Improving the use of principal component analysis to reduce physiological noise and motion artifacts to increase the sensitivity of task-based fMRI

David A. Soltysik[a,*], **David Thomasson**[b], **Sunder Rajan**[a], and **Nadia Biassou**[c]

[a]Division of Biomedical Physics, Office of Science and Engineering Laboratories, Center for Devices and Radiological Health, Office of Medical Products and Tobacco, U.S. Food and Drug Administration, 10903 New Hampshire Ave, Silver Spring, MD 20993, USA

[b]Integrated Research Facility, National Institute of Allergy and Infectious Disease, Ft Detrick, MD 21702 USA

[c]Radiology and Imaging Sciences, Warren Grant Magnuson Clinical Center, National Institutes of Health, Bethesda, MD 20892 USA

Abstract

Background—Functional magnetic resonance imaging (fMRI) time series are subject to corruption by many noise sources, especially physiological noise and motion. Researchers have developed many methods to reduce physiological noise, including RETROICOR, which retroactively removes cardiac and respiratory waveforms collected during the scan, and CompCor, which applies principal components analysis (PCA) to remove physiological noise components without any physiological monitoring during the scan.

New Method—We developed four variants of the CompCor method. The optimized CompCor method applies PCA to time series in a noise mask, but orthogonalizes each component to the BOLD response waveform and uses an algorithm to determine a favorable number of components to use as "nuisance regressors." Whole brain component correction (WCompCor) is similar, except that it applies PCA to time-series throughout the whole brain. Low-pass component correction (LCompCor) identifies low-pass filtered components throughout the brain, while high-pass component correction (HCompCor) identifies high-pass filtered components. Comparison with existing method: We compared the new methods with the original CompCor method by examining the resulting functional contrast-to-noise ratio (CNR), sensitivity, and specificity.

*Corresponding Author. Division of Physics, Office of Science and Engineering Laboratories, Center for Devices and Radiological Health, U.S. Food and Drug Administration, 10903 New Hampshire Ave, Silver Spring, MD 20993, USA, david.soltysik@fda.hhs.gov, Telephone: 301-796-5278, Fax: 301-796-9927.

Results—1) the optimized CompCor method increased the CNR and sensitivity compared to the original CompCor method and 2) the application of WCompCor yielded the best improvement in the CNR and sensitivity.

Conclusions—The sensitivity of the optimized CompCor, WCompCor, and LCompCor methods exceeded that of the original CompCor method. However, regressing noise signals showed a paradoxical consequence of reducing specificity for all noise reduction methods attempted.

Graphical abstract

Keywords

BOLD fMRI; physiological noise; principal components analysis

1. Introduction

Blood oxygenation-level dependent (BOLD) functional magnetic resonance imaging (fMRI) is plagued by raw intrinsic noise, subject motion, and physiological noise. The presence of noise reduces the sensitivity of BOLD fMRI studies and reduces the efficacy of fMRI as a biomarker. If BOLD fMRI were to be used as an endpoint in device trials or for diagnostics, it is essential to understand the value of computational methods for noise reduction upon the sensitivity and specificity of the BOLD signal.

Raw noise, which is independent of the MR signal (Edelstein et al., 1986), is composed of both thermal noise and system noise. Thermal noise derives from the random motion of electrons in the radiofrequency (RF) coil and the tissue being imaged (Haacke et al., 1999). A rise in temperature increases the motion of electrons, increasing thermal noise. Imperfection of the hardware leads to system noise, including low frequency drift (Smith et al., 1999), static field inhomogeneities due to imperfect shimming, nonlinearities in the gradient fields, and irregularities in the performance of the RF coil (Huettel et al., 2004).

Subject motion constitutes a major source of noise in fMRI (Friston et al., 1996). Even head motion of a few millimeters increases the variability of voxel signal intensity as the relative proportion of different tissue types changes inside each voxel. Rigid-body registration methods have been developed to correct for intrascan motion (Cox and Jesmanowicz, 1999). Stimulus- and task-correlated motion can increase the presence of false positives or false negatives in fMRI activation maps (Yetkin et al., 1996). To overcome correlated motion effects, one can use event-related designs (Birn et al., 1999), optimized block durations (Birn et al. 2004), or post-processing methods (Bullmore et al., 1999; Soltysik and Hyde, 2006) to separate true BOLD responses from motion artifact responses.

Physiological noise consists of signal variation in images caused by various processes of the human body. Weisskoff et al. (1993) first reported the presence of cardiac and respiratory waveforms in the power spectra of cortical voxel time series. Cardiac-driven signal changes are mostly due to motion from vascular pulsations in voxels near arterial and venous structures (Dagli et al., 1999). Bulk susceptibility variations in the lungs during respiration leads to systematic variations in the static magnetic field within brain tissue (Raj et al., 2001). These field strength variations lead to image shift, signal changes in the phase encoding direction, and signal variation due to intravoxel dephasing. Mitra et al. (1997) found vasomotor oscillations (0.1 Hz). Furthermore, researchers have discovered low-frequency fluctuations (0.03 Hz) in fMRI data that result from small fluctuations in end-tidal CO_2 that occur naturally during normal breathing (Wise et al., 2004). Fluctuations in both the respiratory volume per time (RVT) (Birn et al., 2006) and the cardiac rate (Shmueli et al., 2007) are also present in fMRI data. Spontaneous BOLD fluctuations that occur without a designated stimulus or task also represent a source of structured noise in fMRI data (Biswal et al., 1995). Because physiological noise results from physiological-dependent fluctuations in the baseline signal, it is proportional to the MR signal, S (Kruger et al., 2001):

$$\sigma_P = \lambda S$$

where λ is a tissue-dependent parameter. Physiological noise will increase with the MR signal, which, in turn, will increase with flip angle (up to the Ernst angle) (Haacke et al., 1999), magnetic field strength, or voxel volume (Edelstein et al., 1986). With increasing field strength, physiological noise limits the achievable image signal-to-noise ratio (SNR) (Kruger et al., 2001), but not the BOLD contrast (Gati et al., 1997). However, other factors may limit BOLD contrast above 7 T (Seehafer et al., 2010).

Many retrospective methods have been developed to reduce the cardiac and respiratory aspects of physiological noise. Biswal et al. (1996) used digital notch filters to remove the frequency components of cardiac and respiration noise. This technique fails, however, when the noise is aliased into the frequency spectrum of the task, as fMRI data is generally acquired with a temporal resolution of 2–4 s. Hu et al. (1995) developed a method called RETROKCOR that fits a low-order Fourier series to the k-space time-series data using phase information from the respiratory or cardiac cycles. However, only low-order corrections are possible, and the method introduces unwanted correlations between voxels. Glover et al.

(2000) developed a method called RETROICOR that was similar to the method of Hu et al. (1995) but operates in image space. Cardiac and respiratory signals are monitored and recorded during the scan. Physiological noise is modeled as a low-order Fourier series, which can then be subtracted from voxel time series.

Thomas et al. (2002) used principal component analysis (PCA) and independent component analysis (ICA) methods to isolate and remove structured noise (cardiac and respiration) and random noise (white noise) from fMRI time series. After component decomposition, the method involved spectral analysis, component identification and deletion, and signal reconstruction. Thomas et al. found that ICA was a better method to remove structured noise, while PCA was better at removing random noise. Both methods resulted in increased BOLD contrast sensitivity.

Behzadi et al. (2007) developed a method called CompCor, which applied PCA only to voxel time series exhibiting the highest temporal standard deviations. These voxels were believed to be contaminated with cardiac and respiratory noise. The top six components resulting from this PCA analysis, believed to represent cardiac and respiratory noise, were regressed from the entire data set. The reduction in noise achieved with CompCor was found to be greater than that achieved with RETROICOR with the extra advantage in that physiological monitoring was not required. However, for subjects with especially severe motion artifacts, CompCor identified signal components associated mostly with motion.

Early papers suggested that respiration accounted for 10–20% of the temporal variance at 1.5 T (Raj et al., 2001) or that cardiac and respiration accounted for as much as 30–36% of the noise at 4 T (Thomas et al., 1998). Therefore, existing methods to remove physiological noise have focused predominantly on removing cardiac and respiratory noise. However, recent evidence suggests that spontaneous BOLD activity represents a much larger portion of physiological noise. Fox et al. (2006) demonstrated that coherent spontaneous fluctuations in human brain activity account for a significant fraction of the variability in BOLD fMRI data. Moreover, the authors were able to reduce trial-to-trial variability of BOLD responses in the left hemisphere motor cortex by subtracting scaled signals from the right hemisphere motor cortex, which only contained noise and spontaneous BOLD fluctuations. A study by Chia-Shang et al. (2006) revealed that the vasculature makes a large contribution to the physiological noise during gradient-echo echo-planar imaging (GE-EPI). The authors showed that physiological noise exhibits an echo time (TE) dependence similar to BOLD signal changes. Furthermore, a PCA of complexity showed that the physiological noise had a reduced complexity compared to white noise. The structured nature of this physiological noise is likely to represent spontaneous BOLD fluctuations from active networks in the brain.

