CHAPTER

5

Recording the Spatial Mapping Cells: Place, Head Direction, and Grid Cells

Kate J. Jeffery[*,1], *Roddy Grieves*[*], *Jim Donnett*[§]

*Institute of Behavioural Neuroscience, Division of Psychology and Language Sciences, University College London, London, United Kingdom; [§]Axona Ltd., St Albans, United Kingdom

[1]Corresponding author

1. INTRODUCTION

The neural encoding of space has been a strong theme in cognitive neuroscience ever since the discovery by O'Keefe, in the 1970s, that cells in the rat dorsal hippocampus selectively respond when the animal goes to particular locations in the environment (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976). The discovery of these so-called *place cells*, which motivated the proposal that the brain—and in particular, the hippocampus—makes an internal map of space (O'Keefe and Nadel, 1978), was enabled by emerging techniques for recording single neurons from freely moving animals. This method, using flexible microwires, has been in use ever since, with some refinements. Although new optical techniques for large-scale neuronal ensemble recording are now on the horizon, none of them yet match electrical methods for stability and temporal precision, and so this technique remains the method of choice for many investigators interested in neural encoding during behavior. The purpose of this chapter is to review the methodology and point out some of its potential pitfalls. We have been both recording spatial neurons (KJ, RG, and JD) and making and selling recording systems (JD) for around 25 years, and based on our experience we have written this chapter with the aim of providing information to help the novice investigator get started. It is by no means an exhaustive description of a large and complex area but presents the basic principles, many of which pertain not just to the spatial neurons but to other brain systems, and we hope it will be useful.

1.1 The Challenges of Chronic Single-Neuron Recording

The discovery of place cells came about as a result of two important technical advances made at around the middle of the last century. The first was the ability to record single neurons *in vivo*: Using such methods, Hubel and Wiesel were able to make their groundbreaking observations of visual responsiveness in anesthetized cats (Wurtz, 2009), thus opening the door to investigation of the neural basis of sensory perception. The second was the adaptation of this method to awake, freely moving animals, which is called chronic recording, and which posed a whole set of new challenges. One was that as the animal moves, the brain moves slightly relative to the cranium and so the sharp and rigid glass or tungsten electrodes used in anesthetized animals would cause damage if the animal moved. Additionally, the sensory inputs to moving animals change rapidly and unpredictably, making it hard to associate specific sensory stimuli with neuronal responses. There are also other problems: the body tends not to prefer chronically implanted foreign objects and tries to wall them off with fibrosis (called, in the brain, gliosis), a process that ensheaths electrodes and reduces their sensitivity (Cheung, 2007), while the implant itself forms a sheltered habitat for bacteria and thus an ongoing potential reservoir of infection. Furthermore, when an animal moves through space, electrical currents are induced in the metallic parts of the implants (especially the electrodes and signal transmission wires), producing

Handbook of In Vivo Neural Plasticity Techniques
https://doi.org/10.1016/B978-0-12-812028-6.00005-7

electrical noise that swamps the tiny neuronal signals, and if the animal stops and grooms or chews then volume-conducted electromyographic signals add their own noise.

Given all these obstacles, it's amazing that chronic neuronal recording works at all. However, over the ensuing decades the various problems have been tackled and largely solved, or at least mitigated. The problem of associating sensory stimuli to responses more or less solved itself when O'Keefe happened, serendipitously, to implant his electrodes in a part of the brain where responses are not tied to specific sensory inputs—the hippocampus. Hippocampal neurons are higher order processors that receive already integrated sensory signals from the periphery and combine them further to form a representation of place, independent of the exact sensory stimuli present at that place. Thus, a place cell fires whenever an animal ventures into, say, a particular region of an enclosure, no matter which direction it approaches from, whether the lights are on or off, etc., and so the signal is remarkably stable no matter what the animal is doing or perceiving. Because of this flexibility, the place specificity of place cells is easy to see. With this methodology, O'Keefe made his seminal and ultimately Nobel Prize–winning (Morris, 2015) discovery.

The other problems, of stability and biocompatibility, were tackled both by O'Keefe et al. and by their successors, who over the last few decades have developed refinements to allow recording of large numbers of neurons under challenging experimental conditions. Below, we first present the basics of single-neuron recording and then expand the discussion into more complex issues to do with behavior and recording in complex environments.

2. RECORDING AND SIGNAL PROCESSING

The aim of a recording system is to detect the occurrence of ion flow across cell membranes, occurring as a result of signal processing by neurons. To really accurately capture these flows, it is necessary to record from inside the cell, but intracellular recording is impractical in an animal that moves, so the next best thing is extracellular recording. Here, an electrode is able to detect current sinks (currents flowing into a cell) and sources (currents flowing out of a cell) in the form of local voltage changes. The most prominent change is that induced by an action potential, and so most of what follows in the remainder of this chapter concerns how action potentials are detected and analyzed.

2.1 A Typical Recording System

The components of a typical chronic single-neuron recording setup are as follows (Fig. 5.1):

- A microwire electrode, located near the neuron and able to detect electrical voltage changes.
- A headstage, located as near to the electrode as possible, which isolates the neural signals and buffers them to protect them against noise pickup.
- A recording cable, which conducts the buffered signals to the preamplifier.
- An analog-to-digital (A–D) converter, which takes the raw, continuously varying analog voltage and digitizes it (converts it to 0s and 1s).
- Filters, which removed unwanted components of the electrical signal based on their frequency.
- An amplifier, which enlarges the signals to make them easier to analyze.

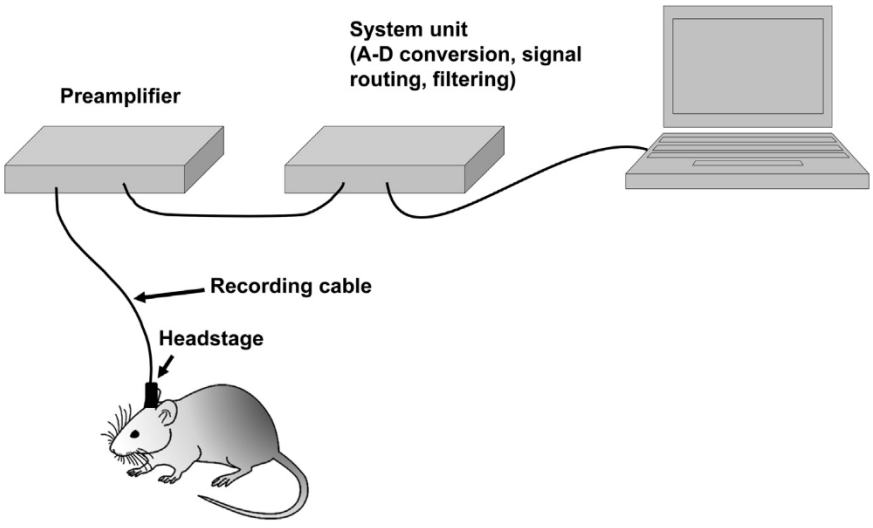

FIGURE 5.1 A typical single-neuron recording system showing the basic elements.

- A referencing system, which allows signals on different recording channels to be compared and subtracted, thus removing some of the noise.
- A storage system of some sort to capture and store the data.

Recording systems may vary in the order in which these events occur. For example, some digitize the signal before filtering, while others might apply analog filters to the signal before digitizing, or even filter the analog signal and then filter the digital signal further on. A diagram of a typical recording system is shown in Fig. 5.1.

2.2 Recording Electrodes

A particularly important development in chronic single-neuron recording was the move to using flexible microwire electrodes with blunt ends, which move as the brain moves and thus do not damage neurons and also can remain in place for many weeks or months. The materials initially chosen for these electrodes—soft metal alloys such as platinum–iridium (Pt–Ir), coated with a microscopic layer of insulating material such as ceramic or (polytetrafluoroethylene, or "Teflon") PTFE—turned out to be remarkably inert (noncorroding) and also biocompatible, inducing relatively little gliosis even when implanted for long periods of time. Although various alternative materials have been tried over the years, including tungsten, nichrome, and stainless steel, Pt–Ir electrodes remain widely used due to their inertness and recording stability.

A characteristic of microelectrodes that is critical for success is their tiny size: A typical hippocampal pyramidal neuron soma is only around 10–30 μm in diameter, and together with the dendrites the whole cell is only around 200 μm, so for an electrode to be able to nestle close enough to a cell to detect its voltage changes (which decline as the square of the distance from the cell), it needs to be equally small. For this reason, wires of between 12 and 25 μm are typically used in chronic recording. These days these are usually wound together in bundles of four, called tetrodes, to allow for better spike separation (assignment of spikes to a given neuron; see below and also Chapters 2 (Souza et al.) and 8 (Tsanov)), but this enlarges the electrode to around 100 μm and is less useful in brain areas where neurons are small, such as the cerebellum. If a large number of neurons are desired, as in modern ensemble recordings, then the number of tetrodes can grow to be very large—up to 20 or 30—causing considerable mechanical damage to the brain. For this reason, there is a move now toward microscopic activity imaging of neurons instead, but this method has its own problems (detailed in Chapter 13 (Dylda et al.). Most notable among these is the slow kinetics of chemical readouts of neural activity, measured in tens of milliseconds (whereas a spike is around 2 ms); so for now, high-temporal resolution recordings continue to rely on direct electrical voltage–change detection.

Another reason for having extremely thin electrode wires is that they have a very high impedance—this is typically measured as the resistance offered against a 1 kHz alternating current and is important because it means that the tiny currents flowing across the cell membranes are converted into large voltages, which can be more easily transmitted and detected by the recording system. However, high impedance means that electrical noise is also amplified; there is thus an optimum range between high and low. As wires have become thinner, enabling more of them to be implanted and cell yield to increase, problems have developed with impedance being too high, which increases thermal noise and also pickup of noise from nearby electrical sources (Loeb et al., 1995), so for very thin electrodes (20 μm or less) impedance is usually lowered to around 250 kΩ by electroplating (Ferguson et al., 2009). This is a process by which a metal is deposited in lumps at the ends of the electrodes by passing a current through the electrodes while they are dipped in metal (usually gold) ion solution. This needs to be done carefully so as not to deposit so much metal as to bridge between electrodes and cause short circuits. When done correctly, the end result is a wire that is very thin but has a high surface area and lower-impedance tip, thus optimizing the balance between size and signal detection.

2.3 Referencing

The noise picked up by electrodes has several sources: some of it, as mentioned above, is artifact—spurious signals generated by movement or thermal noise—and some of it is background—that is, because of brain signals emitted by neurons that are some distance away such that the action potentials cannot be distinguished from the background. One way to reduce these types of noises is to compare the signal from the electrode of interest against a signal from a very similar electrode that is not near any cells. Much noise will be common to both electrodes and so if the signals are subtracted ("common mode rejection") then the noise is subtracted out and the resulting trace is much cleaner and contains mostly the signal of interest. For recording single neurons, it is usual to use a reference electrode. Because noise pickup is impedance-dependent, it is best to have signal and reference electrodes that are impedance-matched. The most common method is to use an electrode from a different but nearby tetrode, but it needs to be one on which few or no spikes are being detected, or else these will appear upside-down (subtracted) on all four recording channels.

2.4 Grounding

One of the most important concepts to an electrophysiologist is the concept of *ground* and related terms such as grounding and the dreaded *ground loop*. Ground refers to the reference electrical potential against which other potentials are measured. For signals being picked up from electrodes, the ground voltage against which the values are measured is established by placing a (relatively) large metal insert in contact with the brain at some distance from the recording site. The most common way of doing this is to insert a jeweler's screw through the skull, far enough to contact the dura underneath. In addition to referencing the signal against the animal's electrical potential, it is usually also necessary to ground the animal as a whole, by having its feet contact a material that is at the potential of the global ground, or Earth potential. This is because otherwise the animal may accumulate charge (particularly if its fur rubs against the environment) and acquire a DC offset, which discharges periodically as static, creating electrical noise. The usual way to do this is to ground the apparatus—this does not work, however, if the apparatus is acrylic or some other insulator. In practical terms, the environment is usually connected by a thick wire to something in the actual ground, such as the grounded wire in a mains socket, which in turn is connected to plumbing pipes or some other highly conductive material embedded in the Earth.

Note that the Earth is not exactly at the same potential everywhere and two wires independently grounded a few feet apart may in fact "see" slightly different potentials. If these two wires are now connected, then even though they are supposed to both be at ground potential, current will flow between them, a situation known as a ground loop. These currents are picked up by the recording equipment and create noise that can completely obscure neuronal recordings, so they need to be eliminated—this can be done by ensuring that everything only connects to the Earth via a single common ground. Ground loops can also occur in other parts of the setup, such as in parts of the apparatus that were supposed to be *shielding* noise (e.g., a copper Faraday cage that completely surrounds the setup), and are one of the commonest causes of intractable noise in recording systems.

2.5 Microdrives

Although electrodes are very thin, they do have a finite size, which limits how many can be implanted without damaging the brain too much. An alternative to filling the brain with microelectrodes is to use fewer of them but to move them slowly through the brain so that a larger number of neurons can be recorded, albeit in sessions that are separated across time. This method, developed by Ainsworth et al. (1969), relies on tiny devices called microdrives, which house a fine-threaded screw that an experimenter turns—usually by hand, but sometimes with an electric motor—to move the electrodes down in tiny movements of as little as 25 μm. However, as tetrode numbers increase, the number of drive elements needed also increases, until the combined size or weight exceeds what an animal can comfortably carry on its head. This, combined with the damage done to the brain by large numbers of tetrodes, limits the number of neurons that can be simultaneously recorded by tetrodes.

Most microdrives in use today are based on a central screw, this being the most efficient way of translating a large (angular) movement into a very small (linear) one. When a screw turns in a captive nut then it will move forwards, whereas if the screw is captive then the nut will move instead. Thus, microdrives fall broadly into two categories: those in which the screw drives the electrodes down and those in which the electrodes are attached to the nut and advanced by turns of the captive screw (see Fig. 5.2 for an example of the latter). For ordinary right-handed screws, in the former case a clockwise turn of the screw will advance the electrodes, whereas in the latter case a counterclockwise turn is needed.

