Received: 18 September 2018 | Revised: 6 December 2018 | Accepted: 8 January 2019

DOI: 10.1002/hbm.24528

WILEY

RESEARCH ARTICLE

Modular preprocessing pipelines can reintroduce artifacts into fMRI data

Martin A. Lindquist[1] | Stephan Geuter[1,2] | Tor D. Wager[2] | Brian S. Caffo[1]

[1]Biostatistics, Johns Hopkins School of Public Health, Baltimore, Maryland

[2]Psychology and Neuroscience, University of Colorado Boulder, Boulder, Colorado

Correspondence
Martin A. Lindquist, Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, 615 N. Wolfe Street, E3634 Baltimore, MD 21205
Email: mlindqui@jhsph.edu

Funding information
NIH, Grant/Award Numbers: R01 EB016061, R01 EB026549

Abstract

The preprocessing pipelines typically used in both task and resting-state functional magnetic resonance imaging (rs-fMRI) analysis are modular in nature: They are composed of a number of separate filtering/regression steps, including removal of head motion covariates and band-pass filtering, performed sequentially and in a flexible order. In this article, we illustrate the shortcomings of this approach, as we show how later preprocessing steps can reintroduce artifacts previously removed from the data in prior preprocessing steps. We show that each regression step is a geometric projection of data onto a subspace, and that performing a sequence of projections can move the data into subspaces no longer orthogonal to those previously removed, reintroducing signal related to nuisance covariates. Thus, linear filtering operations are not commutative, and the order in which the preprocessing steps are performed is critical. These issues can arise in practice when any combination of standard preprocessing steps including motion regression, scrubbing, component-based correction, physiological correction, global signal regression, and temporal filtering are performed sequentially. In this work, we focus primarily on rs-fMRI. We illustrate the problem both theoretically and empirically through application to a test–retest rs-fMRI data set, and suggest remedies. These include (a) combining all steps into a single linear filter, or (b) sequential orthogonalization of covariates/linear filters performed in series.

KEYWORDS

preprocessing, resting-state, fMRI, artifacts, motion

1 | INTRODUCTION

In the past decade, resting-state functional magnetic resonance imaging (rs-fMRI) data have been increasingly used to study intrinsic functional connectivity in the human brain (Biswal, Zerrin Yetkin, Haughton, & Hyde, 1995). Using rs-fMRI, it has been shown that fluctuations in the blood oxygen level dependent (BOLD) signal in spatially distant regions of the brain are strongly correlated (Beckmann, DeLuca, Devlin, & Smith, 2005; De Luca, Beckmann, De Stefano, Matthews, & Smith, 2006; Yeo et al., 2011). While the exact mechanisms driving these correlations remain unclear, it has been hypothesized that it may be due to fluctuations in spontaneous neural activity. Neuroscientists have become increasingly interested in studying the correlation between spontaneous BOLD signals from different brain regions in order to learn more about human brain function (Van Den Heuvel & Pol, 2010).

The analysis is complicated by the fact that the measured BOLD signal consists of both changes induced by neuronal activation, as well as nonneuronal fluctuations. Here, the former is the signal of interest, while the latter is considered nuisance signal. Examples of such nonneuronal fluctuations include drift, spiking artifacts, motion-related artifacts, and fluctuations due to physiological sources such as heart rate and respiration. Failure to properly control for these types of noise can have significant impact on subsequent analysis. For example, head motion has been shown to have systematic effects on resting-state functional connectivity measures (Power et al., 2014; Van Dijk, Sabuncu, & Buckner, 2012), and comparisons between groups of subjects with different levels of head motion have yielded difference maps that could be mistaken for interesting neuronal effects (Van Dijk et al., 2012). Furthermore, it has been shown there are significant correlation between changes in cardiac and respiratory rates and the BOLD signal (Birn, Diamond, Smith, & Bandettini, 2006;

Shmueli et al., 2007; Wise, Ide, Poulin, & Tracey, 2004). These nuisance fluctuations risk artificially inflating functional connectivity measures, or even creating spurious findings (Murphy, Birn, & Bandettini, 2013).

In general, the relative influence of these different nonneuronal fluctuations depends on a number of factors (Caballero-Gaudes & Reynolds, 2017). However, it is clear that if not treated properly, these sources of variation can induce spurious functional connectivity between different brain regions. Thus, it is of great interest to reduce their effects on the analysis of rs-fMRI data.

For these reasons, rs-fMRI data are subjected to a series of preprocessing steps (Caballero-Gaudes & Reynolds, 2017; Weissenbacher et al., 2009; Yan & Zang, 2010) prior to analysis. Typically, each step consists of a separate algorithm often a linear regression of data on nuisance covariates, with residuals used in subsequent analysis, or a linear filtering operation designed to remove a specific type or class of artifacts. These steps can include motion regression (Fox, Zhang, Snyder, & Raichle, 2009; Weissenbacher et al., 2009), scrubbing (Power et al., 2014; Power, Barnes, Snyder, Schlaggar, & Petersen, 2012) or spike regression (Lemieux, Salek-Haddadi, Lund, Laufs, & Carmichael, 2007; Satterthwaite et al., 2013), nuisance regression (e.g., the removal of signal from white matter (WM) and ventricular cerebrospinal fluid (CSF) tissues; Behzadi, Restom, Liau, & Liu 2007; Muschelli et al., 2014), nonseed-based physiological regression methods (e.g., retrospective image correction [RETROICOR]/respiration variation and heart rate correction [RVHRCOR]; Glover, Li, & Ress 2000; Chang, Cunningham, & Glover 2009), data-derived nuisance component removal using ICA (e.g., ICA-based automatic removal of motion artifacts [ICA-AROMA]; Pruim et al. (2015)), the removal of global signal (Fox et al., 2005), and temporal filtering including low- and high-pass filters (Biswal et al., 1995; Cordes et al., 2001). The most popular versions of all of these can be expressed as linear filtering operations.

There are a large variety of rs-fMRI preprocessing pipelines (i.e., combinations of such modular preprocessing steps) that vary in the specific operations they perform on the data, as well as the order in which they are performed. In fact, a paper by Carp (2012) showed that there are nearly as many unique analysis pipelines in the literature as there were studies.[1] While most of the preprocessing steps that are performed are essential, there is relatively little understanding of the effects they have on both the spatial and temporal correlation structure of the resulting data. Importantly, there is a general lack of knowledge regarding potential interactions among the individual preprocessing steps, and how the order in which they are performed impacts the resulting analysis. To date, there is no consensus standard of what steps should be included in a pipeline, or in which order they should be performed. Indeed, investigations into the impact of the order of preprocessing steps have been largely absent from the literature (notable exceptions include Jones, Bandettini, & Birn 2008 and Hallquist, Hwang, & Luna 2013), though a number of groups have suggested that any preprocessing approach should be adapted on a data-dependent basis (Churchill, Spring, Afshin-Pour, Dong, & Strother, 2015; Kay, Rokem, Winawer, Dougherty, & Wandell, 2013; Salimi-Khorshidi et al., 2014).

[1]Note this survey also included task-based fMRI studies.

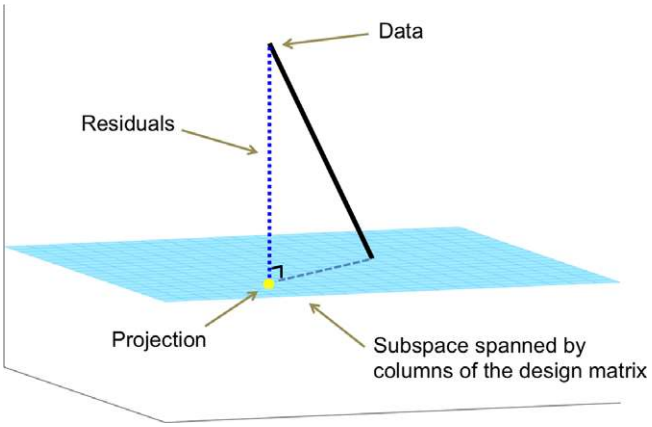

FIGURE 1 An illustration of the geometry of linear projections. The *n*-dimensional data are projected onto the *p*-dimensional subspace spanned by the columns of the design matrix (shown in light blue). The projection (shown in yellow) corresponds to the fitted value. The residuals (i.e., the data minus the fitted value) lie in a subspace that is orthogonal to the columns of the design matrix [Color figure can be viewed at wileyonlinelibrary.com]

Whatever the procedures performed, most popular pipelines have gravitated toward removing nuisance covariates in a series of sequential steps, which simplifies the analysis, without considering the order in which they are performed. However, we show here that these linear filtering operations are not commutative, and furthermore, that performing them in series can reintroduce nuisance signal removed in previous steps.

1.1 | A geometric approach

When working with linear models, it is often fruitful to take a geometric approach toward understanding their behavior by viewing them as linear projections. Here, the fitted value in the regression is seen as the orthogonal projection of the data onto the subspace spanned by the columns of the design matrix. Similarly, the residuals are the projection onto the subspace that is orthogonal to the columns of the design matrix, and are thus uncorrelated with the nuisance components that make up these columns. This is illustrated graphically in Figure 1, where the data (represented as an *n*-dimensional point) are projected onto the *p*-dimensional subspace spanned by the columns of the design matrix (shown in light blue).

In this article, we take a geometric approach toward analyzing modular rs-fMRI preprocessing pipelines. We express commonly used processing steps, such as motion regression, spike regression, nuisance regression, and temporal filtering, as projections onto a subspace of the full *n*-dimensional space (where *n* represents the number of time points) in which the data resides. This allows us to evaluate how different processing steps interact with one another, and how they in many settings actually counteract one another. In particular, it allows us to illustrate that both the order and the manner in which preprocessing steps are performed is critically important for being able to properly interpret subsequent analysis. Using our geometrical approach, we illustrate how, if not performed carefully, certain preprocessing techniques have the effect of reintroducing previously removed artifacts back into the signal.

We illustrate this issue primarily in the context of motion regression and temporal filtering, both theoretically and using test–retest rs-fMRI data. We show how performing high-pass filtering after motion regression reintroduces motion artifacts into the data. Similarly, performing motion regression after high-pass filtering reintroduces unwanted frequency components into the signal. The latter issue has been discussed in a number of papers, including work by Hallquist et al. (2013). Here, we take their results and place it into a broader geometrical framework. While the problems outlined in this article apply for any choice of temporal filter, we use a high-pass filter rather than the more commonly used band-pass filter throughout. The motivation for this choice comes from recent work by Shirer, Jiang, Price, Ng, and Greicius (2015) showing that using a broader filter improves signal-noise separation, test–retest reliability, and group discriminability for rs-fMRI data, both in the context of region-of-interest functional connectivity analysis and dual regression.

Although we focus on a few specific cases, these issues arise anytime multiple preprocessing steps are used that act as projections onto a subspace of the original data-space. In addition, though we focus on rs-fMRI data in this article, we stress that these issues can potentially arise in task fMRI as well if modular preprocessing is performed. To avoid these issues, we recommend that researchers either perform a simultaneous regression on all nuisance covariates, or alternatively orthogonalize later covariates with respect to the ones removed in earlier stages of the pipeline. This latter point implies one must orthogonalize both the data, and all subsequent projections, in order to maintain data orthogonality with the current projection. Interestingly, these approaches have become standard practice in the preprocessing of task-fMRI data.

We believe our framework has the potential to simplify the critical evaluation of preprocessing pipelines, and identify areas where problems can occur. In this article, we illustrate that the issues discussed in this work can have significant effect on subsequent analysis, and we therefore urge that special care be taken when performing preprocessing on rs-fMRI data. We further highlight the need to critically revisit previous work on rs-fMRI data work that may not have adequately controlled for these types of effects.

2 | THEORETICAL BACKGROUND

In this section, we show theoretically how the use of modular preprocessing steps can reintroduce artifacts that were removed in a previous step. We also provide recommendations for circumventing these issues.

2.1 | A geometric approach for evaluating pipelines

The issue we discuss in this article can potentially arise anytime one uses a preprocessing step that projects the data onto a subspace of the n-dimensional space in which it resides. This includes any technique that utilizes a linear model framework to remove artifacts from the signal. To illustrate the problem, let us consider the case of performing motion regression and temporal filtering sequentially. Hallquist et al. (2013) previously showed that in this setting nuisance-related variation can be reintroduced into frequencies that were previously suppressed by the filter. Here, we generalize their findings to incorporate a wider array of preprocessing steps, and place it into a general mathematical framework.

Let $\mathbf{y}$ be an n-dimensional vector containing the rs-fMRI signal from a specific voxel in the brain. Furthermore, let $\mathbf{X}$ be an $n \times p$ design matrix containing the nuisance regressors we seek to remove from $\mathbf{y}$. In our illustration, let us assume that $p = 24$, corresponding to (a) the six motion regressors obtained after rigid body transformation; (b) the regressors squared; (c) their first-order difference; and (d) the first-order difference squared.

To remove the effects of these nuisance components, the next step is to regress them out of $\mathbf{y}$. To do so, we fit a linear model on the form:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon} \tag{1}$$

It is well known that the least-squares estimate of $\boldsymbol{\beta}$ is given by

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}, \tag{2}$$

and the fitted value can be expressed as:

$$\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}. \tag{3}$$

Here, the term $\hat{\mathbf{y}}$ represents the estimated motion in the voxel and corresponds to the nuisance signal that we seek to remove from $\mathbf{y}$.

