

## Online Chapter 16 Setting Up and Running an ERP Lab

### Overview

This chapter describes how to set up an ERP lab, including advice on selecting equipment and software for stimulus presentation, data collection, and data analysis. There are certainly many ways to set up an ERP lab, and some factors will depend on the topic area of your research. However, the suggestions in this chapter will provide a good starting point for almost anyone who is setting up an ERP lab for the first time. Also, many of the suggestions in this chapter are based on my experience, but there are often other good ways of achieving the same goals, and it is worth asking a variety of ERP researchers how they set up their labs. If you already have access to an ERP lab, you may find some good ideas for improving the lab in this chapter.

This chapter begins with some general advice about collecting data and interacting with subjects. It then describes the overall design of an ERP lab, including mundane factors such as seating and lighting. The next sections provide advice about choosing an EEG acquisition system and a data analysis system, and these are followed by a section that describes some of the specific challenges in stimulus presentation. The chapter ends with a section on finding and eliminating sources of electrical noise.

Because some of the advice in this chapter might influence what you buy for your lab, you should consider whether I have any financial conflicts of interest that might bias my advice. Box 5.1 in Chapter 5 provides a financial disclosure that can help you assess my biases (which I hope are minimal).

### Tips for Data Collection

#### *Monitoring the EEG*

Here's the single most important tip for collecting high-quality EEG data (and therefore increasing your chance of getting statistically significant effects that you can publish in high-impact journals, leading to fame and fortune): Carefully watch the EEG throughout the entire recording session so that you can detect problems with the recordings (e.g., loose electrodes, blinks, eye movements, muscle noise). If you are recording from 16 channels, there is a reasonable chance that some kind of problem will occur in a given session. If you are recording from 32 channels, the chance of a problem doubles. If you are recording from 64 channels, the chance of a problem is 4 times higher than with 16 channels. Not only is the chance of a problem higher with more channels, it's harder to detect problems because the EEG is more compressed on the display and because your attentional capacity is exceeded by so much information. So watch the data like a hawk!

Do not try to read journal articles while running subjects. Do not do email. Do not surf the web. Just watch the data (and the subject and the stimuli). It will be very boring, but it's worth it in the long run. And then some day you'll be able to tell your students (as I tell my students) that you spent thousands of boring hours watching the EEG in your first several years of ERP research, and now it's their turn.

#### *Typical Recording Parameters*

The recording parameters in your research will depend on the nature of the ERPs you are planning to examine. However, almost anyone who is looking at basic cognitive or affective ERP components between 100 and 1000 ms can use a sampling rate of between 200 and 1024 Hz, along with a low-pass antialiasing filter with a half-amplitude cutoff somewhere between 1/3 and 1/5 of the sampling rate. If you sample higher than approximately 250 Hz, you will probably want to make a downsampled copy of your data at approximately 250 Hz (with an appropriate antialiasing filter) for use in the subsequent data processing steps. This will minimize disk space and speed up data processing. Why bother sampling at a high rate and then immediately downsampling? Doing the initial recording at a fairly high sampling rate (e.g., 500 or even 1000 Hz) will make you sleep better at night, because you will know that you can always reanalyze your data with the greater temporal resolution if the need arises.

As described in the chapter on filtering, you should not ordinarily filter out low frequencies during data acquisition if your data system has 20+ bits of voltage resolution. Instead, you should filter offline instead

(usually with a half-amplitude cutoff of 0.1 Hz). However, if your system has <20 bits of resolution, you should use a high-pass filter with a half-amplitude cutoff between approximately 0.01 and 0.1 Hz during data acquisition to avoid saturating the system.

If the EEG you are recording has a lot of high-frequency noise, you may want to use a lower low-pass cutoff (e.g., 30 Hz) or a notch filter during data acquisition so that you can more easily monitor overall data quality. This mainly arises in settings where you have a great deal of line noise that cannot be avoided, and this noise is so large that it obscures everything else.

Some EEG recording systems (e.g., BioSemi) have a wonderful feature that allows you to apply filters to the EEG being displayed without actually filtering the data being stored on disk. That way you can easily monitor data quality during the recording but can apply more flexible filters to your data offline. This is truly a great feature, but it has a significant downside: It's not obvious that the system is saving the unfiltered data, and many researchers think that the filtered data they are seeing on the screen are also being saved as filtered data in their data files. So keep this in mind: **With the BioSemi ActiveTwo system, the data saved to disk are filtered only with an antialiasing filter, not with any other low-pass or high-pass filters.** This may also be true of other systems, so check your hardware and software manual to see what is being saved to disk.

You can use any single site you wish as the reference electrode during data acquisition, because you can always re-reference later. For example, EGI systems typically use Cz as the reference electrode, which is not a very good final reference site for most experiments, but this is not a problem in practice because the data can be easily re-referenced offline. However, keep in mind the following two facts: (a) any artifacts that are present in all channels are probably coming from the reference electrode, and (b) the signals from sites near the reference electrode will probably look very small.

The BioSemi system allows you to view the data during recording with any reference you like, **but this is not how the data are saved to disk.** This is a really nice feature as long as you realize that the reference chosen for display is not the same thing as the reference used for saving the data.

### *Keeping Subjects Happy*

ERP experiments tend to be long and boring, with trial after trial of the same basic task. To ensure that you are collecting the highest quality data possible, it is important to keep your subjects happy and relaxed. If they are unmotivated or become bored, they may not pay close attention to their performance, weakening your effects. Moreover, bored subjects tend to become tense and uncomfortable, leading to muscle noise and movement artifacts. By keeping your subject happy and relaxed, you will get larger and more consistent effects, which will save you lots of time in the long run. Here are four suggestions for accomplishing this.

First, talk to the subject. As you are applying the electrodes, chat with the subject, asking about school, work, hobbies, family, sports, etc. (but stay away from potentially inflammatory topics such as politics). Make sure that the conversation revolves around the subject rather than yourself. By doing this, you will ingratiate yourself to the subject, and the subject will be more likely to try to please you during the experiment. In addition, when there are breaks in the experiment, you should continue to converse with the subject to keep the subject alert and interested. Some subjects just want to finish the experiment as rapidly as possible, and you should respect this. But if the subject is at all interested in talking, you should do so. This is also a good time to provide some positive feedback to the subject about behavioral performance, blinks, and eye movements, as well as encouraging the subject to improve if necessary. If a subject isn't doing well (e.g., in terms of behavioral accuracy or artifacts), don't be shy about telling them that, because poor performance and artifact-laden data won't be of much use to you. You should also make sure that the subject understands exactly what is going to happen and when (especially during the electrode application process). People are much more relaxed when they know what to expect.

Second, make sure that the blocks of trials are a reasonable length. If the blocks are too long, the subject's attention is likely to wane toward the end. I find that trial blocks of 5-7 minutes, separated by 1-2 minutes of rest, are optimal in most experiments. In addition, you will usually ask the subject to suppress blinks and eye

movements during the recordings, and it is difficult to do this for a long period of time. Some experiments consist of brief trials separated by an intertrial interval of 1-3 seconds, and this provides a time for subjects to blink (although the offset of the blink during the intertrial may contaminate the early portion of the ERP on the next trial). If this is not possible, I would recommend providing short breaks of 15-20 seconds after every 1-2 minutes.

Third, due to the long duration of a typical ERP experiment, it is helpful to provide snacks and drinks, usually 30-50% of the way through the session. Drinks including caffeine can be particularly useful in helping subjects to maintain alertness (although this may be contraindicated in some experiments). This helps to keep the subjects awake, alert, and friendly. Sometimes people ask me if this is a confound. It's not (according to the technical definition of a confound) as long as it does not differ systematically across groups or conditions. It may, however, impact the results, and other labs might not get the same results if they don't also provide caffeine, hydration, snacks, etc. But the main difference between your results and those of other labs is that your p values are likely to be better than their p values, because you will have cleaner data by making sure that the subjects are happy and awake. In some types of research, you would want to indicate that you are providing food and beverages in your methods sections. However, people don't ordinarily indicate in their methods sections that their subjects were tired, hungry, and unhappy, so ordinarily it shouldn't be necessary for you to indicate in your methods sections that you made sure that your subjects were not tired, hungry, or unhappy.

