

# EEG BASIC PRINCIPLES—PART I

Mark Libenson, MD

Senior Associate, Department of Neurology; EEG  
Laboratory Medical Director, END Technology Program  
Assistant Professor of Neurology, Harvard Medical  
School

## Introduction

Brain waves are difficult to record. For that reason, you should never take it for granted when a page of electroencephalogram is delivered to you for interpretation. The first published brain waves were recorded in curarized dogs and reported in 1913 and were met with controversy. There are two main challenges to getting a good brainwave tracing: the various insulators that intervene between the recording electrode and the electrical generator of the wave we are trying to record (soft tissues, CSF, not to mention the skull bone), but also the intrinsically low voltage of the brain potentials themselves. While these challenges may simply seem like quaint historical curiosities in today's modern world of digital EEG recording, even today these original challenges continue to echo throughout the neurophysiologist's workday. Individuals outside the field of clinical neurophysiology are often surprised to learn that one of the most challenging aspects of interpreting the EEG is not simply recognizing each wave and its meaning, but rather deciding whether or not to reject a large number of the waves in the tracing as artifacts—waves that do not emanate from the brain. The dominance of the daily question, “is this an artifact or is this ‘real?’” relates back to the considerable challenges of recording this very low voltage activity through the skull and the high amplification necessary to create an interpretable waveform.

The first step in a fellow's developing this all-important skill of recognizing EEG artifacts is learning polarity and localization skills. Is the electrical event I am looking at widespread or highly focal? Is it all positive or all negative or are parts both? In short, is the electrical contour of this wave consistent with a biological potential from the brain, or perhaps just related to one or more poor electrode contacts? You will want to become an absolute expert in this thought process on your way to becoming an absolute expert in EEG!

---

## The Basics: EEG Waves are Voltages

It is important to keep in mind that voltages are not absolute measures at a single point, but a property that is always measured *between* to two different points. It is not the situation that a single point on the head is measured as having, for example, 7 microvolts ( $\mu\text{V}$ ). When voltages are measured and displayed, they always refer to a difference between two points. The two prongs of a voltmeter touching on any two points measure the “urge” an electron would experience to move away from the first point and toward the other (consider that a synonym for voltage is *electromotive force* or EMF). If the voltage drop is zero between two points, an electron on a wire attaching the two points would have no urge to move one way or the other along the wire—it would be happy enough to just stay put. The first concept to take away here is that every line of EEG is a readout of the voltage drop—an electron’s “urge” to move—between the two electrode locations you have decided to pair in a channel of an EEG montage. (Occasional exceptions include when one of the electrode points is replaced by an average of two or more electrodes, but we will get to that later). Whenever you look at a line of EEG, you should always be keenly aware of what the two places are that you are comparing in the recording: there is always “Place 1” and “Place 2.” Unfortunately, there is never an EEG line that reads out the EEG activity, for example, just at Fp1—it is always Fp1 compared to something else.

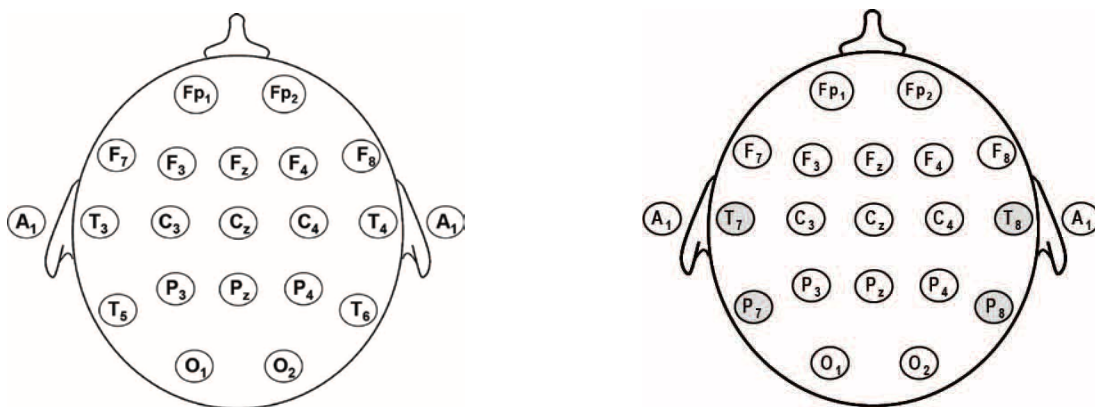
The second point is that each line of EEG essentially represents the steepness of the voltage between the two points measured, and how it changes in time. The steepness metaphor works well if you want to imagine this as an electron’s urge to roll up or down a hill. As you know, bipolar montage chains consist of successive pairs of electrodes with a common electrode in each successive pair (e.g., E1—E2, E2—E3, E3—E4, etc.). We will learn to interpret bipolar montages as successive “steepness” measurements in order to visualize the contours of the electric field they represent on the scalp. A similar, but different analysis can be done with referential montages to generate the same field contours.