Here, it is useful to take a projection approach toward performing linear regression. To do so, we define

$$\mathbf{H} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}' \tag{4}$$

and thus we can write $\hat{\mathbf{y}} = \mathbf{H}\mathbf{y}$. In this setting, $\mathbf{H}$ is referred to as a projection matrix and its application has the effect of projecting the data onto the space spanned by the columns of the design matrix $\mathbf{X}$. Importantly, projection matrices are both idempotent ($\mathbf{H} = \mathbf{H}^2$) and symmetric ($\mathbf{H} = \mathbf{H}'$).

To remove the effects of motion, we subtract $\hat{\mathbf{y}}$ from the data and obtain the residual $\mathbf{e} = \mathbf{y} - \hat{\mathbf{y}}$. In terms of the projection matrix, we can write this as $\mathbf{e} = (\mathbf{I} - \mathbf{H})\mathbf{y}$, where $\mathbf{I}$ represents the $n \times n$ identity matrix. This term is now our signal of interest. It is important to note that $\mathbf{P}_1 = \mathbf{I} - \mathbf{H}$ is also a projection matrix (here, the subscript is simply an identifier to discriminate between the various projection matrices we will define), and it projects the data onto a subspace that is orthogonal to the columns of the design matrix $\mathbf{X}$. Thus, $\mathbf{e}$ resides in a subspace orthogonal to $\mathbf{X}$. Hence,

$$\begin{aligned} cov(\hat{\mathbf{y}}, \mathbf{e}) &= \langle (\mathbf{I} - \mathbf{P}_1)\mathbf{y},\ \mathbf{P}_1\mathbf{y} \rangle \\ &= \mathbf{y}'(\mathbf{I} - \mathbf{P}_1)\mathbf{P}_1\mathbf{y} \\ &= 0 \end{aligned} \tag{5}$$

and the data are uncorrelated with the motion, as required. The last equality holds as $\mathbf{P}_1 = \mathbf{P}_1^2$. (Here, $\langle .,. \rangle$ represents the Euclidean inner product.)

At this point, we have, to the extent possible by linear projections, successfully removed the effects of the nuisance components that make up the columns of $\mathbf{X}$ from the data. However, it is important to realize that researchers at this stage typically perform additional modular preprocessing steps that act as projections of the data into new subspaces. For example, this would occur if we were to continue

by next performing a preprocessing step such as spike regression, component-based correction (CompCor), RETROICOR, RVHRCOR, ICA-AROMA, global signal regression, or temporal filtering. This can have the unfortunate side effect of projecting the data back into the space spanned by the motion regressors in $\mathbf{X}$, and thus reintroducing the effects of components that had previously been removed. This will happen if the two subspaces are not orthogonal to each other.

To illustrate, suppose we perform high-pass filtering on the motion-regressed data. Filtering can be expressed in a linear model using a series of sine and cosine functions as regressors. Fitting this model projects the data onto a space orthogonal to the frequency components one seeks to remove. We can express this operation as $\mathbf{e}_f = \mathbf{P}_2\mathbf{e}$, where $\mathbf{P}_2$ represents a second projection matrix that differs from $\mathbf{P}_1$. Unfortunately, this operation projects the data into a subspace that is no longer constrained to be orthogonal to the space spanned by $\mathbf{X}$. This has the effect of partially reintroducing a correlation between the motion and the signal being studied. This can be noted by observing that

$$\begin{aligned} cov(\hat{\mathbf{y}}, \mathbf{e}_f) &= \langle (\mathbf{I} - \mathbf{P}_1)\mathbf{y},\ \mathbf{P}_2\mathbf{P}_1\mathbf{y} \rangle \\ &= \mathbf{y}'(\mathbf{I} - \mathbf{P}_1)\mathbf{P}_2\mathbf{P}_1\mathbf{y} \end{aligned} \tag{6}$$

is no longer guaranteed to be zero, and indeed in most situations will in fact be nonzero. This illustrates the problem from a theoretical point of view.

Empirically, we illustrate the problem on real resting-state data in Section 3. Figure 2d–f summarizes some of the results. Panel d shows the rs-fMRI time course and the estimated motion from a specific voxel of interest. The two time courses show a large positive correlation ($r = 0.7622$). After motion regression, the processed data and the estimated motion are now uncorrelated ($r = 0$), as would be expected from Equation 5. Finally, after high-pass filtering, the data and the estimated motion are now negatively correlated ($r = -0.3458$). Clearly, high-pass filtering has partially reintroduced a correlation between the data and motion.

Equation 6 highlights that a requirement for the data and motion to be uncorrelated is that $(\mathbf{I} - \mathbf{P}_1)\mathbf{P}_2\mathbf{P}_1 = \mathbf{0}$. Apart from trivial cases (e.g., $\mathbf{P}_1 = \mathbf{0}$, $\mathbf{P}_2 = \mathbf{0}$, $\mathbf{P}_1\mathbf{P}_2 = \mathbf{0}$, or $\mathbf{P}_2 = \mathbf{A}\mathbf{P}_1$ for some $\mathbf{A}$), this holds if $\mathbf{P}_1 = \mathbf{P}_2\mathbf{P}_1$, or equivalently that $\mathbf{P}_1(\mathbf{I} - \mathbf{P}_2) = \mathbf{0}$. This implies that the filter projects the data onto the space orthogonal to that spanned by $\mathbf{X}$. This will almost certainly not hold unless the filter is explicitly designed for this purpose, for example, if the regressors forming $\mathbf{P}_2$ are orthogonalized with respect to those forming $\mathbf{P}_1$.

Empirically, this issue is illustrated in Figure 2. Panels a–c show examples of the matrices $\mathbf{I} - \mathbf{P}_1$, $\mathbf{I} - \mathbf{P}_2$, and $(\mathbf{I} - \mathbf{P}_2)\mathbf{P}_1$, respectively, corresponding to motion regression, high-pass filtering, and their product. Clearly, the last term is not 0, and hence the covariance between $\mathbf{e}_f$ and $\hat{\mathbf{y}}$ will not be equal to 0, as shown empirically in Panel f.

Note a related problem arises if temporal filtering is first applied, followed by motion regression. Here, the later step reintroduces signal into frequency bands which had previously been removed. The reason is similar to that described above as the space orthogonal to $\mathbf{X}$ need not be orthogonal to the space spanned by terms corresponding to the retained frequency components. This issue was discussed in detail in Hallquist et al. (2013).

The issue is further illustrated in Figure 3. Here, we show graphically what happens when a point in n-dimensional space, represented by the black vector, is repeatedly projected onto different subspaces. Suppose the point is first projected onto the light-blue subspace. Note that though shown in two dimensions, this subspace is actually $(n - p_1)$-dimensional, where p_1 is the number of columns in the design matrix used to create the first projection matrix. The projection is shown in yellow. Further suppose, this point is subsequently projected onto the gray subspace. This subspace is $(n - p_2)$-dimensional, where p_2 is the number of columns in the design matrix used to create the second projection matrix. The new location is now shown in green. Here, we note that in the plot to the left, where the two subspaces are orthogonal to one another, the point simultaneously lies in both subspaces. In the plot to the right, where the subspaces are not orthogonal, the point is not located in the light-blue subspace.

To tie this example to the preprocessing setting, let us consider the light-blue subspace to be the space orthogonal to the motion regressors that one seeks to remove, and the gray subspace to be the space orthogonal to the frequency components one seeks to remove. The first projection moves the data into the space orthogonal to the motion components. Thus, the influence of motion has been effectively removed and the data, represented by the yellow point, are uncorrelated with motion. The second projection moves the data into the space orthogonal to the frequency components one seek to remove, thereby removing the effects of these components. The data, represented by the green point, are in the frequency band of interest. If the spaces are orthogonal, as in the left plot, the data continue to be uncorrelated with motion and in the frequency band of interest, regardless of which order the projections are performed. However, if these two subspaces are not orthogonal to one another, the second projection actually moves the data back into the space spanned by the motion regressors. This has the unfortunate effect of reintroducing their effects into the data. Doing the projections in the reverse order would leave a point uncorrelated with the motion regressors, but not in the frequency band of interest.

2.2 | Preprocessing steps as projections

While we have illustrated the problem in the context of motion regression and temporal filtering, we stress that it is not limited to these particular cases. The problem potentially arises in any setting when multiple preprocessing steps that act as linear projections are used sequentially. Preprocessing steps that fit this description include the removal of WM and CSF, CompCor, RETROICOR, RVHRCOR, ICA-AROMA, global signal regression, and spike regression.

In each of these steps, one can define a design matrix $\mathbf{X}_i$ consisting of the nuisance components one seeks to remove from the BOLD signal. These design matrices are then implicitly used to create a projection matrix

$$\mathbf{P}_i = \mathbf{I} - \mathbf{X}_i(\mathbf{X}_i'\mathbf{X}_i)^{-1}\mathbf{X}_i' \tag{7}$$

such as the ones described above. The matrix is then used to project the data onto a subspace that is orthogonal to the space spanned by the nuisance components, thus removing their effect.

For example, filtering can be performed within a linear framework by including sine and cosine functions of the appropriate frequencies into the design matrix. Spike regression can be performed by including delta functions corresponding to each time point one seeks to remove

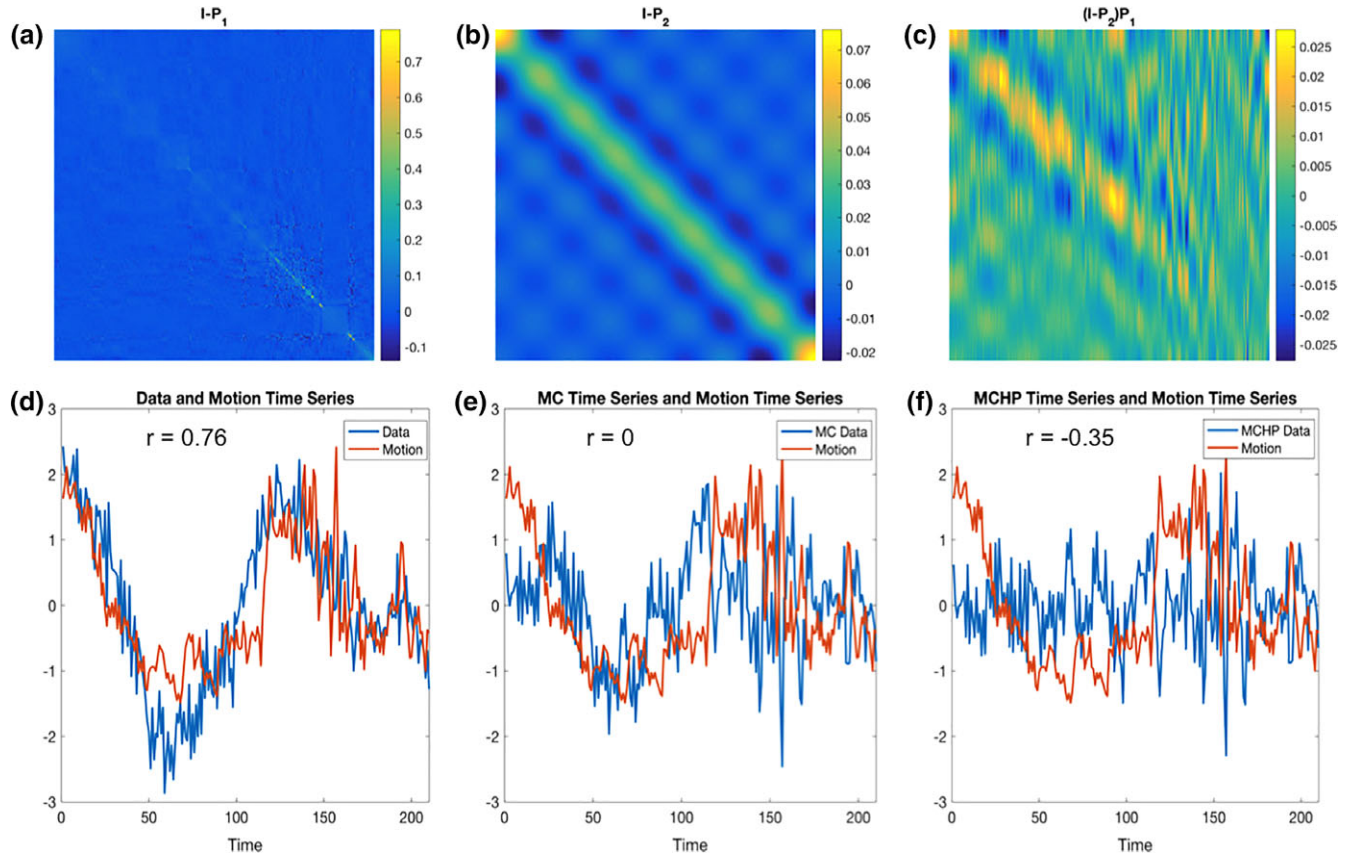

FIGURE 2 Examples of (a) a projection matrix $\mathbf{I} - \mathbf{P}_1$, removing empirically computed motion artifacts, (b) a temporal filtering projection matrix $\mathbf{I} - \mathbf{P}_2$, and (c) the product of both projection matrices $(\mathbf{I} - \mathbf{P}_2)\mathbf{P}_1$, which is clearly nonzero. (d) The resting-state functional magnetic resonance imaging time course (blue) and the estimated motion (red) from a voxel of interest. The two time courses are positively correlated ($r = 0.7622$). (e) The time course after motion regression (blue) together with the estimated motion (red). The two time courses are uncorrelated ($r = 0$). (f) The time course after sequential motion regression and high-pass filtering (blue) together with the estimated motion (red). The two time courses are negatively correlated ($r = -0.3458$). Hence, high-pass filtering has reintroduced a correlation between the data and motion [Color figure can be viewed at wileyonlinelibrary.com]

into the design matrix. Nuisance regression and CompCor can be performed in a similar manner as outlined for motion regression.

