Removal of Ocular Artifacts from EEG Using Learned Templates

Max Quinn[1], Santosh Mathan[2], and Misha Pavel[1]

[1] Oregon Health and Science University, Portland, OR 97239, USA
quinnma2013@alumni.ohsu.edu, pavelm@ohsu.edu
[2] Honeywell Laboratories, Redmond, WA 98052, USA
santosh.mathan@honeywell.com

Abstract. Electroencephalogram (EEG) data can provide information on cognitive states and processes with high temporal resolution, but to take full advantage of this temporal resolution, common transients such as blinks and eye movements must be accounted for without censoring data. This can require additional hardware, large amounts of data, or manual inspection. In this paper we introduce a greedy, template-based method for modeling and removing transient activity. The method iteratively models an input and updates a template; a process which quickly converges to a unique and efficient approximation of the input. When combined with standard source separation techniques such as Independent Component Analysis (ICA) or Principal Component Analysis (PCA), the method shows promise for the automatic and data driven removal of ocular artifacts from EEG data. In this paper we outline our method, provide evidence for its effectiveness using synthetic EEG data, and demonstrate its effect on real EEG data recorded as part of a minimally constrained cognitive task.

Keywords: EEG, EOG, ICA, PCA, BCI, matching pursuit.

1 Introduction

1.1 EOG Removal from EEG for Cognitive State Estimation

EEG is being actively investigated as a method for creating advanced human-computer interfaces, in which cortical activity patterns are used to infer the cognitive state of users or to control external devices.

Cognitive state estimation systems allow for estimates of latent cognitive properties such as mental workload, estimates which can then be used to regulate information provided to a user as part of an augmented cognition system, or for the development of new methods for assessing specific cognitive deficits and tailoring cognitive rehabilitation. Brain-computer interface (BCI) systems may allow patients with motor disabilities to interact with the world by controlling external devices through cortical activity. Devices include traditional computer control devices (such as mice), text communication systems, and robotic limbs.

D.D. Schmorrow and C.M. Fidopiastis (Eds.): AC/HCII 2013, LNAI 8027, pp. 371–380, 2013.

In both cases, granular measures and high temporal resolution are desirable factors, as they allows for the production of more sensitive control mechanisms and a more detailed investigation of cognitive processing. Temporal resolution is easily hindered by physiological contamination from ocular and muscular activity. Recorded blinks and eye movements can be particularly disruptive due to their high amplitude and impulse-like shapes. These impulses disrupt both segment comparisons and local frequency estimates common in BCI and cognitive state estimation systems. The frequency with which these events occur makes the development of automatic systems for removing artifacts (rather than identifying and rejecting them) an important step in developing new tools built on EEG.

1.2 Existing EOG Removal Methods

A 2007 review by Fatourechi, et al. documents the common methods for removing such contaminants from EEG in BCI systems, their benefits and shortcomings, and the frequency with which they are employed in published work [1]. Although most studies neglect to discuss how ocular artifacts are treated (53.7%), the majority of those that employ automatic ocular artifact removal employ an electrooculogram (EOG) paired with a simple linear model (69.7%). This combination consists of dedicated electrodes placed around the eyes and a linear regression model that relates some portion of the signal recorded by each electrode of the EEG to an EOG component. In addition to requiring additional hardware and a more involved setup process, the relationship between the recorded EOG signal and the EEG is not entirely linear. The activity associated to the EOG during different eye movements varies depending on the type of eye-movement [2].

The second most common method for ocular component removal, being reported in 9.1% of papers, is blind source separation, usually in the form of independent component analysis (ICA). ICA is a statistical method that separates additively mixed, independent signals through the optimization of a measure of non-normality in the resulting component signals. ICA may be used with or without an EOG [3,4], and depends on both a large amount of data (for the optimization process) and a manual identification step (to identify components associated with the type of contamination in question).

