

# Tracking eye fixations with electroocular and electroencephalographic recordings

CARRIE A. JOYCE,<sup>a</sup> IRINA F. GORODNITSKY,<sup>b</sup> JONATHAN W. KING,<sup>c</sup> AND MARTA KUTAS<sup>b,d</sup>

<sup>a</sup>Department of Computer Science, University of California, San Diego, California, USA

<sup>b</sup>Department of Cognitive Science, University of California, San Diego, California, USA

<sup>c</sup>Department of Psychology, University of Missouri, Columbia, Missouri, USA

<sup>d</sup>Department of Neurosciences, University of California, San Diego, California, USA

## Abstract

We describe a method, based on recordings of the electroencephalogram (EEG) and eye movement potentials (electrooculogram), to track where on a screen ( $x, y$  coordinates) an individual is fixating. The method makes use of an empirically derived beam-forming filter (derived from a sequence of calibrated eye movements) to isolate eye motion from other electrophysiological and ambient electrical signals. Electrophysiological researchers may find this method a simple and inexpensive means of tracking eye movements and a useful complement to scalp recordings in studies of cognitive phenomena. The resolution is comparable to that of many commercial systems; the method can be implemented with as few as four electrodes around the eyes to complement the EEG electrodes already in use. This method may also find some specialized applications such as studying eye movements during sleep and in human–machine interfaces that make use of gaze information.

**Descriptors:** Eye tracking, Electrooculogram, Event-related brain potentials

There is ample evidence that much can be learned about human information processing by monitoring where the eyes look and for how long. Indeed, many different methods for tracking eye movements have been developed over the past 30 years. These include electrooculography (EOG; e.g., Mowrer, Ruch, & Miller, 1936; Ong & Harman, 1979), corneal reflection (e.g., Monty, 1975; Muller, Cavegn, d'Ydewalle, & Groner, 1993), limbus and pupil tracking (e.g., Eadie, Pugh, & Heron, 1994; Muller et al., 1993), contact lenses (e.g., Ditchburn & Ginsborg, 1953), and Purkinje reflection imaging (e.g., Cornsweet & Crane, 1973; Muller et al., 1993). See Young and Sheena (1975) for a review of methods. These methods vary greatly with respect to their resolution, range, tolerance to head motion, ease of use, invasiveness, and cost. Thus the method of choice is clearly dependent upon the purpose for which it is intended.

In this article, we describe a method that can be used to track an individual's eye fixations by recording activity from as few as four electrodes around the eyes and a modest complement of electroencephalogram (EEG) electrodes on the scalp. The motivation behind this method was to develop an unobtrusive, inexpen-

sive way of simultaneously tracking eye fixations and recording event-related brain potentials (ERPs). Commercial eye tracking systems tend to be costly, and many rely on head-mounted devices that interfere with the electrode placement, especially when caps and nets are used in the collection of ERP data. The method we propose is simple, and requires very little extra time to set up, and only a few spare amplifiers beyond that used in a standard ERP experiment. The method provides good resolution, with mean error on the order of  $1\text{--}2^\circ$  for patterns spanning  $30^\circ$ . This is within the resolution limit of EOG data. According to Young and Sheena (1975), the precision with which EOG output reflects the actual angle of gaze is within  $\pm 1.5^\circ\text{--}2^\circ$ . The precision of our method is within this limit, meaning that it provides optimal resolution under the EOG data constraints. The resolution obtained, nevertheless, is comparable to that of many commercial eye trackers. Such resolution is sufficient for studying eye fixations during typical object, face, or scene viewing and during reading of sufficiently large texts.

Using the EOG to track eye movements offers additional advantages. EOG recordings can be made relatively unobtrusively, and can be easily done even if the individual is wearing glasses, contacts, and other special eye wear, such as pilot goggles. As such, EOG-based eye tracking may stand out in certain specialized applications where other techniques are hard to use. One example is recording eye movements during sleep. The technique's natural coupling with EEG monitoring of cognitive activity likewise renders it a good candidate for applications such as human–machine interfaces and on-line alertness monitoring. Finally, because EOG does not require visualization (i.e., camera recording) of the eye

---

C.A.J. and M.K. were supported by a research grant from the NIA (08313); C.A.J. also was supported by a research grant from the McDonnell Foundation (15573-S6); M.K. was also supported by a research grant from the NICHD (22614). I.F.G. was supported by a research grant from the NSF (IIS-0082119).

Address reprint requests to: Marta Kutas, University of California, San Diego, Department of Cognitive Science, 9500 Gilman Dr., San Diego, CA 92093-0515, USA. E-mail: mkutas@ucsd.edu.

itself, it provides a considerably larger range than many other eye tracking methods. EOG recordings have a range on the order of  $\pm 70^\circ$ . That said, the relationship between EOG output and the angle of gaze is linear only in a limited range. Estimates for this range vary from  $\pm 15^\circ$  to  $30^\circ$ , depending somewhat on the direction of the gaze, that is, horizontal versus vertical (e.g., Moses, 1975). Because the method we describe is based on a learned relationship between EOG output and the changes in the angle of gaze, it is expected to apply in the range of learned response only.

In brief, the method determines where (i.e.,  $x, y$  coordinates on a display in front of them) an individual is fixating at any given moment by monitoring electric field changes generated by eye movements. The electric field is created by a charge differential in the eyeball. The cornea is positively charged relative to the retina, which amounts to having a steady retino-corneal charge of between 0.4 and 1.0 mV that approximates a dipole in both eyes. The dipole analogy is approximate because of distortions in the dipolar properties caused by tissue irregularities in the eyes. As the retino-corneal axis rotates, the orientation of this dipole in three-dimensional space also rotates, resulting in changes in the electric field that can be picked up by electrodes placed around the eyes.

If we could determine the mapping between an eye movement and the associated change in the electric surface potentials, then we could use these potentials as a tool to estimate the loci of eye fixations. As it has proven difficult to characterize this correspondence theoretically, here we propose instead an empirical method in which the correspondence is learned as participants perform a calibrated series of eye movements. The calibrated series of movements provides us with the spatial distribution of the associated EOG and coincident EEG activity across the scalp for a known set of eye motions. From these we compute the correspondence between arbitrary eye movements and changes in the electrical potentials at the electrode sites and then relate those to the  $x, y$  coordinates of the eye fixations. The correspondence is computed in the form of a linear filter and can be applied to new EOG/EEG data (i.e., data not used in deriving the filter) in order to extract arbitrary, unknown eye movements or the corresponding  $x, y$  coordinates of the eye fixations.

The accuracy of bioelectric-based measurement methods depends greatly on the attention paid to practical issues such as the preprocessing needed to filter out noise and to compensate for various distortions in the data. A challenging practical issue we encountered in implementing our method stems from the fact that our laboratory uses AC amplifiers in data recording. Because AC amplifiers are commonly used for EEG and EOG recordings, this is a relevant issue to discuss here. AC amplifiers attenuate the DC and slowly varying components in the data at the rate determined by the amplifier time constant. Hence EOGs recorded via AC amplifiers are visibly distorted. Some methods for correcting AC distortions in EEG recordings can be found in Elbert and Rockstroh (1980) and Ruchkin (1993). We conducted an independent study on this issue, which is presented in Joyce, Gorodnitsky, Teder-Sälejärvi, King, and Kutas (2002). We find that the quality of the amplifiers in various EEG systems can vary significantly, and in some EEG systems, the filter characteristics of AC amplifiers can vary considerably across individual channels, even when each houses an identical AC amplifier circuit. In these instruments, the amplifier responses also differ from the manufacturer-supplied specifications. It is therefore worthwhile to experimentally check amplifier responses in any given EEG setup and verify them against the manufacturer-provided specifications. The true amplifier response can be measured in several ways, some of which we

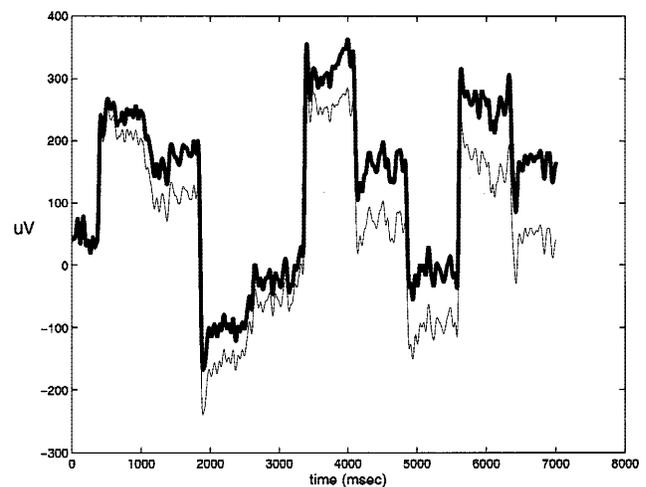
discuss in Joyce et al. If one finds amplifier response characteristics to be inconsistent among system channels or with respect to the manufacturer rating, we suggest compensating for the AC amplifier distortions by estimating the distortion function empirically for each amplifier and for the given filter setting and then applying the estimated correction factors to compensate for the drift. In particular, using single-stage amplifiers, as we have done in our study, we can set the frequency cutoff sufficiently low so that we generate a linear amplifier response in our studies. The linear response is easy to estimate and rectify. We find our method alleviates the problem of signal distortion considerably. The effectiveness of this correction procedure is illustrated in Figure 1. On the other hand, one can use EEG amplifiers specifically rated to provide true DC measurements by some means of DC coupling or DC restoration amplifiers, as was done in Morgan, Patterson, and Simpson (1999), to record saccadic eye movements.