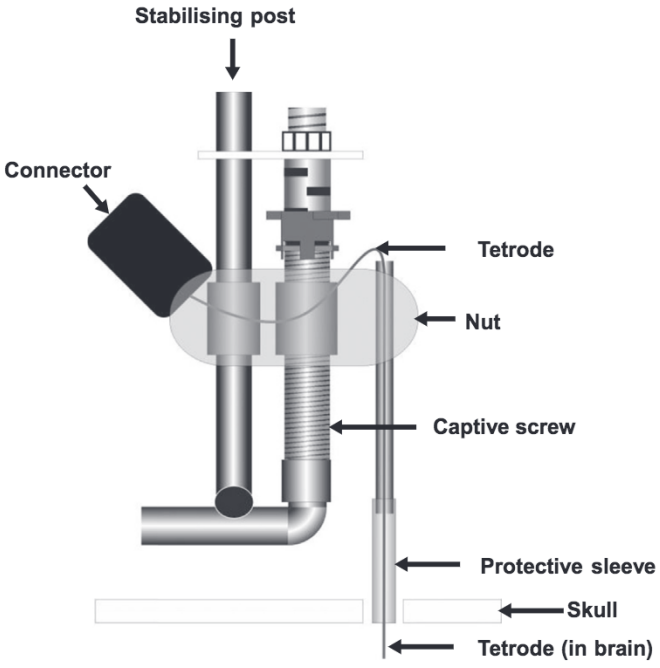

FIGURE 5.2 Diagram of an Axona microdrive based on the "captive screw" principle. The screw is held fixed at both ends, and the nut, to which the electrical connector and electrodes are both attached, moves up and down as the screw turns. Rotational movement of the nut around the screw is prevented with an additional, parallel metal post on which the nut also moves. After positioning, the assembly will be attached to the skull surface with dental acrylic that covers the base of the drive but not the screw. *Figure by J. Huxter, available at https://doi.org/10.6084/m9.figshare.6407504 under a 1029 CC-BY 4.0 license.*

The features of microdrives that are of importance are:

- Weight. A rat can comfortably carry 1–2 g on its head and can, with difficulty, carry 4–5 g, so it is important that the drives be made of lightweight components. This is even truer with mouse microdrives, which are now commonly used in the study of genetically modified animals. Some investigators mitigate the weight problem by using a pulley system to counterbalance it—this is especially necessary with multi-tetrode drives, which can sometimes weigh as much as 20 g and be very difficult for a rat to support.
- Rate of electrode movement. The cell body of neurons varies in size from 5 to 100 μm, so it is necessary to be able to move the electrodes by very small amounts—tens of microns—to optimize recording. The rate of advancement is specified as microns per turn, or "pitch" of the screw. A typical microdrive will have a core screw system with a 200 μm pitch or thereabouts.
- Stability. As investigators usually wish to record a given set of cells for hours, or sometimes days, it is important that the drives do not allow for movement of the electrodes, even if the animal shakes its head vigorously.
- Number of independently moveable units. For some experiments, it is sufficient that the electrode bundle moves *en bloc*, in which case a single screw can drive the bundle. For other experiments, where investigators want to be able to move electrodes independently to maximize the cell yield from each electrode, it is necessary to have several screws. This necessarily adds to the weight of the drive and also the complexity and expense of the mechanism. It also maximizes cell yield, however, which is important for ensemble recording in which as many simultaneously recorded cells as possible are needed.
- Reversibility. Some investigators wish to be able to retract as well as advance the electrodes, in which case the drive needs to be reversible. This feature, while useful, introduces a problem known as backlash, in which turning of the screw does not result in retraction of the electrodes until the slack in the drive (due to the space between screw and nut) has been absorbed. If the backlash of a given drive is not known, this will introduce an uncertainty in the electrode position, which may cause problems if the movements of the drive were being used to estimate location relative to cell layers.
- Method of wiring. It is necessary to attach the electrode wires to the electrical contacts on the drive, and this is often done by hand. For very large numbers of electrodes, as with ensemble recording, this can make preparing the drives very arduous and time consuming.

As mentioned earlier, when tetrodes become too numerous, their bulk starts to produce mechanical brain damage. One solution to this problem, which is becoming increasingly popular, is to use silicon probes. These are micromachined arrays of conductive and insulating material, which allow many, sometimes hundreds, of recording sites to be monitored at once (Cheung, 2007; Buzsáki et al., 2015). Typically these are implanted into the structure of interest and then left in situ, but they can also be attached to microdrives and moved. Although these have been very important in local field potential (LFP) studies of the spatial distribution of current sinks and sources (e.g., Montgomery et al., 2009), their utility in long-term single-neuron studies has so far proved to be more limited although their use is increasing. One reason seems to be their lower biocompatibility—a small industry is engaged in trying to coat the shanks with chemicals that hide the probe from the immune system to prevent gliosis and signal loss (Cheung, 2007), but so far relatively few long-lasting high-yield studies have been published using silicon probes. This technology, however, is steadily advancing and may ultimately be the best solution for so-called high-density (large-ensemble) recording, at least until optical methods improve.

2.6 Signal Buffering, Transmission, and Amplification

Detecting voltage changes at source is one thing, but transmitting this signal in a noise-free way to a computer creates a whole new set of problems, such as signal loss and noise pickup. It is very hard for the tiny brain signals (a few tens of microvolts) to travel as microcurrents up a long recording wire, so the signals need to be conditioned first. This is done with a headstage, which is a set of electronic components situated as close to the recording electrodes as possible, on the head of the animal. Headstages have three important properties: buffering, digital versus analog signal type, and AC versus DC coupling.

Signal buffering by a headstage amplifier is undertaken so that a signal detected by a high-impedance electrode is able to be carried by very low-impedance transmission wires, making it dramatically more noise immune. Additionally, because the currents occurring in the brain are detected but then physically isolated from the rest of the recording system, the system does not itself cause current flow in the brain. Headstage amplifiers are almost always unity gain, which means that they transmit the signals at the same amplitude as they

were detected—this reduces hardware and hence weight on the head of the animals. Some headstages then digitize the signal, i.e., convert the continuously varying voltage signal into discretized large versus small voltages that represent the 0s and 1s of the signal. This means the signal is relatively immune from noise, and it also allows the signal to be multiplexed so as to allow transmission of more than one electrode signal up a single wire. While multiplexing dramatically reduces the cable requirements, current commercially available chips such as those from Intan add the problem that the reference signals become slightly uncoupled in time from other channels because the channels are sampled sequentially, which reduces the common mode rejection efficiency. Analog headstages transmit the continuously varying voltage stream instead of digitizing it—this makes them smaller and lighter in weight, and also less expensive, and for this reason they have been very popular for recordings using smaller numbers of channels (up to around 64).

The final headstage amplifier characteristic concerns whether it is AC or DC coupled, which refers to the presence or absence of a capacitor in series with the input. Because a capacitor charges and discharges periodically, it allows transmission of a high-frequency oscillating signal and thus acts as a high-pass filter. This means that even if the signal acquires a DC offset (a constant voltage added to the time-varying signal because of charge buildup; see Fig. 5.3), this does not affect the signal capture—the system records between the highs and lows of the signal regardless of their actual values. A DC-coupled headstage lacks a capacitor and the associated filtering, allows an experimenter to see DC levels, and is used, for example, in intracellular membrane potential recordings, but charge buildup means the signal can often slowly drift out of range .

A current challenge is how to adapt the methods to recording in complex environments. One particular difficulty is the wire connecting the headstage to the preamplifier, which not only adds weight for the animal to carry but also impedes its ability to walk under overhanging features. This limits the types of observations that can be made. One solution has been to develop wireless recording, in which signals are either transmitted via radio to a receiver unit (telemetry), or else stored locally on a memory card and offloaded at a later time (data logger). The advantage that telemetry has over local storage is that the user can monitor the signal and in some cases intervene experimentally as a function of the signal (e.g., to stimulate when an animal is in a place cell's place field). A disadvantage is that such devices need to be self-powered. This is typically done by mounting a battery either on the animal's head (which adds weight) or body (via a harness, which can annoy and distract the animal). However, bandwidth is traded off against battery life, meaning that for large enough data rates the trial time becomes infeasibly short. Work is under way to solve the battery life problem by using inductive coils to power the devices as the animal moves through space—such methods so far work only in limited situations with limited signal range or limited numbers of electrode channels, or both, but more widely usable products are on the horizon.

Another disadvantage is that there may be signal loss due to interference with the radio transmissions carrying the signals. Among other problems, reflection of the transmitted signal from metallic objects in the environment creates "multipath" interference (reflections canceling out the direct signal); such reflections preclude electrical shielding, which might otherwise be used in a wired recording setup, reducing noise management options. For this reason, hybrid methods are sometimes used, whereby a small portion of the signal is transmitted to allow user monitoring, but the bulk is stored on a memory card.

2.7 Signal Filtering and Referencing

The raw signal coming from the brain typically looks quite noisy, even after referencing, and it is necessary to remove the unwanted components by filtering them off so that the component of interest can be isolated and studied. For the spikes from single neurons, a minimum workable signal:noise ratio is about 2:1—much below that, and the spikes become mixed in with background

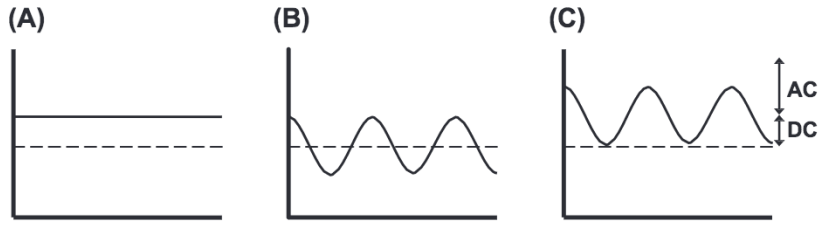

FIGURE 5.3 Different types of voltage variation. *Dotted line* = 0 V, *solid line* = signal. (A) DC offset, in which the signal has a steady deviation (in this case positive) from zero. (B) AC (alternating current) signal, which oscillates above and below zero. (C) AC signal, which oscillates above and below a constant DC offset.

noise, which degrades the data quality. For brain studies, two types of filtering are commonly used: high pass, which filters off the slow components (the LFP) and leaves the high-frequency ones (the spikes), and low pass, which filters off the high-frequency spiky parts of the signal and leaves the slow oscillations.

2.7.1 How Filtering Works

Any signal, even an irregular one, can be expressed as the sum of a number of sine waves of varying frequency, amplitude, and phase. Decomposing an irregular signal into its component sine waves is known as Fourier analysis, and its output is a graph called a power spectrum (Fig. 5.4), which depicts how much of the signal receives contributions from sine waves of various frequencies. A very spiky signal will have a preponderance of component sine waves at high frequencies, whereas a very slow, smooth signal will be made predominant from low-frequency sine waves. For brain signals, the component that is generated by action potentials—called "spikes" for a good reason—will have a high-frequency component of 1 kHz or above, whereas the component contributed by slow oscillations will have frequency components mainly in the range of a small number of Hertz. One of the most important of these is theta, a sinusoidal pattern of about 8–12 Hz in rats (Fig. 5.4) and a little slower (around 4 Hz) in primates. Theta is an important phenomenon in hippocampal electrophysiology, and one we will come back to.

In addition to band-pass filters, most recording systems allow filtering between a very small range of frequencies on either side of 50 or 60 Hz (depending on local mains AC frequency). This "notch" filter thus removes mains noise without damaging components of the brain signal at nearby frequencies.

2.8 Sampling and Analog-to-Digital Conversion

The electrical potentials produced by neurons are analog events, which vary continuously. Computers, on the other hand, are digital and work with binary numbers, which are discrete entities, and so for a computer to capture, store, and analyze neural activity, the smoothly varying analog signal has to be digitized. This is rather like taking a tune that somebody is humming and playing it on a xylophone: It entails sampling the values of the signal at various points along the continuous stream of data and expressing the values as numbers (Fig. 5.5). The circuit that does this is known as an analog-to-digital (A–D) converter.

The sampled and digitized signal has less information than the continuous original and is thus a degraded version of it. This degradation takes two forms, corresponding to the two dimensions of the signal: (1) degradation because of the quantization (discreteness) of the temporal sampling (Fig. 5.6A) and (2) degradation because of quantization of the numbers used to represent the amplitude (Fig. 5.6B).

2.8.1 Loss of Information due to Discrete Timing of Samples

Because the samples are obtained at discrete time intervals, information about what happened in the spaces between the samples is lost. Clearly, then, it is necessary to make sure that samples are taken sufficiently often to

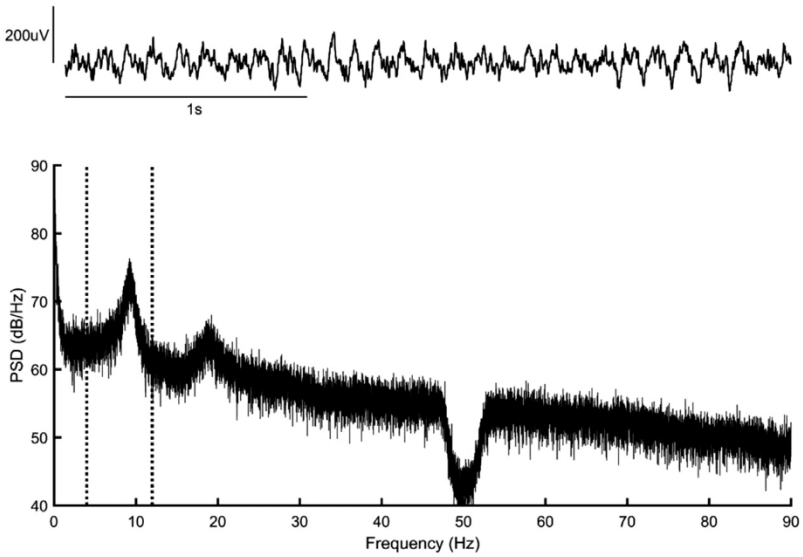

FIGURE 5.4 Power spectrum recorded from the hippocampus of a freely moving rat (theta band-pass filtered trace is at the top). Note the clear theta peak at 8 Hz, with a second harmonic at 16 Hz (Sheremet et al., 2016) and a notch at 50 Hz due to the notch filter (bottom). *PSD*, power spectral density.

FIGURE 5.5 Neural signals are produced continuously (left) but must be sampled at intervals and converted to a sequence of numbers (right) by an A–D converter.

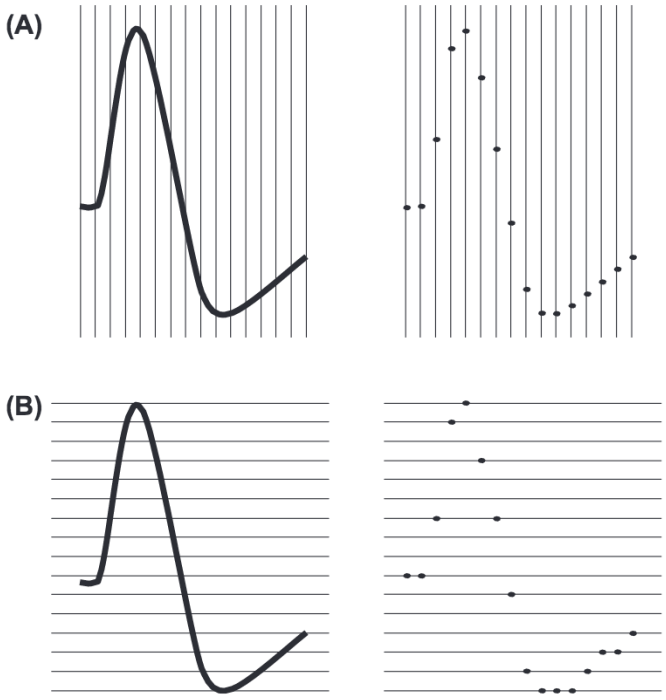

FIGURE 5.6 A typical neuronal spike, around a millisecond in duration, originally sampled at 32 kHz. The raw signal is shown intersected by the points where the samples were taken. (A) Discretization in the time domain at 8 kHz means that only some of the original points are captured (occurring every 0.125 ms); (B) discretization in the amplitude domain means that each sampled point can only be expressed at a limited number of amplitude values.

capture the important aspects of the signal but not so often that the processor struggles to keep up and the computer's hard disk is swamped. Additionally, failing to take samples sufficiently often will cause alterations in both frequency and amplitude of the resulting signal, a process known as *aliasing*, which distorts the reconstructed signal. An example of how aliasing happens is shown in Fig. 5.7.

