

input amplitude, and the previous history of stimulation decides which of the two outputs is observed. This phenomenon is a form of memory.

Evoked potentials recorded from the tectum of goldfish provide a dramatic example of a physiological multivalued characteristic.²⁰⁴⁸ The goldfish eye was stimulated by deep red (650 nm) light flickering at 10 Hz with 100% modulation depth while EPs were recorded simultaneously from the tectum (i.e., the goldfish brain) and from an electroretinogram (ERG) electrode. Light intensity (mean luminance) was slowly increased by almost 100 times, and then slowly decreased to the initial value with a triangular waveform, the whole cycle taking 54 sec. Figure 1.7A shows that the 10-Hz component of the ERG followed the light intensity rather accurately, and reached peak amplitude when the light was brightest. Two 54-sec records are superimposed to illustrate the reliability. The curious behavior of the simultaneously recorded tectal evoked response is shown in Figure 1.7B. EP amplitude did not follow the light intensity accurately at all; the relationship between EP amplitude and stimulus intensity was not even monotonic. As stimulus intensity was progressively increased, response amplitude first rose to a maximum, then fell to a minimum and then rose again. A progressive reduction of light intensity also produced nonmonotonic change in EP amplitude (Fig 1.7B, right side of trace). The left and right sides of the trace in Figure 1.7B are quite different; the EP showed marked hysteresis. This curious behavior was accurately reproducible; Figure 1.7B shows two superimposed traces.

The second harmonic (20-Hz) component of the EP showed even more dramatically nonlinear behavior. Figure 1.7D shows that EP amplitude actually fell to zero twice while light intensity was increasing, and fell to zero twice while light intensity was falling. At the same time, the ERG's behavior was much less nonlinear (Fig 1.7C).

This very nonlinear behavior of the goldfish tectal EP is evident when stimulus intensity is continuously changing; it is a dynamic nonlinearity and the relation between tectal EP amplitude and light intensity is quite different for different rates of change of intensity.

The ERG and tectal EP to abrupt changes of intensity are shown in Figures 1.7E and F. Figure 1.7E shows the 10-Hz component of the ERG and of the tectal EP recorded while light intensity was abruptly increased by a factor of 4, held constant for 240 sec, and then abruptly returned to its initial value. The ERG gave a strong sustained response, and ERG amplitude approximately doubled. A weaker transient response is also evident. The tectal EP, on the other hand, gave a predominantly transient response. The tectal EP's sustained response was much weaker than the ERG's; a fourfold intensity change produced only a small fractional sustained

change in EP amplitude. (Note that these responses are not sustained responses to steady light, but rather the amplitude of the 10- or 20-Hz response to 10-Hz flicker.)

Continuously changing light intensity is by no means a contrived laboratory situation; the everyday visual world comprises areas of very different luminances, so that retinal illumination continuously changes as the eye scans the visual scene. It is interesting that the nonlinear behavior shown in Figure 1.7 was much less evident when the goldfish eye was stimulated by a 10-Hz alternation between different colors rather than a 10-Hz alternation between different intensities, suggesting that red, green, and blue channels have matched nonlinearities so that the *relationship* among red, green, and blue signals depends much less on the rate of change of intensity than does either signal alone. The particular phenomenon illustrated in Figure 1.7 has not been found in the human EP,¹⁹⁹¹ but there is a hint that nonlinear resonance may occur in the human VEP. The 16-Hz VEP resonance curve in Figure 2.144A has a low-frequency falloff that is steep, but not vertical. However, the published curve incorporates data from many separate recording sessions, and in several individual sessions the falloff was effectively vertical; the many data points in Figure 2.144A blur this effect.

Systems that exhibit the jump phenomenon may also generate subharmonic frequencies in response to an input sinusoid. For example, a sine wave input of frequency F Hz may generate subharmonic components of frequency $F/2$, $F/3$, or $2F/3$.⁹⁹⁰ In principle, the particular subharmonic and harmonic components generated can provide a clue to the kind of nonlinearity being investigated.⁹⁹⁰ Methods for recording subharmonic components of EPs are described in Sections 1.8.3 and 1.8.6.

1.2 Electrodes and Amplifiers

1.2.1 Electrode Application

Although electrodes are cheap and at first sight seem simple, the electrical connection between the subject and the analyzing equipment should not be regarded lightly. If electrode connections are inadequate, one's sophisticated analyzing equipment is likely to display sophisticatedly misleading artifacts. Several practical accounts of electrodes, electrode applications, and amplifiers have been published.^{447,448,816,838,867,1974}

When thinking about electrodes and amplifiers it is helpful to regard EP recording under two headings: AC (i.e., alternating current) recording, and DC (i.e., direct current) recording. Difficulties with DC drift in the electrodes (and amplifier) can be severe when recording low-

level signals with a DC-coupled amplifier. On the other hand, an AC-coupled amplifier with a long time constant recovers slowly from transient overload, whereas a DC-coupled amplifier recovers quickly. True DC recording has been advocated when recording slowly changing potentials such as the contingent negative variation (CNV).¹⁰⁴³

An electrode is the point of connection between the complex physiological electrolytes of living tissue and the metallic recording equipment. The metal-to-electrolyte junction itself gives rise to potential differences that can be large relative to the EP signal and in addition can drift with time. This standing potential and especially its drift can be a serious problem in recording slowly varying brain potentials. If not treated with respect, it can even present problems in recording steady-state EPs because some apparently AC-coupled amplifiers are DC-coupled at the input stage, and can be blocked by a sufficiently large DC input voltage. Standing potential differences between electrodes comprise a bias voltage⁶⁷³ and a true polarization voltage. A bias voltage results when differences between electrodes cause an imbalance in the net electrode–electrolyte ionic transfer so that voltages e_1 and e_2 are different (Fig 1.8). In effect, the subject's head becomes a primitive electric cell that will generate current if the two electrodes are connected together. The remedy is that all the electrodes used should be identical, preferably with surfaces of the same pure contaminant-free metal. This consideration and the considerations of inertness and harmlessness to living tissue are the reasons for the popularity of silver, gold, and platinum electrodes. Silver electrodes can become tarnished and this surface contamination leads to bias

potentials, so gold electrodes are often preferred in routine EP recording.

Nevertheless, silver, gold, and platinum can be *polarized*. Polarization occurs when a current passes between a pair of electrodes and causes electrolysis. (The effect is similar to the familiar polarization effect in a primitive, e.g., Leclanché, electrical cell.) Polarization produces a standing voltage between electrodes, and variations in this voltage are a source of noise. Polarization also creates a surface layer whose capacitance shunts the EEG signal and can affect the electrode's frequency response (Fig 1.8).^{447,816,867,2222,2223,2509} Modern amplifiers do not draw sufficient current to cause appreciable polarization, especially when electrodes are of sufficiently large area that current density is minimal.⁸³⁸ However, a substantial polarization effect can be produced by measuring electrode impedance with an ohmmeter that forces a direct current through the electrodes. (This is quite evident: the resistance reading will gradually rise as polarization builds up, and reversing the leads will produce a different reading.) Subsequent changes in this polarization voltage can cause problems even when AC-coupled amplifiers are used. *Electrode impedance should never be measured with a DC ohmmeter.* Artifacts caused by changes in bias potential during a recording are minimized by low interelectrode impedance.⁸³⁸

Polarizable electrodes cannot be used for DC recording. Nonpolarizable or “reversible” electrodes must be used for DC recording and also when frequencies below about 2 Hz are of interest (e.g., in CNV recording). A nonpolarizable electrode is one whose properties do not change if current is passed, and any ion transfer that does occur is completely reversed by reversing the current. A silver–silver chloride (Ag–AgCl) electrode is the most common. Silver electrodes can be made reversible by placing them in pure saline solution and passing a low DC current between them in the dark. Details of the procedure for chloriding and maintaining electrodes are given by Cooper⁴⁴⁷ and other authors.¹⁴⁹⁶ Reversible electrodes that reduce current density by providing a larger surface area than the simple disk are commercially available (Fig 1.9A). These electrodes should be attached firmly with Collodion and left about 15–30 min to stabilize. It is, of course, essential not to mix nonpolarizable and polarizable electrodes.

Platinum needle electrodes can be applied quickly and are sometimes used to obtain EPs in difficult recording situations, for example, with comatose patients.⁸⁶⁹ Careful sterilization is, of course, essential; viral hepatitis, AIDS, and other diseases might be transmitted by an inadequately sterilized needle electrode.^{443,869} Grass Instruments Company provides sterilization instructions. It is widely supposed that needle electrodes produce more artifact than surface electrodes, but this supposi-

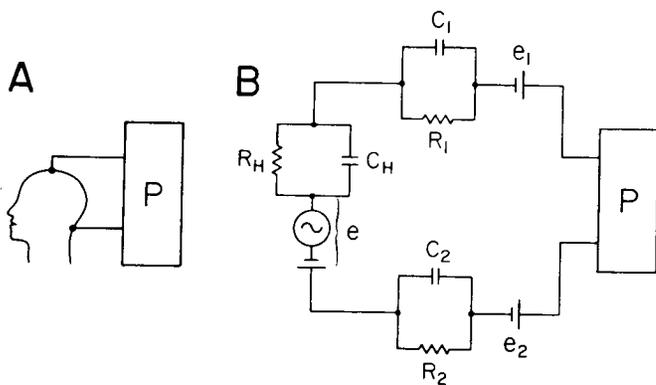


Figure 1.8
Simple Equivalent Circuit of Two Recording Electrodes Connected to a Subject

e , EEG voltage; R_H and C_H , resistive and capacitive components of head impedance; e_1 and e_2 , metal–scalp voltages (bias plus polarization voltages); R_1, R_2, C_1, C_2 , resistive and capacitive components of electrode impedance.

tion has been challenged:²⁷³⁸ contamination by potentials generated at the skin surface may even be less than for surface electrodes. Although it has been found that impedance varies with frequency, this does not seem to be a serious problem if amplifier impedance is greater than 1 million ohms (1 MΩ). Electrode impedance, however, tends to be higher (up to 10 times higher⁸³⁸) for needle than for disk electrodes.

Surface disk electrodes are more commonly used than needle electrodes, and when EEG components below about 2 Hz are not of interest the gold disk or cup electrode is quite satisfactory. In choosing the type of surface electrode and method of application, a compromise must be made between, on the one hand, speed of application and, on the other hand, the length of time for which the electrodes will remain securely fixed while maintaining a low resistance. There are two common types of cup electrode. The first kind is a hemispherical

dome with a hole through the top, made either of gold or of silver (Fig 1.9B). The second kind is also a gold or silver disk, but it has no hole through the top. Electrode application involves four stages: skin preparation, electrode attachment, insertion of conducting jelly, and verification that the electrode resistance is about 5 kΩ or less. If gold electrodes are used, each electrode should first be rubbed with cotton wool soaked in ethyl alcohol to remove any dirt or grease. (Sintered electrodes are best cleaned with ultrasound and mild detergent [W. W. Dawson, personal communication, 1986].) Then a small area of scalp should be cleaned by rubbing it with a piece of cotton gauze damped with ethyl alcohol, or a detergent (e.g., Liqui-Nox, Alconex Inc, 215 Park Avenue South, New York, NY 10003), or one of the special preparations for cleaning skin.

Every laboratory has its own way of applying electrodes; for long-term recordings (1–15 hr), I have found the following procedure satisfactory. Holed-dome gold electrodes are applied using Collodion glue (Maridon Corp, 3953 Bruner Terrace, PO Box 627, Stuart, FL 33495-0627; S.L.E. Limited, Campbell Road, Croydon, Surrey, CR0 2SQ, England). The electrode is held onto the cleaned area of scalp with a small rod or with the end of a commercially available air gun (S.L.E. Limited). Then the glue is either painted around the edge of the electrode with a brush or squeezed from a tube. Drying can be speeded up by the use of a hair dryer or with the air gun. The hole in domed electrodes allows jelly to be

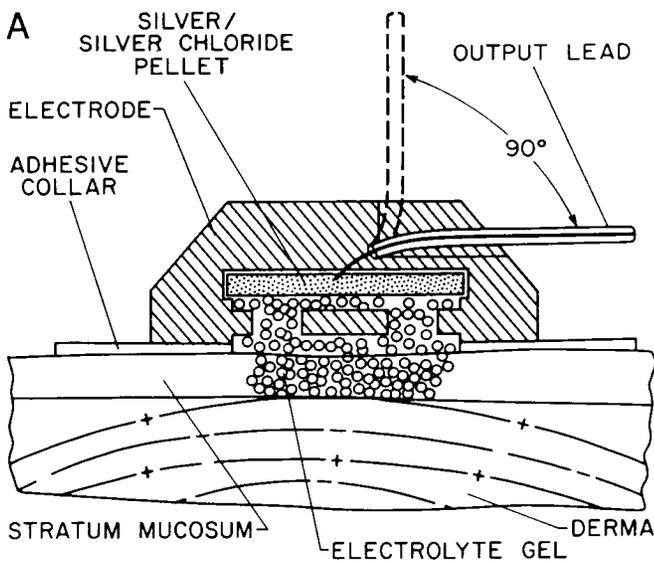
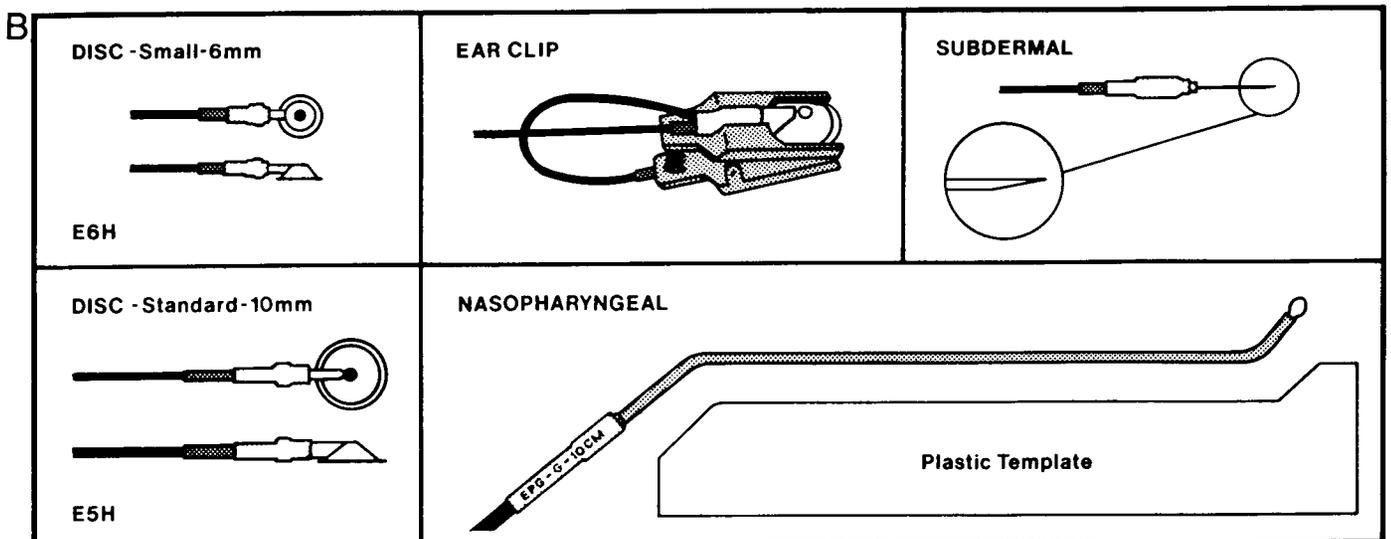


Figure 1.9 Electrodes

(A) Nonpolarizable or reversible electrode for DC and low-frequency recording. (Reproduced by permission of Beckman Instruments Inc, Palo Alto, CA.) (B) Commonly used electrodes for EP recording (Reproduced by permission of Grass Instruments Co, Quincy, MA.)



injected through the blunted needle of a syringe. An obstinately high interelectrode resistance is more often due to a faulty or dirty electrode or to inadequate cleaning of the scalp than to "skin resistance." Care in applying electrodes is worthwhile, if only to avoid the annoyance of having an electrode fall off at a crucial stage of an experiment. The likelihood of such a disaster can also be reduced by wrapping a piece of gauze around the electrode wires and firmly clipping them to the subject's collar.

It is essential to keep electrodes, needles, and syringes clean and properly sterilized and to be clear whose day-to-day responsibility this is. This was always so, but with the emergence of AIDS the most stringent precautions are now mandatory for control subjects as well as for patients. A useful handbook, "Electrode Maintenance and Infection Control in the EEG Laboratory," is available from Grass Instruments Company (Quincy, MA) extracts from which are included in this book as Appendix 1.3. Although this pre-AIDS advice concerns mostly skin infections and hepatitis, it includes the following ground rule: "It is necessary that all patients be regarded as possible sources of infection." The emergence of AIDS indicates that, today, the basic ground rule that must always be obeyed is that "*All subjects, whether patients or controls, should be regarded as possible sources of infection.*" An important point made by Grass and Grass is that the effectiveness of all practical methods for sterilization and disinfection can be reduced by dirt and contamination (e.g., by dried blood). Therefore, electrodes should be thoroughly cleaned before they are autoclaved. A useful addition to hand cleaning is the use of an ultrasonic bath containing water and detergent.

If an earlobe reference is used, then some discomfort may be produced as the electrode glue gradually dries and contracts. This discomfort may be avoided by fixing earlobe electrodes by means of commercially available sticky rings. An ear clip electrode (Fig 1.9B) is convenient, especially for short-term recording.

