Published in final edited form as:

IEEE Trans Biomed Eng. 2009 February ; 56(2): 200–204. doi:10.1109/TBME.2008.2005952.

Removal of BCG artifacts using a non-Kirchhoffian overcomplete representation

Mads Dyrholm, Robin Goldman, Paul Sajda [Senior Member, IEEE], and Truman R. Brown
The authors are with Columbia University, New York.

Abstract

We present a nonlinear unmixing approach for extracting the ballistocardiogram from EEG recorded in an MR scanner during simultaneous acquisition of fMRI. First, an overcomplete basis is identified in the EEG based on a custom multi-path EEG electrode cap. Next, the overcomplete basis is used to infer non-Kirchhoffian latent variables which are not consistent with a conservative electric field. Neural activity is strictly Kirchhoffian while the BCG artifact is not, and the representation can hence be used to remove the artifacts from the data in a way that does not attenuate the neural signals needed for optimal single-trial classification performance. We compare our method to more standard methods for BCG removal, namely ICA and OBS, by looking at single-trial classification performance for an auditory oddball experiment. We show that our overcomplete representation method for removing BCG artifacts results in better single-trial classification performance compared to the conventional approaches, indicating that the derived neural activity in this representation retains the complex information in the trial-to-trial variability.

I. INTRODUCTION

The ballistocardiographic (BCG) artifact is a very strong contaminant of EEG that is acquired simultaneously with fMRI. This artifact is believed to arise via electro-magnetic induction from heart beat related movement of the head/electrodes and also from flow related phenomena when a subject is situated inside the extremely strong static magnetic field of an MRI scanner, see e.g. [1]. Several methods have been proposed for removing BCG artifacts from EEG. These include averaging [2], independent component analysis (ICA) [3], adaptive filtering [4] and PCA ‘optimal basis sets’ (OBS) [5]. These methods can provide cleaner looking EEG with reduced power at the BCG frequencies, however, in our hands, they reduce single-trial EEG classification performance when compared to not removing BCG at all. That is, these methods attenuate or distort that part of the underlying neural signal informative for single-trial classification. This finding may be specific to situations where the number of trials is extremely limited. However, for such cases, single-trial classification methods are appealing because they are supervised and multivariate (e.g. can integrate spatially), and are hence applicable to recordings with very low SNR [6], [7]. Template or regression based BCG removal approaches assume that the BCG artifact is reproducible, and that a reference electrocardiogram (ECG) recording is available. ICA approaches, on the other hand, are appealing because they do not rely on an ECG regressor signal, and they do not assume that the artifact is exactly reproducible. However, for ICA the artifact is extracted by linear unmixing using matrices that are estimated blindly from the data. Hence, the ICA approach can only unmix to the extent that the mixture model is valid in combination with the independence assumption. For linear unmixing, the rank of the EEG and BCG matrices are required to add up to (and not exceed) the data rank. But since the EEG is in itself of full rank, such a rank constraint produces deficient ‘source’ matrices. In fact, in order to avoid deficiency of the cleaned source matrix, one should extract more sources than the number of sensors thus obtaining an overcomplete representation. To remove the BCG, it turns out that we do not have to learn the overcomplete basis as was

generally suggested in [8]; instead we show that the artifact removal problem can be tackled in two steps: 1) Spatial information from a bipolar EEG electrode configuration, made possible by use of a custom EEG cap, results in an overcomplete basis which is *known*. The latent variables of the mixture are inferred by maximizing the posterior assuming a Laplace prior for the non-Kirchhoffian embedding and a uniform prior on the Kirchhoffian part. Such nonlinear source inference allows us to extract a latent variable matrix which has more rows and higher rank than the original data matrix. 2) We show how temporal information can be incorporated by solving an ICA problem in the latent variable matrix. In this way, the artifact components are subtracted to produce the underlying EEG. To determine the quality of the resultant EEG signals we compare our new method to other artifact removal methods, namely Infomax ICA and OBS, by considering EEG data collected simultaneously with fMRI for an auditory oddball task. Using signal detection theory, we compute the area under the curve (AUC) of the receiver operating characteristic for single-trial EEG classification of target vs. standard trials for each of the different BCG removal methods, and can thus determine which method minimally distorts the underlying neural signal.

