulation, but to stimulation in general (Banghart et al. 2004; Boyden et al. 2005; Lima and Miesenböck 2005; Zemelman et al. 2002, 2003). Termed "optogenetics," cells are infected with a viral vector enabling them to produce light gated channels in the neuronal membrane. Of the aforementioned optogenetic approaches, it is the discovery of channelrhodopsins (Nagel et al. 2002; Nagel et al. 2003), and their introduction to electrophysiology (Boyden et al. 2005)

Figure 2.3: The Channelrhodopsin-2 structure (Image: Bamberg 2016): A 7-transmembrane protein, covalently linked to a retinal chromophore through a protonated Schiff base. Blue light ($\approx$460 nm wavelength) induces a change in the retinal complex which, in turn, opens the channel's pore. Channelrhodopsin-2 is a non-specific cation channel, conducting Na^+, K^+, H^+ and Ca^2+. The C-terminal end of Channelrhodopsin-2, which reaches into the cell, can be replaced with e.g. a yellow fluorescent protein (YFP) (Boyden et al. 2005; Nagel et al. 2003).

that has arguably had the greatest impact (Reiner and Isacoff 2016). Channelrhodopsin-2 is a non-specific cation channel (Fig 2.3), which, when exposed to blue light, opens its channel pore. Boyden et al. (2005) were able to show that, compared to previous optogenetic methods (e.g. Banghart et al. 2004; Lima and Miesenböck 2005; Zemelman et al. 2002, 2003), Channelrhodopsin-2 is able to offer sub-millisecond activation times, significantly faster than the seconds or minutes reported previously. Since their initial discovery, a range of new channelrhodopsin variants and approaches has been introduced, providing e.g. temporally stationary spike trains (Gunaydin et al. 2010), transgenic animals (Madisen et al. 2012), longer open states (Berndt et al. 2009; Dawydow et al. 2014), two-color activation (Klapoetke et al. 2014) and even faster activation times (Chronos: Klapoetke et al. 2014).

Channelrhodopsin-2 (ChR-2) was a key component in the implementation of this work. One important property of ChR-2 is its behaviour as a low pass filter. Namely, a light stimulus has been shown, under voltage clamp, to induce currents that are a low pass filtered version of the stimulus (Neef et al. 2013; Tchumatchenko and Newman 2013). ChR-2 acts as a low pass filter with a cutoff of $\approx$70 Hz at 36ºC (Fig. 2.4a, Tchumatchenko and Newman 2013). At 24ºC, ChR-2 acts as a low pass filter with a cut off of 25 Hz (Fig. 2.4b; Ricardo Merino, Max Planck Inst. for Experimental Medicine).

The hardware arrangement for these experiments (Fig. 2.1) facilitates the illumination of the entire electrode field of the multi-electrode array, also called "whole field" illumination. Whole field illumination allows the experimenter to apply the same input to the entire culture simultaneously. Of course, variations in channel density between cells, their respective morphologies and the amount of shade that they are exposed to have some influence on the amount of stimulus each cell receives. In the latter case, shading can be caused by cells overlapping, or from the electrodes and tracks (Fig. 2.5).

(a)

(b)

Figure 2.4: The filtering properties of Channelrhodopsin-2 variants at room and body temperatures. *a)* Plotted here are the theoretical curves generated from a Markov model and experimental data. "ChR2" is Channelrhodopsin-2, "$ChR2_R$" is the H134R mutant, "$ChR2_A$" is the E123T/H134R mutant also known as "ChETA" (Image: Tchumatchenko and Newman [2013]). This plot shows that a cutoff for Channelrhodopsin-2 of $\approx 70\,\mathrm{Hz}$ at 36°C. *b)* The filtering properties of Channelrhodopsin-2 and Chronos variants fit to single pole low pass filter responses (Image: Ricardo Merino, Max Planck Inst. for Experimental Medicine). This plot shows that the cutoff for Channelrhodopsin-2 is $\approx 24\,\mathrm{Hz}$ at 25°C.

2.1.4 Cell Cultures

In vitro experiments employed cortical cultures. These were prepared by either Manuel Schottdorf or Sabine Klöppner in our laboratories at the Max Planck Institute of Experimental Medicine. The protocol used was similar to those employed previously in our laboratory (e.g. in Samhaber et al. [2016]), following Brewer et al. ([1993]). Cortical cells were extracted from embryonic rats (Wisteria WU) either 18 or 19 days after gestation (E18/E19). The impregnated rat was first anaesthetised using CO_2, then euthanized using cervical dislocation. The embryos were then removed via caesarean section and decapitated. The heads were placed in cooled petri-dishes. Each brain was then removed from its seat in the skull cavity. Their cortices were surgically separated and transferred to a HEPES[1] buffered Neurobasal medium (10 mM HEPES). Once all cortices were extracted and settled in the buffer solution,

---

[1] 4-Hydroxyethyl-1-piperazineethanesulfonic acid, Invitrogen, Germany

Figure 2.5: A fluorescence image of a cultured multi-electrode array as used in this work (Image: Manuel Schottdorf, Max Planck Inst. of Dynamics & Self-Organisation). Bright spots represent Channelrhodopsin-2 infected neurons. The dark points and lines are the electrodes and tracks of the multi-electrode array.

the supernatant was poured off and the cortices were trypsinised in a Trypsin/Ethylenedi-aminetetraacetic acid (EDTA)[2] buffer for 15 minutes at body temperature (37 °C). After trypsinisation, cells were moved to a solution of foetal bovine serum (FBS, 10%, Gibco), then thoroughly triturated. This was performed using a syringe with a 1 mm diameter needle. Cells were then separated from the solution through centrifugation for 2 min at a speed of 1200 rpm. After removal of the supernatant, the pellet was resuspended in a liquid medium comprising 2 ml of serum-free B27/Neurobasal[3] medium with an added 0.5 mM of Basic Fibroblast Growth Factor (bFGF) and glutamine. A Neubauer improved counting chamber was used to count the cells. The density was then adjusted to achieve 10,000,000 cells per millilitre, a 100 μl drop of which was placed atop the electrode field (area $\approx 1\,mm^2$) of a poly-d-lysine coated multi-electrode array (MEA). This resulted in a seed density of approximately 10,000 mm^{-2}. This was chosen to ensure good coverage of the electrode field. Lower densities resulted in the clumping of cells, such that they formed islands which may not sit over electrodes. The populated MEAs were then placed in a 37 °C incubator for 4 hours in a 5 % CO_2 environment. This gave the cells time to settle, after which the culture was supplemented with 1 ml of the B27/Neurobasal medium. After 2 days *in vitro*, cultures were rendered sensitive to blue light through infection with 1μl of a Channelrhodopsin-2 (ChR2) delivering viral vector, AAV2.9-hSyn-hChRII(H134R)-eYFP-WPRE-hGH (Penn Vector Core, Uni. of Pennsylvania). Cells expressing ChR2 were tagged with the fluorescent marker, YFP (Fig. 2.5). Medium changes took place every 7 days and were typically performed by either Manuel Schottdorf, Sabina Klöppner, Julian Vogel, or on occasion by me. This entailed removing 500 μl (half) of the medium and replacing it with fresh medium.

All animals used for this work were both bred and raised according to German and European guidelines for experimental animals. This work was performed by colleagues at

---

[2]Trypsin: 0.05%; EDTA: 0.02%; Sigma Aldrich, Taufkirchen, Germany
[3]2% B27, Gibco

the Animal House, part of the Max Planck Institute of Experimental Medicine. All animal experiments were conducted as sanctioned by the state authority responsible for animal welfare in Lower-Saxony, Germany.

2.1.5   Blocking Synaptic Transmission

During experiments, it was important to sever synaptic communications between cells within the cortical culture. To this end, an array of toxins was employed to block both excitatory and inhibitory synaptic transmissions. The blockade comprised $100\,\mu$M APV[4], $100\,\mu$M Picrotoxin[5] and $50\,\mu$M NBQX[6]. These quantities were arrived at through laboratory trials performed by Manuel Schottdorf.

Both APV and NBQX hinder excitatory synaptic transmission. APV targets NMDA[7] receptors, acting as an antagonist, i.e. blocking the binding of agonists to the receptor, which would normally provoke a biological response constituting a transmission (Davies and Watkins 1982; Olverman et al. 1984). Similarly, NBQX targets AMPA[8] receptors, also acting as an antagonist (Sheardown et al. 1990). Picrotoxin hinders inhibitory synaptic transmission by acting as a channel blocker. Its target is the GABA_A[9] receptor chloride channels. Blockers bind inside the pore of the target channel, blocking the flow of ions (Boullay 1812; Newland and Cull-Candy 1992).

2.2   Simulations

Numerical studies were performed using NEURON (Carnevale and Hines 2006) version 7.4 and Python 2.7. Simulations were run either on a laptop (Lenovo ThinkPad, Intel Core i7) or on our clusters at the Max Planck Institute for Dynamics and Self-Organization, Dept. of Non-Linear Dynamics (director, Theo Geisel). The clusters are maintained by York-Fabien Beensen and Denny Fliegner. NEURON was used to simulate a multi-compartment neuron model, and channelrhodopsin-2 model and current clamp as described below.

2.2.1   The Neuron Model

The goals of this work are to examine the relationship between an subcellularly targeted stimulus, the neuronal population response, constituent onset rapidness and the effective passive bandwidth of constituents. Any simulated population response should therefore reproduce the high frequency response characteristics observed in experiments (Sect. 1.3). The model should also capture the morphology of a real dendritic tree and soma. To this end, I have chosen a model from the Blue Brain Project (Giles 2005; Markram 2011;

---

[4]2-amino-5-phosphonopentanoic acid
[5]An equimolar mixture of picrotin ($C_{15}H_{18}O_7$) picrotoxinin ($C_{15}H_{16}O_6$)
[6]2,3-Dioxo-6-nitro-1,2,3,4-tetrahydrobenzo[f]quinoxaline-7-sulfonamide
[7]N-Methyl-D-aspartic acid
[8]2-Amino-3-(3-hydroxy-5-methyl-isoxazol-4-yl)propanoic acid
[9]Named for the receptor's endogenous ligand GABA, γ-aminobutyric acid.

Figure 2.6: A morphologically realistic NEURON model of the layer V thick tufted pyramidal neuron, type I (Image: Hay and Van Geit 2015).

Markram 2006), whose key aim is to collect and unify experimental neuroscientific results in a single model-building system (Markram 2011). The model used was a multicompartment conductance based model of the layer V thick tufted pyramidal neuron, type I (Fig. 2.6, Hay and Van Geit 2015; Markram et al. 2015), which incorporates electrophysiological properties captured from *in vitro* slices of the somatosensory cortex of a 2 week old Wistar Han rat. The model contains a 3D morphological model of the soma as well as apical and basal dendrites reconstructed from said rat (Markram et al. 1997), the axon is a $60\,\mu$m long stub. The electrophysiological properties of the model were established following protocols described in Le Bé et al. (2007) and Wang et al. (2002, 2004). The active mechanisms implemented in the model capture the dynamics of Ca^{2+} concentration (Destexhe et al. 1994), low-voltage- and high-voltage-activated Ca^{2+} currents (Avery and Johnston 1996; Randall and Tsien 1997; Reuveni et al. 1993), hyperpolarisation-activated cation currents (Kole et al. 2006), M-currents (Adams et al. 1982), persistent and transient K^+ currents (Korngreen and Sakmann 2000), persistent Na^+ currents (Magistretti and Alonso 1999), Na^+ channel dynamics (Colbert and Pan 2002), Ca^{2+}-activated K^+ current (Kohler et al. 1996), Shaw-related K^+ channels (Rettig et al. 1992). AMPA, NMDA and GABAA receptors with presynaptic short-term plasticity are also modelled (Fuhrmann et al. 2002).

2.2.1.1   Passive Operation

To measure the effects of the purely passive components of the neuron, all gated ion-channels were deactivated by setting their respective conductances to $0\,S\,cm^{-2}$. The Channelrhodopsin-2 model was, of course, kept active, where the simulations necessitated a light stimulus.

2.2.2   Modelling Channelrhodopsin-2

To perform the numerical experiments described in Sect.2.5, it was important to numerically model the behaviour of Channelrhodopsin-2. To this end, I employed the hidden Markov
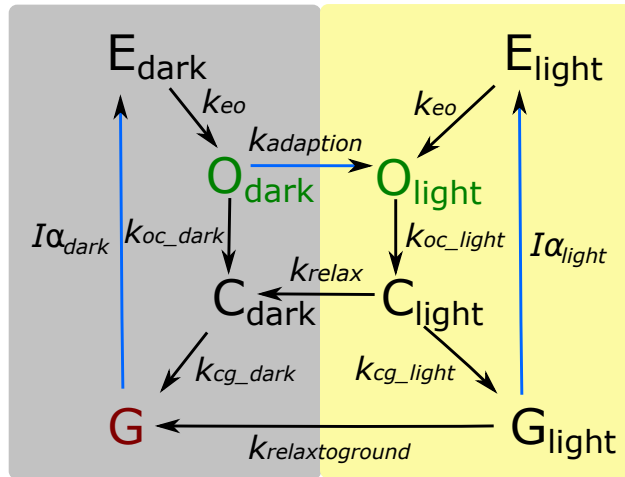
Figure 2.7: A model of Channelrhodopsin-2 kinetics (Image: Fromme 2016). This 8-state hidden Markov model comprises 2 branches, each with 4 states. The model depicts open states, O_{light} and O_{dark}; and closed states C_{light} and C_{dark}. The '$_{light}$' (grey box) and '$_{dark}$' (yellow box) refer to the two branches of the model which allow for differing light- and dark-adapted behaviours respectively, i.e. whether or not the channel has recently been exposed to light. The E_{light} and E_{dark} states represent the excitation of the channel before opening. The G and G_{light} states are the 'ground' states of the two branches. The G state is the ground state in which the model would start and eventually revert to if left in the dark for a sufficient amount of time.

model (HMM) developed by Fromme (2016) that captures Channelrhodopsin-2 kinetics. Fromme (2016) used genetic algorithms to fit HMMs of increasing complexity to data he recorded experimentally from Channelrhodopsin infected cultures. He concluded that an 8 state model with 2 branches (Fig. 2.7) best captures the kinetics of Channelrhodopsin-2 as well as Chronos, a faster Channelrhodopsin variant. Fromme (2016) kindly provided the MATLAB code for the model, which I translated to NMODL[10] for use in the NEU-RON simulation environment, a numerical tool for modelling neurons and networks thereof (Carnevale and Hines 2006). The transition rates of the model that I used are the same as those found in Fromme (2016) for Channelrhodopsin-2.

For the numerical experiments, it was important to ascertain the frequency response of the Channelrhodopsin-2 (ChR-2) model. To this end, the ChR-2 model was applied to a simple soma using the NEURON software package, and stimulated with a small light step. This was simulated at room temperature. Calculating the derivative of this step response gave the impulse response of the model, whose Bode magnitude closely fits that of a single pole low pass filter with cutoff frequency of 13.6 Hz (Fig. 2.8). This is lower than the cutoff of 25 Hz measured for ChR-2 in *in vitro* experiments (Sec. 2.1.3, Fig. 2.4b).

2.2.3 Current Clamp

In their series of seminal works (Hodgkin and Huxley 1952a,b,c,d; Hodgkin et al. 1952) Hodgkin, Huxley and Katz not only characterised the flow of sodium and potassium ions

---

[10]NEURON Modelling Description Language
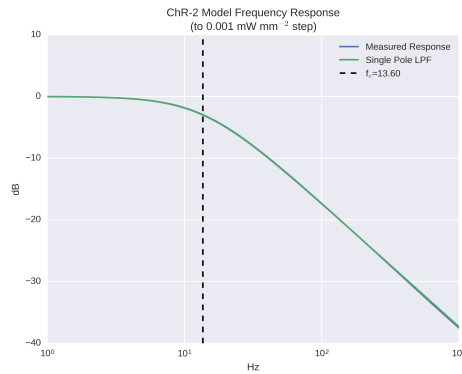
Figure 2.8: The Bode magnitude plot of the Channelrhodopsin-2 model for small light signals. This is plotted against a single pole low pass filter with a cut off of 13.6 Hz.

through a nerve membrane, but also demonstrated the power of the "voltage clamp" technique. The voltage clamp, whose invention is attributed to Kenneth Cole (Purves et al. 2004), is an approach to measuring currents through the membrane of a cell, while using an electrical feedback circuit to hold or "clamp" the membrane to a voltage or a series of voltages. A later extension to this idea, called "patch clamp" (Neher and Sakmann 1976), described a method by which a glass pipette containing an ionic solution and a recording electrode, sealed by suction ("cell attached") to a small point on the cell membrane, can be used to record the activity of a channel or channels in the membrane. Stronger suction can be used to puncture the cell and allow the insertion of the pipette, thus forming a continuous connection between the ionic solution in the pipette and the cell's cytoplasm. Recordings performed in this manner are called "whole cell" recordings (Hamill et al. 1981). Hamill et al. (1981) also describe so called "outside-out" and "inside-out" patching techniques which are not described here. Similarly, Hamill et al. (1981) mention puncturing the cell membrane, inserting the pipette tip, and recording in "current clamp" mode, whereby a current can be injected into the cell while the potential is measured. Magistretti et al. (1996) give simplified examples of circuitry which may be used to perform patch-clamp in voltage clamp mode, in current clamp mode and a current clamp specific circuit (Fig. 2.9). For numerical simulations, I applied current clamp using NEURON's built in IClamp mechanism (Carnevale and Hines 2006).

2.3   Stimuli

In both numerical and laboratory experiments, a correlated, stochastic stimulus was required. For population response measurements, and additional modulating change in the mean of the process was also required. The following sections describe both the stochastic process, the Ornstein-Uhlenbeck process, and the modulating, preconditioned steps.
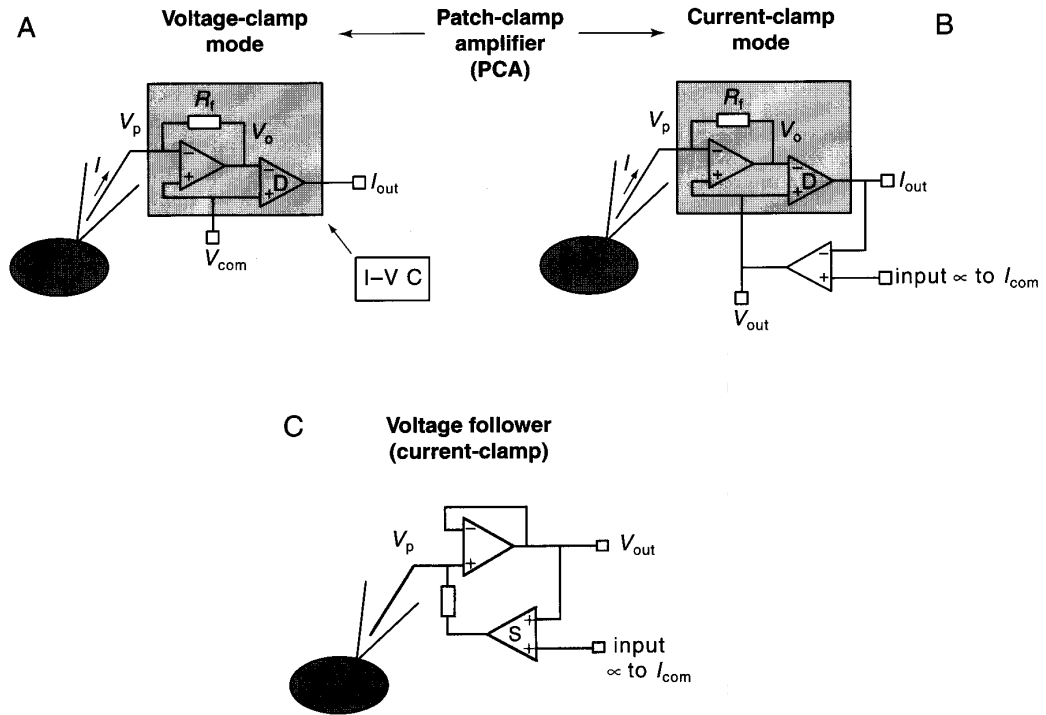
Figure 2.9: Patch-clamp, simplified circuit diagrams (Image: Magistretti et al. 1996). *A)* Patch-clamp circuit, voltage clamp mode. *B)* Patch-clamp circuit, current clamp mode. *C)* Current clamp specific (voltage follower) circuit.

2.3.1 The Ornstein-Uhlenbeck Process

The Ornstein-Uhlenbeck process (OU process) is a stochastic process named after the Dutch physicist Leonard Ornstein and the Dutch-American physicist George Eugene Uhlenbeck. It can be thought of as describing a massive Brownian particle's velocity when influenced by friction. It is a stationary[11] Markovian[12] process that is also Gaussian (Uhlenbeck and Ornstein 1930). The realisation of the OU process used in this work is

$$\frac{dy}{dt} = \frac{\mu - y}{\tau} + \sqrt{\frac{2\sigma^2}{\tau}}\xi(t), \tag{2.1}$$

where μ is the mean of the process, σ^2 is the variance, τ is the correlation time and $\xi(t)$ is a Gaussian white noise process with a mean of 0 and a standard deviation of 1.

In neuroscience, the OU process has been shown to approximate the input background 'noise' that a single cell observes from a large number of upstream counterparts. In essence, the sum of their uncorrelated outputs. To this end, the OU process is used in this work to simulate a cell receiving inputs from thousands of connected cells (Destexhe and Rudolph-Lilith 2012; Lánský and Rospars 1995; Ricciardi and Sacerdote 1979). The idea being to test the cells response to a subtle change in the mean of this process, to see how quickly

---

[11]When shifted in time, there is no change observed to the joint probability distribution.

[12]'Memoryless,' i.e. a prediction about the process' future based on its entire history is as good as a prediction based only on the current state.

cells can respond as a population as done in much of the work discussed in Sect. 1.3.

2.3.1.1   Real-Time Digital Realisation

To implement an Ornstein-Uhlenbeck process of arbitrary length, in real-time, an Euler-Maruyama solution was implemented (Kloeden and Platen 1992), similar to that described in Listing 2.1. For consistency, this algorithm was employed in both laboratory and numerical experiments. This was implemented in C^{++} in the former case and in python in the latter case. As the resulting process ultimately maps to values of light intensity, it was important to avoid the generation of negative values. The process was therefore clipped at 0, which, of course, lead to a shortening of the left tail of the OU process' distribution. To reduce the impact of this clipping, the standard deviation of the process (σ) was scaled relative to the mean (μ) such that $\sigma = \mu/2$, thus ensuring a spacing of two standard deviations between the OU process' mean and 0.

Listing 2.1: A python implementation of the Ornstein-Uhlenbeck process using the Euler-Maruyama method

```
# An OU process: Euler-Maruyama solution
# dt:    sample period
# mu:    mean
# sigma: standard deviation
# tau:   correlation time
# y:     the OU process
# y_out: the trimmed OU process
#
# Note:
# np.random.normal(0., 1.): value drawn from normal dist.
#                           with mean=0 and std=1.
sqrtdt = np.sqrt(dt)
sigcoeff = np.sqrt(2.*(sigma**2.)/tau)
y[i] = y[i-1] \
       + dt*((mu-y[i-1])/tau) \
       + sqrtdt*sigcoeff*np.random.normal(0., 1.)
# Light cannot be negative
y_out[i] = y[i] if y[i] >= 0. else 0.
```

2.3.2   Preconditioned Step

The population responses analysed in this work were from cortical cultures exposed to a correlated, stochastic stimulus mean modulated with a step. In the photostimulus case, the step was first preconditioned (Fig. 2.10) to account for the low pass filter-like behaviour of Channelrhodopsin-2 (Sect 2.1.3). In experiments, this preconditioned step was first scaled to meet the required step size, then added to the Ornstein-Uhlenbeck stimulus.