Only 6.1% of automatic EOG artifact removal is performed with principal component analysis (PCA), an eigendecomposition based method for identifying uncorrelated components. The rarity with which PCA is used is probably due to EOG components and cortical activity not meeting the condition that components be orthogonal. PCA decompositions will produce components that contain combinations of activity from unrelated sources (such as ocular activity and frontal cortical activity). Yet, PCA remains a tempting method because the decomposition is deterministic, fast, and requires less data (when compared to ICA).

ICA and PCA perform a similar task, using statistical properties to identify spatial components that have some sort of coherent activation pattern. However, they differ in terms of assumptions about the data and the difficulty with which projections can be derived from the data. PCA alone is a poor match for the removal of EOG artifacts due to an orthogonality assumption not met by EEG data [5]. However, dropping this assumption and further optimizing for independence makes ICA a more difficult method to employ, in terms of data requirements and derivation properties. For both methods, it is likely the case that the estimated EOG component also contains cortical activity, possibly due to unmet requirements for full separation under PCA, and to suboptimal projections being derived through the optimization for non-normality [5,6].

1.3 Using Temporal Shape for Artifact Removal

In this paper, we describe an extension of ICA-based or PCA-based EOG artifact removal that further constrains the attribution of signal to ocular activity by leveraging the conserved temporal shape of common transients. Our method uses a modified version of matching pursuit, a method by which signals are represented as a linear combination of elements from an over-complete dictionary [8]. Matching pursuit results in a sparse signal representation which has been found to be useful in signal classification and compression contexts. However, the modeling process is prohibitively slow when using common dictionaries, making it unsuitable for online applications [9]. Although applied to EEG soon after its development [7], matching pursuit is rarely applied in the context (not appearing in Fatourechi's survey of artifact removal techniques) due to this poor runtime performance. Methods for accelerating matching pursuit rely on efficient implementation and careful pruning of the dictionary elements considered during signal modeling. Although the gains from these approaches can be substantial, the runtime remains bound to the cardinality of the dictionary being used, which is usually large [9]. We find that artifacts in EEG are sufficiently modeled using a very simple dictionary, consisting of only a few entries. Our approach modifies the matching pursuit process by comparing dictionary elements to the signal at all possible offsets, reducing the dictionary size to only a few elements that can be learned from the data, and are constructed to model an individual subjects ocular artifacts for removal.

Evaluation of EEG processing methods presents an inherent challenge, as we lack a ground truth signal against which processed signals can be compared. For this reason, processing methods are sometimes evaluated using synthetic data, where conclusions regarding a method can only be drawn in so far as the synthetic data is a meaningful approximation of real data. We will demonstrate our method in this synthetic context first, where assumptions will be made explicit and efficacy can be demonstrated clearly, and then we will demonstrate the method using real EEG data, which can not be evaluated for correctness, but compares favorably to the synthetic results.

2 Methods

2.1 Data Collection

We developed our technique for removing ocular artifacts while investigating EEG data collected during performance of a naturalistic reading task. Subjects read passages under a high-workload condition, where text came from sources such as the New Yorker and a challenging time constraint was imposed, or under a low-workload condition, which used easier passages and little time constraint. The nature of the task made it unreasonable to discourage eye-movements, and frequency with which they occurred made rejection unreasonable, so we had to remove the EOG from the recorded data.[1]

The frequency of eye movements made ICA quite effective in producing an independent ocular component using only a few minutes of recorded data. The component was found and identified using the methods described in Jung, et al. [4]. To further isolate ocular activity, we modeled the activation time-course of this channel using our template based method. The modeled ocular activity was then subtracted from this component, leaving other activity as a residual that could be re-integrated into the data at large. A mathematical description of the modeling process follows.

2.2 Processing

The component recognized as containing ocular activity, $\mathbf{x} = (x_1, x_2, ..., x_n)$, is modeled using a shape template $\mathbf{h} = (h_1, h_2, ..., h_m)$ and a scaling coefficient sequence $\mathbf{c} = (c_1, c_2, ..., c_n)$.