My fourth recommendation is something that I started doing when I first set up my own ERP lab, and I don't know if anyone else does it. Specifically, we play background music in the chamber while the subject is doing the task. In fact, we suggest to the subjects that they bring CDs of their favorite music with them (and the musical genres have included classical, pop, rock, metal, rap, country, electronic/ambient, and just about everything else imaginable). Of course, the music produces some distraction from the task, and the sounds will generate ERP activity. However, I believe that the music is less distracting than the alternative, which is mind-numbing boredom. And any ERP activity generated by the music will be unrelated to the stimuli and will just add a small amount of noise to the EEG. I suspect that any additional EEG noise created by the music is more than balanced by a reduction in other forms of EEG noise, such as alpha waves. Of course, there are situations in which background music is a bad idea. For example, it would be problematic in many studies using auditory stimuli, and it might cause too much distraction in some types of subjects. But if these circumstances don't apply to you, I would definitely recommend playing background music. I've never done a direct comparison of data collected with versus without background music, but I strongly suspect that background music leads to a moderate increase in the overall quality of the data. If you play music, make sure that you use shielded speakers and that the cables leading from the stereo to the speakers are shielded. Speaker cables are not usually shielded, but you can easily encase them in some form of shielding, such as Wiremold.

## **Overall Lab Design**

### *Basic Setup*

A slightly simplified diagram of my own lab setup is shown in Figure 16.1. It has several components:

- A recording chamber or recording room/area. As will be discussed later in this chapter, a shielded chamber is not always necessary, especially if your system has active electrodes to minimize line noise. My lab uses chambers for our recordings, but we have a sophisticated temperature control system to ensure a relatively constant and cool temperature. However, if at all possible the subject and experimenter should be in separate rooms so that the subject is not distracted by the experimenter.
- A stimulus presentation computer. For visual stimuli, you want to be able to see exactly what the subject is seeing, so I recommend taking the main video output and running it through a video splitter to two monitors, one in the subject area and one in the experimenter area. A good video splitter will also amplify the video signal so that it will not degrade if the cable between the splitter and the monitor is long. In my lab, the experimenter also has a separate monitor that can be used to display additional information, such as trial-by-trial subject responses. This is very useful if your stimulus presentation

supports it. To minimize electrical noise, the subject's video display may be placed inside a Faraday cage, as described in a later section. Auditory stimuli will also be discussed later.

- An EEG data acquisition system. Traditionally, this system would have a *head box* (where the electrodes plug in to the system), an amplifier (which also contains filters), and an analog-to-digital converter connected to a computer. Many systems now combine the head box, amplifier, and analog-to-digital converter in a single data acquisition module, which then connects via USB or a fiber optic cable to a computer that displays and stores the data.
- A communication system that allows you to talk to the subject. There are many possible ways of setting this up. In Figure 16.1, the experimenter talks into a microphone, which is connected to a mixer and audio amplifier that are in turn connected to speakers inside the subject area. A microphone is then mounted on a table or on the ceiling in the subject area, and this signal is sent to a powered speaker near the experimenter. With this arrangement, auditory stimuli can also be sent from the stimulus presentation computer to the mixer so that the subject can hear the stimuli and can also hear the experimenter's microphone (which can be turned off during stimulus presentation). It is also extremely useful to have a video camera in the subject area so that you can verify that the subject is sitting still, paying attention, not chewing gum, etc. It would be possible to use an inexpensive webcam instead of a microphone and video camera in the subject area.
- A response device. Recommendations for this are provided later in the chapter.

The stimulus presentation system needs a means of communicating with the EEG data acquisition system in real time to send *event codes* whenever an event of some sort occurs (e.g., a stimulus or a response). These event codes are used as the time locking points for averaging, so the timing must be precise. A consistent delay is not a big problem (as long as you know what it is), because the averages can be shifted in time to compensate for the delay. A variable delay, however, is usually difficult to compensate for, and it will have the effect of smearing out the ERP waveforms in time, distorting the onset and offset times of the components and experimental effects (just like a low-pass filter). Unfortunately, the manufacturers of EEG systems and stimulus presentation systems continue to rely on a 1980s technology for event codes, the *parallel port*. I'm hoping that they will soon switch over to something more robust and flexible, like Ethernet.

In most systems, the event codes need to last long enough so that they are present for at least one full sample period. If, for example, you are digitizing at 250 Hz, your sample periods are 4 ms, and you must make sure that your event codes always last at least 4 ms. Otherwise, some event codes will be missed by the digitization system. In addition, some systems require that the parallel port be set to zero for at least one sample period between event codes.

### *Seating*

In experiments using visual stimuli, the subject is seated at a specific distance from the video monitor. Even with a Faraday cage, some electrical noise may be picked up from the monitor, so the subject should not be too close to the monitor. I would suggest a minimum distance of 70 cm, and 100-200 cm is even better if the chamber is large enough.

It is very important for the subject to be seated in a comfortable position. In the early days of ERP recordings, very few electrodes were used, and auditory stimuli were more common than visual stimuli. A large, padded recliner was the most sensible chair for such experiments. This kept the subject comfortable, and minimized any need for subjects to use their neck muscles to support their heads, which in turn minimized EMG noise (see Chapter 5). However, when electrodes are placed over the back of the head, a recliner doesn't work as well: the head puts pressure on the electrodes, and any small head movement will cause the electrodes to move on the scalp, producing a large and sudden voltage shift.

If your experiments will require recording from electrodes on the back of the head, I would not recommend using a recliner. Instead, I would recommend a high-quality office chair that provides good lumbar support and

has an easy mechanism for height adjustment. The only problem with these chairs is that subjects may swivel from side to side. You can just drill a hole through the base and insert a screw to prevent this (which may prevent height adjustments), or you can just watch the subjects carefully to make sure they don't do this. It is also preferable to use a chair that is on *glides* rather than casters so that subjects do not move the chair around inside the chamber. The position of the chair on the floor should be marked so that the chair can be correctly repositioned if it gets moved. I would also recommend against using any kind of chin rest or other head stabilization. Although you might think that this would reduce strain on the neck muscles, my experience has been that chin rests become uncomfortable after 15 minutes or so (and most ERP experiments require at least 1 hour of recording time).




Figure 16.1. Basic setup of a typical ERP lab.

### Behavioral Responses

A small set of behavioral response alternatives is used in most cognitive ERP experiments, and only 2-4 buttons are usually necessary. The most convenient device for responses is usually a small, lightweight, hand-held device of some sort. My lab uses computer game controllers. It is easy to find game controllers that are very comfortable and lightweight, and they usually offer at least 4 response buttons (which is enough for most experiments). Various custom response devices are also possible, but they are not usually as easy to hold and

use as a game controller. A computer keyboard is not usually very good, because it is too big to rest on the subject's lap, and if it's on the table in front of the subjects, EMG noise from the neck and shoulders is likely to occur as the subjects' arms are held up to the keyboard. If a keyboard is necessary because of a large number of response alternatives, the data around the time of the response may be contaminated by EMG noise, and eye movement artifacts may be seen if the subject needs to look at the keyboard to find the right key to press. However, it may be possible to design the experiment so that the data around the time of the response are not important. Also, it may be possible to use a keyboard tray to minimize the need to hold the hands upward and outward to reach the keyboard.

Devices such as keyboards and game pads are not designed for real-time electrophysiological recording, and a delay may occur between the time of the buttonpress and the time that the buttonpress is recorded by the computer. Response times are usually so variable that this small error will not matter. However, it may be a significant problem if you are doing response-locked averaging. Special-purpose *button boxes* with more precise timing can be obtained for some stimulus presentation systems. For really precise response timing, however, it's usually best to record the EMG from the muscles that make the response; the EMG onset can then be used as the time-locking point for response-locked averages.

### *Lighting*

AC-powered lights near the subject can be a major source of electrical noise, and you should make sure that all lights near the subject are DC-powered. There are several types of DC lights. The simplest to install are battery-powered LED lights, which are used in homes to provide lighting for closets, cabinets, etc. They are very easy to install, because they require no wiring, but they provide uneven lighting and require you to change the batteries regularly. A second alternative is to use the kind of DC lights that are found in the interior of recreational vehicles, which can be powered by a 12-volt DC power supply (which must be located far from the subject if you don't have a chamber and outside the chamber if you do). These work fairly well, but they tend to heat up the recording environment. A warm environment tends to lead to skin potentials, and this may be especially problematic if you are using a small chamber (because a small chamber will heat up quickly). In my experience, the best solution is to use LED light strips powered by a DC power supply. A circuit diagram and parts list for the LED lighting used in my lab can be found at <http://erpinfo.org/educational-materials/led-chamber-lighting>. This system is fairly old now, so you may be able to create something cheaper and better with newer technologies.

Rapid progress is being made in using LED for indoor lighting, and I suspect that excellent new solutions will become available over the next few years. However, LED lights that plug into an AC lighting fixture will probably not solve the problem of line noise because the AC-DC converter in the lighting fixture will generate a lot of electrical noise. To eliminate line noise, a DC power supply (or batteries) will be needed (and must be far from the subject).