EOG recordings are also sensitive to muscle activity, EEG activity, and ambient electrical noise as well as movements of the eyelids and eyebrows. In the second half of this article, we discuss our approaches to addressing these practical concerns. We settle on a set of preprocessing steps that are easy to automate and implement as a single routine. Although we want to bring to readers' awareness the implementation issues related to using electrophysiological data, these issues are independent of the eye tracking filter that we propose here. We thus present these practical considerations in a separate section.

## Background

Before we present our empirical beam-forming method for tracking eye motion, we review the existing algorithms as background for why we choose an empirical approach here. Using the EOG signal to track eye fixations involves isolating the electrical signals generated by eye movements from all other electrical sources. However, identifying sources of electrical activity using surface-recorded potentials is a difficult problem.

Electric potentials recorded at surface sites represent the sum of many electric signals that propagated to those surface sites and that were generated not only by movement of the eyes, but also that of



**Figure 1.** The thin line shows one channel of EOG data prior to AC-to-DC filtering. The thick line shows the same channel after restoring the DC signal that gets warped by the AC amplifier.

eyebrows and eyelids, by brain activity, twitching muscles, and sometimes the heartbeat, among other things. The way these signals propagate to the scalp is not trivial. Electric signals travel through the various tissues of the head and face before reaching the scalp surface and thus are attenuated to varying degrees depending on the conductive properties of the tissues. The exact distribution and shape of the tissue through which the signals travel are particular to each individual, and without knowledge of the exact tissue geometries, we cannot estimate the effects these tissues have on the signals. Thus, short of obtaining a detailed map of tissue distribution in a person's head, we cannot know exactly how the signals attenuate to yield the data that we record. Moreover, in strict mathematical terms, the problem of converting surface potentials into information concerning their underlying sources has long been known to be ill posed and thus to possess no unique solution. In other words, even with complete knowledge of electrical and geometrical properties of the head, we cannot unequivocally determine from surface potential information alone the sources that generate these potentials.

Although numerous methods have been proposed to resolve the problem either by restricting the form of the solution or by looking beyond the time domain information, for example, differentiating signals based on frequency and phase information, to date no technique can claim complete and accurate source identification. Because the signal of interest and the interfering signals have overlapping frequency content, the latter set of methods does not work well for finding eye movements, that is, there is no bandpass that will leave the signal of interest unabated while excluding the unwanted electrical activity.

Among the first category of methods, those that work by restricting the particular form of the solution, there are: (a) source modeling, the most well known of which are dipole modeling methods (e.g., Berg & Scherg, 1991); (b) solution modeling via optimization constraints, for example, minimum norm constraints; and (c) component decomposition, most notably principal components analysis (PCA; Berg & Scherg, 1991) and blind source separation (BSS), with independent component analysis (ICA; e.g., Vigario, 1997) being a subclass of BSS methods. The latter group of methods decomposes data into some set of mathematically defined components. The references cited here are those that relate to analysis of ocular activity rather than to EEG analysis, or to the methodology per se.

As the eyes can be considered two dipolar sources, one can imagine applying source localization methods from the first two groups of methods to extract the ocular sources and then identifying changes in the source orientation from the changes in the field strength produced by these sources. The most relevant of these methods are those that can model sources as activations of neurons located within a small, compact area that in a limiting case can approximate a dipole. For a basic discussion of source localization methods see, for example, Koles (1998), and Münte, Urbach, Duzel, and Kutas (2000, pp. 238–242).

To localize sources with high accuracy, one must model the exact conductivity properties of the head through which the electric signals propagate. This requires detailed knowledge of the distribution of tissues, including the geometry of the bone structures for each individual's head and face. Although methods now exist for obtaining realistic head models for localizing EEG sources, the processing time and budget required to implement such approaches renders them impractical for many cognitive studies where the data from many individuals must be assessed. An especially important consideration regarding the method of choice in

this particular case is that the eye tracking application requires precision in identifying dipole orientation. Source localization methods are most commonly evaluated and compared in terms of their degree of spatial resolution, which is known to be limited. EEG/EOG resolution of the strength and field orientation of individual sources is even worse than the spatial resolution. This is because (a) the ability to detect the signal itself depends on dipole orientation, with the electric potential falling off rapidly as a dipole rotates to a more tangential position relative to the head/face surface; (b) contributions to EEG/EOG measurements from neighboring sources can cancel their fields to an unknown extent; and (c) the errors in locating the source translate directly into errors in the estimate of source amplitude/orientation. The performance of such source localization methods is severely hampered, for example, by the fact that the eye movement signals of interest overlap spatially with nearby eye signals generated by eyelid and eyebrow movements. In all, source localization methods do not appear to provide sufficient resolution to be practical for eye tracking.

PCA and BSS do not attempt to localize sources, but instead decompose data into separate components. PCA decomposes data into orthogonal components. Use of an orthogonal decomposition to separate ocular and frontal neural sources is supported by the hypothesis that the dipolar charges associated with these two types of sources are approximately orthogonal to each other. The physical orthogonality of dipole vectors, which are a function of the location and orientation of each source, translates into algebraic orthogonality. Lagerlund, Sharbrough, and Busacker (1997), however, showed that PCA decomposition does not work well for isolating ocular sources. This follows because the orthogonality of sources assumption is weak and brain activity is unlikely to have the form of simple dipoles. ICA and more general BSS methods have only recently been applied to EEG data (see Stone & Porrill, 1998, for a simple introduction to ICA methods). ICA methods attempt to decompose data into statistically independent components. The crucial underlying assumption in applying these methods is that the individual active cortical regions of excitation are assumed to be statistically independent. This assumption, which many find objectionable, precludes feedback between the active areas. Furthermore, the independence assumption clearly does not apply to eye sources; the two eyes (i.e., two different sources) tend to move conjunctively in typical behavioral experiments (i.e., are not independent), and the neuronal activity involved in generating eye motion also is highly likely to be correlated with eye movements. However, some BSS methods, for example, Second Order Blind Inference (SOBI) (Belouchrani et al., 1997), do not require statistical independence of sources and we find that these work better than ICA at separating ocular sources. It is difficult to evaluate the quality of separation by ICA/BSS algorithms because one cannot verify by direct measurements the individual sources that give rise to EEG data. Still, recent validation studies of ICA and BSS methods for electrophysiological source separation can be found in Gorodnitsky (2001) and Gorodnitsky and Belouchrani (2001). Those studies demonstrate that currently available ICA/BSS algorithms cannot extract the true physiological sources from EEG/EOG data.

Even if we were able to reliably resolve physiological sources via BSS, there are two other major drawbacks to using BSS techniques for detecting eye movements. First, BSS techniques can only estimate the size of a particular eye movement relative to the sizes of the rest of the eye movements; in other words, they cannot extract the absolute size of an eye movement. This then makes it difficult to find the actual sizes of the movements without some

kind of reference. Second, the independent components identified by BSS are not ordered in a way that reliably specifies which ones are the ocular sources in any principled fashion. Although both of these limitations can be overcome in practice by adopting various implementation strategies, we found such extensions of BBS to be less practical and more unwieldy than the method we describe below.

