Title

Global signal regression acts as a temporal downweighting process in resting-state fMRI.

Authors

Nalci, Alican
Rao, Bhaskar D
Liu, Thomas T

Publication Date

2017-05-01

DOI

10.1016/j.neuroimage.2017.01.015

Peer reviewed

Global Signal Regression Acts as a Temporal Downweighting Process in Resting-State fMRI

Alican Nalci[a,b,*], Bhaskar D. Rao[b], Thomas T. Liu[a,c,**]

[a] *Center for Functional MRI, University of California San Diego, 9500 Gilman Drive MC 0677, La Jolla, CA 92093*
[b] *Department of Electrical and Computer Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093*
[c] *Departments of Radiology, Psychiatry, and Bioengineering, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093*

Abstract

In resting-state functional MRI (rsfMRI), the correlation between blood oxygenation level dependent (BOLD) signals across different brain regions is used to estimate the functional connectivity of the brain. This approach has led to the identification of a number of resting-state networks, including the default mode network (DMN) and the task positive network (TPN). Global signal regression (GSR) is a widely used pre-processing step in rsfMRI that has been shown to improve the spatial specificity of the estimated resting-state networks. In GSR, a whole brain average time series, known as the global signal (GS), is regressed out of each voxel time series prior to the computation of the correlations. However, the use of GSR is controversial because it can introduce artifactual negative correlations. For example, it has been argued that anticorrelations observed between the DMN and TPN are primarily an artifact of GSR. Despite the concerns about GSR, there is currently no consensus regarding its use. In this paper, we introduce a new framework for understanding the effects of GSR. In particular, we show that the main effects of GSR can be well approximated as a temporal downweighting process in which the data from time points with relatively large GS magnitudes are greatly attenuated while data from time points with relatively small GS magnitudes are largely unaffected. Furthermore, we show that a limiting case of this downweighting process in which data from time points with large GS magnitudes are censored can also approximate the effects of GSR. In other words, the correlation maps obtained after GSR show a high degree of spatial similarity (including the presence of anticorrelations between the DMN and TPN)

*Corresponding Author
**Principal Corresponding Author
Email addresses: analci@ucsd.edu (Alican Nalci), ttliu@ucsd.edu (Thomas T. Liu)

with maps obtained using only the uncensored (i.e. retained) time points. Since the data from these retained time points are unaffected by the censoring process, this finding suggests that the observed anticorrelations inherently exist in the data from time points with small GS magnitudes and are not simply an artifact of GSR.

Keywords: fMRI, global signal regression, temporal downweighting, censoring, anti-correlations

1. Introduction

In resting-state functional magnetic resonance imaging (rsfMRI), the correlation between the blood oxygenation level dependent (BOLD) signals from different brain regions is used to estimate the functional connectivity of the brain in the absence of an explicit task (Fox and Raichle, 2007). A set of regions that shows a high degree of mutual correlation is referred to as a resting-state network (RSN). Since the initial observation by Biswal et al. (1995) of resting-state correlations in the motor RSN, many other RSNs have been identified, including visual, auditory, and language networks (Cordes et al., 2000, 2001; Hampson et al., 2002), the default mode network (DMN) (Raichle et al., 2001), and the task positive network (TPN) (Fox et al., 2005). The DMN and the TPN consist of those brain regions that exhibit decreases and increases, respectively, in metabolic activity during the execution of attention demanding tasks (Raichle et al., 2001). In the resting-state it has been observed that signals in the DMN are negatively correlated with signals in the TPN, giving rise to the notion of the DMN and TPN as anti-correlated networks (Fox et al., 2005).

Many rsfMRI studies employ a pre-processing step known as global signal regression (GSR) in which a global mean time course is regressed out of each voxel time course prior to the computation of the correlations. However, the use of GSR is controversial as it has been shown that the process can create artifactual negative correlations (Fox et al., 2009; Murphy et al., 2009; Weissenbacher et al., 2009; Anderson et al., 2011; Saad et al., 2012). In particular, it has been been claimed that the observed anticorrelations between the DMN and TPN are largely a mathematical artifact introduced by GSR (Murphy et al., 2009). In response to these concerns, Fox et al. (2009) acknowledged the validity of the mathematical argument put forth by (Murphy et al., 2009), but argued that the characteristics of the correlation maps obtained with GSR could not be solely determined by the mathematical constraint. For example, they showed that negative correlations between the DMN and TPN were present even without the application of GSR, an observation supported by subsequent studies (Chai et al., 2012; Chang and

Glover, 2009; Wong et al., 2012; Liang et al., 2012). To address some of the concerns regarding GSR, alternative approaches for either reducing global signal effects or determining when the application of GSR is valid have been proposed (Chai et al., 2012; He and Liu, 2012; Carbonell et al., 2014). Related studies have demonstrated that there are significant neural contributions to the global signal (Schölvinck et al., 2010; Wong et al., 2013; Wen and Liu, 2016) and have also examined the potential diagnostic value of the signal (Yang et al., 2014).

However, despite the growing understanding of the global signal and the potential limitations of GSR, a consensus regarding the use of GSR is lacking. Because the application of GSR greatly improves the spatial specificity of functional connectivity maps (Fox et al., 2009; He and Liu, 2012), it is still widely used (Power et al., 2015; Li et al., 2015) even in the face of the strong concerns that have been raised (Hahamy et al., 2014; Gotts et al., 2013). From our perspective, this lack of consensus partly reflects the difficulty in thinking clearly about the process of GSR, even though it is a relatively straightforward and compact mathematical operation. In the prior work, mathematical proofs that describe the general characteristics of GSR (e.g. the sum of correlation values must be negative) or simplified models based on a few brain regions have been used to examine the effects of GSR (Fox et al., 2009; Murphy et al., 2009; Saad et al., 2012). While these approaches have provided useful perspectives on GSR, it has been difficult to leverage them to develop an intuitive understanding of the specific effects of GSR when it is applied to a typical rsfMRI dataset that has hundreds of time points and tens of thousands of voxels.

In this paper, we introduce a simple framework for understanding the effects of GSR. In particular, we show that the main effects of GSR can be well approximated by a temporal downweighting process in which the data from different time points are attenuated as a function of their global signal magnitude (i.e., absolute value), such that the time points with the highest global signal magnitude experience the greatest amount of attenuation. We also show that a limiting case of the downweighting process in which time points with a global signal magnitude greater than a specified threshold are censored (i.e. excluded) can provide a good approximation to GSR. Furthermore, we demonstrate that the ability to model the effects of GSR as a temporal downweighting or censoring process can provide insight into the controversy regarding the existence of anti-correlated networks.

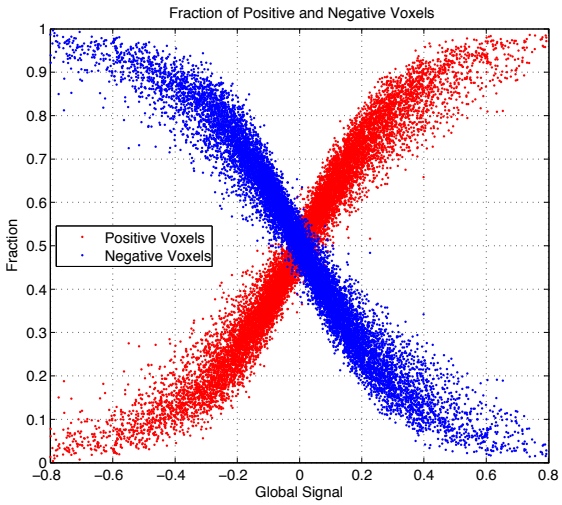

Figure 1: Fractions of positive (red) and negative (blue) voxels versus GS value for each time point in the dataset described under Methods. The GS magnitude is higher when a large fraction of the voxels have the same sign, and tends towards zero when there are roughly equal proportions of positive and negative signal values.

2. Theory

2.1. The global signal as a time-varying measure of spatial homogeneity

As a starting point, it is useful to review the basic properties of the global signal (GS). In describing these properties we focus on providing a reasonable description of the average behavior that is observed empirically in the experimental data, and acknowledge that this "average" description is not intended to cover all possible cases.

The value of the GS at each time point is simply the average of the BOLD percent signal change values across all voxels in the brain. In considering the empirically observed properties of the GS, we find that the value of the GS will be positive when the majority of the voxels have a positive BOLD signal value and will be negative when the majority of voxels have a negative signal value. At time points where there are roughly equal proportions of positive and negative signal values, these signal components will tend to cancel out and the GS magnitude will tend towards zero. Thus, the GS magnitude will be higher at time points when there is a relatively high level of spatial homogeneity (e.g. most voxels have a positive BOLD signal) and lower for time points where the data exhibits spatial heterogeneity (e.g. a roughly equal mix of positive and negative values across the brain). These properties are demonstrated

in Figure 1, where we plot the fraction of positive (red) and negative voxels (blue) versus the GS value for each timepoint in the dataset described under Methods. Large GS magnitudes (i.e large positive or negative values) correspond to zones where there is either a large fraction of either positive or negative voxels, respectively, whereas low GS magnitudes (i.e. small positive or negative values) correspond to zones where there are roughly equal fractions of positive and negative voxels. The GS magnitude can also be relatively low for time points when the BOLD values across voxels are uniformly small.

Another view of the basic properties of the GS is provided in Figure 2, where the BOLD images from a representative subject and slice are shown over 45 time points in panel (a), along with colored bars below them indicating the value of the GS at each time point and histograms (in panel (b)) of the BOLD signal values at 3 selected time points. At time point 72, the image is relatively heterogeneous with a roughly equal mix of negative and positive values. The corresponding histogram is centered about zero and the GS magnitude is close to zero at this time point. In contrast, at time point 82 the image is relatively homogeneous with negative values, leading to a large negative value for the GS. The corresponding histogram clearly shows that the overwhelming majority of voxels have a negative value at this time point. At time point 94 the image is mostly positive with a large positive GS value. Taken together, the observations in Figures 1 and 2 suggest that the magnitude of the GS can be thought of as a time-varying measure of spatial homogeneity.

2.2. Global Signal Regression primarily affects time points where the GS magnitude is high

In this section we take a closer look at the process of global signal regression (GSR). We should note that the simple examples and arguments presented in this section are meant primarily to provide background and motivation for the experimental measures (described in Methods and Results) that empirically demonstrate the temporal downweighting effects of GSR.

Our basic observation is that GSR largely affects those time points where the magnitude of the GS is high. To see why this is the case, we first consider the toy example shown in the left column of Figure 3, where we have designed a set of three signals to highlight the properties of GSR. Each of the signals is constructed as the sum of six Gaussian waveforms of varying magnitudes and time shifts. In Figure 3 (a), the three signals are designed such that the Gaussian waveforms perfectly cancel out in the intervals centered about time points 30, 70, 90, and 110 and sum in a constructive fashion around time points 10 and 50. As a result, the resulting GS (computed as the mean of the three signals) in panel (d) has a large positive peak for the interval around time point 10 and a large negative peak for the interval around time

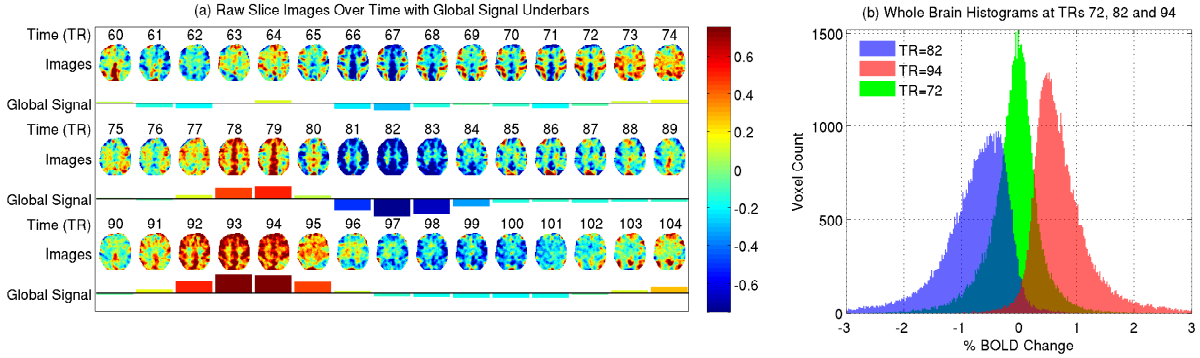

Figure 2: (a) Examples of brain images from a representative subject and slice with GS values represented by colored underbars for each image. A total of 45 consecutive time points are used to demonstrate a range of GS values. Large GS magnitudes correspond to brain images with greater spatial homogeneity, while small GS magnitudes correspond to images where there are similar proportions of negative and positive voxel values. (b) Histogram of whole brain voxel values corresponding to three selected time points with small (TR = 72) and large (TR = 82 and 94) GS magnitudes. The centers of mass of these histograms correspond to the GS values at their respective time points.

point 50 (indicated by the thick black bars), but is otherwise equal to zero. Panel (g) shows the results of applying GSR to each of the original signals from panel (a). Because the GS consists solely of Gaussian waveforms centered about times points 10 and 50, the process of regression completely eliminates the signal components in these intervals. In contrast, regression has no effect on time points where the GS is equal to zero, so that the Gaussians centered about time points 30, 70, 90, and 110 are untouched by GSR.

In panel (b), we have slightly modified the signals from panel (a) so that there is no longer perfect cancellation of the Gaussian waveforms for the intervals centered about time points 30 and 70. As a result, the GS in panel (e) has small but non-zero values in these intervals. For this case, GSR still largely attenuates the waveforms at time points 10 and 50 with minimal attenuation of the waveforms at time points 30 and 70 (as shown in panel (h)). This reflects the fact that the computation of the regression coefficient in GSR is dominated by those time points with larger GS magnitudes, so that scaled versions of the GS will have a better fit to the voxel time series at these time points. From a mathematical point of view, those time points with larger GS magnitude have greater leverage on the regression (Hoaglin and Welsch, 1978; Draper and Smith, 2014). Denoting the global signal as the column vector $\mathbf{g}$, the leverage at each time point is given by the corresponding diagonal term of the projection matrix $\mathbf{g}(\mathbf{g}^T\mathbf{g})^{-1}\mathbf{g}^T$.

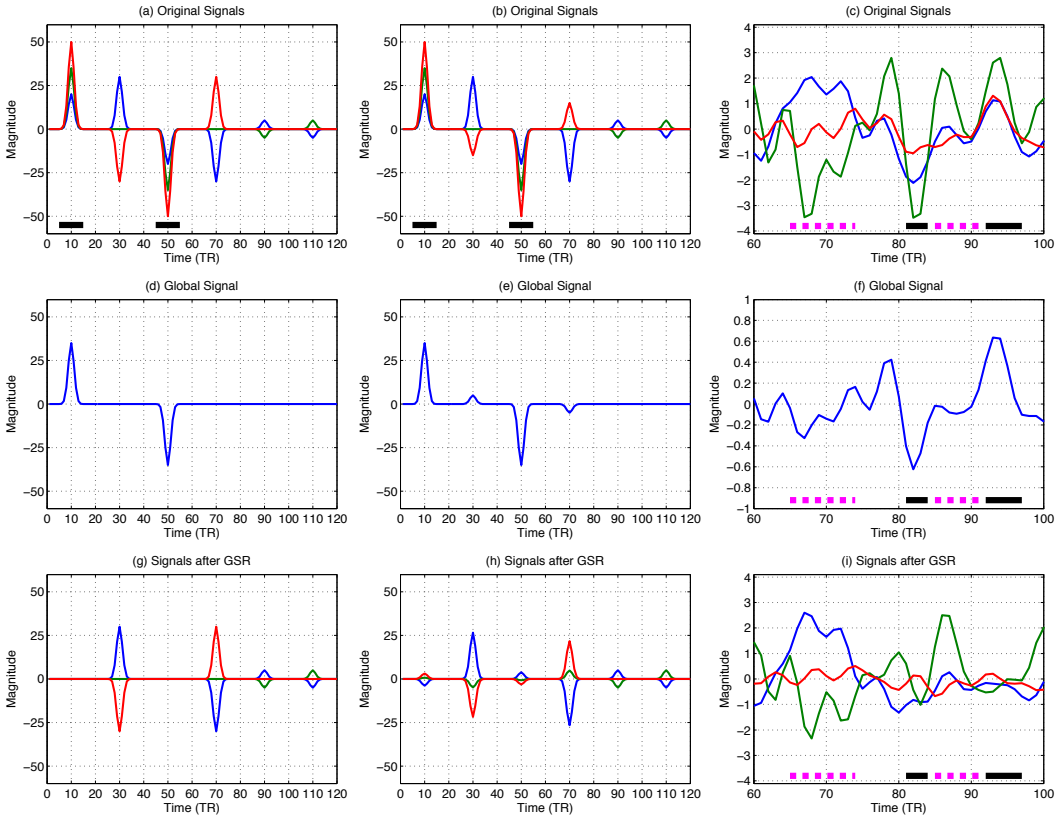

Figure 3: (a) Three voxel time series are constructed such that the Gaussian waveforms at time points 30, 70, 90 and 110 sum to zero, whereas the waveforms at time points 10 and 50 (indicated with thick black bars) sum constructively. This summation results in the large peaks at time points 10 and 50 in the GS shown in (d). Panel (g) shows that the original waveforms at time points 30, 70, 90 and 110 are preserved after GSR, whereas the waveforms at times points 10 and 50 are completely eliminated. Panels (b,e,h) demonstrate a more realistic case where the cancellation of signals at time points 30 and 70 is not complete and the GS has small values at these points. GSR greatly attenuates the waveforms at time points 10 and 50 and has a slight effect on the waveforms at time points 30 and 70. Panels (c,f,i) show three example BOLD time series where the GS in (f) is computed over the entire brain. The magenta and black lines indicates regions with low and high GS magnitudes, respectively. GSR greatly attenuates signals in the high GS regions and has a smaller effect on signals in the low GS regions.

