Exploratory Identification of Cardiac Noise in *f*MRI Images*

Lilla Zöllei[1], Lawrence Panych[2], Eric Grimson[1], and William M. Wells III[1,2]

[1] Massachusetts Institute of Technology,
Artificial Intelligence Laboratory, Cambridge, MA 02139, USA
{lzollei,sw,welg}@ai.mit.edu, panych@bwh.harvard.edu
[2] Department of Radiology, Harvard Medical School and
Brigham and Women's Hospital, Boston, MA 02115, USA

Abstract. A fast exploratory framework for extracting cardiac noise signals contained in rest-case fMRI images is presented. Highly autocorrelated, independent components of the input time series are extracted by applying Canonical Correlation Analysis in the time domain. A close correspondence between some of these components and cardiac noise contributions is established. Our analysis is carried out without using any external monitoring of the subject or any modification applied to the standard image acquisition protocol. Using the results as *a priori* information about the presence of corrupting cardiac noise, several approaches are suggested that could improve the performance of activation detection algorithms on non-rest-case datasets.

1 Introduction

1.1 Functional MRI and Noise

The analysis of functional Magnetic Resonance Images (*f*MRI) has been engaging an increasing number of researchers in both neuro and computational sciences. This new type of image modality may provide a way of understanding and demonstrating the fascinating and fundamental problem of structural and functional relationships in the human brain.

Analyzing *f*MRI time series signals is, however, a challenging task; the scale of signal variation to be recovered is only slightly different from the standard variation of the time series. Moreover, there are several noise sources that corrupt the signals. These corrupting signals could originate, for example, from physical patient movement, the imaging device, and physiological (e.g.: cardiac and respiratory) noise. In the present paper, we are interested in the characterization of the cardiac contributions to physiological noise in the signal. Knowing that this type of noise is always present in the acquisitions, its identification, labeling or potential elimination could lead to a significant performance increase in further data processing.

* This work has been supported by the Whiteman Fellowship and The Harvard Center for Neurodegeneration and Repair.

R.E. Ellis and T.M. Peters (Eds.): MICCAI 2003, LNCS 2878, pp. 475–482, 2003.

1.2 Noise Characterization

Low-frequency noise, which corrupts *f*MRI acquisitions, is often referred to as $1/f$ (or flicker) noise. That nomenclature refers to the shape and position of a low-frequency component in the frequency domain which is always detectable in the images. Although its presence is widely accepted, its exact composition still poses numerous questions. Slowly varying signals, aliasing artifacts, slow changes in patient movement or position, and machine noise have all been attributed to it. Eliminating this frequency range from *f*MRI images is generally not feasible given the fact that it overlaps the standard protocol frequency range.

Cardiac contributions to *f*MRI signals are physiological noise effects that could significantly increase uncertainties in activation detection results. A major problem in examining their influence stems from the fact that the standard sampling frequency (acquiring image volumes at every 2.5 sec, for example) is insufficient to represent them with high fidelity. In addition to the aliasing problem, these signals have a great variance even within a single 4D acquisition of the same subject, leading to high contamination of a wide range of frequency bands. Thus constructing a general noise model predicting their frequency domain occurrence would be extremely difficult.

1.3 Background

Many existing methods that treat confounders in *f*MRI carry out their analysis in *k-space*, where the physiological contributions are found to be quite uncontaminated. In that domain, the physiological cycles can be recovered by re-ordering the image slices according to time of acquisition [4,5] or by combining phase and projection information of the navigator echo [11].

Retrospective image analysis methods often use external monitoring in order to obtain a model for the physiological noise contributions. They then reduce or completely eliminate the contents of the corrupted frequency bands from the acquisitions [1,7]. Many algorithms favor examining phase data contained in the time series, providing a robust descriptor of the signals of interest [7,11,4,5,2, 12], while some also require fast imaging protocols, in order to avoid the critical sampling (aliasing) problem [2].

Another technique used to eliminate cardiac noise from the *f*MRI datasets is cardiac gating [8]. This refers to a special imaging procedure that eliminates cardiac confounders by incorporating EKG data into the image acquisition stage.

Research has also been conducted with respect to spatially localizing this type of noise source in the brain. It has been established that cardiac-induced signals are more localized compared to respiratory contributions and can be spatially restricted to the neighborhood of major blood vessels, CSF pools and the medial and temporal lobes [3].

The major drawback of the majority of these methods is that they either require specialized image acquisition parameters (e.g., short Time of Repetition (TR)), they heavily exploit external monitoring measurements or they need to convert the data back into *k-space* for analysis purposes.