II. METHODS

A. Simultaneous EEG and fMRI acquisition

We recorded simultaneous EEG/fMRI for 12 subjects performing an auditory oddball task. A sequence of short tones were played, and subjects were asked to stay alert and press a button when they heard the ‘target’ tone, which was a pure tone with a higher pitch (500 Hz) than the ‘standard’ tone (350 Hz). Pitch and timing parameters were meant to match previous auditory oddball fMRI-only experiments, see [9]. Each subject was presented 50 target and 200 standard tones. Stimulus intensity was set to 85 dB as measured at the headphones. Tones were presented for 200 ms with an inter-stimulus interval (ISI) chosen from a uniform distribution between 2 and 3 seconds in increments of 200 ms. Whole brain fMRI data were collected on a 1.5T scanner (Philips Medical Systems, Bothell, WA). Functional EPI data were acquired using 15 slices of 64×64 voxels with in-plane resolution of 3.125mm and slice thickness of 8mm. Repetition time (TR) was 3000ms with an echo time (TE) of 50ms. EEG was collected simultaneously using a custom-built MR-compatible system consisting of a cap which is an array of 36 custom Ag/AgCl electrodes (BioMed Products LLC, HI, USA) and 43 custom differential amplifiers [10]. Each electrode has either 2 or 3 leads to produce a hard-wired bipolar montage. In this way, leads from neighboring electrodes can be twisted together to minimize inductive noise thereby reducing gradient induced artifacts which are an additional source of noise in the EEG [2]. For safety, non-magnetic 10k Ω surface mount resistors (IMS, Inc., Portsmouth, RI) connect the electrode and lead under the electrode epoxy to reduce induced current. The lead-pairs then connect to separate differential amplifiers and the 36-electrode cap thus forms a chain of 43 bipolar electrode pairs around the scalp, see Fig. 1(a). The EEG (and concurrent ECG) was acquired at 1000Hz sampling rate on an ADC driven by a field programmable array (FPGA) card (National Instruments, Austin, TX). To enable removal of gradient artifacts, this sampling was synchronized to the MR scanner clock, by a transistor-transistor logic (TTL) pulse at the start of each TR. A software-based 0.5Hz high-pass filter was used to remove DC drift. Gradient artifacts were then removed by aligning data for each bipolar channel to the start of each TR and subtracting the mean across TRs. A ten-point (10ms) median filter was then applied to eliminate the minimal remaining RF artifacts. Software-based 60Hz and 120Hz (harmonic) notch filters were applied to remove line noise artifacts. The high-pass and notch filters were first designed as second-order Butterworth filters (4Hz stop-band for the notch filters), then applied twice (forward in time, then backward in time) to be zero-phase to prevent delay distortions. After recording the EEG/fMRI data, the subjects would exit the scanner and perform the experiment again—this time recording the EEG only. Such ‘outside scanner’ recordings were made for 10 of the subjects (Subject 1 and Subject 6 excluded).