A practical illustration of aliasing occurs when a car wheel appears to be rotating backward when shown on a movie screen. This is because the frame rate of the film is too low relative to the rotational speed of the wheel such that the snapshots that are presented to the eye show the wheel at successively earlier stages in the cycle of its rotation. Strangely, aliasing worsens—that is, the frequency of the resulting signal becomes lower—the closer the sampling frequency is to the underlying frequency. If sampling occurred exactly at the

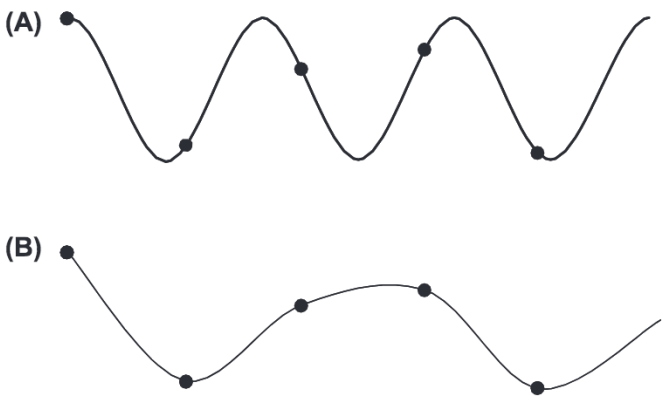

FIGURE 5.7 Aliasing caused by undersampling of a signal. (A) The original signal is a regular sine wave of constant peak amplitude. The *dots* show the points at which samples were taken. (B) When a curve is reconstructed from the samples, using interpolation to fill in the spaces, the reconstructed signal is distorted in both frequency and amplitude relative to the original.

underlying frequency, then the signal would appear flat (e.g., the wheel in our example would appear not to be rotating at all).

The way to remove aliasing is to make sure that the sampling rate is sufficiently high. The Nyquist theorem states that the sampling frequency must be at least double that of the maximum frequency component in the signal so as to capture the peaks and troughs of the fastest component sine wave. Thus, if the spike events in neuronal recordings contain component sine waves of 6 kHz, then the signal needs to be sampled at 12 kHz (at least) to avoid aliasing. In practice, a typical single-neuron sampling rate is usually 24 kHz or higher. Note that even with a high sampling rate such as this, problems could arise if the signal contained, as well as the 6–12 kHz neural frequencies of interest, additional frequencies that approached the sampling frequency—these would, by aliasing, show up as low-frequency components in the signal, which might look like LFP. The way to avoid this problem is to remove these higher frequencies by using an "antialiasing filter," which is a low-pass filter that rejects frequency components above half the sampling rate.

2.8.2 *Loss of Information Because of Discreteness in the Amplitude Domain*

The second way in which digitization causes loss of information is that the numbers representing the

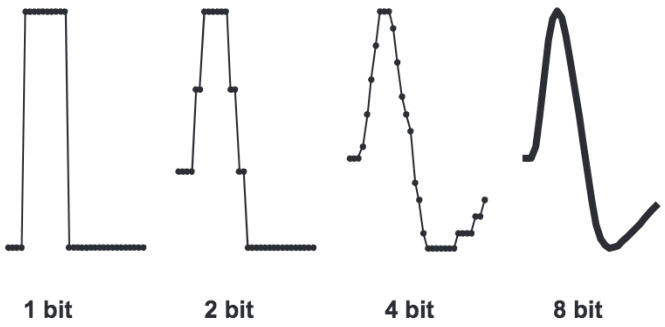

FIGURE 5.8 Representation of the waveform shown in Fig. 5.6, at different resolutions. At 1-bit resolution each sample is either "high" or "low," which seriously degrades the waveform. At 2-bit resolution the shape of the waveform starts to appear, and by 8-bit resolution the digitized waveform now closely resembles the original.

amplitude need to be represented in binary form to be processed by a computer, which requires the decision of how fine grained this amplitude representation should be. This depends on finding the balance between accuracy and storage space and is the issue of *resolution*. In the example in Fig. 5.6B, the spike was captured with 4-bit resolution so that each point could take one of 2^4 (i.e., 16) amplitude values. Fig. 5.8 shows the effect, on the waveform quality, of sampling at 1, 2, 4, and 8 bits. Our experience is that 8-bit resolution is a good balance between waveform quality and bandwidth (how quickly information needs to be transmitted) and storage cost, although 12-bit resolution is also common.

3. ANALYSIS

Having collected and preprocessed the brain signals, the next step is to subject them to the analysis, usually done off-line, that separates out the spikes and determines the firing characteristics of interest, analyzes the LFP, and looks at the interrelationships. Spike sorting spike time analysis and spatial analysis are discussed in the following sections.

3.1 Spike Sorting and Waveform Analysis

The first analysis requirement is to decide what spikes belong together as coming from the same neuron, a process known as spike separation or spike sorting. In the early days of chronic single-neuron recording, spikes were isolated on a simple amplitude basis by picking off the tallest deflections, after appropriate filtering of the signal. By next adding an upper as well as a lower bound, a so-called "window discriminator" could sometimes successfully distinguish more than one neuron. With further developments, the shape of the spikes also came into play, and a template-matching algorithm could distinguish spikes based on their waveforms (Lewicki, 1998), thus successfully separating interneurons (narrow) from principal cells (broad). The hippocampus posed two new problems, however: One was that the large size, tight packing, and high activity level of the neurons meant that it was difficult to attribute spikes to different neurons. The other was that in hippocampus, the principal cells tend to fire complex spike bursts in which the size of successive spikes in the burst steadily decreases, thus making a window discriminator ineffective.

An important step toward solving the spike sorting problem was made by McNaughton et al. who developed the stereotrode (McNaughton et al., 1983b; Fig. 5.9), constructed by tightly twisting two electrode wires together and then cutting the tips flush. Because of the slight separation in space of the two tips, the profile of the waveforms corresponding to a given spike is slightly different on the two electrodes (Fig. 5.9B). The most salient difference is usually in the profile of amplitudes, and so by plotting the amplitudes seen on one electrode against those seen on the other, clusters can be seen in the resulting scatterplot (Fig. 5.9C). This method was then extended from two to four electrodes to form the tetrode (Recce, 1989), which in theory should produce maximum separation, as no two points in the surrounding space can have the same set of distances from all four electrodes. With tetrodes, the clusters from two active neurons will sometimes overlap in the scatterplot from one electrode pair but never from them all.

Since the development of tetrodes, refinements in spike sorting have been made—this became especially important when investigators began to move from hippocampus, where cells form often quite discrete spike clusters that are easily separable, to surrounding regions of neocortex such as the entorhinal region, where cells fire at high rates and clusters are often hard to distinguish by eye. Here, automated methods such as KlustaKwik (Kadir et al., 2014) become necessary. These methods take in as their primary inputs the entire waveform, from which they use principal component analysis to extract parameters such as the peak−peak amplitude, spike width, time to the peak and size of the afterhyperpolarization. Typically, such methods generate a large number of clusters (i.e., they "overcluster"), which an experimenter will then refine by hand. These methods allow not only better spike separation but also larger amounts of data to be handled at once, as the number of tetrodes has been steadily increasing. However, even with these automated tools, which are steadily increasing in capacity (e.g, see Pachitariu et al., 2018), spike sorting remains one of the most time-consuming parts of *in vivo* single-neuron recording.

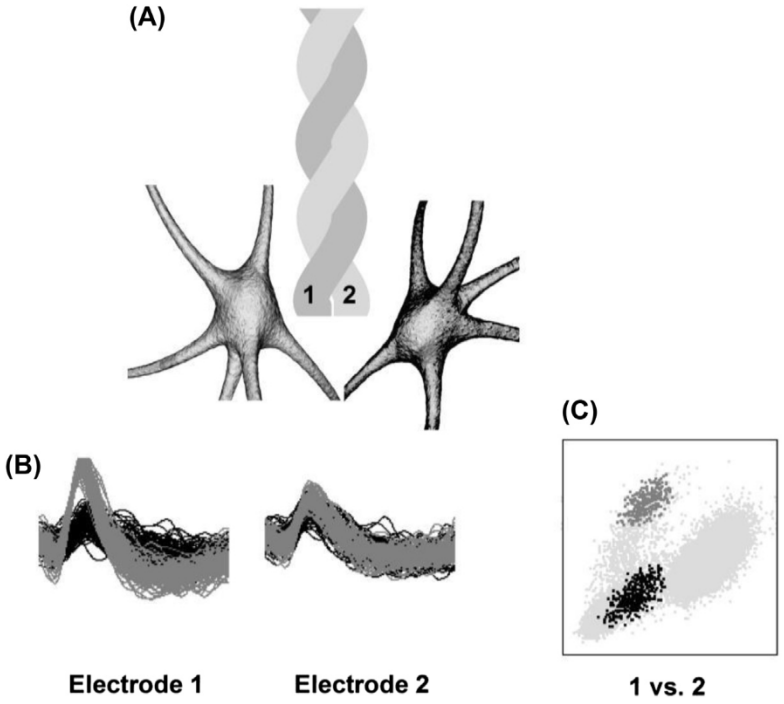

FIGURE 5.9 Using stereotrodes for spike separation by clustering. (A) A stereotrode is a twisted pair of electrodes implanted near neurons. Because the tips of the two electrodes (1 and 2) are spatially separated, they will not lie at the same distance from the nearby neurons. (B) When action potentials emitted by the two neurons are plotted, it can be seen that the neuron nearest electrode 1 (shown in gray) produces larger action potentials on that electrode. By contrast, both neurons are equidistant—and further—from electrode 2, and so the action potentials seen on this electrode are approximately the same amplitude and are smaller. (C) A scatterplot of the amplitude of the action potentials seen on electrode 1 (vertical axis) against electrode 2 (horizontal axis) shows that the two neurons produce different clusters of points, highlighted in black and dark gray (the *light gray dots* are other neurons, not separated from the background). *Figure by K. Jeffery, available at https://doi.org/10.6084/m9.figshare.6409304.*

3.2 Temporal (Spike Train) Analyses

As with the waveform-based sorting methods described above, temporal characteristics of the neuronal activity can be used to decide what kind of neuron is being recorded and whether this corresponds to a single neuron or more than one. Temporal characteristics of possible interest are firing rate, the pattern of interspike intervals (ISIs), periodicity (if there is any), and phase relationships with the LFP. Because a neuronal record consists of a long "train" of spikes, these analyses are sometimes collectively known as spike train analyses.

When the spikes have been collected, the time of each occurrence is logged in the data file: This procedure is known as time stamping, and it enables temporal analysis of the spike train characteristics. The first procedure is usually to plot a spike raster, which is a simple plot of each spike event on a timeline (Fig. 5.10). It is useful to look at such plots before proceeding to further analysis because the eye can sometimes discern patterns that suggest what analyses should be done. In the example shown in Fig. 5.10, a 1-s plot of spike events is shown, together with the associated LFP trace. Even by eye it can be seen that the cell is firing in rhythmic bursts, roughly corresponding to the frequency of LFP oscillations (in this case, hippocampal theta). The raster plot is purely for display purposes however: quantitative analyses such as autocorrelation and cross-correlation are needed to confirm the intuitions suggested by the spike raster.

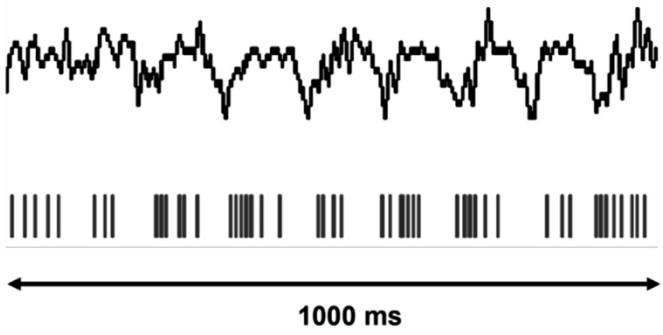

FIGURE 5.10 A plot of local field potential (LFP) theta (top trace) and a spike raster of the action potentials (*vertical bars*). Note that the action potentials appear to arrive in bursts, synchronized with the troughs of the LFP.

3.2.1 *Firing Rate*

The first and simplest spike train analysis is the calculation of the neuron's firing rate. Many neurons fire with

a characteristic rate that can be useful in identification, particularly when used in conjunction with the waveform analysis. Cortical neurons tend to produce tall, broad spikes at a high rate, whereas hippocampal interneurons produce narrow spikes, also at a high rate. Hippocampal pyramidal neurons tend to fire at a low average rate, whereas dentate granule cells fire very little, most of the time.

A distinction needs to be made between average rate and peak rate. While many neurons fire at a fairly constant rate throughout a given recording session, many others will increase (or decrease) their activity according to events that occurred during recording, such as the animal making a sensory perception, entering a particular region of the environment, or performing a particular task. The average rate is found by some measure of central tendency such as the mean or median, whereas the peak rate needs to be extracted by some kind of temporal or spatial chunking process in which the trial is divided into bins (e.g., 1-s bins or 3×3 cm bins) and the rate calculated for each bin, often after smoothing. Deciding what bin size and smoothing parameters to use is a delicate process because the wrong choice can obscure interesting phenomena.

3.2.2 *Interspike Interval Histogram*

After the firing rate, the next interesting feature of spike timing is the inter-spike interval (ISI) distribution: Are the intervals all the same or do they vary, and if so by how much? The ISI histogram, also called a first-order interval histogram, is a plot of the distribution of the intervals between successive spikes from a given neuron in a particular recording session. It is related to the autocorrelation function (see below), but unlike autocorrelation, it only considers the relationships between adjacent spikes and thus carries less information. It is useful for revealing simple temporal relationships. As with all histograms, the x-axis needs to be binned, and the choice of bin width depends on the interval range of interest. For an ordinary spike train, a bin width of about 1–2 ms is appropriate, as neurons cannot fire at a higher rate than this due to the action potential refractory period. Indeed, the ISI histogram can be useful in spike separation because if a neuron has been well isolated then there should be no ISIs less than the refractory period (Fig. 5.11).

3.2.3 *Autocorrelation*

The autocorrelation histogram, sometimes known as an "all-order interval" histogram, is like the ISI histogram except that each spike is compared not just with adjacent spikes but with all other spikes falling within the temporal range of the histogram. It thus contains more information and is useful for revealing multiple periodicities within a spike train. An example is shown in Fig. 5.12. It was plotted by taking each spike in turn (index spike), placing it at time = zero and plotting all other spikes from that neuron (the reference spikes) that fell within 500 ms each side of the index spike. When this procedure was repeated for every spike in the train, a histogram was built up that reflects the most frequently occurring spike intervals. It can be seen in this example that for every index spike, it was likely there would be another spike at about 125 ms, also likely (although fractionally less) that there would be another at 250 ms, and so on. Thus, the neuron was firing rhythmically at about 8 Hz. (In fact, this is the theta-correlated neuron shown in Fig. 5.10.)