Since $(\mathbf{g}^T\mathbf{g})^{-1}$ is simply a scalar, the leverage at the ith time point is proportional to the square g_i^2 of the corresponding global signal value. The concept of leverage is useful for understanding that time points with large magnitudes tend to be more influential in the regression, but the detailed effects of GSR require consideration (and therefore computation) of the entire projection matrix (Draper and Smith,

2014). As a result, while we use theoretical concepts such as leverage to help motivate our approach, we need empirical measures to demonstrate the actual effects.

In Figure 3 (c), we show three BOLD time series from a representative subject (see Methods for details on the data). The GS for the dataset is shown in panel (f), where in contrast to the toy example, the GS is computed over all voxel time series (of which only 3 out of about 50,000 are shown in the plot). The dashed magenta and solid black lines correspond to time intervals in which the GS magnitude is either relatively small or large, respectively. After GSR (see panel (i)), the waveforms are greatly attenuated for those intervals where the GS magnitude is large and minimally affected for the intervals where the GS magnitude is small.

In summary, the examples shown above suggest that the main effect of GSR is to reduce the magnitude of waveforms in those time intervals in which the GS has a relatively large magnitude. In the appendix we present mathematical arguments that provides greater insight into this effect. Specifically, we introduce the concept of splitting the global signal $\mathbf{g}$ into two orthogonal components $\mathbf{g} = \mathbf{g}_H + \mathbf{g}_L$ where the subscripts H and L denote the sets of time points where the global signal is high and low, respectively. A voxel time series $\mathbf{x}$ can be similarly decomposed into the sum $\mathbf{x} = \mathbf{x}_H + \mathbf{x}_L$ of two components. In addition, since GSR is a linear operation, the voxel time series after regression $\tilde{\mathbf{x}}$ can be written as the sum $\tilde{\mathbf{x}} = \tilde{\mathbf{x}}_H + \tilde{\mathbf{x}}_L$. In the appendix we show that $\tilde{\mathbf{x}} \approx \mathbf{x}_L$, such that to first order GSR can be viewed as attenuating the signal component $\mathbf{x}_H$ in temporal regions where the global signal is high while having a relatively small effect on the component $\mathbf{x}_L$ for intervals where the global signal is low. If the global signal is identical to zero for all time points in the set L, the approximation becomes an equality (i.e. $\tilde{\mathbf{x}} = \mathbf{x}_L$), consistent with the observations related to the toy example in the leftmost column of Figure 3.

As noted above, both the toy example and the mathematical perspectives are meant to provide background and insight that will be useful for understanding the extensive empirical findings that are presented in later sections. Our focus has been on presenting examples and arguments that help to explain the average behavior that is empirically observed, and it is certainly possible to come up with exceptions. For example, if the voxel time series is uncorrelated with the global signal, then GSR will have no effect on the time series and therefore there will also be no reductions in the GS magnitude. However, given the widespread correlation observed between the GS and voxel time series (Power et al., 2016), uncorrelated voxels are the exception and do not drive the average behavior.

2.3. Effect of GSR on Seed-Based Correlations

We now consider the effects of GSR on the computation of the correlation coefficients between voxel time series. Without loss of generality, we will assume that the voxel time series have been normalized to be unit norm and zero mean. Then the correlation coefficient r between a seed voxel time series $\mathbf{s}$ and another voxel time series $\mathbf{x}$ is simply given by the inner product $r = \mathbf{s}^T\mathbf{x}$, where the time series are represented as column vectors. Using the orthogonal decomposition described above and in the Appendix, this may be rewritten as $r = \mathbf{s}_H^T\mathbf{x}_H + \mathbf{s}_L^T\mathbf{x}_L$. Similarly, the correlation coefficient after GSR (denoted as $\tilde{r}$) can be written as $\tilde{r} = \beta\tilde{\mathbf{s}}^T\tilde{\mathbf{x}} = \beta\left(\tilde{\mathbf{s}}_H^T\tilde{\mathbf{x}}_H + \tilde{\mathbf{s}}_L^T\tilde{\mathbf{x}}_L\right)$, where the tilde notation indicates the values after GSR and the coefficient $\beta = \left(\|\tilde{\mathbf{s}}\|\|\tilde{\mathbf{x}}\|\right)^{-1}$ is needed to renormalize the signals. Using the approximations discussed in the prior section, the signals after GSR can be expressed as $\tilde{\mathbf{s}} \approx \mathbf{s}_L$ and $\tilde{\mathbf{x}} \approx \mathbf{x}_L$, so that the correlation coefficient can be approximated as $\tilde{r} \approx \beta\mathbf{s}_L^T\mathbf{x}_L$. In other words, the correlation coefficient after GSR is largely determined by those time points where the global signal has low magnitude. This argument can be extended to correlation maps, by defining a data matrix $\mathbf{X}$ consisting of individual voxel time series as columns which are normalized to have unit norm. The correlation map in vectorized form is then given by $\mathbf{R} = \mathbf{s}^T\mathbf{X}$, which can be approximated as $\tilde{\mathbf{R}} \approx \mathrm{diag}\left(\boldsymbol{\beta}\right)\mathbf{s}_L^T\mathbf{X}_L$ after GSR, where $\mathrm{diag}\left(\boldsymbol{\beta}\right)$ denotes the matrix with the vector $\boldsymbol{\beta}$ of renormalizing coefficients along the diagonal. This result suggests the following testable hypothesis: functional correlation maps obtained with GSR should resemble correlation maps obtained from a subset of the original dataset corresponding to time points where the global signal magnitude is low.

To further demonstrate the reasoning behind this hypothesis, we note that a correlation map can also be written as the sum of weighted images $\mathbf{R} = \sum_{t=1}^{T} s_t\mathbf{X}_t$ where s_t and $\mathbf{X}_t$ represent the unit norm seed signal value and the data image, respectively, acquired at time t. (Note $\mathbf{X}_t$ is the t'th row in $\mathbf{X}$). In Fig. 4 we show examples of raw and weighted brain images (both before and after GSR) along with the corresponding correlation maps. The weighted images are obtained by multiplying the raw images by a seed voxel time series (represented by the colored bars) from the posterior cingulate cortex (PCC; see Methods). Due to space constraints, we only show representative images where the GS magnitude is very small (time points 66 to 68) or relatively large (time points 81 to 84 and 92 to 95). However, the correlation map is obtained from the sum of all weighted images (most of which are not shown). Because GSR reduces the magnitudes of both the PCC time series and the raw images at time points where the GS is large, the contribution of the weighted images at these time points to the overall correlation map is greatly diminished. As a result, the correlation map will be dominated by those time points where the

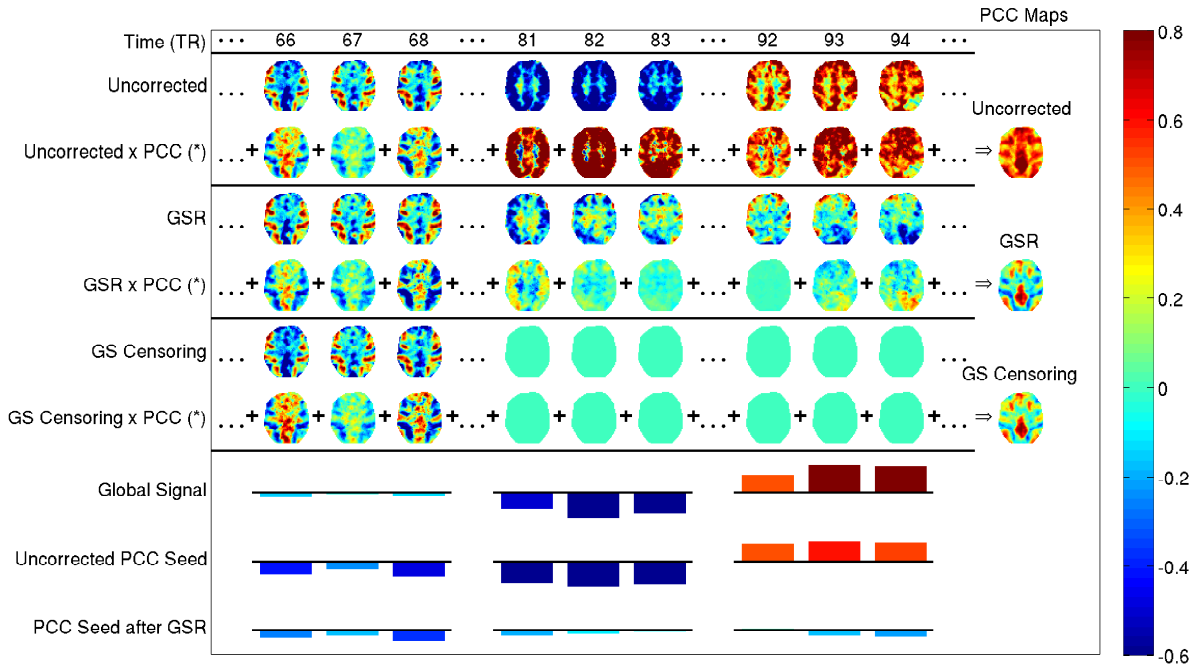

Figure 4: Construction of correlation maps for uncorrected data (rows 1 to 2), data after GSR (rows 3 to 4) and data with GS censoring (rows 5 to 6). Images are shown for representative time points where the GS magnitude is either very small (TR= 66-68) or large (TR = 81-83 and 92-94). In each group, the upper and lower rows show the images both prior to and after multiplication by the unit norm PCC seed signal (referred to as unweighted and weighted), respectively. Note that for the weighted images the voxel time series have been normalized to unit norm prior to multiplication by the unit norm PCC seed signal in order to be consistent with the mathematical description of the correlation sum provided in the text. The PCC correlation maps for each condition are obtained by summing all of the weighted images across the duration of a scan (note only 9 time points are shown here due to space limitations). For display purposes, the PCC weighted images in rows marked with (*) have been scaled to match the intensity scales of the images in the upper rows. The symbol ($\Rightarrow$) indicates that the summation of the images yields the correlation maps on the left, after taking into account the display scaling. The values of the GS and the PCC signal both prior to and after GSR are indicated by the colored bars at each time point, where the PCC-related signals in the bottom two rows have been divided by 2.25 for display purposes. Both the unweighted and weighted images are largely unaffected by GSR for time points where the GS magnitude is small, whereas the images at time points with large GS magnitude are greatly attenuated. As a result, the correlation map after GSR largely reflects the contribution of the weighted images from the time points with small GS magnitude. In the GS censoring approach, the unweighted images with large GS magnitudes are censored (e.g. multiplied by zero) while all other images are unaffected. The resulting correlation map is similar to that obtained with GSR.

GS is small in magnitude, i.e. $\tilde{\mathbf{R}} \approx \mathrm{diag}\,(\boldsymbol{\beta})\,\mathbf{s}_L^T\mathbf{X}_L$ as noted above. This initial result suggests that one might obtain maps similar to those with GSR by simply censoring those time points where the GS is large. An example correlation map obtained using this approach is shown in Figure 4.

As with the temporal downweighting arguments stated in the prior section, the examples and arguments provided in this section are meant to provide background and motivation for the experimental measures (described in Methods and Results) that empirically demonstrate the effect of GSR on correlation maps. In addition, it is important to note the fundamental difference between the temporal downweighting process, which is a time-varying modification of the signal magnitudes, and a simple scaling of the overall amplitude of the entire time series, which represents a time-invariant scaling of the signal magnitudes. To help make the distinction clear, we use term **magnitude** to refer to the absolute value of the GS at each time point and reserve the term **amplitude** to refer to the standard deviation of the GS over the course of the scan. Whereas the correlation coefficient is invariant to a simple scaling of the overall amplitude of the time series, it can be greatly altered by a temporal downweighting process. For example, censoring represents a limiting case of temporal downweighting and can have a significant effect on the correlation coefficient through the zeroing of data at the censored time points.

3. Methods

3.1. Subjects and Data Acquisition

We used data originally analyzed by Fox et al. (2007) and downloaded from www.brainscape.org (dataset BS002). The data were acquired from 17 young adults (9 females) using a 3 T Siemens Allegra MR scanner. Subjects underwent 4 BOLD-EPI fixation runs (32 slices, TR=2.16 s, TE=25 ms, 4×4×4 mm), each lasting 7 minutes (194 frames). Subjects were instructed to look at a cross-hair, and asked to remain still and awake. High-resolution T1-weighted anatomical images were also acquired for the purpose of anatomical registration (TR=2.1 s, TE=3.93 ms, flip angle=7 deg, 1×1×1.25 mm).

3.2. Data Processing

Standard pre-processing steps were conducted with the AFNI software package (Cox, 1996). The first 9 frames from each EPI run were discarded to minimize longitudinal relaxation effects, leaving 185 frames for further analysis. Images were slice-timing corrected and co-registered and the six head-motion parameter time series were retained. The resultant images were converted to coordinates of Talairach and

Tournoux (TT), resampled to 3 mm cubic voxels, and spatially smoothed using a 6 mm full-width-at-half-maximum isotropic Gaussian kernel. Nuisance terms were removed from each voxel's time course using multiple linear regression. These were: (1) a constant term to model the temporal mean of each voxel, (2) a linear trend, and (3) six motion parameters obtained from the head motion correction algorithm. A temporal low pass filter was applied to the remaining time course using a cutoff frequency of 0.1 Hz (Cordes et al., 2000; Yan et al., 2009). Each detrended and filtered voxel time series was converted into a percent change BOLD time series through division by its mean value. The GS was then calculated as the average of the percent change time series across all voxels within the brain for each scan. GSR was implemented by regressing the GS out of each voxel time series (Macey et al., 2004).

Correlation maps were formed by correlating a seed signal with the time-course from every voxel in the brain. The seed signal was computed as the average signal from a region selected with spheres of radius 9 mm (3 voxels) centered about the regions of interest. We obtained the following sphere centers by converting the MNI coordinates obtained from (Van Dijk et al., 2010) to TT coordinates using the MNI to TT conversion algorithm (Lacadie et al., 2008). For primary analysis we focused on the posterior cingulate cortex (PCC) with a seed region centered in TT coordinates [0,-51,26]. For the left and right auditory network (AUD) we used [-41,-26,14] and [41,-26,14]. For left and right motor networks (MOT) we used [-36,-25,57] and [36,-25,57], and for the medial prefrontal cortex we used [0,46,-7]. For MOT and AUD seeds we merged two spheres from the left and right hemispheres of the brain to obtain a final region of interest for each seed. For right intraparietal sulcus we used the coordinates [27,-58,49] from (Fox et al., 2006) and for the white matter seed we defined the seed center as [31,-28,32]. Correlation maps were computed for (a) the pre-processed time series (referred to as uncorrected maps) and (b) the pre-processed time series after the application of GSR or one of the weighting schemes described below.

3.3. GSR Ratio: Characterizing the average effect of GSR

In Section 2.2 we made the observation that GSR largely attenuates the signals in time intervals where the GS is large and has a minimal effect in those intervals where the GS is small. To further characterize this effect, we calculated a GSR ratio $\mathbf{g}_r(t)$ that captures the average effect of GSR at each time point t. This is defined as

$$\mathbf{g}_r(t) = \left\langle \frac{\mathbf{v}_{i,a}(t)}{\mathbf{v}_{i,b}(t)} \right\rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{\mathbf{v}_{i,a}(t)}{\mathbf{v}_{i,b}(t)} \tag{1}$$

where $\mathbf{v}_{i,b}(t)$ and $\mathbf{v}_{i,a}(t)$ are the values of the i'th voxel's time course before and after GSR, respectively, and N is the number of voxels. Because the ratio blows up for small values of $\mathbf{v}_{i,b}(t)$, outliers were removed prior to the computation of the mean. The threshold for outlier detection was defined as the median plus 2.5 times the median absolute deviation of the central 95% of the data, where initial censoring of the data in the extreme tails of the distribution was performed to avoid excessive skewness prior to the computation of the statistics (Leys et al., 2013; Reimann et al., 2005). This process resulted in a threshold value of 2.25 that was applied to the magnitude of $\mathbf{g}_r(t)$.

By definition, the GSR ratio should be less than 1 when the average attenuation (across voxels) due to GSR is large and be close to 1 when the average effect of GSR is minimal. In Figure 5 (a,b) we show the GSR ratio plotted along with the magnitude of the GS from two representative scans. The GSR ratio exhibits local minima at time points where the GS magnitude has relatively large values. In contrast, when the GS magnitude has small values the GSR ratio is close to 1, indicating minimal perturbation on average of the original voxel values.