## Electrode Nomenclature and Positions

By now, you are all fairly familiar with the 10–20 international electrode system for electrode placement and naming. Before the 10–20 system was adopted, some electrode application systems assumed that measurements were being made on a standard adult-sized head and used absolute measurements in

inches or cm from certain anatomical landmarks for electrode placements. The 10–20 system improved on this by making certain basic head measurements at the beginning of the application process (e.g., the distance from nasion to inion, between the two pre-auricular points, etc.) and then marking electrodes positions either 10 or 20 percent along these standardized head axes. This is all to say that a system that works by percentages functions considerably better with varying head sizes, especially those of adults with large heads, children, babies, or even premature infants. The actual resulting 10–20 positions were actually confirmed by drilling holes through the skulls and into the brain surface of carefully measured cadavers to see what part of the brain the electrode position overlaid!

A later, minor modification of the 10–20 system was made, renaming two of the standard electrode positions: T3 and T4 became T7 and T8; T5 and T6 became P7 and P8. There is currently a mildly aggravating problem where some labs have made this last switchover to the newer 10-20 system and some have retained the original 10-20 system. For that reason, it would be reasonable for you to have some acquaintance with the 10-20 version your lab does not use, including the four electrode positions whose names change between the initial 10–20 system and the revised 10–20 system.



While I have included maps of the 10–20 electrode systems here, your goal is not to need these. Instead, just as you really need to memorize the addition and subtraction tables to do math, you should have these electrode positions/names perfectly memorized. The only way you can know that you have memorized the

---

maps perfectly is to demonstrate to yourself that you can take a blank piece of paper and draw all the electrode positions onto the page from memory. If it takes you any time at all to produce this map on paper, try again and keep practicing. Once you have mastered the map, the next key skill is the ability to be given a single electrode position (e.g., C4), and to be able to immediately state all of the surrounding electrodes. When you hear C4, you should be able to rattle off F4, Cz, T8, and P4 as the surrounding contacts. This is a consistently useful skill. For instance, if you see a deflection in C4, the first question you might ask yourself in order to distinguish C4-electrode artifact from a brain wave is whether there is a surrounding field. To check this, you will want to look at the F4, Cz, T8, and P4 positions to see if some of these locations also express the deflection.


It is also important to learn the “common names” for each electrode position. In short, in your technical description of an EEG, it is reasonable to use the 10–20 electrode names, such as F7 or T8 (e.g., “a low voltage spike was seen in F8”). In the conclusion portion of the report where you give your clinical impressions and correlations, it is better to use the more common English names for electrode positions, as this part of the report should be readable by a non-neurophysiologist. Many of the common names are obvious, but some may be a bit surprising. The occipital and frontopolar contacts are as expected. The P7 (T5) contact, although it begins with ‘P,’ is the left posterior temporal contact. The C3 and C4 contacts are named for their proximity to the central sulcus. Perhaps the most confusing are F7 and F8 which are properly referred to as anterior temporal. These were initially named for coverage of the inferior frontal gyrus, but they probably pick up anterior temporal activity even better and anterior temporal spikes are often well recorded here. The midline positions, Fz, Cz, and Pz were initially called mid frontal, mid central, and mid parietal, but I think this can be somewhat confusing (sounding like, for instance, the middle of the parietal lobe) and I prefer frontal midline, central midline, and parietal midline. Left-sided electrodes always use odd numbers and the right-sided contacts use even numbers. The ‘z’ subscript of Fz, Cz, etc., stands for zero and (even though zero is an even number) refers to the midline. When discussing locations in your report using common names, be sure to be adequately specific. For instance, you would not want to refer simply to a “left temporal spike” since this could refer to a spike under any one of the F7, T7, P7, or T1 (aka FT9) contacts, all of which are temporal electrodes in their own right.