(a)

(b)

(c)
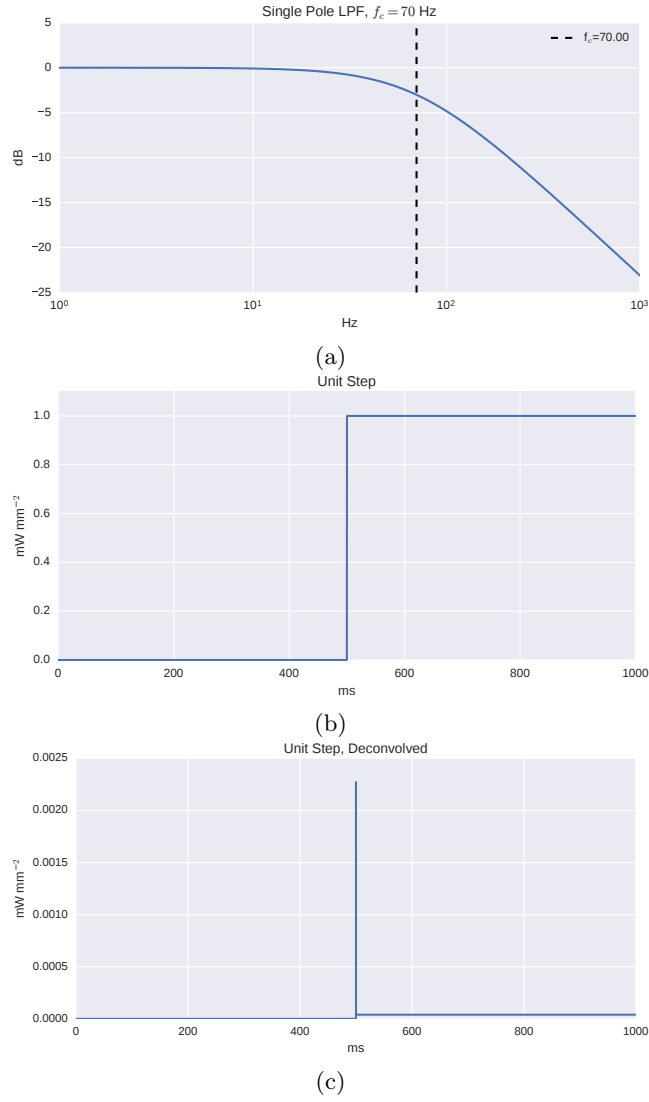
Figure 2.10: The preconditioned step used in *in vitro* experiments. This is done to account for the low pass filter behaviour of Channelrhodopsin-2 at body temperature (Fig. 2.4a). *a)* A single pole low pass filter with a cut off of 70 Hz. *b)* The unit step to be preconditioned. *c)* The preconditioned unit step. The result of deconvolving *b* with *a*.

2.3.3 Preconditioned Step (Model)

Similar to the laboratory experiments, the numerical experiments also required a preconditioned unit step. The frequency response of the Channelrhodopsin-2 (ChR-2) model (Sect. 2.2.2) for small light signals was smaller than that for ChR-2 measured *in vitro*: 13.6 Hz (Fig. 2.8) as opposed to 25 Hz at room temperature (Fig. 2.4b), therefore a separate preconditioning process was followed for numerical experiments (Fig. 2.10). Applying this preconditioned step (Fig. 2.11a) to the ChR-2 model induced a conductance (Fig. 2.11b) similar to a step, with a rise time of approximately 4.28 ms (Fig. 2.11c).

(a)

(b)

(c)
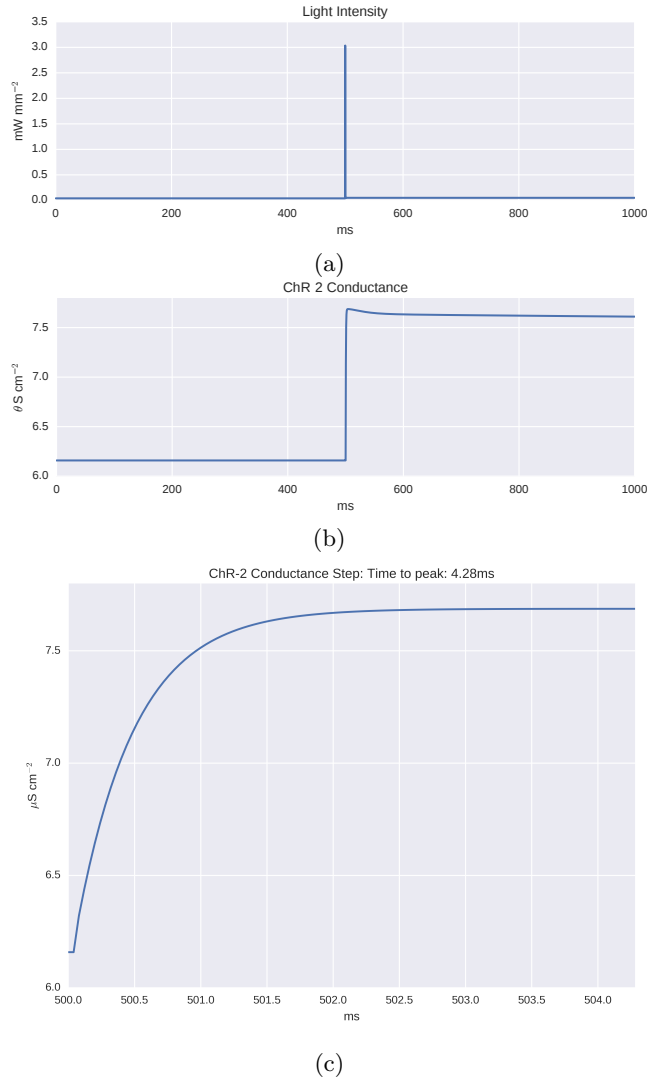
Figure 2.11: The preconditioned step used in the numerical experiments, and the resulting Channelrhodopsin-2 (ChR-2) induced conductance. *a)* The preconditioned step. *b)* The conductance induced when applying the preconditioned step to the Channelrhodopsin-2 model. This is not a perfect step, but sufficiently close for this work. *c)* A zoom of *b* showing the rise time of the step to be $\approx 4.28\,\text{ms}$

2.4 Analyses

To examine the results of both laboratory and numerical experiments, several key analyses were performed. These include local variation analysis, probability of detection analysis and bootstrapped confidence intervals and bands, all briefly summarised below.

2.4.1 Local Variation of Inter-Spike Intervals

Local variation (L_v) is a measure of a distribution's dispersion, i.e. how squeezed or stretched a distribution is. One can also think of L_v as describing how variable a distribution is relative to its mean. In neuroscience, measures of this form are typically used to analyse a spike-train's inter-spike intervals. That is, the times between adjacent action potentials, as detected in a single voltage trace. L_v can be computed as
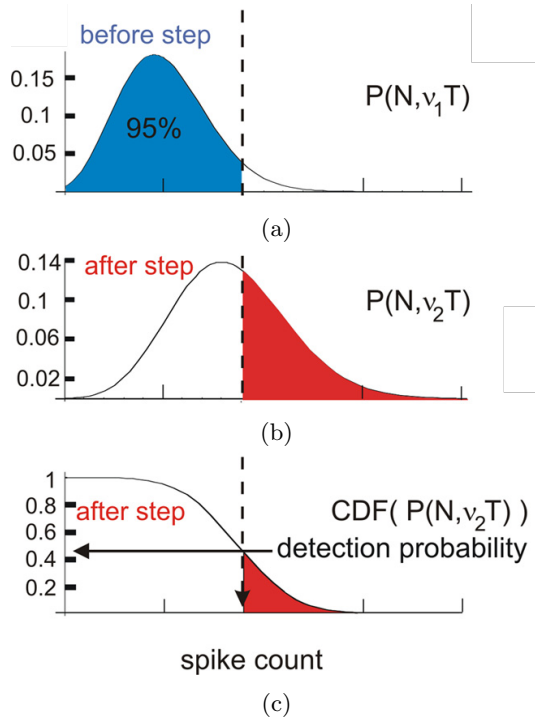
Figure 2.12: Computing the probability of detection of a step onset given the peri-stimulus time histogram (PSTH) of the neuronal population response and the step onset time (Image: Tchumatchenko et al. 2011). *a)* The distribution of PSTH bin-counts before the onset of the step. *b)* The distribution of PSTH bin-counts after the step-onset. *c)* The reverse cumulative distribution of *b*, which gives the probability of detection.

$$L_v = \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{3\left(T_i - T_{i+1}\right)^2}{\left(T_i + T_{i+1}\right)^2}, \tag{2.2}$$

where, n is the number of inter spike intervals (ISIs) and T_i is the i^{th} ISI (Shinomoto et al. 2003). The factor of three normalises the L_v value to that of a Poisson process. As such, a spike-train with an L_v value that is close to 1 behaves as a Poisson process. An L_v of 0 indicates a high, clock-like regularity.

2.4.2   Probability of Detection

To compare the speed of neuronal populations in both laboratory and numerical experiments, a theoretical detector was implemented as per Tchumatchenko et al. (2011). Given a known step onset time, this detector uses the peri-stimulus time histogram (PSTH) of the neuronal population response to calculate the probability of detecting said step. The algorithm computes and compares two empirical bin-count distributions, one before the step onset (Fig. 2.12a) and one after (Fig. 2.12b). The 95^{th} percentile of the pre-onset distribution is then computed and compared with the reverse cumulative distribution (Fig. 2.12c) of post-onset bin-counts to obtain the probability of detection.

2.4.3   Bootstrapping Confidence Intervals

For measurements made in both laboratory and numerical experiments, bootstrapped confidence intervals were computed, using the 'percentile interval' method (Efron 1979; Efron and Tibshirani 1994) as implemented in the python library, 'scikit.bootstrap' (Evans 2014).

2.5   Experimental Protocols

The hardware, software and experimental methods described in this chapter were employed in experimental protocols implemented in the laboratory and in simulations. The following sections detail the experimental protocols used to attain the results presented in Chapters 3 and 4.

2.5.1   Feedback Control & Population Response

Laboratory experiments were performed to examine the response of cortical neurons to a spike-rate controller and the response of a population of cortical neurons to subtle changes to the mean of a correlated, stochastic process. Both experiments were performed serially under the following protocol. Cortical cultures (Sect. 2.1.4) grown on multi electrode arrays (MEA) comprising cells longer than 21 days *in vitro* were taken from their incubator to the experimental set up situated in a dark room. The cultures were placed one at a time in the experimental set-up (Sect. 2.1.1) and examined for both tagged fluorescence and reactivity to light. Cultures could display anything from between zero to (rarely) 40 reactive channels. The typical active culture displayed between 10 to 30 active channels. Fluorescing, reactive cultures were then marked for use, one of which was placed in the experimental set up (Sect. 2.1.1) the remainder were placed back in the incubator.

The response of the chosen culture was then more closely examined. This involved first applying an Ornstein-Uhlenbeck (OU) process with a correlation time, $\tau = 15\,\text{ms}$ and a mean (μ) that increased linearly from 0 to $\approx 0.15\,\text{mW}\,\text{mm}^{-2}$ over a period of 10 seconds, followed by 5 seconds of zero light, then the equivalent OU process mirrored in time, i.e. with a μ decreasing linearly from ≈ 0.5 to $0\,\text{mW}\,\text{mm}^{-2}$ followed again by 5 seconds of zero light. The standard deviation, σ of the OU process was scaled with the mean such that $\sigma = \mu/2$. The process was of course clipped at $0\,\text{mW}\,\text{mm}^{-2}$. The protocol was typically repeated 10 times, where each OU process was trial-identical. The response at all channels of the MEA was recorded. Following this, a pharmacological, synaptic blockade (Sect. 2.1.5) was applied comprising $100\,\mu\text{M}$ APV, $100\,\mu\text{M}$ Picrotoxin and $50\,\mu\text{M}$ NBQX. This was done by first removing $100\,\mu\text{L}$ of medium (Sect. 2.1.4) from the culture, mixing this with the aforementioned quantities of blocker, then slowly pipetting this medium / blocker solution back into the culture. This was done to avoid clouds of concentrated blocker from reaching the cells, which was observed to result in cell death. After the application of blockers the same OU process was applied and the response at all channels recorded. Having recorded

both pre- and post-synaptic blockade responses, the trial-mean response per channel was then computed for each case. Based on this response, channels were chosen for closed loop control. Those channels that displayed a more linear response to the OU process "ramps" and reached rates significantly higher than 2 Hz were observed to respond best to the adaptive control protocol. One of these channels was chosen to be the probe point within the culture from which to perform the feedback control, the remainder were listed to be recorded alongside.

Using the synaptically blocked culture, the feedback control and population response protocols were applied in the following order. First the adaptation rate, γ (Sect. 1.2) of the controller was set, this was typically a value between 0.01 and 1, then the protocol was started. The controllers output set the mean (μ) of an Ornstein-Uhlenbeck (OU) process used as a stimulus. The OU process had a correlation time, $\tau = 15$ ms, the standard deviation, σ was set to $\sigma = \mu/2$. The feedback controller test protocol (Fig. 3.11, Chapter 3) began with 200 seconds of adaptation, where the controller was supplied a target spike rate of 2 Hz. After these 200 seconds of adaptation, the adaptor was turned off, and the PID controller was left to run. The behaviour of the controller and culture was then observed for approximately 300 seconds, if in this time there were indications that the controller was reacting too slowly (undershoot) or too quickly (overshoot), then the protocol was ended, and a new value of γ was chosen. An increase in γ was observed to increase the response speed of the controller. Once an appropriate value for γ was found, the protocol was left to run. As said, this comprised first the aforementioned feedback controller test protocol, followed by the population response protocol, which required the controller to hold the spike rate of the chosen channel at 1 Hz. Trial-based subtle changes to the mean were then applied as described in Sect. 2.3.2. The controller was then left to run for approximately three hours, the closed loop software (Sect. 2.1.2) was left to record the response from the controlled channel and up to 20 other channels, as well as the input into the LED driver, the photodiode output and a sync signal output by the DAC, which was later used to split the recording into its constituent trials. To avoid culture degradation (thus allowing reuse), cultures were not left longer than 4 hours under synaptic blockade. This was the limiting factor in the length of experiment. Following the completion of the experimental protocol, the culture was flushed with fresh medium. This involved first removing $500\,\mu$L of medium / blocker solution from the culture and adding $500\,\mu$L of fresh medium. This was repeated three times, then culture was placed back in the incubator.

2.5.2 Numerical Experiments

Numerical experiments were performed to examine the relationship between the subcellular targeting of stimuli, the response speed of a population, the onset rapidness of the constituent neuron's action potential and the effective passive bandwidth of a constituent neuron. The numerical experiments employed a morphologically realistic layer 5 pyrami-

dal neuron model, implemented using the NEURON software package (Sect. 2.2.1). Experiments examined this cell under both current clamped (Sect. 2.2.3) stimulation and Channelrhodopsin-2 (Sect. 2.2.2) photostimulation. Current clamping was implemented using NEURON's supplied IClamp mechanism. Channelrhodopsin-2 (ChR-2) was implemented using the model described in Sect. 2.2.2. The current clamped stimulus was applied only at the soma, whereas ChR-2 stimuli were applied either globally, to the soma or to the basal dendrites. This subcellular targeting was achieved by implementing ChR-2 channels only in the targeted portion of the cell's structure. In laboratory experiments, one would infect the entire cell, and modify the illumination. The distribution of light gated ion channels was effectively uniform across the target area of the membrane.

For the population response simulations, first a spike rate of 1 to 2 Hz was attained. This was achieved by "hand-tuning" (Table 2.1). In the current clamped somatic stimulation case, hand-tuning required adjusting the size of the stimulus, which, as in the laboratory experiments, was an Ornstein-Uhlenbeck (OU) process with a correlation time $\tau = 15$ ms, mean μ and a standard deviation $\sigma = \mu/2$. In the photostimulation cases, the same OU process was used; however, μ was fixed, and it was rather the effective density of the light gated ion channels that was adjusted. More detail on this process is provided in the following section. Once these model parameters were fixed, the population response protocol was applied. As in the case of the laboratory experiments, this was comprised of the aforementioned OU process with its mean modified by a single step protocol similar to that discussed in Sect. 2.3.3. The step protocol had a half second silence and a half second preconditioned step. An additional half second was prepended to the step to allow the modelled dynamics to settle, the traces capturing this settling time were then removed for analysis. This single step protocol was run 50,000 times on cluster computers. In each case, the OU process was uniquely generated. As mentioned, the step protocol was similar to that used in the laboratory, but not the same. The shape of the step protocol was computed by first estimating the transfer function of the Channelrhodopsin-2 (ChR-2) model and deconvolving a step function with this transfer function. The ChR-2 model's transfer function, at low light levels, was very close to a single pole low pass filter with a cut off of 13.6 Hz. Somewhat lower than the 25 Hz cut off measured in the laboratory (Sect. 2.1.3). During the application of the step protocol, voltage traces were captured from the membrane at the soma and axon to allow spike detection for both population response and onset rapidness measurements. In the photostimulation cases, traces capturing the induced ChR-2 conductance and the induced current were also captured. In the somatic and global photostimulation cases this measurement was made at the soma. In the basal dendritic stimulation case this measurement was made at the base of the basal dendritic tree. Traces capturing the stimulus were also recorded. In the photostimulus case, this was the light signal, in the current clamped case, this was the clamp current.

Table 2.1: The mean of the Ornstein-Uhlenbeck process used in each of the four stimulus cases

Stimulus	μ
Global photostim.	$0.4\,\mathrm{mW\,mm}^{-2}$
Basal dendrites photostim.	$0.4\,\mathrm{mW\,mm}^{-2}$
Soma photostim.	$0.4\,\mathrm{mW\,mm}^{-2}$
Soma curr. clamp	$171\,\mathrm{pA}$

2.5.2.1 Channelrhodopsin-2 Channel Densities

As mentioned above, in the photostimulus cases, the same light intensity was employed in each case, and the Channelrhodopsin-2 (ChR-2) channel density was adjusted to achieve the desired spike rate of $1-2\,\mathrm{Hz}$. This was done, as using a fixed channel density and adjusting the light intensity led to difficulties with the ChR-2 model. Namely, I was unable to find a single channel density that would provide similar behaviour across all three photostimuli. Typically either the global stimulus would demonstrate significantly less linearity than the other cases, or when the global stimulus demonstrated more linear behaviour, the somatic stimulus was unable to produce spikes. This is further demonstrated in some sense by the range of ChR-2 channel densities required (Table 2.2), which span an order of magnitude. This is of course congruent with the idea that a smaller membrane requires a higher channel density to achieve similar levels of conductance.

Table 2.2: Channelrhodopsin-2 channel densities, conductance and membrane area figures

	ChR-2 Density ($\mathrm{S\,cm}^2$)	Membrane Area (μm^2)	Membrane Area (%)	ChR-2 Conductance (nS)
Global	0.0014	48382	100.00	677
Basal dend.	0.0025	19371	40.04	484
Soma	0.035	1289	2.66	451
Axon	–	309	0.64	–
Apical dend.	–	27411	56.66	–

Chapter 3

Closed Loop Feedback Control of a Neuronal Culture

This thesis presents a novel, closed loop electrophysiological approach to controlling the rate of spiking of a neuron or neurons under a correlated, stochastic photostimulus. This closed loop system was then used to examine the response of an *in vitro* neuronal population to subtle changes in the mean of a correlated, stochastic photostimulus. An adaptive PID (proportional-integral-derivative) controller was used to control the neuronal activity at a single electrode within a neuronal culture using a correlated, stochastic photostimulus. From this probe point, a windowed spike rate is computed and passed to an adaptive PID controller. The controller outputs, in turn, a value used to control the mean of an Ornstein-Uhlenbeck process used to stimulate the culture. This chapter describes the results of using this controller in anger. Section 3.2 begins by describing the implementation of the adaptive PID controller in detail. This is followed by Sect. 3.3, which presents the results of measurements made to characterise the closed loop system. Section 3.4 presents the results of controlling a circuit, which emulates the behaviour of a culture. This is followed by Sect. 3.5 which describes controlling cortical cultures. Finally, Sect. 3.6 presents the results of using this closed loop controller to measure the response of an *in vitro* neuronal population to subtle changes in the mean of a correlated, stochastic photostimulus.

3.1   Motivation

During the course of experiments performed in our laboratories (Dept. Molecular Biology of Neuronal Signals, Max Planck Inst. for Experimental Medicine), my colleagues and myself observed that spike rates vary dramatically over the course of experiments (Fig. 3.1). Proposed causes include differences in the laboratory versus incubator environment and movement of the cells, although these have yet to be thoroughly investigated. In the incubator, the cell culture is exposed to a completely dark, high humidity, high CO_2 environment. During experiments, the culture is moved to the experimental setup in a lower humidity, lower
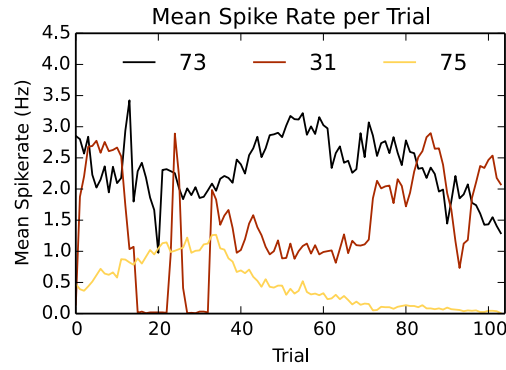
Figure 3.1: Spike rates captured during a trail based experiment. Each line follows spikes captured at a single electrode in a multi-electrode array. Each point plots the mean rate over a single trial. Trials are approximately 2 minutes in length, and the stimulus is identical in every trial. Activity changes significantly over the 100 trails shown (approximately 200 minutes).

CO_2 environment with light stimulation. The lower humidity likely contributes to increased evaporation of the medium and thus higher ionic concentrations. The introduction of light is also a change in environment to which the cells may adapt. Lastly, cell movement may be a factor that influences the effective contact between electrodes and neurons.

For experiments which run over several hours, the variation in spike rates ultimately influences the overall throughput of experiments. For example, the experimentalist configures an experiment to produce a trial rate of $2\,\mathrm{Hz}$, only to return several hours later, to find that the average response rate has dropped to e.g. $0.25\,\mathrm{Hz}$. The experimentalist could, of course, remain by the running experiments, and adjust the stimulus every few trials to achieve the desired average spike rate; however, as a software engineer interested in control algorithms, I thought it possible to automate such process and thus alleviate the experimentalist of this work. To this end, I developed a closed loop controller capable of regulating the spike rate of neurons, to reduce such trial-to-trial variations, without adversely affecting the trial average response.

3.2   PID Implementation

In order to maintain stable spike rates from a probe point within a cortical culture during experiments, a feedback controller was implemented. This feedback controller scaled the size of the applied stimulus in response to the live, windowed spike rate it received.

The feedback controller design used (Fig. 1.2) was first proposed by Lin et al. (2000). For this work, their design was extended upon and implemented digitally using MATLAB Simulink (Fig. 3.2). The PID portion of the implementation is based upon the MATLAB parallel, discrete time implementation (Mathworks Inc. 2015c)

$$K_p + K_i \frac{T_s}{z-1} + K_d \frac{z-1}{zT_s},$$

where K_p, K_i & K_d are the proportional, integral and derivative coefficients respectively; T_s is the sample period of the controller; $1/_{(z-1)}$ is the discrete integral and $^{(z-1)}/_z$ the discrete derivative. In addition to the original design, my implementation includes a clamping circuit in the integration path to avoid integrator wind-up, a low-pass filtered error signal, and a resettable integrator in the integrator (I) path.

This section describes some of the implementation details including the form of the integrators and derivatives, the adaptation rule and the windowed spike rate. The reasons for design decisions i.e. the clamping circuit, filter and resettable integrator, are also given.

3.2.0.1 The Discrete Derivative

The PID implementation (Fig. 3.2) employs a discrete derivative, labelled '$\frac{K(z-1)}{Tsz}$' (Mathworks Inc. 2015a). This is computed as per the equation

$$y(t_n) = \frac{Ku(t_n)}{T_s} - \frac{Ku(t_{n-1})}{T_s},$$

where t_n is the time at sample n, u is the input into the derivative, K is an arbitrary gain factor set at design time (in this case, 1) and T_s is the sample period of the derivative.

3.2.1 The Discrete Integral

The PID implementation (Fig. 3.2) employs a discrete integrator, labelled '$\frac{KTs}{z-1}$' (Mathworks Inc. 2015b). This is computed using the forward Euler method,

$$y(t_n) = y(t_{n-1}) + K(t_n - t_{n-1})u(t_{n-1}).$$

Here, t_n is the n$^{\text{th}}$ sample time, u is the integrator input, and K is an arbitrary gain, set, in this case, to 1. The actual computation is performed in two steps. First the next integrator state, $x(t_{n+1})$ is computed

$$x(t_{n+1}) = x(t_n) + KT_su(t_n).$$

The output is the current integrator state, as computed in the previous time step.

$$y(t_n) = x(t_n).$$

The initial state of the integrators, $x(t_0)$, is set at design time, in this case, to 0.

