

Control of neuronal circuits

from biology to robotics



Raphael Schmetterling

Supervisor: Prof. R. Sepulchre

Department of Engineering
University of Cambridge

This thesis is submitted for the degree of
Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or, is being concurrently submitted, for any degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Raphael Schmetterling

October 2024

Acknowledgements

I would like to thank my supervisor, Rodolphe Sepulchre. His vision inspired me to pursue this path, while his encouragement and positivity kept me going when progress seemed most difficult. Kind and empathic, while also challenging me to fulfil my potential, I am truly grateful to have been his student.

I would also like to thank Tim O’Leary, whose deep intuition and pragmatic advice helped shape the direction of this thesis. Thanks also to Fulvio Forni for his unwavering support and his attention to detail, and for teaching me to work through research questions in a more systematic way.

It has been a joy to work in the Control Group, surrounded by curious, down-to-earth colleagues and friends. Particular thanks to Luka Ribar and Thiago Burghi, who assisted me with many a technical question and who helped me find my feet in those early years.

Finally, I would like to thank my parents for their boundless love, support and encouragement.

Collaborations

Chapter 3: Thiago B. Burghi and Rodolphe Sepulchre

Chapter 4: Thiago B. Burghi and Rodolphe Sepulchre

Chapter 5: Fulvio Forni, Alessio Franci and Rodolphe Sepulchre

Abstract

Rhythmic robotic controllers often take inspiration from neuroscience. Due to the complexity of biophysical neurons, neuro-inspired controllers are frequently founded on abstract models such as non-linear oscillators. Such controllers retain some of the desirable properties of neuronal circuits, but lack the physical embodiment that is characteristic of biological systems and that is key to their effectiveness. In this thesis we take a step towards reconciling biophysics with bio-inspired control, using the language of control theory to do so. The proposed controller is event-based, and has the physical realisation of an analogue neuromorphic circuit. The event-based behaviour is rooted in the presence of mixed feedback. The behaviour is regulated using output feedback and also an adaptive control that tunes the gains of the positive and negative feedback loops. Adaptive control is aligned with neuromodulation, which is central to the adaptation and robustness of animal nervous systems. We illustrate the potential of the event-based neuromorphic approach on the simple mechanical model of a pendulum.

In addition to the pendulum controller, we also propose a methodology for the control of biological or neuromorphic neuronal circuits. In particular, we explore the classical paradigm of indirect adaptive control to design neuromodulatory controllers in biophysical neuronal models. This provides a methodology that aligns with impedance control in robotics. The method relies on parameter estimates obtained with a recently-proposed adaptive observer that implements a centralized recursive least squares algorithm. Inspired by biology, we show that decentralization and redundancy help recover the performance of this algorithm in the presence of uncertainty and mismatch on the internal dynamics of the model.

Table of contents

List of figures	xi
List of tables	xv
1 Introduction	1
1.1 Outline	2
2 Background	5
2.1 Neuromorphic control	5
2.2 Technical preliminaries	7
2.3 Conductance-based modelling	7
2.3.1 Conductance-based model of a neuron	7
2.3.2 Conductance-based model of a neural network	13
2.4 A minimal neuromorphic circuit	14
2.5 Rhythmic neuronal networks	17
2.6 Adaptive observers for conductance-based models	19
Appendices	23
Appendix 2.A Gating Dynamics	23
3 Adaptive conductance control	25
3.1 Introduction	25
3.2 Model reference adaptive conductance control	26
3.2.1 Adaptive reference tracking	27
3.2.2 Adaptive disturbance rejection	29
3.2.3 Network neuromodulation	34
3.3 Discussion	36
3.4 Conclusion	37

Appendices	39
Appendix 3.A Model Parameters	39
4 Robust online estimation of neural circuits	41
4.1 Introduction	41
4.2 Biophysical neuron model	42
4.3 Adaptive recursive least squares estimation	43
4.4 Robustness of the observer	46
4.5 Distributed observer	47
4.6 Redundancy	50
4.7 Discussion	53
5 Neuromorphic control of a pendulum	55
5.1 Introduction	55
5.2 The automaton of a pendulum	56
5.3 Event-based control architecture	57
5.4 Neuromorphic implementation	60
5.5 Control by entrainment	61
5.6 Phase control	63
5.7 Neuromorphic adaptive control	64
5.8 Discussion	65
6 Conclusions	69
6.1 Summary	69
6.2 Future work	70
References	75

List of figures

- 2.1 Circuit representation of a neuron with voltage v that is coupled with synapses to presynaptic neurons i , each with voltage $v_{p,i}$ 8
- 2.2 Two bursts produced by the neuron model in Example 2, showing the membrane potential v (top left) and three selected ionic currents. 11
- 2.3 Fast and slow I-V curves of a spiking circuit. The fast curve is ‘N-shaped’. The slow curve is monotone and its intersection with the line $I = I_{app}$ determines the system’s equilibrium (V_e). The neuron spikes when $V_1^f < V_e < V_2^f$. On the slow curve we see two example values for the applied current, and the corresponding values for V_e . Figure and caption reproduced from [89]. 16
- 2.4 The I-V curves (top) show a system with a double hysteresis: both the fast and the slow I-V curves are ‘N-shaped’. By having the ‘up’ state of the slow curve correspond to the unstable region of the fast system, the system experiences rest-spike bistability, given that the slow threshold is at lower voltage than the fast one, i.e. $V_1^s < V_1^f$. The system has two pairs of threshold and range voltages corresponding to the fast and the slow I-V curves. These variables characterise the circuit’s behaviour, and are adjusted by control of the gains g_x^\pm . The ultraslow curve (not shown) is monotonically increasing. Figure and caption reproduced from [89]. 18
- 2.5 The behaviour of an HCO circuit following the model of Section 2.4. *Left*: the network diagram, showing the two neurons and the inhibitory synapses between them. *Right*: the output behaviour of each neuron. Reproduced from [88]. 19
- 3.2.1 Block diagram representation of the adaptive reference tracking and disturbance rejection problems. Adapted from [6, Chapter 1]. 27
- 3.2.2 Circuit diagram of the adaptive reference tracking problem. 29
- 3.2.3 Simulation of single-neuron reference tracking. 30

3.2.4 Circuit diagram of the disturbance rejection problem, including the required controller.	32
3.2.5 Simulation of single-neuron disturbance rejection.	33
3.2.6 Network disturbance rejection.	35
4.2.1 Fixed scenario for the illustrations of this chapter. <i>Top</i> : membrane potential. <i>Middle</i> : input current. <i>Bottom</i> : maximum conductance of L-type calcium and calcium-activated potassium (the other parameters remain constant). The fluctuating input is fixed in that this particular input signal is used throughout the chapter. It was originally generated by low-pass filtering Gaussian noise. The properties of the filter were changed partway during the simulation to reflect the difference between spike and burst excitability.	44
4.3.1 RLS estimation in the absence of model error. Parameter and output estimates converge to their true values, and remain there barring transient errors while μ_{CaL} and μ_{KCa} are modulated. <i>Top</i> : true voltage v and its estimate \hat{v} . <i>Middle</i> : absolute observer error $ v - \hat{v} $. <i>Bottom</i> : L-type maximal conductance μ_{CaL} , its estimate $\hat{\theta}_{CaL}$, calcium-activated potassium maximal conductance μ_{KCa} , and its estimate $\hat{\theta}_{KCa}$	45
4.4.1 The centralized observer is fragile to model uncertainty. <i>Top</i> : true voltage and its estimate. <i>Middle</i> : absolute observer error. <i>Bottom</i> : time-varying maximal conductances and their estimates.	47
4.5.1 Distributed parameter estimation in the scenario of Fig. 4.2.1. <i>Top</i> : true voltage and its estimate. <i>Middle</i> : absolute observer error. <i>Bottom</i> : μ_{CaL} and μ_{KCa} and their estimates. Although there is some observer error, it is significantly smaller as reflected in the v and \hat{v} spikes being almost indistinguishable.	49
4.6.1 Distributed and redundant online estimation in the scenario of Fig. 4.2.1, in the presence of model error and with $N = 3$. <i>Top</i> : true voltage and its estimate. <i>Upper middle</i> : absolute observer error. <i>Lower middle</i> : the empirical mean estimate $\bar{\theta}_{CaL}$; for comparison, the scaled true parameter μ_{CaL}/N is also shown. We plot the <i>scaled</i> parameter to emphasise that we have replaced each conductance with N separate conductances. <i>Bottom</i> : the same for μ_{KCa} . Labels for the parameter estimates are omitted.	52
5.3.1 The two distinct configurations of event actuation. <i>Left</i> : in-phase identical actuating events (IN-PHASE). <i>Right</i> : anti-phase actuating events of opposite sign (ANTI-PHASE). Sketch courtesy of @artjoy2015.	58

- 5.3.2 The anti-phase rhythm of a Half-Centre Oscillator (HCO) circuit. The size of the bursts is modulated midway. *Top*: the voltages of the two neurons. The dashed horizontal line denotes the voltage threshold, above which the motor is active. *Bottom*: the corresponding actuating events. 59
- 5.3.3 Simulation of the HCOs and actuator events. *Top*: the total torque I . *Middle*: the voltages of HCO 1. *Bottom*: the voltages of HCO 2. At $t = 8$ s (indicated by the dashed black line), the system's configuration is switched from ANTI-PHASE and, after a transient period, it settles on IN-PHASE. 59
- 5.3.4 Block diagram of the complete architecture, including the event-based feedback loops introduced in Sections 5.6 and 5.7. Small arrows over signal transmission lines indicate event-based communication as described in Section 5.3. The HCO block architecture is described in Sections 5.3 and 5.4. 60
- 5.5.1 Small oscillations in the overdamped regime (damping $\alpha = 1.4$). As the burst size of the neural oscillation increases (from top row to bottom row), so does the amplitude of the pendulum's oscillation. 62
- 5.5.2 Entrainment of the pendulum at a fixed frequency in the overdamped regime ($\alpha = 1.4$). The amplitude of small oscillations increases monotonically with g_s^- and g_{us}^+ 62
- 5.5.3 Entrainment with an IN-PHASE neural oscillation in the underdamped regime (in this example, $\alpha = 0.14$), showing energy transferred by the motors each burst. The amplitude of each oscillation is also shown (with an amplitude of π indicating a complete swing). Depending on the initial conditions, the steady-state oscillations are either large or small (left and right plots respectively). 63
- 5.6.1 *Left*: A small inhibitory pulse (green) causes a phase advance or delay of the subsequent burst. The blue curve shows v_A for an isolated HCO. The blue dashed curve shows the behaviour of v_A in the absence of the pulse. *Right*: Phase response curve of an HCO. The curve indicates the phase shift resulting from a fixed pulse ($w = 0.05$ and $P = 0.3$) as a function of the timing of the pulse along the limit cycle. 64
- 5.6.2 The simulation of Fig. 5.5.3 (right plot) with the addition of phase control, forcing convergence to the high- rather than the low-energy state. *Left*: the initial response. The current pulses to the A neurons are shown in purple, and are scaled for clarity. Each pink marker indicates the start of an actuation event. *Right*: the energy transferred per event, E_i , showing convergence to a steady-state oscillation. The pink markers match those of the left plot. 65

5.7.1 The gains tuned by the adaptive controller, when used to obtain the parameters in Fig. 5.5.2. The legend gives the values of A_{ref} in radians. In each case, ω_{ref} is the fixed frequency chosen in Section 5.5.	66
6.2.1 An illustration of the ‘inverted pendulum’ walking down a slope. Reproduced from [39].	71
6.2.2 A single step of the inverted pendulum model walking down a slope. Reproduced from [39].	72
6.2.3 A biped driven by a single bursting neuron that provides torque to the stance leg. At the end of each step the stance and swing legs are swapped over (hence the sudden jumps in angle).	73

List of tables

3.A.1 Reversal potentials and neuron capacitance.	40
4.6.1 Comparison of all three observers, showing the mean and standard deviation of the rms voltage error $e_{v,\text{rms}}$ (in mV) across twenty trials.	52

Chapter 1

Introduction

Many robotic tasks are fundamentally rhythmic in nature. This is best illustrated in the context of animal-like movements such as walking or swimming, which exhibit a repeating pattern while varying continuously in response to a changing environment [99]. There is a rich literature of controllers that draw inspiration from neuroscience to solve such tasks [51]. Biophysical neuron models, however, are highly intricate and not practical for design purposes. Neuro-inspired controllers, therefore, often employ abstract models such as non-linear oscillators in order to generate the rhythmic control signal. While these controllers retain some of the desirable properties of neuronal oscillators, such as robustness, they lack a physical realisation whereas in biology the physical embodiment of the control is key to its effectiveness [84]. The aim of the thesis is to close the gap between the physics of neuronal circuits and the abstractions of bio-inspired controllers, using the language of control theory to do so.

Control tools have previously been used to explore the regulation of bio-inspired rhythms, for example the work of [59] that utilises classical linear methods. Our approach begins by highlighting that biophysical neuron models have structural properties that make them susceptible to feedback control. We then employ a simpler neuron model [89] that retains these properties, but that is amenable to design. The simpler model has a physical realisation as an analogue “neuromorphic” circuit; this permits us to propose a control methodology that exploits the physics of the neuromorphic controller and its interactions with the controlled system.

The biophysical nature of our controller also allows us to utilise important control principles from neurophysiology. In particular, we connect the design question of how to ‘tune’ a rhythmic controller with the key biological concept of *neuromodulation* [104, 28]. Neuromodulators are chemicals that vary the rich dynamical properties of neurons. They are present in all nervous systems, and are central to their function [72, 71]. In a model of a neuron, the effect of

neuromodulation is captured as a change to certain parameters, known as maximal conductances, that appear linearly in the model. In the language of control theory this is *adaptive control*, a well-established methodology for tuning the (typically unknown) parameters of a controller in response to changes in the environment [6]. As we will see in Section 2.3, biophysical neuron models have a structure that meets the technical assumptions required for a system to be adaptively controlled, a property that is retained by our choice of neuromorphic circuit. This permits us to use neuromodulation to regulate the behaviour of the overall system.

Increasingly, the concept of *embodied intelligence* is taking a prominent role in bio-inspired robotics. This recognises that to understand the control of an animal's motion, it is not sufficient to study its nervous system. Rather, factors such as the animal's morphology, material properties and interactions with the environment all play essential roles in determining its behaviour [84]. Embodied intelligence underpins the adaptability and energy-efficiency of animal motion and therefore has great potential for robotics, including in the context of soft and compliant systems [77, 43].

In this thesis we explore the potential of neuromorphic control for embodied systems, doing so with the simple mechanical model of a pendulum. The proposed architecture is well-suited to embodied control because it is *event-based*. The neuromorphic circuit controls the pendulum via pulses of motor torque that force the desired rhythm. This actuating signal is sparse, and so the pendulum's motion is primarily governed by its own dynamics, and the control thus leverages the physics of the pendulum itself. There is a rough timescale matching between the pendulum and the circuit, a common theme in embodied systems [49]. The sensors are also event-based and capture only minimal information; this is another constraint exploited by the control architecture.

1.1 Outline

The contributions of the thesis divide into two parts, the second of which is the rhythmic robotic controller discussed above. The first part proposes a methodology for the neuromodulatory control of neuronal circuits, once again exploiting the biophysics of those circuits. We envision applications for this in the field of neurotechnology, which interfaces neural tissues with electronic devices.

Background material is provided in Chapter 2. The novel results are then laid out in the subsequent three chapters.

Adaptive Conductance Control

In Chapter 3 we consider how to control networks of biophysical neurons. Specifically, we

solve for such circuits the adaptive reference tracking and disturbance rejection problems. We use a port interconnection approach that shapes the conductance of the controlled circuit, and that has similarities to impedance control in robotics. This control methodology is novel in that it acknowledges the dynamics of spiking signals, in contrast to contemporary methods that rely on classical linear techniques.

Robust Online Estimation of Neural Circuits

The conductance controller of Chapter 3 relies on the adaptive observer of [16], the robustness of which is explored in Chapter 4. In particular, we investigate the robustness of this observer to model error. We then demonstrate that robustness can be improved by replacing the observer with the distributed version of [13], and also by introducing redundancy to the observer's model of the neuron.

Neuromorphic Control of a Pendulum

In Chapter 5, we use a pendulum to illustrate the potential of neuromorphic electronics for the event-based control of rhythmic robots. The pendulum is regarded as an automaton with two states: large oscillations (complete swings of the pendulum) and small oscillations (less than a full swing). The controller sets the state of the pendulum, and also regulates the continuous properties of each state (the frequency of large oscillations, and the frequency and amplitude of small oscillations). This is accomplished using open-loop entrainment, an output feedback phase control, and an adaptive control. Such an approach seems novel, and it promises to overcome the limitations of existing bio-inspired controllers which rely either on linear techniques, or on non-linear oscillators that are difficult to analyse and tune.

The code underpinning these results can be found at the links in the footnotes, one for each chapter respectively.^{1 2 3}

The work of Chapter 3 has been published with fellow authors Thiago B. Burghi and Rodolphe Sepulchre in [95]. The work of Chapter 4 has been published with the same authors in [96]. The work of Chapter 5 has been published with Fulvio Forni, Alessio Franci and Rodolphe Sepulchre in [97].

¹Chapter 3: <https://github.com/RJZS/adaptive-conductance-control>

²Chapter 4: <https://github.com/RJZS/robust-neuron-estimation>

³Chapter 5: <https://github.com/RJZS/neuromorphic-pendulum-control>

Chapter 2

Background

2.1 Neuromorphic control

Neurons are *event-based systems*. They output a voltage that varies continuously yet is made up of distinct, countable spikes. In this way, they exhibit characteristics of both analogue and digital systems, combining the advantages of each. They are at the same time a physical system that can be tuned continuously, and a discrete automaton that emits events reliably. Control of a neuron therefore involves setting its discrete state (for example, whether it is spiking or not), and also tuning continuous properties that exist within a state (for example, the spiking frequency). Such behaviour maps on to rhythmic tasks in robotics. For example, to control a robot we may desire to switch its state at will between running and walking [41], or between swimming and walking [53], while also controlling continuous properties such as velocity.

Event-based control can be physically realised with neuromorphics, a branch of electronics that was pioneered in the late 1980s by Carver Mead [74] and that now has a rich literature [98, 38]. Mead observed that the physical principles of a transistor, when operating in the analogue regime, are similar to those of a neuron. A transistor in this regime (that is, with its gate-source voltage below threshold) is therefore a natural building block for spiking electronics. Furthermore, an analogue transistor can dissipate as little as 10^{-12} Watts, an amount comparable to that of a biological neuron [74]. Digital circuits, by contrast, require several orders-of-magnitude more power.

Neuromorphic circuits have potential in robotics because they provide a way to solve reference tracking control tasks for rhythmic systems, even when the system model is uncertain and the reference trajectory is unknown. In the classical approach to reference tracking control, one first defines the behaviour of the system in terms of a real-valued state vector that evolves over time. One then designs a controller that can induce the system to track the desired

(“reference”) state behaviour. This method, however, requires a detailed knowledge of the system model [101, 8]. It can therefore be fragile to uncertainty, a challenge faced in Chapters 3 and 4 below. Furthermore, this classical approach assumes that the reference signal is known. But the design of a reference signal (for example, one that maximises energy-efficiency) is itself a major challenge for many non-linear systems. This is the “trajectory planning” problem, well-known in robotics, and a key motivation for the switch to optimisation-based methods such as reinforcement learning [111]. Such methods are a powerful alternative to tracking controllers, but they are computationally-expensive and it is often difficult to demonstrate their robustness [78, 45, 24].

To reconcile rhythmic control with reference tracking, we can change our characterisation of the system behaviour. Instead of a real-valued trajectory, we can use a temporal sequence of events, specifically of sensor and actuator events. Rather than building and tuning a non-linear model we simply assume that, in the absence of an actuator event, the system is free to act under its own dynamics. The reference signal is then the desired sequence of events, and the reference tracking controller is a neuromorphic circuit capable of generating this sequence. We show this in Chapter 5 for global control of the pendulum. Another example is given by the crawling robot of [5].

The event-based behaviour of neurons is grounded in *mixed feedback* [90]. Neurons are constructed out of positive and negative feedback loops, acting in parallel and each with its own timescale and voltage range [103]. Positive feedback promotes discrete events, and negative feedback promotes continuous regulation. Examples of this mixed-feedback design are provided later on in the chapter. Neuromodulation, which is central to neuron control, consists of tuning the gains of these feedback loops in order to regulate the neuron’s behaviour [104].

In the rest of this chapter we introduce necessary preliminaries to understand the research contributions of the thesis. In Section 2.2 we provide some technical definitions. Section 2.3 lays out the traditional conductance-based methodology for modelling neurons. In Section 2.4 we introduce a neuromorphic circuit [89] that represents a minimal model of neuronal behaviour. Section 2.5 discusses small rhythmic networks of neurons. Finally, in Section 2.6 we outline an adaptive observer [16] that can be used to estimate certain parameters of neuron circuits.

2.2 Technical preliminaries

Neuron models meet the three well-known assumptions required for a system to be amenable to adaptive control. One of these is that the controlled parameters appear linearly in the model. The other two are that the system is *minimum phase* and *relative degree one*. These properties will now be defined.

Definition 1 *The relative degree of a system is the number of times that one must differentiate the output until the input appears. As an example, consider the system $\dot{v} = -kv + u$, with input u and output v . This system is relative degree one.*

To define the minimum phase of a non-linear system such as the neuron, we must first define its zero dynamics. The zero dynamics are the system's internal dynamics when the input is such that the output is forced to be at zero [64]. In neurophysiology, the zero dynamics of a neuron are obtained from the voltage clamp experiment [46].

Definition 2 [64] *A non-linear system with affine input is minimum phase if its zero dynamics are globally asymptotically stable.*

2.3 Conductance-based modelling

Since the seminal work of Hodgkin and Huxley [46], biophysical models in neurophysiology have been founded on nonlinear electrical circuits known as *conductance-based models*. A detailed introduction to such models can be found in [54, 57]. In this section, we extend the system-theoretic conductance-based modelling framework found in [16].

2.3.1 Conductance-based model of a neuron

A conductance-based model of a neuron is a one-port electrical circuit with the architecture shown in Figure 2.1: a capacitor of capacitance $c > 0$ in parallel with a *leak current* I_{leak} and several *intrinsic ionic currents*. The capacitance and leak current model the neuron's cell membrane, and the ionic currents model the flow of ions across this membrane. Additional *extrinsic* currents model the *synaptic currents* $I_{\text{syn,p}}$, due to interconnections with other neurons, as well as an *input (or applied) current* u . The latter represents the current injected into the circuit by an intracellular electrode. The capacitor voltage v , which models the neuronal




Fig. 2.1 Circuit representation of a neuron with voltage v that is coupled with synapses to presynaptic neurons i , each with voltage $v_{p,i}$.

membrane potential, evolves according to Kirchhoff's law, that is

$$c \dot{v} = -I_{\text{leak}} - \sum_{\text{ion} \in \mathcal{I}} I_{\text{ion}} - \sum_{\text{syn} \in \mathcal{S}} \sum_{p \in \mathcal{P}} I_{\text{syn},p} + u \quad (2.1a)$$

$$= -I_{\text{leak}} - I_{\text{int}} - I_{\text{ext}} + u \quad (2.1b)$$

where I_{int} is the parallel combination of intrinsic ionic currents, I_{ext} is the parallel combination of synaptic currents, \mathcal{I} is the index set of intrinsic ionic currents, \mathcal{S} is the index set of synaptic neurotransmitter types, and \mathcal{P} is the index set of *presynaptic neurons*. Note that the circuit is relative degree one.