B. An overcomplete method for BCG removal

Our custom EEG cap, with its hard-wired bipolar montage of multi-lead electrodes, forms multiple chains of bipolar electrode pairs (see Fig. 1a) so that each electrode can be reached by multiple pathways. In this section we show that the multi-path cap provides the necessary redundancy for building an overcomplete representation by using the multiple paths between any two electrodes. Let u_p denote the (signed) voltage measured by a differential amplifier at an arbitrary point in time; here p is the channel and hence also defines which two electrodes are involved in the measurement. Assume that we have more than two electrodes available, and that every electrode is involved in at least two differential pairs. The ‘bipolar’ configuration of our cap is shown in the electrode wiring diagram in Fig. 1(a) where the lines between electrodes represent their connection to a differential amplifier via twisted lead pairs. Using the concept of Kirchhoffian loops — that voltages at each electrode are due to a current flowing through a resistance, as is the case for true EEG voltages — we can split any differential measurement in two parts $u_p = e_p + n_p$ where e_p is a voltage that sums to zero over closed loops as is necessary for resistively generated voltages and n_p is everything else. Hence, both the true EEG and the multi-path consistent (Kirchhoffian) part of BCG artifact assigns to e_p , but, as the loops cover different areas and are perturbed in slightly different ways, the remaining part of the induced signal will not be multi-path consistent and will thus be assigned to n_p . Note the fundamental property of Kirchhoffian loops is that no matter which route we choose between the two electrodes of channel p , adding the piecewise e -contributions along that route will always add up to exactly e_p , i.e. $\pm e_x \pm \dots \pm e_y = e_p$ where the sum can be taken over any path that connects the electrodes of channel p , and the individual ‘ $\pm$ ’ can be either + or – depending on the amplifier polarities in that loop relative to the amplifier polarity of e_p . Also note that this property does not hold for the n -contributions. Using this property, we can now build a linear system involving the known 43 measurements $\{u_p\}$ and the unknown 86 latent variables $\{e_p, n_p\}$ by considering all possible routes on the wiring diagram leading from the electrodes of channel p without adding the same measurement twice

$$\begin{aligned} u_p &= e_p + n_p \\ \pm u_x \pm \dots \pm u_y &= e_p \pm n_x \pm \dots \pm n_y \\ &\vdots \end{aligned} \tag{1}$$

where the pattern of +’s and –’s for the u -terms is the same as for the n -terms. These equations all include one single e -term, namely e_p , while adding different combinations of n -terms. The resulting set of equations can be formulated as a matrix (of zeros, ones, and minus ones), denoted by M_p , so that in matrix notation

$$\mathbf{M}_p \mathbf{u} = \begin{bmatrix} 1_p & \mathbf{M}_p \end{bmatrix} \begin{bmatrix} e_p \\ \mathbf{n} \end{bmatrix} \tag{2}$$

where $\mathbf{u}$ is a 43-dimensional vector of all the differential amplifier measurements, $\mathbf{n}$ the corresponding vector of n -terms, and 1_p a column vector of implicit length filled with ones. The matrix M_p can be determined by exhaustive search knowing the electrode wiring diagram. For instance, given the diagram in Fig. 1(a), there are 70 different routes between frontal electrodes FP1 and FP2. Fig. 1(b) shows some of the rows of $M_{FP1-FP2}$. The joint linear system involving all measurements and latent variables is found by consolidating (2) for all channels p , however it cannot be inverted since the rank is lower than the number of latent variables. With our particular cap, the rank turns out to be 51. Hence, the problem of inferring the latent variables becomes that of maximizing the posterior under an appropriate prior distribution. The Laplace distribution has been suggested as a generically useful prior in overcomplete

representations, and the inference can be carried out by a linear program which minimizes the 1-norm of the latent variables [8]. However, it is clear that minimizing the 1-norm of $\{\hat{e}, \hat{n}\}$ will not uniquely determine the balance between e_p and n_p . In particular, there is the trivial solution $(\hat{e}, \hat{n}) = (0, u)$ from which we learn nothing about the Kirchhoffian embedding. In fact, by our initial definition of e , we seek a solution which is as far from this trivial solution as possible. This can be achieved by minimizing $\|n\|_1$ without regard to the norm of e . We solve this problem in the electrode domain, i.e.