3.2.2 The Adaptation Rule

For the purposes of this work, the adaptation rule (Eq. 1.2) is applied only to initialise the controller, for a period of 200 seconds, before control begins. The training target is a simple step from 0 to 2 Hz, where 2 Hz is held for the entire 200 seconds.
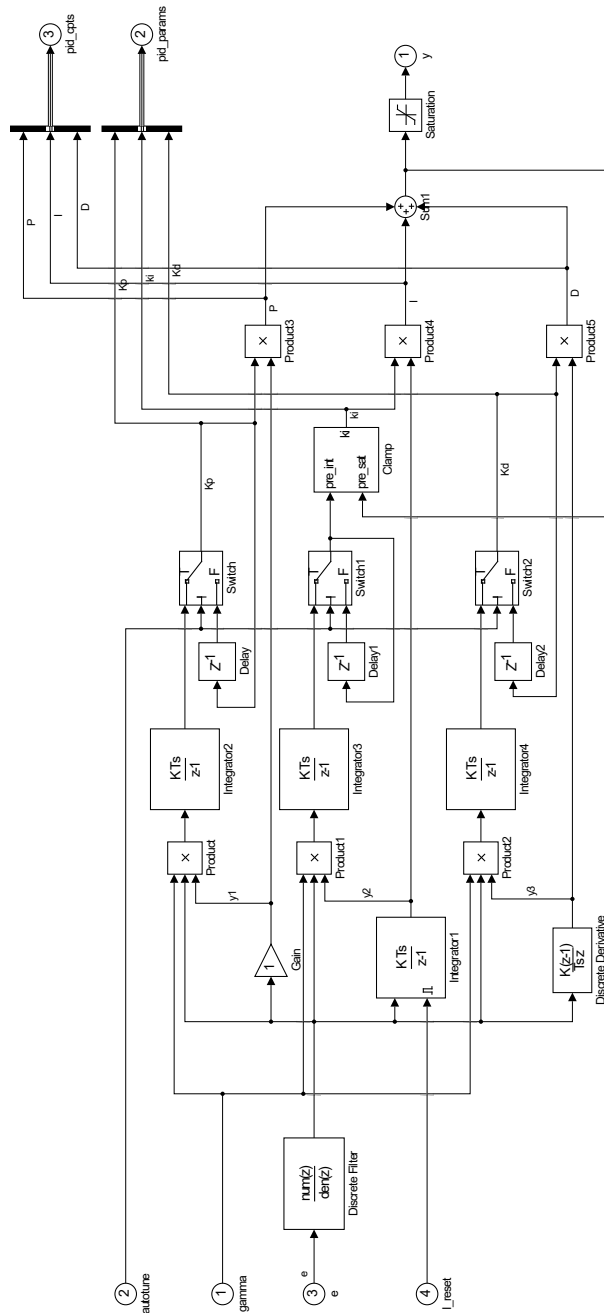
Figure 3.2: A digital realisation of the controller, as implemented in MATLAB's Simulink. The error signal is passed to input 3, input 1 accepts a potentially varying adaptation rate, input 2 controls the 'on'-state of adaptation, input 4 the reset-state of the integrator. Output 1 is the control signal. Outputs 2 and 3 give the PID components (outputs of the P, I & D signal paths) and the controller coefficients, respectively. Discrete derivatives and integrals replace their continuous counterparts depicted in Fig. 1.2. To further facilitate control, a clamping circuit is added to the integrator path; the error is also low pass filtered and the integration (I) path integrator is resettable.

3.2.3 Integrator Clamp & Reset

The adaptive PID controller implementation (Fig. 3.2) shows a clamping circuit and reset switch in the integration (I) path. The purpose of these components is to avoid the phenomenon known as *integral wind-up*. Integral wind-up causes excessive overshooting of the target, which typically results from a jump in the control target. A jump produces an initially large error which accumulates a large value in the integral path of a PID controller. This not only pushes the system into overshoot, it also saturates the actuator, in this case the LED. The result of which is slow controller response, as the controller takes time to 'unwind' the accumulated error (Åström and Hägglund 2006). In the case of this system, integral wind-up is reduced using two strategies 1) a clamping circuit for small jumps in the control target and 2) integrator reset for large jumps. These two approaches essentially disengage the integrator under different conditions.

For small changes, within 50 % of the new target value, an integrator clamping circuit controls the 'on'-state of the integrator. The clamping circuit (from Matlab's Discrete PID controller Mathworks Inc. 2015c), implements logic that switches integration off, if the control signal saturates. This is done by switching the integral (I) coefficient to $K_i = 0$. Once the control signal moves out of the saturation region, integration is switched back on i.e K_i is set to its original value.

For larger control errors (more than 50 % of the target value), the integrator is reset. The state of the integrator is held at 0, thus disengaging the integrator. In this case, proportional and differential components take control, bringing the error within 50 % of the target before the integrator re-engages.

3.2.4 Error Filtering

As the rate window, and therefore error signal can be somewhat jumpy, i.e. at the arrival of a spike, a low pass filter is applied to the error signal (Fig. 3.2). Without this filter, both proportional and derivative paths of the PID controller would produce correspondingly jumpy outputs. The filter employed is a fourth order, low-pass Butterworth with a cutoff frequency of 0.5 Hz (Fig. 3.3). The low cutoff of the filter introduces, of course, some delay to the system.

3.2.5 The Windowed Spike Rate

To compute the on-line spike rate, an exponential window of the form

$$\frac{1}{10}e^{-t/_{10}} \tag{3.1}$$

is used. This spreads the impact of a spike over several seconds after its occurrence. To generate the spike rate, the rate window is essentially convolved with a train of delta functions, where each delta represents a spike. The resulting train has a jump of $^1/_{10}$ as a spike

Figure 3.3: The frequency and phase response of the filter applied to the error signal at the input of the PID controller. This is a fourth order, low-pass Butterworth filter with a cutoff of 0.5 Hz.

arrives, which is then smeared out over the following seconds.

3.3   System Characterisation

3.3.1   System Response Speed

The response speed of the system is an important limit to measure. It gives us an understanding of e.g. the lag that a feedback controller must compensate for and also sets our expectations for the types of experiments that can be performed. The measurement made, in this case, tells us how quickly the system can detect and react to an action potential. This entailed generating artificial spikes at the input of the multi-electrode array amplifier. A modified version of the 'closed-loop' software was used to detect these spikes and to respond, in turn, by sending a pulse to the LED via the digital-to-analogue converter and the LED driver. This test measures the time between a spike recorded from the multi-electrode amplifier and a pulse recorded from the photodiode mounted in the condenser of the inverted microscope. To generate artificial spikes, a function generator was connected to a blanking circuit (Multi Channel Systems GmbH 2007) and placed in the multi-electrode array amplifier. The function generator was configured to generated 2 ms pulses of -1 mV once per second.

The closed-loop software recorded both the spikes measured at the multi-electrode array amplifier and the corresponding light pulses recorded from the photodiode (Fig. 3.4a). Two Gaussian pulse trains were then computed, where the pulses of one train correspond to the artificial spikes and the pulses of the other train correspond to the light pulses (Fig. 3.4b). These two pulse trains were then correlated to find the lag (Fig. 3.4c). The peak of the
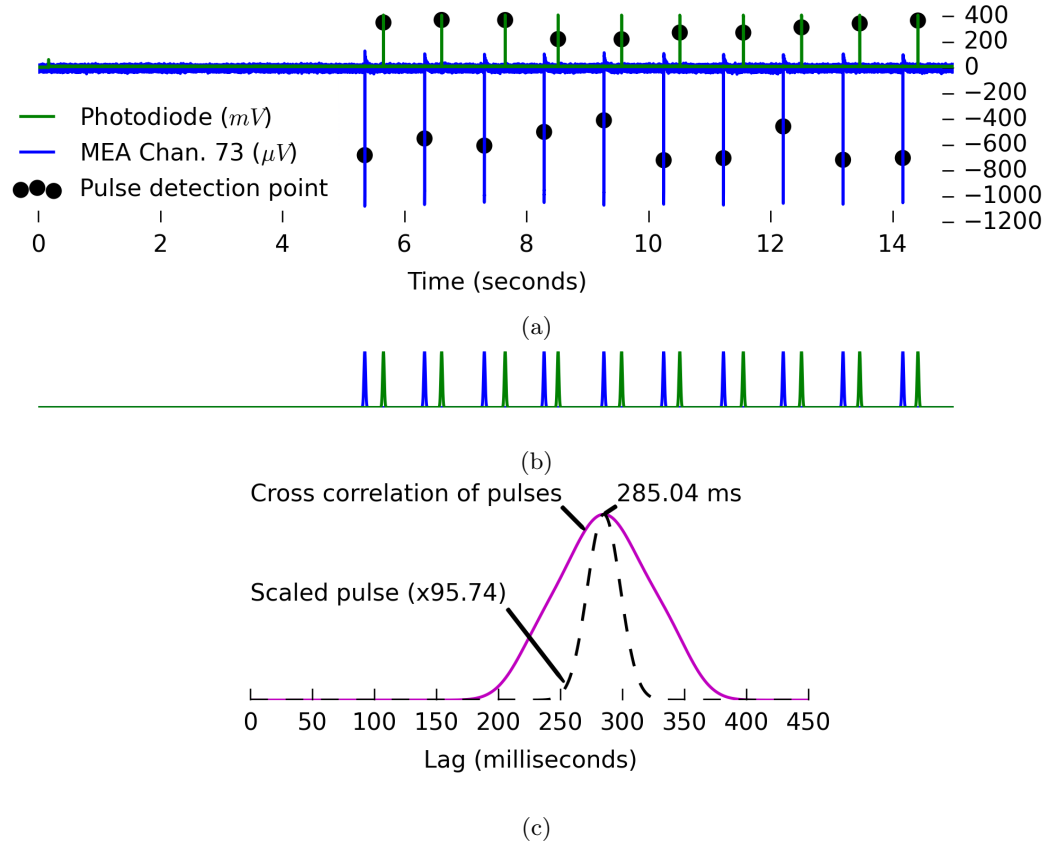
Figure 3.4: Measuring the system response speed. *a)* Spikes generated by a function generator were detected and recorded by the closed loop system (blue). In response to a detected spike, a pulse of light was triggered by the system. This was measured and recorded from the photodiode by the system (green). Offline, the recorded 'spikes' and light pulses were detected (black dots). *b)* The 'spike' (blue) and light pulse (green) times detected in *a* convolved with a Gaussian to create two trains of Gaussian pulses. *c)* These were then cross correlated to produce a lag plot, giving a lag centred about ≈285 ms.

resulting correlation function indicates a response lag of ≈285 ms.

3.3.2 Light Intensity Measurements

During light stimulation, cortical cultures are exposed to a range of light intensities. To measure the range of these intensities, a calibrated photodiode (S2386-8K, Hamamatsu Photonics Deutschland GmbH 2015) and multimeter (Keithley) were used. The photodiode was positioned on the stage of the inverted microscope (Sect. 2.1.1) where the MEA would normally sit. This was connected to the multimeter measuring in DC current mode. Commercial software, MC_Stim (Multi Channel Systems GmbH 2008), was used to send a voltage to the digital-to-analogue converter (DAC), the corresponding current produced by the photodiode was measured (Table 3.1). Using these measurements, along with the conversion factor of $274\,\mathrm{mA\ W^{-1}}$ for blue light (≈470 nm), as provided in the calibration datasheet (Hamamatsu Photonics Deutschland GmbH 2015), the corresponding light intensity values at the electrode field were computed (Table 3.1, Fig. 3.5).

Table 3.1: The relationship between the input into the digital-to-analogue converter (DAC) and the light exposed to a cortical culture.

DAC input (mV)	Calibrated photodiode output (μA$\pm$1)	Light intensity (mW mm^{-2} $\pm10^{-4}$)
50	0	0
75	8	0.0009
100	25	0.0027
150	57	0.0062
200	104	0.0113
500	328	0.0356
1000	680	0.0738
2000	1310	0.1421
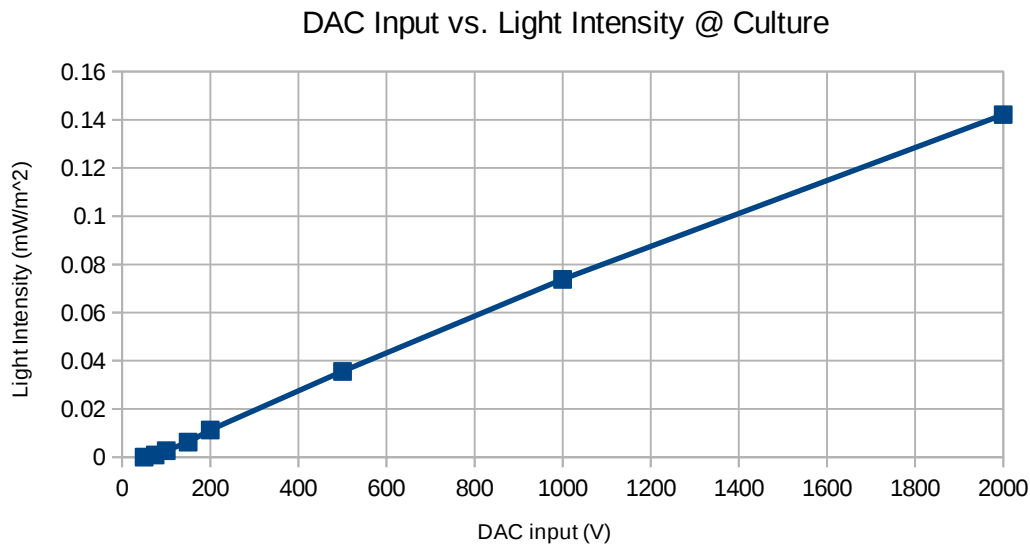
Figure 3.5: The relationship between digital-to-analogue converter (DAC) input and the light intensity that cell cultures are exposed to during stimulation.

3.3.3   LED Driver Calibration

The above light intensity measurements reveal a significant non-linearity near the switch-on point of the LED/LED driver (Figs. 3.6a & 3.6b). To reduce the impact of this non-linearity on the stimulus, the 'switch-on' point of the LED/LED driver was determined. This was done using the software MC_stim (Multi Channel Systems GmbH 2008) to increase the voltage, step-by-step, fed to the digital-to-analogue converter until blue light was observed at the LED. The switch-on point of the LED/LED driver was measured to be 49 mV. All stimuli used in experiments were offset by this amount (e.g. Fig. 3.6c & 3.6d), which prevented the heavy clipping seen before calibration (Fig. 3.6b).

3.3.4   Photoelectric Effect

To ascertain whether a photoelectric effect is present and could influence the results of multi-electrode array recordings, I devised an experiment whereby a multi electrode array (MEA),

(a) Before calibration: DAC output

(b) Before calibration: Measured at photodiode

(c) After calibration: DAC output

(d) After calibration: Measured at photodiode

Figure 3.6: Calibrating the photodiode. Due to 'switch-on' non-linearities in the LED controller / LED, a small offset must be added to any signal sent to the digital to analogue converter (DAC). An Ornstein-Uhlenbeck process (clipped at 2 standard deviations) was sent to the LED and recorded at the photodiode. The offset was increased until no clipping of the input signal was measured at the photodiode. This occurred at an offset of 49 mV. *a & b)* Before Calibration. *c & d)* After Calibration. *a & c)* An Ornstein-Uhlenbeck stimulus as measured at the input to the LED controller (mV, see Sect. 2.1 for details on the experimental set-up). *b & d)* The result of passing *a & c* through the LED controller and LED as measured at the output of the photodiode (mV).

filled with cell medium (Sect. 2.1.4), is placed in the experimental setup (Sect. 2.1.1) and exposed to pulses from the stimulus LED. The stimulation protocol comprised a half second pulse of light and a half second of darkness, which was repeated 40,000 times. Pulses had an amplitude of 4 V at the DAC input, which corresponds to approximately $0.3\,\mathrm{mW\,mm^{-2}}$ at

(a)

(b)
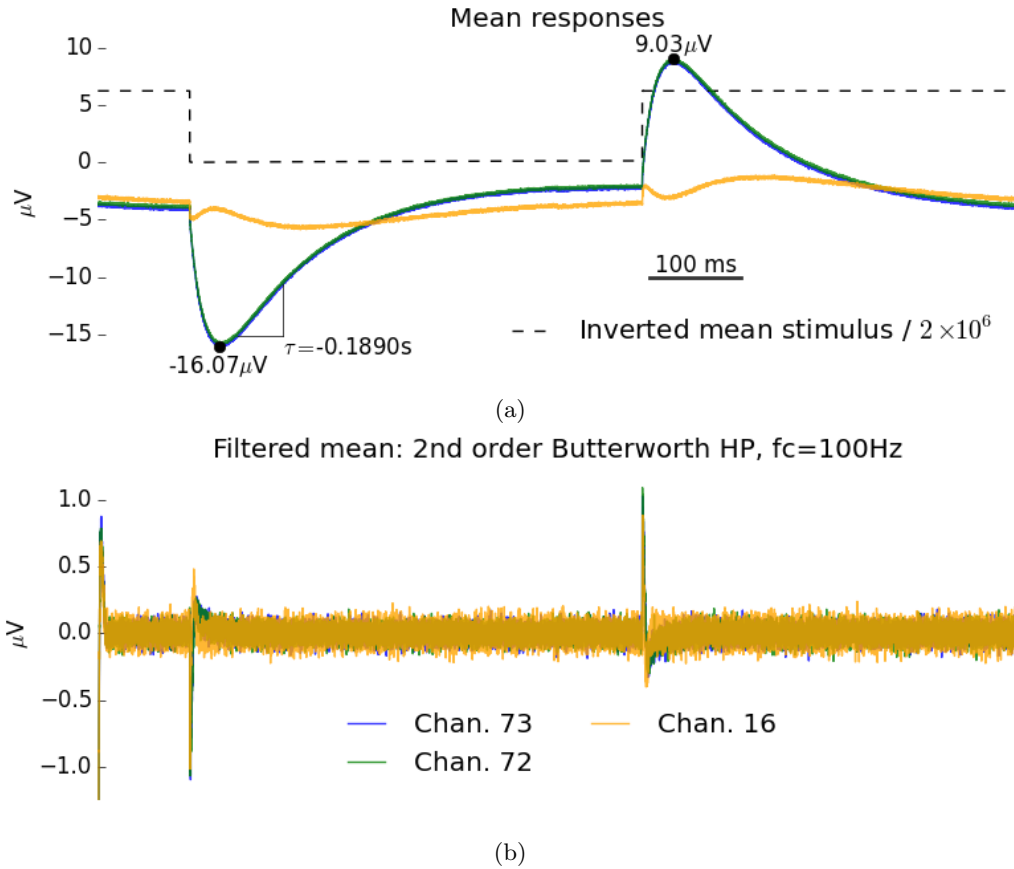
Figure 3.7: The photoelectric effect observed from three different channels on a single MEA. The results are computed from 40,000 half second pulses applied to a multi electrode array covered with $1000\,\mu$L of medium. Shown are the results from channels 72, 73 and 16. These are representative of the range of responses observed. *a)* The mean response measured at the electrodes. *b)* The mean filtered response. The filter applied is a second order high pass Butterworth filter with a cutoff frequency of 100 Hz. The same is used during experiments to reject low frequency oscillations.

the electrode field of the MEA. This is at the upper extreme of light intensity required in our experiments. The results (Fig. 3.7a) demonstrate significant swings in the mean potential measured at the electrodes. However, when one applies a second order high pass Butterworth filter with a 100 Hz cutoff, as is done during experiments to remove low frequency oscillations, the impact of these swings is less evident (Fig. 3.7b). The mean filtered response demonstrates at most a $1\,\mu$V change in potential, which is more than an order of magnitude below the typical voltage threshold used for spike detection. To test that the effect is indeed caused by the photostimulus, the experiments were repeated with card blocking the path of the photostimulus to the multi-electrode array. With this card in place, no effects were observed.

3.4 The *in Silico* Behaviour

Due to the huge range of dynamics observed in the neuronal cultures, it was important to first verify the behaviour of the controller *in silico*. Thus, a test circuit was built that converts light intensity to shaped pulses resembling action potentials, whose frequency increases relative to the light intensity. This circuit, built by Walter Stühmer (Director, Dept. for the Molecular Biology of Neuronal Signals, Max Plank Institute for Experimental Medicine), was first placed in the experimental set-up as a cell culture ersatz, then configured to produce spikes at approximately 2 Hz at a (DAC[1]) voltage of 100 mV. The controller was then trained and tested against a varying target rate.

The system was tested in two stimulus modes 1) using a deterministic stimulus, where the controller's output was fed directly to the LED driver and 2) using the stochastic Ornstein-Uhlenbeck (OU) process, where the controller's output scaled, simultaneously, the OU process' mean (μ) and standard deviation ($\sigma = \mu/2$). In both cases, the LED input was limited to a range from $0 - 4$ V (Figs. 3.8 – 3.9).

The stimuli (Figs. 3.8a & 3.9a) appear to scale largely linearly with the target (Figs. 3.8b & 3.9b) with some exceptions. Given the linear behaviour of the culture ersatz circuit, this is to be excepted. The notable exceptions are at step changes in the target and the noise in the stochastic stimulus (Fig. 3.9a). The step changes in the target result in an overshoot in the stimulus. This is likely due to the corresponding step change in the error for which the controller attempts to compensate. Also evident is the relative noisiness of the stimulus in the stochastic stimulus case (Fig. 3.9a), this is likely due to the stochasticity of the stimulus, which contributes to spiking that is somewhat non-regular when compared to the non-stochastic stimulus. This is also observed in the contrasting local variation (LV) values of the two cases, where the non-stochastic stimulus (Fig. 3.8c) results in a narrower distribution of LV values than in the stochastic case (Fig. 3.9c). As the culture ersatz circuit measures and accumulates the light intensity linearly, a clock-like spiking was expected, as is indicated by the near-zero LV values computed in both cases.

In both cases, the target rate (Figs. 3.8b & 3.9b) appears to be closely followed, with the exception of some overshoot, although in the case of the stochastic process stimulus, the controller produces a somewhat less flat response than in the deterministic case. The observed overshoot is a sign of sub-optimal control, i.e. the controller reaches the target value quickly, but more quickly than the controller can compensate for, resulting in an overshoot in the controller rate. During laboratory trials, lowering the learning rate, γ, of the controller (Sect. 1.2) did reduce overshoot, but also slowed convergence. As the observed level of overshoot was not of dire consequence in this situation, a faster convergence with overshoot was favoured. The less flat response of the controller in the stochastic stimulus case is, as discussed above, likely a result of the stochasticity of the stimulus and the corresponding relative irregularity (Figs. 3.8c & 3.9c) of spiking.

---

[1]The voltage sent to the digital to analogue converter (DAC)

(a)

(b)

(c)

(d)

(e)
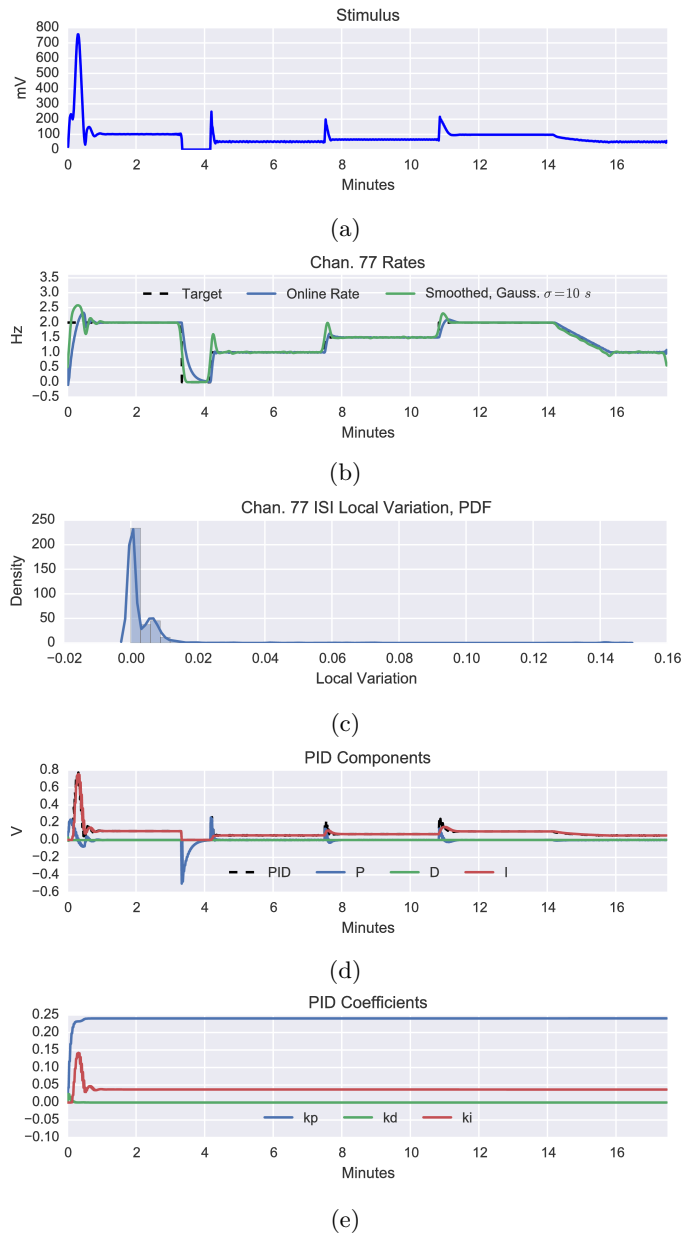
Figure 3.8: Following a moving target. Controlling the cell culture ersatz with a deterministic stimulus. *a)* The downsampled controller output (Fourier method). Except for spikes at step changes, the stimulus appears to change linearly with the target. *b)* The target, online and offline calculated spike rates. Except for overshoots at step changes, the rates seem to follow the target closely. A small offset error is evident in the downward ramp. *c)* The distribution of ISI local variation (Eq. 2.2, 10 s moving window). This is very close to 0, indicating a strong clock-like behaviour, expected from the electronic ersatz culture. *d)* The contribution of P, I & D components to the control signal (see *a*). The integral component contributes most to the controller output, except where it is disengaged. The proportional component provides the next most significant contribution. The derivative component appears to make no contribution. *e)* The P, I & D weights. The first 200 s show adaptation of the coefficients, after which adaptation is disengaged. The proportional coefficient is largest, the integral coefficient is approximately a fifth smaller. The derivative coefficient appears is or is near zero. The adaptation appears to have completed within 60 s.