To further evaluate the properties of the GSR ratio, we multiplied each of the pre-processed voxel time series $\mathbf{v}_{i,b}(t)$ by the GSR ratio to obtain a new time series for each voxel of the form

$$\tilde{\mathbf{v}}_{i,GSRW}(t) = \left\langle \frac{\mathbf{v}_{i,a}(t)}{\mathbf{v}_{i,b}(t)} \right\rangle \cdot \mathbf{v}_{i,b}(t) = \mathbf{g}_r(t) \cdot \mathbf{v}_{i,b}(t) \tag{2}$$

where the subscript $GSRW$ indicates that the time course is GSR ratio weighted. For time points where GSR has the same effect for all voxels, then $\left\langle \frac{\mathbf{v}_{i,a}(t)}{\mathbf{v}_{i,b}(t)} \right\rangle = \frac{\mathbf{v}_{i,a}(t)}{\mathbf{v}_{i,b}(t)}$ and multiplication by the GSR ratio acts in the same manner as GSR so that $\tilde{\mathbf{v}}_{i,GSRW}(t) = \mathbf{v}_{i,a}(t)$. On the other hand, for time points where the effect of GSR varies over voxels, multiplication by the GSR ratio only captures the average spatial effect of GSR so that the weighted values can only approximate the values obtained with GSR, such that $\tilde{\mathbf{v}}_{i,GSRW}(t) \approx \mathbf{v}_{i,a}(t)$.

The GSR ratio weighted time series $\tilde{\mathbf{v}}_{i,GSRW}(t)$ were used to compute an additional set of PCC correlation maps. In Figure 5 (c), (d), we compare these maps to PCC correlation maps obtained using the pre-processed voxel time series both before GSR (i.e. $\mathbf{v}_{i,b}(t)$) and after GSR (i.e. $\mathbf{v}_{i,a}(t)$). There is a strong similarity between the GSR and GSR Ratio weighted maps, suggesting that the main effects of GSR are reflected by its average attenuation (at each time point) of the voxel values. Differences in the maps reflect the fact that GSR performs a unique regression for each voxel time series, whereas the GSR ratio is constrained to use the same weighting function for each voxel time series. Additional examples of GSR ratio weighted maps are shown in the Results section.

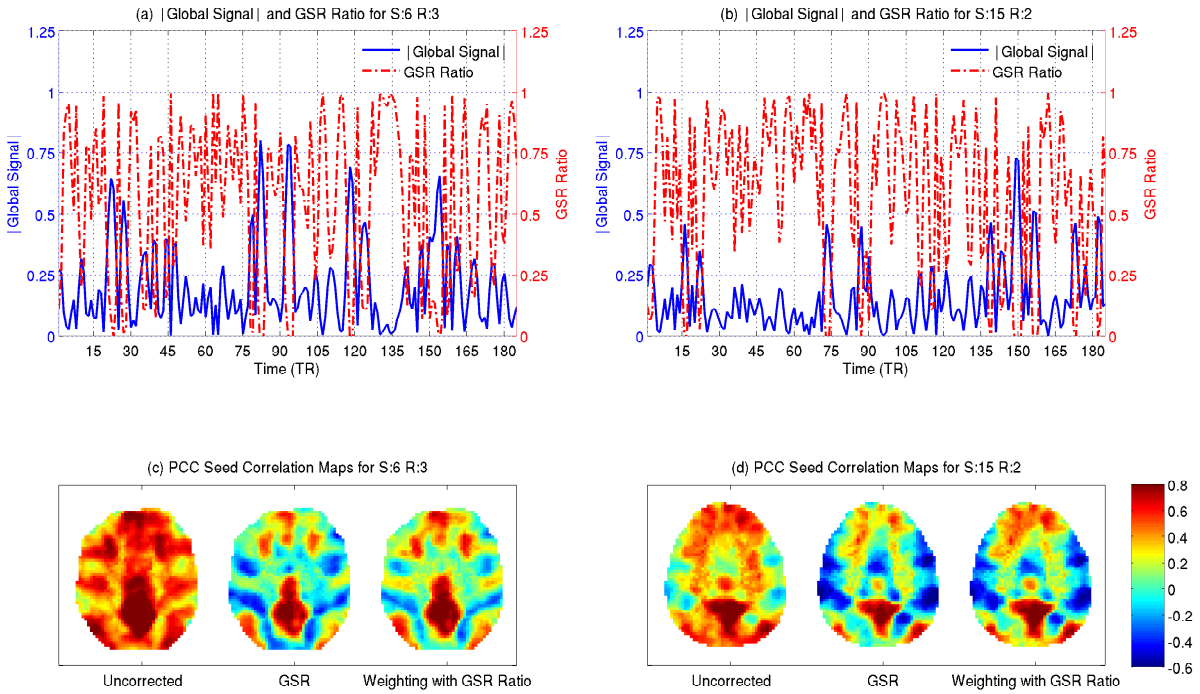

Figure 5: (a), (b) GSR Ratio time series (red dashed lines) and GS magnitudes (solid blue lines) from two representative scans. (c,d) PCC seed correlation maps obtained after: (1) preprocessing only, (2) application of GSR after preprocessing, and (3) GSR ratio weighting after pre-processing.

3.4. Global Signal Weighting: An approximation for the GSR Ratio

In the previous section we presented preliminary evidence for a relation between the GSR ratio and the magnitude of the GS in two representative scans. We further evaluated the generality of this relation by computing the GS and GSR ratio for all scans. As shown in Figure 6, we found a strong inverse dependence of the GSR ratio on the GS magnitude. To quantify this dependence, we fit the data to a piece-wise linear model of the form

$$f\left(GS(t)\right) = \begin{cases} 1 - \alpha \mid GS(t) \mid & \text{for } \mid GS(t) \mid \leq g, \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

To calculate the optimal fit we used a robust regression method based on an iteratively reweighted least-squares (IRLS) algorithm with bisquare weights (Martinez et al., 2010; Wager et al., 2005) as implemented in the MATLAB Curve Fitting Toolbox. This resulted in parameter values of $\alpha = 2.7$ and $g = 0.37$ and

a model fit (shown by the dashed black line) with $R^2 = 0.94$. We found that other models provided similar fits. For example, a rectified Gaussian model (shown by the dashed magenta line) provide a fit with $R^2 = 0.95$. The similarity in the performance of the fits reflects the fact that the main deviation between the models occurs in a region ($|GS| > 0.30\%$) where the density of data points is very low. Due to the simplicity of its form, we will use the piece-wise linear model for the remainder of this paper.

The explanatory power of the piece-wise linear model indicates that the average voxel-independent effect of GSR can be well approximated by a simple function of the GS. In other words, the average weighting effect of the GSR operation can be achieved without actually needing to perform the regression on each voxel time series. Instead, we can multiply each voxel time series by a 'GS weighting' function to obtain a new time series $\tilde{\mathbf{v}}_{i,GSW}(t) = f\left(GS(t)\right) \cdot \mathbf{v}_{i_b}(t)$ where the subscript GSW indicates that the time

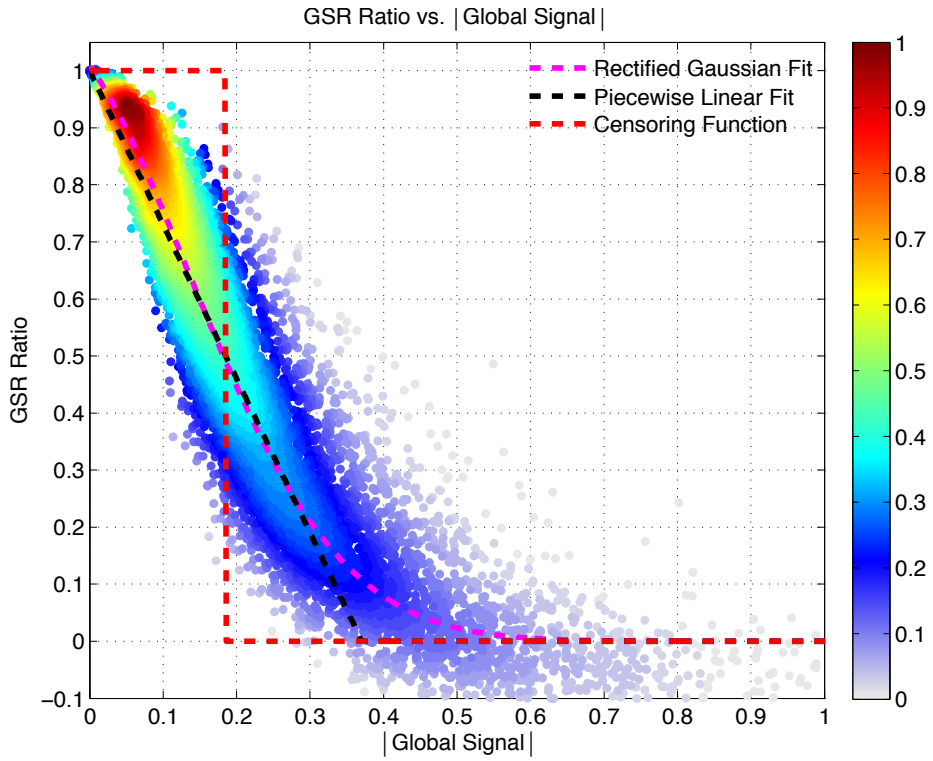

Figure 6: Density weighted scatter plot of GSR Ratios versus GS magnitude (% change) for all subjects and runs, i.e. a total of 12580 data points. Maximum data concentration is indicated by the red color and a density value of 1. Piecewise linear (black dashed line) and rectified Gaussian fits (magenta dashed line) are shown. In addition a censoring function with threshold $g_C = 0.5$ is shown (red dashed line).

series is GS weighted with the piecewise linear function. Correlation maps computed using $\tilde{\mathbf{v}}_{i,GSW}(t)$ will be referred to as GS weighted maps.

3.5. Global Signal Censoring

When using the piece-wise linear function $f\left(GS(t)\right)$ from Equation 3 as a weighting function, the voxel time series values are scaled by a linearly decreasing function that gradually decreases to zero. This can be characterized as a soft censoring operation. We also considered a hard censoring function of the form

$$\mathbf{C}\left(GS(t)\right) = \begin{cases} 0 & \text{for } (1-\alpha|GS(t)|) \leq g_C, \\ 1 & \text{otherwise.} \end{cases} \tag{4}$$

This function censors those time points for which the predicted average weighting due to GSR is below a threshold value g_C. For example, a threshold value of $g_C = 0.5$ would censor those data points for which GSR would scale the data on average by a factor of 0.5 or less. This is equivalent to censoring those time points for which the GS magnitude is greater than 0.18% (i.e. $|GS| > (1-g_C)/\alpha$ using the value of $\alpha = 2.7$ estimated above). A plot of this censoring function is shown by the red dashed line in Figure 6.

The rationale behind the hard censoring approach relies in part on the preliminary observation (see Section 2.3) that the correlation maps after GSR are expected to be largely determined by those time points that are largely unaffected by GSR. The choice of g_C quantifies the degree to which a time point is considered to be "unaffected" by GSR. For example, for the threshold value of $g_C = 0.5$, those time points for which the average scaling due to GSR is predicted to range between 0.5 and 1.0 are considered to be unaffected and will not be censored. As the threshold is increased, the number of time points censored increases as fewer time points are considered to be unaffected by GSR (see Figure 13 in Results).

As with the previously discussed weighting functions, we can multiply each preprocessed voxel time series by the hard censoring function to obtain a new time series $\tilde{\mathbf{v}}_{i,GSC}(t) = \mathbf{C}(GS(t)) \cdot \mathbf{v}_{i_b}(t)$ where the subscript GSC indicates that the signal has been GS censored. Correlations maps computed using $\tilde{\mathbf{v}}_{i,GSC}(t)$ will be referred to as GS censored maps. Figure 7 summarizes the various weighting and censoring approaches and the signal paths used to compute the correlation maps shown in the Results sections.

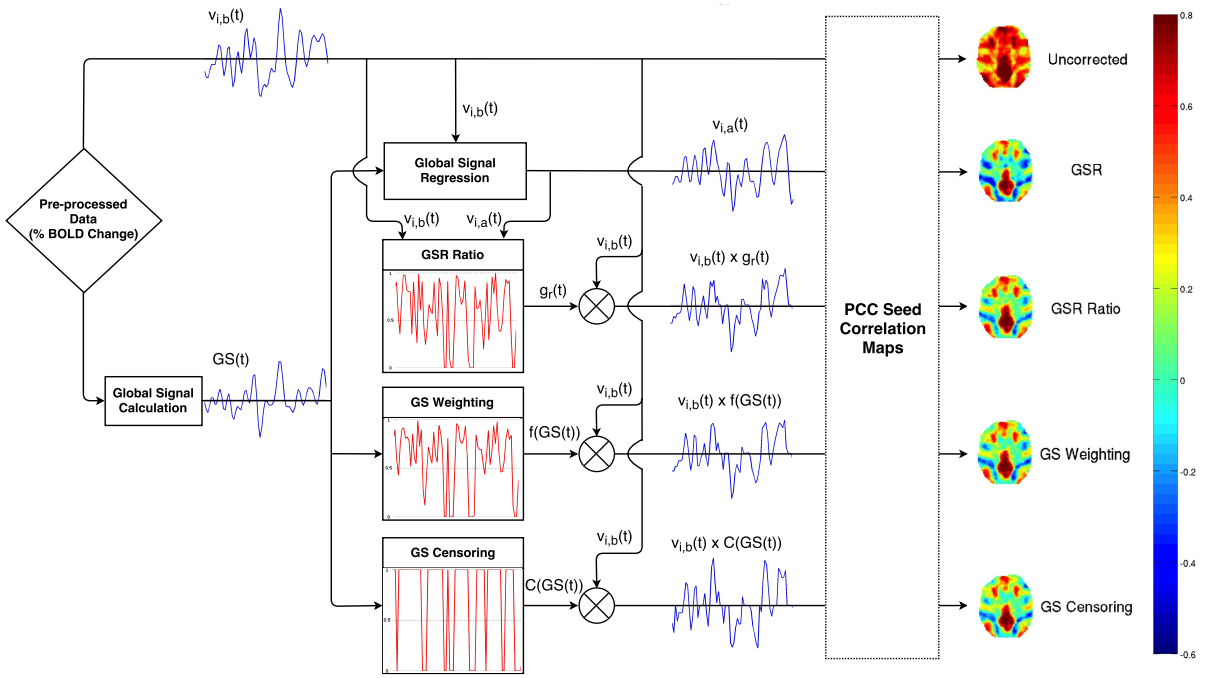

Figure 7: Diagram of the proposed GSR approximations and the data processing pipeline. The average effects of GSR can be approximated by multiplying the pre-processed voxel time courses $v_{i,b}(t)$ with GSR ratio, GS weighting or GS censoring functions (shown in red). The resulting time series are used to form PCC correlation maps shown on the righthand side. The GSR ratio approach uses all the voxel time courses prior to GSR ($v_{i,b}(t)$) and after GSR ($v_{i,a}(t)$) to determine a weighting function to approximate GSR. In contrast, the GS weighting and censoring operations use only the GS as an input to determine the weighting function.

3.6. Motion Censoring

As motion censoring approaches are sometimes used in the analysis of rsfMRI data, we wanted to compare the GS censoring approach with motion censoring. To do so, we implemented the motion censoring method described in (Power et al., 2012), which is based on the frame-wise displacement (FD) and DVARS time series. An FD threshold value was computed for each scan as the mean plus one standard deviation of the FD time series, and an overall threshold for the sample was then computed by averaging over the individual scan thresholds. This resulted in a threshold value of $T_{\mathrm{FD}} = 0.23$ mm. The DVARS threshold value was computed in the same fashion, yielding $T_{\mathrm{DVARS}} = 0.38\%$. Individual FD and DVARS masks were formed by thresholding the FD and DVARS time series with the pre-computed thresholds. The FD and DVARS masks were then dilated to include one frame prior to and one frame

immediately after each of the time frames identified by thresholding. Finally, the intersection of the FD and DVARS masks was computed to form a combined motion censoring mask for each run. Correlation maps were then computed using those time points that survived the censoring process with the combined mask.

3.7. Similarity Measures

We used cosine similarity to quantify the similarity of the correlation maps obtained with GSR and the various approaches (Karahanoğlu and Van De Ville, 2015; Walther et al., 2015). As cosine similarity preserves the means of the correlation maps, it is suitable for comparing approaches where visible differences in the associated maps reflect a strong dependence on the mean level (e.g. compare the first and second rows of Figure 9). We should note that cosine similarity would be equivalent to Pearson correlation if the mean were subtracted prior to the computation of the normalized inner product (Walther

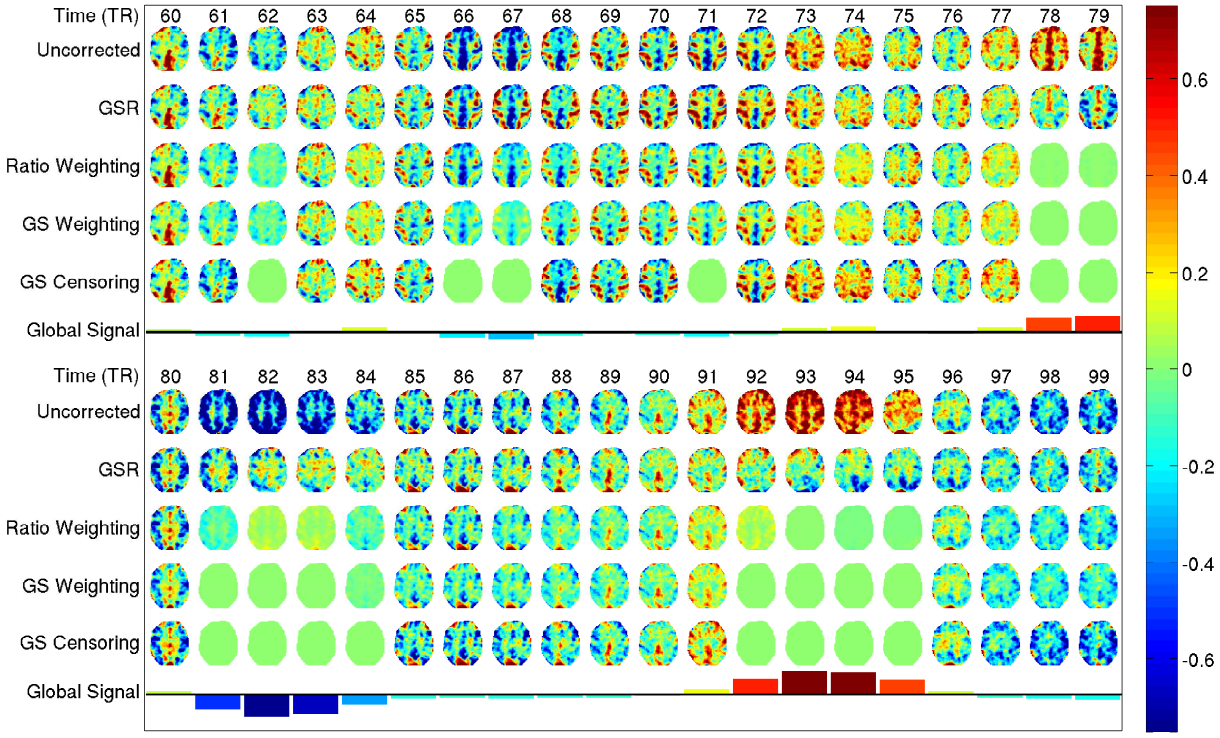

Figure 8: Images from a representative slice and scan obtained before GSR and after the application of GSR, GSR Ratio weighting, GS Weighting, and GS Censoring. The colored bars indicate the value of the GS at each time point.

et al., 2015). However, the mean subtraction procedure can introduce potential confounds related to post hoc centering of the correlation values, which biases correlation maps towards the GSR correlation map (Fox et al., 2009).