2.3 | Interaction among preprocessing steps

Each projection matrix $\mathbf{P}_i$, as defined in the previous section, projects the data onto a subspace that is orthogonal to the space spanned by the nuisance components included in the design matrix $\mathbf{X}_i$. As illustrated in Section 2.1, if a series of projections are applied sequentially there is a risk that the data are projected back into a subspace that one seeks to avoid.

Suppose we seek to perform m different modular preprocessing steps. The potential interaction between preprocessing steps can be explored by computing the product of the projection matrices $\mathbf{P}_1\mathbf{P}_2\cdots\mathbf{P}_m$ and evaluating whether or not it lies in the subspaces spanned by the columns of $\mathbf{X}_i$. Here, in order to not reintroduce nuisance signal contained in $\mathbf{X}_i$, we require

$$\mathbf{P}_1\mathbf{P}_2\cdots\mathbf{P}_m\mathbf{X}_i = 0 \tag{8}$$

for all $i = 1, \ldots m$.

If this condition does not hold, nuisance components will be reintroduced into the signal. Here, the order in which the preprocessing steps are performed becomes important. Consider the following two preprocessing pipelines: (a) motion regression followed by temporal filtering; and (b) temporal filtering followed by motion regression. In the first case, motion is reintroduced by temporal filtering, while in the second case, the filtered frequency components are reintroduced by motion regression. Which of these is more detrimental to the subsequent analysis can be debated, but regardless the reintroduction of unwanted nuisance components will ultimately change the interpretation of the findings.

As a general rule-of-thumb, the nuisance components related to the last preprocessing step performed should be adequately removed from the signal. This can be seen by noting that $\mathbf{P}_m\mathbf{X}_m = 0$ is always true, and thus Equation 8 holds. However, components corresponding to earlier steps are potentially reintroduced if not handled properly.

2.4 | Recommendations

The problem of reintroducing nuisance regressors can be avoided in two ways. First, one can define an omnibus projection matrix consisting of all nuisance variables one seeks to remove from the data. This entails performing motion regression, CompCor, temporal filtering, and so on in a common joint model. This is simple to do in a linear model framework by concatenating the m design matrices into a large omnibus design matrix:

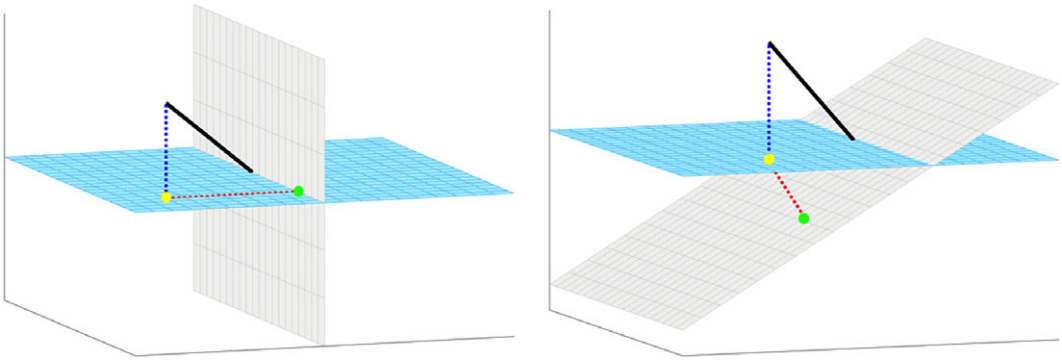

FIGURE 3 Illustration of the effect of multiple projections. Consider a point in *n*-dimensional space represented by the black vector. The point is first projected onto the light-blue subspace; see yellow point. The first projection is subsequently projected onto the gray subspace; see green point. In the plot to the left the two subspaces are orthogonal to one another and thus the second projection simultaneously lies in both subspaces. In the plot to the right, the green point is no longer located in the light blue subspace and thus not orthogonal to the confounds defining that subspace [Color figure can be viewed at wileyonlinelibrary.com]

$$\mathbf{X} = [\mathbf{X}_1 \quad \mathbf{X}_2 \quad \ldots \quad \mathbf{X}_m]. \tag{9}$$

Using this design matrix ensures that the data are projected into a subspace that is orthogonal to all nuisance components contained in **X**.

Another approach is to create the design matrices in such a manner that they project the data onto a series of orthogonal subspaces. For example, spike regression projects the data onto a subspace orthogonal to the removed time points. If motion regression is subsequently performed then it is necessary to remove the same time points as in the spike regression from the motion design matrix. If not, the contribution of these time points is reintroduced into the signal. Another example is the application of temporal filtering before motion regression. Here, both the data and the motion regressors must be filtered to avoid reintroducing the filtered bands back into the data after motion regression. This general point can be summarized as: one must orthogonalize a linear filter or set of nuisance covariates with respect to all previously removed sets of covariates to avoid reintroducing nuisance signals removed in previous steps.

3 | METHODS

While in Section 2 we focused on showing theoretically how the use of sequential modular preprocessing steps can reintroduce previously removed nuisance signal, in this section we focus on showing this empirically. We show this primarily in the context of motion regression and high-pass filtering applied to test–retest rs-fMRI data. We explore how the use of modular preprocessing pipelines reintroduces artifacts that have been previously removed. We also seek to show that using a joint (nonmodular) pipeline circumvents these issues.

3.1 | Data collection

We used the Multi-Modal MRI Reproducibility Resource (Kirby) from the F.M. Kirby Research Center[2]; see Landman et al. (2011) for a detailed explanation of the acquisition protocol. Briefly, it consists of data from 21 healthy adults scanned on a 3T Philips Achieva Scanner. Each participant completed two scanning sessions on the same day. Between sessions the participants briefly exited the scan room and a full repositioning of the participant, coils, blankets, and pads occurred prior to the second session. A T1-weighted MPRAGE structural run was acquired during both sessions (acquisition time = 6 min, TR/TE/TI = 6.7/3.1/842 ms, resolution = $1 \times 1 \times 1.2$ mm^3, SENSE factor = 2, flip angle = 8°). A multislice sensitivity encoded echo planar imaging (SENSE-EPI) pulse sequence (Pruessmann, Weiger, Scheidegger, Boesiger, et al., 1999; Stehling, Turner, & Mansfield, 1991) was used to acquire one rs-fMRI run during each session, where each run consisted of 210 volumes sampled every 2 s at 3 mm isotropic spatial resolution (acquisition time: 7 min, TE = 30 ms, SENSE acceleration factor = 2, flip angle = 75°, 37 axial slices collected sequentially with a 1 mm gap). Participants were instructed to rest comfortably while remaining as still as possible, and no other instruction was provided. We will refer to the first rs-fMRI run collected as Session 1 and the second as Session 2. One participant was excluded from data analyses due to excessive motion.

[2]Publicly available at http://www.nitrc.org/projects/multimodal.

3.2 | Initial preprocessing

Statistical Parametric Mapping (SPM) 12 (Wellcome Trust Centre for Neuroimaging, London, UK) and custom MATLAB (The MathWorks, Inc., Natick, MA) scripts were used to preprocess the Kirby data. To allow for the stabilization of magnetization, four volumes were discarded at acquisition, and an additional volume was discarded prior to preprocessing. Slice timing correction was performed using the slice acquired at the middle of the TR as reference, and rigid body realignment parameters were estimated to adjust for head motion. Structural runs were registered to the first functional frame and spatially normalized to Montreal Neurological Institute (MNI) space using SPMs unified segmentation–normalization algorithm (Ashburner & Friston, 2005). The estimated rigid body and nonlinear spatial transformations were applied to the rs-fMRI data.

3.3 | Creation of motion images

Before proceeding with further preprocessing steps, we estimated the motion at each voxel of the brain. To do so we used a design matrix **X**

consisting of 24 regressors. These included the six motion regressors obtained after rigid-body transformation, the regressors squared, their first-order difference, and the first-order difference squared. After estimating the parameters $\boldsymbol{\beta}_v$ corresponding to these regressors at each voxel v, we proceeded to create "motion images." This was done by computing the fitted values $\mathbf{X}\boldsymbol{\beta}_v$ at each voxel. Thus, we were able to create 4D images of the estimated contribution of motion to the signal at each voxel in the brain. Ultimately, the goal of preprocessing is to remove the effects of motion. However, at this stage, we do not remove this component, but simply compute it to use as a baseline to evaluate the motion-related components left in the signal after performing specific preprocessing steps.

3.4 | Secondary preprocessing pipelines

In the next stage of preprocessing, we constructed three different preprocessing pipelines. In the first pipeline, we first perform motion regression (as described above) and remove the estimated motion from the data. Thereafter, we apply a high-pass filter using a cutoff frequency of 0.01 Hz. We used the high-pass filter implemented in SPM, which uses a set of discrete cosine transformation basis functions (Frackowiak, 2004). We refer to this pipeline as "motion regression followed by high-pass filtering (MCHP)." In the second pipeline, we begin by applying the high-pass filter, and thereafter perform motion regression to remove the estimated motion from the data. We refer to this pipeline as "high-pass filtering followed by motion regression (HPMC)." Finally, in the third pipeline, we perform motion regression and high-pass filtering jointly using a combined design matrix. We refer to this pipeline as "Joint."

While we will primarily study pipelines consisting of two preprocessing steps, high-pass filtering, and motion regression, we note that errors can propagate further when including additional steps. We illustrate using two different three-step pipelines: (a) high-pass filtering, followed by motion regression, followed by nuisance regression (denoted "HPMCNR"); and (b) motion regression, followed by nuisance regression, followed by high-pass filtering (denoted "MCNRHP"). Here, nuisance regression involves removal of the average WM and CSF signal obtained from each subject using eroded masks. These images are based on the SPM8 a priori tissue probability maps that have been cleaned up and eroded to ensure that the WM and CSF compartments are unlikely to contain significant gray matter.

3.5 | Evaluation of pipelines

Next, we sought to evaluate the interaction between motion regression and high-pass filtering in each pipeline. For each pipeline, and after each modular step, we computed the correlation between the preprocessed data and the motion images at each voxel. This was done to evaluate the residual contribution of motion after performing each preprocessing step.

We also parcellated the data at each stage into 268 regions using the Shen atlas (Shen, Tokoglu, Papademetris, & Constable, 2013), and computed the correlation matrices across regions as commonly done in rs-fMRI studies on whole brain functional connectivity. This was repeated for both sessions for each of the 20 subjects. For each session and pipeline, we computed the average correlation matrices at the group level. We also performed t tests to determine whether the Fisher-transformed correlation between the motion time courses and the data after MCHP was significantly different from 0.

In addition, for each pipeline, after each step, we estimated the spectral density for every voxel time course in order to evaluate the relative contribution of different frequencies components. In particular, we focused on the contribution of frequencies lower than 0.01 Hz, as these are the ones we sought to remove from the signal. We will quantify this contribution by computing the proportion of the power that lies in the frequency band below 0.01 Hz in each voxel. This is a measure that is similar in spirit to fractional amplitude of low frequency fluctuations (Zou et al., 2008). To evaluate whether the reintroduced signal is related to motion, we computed the spectral coherence between the estimated motion time courses and the data at each voxel after performing each of the different preprocessing pipelines. The coherence takes values between 0 and 1, and measures the relationship between two signals at each frequency. Here, values close to 1 indicate strong coherence, while values around 0 indicate low coherence.

4 | RESULTS

To properly understand the effects of modular preprocessing pipelines, we illustrate its effects on a single voxel, a single subject, and group analysis.

4.1 | Illustration of voxel-level effects

Figure 4 illustrates the problem at a single voxel of the brain. It shows an example fMRI time course and the estimated motion time course, which are highly correlated with one another ($r = 0.7622$); see Panels a and b. Here, Panel a shows plots of the two time courses, while Panel b shows the relationship using a scatterplot. After motion regression, the resulting times series (Panel c) is uncorrelated with the motion (Panel d), which is to be expected. However, high-pass filtering has the effect of reintroducing the correlation between the data and the motion, which is now $r = -0.3458$ (Panels e and f). This illustrates how the two preprocessing steps interact with one another. The high-pass filter has the effect of projecting the data back into the space spanned by the motion regressors, thus reintroducing a correlation between the signal and the motion.