The signal model $\mathbf{m}$ is constructed by the convolution $\mathbf{m} = \mathbf{h} * \mathbf{c}$. From the signal $\mathbf{x}$ and this model $\mathbf{m}$, we generate the residual $\mathbf{r} = \mathbf{x} - \mathbf{m}$.

For convenience, let $N_m(x_i) = (x_{i-\frac{m}{2}}, ..., x_{i+\frac{m}{2}})$, that is, an m unit neighborhood of x_i. We initialize $\mathbf{h}$ by taking the point-wise average of signal segments that are centered on local amplitude extrema.

$$W = \{w_i : max(|N_m(x_i)|) = |x_i|\},$$
$$\text{where } w_i = \frac{N_m(x_i)}{stdev(N_m(x_i))}$$
$$\mathbf{h} = \mathrm{E}(W)$$

The coefficient sequence $\mathbf{c}$ is initialized with $c_i = 0$ for all i. We greedily add single coefficients using the following update procedure.

$$i^* = \arg\max_i cov(\mathbf{h}, N_m(r_i))$$

[1] EEG data was recorded from standard scalp locations using a 64-channel Biosemi™ ActiveTwo™ system. Data was down-sampled to 256Hz during recording and was high-pass filtered at 1 Hz using an 8th order Butterworth filter. Additional details regarding the experiment and the collection of this data can be found in the Engineering in Medicine and Biology 2010 conference proceedings [10].

$$a^* = \arg\min_a \sum (a\mathbf{h} - N_m(r_{i^*}))^2$$

$$c_{i^*} = a^*$$

$$\text{Update } \mathbf{m} = \mathbf{h} * \mathbf{c} \text{ and } \mathbf{r} = \mathbf{x} - \mathbf{m}$$

Each iteration selects a location from the current residual, fits a scaling coefficient, and updates the model and residual. These steps are repeated until a termination condition is met, which can be based on a limit on the number of non-zero elements of $\mathbf{c}$, a threshold on a^*, or reaching a target residual variance.

After the signal has been modeled, we update the template.

$$\text{Let } \mathbf{r}_i^\circ = \mathbf{x} - \mathbf{h} * (c_1, c_2, ..., c_{i-1}, 0, c_{i+1}, ..., c_n),$$

$$W = \{N_m(r_i^\circ) : c_i \neq 0\},$$

$$\mathbf{h} = \mathrm{E}(W).$$

$\mathbf{r}_i^\circ$ shows us our current residual modified such that a single transient remains unmodeled. This allows us approximate the event at i in isolation, reducing the effect of overlapping occurrences of the transient.

Repeating this process of modeling and template updating based on covariance guides the model toward a group of transients with a similar shape, producing a template increasingly fitting the transient of interest. A final signal estimate $\mathbf{m}_{final}$ is built using the same matching pursuit process that is used during each of the update steps. Additional transients may be modeled by repeating the entire process, starting with $\mathbf{x} - \mathbf{m}_{final}$.

3 Results

3.1 Synthetic Data

To evaluate how well our method for modeling transients decomposes an EEG signal containing both cortical and EOG components, we generated synthetic EEG data with an additive combination of a stable random signal and a sequence of transients generated from a conserved temporal shape. That is, $\mathbf{x_{EEG}} = \mathbf{x_{stable}} + \sum_i \mathbf{h_i} * \mathbf{c_i}$, where $\mathbf{h_i}$ is a transient, and $\mathbf{c_i}$ is a sparse activation sequence for $\mathbf{h_i}$.

Transients are often impulse like (and therefor broadband), so their frequency distributions overlap with that of the stable signal. We generated a transient by producing a random .5 second signal in the 1-8Hz band, and scaling it to have unit variance. There were 30 activations with coefficients in the range .75 to 1.25. The random signal was from the 1Hz to 20Hz frequency range and had variance of .25. This process produced the signal seen in Figure 1. Although not a sophisticated model of EEG activity, the components overlap in frequency distribution, and are reminiscent of contaminated EEG components.