Note that the autocorrelation function is always symmetrical around time = zero; this is because the histogram plots every spike pair interval twice: once with one of the spikes at the center and the second as the

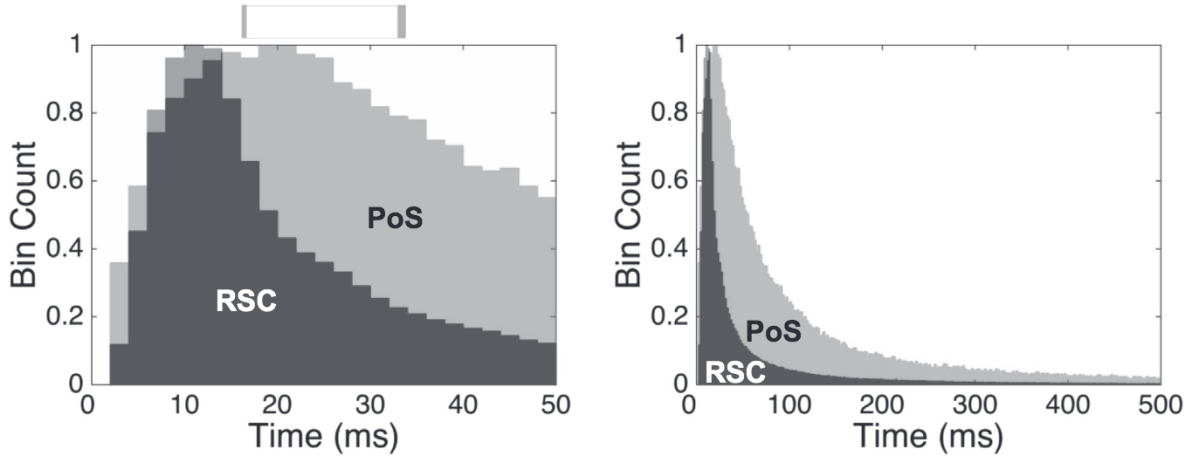

FIGURE 5.11 Interspike interval (ISI) histograms for two neurons, one from retrosplenial cortex (RSC) and one from postsubiculum (PoS), at short (left) and long (right) time intervals, both normalized to the maximum bin count. Note the preponderance of longer ISIs in the PoS neuron, reflective of its slower overall firing rate. The 2 ms refractory period is visible in the left plot. *Plots courtesy of Hector Page.*

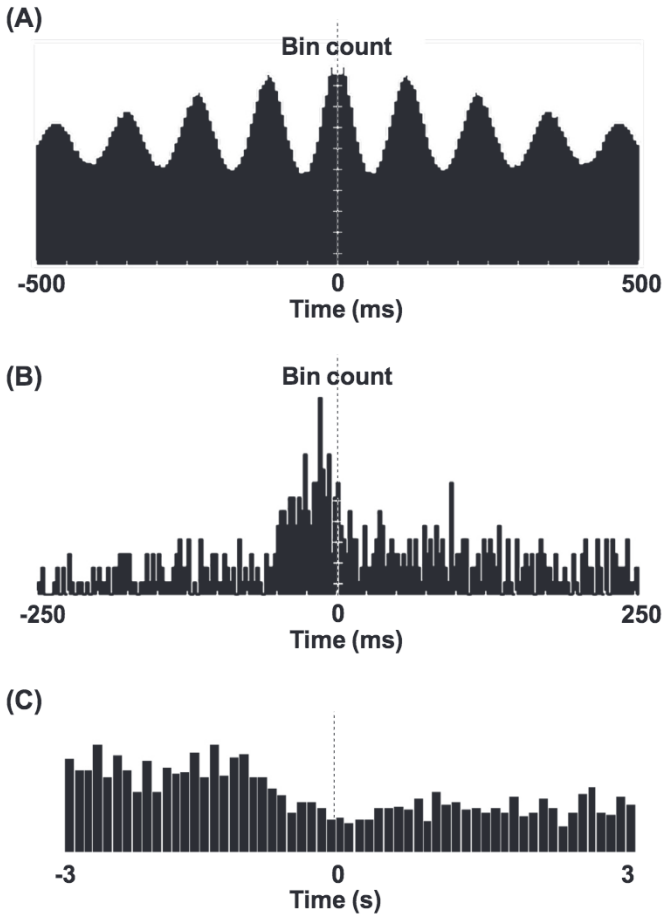

FIGURE 5.12 Temporal correlograms. (A) Temporal autocorrelogram of the cell in Fig. 5.10, compiled by taking each spike as a central reference and plotting the time of firing all the other spikes 500 ms before and after it. It can be seen that peaks of spike firing occurred at about 125 ms intervals: This is because the cell fired in bursts at the theta frequency (~8 Hz). (B) Temporal cross-correlation between the two cells shown in Fig. 5.9. Note the asymmetry of the plot, with a clear peak preceding the midline, indicating that firing of the index cell often tended to be preceded by firing of the reference cell, with a typical interval of about 25 ms. (C) Perievent time histogram of spiking relative to an event (lifting of a food flap) at time 0. The cell's behavior anticipated the event (peak to the left of 0) due to a change in the animal's behavioral state as it approached the reward.

reference, and once with the second as the index and the first as the reference, with the interval thus reversed.

3.2.4 *Cross-Correlation*

Cross-correlation is similar to autocorrelation except that instead of comparing spikes from a single neuron, each instance of a spike from one neuron is compared with all instances of spikes from a different neuron. It is thus not symmetrical: Only one of the neurons supplies the index spikes and only the other supplies the reference spikes, so each interval only occurs once. The cross-correlation histogram is useful for revealing neurons whose activity is correlated, such as the two in the example shown in Fig. 5.12B. It can be seen in this example that there is a peak preceding the central point, which means that firing of an index spike is often preceded, by about 25 ms, by a spike from the other cell. Note that such correlations never prove causation: While one cell may be driving the other, it is also equally possible that both merely have a common cause, e.g., both are related to the theta phase, but slightly different.

3.2.5 *Event Correlation*

Correlations of neural activity with events are plotted in the same way as for correlations with other kinds of neural activity, that is, the event is placed at time zero and occurrences of the neural activity are plotted at the corresponding time interval before or after the event (Fig. 5.12C). Such a plot is called a perievent time histogram or peristimulus time histogram.

3.3 Local Field Potential Analysis

The LFP is generated by synchronous activity of large numbers of neurons active at once, and this activity is oscillatory, often rhythmically so. The LFP signal, obtained by low-pass filtering the raw neural signal, is usually analyzed using a fast Fourier transform (FFT), wavelet analysis, or multitaper spectral density estimation. It can be categorized according to various oscillation frequencies: delta (<4 Hz), theta (4–12 Hz), beta (12–20 Hz), and gamma that is sometimes divided into slow (20–80 Hz) and fast (80–120 Hz; Fig. 5.13).

3.3.1 *Sharp Wave Ripples*

In addition to these generally continuous signals, large-amplitude irregular activity (LIA) events can also be observed. These LIA events include sharp waves that are large, negative polarity events occurring in the CA1 stratum radiatum region of the hippocampus and are often observed when animals are alert but immobile and/or asleep. Often these are accompanied by high-amplitude, high-frequency oscillations in the CA1 stratum pyramidale, known as ripples (Buzsáki et al., 2015), the combination being known as a sharp wave ripple (SWR). Many of these ripples are further associated with spiking activity, whereby neurons such as place cells may fire a small number of spikes each. A closer look at this spiking activity has revealed that it is often sequenced, with cells firing in the same order as they did sometime before the ripple ("replay") or indeed, firing in the same order as they will during the animal's next action ("forward replay"; Foster, 2017). Ripples are usually detected in the normalized squared or root mean squared LFP after it has been filtered to include only the ripple frequency band. Beginning, end, and middle peaks are defined in terms of standard

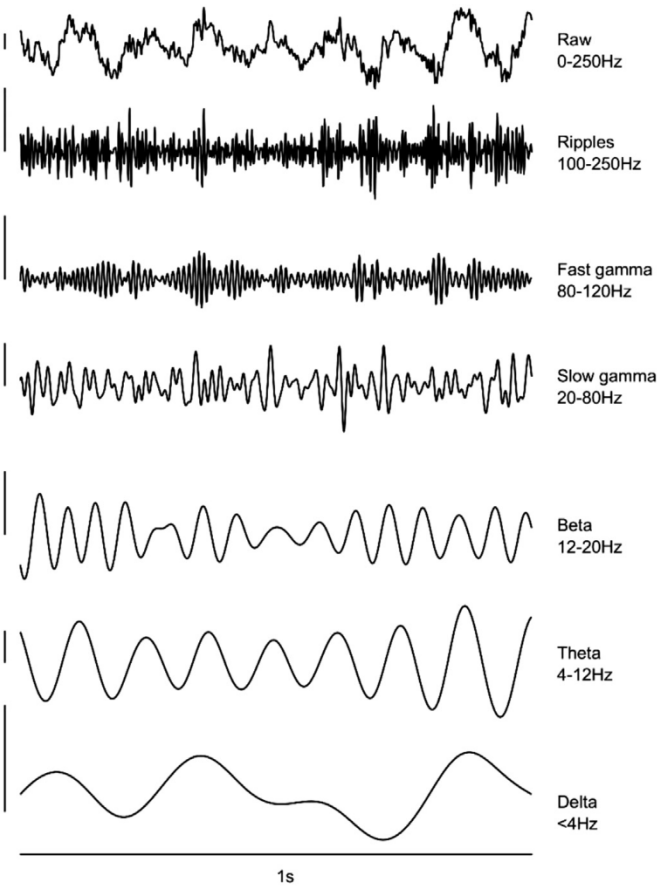

FIGURE 5.13 Illustration of the multiple components of a CA1 hippocampal local field potential trace, showing the original raw trace (top) filtered to reveal the components at various frequencies. Each trace shows 1 s of data, *vertical lines* to the right indicate the scale of 25 μV. *Data courtesy of Eleonore Duvelle.*

deviations above the mean, and groups of peaks that meet these criteria and are not too long in duration are considered candidate ripple events (Maingret et al., 2016; Fig. 5.14).

3.3.2 *Spike–Local Field Potential Relationship*

In addition to analyzing spikes and LFP in isolation, the question arises as to how these relate. This is a complex and not fully answered question because the LFP both arises from spikes and helps to drive them. Part of the neuronal contribution to the LFP is likely to be from the local network of interneurons, but a component also derives from the principal cells, and this leads to a complex relationship, the functional importance of which is much discussed but still unproven.

In the spatial system, the most intensive focus is on the relationship between spiking and theta rhythm, following on from the discovery by O'Keefe and Recce (1993) that place cells, which have an approximately theta frequency burst pattern, in fact shift their phase relationships with theta systematically across the theta cycle. Thus, while a place cell's burst of complex spikes will tend to align with the trough of the local theta LFP, each successive burst occurs slightly earlier until the cell has "precessed" almost a full cycle—360 degrees—of theta. This corresponds to the beginning and end of the place field, and interestingly, precession occurs more slowly the larger the place field so that it is close to 360 degrees every time. The theta phase precession phenomenon has led to much theorization, but whether it has a functional role or is a byproduct of local membrane dynamics is still unproven.

There are two main ways to reveal phase precession: one direct and one indirect. The direct way (Fig. 5.15) is to plot each spike as a function of both location (on one axis) and theta phase (on the other)—the resulting scatterplots are markedly slanted, showing the systematic change in phase with change in position. The indirect way is to simply measure the intrinsic spiking frequency of the neuron and compare it with the LFP theta frequency—if precession is occurring then the intrinsic frequency will be slightly higher.

As well as considering LFP in relation to spiking activity, there is increasing interest in the degree to which the LFP in different brain areas might become synchronized—this is known as cross-frequency coupling (Canolty and Knight, 2010) and is determined by cross-correlating the two LFP traces. It is thought that cross-frequency coupling may allow faster information transfer between the two synchronized structures.

4. BEHAVIORAL MONITORING

Because the purpose of chronic single-neuron recording is to sample activity in the thinking, behaving brain, the control and monitoring of behavior is a critical component of the experimental paradigm. It took some time to appreciate that place cells are exquisitely sensitive to behavioral state (McNaughton et al., 1983a), switching between operating modes when an animal is running versus when it is paused, resting, or asleep. When it is locomoting, its hippocampus is in a theta state—the LFP pattern is dominated by theta, and a place cell will fire complex spike bursts when the animal runs through the region of space that corresponds to its place field. When the animal pauses to groom, rest, consider its options (e.g., in a navigation task) or sleep, the LFP pattern switches into the SWR state, possibly as part of a memory consolidation or planning process (Buzsáki, 2015). Other behavioral states also affect place cell activity, so for this reason it is important to have a method for monitoring the behavioral state of the animal. Below, we first consider the most basic behavior of all, which is where the animal is and how far it is

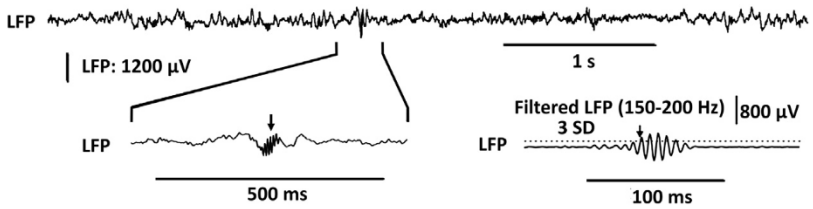

FIGURE 5.14 Ripples in the local field potentials (LFPs) recorded in the left CA1 region of the awake mouse, adapted from Liu et al. (2017). The arrow in the zoomed-in region (bottom left) points at a ripple. The trace on the bottom right shows a high-pass filtered versions of the LFPs to emphasize the high-frequency ripple components of CA1 activity. The horizontal dashed line above the filtered LFP trace represents the mean filtered LFP amplitude plus 3 standard deviations (SD) from a continuous recording of 60 seconds, which is typically used as a threshold (mean ± 3 SD) for automatic detection of sharp-wave ripple activity in the CA1 region. The arrow marks the beginning of ripple activity.

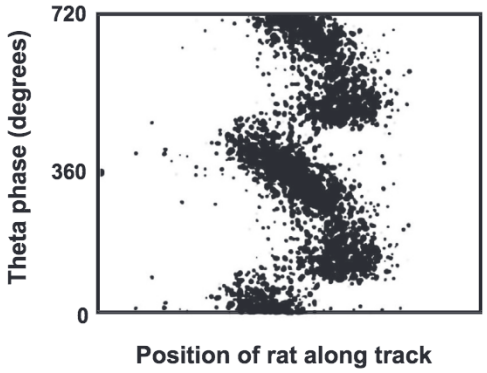

FIGURE 5.15 Phase precession of a place cell showing the spikes from a neuron (*black dots*) as a function of location as a rat ran from left to right along a triangle track (x-axis) and phase of spiking relative to the theta cycle (y-axis). Each spike has been plotted twice, 360 degrees apart, for ease of visualization. Note the downward progression of spike density from left to right, showing the precession in phase. *Adapted from Skaggs et al. (1996).*

moving in what direction; we then consider how to evaluate more complex behaviors.