We adopted an exploratory framework for analyzing *f*MRI images. By examining rest-case datasets, we aim to determine how much additional information we can gain about the presence of corrupting cardiac signals and how such information could be integrated into vessel segmentation or activation detection methods applied to subsequent, stimuli-activated images. The task of identifying noise components in rest-case acquisitions should be less complex than in the activation cases. That is because there are no response functions that would have to be explicitly separated from the noise signals. For the analysis, we do not require any special imaging parameters or the use of external monitoring. Our method is a quick technique that can extract cardiac contributions from the signals when the aliasing properties are favorable. It is also used to support our argument against using band-pass filtering for physiological noise removal.

2 Motivation

During one imaging session, a collection of experiments are run, and this series is preceded and followed by a set of *rest-case* scans. For these rest-case acquisitions, there are no stimuli presented to the subject and there are no tasks to be executed. While response signals are missing, physiological contributions still manifest in these images. If these are detectable, our goal is to gain information about their nature and spatial origin. That information could be very valuable in subsequent non-rest-case data analysis.

Ad hoc experiments suggest the presence of two principal components contributing to the rest-case *f*MRI signals. Displayed in the frequency domain, a higher (appearing in the .08-.2 Hz frequency range) and a lower, $1/f$-type frequency component can be often distinguished. We conjecture that the higher frequency components can be associated with aliased cardiac signals. Whenever they are *absent*, it means that the aliased cardiac contributions occupy the same lower frequency range as the flicker noise. We verify this hypothesis by using external cardiac measurements and by characterizing the noise sources with respect to the corresponding anatomical dataset.

2.1 Approach

For our data-analysis, we propose the use of a multivariate approach. As we make the assumption that a collection of independent signal sources make up the dataset, we aim to find linear transformations to recover mutually uncorrelated components. Because in rest-case datasets the *interesting* (the cardiac) signals are highly autocorrelated, we impose an additional requirement on the components: they have to be maximally autocorrelated. To achieve this task, we adapted the idea of using Canonical Correlation Analysis (CCA) [9,10], and the maximization task to be solved is displayed in Eq. (1):

$$\max_{w_x, w_y} \rho(w_x, w_y) = \frac{w_x^T C_{xy} w_y}{\sqrt{\left(w_x^T C_{xx} w_x\right)\left(w_y^T C_{yy} w_y\right)}}, \tag{1}$$

where x and y are sets of voxel time series, ρ is the correlation coefficient, C_{xx}, C_{yy} and C_{xy} are covariance matrices on x, y and (x, y), and w_x and w_y are the weight vectors that would maximize the correlation coefficient between $w_x x$ and $w_y y$. (According to our formulation, y is merely a delayed version of x, as we are interested in obtaining highly-autocorrelated components.)

We used external heart-beat measurements to verify a close correspondence between the CCA and cardiac signals. Also, in order to demonstrate that the noise contributions mostly originate from regions where anatomy would explain their presence, we constructed correlation maps between the CCA components and the original input time series. We then registered these maps to an anatomical scan of the patient.

A data-driven exploratory algorithm using CCA as a tool for running activation detection on non-rest-case fMRI data has recently been introduced [6]. There, some of the extracted components were attributed to stimulus response, but the residuals were not identified. In our experiments, we target the decomposition of rest-case input images according to this framework and hope to use the results as a priori information in later activation detection studies. For the interpretation of stimulus-activated acquisitions, we would argue for using more sophisticated techniques (for example, ones that explicitly incorporate knowledge about the experimental protocol).

2.2 Algorithm Outline

We execute all of our analysis in the temporal domain. Nevertheless we often present and characterize our results in the frequency domain.

Our preprocessing step first forms a mask to select the volume elements with the highest variation in their time series. This preprocessing step is necessary because of a limited sample size in the temporal domain.

After we apply this mask and detrending to the set of voxel time series, we execute a standard Principle Component Analysis (PCA) on the preprocessed data set. This pre-whitening step is required for the CCA computations. The number of retained components after the PCA decomposition is algorithm-dependent. We currently fix it to be a constant, four. Based upon experimental analysis of the data, we found this number sufficient to consider signals representing 98% of the total variation in the dataset. That many components were in general also enough to contain the higher frequency components, if they were present.

The CCA algorithm is executed next on the output components of the PCA analysis. The input signals are compared to a time-shifted version of themselves (in order to test for autocorrelation properties). A de-mixing matrix is computed that produces the most highly autocorrelated components. While the PCA components are ordered by the magnitude of the component variance, the components obtained by CCA are ordered by the amount of autocorrelation that they have (at time lag = 1). Two of the components that have been extracted from a single rest-case session are shown in Fig. 1 (a) and (b). Fig. 1 (a) represents a lower frequency contribution, while Fig. 1 (b) shows a higher frequency component (peaking at .19-.2 Hz).