Note that if high-pass filtering is performed before motion regression (i.e., the HPMC pipeline), the resulting time course is uncorrelated with the motion (Panels g and h). However, this leads to the reintroduction of frequency components that had previously been removed. This can clearly be seen studying the periodograms shown in Figure 5. Prior to the secondary preprocessing steps, there is clear signal in the range below 0.01 Hz (see Panel a), which has strong coherence with the motion time series at these frequencies (Panel b). If motion regression is performed followed by high-pass filtering (MCHP), then the contribution of frequencies in this range disappears as seen in Panel c. However, if the order of these preprocessing steps is reversed it is clear from Panel e that signal is reintroduced into these frequencies. That the reintroduced signal is related to motion can be seen in Panel f,

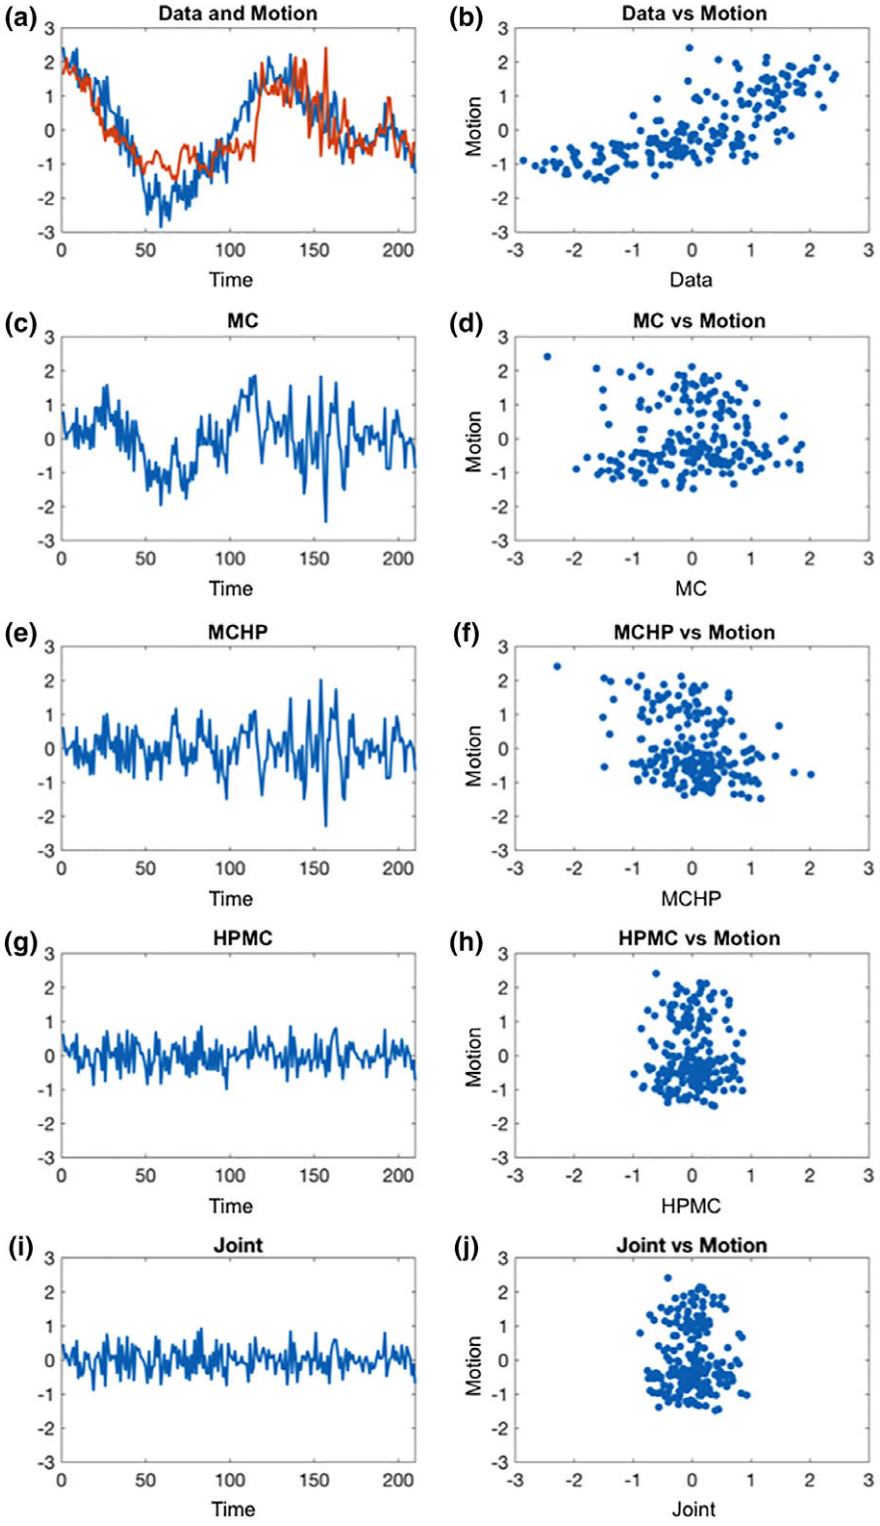

FIGURE 4 Results from an example voxel. (a) The resting-state functional magnetic resonance imaging time course before secondary preprocessing (blue) and the estimated motion (red) from the voxel of interest. (b) A scatterplot of these two time courses show a positive correlation ($r = 0.7622$). (c) The data after motion regression. (d) A scatterplot of this time course and the estimated motion show no correlation ($r = 0$). (e) The data after motion regression followed by high-pass filtering (MCHP). (f) A scatterplot of this time course and the estimated motion show a negative correlation ($r = -0.3458$). (g) The data after high-pass filtering followed by motion regression (HPMC). (h) A scatterplot of this time course and the estimated motion show no correlation ($r = 0$). (i) The data after joint high-pass filtering and motion regression (Joint). (j) A scatterplot of this time course and the estimated motion show no correlation ($r = 0$) [Color figure can be viewed at wileyonlinelibrary.com]

which shows the coherence between the motion time series and the signal after high-pass filtering and motion regression. Clearly, it is increased compared to the results seen in Panel d.

One approach toward handling this problem is to perform both operations within a joint model. As can be seen from Figures 4i,j and 5g,h, performing joint motion regression/temporal filtering provides a

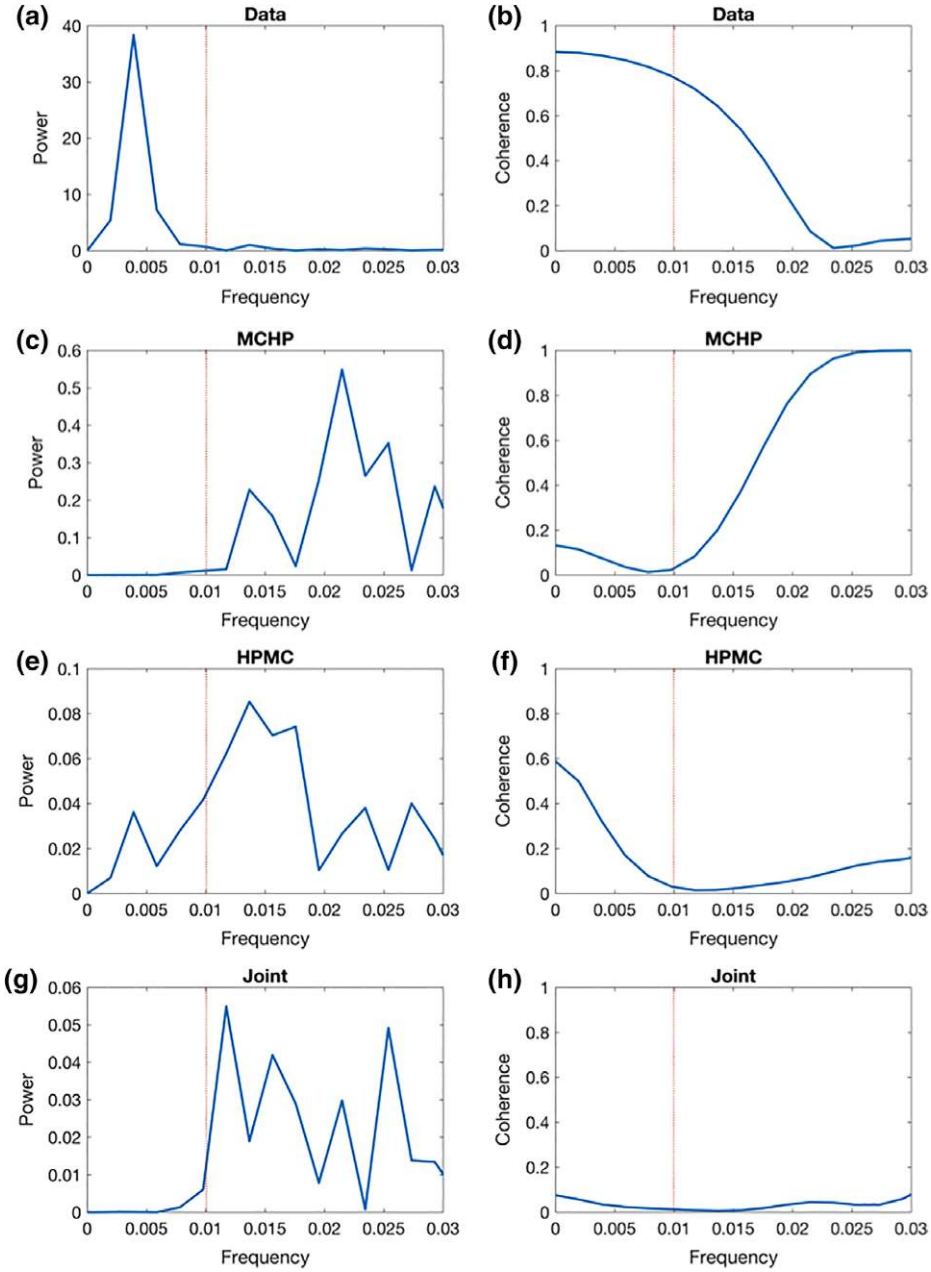

FIGURE 5 Results from an example voxel. (a,b) The periodogram and coherence, respectively, in the range [0 0.03] Hz of the resting-state functional magnetic resonance imaging time course before secondary preprocessing. (c,d) The periodogram and coherence, respectively, of the data after motion regression and high-pass filtering (MCHP). (e,f) The periodogram and coherence, respectively, after high-pass filtering and motion regression (HPMC). Note these two plots show that signal has been reintroduced into frequencies below 0.01 Hz, and that this signal is related to the motion time course. (g,h) The periodogram and coherence, respectively, after joint HMPC (Joint). In each panel, the red dotted line reflects the frequency cutoff of the high-pass filter [Color figure can be viewed at wileyonlinelibrary.com]

time course that is uncorrelated with motion, while only containing information in the frequencies of interest. Thus, in this approach, the effects of both motion regression and temporal filtering are retained in the processed signal as intended.

4.2 | Illustration of whole-brain effects

The point illustrated above is reinforced by looking at each voxel from a randomly chosen subject. First, we focus on studying the reintroduction of motion; see Figure 6. Panel a shows the correlation between the motion regressed data and the motion at each voxel of the brain. Clearly, the correlation is negligible ($<2 \times 10^{-6}$) across all voxels. However, after high-pass filtering, the correlation is reintroduced; see Panel b. Interestingly, the correlation between the data and motion is now mostly negative. Panel c shows these correlations superimposed onto the brain. Clearly, there are patterns consistent with motion artifacts, including a ring-like shape at the edge of the brain. There are also apparent patterns in the frontal cortex.

To further illustrate the impact of the reintroduction of motion, we study the correlation matrices across the 268 regions of the Shen atlas (Shen et al., 2013) both when (a) performing MCHP; and (b) performing HPMC. In the first case, the motion is reintroduced by