4.1 Spatial Monitoring and Firing Profile Analysis

The move to recording from freely moving animals brought a new technical challenge, which is how to record the animal's location. At first, this entailed making a Grass polygraph trace of the spiking activity of a cell, by converting the voltage changes seen by the electrode into deflections of a pen on a continuously scrolling sheet of paper, with time on the x-axis and the visits of the rat to various regions of the chamber shown by letters. When it became apparent that location was a critical factor in the activity of these neurons, a more sophisticated place recording method was needed and so location came to be recorded by an overhead camera (Muller et al., 1987). This is still the most frequent method used today (Fig. 5.16) and replaces the temporal coordinate reference frame with a spatial one. However, it is slowly becoming apparent that hippocampal neurons encode both time and space (Eichenbaum, 2014); so new, higher dimensional forms of data display will soon be needed.

Below, we look at some of the basic analyses that are used to extract the spatial correlates of neuronal firing, including position, directional, and spatial patterning.

4.1.1 *Analyzing Place Cells*

The problem with an overhead camera is how to capture the location of the rat in quantifiable form—the method that was originally developed, still widely used today, was to attach a small LED (color, white, or infrared) to the headstage (the part of the recording system that attaches to the animal's implant) and use an electronic method of detecting the bright spot, finding the x−y coordinates of its centroid and concatenating these points to generate a path to which spikes can then be associated. More recently, methods for extracting position and heading from the entire video image have been developed (e.g., Hédou et al., 2001), but these are expensive in their use of computing resources and time, so the spot-tracking method is still most popular.

An example of the data yielded by position tracking during recording of a place cell is shown in Fig. 5.16, where the path of the animal (the temporally ordered, concatenated x−y positions) is shown as a black line, accumulated across a trial (Fig. 5.16A). The spikes (Fig. 5.16B) are extracted and collected by a computer, for later offline analysis. This entails separating the spikes based on their differential waveforms on the different electrodes (Fig. 5.16C) and using the waveform parameters to identify the clusters belonging to each cell (Fig. 5.16D). The spike and path data are then combined (Fig. 5.16E). The path is typically slightly smoothed with a boxcar algorithm, which entails taking each point and replacing it with the average of itself plus a certain number of its neighbors (typically one or two) on either side, to remove positional noise. An alternative to a boxcar algorithm is a Gaussian one in which the contribution of successively remoter bins is weighted by a Gaussian function of their distance from the central bin. After smoothing, the spikes are then overlaid on the path at

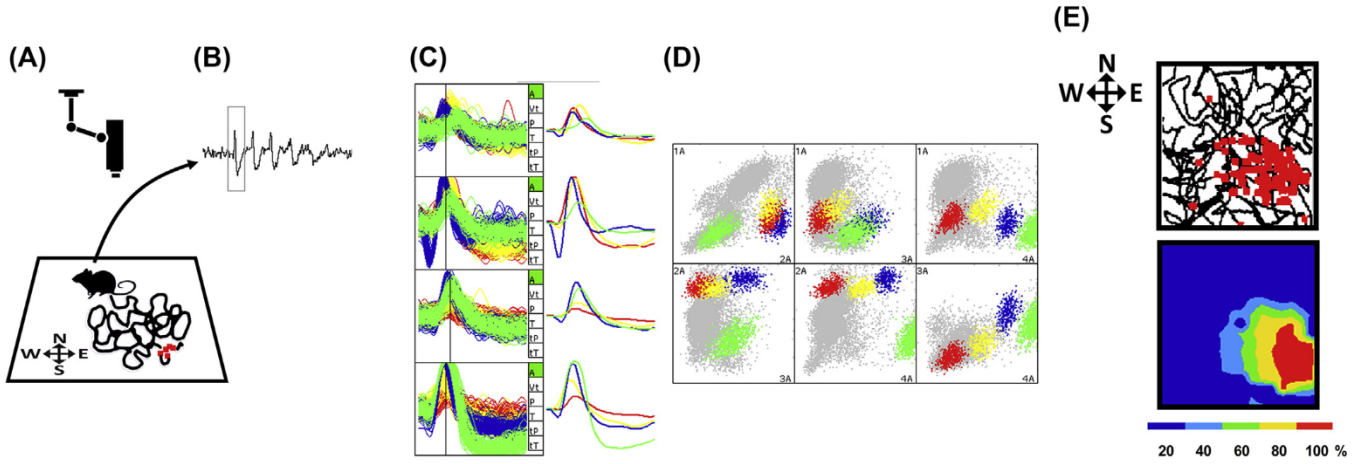

FIGURE 5.16 Recording of a place cell made from a rat exploring a square arena while being monitored by an overhead camera (A). Spikes appearing in the recording trace are collected and time-stamped (B). In offline analysis they are separated on the basis of their differential waveform shapes on the four different electrodes (C) using a clustering method (D). The spikes are then overlaid as points onto the path of the animal and resulting plot is then binned, smoothed and analysed (E). *Figure by K. Jeffery et al., 2015, available at https://doi.org/10.6084/m9.figshare.6429470.*

the point where each occurred. The resulting "spike-plot" makes the location specificity of the cell's activity very obvious.

Fig. 5.16E shows how these data are then postprocessed, using the method originally developed by Muller et al. (1987)—the environment is binned into small squares, and the firing rate of the neuron in each square computed; the resultant ratemap is then smoothed, again usually with a boxcar or Gaussian algorithm except this time in two dimensions. The map is then expressed as a heatplot. Key parameters of the spatial firing can then be computed, of which the most common are as follows:

- Spatial information score
- Coherence
- Aspect ratio
- Similarity score

The spatial information score is a measure of how much localizing information the spikes of a particular place cell carry—that is, given that a spike occurred, how accurately can one infer the rat's position? It is measured in bits per spike. One bit localizes position to within one half of the environment (Fig. 5.17A): Note, however, that this measure does not take into account the spatial colocalization of the activity, which relies on a measure of coherence—high coherence means the activity is concentrated in one place, whereas low coherence means it is distributed in multiple fragments.

Another measure of firing field characteristics is the aspect ratio, which is the ratio of the long to the short axis of the firing field. Place cells do not always (or even usually) produce smooth, symmetric fields, but grid cells (see Section 4.1.2) do, and so the aspect ratio becomes of interest when examining deformation of

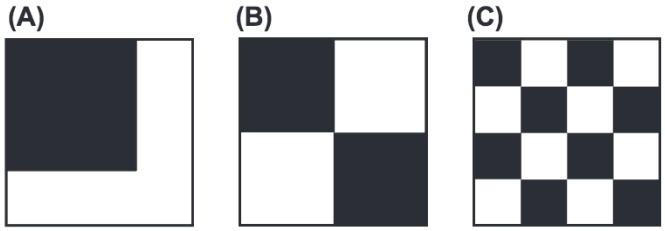

FIGURE 5.17 These patterns of activity all have 1 bit of spatial information content—the activity of the neuron (shown in black) occurs over exactly half the environment, and so if the neuron is active then the place the animal can be is constrained to half of the total area. The difference between the patterns reflects the difference in *coherence*, ranging from high coherence (A) through medium coherence (B) to low coherence (C).

fields resulting from environmental deformation. An example of such deformation is found in Hayman et al. (2011), where on a vertical wall, both place and grid cells were found to be elongated (increased aspect ratio) in the vertical versus horizontal dimensions.

The final measure of firing fields that is commonly used is the similarity, in location and firing rate, to the field from a different situation—either a different cell or the same cell on a different trial or different parts of the same trial. We call this a similarity score, but it is also sometimes called a remapping score, to reflect the fact that place and grid cells often change their firing—remap—when something about the situation is changed. The score is taken by computing the Pearson's correlation between each pixel in the first ratemap and its homolog in the second, ignoring pixels where neither cell fired. For place fields that look by eye more or less the same to an experimenter, the similarity score can range between 0.4 and 1.0 (or close to it), whereas below 0.4, they tend to look as though they have remapped.

Note that the score can go below zero: Nonoverlapping fields will have a negative score. Although the 0.4 score threshold is in common use, a preferable method is to determine the chance value empirically, usually by a shuffle method in which every field in the data set is compared randomly with the others. This procedure will yield the random similarity score, which is generally not zero because for a typical environment, and given typical place field sizes, there is usually a degree of overlap between randomly placed fields.

4.1.2 Analyzing Grid Cells

Three decades after place cells were discovered, a new cell type appeared on the scene, the properties of which necessitated some new types of analysis. These so-called "grid cells," found in entorhinal cortex (Hafting et al., 2005) and pre- and parasubiculum (Boccara et al., 2010), generate place fields that are more circular than those of hippocampal place cells; they are additionally all more or less the same size and evenly spaced (Fig. 5.18). The even spacing, present even in a novel environment, immediately suggested some type of intrinsic distance-measuring process, the nature of which is still unknown.

Many entorhinal cells produce firing patterns that have a degree of spatial inhomogeneity (Hardcastle et al., 2017), but it is generally accepted, rightly or wrongly, that for a cell to be a grid cell, its spatial firing needs to be gridlike, that is, regular. This has led to a need to quantify "gridness." There is no completely satisfactory way of doing this: Every method has the effect of excluding cells that look to be grid cells by eye but which do not meet the analysis criteria, and simultaneously including cells that meet the criteria but which show messy or irregular firing, leaving the experimenter with the difficult task of deciding where to draw the boundary. The most common selection criterion has been the grid, or gridness, score, which is derived by a double autocorrelation procedure, one linear and one rotational (Sargolini et al., 2006). First, the ratemap is autocorrelated by taking the map, shifting it one pixel at a time in every direction and recorrelating; the resulting plot has a periodicity that reflects the underlying regularity of the firing pattern (Fig. 5.18C). The autocorrelogram so produced is then subjected to the same procedure, but this time it is rotated by 1 degree at a time and recorrelated: This produces an undulating line that reflects the 60-degree periodicity of the grid. The grid score is usually taken to be the minimum value at 60 and 120 degrees minus the maximum value at 30, 90, and 150 degrees.

4.1.3 Direction Tracking

Around 15 years or so after place cells were reported, it emerged that some cells, called head direction cells, are responsive to the facing direction of the rat's head (Ranck, 1984; Taube et al., 1990a, 1990b; see also Chapter 9 (Mehlman & Taube), this volume)—it thus became necessary to record direction as well as location. A crude way to do this is to look at the direction of travel of the rat and assume that its head is facing in the same direction—a more direct method was introduced by Taube et al. (1990a) and still widely used, employing two lights separated by a few centimeters, and distinguished by their size or color. A line drawn between the lights is taken to be a proxy for head direction. This method works very well when an animal is foraging over a horizontal surface with its head at a fixed angle to the floor, as when pellet chasing but starts to work less well when the animal can move its head through vertical space as when rearing or climbing, or when an environment is complex and the LEDs sometimes occluded. Solutions to these problems are described later in Section 4.2 on complex behaviors.

Often, an investigator will be interested in both the directional and spatial firing properties of a neuron. Even hippocampal place cells sometimes show directional firing (McNaughton et al., 1983a), but conjunctive spatial/directional firing is frequently observed in parahippocampal cortical areas such as presubiculum

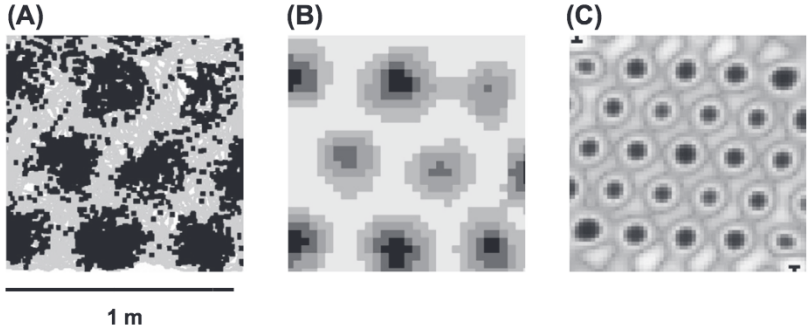

FIGURE 5.18 Activity pattern of an entorhinal grid cell. (A) The spike plot, with spikes as *black dots* superimposed on the gray trace, representing the rat's path as it foraged over a 1 m square area. (B) Ratemap, with dark colors as high rate. (C) Autocorrelogram, revealing the periodicity of the pattern. *Adapted from Barry et al. (2007).*

(Cacucci et al., 2004) and entorhinal cortex (Sargolini et al., 2006). One problem with determining the conjunctive firing properties of a neuron, which was noted early on by Muller et al. (1994) and investigated in detail by Cacucci et al. (2004), is that if a neuron is directional, then in a small arena its firing will be spatially inhomogeneous. This is because not every head direction can be equally sampled in every spatial location, and so spikes will tend to accumulate on those regions of the environment where the animal is able to face in the direction preferred by the cells. When recording place and direction together, it is thus necessary to compensate for the inhomogeneity of place-by-direction firing (Burgess et al., 2005).

4.1.4 *Speed Tracking*

As position and direction, soon after place cells, were discovered, it became apparent that speed is an important modulator of place cell firing particularly because of the switch from theta mode to SWR mode, which is highly speed dependent (Buzsáki et al., 1983; Buzsáki, 2015). One consequence of this switch is that clusters of spikes from a place cell will appear at pause points, generating what looks like a place field at that location, which is generally a location that is of significance to the animal. This may lead an observer to erroneously conclude that the cell is encoding the feature of interest (e.g., a reward location). For this reason, investigators who are interested in pure place coding generally "speed filter" their data so as to exclude epochs of a trial in which an animal's speed falls below a given threshold. In early studies, however, this was frequently not done, and early reports of behavioral correlates of place cell firing may have been looking at SWR-associated activity. Given that not every cell is active during an SWR; however, it becomes of interest to determine under what condition a place cell will also activate in an SWR burst.

4.2 Complex Behavior Tracking

The foregoing discussion examined tracking of foraging behavior in simple environments, but often the environments and/or the behavior are not simple, and there is a growing interest in studying more complex, naturalistic behaviors. Complex environments will be discussed first, and then complex behaviors.

4.2.1 *Complex Environments*

The standard method for constructing a ratemap was described in Section 4.1.1—sometimes, however, sampling of an environment is so patchy that this approach does not work well. For example, if a bin contained 1 spike and 0.5s of position data, the firing rate for that bin would equal 2 Hz. This is correct but intuitively seems less convincing than if there were 200 spikes and 100s of data, which would also equal 2 Hz. Traditionally, cutoffs are employed to compensate for this, such as a minimum amount of time an animal has to spend in a bin for it to be considered "visited." To retain as much information as possible, more advanced approaches can also be used. For instance, Skaggs et al. (1996) introduced an "adaptive binning" method. Rather than using a fixed grid spacing for the two-dimensional (2D) histogram, this method resizes each bin until it contains at least 1s of position data. Firing rate is then calculated as the number of spikes contained inside this new bin divided by the time spent there. This technique is often employed in grid cell experiments where unvisited bins can drastically interfere with the metrics used for assessing grid quality (Wills et al., 2012). It has also been used for processing 3D data from freely flying bats, where there are often gaps of unvisited bins between flight paths (Yartsev and Ulanovsky, 2013).

Place and direction in two dimensions have been the standard parameters of interest for many years, but recently attention has been turning to the recording of more complex variables, or recording in larger spaces or in complex spaces from where a head-mounted LED may not always be visible. An example is our study of a multicompartmented environment (Spiers et al., 2015), for which a single overhead camera was unable to see into every corner because of occlusion by the inner walls. This problem was solved by recording with two cameras and concatenating the path data.