$$\hat{v} = \arg \min \|n\|_1 = \arg \min \|u - Av\|_1 \quad (3)$$

where v represents electrode potentials, and A represents the electrode-amplifier connections. This problem is equivalent to inferring $\{e, n\}$ assuming a Laplace prior on n (i.e. a sparse assumption for the non-Kirchhoffian activity) and a uniform prior on e . We solve (3) using [11] and finally let $\hat{e} = A\hat{v}$ and $\hat{n} = u - A\hat{v}$. This solution has the desired property that, in the event that u is strictly Kirchhoffian, the data will remain unaltered ($\hat{n}$ becomes zero). Now, since $\hat{e}$ and $\hat{n}$ are distinguished by their Kirchhoffian-to-non-Kirchhoffian content ratio, the non-Kirchhoffian BCG activity in $\hat{n}$ can be used to remove BCG activity from u . To do so we use ICA in the latent space of $\{\hat{e}, \hat{n}\}$. To be directly comparable to the ICA method of [3] we perform the ICA step on temporally concatenated epochs, and for the specific purpose of removing BCG artifacts, we inspect the component time courses visually, and remove all components which are dominated by BCG activity. This leads to a varying numbers of components which are then removed from the data by subtracting from u its contributions to $\hat{e} + \hat{n}$. Since the inference is carried out with a temporal i.i.d. assumption the components are smoothed (lowpass filtered at 15Hz) before subtraction. This smoothing/filtering is done by projecting each epoch temporally onto a filtering basis, setting the transition frequency to 15Hz by using a 36 dimensional Discrete Prolate Spheroidal Sequence basis known from multitapered spectral analysis [12]. This choice of transition frequency is simply based on visual inspection and suggests that ~ 15 harmonics, having a fundamental frequency of 1Hz, will reproduce a realistic BCG waveform. Note that this transition frequency is not adjusted in cross-validation.

C. Other methods for BCG removal

1) BCG removal using ICA on the raw data—The ICA method of [3] first extracts epochs of raw data, then concatenates them temporally and performs an ICA decomposition. Here we use Infomax ICA as it is implemented in the EEGLAB toolbox [13]. The components are inspected visually, and components with strong BCG contents are easily identified. If too many ‘BCG components’ are subtracted then eventually the neural signal will also be attenuated, and therefore we use the method as conservatively as possible and subtract only *the strongest* ‘BCG component’ from the ICA decomposition (strongest in terms of power). The resulting cleaned data matrix thus has rank 42 (43 minus one). Subtracting more components results in worse single-trial classification.

2) BCG removal using OBS on the raw data—For comparison with an ECG-based, and easily available method, we use the FMRI plug-in for EEGLAB, provided by the University of Oxford Centre for Functional MRI of the Brain (FMRI) [5], [14]. The method uses information from an ECG channel to subtract a template based BCG artifact estimate. We used the default settings of the plug-in for our comparisons.

D. Classification using BDCA

As mentioned, we have found a significant reduction in single-trial classification performance when BCG removal methods are applied, relative to classification on the data without BCG

removal regardless of our choices of classification methods. We present here the results from the Bilinear Discriminant Component Analysis method of [7]. The data is first epoched from 200ms before to 1000ms after each target or standard event. The classifier integrates each epoch spatially and temporally using a number of bilinear ‘classifier component’ projections. For simplicity we set the number of classifier components to one and we use temporally short integration, spanning the latencies in [200ms, 500ms] following stimulus onset. The classifier hyper-parameters (for smoothness regularization) are tuned on the raw data from a single subject (Subject 1) using five-fold cross validation, and those hyper-parameters are used for training the classifier in all other situations. The classifier is then trained on individual subjects (with the hyper-parameters fixed), five times each leaving 20% of the data out for cross validation, and performances are measured and reported using ‘area under the ROC curve’, i.e. the AUC is based on validation data that was not used for training.