As already mentioned, electrode resistance should be measured using an AC source rather than a DC source to avoid polarization and ion transfer. The AC frequency should be in the EEG range, for example, 10–100 Hz for cortical EPs and about 1,000 Hz for brainstem recording. If the Collodion solvent, acetone, is used in removing electrodes it can be applied with a gauze pad. Application should be sparing because acetone can irritate a sensitive skin and is suspected of being damaging. It is very difficult to remove every particle of glue; one hopes for tolerant subjects: *Remember that ethyl alcohol and acetone are flammable, both in liquid and in vapor form, and that vapor can drift a considerable distance. Smoking should be strictly forbidden in the EP laboratory.*

A more rapid method for attaching an electrode is as

follows: First clean a small area of scalp, rub some electrode paste onto the scalp, and then fill the hollow of an electrode with electrode paste. The electrode can be held on the scalp by crossing two pieces of tape over it or by means of a Velcro headband. Alternatively, a 2 × 2-in. square of cotton gauze soaked in electrode paste can be placed over the electrode and allowed to dry. This method is more suitable for short than for long recording sessions. Electrode impedance does not reach a stable value until about 30 min after application.⁸¹⁶ Electrode jelly tends to dry after several hours, and fresh jelly should be carefully injected through the hole in domed electrodes during long recording sessions.

The use of nasopharyngeal electrodes has been reviewed by several authors.¹⁵³⁶ (See the Grass Instruments handbook mentioned earlier.)

Even when no more than three or four electrodes are used, the process of attaching electrodes can take up an appreciable proportion of the subject's total available time, and this wastage will be repeated in every experiment. It is therefore well worthwhile to develop an efficient routine procedure for applying electrodes. Electrode accessories are available from several companies, including S.L.E. Limited, Campbell Road, Croydon, Surrey CR0 2SQ, England (domed-hole electrodes, Collodion glue in tubes, air gun and compressor, electrode jelly, needles); Grass Instruments Company, 101 Old Colony Ave, Quincy, MA 02169 (domed-hole electrodes, electrode jelly and pastes). Collodion in ether/ethyl alcohol can be made up in hospital or chemistry stores.

1.2.2 The Ten-Twenty and the Queen Square Electrode Systems

The ten-twenty electrode system stems from an attempt to place particular electrodes over particular brain areas independently of skull size. The anterior–posterior (A–P) measurements are based on the distance between the nasion and the inion² as measured along the midline over the vertex (Fig 1.10B). Five points are marked along this line, designated frontal pole (F_p), frontal (F), central (C), parietal (P), and occipital (O). Point F_p is 10% of the nasion–inion distance above the nasion. F is 20% of this distance back from point F_p , and so on in 20% steps for points C, P, and O (hence the name ten-twenty system). Lateral measurements are based on the distance between the left and right preauricular points² measured through point C at the vertex (Fig 1.10A). Ten percent of this distance from the preauricular points gives the left and right temporal points (T), and the central points lie 20%

² The nasion is the delve at the top of the nose, level with the eyes, and the inion is the bony lump at the base of the skull on the midline at the back of the head. The preauricular points are where the external ear flaps merge with the scalp on the face side.

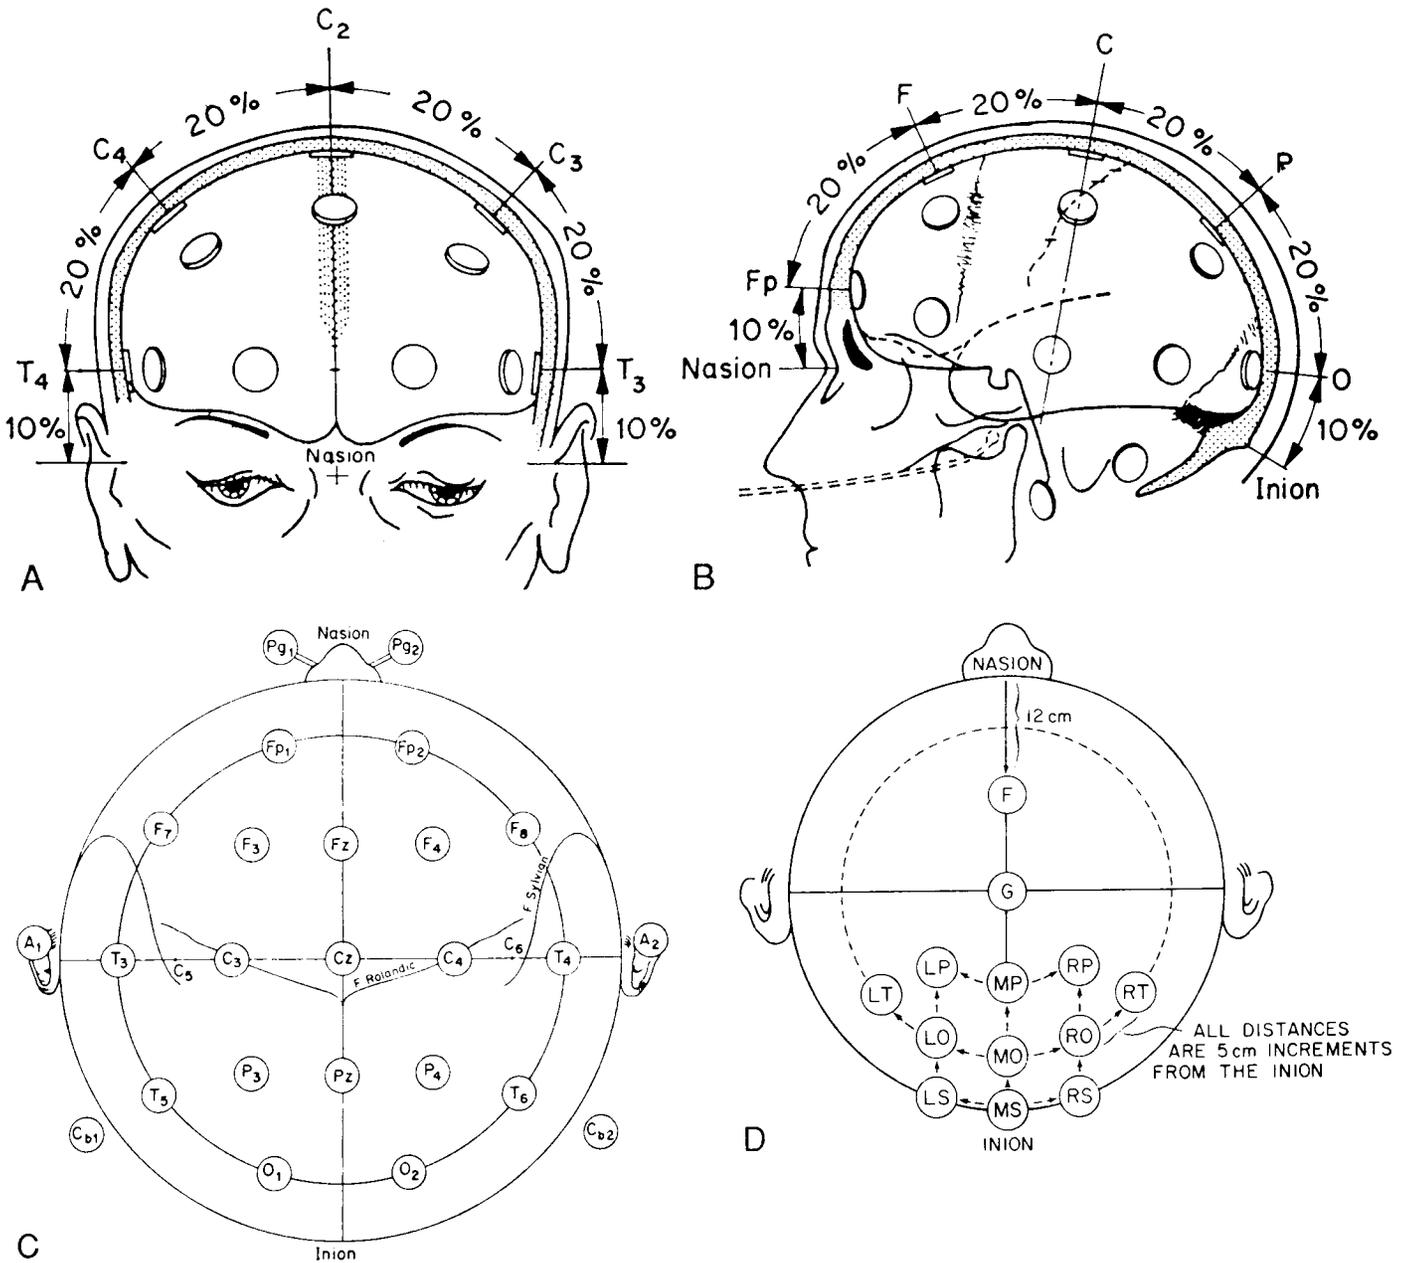


Figure 1.10
Ten-Twenty and Queen's Square Systems of Electrode Placement

(A) Frontal view of the skull showing the ten-twenty method of measurement for the central line of electrodes. (B) Lateral view of skull to show the ten-twenty method of measurement from nasion to inion at the midline. F_p, frontal pole position; F, frontal line of electrodes; C, central line of electrodes; P, parietal line of electrodes; O, occipital line. Percentages represent proportions of the measured distance from the nasion to the inion. Note that the central line is 50% of this distance. The frontal pole and occipital electrodes are 10% from the nasion and inion, respectively. Twice this distance, or 20%, separates the other line of electrodes. (C) A single-plane projection of the head, showing all ten-twenty standard positions and the location of the Rolandic and Sylvian fissures. The outer circle was drawn at the level of the nasion and inion. The inner circle represents the temporal line of electrodes. (D) Queen Square system. F, frontal; O, occipital; P, parietal; S, suboccipital; T, temporal; G, ground; L, left; R, right; M, midline. These diagrams can be made into useful stamps for the indication of electrode placement in routine recording. (Panels A–C are from Jasper HH: *The ten-twenty electrode system of the international federation*. *Electroenceph Clin Neurophysiol* 1958;10:371–375. Reproduced by permission.)

of the distance above the temporal points. The A–P line of electrodes over the temporal lobe, frontal to occipital, is determined by measuring the distance between point F_p through the T position of the central line and back to point O. Even numbers are used as subscripts for the right hemisphere, and odd numbers for the left hemisphere. Electrodes in the midline have z (zero) as subscript. This provides a total of 21 electrode positions. Intermediate positions (e.g., F_2 , C_2 , C_6 , etc.) can be added. Figure 1.10C shows additional pharyngeal electrodes Pg_1 and Pg_2 , and electrodes over the posterior fossae Cb_1 and Cb_2 .

A useful practical guide on the ten-twenty system and the attachment of electrodes is available from Grass Instruments Company (P. F. Harner and T. Saint, "A Review of the International Ten-Twenty System of Electrode Placement").

The Queen Square system of electrode placement has been proposed as standard in recording pattern EPs in clinical testing.²¹¹ All distances are in 5-cm increments from theinion (Fig 1.10D). The midline occipital (Mo) electrode is placed 5 cm above theinion, the right (Ro) and the left (Lo) occipital electrodes are placed 5 cm lateral to the midline electrode, and the frontal (F) is placed 12 cm above the nasion. The ground electrode is placed at the vertex. The lateral occipital electrodes are located in a more favorable position to record responses to hemifield stimulation than are the corresponding O_1 and O_2 of the ten-twenty system.

1.2.3 Electrostatic Field, Potential, and Potential Difference

The electric potential V at any point is a scalar quantity; it has magnitude but no direction. It is commonly measured in volts. In Figure 1.11 the quantities V_1, V_2, \dots, V_i are absolute potentials, that is, expressed relative to electrical infinity. The lines are equipotentials, that is, lines of constant potential. The electric field \vec{E} is a vector quantity; it has both magnitude and direction. It is commonly measured in volts (V) per centimeter. If the potential difference between two closely adjacent points is ΔV V, and the points are very close together, then the magnitude of the electric field component passing through the two points is given by $\Delta V/\Delta d$ V/cm at that location in space where Δd is the separation of the two points in centimeters. In Figure 1.11, $E \cos \theta$ is the magnitude of the electric field component at point M that is at an angle θ to the direction of \vec{E} .

1.2.4 Bipolar, Monopolar, and Referential Recordings

Section 1.2.9 notes that the chief reason for using differential amplifiers in EP work is that a differential amplifier rejects common-mode interference; that is, if the two

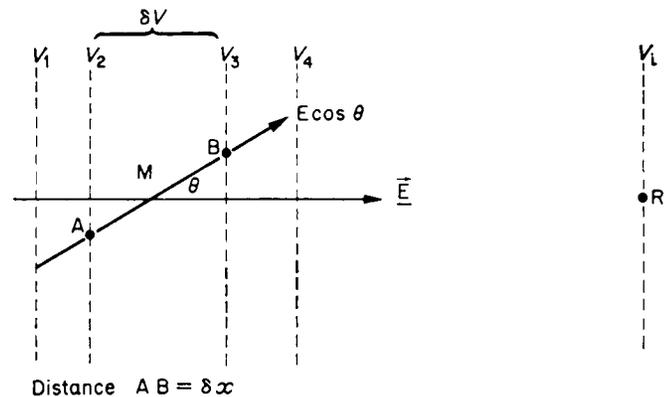


Figure 1.11
Electrostatic Field and Equipotentials

V_1, V_2, \dots, V_i , equipotentials; \vec{E} electric field vector at point M; $E \cos \theta$, the component of the electric field at point M that is at an angle θ to the direction of \vec{E} . (From Regan D: *Evoked Potentials in Psychology, Sensory Physiology and Clinical Medicine*. London, Chapman & Hall, 1972. Reproduced by permission.)

active inputs change by the same voltage with respect to ground, then this change does not appear at the amplifier output.

In EP work one electrode is connected to one of the differential amplifier's inputs and a second electrode to the other input. A differential amplifier measures the difference in potential between the two input electrodes. The voltage and EP waveform recorded generally depend on the locations of *both* electrodes. Some authors refer to a closely spaced electrode pair as bipolar, recognizing that both electrodes are affected by the activity of the underlying brain, and distinguish this arrangement from the so-called "monopolar" arrangement where one electrode (the so-called "indifferent reference") is supposed to be unaffected by the activity of the underlying brain tissue. However, it is important to remember that there is *no location on the body that can be regarded as being at electrical infinity for all source configurations* in the sense that the EP recorded from the active electrode is unaffected by the reference electrode. In a vain attempt to achieve electrical infinity, some authors have placed one of the two input electrodes physically distant from the head, for example, on the arm. The term "monopolar" is best forgotten. In reality all EP recordings are bipolar.

Nevertheless, there is a practical distinction between a closely spaced electrode pair and a widely separated pair. A closely spaced pair is sensitive to the location and orientation of nearby source(s), whereas a widely separated pair is comparatively insensitive to source location. Figure 1.12A illustrates one consequence: In principle, a closely spaced electrode pair could isolate the contribution of a weak source that gave opposite-polarity voltages at the two input electrodes from the contri-

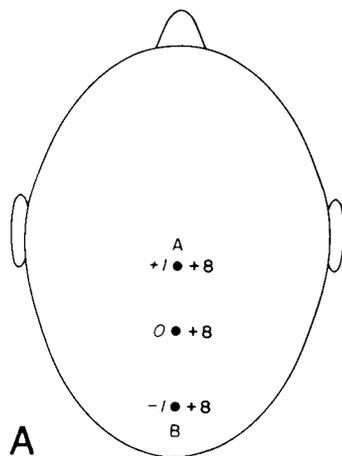


Figure 1.12A
Closely Spaced-Versus-Widely Separated Input Electrodes
 Diagram of head, showing at three points the instantaneous potential (in microvolts) due to two notional intracranial generators. The potentials due to one source are marked in italics and the potentials due to the other source are marked in boldface. The net values of potential are (reading downward) +9, +8, and +7 μV . Electrodes at A and B, each connected with respect to a distant reference electrode, would assume potentials whose amplitudes differed by only roughly 20%, so that the contribution of the italicized source could be assessed only imprecisely. However, recordings made *between* A and B would cancel the contribution of the stronger boldface source and give a more precise measurement of the small italicized source. (From Regan D: *Evoked Potentials in Psychology, Sensory Physiology and Clinical Medicine*. London, Chapman & Hall, 1972. Reproduced by permission.)