In an fMRI study using a repetition time (TR) of 300 ms (yielding a Nyquist frequency of 1.667 Hz, sufficient to sample both respiration (~0.3 Hz) and cardiac (~1 Hz) fluctuations), Petridou at el. (2009) found that time-series physiological fluctuations occurred mainly at frequencies below 0.13 Hz and exhibited a BOLD-like TE dependence. This physiological noise was localized to gray matter and the vasculature. A frequency analysis revealed that the primary respiration and cardiac signals had no significant contribution in white or gray

matter, contradicting earlier studies. A study of noise in fMRI data by Bianciardi et al. (2009a) demonstrated that thermal noise is the most significant source of fMRI signal fluctuations (39% of noise in gray matter voxels), followed by low-frequency drifts (35% of noise in gray matter) and spontaneous BOLD activity (18% of noise in gray matter). Variations in the rates of cardiac pulsation and respiration volume change accounted for a much smaller fraction of noise (2.8% and 2.6% of noise in gray matter, respectively). Cardiac and respiratory fluctuations accounted for by RETROICOR represented a similarly small fraction of the noise (2.5% of noise in gray matter). Low-frequency drift can be accounted for by including polynomials in the regression analysis (Smith et al., 1999). That leaves spontaneous activity as not only the largest source of physiological noise in fMRI data but also as the largest source of fluctuations besides thermal noise that cannot be easily removed. From these results, we conclude that retrospective methods designed to reduce physiological noise should focus predominantly on removing spontaneous BOLD activity.

A recent study has shown that removing spontaneous BOLD fluctuations from fMRI data can enhance sensitivity. Bianciardi et al. (2009b) used a method developed by de Zwart (2008) to identify a reference region of interest (ROI) that correlated with the activation ROI in a separate resting state scan. Bianciardi (2009b) then performed a PCA on the time series in the reference ROI acquired during the fMRI run to identify regressors related to spontaneous BOLD fluctuations. A multiple number of PCA components were used as regressors of no interest in the regression analysis, which resulted in an improvement in the estimated precision.

We developed four new variants of the CompCor method (Behzadi et al., 2007) and compared their ability to reduce physiological noise and increase sensitivity, specificity, and functional contrast-to-noise ratio (CNR). The first variation is the optimized CompCor method, which applies PCA to time series in a noise mask, but orthogonalizes each component to the BOLD response waveform. Next, an algorithm is used to determine the favorable number of components to use as nuisance regressors. Whole brain component correction (WCompCor) applies PCA to time-series throughout the whole brain. Once again, each component is orthogonalized to the BOLD response waveform. Following this, an algorithm chooses the strongest components to use as nuisance regressors.

Low-pass component correction (LCompCor) works like WCompCor, but only identifies low-pass filtered (< 0.1 Hz) components throughout the brain. The low-pass filter is applied to focus our search exclusively on components in the low-frequency domain, where spontaneous BOLD activity is known to dominate (Cordes et al., 2000). Importantly, spontaneous BOLD fluctuations will not be aliased during fMRI acquisitions, as typical pulse sequences have a TR less than 5 s (or a Nyquist frequency greater than 0.1 Hz). Similarly, high-pass component correction (HCompCor) works like WCompCor, but only identifies high-pass filtered (> 0.1 Hz) components. These last two methods were designed specifically to investigate whether the removal of low-frequency or high-frequency noise components had a greater impact on the resulting sensitivity to detect fMRI activation.

Each of the principal component correction methods works on the assumption that a small number of physiological noise fluctuations appear in many voxels distributed throughout the

brain. For example, if there were one physiological noise signal that appeared in most voxels throughout the brain, PCA would be able to identify it as the main component. If physiological noise fluctuations were completely different in every voxel, then regressing out the strongest component identified by PCA would only remove physiological noise for one voxel. The resulting improvement in the sensitivity would be negligible. However, resting state studies have shown that a few distinct components are present in many voxels throughout the brain (De Luca et al., 2006). Therefore, removing the strongest few components discovered using PCA yields the ability to reduce physiological noise in many voxels.

In this study, we collected fMRI data from subjects performing a block-design finger-tapping task. We analyzed the fMRI data using a standard regression analysis, and regression analyses using original CompCor, optimized CompCor, WCompCor, LCompCor, and HCompCor. Our first hypothesis was that regression with optimized CompCor would lead to a greater improvement in the sensitivity to detect fMRI activation compared to regression using original CompCor because the components are orthogonalized to the BOLD response waveform. Our second hypothesis was that regression with WCompCor would lead to a greater improvement in the sensitivity compared to regression using optimized CompCor because components are found using the whole brain. Our third hypothesis was that regression with LCompCor would lead to a greater improvement in the sensitivity compared to regression using HCompCor because low-frequency signals are believed to be the dominant source of noise. Our fourth hypothesis was that regression using LCompCor would lead to a greater improvement in the sensitivity compared to regression using WCompCor because we believe that regressing noise signals of all frequencies should offer no additional benefit and may actually reduce sensitivity. We also examined the functional CNR across the noise reduction methods. Improving the sensitivity and functional CNR of fMRI has the potential to increase the effectiveness of functional brain mapping for clinical applications such as presurgical planning.

2. Materials and methods

2.1 Subjects

Nine healthy subjects (five men, four women, aged 43 ± 10 years) were recruited from the local community. These subjects were assessed by the Edinburgh handedness inventory (Oldfield, 1971) to be strongly right-handed. All volunteers were scanned under a protocol approved by the National Institutes of Health (NIH) Internal Review Board (IRB).

2.2 Experimental design

Subjects were asked to perform a sequential, oppositional finger-tapping task alternating between both hands. At the beginning of each functional run, subjects were told to start tapping each of their four fingers to their thumb on the right hand in succession and continue tapping at a rate that was constant, fast, and comfortable. Every 12 s, an auditory cue (the word "switch") told the subjects to switch hands. Subjects were visually monitored for performance. After each run, subjects were asked if they were able to perform the task well.

Fifteen blocks of 12 s each resulted in exactly three minutes of finger-tapping. This task was expected to produce a robust activation in the primary sensorimotor cortex (SMC). The right and left motor cortices were expected to have opposite activation patterns due to the alternating design of the task.

2.3 Data acquisition

Data were acquired on a Philips Achieva 3.0 T scanner (Philips Medical Systems, Best, Netherlands) using a 6-channel receive-array RF head coil. Subjects were instructed not to move while inside the scanner. Small cushions were used to comfortably secure each subject's head inside the head coil. After a survey scan, a magnetization-prepared rapid gradient echo (MPRAGE) scan (TR/TE/FA = 9.894 ms/4.60ms/8°, FOV = 240 mm, number of slices = 140, slice thickness = 1 mm) was acquired to obtain a high spatial resolution, T_1-weighted image of the brain. Next, nine EPI runs (TR/TE/FA = 2000 ms/35 ms/90°, matrix = 64×64, FOV = 256 mm, number of slices = 32, slice thickness = 4 mm) were acquired. Resulting voxel dimensions were 4.0 mm × 4.0 mm × 4.0 mm. The scanner automatically discarded the first two acquisitions to allow the equilibration of magnetization. An image volume consisting of thirty-two axial slices was acquired during each repetition time to cover most of the brain, focusing on the superior cortex. Over the course of three minutes, the EPI run acquired 90 such image volumes.

Subjects performed the finger-tapping task during the EPI sequences. Each subject was instructed to stay still inside the scanner and to not move his/her head or body during the task. Nine finger-tapping runs were acquired. As this data was originally collected for a reproducibility study on head repositioning (Soltysik et al., 2011), the subject's head was repositioned for the last four finger-tapping runs. The study found that small displacements due to head repositioning in between EPI scans were not a significant source of variability in fMRI activation results.

2.4 Data analysis

MR images were exported from the Philips scanner as NIfTI files. Data were analyzed using AFNI (Cox, 1996), R (R Foundation for Statistical Computing, Vienna, Austria), locally written programs in MATLAB (The Mathworks, Inc., Natick, MA), and Microsoft Excel (2010).

2.5 Pre-processing

The MPRAGE image volume was skull-stripped and warped to the coordinates of Talairach space (Talairach and Tournoux, 1988) using AFNI. Next, the first image volume of each EPI run was aligned to the skull-stripped MPRAGE image volume in subject space using a python script in AFNI (align_epi_anat.py). This script uses a weighted local Pearson coefficient to align a T_2*-weighted image volume to a T_1-weighted image volume (Saad et al., 2009). Alignment results were visually inspected. If corrections were needed, spatial shifts were manually applied prior to rerunning the alignment script.

The EPI runs underwent a six-parameter, rigid-body volume registration routine (Cox and Jesmanowicz, 1999) to correct for any small motion that occurred during the runs. The first

image volume, which had been aligned to the skull-stripped MPRAGE, was used as the base image volume. To assess motion during the runs, we calculated the framewise displacement as performed by Power et al. (NeuroImage, 2014).

EPI intensity information was used to create brain masks (using AFNI's 3dAutomask). Prior to statistical analysis, the voxel intensities of each EPI run were normalized by dividing each voxel value by the mean of its time-series and multiplying by 100.

The data was spatially smoothed in three dimensions using a Gaussian filter with a full width at half maximum (FWHM) of 5 mm (using 3dBlurInMask). The image data were only spatially smoothed inside the brain mask to prevent voxel intensities from being relocated to positions outside the brain.

A rectangular stimulus waveform, representing the right-handed finger-tapping task, was convolved with a standard gamma-variate function (Cohen, 1997) to yield a BOLD response waveform. Using this BOLD response waveform, the activation in the left hemisphere was analyzed.

2.6 Standard regression analysis

A multiple linear regression analysis was performed on the voxel time-series within the brain mask. The baseline data were fit to a second-order polynomial to account for the mean as well as any linear or quadratic trend. The motion parameter files were not used as regressors. The regression analysis was performed using the BOLD response waveform corresponding to the right-handed finger-tapping task as the parameter of interest.