Each current in the circuit is ohmic in nature, but with a conductance that can be nonlinear and voltage-dependent. The leak current has a constant conductance and is given by

$$I_{\text{leak}} = \mu_{\text{leak}}(v - E_{\text{leak}}),$$

whereas the intrinsic ionic currents are modelled by

$$I_{\text{ion}} = \mu_{\text{ion}} m_{\text{ion}}^{p_{\text{ion}}} h_{\text{ion}}^{q_{\text{ion}}} (v - E_{\text{ion}}) \quad (2.2a)$$

$$\tau_{m,\text{ion}}(v) \dot{m}_{\text{ion}} = -m_{\text{ion}} + \sigma_{m,\text{ion}}(v) \quad (2.2b)$$

$$\tau_{h,\text{ion}}(v) \dot{h}_{\text{ion}} = -h_{\text{ion}} + \sigma_{h,\text{ion}}(v) \quad (2.2c)$$

The constants $\mu_{\text{ion}} > 0$ and $E_{\text{ion}} \in \mathbb{R}$ are called (intrinsic) *maximal conductances* and *reversal potentials*, respectively. The static *activation functions* $\sigma_{m,\text{ion}}(v)$ and $\sigma_{h,\text{ion}}(v)$, and *time-constant functions* $\tau_{m,\text{ion}}(v)$ and $\tau_{h,\text{ion}}(v)$, model the nonlinear gating of the ionic conductance. Because $\sigma_{m,\text{ion}} : \mathbb{R} \rightarrow (0, 1)$ and $\sigma_{h,\text{ion}} : \mathbb{R} \rightarrow (0, 1)$ are monotonically increasing and decreasing, respectively, the states m_{ion} and h_{ion} are called *activation* and *inactivation gating variables*, respectively. The time-constant functions vary in shape, but always respect the bounds

$$0 < \underline{\tau}_{\text{ion}} \leq \tau_{m,\text{ion}}(v), \tau_{h,\text{ion}}(v) \leq \bar{\tau}_{\text{ion}}$$

for all $v \in \mathbb{R}$ and some $\underline{\tau}_{\text{ion}}, \bar{\tau}_{\text{ion}} > 0$. The exponents p_{ion} and q_{ion} in (2.2a) are natural numbers (including zero). Each gating variable can be thought of as “opening” or “closing” a particular ionic channel; exponents greater than one represent multiple identical “gates” in series [20].

To develop an intuition for the internal dynamics (2.2b) and (2.2c), consider the case where the activation and time-constant functions are replaced by constant values. The dynamics of each gating variable would then be those of a linear low-pass filter, a simplification used by the neuromorphic circuit in Section 2.4. In both the linear and non-linear cases, the initial conditions of the gating variables exponentially decay and so will be neglected in this thesis.

Note that the boundedness of the activation functions ensures that the set $[0, 1]$ is a positively invariant set for each gating variable, uniformly in v . Similarly, the reversal potentials E_{ion} ensure the presence of a positively invariant set for v . The internal dynamics are globally asymptotically stable when $v = 0$, and so the neuron is a minimum phase system. See [16] for proofs of these statements.

The activation and time-constant functions, as well as the exponents p_{ion} and q_{ion} , can be obtained via the “voltage clamp” experiment of neurophysiology, whereby high-gain inverting feedback is used to obtain the step response of the internal dynamics. This is repeated for voltage steps of different amplitudes, resulting in a family of step response curves. A model is then selected that best fits these results. The voltage clamp method was first introduced by Hodgkin and Huxley [46].

Example 1 *Conductance-based modelling originated in the seminal paper of Hodgkin and Huxley [46]. The Hodgkin-Huxley (HH) model includes two intrinsic ionic currents: a transient sodium current I_{Na} and a potassium current I_K , so that $\mathcal{I} = \{\text{Na}, \text{K}\}$. The voltage dynamics of a single, isolated HH model (no synaptic currents) are given by*

$$c \dot{v} = - \underbrace{\mu_{Na} m_{Na}^3 h_{Na} (v - E_{Na})}_{I_{Na}} - \underbrace{\mu_K m_K^4 (v - E_K)}_{I_K} - \underbrace{\mu_{leak} (v - E_{leak})}_{I_{leak}} + u$$

and the gating variable dynamics are of the form (2.2b)-(2.2c) for $\text{ion} \in \{\text{Na}, \text{K}\}$.

The HH model only includes two voltage-dependent conductances to parameterise the intrinsic current. Those types of currents are necessary and sufficient to model an action potential, or spike [54]. This is due to the presence of fast-activating negative conductance and slower-activating positive conductance, which act as sources of positive and negative feedback respectively [27]. The fast negative conductance is provided by the activation of the inward current I_{Na} (due to the activation gating variable m_{Na}) and the slower positive conductance is provided by both the inactivation of I_{Na} (due to h_{Na}) and the activation of the outward current I_K (due to m_K) [102, 36]. This pairing of fast negative and slower positive conductance, or equivalently fast positive and slower negative feedback, is the source of the spiking behaviour and is the simplest example of the “mixed-feedback motif”, a pattern that is fundamental to neuronal behaviour [100]. For an intuition of the role played by mixed feedback, see Section 2.4 below.

Biophysical neurons may exhibit much richer behaviours, including for instance transitions between spiking and bursting patterns. The conductance-based models of such neurons include more conductances, leading to a plethora of single-neuron models differing from each other by the kinetics and activation ranges of their gating variables.

Example 2 *For the sake of illustration, all the conductance-based neurons in the rest of this chapter and in Chapter 3 follow the same model, similar to the model in [25]. This model includes eight distinct ohmic current sources, each modelling representative currents of the experimental neurophysiological literature. Conductance control of such a model refers to adjusting the eight maximal parameter conductances of the model to shape the total intrinsic current of the neuron. We label those distinct currents with the index set $\mathcal{I} = \{\text{Na}, \text{H}, \text{T}, \text{A}, \text{K}, \text{L}, \text{KCa}, \text{KIR}\}$. Some of these currents share the same type of ion, and so*

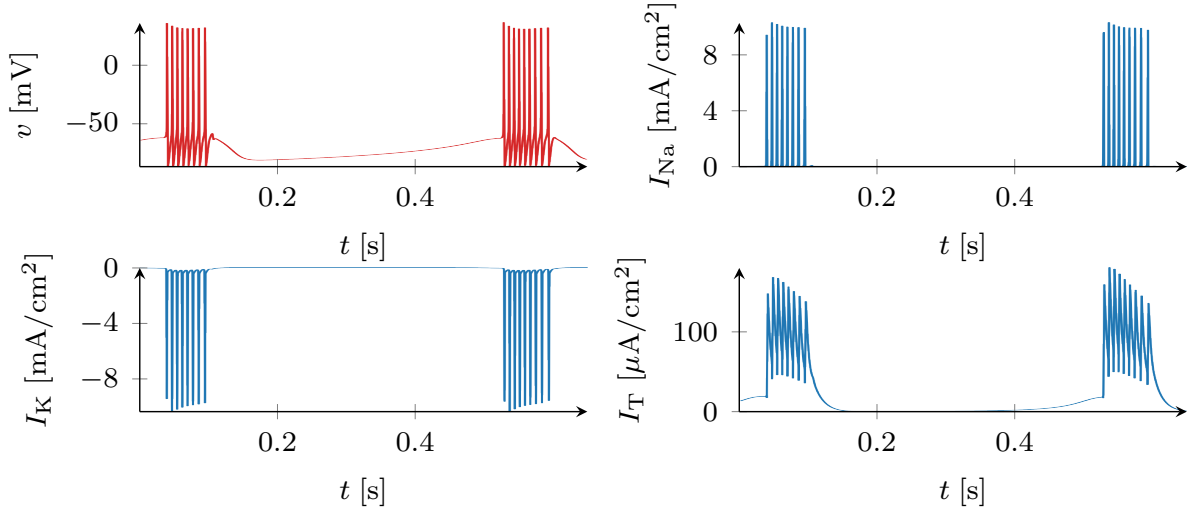


Fig. 2.2 Two bursts produced by the neuron model in Example 2, showing the membrane potential v (top left) and three selected ionic currents.

share the same reversal potential. The voltage dynamics of this model are as follows:

$$\begin{aligned}
 c\dot{v} = & -\mu_{Na}m_{Na}^3h_{Na}(v - E_{Na}) - \mu_Hm_H(v - E_H) \\
 & - \mu_Tm_T^2h_T(v - E_{Ca}) - \mu_Am_A^4h_A(v - E_K) \\
 & - \mu_Km_K^4(v - E_K) - \mu_Lm_L(v - E_{Ca}) \\
 & - \mu_{KCa}\left(\frac{[Ca]}{15 + [Ca]}\right)^4(v - E_K) - \mu_{KIR}\sigma_{m,KIR}(v - E_K) \\
 & - \mu_{leak}(v - E_{leak}) + u
 \end{aligned}$$

Note the presence of the activation function $\sigma_{m,KIR}$, implying that the gating variable m_{KIR} is a static function of the voltage. Note also that the conductance of the current I_{KCa} depends on the calcium concentration $[Ca]$, rather than on the voltage. The calcium concentration is modelled as a low-pass filtered version of the voltage:

$$[\dot{Ca}] = -0.01m_L(V - E_{Ca}) - 0.0025[Ca] \quad (2.3)$$

The remaining intrinsic gating variables evolve according to equations of the type (2.2b)-(2.2c). Figure 2.2 illustrates the simulated behaviour of this model. The precise dynamics of each gating variable are specified in Appendix 2.A.

Recall that the neuron in Example 1 exhibits spiking due to the mixed-feedback motif, that is the presence of both fast-activating negative conductance and slower-activating positive conductance. The model in Example 2 exhibits robust bursting because this pairing is replicated at the timescale of the burst [35]. The T- and L- type calcium currents I_T and I_L provide negative conductance in the slower timescale of calcium activation, and even slower positive conductance is provided by the calcium-activated potassium current I_{KCa} . The importance of slow negative conductance for robust bursting is commonly overlooked [35].

Notice that the functional role of an ionic current is determined solely by whether it contributes positive or negative conductance at a particular timescale. This depends on two things: whether the current is inward or outward, given by the sign of $(v - E_{ion})$, and whether its gating variable at that timescale is activating or inactivating. Thus the role of an activating inward current, for example, is identical to that of an inactivating outward current. This is an essential point, and key to navigating the vast literature of conductance-based models [36].

Synaptic currents arise from electrochemical connections between neurons [31, Chapter 7]. They are modelled as current sources in the same way as intrinsic currents; the only difference is that the voltage dependence of their conductances is on the presynaptic neuron, noted v_p in the example below.

Example 3 *For the sake of illustration, all the conductance-based neurons in the rest of this chapter and in Chapter 3 use one type of (inhibitory) synapse, called a GABA synapse in neurophysiology, obeying the standard model reproduced from [25]:*

$$I_{syn,p} = \mu_{syn,p} s_{syn,p} (v - E_{syn}) \quad (2.4a)$$

$$\dot{s}_{syn,p} = 0.53 \sigma_{syn}(v_p) (1 - s_{syn,p}) - 0.18 s_{syn,p} \quad (2.4b)$$

with a synaptic activation function σ_{syn} given by a sigmoid function of the form

$$\sigma_{syn}(v_p) = \frac{1}{1 + \exp(-(v_p - 2)/5)}. \quad (2.5)$$

Here, $s_{syn,p}$ is the synaptic gating variable and v_p is the membrane voltage of the presynaptic neuron. The constants $\mu_{syn,p} > 0$ and $E_{syn} \in \mathbb{R}$ are (synaptic) maximal conductances and reversal potentials, respectively. This model can represent excitatory or inhibitory synapses, depending on the value of E_{syn} . In the present chapter we set $E_{syn} = -90$, which models the inhibitory synapses encountered in central pattern generators.

2.3.2 Conductance-based model of a neural network

A conductance-based neural network is given by the interconnection of $n_v \in \mathbb{N}$ single neurons. We denote the index set of neurons in the network by $\mathcal{N} := \{1, \dots, n_v\}$ and we describe the dynamics of the i^{th} neuron in the network by attaching an $i \in \mathcal{N}$ subscript to all the variables in (2.1)-(2.4), except for reversal potentials, time-constant functions, and activation functions.

The network state-space model gathers all neuronal membrane voltages in the vector

$$v = (v_1, \dots, v_{n_v})^\top,$$

and all other state variables in the vector

$$w = \text{col}(w^{(1)}, \dots, w^{(n_v)}),$$

where the vector $w^{(i)}$ collects the intrinsic and synaptic gating variables of the i^{th} neuron, that is, $m_{\text{ion},i}$, $h_{\text{ion},i}$, and $s_{\text{syn},p,i}$, as well as the calcium concentration $[Ca]_i$. Notice that with this notation, $\mathcal{P} \subseteq \mathcal{N}$.

In addition to synaptic current interconnections, neuronal networks may include electrical gap junctions, modelled as ohmic resistive currents flowing between neurons. Gap junction currents are thus passive components of the network, just as leak currents are passive components in a single neuron. It follows that the model of a conductance-based neural network has a form completely analogous to that of a single neuron, given by (2.1). The network model is given by the vector equations

$$c \dot{v} = -I_{\text{leak}} - I_{\text{int}} - I_{\text{ext}} + u \quad (2.6a)$$

$$\dot{w} = g(v, w) \quad (2.6b)$$

where I_{leak} is the overall leaky current, given by

$$I_{\text{leak}} = \begin{pmatrix} \mu_{\text{leak},1}(v_1 - E_{\text{leak}}) + \sum_{i \in \mathcal{N} \setminus \{1\}} \mu_{\text{gap},1,i}(v_1 - v_i) \\ \vdots \\ \mu_{\text{leak},n_v}(v_{n_v} - E_{\text{leak}}) + \sum_{i \in \mathcal{N} \setminus \{n_v\}} \mu_{\text{gap},n_v,i}(v_{n_v} - v_i) \end{pmatrix}$$

and $I_{\text{int}} \in \mathbb{R}^{n_v}$, $I_{\text{ext}} \in \mathbb{R}^{n_v}$ and $u \in \mathbb{R}^{n_v}$ are formed by gathering the intrinsic and extrinsic currents and inputs of each neuron in the corresponding n_v -dimensional vectors. The addition of gap junction currents extends the system-theoretic modelling framework of [16].

The corresponding vector function $g(v, w)$ in (2.6b), which represents the *internal dynamics* of the neural network, collects the dynamics of all gating variables and calcium concentrations in the network, respecting the order in which those variables were included in w . For each neuron in the network, these dynamics are easily found from (2.2b), (2.2c), (2.3) and (2.4b).

Example 4 *A Half-Centre Oscillator (HCO) is a network of two neurons mutually coupled by inhibitory synapses. This elementary network is the simplest example of a central pattern generator, a type of neural network capable of generating autonomous rhythms for motor control [69].*

The HCO model used in Chapter 3 interconnects two identical bursting neurons modelled according to Example 2 with the synapse model given in Example 3. The synapse index set is given by $\mathcal{S} = \{G\}$. The state of the internal dynamics is given by $w = \text{col}(w^{(1)}, w^{(2)})$, with

$$w^{(i)} = \text{col}(m_{Na,i}, h_{Na,i}, m_H,i, m_T,i, h_T,i, m_A,i, m_K,i, m_L,i, s_j, [Ca]_i)$$

The HCO model is a building block of more complex central pattern generators such as the example considered in Section 3.2.3, which includes gap junctions as well as synapses.

Biophysical central pattern generators such as the HCO of Example 4 have helped inspire a rich literature of nonlinear oscillating robotic controllers, founded for example on the Matsuoka [73] or Hopf [75] oscillators. See [51] for a review of this field.

2.4 A minimal neuromorphic circuit

Conductance-based models, while biologically faithful, are not well-suited to neuromorphic control. In particular, they are not intuitive to tune because an ionic current may have two gates but only has one tunable maximal conductance parameter. This means that it is not possible to independently control the gains of the feedback loops. By contrast, the work [89] introduces a bursting neuromorphic circuit that implements the mixed-feedback motif such that each current source corresponds to a single feedback loop. This permits a graphical approach to tuning the circuit behaviour by shaping the circuit's static current-voltage curves (one at each timescale). This tuning method is outlined below. Like the conductance-based models, the circuit is minimum phase, relative degree one, and linear in the controlled parameters. In chapter 5 we use a network of such neuron circuits to control the pendulum.

The neuron of [89] is an RC circuit in parallel with a bank of voltage-controlled current sources. Every current source provides either positive or negative conductance, at a particular

timescale. A positive conductance is a source of negative feedback, and vice versa. The gains of the current sources are the control parameters. Each neuron i has three state variables: a voltage v_i (the output) and two low-pass-filtered voltages $v_{s,i}$ and $v_{us,i}$ (respectively ‘slow’ and ‘ultra-slow’ voltages). The dynamics are:

$$\begin{aligned}\tau_f \dot{v}_i &= -v_i + g_f^- \tanh(v_i) - g_s^+ \tanh(v_{s,i}) \\ &\quad + g_s^- \tanh(v_{s,i} + 0.9) - g_{us}^+ \tanh(v_{us,i} + 0.9) \\ &\quad + I_{\text{syn},ij} + i_{\text{app},i} \\ &= -v_i - I_{f,i}^- - I_{s,i}^+ - I_{s,i}^- - I_{us,i}^+ + I_{\text{syn},ij} + i_{\text{app},i} \\ \tau_s \dot{v}_{s,i} &= v_i - v_{s,i} \\ \tau_{us} \dot{v}_{us,i} &= v_i - v_{us,i}\end{aligned}$$

for $i \in \{A1, B1, A2, B2\}$. The current $I_{f,i}^- = -g_f^- \tanh(v_i)$, for example, is a source of negative conductance for neuron i acting at the fast timescale. The timescales τ_f , τ_s and τ_{us} are chosen such that they are all sufficiently separated [89].

For the synapses we follow [88]: the synaptic current from neuron j to neuron i obeys

$$I_{\text{syn},ij} = g_{\text{syn},ij} / (1 + \exp(-2(v_{s,j} + 1)))$$

where $g_{\text{syn},ij}$ is positive for an excitatory synapse, and negative for an inhibitory one.

The neuron’s input is the applied current $i_{\text{app},i}$. All quantities are non-dimensional; see [89] for the circuit implementation.

Following the design procedure in [90], a neuron circuit can be manually tuned using its three current-voltage curves; these curves also demonstrate that the behaviour is quite robust to parameter uncertainty [89]. Each I-V curve is defined as

$$\mathcal{I}_x = V + \sum_{\tau_y \leq \tau_x} I_y^\pm(V) \quad (2.7)$$

where I_y^\pm is the equation of a particular current source, for example $I_f^- = -g_f^- \tanh(V)$ for the fast negative conductance. Notice that each curve includes all elements that act at the timescale of that curve, or faster.

To understand the intuition of the I-V curve approach, consider a reduced version of the neuron that includes only the current sources I_f^- and I_s^+ , that is the fast negative and slow positive conductances. Such a neuron is able to spike but not burst. The fast curve is then defined as $\mathcal{I}_f = V + I_f^-(V)$, where I_f^- is the fast negative conductance element. This conductance




Fig. 2.3 Fast and slow I-V curves of a spiking circuit. The fast curve is ‘N-shaped’. The slow curve is monotone and its intersection with the line $I = I_{app}$ determines the system’s equilibrium (V_e). The neuron spikes when $V_1^f < V_e < V_2^f$. On the slow curve we see two example values for the applied current, and the corresponding values for V_e . Figure and caption reproduced from [89].

creates a local range of negative gradient, and so a region of instability, at the fast timescale. The fast curve is therefore ‘N-shaped’, indicating a bistability at that timescale. The slow curve $\mathcal{I}_s = V + I_f^-(V) + I_s^+(V)$ is monotonically increasing everywhere ($d\mathcal{I}_s/dV \geq 0 \forall V$) due to the additional presence of the slow positive conductance. Spiking occurs when the circuit is unstable at the fast timescale but asymptotically stable at the slow timescale. When this is the case, the circuit is driven periodically ‘around’ the bistability, switching repeatedly between the two stable states. The fast bistability has two voltage thresholds V_1^f and V_2^f which define the unstable region, that is the region of negative gradient on the fast curve. If the circuit receives a constant applied current such that the corresponding equilibrium voltage V_e on the slow curve satisfies $V_1^f < V_e < V_2^f$, and assuming sufficient timescale separation, then the circuit exhibits stable spiking behaviour. Outside of this voltage range, the neuron converges to the equilibrium V_e . Fig. 2.3 shows the circuit’s two I-V curves.

To introduce bursting behaviour, the mixed-feedback motif of the spiking circuit is replicated at the slow and ultra-slow timescales. The addition of a slow negative conductance introduces a bistability to the slow curve, and the addition of an ultra-slow positive conductance restores monotonicity for the ultra-slow curve. The unstable region of the slow curve is defined by the voltage thresholds V_1^s and V_2^s . The slow and fast unstable ranges are designed to overlap such that $V_1^s < V_1^f < V_2^s < V_2^f$. This ensures that the ‘up’ state of the slow bistability corresponds to

the unstable region of the fast curve, and that the slow threshold V_1^s is at a lower voltage than the fast one. As a result there is a region of ‘rest-spike’ bistability in which the circuit bursts, that is, it switches periodically between resting and spiking. Fig. 2.4 shows the corresponding I-V curves.

Neuromodulation with the I-V curves is performed qualitatively by varying the gains g_x^\pm of the conductance elements and thus of the feedback loops. For example, increasing the gains of the slow and ultra-slow positive conductances increases the spiking and bursting frequencies respectively. A key advantage of this approach is that the tuning bounds are evident from the curves. For example, g_f^- can be tuned within the range of values that yields a region of negative gradient on the fast curve. For further details, the interested reader is referred to [89].

2.5 Rhythmic neuronal networks

Having introduced biophysical and minimal models of neuron circuits, we now turn to small rhythmic networks of neurons. Such networks, known as “central pattern generators” (CPGs), are endogenous sources of rhythm and in biology are responsible for directing motor neurons to generate rhythmic movements such as walking or breathing [69].