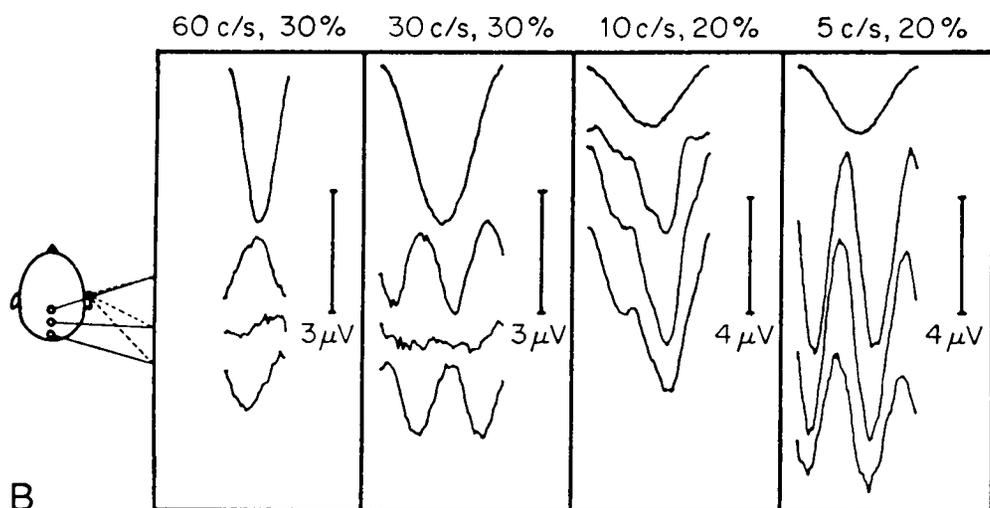


Figure 1.12B

Practical Comparison of Closely Spaced and Widely Separated Electrodes

This is a practical example of the hypothetical situation shown in Figure 1.12A. The uppermost trace in each panel is the stimulus light waveform recorded by a photocell. The traces on the far left are averaged VEPs to 60-Hz flicker, and were generated by the brain's so-called "high-frequency" subsystem. Note the opposite polarity of the top and bottom VEP traces, and compare with the italicized potentials in Figure 1.12A. The traces on the far right are averaged VEPs to 5-Hz flicker, and were generated by the brain's so-called "low-frequency" subsystem. Note their comparatively large amplitudes and constant polarity, and compare with the boldface potentials in Figure 1.12A. Recordings between inion and vertex would emphasize the "high-frequency" subsystem and reject the "low-frequency" subsystem (see Fig 1.12A). These recordings were made in Holland, where the mains frequency is 50 Hz. (From Spekreijse H: *Analysis of EEG Responses in Man*, Ph.D. thesis, The Hague, Junk Publishers, 1966. Reproduced by permission.)

bution of a stronger source that gave closely similar voltages at the two electrodes. Potentials resulting from the stronger source are shown in boldface, and potentials resulting from the weaker source are shown in italic. All potentials are with respect to a distant (ear) reference. Figure 1.12B gives a practical illustration. The net potential at A and B with respect to the ear is the algebraic sum of the potentials resulting from the two sources (+9

and +7 μV , respectively). Only the *difference* (2 μV) reflects the presence of the weaker source, and this would be difficult to quantify with precision if it were measured as the difference of two large quantities (because constant percentage errors of $\pm x\%$ is measurement mean large absolute errors when large quantities are measured). On the other hand, when the potential difference between A and B is measured directly between electrodes

located at A and B, then the absolute error in the difference is reduced to $\pm x\%$ of $2 \mu\text{V}$. This is achieved by virtue of the differential amplifier's common-mode rejection.

1.2.5 The Reference Electrode: A Difficult Choice

There are several alternative views as to the best choice of reference site including cephalic electrodes (nose, ear, linked ears) and extracephalic electrode (e.g., Erb's point). It is known that the ear is not always indifferent, at least for VEPs, to pattern reversal⁹²⁸ (Fig 2.169).

A linked ear reference has several drawbacks. Connecting the ears does not average the preexisting potentials of the two ears, but rather modifies the current flow (and thus the potential distribution) over the whole scalp (Figs 1.13A and B). Potential gradients between the ears are diminished, and in some cases hemispheric asymmetries may be less evident [1748, p 191]. A variant on linked ears is the variable voltage divider shown in Figure 1.13C. If the resistance R is large compared with scalp resistance, then the current through R will be sufficiently small that it does not distort the potential distribution over the scalp. Artifact from the heart's electrical activity can be eliminated by adjusting the potentiometer, but the reference potential thus chosen is unlikely to be optimal for all electrode locations [1748, p 193].

The "average reference electrode"¹⁷⁶⁰ has been advocated in both visual¹²²⁴ and auditory²¹⁹⁵ recording and has the advantage that it does not favor any particular electrode site.¹³⁸³ For spontaneous EEG recording, Offner¹⁷⁶⁰ justified its use by the assumption that the average reference signal is generated by a large number of dipoles that are randomly oriented within the brain, so

that the mean voltage from many electrodes over the entire scalp is near zero and constant. However, this assumption does not hold for EPs. Bertrand et al¹⁷³ stated that the average reference should be (ideally) derived from an array of electrodes spaced regularly around the entire head rather than merely over the scalp. Recognizing this as impracticable, Bertrand et al¹⁷³ stated that the situation can be improved by computing the average reference by means of an appropriate numerical surface integration formula rather than merely taking the mean potential over all electrodes, and that this process can be aided by including a pharyngeal electrode. Nevertheless, a problem with the average reference is that the potential recorded from any given electrode depends on where *all* the other electrodes are located. In practice, whether or not the experimenter gains physiological insight by connecting each amplifier to closely spaced electrodes rather than using a common reference depends very much on the spatial organization of the neuronal EP generators.

In his discussion of the reference electrode question, Nunez gives the following illustration of a major pitfall [1748, pp 178–182]. Auditory evoked potentials were recorded at many scalp locations with respect to a reference electrode on the nose (Fig 1.14A). It was suggested that the generator location for the N1–P1 potential should be immediately underneath the scalp region where EP polarity changed sign, that is, the heavy zero-potential line in Figure 1.14A. This seemed reasonable because the polarity reversal occurred directly over auditory cortex. But Nunez points out that, although this conclusion may have been correct, the line of reasoning was not. For example, if a forehead instead of a nose reference had been chosen, EP polarity would reverse

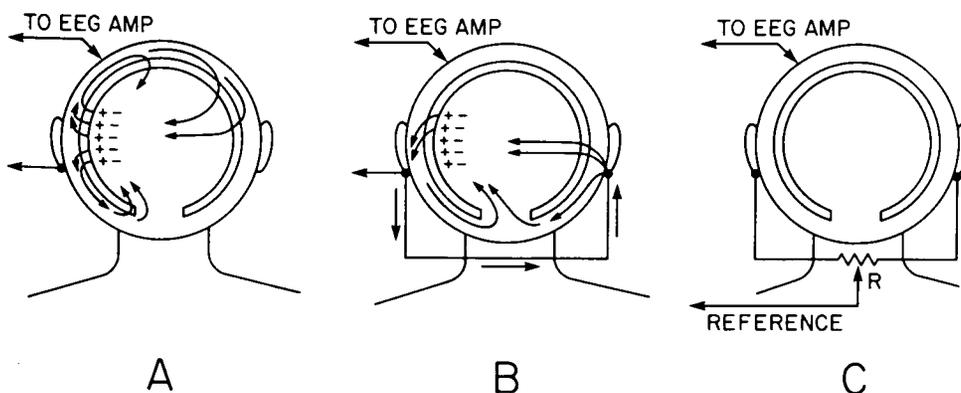


Figure 1.13

(A) Schematic view of a possible generator configuration and the resulting volume-conducted currents. (B) Connecting the ears may change the potential distribution by providing a low-resistance path between the ears. (C) A voltage divider reference. (From Nunez P: *Electric Fields of the Brain: The Neurophysics of EEG*. New York, Oxford University Press, 1981. Reproduced by permission.)

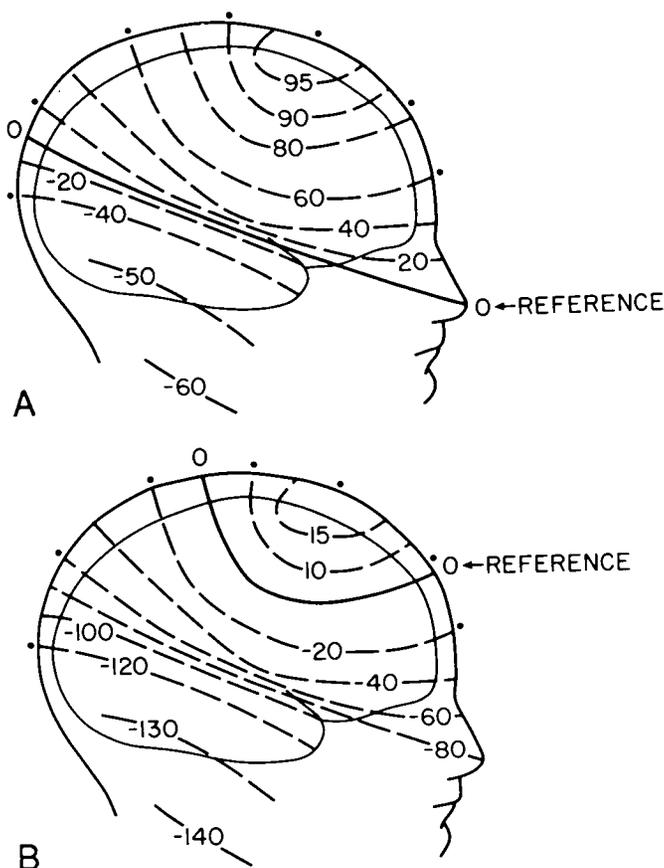


Figure 1.14
Polarity Reversal Does Not Necessarily Indicate an EP Source
 Panels A and B illustrate that the same EP distribution might be taken to indicate different source locations depending on the choice of reference electrode. (From Nunez P: *Electric Fields of the Brain: The Neurophysics of EEG*. New York, Oxford University Press, 1981. Reproduced by permission.)

across the zero-potential line in Figure 1.14B, thus indicating a quite different generator site. Again, if the reference electrode had been placed on the neck, there would be no polarity inversion over the entire head, so that the same line of reasoning would indicate that the EP generator site was not in the head at all! The logical error is that polarity reversal is, in this example, not a valid indication of source location, as it is not *reference invariant* [1748, p 180]. Section 2.3.10 reviews recent evidence that there are multiple contributors to the N1 wave rather than a single source.

1.2.6 Near Fields and Far Fields

The unhappy introduction of the terms “near field” and “far field” into the evoked potential literature has led to some confusion and misunderstanding.

In electromagnetic theory the near field/far field distinction has a clear meaning that can be understood as

follows.²⁰⁰ Consider an electric dipole in a vacuum. (An electric dipole is a pair of equal-and-opposite charges, $+q$ and $-q$, separated by a very small distance d . The strength of the dipole is defined as qd .) An electric dipole generates an electric field in the surrounding space whose magnitude (E) at any point in the surrounding space falls off with the cube of distance (r), provided that r is large compared with d . A static dipole produces no magnetic field, but suppose that the dipole strength oscillates at frequency f Hz—one might visualize the $+q$ and $-q$ charges exchanging places repetitively. This oscillation of dipole strength causes the electric field strength E to oscillate at frequency f Hz and, in addition, an oscillating magnetic field is created in the surrounding space. The oscillating dipole will emit energy that will be carried away in the form of an electromagnetic disturbance traveling at the speed of light c . Let us consider two regions of space: (1) the *near-field* region, defined as being the region where the distance from the dipole is much less than the wavelength of the electromagnetic disturbance ($\lambda = c/f$); (2) the *far field* region, defined as being the region where the distance from the dipole is much greater than λ . In brief, in the near-field region the magnitude of the electric vector \vec{E} falls off with r^3 and the magnitude of the magnetic vector \vec{B} falls off with r^2 . In the far-field region, the magnitudes of both \vec{E} and \vec{B} are interlocked and fall off with r .

The crucial point is that the strength of the electromagnetic *far field* falls off very slowly with distance. This allows appreciable electromagnetic energy to be transported over large distances. Far-field energy propagation is the way in which energy from the sun travels through space to reach our planet. It is the basis of radio and TV communication.

When frequency f is near zero, the far field disappears entirely, as does the near-field \vec{B} vector, leaving only the near-field \vec{E} vector with which we started this discussion. *The far field is essentially a high-frequency phenomenon, and is irrelevant to the recording of EEG or evoked potentials.*

It seems that the original reason for introducing the near field/far field dichotomy in the EP literature was to distinguish between scalp-recorded fields resulting from sources in cortex and sources in midbrain. For example, because the auditory brainstem response is less affected by electrode position on the scalp than are cortical EPs, it was thought to be analogous to a genuine far field. But, in point of fact, the scalp-recorded EPs generated anywhere in the brain are all near fields.

1.2.7 Source Density Analysis

The chief aim of many EP studies is to identify and locate the generator sources of different kinds of EP and EP components. The usual way of going about this is to

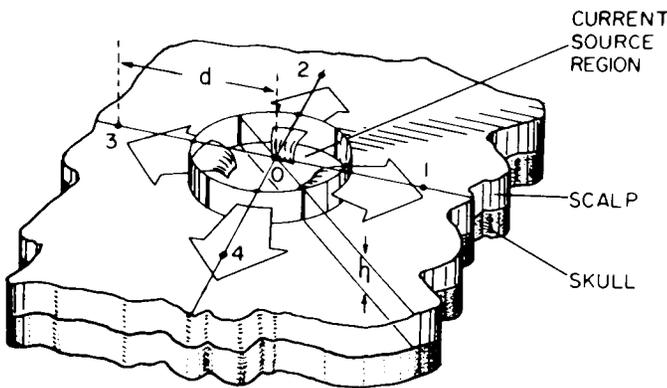


Figure 1.15
The Laplacean

Current flows through the skull into the infinitesimally small, pillbox-shaped volume of scalp, and exits through the sides of the pillbox. By measuring the current exiting the pillbox parallel to the skull, an estimate can be made of the current flow normal to the skull into the pillbox. (From Nunez P: *Electric Fields of the Brain: The Neurophysics of EEG*. New York, Oxford University Press, 1981. Reproduced by permission.)

analyze the data from many electrode sites with respect to a reference. Unfortunately, as noted above, the interpretation of such multichannel records can be complicated by the unknown contribution of activity at the reference site to the measured waveforms. For this reason, source density analysis has attracted attention as an alternative technique because it is reference free.

Source density analysis is a method of deriving the distribution of sources from the distribution of potential. The method has been used to analyze single-unit activity in the cerebellum¹⁷³¹ and cerebral cortex.¹⁶³⁰ A limitation to this technique in depth recording is that the use of one-dimensional electrodes assumes that the brain's properties do not vary rapidly at right angles to the electrode, and full three-dimensional arrays tend to cause significant damage to brain tissue.

Source density analysis has also been applied to EEG analysis,^{1062,1063,1684,2620} and more recently in evoked potential research.^{1470,1472-1474,2389,2506} The basic mathematics is reviewed by Nunez [1748, pp 196-203].

Poisson showed that the source density of a potential field distribution is proportional to $\nabla^2 V$, where V is the potential at any point in the field and ∇^2 is the Laplacean operator. In words, the source density is proportional to the sum of partial second derivatives of the potential because $\nabla^2 V$ is defined as equal to

$$\frac{\partial^2 V}{\partial x^2} + \frac{\partial^2 V}{\partial y^2} + \frac{\partial^2 V}{\partial z^2}$$

The Laplacean can be intuitively understood as an operator that subtracts from the potential at any given

point the average potential in its neighborhood.¹⁴⁷⁴ The Laplacean derivative of the EP, $\nabla^2 V$, is reference free because it is based on differences of voltage. It is a measure of the component of electric current flowing normal to the surface of the scalp.²³⁸⁸ This can be understood by reference to Figure 1.15. Electric current from generators within the brain is injected outward through the skull at some locations and returns through other locations (not shown). In Figure 1.15 all the current entering the bottom of the imaginary pillbox must leave through the sides of the pillbox. For the two-dimensional scalp distribution of the EEG potential, the Laplacean can be approximated by recording from a rectangular array of equally spaced electrodes, and taking the difference between the potential at each electrode and the average potential of its nearest four neighbors.¹⁴⁷⁴ (Hjorth's¹⁰⁶² method was lengthier. Rather than considering the four nearest neighbors only, he subtracted a weighted average of all other electrodes from the signal derived from any one electrode.) Unfortunately, this rectangular-array technique requires a forbiddingly large number of electrodes: a single sample requires 5 electrodes, two neighboring samples require 8 electrodes, a square of four requires 12 electrodes, and so on (Fig 1.16A).

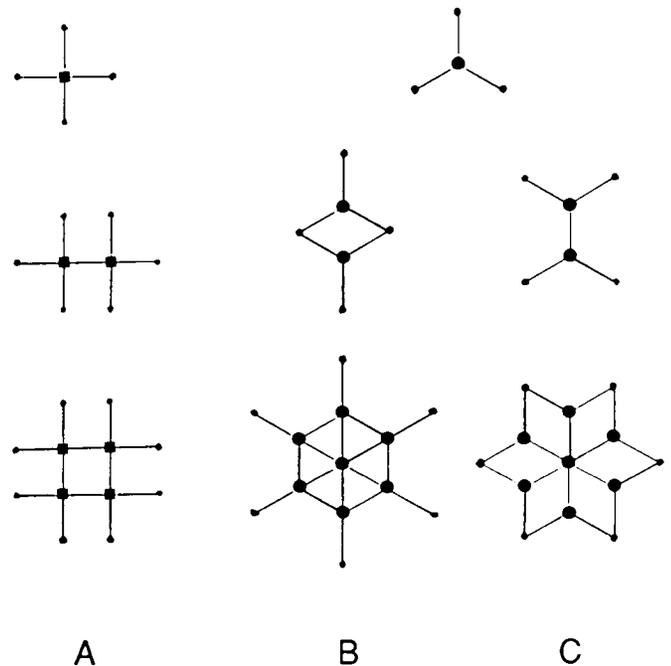


Figure 1.16
Source Density Analysis: Electrode Arrays
(A) Electrode requirements for "Laplacean" computations on a rectangular grid. (B, C) Reduced requirements using a triangular grid and a hexagonal array. (From MacKay DM: Source density analysis of scalp potentials during evaluated action. I. Coronal distribution. *Exp Brain Res* 1984;54:73-85. Reproduced by permission.)