Regarding PID coefficients (Figs. 3.8e & 3.9e), these converged to stable values within 100 s, suggesting that the adaptation period could be cut shorter. The proportional coeffi-

(a)

(b)

(c)

(d)

(e)
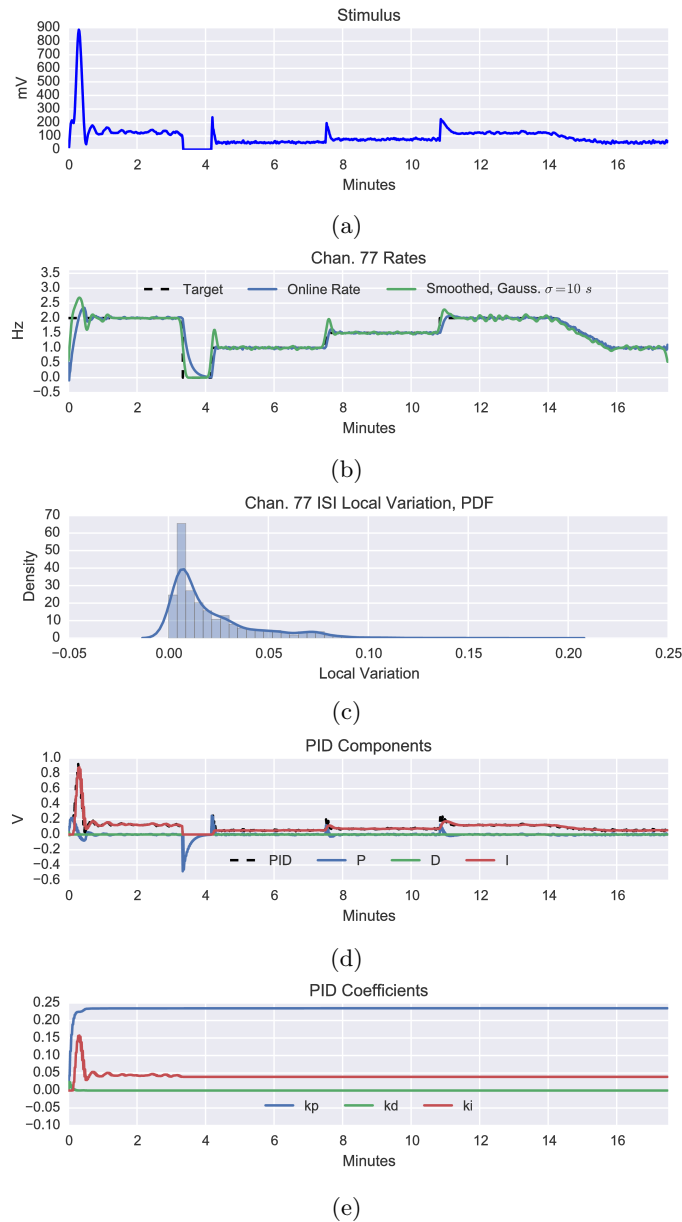
Figure 3.9: Following a moving target. Controlling the cell culture ersatz with a stochastic stimulus. *a)* The downsampled controller output (Fourier method). Except for spikes at step changes, the stimulus appears to change linearly with the target, more ripple is present than in Fig. 3.8a. *b)* The target, online and offline calculated spike rates. Except for overshoots at step changes, the rates seem to follow the target closely. A small offset error is evident in the downward ramp. Again, more ripple is present than in Fig. 3.8b. *c)* The distribution of ISI local variation (Eq. 2.2, 10 s moving window). This is broader than Fig. 3.8c, but still very close to 0, indicating a strong clock-like behaviour, expected from the electronic ersatz culture. *d)* The contribution of P, I & D components to the control signal (see *a*). The integral component contributes most to the controller output, except where it is disengaged. The proportional component provides the next most significant contribution. The derivative component appears to make no contribution. *e)* The P, I & D weights. The first 200 s show adaptation of the coefficients, after which adaptation is disengaged. The proportional coefficient is largest, the integral coefficient is approximately a fifth smaller. The derivative coefficient appears is or is near zero. The adaptation appears to have completed within 60 s.

(a)

(b)

Figure 3.10: The rise and settling time of the controller. Zoomed versions of *a)* Fig. 3.8b and *b)* Fig. 3.9b showing the first step after adaptation. Rise time is ≈7 s for the online calculated rate (blue, exponential window Eq. 3.1). For the Gaussian windowed rate, the rise time is ≈2 s, settling time ≈12 s.

cient dominates, while the integral component plays a lesser, however important role, the derivative component is near zero, suggesting that a PI controller would probably suffice.

Observing the PID components (Figs. 3.8d & 3.9d), the proportional component brought, in both cases, the spiking rate to the target value within 2 seconds, ultimately settling within ≈ 15 s (Fig. 3.10). Newman et al. (2015) report a rise time of 6 s. The integral component then gradually takes over, dominating the steady-state control. During the 'switch-off' at 200 s, the proportional (P) component takes, in both cases, a jump down to -0.5 V, and slowly increases to 0 V. The large jump can be attributed to the sudden -2 Hz error that appears, as the target is switched from 2 Hz to 0 Hz. The jump would likely be more dramatic, were a low pass filtering of the error signal not employed. The following slow convergence of the 'P' output to 0 is due to the long roll-off of the rate window, which is an exponential decay with a time constant of 10 s.

3.5   Feedback Control *in Vitro*

Using the experimental set-up in anger, the adaptive PID controller was observed to control the spiking rate of a single probe point in a Channelrhodopsin-2-infected cortical culture, the results of which are plotted in Fig. 3.11[2]. The cultures were stimulated using the stochastic Ornstein-Uhlenbeck (OU) process. Their responses show some similarities to those observed in the *in silico* case, with the notable exception of the local variation measures, the shape of the stimulus and the noise in the resulting rate measures.

The stimulus (Fig. 3.11a) scales somewhat proportionally to the the target (Fig. 3.11b); however, compared to the non-stochastic *in silico* case (Fig. 3.8a), the relationship between

---

[2]See Appendix A, Figs. A.1 & A.2 for additional examples. Due to their similarity to Fig. 3.11, they are not mentioned explicitly in the results.

(a)

(b)

(c)

(d)

(e)
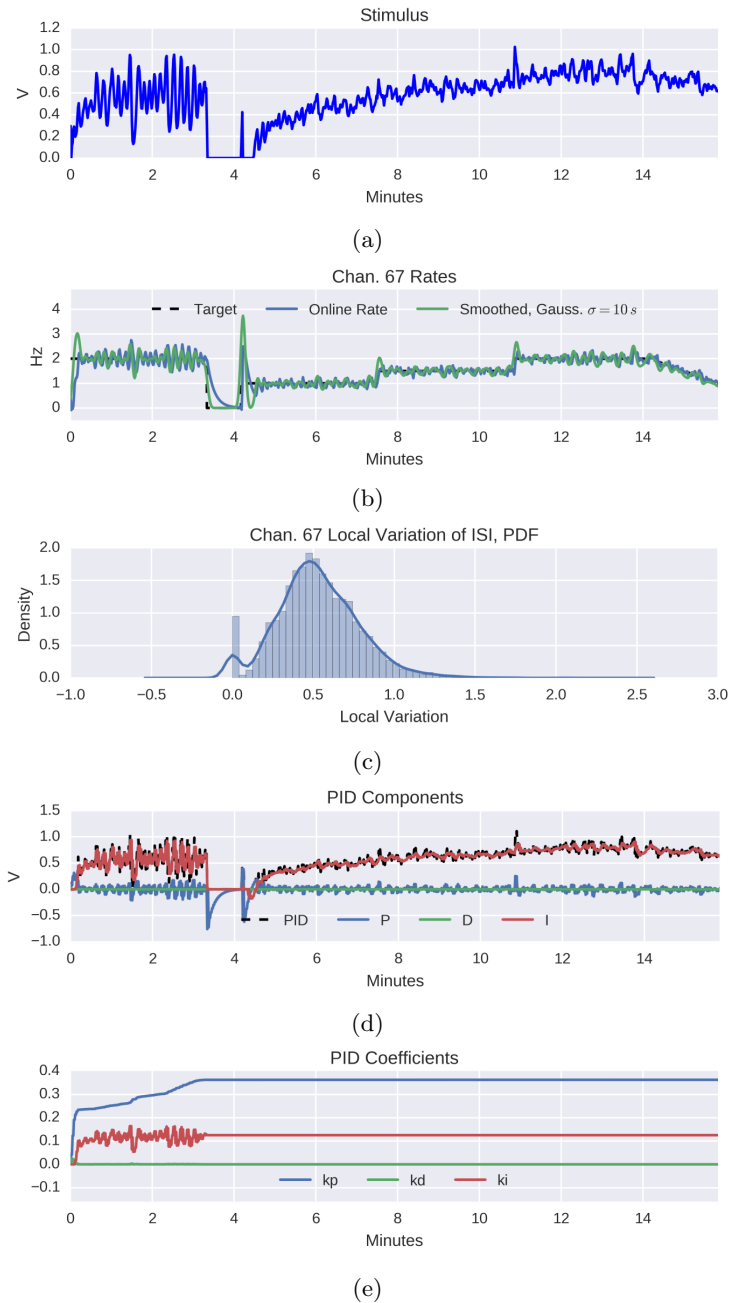
Figure 3.11: OU-process mode stimulus: Step response. *a)* The downsampled controller output (Fourier method). Unlike in the case of the culture ersatz, the stimulus is not perfectly proportional to the target, suggesting some non-linearities in the response of the cell or cells. *b)* The target, online and offline calculated spike rates. Although somewhat jumpy during the adaptation period (first 200 s), the cell or cells appear to follow the target, with some overshoot at steps and some noise. *c)* The distribution of ISI local variation (Eq. 2.2, 10 s moving window). This is distributed about 0.5, indicating significantly non-clock-like behaviour. *d)* The contribution of P, I & D components to the control signal (see *a*). The integral component contributes most to the controller output, except where it is disengaged. The proportional component provides the next most significant contribution. The derivative component appears to make no contribution. *e)* The P, I & D weights. The first 200 s show adaptation of the coefficients, after which adaptation is disengaged. The proportional coefficient is largest, the integral coefficient is approximately a third smaller. The derivative coefficient appears is or is near zero.

stimulus and target is not nearly as linear. This speaks to the relative properties of the cell culture ersatz circuit and a real culture. Despite choosing a channel that reacts relatively linearly in response to the stimulus (see Sect. 2.5), the culture displays significantly less linear behaviour than the cell culture ersatz circuit.

The target rate (Fig. 3.11b) is followed, with the exception of some overshoot in response to step changes in the target rate and some noise in the flat regions. As in the *in silico* case (Fig. 3.9b), the overshoots are likely a result of the adaptation rate γ and the resulting contribution of the proportional (P) component (Fig. 3.11d) of the controller. The increased noise in the flat-regions of the rate (Fig. 3.11b) is most likely due to the natural, non-clock-like spiking of the neurons (Fig. 3.11c), which produce in turn non-regular increases in the rate as a spike arrives. Although the *in silico* case indicates that, if the neuron were to respond with a clock-like rate, the stochasticity of the stimulus would also contribute to noise in the flat regions (Fig. 3.9b).

The local variation of inter-spike intervals has a broad distribution centred about 0.5 (Fig. 3.11c), whereas a Poisson process would have a distribution centred about 1. This implies that, although the cells are not spiking in a clock-like manner, their activity is more regular than that of a random Poisson process. As in the *in silico* cases, the proportional (P) and integral (I) components are the main contributors to the control signal (Fig. 3.11d), with a large negative jump observed in the P component as the target switches from two to zero Hertz. The integrator is deactivated at this point and is first reactivated after the P component brings the error signal to within $50\,\%$ of the target.

3.6   Population Response: Cortical Cultures

Experiments were conducted to examine the *in vitro* response of a cortical culture to subtle changes in the mean of a correlated, stochastic stimulus. Data was collected from seven channels of a multi-electrode array during one recording, made using one cortical culture. From the recording, of approximately 3 hours and 45 minutes in length, the controlled channel was used for analysis as well as selected portions of the six remaining channels, where their rates were maintained at approximately $1\,\text{Hz}$. The results (Fig. 3.12) give the response of cortical neurons to a repeated, one second, whole field stimulus comprising a trial-unique Ornstein-Uhlenbeck (OU) process with a small ($\approx 6.4\,\%$), trial-identical step change (Fig. 2.10c) to its mean. This step size was the smallest trialled that produced an observable response.

The results are given as a peri-stimulus time histogram (PSTH, Fig. 3.12a), which shows the response of a neuronal stimulus to the subtle step; a probability of detection curve (Fig. 3.12b), which gives the probability of detecting the step after its onset; and the trial average controller output (Fig. 3.12c), which indicates how the controller reacts to the population response.

(a)

(b) (c)
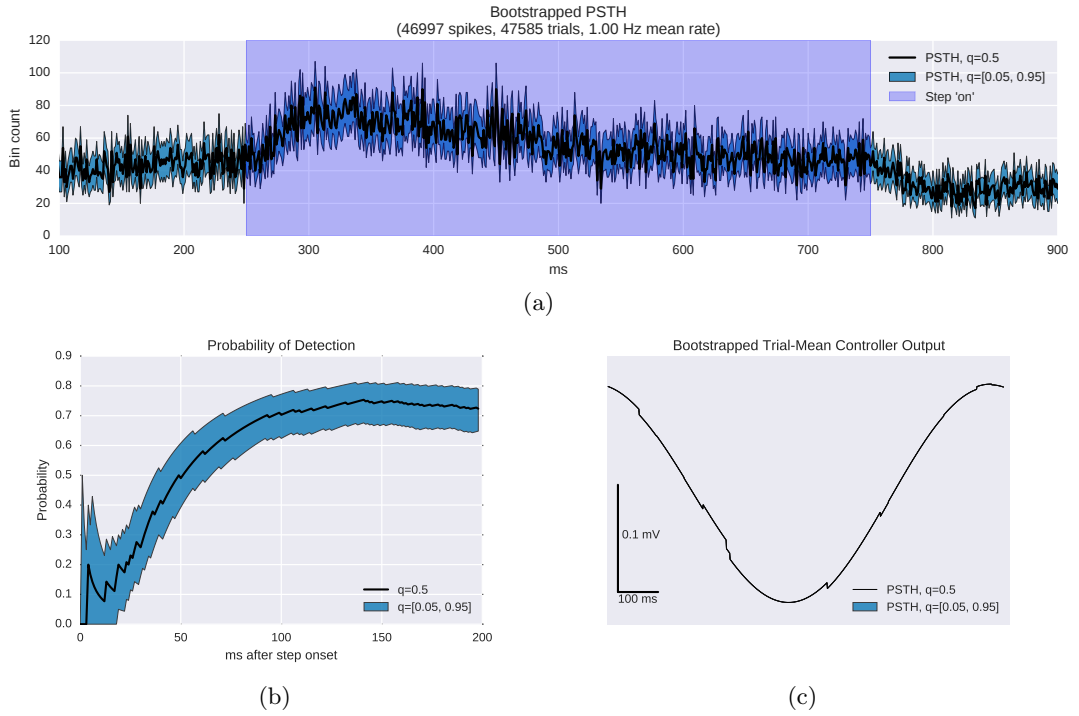
Figure 3.12: The response of an *in vitro* population to a trial-unique Ornstein-Uhlenbeck process with subtle (6.4 %) trial-identical step changes in the mean. Plots are computed from 47,585 trials. Plotted are the bootstrapped median and 95 % confidence bands. These were computed from 10,000 bootstrap samples with a sample size of 47,585, thus they simulate the response of a population of 47,585 neurons. *a)* The peri-stimulus time histogram (PSTH) of the population response, drawn here as lines, rather than bars, to better highlight the confidence bands. The light blue block indicates when the mean step was activated. *b)* The probability of detection as computed using the detector described by Tchumatchenko et al. (2011) also described in Sect. 2.4.2. The response first exceeds chance after 50 ms, and reaches a peak at 150 ms. This is significantly slower than the 20 pA step employed by Tchumatchenko et al. (2011) which comes very closed to a probability of detection of 1 within 2 ms. *c)* The trial average controller output. This demonstrates the influence of the trial-average population response on the controller. This results in a small (≈ 0.1 mV) oscillation. Note: The confidence bands are so narrow that they dot not appear visually in this plot.

The PSTH (Fig. 3.12a) demonstrates a clear albeit slow response to the step. The population appears to first react after several milliseconds, increasing to a peak after approximately 50 ms. Following this peak, the population response decays. After the light switches off, the population response reduces in kind, reaching its 'off' state after approximately 50 ms.

The probability of detection curve (Fig. 3.12b) confirms the clear albeit slow response of the neuronal population to a subtle step. The probability of detection first exceeds chance (0.5) after approximately 50 ms, and reaches a peak in the median of approximately 0.75 after approximately 150 ms. The probability of detection then begins to decay, likely due in part to the decay of the population response (Fig. 3.12a) after reaching its peak, but also to the behaviour of the detector (Sect. 2.4.2), which compares an ever increasing post-onset PSTH window with a static pre-onset PSTH window. This increasing window size is also likely responsible for noisiness at the beginning of the probability of detection curve, where the post-onset PSTH window consists of only several bins (i.e. 1 bin 1 ms after onset, 2 bins

2 ms after onset etc.).

Although slow, the response is overall rather robust. Especially when one considers what the signal to noise ratio (SNR) of a 6.4 % change in the mean of a stochastic process is. If the SNR is roughly estimated as the size of the step (0.064μ) divided by the standard deviation of the background stimulus $(\mu/2)$, then the SNR can be estimated as 0.128 or $20\log_{10}(0.128) \approx -18$ dB. A 0.75 probability of detection in such noise conditions is indeed a robust response.

Finally, the experimental protocol calls for the use of a controller to fix the average rate of spiking. The controller does this by scaling an Ornstein-Uhlenbeck (OU) process, or more specifically, sets its mean (μ) and scales its standard deviation $(\sigma = \mu/2)$. This is done by reacting to the spike rate measured at a probe point within the culture; thus, it is possible that the neuronal population, on average, influences the output of the controller significantly enough to modify the trial stimulus. This influence can by measured by examining the trial-mean response of the controller (Fig. 3.12c), which oscillates with a period of 1 s (also the length of the protocol), which indicates some trial-based influence on the controller behaviour. This influence, however, has an amplitude of approximately 10^{-4} V, significantly small in comparison to the size of the stimulus (e.g. Fig. 3.11a), which is more than two orders of magnitude larger.

Chapter 4

Subcellular Targeting

This thesis looks into the effects of applying stimuli to specific elements of a neuron's structure, called subcellular targeting. This chapter examines the relationship between an subcellularly targeted stimuli, the response of a neuronal population, the onset rapidness of its constituent neurons and their passive bandwidth. This is done by numerically modelling the application of an subcellularly targeted photostimulus to a morphologically realistic layer V pyramidal neuron. There are three subcellular targets: the soma, the basal dendrites and the entire cell, referred to as "global" stimulation. For comparative purposes, a current clamped stimulus is also applied to the cell. The stimuli are applied in a trial-based manner, with a trial-unique, correlated, stochastic process and a subtle trial-identical step change added to the mean of the stochastic process. Subcellular targeting is implemented using a Channelrhodopsin-2 model, as may eventually be used in future laboratory experiments. Because of this, the trial-identical steps are first pre-conditioned to account for the low pass filtering behaviour of the Channelrhodopsin-2 model. The effects of this stimulus are examined with respect to neuronal population response, the onset rapidness of constituent neurons and their passive bandwidth. Section 4.1 begins by examining the neuronal population response to global, somatic and basal dendritic photostimulation, as well as somatic current clamped stimulation. This is followed by Sect. 4.2 which looks into the onset rapidness as measured at the soma and axon of cells respectively. Finally, Sect. 4.3 compares the effective passive transfer functions of each of the stimulus cases.

4.1   Population Response

To examine a neuronal population's response to subcellular targeting, trial-based numerical experiments were performed. A morphologically realistic layer V pyramidal neuron model was exposed to a trial-unique Ornstein-Uhlenbeck (OU) process with a small (6.4 %), trial-identical step change to its mean. The three photostimulus targets chosen were the soma ($\approx$2.66 % of the membrane), basal dendrites ($\approx$40 % of the membrane) and the entire cell, referred to as global stimulation. A somatic current clamp stimulus was also applied for

comparison. The results (Fig. 5.1) present the response of a population comprising 1,000,000 neurons. The peri-stimulus time histograms (PSTH, Fig 5.1a) display a clear response to the step onset in all cases. This is despite the very subtle change in the mean. The response is also somewhat noisy, which is probably due to the size of the population, and the overall spike rate of 1 Hz. Increasing either of these would probably smooth the results to some degree. Examining the rise in the PSTHs, after $t = 0$, the somatic photostimulus displays some speed advantage over the basal dendritic and the global cases, the response to the global stimulus appears slowest. This is similarly reflected in the probabilities of detection (Fig. 5.1b), which show some distinct differences between the 4 cases, especially within the first 100 ms after the onset of the step. Let us first examine the photostimuli. Here, the bootstrapped median detection probability shows a clear contrast between the global, somatic and basal dendritic stimulation cases. The former, global photostimulation, shows the slowest response, followed by basal dendritic photostimulation. Somatic photostimulation rises fastest. If one assumes that subcellular targeting is equivalent to reducing the sinking effect of the target, then Eyal et al. (2014) predict this. There are, however, clear overlaps between the 95 % confidence bands, especially between the global and basal dendritic stimulation cases and basal dendritic and somatic stimulation cases. This is perhaps more a question of sufficient data, as also illustrated by the noisiness of the PSTH plots (Fig. 5.1a). Looking at the probability of detection under a somatic current clamped stimulus (Fig. 5.1b), we can see significant overlap between the current clamped somatic stimulation and basal dendritic stimulation cases, the overlap between the two somatic stimulation cases is less significant, especially within the first 100 ms after the onset of the step. This does not appear to fit with what Eyal et al. (2014) would predict, if it is assumed that subcellular targeting is equivalent to reducing the sinking effect of the target. Looking at the probability of detection results as a whole, only the confidence bands of the somatic photostimulus exceed chance level at any stage in the plot. This speaks to the difficulty of detecting such a subtle step. Indeed, if one estimates the signal to noise ratio (SNR) of the stimulus as the step size (0.064μ) divided by the standard deviation of the OU process $(^{\mu}/_2)$, then the SNR is approximately $20\log_{10}(0.128) \approx -18$ dB, which is exceedingly low.

4.2   Onset Rapidness

For each of the four stimulus cases, the onset rapidness was computed from measurements made at the soma. This was done by first aligning action potentials at -50 mV (Fig. 4.2a)[1] then by computing their corresponding phase plots (Fig. 4.2b & reffig:sim ap soma phase zoom). The onset rapidness was then estimated by first computing the gradient of the phase plane plot (Fig. 4.2d) over a range of values at the beginning of the phase plane plot (Fig. 4.2c). Similar to the approach taken by Eyal et al. (2014), a point was chosen at which

---

[1]Spikes were first up-sampled by a factor of 20, aligned at the point at which they cross above -50 mV, then downsampled by a factor of 20.