In summary, the existing methods for inferring sources from scalp potentials all rely on some theoretical assumptions about the sources themselves. By contrast, the method we propose does not rely on making *any* assumptions about sources. Instead, we ascertain the correspondence between the surface potentials and the behavior of ocular sources empirically. We accomplish this by asking each subject to make a certain set of eye movements by following a visual pattern (template) on a computer screen and then relating these eye movements to the changes in the measured potentials. In so doing, we “learn” a general correspondence, in the sense that it can be applied to data generated by arbitrary eye movements and not just those used to determine the correspondence initially. The current method goes beyond merely correlating ocular movements with EOG measurements as it also adapts so to ignore signals coming from sources other than the eye socket. To facilitate creating a filter that disregards electric signals due to brain activity, we use a complementary set of EEG electrodes measuring neural activity. As we later show, this is an important component that helps our method achieve its high accuracy. However, although the filter “learns” to disregard nonocular sources, such as neural sources and heartbeat, that are present during the template viewing, it does not adapt to those that are not present during the template viewing, such as muscle twitches, that might appear spontaneously during an experiment.

## Methods

To derive the filter, we start with the standard model relating brain sources to the scalp potential data. Conceptually, we can model the slow bioelectric currents as distributions of infinitely small quasi-stationary electric dipoles embedded within the conducting medium of the head. Some bioelectric sources, such as the retino-corneal poles, can be assumed to be approximately single dipoles. The exact details of such a model are irrelevant to our purposes here. In mathematical terms, given  $N$  scalp electrodes and a single quasi-stationary dipole somewhere in the head carrying a charge  $S_i$ , we can write the received signal  $D_j$  at these electrodes as

$$D_j = \mathbf{F} * S_i, \quad (1)$$

where  $\mathbf{F}$  is a so-called propagation matrix. Each row of  $\mathbf{F}$  is called a Green’s function.  $F_i$  tells how much of a signal at the sensors is generated from a charge  $S_i$ .

The contribution from all the dipole charges in the head is additive, meaning that we can write the total signal  $\mathbf{D}$  measured by the scalp electrodes as

$$\mathbf{D} = \mathbf{F} * \mathbf{S}, \quad (2)$$

where  $\mathbf{S}$  denotes a  $K \times N$  matrix of active dipole charges. The columns of  $\mathbf{S}$  are a series of snapshots of activity over time. The  $M \times K$  matrix  $\mathbf{F}$  describes the conductance effect of the medium, that is, an individual head, on the propagating electric signals.  $\mathbf{D}$  is an  $M \times N$  matrix whose columns are snapshots of surface potentials at  $M$  electrodes over  $N$  time points. Equation 2 is the well-

known Lead Field equation, details of which can be found in Burger and van Milaan (1946). Here it simply serves as a starting point for our derivation.

We are interested in finding a filter,  $P$ , that can transform the surface potential data  $\mathbf{D}$  of Equation 2 into the  $x, y$  coordinates of the fixation points on some surface that the subject is viewing. Let the  $x, y$  coordinates of the eye fixations over a period of  $N$  time points be designated by a  $2 \times N$  matrix  $\mathbf{O}$ . The first row of matrix  $\mathbf{O}$  contains the sequence of  $x$  coordinates (horizontal movement) and the second row the corresponding sequence of  $y$  coordinates (vertical movement) of eye fixations, which are sampled over  $N$  time points. Then the filter  $P$  that we seek satisfies the following equation:

$$\mathbf{O} = P * \mathbf{D}, \quad (3)$$

where  $\mathbf{D}$  is a matrix of data as described above.

The rows of the matrix  $\mathbf{O}$  are simply a subset of the rows of  $\mathbf{S}$ , converted into  $x, y$  coordinates. Therefore, the filter  $P$  is inversely related to the matrix  $\mathbf{F}$  of Equation 2. The significance of this is that  $P$ , like  $\mathbf{F}$ , depends only on the physical properties of the head and not on the sources themselves. Once learned, it can be used to find any eye movements.

We find filter  $P$  as follows. We first present calibration patterns (predetermined sequences of fixations) on the viewing screen; we ask participants to track these with their eyes. Fixation points are displayed one at a time. As a new fixation appears, the previous one disappears and participants move their eyes to the new fixation location. Approximately 1–2 s are spent fixating on each point before it disappears and the next one appears in a new location. As these known eye movements are performed, we record the EOG and EEG data, denoted here by a matrix  $\mathbf{Dm}$ . The  $x, y$  coordinates of the calibration patterns are also stored in the  $\mathbf{O}$  matrix, as specified above. The filter  $P$  is then computed as

$$P = \mathbf{O} * \mathbf{Dm}^+, \quad (4)$$

where  $\mathbf{Dm}^+$  is the Moore–Penrose generalized inverse or pseudo-inverse of  $\mathbf{Dm}$  (Barnett, 1990). For matrices that have full rank and contain more columns than rows, such as the  $\mathbf{Dm}$  matrix here, the pseudoinverse is defined as  $\mathbf{Dm}^+ = \mathbf{Dm}'(\mathbf{Dm} * \mathbf{Dm}')^{-1}$ , where  $(\mathbf{Dm} * \mathbf{Dm}')^{-1}$  denotes the regular matrix inverse. This expression, however, does not provide a practical way of finding the pseudo-inverse of a large matrix that also contains noise. In the appendix, we discuss a practical way of computing the pseudoinverse that also regularizes (i.e., prevents noise from dominating) the solution.

The filter  $P$  then can be applied to any new data in order to find the  $x, y$  coordinates of the fixations. Let  $D_{\text{new}}$  denote some new data, digitized over  $N$  time points. The  $2 \times N$  matrix  $\mathbf{XY}$  containing the eye fixation coordinates over  $N$  time points is then simply found as

$$\mathbf{XY} = P * D_{\text{new}} = \mathbf{O} * \mathbf{Dm}' * D_{\text{new}}. \quad (5)$$

Matrix  $\mathbf{XY}$  is analogous in structure to the  $\mathbf{O}$  matrix in that the first row contains the  $x$  (horizontal movement) and the second row contains the  $y$  (vertical movement) coordinates of the eye fixations on the viewing screen.

The experimental results we present in the next section validate our proposal that the propagation filter  $P$ , derived with only one training pattern, is a general filter capable of extracting arbitrary eye fixations from the recorded EOG and EEG data.

Note that the propagation filter  $P$  is specific to both the electrode arrangement and to the structure of an individual's head. If the electrodes are reapplied, their positions may change and the propagation factors for these new electrode locations would need to be determined. Similarly, because individual head structures vary, propagation factors that depend on the individual tissue geometries also vary. Thus,  $P$  must be calculated for each individual and for each session in which electrodes are reapplied. The method presented here, however, can be automated so that computing the filter should not present a major difficulty.

**Results**

An example of the pattern used in our study is shown in Figure 2. The pattern is a sequence of 10 fixation points (small crosshairs), each displayed sequentially on the screen for a duration of 1.5 s. The distance between any sequential fixation points does not exceed 21° vertically, 18° horizontally, and 14° diagonally in this pattern. Such a calibration sequence yields good resolution, but we are currently investigating how performance may be improved were we to use an even greater range of movement directions during filter learning. A detailed discussion of the issues concerning calibration sequence design is given in the next section.

With the appearance of each new fixation cross, the subject is instructed to saccade to the new location at his/her own rate. Viewing is repeated several times with the same pattern. The procedure can then be repeated using one or more different patterns. Although it is not necessary to use more than one calibration pattern, doing so may produce a more robust filter estimate. As the subject tracks a pattern, we calculate the filter by associating the changes in the  $x, y$  coordinates of the pattern with the recorded changes in the EEG and EOG, as explained in the section above. The filters for each individual trial are then averaged.

We can estimate the accuracy of the propagation filter by determining how well it generalizes to the analysis of random sequences of eye movements. Specifically, after creating an average filter using a pattern such as that shown in Figure 2, filter accuracy can be evaluated on data from different patterns that were

not used to create the filter. The known coordinates of the new templates can be compared to the fixation coordinates extracted using the filter to give us an estimate of the error we can expect in fixation localization.