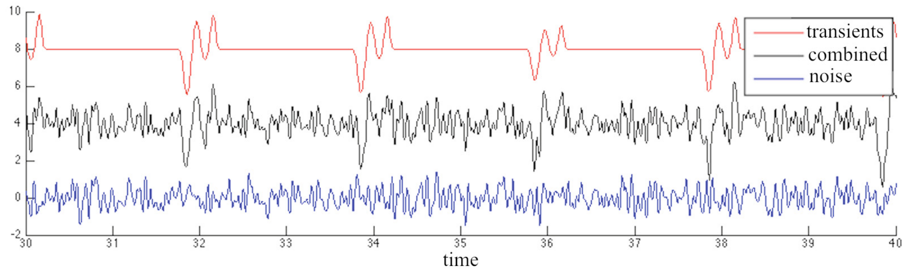

Fig. 1. We model EEG data as the additive combination of a stable random signal between 1 and 20 Hz and repeated activations of a transient in the range 1 to 8 Hz. Our model includes substantial activity not associated with the transient, which would indicate poor separation of signal components when applying a source separation technique to real EEG data.

Initializing the template using signal segments of high local variance produces a fairly poor result. However, applying our template updating procedure several times produces a reasonable approximation, as can be seen in Figure 2. Using the resulting template approximation, the signal is separated into a transient activation sequence and a residual signal, shown in Figure 3

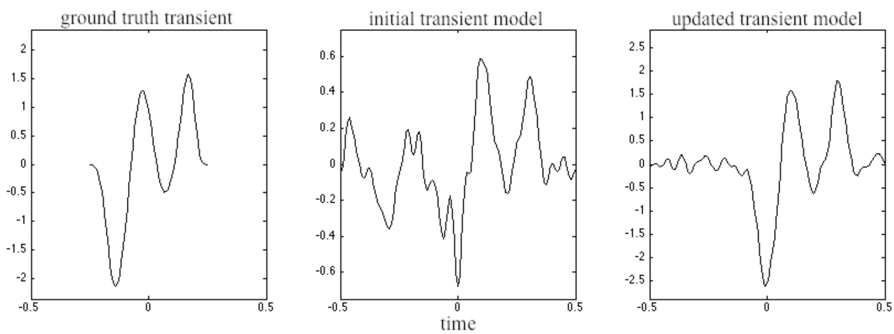

Fig. 2. Our method for modeling a repeated transient produces a good approximation of the generating transient signal, despite a poor initial estimate. The ground truth transient shape is shown on the left, the initial estimate is shown in the middle, and the final approximation is shown on the right.

As we iterated between signal modeling and template updating, we see in Figure 4 that the residual error, the squared sum of the difference between our model residual and the stable random signal, quickly drops off to a constant near the original variance. Although the error cannot be monitored during application to a real signal, the residual variance can be monitored. Figure 4 also shows that when the remaining residual variance stabilizes, there is also little change in the

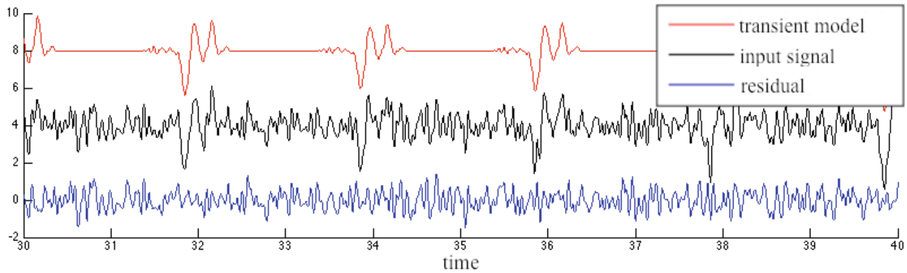

Fig. 3. The modeled signal produces a close approximation of the two contributing signal sources (seen in Figure 1) from a single channel. In conjunction with source separation techniques that use spatial information, our method may produce a more complete separation of contributing sources.