The simplest CPG is the half-centre oscillator (HCO). The HCO, for which an example was given in Section 2.3, has a long history in the neuroscience literature [9]. It is made up of two reciprocally-inhibiting neurons that emit bursts of spikes in anti-phase, as shown in Fig. 2.5. Observe that the nature of the synapse dictates the phase relationship between the neurons: here, inhibitory synapses ensure an anti-phase behaviour. The HCO rhythm can be switched on or off by neuromodulation [21].

The HCO is the building block of many motor patterns. For example, a chain of HCOs connected in series generates the undulatory swimming motion of the lamprey and salamander, a fact that has inspired the robot of [53]. In Chapter 5 of this thesis, a pair of HCOs is used to actuate two motors for the control of a pendulum, with one HCO controlling each motor. These HCOs (and therefore the motors) act in phase when connected by excitatory synapses, and in anti-phase when connected by inhibitory ones. This distributed architecture can scale to large numbers of actuators, and so may have wide applications for network control problems.

We can also connect two HCOs, one fast and one slow, via a central ‘hub’ neuron to obtain a five-neuron network. The work [26] demonstrates that manual neuromodulation can be used to determine whether the hub neuron expresses the fast rhythm, the slow rhythm, or a mix of the two. In Section 3.2.3 of this thesis we solve the same task using adaptive control instead of manual tuning. In order to adaptively control the parameters of a neuronal network, we must

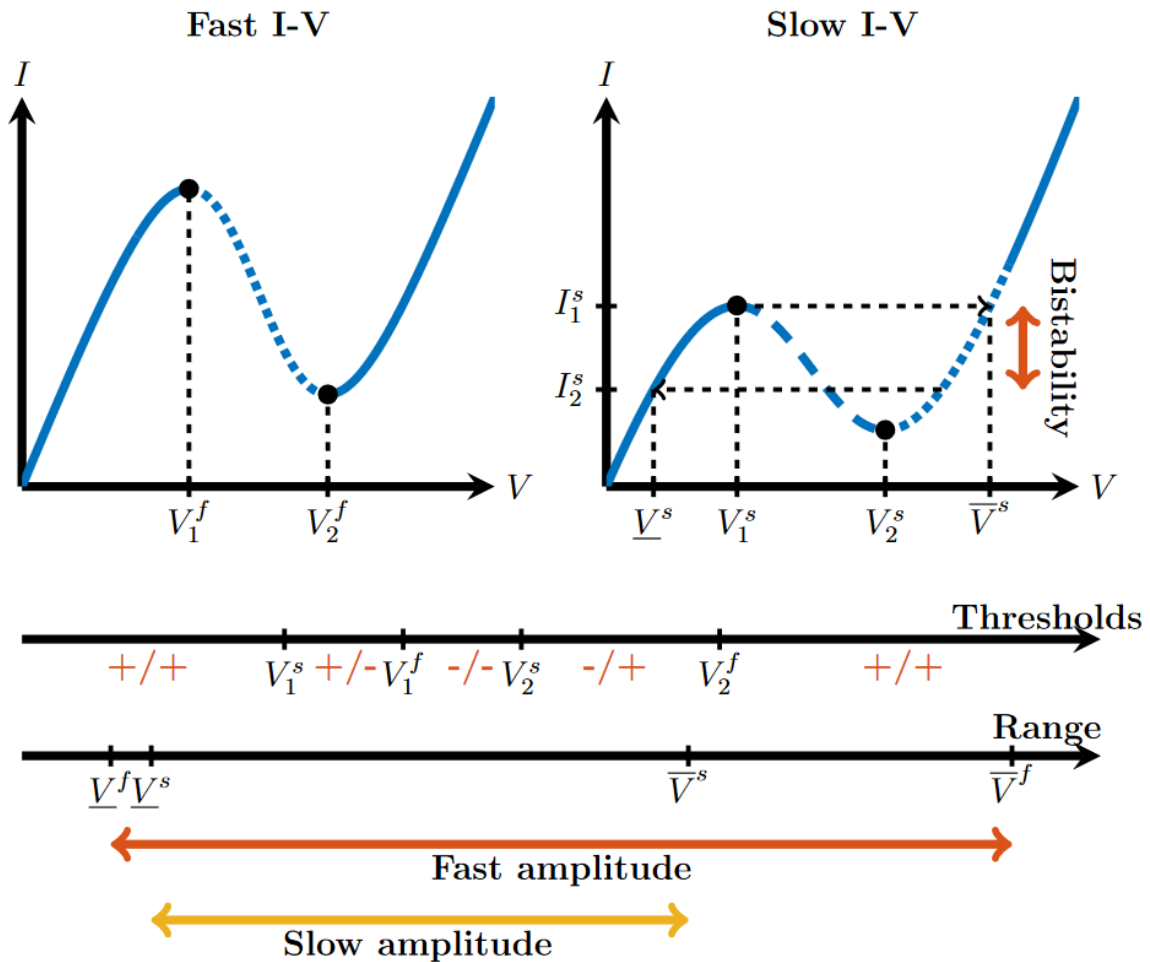


Fig. 2.4 The I-V curves (top) show a system with a double hysteresis: both the fast and the slow I-V curves are ‘N-shaped’. By having the ‘up’ state of the slow curve correspond to the unstable region of the fast system, the system experiences rest-spike bistability, given that the slow threshold is at lower voltage than the fast one, i.e. $V_1^s < V_1^f$. The system has two pairs of threshold and range voltages corresponding to the fast and the slow I-V curves. These variables characterise the circuit’s behaviour, and are adjusted by control of the gains g_x^\pm . The ultraslow curve (not shown) is monotonically increasing. Figure and caption reproduced from [89].




Fig. 2.5 The behaviour of an HCO circuit following the model of Section 2.4. *Left*: the network diagram, showing the two neurons and the inhibitory synapses between them. *Right*: the output behaviour of each neuron. Reproduced from [88].

typically first estimate the current values of those parameters. We do so using the adaptive observer of [16], introduced below. This observer is used in Chapters 3 and 4 to solve control tasks for conductance-based neurons. A form of adaptive control that does not rely on the observer is then used in Chapter 5 as part of the neuromorphic pendulum controller.

2.6 Adaptive observers for conductance-based models

Following the notation of [16], the conductance-based network model (2.6) can be written in the form

$$\dot{v} = \Phi(v, w, u)\theta + b(v, w, u) \quad (2.8a)$$

$$\dot{w} = g(v, w) \quad (2.8b)$$

$$\dot{\theta} = 0 \quad (2.8c)$$

where $v(t) \in \mathbb{R}^{n_v}$ is a vector of measured output membrane voltages, $u(t) \in \mathbb{R}^{n_u}$ is a vector of input currents, $w(t) \in \mathbb{R}^{n_w}$ is a vector of unmeasured internal states, and $\theta \in \mathbb{R}^{n_\theta}$ is a vector of biophysical parameters to be estimated. The vector of unknown constant parameters θ is included in the state-space model via (2.8c).

In this thesis the vector θ only includes maximal conductances.

Example 5 Consider the HCO described in Example 4. Let

$$\mu^{(i)} = (\mu_{Na,i}, \mu_{H,i}, \mu_{T,i}, \mu_{A,i}, \mu_{K,i}, \mu_{L,i}, \mu_{KCa,i}, \mu_{KIR,i}, \mu_{\text{syn},p,i}, \mu_{\text{leak},i})^T \quad (2.9)$$

for $i, p \in \mathcal{N} = \{1, 2\}$ and $p \neq i$. The estimated parameters of the HCO are chosen as

$$\theta = \text{col}(\mu^{(1)}, \mu^{(2)}) \quad (2.10)$$

with $\mu^{(1)}$ and $\mu^{(2)}$ given by (2.9). Letting $v = (v_1, v_2)^\top$ and $w = \text{col}(w^{(1)}, w^{(2)})$, the voltage dynamics of the model can then be written as (2.8a), where

$$\Phi(v, w) = \begin{bmatrix} \varphi(v_1, w^{(1)}) & 0 \\ 0 & \varphi(v_2, w^{(2)}) \end{bmatrix}$$

$$b(t) = (u_1(t)/c_1, u_2(t)/c_2)^\top$$

with

$$\varphi(v_i, w^{(i)}) = -\frac{1}{c_i} \begin{pmatrix} m_{Na,i}^3 h_{Na,i}(v_i - E_{Na}) \\ m_{H,i}(v_i - E_H) \\ m_{T,i}^2 h_{T,i}(v_i - E_{Ca}) \\ m_{A,i}^4 h_{A,i}(v_i - E_K) \\ m_{K,i}^4 (v_i - E_K) \\ m_{L,i}(v_i - E_{Ca}) \\ \left(\frac{[Ca]_i}{15 + [Ca]_i}\right)^4 (v_i - E_K) \\ \sigma_{m,KIR}(v_i)(v_i - E_K) \\ s_{\text{syn},p,i}(v_i - E_G) \\ v_i - E_{\text{leak}} \end{pmatrix}^\top$$

for $i = 1, 2$ and $p \neq i$.

An important property of the parametrisation in Example 5 is that it is decentralised: the network estimation problem decouples into independent single-neuron estimation problems. This decoupling allows the estimation problem to be scaled to a possibly high-dimensional network.

The recent work [16] provides an adaptive observer to estimate the parameters of the system (2.8) in real-time. This observer has global convergence properties and is based on the recursive

least squares algorithm. The state-space realisation of the observer is

$$\dot{\hat{v}} = \Phi(v, \hat{w}, u)\hat{\theta} + b(v, \hat{w}, u) + \gamma(I + \Psi P \Psi^\top)(v - \hat{v}) \quad (2.11a)$$

$$\dot{\hat{w}} = g(v, \hat{w}) \quad (2.11b)$$

$$\dot{\hat{\theta}} = \gamma P \Psi^\top (v - \hat{v}) \quad (2.11c)$$

where $\gamma > 0$ is a constant gain, and the matrices P and Ψ evolve according to

$$\dot{\Psi} = -\gamma\Psi + \Phi(v, \hat{w}, u), \quad (2.12a)$$

$$\dot{P} = \alpha P - \eta P \Psi^\top \Psi P, \quad P(0) \succ 0 \quad (2.12b)$$

where $\alpha > 0$ is a constant forgetting rate, required to discount the initial error between $w(0)$ and $\hat{w}(0)$. We set $\eta = 1$ in Chapter 3 and $\eta = \gamma$ in Chapter 4; they are equivalent.¹ It can be shown that this adaptive observer recursively solves a least-squares regression problem with exponential forgetting, where the regression error is defined by filtering the derivative of v with the first-order filter $H = \gamma/(s + \gamma)$. Notice that Ψ is a low-pass filtered version of Φ , while P can be interpreted as a running estimate (with a forgetting factor) of the parameter covariance matrix [7, Chapter 2]. Without loss of generality, we assume $\Psi(0) = 0$.

The convergence of the above adaptive observer relies on the key property that the internal dynamics (2.8b) are contracting in w on a positively invariant compact set, uniformly in the voltage v . This property is also sufficient to establish that the neuron is minimum phase. We refer the reader to [16] for further details, including a contraction-based proof of convergence.

¹To replace $\eta = 1$ with $\eta = \gamma$, replace $P(t)$ with $\gamma P(t)$ [12].

Appendix

Appendix 2.A Gating Dynamics

All neurons share the same gating variables (2.2). The activation and time-constant functions for each gate are provided below. First note that some of these functions have been defined in terms of functions $\alpha_{\text{ion}}(v)$ and $\beta_{\text{ion}}(v)$:

$$\alpha_{m_{\text{Na}}} = \frac{-0.025(v+40)}{\exp(-(v+40)/10) - 1}$$

$$\alpha_{h_{\text{Na}}} = 0.0175 \exp(-(v+65)/20)$$

$$\alpha_{m_{\text{K}}} = \frac{0.0025(v+55)}{1 - \exp(-(v+55)/10)}$$

$$\alpha_{m_{\text{H}}} = \exp(-14.59 - 0.086v)$$

$$\beta_{m_{\text{Na}}} = \exp(-(v+65)/18)$$

$$\beta_{h_{\text{Na}}} = \frac{0.25}{1 + \exp(-(v+35)/10)}$$

$$\beta_{m_{\text{K}}} = 0.03125 \exp(-(v+65)/80)$$

$$\beta_{m_{\text{H}}} = \exp(-1.87 + 0.0701v)$$

The activation functions are as follows:

$$\sigma_{m,\text{Na}} = \frac{\alpha_{m_{\text{Na}}}(v)}{\alpha_{m_{\text{Na}}}(v) + \beta_{m_{\text{Na}}}(v)}$$

$$\sigma_{m,\text{K}} = \frac{\alpha_{m_{\text{K}}}(v-10)}{\alpha_{m_{\text{K}}}(v-10) + \beta_{m_{\text{K}}}(v-10)}$$

$$\sigma_{m,\text{A}} = \frac{1}{1 + \exp(-(v+90)/8.5)}$$

$$\sigma_{m,\text{T}} = \frac{1}{1 + \exp(-(v+57)/6.2)}$$

$$\sigma_{m,\text{L}} = \frac{1}{1 + \exp(-(v+55)/3)}$$

$$\sigma_{h,\text{Na}} = \frac{\alpha_{h_{\text{Na}}}(v)}{\alpha_{h_{\text{Na}}}(v) + \beta_{h_{\text{Na}}}(v)}$$

$$\sigma_{m,\text{H}} = \frac{\alpha_{m_{\text{H}}}(v)}{\alpha_{m_{\text{H}}}(v) + \beta_{m_{\text{H}}}(v)}$$

$$\sigma_{h,\text{A}} = \frac{1}{1 + \exp((v+78)/6)}$$

$$\sigma_{h,\text{T}} = \frac{1}{1 + \exp((v+81)/4.03)}$$

$$\sigma_{m,\text{KIR}} = \frac{1}{1 + \exp((v+107.9)/9.7)}$$

The time-constant functions are as follows:

$$\begin{aligned}\tau_{m,\text{Na}} &= \frac{1}{0.2(\alpha_{m_{\text{Na}}}(v) + \beta_{m_{\text{Na}}}(v))} & \tau_{h,\text{Na}} &= \frac{1}{0.2(\alpha_{h_{\text{Na}}}(v) + \beta_{h_{\text{Na}}}(v))} \\ \tau_{m,\text{K}} &= \frac{1}{0.2(\alpha_{m_{\text{K}}}(v - 10) + \beta_{m_{\text{K}}}(v - 10))} & \tau_{m,\text{H}} &= \frac{1}{\alpha_{m_{\text{H}}}(v) + \beta_{m_{\text{H}}}(v)} \\ \tau_{m,\text{L}} &= 72 \exp(-(v + 45)^2/400) + 6\end{aligned}$$

$$\begin{aligned}\tau_{m,\text{T}} &= 0.612 + \frac{1}{\exp(-(v + 131.6)/16.7) + \exp((v + 16.8)/18.2)} \\ \tau_{m,\text{A}} &= 0.37 + \frac{1}{0.2(\exp((v + 35.82)/19.697) + \exp((v + 79.69)/-12.7))}\end{aligned}$$

We also have

$$\tau_{h,\text{A}} = \begin{cases} \frac{1}{0.2(\exp((v+46.05)/5) + \exp((v+238.4)/-37.45))} & \text{if } v < -63, \\ 19 & \text{otherwise} \end{cases}$$

and

$$\tau_{h,\text{T}} = \begin{cases} \exp((v + 467)/66.6) & \text{if } v < -80, \\ \exp(-(v + 21.88)/10.2) + 28 & \text{otherwise.} \end{cases}$$

Chapter 3

Adaptive conductance control

3.1 Introduction

Due to the very nature of nervous electrical signals, brain medicine and brain-inspired computing put an increasing demand on completely novel control design methodologies. Neurotechnology interfaces neural tissues with electronic devices via the exchange of electrical signals. While of the same physical nature, those signals are radically different in animals and in machines: neural signals are analogue and spiky, whereas the information processing of actuator and sensor signals is digital and quantised. Current control systems in neurotechnology use classical linear filtering techniques in combination with conventional analogue-to-digital converters from robotics and electronics [108]. They fail to acknowledge the excitable dynamics underlying the generation of spiky signals. Future developments will call for more compliance and more integration between the biological and technological domains. Such requirements will impose control interfaces that interconnect (or, in the behavioural language of Willems [118], *share*) signals of the same nature.

The significance of designing a controller that shares the input-output properties of the controlled physical system has a long and rich history in robotics. Passivity-based control is rooted in the concept of *port interconnection*, which wires the physical terminals of a passive plant and a passive controller both modelled as mechanical ports [82]. In robotics, impedance control is rooted in the design of a controller that shapes the mechanical impedance of the closed-loop system to comply with the environment, itself modelled in terms of a mechanical impedance [47].

The present chapter adopts the same philosophy for the control of a neuronal system. We use the framework of conductance-based modelling both for the controller and for the system to be controlled. We consider conductance-based neural networks in which each neural node is

a one-port circuit composed of one leaky capacitor (the ‘passive membrane’) in parallel with a bank of ohmic current sources of variable conductance. The controller is itself an additional set of ohmic current sources connected in parallel to those of the neuron. Each conductance is voltage-dependent, gating the current flow in a specific temporal and amplitude window.

A key emphasis in the present chapter is on the adaptive control of the maximal conductances of a conductance-based model. Each maximal conductance modulates the relative importance of a specific current source. Similar to the parameters of a conventional controller, the maximal conductances shape the total conductance of the controlled neurons. We wish to demonstrate that the online adaptation of maximal conductances provides a versatile framework to control the behaviour of a neuronal system. Online adaptation of the maximal conductances is aligned with the concept of *neuromodulation*, which as we have seen is of key importance in the biological control of neuronal systems [71, 104, 28]. All nervous systems are subject to neuromodulation, and each neuron is potentially the target of multiple neuromodulators [68]. Furthermore, each modulator can act on multiple ionic currents in the same neuron [72].

Our methodology exploits the classical framework of model reference adaptive control [6]. Recall that this framework relies on key physical properties of electrical and mechanical circuits: a relative degree one between the two terminal variables of the port (current and voltage in the present chapter), and contracting internal dynamics. We show that these properties can be exploited in conductance-based electrical circuits in the same way as they have been exploited in the impedance-based mechanical circuits of robotics. A core element of the proposed adaptive control design is the adaptive observer recently proposed for real-time estimation of conductance-based models in [16] and described above in Section 2.6. We assume in this chapter that the observer’s model (2.11b) of the internal dynamics is a perfect copy of the true dynamics (2.8b). This assumption will be relaxed in Chapter 4.

The chapter is organised as follows. Section 3.2 employs the adaptive observer, introduced above in Section 2.6, to solve the basic problems of adaptive reference tracking and adaptive disturbance rejection, as well as showing the relevance of such problems in the control of a simple biophysical neural network. Section 3.3 discusses the idealised assumptions of the present chapter and possible routes to make the theoretical methodology amenable to practical solutions of control problems in electrophysiology or in neuromorphic engineering.

3.2 Model reference adaptive conductance control

Adaptive observers are instrumental to the classical design approach called *model reference adaptive control (MRAC)* [6]. In this section we illustrate the application of model refer-

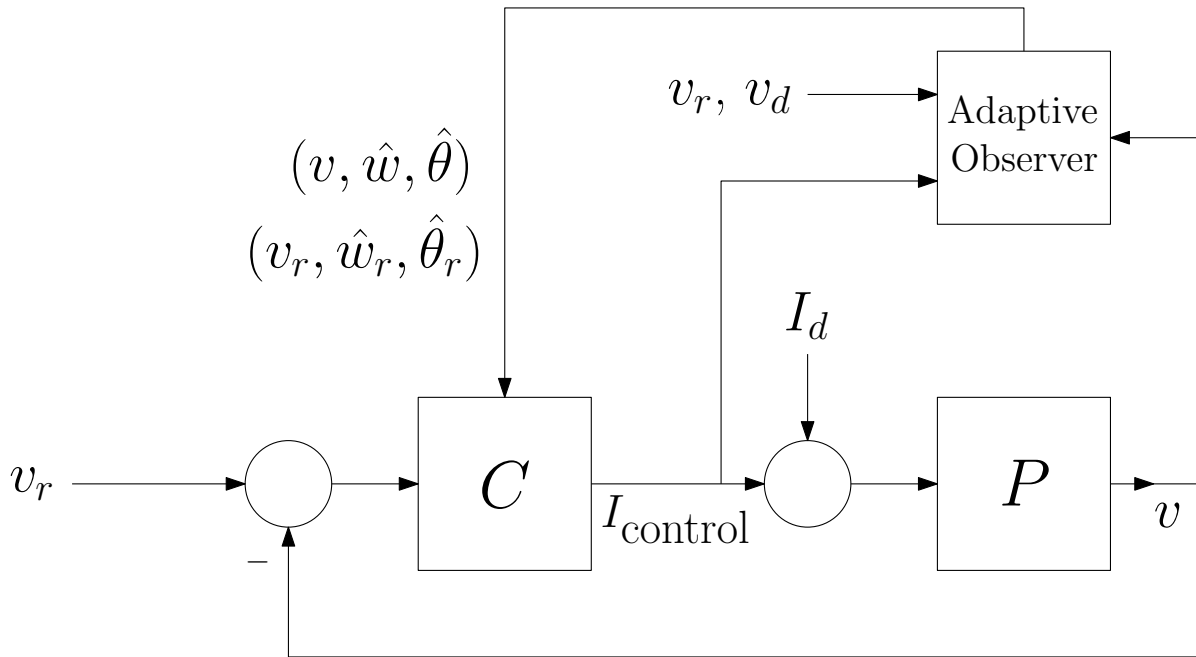


Fig. 3.2.1 Block diagram representation of the adaptive reference tracking and disturbance rejection problems. Adapted from [6, Chapter 1].

ence adaptive control to conductance-based models. We regard a single neuron as a voltage-controlled circuit. We review the two canonical control problems of adaptive control: the adaptive tracking of a reference signal v_r (Section 3.2.1), and the adaptive disturbance rejection of an external current I_d (Section 3.2.2). We then illustrate the relevance of those elementary control problems in a network example (Section 3.2.3). See Figure 3.2.1 for a block diagram representation of the two problems. Convergence to the desired solution is guaranteed in each case as we are assuming a perfect knowledge of the ion channel kinetics $g(v, w)$, an assumption that we relax in the next chapter. Note that convergence occurs regardless of the initial conditions, as they decay due to the contracting internal dynamics.

The code used to generate the simulations of this chapter can be found at the attached link.¹

3.2.1 Adaptive reference tracking

The classical problem of reference tracking is to design a controller such that the voltage output v asymptotically converges to the voltage reference v_r . If we assume that both the output v and the reference v_r are solutions of identical conductance-based models, this tracking problem reduces to the classical problem of *synchronisation*.