2.7 Original CompCor

We calculated the standard deviation for each voxel time-series within the brain mask. Voxels with a standard deviation in the top 2% were identified to create a non-contiguous noise ROI. Voxels were removed from the noise ROI if the time series of the voxels correlated with the BOLD response waveform ($r > 0.136$, $p < 0.2$). PCA was then run on the time series of the remaining voxels within the noise ROI (using AFNI's 3dpc), extracting the top six ranked PCA components. A multiple linear regression analysis was then performed on the voxel time-series within the brain mask, using the PCA components as regressors of no interest and the BOLD response waveform as the parameter of interest.

2.8 Optimized CompCor

As in the original CompCor method, voxels with a standard deviation in the top 2% were identified to create a noise ROI. Unlike the original CompCor method, voxels were not removed if they correlated with the BOLD response waveform. PCA was then run on the voxel time series within the complete noise ROI (using AFNI's 3dpc), extracting the top fifty PCA components. Next, we orthogonalized each component to the BOLD response waveform so there was no correlation between the two time series. Finally, we used an algorithm to determine the favorable number of PCA components to use as nuisance regressors. Specifically, we calculated the number of time series from voxels in the brain that were significantly correlated ($r > 0.207$, $p < 0.05$) to each PCA component. As an empirical

threshold, we determined the lowest-ranking component that was correlated to at least 10% of the voxels in the brain. We then chose this component and all higher ranking components to use as nuisance regressors in a regression analysis.

2.9 WCompCor

PCA was performed on the EPI voxel time series inside the brain mask (using AFNI's 3dpc). Fifty PCA components were identified. Next, we orthogonalized each component to the BOLD response waveform so there was no correlation between the two time series. The same algorithm described for the optimized CompCor method was used to determine the favorable number of PCA components to use as nuisance regressors in a regression analysis.

2.10 LCompCor

A low-pass (< 0.1 Hz) filter was applied to the time series of the voxels inside the brain mask. PCA was run on these low-pass filtered voxel time series inside the brain mask (using AFNI's 3dpc), extracting fifty PCA components. Following this, we orthogonalized each component to the BOLD response waveform and used the same algorithm to determine the favorable number of PCA components to use as nuisance regressors in a regression analysis.

2.11 HCompCor

A high-pass (> 0.1 Hz) filter was applied to the time series of the voxels inside the brain mask. PCA was run on these high-pass filtered voxel time series inside the brain mask (using AFNI's 3dpc), extracting fifty PCA components. Following this, we orthogonalized each component to the BOLD response waveform and used the same algorithm to determine the favorable number of PCA components to use as nuisance regressors in a regression analysis.

2.12 CNR analysis

The percent signal change for each voxel was calculated by taking the coefficient of fit (the beta value), dividing by the mean of the time series used in the regression analysis, and multiplying by 100%,

$$PSC = \frac{\beta}{\mu} \cdot 100\%.$$

The functional contrast-to-noise ratio (CNR) was calculated as the percent signal change divided by the standard deviation of the residuals resulting from regression analysis,

$$CNR = \frac{PSC}{\sigma_{i\,\text{psu}}}.$$

The CNR was calculated for all voxels inside the true positive activation mask for the last four runs of each subject. An analysis of variance (ANOVA) test was performed to see if the CNR resulting from the different noise reduction methods were different. Post hoc, two-

tailed, pairwise *t*-tests were performed to compare the CNR (averaged across runs) across the different noise reduction methods.

2.13 Sensitivity, specificity, and ROC analysis

For each subject, we used the first five finger-tapping runs to create a true positive activation mask. For a liberal true positive activation mask, we first selected every voxel that was active (using a Bonferroni-corrected threshold of $t > 5.05$, $p_{Bon} < 0.05$) in at least two of these five runs. This activation map was then masked by the expected activation ROI, which included the left hemisphere sensorimotor cortex (LH SMC), the right hemisphere superior cerebellum, and the supplementary motor area (SMA). The remaining active voxels were designated the true positive activation mask. It was assumed that the use of this true positive mask would yield a good approximation for analyzing sensitivity and specificity. Voxels appearing active outside of the true positive activation mask were assumed to be false positives. Voxels not appearing active within this true positive activation mask were assumed to be false negatives.

For the last four runs of each subject, we plotted sensitivity versus threshold, specificity versus threshold, and receiver operator characteristic (ROC) curves (true positive fraction versus false positive fraction). For each run, the area under the sensitivity curve was calculated between *t*-statistic thresholds of 0 and 15. The area under the specificity curve was calculated between *t*-statistic thresholds of 0 and 5. We also calculated the area under the ROC curve. Distributions of the area under the sensitivity curve, the specificity curve, and the ROC curve were created for each method. Analysis of variance (ANOVA) tests were performed to see if the mean area under the curve for different noise reduction methods were different. Following this, two-tailed, pairwise *t*-tests were used to evaluate whether the area under the curve for one noise reduction method was significantly greater than the area under the curve for another method.

We tested four hypotheses regarding sensitivity:

Hypothesis 1: The mean area under the sensitivity curve will be greater when performing a regression analysis with the optimized CompCor method than when performing a regression analysis with the original CompCor method. This is because the components from the optimized CompCor method will be forced to be uncorrelated to the BOLD response waveform, preventing us from regressing out true positive signal changes.

Hypothesis 2: The mean area under the sensitivity curve will be greater when performing a regression analysis with the WCompCor method than when performing a regression analysis with the optimized CompCor method. This is because more accurate noise components are likely to be found using data from the whole brain.

Hypothesis 3: The mean area under the sensitivity curve will be greater when performing a regression analysis with the LCompCor method than when performing a regression analysis with the HCompCor method. This is because low-frequency signals are believed to be the dominant source of noise.

Hypothesis 4: The mean area under the sensitivity curve will be greater when performing a regression analysis with the LCompCor method than when performing a regression analysis with the WCompCor method. This is because using regressors of all frequencies is believed to offer no additional benefit and may actually reduce sensitivity.

3. Results

3.1 Task assessment

Based on visual monitoring, subjects appeared to maintain a constant finger-tapping rate. Upon being asked after each run, subjects said they were able to perform the task easily.

3.2 Motion

Intra-run motion was evaluated by examining the motion parameters resulting from the volume registration process. Across the nine subjects, the average number of time points per run with FD > 0.2 mm varied from 28% to 85%. Based on the framewise displacement analysis and the limit of 0.2 mm suggested by Power et al. (2014), it is likely that motion play a significant role in creating signal changes within the runs.

3.3 Component selection

For each run, we optimized the number of PCA components to use as nuisance regressors separately for four of the noise reduction methods. For optimized CompCor, WCompCor, LCompCor, and HCompCor, we plotted the number of time series in the brain that correlated significantly ($r > 0.207$, $p < 0.05$) with each PCA component. Examples of these plots for one run of one subject are shown in Fig. 1. For optimized CompCor and WCompCor, the first two PCA components were significantly correlated with time series from more than 80% of the voxels in the brain. For LCompCor, the first three PCA components were significantly correlated with time series from more than 60% of the voxels in the brain. For HCompCor, the first four PCA components were significantly correlated to time series from nearly 40% of voxels in the brain. In all cases, the number of time series that were significantly correlated fell quickly as the PCA rank increased.

Across runs, the favorable number of PCA components for optimized CompCor ranged from 9 to 26 (average: 15.10 ± 3.29). For WCompCor, it ranged from 9 to 18 (average: 13.75 ± 2.09). For LCompCor, it ranged from 12 to 37 (average: 16.89 ± 5.51). And for HCompCor, it ranged from 11 to 25 (average: 17.06 ± 3.05).

3.4 Activation

Fig. 2 shows activation in the left hemisphere primary sensorimotor cortex for one run of one subject. Activation was thresholded at $p < 10^{-10}$ uncorrected for multiple comparisons. Compared to the standard regression analysis (Fig. 2A) and original CompCor (Fig. 2B), application of the optimized CompCor (Fig. 2C), WCompCor (Fig. 2D), and LCompCor (Fig. 2E) increased the statistical significance of voxels in the motor cortex. In contrast, application of HCompCor (Fig. 2F) decreased the statistical significance of many voxels across the motor cortex.

3.5 CNR analysis

The functional CNR of the finger-tapping runs resulting from the application of each noise reduction method was averaged across all voxels in the true positive activation mask for the last four runs of all nine subjects. Over all subjects, this included 7,480 different voxels. The average CNR resulting from each noise reduction method is shown in Fig. 3. The ANOVA revealed a significant difference ($p < 0.05$) among the CNR values for the different noise reduction methods. Furthermore, a pairwise t-test confirmed that the CNR for the optimized CompCor method was significantly greater ($p < 0.05$) than that for the original CompCor method, revealing the advantages of the optimization. A pairwise t-test also confirmed that the CNR for the WCompCor method was significantly greater ($p < 0.05$) than that for the optimized CompCor method, revealing the advantages of applying CompCor to all brain voxels. In addition, a pairwise t-test confirmed that the CNR for the LCompCor method was significantly greater ($p < 0.05$) than that for the HCompCor method, revealing that removing low-frequency noise signals increased CNR more than removing high-frequency noise signals. Contrary to our expectations, the CNR for the WCompCor method was found to be significantly greater ($p < 0.05$) than that for the LCompCor method, revealing that the removal of noise signals of all frequencies increased CNR more than the removal of only low-frequency noise signals.