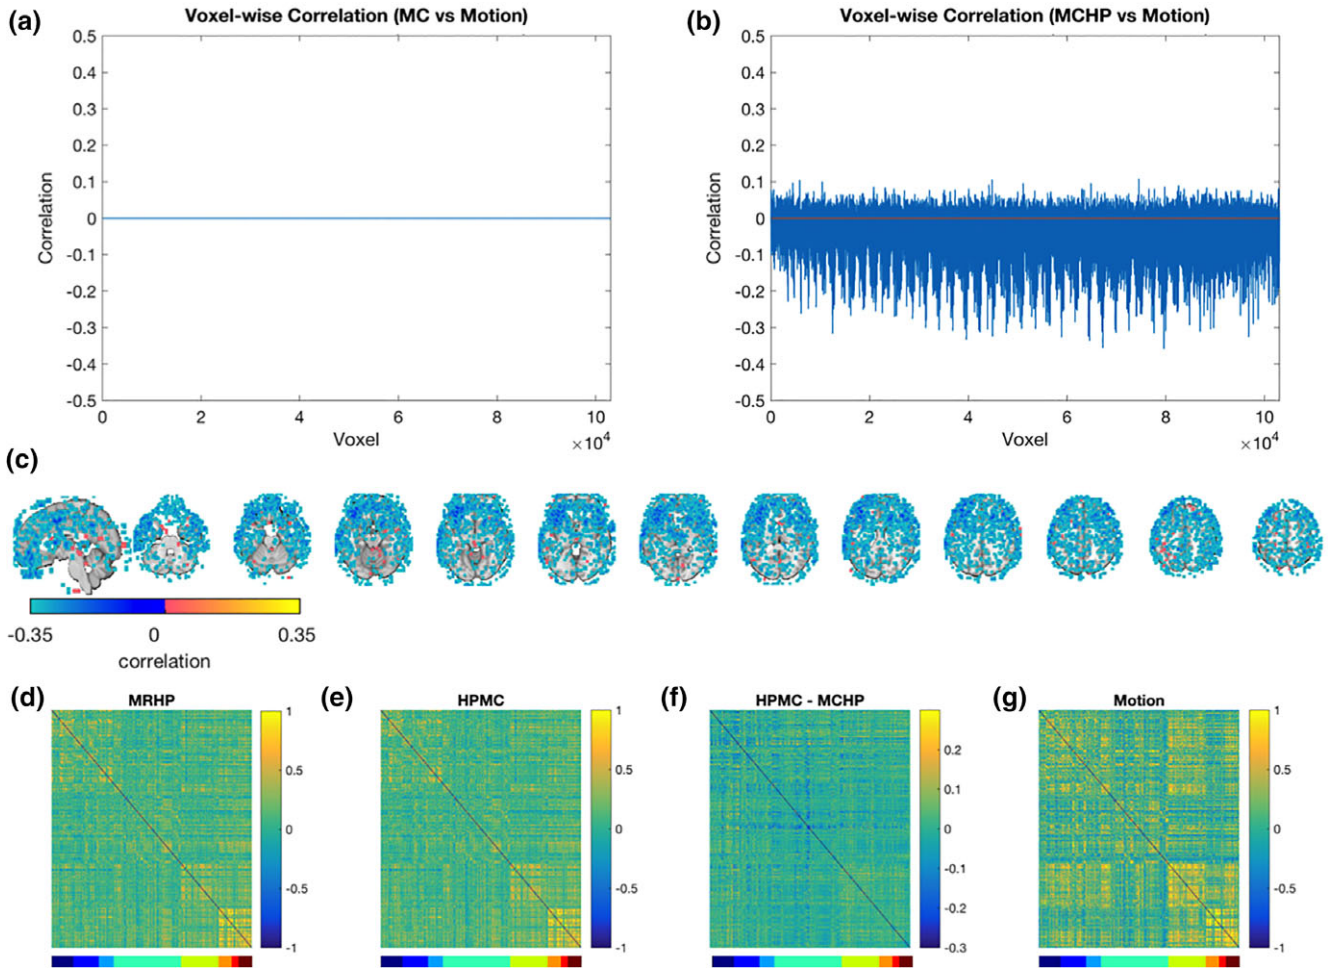

FIGURE 6 Results from an example subject. (a) The correlation between the motion regressed data and the motion at each voxel of the brain is negligible ($<2 \times 10^{-6}$). (b) The correlation is reintroduced after high-pass filtering (MCHP). (c) Correlations from (b) superimposed onto their brain locations. They are thresholded at an arbitrary value of ±0.05 for visualization purposes. (d) The correlation matrix between the 268 regions of the Shen atlas computed on data where motion regression is performed followed by high-pass filtering (MCHP). (e) The same correlation matrix on data where high-pass filtering is performed followed by motion regression (HPMC). (f) the difference between the correlation matrices in e and f. (g) The correlation matrix between the motion time courses over the same regions. Note regions are further split into eight different networks: Medial frontal (dark blue), frontoparietal (blue), default mode (light blue), subcortical-cerebellum (cyan), motor (green), visual I (orange), visual II (red), and visual association (brown) [Color figure can be viewed at wileyonlinelibrary.com]

temporal filtering, while in the second case, it should be removed. Note that in the second case, filtered frequency components are instead reintroduced (see Section 2.3). Panels d and e show the results, and f shows the difference between the two correlation matrices. Clearly, the differences are substantial as it ranges between values of ±0.3. Interestingly, the patterns seen in the difference (Panel f) show similarities to the correlation between the motion time courses over the same regions (Panel g). For interpretation purposes, the 268 regions are further separated into eight different networks (Finn et al., 2015), as illustrated by the color bars shown below the plots in Panels d–g. These include: medial frontal (dark blue), frontoparietal (blue), default mode (light blue), subcortical-cerebellum (cyan), motor (green), visual I (orange), visual II (red), and visual association (brown). The largest differences between pipelines lie in the motor and visual networks, and in the subcortical-cerebellum network.

Note that if high-pass filtering is performed before motion regression, we deal with a different problem as illustrated in Figure 7. Panel a shows the proportion of the power that lies in the frequency band below 0.01 Hz in each voxel after MCHP. Here, we can note that as expected there is a negligible contribution as all values are $<3 \times 10^{-3}$. However, Panel b shows the same plot for data where motion regression is performed after high-pass filtering. The plot shows there has been a significant reintroduction of information from frequencies that had previously been removed. In certain voxels up to 12% of the power lies in frequencies that should ideally be 0. The spatial position of these voxels is shown in Panel d, where the results are superimposed onto the brain. The results are presented using an arbitrary threshold of 2% to better be able to identify voxels with a high proportion of power in the low frequency band. Finally, Panel c shows the same plot for data preprocessed jointly. Much like in Panel a there a negligible contribution of signal in frequencies below 0.01 Hz.

To further explore whether the reintroduced signal is related to motion, we computed the coherence between the estimated motion times courses and the time series obtained using each pipeline at each voxel. The results are shown in Figure 8. Panel a shows the average

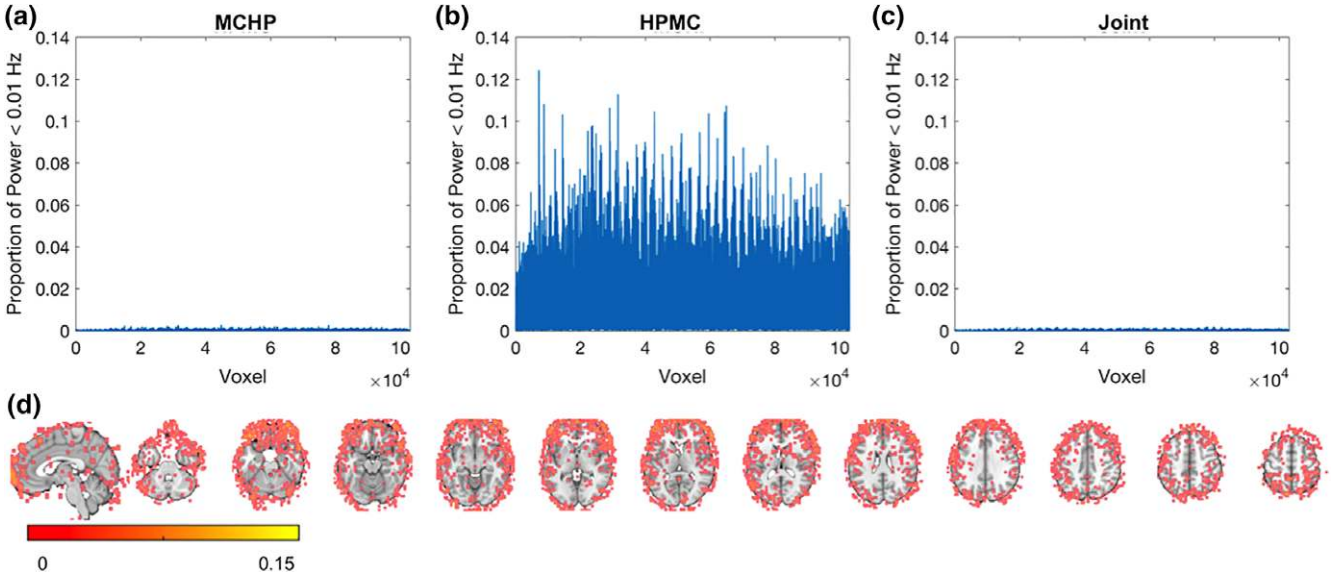

FIGURE 7 Results from an example subject. (a) The proportion of the power that lies in the frequency band below 0.01 *Hz* in each voxel after motion regression followed by high-pass filtering (MCHP). Note there is a negligible contribution. (b) Same plot for data where motion regression is performed after high-pass filtering (HPMC). In certain voxels up to 12% of the power lies in frequencies that should ideally be 0. (c) Same plot for data processed jointly (Joint). Much like in (a) there a negligible contribution. (d) the HPMC results shown in (b) superimposed onto the brain. Note the results are presented using an arbitrary threshold of 2% for visualization purposes [Color figure can be viewed at wileyonlinelibrary.com]

coherence between the signal and motion time series in the frequency band below 0.01 Hz for each voxel after MCHP, while Panels b and c show the same plots for data where motion regression is performed after high-pass filtering (HPMC) and processed jointly (Joint), respectively. Note the coherence for the signal obtained through the HPMC pipeline is heightened compared to the other preprocessing pipelines, illustrating that the signal present in the low-frequency band is more highly related to motion. Panel g shows the results of the HPMC pipeline superimposed onto the brain. The results are presented using an arbitrary threshold of 0.3 for visualization purposes. Panels d–f show similar results corresponding to the average coherence in the frequency band above 0.01 Hz. Note the average coherence for the signal obtained through the MCHP pipeline (Panel e) is slightly larger than the other preprocessing pipelines. Panel h shows the results from this pipeline superimposed onto the brain.

While we have illustrated the issues using pipelines consisting of two preprocessing steps, high-pass filtering, and motion regression, we note that errors can propagate further when including more steps. We illustrate this using two different three-step pipelines: (a) high-pass filtering, followed by motion regression, followed by nuisance regression (HPMCNR); and (b) motion regression, followed by nuisance regression, followed by high-pass filtering (MCNRHP). Figure S1 (Supporting Information) shows a scatterplot of the estimated correlation between the MCHP data and the motion at each voxel plotted against similar values for the MCHRHP pipeline. The plot illustrates that the negative correlation between signal and motion induced in the MCHP pipeline is strengthened by including nuisance regression to the pipeline. Panel b shows the voxel-wise difference between the correlations from the two pipelines superimposed onto the brain. Clearly, there are heightened correlations around the ventricles, as might be expected, as well as also other parts of the brain.

Panel a of Figure S2 (Supporting Information) shows the proportion of the power that lies in the frequency band below 0.01 Hz in each voxel using the HPMC plotted against the same values for the HPMCNR pipeline. Clearly, HPMCNR provides higher values, indicating that the pipeline introduces additional signal into this frequency band compared to the HPMC pipeline. Panel c shows the average coherence between the signal and motion time series in the frequency band below 0.01 Hz in each voxel using the HPMC pipeline plotted against the same values using the HPMCNR pipeline, illustrating that the coherence between the signal and motion is higher when using the more complex pipeline. Panels b and d show similar results comparing MCHP and MCNRHP. As expected here the difference is not large, as high-pass filtering is the last operation performed in each pipeline. However, the coherence between the signal and motion is again higher when using the more complex pipeline.

Finally, Panel a of Figure S3 (Supporting Information) shows the difference between the proportion of the power that lies in the frequency band below 0.01 Hz in each voxel for the HPMCNR and HPMC pipelines. Panel b shows the difference between the average coherence between the signal and motion time series in the frequency band below 0.01 Hz in each voxel for the HPMCNR and HPMC pipelines. Finally, Panel c shows the same results showing the difference between the MCNRHP and MCHP pipelines. Together, Figures S1–S3 (Supporting Information) illustrate that as the number of sequential preprocessing steps increase, the issues related to the reintroduction of signal increases as well.

4.3 | Illustration of group-level effects

Finally, we turn our attention to group-level analysis. Here, we analyze the two sessions of the Kirby data set separately for comparison purposes. Figure 9 shows *t* maps for both sessions that test whether the

