1

Introduction

Our brains evolved to control a complex biological device: our body. As we are finding out today, many millennia of evolutionary tinkering has made the brain a surprisingly versatile and adaptive system, to the extent that it can learn to control devices that are radically different from our body. Brain-computer interfacing, the subject of this book, is a new interdisciplinary field that seeks to explore this idea by leveraging recent advances in neuroscience, signal processing, machine learning, and information technology.

The idea of brains controlling devices other than biological bodies has long been a staple of science-fiction novels and Hollywood movies. However, this idea is fast becoming a reality: in the past decade, rats have been trained to control the delivery of a reward to their mouths, monkeys have moved robotic arms, and humans have controlled cursors and robots, all directly through brain activity.

What aspects of neuroscience research have made these advances possible? What are the techniques in computing and machine learning that are allowing brains to control machines? What is the current state-of-the-art in brain-computer interfaces (BCIs)? What limitations still need to be overcome to make BCIs more commonplace and useful for day-to-day use? What are the ethical, moral, and societal implications of BCIs? These are some of the questions that this book addresses.

The origins of BCI can be traced to work in the 1960s by Delgado (1969) and Fetz (1969). Delgado developed an implantable chip (which he called a "stimoceiver") that could be used to both stimulate the brain by radio and send electrical signals of brain activity by telemetry, allowing the subject to move about freely. In a now-famous demonstration, Delgado used the stimoceiver to stop a charging bull in its tracks by pressing a remote-control button that delivered electrical stimulation to the caudate nucleus in the basal ganglia region of the bull's brain. At around the same time, Fetz showed that monkeys can control the activity of single brain cells to control a meter needle and obtain food rewards (see Section 7.1.1). Slightly later, Vidal (1973) explored the use of scalp-recorded brain signals in humans to implement a simple noninvasive BCI based on "visually evoked potentials" (Section

6.2.4). The more recent surge of interest in BCIs can be attributed to a confluence of factors: faster and cheaper computers, advances in our knowledge of how the brain processes sensory information and produces motor output, greater availability of devices for recording brain signals, and more powerful signal processing and machine-learning algorithms.

The primary motivation for building BCIs today is their potential for restoring lost sensory and motor function. Examples include sensory prosthetic devices such as the cochlear implant for the deaf (Section 10.1.1) and retinal implant for the blind (Section 10.1.2). Other implants have been developed for deep brain stimulation (DBS) to treat the symptoms of debilitating diseases such as Parkinson's (Section 10.2.1). A parallel line of research has explored how signals from the brain could be used to control prosthetic devices such as prosthetic arms or legs for amputees and patients with spinal-cord injuries (e.g., Section 7.2.1), cursors and word spellers for communication by locked-in patients suffering from diseases such as ALS (amyotrophic lateral sclerosis) or stroke (Sections 7.2.3 and 9.1.4), and wheelchairs for paralyzed individuals (Section 12.1.6). More recently, researchers have begun exploring BCIs for able-bodied individuals for a host of applications (Chapter 12), ranging from gaming and entertainment to robotic avatars, biometric identification, and education. Whether BCIs will eventually become as commonplace as current human accessories for sensory and motor augmentation, such as cellular phones and automobiles, remains to be seen. Besides technological hurdles, there are a number of moral and ethical challenges that we as a society will need to address (Chapter 13).

The goal of this book is to serve as an introduction to the field of brain-computer interfacing. Figure 1.1 illustrates the components of a generic BCI. The aim is to translate brain activity into control commands for devices and/or stimulate the brain to provide sensory feedback or restore neurological function. One or more of the following processing stages are typically involved:

1. **Brain recording**: Signals from the brain are recorded using either invasive or noninvasive recording techniques.
2. **Signal processing**: Raw signals are preprocessed after acquisition (e.g., by bandpass filtering) and techniques for artifact reduction and feature extraction are used.
3. **Pattern recognition and machine learning**: This stage generates a control signal based on patterns in the input, typically using machine-learning techniques.
4. **Sensory feedback**: The control signal from the BCI causes a change in the environment (e.g., movement of a prosthetic arm or a wheelchair, change in the grip of a prosthetic hand). Some of these changes can be seen, heard, or felt by the user but in general, one can use sensors in the environment such as tactile sensors, force sensors, cameras, and microphones, and use the information from these sensors to provide direct feedback to the brain via stimulation.