## The Pen-Up/Pen-Down Convention

---

Imagine that you are asked a very simple question: what does it mean in an EEG channel if the pen goes up (imagine a single upgoing spike)? When I ask new EEG'ers this question, the most common response is "it's negative." This is the "right wrong answer," as we shall see. The type of EEG amplifiers that we use on our instruments use a clever strategy. They are *common mode rejection* (CMR) amplifiers which means they essentially subtract Input 2 from Input 1 and amplify the results. Imagine that Input 1 consists of some crazy/noisy signal and Input 2 has the *same* crazy/noisy signal but with the addition of a spike in the middle. The interesting by-product of this subtraction is that any information or wave common to Input 1 and Input 2 subtracts and get cancelled out—the "common mode" or common part of the two waves is rejected. What remains is the difference (the spike), and that difference signal is amplified. Of course, someone had to decide on a convention where, if Input 1 is more negative than Input 2 at any given moment, whether the amplifier should be set up so that the pen deflects upward or downward. The decision was pretty much arbitrary: the original engineers who built EEG machines decided that if Input 1 was more **negative** than Input 2, the pen goes **up**. I personally find this somewhat counterintuitive. Because the electrode attached to Input 1 is very often "the electrode of interest" compared to whatever is attached to Input 2 (and this is particularly true with referential montages wherein Input 2 is usually attached to a reference electrode that is hopefully quiet, such as Fp1-ref), I would have expected that if Input 1 were more positive, the pen would go up (since we generally graph positive numbers up and negative numbers down). As it happens, the original engineers felt otherwise and the pen goes down if Input 1 is more positive! It is too late to complain. You may even decide to use the counterintuitive nature of this convention to help you memorize this setup.

So why is "it's negative" if the pen goes up or "it's positive" if the pen goes down the "right wrong answer?" Counterexamples are easy to find. Imagine you are reading and come upon a *downward* deflection in Fp1–F3.

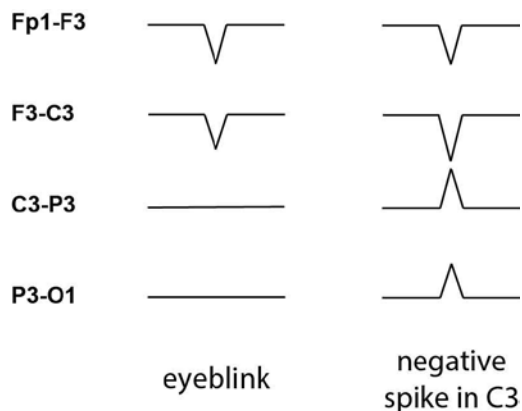


Considering this deflection alone, if you answer "it's positive" the first question you would be asked is, "well which one do you think is positive, Fp1 or F3?" and you would likely answer Fp1. Yes, Fp1 could be positive in this scenario, but it is equally possible that Fp1 is neutral and that F3 is negative. So, a more complete

answer to this downgoing deflection question would be that Fp1 could be more positive or F3 could be more negative.

