1 Introduction

Understanding the principles behind how the brain works is a challenging scientific quest, driven in part by curiosity and in part by the urge to develop drugs and other treatment strategies for brain disorders.

We know today that the basic processing units of the brain are neurons, or nerve cells, that communicate via what is known as *action potentials* (APs). An AP is a brief, high-amplitude fluctuation in the neural membrane potential, and it can be recorded with an intracellularly inserted electrode. An AP also causes a brief fluctuation, or *spike*, in the electric potential in the extracellular space surrounding the active neuron. Extracellular spikes are typically obtained by high-pass filtering the extracellular potentials recorded inside the brain tissue. The signal used to detect spikes is often called the *multi-unit activity* (MUA) and is usually believed to reflect the aggregate AP firing of a relatively small number of nearby neurons.

Depending on what electrode one uses, where one positions it, and how one filters the recorded signals, one can – instead of recording spikes – choose to record more mesoscopic or macroscopic extracellular signals, reflecting the activity of larger populations of neurons. In addition to the spikes, commonly recorded extracellular signals include the *local field potential* (LFP), which is the low-frequency part of extracellular potentials recorded inside brain tissue;[1] the *electrocorticographic* (ECoG) signals, which are electric potentials recorded by electrodes placed on the cortical surface; the *electroencephalographic* (EEG) signals, which are electric potentials recorded by electrodes placed on the scalp; and the *magnetoencephalographic* (MEG) signals, which are magnetic fields recorded immediately outside the head. These different electric and magnetic measurement modalities are illustrated in Figure 1.1. This figure does not show real experimental data, but instead shows the results from a computer simulation based on the methodology and computational schemes that will be presented throughout this book. The simulation illustrates how a 200 ms period of activity in a single biophysically detailed model neuron contributes to the various brain signals.

In addition to the electric and magnetic measurement modalities illustrated in Figure 1.1, brain activity can be probed with optical methods such as *voltage-sensitive dye imaging* (VSDI) measuring membrane potentials, *intrinsic optical signal imaging*

[1] While others have defined the LFP to simply mean the extracellular potential recorded outside neurons, this book defines the LFP to be the low-frequency part of the potential. With this definition, LFP could thus more aptly be thought of as an acronym for "low-frequency potential" (Głąbska et al. 2017).

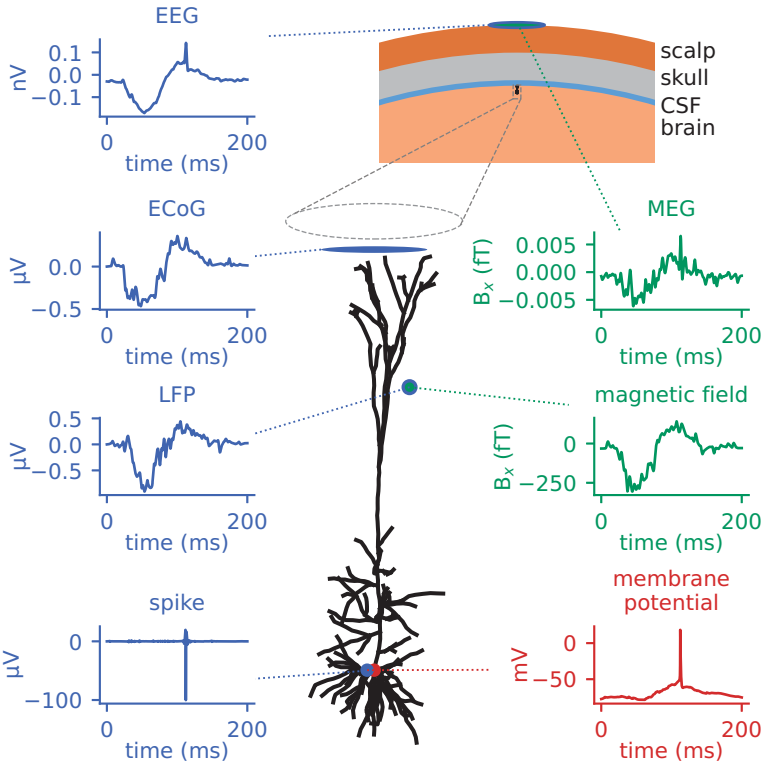

Figure 1.1 Multimodal modeling of a neuron receiving synaptic inputs and firing an action potential. Simulations were done on a multicompartment model of a layer 5 pyramidal cell developed by Hay et al. (2011). Extracellular signals (produced by this single neuron) were computed using methods that will be presented throughout this book. Code available via www.cambridge.org/electricbrainsignals.