III. EXPERIMENTAL RESULTS

To assess our method, we compared single-trial classification performance for three different BCG removal techniques (including ours) and for data with no BCG removed (raw data). To avoid comparing across differences in the ICA learning we extracted 43 components as would be the case with the method of [3]. The number of components that were subtracted for each of the 12 subjects were respectively (6, 7, 6, 0, 1, 2, 3, 7, 9, 6, 10, 4). These numbers were determined before computing any AUC values and kept fixed for the entire following analysis. Table-I summarizes the classification performance results for the data acquired across the 12 subjects. Our new method results in the best average single-trial classification performance across the different BCG methods. The p value shown in the table is the result of a paired, two-sided, signed rank test (Wilcoxon test); the null hypothesis is that the performance is identical to that of classifying the raw data. The result indicated that the neural activity had not been attenuated or distorted by our method, but had been so with the other methods ($p < 0.05$). Some EEG epochs for the raw data and for the new method are shown in Fig. 2. Looking at the raw versus the cleaned EEG we see that the cleaned data had visibly reduced BCG artifacts. Fig. 3 shows the Event-Related Potentials (ERP) for the target condition — i.e. the average target-epoch across all subjects. The inside-scanner ERPs are shown in blue and red, where blue represents the raw data and red represents the cleaned data. The outside-scanner ERPs are shown in green. It is clear that the peak structure of the “cleaned” inside the scanner ERPs resemble those of the outside the scanner ERPs much more than does the “raw” inside the scanner ERPs. Hence, removing the BCG artifact is necessary for reliable estimates, and by using our new method, the BCG removal can be done without degrading the single-trial classification performance.

IV. DISCUSSION / SUMMARY

The multi-path based nonlinear unmixing method we have described here is a novel method for extracting BCG artifacts in EEG acquired in an MR scanner. The approach is particularly attractive in that the attenuation of the neural signal is minimized since all activity that is consistent with a changing electric field is unperturbed. Other methods, particularly those relying on general subspace projections, have the problem that, even though BCG is substantially reduced, the neural signal has significant power in the identified BCG subspace and thus removal of that subspace reduces the neural signal. This is particularly problematic if one is trying to perform single-trial analysis on the EEG, where every fraction of a dB counts. While our method requires a new type of multi-path EEG cap that oversamples the electrode space, the significant improvement in discrimination we see with our method over other BCG removal methods may indicate that in the case of single-trial analysis the use of such novel hardware is greatly beneficial.

Acknowledgements

This work was supported by the National Institutes of Health (NIH) under Grant EB004730.