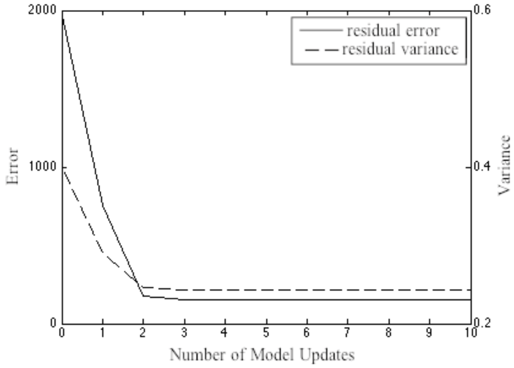

Fig. 4. The residual error closely tracks the total remaining variance in the signal. This provides useful information for constructing an appropriate termination condition when using real data, where error cannot be monitored.

residual error. If this relationship between residual variance and error remains consistent when using real data, it contributes to a natural termination condition for the signal modeling procedure.

3.2 EEG Data

To demonstrate our method for modeling ocular transients, we use a component from a PCA decomposition of the aforementioned data recorded during a reading task. Although our method works well in conjunction with ICA, PCA quickly produces a robust decomposition from which components associated with ocular activity might be more easily identified through automated processes. However, the ocular activity is less isolated when using PCA. Our template approach addresses this issue by introducing the additional shape constraint.

PCA was applied to our recorded EEG data. We selected the component containing blink activity based by examining the scalp distribution of PCA weights. Our signal modeling technique was applied to the the activation time-course of this component. As before, the template modeling method converges on a reasonable approximation of the time-course of the most common transient, as can be seen in Figure 5. Using this template, the activation time-course of the selected PCA component is decomposed into two parts: the modeled transient activation and the residual activity, which is presumably cortical in nature. This further decomposition can be seen in Figure 6. Again, the modeled signal closely tracks what appear to be instances of the transient activity, without introducing obvious artifacts or removing additional activity.

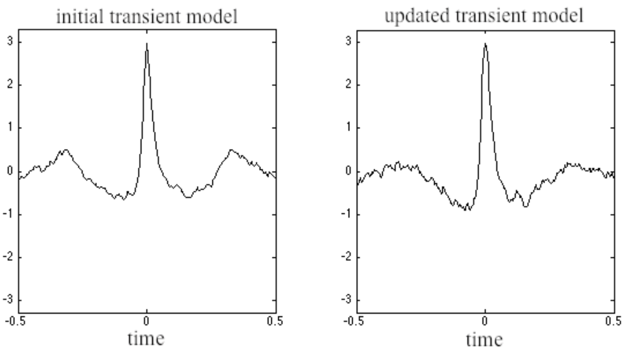

Fig. 5. The estimated transient converges quickly to a plausible shape for common ocular artifacts. The frequency of blinks in rapid succession contributes to an elevated amplitude leading and following the main amplitude spike by about .25 seconds. These sorts of artifacts associated with overlapping transients are reduced by the local modeling used when learning the transient shape.

Using this method allowed us to perform our analyses of data recorded during the reading-task, which included the extraction of local frequency features, without corruption from impulse-like activations.

4 Discussion

In this paper we introduced a template based approach to modeling transients in EEG data. We provided evidence for the efficacy of the method in the context of synthetic data, and demonstrated the result of applying the process to real EEG data. The method provides the benefit of being able to remove common ocular contaminants from recorded EEG signals without rejecting data; without the use of additional hardware; and, when combined with PCA, without requiring large amounts of data or a detailed analysis of separated components.