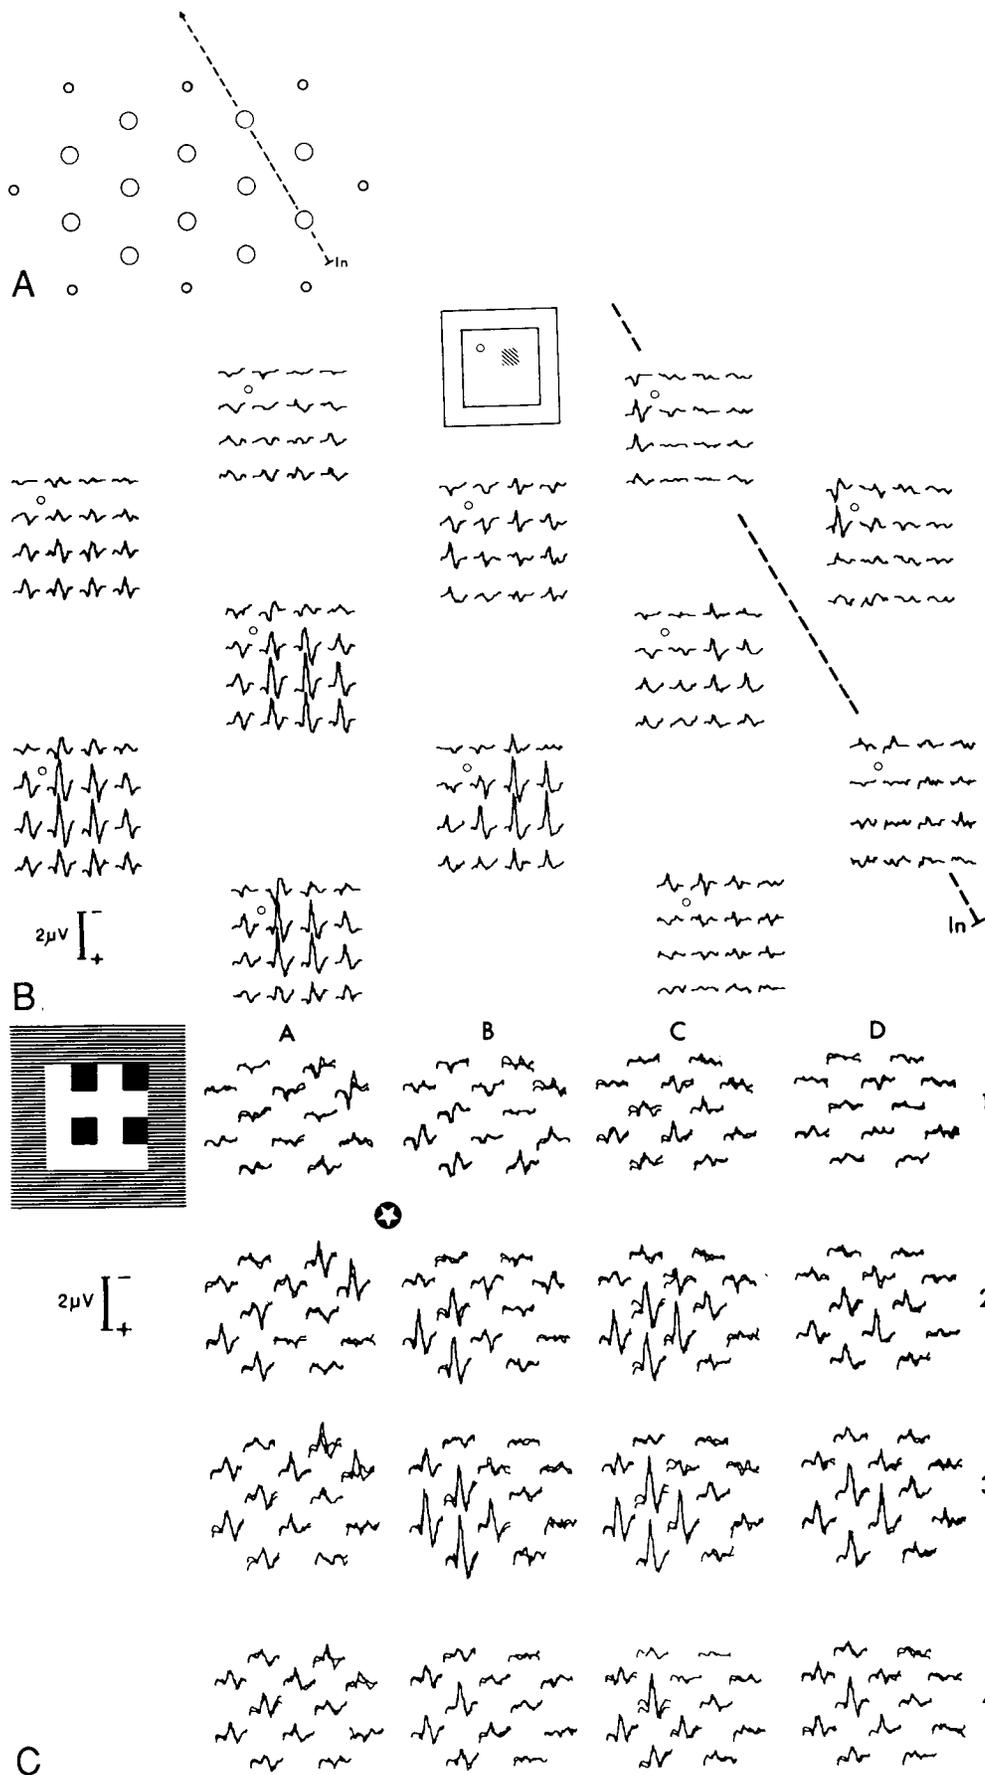


Figure 1.17
Source Density Mapping of Visual Receptive Fields
 (A) Electrode array used to derive current source densities. The dotted line indicates the midline, the arrow pointing from the inion (In.) toward C_z . The electrode spacing was 1.5 cm. (B) Receptive field plots for the 12 scalp locations represented by larger circles in panel A. The fixation point is indicated by a circle and the $0.5^\circ \times 0.5^\circ$ scanning stimulus is shown in the inset. The hatched area represents the $0.5^\circ \times 0.5^\circ$ scanning stimulus presented for 30 msec in one of the 4×4 locations used. (C) Two superimposed sets of scalp field plots for the 16 retinal locations. (From MacKay DM: Source density analysis of scalp potentials during evaluated action. I. Coronal distribution. *Exp Brain Res* 1984;54:73-85. Reproduced by permission.)

MacKay's¹⁴⁷⁰ modified technique is considerably more practical: substantially fewer electrodes are required, and the approximation to the ideal Laplacean is only slightly cruder. MacKay's basic idea is to take the average of the potentials at three rather than four points around each electrode, and to use an array on a triangular rather than a rectangular lattice. Using this system, a single sample requires 4 electrodes, two neighboring samples require 6 electrodes, and as many as seven samples can be derived from only 13 electrodes in a hexagonal array. Figures 1.16B and C illustrate two forms of such an array.

MacKay gives two reasons to prefer a Laplacean derivation rather than a monopolar or bipolar derivation in studies of EP source location: (1) Especially with slow transient evoked potentials, visual inspection does not easily detect small changes in mean level or slope of waveforms containing the large irregular components of higher frequencies that are often widely distributed over the scalp. (2) The Laplacean derivation is differentially sensitive to sources of potential that lie within its "triangle of reference," while being insensitive to sources outside this triangle. (This is because potential gradients caused by remote sources usually have the same polarity and a similar magnitude on both sides, and thus will tend to cancel out when the Laplacean is computed.) The gratifying consequence is that clean waveforms can be recorded after only a few sweeps.

Figure 1.17 illustrates the application of this technique to mapping human visual receptive fields near the fovea. Electrodes were placed only 1.5 cm apart as illustrated in Figure 1.17A, and the total visual field area scanned was $2^\circ \times 2^\circ$. The stimulus was a pattern subtending $0.5^\circ \times 0.5^\circ$ that could adopt any one of 16 (4×4) positions within the $2^\circ \times 2^\circ$ area. Subjects fixated a black dot at the upper left of the $2^\circ \times 2^\circ$ area, and the $0.5^\circ \times 0.5^\circ$ target was presented at different locations in pseudorandom order. Results were plotted in two alternative formats. The first format was an array of 12 "visual field maps," each comprising a montage of 4×4 averaged EPs, the 12 waveforms being laid out to correspond with the scalp locations to which they related. The alternative format was an array of 16 "scalp maps" laid out in a 4×4 array corresponding with the retinal areas stimulated. Each comprised a montage of 12 averaged EPs set out to correspond with the 12 scalp electrode locations, thus showing the scalp distribution of source densities elicited by stimulating one $0.5^\circ \times 0.5^\circ$ area of retina.

Figure 1.17B shows that under source density analysis the retinal receptive fields of points only 1.5 cm apart on the scalp are not only distinguishable, but in some cases radically different. For locations with well-defined receptive fields, the field diameters were only 1° to 2° . The resolution of MacKay's data was sufficiently high to sug-

gest that, at any given location, receptive field diameter may be different for the three main peaks of the pattern EP. The ability of source density analysis to reveal such "fine grain" is particularly interesting in view of the finding of substantial differences in VEPs from closely adjacent sites in monkey brain.⁵⁰¹ The source density distributions in Fig 1.17C are approximately consistent with the known anatomy of the visual projections in the human (though one should bear in mind the considerable intersubject variability of human visual cortex [Fig 2.168]).

1.2.8 Evoked Potential Sources and the Equivalent Dipole Approximation

Is It Possible to Infer the Sources of Scalp Evoked Potentials from Scalp Distribution Alone?

Given the distribution of EP amplitude over the scalp, there are two methods for inferring the location(s) of the responsible intracranial generators. The "forward" method is to make an informed guess as to the electrical characteristics, location, and geometry of the generator(s), and then calculate the scalp distribution. This calculation is relatively straightforward.^{1748,2439a,2589,2707} Alternative source geometries are then compared with the experimentally measured EP topography. The other method, the so-called "inverse" method, is to calculate the location(s) of the intracranial generator(s) directly from the observed EP scalp distribution. Unfortunately, as was already shown by Helmholtz in 1853, there is a major problem with this otherwise attractive method: *any given surface field can be generated by more than one source configuration, so that unique solutions are in general unobtainable.*¹⁰⁰⁵ Therefore the forward method is more useful.

The most popular implementation of the forward method has been in terms of the so-called "equivalent dipole." But before discussing this approach we must digress.

Holes in the Skull and the Distribution of Volume Currents

When attempting to interpret recordings from scalp electrodes the experimenter must confront a knotty problem: What fraction of the recorded signal comes from sources close to the electrodes as opposed to the contributions of more remote sources? The question is lucidly reviewed by Nunez.¹⁷⁴⁸

A popular approach is to use the so-called reciprocity theorem, borrowed from physics and applied to biophysics by Helmholtz [1748, p 182]. The basic idea is that a knowledge of the current density (or electric field) throughout a volume conductor caused by an injection of current between two stimulating electrodes completely specifies how these same electrodes, when serving

as recording electrodes, will record signals produced by dipole sources at any place in the volume conductors [1748, pp 182–188]. For example, the “three-concentric-sphere” model assumes the brain to be equivalent to an isotropic homogeneous sphere that is surrounded by a homogeneous, isotropic spherical skull, and covered by the scalp, the resistivity of the skull being 80 times greater than that of the scalp or of the brain.¹⁷⁴⁸ Electrode sensitivity for different dipole locations can be calculated by means of the reciprocity theorem.

Nunez points out that a limitation of the three-concentric-sphere model of the head is that real skulls contain several holes, and areas of thinning (Fig 1.18). The ratio between skull circumference and thickness is roughly the same as the ratio of skull to brain resistivity (80:1). This coincidence means that long current paths through skull openings and around the scalp have comparable resistances to the more direct paths through the thickness of the skull (Fig 1.18). For this reason alone, any given electrode may receive significant contributions from distant as well as nearby sources.

Dipoles, Quadrupoles, and Multipoles

In the general case, for any combination of sources and sinks, and provided that r is sufficiently large, the potential (V) at distance r is given by the multipole expansion

$$V(r) = (\text{dipole term}/r^2) + (\text{quadrupole term}/r^3) \\ + (\text{octupole term}/r^4) + \dots$$

The Single Equivalent

Dipole: A Valid Approximation?

Section 1.2.4 noted that an electric charge dipole comprises two closely adjacent point charges, $+q$ and $-q$. In the EP context, a single pyramidal cortical cell viewed from a scalp electrode can closely approximate a dipole. Even with sources of finite size, however, a dipole approximation may be valid, providing that we are dealing with (1) a small region of active cortex containing (2) parallel neurons whose activities are (3) synchronous and coherent. But in EP recording there are physiologically plausible source configurations that by no means satisfy these requirements. Consider, for example, a large-scale source comprising many parallel dipoles whose activities are highly correlated, that is, a “deterministic dipole layer.” Whereas the potential falls off with $1/r^2$ for a single dipole, there is essentially no falloff with distance for a dipole layer until the vertical distance becomes comparable to the linear size of the dipole layer. Consequently, the potential at the scalp would be determined by the electrical conductivity of the intervening brain, skull, and skin rather than the distance of the dipole sheet.¹⁷⁵³ A salutary point here is that skull thickness varies about threefold over its surface,²⁵²² quite

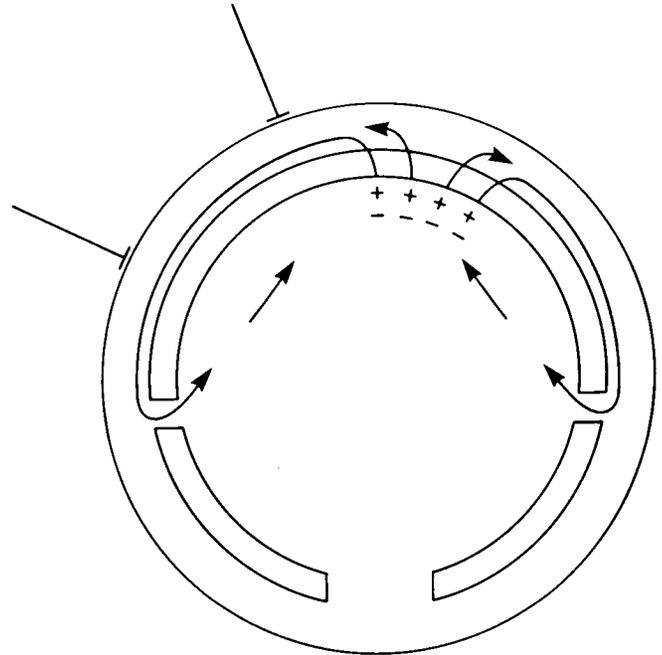


Figure 1.18
Holes in the Skull Create Difficulties in Localizing EP Sources
The resistivity of the skull is roughly 80 times that of scalp or brain tissue. Consequently, long current paths through skull openings are likely to create potentials at scalp sites distant from sources. (From Nunez P: *Electric Fields of the Brain: The Neurophysics of EEG*. New York, Oxford University Press, 1981. Reproduced by permission.)

apart from the naturally occurring holes in the skull (Fig 1.18).

The validity of the single equivalent dipole description is called into question when there are several distinct sources.^{956,1752} In such cases the equivalent dipole may not even be located within the region of active tissue! The possibility of multiple sources must be taken seriously when we consider the evidence that, in primate brain, there are multiple visual (Fig 2.84), auditory (Fig 2.53), and somatosensory (Figs 2.30, 2.70) areas that have quite different retinotopic, tonotopic, or somatotopic projections. Furthermore, there is ample evidence that any given cortical region is connected to other cortical regions many centimeters distant (Fig 2.3). These facts suggest that scalp EPs commonly represent the activities of multiple discrete sources, some of which lie considerable distances apart. Even though each individual source might be described in terms of a single equivalent dipole, a single equivalent dipole is unlikely to provide a fair approximation of the resultant activities of the multiple simultaneously active sources. A final caveat is that the relative activities of multiple visual (or auditory or somatosensory) areas may well change during (say) the first 500 msec after stimulation, and the resulting shifts of

location, orientation, and strength of a single equivalent dipole may well be difficult to interpret in physiological terms.

1.2.9 Amplifiers and Amplifier Bandpass

Amplifiers with a differential (i.e., “push–pull”) input stage⁸¹⁶ are used for EP recording because of their ability to reject “common-mode” noise signals, that is, noise common to the two input leads. The “common-mode rejection ratio” specifies the amplifier’s ability to discriminate between (a) signals in antiphase at the two inputs and (b) signals that are in phase at the two inputs, that is, common mode signals. Thus, an amplifier with a rejection ratio of 100,000:1 (100 dB) can have a gain that is 100,000 times greater for antiphase signals than for inphase signals (when the common-mode rejection control is carefully trimmed). High common-mode rejection is essential because, in most recording situations, common-mode mains interference is much larger than the brain signals that one wishes to record, so an amplifier should be selected whose technical description specifies a high rejection ratio over the frequency range of interest. But this is not enough. In practice, the maximum obtainable ratio can be limited by the imbalance of the electrode impedances. The impedances to ground (Z_1 and Z_2 Ω , respectively) of the amplifier’s two input electrodes will, in general, be different. Suppose that the amplifier’s input impedance is Z_{IN} Ω . Then the maximum obtainable rejection ratio is $2Z_{IN}/|Z_1 - Z_2|$. For example, if $|Z_1 - Z_2| = 2$ k Ω (a typical figure), and $Z_{IN} = 10$ M Ω , the maximum obtainable ratio will be 10,000 (i.e., 80 dB). As a practical illustration, Grass Instruments Company specifies the rejection ratio of its widely used Model P511 amplifier as adjustable to 10,000 with an input impedance of 44 M Ω at 100 Hz. In practice, a high rejection ratio is attained by combining three factors: (a) high rejection ratio specification, (b) high input impedance, (c) minimal difference $|Z_1 - Z_2|$ between electrode impedances, and (d) careful adjustment of the amplifier’s common-mode rejection to optimal. This last step is essential. Condition (c) is effectively met by careful electrode application to ensure that electrode impedances are small, so that the difference $|Z_1 - Z_2|$ cannot be large.