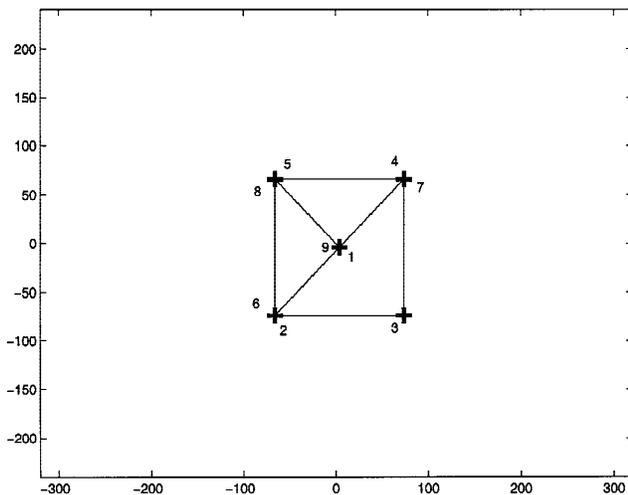
The following procedure was used to compute experimental error bounds. A participant viewed the Figure 2 pattern 30 times and another pattern (similar in dimension but diamond rather than square shaped) 10 times. A propagation filter was computed using 30 square pattern trials. This filter was then applied to each of the 10 diamond pattern trials. Differences between the fixation coordinates marked on the screen and those extracted from the bioelectric data were computed for each trial. This procedure was applied to the data of 13 participants. For ~18° horizontal eye movements, we obtained an error of ~1.2° ( $SD$ : ~0.34°, range: 0.74–3.4°). For ~21° vertical eye movements, we computed an error of 2.2° ( $SD$ : ~0.68°, range: 1.8–5.2°). Note that error rates tend to increase with larger eye movements, and that errors tend to be larger for eye movements along the vertical axis than the horizontal axis.

The calculated error represents the sum of four errors—one coming from the eye fixations themselves not following the pattern exactly, the second occurring because the method of head stabilization we use still permits small head movements, and the third coming from inconsistency in the relationship between the EOG output and the angle of gaze. For example, the relationship between EOG output and the horizontal angle of gaze for a ±30° arc is accurate to within ±1.5–2.0°. The relationship between EOG output and other directions of gaze is accurate within approximately the same limits (Young & Sheena, 1975). The fourth error stems from multiple sources of noise in the data, including new sources such as spontaneous muscle twitches that may arise during an experiment, but not when the template is viewed. Naturally, these four sources of error are also present when we first compute the filter. Thus one can think of these errors as contributing twice to performance analysis, once when we obtain the filter and a second time when we calculate fixation points from data for a new pattern. Given that the resolution of EOG output itself is only within ±1.5–2.0°, we find the performance obtained with this method is rather impressive.

**Practical Issues in Tracking Eye Fixations with Bioelectric Data**

A number of practical issues must be considered before any method based on electric potentials can be put into operation. In the Introduction, we addressed the issue of recording EOG signals via AC amplifiers. In this section, we begin by discussing issues related to data collection, including the number and placement of both EEG and EOG electrodes, the consequences of head motion, and practical issues concerning the creation of a good calibration pattern. Next, we provide our approach to preprocessing the data, including prefiltering, artifact correction, and changes in baselines across the course of an experiment. Finally, we offer some suggestions on how one might improve the accuracy and robustness of the eye tracking filter, which can be affected by intertrial variability.

Although we wish to bring to readers an awareness of the implementation issues related to using electrophysiological data, it is important to note that our choices for dealing with these issues, whether it be the choice of amplifiers to record DC signals or the methods of prefiltering the data, are not an integral part of the eye tracking filter that we propose here. Thus it is important to distinguish between these two different aspects of our proposal and we



**Figure 2.** The pattern viewed by subjects during the EOG calibration trials. Numbers indicate the order in which fixations locations were sequentially presented.

ask that readers consider the preprocessing procedures in view of their relevance to their own setups. Naturally, implementation procedures can be improved further. Some of these problems can be removed altogether, for example, those due to AC amplifier distortions, simply by using better equipment. That said, that our method performs within the limits of the resolution of EOG data says that any further gains from preprocessing procedures will at best be minimal.

### Data Collection

In this section, we discuss issues related to data collection, namely, (a) the number and placement of the EEG and EOG electrodes, (b) the problem of head movement, and (c) the characteristics of the calibration pattern needed for computing a reliable propagation filter.

*Practical aspects of creating patterns for learning the empirical filter.* The basis for the proposed method is to design an empirical filter by recording EOG/EEG activity as subjects track a known pattern presented to them on the screen with their eyes. A key consideration here is to use patterns that the subjects can track cleanly, in other words, without making extraneous eye movements. For instance, we found that tracking a point being slowly traced on the screen is very difficult and leads to small, jittery movements of the eye and many overshoots beyond the point that is being traced. This is true both when subjects make large movements (greater than 20° horizontally and 25° vertically) and very small movements. People's eyes have a tendency to overshoot (and in the case of large movements also undershoot) the fixation endpoint, going a little past it, followed by an adjustment movement back to it. This is presumably because moving over small distance does not provide sufficient time for the eyes to accelerate and decelerate to a stop and in the case of making large movements, having a new fixation point a large distance away from the old one makes it difficult to estimate when to stop. Also, over large distances, people do not move their eyes in a straight line, tracing a small arc instead. Large eye motion is also not ideal because it tends to elicit accompanying head motion. Finally, large eye movements made while the head motion is restricted tend to cause eyestrain, which in turn leads to minute, jittery eye movements around fixation points and elsewhere. Any deviation from what the eye should be fixating on during training contributes to the error in the computation of the filter  $P$ . Even minute errors in the eye movement angle that are almost immeasurable nonetheless translate into significant errors in fixations by the time they are converted into coordinates on a screen that is located a significant distance away from the eyes.

Taking these factors into consideration, optimal patterns should elicit moderately small eye movements where the subjects are allowed to saccade from one point in the pattern to the next at their own pace rather than having the movement trajectory traced for them on the screen. We verified through a series of test that this design is the easiest for a subject in terms of being able to move the eyes along a straight line. With this design, we can then extrapolate movements between the fixation points along a straight line.

To assure good filter precision, the patterns also need to cover all the basic movement directions. Directions included in our patterns are up, down, right, left, and diagonally across the screen, with only one or two repetitions of each direction. Figure 2 shows an example of the pattern used in our study.

A final note on the design. The timing of the presentation of the successive fixation points needs to be moderately paced. If the

sequence of fixation points occurs too quickly, subjects may not be able to track them cleanly. On the other hand, if the fixations need to be held for a long duration, then the minute, uncontrollable adjustments of the eye needed for maintaining fixation will introduce unwanted noise into the signal. As a rule, it is impossible to keep the eyes completely stationary in any one position for any length of time.

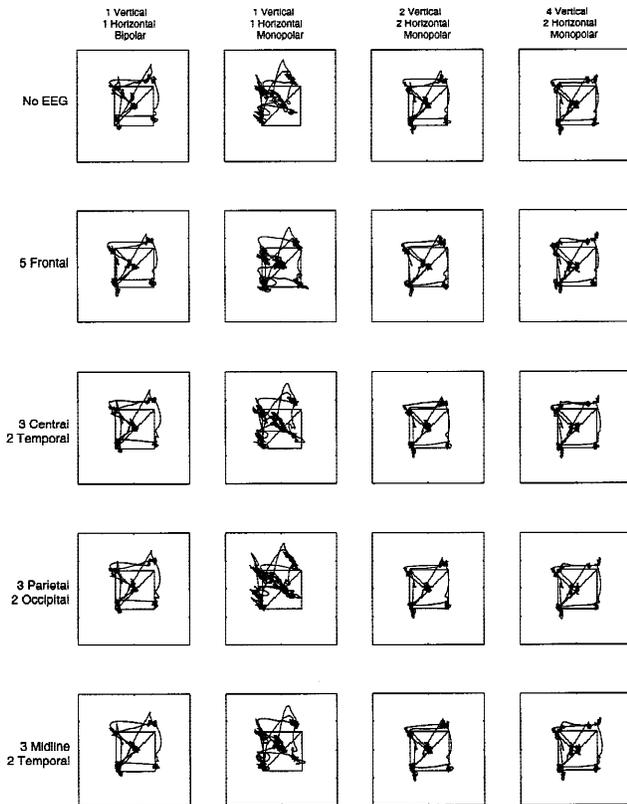
*Electrodes.* The empirical filter,  $P$ , not only learns to map the eye movements from the EOG recordings but also to ignore all signals that come from directions other than the eye socket. As discussed above, to improve the filters ability to disregard electrical signals due to brain activity, clean measurements of this activity should be available for calculating the filter. The filter can be optimized in this regard if EEG electrodes are placed to maximize the sampling of brain activity that is likely to be picked up by the EOG electrodes around the eyes.