To assess the significance of the similarity between the maps obtained with GSR and each of the approaches on a per-scan basis, we constructed an empirical null distribution of the corresponding similarity measure for each scan. This was accomplished through random permutation (10,000 trials) of the temporal ordering of the calculated weights (i.e. GSR ratio, GS weight, or 0 or 1 for GS censoring). The cosine similarity was calculated between correlation maps obtained with GSR and the maps obtained after application of the permuted weights, and this process was repeated over all permutations in order to form a null distribution of similarity values. Note that for GS censoring the randomized approach is similar to the random censoring procedure used by Siegel et al. (2014). In addition, we performed a paired t-test to assess whether there was a significant difference between the similarity of the correlation maps obtained with GSR and the GS weighting approach versus those obtained with GSR and the GS censoring approach.

4. Results

The application of the various weighting approaches to data from a representative scan is demonstrated in Figure 8, where the raw images and global signal from Figure 2 are shown again for comparison. In addition, the images after the application of GSR, GSR ratio weighting, GS weighting, and GS censoring (with threshold $g_C = 0.5$) are shown. For time points where the GS magnitude is low (such as time point 75), the images are largely unaffected by GSR. In contrast, for time points where the GS magnitude is high (such as time points 82 and 93), GSR results in an average decrease in the magnitude of the images. These average effects are reflected in the uniform downweighting of the images with either GSR ratio weighting or GS weighting. Note that the attenuation introduced by GSR can vary across voxels whereas the attenuation due to GSR ratio weighting and GS weighting is by definition the same across voxels, as these approaches are designed to capture the average effect of GSR across all voxels. In the bottom row, the GS censoring approach censors (i.e. multiplies by zero) images at those time points where there is significant attenuation introduced by GSR. With the specific threshold used in this example, the censored time points occur for images where the average attenuation is large (i.e. scaling factor less than $g_C = 0.5$) in the GS weighted images.

Correlation maps obtained with the PCC seed and using the proposed approximation approaches are shown in Figure 9 for a representative slice from each of 20 different scans. Maps obtained with the uncorrected data (i.e. after preprocessing but before GSR) are shown in the top row. From a qualitative viewpoint, the correlation maps obtained after GSR are very similar to those obtained after GSR ratio, GS weighting and GS censoring. This suggests that the main effects of GSR on the correlation maps are well approximated by using voxel-independent weighting and censoring approaches. Maps using other seed voxels are provided in Supplementary Figures 1 to 4.

To quantitatively assess the similarity of the PCC correlation maps obtained with GSR and the various approaches, we tested the similarity values against the empirical null distributions obtained with randomized weighting or censoring, as described in the Methods section. This was done for each scan.

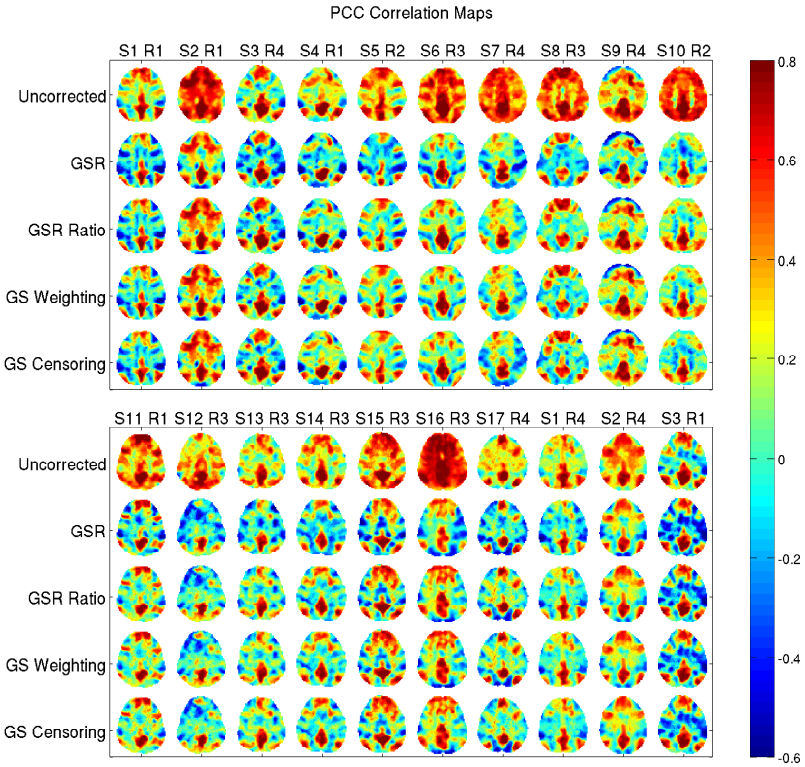

Figure 9: PCC seed correlation maps obtained before GSR, after GSR, and after application of GSR ratio weighting, GS weighting, and GS censoring. Maps are shown for 20 different scans, with at least one scan from each subject. For GS censoring a threshold of $g_C = 0.50$ was used.

For the similarities between the maps obtained with GSR and either the GSR ratio or GS weighting approaches, the empirical p-values were in the range $p < 1 \times 10^{-4}$ (i.e. less than the smallest empirical p-value that could be assessed with 10,000 random trials). For the similarities between maps obtained with GSR and GS censoring, the empirical p-values were in the range $p < 3 \times 10^{-4}$. These results indicate that the observed similarities were significantly different from those that could be obtained with randomized weighting or censoring of the data. To provide a complementary view, we show in Figure 10 the actual similarity values between the various approaches versus the maximum similarity values obtained from the respective null distributions. These maximum values from the null distributions represent highly unlikely values ($p = 1 \times 10^{-4}$). The absolute differences in the means of these maximum null distribution values and the actual similarity values ranged from 0.2191 to 0.2326, and the effect sizes (Cohen's d) ranged from $d = 1.46$ to $d = 1.50$. These results indicate that the effects of randomized weighting and censoring greatly differ from the proposed approaches and the similarity of the correlation maps obtained with these methods to GSR maps cannot be explained by randomized weighting and censoring of the data.

In Figure 11 we compare the similarity between the GSR and GS weighted maps with the similarity between the GSR and GS censored maps. The similarity values were not significantly different ($p = 0.92$), with an absolute difference in similarity means of 0.0003 and effect size $d = 0.0040$. This result suggests that the hard censoring approach (with weights of either 0 or 1) can capture the important spatial features of the GSR maps about as well as the soft censoring approach that has linearly varying weights.

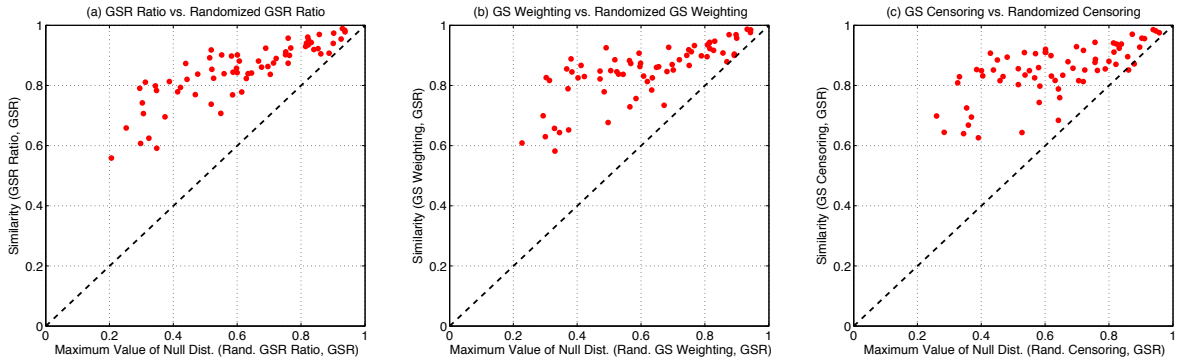

Figure 10: The similarity of the maps obtained with the proposed approaches to the GSR versus the maximum similarity in the respective null distribution for each scan (derived from randomized weighting or censoring). In (a) the absolute difference in the means was 0.2326 and effect size was $d = 1.46$, in (b) the difference in the means was 0.2318 and effect size was $d = 1.50$ and in (c) the difference in the means was 0.2191 and the effect size was $d = 1.49$.

The advantage of the hard censoring approach is that the maps obtained can be simply interpreted as estimates of the functional connectivity for the subset of time points that survive the censoring operation.

As motion censoring (or scrubbing) is finding increasing use in the analysis of rsfMRI data, we compared the effects of motion censoring and GS censoring. First we looked at the percent overlap of

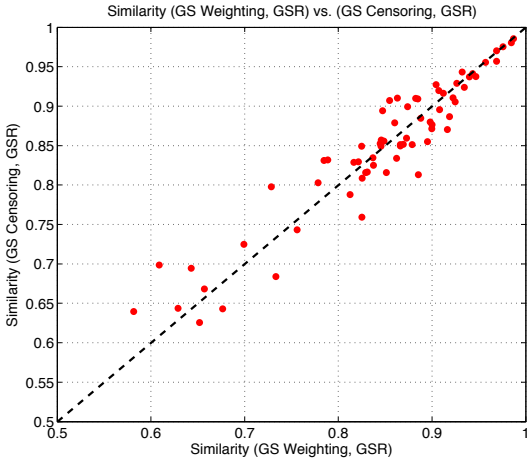

Figure 11: Comparison of similarity values between GS censoring and GSR maps versus similarities between GS weighting and GSR maps. There was not a significant difference ($p = 0.92$) between the values.

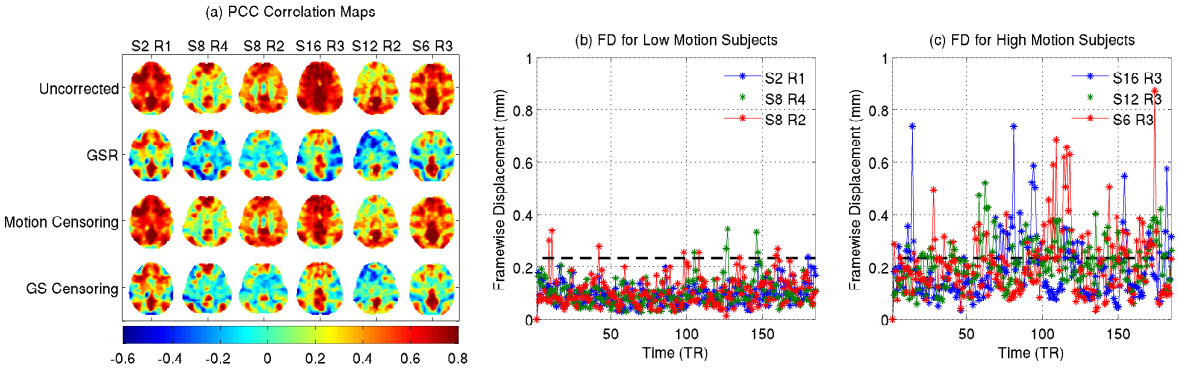

Figure 12: (a) PCC correlation maps for three low head motion runs (S2 R1, S8 R4, S8 R2) and three high motion runs (S16 R3, S12 R2, S6 R3). The maps shown were obtained before GSR and after the application of GSR, motion censoring, and GS censoring. FD time series are shown for (b) low motion and (c) high motion subjects, with the censoring threshold shown by the dashed black lines. The differences between the motion censoring and GS censoring maps indicate that the two approaches are fundamentally different.

time points that were censored by both approaches. This was calculated by dividing the number of time points that were censored by both approaches (i.e. the intersection) by the total number of time points that were censored by either approach (i.e. the union) and then multiplying the ratio by 100. We found that there was relatively low overlap between time points censored by the two approaches, with a median percent overlap of 5.67%. We then investigated temporal censoring based on individual FD and DVARS masks instead of their intersection. Application of the FD only and DVARS only masks resulted in median percent overlaps of 12.73% and 24.83%, respectively. The FD only censoring had a much smaller overlap as compared to DVARS, suggesting that the results obtained with the intersection mask were largely determined by the FD mask. This most likely reflects a closer link between the GS and DVARS time courses (i.e. root-mean square of the temporal derivative of the brain images), as compared to the link between the GS and FD time courses.

Next we compared PCC correlation maps obtained using the two different censoring approaches. Figure 12 shows PCC correlation maps for scans from 6 representative subjects with three low head-

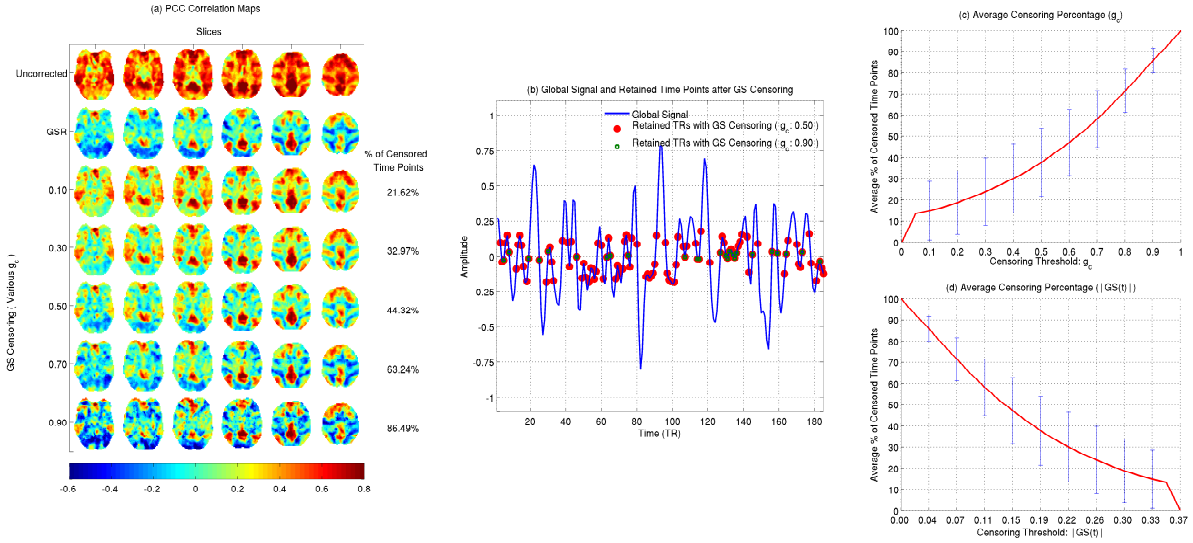

Figure 13: (a) PCC maps from a representative subject obtained before GSR and after the application of GSR and GS censoring approaches using different thresholds. (b) GS of the same subject with time points retained after censoring with $g_C = 0.50$ (red dots) and $g_C = 0.90$ (green circles). (c) Average percentage of censored time points versus censoring threshold g_C, with standard deviation across all scans indicated by the bars. (d) Average percentage of censored time points versus percent change GS magnitude threshold, where the threshold is computed using $|GS| = (1 - g_C)/\alpha$ with $\alpha = 2.7$.

motion scans (S2 R1, S8 R4, S8 R2) and three moderate to high head-motion runs (S16 R3, S12 R2, S6 R3) that were selected according to the RMS values of their respective FD time series. FD time series for low-motion and high-motion runs are shown in panels (b) and (c), respectively, with the censoring threshold for FD indicated with dashed black lines. Low motion runs incur little or no motion censoring and the resultant motion-censored scans are similar or identical to the uncorrected maps. On the other hand, high motion runs incur a large degree of motion censoring, resulting in some clean-up of the correlation maps. For both low motion and high motion runs, the maps obtained with GS censoring are similar to those obtained with GSR but are markedly different from those obtained with motion censoring.