Multiple cameras are also useful for recording location in three dimensions. This is becoming increasingly necessary as investigators start to move from the study of place encoding in two dimensions, to the encoding of three (Jeffery et al., 2013). For a number of years this has been possible through the use of paired cameras, essentially using stereoscopic information such as binocular disparity to infer scene depth, in the same way the brain infers depth using a pair of eyes. This technique is effective but largely relies on brute-force search algorithms to determine the disparity between each pair of camera pixels, which is computationally intensive. The positions of the cameras, like eyes, are also limited; they must be aligned and cannot diverge too drastically from one another. A faster and more flexible approach uses linear transformation algorithms, essentially building a map of pixel correspondences between two camera views. To reconstruct scenes in this way, it is necessary to know the position of each camera relative to the other. This can be achieved by having known pairs of corresponding points: For instance, if one images a checkerboard pattern that has a specific orientation and size through both cameras, it is possible to infer their distance and orientation relative to it and thus to each other. These camera matrices are used to construct a fundamental matrix. If x is a set of points viewed by camera

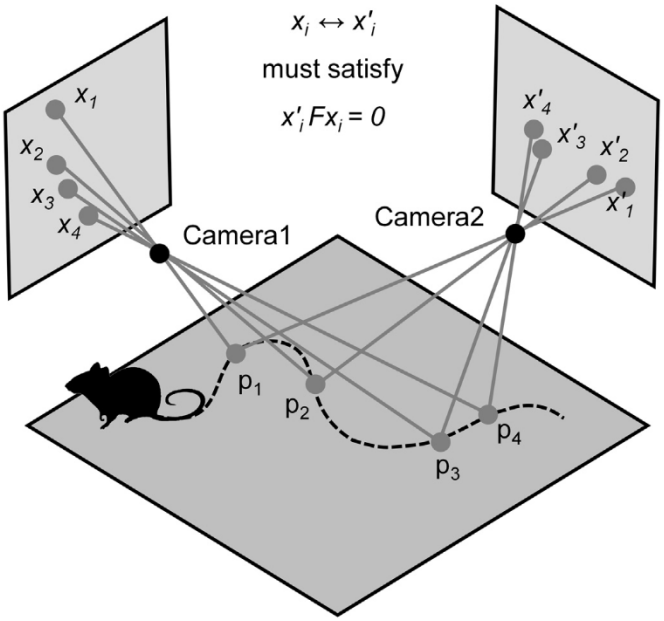

FIGURE 5.19 Two cameras view the path of an animal; four points are projected through the cameras, which for simplicity are imagined here as perfect pinhole objectives, to the camera sensors. The arrangement of these points on the two sensors can be used to infer the real, 3D points. *Figure by K. Jeffery and R. Grieves, available at https://doi.org/10.6084/m9.figshare.6429548.*

1 and x' is the same set viewed by camera 2 (Fig. 5.19), the fundamental matrix, F, represents the relationship between points x and x':

$$x'Fx = 0$$

We can do this because we know they must be the same points in 3D space. Once we have calculated the fundamental matrix, this can be used to triangulate any given pair of points imaged by the cameras into 3D space. In this way, the two cameras can be positioned anywhere at any rotation or offset to one another.

This approach has become readily available only recently, driven in part by advancements in machine vision and robotics (Hartley and Zisserman, 2003). It is more flexible than the optical disparity methods; however, performance improves further using multiple cameras.

Unfortunately, extending these linear algorithms to incorporate more than two cameras is difficult and with every added camera the difficulty increases. An easier approach is simply to perform a 3D reconstruction using every possible pair of cameras and then combine these reconstructions into one single space using a statistical approach, such as the weighted mean of each point where the weighting is the certainty of the point's location. With this setup, an object only ever has to be viewed by two cameras for a successful reconstruction, and because there is no limit to the number of possible cameras, continuous tracking is possible even in cluttered, complex environments. For each of these reconstruction approaches it is imperative to know corresponding points between cameras, i.e., the 3D points being viewed by two cameras are the same at each time point; without this certainty any reconstruction is impossible.

Head direction in 3D is also complicated. Finkelstein et al. (2015) used a tetrahedral array of colored LEDs to track heading in flying bats. In a cluttered environment, however, such methods fall prey to issues with occlusion: A possible solution is to add accelerometer information to the recordings to fill in the gaps when the visual tracking fails (indeed, this is how the brain keeps track of its own heading).

We turn now to the issue of tracking complex behaviors. As the field has advanced, other types of behavior in addition to simple speed have become of interest; such behaviors include exploratory actions such as rearing, planning actions such as vicarious trial and error, self-care behaviors such as grooming or eating/drinking, and social behaviors such as submission, threats and fighting, courtship and copulation, and suckling and infant care. Now, instead of simply tracking location via a head-mounted LED, it becomes necessary to capture the entire image, and then find a way to identify and measure the behavior. In the early days, this was done by storing the recordings on videotape and then replaying the tape while it was observed and scored, ideally by an observer blind to the experimental condition. These days, images are captured digitally but this introduces problems of data rate/storage; while algorithms exist to compress the images, data rate remains a problem in behavioral tracking. In addition, as data rates increase, it becomes increasingly important to find less labor-intensive ways to score behavior; in future it is likely that such tasks will be accomplished using machine learning/artificial intelligence.

5. BEHAVIORAL CONTROL

The importance of spatial coding lies in its relevance for action and so to investigate the functioning of the spatial system it is important to be able to study behavior experimentally; this means not just monitoring it but also controlling it. Spatial coding per se has generally been studied by trying to persuade an animal to explore the entirety of an environment as uniformly as possible, using a "task" known as pellet chasing, introduced by Muller et al. (1987). By scattering the food uniformly over the environment, rats and mice can be persuaded to abandon their natural tendency to stay close to the edges of an arena (called thigmotaxis) and venture out into the center and to

cover the environment homogeneously. These days, a variety of different food rewards are used, particularly sweetened cereals such as Coco Pops, sweetened rice, or liquids such as chocolate milk. Additionally, sometimes investigators will use electrical stimulation of the reward system (medial forebrain bundle, ventral tegmental area, or nucleus accumbens) to deliver reward quickly and easily and without satiating the animals. This method has the complication of exogenously driving neural reward systems, which are linked to the hippocampal system (Oades and Halliday, 1987) and may thus affect it, and so it is not widely used, but it has proved to be very useful in some settings (e.g., see Knierim and McNaughton, 2001, who ingeniously used it to reward animals in microgravity during space flight).

In a pellet-chasing task it is necessary, as well as controlling the environment coverage, to constrain the animal within a restricted region of the environment. This is usually done using walls or edges. Rodents prefer opaque walls, which cater to their natural tendency towards thigmotaxis. The height of the wall, or the height of the edge above the floor, needs to be enough that a normal laboratory animal would not attempt the jump—we typically use around 50 cm for rats and 30 cm for mice, although our enriched-housing animals are bolder and will jump further/higher. Jumping can also be thwarted by making the walls thin and slightly unstable.

Walls have the disadvantage that the headstage bumps against them when the animal is moving and generates noise, so some investigators prefer to leave the environment open and constrain the animal by the edge. If an edge is used instead, be aware that if the surface is too close to the floor otherwise the animal may try to jump off. Even if not, the animal will peer over the edge, a pose known as attentive scanning, and in doing so generate new place fields (Monaco et al., 2014).

Occasionally, an experimental question requires the absence of boundaries—for example, the question of whether place cells need boundaries to form place fields or whether grid cells need developmental experience of boundaries to form grids. To address the former question, we devised a nonphysical, virtual boundary using white noise (Hayman et al., 2008); when the animal encountered the boundary and triggered the loud noise it would retreat away from it again. To avoid the animal using the virtual boundary as a resetting marker for path integration, we also varied the exact location randomly so that the boundary was indistinct—this so-called "fuzzy-boundary arena" proved able to keep the animals corralled in the center of the room in a way that did not provide spatial information.

Pellet chasing is an important behavior, but in cognitive terms, it is rather simple because an animal only needs to move from one food item to the next, a few centimeters away, and does not need to track its location. That place, head direction, and grid cells update their activity anyway was an important demonstration that self-localization is an automatic process, occurring even when not explicitly required for the task at hand. It was the pellet-chasing task that allowed researchers to discover grid cells—without the abnormally uniform coverage of a large (2-m diameter) open arena that it enabled, the grid pattern would likely not have been recognized (indeed the first author, to her lasting chagrin, recorded but failed to recognize—due to the small arena size—a grid cell, in 1994). It should be appreciated, however, that such behavior is not natural for rodents, and it remains unknown how a typical grid cell would behave during the lifetime of a normal animal—it is unlikely that a regular grid is a normal part of a grid cell's repertoire.

In pellet chasing an animal need not track its location (although it probably does anyway as mentioned above); the next level of behavioral complexity is place recognition, where an animal demonstrates that it knows where it is, even if it did not need to navigate there. This means experimentally defining one or more unmarked parts of the environment to have significance and then monitoring behavior at that location. An early use of this was in a water maze task where rats with hippocampal lesions were shown to be unable to recognize the location of a hidden escape platform (Hollup et al., 2001a); using this task, it was subsequently shown that in intact animals, place cells with fields near the platform would migrate their fields toward its location (Hollup et al., 2001b). In another popular paradigm, this time using negative rather than positive emotional significance ("valence"), an animal demonstrates its recognition of the unpleasantness of a place by using avoidance and/or escape. This method has been used to show that a population of place cells can express their fields in more than one reference frame, switching from one to the other (Bures et al., 1997).

Moving on in cognitive complexity, we come to tasks where an animal needs to actively make its way to a goal. The simplest type of goal-directed navigation is beacon navigation, where an animal heads toward a visible (or otherwise detectable) goal, a task which does not involve the hippocampus. As mentioned earlier, goals—reward zones—introduce problems because locomotor behavior becomes interrupted at a goal, and the LFP state switches back and forth between theta and SWR modes. However, there is obviously a great theoretical significance in the question of whether place and other cells encode goal information, and so it has been necessary to overcome these difficulties. One way of accomplishing this is to isolate periods of time in a goal-directed task when the animal is still actively locomoting and therefore not generating SWRs.

For example, Wood et al. (2000) studied place cell activity on the central stem of a reentrant T-maze (one where there is a path from the goal back to the start) and found that the cells changed their firing rates—a phenomenon now known as rate remapping (Leutgeb et al., 2005)—depending on whether the animal was on a left-bound or right-bound trajectory. Thus, there were signatures of goal encoding in the navigation part of the trial where the animal was still in theta mode. An alternative is just to accept the change in LFP state and look at cell activity when the animal is in SWR mode while stopped at a goal (e.g., Wood et al., 1999).

More dynamic behavioral interactions occur when an experimenter uses the animal's own behavior to modulate the state of the environment. A good example of this is the place preference task (Rossier et al., 2000), where an animal needs to visit an unmarked goal location, defined by its place in the room (much like a water maze platform is) and reward is delivered when the animal has waited there for a period of time of usually a few seconds—thus indicating that it knows where it is. The position of the animal is tracked by computer, and the reward is dropped from a height and bounces randomly across the arena so that the animal now needs to search for it, thereby switching into theta mode and undertaking a foraging exercise. Such tasks are hard for rodents because waiting across delays does not come naturally to them. However, with this method it has been possible to show that place cells show a degree of goal encoding (Hok et al., 2007).

Recent years have seen the development of an ingenious new paradigm, virtual reality for rodents, in which a simulation of the environment is presented to the rat that moves past its eyes at a rate dictated by the locomotor rate of the animal. In this apparatus, pioneered by the Mallot group (Hölscher et al., 2005), the animal runs on top of a suspended lightweight ball, turning it with its feet at a rate that is detected by a computer and used to generate a dynamic visual display appropriate to the movement. The hope is that the combination of locomotion and the world appearing to move will fool the animal into thinking it is running through real space; however, the head fixation means that more sophisticated physiological recording, such as patch clamping, can be undertaken. Although this paradigm has produced some extremely interesting findings (e.g., Harvey et al., 2009), it is clear that the experience of the animal is not completely natural (the animal is not fooled) because grid and head direction cells have not yet been reported in head-fixed situation (although, interestingly, firing fields appear, in rescaled form, if the animal can turn its head; Aronov and Tank, 2014). Nevertheless, this method allows dissection of some of the subcomponents of spatial processing and has yielded important insights, such as that even primary sensory processes can be influenced by locomotion (Saleem et al., 2013).

6. PITFALLS FOR THE UNWARY

Place cells have been known about for approaching 50 years now, and the original cohort of place cell researchers is steadily being joined/replaced by the next generation of young (and, these days, not-so-young) scientists. This risks a degree of wheel reinvention, if newcomers to the field do not acquaint themselves fully with an older and sometimes less sophisticated literature. A number of factors that people commonly fail to account for when planning and executing place cell experiments, starting with elementary mistakes that novices often make and then moving to more subtle issues are listed below. The reader is also referred to a recent discussion of how to improve data quality in single-neuron recording by Harris et al. (2016).

6.1 Noise

One of the commonest complaints of a beginner is "My recordings are noisy," and reducing noise is one of an electrophysiologist's greatest skills. Causes of noise were discussed earlier. The diagnostic path requires isolation of the source of the noise, a process of consecutive hypothesis testing and zooming in on the problem that is a microcosm of scientific enquiry in general. A quick guide for tracking down noise is shown in Fig. 5.20.

6.2 Mistaking Noise for Spikes

The commonest error in recording is to collect noise, thinking it is spikes. The first author once spent a happy afternoon recording and clustercutting a place cell with a beautiful place field, only to discover after 2 or 3 h that a temporal autocorrelation yielded a sharp peak at 20 ms—the "cell" was 50 Hz mains interference, which for some reason only occurred in one corner of the apparatus. Electrical noise, when filtered, often looks spike-like, and for this reason, it is important always to (1) listen to the cells you are recording from with a speaker because the human ear is exquisitely sensitive to auditory patterns and easily distinguishes spikes from noise and (2) after a trial or two, run the data through a few simple screening analyses, including autocorrelation, to make sure a putative cell has a 2 ms refractory period and no strange temporal artifacts.

Another common reason for collecting noise is simply that the spike capture threshold was set too low, and

data were collected from distant cells whose spikes barely peeked above the threshold. It is tempting to do this in order not to miss out any data; however, the end result is a large number of spikes with a low signal:noise ratio, which can be difficult to analyze or draw clear conclusions from. It is sometimes possible, with a lot of manual refinement of the cluster, to identify a place field, but experience has shown that these are rarely stable enough over the course of a recording session to yield useful conclusions, and the time spent on the refinement was wasted. As a rule of thumb, if a cell does not have a signal:noise ratio of at least 2:1 it probably is not worth analyzing.