(IOSI) measuring hemodynamics, and *two-photon* Ca^{2+} *imaging* (2PCI) measuring changes in intracellular Ca^{2+} concentrations. Further, *functional magnetic resonance imaging* (fMRI) measures blood flow and blood oxygenation in brain tissue, while *positron emission tomography* (PET) measures metabolic activity.

Taken together, the techniques listed above probe brain activity on a wide spectrum of spatial and temporal scales (Figure 1.2A). The focus of this book will be on the electric and magnetic measurement modalities summarized in Figure 1.1. Other measurement modalities will not be covered further, but for an overview of the physical principles governing other techniques, see Brette & Destexhe (2012).

The last decades have seen the development of better electrode systems, allowing experimentalists to perform electric measurements with steadily increasing numbers of electrode contacts and improved spatial resolution. Cheaper hard drives have also facilitated en masse data storage at higher temporal resolutions. In parallel with the improvements in measurement techniques comes the need for better theoretical models

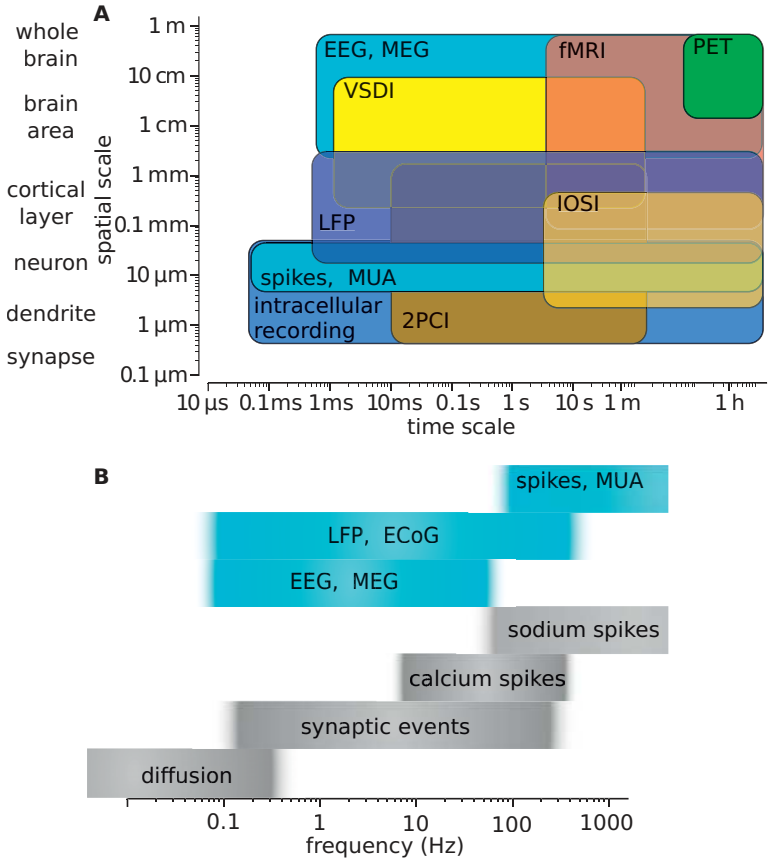

Figure 1.2 Overview of measurements of neural activity A: Spatiotemporal scale covered by various measurements. **B**: Main frequency bands of various extracellular measurements (blue boxes) and various events/processes (grey boxes).

and computational tools that can help us interpret what these measurements tell us. Spike recordings are in principle straightforward to interpret in terms of underlying neural activity: a spike means there is an AP in a nearby neuron. The meso- and macroscopic signals are likely to contain more information about what the brain is doing on the systems level. However, interpreting these systems-level measurements in terms of what they tell us about the underlying network activity is challenging and has largely been done in a qualitative manner only.