¹<https://github.com/RJZS/adaptive-conductance-control>

Here we consider the *adaptive* version of the tracking problem: we assume that the reference generator is a conductance-based model with constant but unknown vector parameter θ_r and that the controlled circuit is a conductance-based model with the same model structure, and with a parameter θ that is also unknown. The problem of adaptive synchronisation has been considered in [33, 34], in the context of secure communication via encryption in chaotic reference generators. Assuming a perfect knowledge of the ion channel kinetics $g(v, w)$, the solution presented here is relevant to track the neuromodulation of a neuron in vivo or to learn the parameters of a silicon neuron from experimental traces.

We will first describe how to solve adaptive reference tracking without synchronisation. This yields an oscillation which follows the reference but with a possible phase shift. The phase shift is then eliminated with an additional resistive element in the controller.

We solve the adaptive reference tracking problem by estimating both the reference model and the controlled model with the observer design of Section 2.6. The observer of the reference generator provides an estimate \hat{x}_r of the reference state vector $x_r = (v_r, w_r, \theta_r)$ whereas the observer of the controlled model provides an estimate \hat{x} of the controlled model state vector $x = (v, w, \theta)$. We use the control law $u(t) = I_{\text{track}}(t) + u_r(t)$, where $u_r(t)$ is the input to the reference neuron, and

$$I_{\text{track}}(t) = \Phi(v, \hat{w})(\underline{\rho}(\hat{\theta}_r) - \bar{\rho}(\hat{\theta})), \quad (3.1)$$

where $\underline{\rho}(x) := \max(0, x)$ is the rectified linear (ReLU) function, and $\bar{\rho}(x) := \min(x, \beta)$. Both functions are applied to their arguments elementwise. Together, they ensure that the solutions of the controlled neuron remain in a positively invariant set, as required for convergence of the observer [16]. We require β to satisfy $\max_j \{\theta_j\} \leq \beta \leq \bar{\beta}$, where $\bar{\beta}$ is the largest value which preserves the set. We empirically choose β by setting it to a large value (relative to plausible values of θ) and reducing it if the membrane potential of the controlled neuron diverges.

Exponential convergence of the estimated parameters ($\hat{\theta}$ and $\hat{\theta}_r$) to the true parameters provides a solution to the tracking problem without synchronisation. The proof relies on the *virtual system* idea of contraction theory [66], and follows the same lines as in [16, Section III.B]: the estimate \hat{x}_r contracts exponentially fast to the reference x_r while the estimate \hat{x} contracts exponentially fast to x . Upon convergence of both the reference and plant observers, we obtain the non-adaptive synchronisation problem between two identical systems.

The solution to this problem is particularly simple for conductance-based models because of their property of output feedback contraction or output feedback incremental passivity [116, 109]. Contraction of the error dynamics is ensured by including the output feedback term $\kappa(v_r - v)$ in the control law, for a sufficiently large gain $\kappa > 0$. The circuit realisation of this feedback term is a resistive wire between the reference and controlled circuit, or a *gap* junction




Fig. 3.2.2 Circuit diagram of the adaptive reference tracking problem.

in the language of neurophysiology. When this wire is introduced, Theorem 2 of [116] applies, and $v(t) \rightarrow v_r(t)$ as $t \rightarrow \infty$. The full control law is thus given by

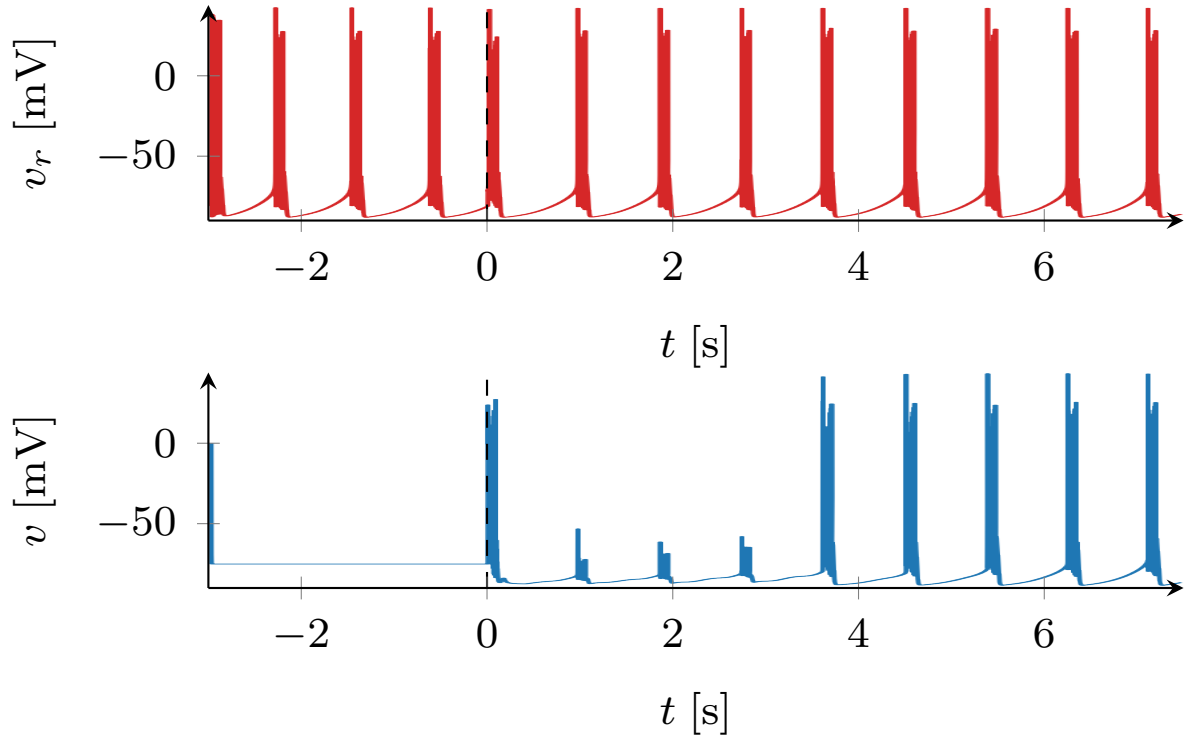
$$u(t) = I_{\text{control}}(t) + u_r(t) = I_{\text{track}}(t) + \kappa(v_r(t) - v(t)) + u_r(t), \quad (3.2)$$

where $I_{\text{track}}(t)$ is given by (3.1).


Figure 3.2.3 illustrates the performance of the adaptive conductance control in a scenario where both the reference and controlled neurons are the bursting neurons of Example 2. Fig. 3.2.3a compares v_r with v , while Fig. 3.2.3b shows the voltage tracking error $v_r - v$. To simplify the exposition, the parameters of the controlled neuron are set at $\theta = 0$, meaning that the 'open-loop' controlled neuron is the model of a passive membrane. The control scheme is shown in Figure 3.2.2. The observers and the controller are switched on at time $t = 0$, after which the parameters of the controlled neuron converge to the parameter values of the reference neuron, and so the voltage behaviour of the controlled neuron converges to that of the reference as well. The presence of the resistive wire ensures that there is no phase difference between the two voltage signals. See Appendix 3.A for the parameters used.

3.2.2 Adaptive disturbance rejection

The classical problem of disturbance rejection has a solution similar to that of the tracking problem. Given a disturbance current I_d generated by a synaptic current, we wish to design a feedback controller that asymptotically rejects that disturbance. A classical solution of the disturbance rejection problem is to include a model of the disturbance in the controller.



(a) *Top* (resp. *bottom*): v_r (resp. v) before and during the running of the observers and controller. The observers, controller and coupling are introduced at $t = 0$.



(b) Reference tracking error, including for $t < 0$ when there is no observer, controller or coupling.

Fig. 3.2.3 Simulation of single-neuron reference tracking.

This problem is of relevance in electrophysiology. Electrophysiologists study the properties of a given neuronal circuit *in vitro* by extracting the circuit from its nervous system and probing its responses to electrical stimuli in an experimental preparation, see e.g. [86]. Classical solutions include pharmacological agents that block specific types of ion channels, thereby reducing the synaptic currents to zero; see e.g. [42]. A downside of using pharmacological agents is that they may affect other properties of the circuit; another downside is the global effect of such agents within a given preparation. These may be undesirable when highly specific ion channel blocking is required.

An interesting alternative for targeted synaptic isolation in an experimental preparation is the design of a conductance-based controller for the rejection of synaptic currents. In this case, synaptic currents to be blocked are regarded as disturbances to be rejected.

Assuming that the pre- and post-synaptic voltages v_d and v are measured, a target synaptic current flowing between the two neuronal membranes can be blocked by means of the adaptive disturbance rejection control scheme shown in Figure 3.2.4. Note that this is the same circuit as in Figure 2.1, but with two additional circuit elements connected in parallel.

The first additional element is the disturbance I_d , which is interpreted as the specific synaptic current to be blocked. This disturbance current is modelled as

$$I_d = -\mu_{\text{syn}}s(v_d)(v - E_{\text{syn}}) \quad (3.3a)$$

$$\dot{s} = a_1\sigma_s(v_d)(1 - s) - a_2s, \quad (3.3b)$$

where σ_s is a monotonically increasing activation function and $a_1, a_2 > 0$ are constant (known) synaptic parameters.

The second additional element is the controller. Its inputs are the measured voltages v and v_d , and it generates a control current I_{control} which is designed to cancel I_d . We require that the behaviour of the closed-loop circuit converges to that of an undisturbed conductance-based model, that is, one where the targeted synaptic connection is absent. We assume that both the pre- and post-synaptic neurons are the bursting neurons of Example 2, and that they are interconnected with the inhibitory synapse of Example 3.

To achieve perfect disturbance rejection, the unknown synaptic maximal conductance μ_{syn} has to be estimated; this estimate is denoted by $\hat{\mu}_{\text{syn}}$. The disturbance rejection controller can then be designed following the *certainty equivalence principle* [6]. In other words, using the estimate $\hat{\mu}_{\text{syn}}$, the controller is designed to perfectly cancel I_d when $\hat{\mu}_{\text{syn}} = \mu_{\text{syn}}$. In this case, the estimation problem therefore requires an estimation method such that $\hat{\mu}_{\text{syn}} \rightarrow \mu_{\text{syn}}$ as $t \rightarrow \infty$. This can be accomplished by the observer given by (2.11)-(2.12), using the parametrisation

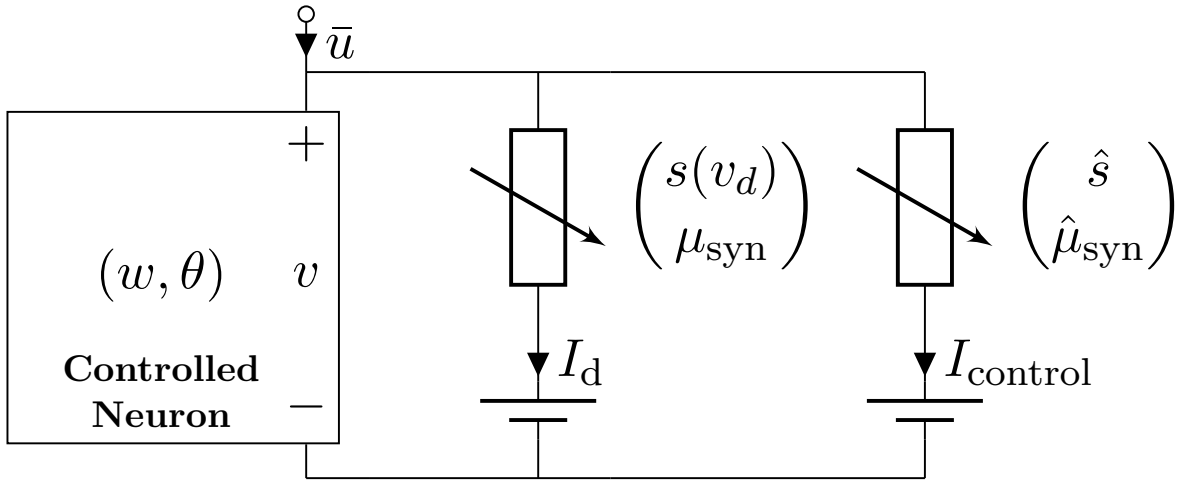


Fig. 3.2.4 Circuit diagram of the disturbance rejection problem, including the required controller.

$\hat{\theta} = \hat{\mu}_{\text{syn}}$. The observer also produces an estimate of the synaptic gating variable such that $\hat{s} \rightarrow s$ as $t \rightarrow \infty$. Given this observer, the input to the neuron is given by $u = I_{\text{control}} + I_d + \bar{u}$, where \bar{u} is an arbitrary new input signal, and

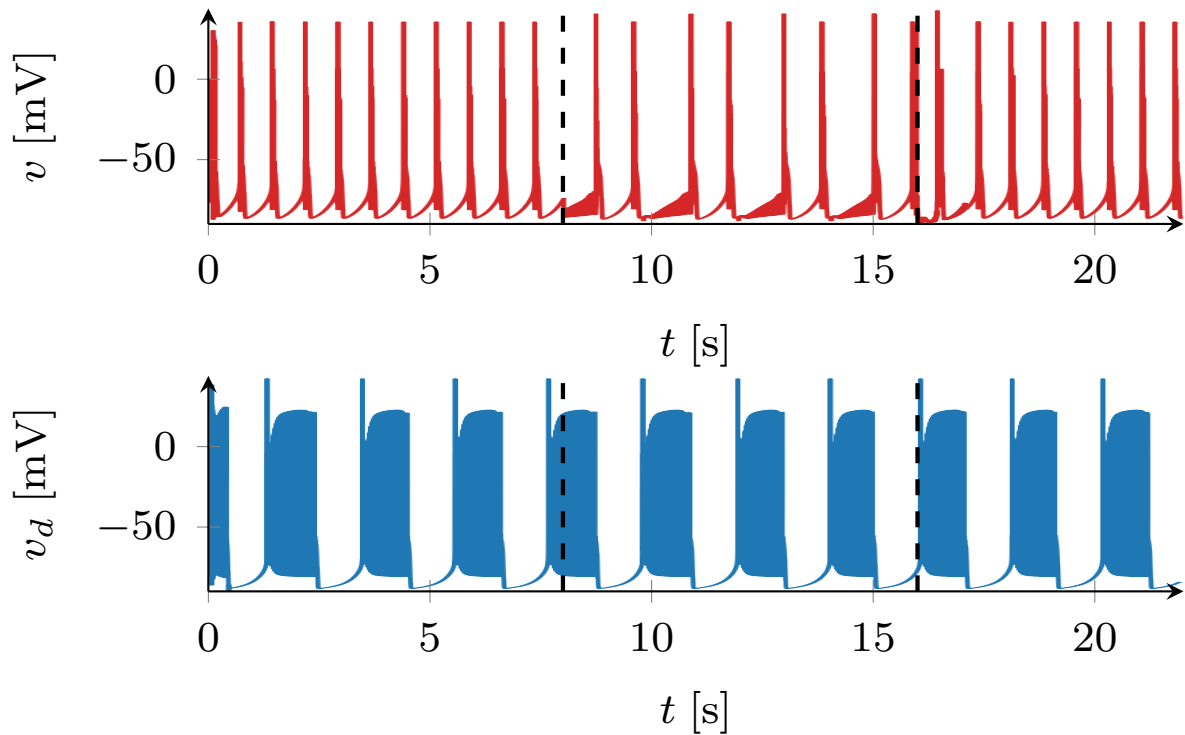
$$I_{\text{control}} = -\hat{I}_d \quad (3.4a)$$

$$\hat{I}_d = -\bar{\rho}(\hat{\mu}_{\text{syn}})\hat{s}(v - E_{\text{syn}}) \quad (3.4b)$$


$$\dot{\hat{s}} = a_1 \sigma_s(v_d)(1 - \hat{s}) - a_2 \hat{s}, \quad (3.4c)$$

where σ_s , a_1 and a_2 are the same as in (3.3b). As in Section 3.2.1, the function $\bar{\rho}$ ensures the positively invariant set.

Figure 3.2.5 illustrates the performance of the disturbance rejection controller. The disturbance is introduced at $t = 8$, and the controller is activated at $t = 16$. Fig. 3.2.5a compares v and v_d , showing the impact on the former of introducing and then rejecting the disturbance. Fig. 3.2.5b plots the disturbance I_d as well as the estimation error $I_d - \hat{I}_d$, which converges to zero. Note the differing ranges for the x-axes. It is striking that I_d is larger after the controller is activated; this is simply because the current is a function of both v and v_d , and these two voltages overlap once the inhibitory disturbance is rejected. The model and observer parameters, and the input currents, are provided in Appendix 3.A.



(a) *Top* (resp. *bottom*): v (resp. v_d) during three stages of the experiment. For $t < 8$, the neuron is undisturbed (the synapse is not present). For $8 < t < 16$, the synapse is present but there is neither observer nor controller. The observer and controller are introduced at $t = 16$.



(b) *Top*: True value of the disturbance synaptic current, I_d . This is larger when $t > 16$ as the bursts of the two neurons have more overlap (I_d is a function of both v and v_d). *Bottom*: The observer's error in estimating I_d , defined by the difference between the true value and the estimate \hat{I}_d . The plot is truncated once the estimate has converged to the true value.

Fig. 3.2.5 Simulation of single-neuron disturbance rejection.


3.2.3 Network neuromodulation

The adaptive conductance controller developed in the previous sections for a single cell is by nature decentralised: it can be applied independently to different neurons (nodes) in a network. Each neuron in the network receives synaptic currents that can be adaptively estimated using the measurements of the presynaptic neuron voltages. As a consequence, an observer can be designed for each neuron of the network, and each of those neurons can be controlled to synchronise to a given neuron or to adaptively reject specific synaptic currents.

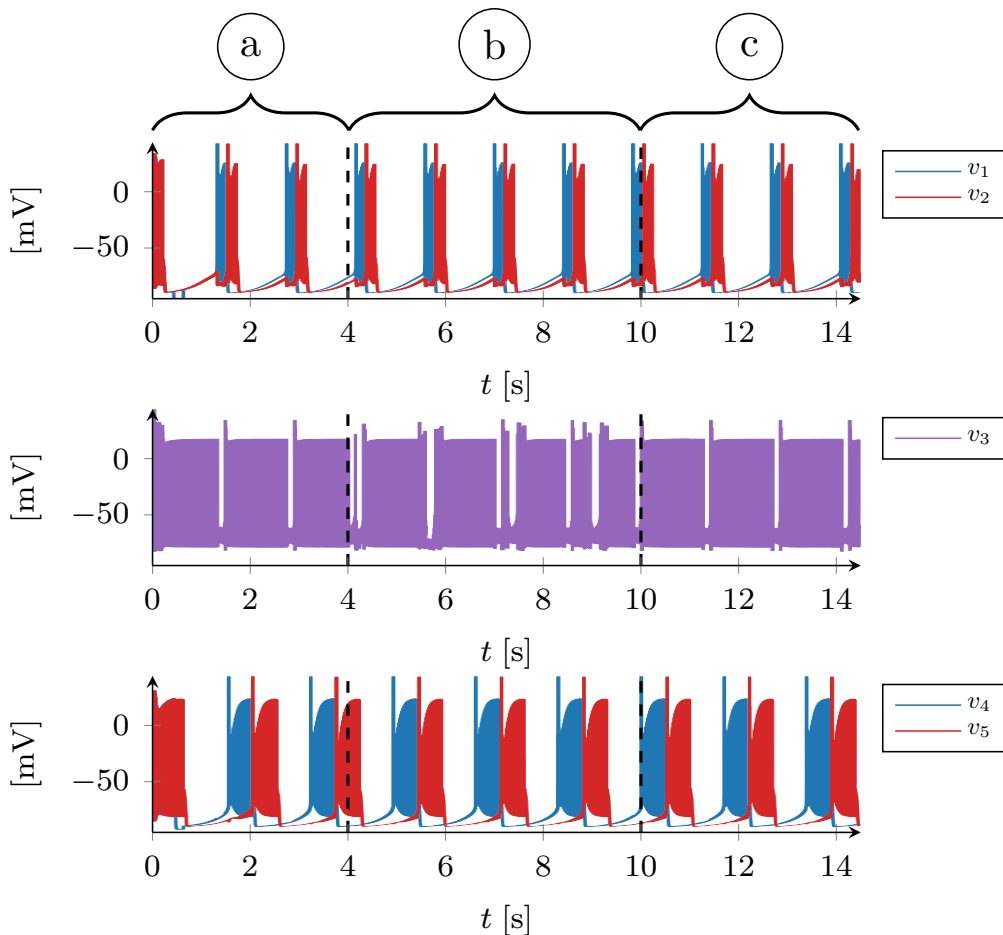
We illustrate the versatility of this adaptive conductance control in a five-neuron network previously analysed in [26] and itself inspired by the Stomatogastric Ganglion, a crustacean central pattern generator [68]. The network interconnects a fast HCO and a slow HCO, both with the model structure of Examples 4 and 5, through a central "hub" neuron. The connectivity diagram of this network is shown in the centre of Figure 3.2.6a. Notice the lack of direct connections between the two HCOs on either side of the hub neuron.

Previous work [26] has shown that this network can be switched between distinct rhythmic states by the modulation of specific internal conductances. In every possible rhythmic configuration, the network is composed of five neurons generated with the model of Example 2 and interconnected with gap junctions and the inhibitory synapses of Example 3. In an application of model-reference adaptive control, the challenge is to adaptively regulate the network by only modulating the maximal conductance parameters. This can involve up to five distinct observers, assuming measurement of the five neuronal voltages.

As a simple illustration, we show how to decouple the two central pattern generators by disconnecting the central hub using the disturbance rejection controller of the previous section. This is achieved by control of the hub (Neuron 3 in Figure 3.2.6a). The controller rejects the inhibitory synapse from Neuron 5, using the same control law as Section 3.2.2. This is illustrated in Figure 3.2.6b by the behaviour of the hub's membrane potential, v_3 (middle plot). During part (a) of the simulation, v_3 expresses a rhythm governed almost entirely by the first HCO (top plot, that is, Neurons 1 and 2). As only Neuron 1 is inhibiting the hub, v_3 is low when v_1 is active and each burst of v_3 is the same length. In part (b), when the disturbance is introduced, v_3 expresses a 'mixed' rhythm governed equally by both HCOs. Bursts of v_3 are interrupted whenever v_1 or v_5 is active. Finally, the observer and controller are introduced in part (c), and v_3 converges to the undisturbed rhythm of (a). The simulation parameters are provided in Appendix 3.A.



(a) The network simulation is divided into three phases: (a) The disturbance (defined as the synapse from Neuron 5 to Neuron 3) is not present. (b) The disturbance is introduced. (c) The observer and controller are introduced (Neuron 3 is coloured grey to indicate that it is the controlled neuron). Diagram adapted from [26].



(b) *Top*: First HCO pair. *Middle*: Hub neuron. *Bottom*: Second HCO pair. The simulation is divided into the three phases defined in Figure 3.2.6a. In phase (a), v_3 is a simple bursting oscillation and is inhibited only when v_1 is active. In (b), which starts at $t = 4$, the hub expresses a 'mixed' rhythm as it is inhibited by both v_1 and v_5 . This causes the bursts of v_3 to vary in length. During (c), starting at $t = 10$, the bursting behaviour converges to the same as in (a).