A second reason for reducing electrode resistance is that the lower the resistance, the lower is the noise voltage at the amplifier inputs produced by capacitive (electrostatic) coupling between the amplifier and nearby AC voltage sources, such as mains-powered lights and mains cables. Several reviews of amplifier design are available including Cooper, Ossleton, and Shaw’s book.⁴⁴⁸

Underrating the importance of an EEG amplifier’s

two bandpass controls is an error that can result in much wasted effort. These two apparently insignificant knobs share with the scalp–electrode interface the capability of wasting the sophistication and expense of subsequent data processing. It is a temptation to use a severe high-frequency cutoff because the resulting EP looks satisfyingly clean and noise free (compare Figs 1.22A and D), but this achievement resembles the ostrich’s sense of security as it stands with its head in the sand. In practice it is best to strive not for overly smooth waveforms, but rather for faithful reproduction of the EP signal and to demonstrate that the important features are reproducible. An “untouched” EP recorded with adequate bandwidth should usually have visible high-frequency noise unless the number of replications is very large indeed; a little high-frequency noise is a healthy feature.

Amplifier noise over a bandwidth of (say) 1–100 Hz is usually less than about 3 μ V peak-to-peak and this is not obtrusive in an averaged EP where most “noise” is of biological origin.

It is a worthwhile exercise to calibrate the bandpass of one’s VEP recording system rather than taking on trust the markings on the low-pass and high-pass filter controls, or for that matter the numbers displayed by the microprocessor. Bear in mind that the pen recorder that writes out the VEP waveform may introduce extra high-frequency filtering if the pen speed is too high. Especially if one plans to record steady-state EPs, it is essential to measure the recording equipment’s phase-versus-frequency calibration. Constant-amplitude sine waves of about 5- to 20- μ V amplitude should be fed into the EEG amplifier, averaged, and written out with the same procedure used for VEP recording. (Some amplifiers will give very distorted outputs unless noise is added to the input sine wave.) Calibration results may be surprising, even with commercial equipment. For example, the author plotted a calibration curve for a popular EP amplifier with the filter controls turned to settings clearly marked as “DC” and “30 Hz.” Although the 30-Hz setting was approximately correct, the so-called “DC” setting was quite misleading. In spite of the label, the equipment was not DC-coupled; the low-frequency 3-dB point was about 0.6 Hz, and the effect of this on broad EP peaks is by no means negligible (see below). In some amplifiers the filter settings specify the frequency for which gain is -3 dB (70.7%) of the maximum gain in the flat part of the bandpass, whereas other amplifiers specify the 50% rather than 70.7% point. Also, the rolloff rate varies between amplifiers. For example, Goff⁸³⁸ compared two popular EEG amplifiers and found that, with a 1.0-Hz setting, the frequency response was flat down to 5 Hz for one but only down to 10 Hz on the other.

Some amplifiers provide a 60-Hz “notch” filter for

eliminating mains interference. In EP recording this should be used only as a last resort and preferably not at all. A notch filter affects phase at frequencies far from the notch and this, combined with its effects on amplitude, can distort the EP.

Figure 1.2 shows the effects of low-pass, high-pass, and bandpass filtering on sine wave input signals. This figure relates more directly to steady-state VEPs than to transient VEPs. The effects of filtering on transient VEPs is brought out in Figures 1.19–1.21. To compare the effect of filtering on sharp (i.e., short-lived) versus broad VEP components, a short (20-msec) and a long (200-msec) positive pulse were fed into the filters; there was a 50-msec gap between the two pulses. Both short and long input pulses are reproduced tolerably faithfully with filter bandpass set at DC–250 Hz (Figs 1.19A, 1.20A, 1.21A). Note that the corners are sharp, that the tops are flat, and that the trace never falls below the baseline of 0 V.

The effect of high-frequency attenuation is shown in Figure 1.19. The sharp corners of the pulses are appreciably rounded by the slight increase in high-frequency attenuation caused by lowering the 3-dB point from 250 Hz (A) to 100 Hz (B). The corners become progressively more rounded as the severity of the high-frequency attenuation is progressively increased (A–D). Note that the short pulse is more severely distorted by high-frequency attenuation than is the long pulse: a 15-Hz setting reduces the peak amplitude of the short pulse as well as rounds the corners, while the peak amplitude of the long pulse is not appreciably reduced (and would not be until high-frequency attenuation were made considerably more severe than in Fig 1.19D).

The effect of low-frequency attenuation is illustrated in Figure 1.20. The two input pulses are reproduced fairly faithfully in (A) with a filter bandpass of DC–250 Hz. Introducing a slight low-frequency attenuation in (B) causes appreciable distortion because DC levels cannot be indefinitely maintained without DC coupling. The distortion consists of a sag during the duration of the pulse and a transient negative excursion at the termination of the pulse. These two effects grow progressively more evident as the severity of the low-frequency attenuation is progressively increased (Figs 1.20A–C). With severe low-frequency attenuation, the long pulse no longer resembles a positive rectangle: only the onset and the offset are signaled; during the middle of the pulse there is virtually no output at all (Fig 1.20C).

Thus, the distortion caused by low-frequency attenuation is quite different from the distortion caused by high-frequency attenuation which, as illustrated in Fig 1.19, does not cause any sag or negative excursions at all, but instead rounds the corners and reduces the maximum deflection. Furthermore, low-frequency attenuation dis-

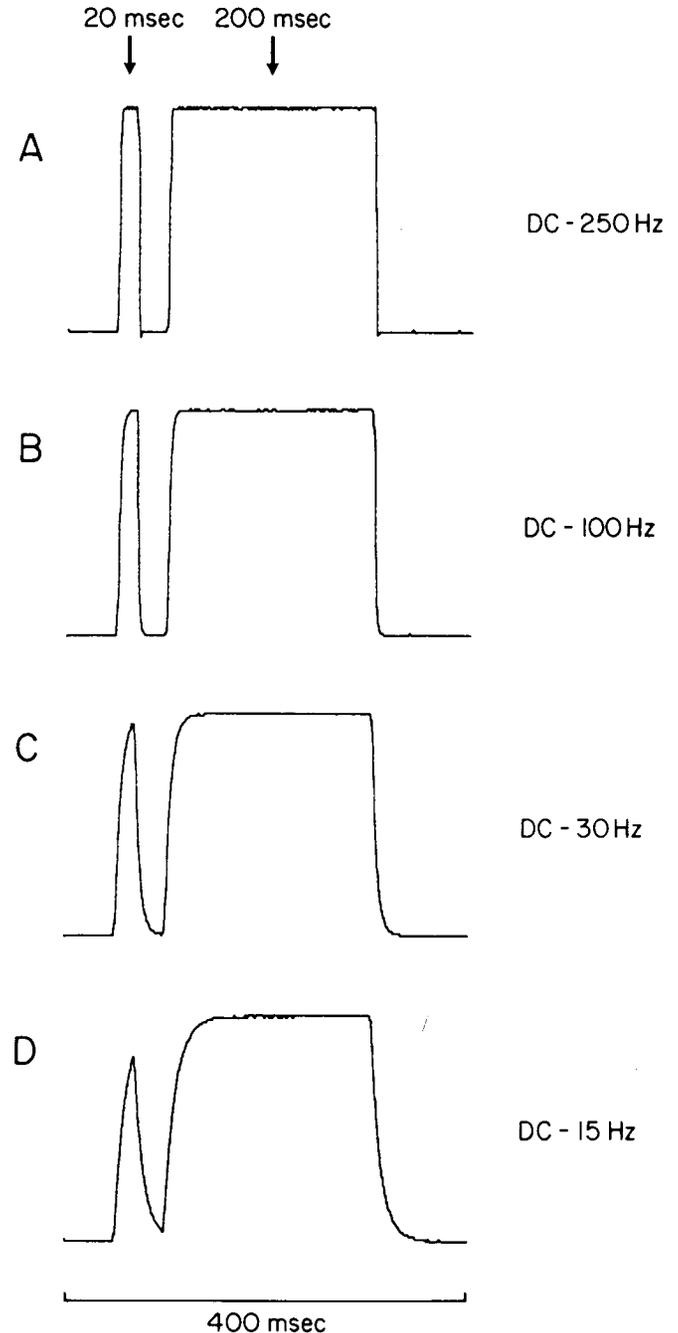


Figure 1.19
Effects of High-Frequency Attenuation on Short and Long Pulses

Both short and long pulses are faithfully reproduced with an amplifier bandpass of DC–250 Hz (A), but as the high-frequency 3-dB point is progressively lowered (A–D), the corners become more and more rounded. Distortion is greater for the short pulse than for the long pulse; setting the upper 3-dB frequency at 15 Hz reduces the peak amplitude of the short pulse as well as rounds the edges (D). The low-pass filters used were similar to those used for Figures 1.2A and B.

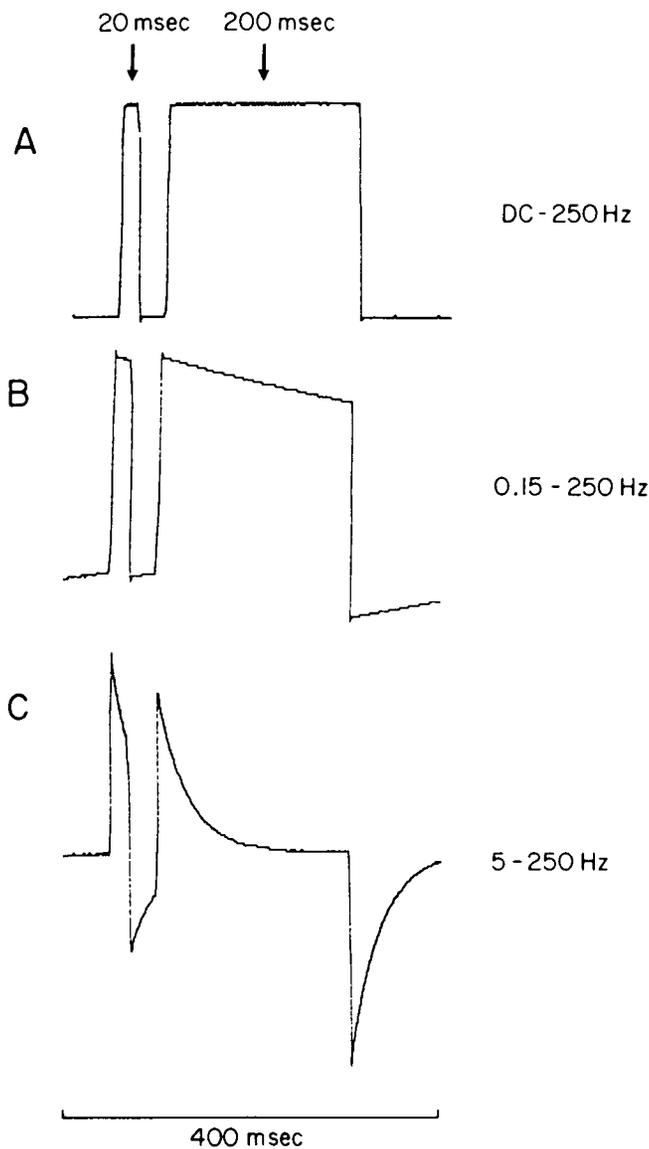


Figure 1.20
Effects of Low-Frequency Attenuation on Short and Long Pulses

Both short and long pulses are faithfully reproduced with an amplifier bandwidth of DC–250 Hz (A), but as the lower 3-dB frequency is progressively raised (A–C), the pulses sag more and more. Distortion is greater for the long pulse than for the short pulse; the trace almost returns to baseline during the pulse, and pulse offset produces a sharp artifactual negative deflection. The high-pass filters used were similar to that used in Figures 1.2C and D.

torts the long pulse more than the short pulse (Fig 1.20C), whereas high-frequency attenuation distorts the short pulse more than the long pulse (Fig 1.19D).

Bandpass filtering combines low-frequency and high-frequency attenuation. The effect of severe bandpass filtering is illustrated in Figure 1.21. Corner rounding, sag, and negative deflections are all evident, and both short

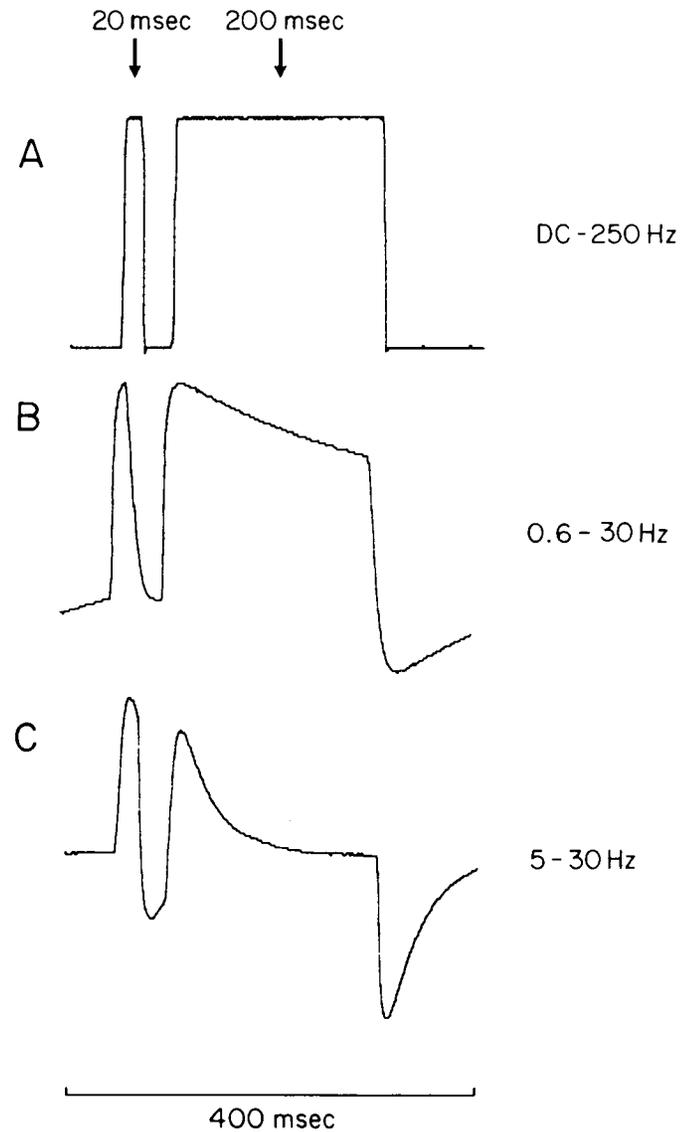


Figure 1.21
Effects of Bandpass Filtering on Short and Long Pulses
The 20- and 200-msec unidirectional pulses shown in panel A are both severely distorted by a 5- to 30-Hz bandpass filter (C). The bandpass filter used was similar to that used in Figures 1.2E and F.

and long rectangular positive pulses (A) are severely distorted (C).

Although the filter settings in Figure 1.21C were considerably more severe than would be used in practical VEP recording, it should be borne in mind that some degree of rounding, sag, and spurious negative deflections are always present with the kind of bandpass filters commonly used for EP recording. Exaggeratedly severe filtering was used in Figure 1.21 to emphasize the point.

So far we have discussed the effect of bandpass filter settings on sine waves and pulse calibration signals. Now we turn to the effect of filter settings on real EPs.