To provide additional insight into the GS censoring approach, we show in Figure 13(a) a representative subject's PCC correlation maps obtained prior to GSR and after the application of either GSR or GS Censoring using different thresholds ranging from $g_C = 0.10$ to $g_C = 0.90$. As we increase the the threshold value, the number of time points censored increases, as indicated by the percentage values listed next to each row. For a threshold value of $g_C = 0.10$, only time points with a GSR ratio below 0.10 and a relatively high GS magnitude ($|GS| > 0.33\%$) are censored, corresponding to about 22% of

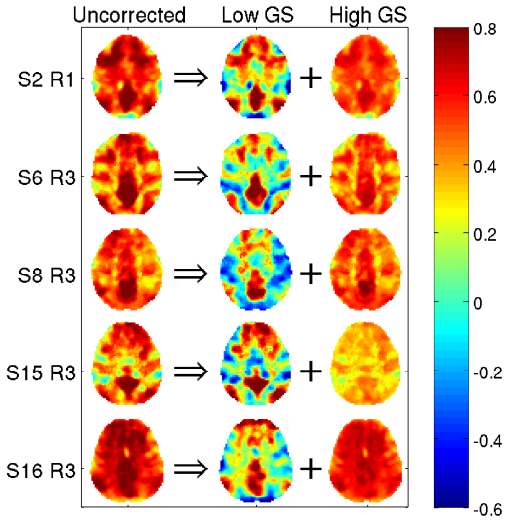

Figure 14: PCC correlation maps expressed as the sum of a map corresponding to retained points with low GS and a map based on the censored points with high GS. For display purposes, the low GS map were scaled to better delineate the features in the map. The symbol (⇒) indicates that the correlation maps is equal to the sum of the low GS and high GS maps, after taking into account the display scaling of the low GS maps.

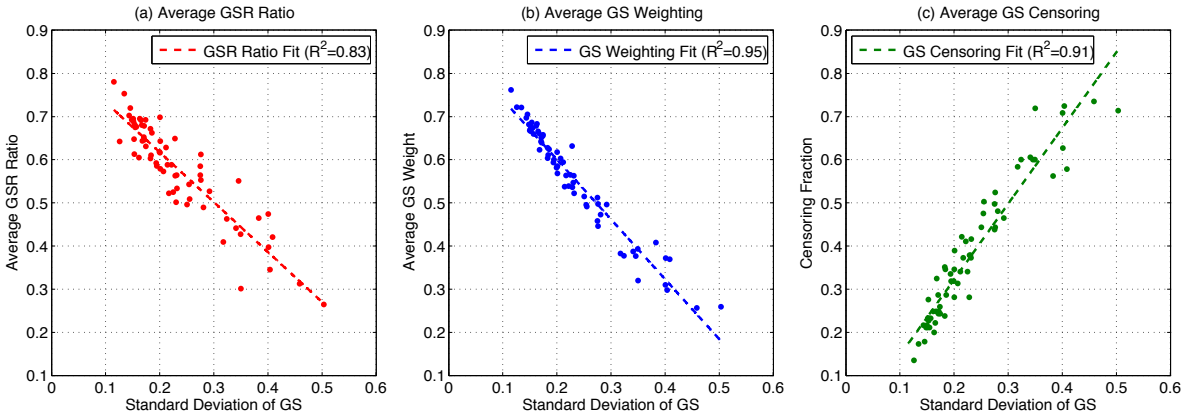

Figure 15: (a) The average GSR ratio for each scan (computed as the mean of GS ratios across all time points within a scan) versus the standard deviation of the GS across the scan. (b) The average of temporal weights obtained with the GS weighting approach for each scan versus the corresponding GS standard deviation. (c) The average fraction of censored time points versus the GS standard deviation for each scan.

the total number of points. At this level of censoring, the DMN and TPN are more clearly delineated as compared to the uncorrected maps, but the value of the correlations in the TPN are more positive than those observed in the after GSR map. As we increase the value of g_C and censor an increasing number of time points, the correlation values in the TPN become more negative and the maps become more similar to those obtained with GSR. In Figure 13(b) we indicate the location of the censored time points for threshold levels of $g_C = 0.50$ (red dots) and $g_C = 0.90$ (green circles). At a threshold value of $g_C = 0.90$ only about 14% of the time points are retained and yet the PCC correlation maps are similar to those obtained with GSR. In Figure 13(c), we show that the average percentage (across all scans) of censored time points increases with the threshold g_C, consistent with the qualitative behavior discussed above. Because of the approximate relation between GSR ratio and GS magnitude (see Figure 6 and Equation 3), we can associate each threshold value g_C with a GS magnitude threshold $|GS| = (1 - g_C)/\alpha$ where the value $\alpha = 2.7$ is obtained from the fit to the data. Only those time points with GS magnitude higher than the GS magnitude threshold are censored. As the GS threshold is increased, the average percentage of censored points decreases, as shown in Figure 13(d).

The GS censoring approach implements a temporal partitioning of the data into a subset of retained points (with relatively low GS magnitude) and a subset of censored points (with relatively high GS magnitude), where the exact delineation between low and high magnitudes is determined by the threshold.

As noted previously in Section 2.3, a correlation map can be decomposed into a sum of weighted images. Using this observation, we can decompose correlation maps into the sum of a map corresponding to the subset of retained points (these are the GS censored maps already shown) and a map corresponding to the subset of censored points. Figure 14 shows examples of this decomposition applied to PCC correlation maps with a censoring threshold of $g_C = 0.5$. As expected, the high GS subset maps exhibit a high degree of spatial homogeneity while the low GS subset maps reveal the spatial structure of the anti-correlated networks.

By averaging the GSR ratio across the duration of a scan we can obtain a value that represents the average temporal downweighting due to GSR. In Figure 15(a) we show that the average GSR ratio exhibits an inverse relation with the standard deviation of the GS. Thus, scans with greater GS fluctuations undergo more downweighting (on average) by GSR. In Figure 15(b) we show that a similar inverse dependence on GS standard deviation holds for the average GS weighting across the scan (blue dots). The average fraction of time points censored by GS censoring (green dots) increases with GS standard deviation (shown in Figure 15(c)), consistent with the associated increase in downweighting shown in panels (a) and (b) (i.e. more censoring for lower GS ratio and GS weighting values). The implications of these relations are discussed further in the following sections.

5. Discussion

5.1. A new framework for understanding GSR

We have shown that the average effects of GSR can be well approximated by a temporal downweighting of the voxel time series, where the weighting factor varies with time but is uniform across space. The weighting factor decreases with the magnitude of the global signal, so that time points with large global signal magnitudes are greatly attenuated whereas those time points with small global signal magnitudes are largely unaffected. We also introduced GS censoring as a limiting case of the downweighting approach, in which the weighting factor is equal to zero for time points where the GS magnitude exceeds a specified threshold and is equal to one otherwise.

In the prior studies and debates concerning the use of GSR, a combination of mathematical arguments, empirical findings, and simulations have been used to examine the strengths and limitations of the approach (Fox et al., 2009; Murphy et al., 2009; Saad et al., 2012; He and Liu, 2012). However, the process of GSR has remained somewhat mysterious as it is not easy to visualize the process of regression

with a global mean signal in a high-dimensional signal space. The use of differing perspectives in prior studies has complicated the development of a unified view of GSR and made it difficult to resolve the ongoing debate. By showing that GSR can be well approximated by a temporal downweighting operation, we have introduced a simple way of viewing GSR that facilitates a more intuitive understanding of its effects. Instead of having to visualize how regression affects each individual voxel time series, we can simply consider consider how the data from each time point is downweighted prior to the computation of functional correlations. In the limit of GS censoring, the data from each time point is either included in (weighting of 1.0) or excluded (weighting of 0.0) from the computation. The similarity in the maps obtained with GSR and GS censoring suggests that the first order effects of GSR are well approximated by censoring a large fraction (i.e. greater than 50% or more) of the data.

5.2. Anti-correlated Networks

With respect to the highly debated anti-correlated networks, our work suggests that the negative correlations between the DMN and the TPN are not simply an artifact of GSR. Specifically, the application of the GS censoring approach shows that negative correlations can be observed when computing the functional correlation for a temporal subset of the data. Within this subset, there is no mathematical constraint that forces the existence of negative correlations, because the images retained in the subset are not modified by the censoring operation. In other words, the negative correlations are inherent in the data when considering time points with low GS magnitudes.

Our results suggest that the presence of anti-correlated networks is often obscured by the contributions of images at those time points where the GS magnitude is high (see for example Figure 4). Because these peaks in GS magnitude occur when there is a high degree of spatial homogeneity in the images (e.g. either largely positive or negative), the inclusion of these images gives rise to spatial homogeneity in the correlation maps. By downweighting the contribution of these images, GSR and the various associated approaches presented in this paper leads to a reduction in the spatial homogeneity of the correlation maps.

Prior work has shown that the application of GSR (Fox et al., 2009; Murphy et al., 2009) forces the sum of correlation values to be negative, thus potentially introducing artifactual negative correlations. The present work does not contradict these prior findings, but rather presents a new way of looking at the effects of GSR. In essence, the process of GSR can be viewed as approximately partitioning the data into two sets – a set of images in which the GS magnitude is relatively high and a set in which the

magnitude is relatively low. GSR greatly attenuates the images in the high magnitude set while having a smaller effect on the low images in the low magnitude set. When viewed from this perspective, GSR helps to reveal the inherent negative correlations that exists in the low magnitude set by reducing the contributions of the images in the high magnitude set, which have a high degree of spatial homogeneity.

5.3. Differences in Global Signal Magnitude

What differentiates images with high GS magnitude from those with low GS magnitude? While the origins of the GS are not completely understood, there is growing evidence that a significant component of the GS is due to neural fluctuations (Schölvinck et al., 2010; Wong et al., 2013; Wen and Liu, 2016). In particular, Wong et al. (2013) found that the average amplitude of the global signal (computed as its standard deviation over the course of a scan) exhibited an inverse relation to EEG measures of vigilance, with higher amplitudes corresponding to lower states of vigilance. In a recent preliminary study, Falahpour et al. (2016) demonstrated a negative correlation between the GS and EEG vigilance time courses over the course of a scan. Thus, at time points where the GS signal exhibited positive peaks, the EEG vigilance was found to be low, whereas negative peaks in the GS corresponded to high EEG vigilance values. Chang et al. (2016) found evidence for a similar negative correlation between the GS time course and LFP measures of arousal in non-human primates. Furthermore, a recent study identified characteristic events in global electro-cortical activity that were related to drops in arousal (Liu et al., 2015). These temporally distinct events may contribute to the appearance of high magnitude events observed in the GS. The link between global activity and dynamic changes in arousal finds further potential support in the observations of Pisauro et al. (2016), who used pupilometry to measure arousal states in mice. They found that increases in arousal were associated with decreases in an optical measure of global hemodynamic activity.

While the neurobiological mechanisms linking vigilance to global fluctuations are still not fully understood, it is thought that the global nature of the fluctuations may reflect widespread projections of various arousal systems onto the cortex (Jones, 2005; Picchioni et al., 2013). Although further work is needed, the existing evidence suggests that images with high GS magnitudes (i.e. uniformly positive or negative) are associated with temporal peaks and valleys in the state of vigilance or arousal. Because GSR and its variants downweights these images, it is likely that these approaches are minimizing the contribution of vigilance fluctuations to the resulting correlation maps. In addition, the results in Figure 15 demonstrate that scans with higher GS standard deviation undergo a higher degree of downweighting

or censoring, suggesting that GSR and its variants may have a greater effect on scans with low mean vigilance levels. Further studies are needed to better understand the origins of vigilance-related signal components in fMRI and to assess the effects of GSR (and its variants) on these components.

As shown in Figures 2, 4, and 8, images from time points with low GS magnitudes can contain spatial patterns resembling the DMN and TPN. For example at time point 68 in Figure 4, the GS is relatively small and the image values in the DMN and TPN are largely negative and positive, respectively. The opposite signs in these two networks (that are both fairly spread out across space) results in a cancelation of voxel values when computing the global average, consistent with the low GS magnitude. When multiplied by the value of the PCC signal (which is negative because it is in the DMN), the resulting weighted images have positive values in the DMN and negative values in the TPN. The inclusion of these weighted images in the correlation sum will tend to enhance the presence of anti-correlations in the resulting correlation map. Thus, the presence of anti-correlated DMN and TPN networks after GS censoring largely reflects the fact that the retained images already exhibit this spatial relationship.

5.4. Related Approaches

Our findings with the GS censoring approach complement prior studies, which demonstrated that key features of resting-state functional connectivity maps could be obtained using a fraction (e.g. 5 to 15%) of the original time points (Tagliazucchi et al., 2012; Liu and Duyn, 2013). In particular, Liu and Duyn (2013) found that an average image constructed from a set of timepoints showing the highest (top 15%) PCC signal values exhibited a pattern that almost perfectly matched the group average PCC correlation map, which exhibited anti-correlation between the DMN and TPN. They further decomposed the selected images into a sum of co-activation patterns (CAPs). It is likely that the CAP time points identified by (Liu and Duyn, 2013) partially overlap with the time points that would be retained with a GS censoring approach, since it is these time points that exhibit the DMN and TPN spatial patterns in the individual images. The partial overlap reflects the fact that the CAP time points occur for positive peaks in the PCC signal, whereas the GS censoring approach will pick out time points that can have either positive or negative peaks in the PCC signal. As support for this conjecture, we note that Liu and Duyn (2013) found that GSR did not affect the spatial patterns of the CAPs. This is consistent with our finding that GSR has a minimal effect on those time points for which the GS magnitude is low. In addition, in Supplementary Figure 6, we show that the average of images from those time points that are retained by GS censoring and have a positive PCC signal shows a strong similarity with the PCC

correlation maps obtained after GSR. Further work to elucidate the connection between the GS censoring and CAP approaches would be of interest.

At first glance the downweighting approaches presented here may be considered to be similar to the frame-to-frame intensity stabilization that has been used in prior studies (also referred to as global signal normalization) (Fox et al., 2009). In the stabilization approach, all computations, including calculation of the global signal, are performed prior to the removal of the temporal means of the voxel time courses. At each time point, the image data are divided by the global signal and then a constant term (typically equal to 1.0) is subtracted from the normalized images prior to the computation of inter-voxel correlations. The apparent similarity between the stabilization and downweighting approaches lies in the fact that they both apply a uniform scaling to all voxels at a given time point. A key difference is that the global signal used in the stabilization approach has a large positive mean while the global signal signal in GSR and its variants has zero mean. Taking into account this difference, it can be shown that the stabilization process is equivalent to simply subtracting out the zero mean global signal from each demeaned voxel time series (see Appendix B). Thus, intensity stabilization is more accurately viewed as a global mean subtraction approach (also known as global signal subtraction) as opposed to a downweighting or scaling approach. In comparing intensity stabilization to GSR, the key difference is that GSR finds the optimal fit (through regression) between the voxel time series and the GS prior to removal of an appropriately scaled version of the GS, whereas intensity normalization simply subtracts the GS from each voxel time series, without any voxel-specific scaling. Supplementary Figure 5 provides further examples of the fundamental difference between stabilization, GSR, and downweighting.

In this paper we have focused on the effects of GSR on seed-based correlation analysis. Independent components analysis (ICA) is another popular approach for the analysis of resting-state fMRI data, and some ICA implementations use global signal subtraction as a pre-processing step (Remes et al., 2011). As discussed above and in Appendix B, when the voxel time series is approximately equal to the GS (i.e. voxel-specific scaling is approximately 1.0), then GSR and global signal subtraction would be expected to have similar downweighting effects. However, the validity of this approximation is likely to vary across voxels and scans. In addition, some ICA implementations do not use global signal subtraction. Remes et al. (2011) have reported differences in ICA results obtained with and without global signal subtraction, with slightly better performance observed with subtraction. Additional work would be useful for better understanding the effect of the global signal on ICA-based analyses, including an examination of the effects of GS downweighting and censoring.

In the Results section, we showed that the GS censoring approach differs from the motion censoring methods that have been adopted by some rsfMRI studies, with a relatively small overlap between the censored points. As it is the average of all brain signals, the GS captures the contributions of many potential signal sources, such as motion-related artifacts and the neuronal contributions previously discussed. In computing the GS, we used a "clean" version of the GS in which nuisance terms, such as motion covariates, were regressed out as part of the preprocessing pipeline (He and Liu, 2012). This step minimizes but does not entirely eliminate the contributions of motion to the GS. We also performed our analyses with a version of the GS in which motion covariates were not removed, and still found clear differences in the effects of GS censoring and motion censoring (see Supplementary Figure 7). Nevertheless, further work is needed to better understand the relation between GS and motion censoring.

In prior work, our group presented evidence suggesting that the GS could be viewed as an additive confound (He and Liu, 2012; Wong et al., 2012). The current work provides further support for this view, as it presents a useful temporal partitioning of the GS into sections with low and high signal magnitudes. From this perspective, the additive effect of the GS is primarily confined to the time points with high GS magnitudes, and it is these time points that are largely attenuated by GSR and the other approaches proposed here. It is also important to note that the orthogonality between the temporal segments with low and high signal magnitudes is fundamentally different from the approximate orthogonality between the GS and other principal components discussed previously in (He and Liu, 2012; Carbonell et al., 2011)

5.5. Future Steps

While our results provide insight into the effects of GSR, the question of how best to use GSR or one of the variants proposed in this paper will require further discussion and investigation. Much of the prior concern regarding the use of GSR arose out of the potential introduction of negative correlations. The current work partly mitigates this concern because it provides a means (GS censoring) of defining a temporal subset of the original data in which the negative correlations already exist. Researchers can use this approach to examine group differences in negative correlations within these temporal subsets. For example, it would be instructive to use the censoring approach to compare anticorrelations in studies where global signal differences have been found (e.g. schizophrenic versus healthy controls (Yang et al., 2014)).