### *Electrode-Amplifier-Computer Connections*

In a conventional passive recording system, the EEG is picked up by the electrodes and travels through cables to the amplifier system. Before they are amplified, the signals are miniscule, and any electrical noise that is picked up by the cables will be relatively large compared to the EEG. Consequently, the cables between the electrodes and the amplifier system should be kept as short as is practical in these systems. This is not an issue for systems with active electrodes.

Many EEG amplifier systems have a *headbox* that provides a flexible means of connecting the electrodes to the amplifier. Flexibility is important. For example, you want to be able to select a common reference electrode for most of the channels (e.g., a mastoid or earlobe reference) while allowing other channels to use a different reference electrode (e.g., eye movements are usually recorded as the voltage between electrodes on either side of the eyes). Because the output of the headbox is not amplified, the cable from the headbox to the amplifier should be relatively short. Optimally this cable should be shielded, especially if the actual amplifier is outside the recording chamber. If the amplifier is inside the chamber, an unshielded ribbon cable will probably

be sufficient (a ribbon cable keeps the active, reference, and ground signals close to each other so that they pick up the same noise signals, and this allows the noise to be subtracted away by the differential amplifier).

If the amplifier is connected to AC power, it is usually best for it to be placed outside the recording chamber to avoid inducing noise in the electrodes. If it is battery powered, it can probably be placed inside the chamber, which is good because only amplified signals will leave the chamber, where electrical noise is much greater. Once the signals have been amplified, they are usually much larger than any induced electrical noise, so longer cables may be used to connect the amplifier to the digitization computer. However, some noise may still be picked up at this stage, and some shielding may be necessary even after the signals have been amplified.

Some newer systems have completely eliminated most of these issues (e.g., BioSemi ActiveTwo, Brain Products actiCHamp). In these systems, the electrode cables plug into a single box that amplifies and digitizes the signals, and once the signals have been digitized they will not be degraded as they pass through a cable to the computer that displays and saves the data.

### *Recording Chamber*

One of the most expensive elements of an ERP lab is the electrically isolated chamber. These chambers are usually designed as sound-attenuating chambers for hearing testing, broadcasting, or auditory research, but some manufacturers offer electrical shielding options. I previously used chambers made by Industrial Acoustics Company and Acoustic Systems, and they worked quite well. My current chambers are made by Eckel Industries, and they also work quite well.

Is an electrically isolated chamber really necessary? Not always. If you have active electrodes, you probably won't need a chamber unless your lab is very close to motorized devices (e.g., centrifuges, elevator motors) or transformers. Even if you have passive electrodes, you might be able to get away with using a low-pass filter with a half-amplitude cutoff around 30 Hz to get rid of line-frequency noise.

Chambers also have a down side that is not obvious. Specifically, chambers tend to be small and poorly ventilated, and they tend to heat up as the subject sits inside. As described in Chapter 5, this can lead to skin potentials, which can dramatically reduce your statistical power (especially if you are looking at the slower, later components such as P3, N400, LPP, etc.). Skin potentials are especially problematic when the electrode impedances are high, even if you have a system that is designed for high-impedance recordings (for details, see the online supplement to Chapter 5 and Kappenman & Luck, 2010). A standard laboratory room may be much easier to maintain at a cool and constant temperature, and this may end up giving you much higher statistical power, even if the amount of line noise is a little higher. Thus, chambers have both advantages and disadvantages. So what should you do?

You are probably better off getting a chamber if you are very interested in relatively high-frequency activity that is easily distorted by line noise (e.g., gamma oscillations) and you don't care about slow, late components. You might also want to get a chamber if you will be doing a great deal of ERP research for many years, and you are willing to invest in a special-purpose cooling system for your chamber (as my lab has done). However, if you have active electrodes and you are mainly interested in slower components, you may not want to get a chamber (unless you are willing to invest in a special-purpose cooling system).

If you do use a chamber, don't assume that it effectively eliminates all sources of electrical noise. Chambers often include fans and lights that may cause substantial electrical noise. In fact, the first chamber I ever used induced so much noise that I was able to get cleaner recordings outside the chamber than inside. I eventually found the circuit breaker for the chamber, and turning it off eliminated the noise. In addition, a chamber doesn't help much if you put devices such as CRT monitors into the chamber that create a lot of noise. A later section of this chapter will describe how to minimize noise from CRT monitors and other devices.

### *Attenuating Acoustic Noise*

Most ERP recording chambers are made by companies that mainly focus on acoustic isolation, and a side benefit of a recording chamber is that it reduces noise (e.g., from doors opening, people talking, chairs

squeaking). However, it turns out to be quite difficult and expensive to completely block out noise from outside the chamber (and even more difficult if you are using a regular testing room). A much simpler and less expensive approach is to cover up these outside noises with *masking noise*. For example, you can buy *white noise* generators that create a constant hissing sound. This sound gets quite obnoxious after a few minutes, however, and it's better to get a gradually varying source of noise that emphasizes lower frequencies (which are especially good at covering up the sounds of people talking). For example, you can purchase a device that is intended to help people sleep by producing calming ocean-like sounds. My lab purchased a recording of "sounds of the ocean" that we play through speakers or headphones during our experiments (including our behavioral experiments). This is much more pleasant than white noise and is reasonably effective.

## Choosing an EEG Acquisition System

People often ask for my advice about what EEG acquisition system to buy. I've used only a few commercial systems in my own lab (Neuroscan, BioSemi ActiveTwo, and most recently Brain Products actiCHamp), and I've never done a formal comparison of systems. Consequently, it would be inappropriate for me to recommend specific manufacturers. However, I can provide some general advice about how to go about selecting data acquisition equipment.

First, however, I'd like to provide a caveat. The suggestions here are my own opinions, and they may be different from the opinions of other researchers. Consequently, I would recommend seeking advice from a variety of experienced ERP researchers.

### *How Many Channels?*

Most new ERP researchers assume that they should buy a system with as many channels as possible (e.g., 64 or even 128 channels), but this is typically a mistake. My lab has two 64-channel systems, but we rarely record from more than 40 channels, and we have published dozens of ERP papers in top journals with fewer than 30 channels. The skull smears ERP scalp distributions so much that adding more electrodes beyond 30-40 provides only a small amount of additional information, and it comes at a high cost in terms of reduced data quality and statistical power. I encourage you to read the online supplement to Chapter 5 (available at <http://mitpress.mit.edu/luck2e>), which provides the rationale for this assertion.

For most laboratories and most experiments, 32 channels is sufficient. Some labs can get by with only 16 channels, but very few labs need more than 64 channels. In addition, many systems are expandable and make it easy to add more channels at a later time. In most cases, I therefore recommend starting out with 32 channels. If you already have a lot of ERP experience and plan to make ERPs a central part of your research for the next decade, it would be worth getting 64 channels. If you do get a 64-channel system, however, make sure that you can easily do recordings with only 30-40 of the channels, because using this smaller subset of channels will likely lead to better data, better p values, and more and better publications. Box 16.1 describes my own personal reasons for getting a 64-channel system but recording from only a subset of the electrodes.

### *Electrode Types*

Once you've decided on how many channels you want, the next major decision is whether to go with (a) a passive low-impedance system, (b) a passive high-impedance system, or (c) an active high-impedance system. Chapter 5 describes the different types of systems in detail, and here I will provide a brief reminder and then distill the abstract issues into some specific advice.

The main advantage of high-impedance recordings is that the electrode preparation process will be faster for you and more comfortable for the subjects. It also reduces the (already very low) likelihood of disease transmission (which has the benefit of making it easier to get human subjects research approval). The main disadvantage of high-impedance recordings is an increase in skin potentials, which can dramatically decrease statistical power for longer-latency, slower ERP components (anything after approximately 200 ms, with a progressive loss of power as the latency of the component increases).



The main advantage of active electrodes is that you will pick up much less electrical noise from the recording environment (especially line noise). The main disadvantage is that the electrodes are more expensive and, for some systems, less convenient (e.g., because it is more difficult to implement online impedance checking). By using active electrodes, you may be able to avoid using a shielded chamber, which in turn may make it easier to keep the recording environment cool, and this reduces skin potentials and therefore increases your statistical power. Some active electrode systems also use a driven right leg circuit (see Chapter 5), which further improves the system's ability to minimize environmental electrical noise.

A third factor that differs across systems is the nature of the conductive medium that actually connects the electrodes to the scalp. Most systems use some kind of electrolyte gel, whereas others use saline. The advantage of a gel is that it provides a stable connection that actually tends to improve over the course of a session (as the gel hydrates the skin), minimizing artifacts that arise from small head movements and reducing skin potentials. The disadvantage of a gel is that it takes a little longer to apply at the beginning and clean up at the end. The advantage of saline is that it is fast, and the disadvantage is that it is less stable and therefore tends to yield noisier data.