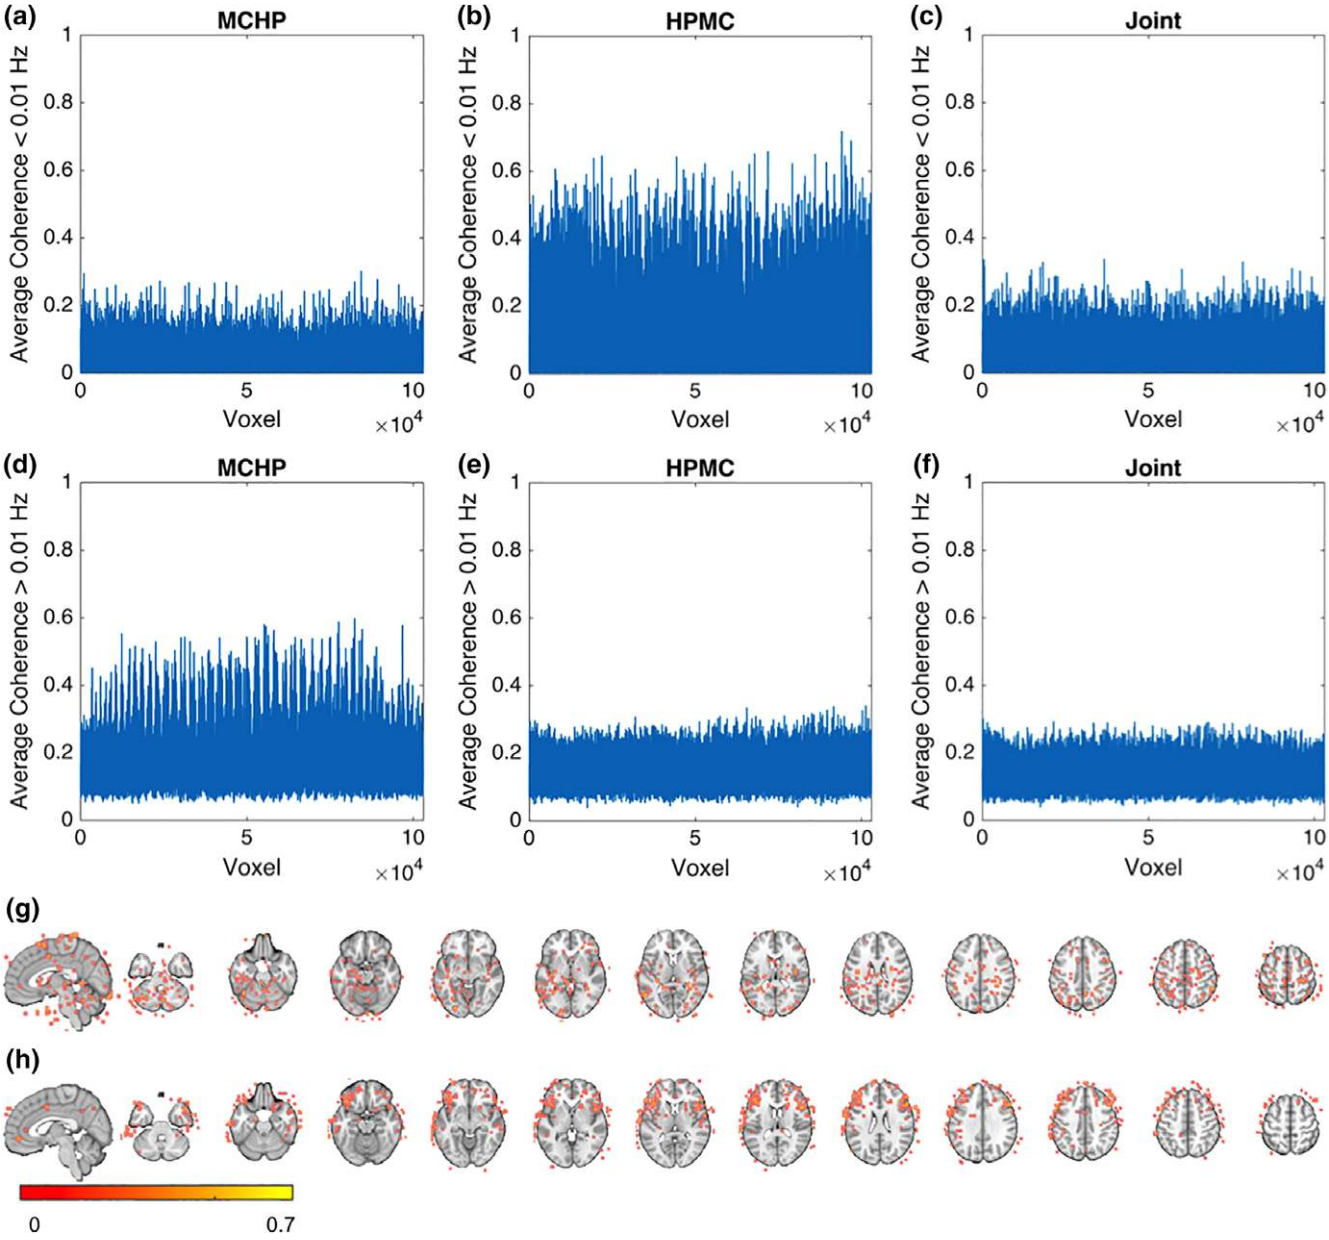

FIGURE 8 Results from an example subject. (a) The average coherence between the signal and motion time series in the frequency band below 0.01 *Hz* in each voxel after motion regression followed by high-pass filtering (MCHP). (b) Same plot for data where motion regression is performed after high-pass filtering (HPMC). (c) Same plot for data processed jointly (Joint). Note the coherence for the signal obtained using the HPMC pipeline is heightened compared to the other preprocessing pipelines. (d–f) Similar results showing the average coherence in the frequency band above 0.01 *Hz*. Note the average coherence for the signal obtained through the MCHP pipeline (d) is slightly larger than the other preprocessing pipelines. (g) The HPMC results shown in (b) superimposed onto the brain. (h) The MCHP results shown in (d) superimposed onto the brain. Note for both (g) and (h), the results are presented using an arbitrary threshold of 0.3 for visualization purposes [Color figure can be viewed at wileyonlinelibrary.com]

Fisher-transformed correlation between the motion time courses and the data after MCHP is significantly different from 0. Results are thresholded at $p < 0.001$. Clearly, the results are very similar across the two sessions, and there is clear correlation in frontal areas of the brain.

Panels (a) and (b) of Figures 10 and 11 show the average correlation matrix for data obtained from the MCHP and HPMC pipelines, for the two sessions. Panel (c) shows the difference between these two correlation matrices. In both sessions, the differences range between values of ±0.08. This illustrates that while averaging across subjects has removed some of the differences between preprocessing streams apparent at the single-subject level (see Figure 6); there remain substantial differences that are consistent in sign across participants, introducing systematic bias in group-level results. Panels (d–f) show the *SD* of the correlation matrices (MCHP, HPMC, HPMC–MCHP, respectively) across subjects. The visual networks show the largest between subject variation for both the MCHP and HPMC pipelines, while the medial frontal, frontoparietal, default mode networks show the largest variation in the difference between pipelines. Figures S4 and S5 (Supporting Information) show the difference

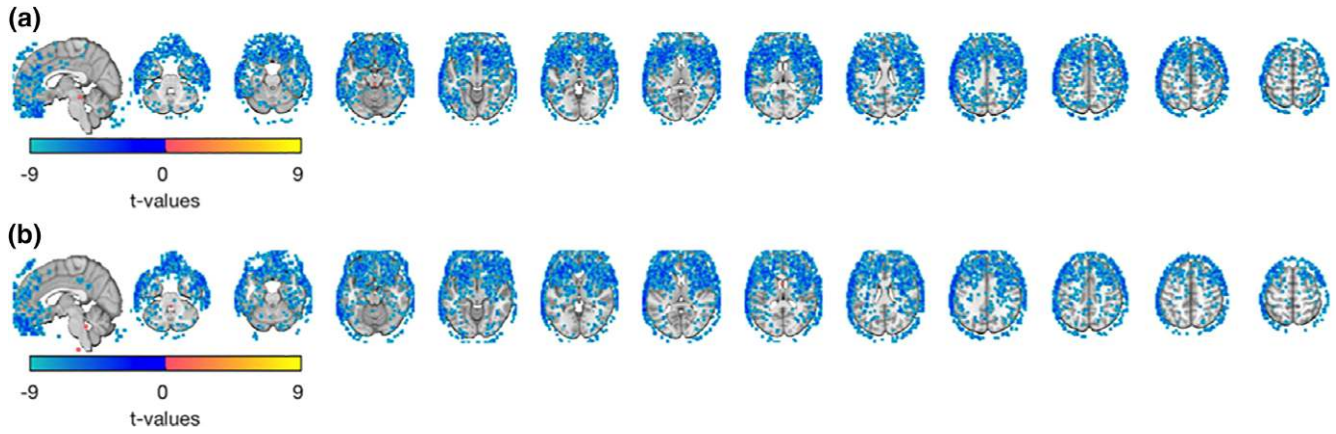

FIGURE 9 Results from the group-level analysis. (a) Group-level *t* maps for data from Session 1 testing whether the Fisher-transformed correlation between the motion time courses and the data after motion regression followed by high-pass filtering (MCHP) is significantly different from 0. Results are thresholded at $p < 0.001$. (B) Same results for data from Session 2 [Color figure can be viewed at wileyonlinelibrary.com]

between correlation matrices for each of the 20 subjects. Clearly, there are substantial differences across subjects in the locations where the largest differences between pipelines occur, which may ultimately counteract large biases from appearing in the group-level results.

It is interesting to see in Figure S6 (Supporting Information) how these differences are enhanced when incorporating an additional preprocessing step. Panels (a) and (b) show the average correlation matrix when using the MCNRHP and HPMCNR pipelines, respectively, for the 20 subjects in Session 1. Panel (c) shows the difference between these two correlation matrices. Comparing these results to the equivalent results in Figures 10 and 11 indicate a larger difference between the two preprocessing pipelines. Panels (d–f) show the *SD* of the correlation matrices across subjects.

Panels (a–c) of Figure 12 show the group average proportion of power that lie in the frequency band below 0.01 Hz at each voxel for both sessions using the three two-step pipelines (MCHP, HPMC, and Joint) and data from Session 1. Note that for MCHP and Joint, there is a negligible contribution in that frequency band. For the HCMP pipeline in certain voxels up to 4% of the power lies in frequencies that should ideally be 0. Panels (d–f) show the same results

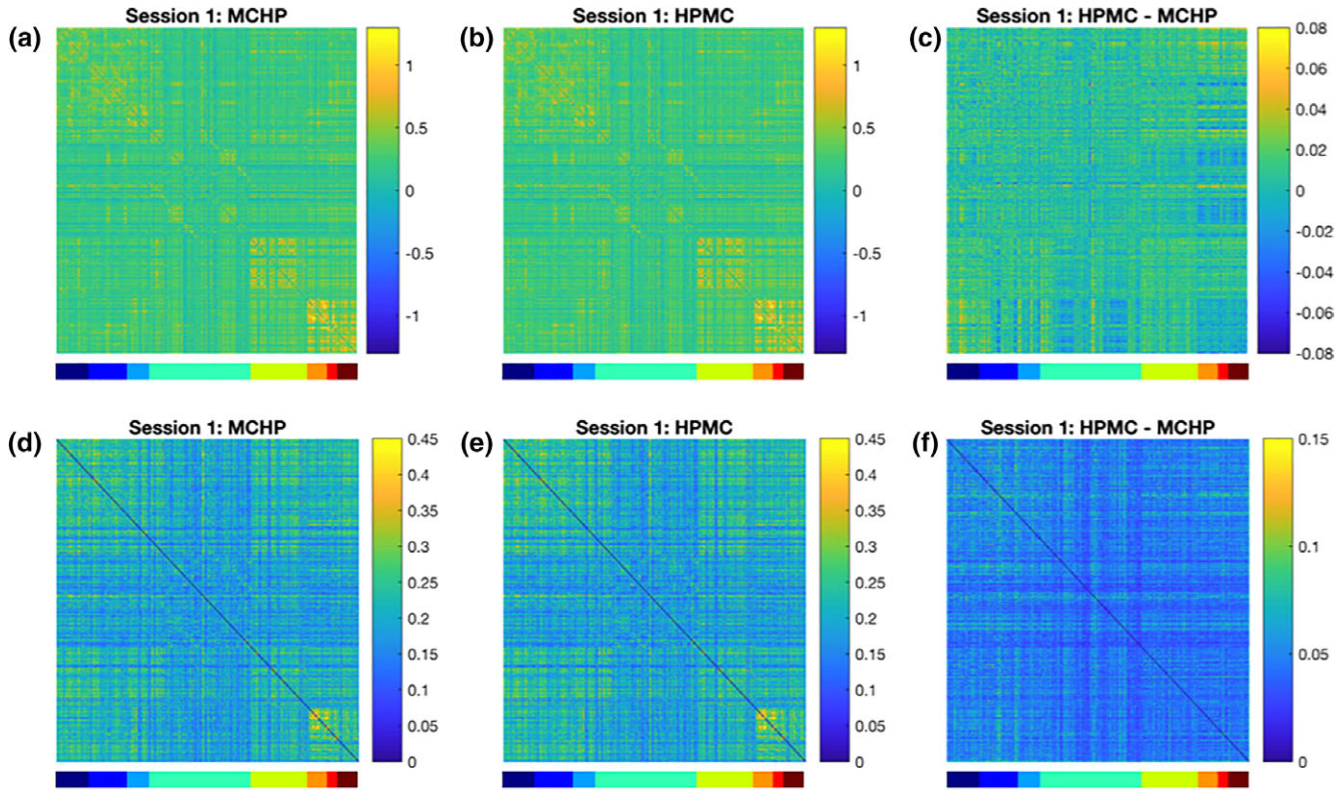

FIGURE 10 Results from the group-level analysis. (a) The average correlation matrix when performing motion regression followed by high-pass filtering (MCHP) for the 20 subjects in Session 1. (b) The average correlation matrix when performing high-pass filtering followed by motion regression (HPMC) for the 20 subjects in Session 1. (c) The difference between the two correlation matrices shown in (a) and (b). (d–f) The *SD* of the correlation matrices (MCHP, HPMC, HPMC-MCHP, respectively) across subjects [Color figure can be viewed at wileyonlinelibrary.com]