The electric field caused by ocular movement is strong in the vicinity of the eyes, but attenuates differentially away from the eyes across different parts of the head. Figure 3, column 2, illustrates that relatively veridical, although somewhat noisy, eye movement signals can be extracted with only two EOG electrodes and a complement of EEG electrodes; namely, using one vertical electrode above the left or the right eye and one horizontal electrode at the outer canthus of the left or right eye. Column 2 shows, in order, eye movement data recorded using two eye electrodes each referenced to the left mastoid together with no EEG channels, five frontal EEG channels, three central and two lateral temporal EEG channels, three parietal and two occipital EEG channels, or three midline (frontal, central, parietal) and two lateral temporal EEG channels. Adding the five frontal EEG channels noticeably improves the extracted eye movements, and does so to a greater degree than any other combination of EEG channels. Adding a complementary vertical electrode (so that one is above the eye and one is below the eye) and a horizontal (one on the outer canthus of each eye) electrode, as in column 3, noticeably cleans up the extracted eye movement signals. In fact, the addition of EEG channels to these four EOG channels does little to improve the extracted eye movements. However, because noise is unpredictable, addition of some frontal EEG channels is advisable to ensure the cleanest possible signal. Column 4 illustrates that the addition of another pair of electrodes vertically, above and below the other eye does little to improve the performance of the propagation filter.

Comparing column 3 with column 1 in Figure 3 further illustrates that better separation of EOG sources is obtained with monopolar (eye to mastoid) than bipolar EOG recordings. This follows from the fact that any offset in vertical or horizontal alignment of EOG electrodes will have no effect on monopolar recordings, whereas it can corrupt the horizontal and vertical measures when respective electrode readings are subtracted from each other.

Thus, we recommend using the following monopolar electrodes: two vertical EOG electrodes (one above and one below the eye), two horizontal EOG electrodes (on the outer canthus of each eye), and five well-spaced frontal scalp electrodes.

*Head movement.* It is natural for the head to move together with eyes in the direction of gaze. These movements are mostly unconscious and individuals therefore cannot restrain from making them to some degree. A slight angle difference in eye movement due to the head adjustment becomes a large error by the time it is



**Figure 3.** The extracted x,y coordinates for one trial of EOG data. The filters in rows 1–5 were created by adding (a) no EEG channels, (b) five frontal EEG channels, (c) three central and two lateral temporal EEG channels, (d) three parietal and two occipital EEG channels, or (e) three midline (frontal, central, parietal) and two lateral temporal EEG channels to the EOG channels. Columns 1–4 illustrate eye movement data extracted when EOG information came from (a) one vertical (upper referenced to lower) and one horizontal (left referenced to right) bipolar EOG channel, (b) one vertical and one horizontal monopolar EOG channel (referenced to mastoid), (c) two vertical and two horizontal monopolar EOG channels, or (d) four vertical and two horizontal monopolar EOG channels. Data are plotted over template coordinates.

translated into the screen fixation coordinates located several inches away.

We chose to reduce head movements by stabilizing each subject’s head via a personalized dental impression palate fixed to a stationary surface, which the subject bites down on for the duration of the experiment. In the long run, correcting for head movement with the aid of a head tracker is a more accurate, less intrusive, and more practical method. Several relatively inexpensive methods, such as optically tracking a point on the head, provide sufficiently high resolution and could be readily integrated with the EEG/EOG recordings proposed here (e.g. Origin Instruments Corporation, 2000).

**Data Processing**

Surface potential data must be preprocessed in several steps. These include, in order, (1) filtering out noise, (2) correcting for blinks, (3) adjusting the data in time to account for the delay in a subject’s response to the fixation stimulus, (4) correcting for the distortions in the data introduced by the AC amplifier circuit, and (5) base-

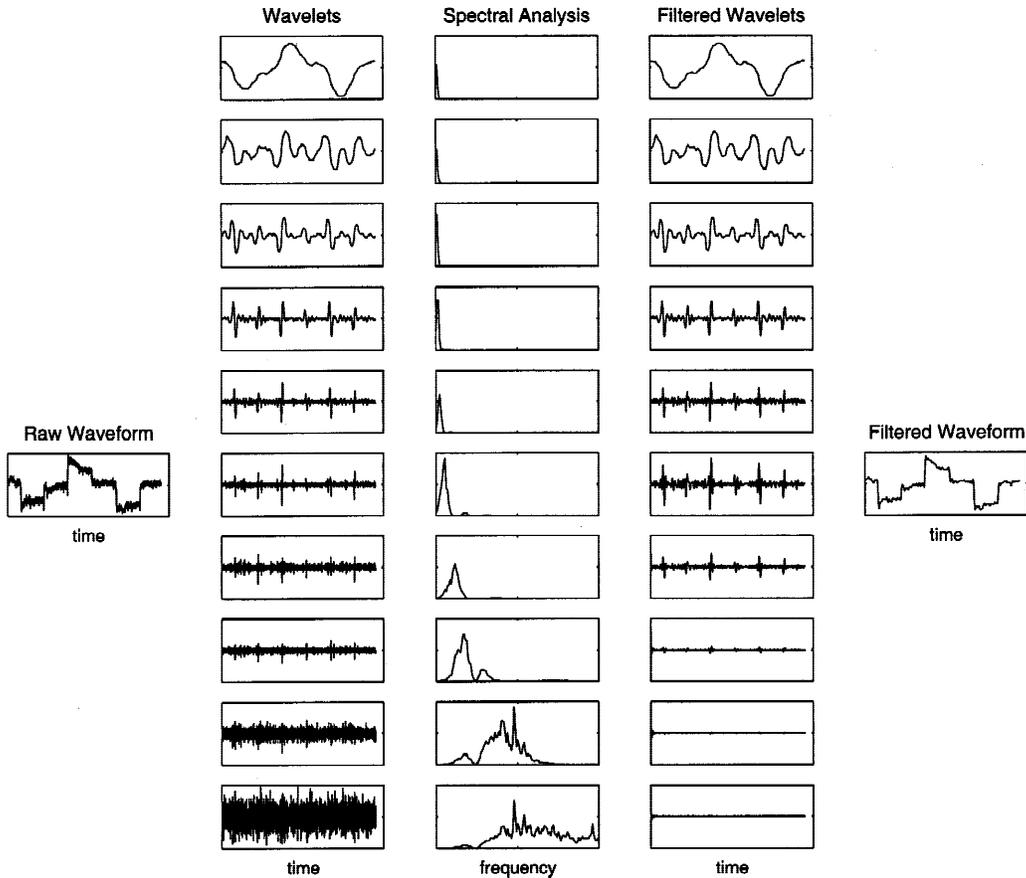
lining the recorded data. Note that steps 1, 2, 4, and 5 must be carried out identically on calibration data prior to using it to create the filter, and on the experimental data prior to applying the filter for extracting eye movements. Step 3, adjusting for time delay, only needs to be performed on calibration data prior to filter creation.

*Filtering.* Electrophysiological data contain a fair amount of high frequency artifacts, generated either by the equipment or within the body. Filtering out some of these artifacts yields a better propagation filter, although with some caveats as detailed below.

Bioelectric data are notoriously difficult to filter because they are nonstationary and because the frequency content of the signal often overlaps with that of the noise. The EOG signals due to ocular movement contain very low frequency, DC-like components, punctuated by bursts of quick amplitude changes of high frequency content that we cannot afford to lose by digital filtering in the frequency domain. Due to this dichotomy in the frequency content specific to the ocular movement signals, the following wavelet-based filtering method was found to work well. Note, however, that this method filters data quite heavily and is designed specifically for these types of signals. It is not appropriate for processing most other EEG signals.