Indeed, both situations are easy to imagine using real-life examples (see below). With eyeblink artifact, Fp1 is more positive than F3 and, indeed, the pen goes down in the Fp1-F3 channel when there is eyeblink artifact—this is because of a positivity present in Fp1 which comes from the positive charge on the cornea which bobs upward toward Fp1 during an eyeblink. We are used to seeing downward pen deflections in Fp1-F3 with eyeblink artifacts. So with the example of eyeblink artifact on the left side of the figure below, it's correct to say that when the pen goes down in Fp1-F3, it's *because Fp1 is positive*.

For real-life example #2, if there is a negative spike in C3, the pen in the Fp1-F3 channel may still go down exactly as it did with the eyeblink artifact example, but for the opposite reason. The successive channels on the right side of the figure below tell the rest of the story. The phase reversal in C3 tells us that the pen went down in Fp1-F3 *because F3 was negative*. These are two nice physiological examples of the same downward pen deflection in Fp1-F3 reflecting either a positivity in Input 1 (as in eyeblink artifact) or a negativity in Input 2 (as with a negative spike in C3 with a field that extends to F3).



Another possible criticism of answering “it’s negative” to the pen going up is that it is possible that both inputs are measuring positive voltages, but that Input 1 is less positive than Input 2. Input 1 could be seeing +3  $\mu\text{V}$  and Input 2 could be seeing +5  $\mu\text{V}$ . The pen would go up, and this would make the answer “it’s

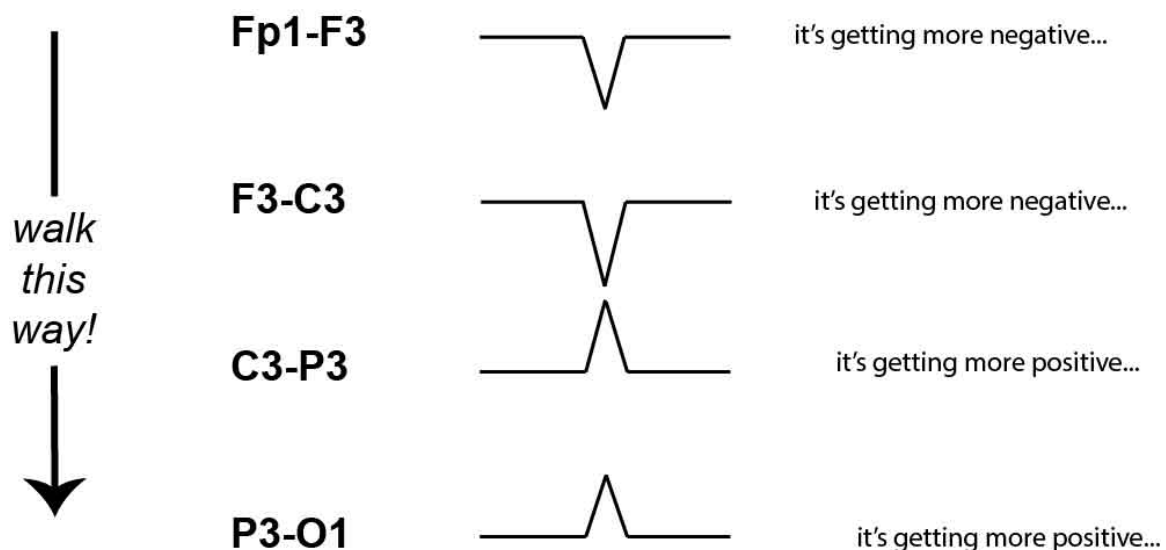


negative” seem rather unsatisfactory since both of the voltages are positive! The reason that “it’s negative” is a *right* wrong answer is that it is true in many cases. When an electrode of interest (e.g., C3) is attached to Input 1 and a relatively neutral reference (e.g., the nose) is attached to Input 2 creating a C3-nose channel, an upgoing deflection does suggest a negative event in C3. (And yes, it could also represent a positive event in the nose!)