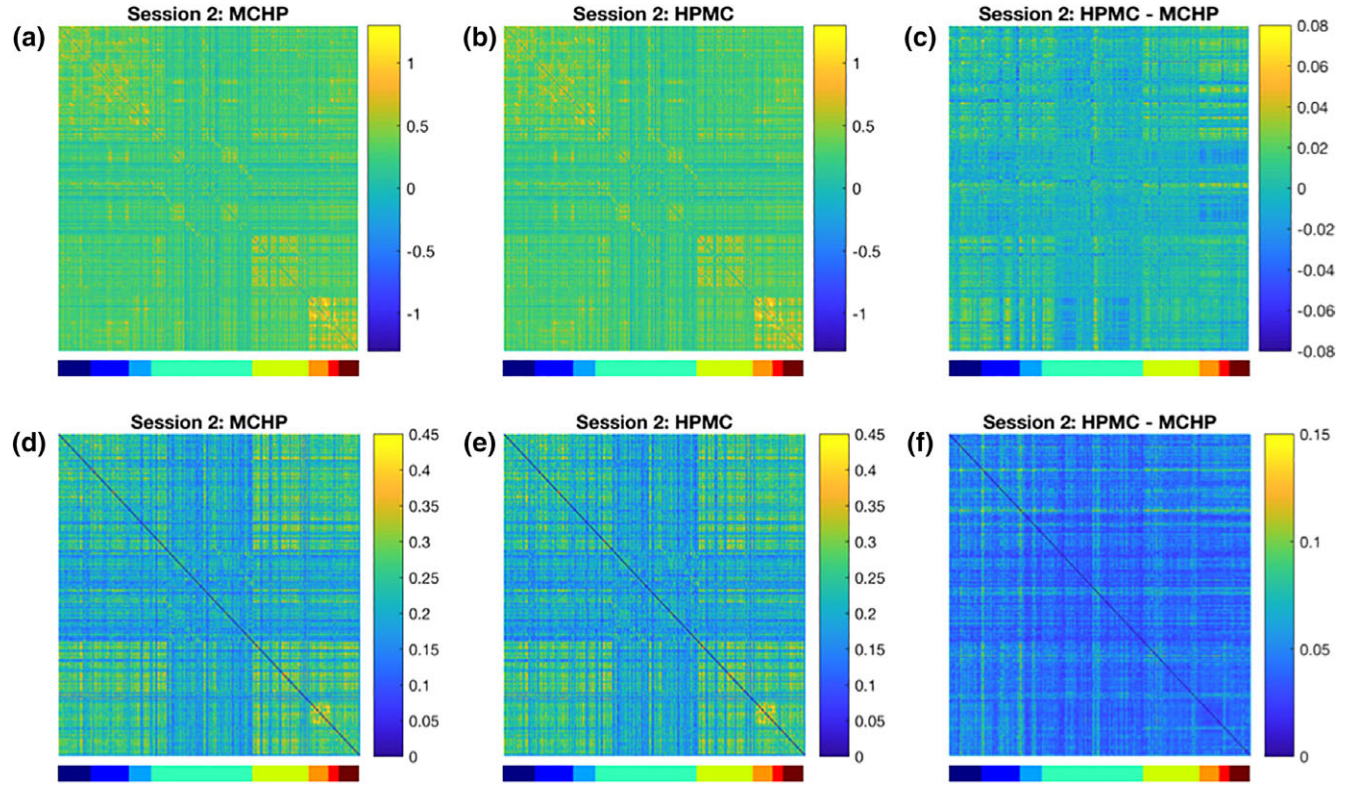

FIGURE 11 Results from the group-level analysis. (a) The average correlation matrix when performing motion regression followed by high-pass filtering (MCHP) for the 20 subjects in Session 1. (b) The average correlation matrix when performing high-pass filtering followed by motion regression (HPMC) for the 20 subjects in Session 2. (c) The difference between the two correlation matrices shown in (a) and (b). (d–f) The *SD* of the correlation matrices (MCHP, HPMC, HPMC-MCHP, respectively) across subjects [Color figure can be viewed at wileyonlinelibrary.com]

for Session 2. The results show clear similarities between the two sessions. Panels (g–f) show the HPMC results from the two sessions superimposed onto the brain. Note the results are presented using an arbitrary threshold of 0.02 for visualization purposes. As seen in previous figures, there are patterns consistent with motion artifacts, including a ring-like shape at the edge of the brain. There are also apparent patterns in the frontal cortex.

Panels (a–c) of Figure 13 show the group-level average coherence between the signal and motion time series in the frequency band below 0.01 Hz at each voxel using the three pipelines and data from Session 1. Note the coherence for the signal obtained through the HPMC pipeline is higher than the other preprocessing pipelines, illustrating that the reintroduced signal is related to motion. Panels (d–f) show the same results for data from Session 2. Panels (g) and (h) show the HPMC results superimposed onto the brain.

Finally, Figure 14 shows similar results related to the group-level average coherence between the signal and motion time series in the frequency band above 0.01 Hz at each voxel using the three pipelines for data from Session 1. Here, the differences between pipelines are more subtle with the MCHP pipeline showing slightly higher values than the other two pipelines. Panels (d–f) show the same results for data from Session 2. Panels (g) and (h) shows the MCHP results superimposed onto the brain.

Together, these results illustrate two important points. First, both the reintroduction of motion and unwanted frequency components are not necessarily canceled out across subjects and can propagate to the group level. Second, the effects are replicable across test–retest data. This shows that improper preprocessing can introduce a systematic error that can bias results in a manner that are replicable, and thereby potentially confused as real effects.

5 | DISCUSSION

It is common practice in the field of rs-fMRI for researchers to piece together modular preprocessing pipelines consisting of a number of separately developed algorithms; each designed to remove a specific type or class of artifacts. These artifacts can be related to scanner drift, signal spikes, motion, and signal fluctuations due to heart rate and respiration. In this article we offer a critic of this modular preprocessing approach, and argue that using such a pipeline can potentially have adverse effects on subsequent statistical analysis (e.g., resting-state functional connectivity).

While all preprocessing steps performed on fMRI data are important, we argue there needs to be a clear understanding about the effects they have on both the spatial and temporal correlation structure. More generally, it is critical to study the interactions among the individual preprocessing steps. In this work, we have proposed a geometrical framework that provides a way to better understand these interactions. The framework can be used to evaluate any preprocessing step that can be expressed within a linear model framework; examples of such include motion regression, scrubbing, removal of WM and CSF signal, RETROICOR, RVHRCOR,

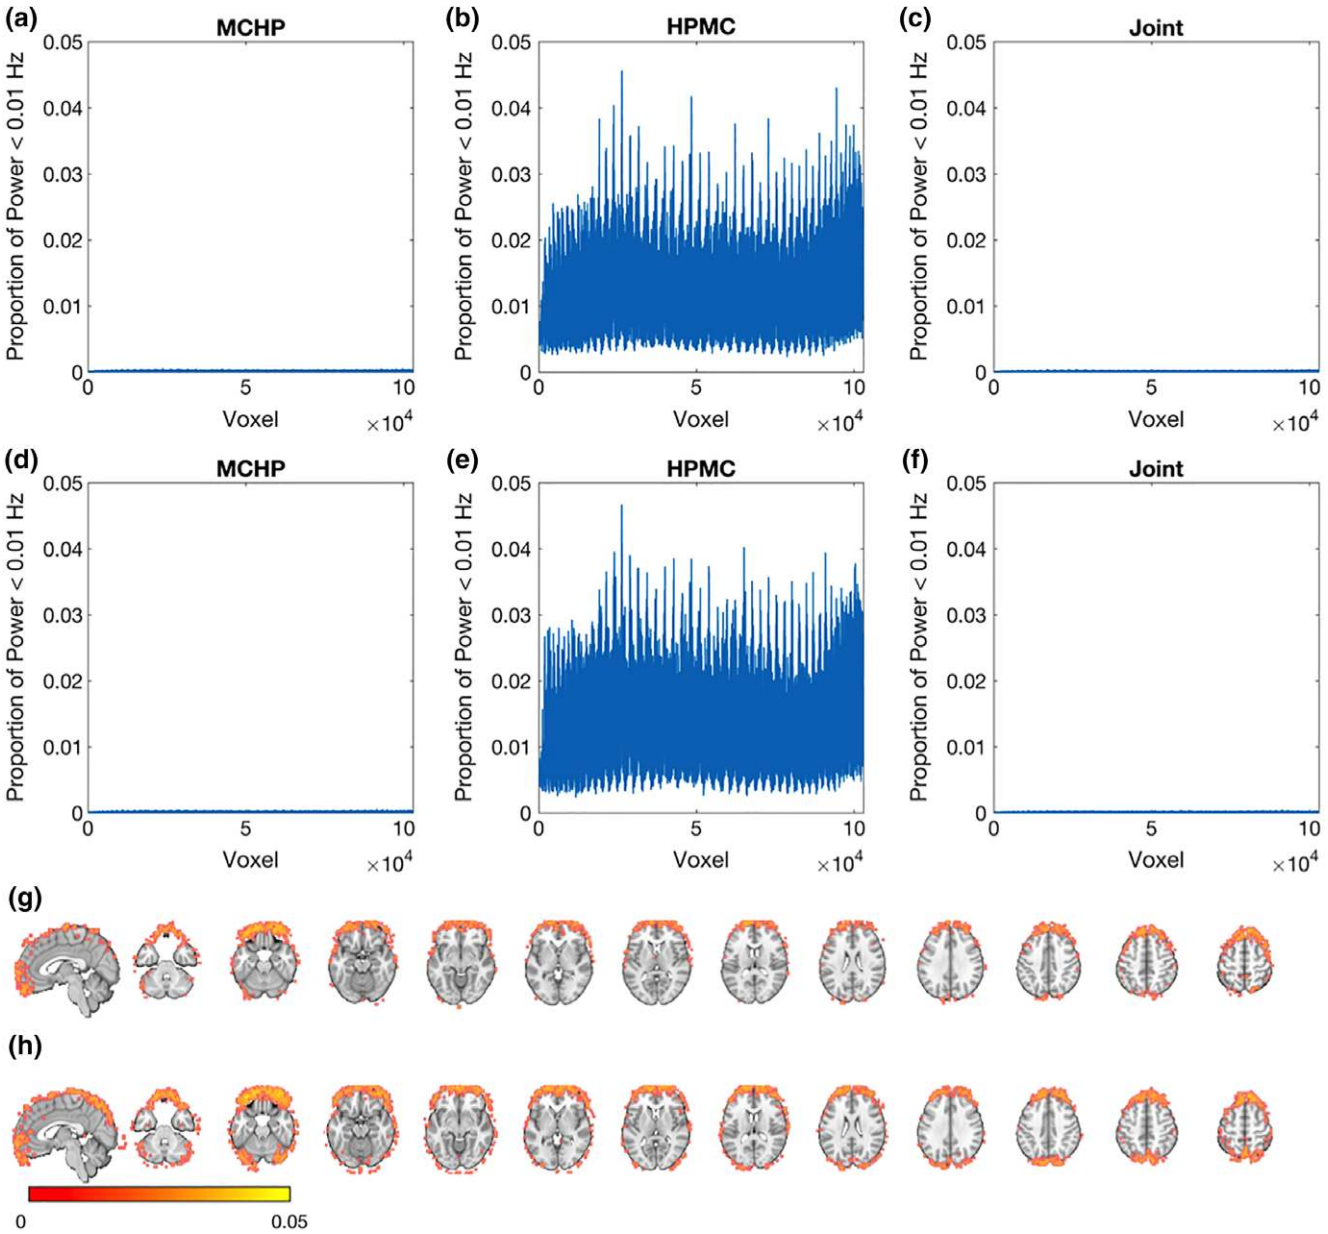

FIGURE 12 Results from the group-level analysis. (a) The group-level average proportion of the power that lies in the frequency band below 0.01 *Hz* at each voxel after motion regression followed by high-pass filtering (MCHP) using data from Session 1. Note there is a negligible contribution. (b) Same plot for data where motion regression is performed after high-pass filtering (HPMC). In certain voxels up to 4% of the power lies in frequencies that should ideally be 0. (c) Same plot for data processed jointly (Joint). Much like in (a) there a negligible contribution. (d–f) Same results for data from Session 2. (g) The HPMC results shown in (b) superimposed onto the brain. (h) The HPMC results shown in (e) superimposed onto the brain. Note for both (g) and (h), the results are presented using an arbitrary threshold of 0.02 for visualization purposes [Color figure can be viewed at wileyonlinelibrary.com]

ICA-AROMA global signal removal, and temporal filtering. Together these account for a large segment of the set of possible preprocessing steps applied to rs-fMRI data.

Using our approach, we were able to illustrate how preprocessing steps performed at a later stage of the pipeline can potentially reintroduce artifacts that had previously been removed from the data in an earlier step. Potential interactions between preprocessing steps are explored by computing the product of their respective projection matrices and evaluating whether or not when applied together they project the data into a subspace spanned by the various nuisance components. If this is the case, these nuisance components will be effectively reintroduced into the signal in a manner consistent with the order in which they were performed. Hence, the order in which the steps are organized in the preprocessing pipeline is critical. As a rule-of-thumb, the nuisance components corresponding to the last preprocessing step performed should be adequately removed from the signal, with components corresponding to earlier steps potentially reintroduced if they are not orthogonal to components removed in subsequent steps. We illustrate these issues both theoretically and using test–retest fMRI data. Empirically, we find that the reintroduced artifacts are consistent across sessions, and can potentially influence findings at both the single subject and group level.