We note that the terms "spike" and "AP" are often used interchangeably in the literature. In this book, we will mostly preserve the term "spike" for the extracellular signature of the intracellular AP. However, we do make some exceptions from this "rule" to accommodate established terminology – for example, when we refer to "spiking neuron models," which is the common term for models that generate APs or at least predict the times of their occurrence, or in reference to a "spike train," which refers to a sequence of APs.

1.1 Forward Modeling of Extracellular Signals

Over the last decades, the field of computational neuroscience has grown both in size and scope, and the use of mathematics is now ubiquitous in neuroscience. The types of modeling pursued typically belong to one of three categories (Dayan & Abbott 2001):

- In *mechanistic* modeling, the goal is to explain a phenomenon in terms of the physical processes that give rise to it. This is the traditional physics approach. Its most prominent application in computational neuroscience is the Hodgkin-Huxley model (Hodgkin & Huxley 1952*a*), where the AP is modeled based on the electrical properties of the cell membrane and ion channels.
- In *descriptive* or *statistical* modeling, the goal is rather to summarize experimental data compactly. While such models may be motivated by neurobiological insights, the goal is to account mathematically for a phenomenon, not to explain it. An example of such a model in neuroscience is the Gabor-filter description of receptive fields of so-called simple cells in the primary visual cortex in mammals (Dayan & Abbott 2001, chapter 2).
- In *normative* or *interpretive* modeling, the goal is to explain a phenomenon in terms of its function. Normative modeling is unique to biological systems and is not pursued in physics. A question of *why* a stone falls to the ground is not fruitful, as the law of gravity is not set up to perform a certain task. Biological systems, however, have developed under evolutionary pressure, and the "why" question is sensible. An example of normative modeling in neuroscience is the use of information theory to explain why salient features of receptive fields in the early visual system are well suited to convey information about the world (Dayan & Abbott 2001, chapter 4). Here the basic idea is that these receptive fields are optimized to convey the maximum amount of information about the natural world.

In this book, we will mostly deal with the mechanistic modeling of extracellular electric and magnetic brain signals based on biophysics-based models of neurons and networks. Although simplistic, physics-based neuron models existed earlier (Lapicque 1907), mechanistic modeling in neuroscience arguably started in full with the pioneering works of Alan Hodgkin and Andrew Huxley on axonal action-potential generation and propagation (Hodgkin & Huxley 1952*a*). Together with the work of Wilfred Rall on how neuronal dendrites integrate synaptic inputs (Rall 1959, Segev et al. 1995), the work by Hodgkin and Huxley established what still today remains the standard formalism for the biophysical modeling of neurons. These works typically form a core in courses in computational neuroscience lectured at universities and are covered in many excellent textbooks (see e.g. Koch (1999), Dayan & Abbott (2001), Gerstner et al. (2014), and Sterratt et al. (2023)).

Despite groundbreaking works by Rall (1962) and Holt & Koch (1999) providing the physical foundation, the mechanistic modeling of the link between neural activity and what is measured extracellularly in different types of experiments has received relatively little attention in the computational neuroscience community. As a consequence, it has

been difficult to compare predictions from network modeling with experimental data other than spikes (the extracellular signatures of neuronal APs).

The type of mechanistic modeling of extracellular signals pursued in this book is often called *forward modeling*, since its purpose is to take us "forward" from a mechanistic model of neural activity to the corresponding extracellular measurement signal. In contrast, *inverse modeling* aims to estimate the underlying activity from a set of measurements. The signals that we focus on, and the typical signal-frequency ranges that they cover, are summarized in Figure 1.2B (blue bands). The figure also indicates which kinds of neural events (grey bands) are expected to contribute within the various frequency ranges.

A key reason for focusing on electric and magnetic extracellular signals is that a (forward) link from neural activity to these signals is well established by the so-called *volume-conductor* theory. Another reason is that the high time resolution of these signals, typically on the order of a millisecond or so, makes their quantitative modeling particularly important since much of the relevant neural activity in the brain takes place on this millisecond time scale. A third reason is the rapid development of new electrode systems and experimental techniques for measuring electric and magnetic brain signals.