## Bipolar and Referential Montages

### Bipolar Montages

There are two major types of montage setups used in conventional EEG interpretation. The first and most commonly used is the bipolar montage. Considering the left parasagittal chain as an example, the method here is to connect Fp1 to F3, F3 to C3, C3 to P3, and P3 to O1. Note that 5 electrode contacts yield 4 electrode pairings or *derivations*, or 4 channels. Recalling our earlier discussion of voltages and contours, each of the readouts of these successive pairs will give us an idea of the “slope” of the voltage change contour between each of the chain’s pairings of these five electrode. Stringing these slopes successively together will give us a nice representation of the voltage contour along the chain. Let’s look at an example:



---

When you look at this image, if you already have some experience in reading EEGs, you may immediately say “Aha, there is a negative spike in C3!”—and you would be correct. This time, though, I would like you to look at these 4 channels in a different way. Sweeping from the top to the bottom of the image, a specific voltage contour is suggested. In order to see the contour well, we want to imagine ourselves standing on the head at Fp1 (at the forehead) and walking down the chain, electrode by electrode, toward the back of the head. At each point we are going to dip our toe in the water, so to speak, and measure the voltage of where we are standing *compared to* what’s coming up next (remember—voltage is always measured between two points). Imagine that you are standing on Fp1 and looking down toward F3. Because the pen goes down in Fp1–F3, it is true (based on the discussion above) that this could imply either that Fp1 is more positive or F3 is more negative. Instead, I would like you to think of this deflection in a different way. Because you are walking in a particular direction, from the front of the head to the back of the head, I would like your take-away perception of this downward deflection simply to be that “things are getting more negative” as you walk from Fp1 to F3. Next, you step forward to F3 and you look at the F3–C3 deflection and again, the pen goes down and again your perception is “things are getting still more negative even as I walk from F3 to C3.” Now you have arrived at C3 and as you look ahead to P3, the C3–P3 wave implies something different is going on. The pen now goes up which tells us that as you walk from C3 to P3, “things are getting more positive.” Finally, walking from P3 to O1, the final upgoing wave tells us that things are getting more positive still. Recapping your journey and looking at all four deflections, the first two downgoing waves tell you that as you are walking from Fp1 to F3 to C3, things are getting continuously more intensely negative. At C3 you must have been standing at the trough of the negativity because now as you move forward past C3, things are getting more and more positive. The first two downgoing waves tell you that you are climbing down a valley of negativity, but once you pass C3, the pen direction has flipped up and things are finally starting to get more positive, and you’re climbing back up out of the valley of negativity.

There are two take-home lessons from this narration. The first is the concept of looking down bipolar montage chains and understanding that they describe a voltage contour. The second is an alternative way of understanding a phase reversal. When the pen direction stays the same as you go down a chain but




---

finally flips direction, this tells us that we are at the peak (or perhaps a valley) of a voltage contour. This corresponds to the basic concept that the location of a phase reversal represents either a maximum positivity or negativity, but envisioned in a slightly different way.

### Referential Montages

Let's imagine the same experience walking down the same parasagittal chain, but now using a referential montage. The first difference is that five electrode contacts generate five channels, since each contact is compared to a reference. Considering this example, let's assume that there is some ideal reference electrode which we'll call "ref." In the case of referential montages, visualizing the mountain range or valley of negativity is even easier. The height of each wave in the chain can be considered an absolute measurement of the height of the voltage, and the dots of the peaks of each waves can be connected to see the contour of the voltage map across the chain. Imagine that the height of each wave in a referential montage is the height of a telephone pole. You can now string a wire connecting the tops of each of the five telephone poles and that wire will mimic the contour of the voltage crossing the scalp, simple as that!



Once you have mastered these concepts, you can look down any bipolar chain or any referential chain and imagine a voltage contour. Of course, there are multiple chains across the head in each montage and you can probably see what's coming next. For instance, in the AP bipolar montage, there are four long chains going down the head from left to right (left temporal, left parasagittal, right parasagittal, right temporal), plus a small chain in the midline (Fz–Cz, Cz–Pz). Each of these contours can be aligned next to one another to get an idea of the whole voltage map across the head. Repeating the same exercise, but now with transverse chains from a transverse bipolar montage, in theory, the contours from the transverse chains can be used to align the five anteroposterior bipolar contours to creating a mesh-like three-dimensional gradient map of the scalp.