Figure 1.22 illustrates how incorrect filter settings can

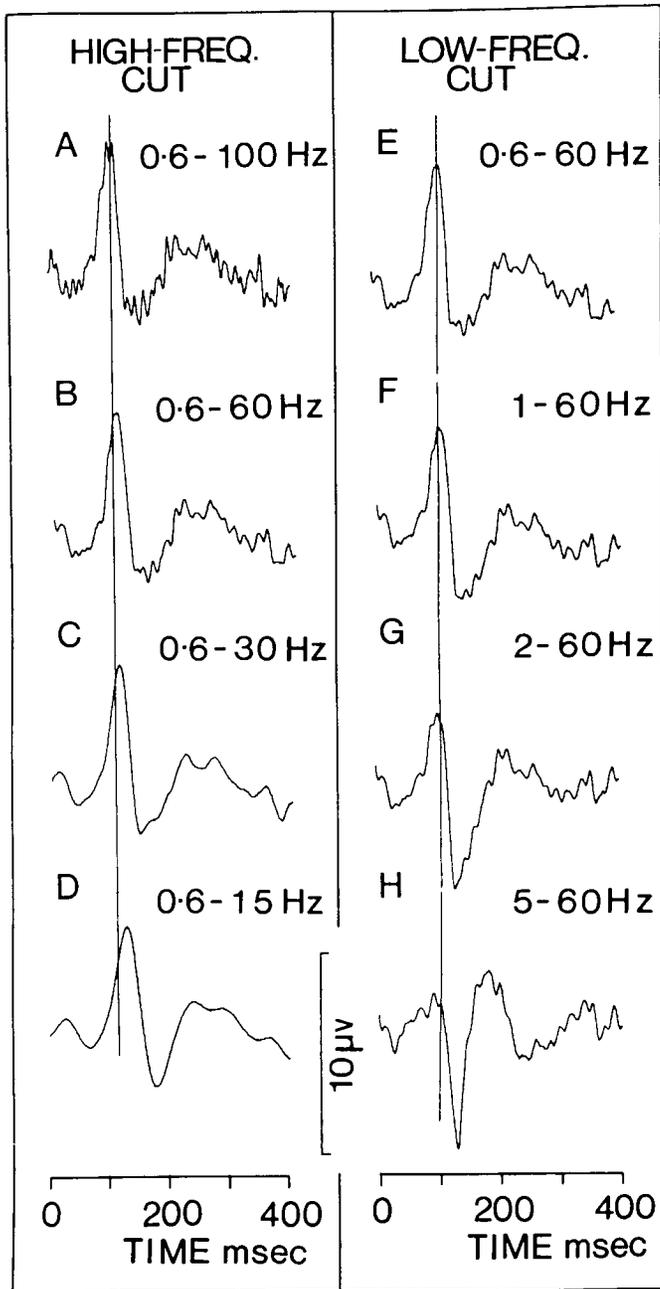


Figure 1.22
Effects of Conventional Analog Filtering on Transient EPs
 (A–D) The effects of progressively increasing the severity of high-frequency cut. The vertical line brings out the spurious delay caused by high-frequency filtering. (E–H) The effects of progressively increasing the severity of low-frequency cut. Compare the waveforms in panels E and H. All the traces were obtained from the same EEG sample. This figure should be compared with Figures 1.19 and 1.20. The stimulus was conventional clinical pattern reversal, 50 min arc checks of near-100% contrast. The field subtended 25°, and the EEG was recorded betweeninion and right ear with vertex grounded.

produce spurious EP delay and also distort the EP waveform. (Note that all the traces in Figure 1.22 were obtained from the same 1-min EEG recording.) First we consider the effects of high-frequency attenuation. Figure 1.22A was recorded with a bandpass of 0.6–100 Hz. The chief feature is a positive deflection at 107 msec. The 107-msec peak is preceded and followed by smaller negative deflections, and there is a broad slow positive deflection peaking at roughly 240 msec. In Figure 1.22A it is not clear whether the double peak at 102 and 112 msec and the shoulder at 75 msec are part of the EP or whether they result from the substantial high-frequency noise. Reducing the upper 3-dB point to 60 Hz (Fig 1.22B) attenuates the high-frequency noise and produces a cleaner-looking EP, but also removes the double peak. This 0.6–60 Hz 100-trial VEP compares closely with the 1,000-trial average of Figure 1.45D. Increasing the severity of the upper cut still further (30 Hz, Fig 1.22C) removes the high-frequency noise and produces an even cleaner-looking EP. On the other hand, fine details such as the shoulder and double peak are now lost and, more seriously, an appreciable spurious delay has been created by the too-severe filtering (the 107-msec peak is delayed to 117 msec). An even more severe high-frequency cut (15 Hz, Fig 1.22D), further delays the 107-msec peak to 133 msec. Compare this 100-trial record in Figure 1.22D with the 1,000-trial average in Figure 1.45D, where noise was reduced by averaging rather than filtering, demonstrating that the 75-msec shoulder is a genuine EP feature and that the positive peak has a latency of 107 msec.

The optimum bandpass depends on the particular kind of EP being recorded, and quite different bandpass settings are needed for contingent negative variation (CNV), VEP, and auditory brainstem response, for example.

Now we consider the effects of low-frequency attenuation. Figures 1.22E–H show how gross waveform distortion is produced by a too-severe low-frequency attenuation. The same EEG record filtered from 0.6 to 60 Hz and from 5 to 60 Hz gave quite different waveforms (compare Figs 1.22E and H). The large positive peak in (E) was strongly attenuated in (H), whereas the immediately following negative peak was much enhanced. The effects of low-frequency attenuation on latency were comparatively small; the positive peak went from 111 msec (E) to 96 msec (H), and the negative peak went from approximately 138 msec (E) to 131 msec (H). This trend can be followed progressively in Figures 1.22E through H. Thus, the main effect of too-severe low-frequency attenuation is to change EP waveform and distort the relative amplitudes of different EP components. *It is also possible to create spurious peaks by excessive low-frequency attenuation.* The reason for these effects is intuitively obvious. The main 100-msec positive wave of

Fig 1.22E is asymmetric about the baseline, and to reproduce this asymmetry the amplifier must sustain a positive DC level during the first 200 msec of the EP. Figures 1.20C and D illustrate that changing the lower 3-dB point from 0.6 to 5 Hz reduces the amplifier's ability to sustain a DC level. That is the reason why the area above the baseline approximates the area below the baseline in Figure 1.22H.

1.2.10 Analog-to-Digital Conversion

Sampling

An EEG amplifier's output signal is a voltage that is a continuous function of time, and as such is often referred to as an "analog signal." Some EP analyzers operate directly on this continuous analog signal. Analog signal analyzers include the analog Fourier series analyzer described in Section 1.8.1, the phase-locked amplifiers (Section 1.8.4), and some early averaging machines that used magnetic tape storage.^{1147,1475} Digital analyzers, however, do not operate directly on the continuous analog EEG signal. The EEG is first "sampled" or "digitized" so that it is represented by a sequence of numbers rather than as a continuous function. The purpose of an analog-to-digital converter (ADC) is to allow a numerical evaluation of an analog signal by first sampling it at regular intervals, and then expressing the resulting sequence of voltage levels as a sequence of numbers. Be-

cause of the way digital hardware operates, these numbers are expressed in binary form in the machine's memory.

Aliasing

Aliasing is a potential problem in any sampled data system. The effect is often illustrated in terms of Western movies. The film camera samples the visual scene at some fixed rate, each sample occupying one frame of the film. If a wagon wheel is rotating slowly enough, the motion is reproduced faithfully. But at higher rates of rotation the wheel may appear to rotate more slowly than its actual rate and may even appear to rotate backward. The false rotation speeds are generated because the movie camera's sampling rate is too low to accurately record high rates of rotation.

Similar considerations apply to the sampled EEG waveform. Errors will result unless the sampling rate is more than twice the highest frequency present in the EEG. This minimum sampling rate requirement is known as the Nyquist criterion. EEG frequencies above half the sampling rate will be represented as spurious low-frequency components that were not present in the original EEG signal; these higher frequencies are said to be "folded back," and they appear as noise of lower frequency.

Figure 1.23 illustrates how a too-low sampling rate (i.e., "undersampling") creates spurious frequency com-

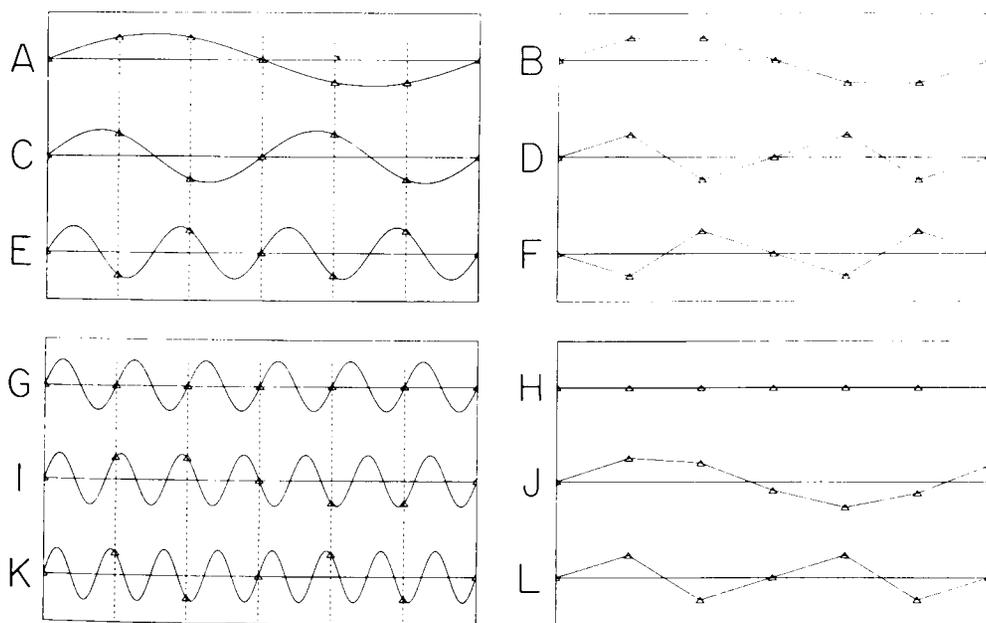


Figure 1.23

Aliasing

Six samples per cycle of a sine wave (A) give a rather faithful representation (B). Three samples per cycle (C) still correctly represent the sine wave's frequency (D). Fewer than two samples per cycle (E-K) incorrectly represent the sine wave's frequency (F-L) as a frequency that is lower than the true frequency or even (G, H) as a constant signal.

ponents. A sampling rate six times higher than the frequency of a sine wave signal (A) correctly represents the sine wave's frequency, and even without low-pass filtering (i.e., when the points are joined with straight lines) the sampled waveform approximates the original sine wave (B). At a sampling rate of less than two per cycle, however, the true frequency of the sine wave signal is not represented, and in addition spurious frequencies are introduced (E–L). For example, although a sampling rate of F Hz represents an $(F/3)$ -Hz sine wave correctly as an $(F/3)$ -Hz waveform (C, D), the same sampling rate also represents a $(4F/3)$ -Hz sine wave as an $(F/3)$ -Hz waveform rather than as the $(4F/3)$ -Hz waveform that it really is (K, L). Undersampling can even represent a sine wave (G) as a straight line (H).

We can now see intuitively that sampling less than twice per cycle will wrongly represent the frequency of a sine wave signal, and that even when sampling at a rate of exactly twice per cycle, the sine wave signal may be represented erroneously as a straight line. (This occurs when samples are taken at alternate zero-crossings.) What may be less obvious is that, in principle, sampling at a rate just higher than twice per cycle and then suitably low-pass filtering the sampled waveform will restore the original sine wave signal. But this is, in fact, the case.

In practical instruments, any input frequencies higher than half the sampling rate are effectively removed by an "anti-aliasing" filter. An ideal anti-aliasing filter would look like Figure 1.24A. It would pass all the desired input frequencies with no attenuation, and completely reject any frequencies higher than half the sampling frequency. Such an ideal filter is, however, theoretically and practi-

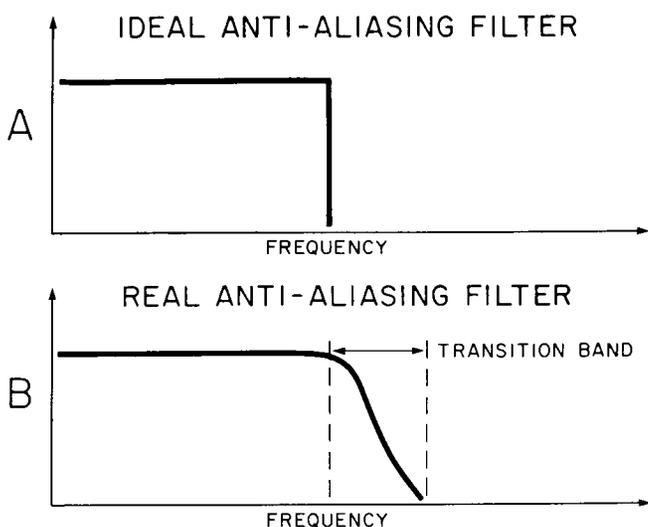


Figure 1.24
Practical Anti-aliasing Filters Require Sampling Faster Than the Theoretical Minimum Requirement
 (A) Ideal anti-aliasing filter. (B) Practical anti-aliasing filter.

cally impossible. Real filters look like Figure 1.24B, with a gradual rolloff and finite rather than infinite attenuation of unwanted signals. Large input signals that are not sufficiently attenuated in the transition band (Fig 1.24B) might still alias into the desired frequency range. To avoid this problem, the sampling rate is in practice set at about twice the highest frequency in the transition band and anti-aliasing filters are designed to produce a very large attenuation beyond the transition band. Typically, the sampling rate is set at 2.5 to 4 times the maximum desired input frequency.

Resolution of the Analog-to-Digital Converter

The resolution of an analog-to-digital converter (ADC) expresses its ability to distinguish different voltage levels. At the minimum possible resolution of one "bit," the input voltage is described by 0 or 1 depending on whether it is greater or smaller than some specified value. This specified value may be zero, in which case the process is called zero-crossing analysis. An 8-bit ADC resolves 256 (i.e., 2^8) different levels of the total voltage range accepted by the converter, whereas a 12-bit converter resolves 4,096 levels. It is important to realize that the ADC's resolution is specified in terms of its acceptable range of input voltages. Thus, an 8-bit ADC whose range is ± 5 V has an effective resolution of only 64 rather than 256 levels if the input voltage is limited to ± 1.25 V. To ensure that an ADC's range is being fully used, it is useful to monitor the input waveform at the input to the ADC by means of a cathode ray oscilloscope (CRO), and to set the CRO's sensitivity so that a full-scale deflection on the oscilloscope face corresponds to the full-range input to the converter.

The conversion of the analog signal into digital form is only as accurate as the resolution of the ADC. The conversion adds noise to the signal, called "quantization error." The root mean square (RMS) value of this noise is $0.29K$, where K is the amplitude between successive analog-to-digital conversion levels.¹⁵³ If a noise-free sine wave is fed into an averaging computer the quantization noise may be quite evident. However, if the signal is accompanied by random noise—as is the case in practical EP averaging—the averaging progressively reduces the effect of quantization, usually to negligible proportions. Picton et al¹⁸⁶⁸ compared quantization noise for ADCs of different resolution. Their data are shown in Figure 1.25. The ratio "r" is the increase of noise level over an ADC of infinite resolution. They found that quantization noise increased the noise level by no more than 3% for ADCs of 6-bit or better resolution, and pointed out that, for most purposes, an 8-bit converter is quite adequate, because when the signal is noisy the averaging process tends to smooth out the analog-to-digital steps. Converters of better resolution will not signifi-

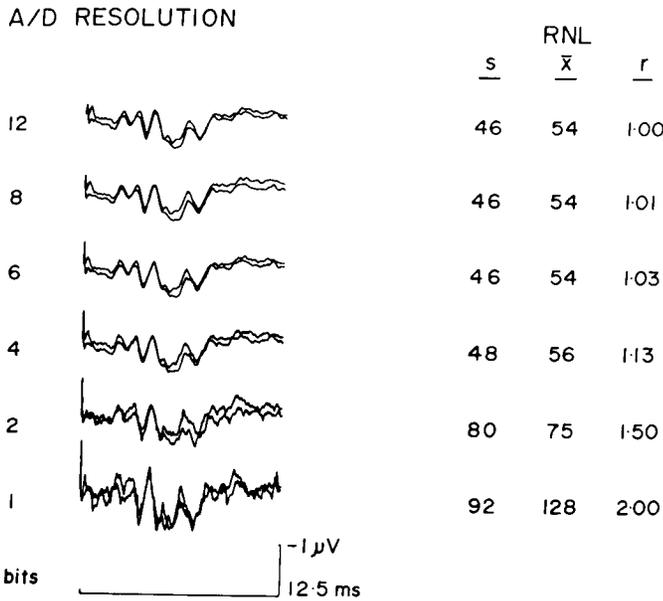


Figure 1.25
Analog-to-Digital Converter Resolution
 On the left are shown the auditory brainstem responses to a 70-dB nHL click presented at a rate of 10 per second. The EPs were recorded between the vertex and the mastoid, with negativity at the vertex being represented by an upward deflection. The bandpass of the recording was 20–2,000 Hz. The ADC was adjusted so as to accommodate an input range of voltages equivalent to $\pm 25 \mu\text{V}$ after amplification. The resolution of the ADC was varied from 12 bits to 1 bit. At each level of resolution, two replicate averages were obtained, each based on 2,000 trials. On the right of the figure are shown the residual noise level (RNL) as estimated using the \pm Reference. The values (s) for this particular subject and the mean values for eight recordings (\bar{x}) are shown together with the theoretical ratio (r) that represents the increase in the noise level over an ADC with infinite resolution. The waveforms recorded with 1-bit resolution represent polarity histograms. The amplitude calibration and the residual noise levels are inaccurate for these polarity histograms. Accurate values would require knowledge of the probability density function of the noise. Note that negative is upward. (From Picton TW, Hink RF, Perez-Abalo M, Linden RD, Wiens AS: Evoked potentials: How now? *J Electro-physiol Technol* 1984;10:177–221. Reproduced by permission.)

cantly improve the accuracy or efficiency of the averaging process, but will allow the efficient processing of signals that span less than the full range of the converter. On the other hand, spectrum analyzers may use ADCs of greater than 8-bit resolution. For example, the instrument used in Figure 1.29 has 12-bit converters.

1.2.11 Digital Filters, Physical Filters, and Phase Distortion

It is possible to design digital filters that do not introduce the phase shifts associated with a physical filter.^{526,936,2629} By use of digital filters instead of conventional physical

(i.e., “analog”) filters, distortion of the EP waveform and spurious displacements of peak latencies caused by phase shifts (Fig 1.22) can be minimized or avoided.