We believe that forward modeling of extracellular signals has many uses, some of which are listed below:

- Forward modeling allows us to run controlled neural simulations to systematically explore how various aspects of neuronal activity are reflected in extracellular signals. Insights from forward modeling can be used (inversely) to interpret experimentally recorded signals in terms of what they tell us about the underlying neural activity, at least on a qualitative level. Many neuroscientists have the notion, for example, that LFPs recorded in a particular layer of cortex necessarily stem from neurons with somas in the same layer, while biophysics-based modeling shows that this is not correct.
- Forward modeling allows for a quantitative comparison between network model predictions and experimentally recorded LFP, ECoG, EEG, and MEG signals. Such comparisons will in turn allow the data to be used (inversely) to estimate model parameters, facilitating the development of models that bridge the scales between the single-neuron level and the systems level in the brain (Einevoll et al. 2019).
- Forward models can be used as starting points when developing new (inverse) methods for data analysis – that is, the estimation of neural activity from experimental data. An example is the development of methods for estimating the current-source density (CSD) from multielectrode LFP recordings based on a corresponding forward model (Pettersen et al. 2006, Potworowski et al. 2012, Cserpan et al. 2017).
- Forward models allow us to produce simulation-based benchmarking data for the validation of data analysis methods, such as spike-sorting methods (Einevoll et al. 2012, Hagen et al. 2015, Buccino et al. 2018) or methods for estimating

current-source densities (CSD) from multielectrode LFPs recordings (Pettersen et al. 2008, Hagen et al. 2017). Another example is test-data generation for the validation of laminar population analysis (LPA) (Głąbska et al. 2016).

1.2 Overview of the Contents of This Book

In brief, this book presents (i) a biophysical theory for computing electric and magnetic brain signals, (ii) computations of such signals for various scenarios (mostly using computer simulations), and (iii) extractions of key insights from these computations.

Throughout the book, what we will refer to as the "MC+VC scheme" will hold a central position. This is the standard scheme for biophysics-based computation of electric and magnetic brain signals. The scheme will be defined properly later on, but in brief, it entails the computation of neural electric activity using multicompartment (MC) neuron models as well as the computation of extracellular signals using volume-conductor (VC) theory. Our main focus will be on electric signals, whereas magnetic signals will be covered more briefly. The reason for this is partly that we, the authors, have our main expertise in electric signals and partly that magnetic brain signals are well covered in the book by Ilmoniemi & Sarvas (2019).

Our focus will be on general principles of the generation of extracellular signals. We will, however, have a slight bias towards extracellular potentials generated by cortical neurons. For spikes (Chapter 7), the principled insights from considering cortical neurons should also directly apply to spikes from neurons in other brain areas, as there is nothing principally different in the action potential and spike generation of cortical neurons compared to other neurons. However, the LFP (Chapter 8) generated by cortical neurons and populations of such can be uniquely large due to the so-called open-field dendritic structure and layered organization of the dominant pyramidal neurons. Example results for cortical LFPs will thus be less representative for the various LFPs that can be recorded around the brain. The EEG (Chapter 9), ECoG (Chapter 10), and MEG signals (Chapter 11) are in any case expected to be dominated by contributions from cortical neurons due to their proximity to the recording devices so, for these signals, our choice of using cortical neurons as examples is warranted.

The book is outlined as follows. In the next five chapters (Chapters 2–6), we establish the theoretical foundation for computing brain signals, focusing mainly on extracellular potentials. In the first of these theory chapters (Chapter 2), we walk the reader through fundamental physical laws and concepts as well as define the MC+VC scheme (Section 2.6.2). The main theory for computing neurodynamics using multicompartment neuron models (the MC part) is introduced in Chapter 3. The volume-conductor theory (the VC part) that takes us (forward) from neurodynamics to predictions of extracellular signals is introduced in Chapter 4.