---

## The Ideal Reference Electrode

In the above example, we are assuming that there is some electrode location that represents an ideal reference which we have called “ref.” It is worth it to pause a moment to consider what would make an ideal reference. In fact, one of the greatest advantages of using the subtraction technique of the common mode rejection amplifier is that, if both of the electrodes being compared have a lot of electrical noise in them but it is *the same electrical noise*, that noise will cancel out and what we’ll see, ideally, is only the difference in brain wave signals between the two points—the noise will be gone. For example, if the foot and the head had the exact same noise signal, Fp1–foot would be a pretty good channel to look at the brain activity in Fp1: subtracting Fp1 from the foot would cancel out all the noise and just leave the brainwave signal of Fp1. The reason Fp1–foot does not work out so well as an EEG channel is that the electrical noise in the foot and the noise in Fp1 do differ quite a bit, so there would still be a big residual noise signal in an Fp1–foot channel—the noise would not cancel out because each point has *different* noise. As you move the reference electrode closer and closer to the head, and indeed closer to Fp1, the electrical noise signal in the reference will begin to resemble the noise signal in Fp1 progressively better and your channel will yield an increasingly cleaner readout. Of course, this goal of excellent noise reduction may tempt us to want to put the reference as close to Fp1 as possible. The pitfall here is that, as you get too close to Fp1, the reference electrode will now have some of the Fp1 brainwave activity in it as well. The extreme example of this noise reduction strategy would be to use a reference electrode which is exactly next to Fp1. Sure, all the noise is now gone from Fp1–ref channel, but the brain wave activity of Fp1 will also be recorded by the reference electrode in full, and everything, noise and brainwaves, will cancel out and you’ll get a flat line—not very useful!

The result of moving the reference closer and closer to Fp1 illustrates the basis of another generality in EEG: the greater the interelectrode distance, in general, the higher the voltage of a channel will be. This is true because both types of signal, the noise (non-cerebral) signal and the electrocerebral/brainwave signal become increasingly unlike as electrodes are moved farther apart. So in general, the signal from an Fp1-F3 channel will be lower than an Fp1-C3 (so-called “double-distance”) channel, both noise and brainwaves. This is the basis for requiring that EEG recordings done to assess for brain death (electrocerebral inactivity) be done with double-distance montages, since these montages would be more sensitive to picking up lower voltage activity.

---

In fact, there is no “perfect” reference electrode. *The perfect reference electrode would have all of the same noise signal as the electrode of interest, but none of the brain wave activity of the electrode of interest.* Alas, this holy grail of EEG can never be perfectly realized, but we can get close. Keep in mind that there is no single solution to this problem. The best reference varies from patient to patient, and even from sleep state to sleep state in the same patient. For instance, in some patients, a reference electrode chosen from the earlobe, the nose, the chin, or an electrode placed on the back of the neck can make a pretty good reference. In other patients, one or more of these contacts may be quite noisy and not suitable. The reference strategy that yields the quietest, cleanest signal is patient- and situation-specific.

## **The Average Reference: Strengths and Weaknesses**

One imaginative possible solution to the problem of finding an ideal reference point is not to use a single point reference point at all, but rather to average the voltage of all the electrodes on the head, and use that average as the comparison reference point. Often this works out pretty well. The underlying assumption that would make the average reference a good choice is that, at any time, there are so many points included in the average that all the noise signals and different waves recorded by this large number of electrodes will cancel out to zero when they are averaged together, and that average will approximate electrical neutrality which—*voilà!*—would make a great reference electrode. In many situations this can actually work out quite well, especially when the EEG is relatively quiet. If at a given point in time many of the electrodes have values of, let’s say, +5 or +9, and many others are at -4 and -8, there is a good chance that the whole group will average to something near zero. A problem arises though if there is a high-voltage event in the EEG, or if a single event involves many electrodes, or the worst case being a high-voltage event that involves many electrodes. In this case, the average of all the electrodes is far from zero and, being attached to Input 2 of the amplifier, that active average will be “subtracted” from all channels and appear in all channels—but upside-down because it is subtracted.