These principles lead to the following recommendations

- If you are interested only in EEG oscillations and/or early ERP components (prior to 200 ms), skin potentials will have little or no impact on your data, and a high-impedance system is fine.
- If you are interested in longer-latency components (especially components that peak 400 ms or later), you should use low-impedance recordings unless you can maintain a consistently cool and dry recording environment. Note, however, that some high-impedance systems make it easy for you to reduce the electrode impedances and perform low-impedance recordings. A good compromise might therefore be to use a high-impedance system with active electrodes and make sure that the impedances are not terribly high (e.g., less than 20-30 K $\Omega$  if the recording environment is warm, and less than 50-60 K $\Omega$  if the recording environment is cool).
- If you are recording from infants or small children, speed of electrode application is extremely important. Unless you are recording from only a few electrodes, you will probably want a high-impedance system in this situation. In the vast majority of such cases, however, I would recommend a gel-based system rather than a saline-based system. It will take a little more time to apply the electrodes, but the improvement in data quality will be well worth it.

#### *Box 16.1- Why I Bought a 64-Channel System For My First Lab*

When I set up my first ERP lab as an assistant professor, I spent a large proportion of my meager startup funds to purchase a 64-channel EEG acquisition system. However, I bought a set of 32-channel caps along with 64-channel caps, because I was aware of the reduction in data quality and statistical power that arises when you record from large numbers of channels. Almost all of the experiments we conducted used the 32-channel caps, and 32 of those 64 channels remained idle about 99% of the time. I also paid a ridiculous amount of money to buy a source localization software package, which was also rarely used.

You might wonder why I spent so much of my precious start-up funds on a 64-channel system and source localization software. The answer is simple: I wanted to be able to write grant proposals saying that I had the capacity to do high-density recordings and source localization. At this point (the mid 1990s), a lot of people were very excited about source localization, and I was concerned that my grants would be criticized for being insufficiently innovative if I didn't have the capacity for high-density recordings and source localization. It's impossible to know if this strategy was necessary, but I was successful at getting grants, and I continue to use this strategy in the grants I write today.

My grad students and postdocs have run a lot of studies using 30-40 of our 64 channels over the past 20 years, and we've published a lot of papers in top journals. We occasionally do 64-channel experiments (mainly when we've found a new effect and want to carefully characterize its scalp distribution "for the record"). I've never regretted this strategy, and I encourage others to follow it when appropriate.

## Amplification and Digitization

In the old days, a typical EEG acquisition system had an amplifier unit (that also included analog filters) and a separate analog-to-digital converter unit. Most newer systems integrate these two functions together, and in many cases the only analog filter is the antialiasing filter (which is automatically set on the basis of the sampling rate you specify). Digital filters are superior to analog filters, so this is usually a good thing. The amplifiers used in almost all systems are of very high quality, and the main parameter that you check is the common mode rejection (defined in Chapter 5), which should be at least 100 dB.

The main parameters of the analog-to-digital converter are the voltage resolution (usually between 12 and 32 bits) and the maximal sampling rate (e.g., 5,000 Hz per channel). The sampling rate will depend on the kind of research you do. Except for research focused on early sensory activity, a sampling rate of 250 Hz is usually adequate, and there is rarely any value in recording faster than 1000 Hz. I recommend at least 20 bits of voltage resolution (unless you are on a very tight budget). This will mean that you can record without first filtering the low frequencies using a hardware filter, and instead you can filter offline with a digital filter. If a system has less than 20 bits of resolution, then you should make sure it provides good control over the high-pass filter (with half-amplitude cutoff options at approximately 0.01 Hz and approximately 0.1 Hz). See Chapter 7 for additional advice on filter settings.

It is very useful if the system allows you to flexibly filter and re-reference the data for the online EEG display during recording. As I described earlier, it's especially nice if the system can save the original (unfiltered, unreferenced) data while at the same time displaying the filtered and re-referenced data. You can always filter and re-reference the data offline, but you can't unfilter data that has already been filtered, and some types of re-referencing can be difficult to undo. It's also important to be able to control the scale of the EEG display, and it's helpful to be able to control the polarity. A convenient system for checking electrode impedances is also very useful (even in a high-impedance system).

It is also useful if the system contains auxiliary inputs that allow you to record other types of signals, such as your computer's audio output. This is especially important for testing the timing of your system (as will be described later in this chapter). These inputs should have a range of at least  $\pm 200$  mV and up to  $\pm 5$  V.

In almost all cases, you will want a system that provides continuous EEG recording. Many ERP paradigms involve stimuli that are separated by long periods of time with no ERP-eliciting events, and it is therefore possible to record just the EEG segments surrounding the stimuli, pausing during intertrial intervals. This is called *epoch-based* recording. However, you may later wish that you had recorded a longer epoch. For example, you may find an effect that is still growing in amplitude at the end of the epoch, or you may want to evaluate prestimulus overlap more thoroughly. Epoch-based recording does not allow this. The alternative (*continuous* recording) is to record the EEG continuously and extract whatever epochs are necessary during the averaging process. All of the data are saved, so it's possible to re-average the data with a different epoch at a later time. This is almost always better than epoch-based recording.

Finally, it is very useful if the system can provide a spectral display (i.e., Fourier-transformed data, showing amplitude or power as a function of frequency). This makes it much easier to quantify the amount of noise coming from different frequency bands (e.g., the 50 or 60 Hz line noise frequency), which can help you track down and eliminate sources of noise.

## The Data Analysis System

### *ERPLAB Toolbox*

I've taken a look at several of the commercial ERP analysis systems that are available, and I haven't been favorably impressed by most of them. They are full of bells and whistles, but they are expensive, lacking in important features, difficult to customize, and inconvenient to use in large-scale research. Some actually encourage the use of methods that are known to be flawed, and they do not provide convenient ways to implement many of the "best practices" in ERP research. Most importantly, commercial packages necessarily focus on existing methods, whereas scientific progress requires the invention of new approaches.

For many years, my lab used a custom software package that was developed by Jon Hansen (with help from several other programmers) in Steve Hillyard's lab at UCSD. This package had many great features, but it was difficult to port to new operating systems and difficult to modify. Eventually, this package became too difficult to maintain, and it lacked some newer analysis methods that were becoming common in the field. I therefore decided that it was time for my lab to create a new package using a software development environment that was platform-independent and easy to learn, namely Matlab. I then hired a very talented engineer, Javier Lopez-Calderon, to do the actual programming, and I obtained a grant from NIH to support the project. The result is ERPLAB Toolbox, an open source, easily extensible ERP data analysis package (see Lopez-Calderon & Luck, 2014 for details). It can be downloaded for free at <http://erpinfo.org/erplab>. ERPLAB works in concert with EEGLAB Toolbox, which was created earlier by Scott Makeig and Arno Delorme at UCSD (Delorme & Makeig, 2004, <http://sccn.ucsd.edu/eeglab/>).

The EEGLAB/ERPLAB combination is both easy to use and powerful. The ease of use comes from a graphical user interface (GUI), and the power comes from the ability of end users to write their own Matlab scripts. Perhaps the most important feature, which ERPLAB inherited from EEGLAB, is that every action the user takes with the GUI has a script equivalent, which is saved in a *history* structure. To create a script for automating a sequence of operations, you simply carry out those operations in the GUI and then copy the history into a script file. This makes it easy for people with no programming experience to learn how to write simple scripts for automating processing (but see Box 16.2 for a discussion of the downside of automated scripts). After gaining experience with writing simple scripts, it's a relatively simple step for a researcher to write somewhat more complex scripts that implement brand new processing routines that move the field in new directions.

Thousands of researchers have downloaded EEGLAB and ERPLAB, and hundreds of papers have been published with data analyzed with these toolboxes. I would recommend the EEGLAB/ERPLAB combination for almost any ERP researcher.

#### *Box 16.2- Scripts are Wonderful and Terrible*

Scripts that automate EEG/ERP data processing can be a truly wonderful thing. Once you've figured out how to analyze the data from a couple subjects, you can avoid a lot of pointing and clicking by writing a script that automatically processes the data from the remaining subjects. Scripts also allow you to see how other people have analyzed their data and then apply the same procedures to your own data. And after you've analyzed your data, you can easily change one parameter (e.g., the prestimulus baseline period) and then your computer can reanalyze all of your data while you eat lunch.

However, scripts can also be terrible. Many times I have seen a grad student blindly running a script without actually knowing exactly what it does or how it works. And many times I have seen a grad student process a subject's data using a script without ever actually looking at the data. It is absolutely essential that you look closely at each subject's EEG, each subject's artifacts, each subject's behavior, etc. One size does not fit all with artifact rejection or correction (see Chapter 6), and there might be other idiosyncratic problems that can be discovered by looking at the data. Data are precious and require your careful attention!

