3.3 Recording Studies

The approaches described so far involve altering brain activity and measuring effects on behavior. A different strategy is to record from brain regions as the organism is engaged in different tasks and infer the hypothesized causes of the brain activity. The initial application of this strategy involved inserting electrodes into the brain and recording the electrical activity generated by the neuron or neurons closest to the electrode. This approach, known as *single-cell recording*, became the workhorse technique of mid–twentieth century neuroscience. A powerful illustration of the approach was Hubel and Wiesel's (1959) exploration of neurons in BA 17, otherwise known as primary visual cortex or V1. By varying the stimulus presented to a neuron in a cat or a monkey while recording from it, they determined that edges elicited the largest response, with different neurons responding to edges at different orientations. Some cells responded when the edges were stationary, others to edges moving in specific directions. Inspired by these findings, researchers undertook numerous studies seeking to identify the features of stimuli that would elicit activity in different visual processing areas, as we will describe further in Section 5.3.

Inserting electrodes into the brain is highly invasive and regarded as morally unacceptable in humans except for patients being evaluated for neurosurgery to remove tumors. In that context, recording from neurons while patients perform tasks allows surgeons to avoid cutting areas regarded as of critical cognitive importance, such as those involved in producing or understanding language. Researchers have often been able to take advantage of these situations and, with the patient's permission, record from neurons as different stimuli are presented. In a widely cited study, Quiroga et al. (2005) identified neurons that responded selectively to different images of well-known people, such as Jennifer Aniston or Bill Clinton, or even to their spoken name (suggesting that the neuron was responding to the person, not their visual appearance).

Given the moral issues in inserting electrodes in order to record brain activity, researchers working on humans often turn to noninvasive techniques for recording neural activity. One of the first to be developed was the electroencephalogram (EEG), which records electrical signals from electrodes placed on the scalp. In the pioneering research with this approach, Berger (1930) detected oscillations that varied in frequency depending on the activity that the subject was performing. When participants were quiet and kept their eyes closed, he detected oscillations of approximately 10 Hz (he referred to these as *alpha waves*). When participants opened their eyes or were asked to perform a cognitive task, the rhythms would increase to between 20 and 30 Hz and the amplitude would decrease (he called these

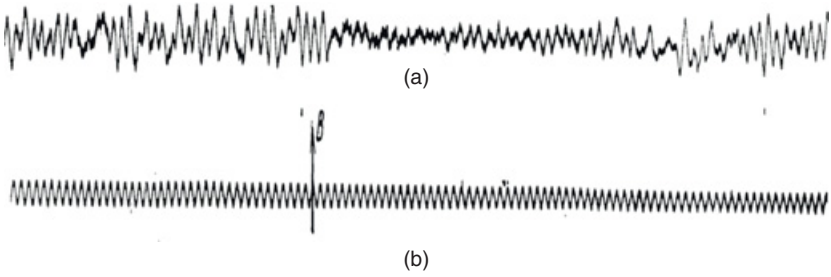

Figure 10 (a) An eight-second recording of a subcutaneous EEG by Berger (1930). (b) A 10 Hz timing signal. Initially, the EEG consists of alpha waves, but shortly after Berger strokes the subject's hand with a glass rod (indicated by the B on the timing bar), a period of low-amplitude, shorter-period beta waves ensues, which then transitions back to alpha waves.

beta waves) (Figure 10). In ensuing years, researchers differentiated yet faster and smaller amplitude gamma rhythms (above 30 Hz), which appear during cognitively demanding tasks, and slower delta (less than 4 Hz) and theta (4–7 Hz) rhythms, associated with the transition to sleep. The current measured and recorded in these oscillations is not due to action potentials but to synchronized fluctuations in the resting potential of a population of neurons (see Section 2.1). EEG has attracted renewed interest as researchers have recognized the importance of oscillations in cortical processing (see Section 9.3).

While EEG can measure brain activity with high temporal resolution, its spatial resolution is very poor, in part because electrical activity disperses widely. Noninvasive techniques that measure blood flow as a proxy for neural activity, such as functional magnetic resonance imaging (fMRI), offer much higher spatial resolution (albeit with a loss of temporal resolution since the rate of blood flow changes slowly). To interpret the neural activity in terms of contributions to behavior, researchers ask subjects to perform different cognitive tasks while lying in a scanner. When there is a large change in blood flow, researchers infer that this particular brain region is selectively contributing to the activity. The challenge is to infer just what an area has contributed to a task. Often this is addressed by exploring different tasks and asking what feature the tasks that elicit greater activity have in common. The challenge of identifying the contribution of a given area is increased when not just one but multiple brain regions exhibit increased activity in a given task. One response has been to shift the focus from individual brain regions to identifying networks of brain areas whose activity changes in a coordinated fashion. The goal is then to correlate

these networks with the type of processing required in the tasks that activate these areas (see Section 6.4 for analysis of networks).

Marcus Raichle, one of the pioneers in fMRI studies, drew attention to the fact that while some brain regions increase their activity when a subject is performing a task, others exhibit reduced activity. Raichle et al. (2001) scanned subjects when they were not assigned a task (a condition referred to as the *resting state*) and identified a network of brain regions whose activity was heightened in this condition and reduced in task conditions. Raichle termed this the *default mode network*. This opened up a new line of inquiry into what brain regions are doing when a person is not specifically challenged to perform a task.