Fig. 3.2.6 Network disturbance rejection.

3.3 Discussion

The results in Section 3.2 suggest that Model Reference Adaptive Control, a classical control paradigm, provides a sound methodological framework for designing online neuromodulation in conductance-based neuronal networks. The control strategy is neuromorphic in the sense that it imitates the continuous adaptation of maximal conductances in biological networks by neuromodulators. It is also classical in the sense that it closely resembles the application of adaptive control in robotics.

The simulations presented in Section 3.2 are of course highly idealised: they assume a perfect modelling knowledge of the voltage-dependent conductances and ideal measurements of all cellular voltages. These assumptions are certainly not met in any practical environment, whether in an experiment of electrophysiology or in the design of analogue neuromorphic hardware. Both *in vivo* and *in silico*, the dynamics of gating variables are both uncertain and variable, that is, they vary with the environment (e.g. temperature) and they vary from cell to cell. Likewise, voltage measurements are noisy, both *in vivo* and *in silico*, which introduces trade-offs and constraints in the observer parameters. For a rigorous study of how noise impacts the system identification of conductance-based models, see [15]. Robustness of the observer to model uncertainty is explored in the next chapter.

We envision at least three distinct angles of attack to increase the robustness and the biological plausibility of the proposed adaptive design:

- **Redundancy and degeneracy** of conductances seem essential in biological neurons to cope with uncertainty. Biological neurons can use vastly different choices of maximal conductances to exhibit the same behaviour [71]. It has been suggested that this degeneracy plays a key role in homeostasis [81]. Viewing the internal dynamics of a neuron as a bank of nonlinear filters that collectively shape the conductance of the total internal current, the adaptive controller functions very much like a one-layer artificial neural network with recursive least-squares estimation of its parameters (this provides a link between biophysical models and the phenomenological models proposed in [14]). The bank of conductances is however not arbitrary in neurophysiology. Ion channel *types* identify specific time scales and amplitude ranges of activation that are critical for the excitability thresholds of the neuron. These specific scales of activation have been well-documented over a long history of voltage-clamp experiments. The resulting dynamic conductances shape the closed-loop behaviour very much like the zeros and poles of a classical LTI controller shape the sensitivity of the closed-loop system [27, 37]. The use of a form of redundancy to improve robustness is explored in Section 4.6 below.

- **Voltage measurements** are assumed to be exact in this chapter, but the spiking nature of the signals suggests that a much coarser information about reference or disturbance signals might be sufficient to modulate the rhythm of a neuron. Future work will explore this possibility to reduce the computational demand of the full observer. Here, neurophysiology will also be a useful guide. For instance, it is well-known that calcium is an essential second messenger involved in neuromodulation. In the present chapter (as in many models of the literature), the intracellular calcium concentration is simply modelled as a low-pass filtered version of the voltage (as in (2.3)). Earlier models have suggested simple yet general homeostatic principles for the adaptation of cellular conductances using the intracellular calcium concentration as a feedback signal [81].
- **Hierarchical adaptation.** The adaptive controller in this chapter is at the cellular scale, and the emphasis is on showing that adaptive controllers at the cellular scale can modulate behaviours at the network level. But there is no doubt that adaptation is multi-scale in biological networks. Adaptive conductance control of conductance-based models is conducive to hierarchical and multi-scale controllers that deserve further research. Adaptive control of a synaptic conductance can be further decentralised and use a conductance observer that only involves the synaptic gating variable, as well as the pre- and post- synaptic voltages. Likewise, adaptive control of a mean-field (or large ensemble) of conductances can be based on an aggregate observer that only involves the mean-field voltage of a population. Here, we also anticipate that the physical nature of the electrical circuit model will allow flexible designs of multi-scale and hierarchical adaptive controllers.

3.4 Conclusion

This chapter has investigated the classical framework of model-reference adaptive control to design neuromodulatory controllers in conductance-based neuronal models. A key message of the chapter is that conductance-based models are linearly parameterised in maximal conductance parameters, and that the adaptive control of individual conductances provides flexible adaptation principles reminiscent of those observed in neurophysiology. The proposed methodology makes use of the adaptive observer recently proposed in [16]. It fundamentally relies on the physical input-output properties of conductance-based models, namely the relative degree one property between currents and voltages, and the contraction property of the internal (gating) dynamics. It is also decentralised, in the sense that the adaptive controller of each node

in a network only estimates local states and parameters based on local measurements, that is, the nodal voltage and the voltage of presynaptic neurons.

We have presented adaptive controllers to solve the two key control problems of reference tracking and disturbance rejection. We have also provided a simple illustration of the role of nodal adaptive control in a conceptual network inspired by the stomatogastric ganglion.

The results of this chapter are preliminary in that they assume full knowledge of the individual conductance models and perfect measurements of the voltage variables. The former constraint is relaxed in the next chapter, and these idealised results suggest a strong potential for classical solutions of adaptive control to provide neuromodulation principles in biophysical or artificial conductance-based networks.

Appendix

Appendix 3.A Model Parameters

The simulation parameters are as follows. In Section 3.2.1, the model parameters are given by

$$\mu_r = \left(120, 0.1, 2, 0, 80, 0.4, 2, 0, 0.1\right)^T$$

The input current $u_r = -2$ and $\kappa = 0.04$. The observer parameters are $\alpha = 0.0008$ and $\gamma = 2$.

In Section 3.2.2 the parameters are given by

$$\mu = \left(60, 0.1, 2, 0, 80, 0.4, 2, 0, 0.12\right)^T$$

and

$$\mu_d = \left(130, 0.1, 3.2, 0, 80, 1, 2, 0, 0.1\right)^T$$

The synapse has parameters $a_1 = 0.53$ and $a_2 = 0.18$. The activation function σ_s has the form given in (2.5). The maximal conductance of the synapse is $\mu_{\text{syn}} = 2.5$. The input currents are $\bar{u} = -2$ and $u_d = -1$, except in the period $t < 400$ when $u_d = -7.5$. This is to delay the first burst of the disturbance neuron to better illustrate its impact. The observer parameters are $\alpha = 0.001$ and $\gamma = 5$, and the controller parameter $\beta = 100$.

In Section 3.2.3 the parameters are given by

$$\mu_1 = \mu_2 = \left(120, 0.1, 1.6, 0, 80, 0.8, 2, 0, 0.1\right)^T$$

and

$$\mu_3 = \left(60, 0.1, 2, 0, 30, 0, 1, 0, 0.1\right)^T$$

The maximal conductances of Neurons 4 and 5, μ_4 and μ_5 , are the same as μ_d above. The synapses all have a_1, a_2 and σ_s as in Section 3.2.2, given above. The maximal conductances of

Table 3.A.1 Reversal potentials and neuron capacitance.

E_{Na}	E_{H}	E_{Ca}	E_{K}	E_{syn}	E_{leak}	c
45	-43	120	-90	-90	-55	0.1

the synapses are $\mu_{\text{syn},2,1} = \mu_{\text{syn},1,2} = 0.8$, $\mu_{\text{syn},5,4} = \mu_{\text{syn},4,5} = 0.6$ and $\mu_{\text{syn},1,3} = \mu_{\text{syn},5,3} = 8$. The gap junctions have conductance $\mu_{\text{gap}} = 0.004$.

The input currents are $\bar{u}_1 = \bar{u}_2 = -3.5$, $\bar{u}_3 = 38$ and $\bar{u}_4 = \bar{u}_5 = -3.2$, except in the period $t < 600$ when $\bar{u}_1 = -8$ and $\bar{u}_4 = -7$. This is as the neurons in the HCO are at rest unless they are inhibited and then released from inhibition. This release from inhibition generates a burst, which in turn inhibits and releases the other neuron in the HCO, and so on. The observer parameters are $\alpha = 0.0004$ and $\gamma = 5$, and the controller parameter $\beta = 100$.

All neurons simulated in this chapter used the same values for the reversal potentials and the neuron capacitance. These values are listed in table [3.A.1](#).

Chapter 4

Robust online estimation of neural circuits

4.1 Introduction

With the recent advancements in our ability to record and manipulate neural activity [18], there is a growing call for control and systems tools that can make use of this new technology [113] for applications including brain-machine interfaces [48] and the treatment of neuronal diseases [79]. The first step in such control tasks is often to obtain model estimates, and there is an extensive literature on fitting neuron models by batch estimation [115, 76, 80, 2]. A downside of such methods is the fact that they are not able to track time-varying model parameters, which is often necessary to characterize neural behavior. By contrast, [16] proposed an adaptive observer for conductance-based neuron models capable of estimating and tracking model parameters in real time. This method was introduced above in Section 2.6 and it was applied in Chapter 3 to solve classical control problems by adaptively controlling the maximal conductance parameters of these models, an approach aligned with the biological concept of neuromodulation [103].

A critical property for practical application of such adaptive methods is to ensure robustness to model uncertainty. Conductance-based models are built from the parallel interconnection of distinct current sources whose specific kinematics and activation range are only approximately known and variable across implementations. The objective of the present chapter is to assess the robustness of adaptive estimation against uncertainty in the internal dynamics of conductance-based models. We investigate the effect of uncertainty in a typical neuronal behavior, namely the neuromodulation of a neuron from spiking to bursting by varying slow calcium conductances.

When the internal dynamics of a neuronal model (also called the channel kinetics) is assumed to be known, [16] showed that a simple adaptive observer equivalent to the Recursive Least Squares (RLS) method can be used to estimate the remaining parameters (maximal conductances). Although [16] also presents a more elaborate scheme for estimating internal

dynamics parameters with local guarantees, here we focus on the basic RLS scheme. In agreement with classical RLS analysis [44], we show empirically that the RLS scheme is sensitive to uncertainty. The main result of this chapter is to show that a good estimation performance can be recovered by the combination of two factors that avoid the need for an internal dynamics estimation: first, by decentralizing the RLS scheme as investigated in [13] with the objective of reducing computational complexity; and second, by introducing redundancy in the estimated model. The positive role of redundancy for robustness and adaptation has been extensively demonstrated in neurophysiology [28, 71]. Here, we mimic biological redundancy by sampling redundant models of ionic currents from a given distribution. This idea can be compared to the method of random features, where random samples of a particular type of basis function are used to solve a regression problem [87]. Our redundant model structure approach can also be related to the ensemble Kalman filter [112], as well as the more general feedback particle filter [120]. Here, however, the gradient of the observed variable is a function not just of the hidden variables (the internal dynamics), but also of the observed variable itself.

Using the rms (root mean square) observer output error as our performance measure, we study the robustness of the different types of adaptive observer algorithms for conductance-based models in the presence of model error. We show that the distributed version of the observer is more robust than the centralized one, and that introducing redundancy to the model structure, according to our proposed approach, further improves this robust behavior.

4.2 Biophysical neuron model

The neuron in this chapter is a conductance-based model following the structure of Section 2.3. It includes five typical ionic currents of a bursting neuron: a transient sodium current I_{Na} , a potassium current I_{K} , a T-type calcium current I_{CaT} , an L-type calcium current I_{CaL} and a calcium-activated potassium current I_{KCa} . We therefore have $\mathcal{S} = \{\text{Na}, \text{K}, \text{CaT}, \text{CaL}, \text{KCa}\}$. The voltage dynamics of a single, isolated neuron (no synaptic currents) are given by

$$\begin{aligned}
 c \dot{v} = & -\mu_{\text{Na}} m_{\text{Na}} h_{\text{Na}} (v - E_{\text{Na}}) \\
 & -\mu_{\text{K}} m_{\text{K}} (v - E_{\text{K}}) \\
 & -\mu_{\text{CaT}} m_{\text{CaT}} h_{\text{CaT}} (v - E_{\text{Ca}}) \\
 & -\mu_{\text{CaL}} m_{\text{CaL}} (v - E_{\text{Ca}}) \\
 & -\mu_{\text{KCa}} \sigma_{\text{KCa}}([\text{Ca}]) (v - E_{\text{K}}) - \mu_{\text{leak}} (v - E_{\text{leak}}) + u,
 \end{aligned}$$

where $[Ca]$ is the calcium concentration, governed by

$$\begin{aligned}\tau_{Ca}[\dot{Ca}] = & -0.03m_{CaT}h_{CaT}(v - E_{Ca}) \\ & - 0.3m_{CaL}(v - E_{Ca}) - [Ca],\end{aligned}$$

with τ_{Ca} a constant. Note that in this chapter the exponents p_{ion} and q_{ion} always take value unity; this simplifies the notation without losing any behaviours of interest. For a full list of parameters used, we refer the reader to the Julia code attached to this chapter.¹

We consider a fixed scenario illustrated in Fig. 4.2.1: the neuron is driven by a known fluctuating input calibrated to expose its excitable behavior. The choice of a fluctuating rather than constant input is made because this is known to lead to more reliable spiking behaviour in biological neurons [67]. During the simulation, the maximal conductance of L-type calcium, μ_{CaL} , and calcium-activated potassium, μ_{KCa} , are ramped up, which results in a modulation from spike excitability to burst excitability. Such neuromodulation is a key cellular mechanism in neurophysiology [25, 26].

4.3 Adaptive recursive least squares estimation

Our starting point is the simple centralized RLS-based observer introduced in Section 2.6 above. For simplicity we will study the single-neuron case $n_v = 1$, but the results generalise easily to networks of arbitrary size. For the example neuron in Section 4.2, we have

$$\theta = \text{col}(\mu_{Na}, \mu_K, \dots, \mu_{leak})$$

and

$$w = \text{col}(m_{Na}, h_{Na}, m_K, \dots, [Ca]).$$

Recall that we use the adaptive observer to estimate the linear parameters θ , in our case the vector of maximal conductances μ_{ion} . We choose maximal conductances for the unknown parameters because the modulation of conductance properties is a key control mechanism of neurophysiology, where it is performed by neuromodulators such as dopamine and serotonin [30]. Neuromodulation is a well-studied phenomenon; it is essential to the function of all nervous systems [68].

¹<https://github.com/RJZS/robust-neuron-estimation>

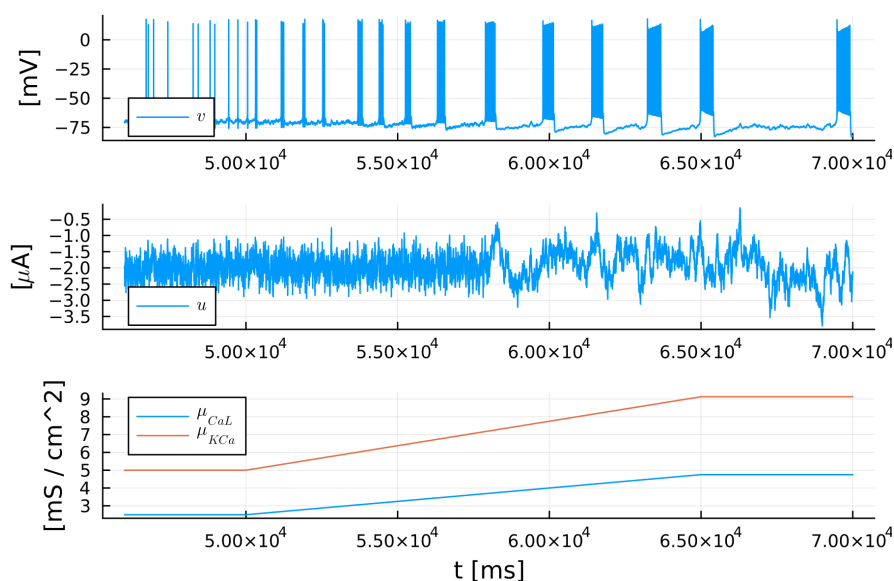


Fig. 4.2.1 Fixed scenario for the illustrations of this chapter. *Top:* membrane potential. *Middle:* input current. *Bottom:* maximum conductance of L-type calcium and calcium-activated potassium (the other parameters remain constant). The fluctuating input is fixed in that this particular input signal is used throughout the chapter. It was originally generated by low-pass filtering Gaussian noise. The properties of the filter were changed partway during the simulation to reflect the difference between spike and burst excitability.

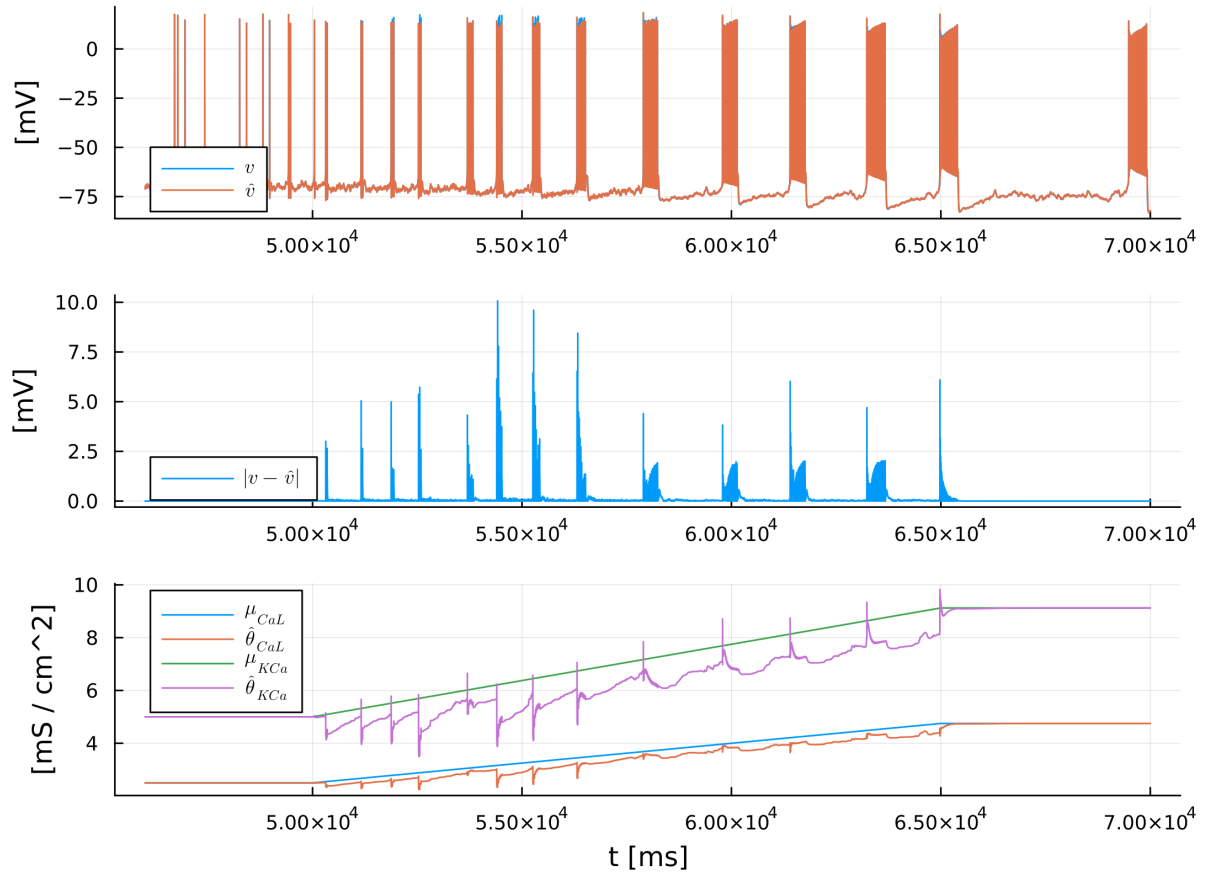


Fig. 4.3.1 RLS estimation in the absence of model error. Parameter and output estimates converge to their true values, and remain there barring transient errors while μ_{CaL} and μ_{KCa} are modulated. *Top*: true voltage v and its estimate \hat{v} . *Middle*: absolute observer error $|v - \hat{v}|$. *Bottom*: L-type maximal conductance μ_{CaL} , its estimate $\hat{\theta}_{CaL}$, calcium-activated potassium maximal conductance μ_{KCa} , and its estimate $\hat{\theta}_{KCa}$.

Recall that P , with dynamics (2.12), can be interpreted as a running estimate (with a forgetting factor) of the parameter covariance matrix [7, Chapter 2]. Note also the use of output injection, that is the injection of true v into the \hat{v} and \hat{w} dynamics.

Fig. 4.3.1 shows that the observer is able to learn the task of Fig. 4.2.1, in the absence of model error. The observer generates a time-varying estimate $\hat{\theta}(t)$ of the vector of maximal conductance parameters $\theta(t)$. The vector $\theta(t)$ includes the two parameters that are modulated, μ_{CaL} and μ_{KCa} , which are shown along with their estimates on the bottom plot. The neuron voltage is also estimated (top plot). As the observer has a perfect model of the internal dynamics $\dot{w} = g(v, w)$, errors in the voltage and parameter estimates are only present transiently, while the maximal conductances vary.

4.4 Robustness of the observer

To investigate the robustness of the centralized observer against variability in the internal dynamics, we introduce the random variables $p \sim U(1-r, 1+r)$ and $q \sim U(-s, s)$, where $U(a, b)$ represents the uniform distribution with support $[a, b]$. We define stochastic versions of the gating variable dynamics (2.2b)-(2.2c) as follows:

$$p_{m,\text{ion}} \tau_{m,\text{ion}}(v) \dot{m}_{\text{ion}} = -m_{\text{ion}} + \sigma_{m,\text{ion}}(v - q_{m,\text{ion}}) \quad (4.1a)$$

$$p_{h,\text{ion}} \tau_{h,\text{ion}}(v) \dot{h}_{\text{ion}} = -h_{\text{ion}} + \sigma_{h,\text{ion}}(v - q_{h,\text{ion}}). \quad (4.1b)$$

The effect of the random variables is to respectively scale and shift the time-constant and activation functions. Recall that each activation function is a monotonic function, specifically a sigmoid or other ‘S-curve’, with range $(0, 1)$. A shift in this function can therefore be thought of as a shift in the gating variable’s ‘half-activation’. By perturbing the half-activation and timescale of each gating variable, we are introducing error into two of its defining properties. Note we continue to assume that the conductance-based modelling approach is fundamentally correct.

We collect the random samples into vectors $p, q \in \mathbb{R}^{n_w}$ and the randomized dynamics into a function $g(v, \hat{w}; p, q)$. We replace the observer’s internal dynamics (2.11b) with

$$\dot{\hat{w}} = g(v, \hat{w}; p, q). \quad (4.2)$$

This introduces mismatch between the true dynamics (2.8b), which remain deterministic, and the observer’s model of these dynamics. In this chapter, we take $r = 0.04$ and $s = 4mV$.

Fig. 4.4.1 illustrates the observer’s performance with one set of samples for model error. The figure is otherwise identical to Fig. 4.3.1. We chose $\gamma = 8$ and $\alpha = 0.005$. Higher γ reduce observer error but are sensitive to noise measurement, hence we use the same value throughout the chapter for a fair comparison. The value for α was tuned manually to optimize performance.