(a) The normalised peri-stimulus time histograms (PSTH)

(b) The probability of the neuronal population detecting a step

Figure 4.1: The figures are generated from 1,000,000 simulated trials that apply a trial-unique Ornstein-Uhlenbeck process with a trial-identical subtle (6.4%) step change in its mean (see Sect. 2.5). Both figures plot the median and 95 % bootstrapped confidence bands. 10,000 bootstrap samples were taken, with a sample size of 1,000,000. Thus, the results simulate the response of a population of 1,000,000 neurons. *a)* The PSTH is normalised to the pre-step-onset mean of the PSTH median. Line plots are used instead of bars to more easily compare the four cases. The somatic photostimulus appears to produce the fastest response, followed by basal dendritic then global photostimulus. *b)* This is computed using the detector described by Tchumatchenko et al. (2011) also described in Sect. 2.4.2. Differences in the rise times of the photostimulus cases are evident and fit with the prediction (based on Eyal et al. 2014), i.e. global is slowest followed by basal dendritic then somatic photostimulation. The somatic current clamped stimulus, however, appears closer to basal dendritic, than somatic photostimulation, which is not as predicted. The overall probabilities of detection attained are low, in some cases, barely exceeding chance (0.5).

to measure the onset rapidness ($^{dv}/_{dt} = 25\,\mathrm{mV\,ms}$, black dashed line in Figs 4.2c & 4.2d). This was chosen because it 1) is near to the onset of the action potential, 2) is a relatively high value and 3) appears less beleaguered by noise. From this point, the onset rapidness for each action potential was measured[2], giving a distribution of onset rapidness values for each of the four stimulus cases (Fig. 4.3a). The distributions show significant overlaps, the exception being the strength of the left-skew present in all cases. Global and basal dendritic photostimulation cases show the strongest, left-most peaks. The current injected somatic stimulus and somatic photostimulus cases, although also left-skewed, display significantly smaller peaks in the skew. Examining the peaks of these distributions (Fig. 4.3b), the stimulation of the smallest section of the membrane (somatic stimulation – $2.66\,\%$) appears to result in the fastest onset rapidness, whereas the stimulation of the largest portion of the membrane (global stimulation – $100\,\%$) results in the slowest onset rapidness.

This result is not, however, reflected in the onset rapidness measured at the axon (Fig. 4.4). The onset rapidness distributions (Fig 4.4a) have significant overlaps, especially about the peak, the largest differences in the distributions appear between onset rapidness values of 5 and $7\,\mathrm{ms}^{-1}$, manifesting as pregnant bulges, especially in the somatic and basal dendritic cases. Short peaks at approximately $11.5\,\mathrm{ms}^{-1}$ are also evident. The overall relative jaggedness of this distribution compared to the *soma* case (Fig. 4.3a) indicates a higher level of noise. Examining the peaks of the distributions (Fig. 4.4b), the somatic current clamped stimulus appears to have a marginally faster onset than the other three cases, otherwise clear differences between the four cases are not easily discerned.

4.3   Passive Transfer Function Estimates

For each of four subcellular targets, the effective transfer function estimates of the passive cell were measured. This was done for two cases: 1) excluding and 2) including Channelrhodopsin-2 dynamics. To do this, a morphologically accurate L5 pyramidal cell (Sect. 2.2.1) was rendered passive by setting all active membrane conductance mechanisms to zero, except, of course, for the Channelrhodopsin-2 mechanism, which was required for photostimulation. An Ornstein-Uhlenbeck stimulus was then applied to the cell via either somatic current clamp; global, basal dendritic or somatic photostimulation. The size of the stimulus was chosen, in each case, to produce a spike rate of approximately $1\,\mathrm{Hz}$ in the active cell.

To estimate the transfer functions excluding Channelrhodopsin-2 dynamics, the induced conductance was taken as the input in the photostimulus cases. This avoided the low frequency filtering behaviour of the Channelrhodopsin-2 model which maps light stimulus to channel conductance. In the current clamp case, the injected current was taken as the input. In all cases, the output was the membrane voltage at the soma. Examining the

---

[2]Linear interpolation was used to compute the onset rapidness at $\frac{dv}{dt} = 25\,\mathrm{mv\,ms}^{-1}$.

(a)

(b)

(c)

(d)
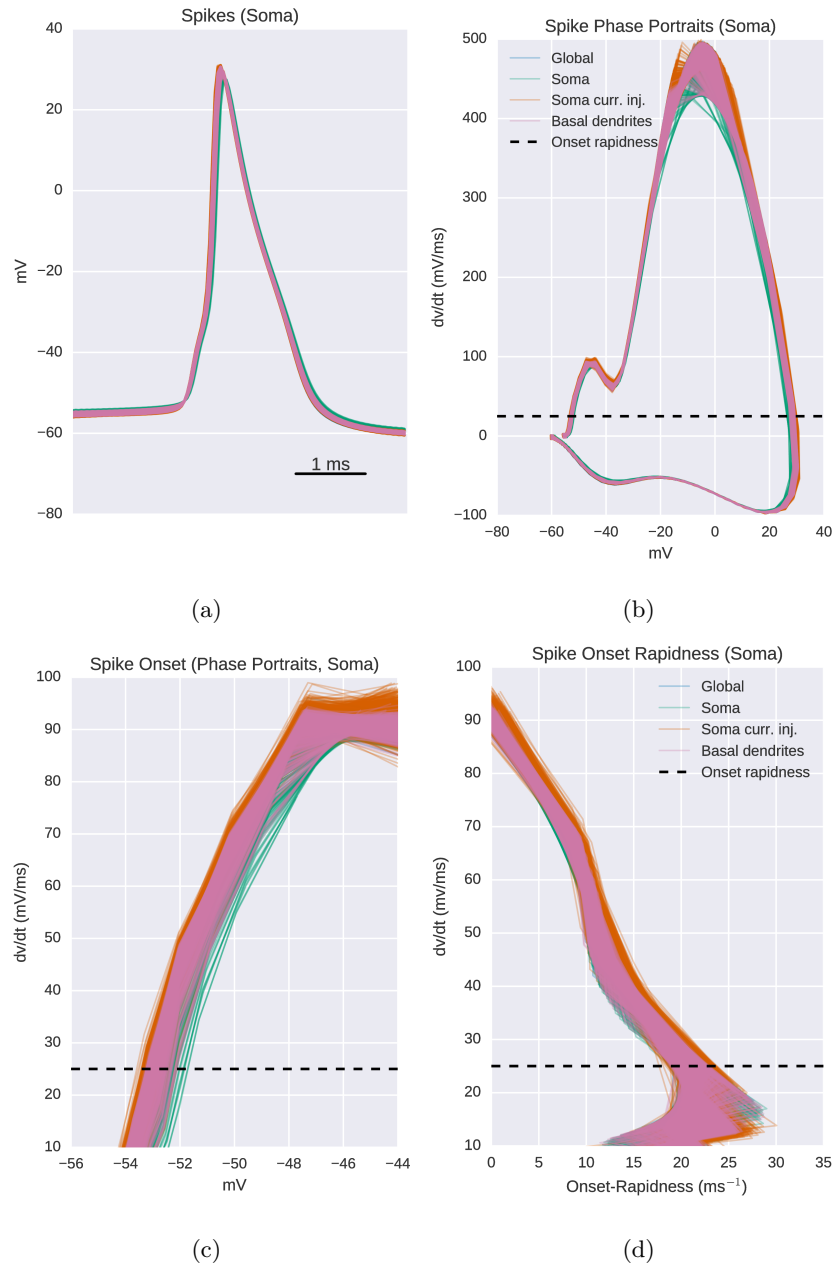
Figure 4.2: The spikes, phase and onset rapidness plots as measured at the *soma*. The figures are generated from 50,000 simulated trials that apply a trial-unique Ornstein-Uhlenbeck process with a trial-identical subtle (6.4%) step change in its mean (see Sect. 2.5). *a)* 1000 randomly selected spikes measured at the *soma* and aligned at -50 mV. *b)* The corresponding phase plane plots. *b, c & d)* The black dashed line shows the point at which the onset rapidness measurement was made. *c)* Zoom to the onset of the phase plane plots in *b. d)* The gradient of *c* plotted against the derivative of the spike (i.e. the vertical axis in *c.*) This shows the range of onset rapidness values from which one can choose. The black dotted line illustrates the onset rapidness value that was chosen. The point chosen was to be 1) early in the onset, 2) relatively high and 3) above the noise.

normalised transfer function estimates (Fig. 4.5a), the similarity in shape between the two somatic stimulation cases is evident, especially at higher frequencies. Looking at the pass-band of the same transfer functions (Fig. 4.5b, Table 4.1), the somatic current clamped case has the lowest bandwidth of all four cases. The global photostimulus case has the lowest

(a)

(b)
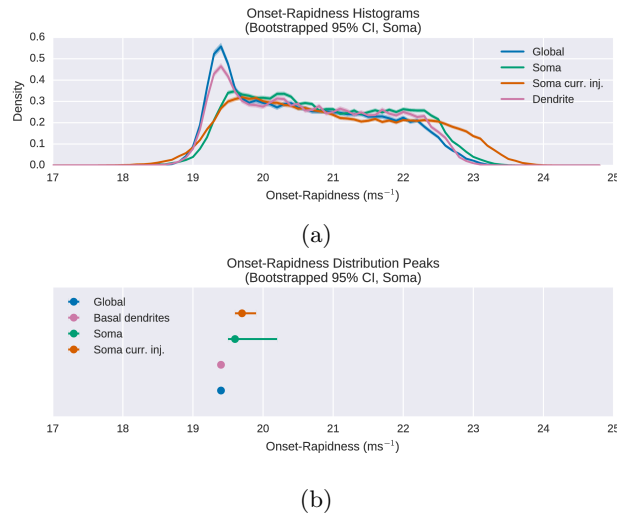
Figure 4.3: *a)* The distribution of onset rapidness values per stimulus case as measured at the *soma. b)* A bootstrapped estimate of the peak of the distributions in *e.* The two stimuli extremes, somatic and global stimulus show correspondingly extreme onset rapidness values.

(a)

(b)
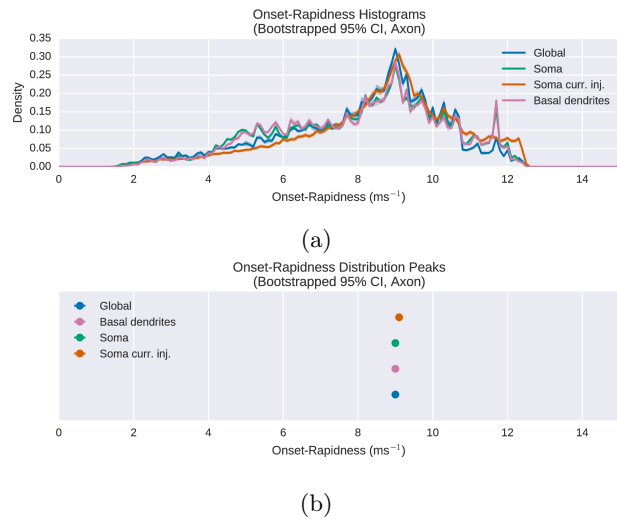
Figure 4.4: *a)* The distribution of onset rapidness values per stimulus case as measured at the *axon. b)* A bootstrapped estimate of the peak of the distributions in *e.* Apart from a slightly larger value for the somatic current clamp case, it is difficult to discern differences between the onset rapidness peaks.

Table 4.1: Median effective passive cell bandwidths

	Bandwidth	Bandwidth (Input: Light)
Soma curr. inj.	3.02	–
Global	3.21	3.13
Basal dendrites	3.62	3.50
Soma	3.64	3.49

bandwidth of the photostimulation cases, the somatic and basal dendritic photostimulus cases have the widest bandwidths with significant overlaps at the cutoff.

To estimate the effect of Channelrhodopsin-2 dynamics on the bandwidth of the passive

(a)

(b)

Figure 4.5: The effective passive transfer functions of a morphologically realistic layer V pyramidal neuron model for 3 different photostimulus targets and soma current clamped stimulus. The stimulus was an Ornstein-Uhlenbeck process scaled to produce a spike rate of approximately 1 Hz. Transfer functions were estimated using the Welch method. Plotted are the bootstrapped medians with 95 % confidence intervals. *a)* Plots the transfer function estimates normalised to the gain at the lowest frequency. The somatic stimuli have similar shapes, with less attenuation at higher frequencies that the global and basal dendritic photostimuli. *b)* A zoom of the transfer functions about the passband. Here the differences in gain are less apparent. *a & b)* The black dashed line is the passband cutoff defined as $|S_{xy}/S_{xx}| = 0.7 \approx -3\,\mathrm{dB}$. The cutoffs under somatic and basal dendritic photo stimulation are highest at $\approx 3.6\,\mathrm{Hz}$, followed by global stimulation ($\approx 3.2\,\mathrm{Hz}$) then somatic current clamp ($\approx 3.0\,\mathrm{Hz}$).

cell, similar estimates were made. Here, however, rather than using the light induced current as the input in the photostimulus cases, the input was taken as the light intensity applied to the membrane. At high frequencies, the photostimulus cases demonstrate marked differences when compared to the "conductance as input" (Fig. 4.5a) transfer functions. This is due to the Channelrhodopsin-2 model which behaves as a low pass filter at small light intensities. Despite this, in the passband (Fig. 4.5b, Table 4.1), the gain curves appear very similar to the "conductance as input" curves. Of the four cases, the somatic current clamped case has the lowest bandwidth, with some confidence band overlap with the global photostimulus case. The global photostimulus case has the lowest cutoff of the photostimulation cases, the somatic and basal dendritic photostimulus cases have the widest bandwidths with significant overlaps at the cutoff.

4.4   Comparison of Numerical Results

Of the three metrics given above, four data points have been collected, three of which use the same stimulus method. Although this is an insufficient amount of data to perform a convincing statistical analysis of their relationship, one can gather some coarse insights in their congruences. To do this, the three metrics are plotted in opposing axes (Fig. 4.7). Here, a surrogate for population response speed is used: the probability of detection at 50 ms. This was chosen, as it is roughly the point at which the fastest response reaches its peak probability of detection (Fig. 5.1b). At a glance, no clear relationship between the three measures is discernible. However, if one looks only at the photostimulus results, one could argue that there is some positive congruence in all three cases. Although, there is little to no separation between the onset rapidness values for the global and basal dendritic stimulation cases (Figs. 4.7a & 4.7b). Similarly, the bandwidth values show little separation and significant overlaps for the basal dendritic and somatic stimulation cases (Figs. 4.7b & 4.7c). When one considers only the two extreme stimulus cases, global and somatic photostimulation, which represent ≈2.66 % and 100 % of the membrane surface area respectively (the basal dendrites represent ≈40 %), there is a strong suggestion of increasing congruence between all three metrics.

The results given here (Fig. 4.7) plot the bandwidth *excluding* Channelrhodopsin-2 (ChR-2) dynamics. If one plots the bandwidth including ChR-2 dynamics, the results are very similar. For this reason, the results are not presented explicitly here, but can be found in Appendix A, Sect A.2.

(a)

(b)

Figure 4.6: The effective passive transfer functions of a morphologically realistic layer V pyramidal neuron model for 3 different photostimulus targets and soma current clamped stimulus. The stimulus was an Ornstein-Uhlenbeck process scaled to produce a spike rate of approximately 1 Hz. Transfer functions were estimated using the Welch method. Plotted are the bootstrapped medians with 95 % confidence intervals. These plots include Channelrhodopsin-2 dynamics. *a)* Plots the transfer function estimates normalised to the gain at the lowest frequency. The photostimulus cases display significantly higher attenuation at higher frequencies than the somatic current clamp case. The global and basal dendritic photostimuli produce the highest attenuation at high frequencies, followed by the somatic photostimulus. *b)* A zoom of the transfer functions about the passband. Here the differences in gain are less apparent. *a & b)* The black dashed line is the passband cutoff defined as $|S_{xy}/S_{xx}| = 0.7 \approx -3$ dB. The cutoffs under somatic and basal dendritic photo stimulation are highest at $\approx$3.5 Hz, followed by global stimulation ($\approx$3.1 Hz) then somatic current clamp ($\approx$3.0 Hz).

(a)

(b)

(c)
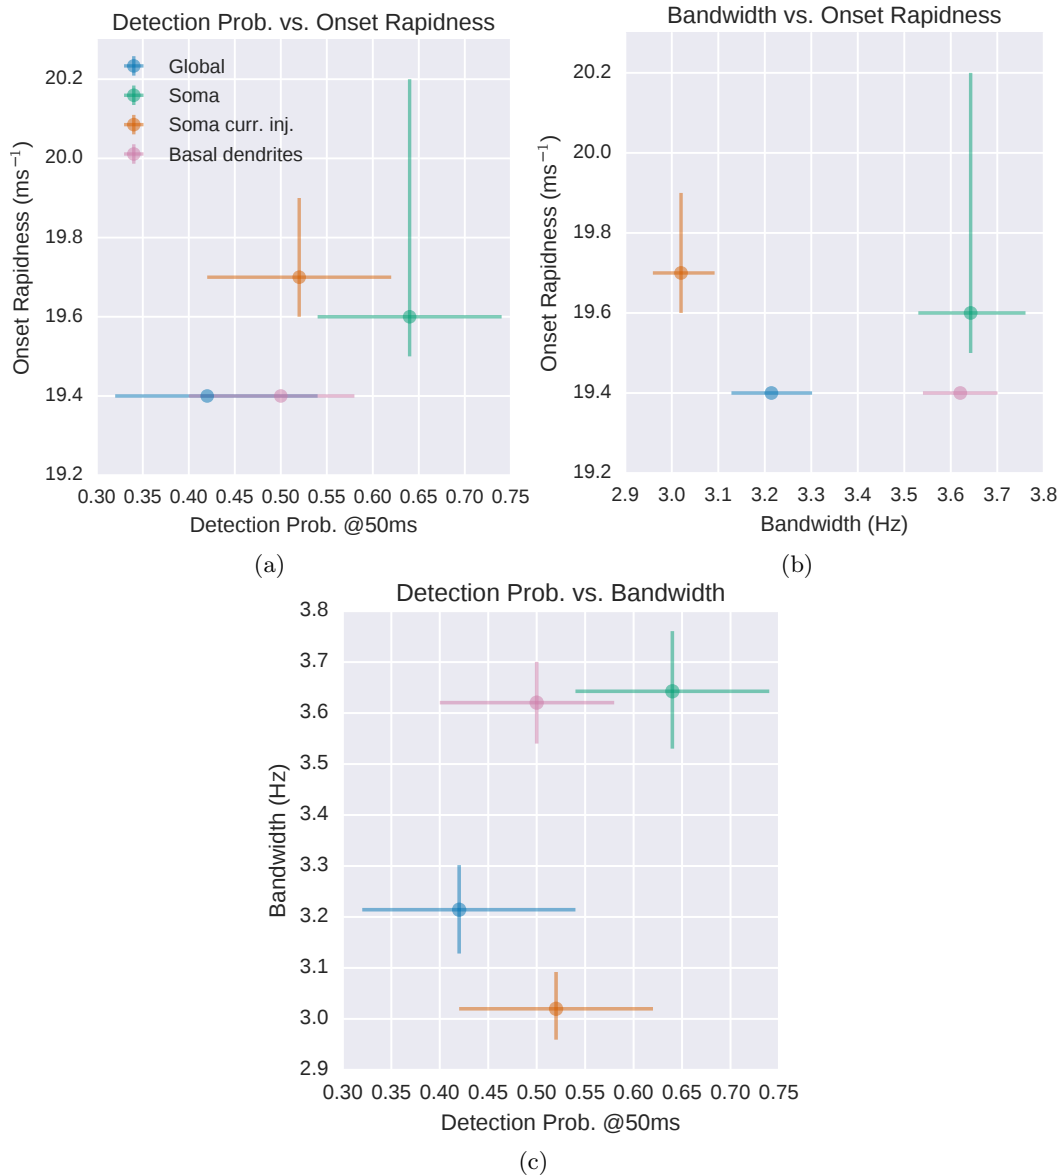
Figure 4.7: A comparison of three metrics: The onset rapidness as measured at the soma (Fig. 4.3), the effective passive bandwidth (Fig. 4.5b) and the probabilities of detection at 50 ms (Fig. 5.1b). The values plotted are the bootstrapped medians with 95 % confidence intervals. *a)* The probability of detection at 50 ms against onset rapidness. *b)* The effective passive bandwidth versus onset rapidness. *c)* The the probability of detection at 50 ms against the effective passive bandwidth. Focussing in on the two extreme cases, somatic and global photostimulation, congruent increases in all three metrics can be seen.

Chapter 5

Discussion & Conclusions

This thesis presents novel results in the areas of closed loop electrophysiology and neuronal population encoding. The work begins with the implementation of a novel closed loop electrophysiological system, capable of controlling a neuron or neurons exposed to a correlated, stochastic photostimulus. This is then used to examine the response of an *in vitro* neuronal population to subtle changes in the mean of a correlated, stochastic photostimulus. Following this, this thesis looks into the application of stimuli to specific elements of a cell's structure, called "subcellular targeting." The effects of subcellular targeting are explored numerically. Namely, its effect on the response speed of a neuronal population, the onset rapidness of constituent neurons and their passive bandwidth. This chapter discusses the results, beginning with the novel closed loop electrophysiological approach in Sect. 5.1. This is followed by Sect. 5.2, which discusses the effect of a truly subtle stimulus on an *in vitro* neuronal population. Section 5.3 discusses subcellular targeting. Finally, future work is presented in Sect. 5.4

5.1 Feedback Control Using a Stochastic Process

Closed loop electrophysiology or closed loop neuroscience describes in essence, the insertion of a silicon-based feedback path between the output of a neuron or neurons and their input (Arsiero et al. 2007; El Hady 2016; Grosenick et al. 2015; Potter et al. 2014). This idea is far from new. From voltage clamp (Kenneth Cole; Hodgkin and Huxley 1952a) and path clamp (Neher and Sakmann 1976) to dynamic clamp (Sharp et al. 1993) and response clamp (Bölinger and Gollisch 2012; Wallach 2013; Wallach et al. 2011), this approach has a long history and has become an essential experimental tool in the electrophysiologist's toolbox. More recently, Newman et al. (2015) demonstrated that it possible to control the rate of spiking with a deterministic stimulus using a PI (proportional-integral) controller. This thesis extends on this work, presenting the first known attempt to control the spike rate of a neuron or neurons with a correlated, stochastic process driven by an adaptive PID controller.

Chapter 3 gives the results of a novel closed loop electrophysiological system used to control the rate of spiking of a cortical culture under a correlated, stochastic photostimulus. The system fulfils three key requirements, which allow the experimentalist to regulate spike rates in trial-based, "whole field" photostimulus experiments. The key requirements being *1)* that the stimulus does not induce 'clock-like' behaviour in the cells, *2)* that the controller has no significant influence on the outcome of a trial and last but not least *3)* that the control results in the desired rate of activity from the culture.

A closed loop system that regulates spike rates in trial-based experiments should not produce stimuli that result in unnatural, clock-like spiking. One indicator that can be used to quantify this type of behaviour is local variation (LV). An LV of 0 indicates perfect, clock-like spiking, whereas an LV of 1 indicates a Poisson process-like behaviour. The LV values of the spike trains generated in these closed loop experiments (Fig. 3.11c and Appendix A, Figs. A.1c & A.2c) are distributed about an LV of approximately 0.5, indicating significantly non clock-like behaviour. This is the result of applying an exponential rate window with a long tail ($\tau = 10\,$s) to smooth spike rate activity, which essentially blurs the effect of each spike over a longer time frame, filling the empty space between spikes, also known as the inter-spike interval. This filling of the inter-spike interval keeps the control error lower than it would be if there were e.g. a smaller window. A larger control error would drive the controller to compensate in a more drastic manner, producing more extreme swings in the control signal and a potentially more clock-like behaviour. Looking at this from another perspective, a more clock-like behaviour is also likely to surface as the target rate approaches or becomes smaller than $1/\tau$, the inverse of the time constant of the exponential window.