Compared to the average CNR resulting from standard regression (175 ± 90), the average CNR resulting from WCompCor (297 ± 152) was 69% higher. Compared to the average CNR resulting from the original CompCor method (216 ± 103), the average CNR for WCompCor was 38% higher.

3.6 Sensitivity, specificity, and ROC analysis

The sensitivity curves, specificity curves, and ROC curves for one run of one subject are shown in Fig. 4. Differences were seen for sensitivity and specificity curves, but the ROC curves were virtually indistinguishable. These results were similar across subjects and runs.

Fig. 5 (left) shows histograms of the area under the sensitivity curve for each analysis method using data from the last four runs in all nine subjects. For reference, the histogram of the area under the sensitivity curve for the standard regression analysis is shown in Fig. 5A (left). A single-factor ANOVA revealed that these distributions were not the same ($p < 0.05$). Using two-tailed, pairwise t-tests, we found that the mean area under the sensitivity curve for the optimized CompCor method was significantly greater ($p < 0.05$) than the mean area under the sensitivity curve for the original CompCor method, suggesting that the optimized CompCor method was better than the original CompCor method at reducing noise and increasing sensitivity (Fig. 5B,C left). The mean area under the sensitivity curve for the WCompCor method was significantly greater ($p < 0.05$) than the mean area under the sensitivity curve for the optimized CompCor method, revealing that applying CompCor to all brain voxels improved sensitivity more than applying CompCor to an ROI of noisy voxels. (Fig. 5C,D left). The mean area under the sensitivity curve for the LCompCor method was significantly greater ($p < 0.05$) than the mean area under the sensitivity curve for the HCompCor method, revealing that removing low-frequency noise signals increased sensitivity more effectively than removing high-frequency noise signals (Fig. 5E,F left). In

contrast to our expectations, the mean area under the sensitivity curve for the WCompCor method was significantly greater ($p < 0.05$) than the mean area under the sensitivity curve for the LCompCor method, revealing that the removal of noise signals of all frequencies increased sensitivity more than the removal of only low-frequency noise signals (Fig. 5D,E left). Average areas under the sensitivity curves are shown in Table 1. Results of the t-tests are summarized in Table 2. In addition, t-tests revealed that the mean area under the sensitivity curve for each noise reduction method was significantly greater ($p < 0.05$) than the mean area under the sensitivity curve for the standard regression method.

Fig. 5 (center) shows histograms of the area under the specificity curve for each analysis method using data from the last four runs in all nine subjects. For reference, the histogram of the area under the specificity curve for the standard regression analysis is shown in Fig. 5A (center). A single-factor ANOVA revealed that the distributions were not the same ($p < 0.05$). Using two-tailed, pairwise t-tests, we found that the mean area under the specificity curve for the original CompCor method was significantly greater ($p < 0.05$) than the mean area under the specificity curve for the optimized CompCor method, suggesting that the optimized CompCor method yielded more false positives (Fig. 5B,C center). The mean area under the specificity curve for the optimized CompCor method was significantly greater ($p < 0.05$) than the mean area under the specificity curve for the WCompCor method, suggesting that the WCompCor method yielded even more false positives (Fig. 5C,D center). The mean area under the specificity curve for the HCompCor method was significantly greater ($p < 0.05$) than the mean area under the specificity curve for the LCompCor method, suggesting that removing low-frequency noise yielded more false positives than removing high-frequency noise (Fig. 5E,F center). Finally, the mean area under the specificity curve for the LCompCor method was significantly greater ($p < 0.05$) than the mean area under the specificity curve for the WCompCor method, suggesting that the WCompCor method yielded more false positives than the LCompCor method (Fig. 5D,E center). Average areas under the specificity curves are shown in Table 1. Results of the t-tests are summarized in Table 2. In addition, t-tests revealed that the mean area under the specificity curve for each noise reduction method was significantly less ($p < 0.05$) than the mean area under the specificity curve for the standard regression method.

Fig. 5 (right) shows distributions of the mean area under the ROC curve for all analysis methods using data from the last four runs of all nine subjects. A single-factor ANOVA revealed that there was no difference between the distributions ($p < 0.05$). Average areas under the ROC curves are shown in Table 1. In general, the mean area under the ROC curve was similar across methodology.

4. Discussion

In this study, we compared the results of five different PCA methods to reduce physiological noise and motion artifacts from EPI data sets in order to improve the sensitivity of detecting task-related fMRI activation. We confirmed our first hypothesis that the average area under the sensitivity curve for the optimized CompCor method was significantly greater ($p < 0.05$) than the average area under the sensitivity curve for the original CompCor method. In addition, the optimized CompCor method yielded a significantly higher functional CNR ($p <$

0.05) compared to the original CompCor method. These results revealed that the methodology added for the optimized CompCor method (orthogonalizing the PCA components to the BOLD response waveform and choosing a favorable number of PCA components to use as nuisance regressors) improved the ability to reduce noise and detect activation compared to the original CompCor method.

Next, we confirmed our second hypothesis that the average area under the sensitivity curve for the WCompCor method was significantly greater ($p < 0.05$) than the average area under the sensitivity curve for the optimized CompCor method. In addition, the WCompCor method yielded a significantly higher functional CNR ($p < 0.05$) compared to the optimized CompCor method. These results revealed that applying PCA to time series from voxels throughout the entire brain mask improved the ability to reduce noise and detect activation compared to applying PCA to a small ROI of noisy voxels.

Next, we confirmed our third hypothesis that the average area under the sensitivity curve for the LCompCor method was significantly greater ($p < 0.05$) than the average area under the sensitivity curve for the HCompCor method. In addition, the LCompCor method yielded a significantly higher functional CNR ($p < 0.05$) compared to the HCompCor method. These results revealed that the use of low-frequency PCA components as nuisance regressors was much better than the use of high-frequency PCA components as nuisance regressors for reducing noise and improving the detection of activation.

Finally, we rejected our fourth hypothesis, instead finding that the average area under the sensitivity curve for the WCompCor method was significantly greater ($p < 0.05$) than the average area under the sensitivity curve for the LCompCor method. In addition, the WCompCor method yielded a significantly higher functional CNR ($p < 0.05$) compared to the LCompCor method. These results revealed that applying PCA to low-frequency filtered time series offered no improvement in sensitivity compared to applying PCA to unfiltered time series.

When comparing the five different PCA methods for reducing noise, we found that the WCompCor method yielded the greatest area under the sensitivity curve and the highest average functional CNR. In addition, the WCompCor method required the lowest number of PCA components to use as nuisance regressors (averaging 13.75 ± 2.09) compared to optimized CompCor (average: 15.10 ± 3.29), LCompCor (average: 16.89 ± 5.51), and HCompCor (average: 17.06 ± 3.05). These results suggest that WCompCor was best-suited for identifying non-task-related sources of noise in the time-series data. Therefore, to reduce noise and improve the sensitivity of task-based fMRI activation maps, we recommend applying the WCompCor method over all other PCA methods tested.

In contrast to the increase in sensitivity, our results showed that each of the regressor methods decreased the specificity of the resulting activation map compared to the standard regression analysis. These results reveal that the increase in sensitivity and functional CNR for each of these noise reduction methods comes with the price of a decrease in specificity.

One explanation for this decrease in specificity is that the increase in functional CNR allowed the regression analysis to more easily identify task-correlated motion signals, which

would increase the presence of false positive activation. The framewise displacement analysis revealed that, across runs, between 28% and 85% of the time points had a framewise displacement greater than 0.2 mm. If this excess motion contributed to task-correlated signal changes, it would have decreased the specificity, as these signals would not have been removed with the PCA methods. It is possible that "scrubbing" motion-contaminated time points may have increased specificity. However, too many time points were affected by motion in these runs to consider removing motion-contaminated data. The potential increase in false positive activation tempers the enthusiasm one would otherwise have for the increase in sensitivity seen for these noise reduction methods. Although the current study used a task duration ideally designed to separate true BOLD responses from task-correlated motion, longer or shorter task durations may lead to an even greater decrease in specificity as task-correlated motion signals become difficult to separate from true BOLD responses. The current study also used a finger-tapping task, which is more likely to yield task-correlated motion artifacts than cognitive tasks not requiring motion. Applying WCompCor to an fMRI study with a cognitive task is expected to result in less task-correlated motion and less of a decrease in specificity.

Previous studies using nuisance regressors did not take into account the effect of the analysis technique on both sensitivity and specificity. The study introducing RETROICOR (Glover et al., 2000) provided no analysis of sensitivity or specificity. The study introducing CompCor (Behzadi et al., 2007) reported an analysis using simulated data that showed improved ROC curves when applying CompCor compared to RETROICOR or uncorrected data. However, our study using human subject data found no significant difference between the ROC curves of CompCor and standard regression. The PCA method used by Bianciardi et al. (2009b) did not analyze specificity. Future work may be needed to more clearly understand the effect of different noise regressor methods on both sensitivity and specificity of fMRI activation.