A digital filter is one in which signals are represented as sequences of numbers. These sequences are known as digital signals. Digital filtering in the time domain is performed by a computer on the digitized time series, in this case the digitized EEG or averaged EP.³

For the purpose of illustration we will consider digital filtering that is performed in the time domain by a procedure called “convolving an unfiltered waveform with a weighting function that extends symmetrically forward and backward in time.” At first hearing, something that extends both ways along the time axis may seem magical, and to many readers a formal mathematical description only deepens the mystery. However, what is actually done to the unfiltered waveform is not magical at all, and is easily understood at an intuitive level.

For example, suppose that EEG voltage is sampled at 200 Hz, and the following samples are recorded. Starting from an arbitrary time t sec, the EEG is sampled successively at the instants t sec, $t + 0.005$ sec, $t + 0.01$ sec, and so forth. Let the amplitudes of these successive samples be (in volts) $V_t, V_{t+0.005}, V_{t+0.01}, V_{t+0.015}, V_{t+0.02}, V_{t+0.025}$, and so forth. Now we calculate the mean of the first three samples, and replace the voltage $V_{t+0.005}$ with this new mean voltage. Then we calculate the mean of the second, third, and fourth samples and replace $V_{t+0.01}$ with this new value, and so on. This will give a new sequence of samples $V'_{t+0.005}, V'_{t+0.010}, V'_{t+0.015}, V'_{t+0.020}$, and so forth (in volts) where

$$V'_{t+0.005} = 1/3(V_t + V_{t+0.005} + V_{t+0.010})$$

$$V'_{t+0.010} = 1/3(V_{t+0.005} + V_{t+0.010} + V_{t+0.015})$$

...

In this new record, the EEG record is still represented by 200 samples per second, just as in the original EEG record, but the EEG has, in effect, undergone low-pass filtering in the sense that high-frequency components have been attenuated relative to low-frequency components. Filtering will be more severe; that is, the upper 3-dB frequency will be lower, if means are computed for successive blocks of five rather than three samples, more severe still for blocks of seven samples, and so on.

As a point of detail, rather than computing the simple means of, for example, successive blocks of three sample voltages, some digital low-pass filters assign less weight to the first and last sample than to the middle sample in any given block. However, the above description brings out the essential point.

³ Digital filtering can also be performed in the frequency domain, but for our present purpose time-domain and frequency-domain procedures give similar results (see Ref 2142).

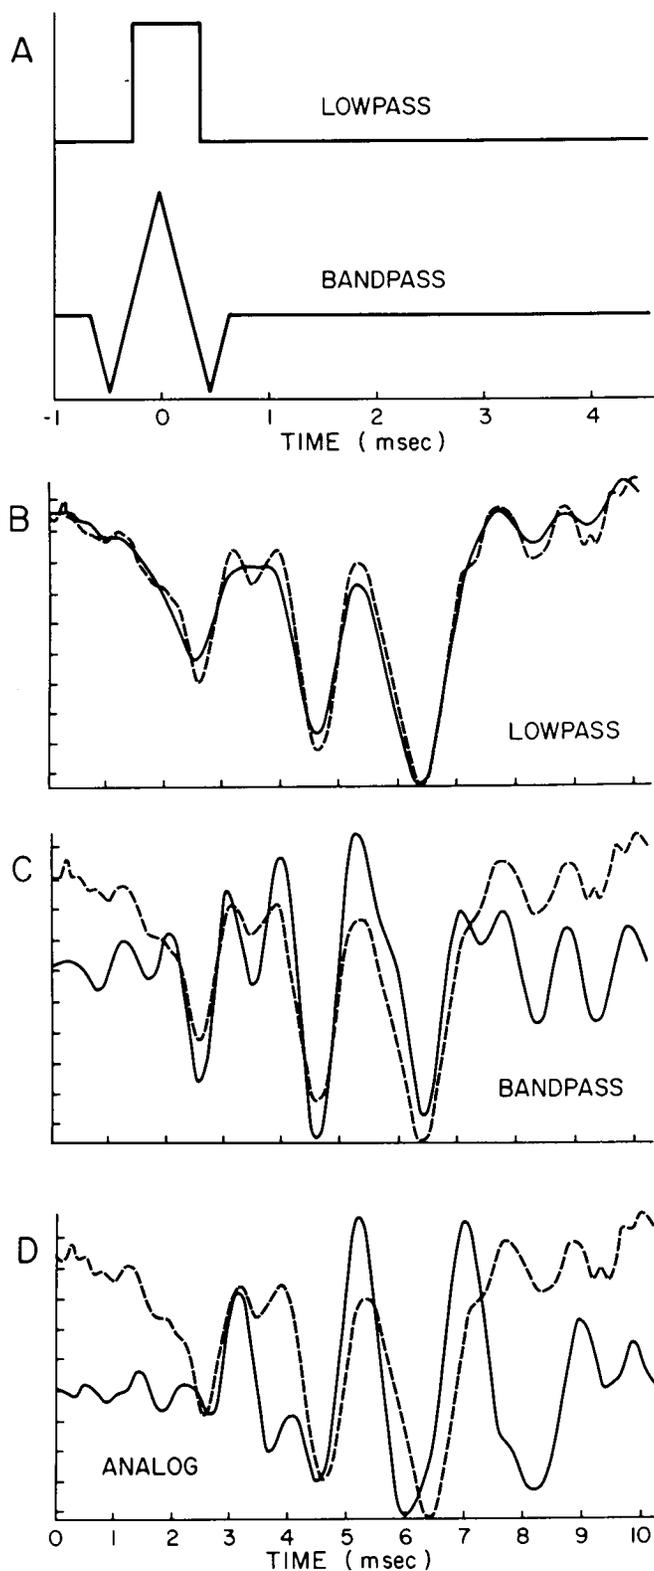


Figure 1.26
Digital Filtering Compared With Conventional Analog Filtering
 (A) Weighting function for a low-pass digital filter and for a high-pass digital filter. The dashed line in panels B–D is an averaged EP prerecorded using an amplifier with 11- to 3,400-

Clearly, the procedure just described takes equal account of voltage samples recorded before and after the midpoint of each time block. Contrast this kind of low-pass filtering with low-pass filtering carried out by a physical filter. A physical filter cannot respond to inputs *before* they occur: this is the reason why physical filters cannot attenuate one frequency relative to another frequency without at the same time causing a relative phase shift and why, for a given relative attenuation, the relative phase shift cannot be less than a certain amount. Figure 1.2 plots these minimum phase shifts for physical filters.

Figure 1.26 compares the effects of digital and analog (physical) filtering on the averaged brainstem auditory EP. Two digital filters were compared: the first had a weighting function that gave low-pass filtering (as in the example just discussed), and the second had a weighting function that gave bandpass filtering (Fig 1.26A). The low-pass weighting function was either zero or positive, and suppressed high frequencies. The bandpass weighting function had equal areas above and below the horizontal zero line, thus suppressing both DC and high frequencies. Note that, because the weighting functions were symmetrical about zero time, the recording of the EP had to start slightly before stimulus delivery. The dashed line in Figures 1.26B–D is a low-noise brainstem auditory EP (BAEP) evoked by summing the responses to 20,000 tone pips each of 5-msec duration over a bandwidth of 11–3,400 Hz. The continuous line in (B) is the same averaged waveform after digital low-pass filtering. High frequencies were removed without appreciable distortion of the main peaks. Figure 1.26C shows the effect of bandpass digital filtering; sustained deflections were preferentially removed, enhancing the oscillatory peaks, without affecting peak latencies (though onset and offset latencies are affected [D. Ruchkin, personal communication]). Figure 1.26D shows the effect of conventional analog filtering with a bandpass filter consisting of three-stage (18 dB per octave) minimal-phase-shift high-pass and low-pass sections. Just as with the digital bandpass filter, the analog filter attenuated both high-frequency and low-frequency components. The important point is that the analog filter also distorted the EP waveform and spuriously altered peak latencies (D), whereas the digital filter did not (C).

Phase errors caused by filtering can be a problem in

Hz bandwidth. The continuous line in panels B and C is the EP filtered by the lowpass and bandpass digital filters, respectively. (D) Effect of conventional analog filtering with 340- to 2,300-Hz bandpass. (Modified from Moller AR: A digital filter for brain stem evoked responses. *Am J Otolaryngol* 1980;1:372–377. Reproduced by permission.)

recording steady-state EPs as well as in recording transient EPs. A case in point is the advantages to be gained by using digital rather than analog low-pass filters when recording sweep EPs (Section 1.9.2).

1.2.12 Errors in Estimating Transmission Time Caused by the Filtering Characteristics of the Recording Equipment

The filtering properties of the recording equipment can introduce errors into the estimation of transmission time from eye to scalp in the case of visual EPs, and from ear to scalp in the case of auditory EPs. This is the case for both transient and steady-state EPs.

Figures 1.19–1.22 illustrate this point for transient responses. In the case of visual EPs, Figures 1.22A–D show that the measured latency of the main positive peak shifted from an initial 107 msec to a delayed 133 msec when the amplifier filter's upper 3-dB frequency was lowered from 100 to 15 Hz. Figure 1.26D illustrates the effect of physical bandpass filters on the measured latencies of auditory brainstem EPs (compare the dashed and continuous lines).

For steady-state EPs, transmission time is estimated from the plots of gain and phase versus frequency (see Section 1.3.3). It is most important, however, to bear in mind that the recording equipment introduces phase shifts and that these phase shifts *must be subtracted* from the EP phase shifts before attempting to estimate the EP's "apparent latency." The phase shifts introduced by the recording equipment are associated with filtering, and depend on the amplifier's filter setting. Perhaps the simplest way to deal with these phase shifts is to calibrate the recording equipment with the filter settings used in the EP recording, and to subtract these phase shifts from the EP phase shifts.^{1973,1975} If this phase calibration is omitted, an appreciable error can result. Published studies of steady-state EPs do not always state that this correction has been applied.

1.2.13 Errors in Estimating Transmission Time Caused by the Filtering Characteristics of the Brain

The previous section emphasized that the filtering properties of the recording equipment can cause a systematic error in estimated EP transmission time. It is easy to overlook that the same holds for the filtering properties of the brain itself. (This frequency filtering is not that due to current flow through passive media, but rather to filtering that takes place in retinal photoreceptors, intervening neurons, and any reverberant feedback pathway in the brain.) In the case of steady-state EPs, for example, unless any phase shifts caused by high-pass or low-pass

filtering within the brain are allowed for, the estimated EP transmission time can be erroneous. Unfortunately, this correction is less easily calculated for the EP system of the brain than for the straightforward case of a linear system described in Section 1.2.12.

One method for dealing with this problem is to treat the brain *as though* it were a linear frequency filter. This approach runs as follows:

1. Select a harmonic component of the steady-state EP and record its amplitude and phase over a range of stimulus frequencies.
2. Plot gain versus frequency and phase versus frequency. If the amplitude of the physical stimulus does not depend on frequency, EP amplitude can be substituted for gain.
3. From the gain-versus-frequency data, calculate the phase shifts that would be produced by a linear minimum-phase system. (Figure 1.2 gives examples of the relationship between gain and phase for a linear minimum-phase system.)
4. Subtract these calculated phase shifts from the measurements of EP phase to give a plot of corrected EP phase versus frequency.
5. Calculate "apparent latency" as defined in Section 1.3.3.

This correction procedure was used by Regan^{1973,1975,1977} and by van der Tweel and Lunel.²⁵⁷⁸ The correction can be substantial; for example, Regan¹⁹⁷⁷ estimated the apparent latency of the steady-state flicker VEP over the range 36–58 Hz to be 75 msec uncorrected, but 62 msec after correction for phase shifts caused by curvature of the amplitude-versus-frequency curve, so that the brain's bandpass filtering introduced an extra 13 msec to the apparent latency of the high-frequency VEP, that is, a 21% error in the estimation of transport time. Again, the apparent latency of flicker VEPs over the range 16–29 Hz was 120 msec uncorrected and 100 msec after correction;¹⁹⁷⁸ in this case, the brain's bandpass filtering introduced an extra 20 msec to the apparent latency of the intermediate-frequency VEP, that is, a 20% error.

Such errors in estimating apparent latency can be seriously misleading when steady-state VEPs are used as an aid in diagnosing multiple sclerosis: phase shifts caused by the shape of the amplitude-versus-frequency curve can be misinterpreted as delay and, conversely, a real increase in transmission time can be masked by phase shifts caused by the amplitude curve.

Figures 1.2C and D illustrate that high-frequency attenuation is associated with a progressive increase of phase lag with increasing frequency. Over a substantial frequency range, phase lag is approximately propor-

tional to frequency, so that the slope of the phase-versus-frequency plot mimics the effect of a real transmission time, though it is nothing of the kind. If the high-frequency attenuation of the VEP increases in steepness, this spurious delay will increase and may be misinterpreted as due to a slowed axonal conduction caused by demyelination. Few reports of steady-state VEPs in multiple sclerosis have explicitly recognized this point, and in several it is difficult to estimate the correction because the amplitude-versus-frequency plot is not provided. In the converse situation the spurious change in apparent latency resulting from the brain's filtering is in the direction opposite the change actually observed, and in this case one can be more confident that a slowing of conduction has occurred. For example, in the center panel in Figure 3.3 of Ref 1623, VEP amplitude shows strong high-frequency attenuation for the right eye, but it is the left eye that shows increased apparent latency.

Turning from steady-state EPs to transient EPs, we have seen earlier how the latency of a peak can be misleadingly increased from 107 to 133 msec by increasing the EEG amplifier's high-frequency attenuation (Fig 1.22). Similarly, even when transmission time from eye or ear to scalp is not changed, a change in the brain's high-frequency attenuation (e.g., caused by pathology or even a reduction of light level) can be expected to alter the latencies of VEP peaks.

1.2.14 Noise and Interference

It is noise (defined as undesired signal) and not the maximum available amplification that determines how small a signal can be detected. Methods of recording and measuring very small signals such as EPs are basically procedures for rejecting noise. In EP recording the three major sources of noise are (1) electrical, magnetic, and radiated interference, (2) biological noise (e.g., alpha activity), and (3) internal instrumental noise (including amplifier noise). Amplifier^{151,2701} noise is generally of negligible importance when low-resistance (i.e., a few thousand ohms) scalp electrodes are used. On the other hand, it is usually necessary in EP research to deal with problems caused by interference. In practice, the investigator must make a compromise between flexibility and a "once and for all" cure of interference problems. The principles of the systematic detective work required to track down sources of interference are discussed by Donaldson⁶²¹ (see also Ref 1659). A source of electrical interference can often be located by soldering a 5-k Ω resistor to the end of a pair of wires which are screened almost up to the resistor. If the screened lead is connected to a high-gain amplifier, then the resistor at its end can be used as a probe to assess the distributions of interfering electric fields.⁶²¹ A loop of wire can be used in a similar manner

to track down a source of magnetic interference.⁶²¹

Correct grounding and screening are essential. This point is discussed by Morrison¹⁶⁵⁹ (see also Refs 621 and 2701). Ground loops are a potent cause of interference. It is therefore essential that the recording equipment be grounded at only one point, and that the connection to ground is of low resistance. On occasion it is necessary to construct a low-resistance ground by connecting a thick wire to a buried, watered metal plate. If such a ground is used, then all connections to the mains ground should be broken (see Refs 621 and 2701).

Electrical interference at mains frequency can be attenuated strongly by surrounding the subject with a grounded Faraday cage made, for example, of chicken wire. Often it proves possible to obtain satisfactory screening while leaving large holes in the cage. Nevertheless, it is good practice to remove electrical interference at the source by using screened mains leads and by enclosing electrical devices in grounded metal boxes. In principle, mains frequency magnetic interference can also be removed by screening, in this case with a number of layers of high-permeability metal. Unfortunately, this solution is generally impractical. Three practical ways of minimizing magnetic interference are (1) to twist the electrode leads together so as to reduce the area of the pickup loop, (2) to move the source of the magnetic interference as far away as possible, and (3) to align the source so that it produces minimum interference.

Radiofrequency interference can often (but not always) be identified by connecting a loudspeaker to the EEG amplifier's output.⁴ Such interference can often be reduced by placing small inductors (e.g., a wire wrapped several times around a resistor) right at the amplifier inputs. Dawson pointed out that, *even though high-frequency interference may not pass right through the recording equipment, it can cause an early stage in an amplifier to saturate and thus produce spurious readings*, so that efforts should be made to prevent high-frequency interference from reaching the amplifier inputs.

1.2.15 Recording Artifacts

The ability of modern equipment to extract a hidden signal from noise is indeed impressive. But it is a dangerous mistake to overestimate the equipment's ability to reject nonsignal. Prevention is by far the best approach when dealing with noise and artifacts: clean inputs are

⁴ G. D. Dawson demonstrated this to his graduate student, the author, in 1961. The radio commentary on an England-versus-Australia cricket match that issued from the loudspeaker connected to an EEG amplifier, Dawson explained, was due to nonlinearity within the amplifier that partially rectified radiofrequency signals within the EEG amplifier.

preferable to noisy inputs. Even if a few commonsense preventive measures do not solve the problem entirely, they often reduce it to manageable proportions. Rather than relying entirely on the computer, it can add enormously to one's efficiency to decrease both the number of artifact-contaminated trials and the degree of contamination by carefully instructing the subject to relax and to blink as little as possible during the recording. These instructions should be reinforced by asking the subject after every recording whether he or she moved or blinked during the recording. Of course, these instructions will be futile unless time and care have been devoted to ensuring that the subject is truly comfortable. If a chair is used, the correct height is crucial. This may mean that several alternative chairs must be available, or that the legs of an old chair are shortened. Some laboratories favor an old-fashioned dentist's chair because of its adjustability; other laboratories use a domestic armchair, office chair, or even a specially made padded wooden box, depending on the stimulus setup. A supine position can help to relax the subject. All this may seem unsophisticated, but it is most certainly worthwhile spending several tedious days to ensure that the physical setup allows one's subjects to be truly comfortable and relaxed.