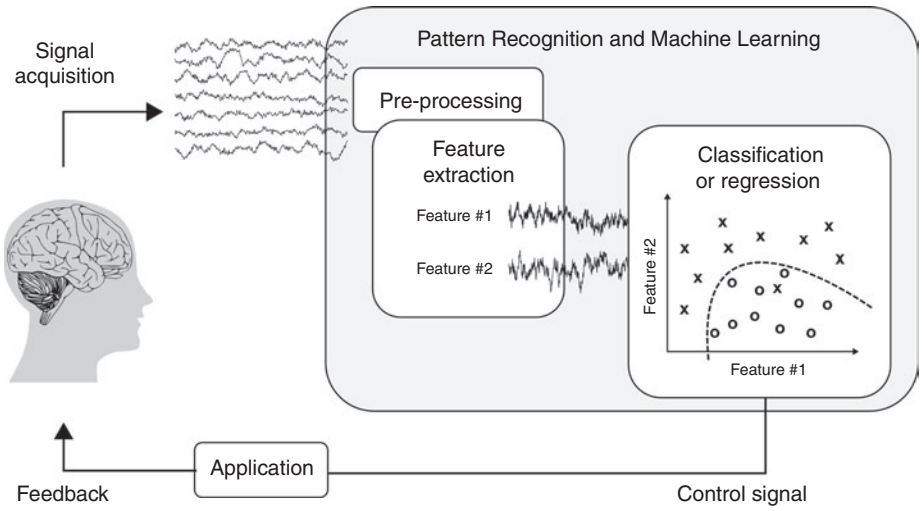

Figure 1.1. **Basic components of a brain-computer interface (BCI).** (Adapted from Rao and Scherer, 2010).

5. **Signal processing for stimulation**: Before stimulating a particular brain region, it is important to synthesize an activity pattern for stimulation that mimics the type of activity normally seen in the brain region and that will have the desired effect. This requires a good understanding of the brain area being stimulated and the use of signal processing (and potentially machine learning) to home in on the right stimulation patterns.
6. **Brain stimulation**: The stimulation pattern received from the signal processing component (5) is used in conjunction with invasive or noninvasive stimulation techniques to stimulate the brain.

It is clear from the stages of processing listed above that to begin building BCIs, one must have a background in at least four essential areas: basic neuroscience, brain recording and stimulating technologies, elementary signal processing, and basic machine-learning techniques. Often, beginners in BCI come with a background in one of these areas but usually not all of them. We therefore begin our journey into the world of BCIs with Part I (Background), which introduces the reader to basic concepts and methods in these four areas.

Part I

Background

2

Basic Neuroscience

Weighing in at about three pounds, the human brain is a marvel of evolutionary engineering. The brain transforms signals from millions of sensors located all over the body into appropriate muscle commands to enact a behavior suitable to the task at hand. This closed-loop, real-time control system remains unsurpassed by any artificially created system despite decades of attempts by computer scientists and engineers.

The brain's unique information processing capabilities arise from its massively parallel and distributed way of computing. The workhorse of the brain is a type of cell known as a *neuron*, a complex electrochemical device that receives information from hundreds of other neurons, processes this information, and conveys its output to hundreds of other neurons. Furthermore, the connections between neurons are plastic, allowing the brain's networks to adapt to new inputs and changing circumstances. This adaptive and distributed mode of computation sets the brain apart from traditional computers, which are based on the *von Neumann architecture* with a separate central processing unit, memory units, fixed connections between components, and a serial mode of computation.

In this chapter, we provide a primer on neuroscience. Starting from the biophysical properties of neurons, we explore how neurons communicate with each other, how they transmit information to other neurons via junctions called synapses, and how synapses are adapted in response to inputs and outputs. We then explore the network level architecture and anatomy of the brain, learning how different areas of the brain are specialized for different functions.

2.1 Neurons

A neuron is a type of cell that is generally regarded as the basic computational unit of the nervous system. As a crude approximation, the neuron can be regarded as a leaky bag of charged liquid. The membrane of a neuron is made up of a lipid bi-layer (Figure 2.1) that is impermeable except for openings called *ionic channels* that selectively allow the passage of particular kinds of ions.