Here are some pieces of advice about using scripts to automate data processing. First, never use someone else's script unless you fully understand every single parameter in every line of the script. Second, never run a script without first looking carefully at the EEG from the entire session. At a minimum, you will need to check the artifact correction and rejection parameters by looking at the EEG. Third, don't start using a script to process the data from a given experiment until you've manually processed at least 5 subjects. Once you've processed the first 5+ subjects manually, you'll have a better idea about how the parameters should be set for the remaining subjects. Finally, after each subject has been processed with a script, check the number of trials rejected and the number of trials accepted in each condition to make sure that you actually have the right number of trials in each condition.

The best time to use a script is after you've already processed each subject once manually. At that point, you will know what parameter values are appropriate for each subject, and you can then use a script to reanalyze the data with slight parameter changes (e.g., different filter settings).

### *Commercial ERP Analysis Packages*

However, if you need to buy a commercial package for some reason, here are some important features that you should consider when you are deciding what to purchase:

1. Easy methods for averaging the data on the basis of sophisticated criteria, such as sequences of stimuli, correct versus incorrect responses, and responses with various different reaction times. Response-locked averages must be possible.
2. A broad set of artifact rejection functions, such as those described in Chapter 6. You should be able to set the parameters on the basis of visual inspection of the EEG. It is also useful to be able to do artifact rejection manually on the basis of visual inspection.
3. A modern approach to artifact correction, such as ICA. The regression-based approach developed by Gratton et al. (1983) was a great step forward in the 1980s, but better methods are available now.
4. The ability to filter the raw EEG and the averaged ERP waveforms with a variety of half-amplitude cutoffs and rolloff slopes. The impulse response functions of the filters must be clearly described.
5. A broad set of component measurement routines should be available. For example, the system should be able to do both true peak and local peak measurements (see Chapter 9).
6. It must be easy to perform mathematical operations on the ERP waveforms, including re-referencing the data and forming difference waves.
7. The ability to import and export data in different formats, particularly text files, is very useful.
8. The ability to automate processing can be very useful. For example, you may find that you need to re-average the data from each subject in an experiment with a longer time epoch, and you will save a lot of time and effort if this process can be automated.
9. A convenient but flexible means of plotting the ERP waveforms is essential. You must be able to specify the line types and line colors used for overlapping waveforms, the placement of these sets of overlapping waveforms, and the formatting of the time and voltage axes. In addition, the system must be able to save the plots in a vector-based file format that can be imported by standard vector-based graphics packages (bitmapped outputs are not usually suitable for creating figures of ERP waveforms for publication).

When you are evaluating a commercial system, think carefully about how well this system will work in a complex experiment with dozens of subjects. In my experience, a package can work very well in a demonstration with data from a single subject in a simple experiment but turn out to be extremely inconvenient in the context of real experiments. For example, some systems create a separate data file for each condition of the experiment. Well-designed experiments often have 20+ conditions per subject because of counterbalancing factors, and a sequence of 8 data processing steps would create 160+ files for each subject!

I evaluated several ERP analysis packages for use at the ERP Boot Camp. (ERPLAB Toolbox was not fully functional for the first several years of the Boot Camp, and our previous custom software did not have a graphical user interface and was too difficult to learn.) Of all the commercial packages available at the time, BrainVision Analyzer did the best job of allowing us to teach the participants what we considered the best practices in ERP data analysis, and we used it for several years. However, I've talked to other researchers who are very fond of other packages, so you will need to determine for yourself which package best meets your needs.


## The Stimulus Presentation System

### *Timing of Event Codes*

In general, the key to a good ERP stimulus presentation is precise timing of the event codes with respect to the stimuli. It turns out to be surprisingly difficult to ensure that an event code occurs at the precise moment that the stimulus appears. There are two reasons for this. First, modern computer operating systems are not designed to operate in real time, and programs are often interrupted briefly so that the operating system can conduct various housekeeping tasks. These interrupts are not usually noticeable during normal computer usage, but they can introduce timing errors that are large relative to the time scale of ERP recordings (as much as 200 ms). Real-time programming and precise timing was actually easier to achieve with more primitive systems, such as the early PCs and the Apple II. A second issue is that several pieces of hardware are typically interposed between the computer's CPU and the device that actually produces the physical stimulus, so the CPU doesn't actually "know" when the stimulus is actually presented. This is especially true for visual stimuli, as will be described in detail a little later. Box 16.3 describes an example of a timing error that I discovered in my own lab. If it can happen in my lab, it can happen in yours!

### *Box 16.3- An Example of a Timing Error*

A few years ago, I was on the road and preparing to give a Mini ERP Boot Camp, and I wanted to show an example of the C1 wave, which is the earliest ERP response to a visual stimulus (see Chapter 3). My lab had been pilot-testing a C1 paradigm, and I asked the graduate student who was running the experiment to email me a grand average of the 6 pilot subjects. The waveform is shown below. I immediately emailed back to the student saying that we had a problem. Take a look at the waveform and see if you can figure out what the problem is.



The problem is that the waveform starts to deviate from baseline right around time zero, which is the moment at which the event code was detected by our EEG acquisition system and is supposed to be the moment of stimulus onset. Although ERP responses from the brainstem can be observed for auditory stimuli can be observed within a few milliseconds of stimulus onset (for sudden-onset sounds such as clicks), large ERPs cannot be observed for visual stimuli until the information arrives in primary visual cortex. And it typically takes 40-60 ms for visual information to be accumulated by the retina and project through the LGN to primary visual cortex. So, you should never see a visual response prior to approximately 40 ms (except under unusual circumstances). In this case, the problem arose because the event code was being generated approximately 50 ms after the stimulus was presented.

You should take away two lessons from this example. First, you need to verify that the timing of the stimuli and event codes is synchronized in your own lab. And test your timing frequency, because you can't assume that the timing will work in all experiments just because it works for one experiment. Second, always take a look at the onset times of your earliest components and earliest effects. As I noted in Chapter 8 and repeated in Chapter 10, "...it is important to look at the baseline and be very suspicious about any effects (differences between groups or conditions) that begin near time zero... In addition, you should be highly suspicious about any effects that begin within 100 ms of stimulus onset unless they reflect differences in the stimuli (e.g., larger sensory responses for brighter stimuli)."

Given the difficulty of ensuring precise timing, you should not attempt to create your own stimulus presentation programs unless you really know what you are doing (or have access to software development systems that simplify real-time control). In addition, you should not assume that commercially available stimulus presentation software has the level of precision necessary for ERP recordings. Many years ago, I tested one of the most popular stimulus presentation programs of the time, and I found timing errors of hundreds of milliseconds (it was truly shocking!). Whatever system you use, you should verify the precision of its timing. Keep in mind that a constant, known delay between the stimulus and the event code is not a major problem, because it can be subtracted away by shifting the averaged waveforms by the delay factor. An unpredictably varying delay is not acceptable, however, because there is typically no way to compensate for it.

The best way to test the timing of a stimulus presentation system is to somehow record the stimulus with your data acquisition system. With auditory stimuli, for example, you can place a microphone in front of the speakers and record the microphone's output as if it were the EEG. You can then use your data analysis software to see if the onset of the stimulus occurs exactly at the time of the condition code. This is a bit more difficult with visual stimuli, but you can get a photosensor that creates a voltage that varies according to the intensity of the light coming into it. This sensor can be pointed toward the part of the screen where the stimuli appear, and the output can be recorded by your data acquisition system. You can then see if the onset of light occurs at the same time as the event code. When you do this kind of testing, I would recommend using a higher-than-normal sampling rate to record the signals (1000 Hz or greater). This will make it easier to see these signals, which may contain very fast transitions. In addition, you should record a large number of stimuli so that you can test for the presence of occasional large timing errors. If you find timing errors, you may be able to eliminate them by disabling any unnecessary hardware and software systems on the computer. For example, you can make sure that no other programs or operating system extensions are running in parallel on the computer, and you can disconnect the computer from networks, unnecessary hard drives, printers, etc.

### *Stimulus-Related Artifacts*

Whatever stimuli you use, you should make sure they will not cause electrical artifacts that will be picked up by the electrodes or cause reflexes that will contaminate the data. For example, most audio headphones operate by passing a current through a coil to cause the movement of a membrane with respect to a magnet, and this can induce a current in the electrodes. Thus, shielded headphones may be important for auditory ERP studies. In addition, sudden noises can elicit a muscle twitch called the post-auricular reflex, and this can be picked up by the EEG electrodes. One way to tell the difference between this reflex and induced voltages from the headphones is that there should be zero delay for induced voltages (assuming that you've verified the timing of your event codes).