An important component in VC theory is the conductivity of brain tissue. In Chapter 4, we assume it to be constant since this makes the theory much simpler. In Chapter 5, we discuss this assumption and introduce ways to extend the VC theory to

more general cases where the conductivity depends on the position, direction, or signal frequency. The last theory chapter, Chapter 6, describes ways to implement the MC+VC framework for running numerical simulations on computers, as well as ways to reduce model complexity in order to obtain more efficient simulations.

Following the theory chapters is a series of chapters where we apply the defined theory to simulate the kinds of extracellular signals summarized in Figure 1.2B. In Chapter 7, we consider spikes, the extracellular signatures of neuronal APs. As they can normally just be picked up by electrodes in close proximity to the AP-firing neuron, spikes are the most local extracellular signals that we consider. In Chapter 8, we follow up with another (albeit less) local measure, namely the local field potential (LFP). The LFP is the low-frequency part of electric potential recorded in the grey matter of the brain. In Chapters 9 and 10, we follow up with simulations of extracellular potentials on the largest spatial scale, which are those recorded on the scalp (EEG) and cortical surface (ECoG) respectively. In all these chapters, the main aim is to provide insight into how various features of neuronal structure and activity, such as synaptic positions, neural morphology, firing frequency, and degree of synchrony within a network, are related to various features in the extracellular potential.

The prediction of magnetic brain signals, and the physical theory needed to do so, is covered separately in the rather brief Chapter 11. The focus is on the MEG signal (the magnetic signal recorded immediately outside the head), but we also briefly outline methods for computing magnetic fields inside the brain. The origin of the magnetic signals are the same as for the electric signals, and both can be simulated using the MC+VC scheme.

As implied earlier, we are mainly concerned with the forward computation of extracellular signals resulting from a given neural activity. Computing how extracellular signals affect neural activity is in principle described with the same physics. The governing theory for electric and magnetic stimulation is briefly described in Section 4.6.

Although it is fairly well established that the main source of extracellular brain signals are currents produced by neuronal activity, the signals may in principle contain additional contributions from glial activity or the diffusion of ions along extracellular concentration gradients. These additional sources are, under most circumstances, believed to be relatively small and to predominantly affect the very-slow frequency components of extracellular potentials. As there is no principal difference between neuronal and glial membrane currents, both can be computed with the same MC+VC framework, and when we make general references to neural current sources, they can include all cellular sources. Extracellular diffusive currents do, however, belong to another category and are not accounted for in the MC+VC framework. Theory for computing diffusion potentials is presented in Chapter 12, where we also give some estimates of their relative contributions to extracellular potentials.

In the final Chapter 13, we first describe some common misconceptions about extracellular potentials. Next, we discuss the general reliability of MC+VC-based modeling of electric and magnetic extracellular signals before we round off the book by listing some future applications of forward modeling.

1.3 Guide to Reading This Book

Different readers may have different motivations for consulting this book. It is largely a physics book, focusing on the physical and biophysical processes governing extracellular brain signals. However, we have made some effort to present the material in a way that should make it accessible not only to neuroscientists with strong schooling in physics and mathematics, but also to other neuroscientists trying to make sense of extracellularly recorded brain signals.

All readers are advised to read the basic theory chapter (Chapter 2), as key concepts and variables that will be used throughout the book are introduced there. Parts of this theory chapter will be challenging for readers without a physics background, but understanding all of it will not be necessary. The most important takeaway from Chapter 2 is the description of neural activity in terms of membrane current sources in Section 2.6.1 and especially the definition of the MC+VC scheme in Section 2.6.2. This standard scheme for computing extracellular brain signals will be used throughout most parts of the book.

The MC part, which is about how one uses multicompartment neural models to compute neural membrane currents, is described in further detail in Chapter 3. The first two sections (Sections 3.1–3.2) should cover the material that is most important. The VC part, about how one uses volume-conductor theory to compute the extracellular potentials resulting from the MC simulations, is described in further detail in Chapter 4. The main idea is established already in Section 4.1 of the chapter. Although the theory can be extended to more general cases, readers who understand the sections just suggested should have a good foundation for following the later parts of this book.