The LCompCor method regressed out PCA components in the low frequency range (< 0.1 Hz), where spontaneous BOLD activity is believed to be dominant (Cordes et al., 2000). If spontaneous BOLD fluctuations were the dominant source of physiological noise, we would have expected the LCompCor method to yield equal or better results than the WCompCor method. Unexpectedly, the WCompCor method yielded better improvements in sensitivity compared to the LCompCor method. One explanation for why WCompCor yielded a sensitivity higher than LCompCor is that there were strong sources of noise besides the spontaneous BOLD fluctuations that were reduced using WCompCor. For example, our analysis revealed that for each run, between two and nine PCA components correlated significantly with the framewise displacement vector, revealing that motion-related noise was captured by the principal components analysis. Also, the global mean signal is likely to be composed of a wide range of frequencies, and WCompCor would have been more effective than LCompCor at removing it. A second explanation is that spontaneous BOLD fluctuations might occur above the commonly cited 0.1 Hz threshold as one recent study has shown (Niazy et al., 2011). In either case, our data suggest that to most reliably reduce noise in task-based fMRI data, PCA should be applied to unfiltered time series from voxels throughout the brain.

We did not acquire cardiac and respiratory data during the fMRI acquisitions, so we were unable to directly compare our methods with RETROICOR (Glover et al., 2000). The RETROICOR method is specifically designed to remove only cardiac and respiratory noise components. Behzadi et al. (2009) found that the CompCor method was superior to the RETROICOR method in terms of the activation volume (i.e., sensitivity). In our study, we found that the WCompCor method was superior to the original CompCor method in terms of sensitivity and CNR. As we have shown, the WCompCor method is designed to remove many different types of noise components, including but not limited to cardiac and respiratory fluctuations. Because RETROICOR only removes cardiac and respiratory signals, we believe that WCompCor should be superior to RETROICOR for improving the sensitivity of detecting task-based fMRI activation. Future research could employ the PESTICA algorithm (Beall, 2010) to compare the abilities of RETROICOR and WCompCor to retrospectively remove noise from fMRI data.

For every regressor of no interest that is included in the regression analysis, there is a loss of one in the degrees of freedom, reducing the significance of the resulting test statistic. As a result, there are diminishing returns for regressing out an increasing number of nuisance regressors. In our analysis, the activation of the final regression analysis was evaluated using the adjusted degrees of freedom. Previous papers using PCA appear did not adjust for the change in the degrees of freedom due to regressing out PCA components, which questions the optimization of these studies (Thomas et al., 2002, Behzadi et al., 2007). Because we adjusted the degrees of freedom for every regressor of no interest included in the regression analysis, we believe that our study yielded a more believable optimization of the CompCor method to reduce noise.

A disadvantage of a component correction method like WCompCor is that there is no single ideal number of components to regress from the data. We devised an algorithm to determine a favorable number of components to use as nuisance regressors. However, determining the number of nuisance regressors separately for each run can complicate the data analysis. For our data, we chose between 9 and 18 (average: 13.75 ± 2.09) components to use as nuisance regressors for WCompCor.

Another disadvantage of a component correction method like WCompCor is that it may not work well for short time-series. The degrees of freedom in the regression analysis is equal to the number of time points in the fMRI run minus the number of parameters used in the regression model. The shorter the time-series, the fewer degrees of freedom remaining after applying an identical regression model. Our fMRI runs had 90 image volumes (or time points) each. After fitting the baseline to a polynomial of second order (three parameters) and a BOLD response waveform (one parameter), there were 86 degrees of freedom remaining. If we also regressed out 14 components of no interest, that would have left us with 72 degrees of freedom. If our fMRI runs had been acquired with only 50 image volumes, we would have been left with only 32 degrees of freedom. Such a small number of degrees of freedom would lower the statistical significance, potentially reducing any possible benefit from the WCompCor method.

The PCA method has the ability to identify the global mean signal. Previously, researchers have pointed out that removing the global mean signal can yield spurious negative activation in resting state studies (Murphy et al., 2009). However, our method orthogonalized each PCA component to the BOLD response waveform, preventing any increase in spurious negative activation caused by regressing out the unfiltered global mean signal.

Finally, it should be stressed that the WCompCor method is intended for use on task-related fMRI data and should not be applied to resting state fMRI data where the intent is to find resting state networks. The WCompCor method is likely to remove resting state BOLD fluctuations that form resting state networks, as these fluctuations also contribute strongly to physiological noise. If investigators are interested in removing physiological noise from resting state fMRI data, the only known physiological noise of interest would be non-BOLD physiological noise, such as cardiac and respiratory signal changes. These signals can be removed using either RETROICOR (Glover et al., 2000) or CompCor (Behzadi et al., 2007).

5. Conclusions

We compared the results of five different PCA methods (original CompCor, optimized CompCor, WCompCor, LCompCor, and HCompCor) to reduce physiological noise and motion artifacts and increase the sensitivity of detecting fMRI activation. First, we found that orthogonalizing PCA components to the BOLD response and choosing a favorable number per run (optimized CompCor) greatly improved the sensitivity. Second, we found that applying PCA to voxel time series throughout the whole brain (WCompCor) further improved the sensitivity. Third, we found that applying to PCA to low-pass filtered data (LCompCor) improved the sensitivity much better than when applying PCA to high-pass filtered data (HCompCor). Fourth, we found that applying PCA to low-pass filtered data (LCompCor) offered no improvement over applying PCA to unfiltered time series (WCompCor). Therefore, to reduce noise and improve the sensitivity of task-based fMRI activation maps, we recommend application of the WCompCor method. The disadvantage of this method, as with all of the PCA methods tested, is that the specificity is decreased. This decrease in specificity may be due to the increased sensitivity of detecting task-correlated motion artifacts.

The fact that sensitivity increases while specificity decreases may limit the utility of this noise reduction method. The use of WCompCor may be best suited for an application like presurgical planning, where sensitivity is valued more than specificity. Investigators asking general neuroscience questions that demand both high sensitivity and specificity may wish to avoid using noise regressors in their analysis.

Acknowledgments

References

Beall EB. Adaptive cyclic physiologic noise modeling and correction in functional MRI. Journal of Neuroscience Methods. 2010; 187:216–228. [PubMed: 20096307]

Behzadi Y, Restom K, Liau J, Liu TT. A component based noise correction method (CompCor) for BOLD and perfusion fMRI. NeuroImage. 2007; 37:90–101. [PubMed: 17560126]

Bianciardi M, Fukunaga M, van Gelderen P, Horovitz SG, de Zwart JA, Shmueli K, Duyn JH. Sources of functional magnetic resonance imaging signal fluctuations in the human brain at rest: a 7 T study. Magn Reson Imaging. 2009a; 27:1019–1029. [PubMed: 19375260]

Bianciardi M, van Gelderen P, Duyn JH, Fukunaga M, de Zwart JA. Making the most of fMRI at 7 T by suppressing spontaneous signal fluctuations. NeuroImage. 2009b; 44:448–454. [PubMed: 18835582]

Birn RM, Bandettini PA, Cox RW, Shaker R. Event-related fMRI of tasks involving brief motion. Hum Brain Mapp. 1999; 7:106–114. [PubMed: 9950068]

Birn RM, Cox RW, Bandettini PA. Experimental designs and processing strategies for fMRI studies involving overt verbal responses. NeuroImage. 2004; 23:1046–1058. [PubMed: 15528105]

Birn RM, Diamond JB, Smith MA, Bandettini PA. Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. NeuroImage. 2006; 31:1536–1548. [PubMed: 16632379]

Biswal B, DeYoe AE, Hyde JS. Reduction of physiological fluctuations in fMRI using digital filters. Magn Reson Med. 1996; 35:107–113. [PubMed: 8771028]

Biswal B, Yetkin FZ, Haughton VM, Hyde JS. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magn Reson Med. 1995; 34:537–541. [PubMed: 8524021]

Bullmore ET, Brammer MJ, Rabe-Hesketh S, Curtis VA, Morris RG, Williams SCR, Sharma T, McGuire PK. Methods for diagnosis and treatment of stimulus-correlated motion in generic brain activation studies using fMRI. Hum Brain Mapp. 1999; 7:38–48. [PubMed: 9882089]

Cohen MS. Parametric analysis of fMRI data using linear systems methods. NeuroImage. 1997; 6:93–103. [PubMed: 9299383]

Cordes D, Haugton VM, Arfanakis K, Wendt GJ, Turski PA, Moritz CH, Quiglet MA, Meyerand MA. Mapping functionally related regions of brain with functional connectivity MR imaging. AJNR. 2000; 21:1636–1644. [PubMed: 11039342]

Cox RW. AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res. 1996; 29:162–173. [PubMed: 8812068]

Cox RW, Jesmanowicz A. Real-time 3D image registration for functional MRI. Magn Reson Med. 1999; 42:1014–1018. [PubMed: 10571921]

Dagli MS, Ingeholm JE, Haxby JV. Localization of cardiac-induced signal change in fMRI. NeuroImage. 1999; 9:407–415. [PubMed: 10191169]

De Luca M, Beckman CF, De Stefano N, Matthews PM, Smith SM. fMRI resting state networks define distinct modes of long-distance interactions in the human brain. NeuroImage. 2006; 29:1359–1367. [PubMed: 16260155]

de Zwart JA, Gelderen PV, Fukunaga M, Duyn JH. Reducing correlated noise in fMRI data. Magn Reson Med. 2008; 59:939–945. [PubMed: 18383291]

Edelstein WA, Glover GH, Hardy CJ, Redington RW. The intrinsic signal-to-noise ratio in NMR imaging. Magn Reson Med. 1986; 3:604–618. [PubMed: 3747821]

Fox MD, Snyder AZ, Zacks JM, Raichle ME. Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. Nat Neurosci. 2006; 9:23–25. [PubMed: 16341210]

Friston KJ, Williams S, Howard R, Frackowiak RSJ, Turner R. Movement-related effects in fMRI time-series. Magn Reson Med. 1996; 35:346–355. [PubMed: 8699946]

Gati JS, Menon RS, Ugurbil K, Rutt BK. Experimental determination of the BOLD field strength dependence in vessels and tissue. Magn Reson Med. 1997; 38:296–302. [PubMed: 9256111]

Glover GH, Li TQ, Ress D. Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magn Reson Med. 2000; 44:162–167. [PubMed: 10893535]

Haacke, EM.; Brown, RW.; Thompson, MR.; Venkatesan, R. Magnetic resonance imaging: physical principles and sequence design. New York: Wiley-Liss; 1999.