### *Stimulus Timing on Video Monitors*

The remainder of this chapter will focus on presenting visual stimuli on CRT and LCD video monitors. The majority of ERP experiments use such monitors, and they lead to particularly challenging timing problems that you need to understand if you're going to use them. I will just cover the basics here, but other resources are available if you want to know more (see, e.g., Brainard, Pelli, & Robson, 2002; Elze, 2010). I will start by discussing CRT displays, because they are an older and more uniform technology, and then I will discuss LCDs, which are still rapidly evolving. If you plan to use an LCD monitor, you'll still want to read the section on CRTs, which will introduce several key concepts.

**CRT Basics.** As illustrated in Figure 16.2A, CRTs operate by passing a narrow electron beam through a magnetic field, which directs it to a specific location in a phosphor layer at the front of the monitor. When the electron beam hits the phosphor, the phosphor emits light that is proportional to the intensity of the electron beam. The location that is struck by the electron beam at a given time corresponds to a single pixel on the monitor. In a color CRT, three beams are used, and they are pointed at slightly different locations containing phosphors that emit red, green, and blue light; for the sake of simplicity, however, we will assume a single electron beam except as noted.

Only one pixel is activated by the electron beam at a given moment. The magnetic field is modulated moment by moment to control which pixel is being activated. In an oscilloscope, the electron beam can be moved around in any desired pattern. In a CRT, however, a *raster* pattern is used, and the electron beam is called a *raster beam* (see Figure 16.2B). The raster beam starts in the upper left corner of the screen and illuminates all of the pixels on the top row of the monitor, one by one, with an intensity that is adjusted individually for each pixel. It then shifts down to the next row and draws the next set of pixels, and so on until it reaches the bottom right corner, at which point every pixel on the monitor will have been illuminated by some amount. The beam is then turned off briefly and moves back to the upper left corner of the monitor, and the whole process is repeated. The rate of repetition (the *refresh* rate) is usually between 50 and 100 Hz, corresponding to a time between refreshes of between 20 and 10 ms, respectively.

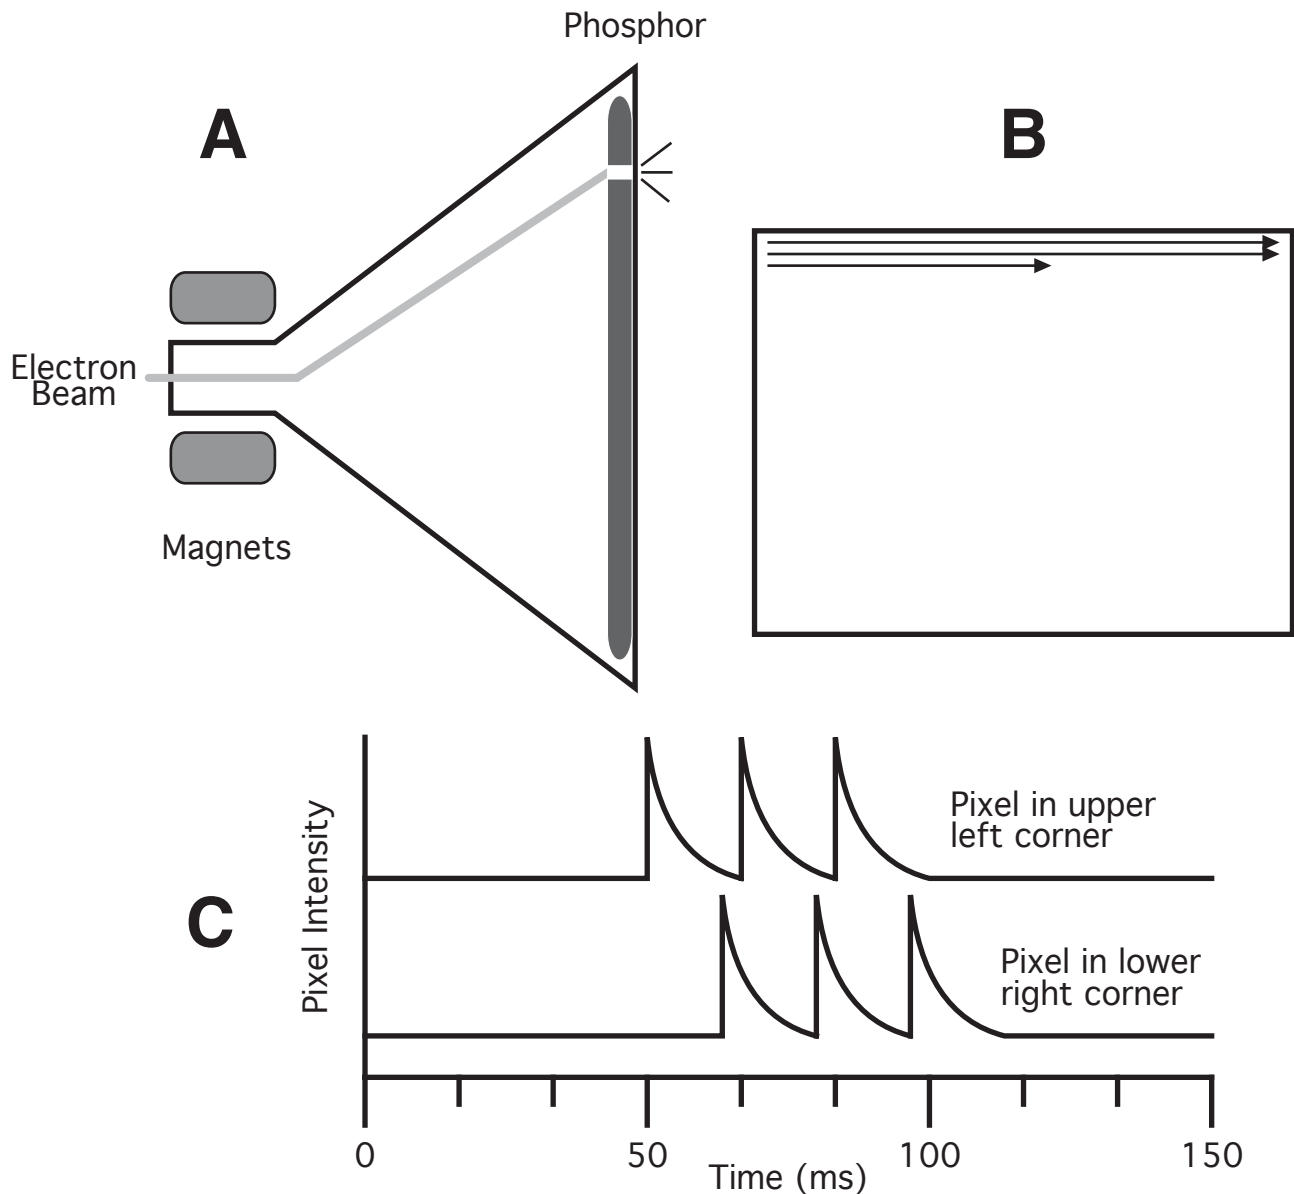


Figure 16.2. Operation of a CRT monitor. (A) The electron beam is directed by a magnetic field to a specific location at the front of the monitor, causing a small spot in the phosphor layer to become illuminated. (B) Raster pattern of the movement of the electron beam over time. (C) Intensity of the luminance from pixels on two different locations on a CRT monitor.

The intensity of the raster beam for each pixel is determined by a series of numbers stored in the video card's *frame buffer*. The frame buffer is just a contiguous set of memory locations, where each location contains the intensity value for a given pixel (or three values in the case of a color CRT, one each for the red,

green, and blue intensities). When a program draws an image to the monitor, it is really just updating values in the frame buffer. No change occurs on the monitor until the raster beam reaches the set of pixels corresponding to the updated values, at which point the values are read from the frame buffer and the pixels are activated accordingly. This is a big part of why timing is so tricky for CRTs: The time at which something actually appears on the screen is determined by both the contents of the frame buffer, which your program directly controls, and the position of the raster beam, which your program does not directly control.

As illustrated in Figure 16.2C, a pixel's intensity increases virtually instantaneously when the raster beam hits the phosphor, and then the intensity falls fairly rapidly toward zero when the raster beam moves away. As an example, consider a scenario in which (a) the refresh rate is set at 60 Hz (16.67 ms per refresh), (b) the raster beam starts at the upper left corner at time zero, and (c) the goal is to turn the entire screen on at 50 ms, leave it on for 50 ms, and then turn it off. The top trace in Figure 16.2C shows the intensity of the pixel at the top left corner of the screen for this scenario. The intensity of the raster beam will be low when it hits this pixel at times 0, 16.67 ms, and 33.33 ms, and then transition to high when it hits the pixel at times 50, 66.67, and 83.33 ms, returning again to a low value at 100 ms. The pixel's illumination will spike upward at these times and then decay downward in the intervening periods. We don't perceive it as flickering, however, because the retina integrates information over time.