REFERENCES

1. Debener S, Mullinger K, Niazy R, Bowtell R. Properties of the ballistocardiogram artefact as revealed by EEG recordings at 1.5, 3 and 7 T static magnetic field strength. *Int J Psychophysiol.* Jul;2007
2. Goldman RI, Stern JM, Engel JJ, Cohen MS. Acquiring simultaneous EEG and functional MRI. *Clin Neurophysiol Nov;2000 111(11):1974–80.* [PubMed: 11068232]
3. Srivastava G, Crottaz-Herbette S, Lau KM, Glover GH, Menon V. ICA-based procedures for removing ballistocardiogram artifacts from EEG data acquired in the MRI scanner. *Neuroimage 2005;24(1):50–60.* [PubMed: 15588596]
4. Bonmassar G, Purdon P, Jaaskelainen IP, Chiappa K, Solo V, Brown BJ, N. E. Motion and ballistocardiogram artifact removal for interleaved recording of EEG and EPs during MRI. *Neuroimage Aug;2002 16(4):1127–41.* [PubMed: 12202099]
5. Niazy RK, Beckmann CF, Iannetti GD, Brady J, Smith SM. Removal of FMRI environment artifacts from EEG data using optimal basis sets. *Neuroimage Nov;2005 28(3):720–37.* [PubMed: 16150610]
6. Parra L, Alvino C, Tang A, Pearlmutter B, Yeung N, Osman A, Sajda P. Linear spatial integration for single-trial detection in encephalography. *Neuroimage 2002;17(1):223–230.* [PubMed: 12482079]
7. Dyrholm M, Christophoros C, Parra LC. Bilinear discriminant component analysis. *Journal of Machine Learning Research May;2007 8:1097–1111.*
8. Lewicki MS, Sejnowski TJ. Learning overcomplete representations. *Neural Computation 2000;12(2):337–365.* [PubMed: 10636946]
9. Friedman D, Goldman R, Stern Y, Brown T. The brain's orienting response: An event-related functional magnetic resonance imaging investigation. *Hum Brain Mapp. May;2008*
10. Sajda, P.; Goldman, R.; Philiastides, M.; Gerson, A.; Brown, T. A system for single-trial analysis of simultaneously acquired EEG and fMRI; *IEEE EMBS Conference on Neural Engineering; 2007.*
11. Abdelmalek NN. Algorithm 551: A Fortran subroutine for the L1 solution of overdetermined systems of linear equations. *ACM Transactions on Mathematical Software 1980;6(2):228–230.*
12. Thomson DJ. Spectrum estimation and harmonic analysis. *Proc. IEEE 1982;70(9):1055–1096.*
13. Delorme A, Makeig S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics. *Journal of Neuroscience Methods 2004;134:9–21.* [PubMed: 15102499]
14. Iannetti G, Niazy R, Wise R, Jezzard P, Brook J, Zambreanu L, Vennart W, Matthews P, Tracey I. Simultaneous recording of laser-evoked brain potentials and continuous, high-field functional magnetic resonance imaging in humans. *Neuroimage 2005;28(3):708–719.* [PubMed: 16112589]

(a) Electrode wiring diagram.

(b) 70 ways to estimate electric field potential FP1-FP2.

Fig. 1.

(a) EEG cap electrode wiring diagram showing the configuration of bipolar electrode pairs. Each electrode is connected to multiple leads. A line between two electrodes indicates a connection via twisted lead pair to a differential amplifier. Thus, each electrode is involved in measuring more than one voltage difference. (b) Matrix representation of the 70 different routes between electrodes FP1 and FP2. Symbol '+' represents a +1 element, symbol 'o' represents a -1. The first row represents the shortest route which is simply the measurement $\mu_{FP1,FP2}$ itself. The differential amplifier is connected with the + terminal on FP1, so the matrix element is +1. The second row represents the path going through electrodes AF3-F3-FC1-FZ-FC2-F4-AF4. The signs are set according to the differential amplifier polarity.

Fig. 2. Epochs for the first five target or standard events, each showing 200ms before to 1000ms after the event, for channels 1–10 are shown for Subject 11 in two cases (same scale): (a) The raw data with BCG artifact. One can clearly see the BCG artifact ‘bumps’ in the signal. (b) The cleaned data using our new method to remove BCG artifact. The BCG artifacts are visibly reduced.

Fig. 3.

ERPs at four different electrodes referenced to the left mastoid. The ERP for the raw data is shown in blue, the ERP for the cleaned data is shown in red, and the ERP for 10 subjects recorded outside the MR scanner is shown in green.

TABLE I
Cross validated AUC for different BCG removal methods.

BCG removal	— Subjects —												p ^I	
	1	2	3	4	5	6	7	8	9	10	11	12		Avg. ^I
New method	.90	.91	.81	.85	.91	.96	.86	.93	.85	.88	.80	.88	.88	.557
Do nothing	.89	.90	.80	.85	.92	.96	.79	.93	.85	.90	.76	.89	.87	≡ 1
ICA	.89	.89	.78	.86	.92	.96	.80	.92	.85	.88	.73	.88	.86	.034
OBS	.89	.91	.75	.75	.87	.86	.77	.87	.87	.89	.77	.88	.84	.034

^IThe average AUC and p values, in the right columns, were computed before rounding the numbers to two decimals.