For a closed loop system regulating spike rates in trial based experiments, it is important that the controller is not significantly influenced by the average outcome of the trail. To examine this, one can look at the trial average *output* of the controller (Fig. 3.12c). This describes the influence of trial-based changes to the spike rate on the output generated by the PID controller. The results show a trial-mean controller response that is clearly periodic, with the same period as the trial-identical stimulus, indicating that trial-based changes in the spike rate do have an influence on the PID controller outputs. However, this effect is small, in the order of $10^{-4}\,$V, and is not significant enough to be of concern, especially when one considers that the applied stimuli typically have amplitudes in the hundreds of millivolts (see e.g. Fig. 3.11a). For those planning to implement a similar system, it is important to mention that the size of this influence is likely a result of the time constant of the exponential rate window and the period of the trial. A longer trial period or a trial stimulus with larger periodic changes would likely require a larger window.

Looking at the controller's ability to follow the target (Figs. 3.11b and Appendix A, Figs A.1b & A.2b), the control results can not be described as smooth, especially when compared to the electronic culture ersatz (Fig. 3.9b), there are also clear overshoots in response to steps. For the purposes of these experiments, however, the level of control

was sufficient and indeed prevented the kinds of large swings in rates observed in other experiments in our laboratories (Fig. 3.1), which was the key reason to employ a controller in the first instance. Of course, the requirements placed upon such a controller would not be the same for each experiment, and the smoothness of the rate observed depends very much on the time horizon over which one looks. For spike trains with a higher local variation, the fewer spikes one observes, the less likely the overall rate matches the desired rate.

Examining the coefficients converged upon by the adaptation mechanism (Fig. 3.11e and Appendix A, Figs. A.1e & A.2e), a clear pattern emerges across the experiments presented: that that the derivative path of the PID controller is mostly likely redundant, meaning a PI controller would suffice, as used in previous work (e.g. Newman et al. 2015). The decision to use a PID controller in this work was made to gain some insight into the adaptation of a PID controller in a closed loop electrophysiological system, i.e. what a gradient descent based algorithm (Lin et al. 2000) would make of a 'laggy' system (Fig. 3.4c, Sect. 3.2.4) controlling a biological 'plant,' which was not discussed in previous work.

The gradual reduction in size of the derivative (D) portion of the PID controller during adaptation (Fig. 3.11e and Appendix A, Figs. A.1e & A.2e) indicates that the trajectory of the spike rate, i.e. where it is headed, is insignificant relative to the current spike rate and its long-term history. Similarly, in all three cases presented, there are not insignificant differences between the P and I portions of the controller, more specifically, that the proportional control coefficient typically reaches a value more than twice that of the integral coefficient, which is an indication of the relative importance of the current spike rate and its history respectively. Of course, the coefficient values reached are not just a result of adapting to the controlled system, but also of the learning rate chosen. The jumpiness of the controlled rate is likely a symptom of putting too much emphasis on the current spike rate (the proportional component – P), which was observed to be the result of a smaller adaptation rate. The result of a strong P component is evident when looking at the controller outputs broken down into P, I and D components (Fig. 3.11d and Appendix A, Figs. A.1d & A.2d), where the large P value produces a large output when the error is high, i.e. at step onsets, thus producing overshoots. A larger learning rate was observed to reduce the relative difference between the P and I controller coefficients, leading to reduced overshoot at the cost of longer rise times.

In summary, a stochastic stimulus can be adjusted by a feedback controller to control the rate of activity of a cell or cells. This can be done such that spiking is not unnaturally regular, and that the response of the cell or cells to a trial based stimulus minimally influences the controller's behaviour. The control is not as smooth as it might be. This may be due to too large a proportional (P) component, which seems to result from a smaller learning rate. However, the control result was sufficient for the population response experiments discussed in the following section. Indeed, for experimentalists who want to adopt this technique, there is plenty of "wiggle" room offered through e.g. the learning parameter of the adaptive

PID and the rate window.

The introduction of this new approach to spike rate control allows the electrophysiologist to take new strides when it comes to trial-based experiments. Using this closed-loop approach, the electrophysiologist now has a rate "knob" that he or she can turn. For those experiments which require e.g. a fixed rate of spiking, a changing rate of spiking or a step change in spiking, this can be automated using this closed-loop approach. Moreover, the experiences in our laboratories have shown that trail-based experiments can be subject to wild trial-to-trial variations in the rates of spiking (Fig. 3.1). With this closed loop system, the experimentalist can ensure that more trials produce a usable response, resulting in more data for the same time in the lab.

5.2   *In Vitro* Population Response: Robustness & Speed

Previous work has shown that a population of neurons, each subject to tens of thousands of inputs, observes, in sum, a stochastic process. When constituent cells are exposed to this kind of correlated, stochastic noise, the population displays a bandwidth proportional to the correlation time of the signal (Brunel et al. 2001). This bandwidth has been shown in both theoretical and experimental work to reach ranges in the hundreds of Hertz, and could be advantageous when it comes to information encoding (Boucsein et al. 2009; Brunel et al. 2001; Fourcaud-Trocmé et al. 2003; Higgs and Spain 2009; Ilin et al. 2013; Köndgen et al. 2008; Naundorf et al. 2005; Ostojic et al. 2015; Tchumatchenko et al. 2011; Wei and Wolf 2011). Indeed, researchers have shown that a neuronal population can respond within 1–2 ms to a 20 pA step in the mean of a stochastic process (Tchumatchenko et al. 2011). This step size is subtle, corresponding to the single strong synaptic input. Relative to the size of the background noise, however, it represents a jump of approximately 40 %. What, then, would be the response of a population to a step that is extremely subtle, relative to the size of the background noise? Does the size of the step matter?

To answer this question, I performed *in vitro* experiments using Channelrhodopsin-2 infected cortical cultures. These experiments exposed a cortical culture to a feedback controlled, trial-unique Ornstein-Uhlenbeck photostimulus with a trial-identical 6.4 % preconditioned step change in the mean. This step size was reached by performing multiple experiments, increasing the step size until a response was observed. The results of this protocol show a population step response (Fig. 3.12a) which reaches a peak within approximately 50 ms and a detection probability (Fig. 3.12b) which peaks at approximately 0.75 within 150 ms. This shows, that a population's response to a step can also be slow, and that the size of the step does matter.

The results show that a subtle jump in the mean of a correlated, stochastic photostimulus can produce a neuronal population response; but, not one that is fast enough to be of use in e.g. sensory systems. Experimental evidence shows that such systems process inputs within

anything from 5 – 200 ms (Stanford et al. 2010; Swadlow and Hicks 1996; Thorpe et al. 1996). A neuronal sub-population whose response peaks at 50 ms (Fig. 3.12a) and whose probability of detection peaks after 150 ms (Fig. 3.12b) is not likely to be of great use in this context. If one were to instead of applying a step, apply a small short pulse of a few milliseconds, roughly the duration of an action potential, the short pulse is not likely to be detected at all. Therefore, these results tell us that, given enough neurons and the right detector, a subtle step change in the mean of a correlated, stochastic signal can be detected; however, the result is too slow to be of use in sensory systems.

Despite the slowness of response, what the results also demonstrate is how robust a detector the population can be. A 6.4 % jump in the mean results in an observable jump in the population response (Fig. 3.12a), and reaches a peak in the median probability of detection of approximately 0.75 (Fig. 3.12b). If we roughly estimate the signal to noise ratio of the stimulus as the step size (0.064μ) divided by the standard deviation of the Ornstein-Uhlenbeck process $(^{\mu}/_2)$, we have $\frac{0.064\mu}{^{\mu}/_2} = 0.128$, which in terms of decibel gain corresponds to $20\log_{10}(0.128) \approx -18\,$dB. Under such signal to noise conditions, a peak probability of detection of 0.75 can be considered a robust response. Of course, the choice of detector plays a significant role in this result, this is important to note, considering that the detector used here was not designed on the principle of biological plausibility, but rather for purposes of comparison with other population detection experiments (e.g. Tchumatchenko et al. 2011).

In summary, a neuronal population can indeed respond to very subtle changes in the mean of a correlated, stochastic stimulus. The response is slow, likely too slow to be of use in sensory systems, showing us that "size does matter". However, this combination of the neuronal population and detector demonstrate that a population can detect subtle changes in a robust manner. The population and detector demonstrated a $\approx 75\,\%$ probability of detection in extremely poor ($\approx$-18 dB) signal to noise conditions.

5.3   Subcellular Targeting

Previous work has shown that neuronal populations exposed to a correlated, stochastic stimulus, are capable of responding to mean-modulated signals with frequencies in the hundreds of Hertz. This high frequency encoding capability is influenced by the onset speed of constituent cell's action potentials (Boucsein et al. 2009; Brunel et al. 2001; Fourcaud-Trocmé et al. 2003; Higgs and Spain 2009; Ilin et al. 2013; Köndgen et al. 2008; Naundorf et al. 2005; Ostojic et al. 2015; Tchumatchenko et al. 2011; Wei and Wolf 2011). The onset speed of an action potential is thought to be the result of fast sodium channels, the channels responsible for the influx of cation current into the cell (hyperpolarisation) at the initiation of an action potential (Fourcaud-Trocmé et al. 2003). An alternative hypothesis, is that morphology plays a role. Eyal et al. (2014) propose that the size of the dendritic tree has a significant influence on an action potential's onset. Under somatic stimulation, the dendritic

tree represents a type of current sink into which onset sodium currents flow. The bigger this sink, the faster the sodium currents flow into it and, therefore, the faster the onset of the action potential. If one were to eliminate the dendritic tree, then the onset should be slower. Imagine then that if, instead of removing the tree, one were to stimulate the tree. Currents would flow through the dendritic membrane first, then into the remainder of the cell, which would, for all intents and purposes, reduce or eliminate the sinking behaviour of the dendritic tree. This is the idea that this work aims to test, under "subcellular targeting," i.e. the application of a stimulus to specific elements of a neuron's structure.

This thesis examines the relationship between subcellular targeting, the response speed of a neuronal population, the onset rapidness of its constituent neurons and their passive bandwidth. Chapter 4 presents the results of numerical experiments performed to examine this relationship. This is done by applying a photostimulus targeting light gated cation channels (Channelrhodopsin-2) at either the soma, basal dendrites or the entire cell, referred to as global photostimulation. Current clamped stimulation is also applied for comparative purposes. The cell used is a morphologically realistic layer V pyramidal neuron. In these numerical experiments, a trial-unique, correlated, stochastic photostimulus is applied to the subcellular target. A subtle 6.4% trial-identical step modulates the mean of the photostimulus. The step is preconditioned to account for the low pass filtering properties of the Channelrhodopsin-2 model used.

5.3.1   Population Response Speed

Let us begin by examining the speed of the neuronal population in response to subcellular targeting of photostimuli. Differences in the population response are discernable from the peri-stimulus time histograms (Fig. 5.1a) that compare the three photostimulation cases. The probabilities of detection (Fig. 5.1b) also demonstrate notable differences: The global photostimulus produces the slowest response, followed by the basal dendritic then the somatic photostimulus. If we assume that the subcellular targeting of stimuli is equivalent to reducing the electrical load represented by the target, then these results are predicted by Eyal et al. (2014). In the case of the slowest response, produced under a global stimulus, the entire membrane is stimulated, which should result in the smallest sink of all three cases, the next slowest response results from stimulating the basal dendritic tree ($\approx 40\%$ of the membrane), which should leave a larger sink (soma, axon and apical dendrite) than a global stimulus. Of the three photostimulus cases, the largest sink should be produced when stimulating only the soma ($\approx 2.66\%$ of the membrane), and indeed, this corresponds to the fastest population response.

Adding the current clamped somatic stimulus into the mix, however, we observe that its probability of detection (Fig. 5.1b) appears more to match that of basal dendritic photostimulation than that of somatic photostimulation, which is not what Eyal et al. (2014) predict. Fourcaud-Trocmé et al. (2003) might say that the somatic photostimulation must

therefore generate larger sodium currents and a faster action potential onset than the current clamped case. Indeed, the application of the Channelrhodopsin-2 model mimics a cation channel current additional to those sodium channels already present in the model and used in the current clamp case. This is, however, not entirely evident from the onset rapidness measures made at the soma (Fig. 4.3b), which show bootstrapped onset rapidness distribution peaks. The bootstrapped median onset rapidness for somatic current clamped stimulation is clearly larger than that of somatic photostimulation. However, the 95 % confidence intervals for somatic photostimulation are both wider and higher than those for somatic current clamped stimulation. Thus, the presence of extra cation channels introduced by the ChR-2 model could lend the photostimuli a speed advantage over the current clamped stimulus which ultimately hinders a direct comparison of current clamped stimulus and photostimuli.

5.3.2 The Role of Onset Rapidness

A wealth of literature predicts that a faster neuronal population response is caused by a faster onset of the constituents' action potentials (Boucsein et al. 2009; Brunel et al. 2001; Fourcaud-Trocmé et al. 2003; Higgs and Spain 2009; Ilin et al. 2013; Köndgen et al. 2008; Naundorf et al. 2005; Ostojic et al. 2015; Tchumatchenko et al. 2011; Wei and Wolf 2011). However, the role of onset rapidness in response to subcellularly targeted stimuli is yet to be examined. To this end, during the numerical experiments testing the response speed of neuronal populations to subcellular targeting, recordings of the membrane potential at both the soma and axon were made and examined about the recorded spike times.

Looking at the onset rapidness distribution peak values at the soma (Fig. 4.2d), there does appear to be some congruence between the onset rapidness and the probability of detection (Fig. 5.1b). Namely, that the global stimulus produces the slowest rise in its probability of detection and a correspondingly low onset rapidness, the somatic photostimuli produce detection probabilities that rise faster and have a correspondingly higher onset rapidness (Fig. 4.7a). For basal dendritic stimulation, one would expect an onset rapidness somewhere between global photostimulation and current clamped somatic stimulation. However, the value appears to be very close to that of global stimulation. Similarly, one would expect the somatic current clamped stimulus to result in an onset rapidness somewhere between that produced by somatic photostimulation and basal dendritic stimulation, yet this is not the case. When excluding the somatic current clamp case[1], one could argue that there is a positive relationship between onset rapidness and probability of detection (Fig. 4.7a). However, it is only when focussing on the two extreme cases, somatic and global photostimulation, that one sees any congruent increases in onset rapidness and population response speed. Any such congruences in the onset rapidness measured at the axon are not evident.

---

[1]See the discussion in Sect. 5.3.1 on hindrances to the direct comparison of photo- and current clamped stimuli.

5.3.3   The Role of Passive Electrical Properties

Eyal et al. (2014) state that a cell's morphology plays a deciding role in the onset speed of its action potential. They claim that the dendritic tree acts as a current sink, and that a larger current sink forces a faster draining of sodium current into the dendritic tree, thus producing spikes with faster onsets, and in turn faster population responses. If one were to assume that subcellular targeting is equivalent to reducing the load represented by the target, then, from a linear systems perspective, one would expect to see differences in the effective input / output transfer functions of the passive cell, especially at higher frequencies. More specifically, that a stimulation approach which results in a faster population response, should demonstrate a faster constituent action potential onset and therefore a stronger passive frequency response at higher frequencies.

To investigate this numerically, a morphologically realistic L5 pyramidal neuron model was rendered passive by setting all channel conductances, except those for Channelrhodopsin-2 (ChR-2), to zero. The ChR-2 channels were of course required to apply the stimulation. An Ornstein-Uhlenbeck process was then applied as stimulus, then effective passive gain plots were computed (Fig. 4.5b). In all cases, the output was taken as the voltage at the somatic membrane. The input, however, differed depending on the case in question. In the current clamp case, the clamped current through the membrane at the soma was taken as the input. In the photostimulus cases, the ChR-2 induced conductance was taken as the input, this was recorded at the soma for the somatic and global photostimulation cases, for the basal dendritic stimulation case, this was measured at the base of the dendritic tree.

At first glance, the results show no obvious relationship between the bandwidth of the effective passive transfer function of the cell and the rate at which the probability of detection rises (Fig. 4.7c). However, when one ignores the somatic current clamp case[2], there does appear to be some congruence between the two measures. This is further emphasised when one focusses on the two extreme photostimulus cases: somatic and the global photostimulus.

5.3.4   Channelrhodopsin-2 Conductance Parameters

The numerical implementation of the subcellular targeting experiments required that I choose conductance parameters for Channelrhodopsin-2 (ChR-2). Due to the behaviour of the model under the three different photostimulus conditions, I chose to fix the light intensity in all three cases, and change instead the effective channel density of the ChR-2 model. This resulted in a range of conductance values spanning an order of magnitude $(1.4 \times 10^{-3} - 3.5 \times 10^{-2}\,\mathrm{S\,cm^{-2}})$, which speaks to the vast differences in membrane surface area exposed to somatic ($\approx 2.66\,\%$), basal dendritic ($\approx 40\,\%$) and global ($100\,\%$) photostimuli, and the different effective channel densities required to achieve similar levels of overall conductance.

---

[2]See the discussion in Sect. 5.3.1 on hindrances to the direct comparison of photo- and current clamped stimuli.

The choice of conductance parameters (Table 2.2) show an increase that is clearly congruent with the population response speed, which is best seen in the probability of detection curves (Fig. 5.1b). This is perhaps one failing of the approach taken, which could lead one to believe that the increasing response speeds are merely an artefact of the parameter choices made. However, the overall conductance values (Table 2.2) and the two sets of gain curves (Figs. 4.5b & 4.6b) speak against this. The two sets of gain curves measure the effective transfer function of the passive cell in the former case and the passive cell plus Channelrhodopsin-2 dynamics in the latter case. Although significant differences between the two sets of curves are apparent at high frequencies, within the pass band the differences appear insignificantly small, i.e. the effect of the parameter choices is not apparent in the pass band of any of the photostimulus cases. Furthermore, the overall conductance values (Table 2.2) are largest for the slowest (global) photostimulation case, followed by the basal dendritic case, then the somatic photostimulation case. If the choice of conductance parameters were the cause of the differences in population response speed, one would expect the overall ChR-2 conductance to be positively correlated with the population response speed; here, any correlation appears to be negative.

The Channelrhodopsin-2 (ChR-2) conductances span an order of magnitude, 1.4×10^{-3} – $3.5 \times 10^{-2} \, \mathrm{S \, cm^{-2}}$. This gives us some idea of the potential difficulties of executing such an experiment in the laboratory. More specifically, that it may be difficult to target a very small portion of the membrane and a very large portion of the membrane and expect to get comparable responses, especially if the two measurements are made with the same cell or cells with similar ChR-2 channel densities. Moreover, that targeting the soma required a ten fold increase in the effective channel density of ChR-2 also points to potential difficulties in producing spikes at all through stimulation of such a small portion of a living neuron's membrane. Finally, the assumption made in the numerical experiments is that channel distribution is even across the entire membrane, which is not in agreement with simple observations of ChR-2 infected and tagged cells made under the microscope. Soma typically appear brighter and have therefore a denser distribution of ChR-2 channels than e.g. the dendritic tree. This indicates that stimulating at the soma and stimulating the same effective membrane surface area in the dendritic tree could produce differences in responses due to the differing ChR-2 channel densities. This would need to be accounted for before one could begin to understand any morphological effects.

5.3.5   Global vs. "Whole Field" Stimulation

One of the key goals of this work is to understand more about the effect of subcellularly targeted stimuli on the response speed of a neuronal population. Ideally, the numerical experiments presented here would therefore be reproduced in the laboratory. The whole field stimulation experiments represent a step in this direction. The population response protocols implemented numerically, were also applied in the laboratory. In this case, global

stimulation or rather "whole field" stimulation was applied to cultures of cortical cells from rat embryos. The cultures had been infected with a Channelrhodopsin-2 (ChR-2) carrying viral vector. As in the simulated global stimulation case, a trial-unique Ornstein-Uhlenbeck process was applied to the cells with a trial-identical preconditioned step.

The results (Fig. 5.1) reveal that the population response of the culture is both larger and slower than the most similar simulated case of global stimulation. There are several potential reasons for this, including the choice of ChR-2 conductance parameters, the presence of shade in the whole field stimulation experiments and the differing morphologies of the cells. The choice of the conductance parameter for the ChR-2 model which provides light gated cation channels likely had an effect on the speed of action potential onset, thus influencing the overall speed of the population's response. Also, cells under whole field stimulation were subject to shade from neighbouring cells, from the electrodes upon or near which they sat and due to the positioning of the photostimulus which came from only one direction. This is in contrast to the simulated global stimulation, where the neuron model was provided with the same intensity of light at every point on the membrane. Finally, the cell model used in simulations, although morphologically realistic, would have had a different morphology to those cells used in the laboratory. Still, that the responses have significant confidence band overlaps within the first 50 ms is encouraging and indicates that the model and parameter choices are within reasonable bounds.

5.3.6   Conclusion

The evidence presented in this thesis shows that the choice of target for subcellular stimulation has an effect on the response speed of a neuronal population (Fig. 5.1b). The somatic current clamp case, however, does not fit the pattern, this could be due to the speed advantages that the photostimuli are afforded through the effective provision of an additional cation current at spike onset. Looking then at the photostimulus cases, there is a clear relationship between the stimulus target, the membrane area (Table 2.2) and the population response speed. Namely that the larger the stimulus area, the slower the population response. According to Eyal et al. (2014), the current sinking properties of a constituent neuron's morphology influence the response speed of the population. If subcellular targeting is equivalent to reducing the electrical load represented by the target, then the observed effect could be due to the change in the effective current sink represented by non-stimulated portions of the cell.

Eyal et al. (2014), Fourcaud-Trocmé et al. (2003), Ilin et al. (2013), Naundorf et al. (2005), and Wei and Wolf (2011) show that there is a relationship between population response speed, onset rapidness and the passive bandwidth of a cell. The evidence presented in this thesis indicates that an additional variable, the area of the membrane stimulated, could also influence the population response speed, onset rapidness and passive bandwidth (Fig. 4.7). Although there is insufficient data to form a strong statistical argument, the re-

(a) The normalised peri-stimulus time histograms (PSTH)

(b) The probability of the neuronal population detecting a step

Figure 5.1: Comparison of neuronal population responses computed numerically and measured from *in vitro* data. The numerical and *in vitro* results are computed from 50,000 and 47,585 trials respectively. A trial-unique Ornstein-Uhlenbeck process with a trial-identical subtle (6.4%) step change in its mean is applied as stimulus (see Sect. 2.5). Both figures plot the median and 95 % bootstrapped confidence bands. 10,000 bootstrap samples were taken, with sample sizes of 50,000 (numerical) and 47,585 (*in vitro*). Thus, the results simulate the response of a population of 50,000 and 47,585 neurons respectively. The stimulus applied in the laboratory was a "whole field" stimulus, most closely akin to the simulated global stimulus, with the exception of some shading in the whole field case. *a)* The PSTH is normalised to the pre-step-onset mean of the PSTH median. Line plots are used instead of bars to more easily compare the five cases. There is a clear rise in the PSTH after the onset of the step. Although, it is difficult to discern any differences in the rise time or peak of the numerical data, the *in vitro* response is clearly stronger, reaching a higher peak. *b)* This is computed using the detector described by Tchumatchenko et al. (2011) also described in Sect. 2.4.2. Again, the response of the *in vitro* population is clearly stronger than the numerical response. It is, however, markedly slower in the first 50 ms.

sults of the two extreme photostimulus cases, global and somatic photostimulation, demonstrate congruent increases and encourage further investigation.

This thesis shows that the choice of subcellular target clearly influences the population response speed and may in turn influence onset rapidness and passive bandwidth. As neuronal inputs are predominantly synaptic, i.e. via dendrites, the effects of subcellular targeting are not likely to have a significant influence on processing within the brain, with perhaps the exception of gap-junctions between cells – channels that traverse the membranes of two neurons, electrically connecting them. These could potentially provide an input at almost any part of the cell membrane. Thus, subcellular targeting may play a role in the point at which two cells electrically couple via gap-junctions.

From an experimental perspective, subcellular targeting has a clear impact. The evidence presented in this thesis shows that the choice of target influences the result that an experimentalist observes. Considering that ubiquitous techniques such as patch clamp typically target the soma and the relatively new but popular technique, optogenetics, often targets the entire cell, it is likely that the experimentalist would see different response speeds for the equivalent protocol applied under these two forms of stimulus. The results speak also to the neuron's response to a somatic or global stimulus performed in experiments compared to the response to synaptic inputs seen in the brain. Of the four stimulus cases presented in this thesis (Fig. 5.1), dendritic stimulation is probably closest to synaptic stimulation, where the dendrites act as inputs. These results indicate that the speed of a neuronal population response to synaptic inputs in the brain likely lies between that of somatic and global stimulation seen in the laboratory.