First, the data at each sensor from each trial are decomposed into several wavelet levels using a biorthogonal wavelet decomposition (Sweldens, 1998). A spectral analysis is then performed on all the wavelet levels to determine for which the maximal contributing frequency exceeds 5 Hz. These levels are then low-pass filtered at 8 Hz using a Butterworth filter. This filter eliminates the high frequency noise at these wavelet levels while retaining their lower frequencies that may contain slow wave information that is crucial for maintaining information about the amplitude and morphology of the eye movements. We determine which of these filtered levels contains meaningful slow wave signals in the following manner. First, the means of each filtered level and the remaining unfiltered levels are computed and divided by the sum of all the means; this results in a number that represents the proportion of the overall signal power that each level contributes. A level is eliminated only if the proportion of power contributed by that *filtered* level is less than 0.0001. Finally, the “clean” data are reconstructed by combining the remaining filtered levels with all unfiltered levels. Note that this procedure does not eliminate all frequency components above 5 Hz from the data, as unfiltered levels that are dominated by lower frequencies still contain high frequency information associated with the sharp onsets and offsets of saccades. Figure 4 illustrates how this method was applied and shows that the filtered signal retains the high frequency information necessary to preserve the morphology of saccades while at the same time eliminating the non-saccade-related high-frequency noise.

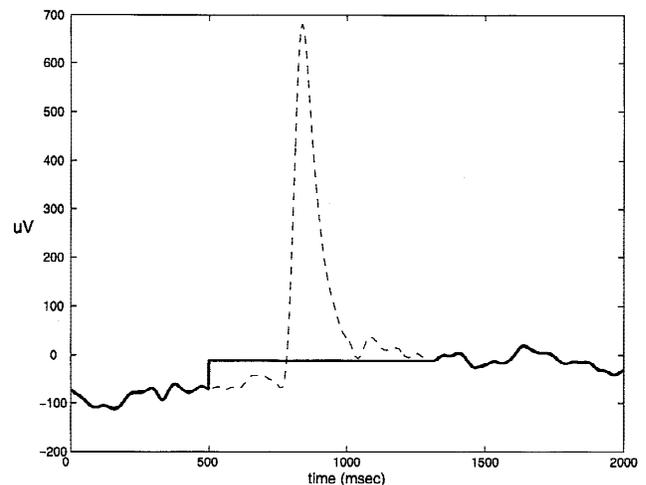
Two points must be noted here. First, the Haar wavelet decomposition is designed to filter boxcar-like ocular motion signals and thus may be constructed as a better filtering method for our eye movement data. The issue, however, is not entirely straightforward for this application. The filtering must be applied very carefully so as not to eliminate the artifactual signals from data which we actually would like the filter to learn to disregard when it later encounters them in the experimental data. The second point is that it is essential to apply the identical filtering process during the learning stage when filter is being created and subsequently when it is applied to new data. We found that the filtering method described here provides the necessary consistency in the filtering process regardless of the



**Figure 4.** Wavelet filtering sequence. The first column contains one channel of unfiltered EOG data. The second column contains the levels of the biorthogonal wavelet decomposition ordered from the lowest to highest frequency bands. The third column contains the spectral analysis of each wavelet level. The peak of the curve lies at the frequency contributing maximally to that level. The fourth column contains the unfiltered low-frequency levels (the top five levels) and the filtered high-frequency levels (the bottom five levels). In this example the bottom four levels were eliminated due to low power contribution to the low-frequency signal. The fifth level was retained in its filtered form (because it contained enough low-frequency content) and recombined with the top five levels to create the filtered data in the fifth column.

taxonomy of the eye movement pattern. The issue of optimal filtering must be investigated further, however.

*Blink correction.* Blinks are difficult to discriminate from ocular motion because both originate, at random times, from essentially the same spatial location, and because blinks produce electric signals analogous to vertical saccades. Thus we eliminate blinks prior to creating and applying the filter. We use a combined polarity inversion detection and amplitude cut-off criterion between one upper and one lower vertical EOG electrode. The onset and offset of each blink are determined by finding the point at which the slope decreases to a certain level before and after the point of maximum polarity inversion, respectively. To ensure that double-peaked blinks will be corrected accurately, we set an additional criterion whereby the amplitude of the onset and offset points has to be within one standard deviation of the data mean. Correction consists simply of flattening the signal between the marked onset and offset points (see Figure 5). One can make the correction either with a sloping line or by putting in a step, as in Figure 5. In general, the exact path of the eye movement



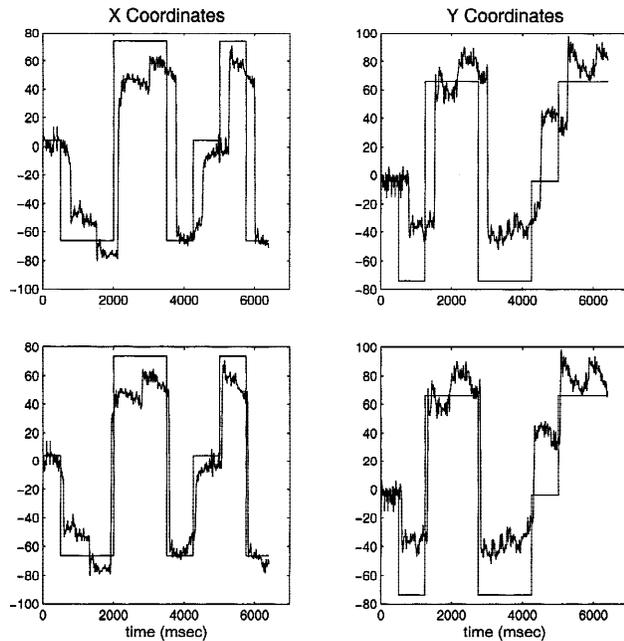
**Figure 5.** One channel of vertical EOG data. The dashed line shows the original blink-contaminated data. The solid line shows the line of correction resulting from the blink correction procedure.

during a blink is neither known nor important to track. What is important is to recover the change in eye fixation after the blink, which our method provides. We choose to use a step function in our implementation, as it is quite likely that a saccade would occur during a blink and a step function is a more appropriate way to reflect this.

*Temporal alignment of the data and the model.* It can take up to a few hundred milliseconds for a person to program and initiate a saccade to each new fixation location in the learning pattern sequence. As a consequence, there is a temporal delay between the eye movements as predicted by the model and the actual movement of the eye (Figure 6, row 1).

We use an automated method to temporally align the two sequences. The time point of the first deflection from the initial fixation coordinates in either the X or Y direction is taken as the beginning of the first “saccade” in the calibration pattern. The algorithm then looks for the largest amplitude change occurring in the data channel within the 400-ms window following that calibration saccade along the corresponding axis (a horizontal channel for X and a vertical channel for Y). The difference in the latency of onset between the first model saccade and the corresponding first saccade in the recorded data is calculated and all data channels are shifted in time by that amount (Figure 6, row 2).

*Baselining the data.* Baselining is a standard step in processing EEG data. The mean amplitude of activity during the prestimulus interval is calculated and subtracted from all of the data points in the recorded epoch. Baselining is performed separately for each sensor for each trial. This is the final adjustment of raw data before the filter is created.



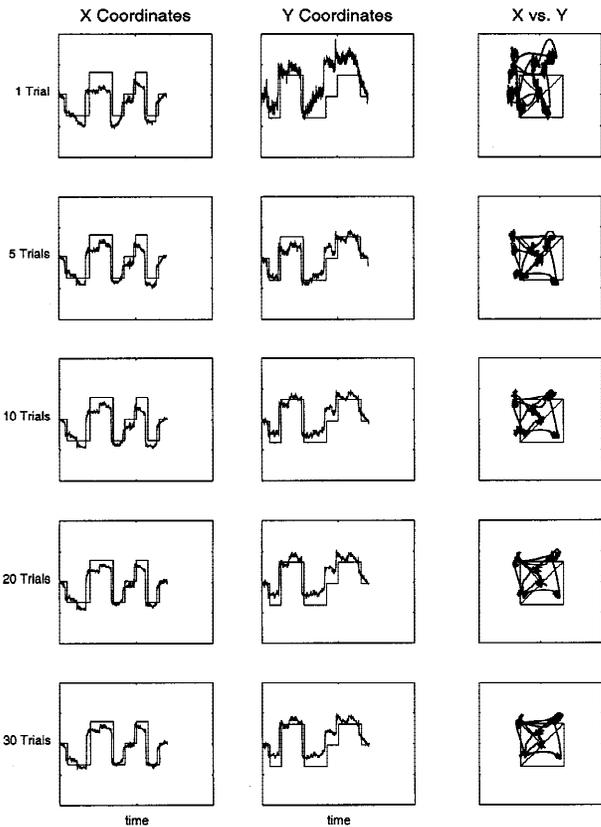
**Figure 6.** X (column 1) and Y (column 2) coordinates extracted from one trial of data and plotted over the model coordinates. The first row shows the temporal difference between the model and the data due to the time required for saccade programming and execution. The second row shows the same data following temporal alignment.