We take our performance measure to be $e_{v,\text{rms}}$, the rms value of the observer error over the duration of the simulation in Fig. 4.2.1. This is computed as $(\sum_{t=1}^T (v(t) - \hat{v}(t))^2 / N)^{\frac{1}{2}}$ where N is the number of simulation time steps; the step size is $\Delta t = 0.1$. The mean value of $e_{v,\text{rms}}$ is computed over twenty trials (Fig. 4.4.1 shows one of these trials) and the result is shown in the first column of Table 4.6.1. Although output injection ensures that the spike and burst estimates align, there is significant error in the output and parameter estimates. The parameter

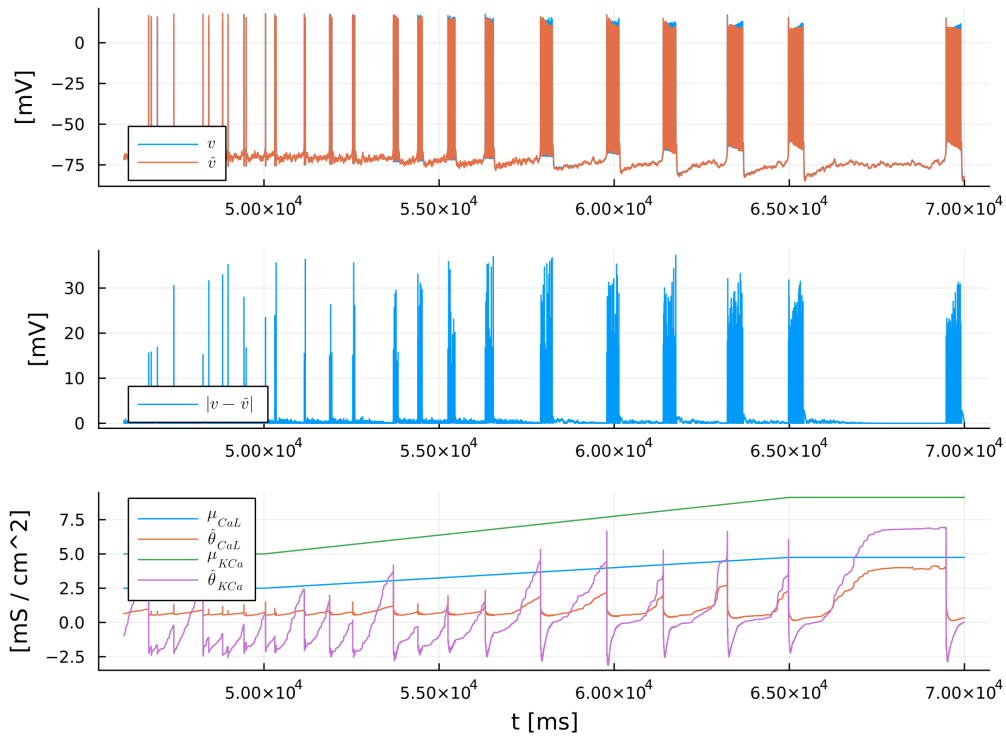


Fig. 4.4.1 The centralized observer is fragile to model uncertainty. *Top*: true voltage and its estimate. *Middle*: absolute observer error. *Bottom*: time-varying maximal conductances and their estimates.

estimates fail to track the modulation of calcium currents, limiting any practical use of the online observer.

4.5 Distributed observer

The centralized observer introduced above is based on the recursive least squares algorithm [16]. It makes use of second-order information, which becomes a source of error when we introduce model mismatch. This information is encoded in matrix P of the observer, which is proportional to the covariance matrix of the empirical estimate of θ under the assumption of zero-mean, independent, identically distributed additive noise [7, Chapter 2]. A first remedy to increase robustness is therefore to neglect the covariances by decentralizing the parameter estimation. The core idea is to approximate the matrix $\Psi\Psi^T$, which appears in the update equation for P , by its (potentially block) diagonal elements to yield a decentralized learning rule. We thus replace the $n_\theta \times n_\theta$ matrix P with n_θ scalars P_i . This idea was already explored

in [13] with the goal of reducing the algorithmic complexity of the estimator from $O(n_\theta^2)$ to $O(n_\theta)$, a significant benefit for more complex neurons and networks with many synapses.

The distributed observer with model mismatch has the form:

$$\dot{\hat{v}} = \sum_{j=1}^{n_\theta} \Phi_j^T(v, \hat{w}^j, u) \hat{\theta}_j + a(v, \hat{w}, u) \quad (4.3)$$

$$+ (\gamma_0 I + \sum_{j=1}^{n_\theta} \gamma_j \Psi_j^T P_j \Psi_j)(v - \hat{v})$$

$$\dot{\hat{w}}^j = g_j(v, \hat{w}^j; p_j, q_j) \quad (4.4)$$

$$\dot{\hat{\theta}}_j = \gamma_j P_j \Psi_j(v - \hat{v}) \quad (4.5)$$

where $\gamma_0, \gamma_1, \dots, \gamma_{n_\theta} > 0$ are constant gains. The matrices P_j and Ψ_j evolve according to

$$\dot{\Psi}_j = -\gamma_j \Psi_j + \Phi_j(v, \hat{w}^j, u)$$

$$\dot{P}_j = \alpha_j P_j - \alpha_j P_j \Psi_j^T \Psi_j P_j \quad P_j(0) \succ 0$$

with $\alpha_j > 0 \forall j$.

For the example neuron in Section 4.2, for instance, we have

$$\theta_1 = \mu_{\text{Na}}, \theta_2 = \mu_{\text{K}}, \dots, \theta_{n_\theta} = \mu_{\text{leak}}$$

and

$$w^1 = \text{col}(m_{\text{Na}}, h_{\text{Na}}), w^2 = m_{\text{K}}, \dots, w^{n_\theta} = \theta$$

Fig. 4.5.1 shows a representative example of the distributed observer's performance in the presence of randomly-sampled model error. We set $\gamma_j = 8$ for all j , the same as with the centralized observer for a fair comparison. We set by hand tuning $\alpha_j = 2 \times 10^{-4}$ for all j . To compare the performance of the two observers, we compute the mean value of $e_{v,\text{rms}}$ across twenty trials. The results are shown in the first two columns of Table 4.6.1. The observer error is indeed reduced by the use of the distributed observer. The parameter estimates have also improved, as they oscillate less and are therefore more meaningful. We do not necessarily expect the parameter estimates to settle near the true values, as the observer error can be reduced by exploiting the biological redundancy between currents. In the next section, we exploit the theme of redundancy to further improve robustness.

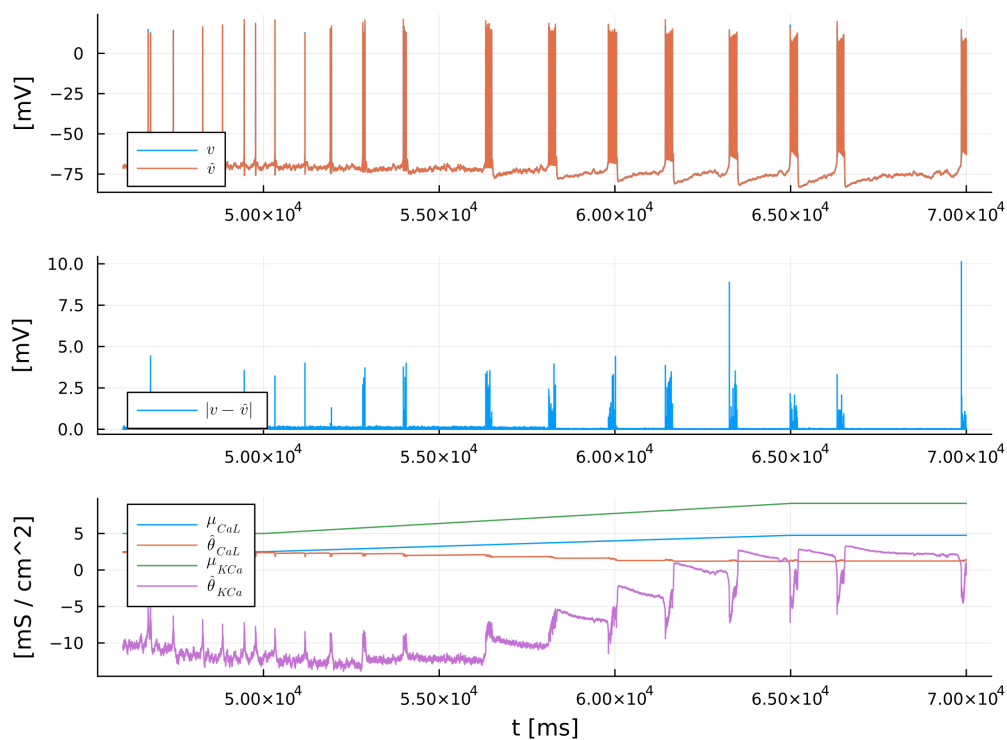


Fig. 4.5.1 Distributed parameter estimation in the scenario of Fig. 4.2.1. *Top*: true voltage and its estimate. *Middle*: absolute observer error. *Bottom*: μ_{CaL} and μ_{KCa} and their estimates. Although there is some observer error, it is significantly smaller as reflected in the v and \hat{v} spikes being almost indistinguishable.

4.6 Redundancy

In biology, redundancy between ionic currents allows them to compensate for each other to achieve desired behavioral properties [28]. Redundancy is known to play a crucial role in the homeostasis of neuronal function in spite of the highly variable ion channel density, both across time and from animal to animal [71].

Redundancy can also be exploited to increase the robustness of estimation to model uncertainty. To test that idea, we consider an augmented observer model that includes for each gating variable N equations of the form (4.1a) or (4.1b), each with its own samples for p and q . That is to say, each element of (4.4) is replaced with N elements

$$\hat{w}^{j,i} = g_j(v, \hat{w}^{j,i}; p_i^j, q_i^j) \quad (4.6)$$

for $i = \{1, \dots, N\}$. We can consider this step as replacing a single estimate of the gating variable m_{ion} or h_{ion} with N particles of the same, and the resulting algorithm analogous to an ensemble observer. Note that we do not change the model of the neuron being observed.

Every ionic current now has N corresponding terms in the \hat{v} dynamics, each with its own gating variable dynamics perturbed by random samples for p_i^j and q_i^j . The observer can exploit these N redundant currents, each with slightly different gating variable dynamics, to find a better fit to the true neuron model. Hence, we replace (4.3) with

$$\begin{aligned} \dot{\hat{v}} = & \sum_{j=1}^m \bar{\Phi}_j^T(v, \hat{w}^j, u) \hat{\theta}_j + a(v, \hat{w}, u) \\ & + (\gamma_0 I + \sum_{j=1}^m \gamma_j \bar{\Psi}_j^T \bar{P}_j \bar{\Psi}_j)(v - \hat{v}), \end{aligned} \quad (4.7)$$

where $j \in \mathcal{S}$ is the set of membrane currents. We now have

$$\bar{\Phi}_j = \text{col}(\Phi_j(v, \hat{w}^{j,1}, u), \dots, \Phi_j(v, \hat{w}^{j,N}, u)) \in \mathbb{R}^N$$

for $j = 1, \dots, n_\theta - 1$, and

$$\bar{\Phi}_{n_\theta}(v, \hat{w}^{n_\theta}, u) = \bar{\Phi}_{n_\theta}(v) = -(v - E_{\text{leak}})$$

with the latter corresponding to the leak current regressor. \bar{P}_j is now an $N \times N$ diagonal matrix with leading diagonal $\text{col}(\bar{P}_j^1, \dots, \bar{P}_j^N)$.

We also have

$$\hat{\theta}_j = [\hat{\theta}_j^1, \dots, \hat{\theta}_j^N],$$

the set of parameter estimates corresponding to a particular maximal conductance of the reference neuron. We define the empirical mean over this set,

$$\bar{\theta}_j = \frac{1}{N} \sum_{i=1}^N \hat{\theta}_j^i,$$

which provides an estimate of the scaled maximal conductance μ_j/N .

We will apply redundancy only to the distributed observer. The diagonal nature of the distributed observer makes it scalable with respect to the increased number of states and parameters.

Redundancy is antagonist to persistency of excitation, the uniqueness condition for recursive least-squares, and indeed a naive implementation of the observer leads to situations where some estimated maximal conductances become negative, causing instability.

To prevent the divergence of a redundant estimator, we modify the $\hat{\theta}_j$ update law (4.5) to include a consensus term, that regularizes the variance of the redundant parameters. The update law for the i th redundant element of $\hat{\theta}_j$ is now

$$\dot{\hat{\theta}}_j^i = \gamma_j \bar{P}_j^i \bar{\Psi}_j^i (v - \hat{v}) - \beta (\hat{\theta}_j^i - \bar{\theta}_j). \quad (4.8)$$

Fig. 4.6.1 illustrates the performance of the redundant estimator in the presence of model error. We set γ_j and α_j as in Section 4.5. We chose by hand tuning $\beta = 5 \times 10^{-5}$, a value low enough that the redundant parameters take on distinct values.

The mean rms error across twenty trials is provided in Table 4.6.1. The third and fourth columns refer to the final algorithm, respectively with 3 and 9 redundant elements per gating variable. As expected, we see significant improvement of the observer error when redundancy is introduced, and a greater improvement with more redundancy.

We see however in Fig. 4.6.1 that the redundant terms do not necessarily directly track modulation. This is as the observer is exploiting redundancy across all terms to minimise voltage error, not just those corresponding to a single ionic current.

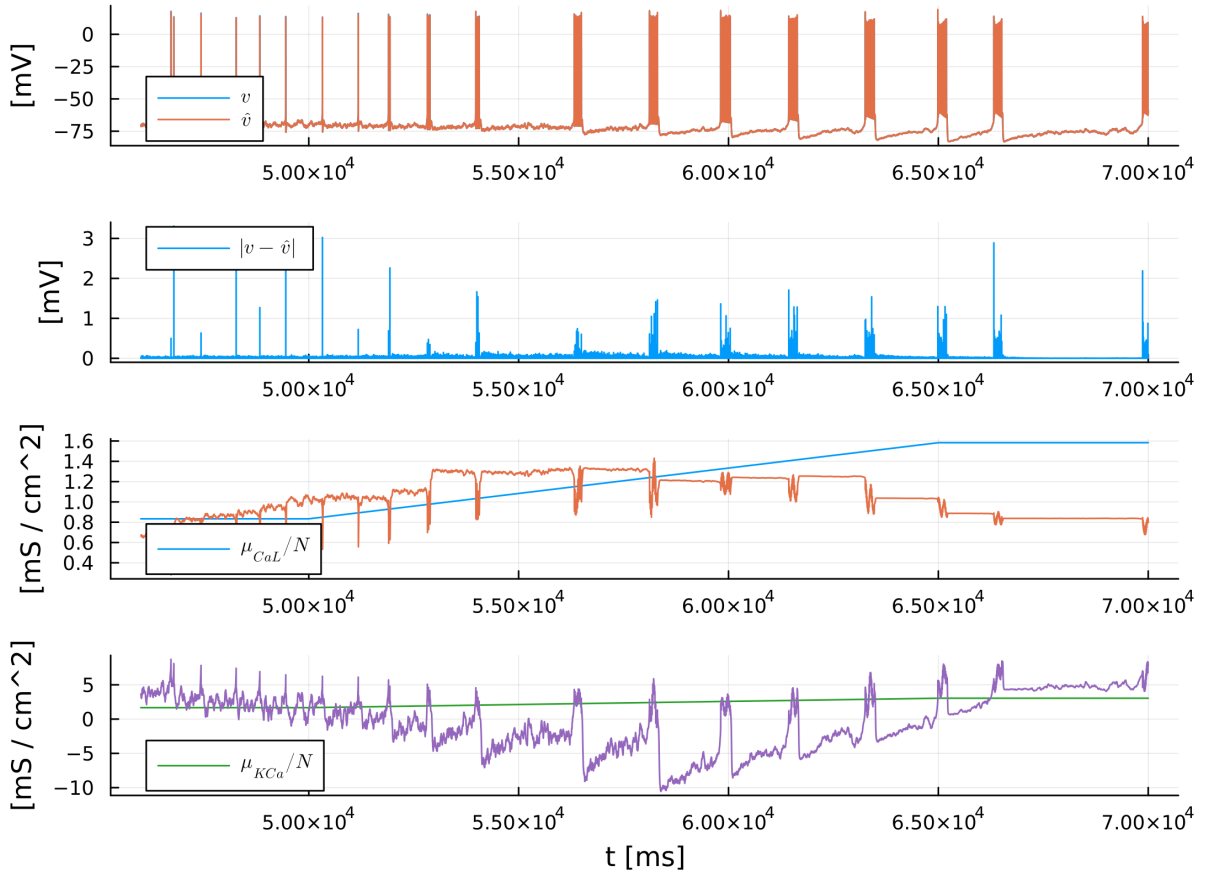


Fig. 4.6.1 Distributed and redundant online estimation in the scenario of Fig. 4.2.1, in the presence of model error and with $N = 3$. *Top*: true voltage and its estimate. *Upper middle*: absolute observer error. *Lower middle*: the empirical mean estimate $\hat{\theta}_{\text{CaL}}$; for comparison, the scaled true parameter μ_{CaL}/N is also shown. We plot the *scaled* parameter to emphasise that we have replaced each conductance with N separate conductances. *Bottom*: the same for μ_{KCa} . Labels for the parameter estimates are omitted.

Table 4.6.1 Comparison of all three observers, showing the mean and standard deviation of the rms voltage error $e_{v,\text{rms}}$ (in mV) across twenty trials.

	Centralized	Distributed	$N = 3$	$N = 9$
Mean	1.15	0.0788	0.0280	0.0241
Standard Deviation	0.14	0.017	0.0056	0.0030

4.7 Discussion

Our results suggest that an adaptive observer can be used to estimate neuronal parameters, even in the presence of modelling error. The chosen model of uncertainty is plausible in a neuromorphic context where a key cause of error is transistor mismatch introduced during manufacturing [105]. This component imprecision is one of the main challenges facing designers of neuromorphic hardware [65]. In a biological context, the cell is of course part of a living system and is therefore time-varying. The online nature of the observer ensures that our results extend to this time-varying case.

In future work, the results above should also be tested against measurement error. Noisy voltage measurements introduce trade-offs in the design parameters. For a rigorous examination of the impact of noise on the system identification of conductance-based models, we refer the reader to [15].

The performance metric of our study was the rms observer. This is a reasonable first step to make a quantitative comparison of different observers. However, it is only a proxy of the practical objective to estimate parameters in order to track neuromodulation in an experimental setup or to use learning experiments for hardware implementations of neuromorphic neurons. This will be the topic of future research.

Chapter 5

Neuromorphic control of a pendulum

5.1 Introduction

Most robotic tasks can be decomposed as sequences of events [60, 119]. This is best illustrated in the context of animal-like movements such as walking [40] or swimming [53]. In spite of spectacular advances in robotics, animals still drastically overperform robots in performing such tasks reliably in changing and uncertain environments. The vision of neuromorphic engineering is that the event-based nature of animal control is key to its superior performance over our current clocked digital technology [74].

The potential of event-based control was recognized from the early days of neuromorphic engineering [22]. The event-based PI control in [1] launched the new and still flourishing field of event-based control; see [4] for a recent survey. While the theoretical benefits of event-based over sampled data control are now clearly established, most of the existing literature has concentrated on emulating the behavior of classical digital control systems with an event-based architecture. In contrast, the design of control systems that interconnect rhythmic systems through sensing and actuating events is still in its infancy [100, 32].

The goal of this chapter is to explore the potential of neuromorphic control with the simple mechanical model of a pendulum. We regard the pendulum as a rhythmic system and we regard the control problem as the design of another rhythmic system able to orchestrate the behavior of the pendulum. The rhythmic controller is designed in such a way that the desired behavior is achieved by a form of synchrony between the two rhythmic systems.

The design of a rhythmic event-based controller starts with the design of an automaton: the desired behavior of the controlled system is represented by a temporal sequence of sensor and actuator events. The automaton controller is a rhythmic system capable of generating the desired sequence of events. The control objective is formulated as a synchronization problem

between the plant and the automaton. The second step is to design a regulator that uses output feedback and that tunes continuous parameters to ensure robustness and adaptability of the controller.

We show that the bio-inspired architecture of neuromorphic controllers such as presented in [90] provides a flexible methodological framework for such a two step design. The neural architecture of the controller consists of simple motifs. The topology of the network defines the automaton, whereas the regulatory properties are achieved by output feedback, and by neuromodulation of the parameters using the classical framework of adaptive control. Despite being rhythmic and non-linear, the closed-loop system is amenable to rigorous analysis using existing theory. We provide no theoretical development in this chapter but point to the relevant references in each section.

The idea of entraining a mechanical system with a neuro-inspired controller has a rich and long history; it underlines much of the literature connecting the central pattern generators of neuroscience and the rhythmic controllers of legged robotics [51]. Nonlinear oscillators have been successfully designed to orchestrate rhythmic behaviors [53, 52], and neuromorphic implementations have received recent attention [3]. A limitation of those approaches is the lack of a systematic modelling framework, which has motivated the design of linear controllers [92]. Neuro-inspired controllers have been implemented experimentally, see e.g. [63, 106]. But the tuning of those architectures has proven difficult. Event-based control has been explored for the linearised pendulum [59], with an approach that utilises classical linear methods. The design of a controller by means of an adaptively tuned rhythmic automaton and an output feedback phase control seems novel.

The chapter is organised as follows. The automaton of the plant is introduced in Section 5.2. Section 5.3 presents the event-based control architecture; its implementation follows in Section 5.4. In Section 5.5, we control the pendulum in open-loop. Sections 5.6 and 5.7 introduce output and adaptive feedback respectively. Section 5.8 discusses the potential of the present approach, with the intention of motivating further research from the community.

5.2 The automaton of a pendulum

We consider the non-dimensionalised dynamical model of a pendulum

$$\ddot{q} + \alpha\dot{q} + \sin(q) = I \tag{5.1}$$

where q is the pendulum's angle from the resting position, α is dimensionless damping and I is dimensionless torque.

The rich dynamical behaviour of this seemingly simple system is part of any textbook of nonlinear dynamics see e.g. the excellent treatment in [110]. With a constant torque, the qualitative behaviour can be comprehensively studied through phase-portrait analysis. The key properties of the pendulum dynamics can be summarized as a function of the two parameters α and I .

The behavior of the pendulum is simple in the so-called overdamped regime ($\alpha > 1$). In this regime, the only attractor of the system is a stable equilibrium, for $I < 1$, or a limit cycle, for $I > 1$. The behavior is more complex in the underdamped regime ($\alpha < 1$). In this regime, the stable equilibrium can coexist with a stable limit cycle, depending on the initial energy of the system. This bistability is the source of complex behaviors, including co-existence between small and large-amplitude oscillations and sensitivity to initial conditions under non-constant forcing. Small variations of the applied torque suffice to switch the behavior between “small” and “large” oscillatory behavior.