Less than a microsecond after the pixel in the upper left corner is illuminated at 50 ms, the next pixel to the right will be illuminated, and then the next one, and so on until the entire top row has been illuminated (drawing one row of pixels will take approximately 120 microseconds with a 640 x 480 resolution). The pixels in the next row will then be sequentially illuminated, and then the next row, and so on until the entire display has been illuminated (which will take approximately 16 milliseconds with a 60-Hz refresh rate). The bottom trace in Figure 16.2C shows the actual illumination that would be observed for a pixel near the bottom of the display. The illumination of this pixel is delayed by approximately 16 ms compared to the pixel in the upper left hand corner of the display. Thus, the display is not drawn in a single instant, but instead is gradually painted over a period of many milliseconds.

**Timing Errors on CRTs.** If your stimulus presentation program draws an object at some random time, it may cause a large, visible artifact known as *tearing*. This happens when the frame buffer is being updated in the same set of locations that are currently being drawn to the display by the raster beam. It is important to avoid this artifact, which can be very distracting.

To provide precise timing and avoid tearing, video cards send a signal to the computer called a *vertical retrace interrupt* or *video blanking interrupt* just after the raster beam has finished a cycle of drawing and just before the next cycle is about to begin. In most cases, the best way to ensure correct timing is to use the following sequence of events in the stimulus presentation program: (a) wait for the interrupt signal, (b) send an event code, and (c) draw the stimuli in a manner that ensures that any changes in the intensity values stored by the frame buffer occur after the interrupt but before those values are read by the raster beam.

By following this sequence, you can be sure that your event code is linked to the position of the electrode beam and therefore to the moment at which the illumination of the monitor actually changes. There will be a delay between the event code and the illumination of a given pixel, and this delay will depend on the location of the pixel (i.e., it will be larger for lower rows than for higher rows). However, this delay is constant for a given location, and you can determine the amount of delay by recording the illumination at a given location by recording the signal from a photosensor at that location. A known, constant delay is a minor annoyance, whereas an unknown, variable delay can be a major problem. Note that an implication of this is that visual ERP latencies will appear to be 5-15 ms longer for lower-field stimuli than for upper-field stimuli if you don't take into account the delay in drawing stimuli toward the bottom of the display.

The tricky part of this sequence is making sure that you do not change values in the frame buffer after the those values have already been used to control the raster beam on a given raster cycle. There are two common ways in which this can occur. First, if you are drawing a complex stimulus or have a slow computer or software system, it may take longer than one raster cycle to draw the stimulus. Second, you may not draw from top to



bottom, and a portion of the stimulus display may be drawn to the frame buffer after the raster beam has already drawn this portion of the display. For example, imagine you are conducting a visual search experiment in which 24 squares are drawn at random locations across the display. If you wait for the vertical retrace interrupt and then start drawing the squares to the frame buffer in random order, you may end up drawing a square near the top of the frame buffer 5 ms after the interrupt. By this point, the raster beam will have already drawn the top portion of the frame buffer, so you may see tearing artifacts, and the square will not appear until early in the next raster cycle.

In most cases, the best way to avoid these problems is to predraw each stimulus display in an offscreen memory buffer. That is, prior to each trial of the experiment, you can predraw the stimulus displays in an area of memory that simulates the frame buffer but is not actually visible. When it is time to display the stimulus, you will wait for the vertical retrace interrupt, send an event code, and then copy the simulated frame buffer into the actual frame buffer. As long as the copying is done from top to bottom and is faster than the raster beam, this will guarantee that your timing is perfectly precise (although pixels near the bottom of the display will still be illuminated later than pixels near the top of the display). Alternatively, some video cards contain multiple frame buffers, and it is possible to predraw the stimuli in one frame buffer while another frame buffer is being displayed. The predrawn frame buffer can then be displayed by telling the video card to switch the displayed location, which is virtually instantaneous. No copying from memory into the frame buffer is necessary in this case.

*LCD Monitors.* At the time of this writing, LCD monitors have completely displaced CRT monitors for general computer use, and it's difficult to buy a high-quality CRT (unless you're willing to spend a whole lot of money for one that is made specifically for visual psychophysics). This is a shame, because CRTs have many benefits over LCDs for ERP experiments, the most important of which is that they all behave in largely the same way. LCD technologies, in contrast, are evolving rapidly, and it's hard to predict exactly how a given LCD monitor will work without extensively testing it. Wang and Nikolic (2011) published a study in which they carefully characterized a specific LCD monitor, concluding that it was sufficient for most visual studies, but by the time the paper was published the monitor was no longer being produced. If you are interested, a lot of technical information about recent LCDs can be found at <http://www.tftcentral.co.uk>.

LCD displays operate by using electrical signals to change the polarization of display elements, which in turn influences the transmission of light. No raster beam is necessary, so LCDs can potentially make instantaneous changes to pixels in any position on the screen. However, LCDs are being integrated into computer systems that were originally designed with CRTs in mind, and information from the video card's frame buffer is typically transmitted serially to the LCD display. Consequently, many of the issues that arise with CRTs also arise with LCDs. Even worse, LCDs typically produce a delay of between 5 and 30 ms between the time the video signal arrives and the time at which the display changes (above and beyond the delay that arises if the stimuli are not synchronized with the vertical retrace interrupt).

The optical properties of LCDs can also be problematic. In the context of ERP recordings, the biggest problem is that the transmission of light by LCDs changes relatively slowly. As a result, stimulus onsets and offsets will be more gradual than on CRTs. This has gotten much better over the past 10 years. However, you need to be careful of displays that advertise superfast timing, because they often achieve this by initially setting a given pixel to the maximum possible value (which causes the luminance to rise very quickly) and then reducing the intensity over a period of a few milliseconds. This causes the luminance to initially overshoot the luminance that you specify and then fall down to the appropriate level. This will not be a problem for researchers who are interested in relatively late cognitive phenomena, but it can be a problem for researchers who are interested in vision per se.

LCDs do have a very significant advantage over CRTs: They don't pump a stream of electrons directly toward the subject's head. Consequently, LCDs should produce less electrical noise than CRTs.

Given the difficulty of finding good CRTs and the likelihood of lower electrical noise in LCDs, many ERP researchers are switching to LCDs. My own lab is in the process of switching right now. If you want to switch,

I would recommend testing the timing with a photosensor immediately after purchasing a given LCD so that you can return it if the timing is insufficient for your purposes. And you certainly need to test the timing before running any experiments to determine the size of the time lag produced by your monitor.

**Software Packages.** Various stimulus presentation programs are available, and some are designed expressly for use in ERP experiments. Before purchasing a program, you should inquire about the nature of the video timing. Every vendor will tell you that the timing is accurate, but you should find out exactly how the program works. In particular, you should find out whether the program (a) synchronizes stimulus presentation and event codes to the vertical retrace interrupt, (b) predraws each display to an offscreen memory buffer, and (c) might be occasionally interrupted by other programs or the operating system, which may cause large timing errors.

If you are inclined to write your own visual stimulus presentation programs, I would highly recommend using MATLAB and a set of routines called the PsychToolbox. This toolbox was developed by two highly respected vision researchers, David Brainard and Dennis Pelli (Brainard, 1997; Pelli, 1997), and they have rigorously tested its timing. MATLAB provides an excellent environment for the relatively rapid development of new experiments, and the PsychToolbox makes it easy to implement the sequence of procedures described here. However, you need to be very careful about timing when you use PsychToolbox, because this package gives you complete control to get it right and to get it wrong.

## **Finding and Eliminating Sources of Electrical Noise**

### *The General Problem of Electrical Noise*

The voltage fluctuations of the scalp EEG are tiny, and the EEG must be amplified by a factor of 1,000-50,000 before it can be accurately measured. There are many sources of electrical activity in a typical laboratory that are much larger than the EEG, and these large electrical sources can produce small voltage fluctuations in the subject, in the electrodes, and in the wires leading from the subject to the amplifier. These induced voltage changes can be quite considerable when they are amplified along with the EEG (although once the EEG has been amplified, the effects of additional induced voltages are usually insignificant). Although some of the induced electrical noise can be eliminated by filters and other postprocessing techniques, it is always best to eliminate noise at the source (this is a corollary of Hansen's Axiom: *There is no substitute for clean data*). In this section, I will describe the major sources of electrical noise and discuss strategies for finding and minimizing sources of noise in ERP recordings.