In addition, by comparing the correlation maps obtained before and after GSR with those obtained after GS censoring, researchers can estimate the extent to which the functional connectivity maps are

dominated by the GS signal in the censored time points. In comparing maps across groups or conditions, the proportion of points that are censored could prove to be a useful metric. Given the prior work relating average GS amplitude to vigilance (Wong et al., 2013) and the results shown in Figure 15, it is likely that scans in subjects with a lower level of vigilance and higher GS amplitudes will experience a greater degree of censoring.

In this work, we introduced the concept of GS censoring using a threshold based upon a piecewise linear function of the GS, as this approach was useful for making a link between GSR and GS censoring. In general, there is no need to use a function of the GS and instead one can censor time points using a threshold on the GS magnitude (see for example Figure 5). However, the choice of threshold may depend on the specifics of the fMRI acquisition (such as echo time and magnetic field strength), as these parameters can affect the range of GS magnitudes. Further work is needed to best determine how to best generalize the censoring approach to a wide range of studies. In addition, there is a direct loss in the temporal degrees of freedom when using GS censoring, and the implications of this loss on between-subject and between-group comparisons will need to be considered. Furthermore, our observation that GSR can be approximated as a temporal downweighting process suggests that there is an effective loss of temporal degrees of freedom when using GSR. Additional work is needed to quantify this loss and assess its impact on subsequent analyses.

As shown in Figure 13, the choice of the threshold has a direct effect on the extent and magnitude of negative correlations, with the anti-correlations becoming more pronounced as a greater number of time points are censored. However, this observation does not tell us how best to select a threshold level. Is a threshold that reveals more anti-correlations preferred to one that yields less pronounced anti-correlations? To answer this question we will need a deeper understanding of the mechanisms that give rise to large GS magnitudes, including whether or not these mechanisms can be viewed as introducing an additive signal confound that obscures the presence of anti-correlations in the censored data points. If this is the case, it may be possible to devise algorithms that use the data in the retained time points to estimate the underlying data in the censored time points. Future studies are needed to address these issues.

Appendix A. Approximating the Effects of GSR

In this appendix we provide a mathematical perspective of the basic observations presented in Section 2.2. Our goal is to provide background and motivation for the empirical findings (see Methods and Results) demonstrating that GSR largely attenuates signals for time intervals where the global signal magnitude is high. To begin, we express the global signal as the sum of two non-overlapping parts in time:

$$\mathbf{g} = \mathbf{1}_H \mathbf{g} + \mathbf{1}_L \mathbf{g} = \mathbf{g}_H + \mathbf{g}_L, \tag{A.1}$$

where $\mathbf{1}_H$ is the indicator function for the set H of points surrounding large values in the global signal and $\mathbf{1}_L$ is the indicator function for the set L of remaining points corresponding to relatively small values of the global signal.

In the example shown in Figure 3, the set H corresponds to the time points surrounding the peaks of the global signal and is indicated by the thick black horizontal bars. Similarly, we can also partition an arbitrary voxel time series in the same way $\mathbf{x} = \mathbf{x}_H + \mathbf{x}_L$. Note that by construction, $H \cap L = 0$ and the time series from different sets are orthogonal, such that $\mathbf{g}_H^T \mathbf{g}_L = 0$, $\mathbf{x}_H^T \mathbf{x}_L = 0$, $\mathbf{g}_H^T \mathbf{x}_L = 0$, and $\mathbf{g}_L^T \mathbf{x}_H = 0$. Note that without loss of generality, we assume that both $\mathbf{g}$ and $\mathbf{x}$ are zero-mean percent normalized time series (i.e. both represent percent BOLD signal changes).

Using the decomposition above, we rewrite the voxel time series after GSR as:

$$\begin{aligned}
\tilde{\mathbf{x}} &= \mathbf{x} - \mathbf{g}\left(\mathbf{g}^T\mathbf{g}\right)^{-1}\mathbf{g}^T\mathbf{x} && \text{(A.2)}\\
&= \mathbf{x}_H + \mathbf{x}_L - \frac{\left(\mathbf{g}_H + \mathbf{g}_L\right)\left(\mathbf{g}_H + \mathbf{g}_L\right)^T\left(\mathbf{x}_H + \mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)} && \text{(A.3)}\\
&= \left(\mathbf{x}_H - \frac{\mathbf{g}_H\left(\mathbf{g}_H + \mathbf{g}_L\right)^T\left(\mathbf{x}_H + \mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)}\right) + \left(\mathbf{x}_L - \frac{\mathbf{g}_L\left(\mathbf{g}_H + \mathbf{g}_L\right)^T\left(\mathbf{x}_H + \mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)}\right) && \text{(A.4)}\\
&= \underbrace{\left(\mathbf{x}_H - \frac{\mathbf{g}_H\left(\mathbf{g}_H^T\mathbf{x}_H + \mathbf{g}_L^T\mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)}\right)}_{\tilde{\mathbf{x}}_H} + \underbrace{\left(\mathbf{x}_L - \frac{\mathbf{g}_L\left(\mathbf{g}_H^T\mathbf{x}_H + \mathbf{g}_L^T\mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)}\right)}_{\tilde{\mathbf{x}}_L}, && \text{(A.5)}
\end{aligned}$$

where we have used the orthogonality relations described above. The first part $\tilde{\mathbf{x}}_H$ in the expression corresponds to the part of $\tilde{\mathbf{x}}$ with points in the set H while the second term $\tilde{\mathbf{x}}_L$ corresponds to points in the complementary set L.

As noted in Section 2.1, large peaks in the global signal occur when the majority of voxel waveforms sum in a constructive fashion. As an approximation, we represent the voxel waveforms surrounding these peaks as a scaled version of the global signal $\mathbf{x}_H \approx \alpha \mathbf{g}_H$. Although we expect that for most voxels $\alpha > 0$ because of the widespread positive correlation that exists empirically between the global signal and the voxel time series (Power et al., 2016), it is important to note that the derivations below do not depend on the sign of α. Using this approximation, we rewrite the first term $\tilde{\mathbf{x}}_H$ in Eq. (A.5) as

$$\tilde{\mathbf{x}}_H \approx \alpha \mathbf{g}_H - \frac{\mathbf{g}_H \left(\alpha \mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{x}_L\right)}{\left(\mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{g}_L\right)} \tag{A.6}$$

$$= \alpha \mathbf{g}_H \left(1 - \frac{\mathbf{g}_H^T \mathbf{g}_H + \left(\mathbf{g}_L^T \mathbf{x}_L\right)/\alpha}{\mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{g}_L}\right) \tag{A.7}$$

$$\approx \alpha \mathbf{g}_H \left(1 - \frac{\mathbf{g}_H^T \mathbf{g}_H}{\mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{g}_L}\right) \tag{A.8}$$

where the last approximation reflects our expectation that $\mathbf{g}_H^T \mathbf{g}_H \gg \mathbf{g}_L^T \mathbf{x}_L$, due to the small magnitude of the signal components in $\mathbf{g}_L$ and the fact that the cancellation of signals across voxels will tend to make $\mathbf{g}_L$ orthogonal to the voxel time courses for time points within the set L.

In the event that $\mathbf{g}_L$ is equal to zero, it follows directly from Eq. (A.7) that $\tilde{\mathbf{x}}_H = 0$. In addition, this will lead to $\tilde{\mathbf{x}}_L = \mathbf{x}_L$ in Eq. (A.5) , so that the remaining signal is simply $\tilde{\mathbf{x}} = \mathbf{x}_L$. In other words, GSR will eliminate the signal at time points where the global signal is large but will not affect the original voxel values $\mathbf{x}_L$ for those time points where the global signal is zero. This is the case for the example shown on the lefthand column of Figure 3, where the peaks in the original voxel time series that coincide with the large peaks in the global signal are completely eliminated by GSR, whereas those signals that correspond to segments where the global signal is zero (due to complete cancelation of the signals across voxels) are unaffected by GSR.

In the more realistic case where the cancellation of signals is not complete, the magnitudes of the signal components in $\mathbf{g}_L$ are still expected to be smaller than the magnitudes in $\mathbf{g}_H$, but $\mathbf{g}_L$ will differ from zero for some points in L. An example of this case is shown in the middle column of Fig. 3, where the cancellation of the signal components centered about time points 30 and 70 is not complete, leading to the appearance of small magnitude peaks in the global signal around these time points. The attenuation term $\left(1 - \frac{\mathbf{g}_H^T \mathbf{g}_H}{\mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{g}_L}\right)$ in Eq. (A.8) can be rewritten as the ratio $\frac{\mathbf{g}_L^T \mathbf{g}_L}{\mathbf{g}_H^T \mathbf{g}_H + \mathbf{g}_L^T \mathbf{g}_L}$ of the energy (i.e. squared norm) of the smaller magnitudes signals in $\mathbf{g}_L$ to the overall energy of $\mathbf{g}$. Based on the partitioning of the signal, this ratio is expected to be much smaller than 1. As a result, the signal components in $\mathbf{x}_H$

will still be greatly reduced by GSR as shown in panels (h) and (i) of Figure 3.

We can approximate $\tilde{\mathbf{x}}_L$ in Eq. A.5 as

$$\tilde{\mathbf{x}}_L \approx \left(\mathbf{x}_L - \frac{\mathbf{g}_L\left(\alpha\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{x}_L\right)}{\left(\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L\right)}\right) \tag{A.9}$$

$$\approx \mathbf{x}_L - \mathbf{g}_L\frac{\alpha\mathbf{g}_H^T\mathbf{g}_H}{\mathbf{g}_H^T\mathbf{g}_H + \mathbf{g}_L^T\mathbf{g}_L} \tag{A.10}$$

$$\approx \mathbf{x}_L - \alpha\mathbf{g}_L \tag{A.11}$$

where we have used the approximations already noted above ($\mathbf{x}_H \approx \alpha\mathbf{g}_H$, $\mathbf{g}_H^T\mathbf{g}_H \gg \mathbf{g}_L^T\mathbf{g}_L$, and $\mathbf{g}_H^T\mathbf{g}_H \gg \mathbf{g}_L^T\mathbf{x}_L$). The perturbation $\alpha\mathbf{g}_L$ in the final expression is proportional to the global signal. Since this occurs in time intervals where the global signal magnitude is assumed to be small (by definition), the magnitude of the perturbation is also expected to be small. This is consistent with the minimal perturbation of the signals at time points 30 and 70 in panel (h) of Figure 3 and in the intervals indicated by the magenta dashed lines in panel (i) of the same figure. As discussed above, when $\mathbf{g}_L = 0$, the waveforms are completely unaffected and $\tilde{\mathbf{x}}_L = \mathbf{x}_L$ (i.e. time points 30, 70, 90, and 110 in panel (g) and time points 90 and 110 in panel (h).

Appendix B. Intensity Stabilization Approach

In this approach, computations are performed prior to the removal of the voxel-wise temporal means or percent normalization. Thus, we can define the unnormalized global signal $\mathbf{g}_U$ as the sum $\mathbf{g}_U = \bar{g}\mathbf{1}_M + \tilde{\mathbf{g}}$ of a mean term $\bar{g}$ and a zero-mean fluctuation term $\tilde{\mathbf{g}}$, where $\mathbf{1}_M$ indicates a $M \times 1$ column vector of ones. An unnormalized voxel time series is similarly defined as $\mathbf{x}_U = \bar{x}\mathbf{1}_M + \tilde{\mathbf{x}}$. The process of intensity stabilization (or GS normalization) is formally defined as

$$\left(\mathrm{diag}\left(\mathbf{g}_U\right)\right)^{-1}\mathbf{x}_U - \mathbf{1}_M \tag{B.1}$$

where $\mathrm{diag}\left(\mathbf{g}_U\right)$ denotes the matrix with $\mathbf{g}_U$ along the diagonal (Fox et al., 2009).

To proceed, let $\tilde{g}[i]$ and $\tilde{x}[i]$ denote the values of $\tilde{\mathbf{g}}$ and $\tilde{\mathbf{x}}$ at the ith time point. Then the normalized

values are

$$\begin{aligned}
\frac{\bar{x}+\tilde{x}[i]}{\bar{g}+\tilde{g}[i]}-1 &= \frac{\tilde{x}[i]-\tilde{g}[i]+\bar{x}-\bar{g}}{\bar{g}+\tilde{g}[i]} && \text{(B.2)}\\
&= \frac{\tilde{x}[i]-\tilde{g}[i]}{\bar{g}+\tilde{g[i]}} && \text{(B.3)}\\
&\approx \frac{\tilde{x}[i]-\tilde{g}[i]}{\bar{g}}\left(1-\frac{\tilde{g}[i]}{\bar{g}}\right) && \text{(B.4)}\\
&\approx \frac{\tilde{x}[i]-\tilde{g}[i]}{\bar{g}} && \text{(B.5)}\\
&= \frac{\tilde{x}[i]}{\bar{x}}-\frac{\tilde{g}[i]}{\bar{g}} && \text{(B.6)}\\
& && \text{(B.7)}
\end{aligned}$$

where we have made use of the fact that the magnitude of the fluctuations in fMRI GS time series are typically only a few percent at most of the overall mean, so that $\bar{g} \gg \tilde{g}[i]$. In addition we assumed that $\bar{x} = \bar{g}$, where this last relation holds because we can always scale the data such that the means of all the voxels are the same. Thus, to first order the process of GS normalization is equivalent to simply taking the difference between the percent change voxel time series $\tilde{x}[i]/\bar{x}$ and the percent change global signal $\tilde{g}[i]/\bar{g}$. Using vector notation, the approximation to GS normalization is expressed as the difference $\mathbf{x} - \mathbf{g}$ of the percent normalized voxel time series and global signals. Thus, GS normalization and GS subtraction are nearly identical methods.

The GS normalization (and subtraction) approach and GSR are expected to give similar results when

$$\mathbf{x} - \mathbf{g} \approx \mathbf{x} - \mathbf{g}\left(\mathbf{g}^T\mathbf{g}\right)^{-1}\mathbf{g}^T\mathbf{x} \tag{B.8}$$

This will occur whenever the fit coefficient $\alpha = \left(\mathbf{g}^T\mathbf{g}\right)^{-1}\mathbf{g}^T\mathbf{x}$ between the voxel time series and the GS is approximately equal to 1.0.

Appendix C. Supplementary Figures

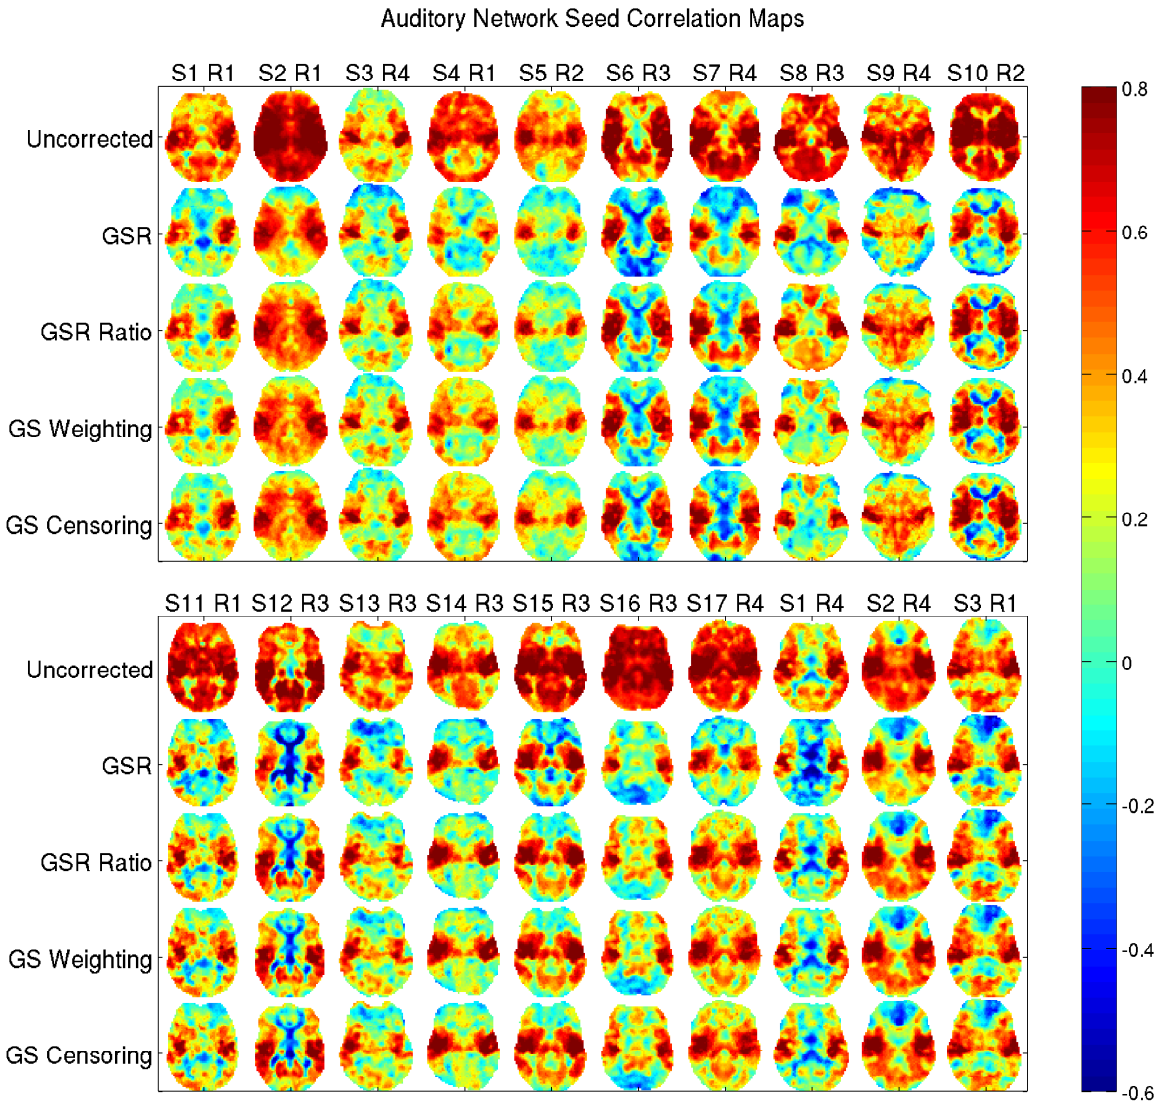

Supplementary Figure 1: Auditory network seed correlation maps obtained before GSR, after GSR, and after application of GSR ratio weighting, GS weighting, and GS censoring. For GS censoring a threshold of $g_C = 0.50$ was used.