**Construction and Generalization of the Filter**

Once the data from pattern tracing have been preprocessed, it is a fairly simple matter to apply Equation 3 to compute the filter. We have found, however, that eye movements in an individual calibration sequence may contain too much jitter because the eyes deviate from a template pattern in a random fashion. To reduce the impact of this random variability, we ask subjects to view the template several times. Separate filters are created for each viewing, using Equation 3, and averaged together to derive the final propagation filter. As a rule, subjects become practiced, and therefore more accurate over time, at tracking the template pattern. Figure 7 illustrates the noise reduction and improvement in eye movement extraction when the filter is computed using several repetitions of the same pattern. Because the filter generalizes to any movement sequence, an even better approach might be to use multiple viewings of several different pattern sequences, create separate filters for each, and average the resulting filters across all templates. This may reduce the contribution of noise in the eye movements due to a particular choice of template.

**Data Registration**

We attempt to control for the appearance of spurious slow drifts in the recorded data at several stages of data processing. Nonetheless, even after such precautions are taken, errors due to head movements and residual amplifier drift can remain in the data. We thus further correct for these by asking subjects to fixate at a known



**Figure 7.** X, Y coordinates extracted using filters created from 1 trial (row 1), 5 trials (row 2), 10 trials (row 3), 20 trials (row 4), and 30 trials (row 5). Extracted X and Y coordinates are shown alone (columns one and two) and plotted together (column three), over the model coordinates.

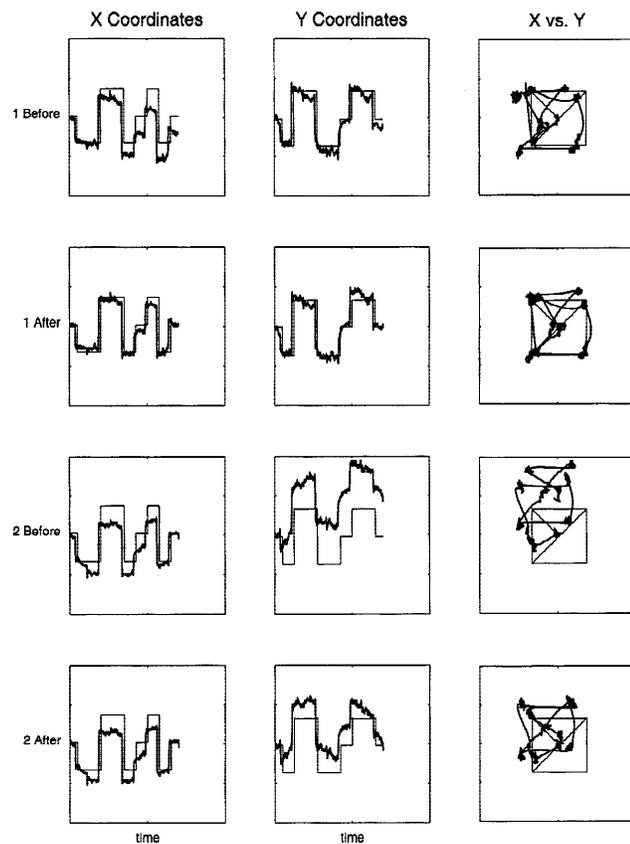
location both at the beginning and the end of the stimulus presentation. Two regression lines are then calculated: one between the initial and final fixation coordinates in the extracted eye movements, and another between the known initial and final fixation coordinates. The difference between these regression lines is subtracted from the extracted eye movement coordinates with the consequence of aligning the extracted fixation coordinates with the actual fixation coordinates.

Application of this procedure is shown in Figure 8. The displacement being corrected is quite large in this example and is primarily due to the head movement. As we state above, this major artifact in our data could be eliminated with the use of a head tracker.

This procedure often corrects the coordinate discrepancies in extracted eye movements (Figure 8, rows 1 and 2), although there are obviously cases where the shifts in head position, for example, do not occur in a consistent manner throughout the experiment, and hence slow potential artifacts cannot be fully corrected (Figure 8, columns 3 and 4).

### Saccade Detection

The steps outlined above attenuate much of the residual noise in the extracted eye movements. To clean up any small, irrelevant perturbations that may remain, we perform one additional step, which involves detecting and delimiting saccades in the eye movements.

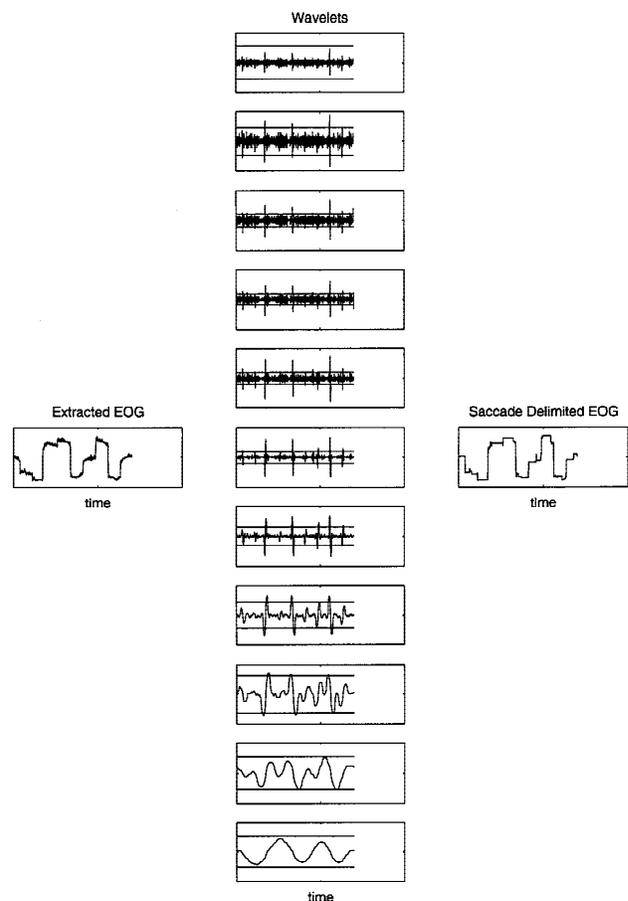


**Figure 8.** Extracted X,Y coordinates plotted over the model. Rows 1 and 2 show the same trial before and after registration, respectively. Rows 3 and 4 show another trial before and after registration. Notice in row 4 that some types of drift error cannot be corrected by performing registration.

Each eye movement source is first decomposed into its wavelet levels (as described above). Saccade onsets and offsets appear as sharp peaks in many of the wavelet levels (Figure 9, column 2); they can be detected automatically by searching for values within each wavelet level that exceed 2 times its standard deviation (horizontal lines in Figure 9, column 2). We use the coordinates of the fixations at the time of these saccades and eliminate everything else in the eye movement sequence to reduce it to a box-like time series that represents only the major saccades (Figure 9, column 3).

### Discussion

This report outlines a method for tracking an individual's eye fixations (the  $x,y$  coordinates) on a viewing screen using electrical signals recorded via EOG and EEG electrodes. The method makes use of an empirically derived filter, which is trained on a known sequence of eye movements and which can then be used to analyze random patterns of eye motions. This method has many advantages. For investigators already recording EEG data, little additional time, effort, and expense are required to simultaneously record and then extract eye fixation coordinates. The method is



**Figure 9.** The saccade delimiting procedure. Column 1 shows one channel (X coordinates) of extracted eye movement data. Column 2 shows the wavelet decompositions of that channel. Notice that the peaks in the wavelet levels correspond to saccade onsets and offsets in the extracted coordinates. Column 3 shows the data reduction resulting from saccade delimiting.

shown to perform within the limits of accuracy with which EOG output reflects the angle of gaze. Such resolution is sufficient for many types of experimental studies. A number of practical issues must be considered when using electric potential data. We present a set of preprocessing steps that are easy to automate and that can be implemented as a single routine along with the proposed filter

to obtain the performance reported in this paper. Further improvements in implementation are feasible. By far the biggest improvement would be to incorporate a high-resolution method for tracking head movement, which would not only correct for head motion artifacts but also make the proposed method support more natural viewing of stimuli.