6.3 Over- and Underclustering

The most common error in spike separation is overclustering—that is, dividing the spikes from a cell into more than one cluster. It commonly occurs in the hippocampus because of the decrescendo complex spike burst, which produces spikes of different amplitudes. These show on the scatterplot as long thin clusters aligned with the origin, which are sometimes broken up into more than one cluster by an automated algorithm, in which case they need to be rejoined by hand. Very occasionally two cells will have clusters that fall on a line, but these can be distinguished either by forming clear clumps (complex spikes do not do this, as they

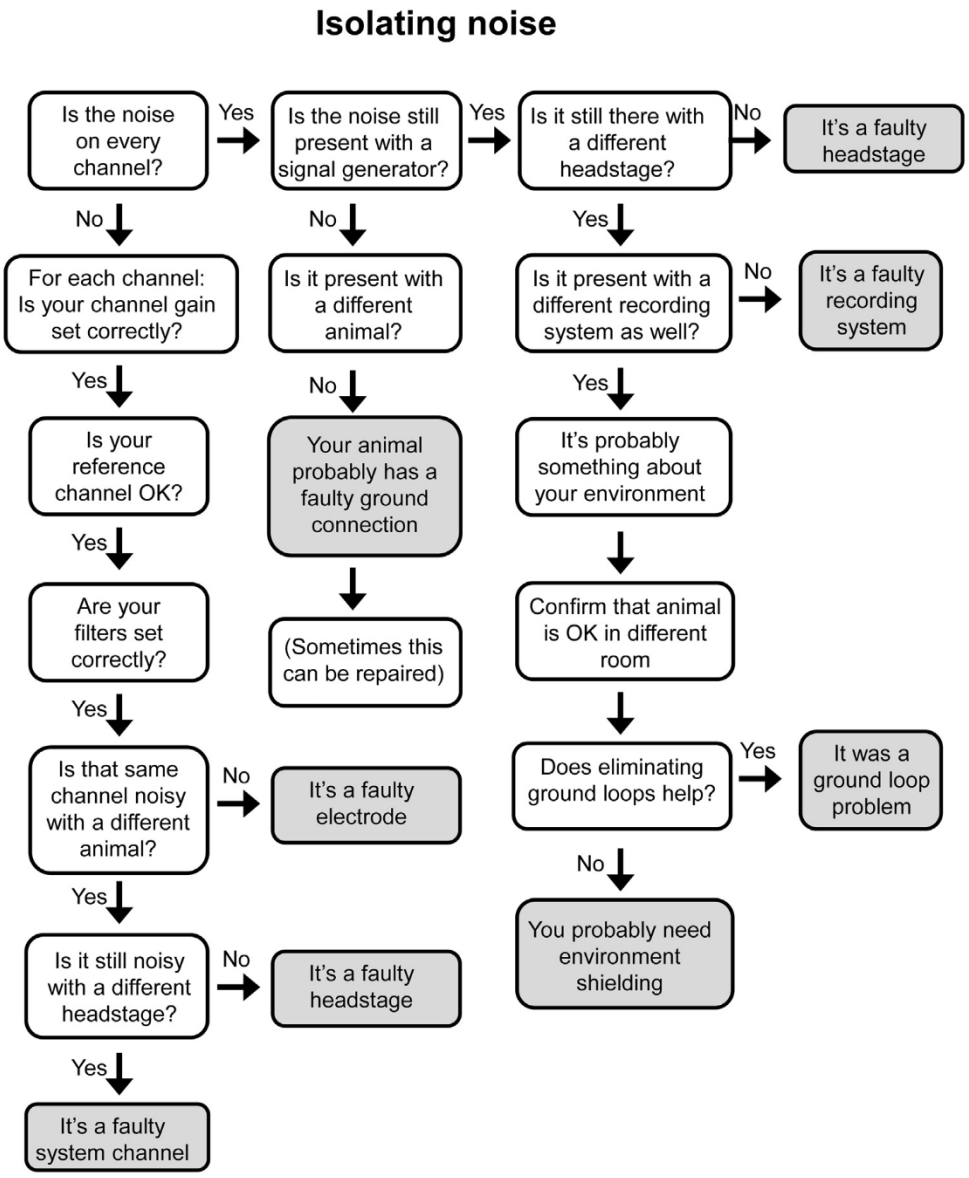

FIGURE 5.20 Flowchart for tracking down sources of noise. *Figure by K. Jeffery and J. Donnett, available at https://doi.org/10.6084/m9.figshare.6409640 under a 1029 CC-BY 4.0 license.*

are continuously distributed in size) or by a temporal cross-correlation, for which the presence or absence of a peak at the complex ISI (around 2 ms) will distinguish one versus two cells. Overclustering also occurs because automated algorithms tend to overcluster, and the results usually need to be manually refined to rejoin clusters belonging to a single cell. In the hippocampus it is easier to see when overclustering has occurred—in cortical regions such as entorhinal cortex, where spikes from multiple cells often form one large cluster, then it is often necessary to use temporal cross-correlation to decide whether nearby clusters actually come from the same cell. For cells with spatial or directional correlates, the firing fields and directional tuning curves can be useful in validating a decision (although not, of course, for making a decision—clustering must always be done using waveform parameters).

Underclustering occurs when two cells are treated as one and is also a common mistake made when an experimenter fails to spike separate cells having highly overlapping clusters in the cluster space. This is less common now that automated clustering methods are available. If two cells have differing place fields then one can usually look at the clusters corresponding to the two fields and see that they are slightly offset relative to each other—but of course, clustering must be done on some waveform parameter that distinguishes the clusters, not on the place fields. If place fields overlap, then as discussed above, a cross-correlation will reveal whether there are spikes occurring in the refractory period that signal the possible presence of a second cell, if its place field overlaps with that of the first. It is, however, sometimes just not possible to perfectly spike separate neurons. These days reviewers frequently ask for quantitative evidence of cluster separation quality such as the isolation distance (Schmitzer-Torbert et al., 2005), although there are not yet agreed standards for acceptable differentiation.

6.4 Cells Sometimes Have Inverted or Complex Waveforms and/or Dual Clusters

Sometimes, inverted spikes will appear on a recorded trace. If these are on all channels then they are almost invariably due to spikes on the reference electrode and are inverted due to the subtraction. However, sometimes they are not on all channels, and inspection of the reference channel does not reveal spikes—in this case these are true inverted spikes. A related issue occurs when a cell has a complex, multipeaked waveform, which sometimes leads to its spikes being separated into two clusters, depending on which peak triggered each spike capture event (this can be diagnosed when two clusters both have strange waveforms, temporal cross-correlation reveals a high degree of correlated firing, and the place fields are identical). And finally, sometimes a cell's spikes seem to be unusually huge. These phenomena occur when an electrode is nestled very close to the cell soma or perhaps between two dendritic branches or the axon hillock—in these cases it is seeing a different mix of current sinks and sources than the usual simple biphasic profile. These cells can be recorded in the usual way (merging two clusters where necessary in postprocessing), but one needs to check that a cell that is, say, inverted on one channel is not also being collected uninverted on another and therefore recorded twice.

6.5 Spurious Place Fields Appear When Environmental Coverage Is Nonuniform

It is natural that interest is now moving from pure place coding to the more complex issues of planning and navigation, but as discussed earlier, almost all such tasks introduce spatial heterogeneity into the location data from the animal. Methods for place field analysis attempt to control for this by normalizing each pixel's firing rate by the time spent there (dwell time normalization), but these do not always completely eradicate the artifacts—for example, one or two spikes occurring in a pixel where the animal spent almost no time can result in a spuriously large pixel rate, which can resemble a place field. These artifacts can mostly be removed by appropriate smoothing and choice of place field coherence measures or by adaptive binning, but some may remain. Thus, a putative place field ideally should be verified by recording it in the same place on more than one occasion to show that the spatial pattern is stable. Sometimes this is achieved by dividing the trial into two halves temporally and then correlating the maps from the first and second halves. It is also highly recommended to bracket each set of trials with a baseline at the start and the end of the session and correlating these maps to check that the field is stable. This method runs into problems, however, if the intervening experience of the rat caused place cells to remap their fields, as may be the case if learning or other experimental intervention was involved. It helps, when presenting place cell data, to always present the spike plots as well as the ratemaps because the experienced eye is more convinced by seeing the actual spike data than by the binned, smoothed, and otherwise manipulated heatplots.

6.6 Spurious Goal Encoding When an Animal Switches Into Sharp Wave Ripple Mode

One of the most common problems plaguing experimental design concerns how to determine whether a cell has a place field, or some other type of

goal-related encoding, when an animal is always stopped at a particular place such as a decision point or reward zone. It is important to record LFP so that one knows what state the animal is in, but this does not remove the problem that it is not possible to record a place field from a stopped animal, and conversely, that a stopped animal generates multiunit activity that resembles place fields. One way around this has been to make a place rewarding only some of the time, with that time being made explicit to the animal, so that recordings of the same place can be made during both goal-directed activity and foraging. For example, we undertook a task in which an animal foraged for food in an arena, and then, on hearing a tone, ran to one part of that arena for food (Jeffery et al., 2003). We did not see clear evidence of goal encoding during the foraging phases; however, such encoding has been observed in other experiments (Hok et al., 2007; Dupret et al., 2010). These differences in experimental outcomes may be due to place cells sometimes remapping when task demands change within the apparatus (Markus et al., 1995).

The other problem with interpreting apparent goal encoding is that place fields are modulated by attention to spatial cues (Kentros et al., 2004; Monaco et al., 2014), and an animal is more attentive to such cues when it reaches a goal, as evidenced by, for example, rearing on a water maze platform (Sutherland and Dyck, 1984). When a place cell shows increased activity near a goal, does this point to a specific role in encoding reward locations or is it just a by-product of the cells receiving heightened spatial information at that place? This may not be a functionally meaningful distinction, but it should be borne in mind when interpreting goal-related firing. There is evidence that place cells can become linked to the reward value of a location—for example, stimulation of place cells concurrently with rewarding stimulation during sleep caused animals to go to that place on awakening (De Lavilléon et al., 2013). Whether such linkage is in the inputs to the place cells, or the outputs, remains to be determined.

6.7 The Stability of Place Fields Across Trials Does Not Prove That Place Cells Have Memory

Because of the role of the hippocampus in memory, and the plastic nature of its synapses, there is naturally much interest in the degree to which place cells can learn about changes to the environment. It is assumed that such learning would involve relocation of place fields (remapping), but if an environmental change is used to induce the remapping, then a problem arises as to how to dissociate the sensory and memory components of the remapping. It is not necessarily the case that just because a place field reappears across successive trials that the cells therefore have learned about the place—it could just be that the same sensory inputs drive the field every time. That said, there is now increasing evidence that in fact some or all fields are learned because new fields appear where they were previously absent when an animal scans the place (Monaco et al., 2014) and can even be artificially generated by electrical stimulation (Dragoi et al., 2003). This may, however, be due to local synaptic plasticity in the immediate inputs to the cell, which is memory of a sort, but not necessarily of the more cognitive episodic type that some investigators assume.

6.8 Place Cells Are Sensitive to Nonspatial Contextual Cues

The most striking aspect of place cells is the spatial localization of their firing (i.e., their place fields), but a large body of literature exists to show that place cells can be modulated by nonspatial aspects of the environment too (this is reviewed in Jeffery, 2007, and a model of the possible modulation mechanism proposed in Hayman and Jeffery, 2008). This modulation occurs either by switching on and off of place fields (sometimes called global remapping, if all cells do it, or partial remapping, if only some cells do it) or a change in firing rate (rate remapping, as seen in Wood et al., 2000 and also in Leutgeb et al., 2005). Equally, however, place cells seem rather unaffected by many changes occurring in the environment. Broadly speaking, it seems that place cells become interested in aspects of the environment that are persistent over long periods of time, particularly if these have some type of behavioral significance (e.g., food is present in condition X but not Y) such that changes in these factors cause remapping.

This context encoding poses problems when testing hypotheses about more complex aspects of place cell encoding such as motivation. Whenever the animal is presented with a cue that, say, signals presence of a reward, this also becomes a stable persistent cue that defines a state of the environment. Thus, it may not be the rewarding aspects of the reward so much as its environment-defining characteristics that drive the cell. To make things more complex still, the aspects of an environmental situation that define its state may even be nonsensory. For example, Markus et al. found that place cells globally remapped when a task was switched from foraging to goal directed, even though the animal was in exactly the same arena (Markus et al., 1995) and thus perceiving the same sensory cues. Similarly, as mentioned earlier, Wood et al. (2000) saw rate remapping when an animal switched from a left-bound to a right-bound trajectory, which again is something

internal to the rat rather than external in the world. Irritatingly, it is sometimes the case that cells remap in some animals and not in others—it thus seems that the hippocampus decides for itself, on a case-by-case basis, what collection of environmental characteristics (both sensory and internal) constitute state A versus states B, C, etc. It may even be that the *function* of the hippocampus is to make these categorizations (often called pattern separation and pattern completion—separation when it judges two states to be different despite many similarities, and completion when it judges them to be the same despite minor differences).

Given this role in pure context encoding, it thus becomes difficult to show that a place cell is interested in the motivational or affective aspects of a stimulus irrespective of its environment defining ones. It may even be that this is an undifferentiable distinction. Nevertheless, when considering how to test for emotional significance, these issues need to be kept in mind.

6.9 Place Cells Are Sensitive to Boundaries, and This Includes Objects

Muller and Kubie (1987), O'Keefe and Burgess (1996), and Gothard et al. (1996a, 1996b) showed some time ago that place cells form relative to boundaries and to some extent can be shifted by moving these. The conditions under which the system considers an environmental feature to be a boundary have not been worked through fully, but it seems that any reasonably large object can be a stimulus for place field anchoring. Thus, when determining whether place cells are interested in objects per se, it needs to be ruled out that they are not just using the barrier properties of the object as another spatial cue. One way to rule this out might be to change the object for a different one with the same barrier properties and see if the cell accepts the substitution without remapping. However, bear in mind that place cells are also modulated by the contextual stimuli in an environment. If the animal's hippocampus decides the environment has two states, sometimes with a barrier of type A and sometimes with a barrier of type B, then it may appear to remap to the sensory nature of the barrier (the object) when in fact it is remapping to the change in context signaled by the barrier. This is a subtle distinction: one way to rule this out would be to see if other cells, distant from the object, also remapped, suggesting that the barrier served a context-defining function. An example of this is to be found in Jeffery and Anderson (2003) in which place cells remapped their fields when features distant from the place fields were changed and also failed to remap when features near to the field were changed, depending (presumably) on whether the changes were large enough to constitute a context change.

This is not to say that place cells are uninterested in objects per se—this has not yet been ruled in or out, but a general point emerges for these considerations, which is that the information that drives a place cell may come in through functionally different pathways. Before concluding that place cells encode a particular kind of information, it needs to be ruled out that they are not using this information in a different way from the one the experimenter had in mind.

6.10 Place Cells Are Also Time Cells

As part of the move toward linking the place cell system to memory, investigators have begun looking not just at more spatially complex navigation tasks but also tasks in which other, nonspatial kinds of cognitive processing are required. In the domain of human functional neuroimaging, there is evidence that the entorhinal-hippocampal system may be involved in making "maps" of other types of cognitive relationship (Constantinescu, 2016) and the question naturally arises as to whether place cells might be involved in nonspatial encoding. We have already seen that such experiments may be confounded by the propensity for place cells to activate en masse when an animal is paused or to encode the contextual as well as spatial aspects of a task. Another potential confound that needs to be taken into account, following from recent findings from the Eichenbaum lab, is that place cells are also responsive to temporal as well as spatial intervals—that is, they may also be "time cells" (MacDonald et al., 2011; Eichenbaum, 2014). This means that in tasks having a temporal component, such as a delay, we might expect to see place cells become transiently active, as the cells track the animal's progress through the time interval. A possible example of this is a recent study in which place cells became active in response to a particular part of a smoothly changing auditory tone (Aronov et al., 2017). Because time and tone were linked in the experimental design, the activity seen may have been due to temporal encoding, with the tone being a marker for time. In looking for other types of coding, therefore, it is important to consider whether the activity being measured is independent of both space and time.