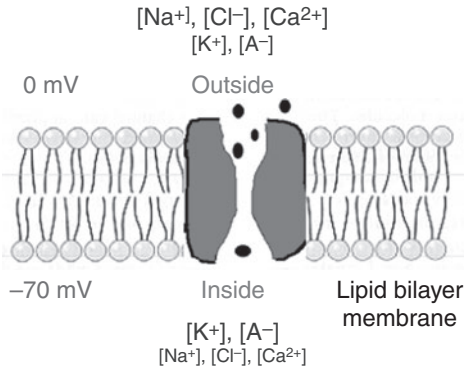

Figure 2.1. **The electrochemical dance of ions in a neuron.** The diagram depicts the larger concentration of sodium, chloride, and calcium ions outside the neuron and the larger concentration of potassium ions and anions inside the neuron (maintained by active pumps), resulting in a "resting" potential difference of approximately −70mv across the lipid bi-layer membrane. Proteins known as ionic channels, which are embedded in the membrane, act as gates regulating the flow of ions into and out of the neuron.

Neurons reside in an aqueous medium with a larger concentration of sodium (Na+), chloride (Cl^-), and calcium (Ca^2+) on the outside of the cell and a greater concentration of potassium (K+) and organic anions (A^-) inside the cells (Figure 2.1). As a result of this imbalance, there is a potential difference of approximately −65 to −70 mV across the neuron's membrane when the neuron is at rest. There exist active pumps that work to maintain this potential difference by expending energy.

2.2 Action Potentials or Spikes

When the neuron receives sufficiently strong inputs from other neurons (see Section 2.4 below), a cascade of events is triggered: there is a rapid influx of Na+ ions into the cell, causing the membrane potential to rise rapidly, until the opening of K+ channels triggers the outflux of K+ ions, causing a drop in the membrane potential. This rapid rise and fall of the membrane potential is called an *action potential* or *spike* (Figure 2.2), and represents the dominant mode of communication between one neuron and another. The spike is an all-or-one stereotyped event with little or no information in the shape of the spike itself – information is thought to be conveyed instead by the *firing rate* (number of spikes per second) and/or the timing of spikes. Neurons are therefore often modeled as emitting a 0 or 1 digital output. Similarly, in extracellular recordings typically done in awake animals (Section 3.1.1), a spike is often represented as a short vertical bar at the time the spike occurred.

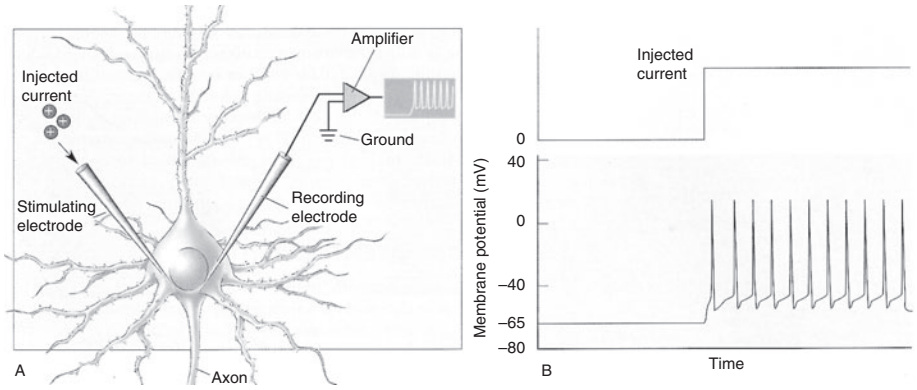

Figure 2.2. **Generation of spikes or action potentials.** (A) depicts the experimental procedure of injecting a current (positive ions) into the cell body of a neuron using a stimulating electrode and recording the change in membrane potential of the cell using a recording electrode. (B) shows the result of injecting a sufficiently large amount of current, which results in a sequence of spikes or action potentials. Each spike has a stereotypical shape that rises rapidly above 0 mv and falls again. After each fall, the constant injection of current causes the potential to ramp up again until a "threshold" of slightly below –40 mv (for this neuron) is reached, which causes the cell to fire again (from Bear et al., 2007).

2.3 Dendrites and Axons

Neurons in different regions of the brain have different morphological structures, but the typical structure includes a cell body (called the *soma*) connected to a tree-like structure with branches called *dendrites* and a single branch called the *axon* that emanates from the soma and conveys the output spike to other neurons (see Figure 2.3). The spike is typically initiated near the junction of the soma and axon and propagates down the length of the axon. Many axons are covered by *myelin*, a white sheath that significantly boosts the speed of propagation of the spike over long distances. The terms *white matter* and *gray matter* correspond respectively to the myelinated axons connecting different brain regions and the regions containing the cell bodies.