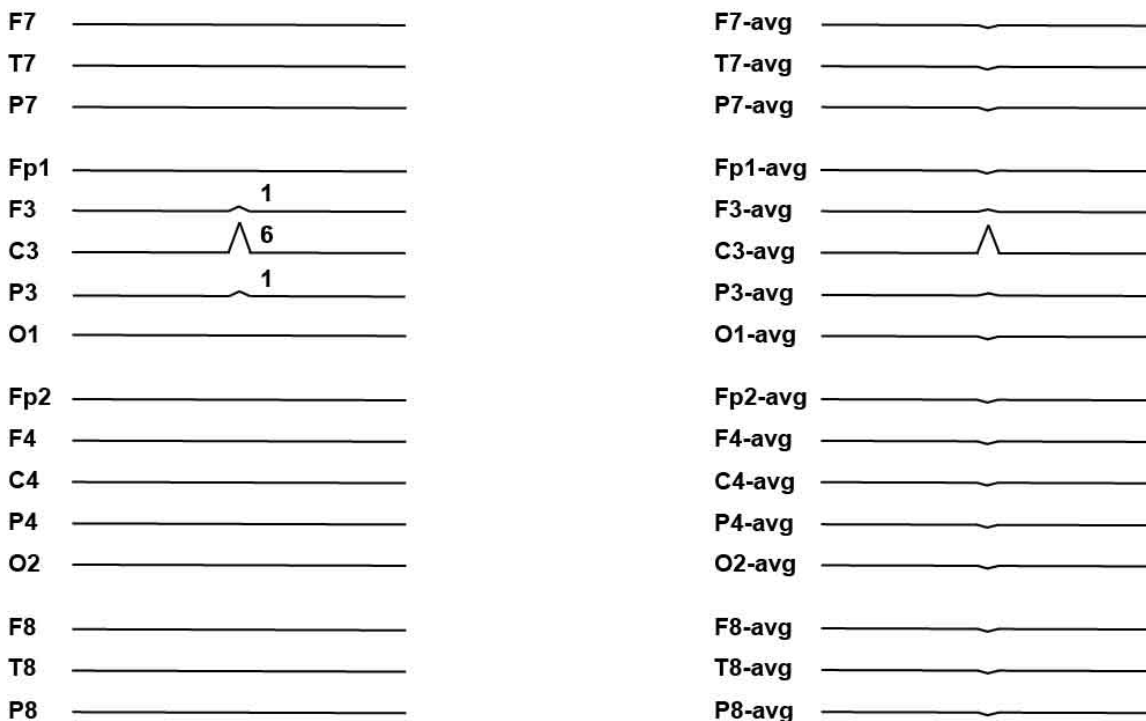
In fact, when you’re looking at any EEG and contemplating switching over to an average reference montage, before making the switch you might practice estimating in your head what the average of all the contacts would look like at any given moment on the page before the switchover. You can start this

---

exercise during a segment that includes eyeblink artifact or vertex waves of sleep, both of which are usually widespread enough and of high enough voltage to have an impact on the average.

In general, there are two types of techniques you might use to estimate the average. One way, the numerical approach, would be to quantify the height of every wave going up and down, add them all together, and then divide the sum by the total number of channels. The quicker and more practical approach would be to do this visually, weighting different heights in your head and saying, at this point, the average of all these wave heights I see going up and down would be approximately such and such a height. In this way, looking at all the ups and downs on the page at a particular moment, you are visually estimating the arithmetic mean.

Let's consider two examples. In the first example there is a well-defined, low voltage spike in one of the 16 channels (C3), with a small amount of adjacent electric field in two of the surrounding electrodes (F3 and P3). The figure below depicts just such a discharge, with the left panel showing the spike in referential montage using an imaginary ideal reference.