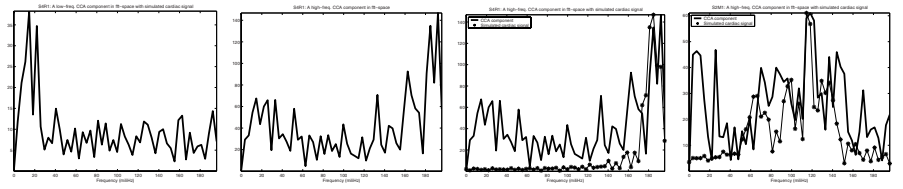

(a) (b) (c) (d)

Fig. 1. CCA components displayed in the Fourier domain. In all images, the x-axis represents the frequency interval of [0, 200] mHz and along the y-axis magnitude values are displayed. (a)-(b): a low- and high-frequency component of a single rest-case acquisition (c): the high-frequency component from (b) with the corresponding cardiac signal overlaid (d): a non-rest-case high-frequency component with the corresponding cardiac signal overlaid .

We construct correlation maps and register the results onto the anatomical scan in order to determine from where these components originate with respect to the underlying anatomy. These correlation maps are created by examining all of the extracted components. The correlation component is calculated between the original time series and the CCA signals in the temporal domain. We use a normalized version of these maps in order to compare the structuredness and relative contributions of the different components to particular anatomical locations of the brain. Four slices of two correlation maps created with respect to high frequency components are shown in Fig. 2. (In these examples, the components highly match the cardiac signal measured throughout the image acquisition procedure.)

(a) (b) (c) (d)

Fig. 2. Four sample slices of registered correlation maps. (a) and (b): maps calculated between a rest-case dataset and a high-frequency component; (c) and (d) maps calculated between a non-rest-case dataset and a high-frequency component. The set of higher correlation values (in white) outline potential cardiac signal sources in the brain.

2.3 Data Description

We carried out analysis on 12 rest-case acquisitions of the same subject. The *f*MRI images were taken in 12 different sessions, the rest-case images being scheduled for the very beginning and either for the middle or the end of the

fMRI imaging series. (Thus the elapsed time between the rest-case scans was either 17 or 34 minutes.) During these acquisitions no stimuli were presented to the subject. The (64x64x24) axial acquisitions were taken by a TR of 2.5 sec. The images were motion corrected before analysis.

The SPGR image that is used as the anatomical reference is of higher resolution. It is a sagittal acquisition of 110 slices and 256 x 256 image size.

Our external monitoring for each of the input datasets resulted from manually recording the cardiac rate of the subject every minute.

3 Results

In this section, we give some numerical indicators and qualitative analysis that demonstrates the usefulness and validity of our algorithm. The results are very promising as they demonstrate that whenever the cardiac signals are aliased to a distinguishable frequency interval, we could correctly identify them. We found that we could obtain distinct, peaked high frequency components (see, for example, in Fig. 1 (a)) in 7 out of 12 of the datasets. What is more, in all these cases, the high-frequency CCA component could be closely matched with the externally-measured cardiac signal (Fig. 1 (c)). In an additional input series a higher frequency component was also present, but it was spread out over a very wide frequency range. We established a direct correspondence between this non-spikyness and a high variation in the recorded cardiac rate during the image acquisition.

We were unable to distinguish high-frequency elements in the remaining four of the rest-case datasets. The absence of these components was not a mistake of the algorithm. We were able to show that, in all those cases, the cardiac contributions were aliased back to the lower frequency range.

In 2/3 of the examined images, we recorded extensive cardiac rate changes within single rest-case acquisitions. This meant that the aliased variation in the subject's cardiac rate exceeded a quarter of the sampling frequency range (> 50 mHz). During a whole series of successive acquisitions or between different imaging sessions the chances for a comparable variation are even higher. Due to this phenomenon, we advise against band-pass filtering or complete frequency range elimination. The band to be removed can get quite large and such a modification in the dataset could result in a significant information loss about other components.

Such a great variance in the heart beat rate even within one imaging session prompted us to explore information that the correlation maps could provide. Even if the cardiac contribution occupies a wide range of frequency domain and it is difficult to predict its exact occurrence beforehand, the location where these signals originate with respect to anatomy should be more stable. That is the reason why we created the correlation maps between the input times series and the CCA components. The results that we present with respect to them are promising but not yet rigorously tested. We first provide some qualitative results and then describe validation experiments by which we intend to make them more robust.

We mapped the correlation values created with respect to the high frequency components onto the anatomical MR image of the subject. By visual analysis, the anatomical regions with the highest scores agreed with our expectation based upon the literature. The indicators corresponded to higher scores mostly along major blood vessels, the temporal lobe and some CSF pools (see Fig. 2 (a) and (b)). This agreement was consistent in the case of all the analyzed datasets.