After following the instructions in the previous paragraphs, readers that consult this book for help in interpreting various kinds of extracellular recordings may choose to go directly to the relevant chapter. The book contains chapters devoted to computing and examining spikes (Chapter 7), LFPs (Chapter 8), EEGs (Chapter 9), ECoGs (Chapter 10), MEGs (Chapter 11), and diffusion potentials (Chapter 12). In many cases, we will be analyzing the extracellular signals by examining their frequency content, presented in terms of amplitude or power spectra. Readers that are not familiar with these concepts can get a brief introduction in Appendix F.

A more biophysically oriented reader will probably choose to delve deeper into the theoretical material in Chapters 2–5. When it comes to the biophysical modeling of neurons, the introduction in Chapter 3 is kept rather brief as this topic is covered extensively in several other books, such as those by Koch (1999), De Schutter (2009), or Sterratt et al. (2023). The focus of this book is rather on extracellular signals and on the tissue (medium) they propagate through. The later parts of Chapter 2 deal with the problems and approximations involved when applying electromagnetic theory to a complex medium such as brain tissue, while the equations for signal propagation in brain tissue are derived in Chapters 4 and 5. In Chapter 5, we review both the experimental and theoretical efforts to understand the properties of brain tissue and its electric conductivity.

Readers who consult the book to learn how to develop computational models of extracellular signals should get a solid theoretical platform through Chapters 2–4. Technical details useful for numerical implementation, as well as methods for computing extracellular signals using more simplified and computationally efficient models of networks of neurons, are found in Chapter 6.

1.4 Guide to Simulations, Codes, and Figures

Much of the insight into the nature of extracellular electric signals of neural origin presented throughout this book has been gained through computer-based simulations, such as those shown in Figure 1.1. This figure, as well as many others in this book, are based on simulations using the MC+VC scheme and using our own software tool LFPy (Hagen et al. 2018). In the following chapters, we use example simulations of this kind both to illustrate the methodology for the forward modeling of extracellular signals as well as to gain insights into how various features of neuronal activity affect various features of spikes and LFP, ECoG, EEG, and MEG signals. As a supplement to this book, the simulation codes used to produce the figures are freely available and hosted in an external code repository at https://github.com/LFPy/ElectricBrainSignals (Hagen & Ness 2023), also accessible via www.cambridge.org/electricbrainsignals. The codes are released under open, free-to-use licenses.

As we shall see later on, a neuron's contribution to the extracellular signal depends on its morphology, as well as on how synapses and ion channels are distributed onto this morphology. There are thus certain salient features in extracellular signals that can only be simulated with rather complex neuron models that account for morphological and biophysical detail. The simulation in Figure 1.1 was done using a biophysically detailed model of a rat cortical layer 5 pyramidal cell developed by Hay et al. (2011). This model, which we will simply refer to as the *Hay model*, is presented in further detail in Section 3.2.5. Throughout this book, the Hay model will serve as our default go-to model when we need a biophysically detailed neuron model. We will, however, also present simulations on simpler, generic neuron models, as well as on biophysically detailed models representing other types of neurons. These other models will be described briefly in the sections where they are first used.

Readers of this book are cordially encouraged to download the simulation codes used throughout this book, re-run them, and modify them in an exploratory manner. Playing around with the codes interactively while consulting this textbook for interpretations and explanations might be a good learning strategy. Demonstrations in this book include the effect of neuronal morphology, the effects of synapse distributions, and the effects of the distribution of different ion channels on extracellular signals. However, there are several uncharted paths for such ventures, and readers that want to pursue particular research questions may use the codes as a starting point for setting up their own simulations. Thus, our hope is that this book and these codes may also serve as a basis for future

research – for instance, aiding the interpretation of experimental recordings of neural activity.

In case the provided codes result in new publications, kindly cite this book, the relevant original publication, and Hagen & Ness (2023).

The main figure-generating files are set up as Jupyter[2]) notebooks and rely on the Python programming language,[3] with a few other package dependencies such as the aforementioned LFPy and neural simulation software such as NEURON (Hines & Carnevale 1997, Hines et al. 2009). The notebook codes may readily be employed in interactive cloud computing services such as the EBRAINS Collaboratory (www.ebrains.eu/tools/collaboratory).

[2] https://jupyter.org
[3] https://python.org