## REFERENCES

- Barnett, S. (1990). *Matrices, methods, and applications*. New York: Clarendon Press.
- Belouchrani, A., Abed Meraim, K., Cardoso, J. F., & Moulines, E. (1997). A blind source separation technique using second order statistics. *IEEE Transactions on Signal Processing*, 45, 434–444.
- Berg, P., & Scherg, M. (1991). Dipole models of eye movements and blinks. *Electroencephalography and Clinical Neurophysiology*, 79, 36–44.
- Burger, H. C., & van Milaan, J. B. (1946). Heart-vector and leads. Part I. *British Heart Journal*, 8, 157–161.
- Cornsweet, T. N., & Crane, H. D. (1973). Accurate two-dimensional eye tracker using first and fourth Purkinje images. *Journal of the Optical Society of America*, 63, 921.
- Ditchburn, R. W., & Ginsborg, B. L. (1953). Involuntary eye movements during fixation. *Journal of Physiology*, 119, 1–17.
- Eadie, A. S., Pugh, J. R., & Heron, G. (1994). The measurement of small eye movements using an infra-red limbus reflection technique. In G. d'Ydewalle and J.V. Rensbergen, (Eds.). *Visual and oculomotor functions: Advances in eye movement research* (pp. 409–421). Amsterdam: Elsevier Science, North Holland.
- Elbert, T., & Rockstroh, B. (1980). Some remarks on the development of a standardized time constant. *Psychophysiology*, 17, 504–505.
- Gorodnitsky, I. F. (2001). A study to validate Blind Source Separation algorithms for EEG analysis. <http://cloe.ucsd.edu/BSS-validation.html>.
- Gorodnitsky, I. F., & Belouchrani, A. (2001). Joint cumulant and correlation based signal separation with application to EEG data analysis. Third International Conference on Independent Component Analysis and Signal Processing, No. 2, San Diego, CA. <http://ica2001.ucsd.edu>.
- Joyce, C. A., Gorodnitsky, I. F., Teder-Sälejärvi, W. A., King, J. W., & Kutas, M. (2002). Variability in AC amplifier distortions: Estimation and correction. *Psychophysiology*, 39, 633–640.
- Koles, Z. J. (1998). Trends in EEG source localization. *Electroencephalography and Clinical Neurophysiology*, 106, 127–137.
- Lagerlund, T. D., Sharbrough, F. W., & Busacker, N. E. (1997). Spatial filtering of multichannel electroencephalographic recordings through principal component analysis by singular value decomposition. *Journal of Clinical Neurophysiology*, 14, 73–82.
- Monty, R. A. (1975). An advanced eye-movement measuring and recording system. *American Psychologist*, 30, 331–335.
- Morgan, S. W., Patterson, J., & Simpson, D. G. (1999). Utilizing EOG for the measurement of saccadic eye movements. In *IEEE: Biomedical Research in the Third Millennium* (pp. 33–36). Melbourne: The Food & Packaging Cooperative Research Centre (CRC).
- Moses, R. A. (1975). *Adler's physiology of the eye* (6th ed.). St. Louis: CV Mosby.
- Mowrer, O. H., Ruch, R. C., & Miller, N. E. (1936). The corneo-retinal potential difference as the basis of the galvanometric method of recording eye movements. *American Journal of Physiology*, 114, 423.
- Müente, T. F., Urbach, T. P., Duzel, E., & Kutas, M. (2000). Event related brain potentials in the study of human cognition and neuropsychology. In F. Boller & J. Grafman, (Eds.). *Handbook of neuropsychology* (2nd ed., Vol. 1, pp. 139–258). Amsterdam: North Holland.
- Muller, P. U., Cavegn, D., d'Ydewalle, G., & Groner, R. (1993). A comparison of a new limbus tracker, corneal reflection technique, Purkinje eye tracking and electro-oculography. In G. d'Ydewalle & J. V. Rensbergen (Eds.). *Perception and cognition: Advances in eye movement research* (pp. 393–401). Amsterdam: Elsevier Science, North Holland.
- Ong, J., & Harman, G. A. (1979). Eye movements simultaneously recorded by electrooculographic and photoelectric methods. *Perceptual and Motor Skills*, 48, 619–624.
- Origin Instruments Corporation. (2000). HeadMouse. <http://www.orin.com>.
- Ruchkin, D. S. (1993). AC-to-DC inverse filtering of event-related potentials. In W. C. McCallum & S. H. Curry (Eds.). *Slow potential changes in the human brain* (pp. 91–97). New York: Plenum Press.
- Stone J. V., & Porrill, J. (1998). Independent component analysis and projection pursuit: A tutorial introduction. Technical report: <http://www.shef.ac.uk/~pc1jvs/index.html>.
- Sweldens, W. (1998). The lifting scheme: A construction of second generation wavelets. *SIAM Journal on Mathematical Analysis*, 29, 511–546.
- Vigario, R. N. (1997). Extraction of ocular artefacts from EEG using independent component analysis. *Electroencephalography and Clinical Neurophysiology*, 103, 395–404.
- Young, L. R., & Sheena, D. (1975). Eye movement measurement techniques. *American Psychologist*, 30, 315–330.

(RECEIVED February 20, 2001; ACCEPTED February 14, 2002)

## APPENDIX: COMPUTATION OF THE MOORE-PENROSE INVERSE

The Moore–Penrose inverse of a matrix can be computed with the aid of a rank-revealing matrix decomposition, such as a singular value decomposition (SVD). The SVD of the  $M \times N$  matrix  $\mathbf{Dm}$  in Equation 3 is a product of three matrices

$$\mathbf{Dm} = \mathbf{USV}', \quad (6)$$

where  $\mathbf{S}$  is a  $M \times N$  matrix containing singular values on its diagonal and zeros otherwise and  $\mathbf{U}$  and  $\mathbf{V}$  are square  $M \times M$  and  $N \times N$  matrices containing the left and right singular vectors of the data, respectively. The pseudoinverse of  $\mathbf{Dm}$  is then defined as

$$\mathbf{Dm}^+ = \mathbf{VS}^{-1}\mathbf{U}' \quad (7)$$

There are two issues at stake here. First, when the matrix  $\mathbf{Dm}$  is very “wide,” that is, having many more columns (time points) than rows (the number of channels), its decomposition requires a large number of computer operations, so that computing the pseudoinverse using Equation 7 directly is not advisable. Second, the matrix  $\mathbf{Dm}$  may be ill conditioned, meaning, among other things, that the pseudoinverse would tend to amplify small noise in data. For this reason, it is best to discard the smallest singular values altogether in the matrix  $\mathbf{S}$ . In this process, the right and left singular vectors from  $\mathbf{U}$  and  $\mathbf{V}$  corresponding to the discarded singular values are also discarded. As a side benefit of this procedure, the dimensions of the matrices can be significantly reduced, simplifying the cost of computation. This method is called truncated SVD (TSVD) regularization and a complete theory has been developed around this subject, including how to select which singular values to discard. For further information see Barnett (1990). To deal with the first issue, we suggest the following algebraic trick to speed up the computation of the Moore–Penrose inverse. Instead of computing SVD of  $\mathbf{Dm}$ , we compute only a

partial SVD of its transpose, namely  $\mathbf{Dm}^t$ , as follows. The SVD of the matrix  $\mathbf{Dm}^t$ , expressed in terms of the matrices appearing in Equations 6 and 7 is

$$\mathbf{Dm}^t = \mathbf{V}\mathbf{S}\mathbf{U}^t. \quad (8)$$

From Equation 7, only the first  $M$  columns of  $\mathbf{V}$  are needed to compute  $\mathbf{Dm}^+$ . In the partial SVD of  $\mathbf{Dm}^t$ , we can compute only these first  $M$  columns of  $\mathbf{V}$  and stop the computation after that. Let

$\mathbf{U}_M$  and  $\mathbf{S}_M$  denote matrices containing the first  $M$  columns of  $\mathbf{U}$  and  $\mathbf{S}$ , respectively. The pseudoinverse is then found as

$$\mathbf{Dm}^+ = \mathbf{V}_M\mathbf{S}_M^{-1}\mathbf{U}^t. \quad (9)$$

Combining this partial SVD computation with the truncated SVD procedure discussed above reduces the number of columns of the  $\mathbf{V}$ ,  $\mathbf{S}$ , and  $\mathbf{U}$  matrices used even further, thus reducing the computational complexity.