A validation procedure that would better prove the agreement between all the correlation maps obtained for a single subject is still in progress. We are to obtain a detailed segmentation of the corresponding SPGR dataset and register the highly correlated areas to some of its specific anatomical classes. In addition, we could test how well some of the major vessels can be located by a similar analysis if we aligned the correlation maps of our high frequency components to Magnetic Resonance Angiography (MRA) images. (A close agreement between those could even facilitate an understanding of how information from MRA's could be fused into the *f*MRI domain.)

For verification purposes, we also ran our algorithm on a set of 12 non-rest-case datasets. Although the high- and low-frequency components did not separate as clearly as in the rest-case inputs (mixed components, containing both low- and high-frequency components occurred more often), the high-frequency elements matched the externally-measured cardiac signals just as closely as in the rest-case analysis (Fig. 1 (d)). And creating the correlation maps with respect to these *interesting* components also proved to be very similar to the maps in the non-rest case. (Notice the similarity between Fig. 2 (a)-(b) and (c)-(d).)

4 Conclusion

We adapted a fast exploratory *f*MRI signal analysis framework to identify cardiac noise sources in rest-case datasets. We decomposed the input rest-case time series into highly autocorrelated independent components and established a high-level correspondence between some of these and cardiac contribution. We verified our hypothesis by utilizing external monitoring recordings and a set of anatomical references.

We suggested several potential applications of our analysis in activation detection studies. We primarily stated that using the correlation measures as confounder space indicators, we could use them as a priori information for hypothesis testing in non-rest-case data. That would be preferred over band path filtering or signal elimination solutions to noise reduction. That is mostly because we showed that the cardiac contribution often appears over a wide frequency range.

In the future, we would like to automate the procedure of extracting the high-frequency components and also to determine whether we could increase the accuracy of our normalized correlation maps by processing our input sequences slice-by-slice instead of as a whole volume. In this case our initial masking step would not have to be as restrictive as it is now.

References

1. B. Biswal, E.A. DeYoe and J.S. Hyde "Reduction of Physiological Fluctuations in FMRI Using Digital Filters" *Magnetic Resonance in Medicine*, 1996, vol 35, p 107–113.
2. K.H. Chuang and J.H. Chen "IMPACT: Image-based Physiological Artifacts Estimation and Correction Technique for *f*MRI " *Magnetic Resonance in Medicine*, 2001, vol 46(2), p 344–353.
3. Dagli, Ingeholm, Haxby "Localization of Cardiac-induced Signal Changes in *f*MRI " *NeuroImage*, 1999, vol 9, p 407–415.
4. L.R. Frank, R.B. Buxton, E.C. Wong "Detection of Physiological Noise Fluctuations From Undersampled Multislice *f*MRI Data" *Proc. Intl. Magn. Reson. Med.*, 1999, p 277.
5. L.R. Frank, R.B. Buxton, E.C. Wong "Estimation of Respiration-Induced Noise Fluctuations From Undersampled Multislice *f*MRI Data" *Magnetic Resonance in Medicine*, 2001, vol 45, p 635–644.
6. O. Friman, M. Borga, P. Lundberg and H. Knutsson "Exploratory *f*MRI Analysis by Autocorrelation Maximization" *NeuroImage*, 2002, vol 16(2).
7. G.H. Glover,T. Li and D. Ress "Image-Based Method for Retrospective Correction of Physiological Motion Effects in *f*MRI : RETROICOR" *Magnetic Resonance in Medicine*, 1999, 99–4512.
8. A.R. Guimaraes, J.R. Melcher, T.M. Talavage, J.R. Baker, P. Ledden, B.R. Rosen, N.Y.S. Kiang, B.C. Fullerton and R.M. Weisskoff "Imaging Subcortical Auditory Activity in Humans" *Human Brain Mapping*, 1998, vol 6, p 33–41.
9. H. Hotelling "The Most Predictable Criterion" *Journal of Educational Psychology*, 1935, vol 26, p 139–142.
10. H. Hotelling "Relations Between Two Sets of Variates" *Biometrica*, 1936, vol 28, p 321–377.
11. T.H. Le, X. Hu "Retrospective estimation and correction of physiological artifacts in *f*MRI by direct extraction of physiological activity from MR data" *Magnetic Resonance in Medicine*, 1996, vol 35(3), p 290–298.
12. U. Ziyan, J.L. Ulmer, T.M. Talavage "Image-Space Based Estimation and Removal of Respiration Noise from *f*MRI Data" *Proc. Intl. Soc. Mag. Reson. Med.*, 2002, vol 10.