2.4 Synapses

Neurons communicate with each other through connections known as *synapses*. Synapses can be electrical but are more typically chemical. A synapse is essentially a gap or *cleft* between the axon of one neuron (called the presynaptic neuron) and a dendrite (or soma) of another neuron (called the postsynaptic neuron) (see Figure 2.3). When an action potential arrives from a presynaptic neuron, it causes the release of chemicals known as neurotransmitters into the synaptic cleft. These chemicals in turn bind to the ionic channels (or receptors) on the postsynaptic neuron, causing these channels to open, thereby influencing the local membrane potential of the postsynaptic cell.

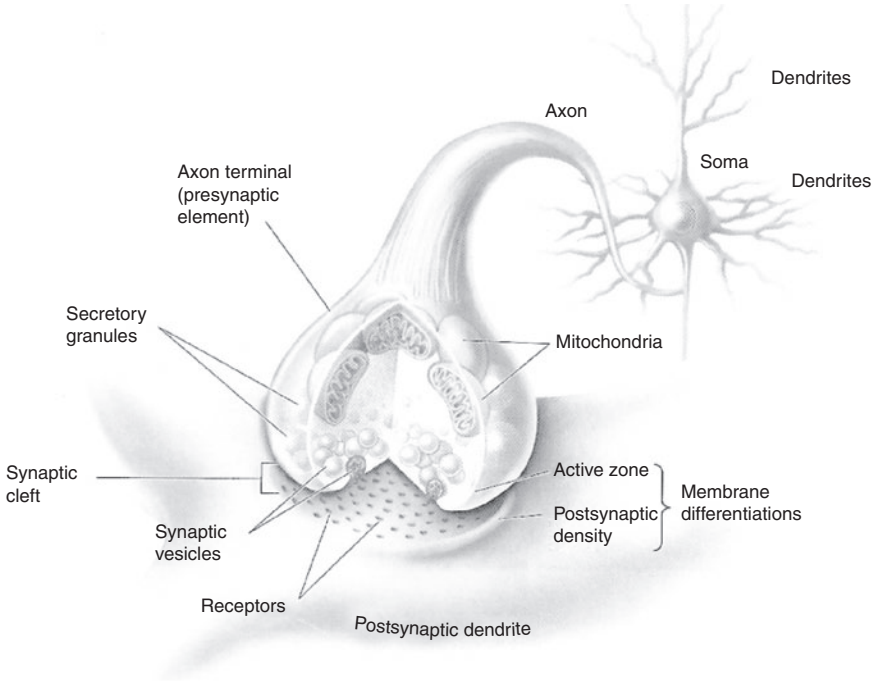

Figure 2.3. **Dendrites, soma, axon, and synapse.** The figure depicts a connection from one neuron to another. The dendrites, cell body (soma), and axon of the first neuron are shown, along with the synapse this axon makes on the dendrite of a different neuron. A spike from the first neuron causes the release of neurotransmitters stored in synaptic vesicles in the "presynaptic" axon terminal. These neurotransmitters bind with receptors in the "postsynaptic" dendrite, causing the ionic channels to open. This results in the influx or outflow of ions, changing the local membrane potential of the postsynaptic neuron (adapted from Bear et al., 2007).

Synapses can be excitatory or inhibitory. As the name suggests, excitatory synapses cause a momentary increase in the local membrane potential of the postsynaptic cell. This increase is called an excitatory postsynaptic potential (EPSP). EPSPs contribute to a higher probability of firing a spike by the postsynaptic cell. Inhibitory synapses do the opposite – they cause inhibitory postsynaptic potentials (IPSPs), which temporarily decrease the local membrane potential of the postsynaptic cell. A neuron is called excitatory or inhibitory based on the kind of synapse it forms with postsynaptic neurons. Each neuron forms only one kind of synapse, and therefore if an excitatory neuron is to inhibit a second neuron, it must excite an inhibitory "interneuron," which then inhibits the desired neuron.

2.5 Spike Generation

The generation of a spike by a neuron involves a complex cascade of events involving sodium and potassium channels as described above. However, in many cases, this