An oscillating voltage in a conductor generates a magnetic field that will induce a small oscillating voltage in nearby conductors, and this is how electrical noise in the environment shows up in the EEG. Most of the oscillating voltages in a typical ERP lab arise from AC-powered devices, such as lights and fans. AC line current consists of sinusoidal oscillations at either 50 Hz or 60 Hz, depending on what part of the world you live in, and this can induce 50 or 60 Hz line noise oscillations in your EEG recordings. Video monitors are a common AC-powered device near the subject, and in addition to creating noise at the line frequency they may also generate noise at their refresh rate (anywhere between 50 and 120 Hz). Noise from video monitors is often spiky rather than sinusoidal. Noise from the video monitor is especially problematic because the stimuli are typically time-locked to the video refresh, so the noise waveform is the same on every trial and is not reduced by the averaging process.

### *Common Sources of Electrical Noise*

No matter whether you are using an electrically shielded chamber or not, there are several electrical devices that may be near the subject. In general, you can deal with noise from these devices by turning them off, placing them far from the subject, or shielding them.

The video monitor is often the worst source of electrical noise. If you have a shielded chamber, you might be able to install a shielded window and place the monitor just outside the window. Alternatively, you can place the monitor inside a *Faraday cage*, which is basically a small electrically shielded box with a shielded window. This is what I've used in my lab for the past 20 years. Figure 16.3 shows the design for our Faraday

cage, which can be built by any reasonably skilled woodworker for a relatively small amount of money (approximately \$400 in the U.S.). This cage consists of copper screen shielding surrounded by a wooden exterior (with ventilation holes). A shielded piece of glass is placed at the front so that the subject can see the front of the monitor. Shielded glass can be purchased from computer accessory suppliers. There are several different types, and you should get one that has a ground wire coming out of it that you can attach to the copper shielding. You should also get a shielded AC power cord, which can be obtained from an electronics supply company. A well-shielded Faraday cage can dramatically reduce electrical noise in the EEG, and it is well worth the modest expense if you are using a CRT.



*Figure 16.3. Faraday cage that can be used to reduce electrical noise from a video monitor. It consists of a plywood box with ventilation holes, completely lined with copper screen (copper sheets may be used on surfaces without ventilation holes). The front of the cage has a glare filter that is covered with a conductive film, which is connected to the copper shielding. A power strip is placed inside the cage (not shown); this is connected to a shielded AC power cord.*

In principle, LCD monitors should generate much less electrical noise than CRT monitors, but they still contain power supplies that could be a problem (especially if your system uses passive rather than active electrodes). We have recently started using an LCD monitor without a Faraday cage, but it's too soon to tell how much noise it generates. It's entirely possible that different LCD monitors will generate vastly different amounts of electrical noise depending on the shielding of the power supply, so it may always be difficult to make a general statement about electrical noise from LCD monitors. In addition, if you are using passive rather than active electrodes, you will likely pick up much more electrical noise from the monitor (whether it is a CRT or an LCD).

Fans (and other devices with electric motors) can also generate very high levels of electrical noise, with a frequency that will depend on the speed of the motor. Although it is good to keep your recording environment cool, you shouldn't try to achieve this with a fan (unless it is very far from the subject and/or outside of a shielded chamber). If you find that you have periods of large electrical noise that come and go unpredictably and last for a few minutes, this noise could be generated from ventilation fans in the ceiling above the subject.

As described earlier in this chapter, AC-powered lights can generate substantial electrical noise, and you should replace them with DC-powered lights.

Cables connected to computers and other electronic devices can also end up bringing noise into proximity with the subject. For example, when I was putting together my first lab as an assistant professor, I was surprised to find that the cables leading from a stereo amplifier outside the chamber to speakers inside the


chamber created significant electrical noise in the EEG. We have recently had a similar problem with USB cables connecting our response devices to the computers.

In many cases (e.g., speaker cables), noise from cables can be reduced or eliminated by placing them inside of a shielded tube. In the U.S., there is a product called Wiremold that is made for this purpose and can be purchased at most hardware stores and building supply stores. It consists of clips that are screwed into a wall, and then a rectangular metal tube containing the cables snaps into these clips. Wiremold is inexpensive, and it's easy to install if you can operate a drill and a screwdriver.

### *Finding Sources of Electrical Noise*

How can you determine how much noise is present in your recording environment and find the sources of noise? Why I was an assistant professor and trying to optimize the noise level in my brand-new ERP lab, I “invented” a device for detecting and measuring electrical noise (I put the term *invented* inside quotation marks because this device was simply three resistors wedged inside some cardboard). I invented this device because it was hard to measure environmental noise with a real subject, because the signal coming from a real subject contains biological noise (e.g., muscle noise and EEG oscillations) as well as induced electrical noise from the environment. To measure environmental noise, I needed to create a nonbiological “subject” that would pick up electrical noise much like a real person but without generating any additional noise.

Figure 16.4 shows the device, which you can easily build yourself. It contains three  $4.7\text{ K}\Omega$  resistors connected into a triangular configuration, with a wire coming from each vertex of the triangle. I call this a *fake head* (or a *phantom EEG source* if I'm trying to sound scientific).



*Figure 16.4. Simple simulated head for use in finding noise sources. It consists of three  $4.7\text{ K}\Omega$  resistors (or similar values), connected to form a triangle. Each corner of the triangle is connected to wires leading to the active, reference, and ground inputs of the EEG amplifier. The set of resistors should be enclosed in a nonconductive substance (such as cardboard) for ease of handling.*

Why use  $4.7\text{ K}\Omega$  resistors? I was trying to simulate the  $5\text{ K}\Omega$  nominal impedance of a subject's scalp, and  $4.7\text{ K}\Omega$  is a standard value for resistors (for some reason). You can use any resistance between  $4.7$  and  $100\text{ K}\Omega$  to make your own fake head. Why use three resistors? I was trying to simulate having an active electrode, a ground electrode, and a reference electrode, each separated by the impedance of the scalp. Each vertex of the triangle simulates a single electrode. A wire is connected to each vertex, and these wires are plugged into the active, ground, and reference inputs of your EEG amplifier. The hardest part is figuring out how to get appropriate connectors to plug these three wires into the EEG recording system. With some systems (especially those that use active electrodes), the simplest thing is to dunk the electrodes and wires into buckets of salt water and let the salt water conduct the signals from the wires to the electrodes.

By connecting the fake head to your amplifier and digitization system, you can see how much noise is present at different locations inside the recording chamber or room. First, place the simulated head where the subject sits and connect the fake head to your amplifier and digitization system using a fairly high sampling rate (e.g., 1000 Hz) and wide open filter settings (e.g., DC-300 Hz). You should see clear signs of electrical noise (including a sinusoidal line-frequency oscillation) on your EEG display. Then, turn off every electrical device in or near the recording chamber, including any AC power entering the chamber, such as the power for the fan and lights. This should yield a much cleaner signal on your EEG display (i.e., a nearly flat line). Once you have obtained a clean signal, you can turn the electrical devices back on one at a time, noting which ones create noticeable noise. You may also find it useful to place the fake head near to wires and devices that are potential noise sources, which will make it easier to see any noise that they generate. However, you shouldn't hold the fake head while doing these measurements, because you will act as a giant antenna for noise.

## References

- Brainard, D.H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-436.
- Brainard, D.H., Pelli, D.G., & Robson, T. (2002). Display characterization. In J. Hornak (Ed.), *Encyclopedia of Imagine Science and Technology* (pp. 172-188). New York: Wiley.
- Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134, 9-21.
- Elze, T. (2010). Achieving precise display timing in visual neuroscience experiments. *Journal of Neuroscience Methods*, 191, 171-179. doi: 10.1016/j.jneumeth.2010.06.018
- Gratton, G., Coles, M.G.H., & Donchin, E. (1983). A new method for off-line removal of ocular artifacts. *Electroencephalography and clinical neurophysiology*, 55, 468-484.
- Kappenman, E.S., & Luck, S.J. (2010). The effects of electrode impedance on data quality and statistical significance in ERP recordings. *Psychophysiology*, 47, 888-904.
- Lopez-Calderon, J., & Luck, S. J. (2014). ERPLAB: An open-source toolbox for the analysis of event-related potentials. *Frontiers in human neuroscience*, 8, 213. doi: 10.3389/fnhum.2014.00213
- Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437-442.
- Wang, P., & Nikolic, D. (2011). An LCD Monitor with Sufficiently Precise Timing for Research in Vision. *Frontiers in human neuroscience*, 5, 85. doi: 10.3389/fnhum.2011.00085