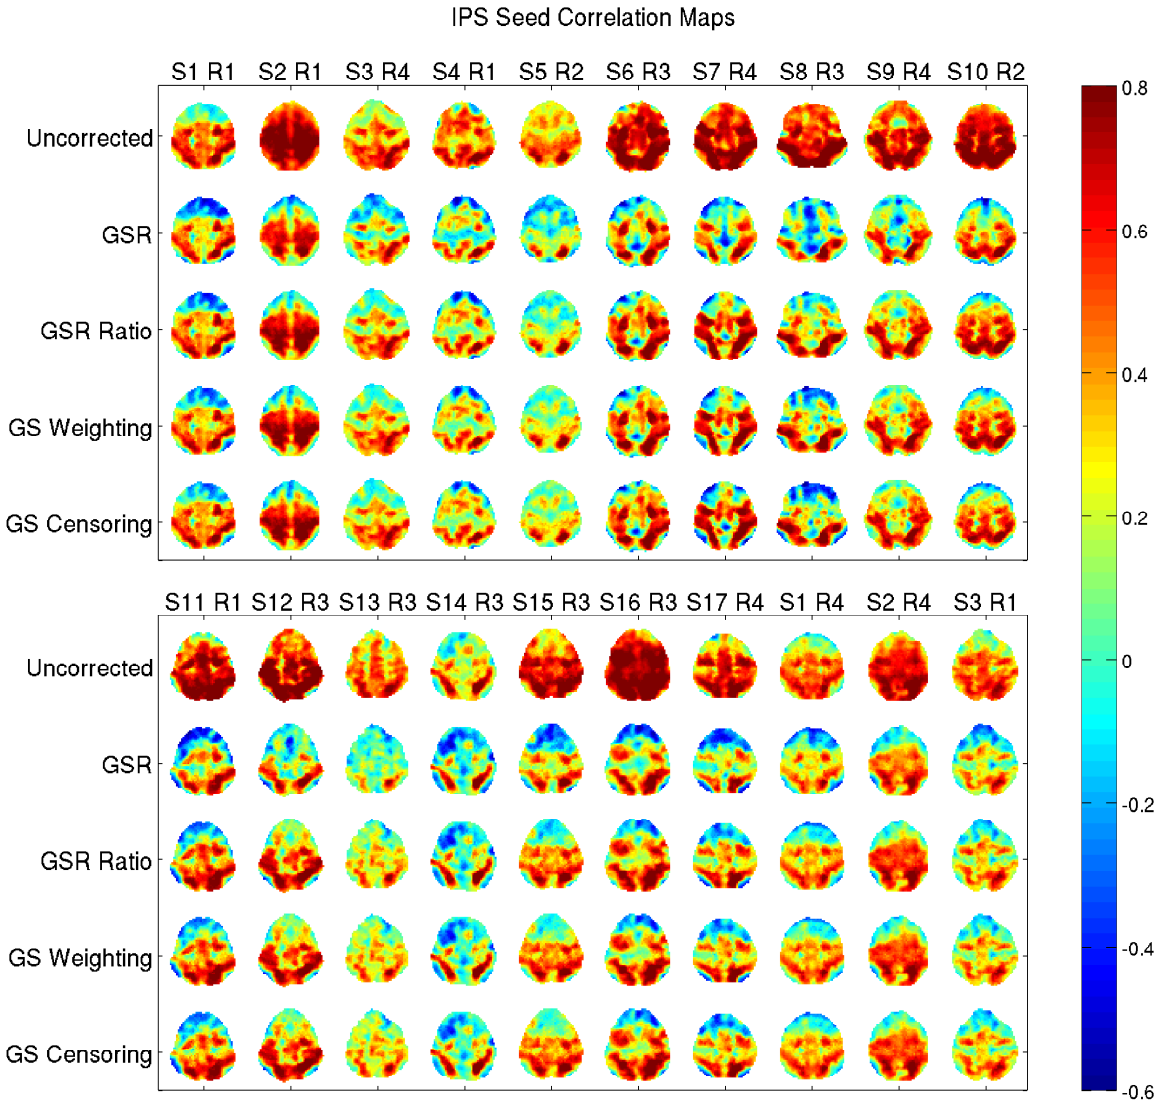

Supplementary Figure 2: IPS seed correlation maps obtained before GSR, after GSR, and after application of GSR ratio weighting, GS weighting, and GS censoring. For GS censoring a threshold of $g_C = 0.50$ was used.

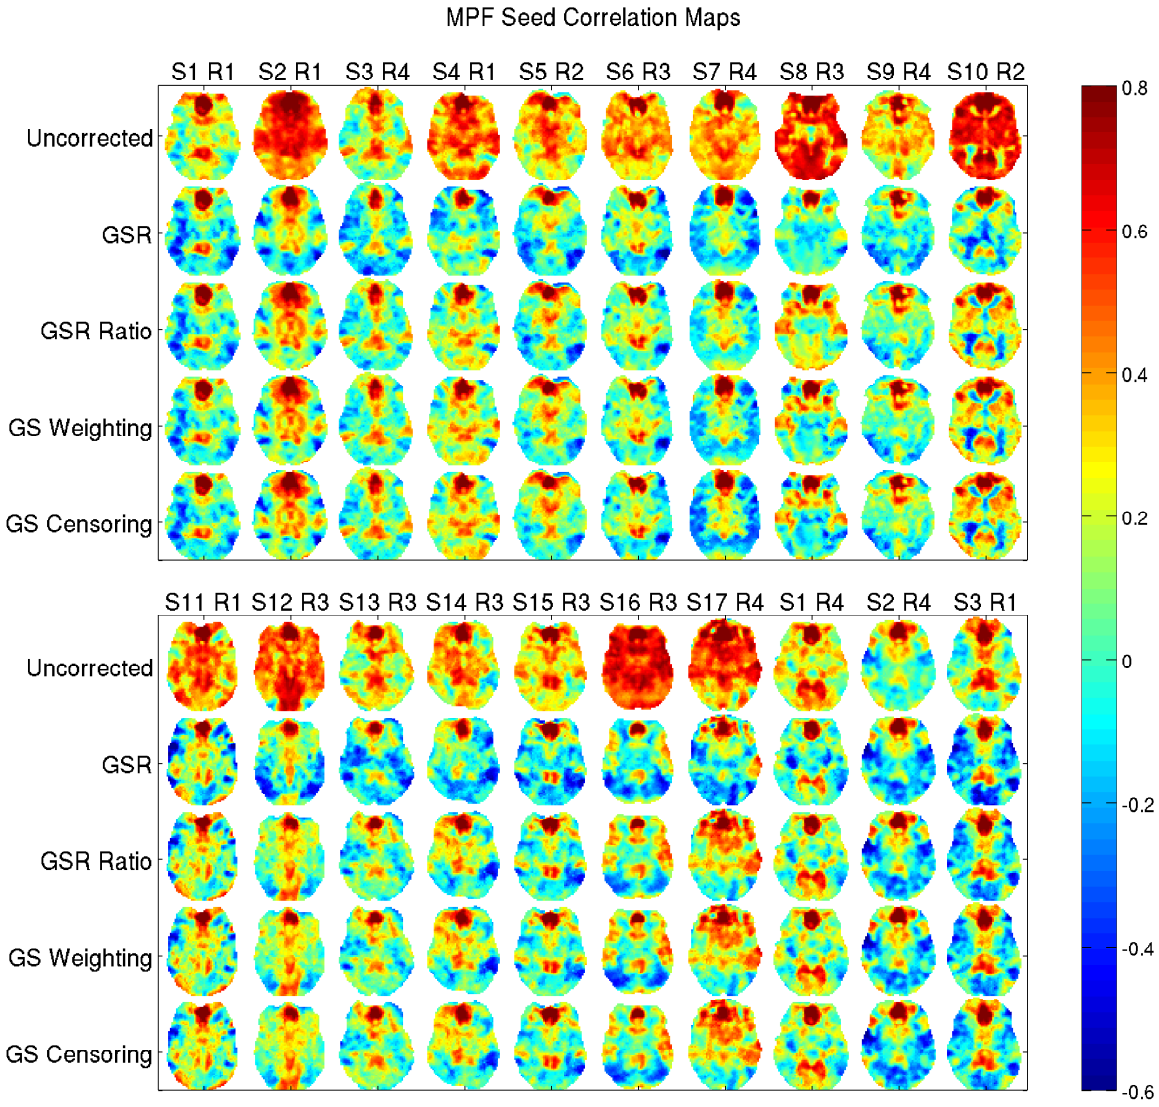

Supplementary Figure 3: MPF seed correlation maps obtained before GSR, after GSR, and after application of GSR ratio weighting, GS weighting, and GS censoring. For GS censoring a threshold of $g_C = 0.50$ was used.

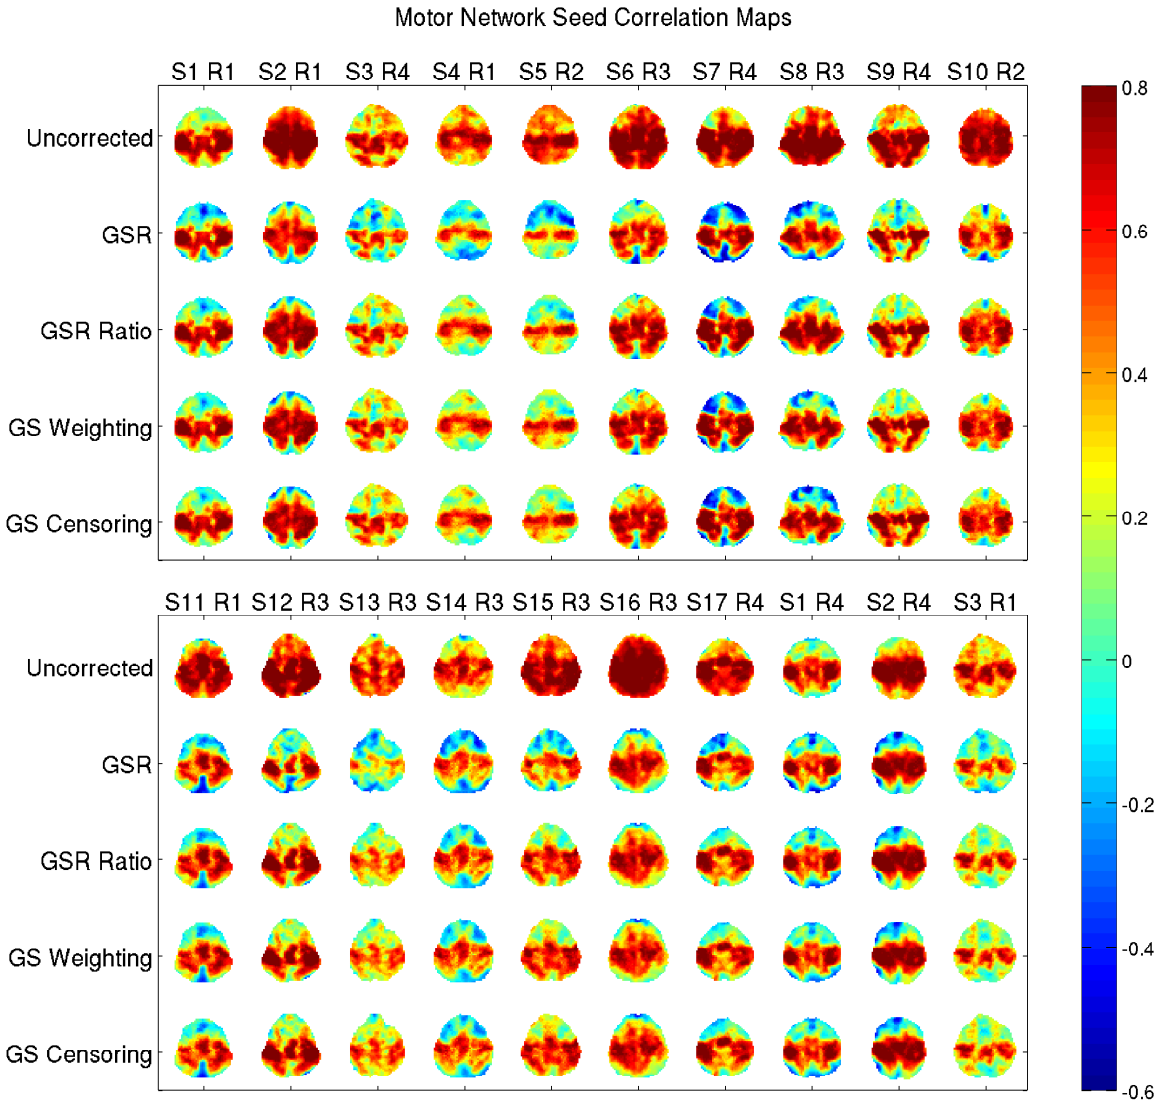

Supplementary Figure 4: Motor network seed correlation maps obtained before GSR, after GSR, and after application of GSR ratio weighting, GS weighting, and GS censoring. For GS censoring a threshold of $g_C = 0.50$ was used.

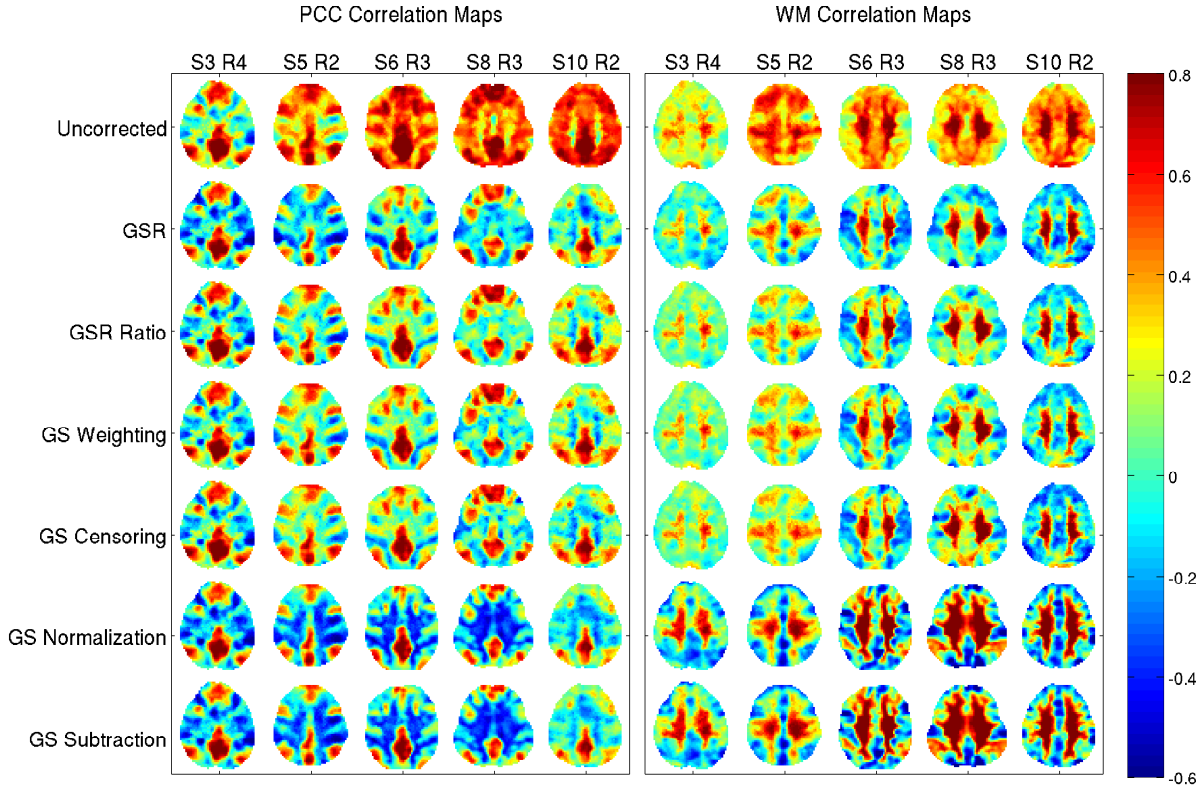

Supplementary Figure 5: Posterior-cingulate cortex (PCC) and white-matter (WM) seed correlation maps obtained prior to GSR and after the application of GSR, GSR ratio weighting, GS weighting, GS censoring, GS normalization (intensity stabilization), and GS subtraction. Consistent with the approximation shown in Appendix B, GS normalization and GS subtraction yield nearly identical maps for all scans. For some scans, the GS normalization and subtraction maps show significant differences with the maps obtained after GSR and the other GS weighting approaches, indicating clear differences in the approaches. These differences are greater when using a WM seed, where the fit between the WM seed voxel time series and the GS is expected to be lower, such that subtraction of the GS is not a good approximation for subtraction of an optimally scaled version of the GS. When the GS normalization and GSR approaches provide similar PCC correlation maps, there is a reasonably good fit between the PCC and GS time series, such that subtraction of the GS is a good approximation to subtraction of a scaled version of the GS (see also Appendix B)