Hu X, Le TH, Parrish T, Erhard P. Retrospective estimation and correction of physiological fluctuation in functional MRI. Magn Reson Med. 1995; 34:201–212. [PubMed: 7476079]

Huettel, SA.; Song, AW.; McCarthy, G. Functional magnetic resonance imaging. Sunderland, Massachusetts: Sinauer Associates, Inc.; 2004.

Krueger G, Kastrup A, Glover GH. Neuroimaging at 1.5 T and 3.0 T: Comparison of oxygenation-sensitive magnetic resonance imaging. Magn Reson. Med. 2001; 45:595–604. [PubMed: 11283987]

Liu CSJ, Miki A, Hulvershorn J, Bloy L, Gualtieri EE, Liu GT, Leigh JS, Haselgrove JC, Elliot MA. Spatial and temporal characteristics of physiological noise in fMRI at 3T. Acad Radiol. 2006; 13:313–323. [PubMed: 16488843]

Mitra PP, Ogawa SO, Hu X, Ugurbil K. The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging. Magn Reson Med. 1997; 37:511–518. [PubMed: 9094072]

Murphy K, Birn RM, Handwerker DA, Jones TB, Bandettini PA. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? NeuroImage. 2009; 44:893–905. [PubMed: 18976716]

Niazy RK, Xie J, Miller K, Beckmann CF, Smith SM. Chapter 17 – Spectral characteristics of resting state networks. Prog Brain Res. 2011; 193:259–276. [PubMed: 21854968]

Oldfield RC. The assessment and analysis of handedness: The Edinburgh Inventory. Neuopsychologia. 1971; 9:97–113.

Petridou N, Schafer A, Gowland P, Bowtell R. Phase vs. magnitude information in functional magnetic resonance imaging time series: toward understanding the noise. Magn Reson Imaging. 2009; 27:1046–1057. [PubMed: 19369024]

Power JD, Mitra A, Laumann TO, Snyder AZ, Schlaggar BL, Petersen SE. Methods to detect, characterize, and remove motion artifact in resting state fMRI. Neuroimage. 2014; 84:320–341. [PubMed: 23994314]

Raj D, Anderson AW, Gore JC. Respiratory effects in human functional magnetic resonance imaging due to bulk susceptibility changes. Phys Med Biol. 2001; 46:3331–3340. [PubMed: 11768509]

Saad ZR, Glen DR, Chen G, Beauchamp MS, Desai R, Cox RW. A new method for improving function-to-structural MRI alignment using local Pearson correlation. NeuroImage. 2009; 44:839–848. [PubMed: 18976717]

Seehafer JU, Kalthoff D, Farr TD, Wiedermann D, Hoehn M. No increase of the blood oxygenation level-dependent functional magnetic resonance imaging signal with higher field strength: implications for brain activation studies. The Journal of Neuroscience. 2010; 30:5234–5241. [PubMed: 20392946]

Shmueli K, van Gelderen P, de Zwart JA, Horovitz SG, Fukunaga M, Jansma JM, Duyn JH. Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI BOLD signal. NeuroImage. 2007; 38:306–320. [PubMed: 17869543]

Smith AM, Lewis BK, Ruttimann UE, Ye FQ, Sinnwell TM, Yang Y, Duyn JH, Frank JA. Investigation of low frequency drift in fMRI signal. NeuroImage. 1999; 9:526–533. [PubMed: 10329292]

Soltysik DA, Hyde JS. Strategies for block-design fMRI experiments during task-related motion of structures of the oral cavity. NeuroImage. 2006; 29:1260–1271. [PubMed: 16275020]

Soltysik DA, Thomasson D, Rajan S, Gonzalez-Castillo J, DiCamillo P, Biassou N. Head-repositioning does not reduce the reproducibility of fMRI activation in a block-design motor task. NeuroImage. 2011; 56:1329–1337. [PubMed: 21406235]

Talairach, J.; Tournoux, P. Co-planar stereotaxic atlas of the human brain. New York: Thieme Medical; 1988.

Thomas CG, Harshman RA, Menon RS. Noise reduction in BOLD-based fMRI using component analysis. NeuroImage. 2002; 17:1521–1537. [PubMed: 12414291]

Thomas CG, Menon RS. Amplitude response and stimulus presentation frequency response of human primary visual cortex using BOLD EPI at 4 T. Magn Reson Med. 1988; 40:203–209. [PubMed: 9702702]

Weisskoff, RM.; Baker, J.; Belliveau, J.; Davis, TL.; Kwong, KK.; Cohen, MS.; Rosen, BR. Power spectrum analysis of functionally-weighted MR data: What's in the noise?; Proceedings of the 12th annual meeting of the society of magnetic resonance in medicine; 1993. p. 7

Wise RG, Ide K, Poulin MJ, Tracey I. Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in BOLD signal. NeuroImage. 2004; 21:1652–1664. [PubMed: 15050588]

Yetkin FZ, Haughton VM, Cox RW, Hyde J, Birn RM, Wong EC, Prost R. Effect of motion outside the field of view on functional MR. AJNR. 1996; 17:1005–1009. [PubMed: 8791907]

Highlights

- We developed four new PCA methods to identify nuisance regressors in fMRI analysis

- We compared these PCA methods with CompCor, an established PCA method

- The best improvement in CNR and sensitivity resulted from the whole brain component correction (WCompCor) method