5.4 Future Work

This thesis presents results that demonstrate the behaviour of a closed loop feedback system driving a stochastic stimulus to control the spike rate of a cell or cells. The system was shown to be capable of controlling the spike rates of neurons by simultaneously scaling the mean and standard deviation of a correlated, stochastic photostimulus. The closed loop system was then applied to trial based experiments where an *in vitro* neuronal population was exposed to subtle step changes in the mean of a correlated, stochastic photostimulus. These experiments could be extended to look at other forms of mean modulation, or perhaps other forms of modulation. For example, one could examine the response of a neuronal population to a range of different step sizes, which would give some insight into how population response speeds change for different input magnitudes. Similarly, instead of steps, different magnitudes of short pulses in the order of e.g. a few milliseconds in length could be applied, to see how the population responds to stimuli similar in length to an action potential. More complicated forms of mean modulation could also be employed providing, perhaps more subtle insights into the responses of neuronal populations e.g. ramp and sinusoidal

modulations, triangle pulses and raised cosine pulses. Finally, the closed loop system was used to control the rate of spiking of a neuron or neurons by scaling the mean and standard deviation of an Ornstein-Uhlenbeck (OU) process. The controller could be modified such that, instead of scaling the mean and standard deviation, it could scale just the mean, or just the standard deviation or perhaps even the correlation time of the OU process and / or any combination of the three. It would be interesting, in this case, to see if the spike rate can be controlled using these kinds of changes to the stimulating process.

The closed loop system described in this thesis controls the spike rate using a single feedback loop to scale the mean and standard deviation of a correlated stochastic process. This control is based on measurements taken at a single probe point within the culture. Given that these cultures are grown on microelectrode arrays, one could ideally perform control experiments at multiple points in a single culture simultaneously, thus significantly increasing the throughput of and data collected from these and similar experiments. Such a project would require an array of independently controllable light sources, each targeting one of the electrodes on the array, and rather than a single controller, one for each of the array electrodes and light sources. Assuming this would also be controlled from a computer or other digital device, a digital to analogue converter (DAC) would also be required for each light source, although a multiplexing DAC might save on hardware costs. Alternatively, the control loops could feasibly be implemented in analogue circuitry.

In addition to this, this thesis presents the first steps taken in understanding the role of subcellular targeting in the speed of a population's response. The results show some congruences between the surface area of the target, the speed of the population response, the onset rapidness of constituent neuron's action potentials and the passive bandwidth of the cell. However, a strong statistical argument is difficult to make. Despite this, the extreme photostimulus cases, somatic and global photostimulus and their relationships do encourage further examination. Specifically, increasing the number of trials could narrow the confidence bands of the passive gain and the probability of detection curves, as well as the peri-stimulus time histograms. Similarly, more trials might sharpen differences between the onset rapidness distribution values. Further to this, a more fine grained choice of surface areas may help elucidate the relationship between the surface area of the stimulus target and e.g. the measured onset rapidness. In the onset rapidness case, for example, the values for global (100 %) photostimulation and basal dendritic ($\approx$40 %) photostimulation are almost identical, suggesting that a higher resolution binning of surface areas between $\approx$40 % and $\approx$2.66 % (somatic photostimulation) may provide a more compelling and statistically convincing argument for a relationship. The same could be said for the case of the effective passive bandwidth for surface areas between 100 % and $\approx$40 %. Last but not least, the subtle step size used in these experiments could be increased to see if these relationships still hold.

Finally, this thesis presents work demonstrating the response of a neuronal population to truly subtle step changes in the mean of a stochastic stimulus under whole field stimulation.

Subcellular targeting could be extended to the laboratory by examining how one can expose individual neurons to a narrow beam photostimulus. Beams suitably narrow to stimulate e.g. a soma are technically possible to generate – the blue light of a CD player, for example, is less than 600 nm wide at its target. Applying the light to the target could be achieved using optical methods e.g. laser light, coupled into optical path of a microscope, or a small stimulus source could be physically positioned next to a target. In the latter case, the positioning tools typically employed when patch clamping should be well suited to the task. Once developed, such an subcellular targeting technique could be combined with internal cell recordings allowing the experimenter to simultaneously capture action potential waveforms for onset rapidness measurements and capture spike times for population response measurements.

Bibliography

Abbott, L. F. and P. Dayan (1999). "The effect of correlated variability on the accuracy of a population code." In: *Neural computation* 11.1, pp. 91–101.

Adams, P. R., D. A. Brown, and A. Constanti (1982). "M-currents and other potassium currents in bullfrog sympathetic neurones." In: *The Journal of physiology* 330, pp. 537–572.

Arsiero, M., H.-R. Lüscher, and M. Giugliano (2007). "Real-time closed-loop electrophysiology: towards new frontiers in in vitro investigations in the neurosciences." In: *Archives italiennes de biologie* 145.3-4, pp. 193–209.

Åström, K. J. and T. Hägglund (2006). *Advanced PID Control*. ISA - The Instrumentation, Systems, and Automation Society; Research Triangle Park, NC 27709.

Åström, K. J. and R. M. Murray (2008). *Feedback systems: an introduction for scientists and engineers*. 1st ed. New Jersey: Princeton University Press.

Averbeck, B. B., P. E. Latham, and A. Pouget (2006). "Neural correlations, population coding and computation". In: *Nature Reviews Neuroscience* 7.5, pp. 358–366.

Avery, R. B. and D. Johnston (1996). "Multiple channel types contribute to the low-voltage-activated calcium current in hippocampal CA3 pyramidal neurons." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 16.18, pp. 5567–5582.

Bamberg, E. (2016). *Research Profile Prof. Dr. Ernst Bamberg*. URL: http://www.biophys.mpg.de/en/bamberg.html (visited on 09/26/2016).

Banghart, M., K. Borges, E. Isacoff, D. Trauner, and R. H. Kramer (2004). "Light-activated ion channels for remote control of neuronal firing". In: *Nature Neuroscience* 7.12, pp. 1381–1386.

Berndt, A., O. Yizhar, L. A. Gunaydin, P. Hegemann, and K. Deisseroth (2009). "Bi-stable neural state switches". In: *Nature Neuroscience* 12.2, pp. 229–234.

Bernstein, J. (1902). "Untersuchungen zur Thermodynamik der bioelektrischen Ströme". In: *Pflüger Archiv für die Gesammte Physiologie des Menschen und der Thiere* 92.10-12, pp. 521–562.

Bölinger, D. and T. Gollisch (2012). "Closed-Loop Measurements of Iso-Response Stimuli Reveal Dynamic Nonlinear Stimulus Integration in the Retina". In: *Neuron* 73.2, pp. 333–346.

Boucsein, C., T. Tetzlaff, R. Meier, A. Aertsen, and B. Naundorf (2009). "Dynamical response properties of neocortical neuron ensembles: multiplicative versus additive noise." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 29.4, pp. 1006–1010.

Boullay, P. F. G. (1812). "Analyse chimique de la coque du Levant, Menispermum cocculus". In: *Bulletin de Pharmacie* 1, pp. 5–34.

Boyden, E. S., F. Zhang, E. Bamberg, G. Nagel, and K. Deisseroth (2005). "Millisecond-timescale, genetically targeted optical control of neural activity". In: *Nat Neurosci* 8.9, pp. 1263–1268.

Brandt, R. D. and F. Lin (1999). "Adaptive interaction and its application to neural networks". In: *Information Sciences* 121.3–4, pp. 201–215.

Brecht, M., A. Roth, and B. Sakmann (2003). "Dynamic Receptive Fields of Reconstructed Pyramidal Cells in Layers 3 and 2 of Rat Somatosensory Barrel Cortex". In: *The Journal of Physiology* 553.1, pp. 243–265.

Brecht, M. and B. Sakmann (2002). "Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex". In: *The Journal of Physiology* 543.1, pp. 49–70.

Brewer, G. J., J. R. Torricelli, E. K. Evege, and P. J. Price (1993). "Optimized survival of hippocampal neurons in B27-supplemented neurobasal(TM), a new serum-free medium combination". In: *Journal of Neuroscience Research* 35.5, pp. 567–576.

Brunel, N., F. S. Chance, N. Fourcaud-Trocmé, and L. F. Abbott (2001). "Effects of synaptic noise and filtering on the frequency response of spiking neurons". In: *Physical Review Letters* 86.10, pp. 2186–2189.

Callaway, E. M. and L. C. Katz (1993). "Photostimulation using caged glutamate reveals functional circuitry in living brain-slices". In: *Proceedings of the National Academy of Sciences of the United States of America* 90.16, pp. 7661–7665.

Carl Zeiss (2001). *Axiovert 200 / Axiovert 200 M Inverted Microscopes.*

Carnevale, N. T. and M. L. Hines (2006). *The NEURON Book.* Cambridge University Press.

Chien, K. L., J. A. Hrones, and J. B. Reswick (1972). "On the automatic control of generalized passive systems". In: *Trans. ASME* 74, pp. 175–185.

Cohen, G. and G. Coon (1953). "Theoretical consideration of retarded control". In: *Trans. ASME* 75, pp. 827–834.

Cohen, M. R. and A. Kohn (2011). "Measuring and interpreting neuronal correlations." In: *Nature neuroscience* 14.7, pp. 811–819.

Colbert, C. M. and E. Pan (2002). "Ion channel properties underlying axonal action potential initiation in pyramidal neurons." In: *Nature neuroscience* 5.6, pp. 533–8.

Cominos, P. and N. Munro (2002). "PID controllers: recent tuning methods and design to specification". In: *IEE Proceedings - Control Theory and Applications* 149.1, pp. 46–53.

Davies, J. and J. C. Watkins (1982). "Actions of D and L forms of 2-amino-5-phosphonovalerate and 2-amino-4- phosphonobutyrate in the cat spinal cord". In: *Brain research* 235, pp. 378–386.

Dawydow, A., R. Gueta, D. Ljaschenko, S. Ullrich, M. Hermann, N. Ehmann, S. Gao, A. Fiala, T. Langenhan, G. Nagel, and R. J. Kittel (2014). "Channelrhodopsin-2-XXL, a powerful optogenetic tool for low-light applications". In: *Proceedings of the National Academy of Sciences of the United States of America* 111.38, pp. 13972–13977.

Dayan, P. and L. F. Abbott (2002). *Theoretical Neuroscience Computational \& Mathematical Modeling of Neural Systems*. 1st. Cambridge, MA: MIT Press.

Destexhe, A., D. Contreras, T. J. Sejnowski, and M. Steriade (1994). "A model of spindle rhythmicity in the isolated thalamic reticular nucleus". In: *Journal of Neurophysiology* 72.2, 803 LP –818.

Destexhe, A. and M. Rudolph-Lilith (2012). "Models of Synaptic Noise". In: *Neuronal Noise*. Boston, MA: Springer US, pp. 67–110.

Ecker, A. S., P. Berens, G. a. Keliris, M. Bethge, N. K. Logothetis, and A. S. Tolias (2010). "Decorrelated neuronal firing in cortical microcircuits." In: *Science (New York, N.Y.)* 327.5965, pp. 584–587.

Efron, B. (1979). "Bootstrap Methods: Another Look at the Jackknife". In: *The Annals of Statistics* 7.1, pp. 1–26. arXiv: arXiv:1306.3979v1.

Efron, B. and R. Tibshirani (1994). *An Introduction to the Bootstrap*. Florida: Chapman and Hall/CRC.

El Hady, A. (2016). *Closed Loop Neuroscience*. 1st ed. Academic Press.

Ermentrout, G. B. and N. Kopell (1984). "Frequency Plateaus in a Chain of Weakly Coupled Oscillators, I." In: *SIAM Journal on Mathematical Analysis* 15.2, pp. 215–237.

Evans, C. (2014). *scikits.bootstrap version 0.3.2*. URL: https://scikits.appspot.com/bootstrap (visited on 08/25/2016).

Eyal, G., H. D. Mansvelder, C. P. J. de Kock, and I. Segev (2014). "Dendrites impact the encoding capabilities of the axon." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 34.24, pp. 8063–71.

Fatt, P. and B. Katz (1953). "The electrical properties of crustacean muscle fibres." In: *The Journal of physiology* 120.1-2, pp. 171–204.

Fork, R. L. (1971). "Laser stimulation of nerve cells in aplysia". In: *Science* 171.3974, pp. 907–&.

Fourcaud-Trocmé, N., D. Hansel, C. van Vreeswijk, and N. Brunel (2003). "How spike generation mechanisms determine the neuronal response to fluctuating inputs". In: *J. Neurosci.* 23.37, pp. 11628–11640.

Fromme, U. (2016). "Investigation of voltage- and light-sensitive ion channels". PhD thesis. Georg-August Universität.

Fuhrmann, G., I. Segev, H. Markram, and M. Tsodyks (2002). "Coding of temporal information by activity-dependent synapses". In: *Journal of Neurophysiology* 87.1, pp. 140–148.

Georgopoulos, A. P., A. B. Schwartz, and R. E. Kettner (1986). "Neuronal population coding of movement direction." In: *Science* 233.4771, pp. 1416–9.

Giles, J. (2005). "Blue Brain boots up to mixed response." In: *Nature* 435.7043, pp. 720–721.

Grosenick, L., J. H. Marshel, and K. Deisseroth (2015). *Closed-loop and activity-guided optogenetic control.*

Gunaydin, L. A., O. Yizhar, A. Berndt, V. S. Sohal, K. Deisseroth, and P. Hegemann (2010). "Ultrafast optogenetic control". In: *Nature Neuroscience* 13.3, 387–U27.

Gutkin, B. S. and G. B. Ermentrout (1998). "Dynamics of membrane excitability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics." In: *Neural computation* 10.Cv, pp. 1047–1065.

Hamamatsu Photonics Deutschland GmbH (2015). *S2386-8K Calibration Report.* Herrsching.

Hamill, O. P., A. Marty, E. Neher, B. Sakmann, and F. J. Sigworth (1981). "Improved patch clamp techniques for high resolution current recording from cells and cell-free membranes". In: *Pflugers Archiv : European journal of physiology* 391.2, pp. 85–100.

Hay, E. and W. Van Geit (2015). *me-type fact sheet - L5_TTPC1_cADpyr.* URL: https://bbp.epfl.ch/nmc-portal/microcircuit%7B%5C#%7D/metype/L5%7B%5C_%7DTTPC1%7B%5C_%7DcADpyr/details (visited on 08/22/2016).

Hazen, H. (1934). "Theory of servo-mechanisms". In: *Journal of the Franklin Institute* 218.3, pp. 279–331.

Higgs, M. H. and W. J. Spain (2009). "Conditional Bursting Enhances Resonant Firing in Neocortical Layer 2-3 Pyramidal Neurons". In: *Journal of Neuroscience* 29.5, pp. 1285–1299.

Hodgkin, A. L. and A. F. Huxley (1952a). "A quantitative description of membrane current and its application to conduction and excitation in nerve". In: *Journal of Physiology-London* 117, pp. 500–544.

— (1952b). "Currents carried by sodium and potassium ions through the membrane of the giant axon of Loligo." In: *The Journal of physiology* 116, pp. 449–472.

— (1952c). "The components of membrane conductance in the giant axon of Loligo." In: *The Journal of physiology* 116, pp. 473–496.

— (1952d). "The dual effect of membrane potential on sodium conductance in the giant axon of Loligo". In: *The Journal of Physiology* 116, pp. 497–506.

Hodgkin, A. L., A. F. Huxley, and B. Katz (1952). "Measurement of current-voltage relations in the membrane of the giant axon of Loligo". In: *Journal of Physiology* 116, pp. 424–448.

Hodgkin, A. L. and A. F. Huxley (1939). "Action Potentials Recorded from Inside a Nerve Fibre". In: *Nature* 144.3651, pp. 710–711.

Ilin, V., A. Malyshev, F. Wolf, and M. Volgushev (2013). "Fast computations in cortical ensembles require rapid initiation of action potentials." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 33.6, pp. 2281–92.

Klapoetke, N. C., Y. Murata, S. S. Kim, S. R. Pulver, A. Birdsey-Benson, Y. K. Cho, T. K. Morimoto, A. S. Chuong, E. J. Carpenter, Z. Tian, J. Wang, Y. Xie, Z. Yan, Y. Zhang, B. Y. Chow, B. Surek, M. Melkonian, V. Jayaraman, M. Constantine-Paton, G. K.-S. Wong, and E. S. Boyden (2014). "Independent optical excitation of distinct neural populations". In: *Nature Methods* 11.3, 338–U333.

Kloeden, P. E. and E. Platen (1992). "Stochastic Differential Equations". In: *Numerical Solution of Stochastic Differential Equations*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 103–160.

Kohler, M., B. Hirschberg, C. T. Bond, J. M. Kinzie, N. V. Marrion, J. Maylie, J. P. Adelman, and A. M. Köhler (1996). "Small-conductance, calcium-activated potassium channels from mammalian brain". In: *Science* 273.5282, pp. 1709–1714.

Kole, M. H. P., S. Hallermann, and G. J. Stuart (2006). "Single Ih channels in pyramidal neuron dendrites: properties, distribution, and impact on action potential output." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 26.6, pp. 1677–87.

Köndgen, H., C. Geisler, S. Fusi, X.-J. J. Wang, H.-R. Lüscher, and M. Giugliano (2008). "The dynamical response properties of neocortical neurons to temporally modulated noisy inputs in vitro." In: *Cerebral cortex (New York, N.Y. : 1991)* 18.9, pp. 2086–2097.

Koninklijke Philips N.V. (2016). *LUXEON Rebel Color Line*. Amsterdam, Netherlands.

Korngreen, a. and B. Sakmann (2000). "Voltage-gated K+ channels in layer 5 neocortical pyramidal neurones from young rats: subtypes and gradients." In: *The Journal of physiology* 525 Pt 3, pp. 621–639.

Küpfmüller, K. (1928). "{Ü}ber die Dynamik der selbstt{ä}tigen Verst{ä}rkungsregler". In: *Elektrische Nachrichtentechnik* 5.11, pp. 459–467.

Lánský, P. and J. P. Rospars (1995). "Ornstein-Uhlenbeck model neuron revisited". In: *Biological Cybernetics* 72.5, pp. 397–406.

Le Bé, J. V., G. Silberberg, Y. Wang, and H. Markram (2007). "Morphological, electrophysiological, and synaptic properties of corticocallosal pyramidal cells in the neonatal rat neocortex". In: *Cerebral Cortex* 17.9, pp. 2204–2213.

Lee, C., W. H. Rohrer, and D. L. Sparks (1988). "Population coding of saccadic eye movements by neurons in the superior colliculus." In: *Nature* 332.6162, pp. 357–360.

Lima, S. Q. and G. Miesenböck (2005). "Remote control of behavior through genetically targeted photostimulation of neurons". In: *Cell* 121.1, pp. 141–152.

Lin, F., R. D. Brandt, and G. Saikalis (2000). "Self-tuning of PID controllers by adaptive interaction". In: *Proceedings of the 2000 American Control Conference. ACC (IEEE Cat. No.00CH36334)* 5.June, pp. 3676–3681.

Madisen, L., T. Mao, H. Koch, J.-m. Zhuo, A. Berenyi, S. Fujisawa, Y.-W. A. Hsu, A. J. Garcia III, X. Gu, S. Zanella, J. Kidney, H. Gu, Y. Mao, B. M. Hooks, E. S. Boyden, G. Buzsaki, J. M. Ramirez, A. R. Jones, K. Svoboda, X. Han, E. E. Turner, and H. Zeng (2012). "A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing". In: *Nature Neuroscience* 15.5, pp. 793–802.

Magistretti, J. and A. Alonso (1999). "Biophysical properties and slow voltage-dependent inactivation of a sustained sodium current in entorhinal cortex layer-II principal neurons: a whole-cell and single-channel study." In: *The Journal of general physiology* 114.4, pp. 491–509.

Magistretti, J., M. Mantegazza, E. Guatteo, and E. Wanke (1996). *Action potentials recorded with patch-clamp amplifiers: Are they genuine?*

Mainen, Z. F. and T. J. Sejnowski (1996). *Influence of dendritic structure on firing pattern in model neocortical neurons.* URL: http://www.ncbi.nlm.nih.gov/pubmed/8684467.

Markram, H. (2011). "Newsmaker interview: Henry Markram. Blue Brain founder responds to critics, clarifies his goals. Interview by Greg Miller." In: *Science (New York.*

Markram, H., J. Lübke, M. Frotscher, A. Roth, and B. Sakmann (1997). "Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex." In: *The Journal of physiology* 500 (Pt 2, pp. 409–40.

Markram, H. (2006). "The blue brain project." In: *Nature reviews. Neuroscience* 7.2, pp. 153–60.

Markram, H., E. Muller, S. Ramaswamy, M. W. Reimann, M. Abdellah, C. A. Sanchez, A. Ailamaki, L. Alonso-Nanclares, N. Antille, S. Arsever, G. A. A. Kahou, T. K. Berger, A. Bilgili, N. Buncic, A. Chalimourda, G. Chindemi, J. D. Courcol, F. Delalondre, V. Delattre, S. Druckmann, R. Dumusc, J. Dynes, S. Eilemann, E. Gal, M. E. Gevaert, J. P. Ghobril, A. Gidon, J. W. Graham, A. Gupta, V. Haenel, E. Hay, T. Heinis, J. B. Hernando, M. Hines, L. Kanari, D. Keller, J. Kenyon, G. Khazen, Y. Kim, J. G. King, Z. Kisvarday, P. Kumbhar, S. Lasserre, J. V. Le Bé, B. R. C. Magalhães, A. Merchán-Pérez, J. Meystre, B. R. Morrice, J. Muller, A. Muñoz-Céspedes, S. Muralidhar, K. Muthurasa, D. Nachbaur, T. H. Newton, M. Nolte, A. Ovcharenko, J. Palacios, L. Pastor, R. Perin, R. Ranjan, I. Riachi, J. R. Rodríguez, J. L. Riquelme, C. Rössert, K. Sfyrakis, Y. Shi, J. C. Shillcock, G. Silberberg, R. Silva, F. Tauheed, M. Telefont, M. Toledo-Rodriguez, T. Tränkler, W. Van Geit, J. V. Díaz, R. Walker, Y. Wang, S. M. Zaninetta, J. Defelipe, S. L. Hill, I. Segev, and F. Schürmann (2015). "Reconstruction and Simulation of Neocortical Microcircuitry". In: *Cell* 163.2, pp. 456–492.

Mathworks Inc. (2015a). *Compute discrete-time derivative.* URL: http://de.mathworks.com/help/simulink/slref/discretederivative.html (visited on 12/18/2015).

— (2015b). *Perform discrete-time integration of accumulation of signal.* URL: http://de.mathworks.com/help/simulink/slref/discretetimeintegrator.html (visited on 12/18/2015).

— (2015c). *PID Controller, Discrete PID Controller*. URL: http://de.mathworks.com/help/simulink/slref/pidcontroller.html (visited on 12/18/2016).

Maunsell, J. H. and D. C. Van Essen (1983). "Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation". In: *Journal of Neurophysiology* 49.5, 1127 LP –1147.

Microsoft Corporation (2016). *General Reference for the .NET Framework*. URL: https://msdn.microsoft.com/en-us/library/sxe8hcf2(v=vs.100).aspx (visited on 08/22/2016).

Miller, C. and M. M. White (1980). "A voltage-dependent chloride conductance channel from torpedo electroplax membrane". In: *Annals of the New York Academy of Sciences* 341.1 Anion and Pro, pp. 534–551.

Minorsky, N. (1922). "Directional stability of automatically steered bodies". In: *Journal of the American Society of Naval Engineers* 34, pp. 280–309.

Multi Channel Systems GmbH (2007). *MEA1060 Amplifiers*.

— (2008). *Stimulus Generator 1000 Series*. Reutlingen, Germany.

— (2012a). *MEA Amplifier for Inverse Microscopes*.

— (2012b). *Microelectrode Array (MEA) Manual*. Reutlingen, Germany.

— (2016). *McsUsbNet.dll*. URL: http://www.multichannelsystems.com/software/mcsusbnetdll (visited on 08/22/2016).

Nagel, G., D. Ollig, M. Fuhrmann, S. Kateriya, A. M. Mustl, E. Bamberg, and P. Hegemann (2002). "Channelrhodopsin-1: A light-gated proton channel in green algae". In: *Science* 296.5577, pp. 2395–2398.

Nagel, G., T. Szellas, W. Huhn, S. Kateriya, N. Adeishvili, P. Berthold, D. Ollig, P. Hegemann, and E. Bamberg (2003). "Channelrhodopsin-2, a directly light-gated cation-selective membrane channel". In: *Proceedings of the National Academy of Sciences of the United States of America* 100.24, pp. 13940–13945.