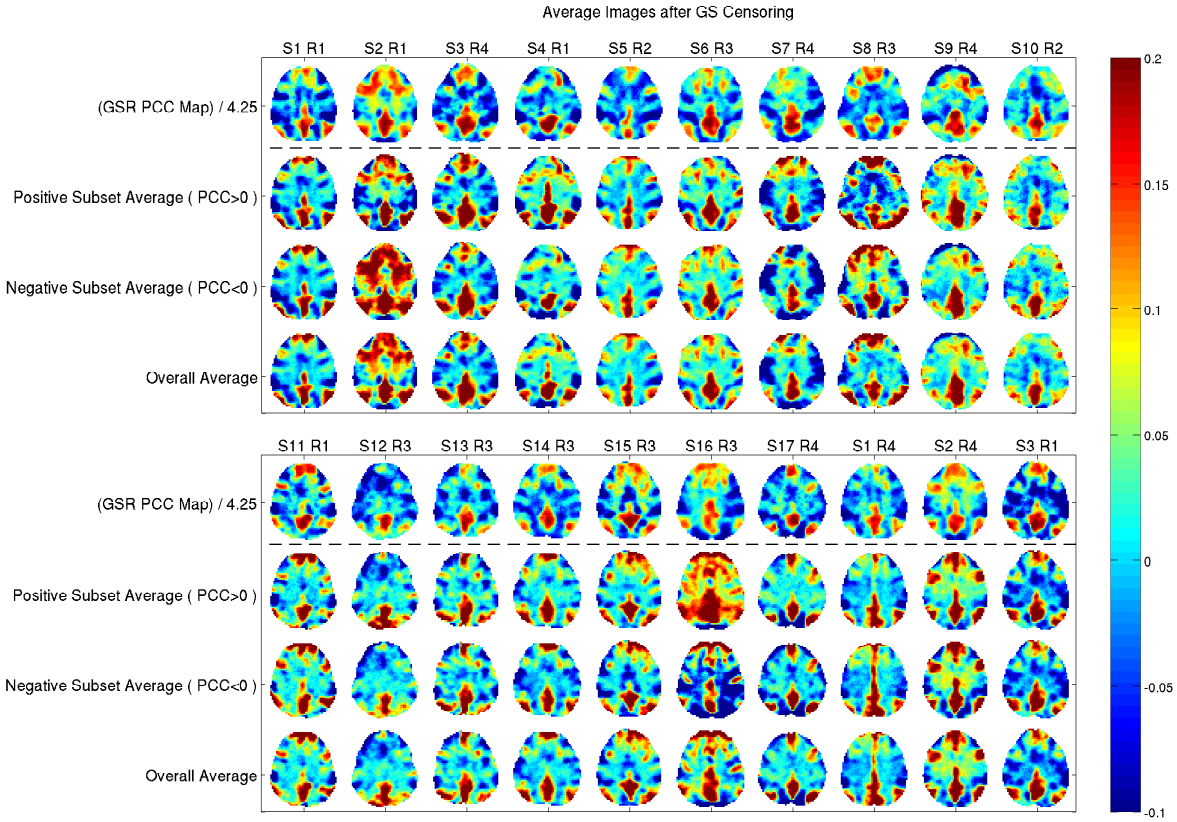

Supplementary Figure 6: The spatial patterns in the PCC correlation maps after GSR are similar to the patterns found in the average of the images retained by the GS censoring approach (with $g_C = 0.5$). To account for the occurrence of positive and negative values of the PCC seed signal, the set of retained images was divided into subsets with positive and negative PCC signal values and averages were computed for each subset. The average for the negative subset was multiplied by -1.0 to account for the sign change. The overall averages of the two subsets (after multiplying the negative subset by -1.0) are also shown. The PCC correlation map is scaled for display purposes.

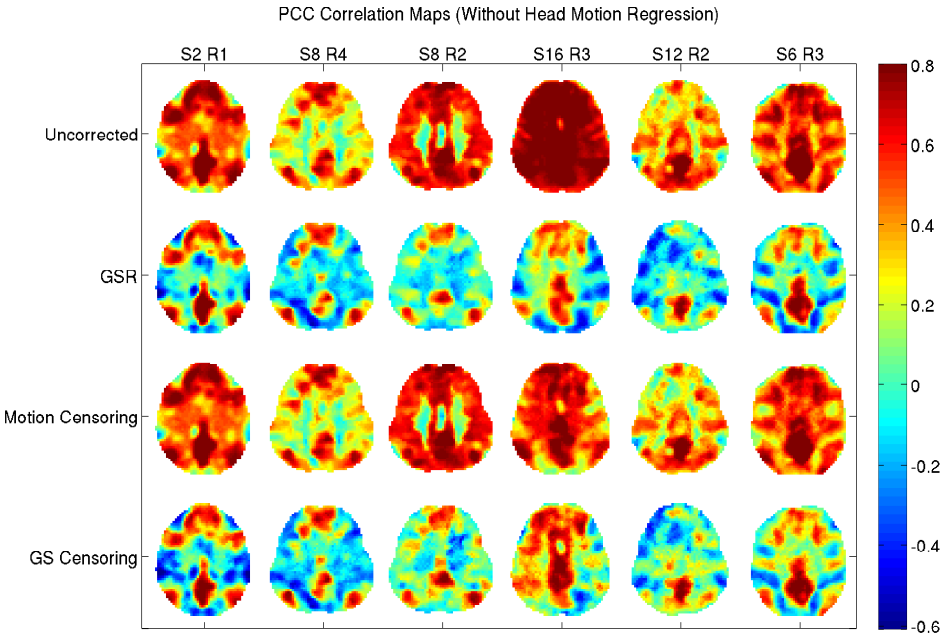

Supplementary Figure 7: PCC correlation maps obtained without regression of head motion covariates in the preprocessing stage. Maps are shown for the same scans used in Figure 12, with three low head motion runs (S2 R1, S8 R4, S8 R2) and three high motion runs (S16 R3, S12 R2, S6 R3). These maps were obtained before GSR and after the application of GSR, motion censoring, and GS censoring. The differences between motion and GS censoring are evident even when motion covariates are not regressed out.

Acknowledgments

This work was partially supported by the UC San Diego Frontiers of Innovation Scholars Program (FISP) Project Fellowship.

References

Anderson, J. S., Druzgal, T. J., Lopez-Larson, M., Jeong, E.-K., Desai, K., Yurgelun-Todd, D., Jun. 2011. Network anticorrelations, global regression, and phase-shifted soft tissue correction. Human Brain Mapping 32 (6), 919–934.

Biswal, B., Yetkin, F. Z., Haughton, V. M., Hyde, J. S., Oct. 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic Resonance in Medicine 34 (4), 537–541.

Carbonell, F., Bellec, P., Shmuel, A., 2011. Global and system-specific resting-state fMRI fluctuations are uncorrelated: principal component analysis reveals anti-correlated networks. Brain connectivity 1 (6), 496–510.

Carbonell, F., Bellec, P., Shmuel, A., Feb. 2014. Quantification of the impact of a confounding variable on functional connectivity confirms anti-correlated networks in the resting-state. NeuroImage 86, 343–353.

Chai, X. J., Castañón, A. N., Ongür, D., Whitfield-Gabrieli, S., Jan. 2012. Anticorrelations in resting state networks without global signal regression. NeuroImage 59 (2), 1420–1428.

Chang, C., Glover, G. H., Oct. 2009. Effects of model-based physiological noise correction on default mode network anti-correlations and correlations. NeuroImage 47 (4), 1448–1459.

Chang, C., Leopold, D. A., Schölvinck, M. L., Mandelkow, H., Picchioni, D., Liu, X., Ye, F. Q., Turchi, J. N., Duyn, J. H., Apr. 2016. Tracking brain arousal fluctuations with fMRI. Proceedings of the National Academy of Sciences of the United States of America 113 (16), 4518–4523.

Cordes, D., Haughton, V. M., Arfanakis, K., Carew, J. D., Turski, P. A., Moritz, C. H., Quigley, M. A., Meyerand, M. E., Aug. 2001. Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. AJNR American journal of neuroradiology 22 (7), 1326–1333.

Cordes, D., Haughton, V. M., Arfanakis, K., Wendt, G. J., Turski, P. A., Moritz, C. H., Quigley, M. A., Meyerand, M. E., 2000. Mapping functionally related regions of brain with functional connectivity MR imaging. American Journal of Neuroradiology 21 (9), 1636–1644.

Cox, R. W., Jun. 1996. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. Computers and biomedical research, an international journal 29 (3), 162–173.

Draper, N. R., Smith, H., Aug. 2014. Applied Regression Analysis, 3rd Edition. John Wiley & Sons.

Falahpour, M., Wong, C. W., Liu, T. T., 2016. The resting state fMRI global signal is negatively correlated with time-varying EEG vigilance. In: Proceedings of the 24th Annual Meeting of the ISMRM. p. 641.

Fox, M. D., Corbetta, M., Snyder, A. Z., Vincent, J. L., Raichle, M. E., 2006. Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. Proceedings of the National Academy of Sciences 103 (26), 10046–10051.

Fox, M. D., Raichle, M. E., Sep. 2007. Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nature Reviews Neuroscience 8 (9), 700–711.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., Raichle, M. E., 2005. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proceedings of the National Academy of Sciences of the United States of America 102 (27), 9673–9678.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Raichle, M. E., 2007. Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. Neuron 56 (1), 171–184.

Fox, M. D., Zhang, D., Snyder, A. Z., Raichle, M. E., Jun. 2009. The global signal and observed anticorrelated resting state brain networks. Journal of Neurophysiology 101 (6), 3270–3283.

Gotts, S. J., Saad, Z. S., Jo, H. J., Wallace, G. L., Cox, R. W., Martin, A., 2013. The perils of global signal regression for group comparisons: a case study of Autism Spectrum Disorders. Frontiers in human neuroscience 7, 356.

Hahamy, A., Calhoun, V., Pearlson, G., Harel, M., Stern, N., Attar, F., Malach, R., Salomon, R., Aug. 2014. Save the global: global signal connectivity as a tool for studying clinical populations with functional magnetic resonance imaging. Brain connectivity 4 (6), 395–403.

Hampson, M., Peterson, B. S., Skudlarski, P., Gatenby, J. C., Gore, J. C., 2002. Detection of functional connectivity using temporal correlations in MR images. Human brain mapping 15 (4), 247–262.

He, H., Liu, T. T., Feb 2012. A geometric view of global signal confounds in resting-state functional MRI. NeuroImage 59 (3), 2339–48.

Hoaglin, D. C., Welsch, R. E., Feb. 1978. The Hat Matrix in Regression and ANOVA. The American Statistician 32 (1), 17–22.

Jones, B. E., Nov. 2005. From waking to sleeping: neuronal and chemical substrates. Trends in pharmacological sciences 26 (11), 578–586.

Karahanoğlu, F. I., Van De Ville, D., July 2015. Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. Nature Communications 6 (7751).

Lacadie, C. M., Fulbright, R. K., Rajeevan, N., Constable, R. T., Papademetris, X., 2008. More accurate Talairach coordinates for neuroimaging using non-linear registration. Neuroimage 42 (2), 717–725.

Leys, C., Ley, C., Klein, O., Bernard, P., Licata, L., 2013. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. Journal of Experimental Social Psychology 49 (4), 764–766.

Li, J. M., Bentley, W. J., Snyder, A. Z., Raichle, M. E., Snyder, L. H., May 2015. Functional connectivity arises from a slow rhythmic mechanism. Proceedings of the National Academy of Sciences of the United States of America 112 (19), E2527–35.

Liang, Z., King, J., Zhang, N., Jan. 2012. Anticorrelated resting-state functional connectivity in awake rat brain. NeuroImage 59 (2), 1190–1199.

Liu, X., Duyn, J. H., 2013. Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proceedings of the National Academy of Sciences 110 (11), 4392–4397.

Liu, X., Yanagawa, T., Leopold, D., Chang, C., Ishida, H., Fuji, i. N., Duyn, J., Aug. 2015. Arousal transitions in sleep, wakefulness, and anesthesia are characterized by an orderly sequence of cortical events . NeuroImage 116, 222–231.

Macey, P. M., Macey, K. E., Kumar, R., Harper, R. M., May 2004. A method for removal of global effects from fMRI time series. NeuroImage 22 (1), 360–366.

Martinez, W. L., Martinez, A. R., Martinez, A., Solka, J., 2010. Exploratory data analysis with MATLAB. CRC Press.

Murphy, K., Birn, R. M., Handwerker, D. A., Jones, T. B., Bandettini, P. A., Feb. 2009. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? NeuroImage 44 (3), 893–905.

Picchioni, D., Duyn, J. H., Horovitz, S. G., Oct. 2013. Sleep and the functional connectome. NeuroImage 80, 387–396.

Pisauro, M. A., Benucci, A., Carandini, M., Jun. 2016. Local and global contributions to hemodynamic activity in mouse cortex. Journal of Neurophysiology 115 (6), 2931–2936.

Power, J. D., Barnes, K. A., Snyder, A. Z., Schlaggar, B. L., Petersen, S. E., Feb. 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. NeuroImage 59 (3), 2142–2154.

Power, J. D., Plitt, M., Laumann, T. O., Martin, A., 2016. Sources and implications of whole-brain fMRI signals in humans. NeuroImage, In Press.

Power, J. D., Schlaggar, B. L., Petersen, S. E., Jan. 2015. Recent progress and outstanding issues in motion correction in resting state fMRI. NeuroImage 105, 536–551.

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., Shulman, G. L., 2001. A default mode of brain function. Proceedings of the National Academy of Sciences 98 (2), 676–682.

Reimann, C., Filzmoser, P., Garrett, R. G., 2005. Background and threshold: critical comparison of methods of determination. Science of the Total Environment 346 (1), 1–16.

Remes, J. J., Starck, T., Nikkinen, J., Ollila, E., Beckmann, C. F., Tervonen, O., Kiviniemi, V., Silven, O., May 2011. Effects of repeatability measures on results of fMRI sICA: a study on simulated and real resting-state effects. NeuroImage 56 (2), 554–569.

Saad, Z. S., Gotts, S. J., Murphy, K., Chen, G., Jo, H. J., Martin, A., Cox, R. W., 2012. Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. Brain connectivity 2 (1), 25–32.

Schölvinck, M. L., Maier, A., Ye, F. Q., Duyn, J. H., Leopold, D. A., Jun. 2010. Neural basis of global resting-state fMRI activity. Proceedings of the National Academy of Sciences of the United States of America 107 (22), 10238–10243.

Siegel, J. S., Power, J. D., Dubis, J. W., Vogel, A. C., Church, J. A., Schlaggar, B. L., Petersen, S. E., May 2014. Statistical improvements in functional magnetic resonance imaging analyses produced by censoring high-motion data points. Human Brain Mapping 35 (5), 1981–1996.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., Chialvo, D. R., 2012. Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. Frontiers in Physiology 3 (15).

Van Dijk, K. R., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W., Buckner, R. L., 2010. Intrinsic functional connectivity as a tool for human connectomics: theory, properties, and optimization. Journal of neurophysiology 103 (1), 297–321.

Wager, T. D., Keller, M. C., Lacey, S. C., Jonides, J., 2005. Increased sensitivity in neuroimaging analyses using robust regression. Neuroimage 26 (1), 99–113.

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., Diedrichsen, J., December 2015. Reliability of dissimilarity measures for multi-voxel pattern analysis. NeuroImage 137, 188–200.

Weissenbacher, A., Kasess, C., Gerstl, F., Lanzenberger, R., Moser, E., Windischberger, C., Oct. 2009. Correlations and anticorrelations in resting-state functional connectivity MRI: a quantitative comparison of preprocessing strategies. NeuroImage 47 (4), 1408–1416.

Wen, H., Liu, Z., Jun. 2016. Broadband Electrophysiological Dynamics Contribute to Global Resting-State fMRI Signal. Journal of Neuroscience 36 (22), 6030–6040.

Wong, C. W., Olafsson, V., Tal, O., Liu, T. T., Oct. 2012. Anti-correlated networks, global signal regression, and the effects of caffeine in resting-state functional MRI. NeuroImage 63 (1), 356–364.

Wong, C. W., Olafsson, V., Tal, O., Liu, T. T., Dec. 2013. The amplitude of the resting-state fMRI global signal is related to EEG vigilance measures. NeuroImage 83, 983–990.

Yan, L., Zhuo, Y., Ye, Y., Xie, S. X., An, J., Aguirre, G. K., Wang, J., Apr. 2009. Physiological origin of low-frequency drift in blood oxygen level dependent (BOLD) functional magnetic resonance imaging (fMRI). Magnetic Resonance in Medicine 61 (4), 819–827.

Yang, G. J., Murray, J. D., Repovs, G., Cole, M. W., Savic, A., Glasser, M. F., Pittenger, C., Krystal, J. H., Wang, X.-J., Pearlson, G. D., Glahn, D. C., Anticevic, A., May 2014. Altered global brain signal in schizophrenia. Proceedings of the National Academy of Sciences of the United States of America 111 (20), 7438–7443.