- However, regressing noise signals showed a paradoxical consequence of reducing specificity for all noise reduction methods attempted
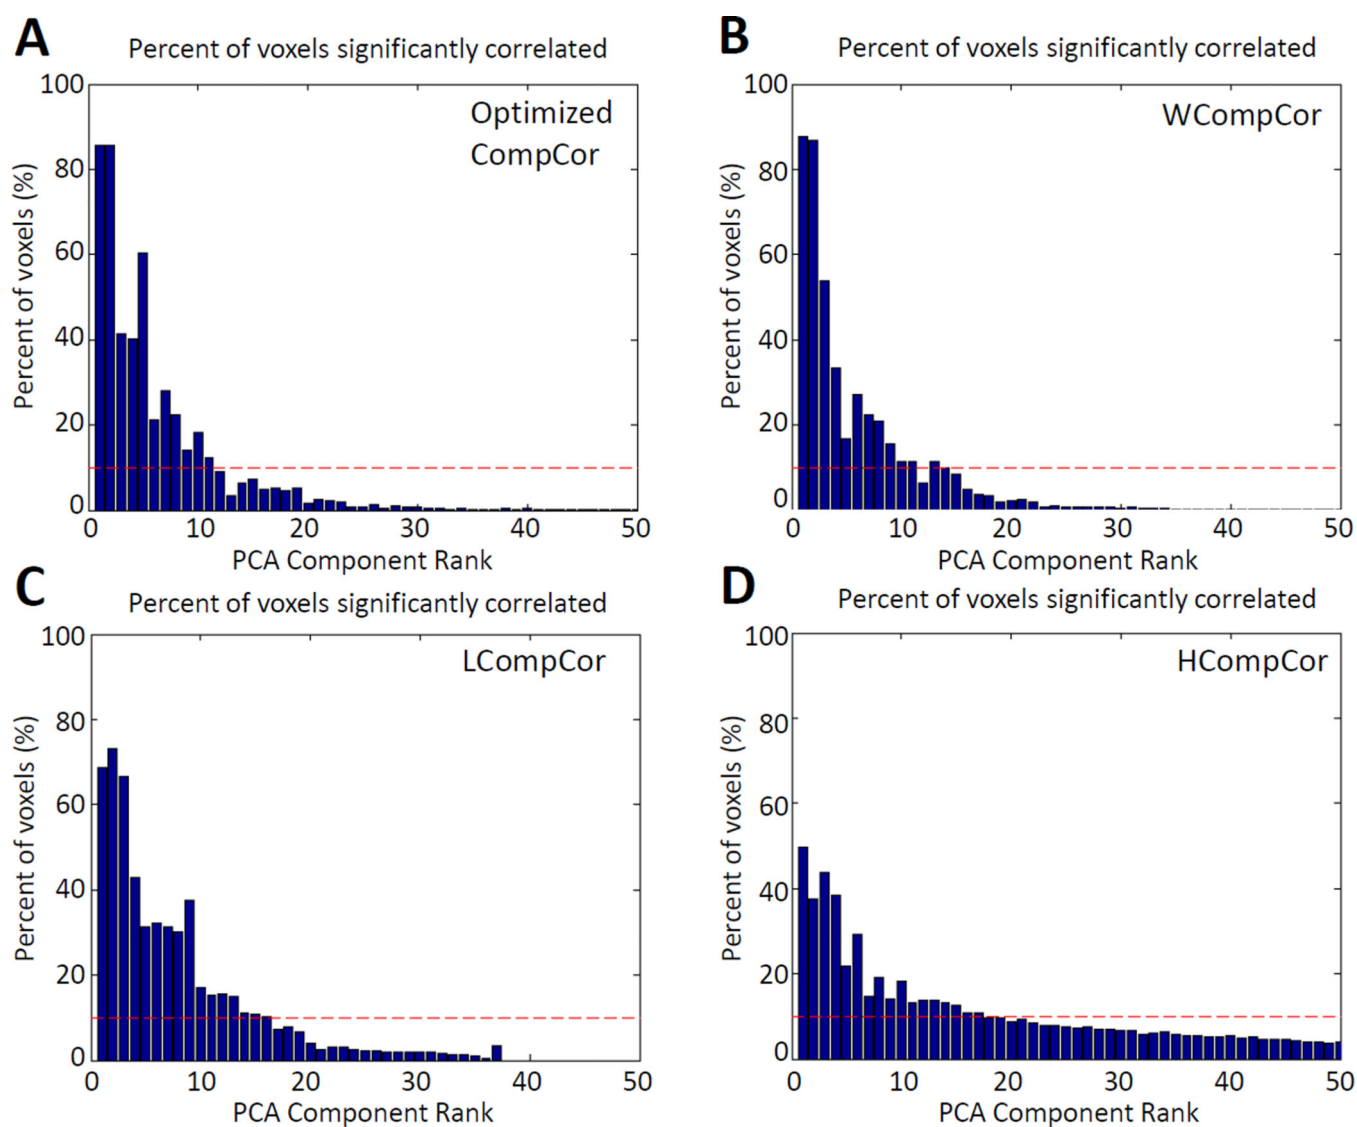
Figure 1.
Plots of the percentage of voxels that were significantly correlated ($p < 0.05$) to each PCA component versus the PCA component rank for one run of one subject. A) Optimized CompCor, B) WCompCor, C) LCompCor, and D) HCompCor. The red dashed line indicates the level of 10%.
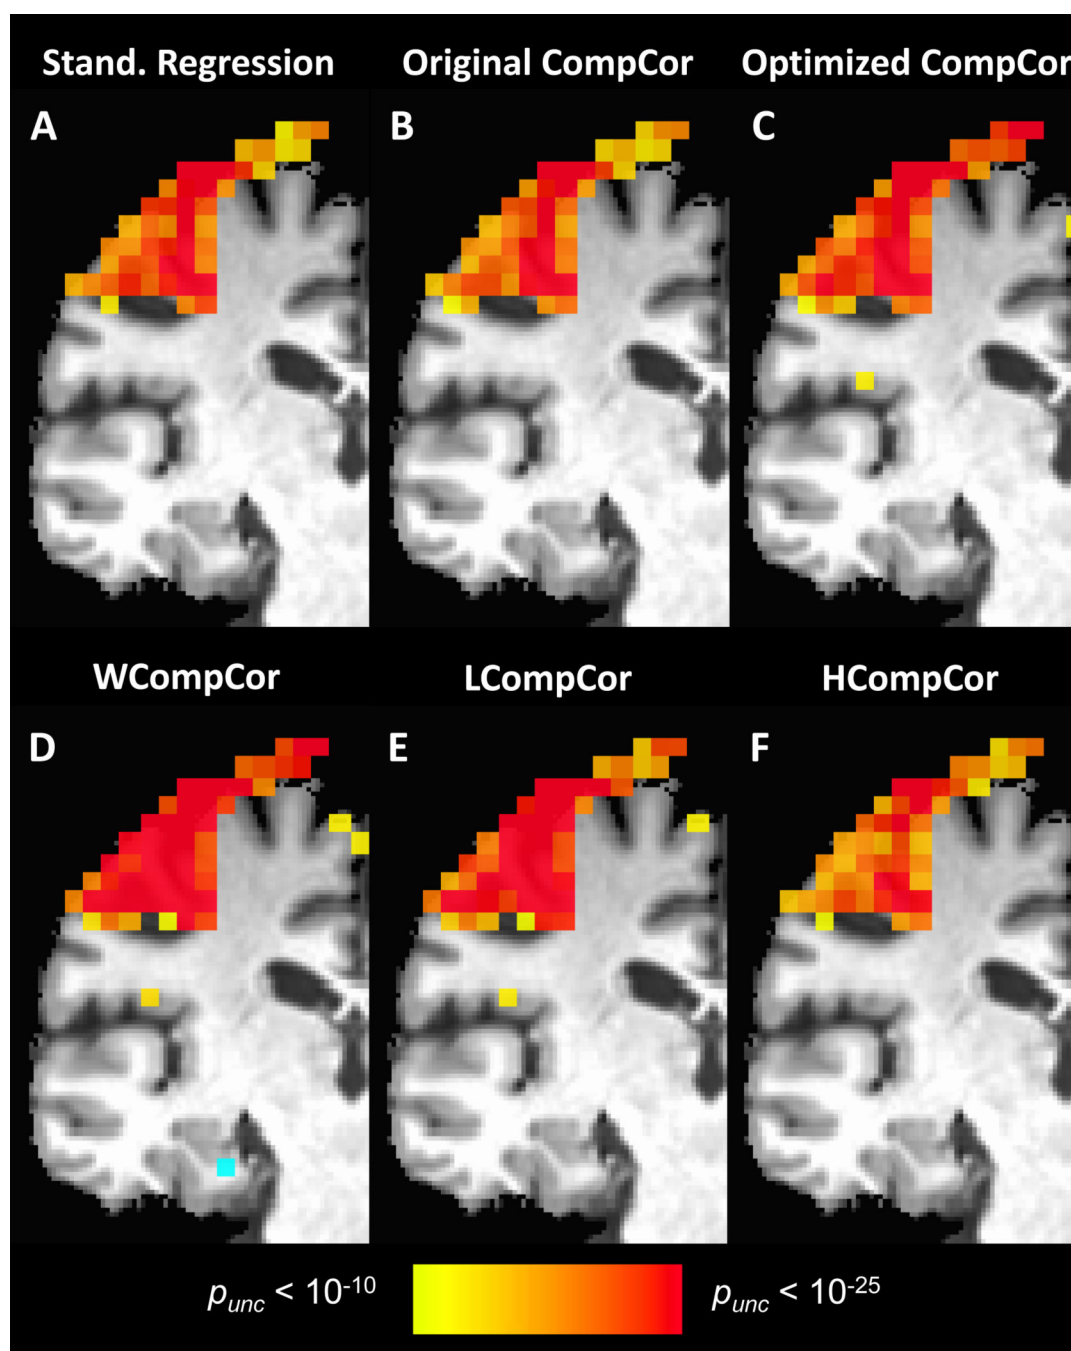
Figure 2.
Coronal activation maps show activation in the left hemisphere primary sensorimotor cortex for A) standard regression, B) regression with original CompCor, C) regression with optimized CompCor, D) regression with WCompCor, E) regression with LCompCor, and F) regression with HCompCor.
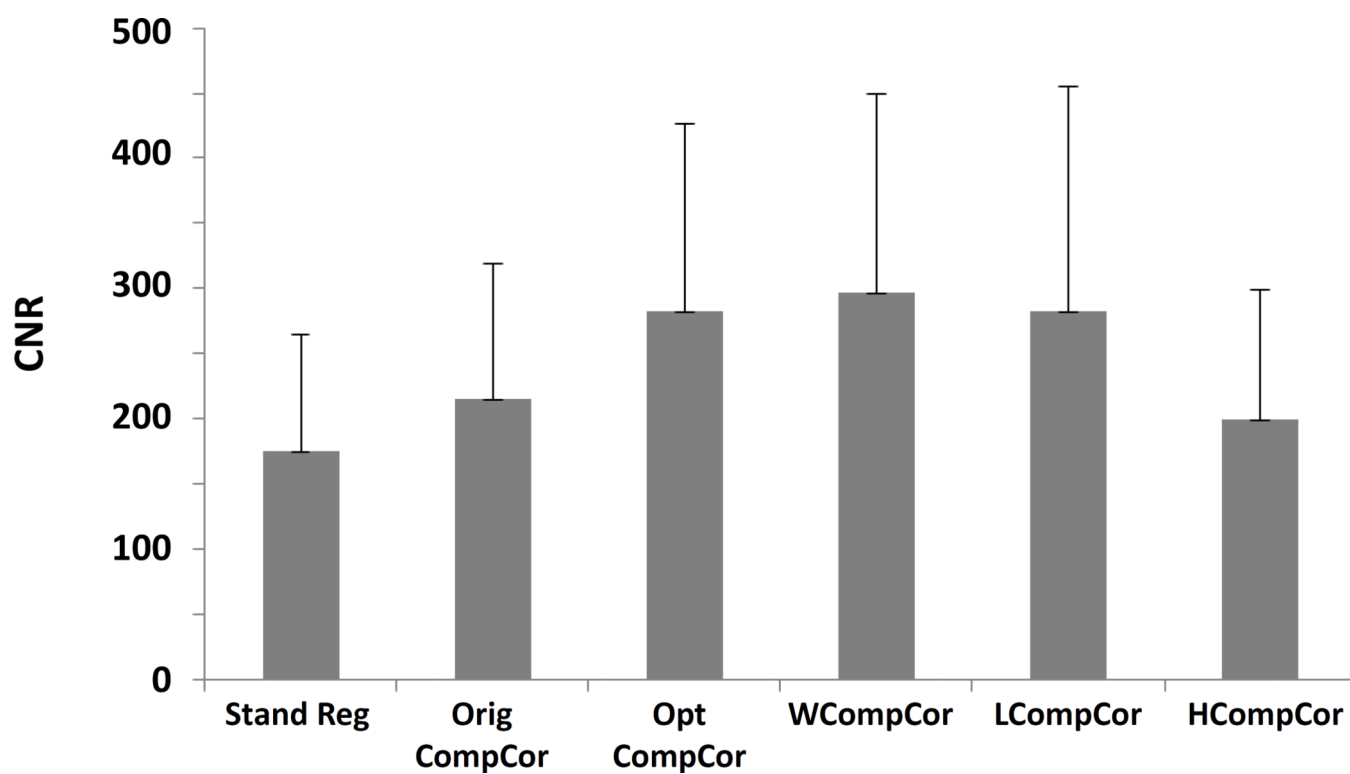
Figure 3.
CNR, averaged across all voxels in the true positive activation mask for the last four runs of all nine subjects, for each noise reduction method.
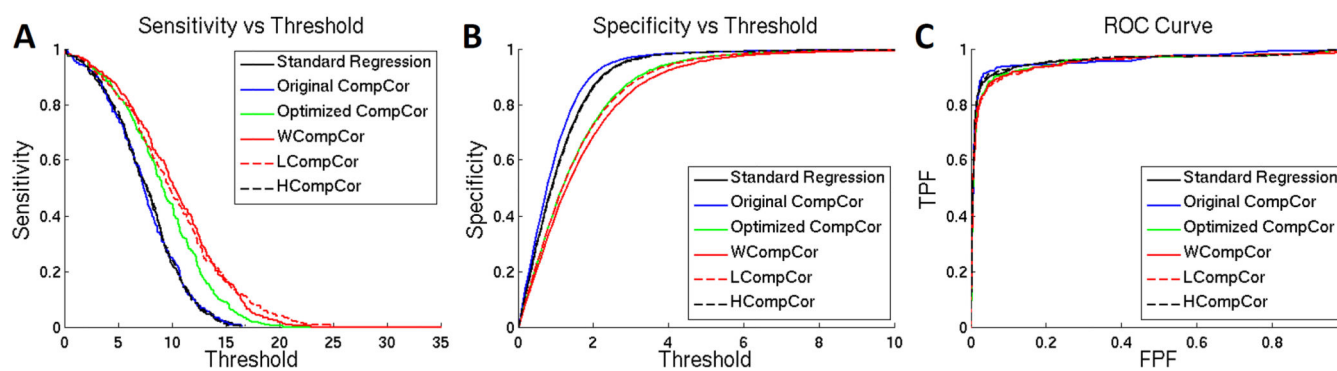
Figure 4.
(A) Sensitivity plotted against threshold, (B) specificity plotted against threshold, and (C) the ROC curves for one run of one subject.
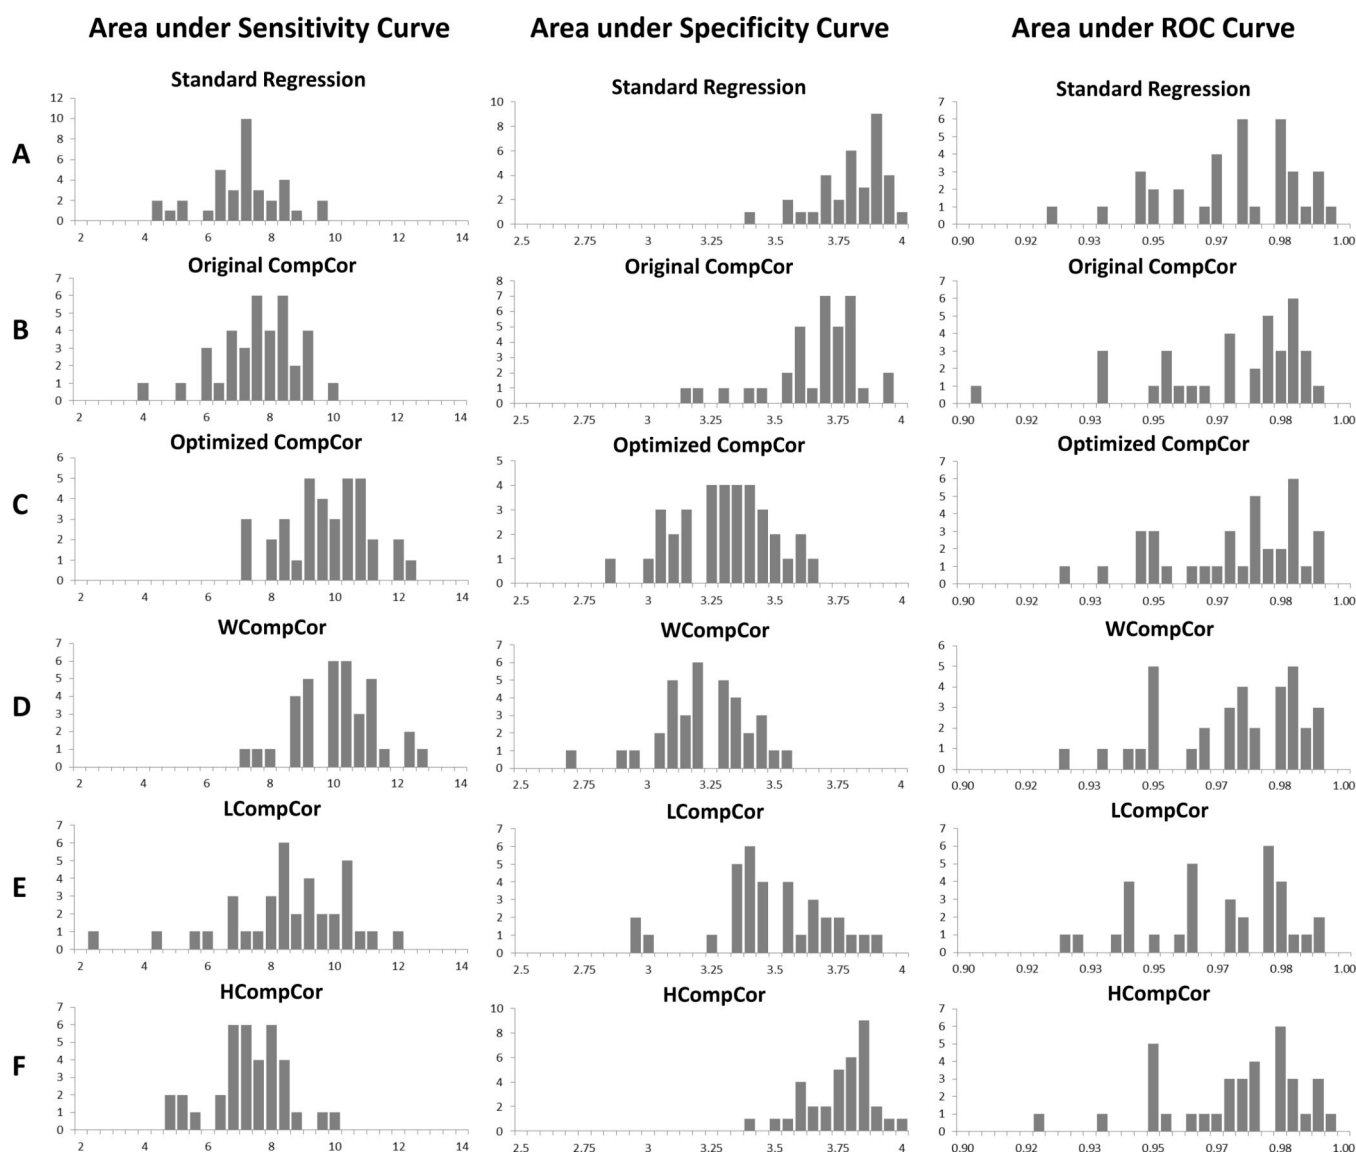
Figure 5.
Histograms of area under the sensitivity curve (left), area under the specificity curve (center), and area under the ROC curve (right) for five different analysis methods: (A) standard regression, (B) original CompCor, (C) optimized CompCor, (D) WCompCor, (E) LCompCor, and (F) HCompCor.