We will then try to predict what this discharge would look like if displayed in an average reference montage. We have the opportunity to calculate a quick arithmetic average:  $6 + 1 + 1 = 8$ . The total of 8 divided by the 16 channels = 0.5. In an average reference montage (e.g., F7-avg, T7-avg,... etc.) each electrode channel is compared to this average reference. Note that, at the time of the spike the average reference has a value of 0.5, and it is attached to Input 2 of the amplifier, so effectively 0.5 will be subtracted from all lines/channels. As the right panel of Figure 7 shows, when this spike is displayed in the average reference montage, there is a very nice representation of what is going on. The 0.5 value is what we consider the “contamination” of the average reference. We see in the right panel that each line that we know to be neutral (including the whole of the right hemisphere) shows a very small downward “notch” of 0.5 units in height. Still, this is a fairly acceptable representation of the left central spike, and the very small notches (and small 0.5 unit loss in height of the truly active electrodes) is not very distracting. It is interesting to consider that, while these small downgoing notches are visible in this idealized example in which all of the lines are completely flat except for the moment of the spike, in a real EEG the background will have a variety of irregular waves in it. Even if these are of low


---

voltage, this little 0.5 unit downward notch may get visually lost in an active EEG background. Therefore, for this particular example, the average reference is a satisfactory choice to demonstrate this spike.

For the second example, we will look at the opposite situation, where there is a discharge in the EEG of sufficient voltage, and involving so many channels, that the average reference does become highly contaminated. We will then see what type of havoc this wreaks on the EEG display when the average reference is used. In the example shown in the figure below, there is a broad spike over the left hemisphere ranging from 6 to 10 units in height. In this idealized example, the right hemisphere is completely quiet. Adding up the values to calculate the average reference we get  $6 + 7 + 6 + 8 + 9 + 10 + 9 + 9 + 0 + 0 + 0 + 0 + 0 + 0 + 0 + 0 = 64$ . 64 divided by 16 channels equals 4. Now when the average reference is used to display this spike, since it will be attached to Input 2 of the CMR amplifier, 4 units will be subtracted from every value/wave-height from the left side of the figure panel. Displayed in the average reference in the right panel of the figure, the spike in the left hemisphere can still be seen with values of 8, 9, and 10 now displayed with wave heights of 4, 5, and 6. The real confusion comes when we look at the right hemisphere. We know from the left panel that the right hemisphere is truly electrically quiet. Yet there is now a downgoing wave of 4 units in every channel over the right hemisphere.


Consider what it would be like to look at this average reference displayed without knowing the original picture. Would it be obvious to you that this is a display of a pure left hemispheric spike? What we have learned is that the average reference can be a clever way to cancel out the average noise (and brainwave activity) in






the scalp electrodes. The downside is that when there is a big voltage event, either a high voltage event, an event involving many electrodes, and especially with high voltage events involving many electrodes, there can be significant contamination of the average and unexpected and confusing results may appear on the page. In fact, the result is not unexpected, but comes “straight out of the math” based on the setup of the average reference. The problem is, when you are reading an EEG, you don’t always happen to see the first image followed by a transformation to the average reference. You may only be looking at the average reference and need to piece together the nature of the electric field from that. While these examples of the “good and the bad” of the average reference are instructive, more subtle versions of errors related to contamination of the reference electrodes can also occur, such as using the earlobe electrodes when looking at midtemporal spikes, etc.

For those who can’t imagine why the average of all the scalp electrodes does not always add up to zero as a matter of “charge balance,” here are two simple schematic counterexamples.



Counterexample #1



Counterexample #2

In example 1, the top half of the head, modeled as a sphere, is covered with electrical negativity. The bottom half happens to be positive at this moment. Because the scalp electrodes are only placed on the “northern hemisphere” of the head, all the recording electrodes could be negative at one time (resulting in an average that is strongly negative). Example 2 may be yet more “physiological.” In this example, the scalp is covered with negative charges because the other ends of the dipoles (the positivities) for this event are deeper within the head, too far away to be measured by the recording electrodes. In this case again, the whole scalp will average out to a negative value at this point in time (with the deeper positivities remaining unrecorded).