7. CONCLUSIONS

Although the spatial neurons have been known about for many decades, it is clear there is an enormous amount yet to be learned about this fascinating system, and scientists who are looking for a research question will find much that is exciting in this field. We hope that this guiding chapter will provide a leg up, as it

were, to novice electrophysiologists who plan to start recording place, head direction, or grid cells. The chapter only skims the surface, however, of a rich and complex and ever-growing methodology, and the literature, large and daunting though it is, remains by far the best resource—particularly the early literature, which is often neglected. The future, however, promises spectacular new methods that will blend the best of both electrophysiological and imaging methods so that it will eventually be possible to record huge ensembles of neurons with a high spatial, temporal, and cell-type resolution. Even with these methods, however, many of the basic principles outlined here remain intact. The most important one is this—when interpreting spatial neuron data one must always keep in mind the question—what are these neurons for? How do they help an animal live a normal life in the real world? After all, neurons were not built by engineers or theoretical neuroscientists, they were built by evolution, and it helps to keep in mind how evolution works when considering functional hypotheses.

Conflict of Interest Statement

KJ and JD are directors of Axona Ltd.

References

Ainsworth, A., Gaffan, G.D., O'Keefe, J., Sampson, R., 1969. A technique for recording units in the medulla of the awake, freely moving rat. J. Physiol. 202, 80P–82P.

Aronov, D., Nevers, R., Tank, D.W., 2017. Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. Nature 543, 719–722.

Aronov, D., Tank, D.W., 2014. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. Neuron 84, 442–456.

Barry, C., Hayman, R., Burgess, N., Jeffery, K.J., 2007. Experience-dependent rescaling of entorhinal grids. Nat. Neurosci. 10.

Boccara, C.N., Sargolini, F., Thoresen, V.H., Solstad, T., Witter, M.P., Moser, E.I., Moser, M.-B., 2010. Grid cells in pre- and parasubiculum. Nat. Neurosci. 13, 987–994.

Bures, J., Fenton, A.A., Kaminsky, Y., Rossier, J., Sacchetti, B., Zinyuk, L., 1997. Dissociation of exteroceptive and idiothetic orientation cues: effect on hippocampal place cells and place navigation. Philos. Trans. R. Soc. Lond. B Biol. Sci. 352, 1515–1524.

Burgess, N., Cacucci, F., Lever, C., O'keefe, J., 2005. Characterizing multiple independent behavioral correlates of cell firing in freely moving animals. Hippocampus 15, 149–153.

Buzsáki, G., 2015. Hippocampal sharp wave-ripple: a cognitive biomarker for episodic memory and planning. Hippocampus 25, 1073–1188.

Buzsáki, G., Leung, L.W., Vanderwolf, C.H., 1983. Cellular bases of hippocampal EEG in the behaving rat. Brain Res. 287, 139–171.

Buzsáki, G., Stark, E., Berényi, A., Khodagholy, D., Kipke, D.R., Yoon, E., Wise, K.D., 2015. Tools for probing local circuits: high-density silicon probes combined with optogenetics. Neuron 86, 92–105.

Cacucci, F., Lever, C., Wills, T.J., Burgess, N., O'Keefe, J., 2004. Theta-modulated place-by-direction cells in the hippocampal formation in the rat. J. Neurosci. 2004, 8265–8277.

Canolty, R.T., Knight, R.T., 2010. The functional role of cross-frequency coupling. Trends Cogn. Sci. 14, 506–515.

Cheung, K.C., 2007. Implantable microscale neural interfaces. Biomed. Microdevices 9, 923–938.

Constantinescu, A.O., 2016. Organizing conceptual knowledge in humans with a gridlike code. Science 352, 1464–1468.

De Lavilléon, G., Lacroix, M., Rondi-Reig, L., Benchenane, K., 2013. Explicit memory creation during sleep: a causal role of place cell on navigation. Nat. Neurosci. 18, 1–39.

Dragoi, G., Harris, K.D., Buzsaki, G., 2003. Place representation within hippocampal networks is modified by long-term potentiation. Neuron 39, 843–853.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., Csicsvari, J., 2010. The reorganization and reactivation of hippocampal maps predict spatial memory performance. Nat. Neurosci. 13, 995–1002.

Eichenbaum, H., 2014. Time cells in the hippocampus: a new dimension for mapping memories. Nat. Rev. Neurosci. 15, 732.

Ferguson, J.E., Boldt, C., Redish, A.D., 2009. Creating low-impedance tetrodes by electroplating with additives. Sens. Actuators A Phys. 156, 388–393.

Finkelstein, A., Derdikman, D., Rubin, A., Foerster, J.N., Las, L., Ulanovsky, N., 2015. Three-dimensional head-direction coding in the bat brain. Nature 517, 159–164.

Foster, D.J., 2017. Replay comes of age. Annu. Rev. Neurosci. 40.

Gothard, K.M., Skaggs, W.E., McNaughton, B.L., 1996a. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J. Neurosci. 16, 8027–8040.

Gothard, K.M., Skaggs, W.E., Moore, K.M., McNaughton, B.L., 1996b. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. J. Neurosci. 16, 823–835.

Hafting, T., Fyhn, M., Molden, S., Moser, M.B., Moser, E.I., 2005. Microstructure of a spatial map in the entorhinal cortex. Nature 436, 801–806.

Hardcastle, K., Maheswaranathan, N., Ganguli, S., Giocomo, L.M., 2017. A multiplexed, heterogeneous, and adaptive code for navigation in medial entorhinal cortex. Neuron 94, 375–387.e7.

Harris, K.D., Quiroga, R.Q., Freeman, J., Smith, S.L., 2016. Improving data quality in neuronal population recordings. Nat. Neurosci. 19, 1165–1174.

Hartley, R., Zisserman, A., 2003. Multiple View Geometry in Computer Vision. Cambridge University press.

Harvey, C.D., Collman, F., Dombeck, D.A., Tank, D.W., 2009. Intracellular dynamics of hippocampal place cells during virtual navigation. Nature 461, 941.

Hayman, R., Verriotis, M.A., Jovalekic, A., Fenton, A.A., Jeffery, K.J., 2011. Anisotropic encoding of three-dimensional space by place cells and grid cells. Nat. Neurosci. 14.

Hayman, R.M.A., Donnett, J.G., Jeffery, K.J., 2008. The fuzzy-boundary arena-A method for constraining an animal's range in spatial experiments without using walls. J. Neurosci. Methods 167, 184–190.

Hayman, R.M., Jeffery, K.J., 2008. How heterogeneous place cell responding arises from homogeneous grids – a contextual gating hypothesis. Hippocampus 18.

Hédou, G., Pryce, C., Di Iorio, L., Heidbreder, C.A., Feldon, J., 2001. An automated analysis of rat behavior in the forced swim test. Pharmacol. Biochem. Behav. 70, 65–76.

Hok, V., Lenck-Santini, P.-P., Roux, S., Save, E., Muller, R.U., Poucet, B., 2007. Goal-related activity in hippocampal place cells. J. Neurosci. 27, 472–482.

Hollup, S.A., Kjelstrup, K.G., Hoff, J., Moser, M.-B., Moser, E.I., 2001a. Impaired recognition of the goal location during spatial navigation in rats with hippocampal lesions. J. Neurosci. 21, 4505–4513.

Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., Moser, E.I., 2001b. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J. Neurosci. 21, 1635–1644.

Hölscher, C., Schnee, A., Dahmen, H., Setia, L., Mallot, H.A., 2005. Rats are able to navigate in virtual environments. J. Exp. Biol. 208, 561–569.

Jeffery, K.J., 2007. Integration of the sensory inputs to place cells: what, where, why, and how? Hippocampus 17, 775–785.

Jeffery, K.J., Anderson, M.I., 2003. Dissociation of the geometric and contextual influences on place cells. Hippocampus 13.

Jeffery, K.J., Gilbert, A., Burton, S., Strudwick, A., 2003. Preserved performance in a hippocampal dependent spatial task despite complete place cell remapping. Hippocampus 13, 175–189.

Jeffery, K.J., Jovalekic, A., Verriotis, M., Hayman, R., 2013. Navigating in a three-dimensional world. Behav. Brain Sci. 36.

Jeffery, K.J., Wilson, J.J., Casali, G., Hayman, R.M., 2015. Neural encoding of large-scale three-dimensional space-properties and constraints. Front Psychol 6, 927.

Kadir, S.N., Goodman, D.F.M., Harris, K.D., 2014. High-dimensional cluster analysis with the masked EM algorithm. Neural Comput. 26, 2379–2394.

Kentros, C.G., Agnihotri, N.T., Streater, S., Hawkins, R.D., Kandel, E.R., 2004. Increased attention to spatial context increases both place field stability and spatial memory. Neuron 42, 283–295.

Knierim, J.J., McNaughton, B.L., 2001. Hippocampal place-cell firing during movement in three-dimensional space. J. Neurophysiol. 85, 105–116.

Leutgeb, S., Leutgeb, J.K., Barnes, C.A., Moser, E.I., McNaughton, B.L., Moser, M.B., 2005. Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. Science 309, 619–623.

Lewicki, M.S., 1998. A review of methods for spike sorting: the detection and classification of neural action potentials. Network 9, R53–R78.

Liu, Y., McAfee, S.S., Heck, D.H., 2017. Hippocampal sharp-wave ripples in awake mice are entrained by respiration. Sci Rep. 7, 8950.

Loeb, G.E., Peck, R.A., Martyniuk, J., 1995. Toward the ultimate metal microelectrode. J. Neurosci. Methods 63, 175–183.

MacDonald, C.J., Lepage, K.Q., Eden, U.T., Eichenbaum, H., 2011. Hippocampal "time cells" bridge the gap in memory for discontiguous events. Neuron 71, 737–749.

Maingret, N., Girardeau, G., Todorova, R., Goutierre, M., Zugaro, M., 2016. Hippocampo-cortical coupling mediates memory consolidation during sleep. Nat. Neurosci. 19, 959–964.

Markus, E.J., Qin, Y.L., Leonard, B., Skaggs, W.E., McNaughton, B.L., Barnes, C.A., 1995. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. J. Neurosci. 15, 7079–7094.

McNaughton, B.L., Barnes, C.A., O'Keefe, J., 1983a. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp. Brain Res. 52, 41–49.

McNaughton, B.L., O'Keefe, J., Barnes, C.A., 1983b. The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. J. Neurosci. Methods 8, 391–397.

Monaco, J.D., Rao, G., Roth, E.D., Knierim, J.J., 2014. Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. Nat. Neurosci. 17, 725.

Montgomery, S.M., Betancur, M.I., Buzsáki, G., 2009. Behavior-dependent coordination of multiple theta dipoles in the hippocampus. J. Neurosci. 29, 1381–1394.

Morris, R.G.M., 2015. The mantle of the heavens: reflections on the 2014 Nobel Prize for medicine or physiology. Hippocampus 25, 682–689.

Muller, R.U., Bostock, E., Taube, J.S., Kubie, J.L., 1994. On the directional firing properties of hippocampal place cells. J. Neurosci. 14, 7235–7251.

Muller, R.U., Kubie, J.L., 1987. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. J. Neurosci. 7, 1951–1968.

Muller, R.U., Kubie, J.L., Ranck, J.B., 1987. Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. J. Neurosci. 7, 1935–1950.

O'Keefe, J., 1976. Place units in the hippocampus of the freely moving rat. Exp. Neurol. 51, 78–109.

O'Keefe, J., Burgess, N., 1996. Geometric determinants of the place fields of hippocampal neurons. Nature 381, 425–428.

O'Keefe, J., Dostrovsky, J., 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res. 34, 171–175.

O'Keefe, J., Nadel, L., 1978. The Hippocampus as a Cognitive Map. Clarendon Press, Oxford.

O'Keefe, J., Recce, M.L., 1993. Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3, 317–330.

Oades, R.D., Halliday, G.M., 1987. Ventral tegmental (A10) system: neurobiology. 1. Anatomy and connectivity. Brain Res. Rev. 12, 117–165.

Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., Harris, K.D., Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv 061481; https://doi.org/10.1101/061481.

Ranck, J.B., 1984. Head-direction cells in the deep layers of the dorsal presubiculum in freely moving rats. Soc. Neurosci. Abstr. 10, 599.

Recce, M., 1989. The tetrode: a new technique for multi-unit extracellular recording. In: Society for Neuroscience Abstracts, p. 1250.

Rossier, J., Kaminsky, Y., Schenk, F., Bures, J., 2000. The place preference task: a new tool for studying the relation between behavior and place cell activity in rats. Behav. Neurosci. 114, 273–284.

Saleem, A.B., Ayaz, A., Jeffery, K.J., Harris, K.D., Carandini, M., 2013. Integration of visual motion and locomotion in mouse visual cortex. Nat. Neurosci. 16, 1864–1869.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B.L., Witter, M.P., Moser, M.-B., Moser, E.I., 2006. Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science 312, 758–763.

Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K., Redish, A.D., 2005. Quantitative measures of cluster quality for use in extracellular recordings. Neuroscience 131, 1–11.

Sheremet, A., Burke, S.N., Maurer, A.P., 2016. Movement enhances the nonlinearity of hippocampal theta. J. Neurosci. 36, 4218–4230.

Skaggs, W.E., McNaughton, B.L., Wilson, M.A., Barnes, C.A., 1996. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6, 149–172.

Spiers, H.J., Hayman, R.M.A., Jovalekic, A., Marozzi, E., Jeffery, K.J., 2015. Place field repetition and purely local remapping in a multicompartment environment. Cereb. Cortex 25, 10–25.

Sutherland, R.J., Dyck, R.H., 1984. Place navigation by rats in a swimming pool. Can. J. Exp. Psychol. 38, 322–347.

Taube, J.S., Muller, R.U., Ranck Jr., J.B., 1990a. Head-direction cells recorded from the postsubiculum in freely moving rats. 1. Description and quantitative analysis. J. Neurosci. 10, 420–435.

Taube, J.S., Muller, R.U., Ranck Jr., J.B., 1990b. Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. J. Neurosci. 10, 436–447.

Wills, T.J., Barry, C., Cacucci, F., 2012. The abrupt development of adult-like grid cell firing in the medial entorhinal cortex. Front. Neural Circuits 6.

Wood, E.R., Dudchenko, P.A., Eichenbaum, H., 1999. The global record of memory in hippocampal neuronal activity. Nature 397, 613–616.

Wood, E.R., Dudchenko, P.A., Robitsek, R.J., Eichenbaum, H., 2000. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. Neuron 27, 623–633.

Wurtz, R.H., 2009. Recounting the impact of Hubel and Wiesel. J. Physiol. 12, 2817–2823.

Yartsev, M.M., Ulanovsky, N., 2013. Representation of three-dimensional space in the hippocampus of flying bats. Science 340, 367–372.