In the event-based framework of the present paper, the torque is not constant but a sequence of impulses of short duration. As in the discussion above, we can distinguish between two types of stationary behaviors: *small oscillations* (less than a full swing), that replace the stable equilibrium obtained with a constant torque, and *large oscillations* (sequences of complete rotations), as in the regime of constant torque with $I > 1$. This discrete distinction leads to the description of the pendulum as a two state automaton, with a “low energy state” corresponding to small oscillations and a “high energy state” corresponding to large oscillations. Both states can be entrained by a periodic sequence of impulses. Event-based control of the pendulum can be thought as orchestrating the trajectory of that automaton (that is, the switches between the high and low states), as well as tuning each of the states: regulating the frequency and amplitude of the oscillation in the low state, and regulating the frequency of oscillations in the high state.

5.3 Event-based control architecture

To conceive the rhythmic automaton of the controller, imagine two children pushing a swing. By pushing harder or more frequently, they can increase the swing's amplitude or frequency respectively. The children can coordinate in two different ways: they can stand on the same side of the swing and push it in synchrony, or on opposite sides and push it in anti-synchrony. We replicate a similar architecture with two independent motors forcing the pendulum. Each




Fig. 5.3.1 The two distinct configurations of event actuation. *Left*: in-phase identical actuating events (IN-PHASE). *Right*: anti-phase actuating events of opposite sign (ANTI-PHASE). Sketch courtesy of @artjoy2015.

actuator periodically ‘kicks’ the pendulum with a pulse of variable duration. The duration and frequency of the pulses control the amplitude and frequency of the pendulum. Like the children standing on the same side, the actuators can produce a torque of the same sign and act in phase. Or, they can produce torque of the opposite sign and act in anti-phase, ‘kicking’ the pendulum back and forth. We call these two options the “in-phase configuration” (IN-PHASE) and the “anti-phase configuration” (ANTI-PHASE) respectively, and we illustrate them in Fig 5.3.1.

In neuromorphic circuits, such rhythms are generated endogenously by simple neuronal motifs, inspired from the central pattern generators of neuroscience [10] that are discussed in Section 2.5. The simplest rhythmic motif is the “half-centre oscillator” (HCO) [70]. An HCO comprises two ‘neurons’ (arbitrarily labelled as ‘A’ and ‘B’) mutually connected by inhibitory ‘synapses’. In this chapter, we use the HCO model of [90] which was developed for the ease of its tunability. Our controller includes one HCO per motor. The HCO drives the motor as follows: when neuron A is spiking, specifically when its voltage is above some threshold, the motor produces a constant torque. When the neuron is below this threshold, the motor is at rest. Fig. 5.3.2 shows the behavior of an HCO and the corresponding motor output, for two burst sizes. The burst size controls the actuator event.

The coordination between the two motors relies on synaptic interactions between the two HCOs. Excitatory synapses between the pair of A neurons, and also between the pair of B neurons, ensure that the two HCOs are in phase. Inhibitory synapses between the same neurons ensure that the two HCOs are in anti-phase.

The frequency and amplitude of the pendulum oscillations are estimated from event-based sensors. In this chapter, we use three photodetectors that detect specific angle crossings: one for




Fig. 5.3.2 The anti-phase rhythm of a Half-Centre Oscillator (HCO) circuit. The size of the bursts is modulated midway. *Top*: the voltages of the two neurons. The dashed horizontal line denotes the voltage threshold, above which the motor is active. *Bottom*: the corresponding actuating events.




Fig. 5.3.3 Simulation of the HCOs and actuator events. *Top*: the total torque I . *Middle*: the voltages of HCO 1. *Bottom*: the voltages of HCO 2. At $t = 8$ s (indicated by the dashed black line), the system's configuration is switched from ANTI-PHASE and, after a transient period, it settles on IN-PHASE.




Fig. 5.3.4 Block diagram of the complete architecture, including the event-based feedback loops introduced in Sections 5.6 and 5.7. Small arrows over signal transmission lines indicate event-based communication as described in Section 5.3. The HCO block architecture is described in Sections 5.3 and 5.4.

the output feedback phase controller ($q = q_p$), and two for the adaptive regulation of frequency and amplitude ($q = 0$ and $q = A_{ref}$).

Fig. 5.3.4 shows the block diagram of the complete architecture.

5.4 Neuromorphic implementation

The circuit realization of our rhythmic controller follows the mixed-feedback motif of spiking control systems [100]. Such systems are minimum phase, relative degree one, and linear in the control parameters. Those properties make the controller ideally suited for output feedback control and for adaptive regulation of the control parameters. See Section 2.4 for the circuit model.

For simplicity, we use identical parameters for the four neurons as well as for each type of synapse. The controller behavior is however quite robust to heterogeneity in the parameters, as shown by the I-V curve method discussed in Section 2.4. The ultra-slow timescale τ_{us} is chosen to (roughly) align with the natural frequency of the pendulum. This nominal choice enables a low-power controller to achieve a wide range of amplitudes by resonance. Recall from Section 2.4 that the other timescales are chosen such that they are all sufficiently separated [89]. We set $i_{app,i}$ to a nominal value $i_{nominal} = -1$. The controller does not oscillate spontaneously, but a brief pulse of current in any of the neurons is sufficient to turn the rhythm ‘on’. A full list of

the parameter values used in each figure is available with the attached code.¹ All quantities in this chapter are non-dimensional.

The only tuning parameters in this chapter are the gain of the slow negative conductance g_s^- and the gain of the ultraslow positive conductance g_{us}^+ . Those two parameters respectively modulate the burst size and the burst frequency. All the other parameters are fixed throughout.

5.5 Control by entrainment

Feedforward control of a rhythm by another rhythm is commonly referred to as *entrainment*. The role of entrainment in biology can be regarded as analog to the role of resonance in mechanics. It is widely observed, and it is robust to uncertainty in that it does not require an exact prior match between the frequencies of the input rhythm and the entrained rhythm [61]. A stable limit cycle is always locally entrainable by any periodic forcing signal of small amplitude and ‘close enough’ frequency, as shown by Section III of [117]. Global entrainment has also been studied for contractive systems [107, 91].

As a first step, we consider entrainment at a fixed frequency (chosen to be close to the natural frequency of the pendulum) by a feedforward controller.

Fig. 5.5.1 shows entrainment of the overdamped pendulum at the fixed frequency, for different choices of burst size and using the anti-phase configuration of the controller. The amplitude increases with burst size, which is regulated by g_s^- . This gain also affects the frequency of the HCO, a variation compensated by the parameter g_{us}^+ . Fig. 5.5.2 illustrates the monotonic relationship between the amplitude of the entrained pendulum and the neural gains g_s^- and g_{us}^+ at the fixed frequency.

The entrainment behavior becomes more complex in the underdamped regime: IN-PHASE entrainment leads to one of two possible stable states (small or large oscillations), depending on the initial conditions. Fig. 5.5.3 shows the energy transferred by the motors during each neural period, $E_i = \int_{t_i}^{t_i+T} \dot{q} I dt$, where each time t_i is at the beginning of a burst event, and T is the time until the next such event. The figure shows that the pendulum can converge to either the high or the low energy state. The high-energy state in Fig. 5.5.3 is a 2:1 oscillation, meaning two complete swings of the pendulum for every actuator event.

¹<https://github.com/RJZS/neuromorphic-pendulum-control>

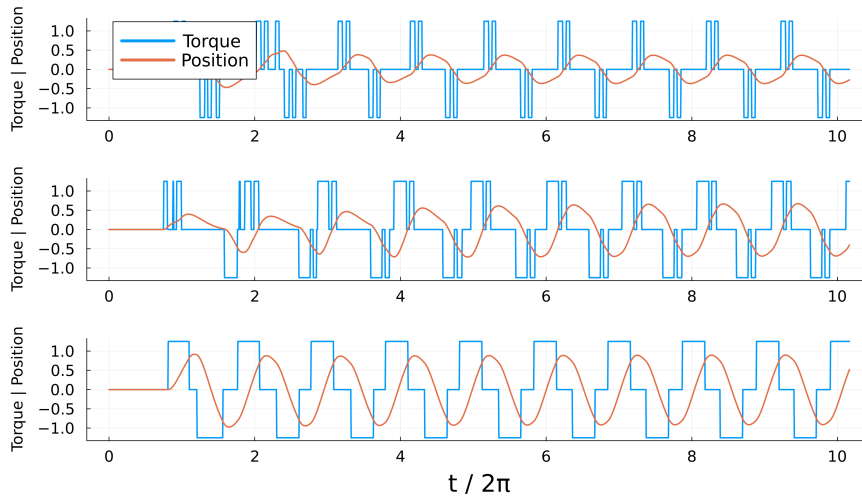


Fig. 5.5.1 Small oscillations in the overdamped regime (damping $\alpha = 1.4$). As the burst size of the neural oscillation increases (from top row to bottom row), so does the amplitude of the pendulum's oscillation.




Fig. 5.5.2 Entrainment of the pendulum at a fixed frequency in the overdamped regime ($\alpha = 1.4$). The amplitude of small oscillations increases monotonically with g_s^- and g_{us}^+ .

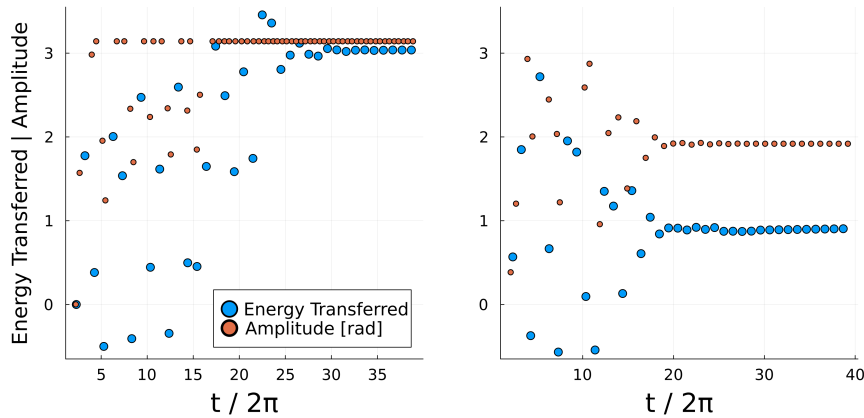


Fig. 5.5.3 Entrainment with an IN-PHASE neural oscillation in the underdamped regime (in this example, $\alpha = 0.14$), showing energy transferred by the motors each burst. The amplitude of each oscillation is also shown (with an amplitude of π indicating a complete swing). Depending on the initial conditions, the steady-state oscillations are either large or small (left and right plots respectively).

5.6 Phase control

In contrast to small oscillations, the entrainment of large oscillations in the underdamped regime is *fragile*. A slight change in the parameters can lead to a non-periodic behaviour, which again is consistent with the chaotic attractor that the pendulum can exhibit with a small sinusoidal input superposed to a constant torque input. In this section, we use feedback control to significantly enlarge the basin of attraction of the large oscillations.

To design an event-based output feedback control, we use the concept of phase control. Phase control consists of advancing or delaying the next actuator event via small negative pulses added to the nominal (constant) input current of each neuron. The phenomenon is illustrated in Fig. 5.6.1 (left). The delay or advance of the next actuator event is a function of the timing of the pulse, a relationship known as the ‘phase response curve’ (PRC) [93]. It is known that a monotone phase response curve allows control of the phase of an oscillator by a simple proportional controller; see [93, 29]. The PRC method is based on linearized analysis, which assumes *small* perturbations. We use perturbations of fixed amplitude (P) and fixed duration (w). Fig. 5.6.1 (right) illustrates the strong monotonicity of the PRC of the HCO oscillator, making the simple proportional phase control effective. The sensory measurements are made with an event-based photodetector placed at $q = q_p$. The phase controller applies a current pulse at each sensory event $q(t) = q_p$. The parameter q_p (selected to $q_p = -1$ rad in the simulation of Figure 5.6.2) can be regarded as the selection of a proportional gain. It does




Fig. 5.6.1 *Left*: A small inhibitory pulse (green) causes a phase advance or delay of the subsequent burst. The blue curve shows v_A for an isolated HCO. The blue dashed curve shows the behaviour of v_A in the absence of the pulse. *Right*: Phase response curve of an HCO. The curve indicates the phase shift resulting from a fixed pulse ($w = 0.05$ and $P = 0.3$) as a function of the timing of the pulse along the limit cycle.

not require a precise tuning. In the IN-PHASE configuration, the feedback pulses are at each sensory event injected alternately into the two A neurons and the two B neurons.

Fig. 5.6.2 illustrates the stabilizing property of a proportional phase control. The initial conditions are those of Fig. 5.5.3 (right plot), but the phase control forces convergence to the high- rather than low-energy state. The steady-state frequency is slightly different to that obtained with feedforward entrainment, and varies with q_p .

5.7 Neuromorphic adaptive control

Online adaptation of the control parameters can be used as well as phase control in order to regulate an oscillation.

Robust adaptive regulation of the control parameters has been studied above in Chapters 3 and 4, as well as in [16, 13]. This work relies on the property that the proposed neuronal architecture inherently satisfies the standard assumptions of adaptive control: the system is minimum phase, relative degree one, and linear in the parameters. Adaptive control is analogous to the biological concept of neuromodulation [71].

We consider the online adaptation of the two parameters g_{us}^+ and g_s^- , which regulate the neural frequency and burst size respectively. We use these to regulate the frequency and amplitude of small oscillations in ANTI-PHASE. The online adaptation is implemented via the integral action of a prediction error computed from the sensory events.




Fig. 5.6.2 The simulation of Fig. 5.5.3 (right plot) with the addition of phase control, forcing convergence to the high- rather than the low-energy state. *Left*: the initial response. The current pulses to the A neurons are shown in purple, and are scaled for clarity. Each pink marker indicates the start of an actuation event. *Right*: the energy transferred per event, E_i , showing convergence to a steady-state oscillation. The pink markers match those of the left plot.

For a reference frequency ω_{ref} , we predict that a zero crossing $q(t) = 0$ will occur every π/ω_{ref} seconds. A correction p_ω is added to the parameter g_{us}^+ at each zero crossing, with a sign determined by the sign of the prediction error and an amplitude proportional to its magnitude.

Likewise, for a reference amplitude A_{ref} for the pendulum's oscillation, we predict that an event $q = A_{\text{ref}}$ occurs $\pi/2\omega_{\text{ref}}$ seconds after each $q = 0$ event. If no such event occurs before the next zero crossing, we apply a fixed correction p_A to the parameter g_s^- . Otherwise, the correction is applied at the event $q = A_{\text{ref}}$ with a sign and amplitude determined by the prediction error.

The simple adaption law described above can be regarded as a phase integral control. In practice, the parameter variation is confined to a finite range such that bursting is preserved.

Fig. 5.7.1 illustrates parameter convergence for the reference values illustrated in Figure 5.5.2.

5.8 Discussion

We have presented a neuromorphic, event-based framework for controlling the oscillations of a mechanical pendulum. A simple neuromorphic controller was used to generate actuator events capable of entraining the pendulum oscillation while regulating its energy. The feedforward controller exhibits bistability between low-energy and high-energy oscillations in the underdamped regime. A proportional phase control was designed to stabilise the high energy oscillation only. Finally, an integral phase control was designed to regulate oscillations




Fig. 5.7.1 The gains tuned by the adaptive controller, when used to obtain the parameters in Fig. 5.5.2. The legend gives the values of A_{ref} in radians. In each case, ω_{ref} is the fixed frequency chosen in Section 5.5.

at a given frequency and amplitude. While no formal convergence analysis was included in the chapter, the design was supported by existing analysis and system properties that enable regulation by simple feedback controllers (minimum phase and relative degree one).

The event-based nature of the designed controller offers a number of potential advantages. Prime and foremost, the energy exchange between the systems is confined to the events. At all other times, the controlled system is free to act under its own dynamics. The control action serves only to entrain these dynamics, and so the entraining controller always aligns with the system's physics. As a consequence, the temporal sparsity of the events is a direct measure of the energy efficiency of the design.

High impedance control during the events is an inherent source of robustness to model uncertainty [62]. Temporal sparsity of the events simultaneously ensures low impedance control when averaged over time. In this way, event-based control could potentially combine the benefits of soft actuation with the high impedance requirements of robust control, suggesting a path to overcoming the classical stiffness-compliance trade-off [85]. Current methods, by contrast, must take one of two distinct approaches: the high-impedance inversion-based methods, or the impedance control method of soft robotics. The considered two-motor architecture also suggests how actuation events can be easily distributed over possibly many actuators, with phase relationships ensured by the architecture of the controller.

The purpose of the chapter was to illustrate the potential of neuromorphic control in a simple mechanical system. We hope that it will stimulate further theoretical and practical research in the control of robotic or biomechanical rhythms. It would seem natural to explore the potential of neuromorphic control architectures in robotic designs that have already exploited rhythmicity, for instance [53].

Chapter 6

Conclusions

Rhythmic control tasks, whether biological, electronic or mechanical, can be solved with neuromorphic controllers using concepts from classical control theory. This is because neurons have structural properties that make them amenable to output and adaptive feedback control. The resulting closed-loop behaviour is event-based, making this approach suited to the design of embodied systems. The event-based behaviour derives from the presence of mixed feedback. Neuromodulation is used to regulate the behaviour by tuning the gains of the feedback loops.

In this thesis we considered two types of control task. When the controlled system was itself a neuronal circuit, the controller was formed of a bank of additional conductances that shaped the conductance of the overall circuit in a manner reminiscent of impedance control. When the plant was a mechanical system, we took a two-step approach. Firstly, we represented the desired behaviour as a sequence of events, and designed an automaton controller capable of generating the desired sequence. Secondly, we designed a regulator that uses adaptive and output feedback to ensure the controller's robustness and adaptability.

The latter approach illustrates the potential of neuromorphic event-based control for robotics. The aim of such an approach is to control a system by interacting minimally with it through localised events, thus producing a controller that is highly robust and energy-efficient, and that exploits the physics of the robot and its sensors.

We now provide a summary of the thesis, followed by a discussion of potential future work.

6.1 Summary

Adaptive conductance control

In Chapter 3 we have solved the two canonical adaptive control tasks (reference tracking and disturbance rejection) for networks of conductance-based neurons. Neurons comprise an RC

circuit in parallel with a bank of current sources of variable conductance. In this chapter, the controller is itself a set of current sources connected to the neuron in parallel. This *port interconnection* philosophy is the same as that of passivity-based control. For the reference tracking task, the controller additionally includes a resistive coupling to ensure synchronization. To control multiple neurons in a network, we exploit the fact that the adaptive conductance controller is decentralised and can therefore be applied independently to each neuron.

Robustness of the adaptive observer

The controller of Chapter 3 relies on model estimates obtained online using an adaptive observer [16]. A version of the observer is used that assumes perfect modelling knowledge of the neurons' internal dynamics, an assumption that we have relaxed in Chapter 4. In order to quantify the observer's robustness to modelling error, we define a fixed scenario in which the observer must learn all of a neuron's maximal conductances, some of which are time-varying. Robustness is measured using the rms value of the observer error over the duration of the scenario. We have demonstrated that robustness can be improved by replacing the adaptive observer with its distributed version [13], which neglects second-order information about the statistics of the parameter estimates. Robustness is further improved by introducing redundancy to the observer's model of the neuron. To ensure stability of the resulting observer, we regularise the variance of the redundant parameter estimates by introducing a consensus term to the parameter update law.

Neuromorphic control of the pendulum

In Chapter 5 we have controlled a pendulum to track an event-based reference signal generated by an automaton. The signal is a temporal sequence of sensor and actuator events that represents the desired behaviour. This synchronization task is solved with an event-based, neuromorphic control architecture. Open-loop entrainment is used to control the pendulum in the overdamped regime (where the behaviour is monotonic) as well as in the underdamped regime (where a bistability is present). In closed-loop, phase control is used to enlarge the basin of attraction of large oscillations, and adaptive control is used to tune the behaviour of small oscillations.

6.2 Future work

The work of this thesis is theoretical and simulation-based. One starting point for future work would therefore be to implement the proposed control frameworks in both neurotechnological and neuromorphic settings.

Neurotechnology

The approach of Chapters 2 and 3 can be implemented using *dynamic clamp*, a well-established




Fig. 6.2.1 An illustration of the ‘inverted pendulum’ walking down a slope. Reproduced from [39].

technology that permits the creation of hybrid neuronal networks comprising both biological and artificial neurons [86]. While this method is only feasible for small networks, emerging technology promises to open up similar capabilities at larger scales [18, 79]. This would permit advances in, among other fields, neuromodulation therapy that treat disorders such as Parkinson’s [56]. To make optimal use of our framework at larger scales, further theoretical progress will be required: as discussed in Section 3.3, exact voltage measurements could be replaced with coarser information, and there is scope to explore multi-scale and hierarchical adaptive controllers.

Neuromorphic Control

A natural next step for the neuromorphic controller of Chapter 5 is to build the proposed system to control a physical pendulum. Beyond this, the approach of the chapter can be applied to a wide range of mechanical systems. One need not stray too far from the pendulum itself to find familiar robotic or biomechanical rhythms. For example, the ‘inverted pendulum’ model lies at the foundation of the literature on walking [39, 58]. It consists of two rigid rods representing the legs, each with a point mass m at the bottom (the ‘feet’) and connected at the top by a point mass M (the ‘hips’). Fig. 6.2.1 shows this model walking unpowered down a slope, and Fig. 6.2.2 shows a single step. At the end of each step, when the swing leg hits the ground (‘heelstrike’), the swing and stance legs are swapped and the motion repeats. Note that, during each step, the swing leg will pass briefly through the ground. This ‘scuffing’ is ignored, and the leg is permitted to pass through without collision. Scuffing can be eliminated by introducing ‘knees’, specifically by adding pin joints along the legs so as to divide each leg into two rigid rods.




Fig. 6.2.2 A single step of the inverted pendulum model walking down a slope. Reproduced from [39].

One way to drive the inverted pendulum along flat ground is with a motor attached to the hips [17, 50]. Such a motor could be controlled by a bursting neuron. Fig. 6.2.3 demonstrates this on a version of the model having feet that are curved rather than point masses (a simple way to improve stability [55]).

Introducing compliance to this system provides additional functionality. If the rigid legs are replaced with springs, one obtains a two-state automaton that can run as well as walk [41]. Indeed, compliance is an important aspect of embodied robotics and therefore a major topic of interest for future work. As pulse-based actuation is locally stiff but on average compliant (because the motors are inactive much of the time), the approach of this thesis suggests a path to overcoming the stiffness-compliance trade-off [85] that limits the accuracy of compliant systems [11]. This could be of particular value in soft robotics [23], where neuromorphics is expected to play a significant role [83].