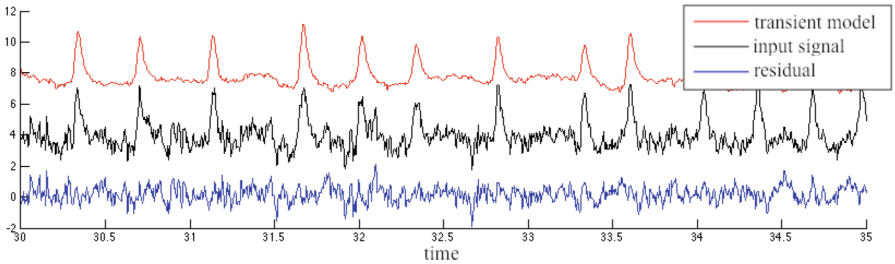

Fig. 6. On ocular and frontal activity component (in black) is split into a transient component (in red) and a locally stationary signal (in blue). It appears that this separation, using a component derived using PCA, creates for a more easily automated method for removing ocular contamination in EEG when compared to ICA based methods, which usually require the attention of a researcher to identify ocular components.

There are several natural concerns associated with the method, including implementation details such as termination conditions and the potential for learning degenerate templates that model more than the intended transient activity. Additionally, the known shortcomings of matching pursuit remain applicable, with the potential for highly suboptimal signal representations to occur when transients overlap frequently. Furthermore, we have only peripherally discussed the situation where multiple transient patterns occur within the same signal. For example, transients associated with left-to-right and right-to-left eye movements usually share a component when identified using PCA or ICA, but may not be appropriately modeled with a single template.

In practice, many of these concerns can be mitigated, without significant effort, by empirically adjusting the few parameters in the system, such as the width of the transient model and the threshold on activation coefficients. The presence of multiple transient shapes requires a bit more engineering, but initial results using clustering of signal segments (rather than a simple mean), or sequential application of the method, are encouraging.

Several additional experiments remain future work, including: comparing our method to EOG signals gathered with dedicated electrodes, further automation through the integration of priors for the spatial distribution of common ocular contaminants, and the modeling and detection of evoked response potentials without the use of a time-lock.

References

1. Fatourechi, M., Bashashati, A., Ward, R.K., Birch, G.E.: EMG and EOG Artifacts in Brain Computer Interface Systems: A Survey. Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology 118(3), 480–494 (2007)
2. Croft, R.J.: The Removal of Ocular Artifact from the EEG. University of Wollongong (1999)

3. Parra, L.C., Spence, C.D., Gerson, A.D., Sajda, P.: Recipes for the Linear Analysis of EEG. NeuroImage 28(2), 326–341 (2005)
4. Jung, T., Humphries, C., Makeig, S., Mckeown, M.J., Iragui, V., Sejnowski, T.J.: Extended ICA Removes Artifacts from Electroencephalographic Recordings. Neural Information Processing Systems 10, 894–900 (1998)
5. Jung, T., Humphries, C., Lee, T., Makeig, S., Mckeown, M.J., Iragui, V., Sejnowskil, T.J.: Removing Electroencephalographic Artifacts: Comparison Between ICA and PCA. Neural Networks for Signal Processing VIII. In: Proceedings of the 1998 IEEE Signal Processing Society Workshop, pp. 63–72 (1998)
6. Castellanos, N.P., Makarov, V.A.: Recovering EEG Brain Signals: Artifact Suppression with Wavelet Enhanced Independent Component Analysis. Journal of Neuroscience Methods 158, 300–312 (2006)
7. Durka, P.J., Blinowska, K.J.: Analysis of EEG Transients by means of Matching Pursuit. Annals of Biomedical Engineering 23(5), 608–611 (1995)
8. Mallat, S.G., Zhang, Z.: Matching Pursuits With Time-Frequency Dictionaries. IEEE Transactions on Signal Processing 41(12), 3397–3415 (1993)
9. Krstulovic, S., Gribonval, R.: Mptk: Matching Pursuit Made Tractable. In: ICASSP 2006 Proceedings, 2006 IEEE International Conference on Acoustics, Speech and Signal Processing, vol. 3 (2006)
10. Mathan, S., Smart, A., Ververs, T., Feuerstein, M.: Towards an Index of Cognitive Efficacy. Engineering in Medicine and Biology (2010)