Before starting a recording one should wait until the subject indicates that he or she is ready. It is important to tell the subject immediately when a recording has ended so that they can be confident of a rest between recordings.

Artifacts fall into two classes: (1) those that are not time-locked to the stimulus and (2) those that are time-locked to the stimulus. If their amplitudes are sufficiently large, artifacts that are not locked to the stimulus, or that occur only once or twice, may be prominent in the averaged EP waveform. Examples of such artifacts are those resulting from mains pulses, eye movements, or eye blinks or movements of the subject. Stimulus-locked artifacts can be insidious. For example, an auditory stimulus can provoke a stimulus-locked eye movement that creates an artifactual potential far from the eyeball. Eye movement artifacts associated with CNV and other cognitive paradigms can be large compared with the signal. Solutions to this special problem are discussed in Section 1.2.16.

Commercial averagers are often provided with an "artifact reject" mode of operation. This mode cuts out any trials with unusually large voltage excursions. The EEG recorded during each sweep is temporarily stored in a buffer memory. If the EEG amplitude does not exceed some preset level during the sweep time, the contents of the buffer memory are added to the previously summed samples, the sweep counter is incremented, and the buffer memory is cleared. On the other hand, if the EEG amplitude exceeds the preset level at some point during

the sweep, the contents of the buffer memory are not added to the previously summed samples, the sweep counter is not incremented, and the buffer memory is cleared. Some commercial averagers allow the artifact reject to be operated either when the EEG amplitude exceeds some limit or when the experimenter pushes a button; this facility can be useful when recording from infants, because it allows the experimenter to reject any sweep during which the subject was not looking at the stimulus. Even though blink artifacts may not be large in recordings from the vertex or inion, a blink artifact can be rejected by recording the larger blink potentials from electrodes near the eye and using this signal rather than the EEG to trigger the artifact reject mode.

A second approach ("clipping") is to allow the ADC to saturate at levels just outside the limits of noise on low-noise trials. Picton et al¹⁸⁶⁸ compared the effectiveness of clipping and rejection in dealing with high-noise trials when recording the auditory brainstem response in 10 subjects (Fig 1.27). Effectiveness was defined as the ratio between the standard deviation of the averaged response and the standard deviation of the \pm reference. Picton et al concluded that, for a fixed recording time, clipping is more efficient than rejection.

Gratton et al⁸⁷¹ have described a computational method for off-line removal of ocular artifacts from stored single-trial EP recordings.

When recording EPs to stimuli of any modality it is essential to control for stimulus-locked artifacts by recording averages under "stimulus-off" conditions. Because it is conceivable that the presence of a stimulus could alter both the background EEG activity and the subject's psychological state, a useful control is to record averages while stimulating at a frequency slightly different from the averager sweep frequency *without previously warning the subject*.¹⁹⁸³ (Note: It is essential to ensure that the frequencies are sufficiently different that any averaged EP is of negligible amplitude.) This procedure gives an estimate of the noise level in the averaged EP; a knowledge of the noise level is essential before two waveforms can be judged to be different.

Even artifacts of very small amplitudes can be troublesome in the averaged EP if the time of occurrence of the artifact is locked to the averaging sweep. Such time-locked artifacts may be biological or instrumental. Although auditory EPs produced by the click of a stroboflash may be of small amplitude, the sound is best eliminated because it is difficult to ensure that the auditory stimulus does not affect either the much slower VEP or the myogenic response to the visual stimulus.

The myogenic potential is, perhaps, the most notorious among artifacts of biological origin. It is more serious for visual and auditory EPs than for somatosensory EPs. The degree of myogenic contamination of an EP

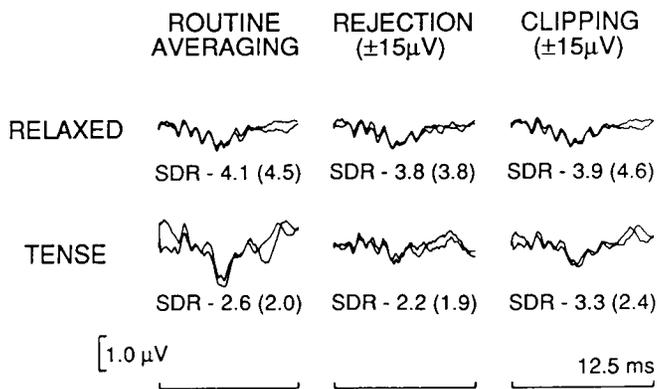


Figure 1.27
Comparison of Rejection and Clipping Procedures for Dealing With High-Noise Trials

In the left column of this figure are shown the auditory brainstem responses to 70-dB clicks presented at a rate of 10 per second. The responses were recorded between vertex and mastoid, with negativity at the vertex being represented by an upward deflection. Each tracing represents the average of 2,000 trials. Recordings were made under two conditions. In one the subject was RELAXED, and in the other the subject intermittently gritted her teeth (TENSE). Beneath each of the replicate averages is given the signal-to-noise measurements for 10 subjects given in brackets. The auditory EPs recorded under the TENSE condition were distorted and showed a much lower signal-to-noise ratio than when the subject was RELAXED. In the middle column are shown the effects of rejecting all trials containing amplitude of greater than $\pm 15 \mu\text{V}$. In the RELAXED condition this resulted in rejecting 21% of the trials for this subject (on average 25% for all subjects). In the TENSE condition the rejection rate was 58% for this subject (59% on average). In the right column is shown the effect of clipping the waveform at $\pm 15 \mu\text{V}$. This has no significant effect on the signal-to-noise ratio in the RELAXED condition and causes a significant improvement in the signal-to-noise ratio in the TENSE condition. (From Picton TW, Hink RF, Perez-Abalo M, Linden RD, Wiens AS: Evoked potentials: How now? *J Electrophysiol Technol* 1984;10:177-221. Reproduced by permission.)

can be assessed by varying muscle tension^{181,183} (Fig 1.28), but perhaps the most satisfactory way of distinguishing myogenic from other potentials is to measure topographical distributions.²⁵⁸⁸ Topographical plots also offer a solution to the problem of recognizing contamination of the EP from the ERG, eye blinks and the AC potentials due to changes in the direction of the DC corneoretinal field during eye movements. Note that eye blinks and eye movements may be triggered by the stimulus, so that quite small blink or eye-movement potentials can be summed by the averager when they are synchronous with the stimulus.

Visual EPs have been reported to be influenced by the phase of the cardiac cycle.^{342,451}

The reported influence of eye position on alpha activity^{1680,1681} has been attributed not to changes in eye posi-

tion per se, but rather to accompanying changes in visual input.³⁸¹ The amplitude of "averaged alpha" can be reduced by averaging only during periods of low alpha activity, although it is not clear whether this maneuver might bias the EPs to any serious extent. According to D. Ruchkin (personal communication, 1986), suitable digital filtering with the first zero at the alpha frequency can be quite effective.

When a particular stimulus is the subject's cue to make some sensory discrimination or to perform some motor action, then *if the time of occurrence of the stimulus is predictable* the subject may alert himself or herself and prepare for mental or motor activity. If the EPs to such cueing stimuli are found to be different from EPs to control stimuli, then the differences may be due in part to the subject's stimulus-locked alerting reaction.¹⁶⁹⁰ A good approach to this problem is to randomize the stimulus delivery so that the subject cannot anticipate the nature of the next stimulus.¹⁶⁹⁰

Photoelectric artifacts can occur if the stimulus light is allowed to fall onto the electrodes during VEP recording, and this artifact is especially troublesome when recording the early receptor potential of the ERG (Section 2.6.3). When shock stimuli are used (e.g., to generate phosphenes), some care is necessary to ensure that electrical signals from the stimulus do not spread to the recording electrodes. Electrostatic or magnetic coupling between stimulus circuits and the recording circuits must be eliminated by such tactics as using separate power supplies, avoiding earth loops, and maintaining a generous distance between the stimulus and the recording leads and circuitry.

It is best to calibrate the recording system by feeding microvolt signals into the headboard from a source impedance comparable to that of the subject's electrodes.

1.2.16 Eye-Movement Artifacts in Cognitive EP Recording

In healthy eyes there is a standing (DC) potential of several millivolts across each eyeball (cornea positive with respect to the opposite side of the eye). Consequently, eye rotation can cause a transient potential at the vertex, and this potential can be large; a voluntary downward rotation of 10° produces a negative shift of about 50 microvolts at the vertex.¹⁰⁴³ The force of this point is that some subjects perform systematic eye movements during cognitive EP paradigms. In CNV experiments, for example, large involuntary eye movements are commonly synchronized with the preparatory interval, especially when the eyes are closed.¹⁴³⁴

Hillyard¹⁰⁴³ compared several methods for eliminating the contamination of CNV by eyeball-rotation potentials.

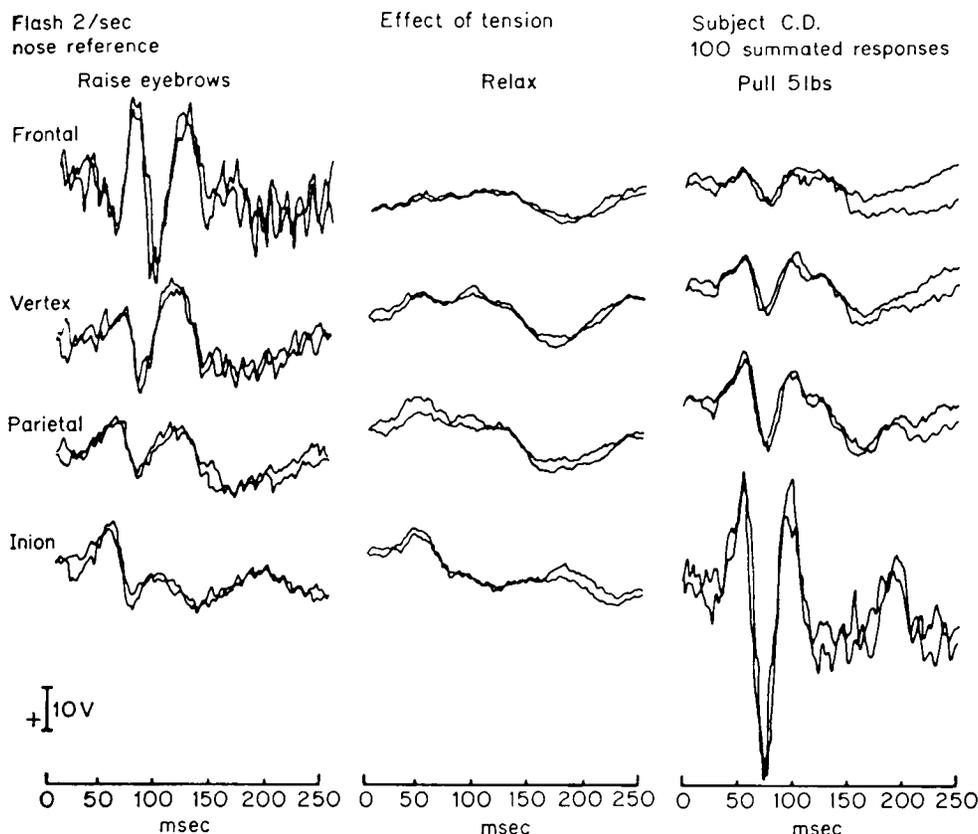


Figure 1.28

Effect of Muscle Tension on Photically Evoked Myogenic Responses

The center section shows visual EPs from midline electrodes for a relaxed subject. When the eyebrows are raised a new response with a latency comparable to an evoked cortical response appears in the frontal region and, to a lesser extent, at other electrodes (left section). The right section shows that, when the neck and posterior cranial muscles are tensed by having the subject support a weight, a large visual EP appears which is maximal in the occipital regions and has latencies of N50, P75, and N100. The responses shown in the figure were produced by a flash stimulus of intensity 8 (Grass photostimulator P2). Photicmyogenic responses can be obtained down to perceptual threshold. (From Bickford RG: In Donchin E, and Lindsley DB (eds): *Average Evoked Potentials*, NASA SP-191, Washington, D.C., U.S. Government Printing Office, 1969. Reproduced by permission.)

Visual Fixation and Elimination of Trials With an Electro-oculogram

With adults the preferred method is to avoid the eye-rotation problem in the first place by providing a fixation target and instructing the subject to fixate it during trial intervals, reserving blinks for intertrial intervals. Residual contamination is $2 \mu\text{V}$ at most¹⁰⁴³ in the vertex-mastoid channel, but if necessary this can be compensated by one of the methods described below.

Subtraction of the Artifact Based on the Calibrated Electro-oculogram

The subject is required to make a series of voluntary mirror-image eye movements in response to paired stimuli. The total vertex potential is the linear sum of two components: the eye-rotation artifact and the CNV pre-

ceding the ocular rotation. Algebraic subtraction of the potentials during upward and downward eye movements eliminates the CNV and gives twice the artifact.¹⁰⁴⁸ The experiment is repeated for small, intermediate, and large eye movements, and a calibration curve is plotted. This function is approximately linear, and gives the amount of artifact per unit electro-oculogram (EOG). Hillyard and Galambos¹⁰⁴⁸ found that CNVs recorded with eyes closed contained 10–50% eye-movement artifact. The drawbacks of this procedure are that it is time consuming, and secondary errors may arise as a result of lid movements.

Potentiometric Subtraction of Artifact^{884,885}

One terminal of a 25-k Ω potentiometer is attached to an electrode above the eyes, and the other end is connected

to the mastoid. The center tap is adjusted until the eye rotation artifact at the tap and at the vertex are identical. The tap rather than the mastoid electrode is then used as reference. The difficulty is that this procedure corrects only one lead, such as the vertex, but in so doing may increase the artifact in other leads.⁴²⁷

1.3 Distinctions Between the Transient and Steady-State Responses of a System

During the following discussion it will be helpful to keep in mind the distinction between two points: (1) Any waveform that exists in the physical world can be completely described in either the time domain or the frequency domain, but in some practical cases one description may be more convenient than the other. (2) Although the response of a linear system to a single transient input contains the same information as the response to a repetitive input, for a nonlinear system the two responses may contain different information.

1.3.1 The Time Domain and the Frequency Domain

To most Western minds, everyday experience seems to be a succession of transient events, each of which has a beginning and an ending. For example, a musical melody may contain several different tone frequencies, but each note lasts for only a brief time rather than enduring indefinitely. Thus, our everyday experience predisposes us to regard a time-domain description of physical events—including oscillations—as natural and intuitive.

In electrical engineering the forward and the inverse Fourier transforms are basic to the theory of communication. Taken together the two transforms assert that a time series can be completely described in terms of frequency, and vice versa—the two descriptions are equivalent. But for the reasons described before, only the time-domain description seems natural. A deep understanding of the frequency domain is less easily attained. One reason why physical insight is so elusive is that Fourier's definition of "frequency" is far removed from our everyday experience of oscillation. A sharp illustration of this point is that, having obtained for the first time (in 1922) the spectrum of a frequency-modulated sine wave and found that the spectrum contains *no energy at the modulating frequency* (see Fig 1.64), Carson³⁶⁷ wrote that "The foregoing solutions, though unquestionably mathematically correct, are somewhat difficult to reconcile with our physical intuitions, and our physical con-

cepts of such 'variable frequency' mechanisms as, for example, the siren." The student of today who feels uneasy with Carson's mathematical result is in good company!

The main reason for this unease with the frequency domain is that Fourier's frequency-domain description is cast in terms of oscillations that persist through infinite time: each sine wave stretches infinitely into the past and future. Thus a "changing frequency" becomes a simple *contradiction in terms*. Fourier defined the "rules of the game" so that what we think of as a changing frequency must be described in terms of constant frequencies. In Fourier mathematics the concept of "frequency" is such as to ensure that the time-domain description and the frequency-domain description are mutually exclusive.

In view of the practical scientist's reactions to this seeming artificiality, it is remarkable that a formal attempt to reconcile the exclusiveness of the time domain and the frequency domain in the context of the theory of communication was not published until Gabor's 1946⁷⁹⁰ paper.⁵ Gabor traced his line of thought back to early work on wave mechanics and especially to Heisenberg's principle of indeterminacy, published in 1927. This principle states that there is a trade-off between the precision with which we may simultaneously know the momentum and the position of a particle. From Heisenberg's insight, Pauli was able to redefine observable physical quantities in such a way that the uncertainty relationships between them appear as a direct consequence of the mathematical identity

$$\Delta f \Delta T \approx 1$$

where Δf is the uncertainty of the frequency of an oscillation of duration ΔT .

Gabor's 1946 paper⁷⁹⁰ was directed toward achieving theoretical understanding of a long-standing practical problem, namely, how to maximize the rate of information transfer along a given radio link. By 1927 it was already known that high-definition television required a wide bandwidth. Gabor showed that the *uncertainty relation* $\Delta f \Delta T \approx 1$ sets a theoretical upper limit to the rate of transmission of "quanta of information" along a channel of given bandwidth.⁶ (See Shannon and Weaver²²⁶⁰ for a parallel development in the United States.)

⁵ Outside the spatial-vision research community Gabor is best known for a different achievement—the invention of holography (some decades before the laser was invented), for which he was awarded the Nobel Prize.

⁶ The uncertainty relation involves a scaling constant that depends on how one chooses to define frequency bandwidth Δf and how one defines the duration of the signal ΔT . Gabor chose the RMS definition of