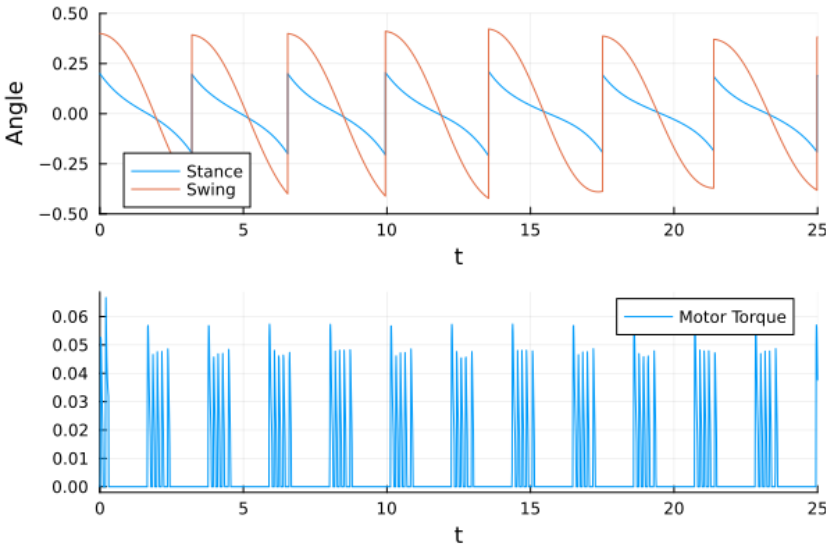


Fig. 6.2.3 A biped driven by a single bursting neuron that provides torque to the stance leg. At the end of each step the stance and swing legs are swapped over (hence the sudden jumps in angle).

References

- [1] Åarzen, K.-E. (1999). A simple event-based pid controller. *IFAC Proceedings Volumes*, 32(2):8687–8692.
- [2] Abu-Hassan, K., Taylor, J. D., Morris, P. G., Donati, E., Bortolotto, Z. A., Indiveri, G., Paton, J. F., and Nogaret, A. (2019). Optimal solid state neurons. *Nature communications*, 10(1):5309.
- [3] Angelidis, E., Buchholz, E., Arreguit, J., Rougé, A., Stewart, T., von Arnim, A., Knoll, A., and Ijspeert, A. (2021). A spiking central pattern generator for the control of a simulated lamprey robot running on spinnaker and loihi neuromorphic boards. *Neuromorphic Computing and Engineering*, 1(1):014005.
- [4] Aranda-Escolastico, E., Guinaldo, M., Heradio, R., Chacon, J., Vargas, H., Sánchez, J., and Dormido, S. (2020). Event-based control: A bibliometric analysis of twenty years of research. *IEEE Access*, 8:47188–47208.
- [5] Arbelaiz, J., Franci, A., Leonard, N. E., Sepulchre, R., and Bamieh, B. (2024). Excitable crawling. *arXiv preprint arXiv:2405.20593*.
- [6] Åström, K. J. and Wittenmark, B. (2008). *Adaptive Control*. Dover Publications, Mineola, NY, 2nd edition.
- [7] Åström, K. J. and Wittenmark, B. (2013). *Adaptive control*. Courier Corporation.
- [8] Bin, M., Marconi, L., and Sepulchre, R. (2023). An internal model principle for open systems. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 1919–1924. IEEE.
- [9] Brown, T. G. (1914). On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. *The Journal of physiology*, 48(1):18.
- [10] Bucher, D., Haspel, G., Golowasch, J., and Nadim, F. (2015). Central pattern generators. *eLS*, pages 1–12.
- [11] Buchli, J. and Ijspeert, A. J. (2008). Self-organized adaptive legged locomotion in a compliant quadruped robot. *Autonomous Robots*, 25:331–347.
- [12] Burghi, T. B. (2024). Personal Communication.

- [13] Burghi, T. B., O’Leary, T., and Sepulchre, R. (2022). Distributed online estimation of biophysical neural networks. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 628–634. IEEE.
- [14] Burghi, T. B., Schoukens, M., and Sepulchre, R. (2020). System identification of biophysical neuronal models. In *59th IEEE Conference on Decision and Control*, pages 6180–6185, Jeju Island, Republic of Korea.
- [15] Burghi, T. B., Schoukens, M., and Sepulchre, R. (2021). Feedback identification of conductance-based models. *Automatica*, 123:109297.
- [16] Burghi, T. B. and Sepulchre, R. (2023). Adaptive observers for biophysical neuronal circuits. *IEEE Transactions on Automatic Control*.
- [17] Byl, K. and Tedrake, R. (2008). Approximate optimal control of the compass gait on rough terrain. In *2008 IEEE International Conference on Robotics and Automation*, pages 1258–1263. IEEE.
- [18] Chen, R., Canales, A., and Anikeeva, P. (2017). Neural recording and modulation technologies. *Nature Reviews Materials*, 2(2):1–16.
- [19] Davies, M., Wild, A., Orchard, G., Sandamirskaya, Y., Guerra, G. A. F., Joshi, P., Plank, P., and Risbud, S. R. (2021). Advancing neuromorphic computing with loihi: A survey of results and outlook. *Proceedings of the IEEE*, 109(5):911–934.
- [20] Dayan, P. and Abbott, L. F. (2001). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT Press, Cambridge, MA.
- [21] Dethier, J., Drion, G., Franci, A., and Sepulchre, R. (2015). A positive feedback at the cellular level promotes robustness and modulation at the circuit level. *Journal of Neurophysiology*, 114(4):2472–84.
- [22] DeWeerth, S. P., Nielsen, L., Mead, C. A., and Åström, K. J. (1991). A simple neuron servo. *IEEE Transactions on Neural Networks*, 2(2):248–251.
- [23] Dou, W., Zhong, G., Cao, J., Shi, Z., Peng, B., and Jiang, L. (2021). Soft robotic manipulators: Designs, actuation, stiffness tuning, and sensing. *Advanced Materials Technologies*, 6(9):2100018.
- [24] Doyle, J. (1996). Robust and optimal control. In *Proceedings of 35th IEEE Conference on Decision and Control*, volume 2, pages 1595–1598. IEEE.
- [25] Drion, G., Dethier, J., Franci, A., and Sepulchre, R. (2018). Switchable slow cellular conductances determine robustness and tunability of network states. *PLoS computational biology*, 14(4):e1006125.
- [26] Drion, G., Franci, A., and Sepulchre, R. (2019). Cellular switches orchestrate rhythmic circuits. *Biological Cybernetics*, 113(1):71–82.
- [27] Drion, G., O’Leary, T., Dethier, J., Franci, A., and Sepulchre, R. (2015a). Neuronal behaviors: A control perspective. In *54th IEEE Conference on Decision and Control*, pages 1923–1944.

- [28] Drion, G., O’Leary, T., and Marder, E. (2015b). Ion channel degeneracy enables robust and tunable neuronal firing rates. *Proceedings of the National Academy of Sciences*, 112(38):E5361–E5370.
- [29] Efimov, D., Sacré, P., and Sepulchre, R. (2009). Controlling the phase of an oscillator: a phase response curve approach. In *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 7692–7697. IEEE.
- [30] Eppinger, B., Hämmerer, D., and Li, S.-C. (2011). Neuromodulation of reward-based learning and decision making in human aging. *Annals of the New York Academy of Sciences*, 1235(1):1–17.
- [31] Ermentrout, G. B. and Terman, D. H. (2010). *Mathematical Foundations of Neuroscience*. Springer, New York.
- [32] Fernandez Lorden, C. (2023). Neuromorphic control of embodied central pattern generators. *Université de Liège, Liège, Belgique*. (Unpublished master’s thesis), <https://matheo.uliege.be/handle/2268.2/18256>.
- [33] Fradkov, A. and Markov, A. (1997). Adaptive synchronization of chaotic systems based on speed gradient method and passification. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 44(10):905–912.
- [34] Fradkov, A., Nijmeijer, H., and Markov, A. (2000). Adaptive observer-based synchronization for communication. *International Journal of Bifurcation and Chaos*, 10(12):2807–2813.
- [35] Franci, A., Drion, G., and Sepulchre, R. (2017). Robust and tunable bursting requires slow positive feedback. *Journal of Neurophysiology*, 119(3):1222–1234.
- [36] Franci, A., Drion, G., Seutin, V., and Sepulchre, R. (2013). A Balance Equation Determines a Switch in Neuronal Excitability. *PLOS Computational Biology*, 9(5):e1003040.
- [37] Franci, A. and Sepulchre, R. (2014). Realization of nonlinear behaviors from organizing centers. In *53rd IEEE Conference on Decision and Control*, pages 56–61.
- [38] Gallego, G., Delbrück, T., Orchard, G., Bartolozzi, C., Taba, B., Censi, A., Leutenegger, S., Davison, A. J., Conradt, J., Daniilidis, K., et al. (2020). Event-based vision: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(1):154–180.
- [39] Garcia, M., Chatterjee, A., Ruina, A., and Coleman, M. (1998). The simplest walking model: stability, complexity, and scaling.
- [40] Garofalo, G., Ott, C., and Albu-Schäffer, A. (2012). Walking control of fully actuated robots based on the bipedal slip model. In *2012 IEEE International Conference on Robotics and Automation*, pages 1456–1463. IEEE.
- [41] Geyer, H., Seyfarth, A., and Blickhan, R. (2006). Compliant leg behaviour explains basic dynamics of walking and running. *Proceedings of the Royal Society B: Biological Sciences*, 273(1603):2861–2867.

- [42] Golowasch, J. and Marder, E. (1992). Ionic currents of the lateral pyloric neuron of the stomatogastric ganglion of the crab. *Journal of Neurophysiology*, 67(2):318–331.
- [43] Hauser, H., Ijspeert, A. J., Füchslin, R. M., Pfeifer, R., and Maass, W. (2011). Towards a theoretical foundation for morphological computation with compliant bodies. *Biological cybernetics*, 105:355–370.
- [44] Haykin, S. S. (2002). *Adaptive filter theory*. Pearson Education India.
- [45] Heaven, D. et al. (2019). Why deep-learning ais are so easy to fool. *Nature*, 574(7777):163–166.
- [46] Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, 117(4):500–544.
- [47] Hogan, N. (1985). Impedance Control: An Approach to Manipulation: Part I—Theory. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):1–7.
- [48] Homer, M. L., Nurmikko, A. V., Donoghue, J. P., and Hochberg, L. R. (2013). Sensors and decoding for intracortical brain computer interfaces. *Annual review of biomedical engineering*, 15:383–405.
- [49] Iida, F. and Giardina, F. (2023). On the timescales of embodied intelligence for autonomous adaptive systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 6(1):95–122.
- [50] Iida, F. and Tedrake, R. (2010). Minimalistic control of biped walking in rough terrain. *Autonomous Robots*, 28:355–368.
- [51] Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks*, 21(4):642–653.
- [52] Ijspeert, A. J., Crespi, A., and Cabelguen, J.-M. (2005). Simulation and robotics studies of salamander locomotion: applying neurobiological principles to the control of locomotion in robots. *Neuroinformatics*, 3:171–195.
- [53] Ijspeert, A. J., Crespi, A., Ryczko, D., and Cabelguen, J.-M. (2007). From swimming to walking with a salamander robot driven by a spinal cord model. *science*, 315(5817):1416–1420.
- [54] Izhikevich, E. M. (2007). *Dynamical Systems in Neuroscience*. MIT Press, Cambridge, MA.
- [55] Jeon, Y., Park, Y.-s., and Park, Y. (2013). A study on stability of limit cycle walking model with feet: Parameter study. *International Journal of Advanced Robotic Systems*, 10(1):49.
- [56] Johnson, M. D., Lim, H. H., Netoff, T. I., Connolly, A. T., Johnson, N., Roy, A., Holt, A., Lim, K. O., Carey, J. R., Vitek, J. L., et al. (2013). Neuromodulation for brain disorders: challenges and opportunities. *IEEE Transactions on Biomedical Engineering*, 60(3):610–624.

- [57] Keener, J. and Sneyd, J. (2009). *Mathematical Physiology*, volume 8/1. Springer, New York, NY, 2 edition.
- [58] Kuo, A. D. (2001). A simple model of bipedal walking predicts the preferred speed–step length relationship. *J. Biomech. Eng.*, 123(3):264–269.
- [59] Kuo, A. D. (2002). The relative roles of feedforward and feedback in the control of rhythmic movements. *Motor control*, 6(2):129–145.
- [60] Lakatos, D., Petit, F., and Albu-Schäffer, A. (2014). Nonlinear oscillations for cyclic movements in human and robotic arms. *IEEE Transactions on Robotics*, 30(4):865–879.
- [61] Lakatos, P., Gross, J., and Thut, G. (2019). A new unifying account of the roles of neuronal entrainment. *Current Biology*, 29(18):R890–R905.
- [62] Lee, H., Lee, J., Keppler, M., and Oh, S. (2024). Robust elastic structure preserving control for high impedance rendering of series elastic actuator. *IEEE Robotics and Automation Letters*, 9(4):3601–3608.
- [63] Lewis, M. A., Tenore, F., and Etienne-Cummings, R. (2005). Cpg design using inhibitory networks. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pages 3682–3687. IEEE.
- [64] Liberzon, D., Morse, A. S., and Sontag, E. D. (2002). Output-input stability and minimum-phase nonlinear systems. *IEEE Transactions on Automatic Control*, 47(3):422–436.
- [65] Liu, S.-C. and Delbruck, T. (2010). Neuromorphic sensory systems. *Current opinion in neurobiology*, 20(3):288–295.
- [66] Lohmiller, W. and Slotine, J.-J. E. (1998). On Contraction Analysis for Non-linear Systems. *Automatica*, 34(6):683–696.
- [67] Mainen, Z. and Sejnowski, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506.
- [68] Marder, E. (2012). Neuromodulation of neuronal circuits: back to the future. *Neuron*, 76(1):1–11.
- [69] Marder, E. and Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, 11(23):R986–R996.
- [70] Marder, E., Kedia, S., and Morozova, E. O. (2022). New insights from small rhythmic circuits. *Current opinion in neurobiology*, 76:102610.
- [71] Marder, E., O’Leary, T., and Shruti, S. (2014). Neuromodulation of Circuits with Variable Parameters: Single Neurons and Small Circuits Reveal Principles of State-Dependent and Robust Neuromodulation. *Annual Review of Neuroscience*, 37(1):329–346.
- [72] Marder, E. and Thirumalai, V. (2002). Cellular, synaptic and network effects of neuromodulation. *Neural Networks*, 15(4-6):479–493.
- [73] Matsuoka, K. (1987). Mechanisms of frequency and pattern control in the neural rhythm generators. *Biological cybernetics*, 56(5):345–353.

- [74] Mead, C. (1989). *Analog VLSI and neural systems*. Addison-Wesley Longman Publishing Co., Inc.
- [75] Mees, A. and Chua, L. (1979). The hopf bifurcation theorem and its applications to nonlinear oscillations in circuits and systems. *IEEE transactions on circuits and systems*, 26(4):235–254.
- [76] Meliza, C. D., Kostuk, M., Huang, H., Nogaret, A., Margoliash, D., and Abarbanel, H. D. (2014). Estimating parameters and predicting membrane voltages with conductance-based neuron models. *Biological cybernetics*, 108:495–516.
- [77] Mengaldo, G., Renda, F., Brunton, S. L., Bächer, M., Calisti, M., Duriez, C., Chirikjian, G. S., and Laschi, C. (2022). A concise guide to modelling the physics of embodied intelligence in soft robotics. *Nature Reviews Physics*, 4(9):595–610.
- [78] Moos, J., Hansel, K., Abdulsamad, H., Stark, S., Clever, D., and Peters, J. (2022). Robust reinforcement learning: A review of foundations and recent advances. *Machine Learning and Knowledge Extraction*, 4(1):276–315.
- [79] Ng, K. A., Greenwald, E., Xu, Y. P., and Thakor, N. V. (2016). Implantable neurotechnologies: a review of integrated circuit neural amplifiers. *Medical & biological engineering & computing*, 54:45–62.
- [80] Nogaret, A., Meliza, C. D., Margoliash, D., and Abarbanel, H. D. (2016). Automatic construction of predictive neuron models through large scale assimilation of electrophysiological data. *Scientific reports*, 6(1):32749.
- [81] O’Leary, T., Williams, A. H., Franci, A., and Marder, E. (2014). Cell types, network homeostasis, and pathological compensation from a biologically plausible ion channel expression model. *Neuron*, 82(4):809–821.
- [82] Ortega, R., van der Schaft, A. J., Mareels, I., and Maschke, B. (2001). Putting energy back in control. *IEEE Control Systems Magazine*, page 17.
- [83] Park, H.-L., Lee, Y., Kim, N., Seo, D.-G., Go, G.-T., and Lee, T.-W. (2020). Flexible neuromorphic electronics for computing, soft robotics, and neuroprosthetics. *Advanced Materials*, 32(15):1903558.
- [84] Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093.
- [85] Pratt, G. A., Williamson, M. M., Dillworth, P., Pratt, J., and Wright, A. (1997). Stiffness isn’t everything. In *Experimental Robotics IV: The 4th International Symposium, Stanford, California, June 30–July 2, 1995*, pages 253–262. Springer.
- [86] Prinz, A. A., Abbott, L. F., and Marder, E. (2004). The dynamic clamp comes of age. *Trends in Neurosciences*, 27(4):218–224.
- [87] Rahimi, A. and Recht, B. (2007). Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.

- [88] Ribar, L. (2020). *Synthesis of neuromorphic circuits with neuromodulatory properties*. PhD thesis.
- [89] Ribar, L. and Sepulchre, R. (2019). Neuromodulation of Neuromorphic Circuits. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 66(8):3028–3040.
- [90] Ribar, L. and Sepulchre, R. (2021). Neuromorphic control: Designing multiscale mixed-feedback systems. *IEEE Control Systems Magazine*, 41(6):34–63.
- [91] Russo, G., Di Bernardo, M., and Sontag, E. (2010). Global entrainment of transcriptional systems to periodic inputs. *PLoS Computational Biology*, 6(4):e1000739.
- [92] Ryu, H. X. and Kuo, A. D. (2021). An optimality principle for locomotor central pattern generators. *Scientific Reports*, 11(1):13140.
- [93] Sacre, P. and Sepulchre, R. (2014). Sensitivity analysis of oscillator models in the space of phase-response curves: Oscillators as open systems. *IEEE Control Systems Magazine*, 34(2):50–74.
- [94] Sarpeshkar, R. (1998). Analog versus digital: extrapolating from electronics to neurobiology. *Neural computation*, 10(7):1601–1638.
- [95] Schmetterling, R., Burghi, T. B., and Sepulchre, R. (2022). Adaptive conductance control. *Annual Reviews in Control*.
- [96] Schmetterling, R., Burghi, T. B., and Sepulchre, R. (2023). Robust online estimation of biophysical neural circuits. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 703–708. IEEE.
- [97] Schmetterling, R., Forni, F., Franci, A., and Sepulchre, R. (2024). Neuromorphic control of a pendulum. *IEEE Control Systems Letters*.
- [98] Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., and Plank, J. S. (2017). A survey of neuromorphic computing and neural networks in hardware. *arXiv preprint arXiv:1705.06963*.
- [99] Sepulchre, R. (2022a). Clocks and rhythms [from the editor]. *IEEE Control Systems Magazine*, 42(5):4–5.
- [100] Sepulchre, R. (2022b). Spiking control systems. *Proceedings of the IEEE*, 110(5):577–589.
- [101] Sepulchre, R. (2023). Principled internal models [from the editor]. *IEEE Control Systems Magazine*, 43(2):3–4.
- [102] Sepulchre, R., Drion, G., and Franci, A. (2018). Excitable Behaviors. In *Emerging Applications of Control and Systems Theory*, pages 269–280. Springer International Publishing.
- [103] Sepulchre, R., Drion, G., and Franci, A. (2019a). Control across scales by positive and negative feedback. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):89–113.

- [104] Sepulchre, R., O’Leary, T., Drion, G., and Franci, A. (2019b). Control by neuromodulation: A tutorial. In *2019 18th European Control Conference (ECC)*, pages 483–497.
- [105] Serrano-Gotarredona, T. and Linares-Barranco, B. (1999). Systematic width-and-length dependent cmos transistor mismatch characterization and simulation. *Analog Integrated Circuits and Signal Processing*, 21(3):271–296.
- [106] Simoni, M. F. and DeWeerth, S. P. (2006). Two-dimensional variation of bursting properties in a silicon-neuron half-center oscillator. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 14(3):281–289.
- [107] Sontag, E. (2010). Contractive systems with inputs. In Willems, J. C., Hara, S., Ohta, Y., and Fujioka, H., editors, *Perspectives in Mathematical System Theory, Control, and Signal Processing: A Festschrift in Honor of Yutaka Yamamoto on the Occasion of his 60th Birthday*, pages 217–228. Springer Berlin Heidelberg, Berlin, Heidelberg.
- [108] Sorrell, E., Rule, M. E., and O’Leary, T. (2021). Brain–Machine Interfaces: Closed-Loop Control in an Adaptive System. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):167–189.
- [109] Stan, G. B. and Sepulchre, R. (2007). Analysis of Interconnected Oscillators by Dissipativity Theory. *IEEE Transactions on Automatic Control*, 52(2):256–270.
- [110] Strogatz, S. H. (1996). Nonlinear dynamics and chaos.
- [111] Sutton, R. S., Barto, A. G., and Williams, R. J. (1992). Reinforcement learning is direct adaptive optimal control. *IEEE control systems magazine*, 12(2):19–22.
- [112] Taghvaei, A., Mehta, P. G., and Georgiou, T. T. (2022). Optimality vs stability trade-off in ensemble kalman filters. *IFAC-PapersOnLine*, 55(30):335–340.
- [113] Tang, E. and Bassett, D. S. (2018). Colloquium: Control of dynamics in brain networks. *Reviews of modern physics*, 90(3):031003.
- [114] Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neural networks*, 111:47–63.
- [115] Van Geit, W., De Schutter, E., and Achard, P. (2008). Automated neuron model optimization techniques: a review. *Biological cybernetics*, 99:241–251.
- [116] Wang, W. and Slotine, J.-J. E. (2004). On partial contraction analysis for coupled nonlinear oscillators. *Biological Cybernetics*, 92(1):38–53.
- [117] Wieland, P., Schmidt, G. S., Sepulchre, R., and Allgöwer, F. (2010). Phase synchronization through entrainment by a consensus input. In *49th IEEE Conference on Decision and Control (CDC)*, pages 534–539. IEEE.
- [118] Willems, J. C. (2007). The Behavioral Approach to Open and Interconnected Systems. *IEEE Control Systems Magazine*, 27(6):46–99.
- [119] Williamson, M. M. (1999). *Robot arm control exploiting natural dynamics*. PhD thesis, Massachusetts Institute of Technology.

-
- [120] Yang, T., Mehta, P. G., and Meyn, S. P. (2013). Feedback particle filter. *IEEE transactions on Automatic control*, 58(10):2465–2480.