Naundorf, B., T. Geisel, and F. Wolf (2005). "Action potential onset dynamics and the response speed of neuronal populations". In: *Journal of Computational Neuroscience* 18.3, pp. 297–309. arXiv: 0411042 [q-bio].

Neef, A., A. E. Hady, J. Nagpal, K. Bröking, G. Afshar, O. O. M. Schlüter, T. Geisel, E. Bamberg, R. Fleischmann, W. Stühmer, F. Wolf, A. El Hady, J. Nagpal, K. Bröking, G. Afshar, O. O. M. Schlüter, T. Geisel, E. Bamberg, R. Fleischmann, W. Stühmer, and F. Wolf (2013). "Continuous Dynamic Photostimulation - inducing in-vivo-like fluctuating conductances with Channelrhodopsins". In: *ArXiv e-prints*. arXiv: 1305.7125 [q-bio.NC].

Neher, E. and B. Sakmann (1976). "Single-channel currents recorded from membrane of denervated frog muscle-fibers". In: *Nature* 260.5554, pp. 799–802.

Newland, C. F. and S. G. Cull-Candy (1992). "On the mechanism of action of picrotoxin on GABA receptor channels in dissociated sympathetic neurones of the rat." In: *The Journal of Physiology* 447.1, pp. 191–213.

Newman, J. P., M.-f. Fong, D. C. Millard, C. J. Whitmire, G. B. Stanley, and S. M. Potter (2015). "Optogenetic feedback control of neural activity". In: *eLife* 4. Ed. by M. Bartos.

Nyquist, H. (1932). "Regeneration theory". In: *Bell System Technical Journal* 11, pp. 126–147.

Okun, M. and I. Lampl (2008). "Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities". In: *Nature Neuroscience* 11.5, pp. 535–537.

Olverman, H. J., A. W. Jones, and J. C. Watkins (1984). "L-glutamate has higher affinity than other amino acids for [3H]-D-AP5 binding sites in rat brain membranes". In: *Nature* 307.5950, pp. 460–2.

Ostojic, S., G. Szapiro, E. Schwartz, B. Barbour, N. Brunel, and V. Hakim (2015). "Neuronal Morphology Generates High-Frequency Firing Resonance". In: *Journal of Neuroscience* 35.18, pp. 7056–7068.

Paradiso, M. A. (1988). "A theory for the use of visual orientation information which exploits the columnar structure of striate cortex". In: *Biological Cybernetics* 58.1, pp. 35–49.

Pearce, J. M. S. (2001). "Emil Heinrich Du Bois-Reymond (1818-96)". In: *Journal of Neurology, Neurosurgery & Psychiatry* 71.5, pp. 620–620.

Potter, S. M., A. El Hady, and E. E. Fetz (2014). "Closed-loop neuroscience and neuroengineering". In: *Frontiers in Neural Circuits* 8, p. 115.

Pouget, A., P. Dayan, and R. Zemel (2000). "Information processing with population codes". In: *Nature Reviews Neuroscience* 1.2, pp. 125–132.

Purves, D., G. J. Augustine, D. Fitzpatrick, W. C. Hall, A.-S. LaMantia, J. 0. McNamara, and S. M. Williams (2004). *Neuroscience*. 4th ed. Sunderland, MA, USA: Sinauer Associates, Inc.

Randall, A. and R. Tsien (1997). "Contrasting biophysical and pharmacological properties of T-type and R-type calcium channels". In: *Neuropharmacology* 36.7, pp. 879–893.

Reiner, A. and E. Y. Isacoff (2016). "The Brain Prize 2013: the optogenetics revolution". In: *Trends in Neurosciences* 36.10, pp. 557–560.

Rettig, J., F. Wunder, M. Stocker, R. Lichtinghagen, F. Mastiaux, S. Beckh, W. Kues, P. Pedarzani, K. H. Schröter, and J. P. Ruppersberg (1992). "Characterization of a Shaw-related potassium channel family in rat brain." In: *The EMBO Journal* 11.7, pp. 2473–2486.

Reuveni, I., A. Friedman, Y. Amitai, and M. J. Gutnick (1993). "Stepwise repolarization from Ca2+ plateaus in neocortical pyramidal cells: evidence for nonhomogeneous distribution of HVA Ca2+ channels in dendrites." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 13.11, pp. 4609–21.

Ricciardi, L. M. and L. Sacerdote (1979). "The Ornstein-Uhlenbeck process as a model for neuronal activity". In: *Biological Cybernetics* 35.1, pp. 1–9.

Salinas, E. and L. F. Abbott (1994). "Vector reconstruction from firing rates". In: *Journal of Computational Neuroscience* 1.1-2, pp. 89–107.

Samhaber, R., M. Schottdorf, A. El Hady, K. Bröking, A. Daus, C. Thielemann, W. Stühmer, and F. Wolf (2016). "Growing neuronal islands on multi-electrode arrays using an accurate positioning-$\{\mu\}$CP device". In: *Journal of Neuroscience Methods* 257, pp. 194–203.

Seung, H. S. and H. Sompolinsky (1993). "Simple models for reading neuronal population codes." In: *Proceedings of the National Academy of Sciences of the United States of America* 90.22, pp. 10749–53.

Sharp, A. A., M. B. O'Neil, L. F. Abbott, and E. Marder (1993). "The dynamic clamp: artificial conductances in biological neurons". In: *Trends in Neurosciences* 16.10, pp. 389–394.

Sheardown, M. J., E. O. Nielsen, A. J. Hansen, P. Jacobsen, and T. Honoré (1990). "2,3-Dihydroxy-6-nitro-7-sulfamoyl-benzo(F)quinoxaline: a neuroprotectant for cerebral ischemia." In: *Science (New York, N.Y.)* 247.4942, pp. 571–4.

Shinomoto, S., K. Shima, and J. Tanji (2003). "Differences in spiking patterns among cortical neurons." In: *Neural computation* 15.12, pp. 2823–2842.

Silberberg, G., M. Bethge, H. Markram, K. Pawelzik, and M. Tsodyks (2004). "Dynamics of population rate codes in ensembles of neocortical neurons". In: *Journal of Neurophysiology* 91.2, pp. 704–709.

Sompolinsky, H., H. Yoon, K. Kang, and M. Shamir (2001). "Population coding in neuronal systems with correlated noise". In: *Physical Review E* 64.5, p. 51904.

Sparks, D. L., R. Holland, and B. L. Guthrie (1976). "Size and distribution of movement fields in the monkey superior colliculus". In: *Brain Research* 113.1, pp. 21–34. arXiv: NIHMS150003.

Stanford, T. R., S. Shankar, D. P. Massoglia, M. G. Costello, and E. Salinas (2010). "Perceptual decision making in less than 30 milliseconds." In: *Nature neuroscience* 13.3, pp. 379–85.

Swadlow, H. a. and T. P. Hicks (1996). "Somatosensory cortical efferent neurons of the awake rabbit: latencies to activation via supra–and subthreshold receptive fields." In: *Journal of neurophysiology* 75.4, pp. 1753–1759.

Tchumatchenko, T., A. Malyshev, F. Wolf, and M. Volgushev (2011). "Ultrafast population encoding by cortical neurons." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 31.34, pp. 12171–9.

Tchumatchenko, T. and J. P. Newman (2013). "Delivery of continuously-varying stimuli using channelrhodopsin-2". In: *Frontiers in neural circuits* 7.December, pp. 1–9.

Texas Instruments Inc. (1994). *OPT101 Monolithic Photodiode and Single-Supply Transimpedance Amplifier - Data Sheet.* Dallas.

Thorpe, S., D. Fize, and C. Marlot (1996). *Speed of processing in the human visual system.*

Uhlenbeck, G. E. and L. S. Ornstein (1930). "On the Theory of the Brownian Motion". In: *Physical Review* 36.5, pp. 823–841.

Urquizo, A. (2011). *A block diagram of a PID controller in a feedback loop.* URL: http://commons.wikimedia.org/wiki/File:PID.svg,%20CC%20BY-SA%203.0,%20https://commons.wikimedia.org/w/index.php?curid=17633925 (visited on 09/22/2016).

Vacučiaková, L. (2016). "The influence of cellular morphology and sub-threshold conductances on the neuronal transfer function". Master's Thesis. Georg August University, Göttingen.

Vogels, R. (1990). "Population coding of stimulus orientation by striate cortical cells". In: *Biological Cybernetics* 64.1, pp. 25–31.

Wagenaar, D. A. (2005). "Controlling Bursting in Cortical Cultures with Closed-Loop Multi-Electrode Stimulation". In: *Journal of Neuroscience* 25.3, pp. 680–688.

Wallach, A. (2013). "The response clamp : functional characterization of neural systems using closed-loop control". In: *Frontiers in Neural Circuits* 7.January, pp. 1–7.

Wallach, A., D. Eytan, A. Gal, C. Zrenner, and S. Marom (2011). "Neuronal response clamp." In: *Frontiers in neuroengineering* 4.April, p. 3. arXiv: 1008.1410.

Wang, Y., A. Gupta, M. Toledo-Rodriguez, C. Z. Wu, and H. Markram (2002). "Anatomical, physiological, molecular and circuit properties of nest basket cells in the developing somatosensory cortex." eng. In: *Cerebral cortex (New York, N.Y. : 1991)* 12.4, pp. 395–410.

Wang, Y., M. Toledo-Rodriguez, A. Gupta, C. Wu, G. Silberberg, J. Luo, and H. Markram (2004). "Anatomical, physiological and molecular properties of Martinotti cells in the somatosensory cortex of the juvenile rat." In: *The Journal of physiology* 561.Pt 1, pp. 65–90.

Wei, W. and F. Wolf (2011). "Spike onset dynamics and response speed in neuronal populations". In: *Physical Review Letters* 106.8. arXiv: 1103.0668.

Wu, S., S.-I. Amari, and H. Nakahara (2002). "Population coding and decoding in a neural field: a computational study". In: *Neural Computation* 14.5, pp. 999–1026.

Zemelman, B. V., G. A. Lee, M. Ng, and G. Miesenböck (2002). "Selective photostimulation of genetically chARGed neurons". In: *Neuron* 33.1, pp. 15–22.

Zemelman, B. V., N. Nesnas, G. A. Lee, and G. Miesenböck (2003). "Photochemical gating of heterologous ion channels: remote control over genetically designated populations of neurons." In: *Proceedings of the National Academy of Sciences of the United States of America* 100.3, pp. 1352–7.

Ziegler, J. G. and N. B. Nichols (1942). "Optimum Settings for Automatic Controllers". In: *Transactions of ASME* 64, pp. 759–768.

Zohary, E., M. N. Shadlen, and W. T. Newsome (1994). *Correlated neuronal discharge rate and its implications for psychophysical performance.* URL: http://www.ncbi.nlm.nih. gov/pubmed/8022482.

Appendices

Appendix A

Additional Results

The results presented in Chapters 3 and 4 are supplemented in this Appendix. The results presented here are similar to those already presented, and were not, therefore explicitly mentioned in the results chapters. They are, however, presented here for completeness. Section A.1 presents additional controller behaviour results. This is followed by Sect. A.2, which presents comparative plots of subcellular targeting metrics.

A.1   Controller Behaviour

The results of applying an adaptive feedback controller *in vitro*, to control the spike rate of a neuron or neurons are presented in Sect. 3.5. The results of two additional experiments are presented here (Figs. A.1 & A.2), which show behaviour similar to that discussed in Sect. 3.5.

A.2   Subcellular Targeting

Section 4.4 compares metrics from numerical subcellular targeting experiments: the population response speed, onset rapidness and passive bandwidth. In Fig. 4.7 the three metrics are compared, where the bandwidth *excludes* Channelrhodopsin-2 dynamics. The results presented here (Fig. A.3), which *include* Channelrhodopsin-2 dynamics in the bandwidth measurement, do not show any significant differences.

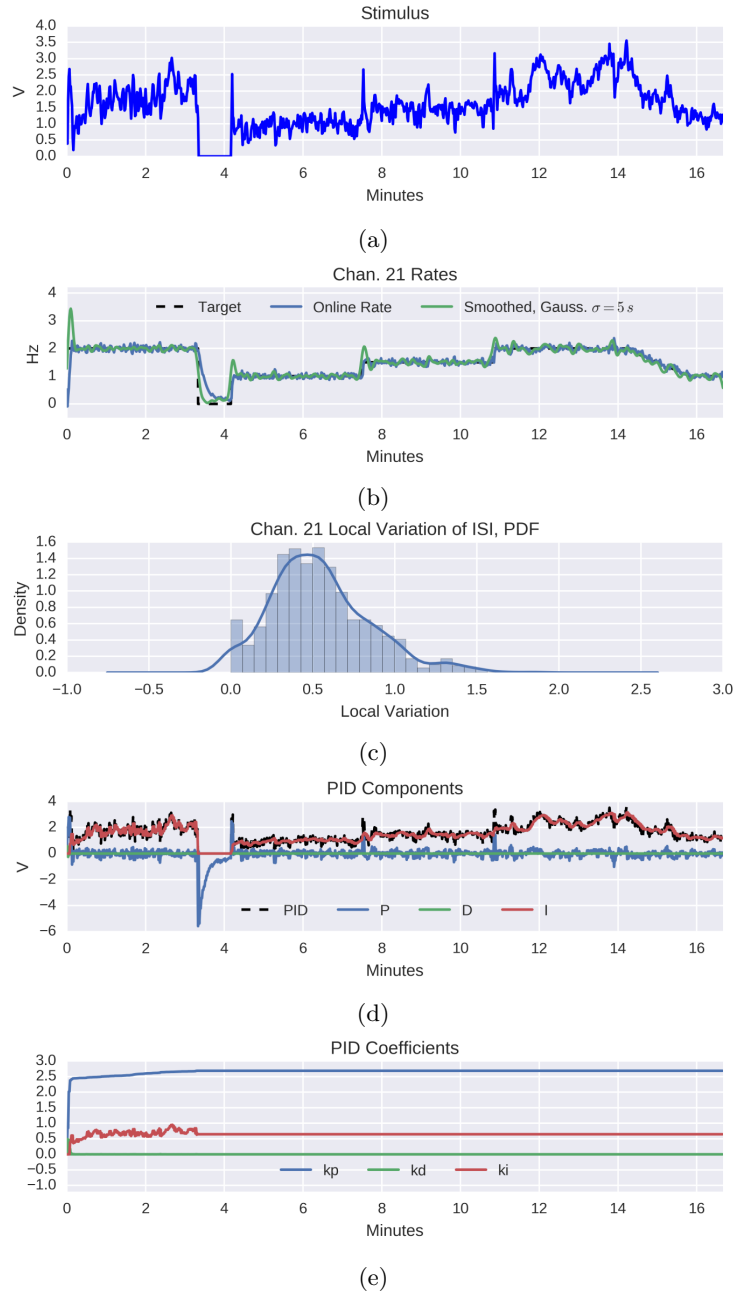(a)

(b)

(c)

(d)

(e)

Figure A.1: OU-process mode stimulus: Step response. *a)* The downsampled controller output (Fourier method). Unlike in the case of the culture ersatz, the stimulus is not perfectly proportional to the target, suggesting some non-linearities in the response of the cell or cells. *b)* The target, online and offline calculated spike rates. Although somewhat jumpy during the adaptation period (first 200 s), the cell or cells appear to follow the target, with some overshoot at steps and some noise. *c)* The distribution of ISI local variation (Eq. 2.2, 10 s moving window). This is distributed about 0.5, indicating significantly non-clock-like behaviour. *d)* The contribution of P, I & D components to the control signal (see *a*). The integral component contributes most to the controller output, except where it is disengaged. The proportional component provides the next most significant contribution. The derivative component appears to make no contribution. *e)* The P, I & D weights. The first 200 s show adaptation of the coefficients, after which adaptation is disengaged. The proportional coefficient is largest, the integral coefficient is approximately a third smaller. The derivative coefficient appears is or is near zero.
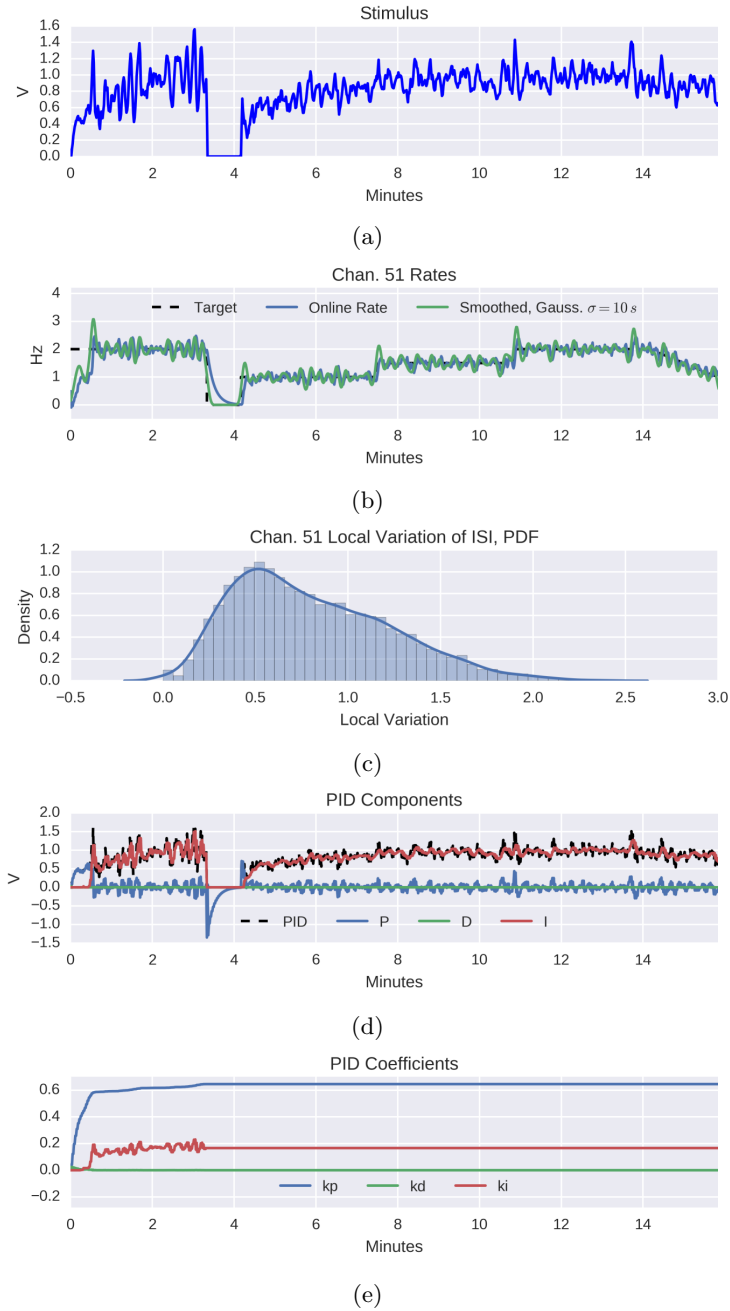
(a)

(b)

(c)

(d)

(e)

Figure A.2: OU-process mode stimulus: Step response. *a)* The downsampled controller output (Fourier method). Unlike in the case of the culture ersatz, the stimulus is not perfectly proportional to the target, suggesting some non-linearities in the response of the cell or cells. *b)* The target, online and offline calculated spike rates. Although somewhat jumpy during the adaptation period (first 200 s), the cell or cells appear to follow the target, with some overshoot at steps and some noise. *c)* The distribution of ISI local variation (Eq. 2.2, 10 s moving window). This is distributed about 0.5, indicating significantly non-clock-like behaviour. *d)* The contribution of P, I & D components to the control signal (see *a*). The integral component contributes most to the controller output, except where it is disengaged. The proportional component provides the next most significant contribution. The derivative component appears to make no contribution. *e)* The P, I & D weights. The first 200 s show adaptation of the coefficients, after which adaptation is disengaged. The proportional coefficient is largest, the integral coefficient is approximately a third smaller. The derivative coefficient appears is or is near zero.
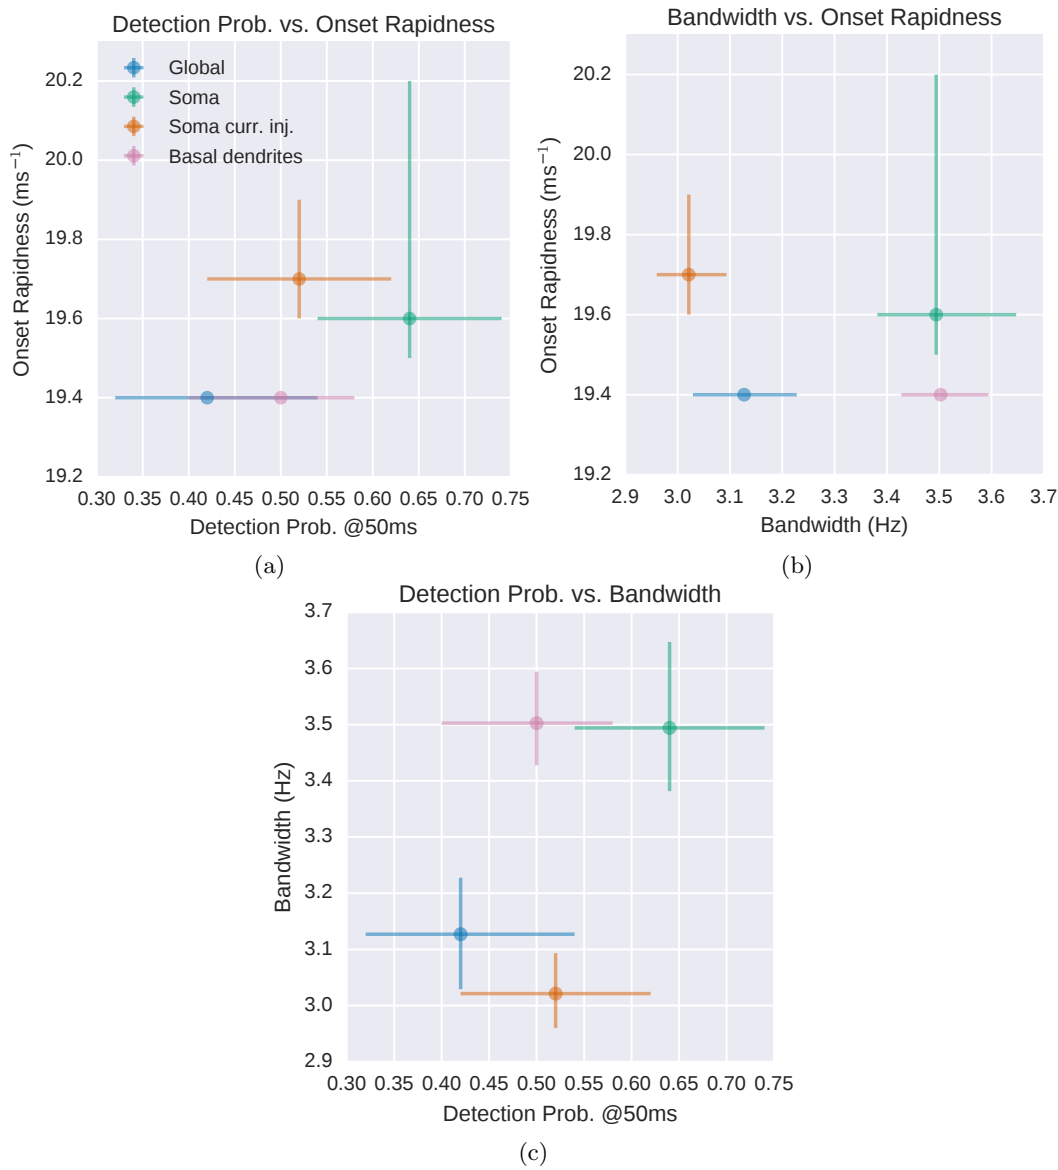
(a)

(b)

(c)

Figure A.3: A comparison of three metrics: The onset rapidness as measured at the soma (Fig. 4.3), the effective bandwidth including Channelrhodopsin-2 (ChR-2) dynamics (Fig. 4.6b) and the probabilities of detection at 50 ms (Fig. 5.1b). The values plotted are the bootstrapped medians with 95 % confidence intervals. *a)* The probability of detection at 50 ms against onset rapidness. *b)* The effective bandwidth versus onset rapidness. *c)* The the probability of detection at 50 ms against effective bandwidth. Measuring passive bandwidth and including the presence of ChR-2 dynamics, there is little change in the relationships between the 3 metrics compared to Fig. 4.7. Focussing in on the two extreme cases, somatic and global photostimulation, congruent increases in all three metrics can still be seen.