Table 1

Average areas under the sensitivity, specificity, and ROC curves (mean ± standard deviation)

Method	Average area under the sensitivity curve	Average area under the specificity curve	Average area under the ROC curve
Standard Regression	6.85 ± 1.25	3.78 ± 0.14	0.969 ± 0.017
Original CompCor	7.42 ± 1.23	3.63 ± 0.18	0.969 ± 0.020
Optimized CompCor	9.55 ± 1.33	3.28 ± 0.19	0.969 ± 0.017
WCompCor	9.94 ± 1.29	3.20 ± 0.18	0.969 ± 0.017
LCompCor	8.39 ± 1.95	3.47 ± 0.25	0.960 ± 0.029
HCompCor	7.08 ± 1.21	3.73 ± 0.13	0.970 ± 0.017

Table 2

Results of the pairwise *t*-tests comparing average areas under the curves

Hypothesis	*p*-value for comparing average areas under the sensitivity curve	*p*-value for comparing average areas under the specificity curve
Optimized CompCor > Original CompCor	3.57×10^{-11} *	4.68×10^{-14} **
WCompCor > Optimized CompCor	1.51×10^{-6} *	8.12×10^{-10} **
LCompCor > HCompCor	6.02×10^{-5} *	2.19×10^{-9} **
LCompCor > WCompCor	1.38×10^{-11} **	1.71×10^{-11} *

*indicates significance and confirmation of hypothesis,

**indicates significance and confirmation of opposite of hypothesis