3.4 Altogether Now

Since each of the methods we have reviewed – lesion, stimulation, and recording – present different epistemic challenges, one strategy for reducing these challenges is to employ all three together. We illustrate this approach with research on an area in BA 19 variously known as V5 or MT. Human patients with lesions in this area are unable to detect motion (rather, they see the world as a sequence of still images). To explore how V5 contributed to detecting motion, Britten et al. (1992) presented monkeys with displays in which objects were either moving in a common direction or in random directions and trained them with food rewards to press a different lever for each condition. They then introduced an ambiguous display in which about half of the objects moved in a common direction, while the rest moved randomly. The monkeys still responded, sometimes selecting the direction in which half the dots were moving and sometimes selecting the lever that indicated that they were moving randomly. By recording from neurons in MT, the researchers identified neurons that corresponded to the response that the monkeys made to the ambiguous display and inferred that activity of these neurons constituted the monkeys' perception of motion. Finally, they microstimulated neurons that corresponded to a specific response, and showed that they could bias the monkeys to make that response. They took these findings as strong confirmation that these neurons are responsible for monkeys' perception of motion. In this research, results from lesion, stimulation, and recording studies all converged, making the case especially compelling. Although one might challenge each of the techniques alone, using them together provides strong evidence that area MT is involved in visual perception of motion.

Underlying these efforts to determine the activity performed by a brain area is the idea that brain areas are modules responsible for specific tasks – for

example, MT is an area for detecting motion. As compelling as the case that MT processes motion information is, subsequent research has shown that this is not all MT does. It is also active in binocular vision, for example. Drawing on such evidence, Burnston (2016) and Anderson (2014) argue against fixed assignments of specialized functions to brain areas, arguing instead for more contextualized accounts in which brain regions perform multiple types of processing. What they do in a given situation depends on features of that situation: in different situations, they form coalitions with different brain areas and perform different functions. If such claims are correct, they can explain the versatility and adaptability of neural processing, but they make the challenge of figuring out what brain areas do much more difficult.

3.5 Computational Modeling

We finish by briefly noting another method utilized in neuroscience – computational modeling. This involves identifying variables thought to describe the changing states of a neural system and developing equations (typically differential equations) that characterize how values of these variables change as the values of other variables change. By starting with values for variables that are thought to describe the nervous system at one point in time, and having a computer iteratively apply these equations to determine subsequent values, researchers seek to simulate the brain system. A successful stimulation is one that generates a succession of values of variables that correspond to those measured in the brain. When the equations formalize what is already hypothesized about how the brain functions, a successful simulation provides evidence that the hypothesized account is correct (also see Section 6.3 on the contribution of computational modeling to neuroscientific explanation).

Computational models can also be developed as part of a discovery process. If one succeeds in developing a simulation that matches the behavior of the brain, one can interpret the equations as hypotheses about the processes actually operative in the brain. In that case, though, one generally seeks independent evidence that there are processes in the brain that correspond to the equations.

3.6 Summary

To investigate brains, researchers use a variety of techniques including lesioning, stimulating, and recording. As we have noted, the inference from these studies to what is happening in the brain is often indirect. Accordingly, researchers must often combine multiple methods. Brain researchers also often invoke computational models, which allow them to determine how the

brain would work if a particular hypothesis was correct or, in some cases, to advance new hypotheses.

4 From Whom Do Neuroscientists Learn about the Nervous System?

To learn about nervous systems, researchers must actually study nervous systems, using methods such as those introduced in the previous section. But whose nervous systems should they study? If the researcher is interested in a specific individual, then they would reasonably choose to study that individual. But science is generally focused on types, not tokens, where types are classes of entities taken to be the same in relevant respects. The goal is to generalize across the members of the type. This is relatively straightforward in the physical sciences. Chemists are not interested in a given specimen of, for example, gold, but in all instances of gold. What they discover in studying one specimen is assumed to apply to all instances. Neural scientists seek similar generalizability, although the scope of generalization is less clear cut.

A variety of characteristics can be used to identify types of organisms. For example, one might be interested in left-handed human beings. One might focus on species, for example, humans. Species membership is not determined in terms of necessary and sufficient conditions, as it is with elements like gold. Instead, what is relevant is the organism's history: Who were its parents? As species themselves originate from other species (as members of a person's family arise from other members of the person's family), these relations are often represented in branching trees. These descent relations correspond to inheritance – genetically based traits that emerge at one node in the tree are generally inherited by the branches. In this respect, evolution is a conservative process: as observed by Ernst von Baer in the decades before Darwin published his account of evolution through natural selection, new traits develop as variations and modifications of existing traits. Accordingly, generalization in biology, including neuroscience, involves applying what is learned about some species to those appearing in a particular clade (descendants of a common ancestor) in the evolutionary tree. One common way in which variation arises in descendants is with a mutation in which part of a chromosome is duplicated, generating multiple copies of some genes. Through further mutations coupled with natural selection, these duplicated genes differentiate and code for proteins that perform specialized tasks. As a result, descendant species retain the same basic traits but give rise to specialized versions. For this reason, biologists often find it useful to look back in the evolutionary tree to where a trait first emerged.