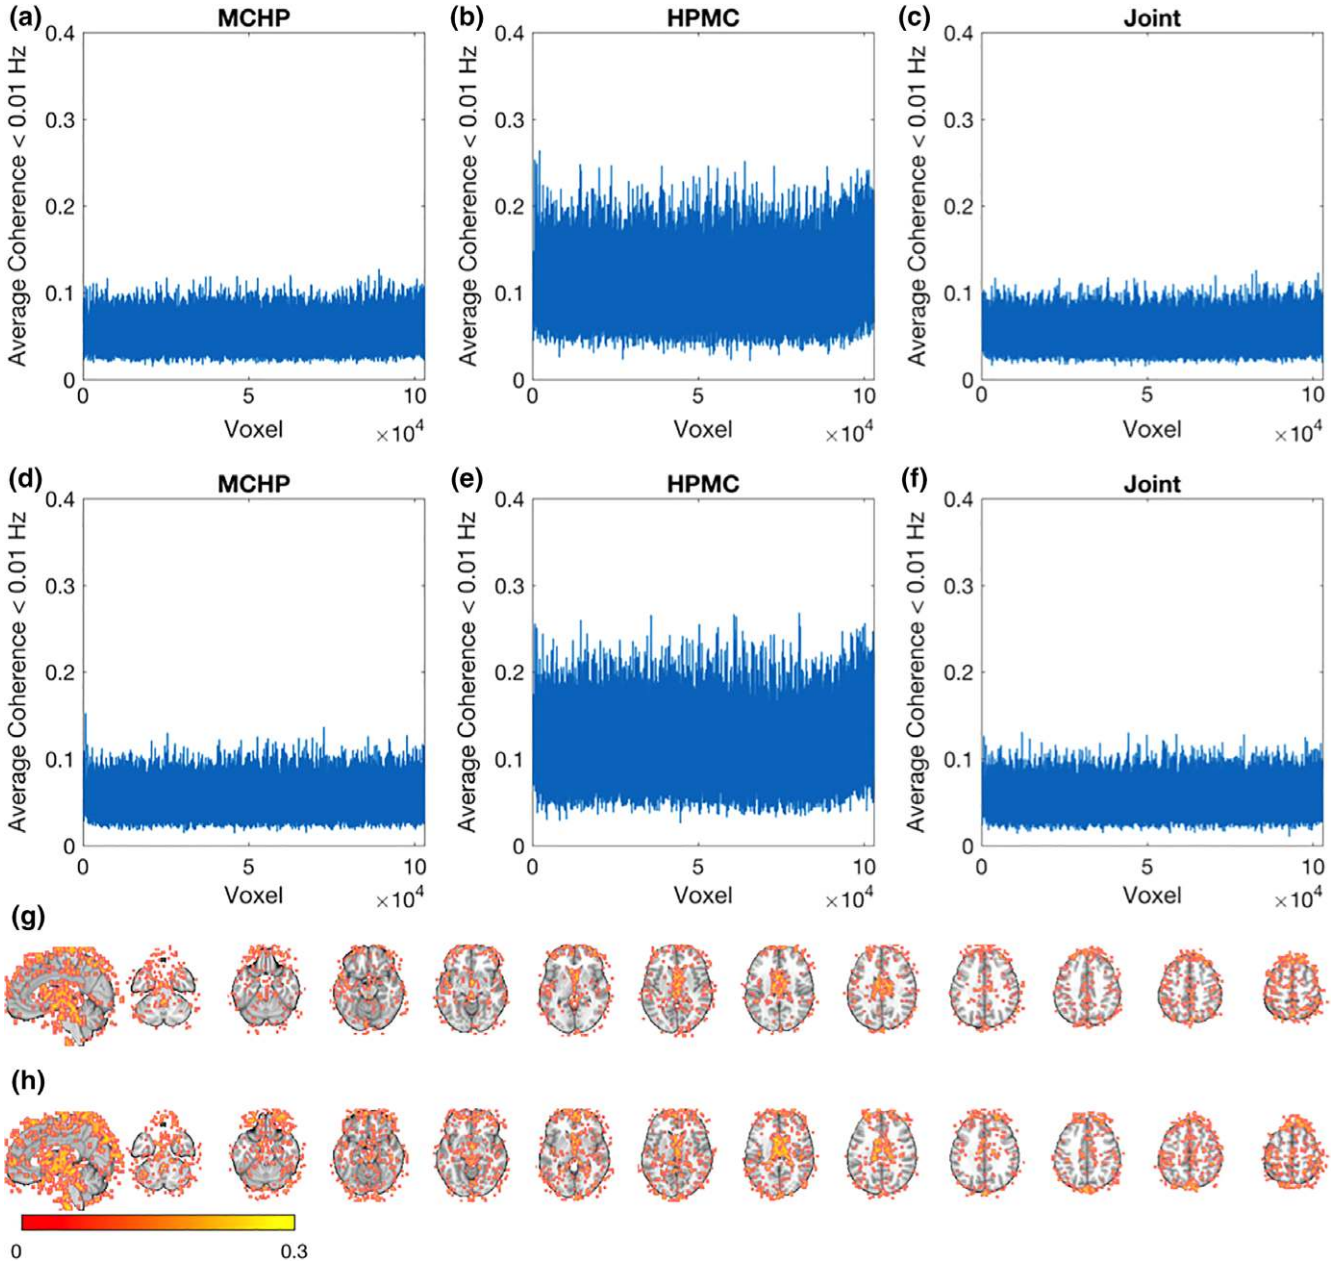

FIGURE 13 Results from the group-level analysis. (a) The group-level average coherence between the signal and motion time series in the frequency band below 0.01 Hz at each voxel after motion regression followed by high-pass filtering (MCHP) for data from Session 1. (b) Same plot for data where motion regression is performed after high-pass filtering (HPMC). (c) Same plot for data processed jointly (Joint). Note the coherence for the signal obtained through the HPMC pipeline (a) is heightened compared to the other preprocessing pipelines (b) and (c). (d–f) Same results for data from Session 2. (g) The HPMC results shown in (b) superimposed onto the brain. (h) The HPMC results shown in (e) superimposed onto the brain. Note for both (g) and (h), the results are presented using an arbitrary threshold of 0.15 for visualization purposes [Color figure can be viewed at wileyonlinelibrary.com]

We note that the issues discussed in this article can be circumvented in two ways. The first approach would be to abandon the modular approach to preprocessing the data, and instead use a joint approach that simultaneously performs the different preprocessing steps within an omnibus framework. For example, it is relatively straightforward to formulate a single linear model that simultaneously performs motion regression, nuisance regression, and temporal filtering. This is an approach advocated by Caballero-Gaudes and Reynolds (2017), and implemented in Analysis of Functional NeuroImages (AFNI) (Cox, 1996) 3dTproject. In general, we believe that the development of models that incorporate multiple preprocessing steps promises to play an important role in the future (Lindquist et al., 2008). The second approach is to formulate the design matrices used in each preprocessing step in such a manner that when applied sequentially they are constrained to project onto orthogonal subspaces. For example, if motion regression is performed after temporal filtering, the columns of the design matrix used in the motion regression should also be temporally filtered in a similar manner to ensure the data are projected onto an orthogonal subspace. Similarly, if temporal filtering is performed after motion correction, then each column

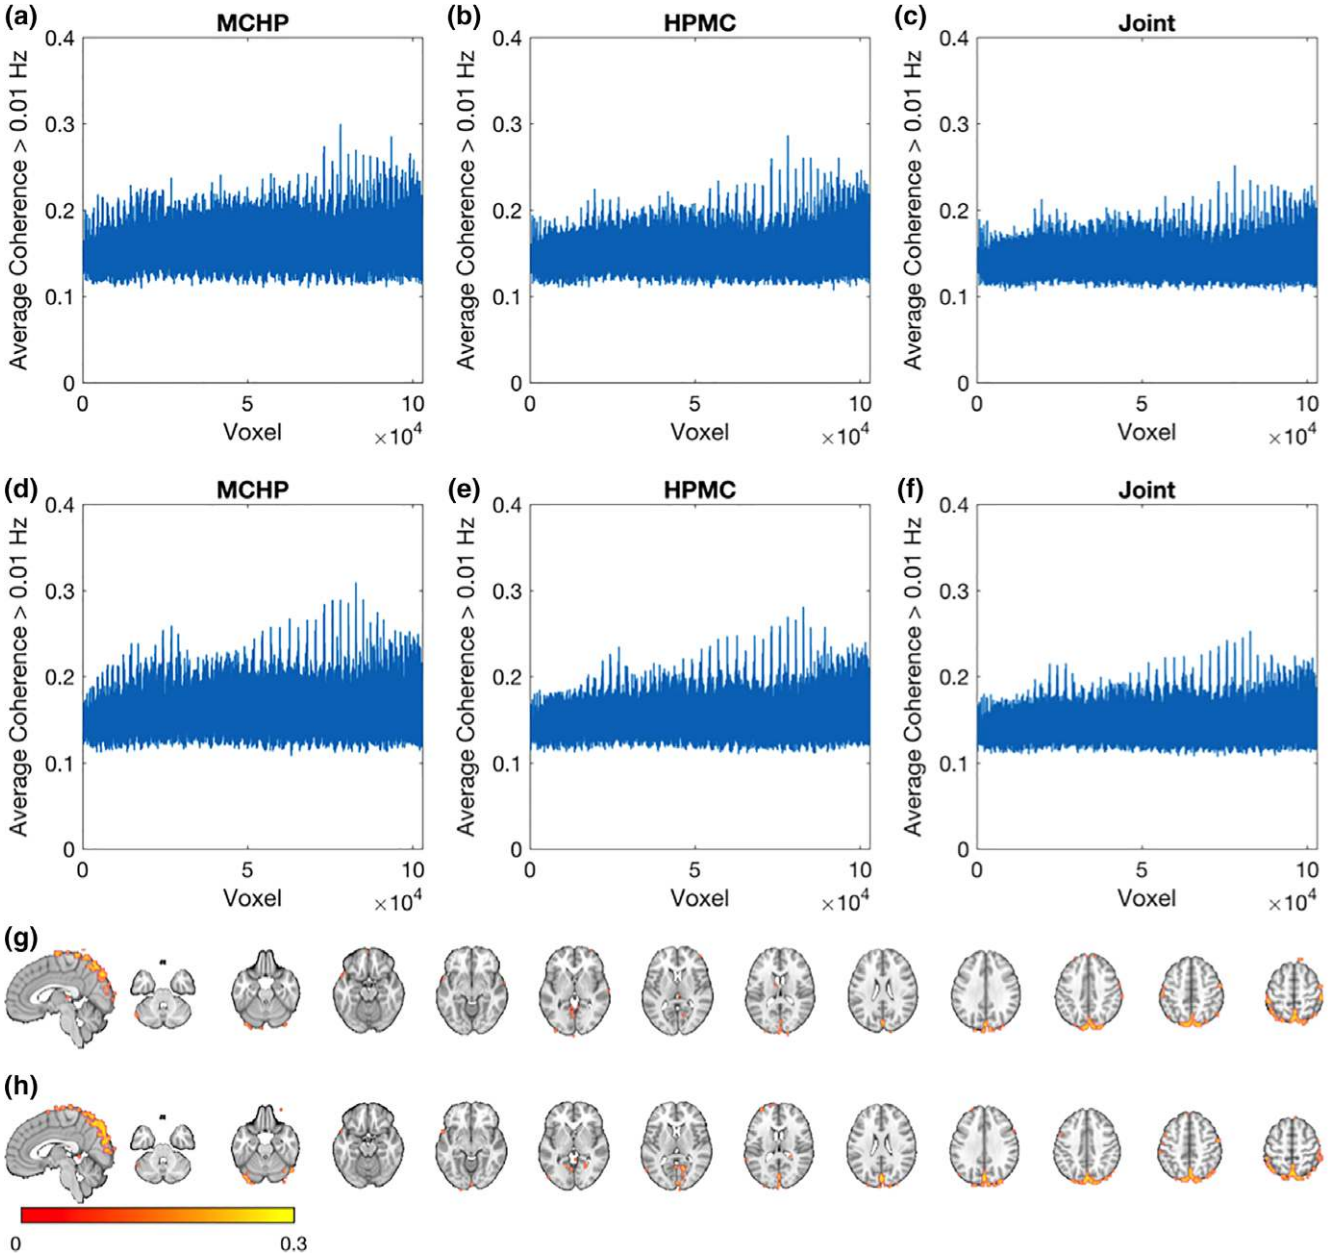

FIGURE 14 Results from the group-level analysis. (a) The group-level average coherence between the signal and motion time series in the frequency band above 0.01 Hz at each voxel after motion regression followed by high-pass filtering (MPHC) for data from Session 1. (b) Same plot for data where motion regression is performed after high-pass filtering (HPMC). (c) Same plot for data processed jointly (Joint). (d–f) Same results for data from Session 2. (g) The MCHP results shown in (a) superimposed onto the brain. (h) The MCHP results shown in (d) superimposed onto the brain. Note for both (g) and (h), the results are presented using an arbitrary threshold of 0.15 for visualization purposes [Color figure can be viewed at wileyonlinelibrary.com]

of the filter needs to be motion corrected. The general rule-of-thumb is that one must orthogonalize both the data, and all subsequent projections to maintain data orthogonality with the current projection. Practically, it makes sense to carefully identify all regressors that one seeks to remove a priori to ensure that they are handled in an appropriate manner.

In the context of temporal filtering we note that though the inclusion of sine and cosine basis elements in the design matrix (as described in this article) is elegant for its simultaneity, it may be suboptimal in various ways related to ringing artifacts and covariance stationarity (Christiano & Fitzgerald, 2003). While this does not undermine the points we seek to make about the value of sequential orthogonalization, we urge caution in encouraging readers to value a tidy design matrix over a filter that may be better suited for fMRI data.

While temporal filtering and nuisance regression are typically performed after other standard preprocessing steps, thereby making it straightforward to implement a joint filter on the data that includes both, other nuisance variables are often handled at other points in the pipeline. For example, this is true for measured physiological signals (e.g., using the RETROICOR method) and data-derived nuisance components (e.g., using ICA-AROMA). For measured physiological signal, removal is often performed in the early stages of the pipeline. For

example, Jones et al. (2008) suggest performing RETROICOR between the motion correction and slice-time correction steps. For ICA-AROMA, these components are derived after spatial smoothing, and removed using a partial regression approach in which voxels are regressed on both putative signal and noise components to avoid removal of parts of nuisance components that correlate with signal. This is generally performed prior to nuisance regression, which includes WM, CSF, and linear trend removal, as well as high-pass filtering (Pruim et al., 2015). In both these cases, latter preprocessing steps have the effect of potentially reintroducing the removed signal components. It is therefore critical that the effects of these components are removed from all subsequently performed preprocessing steps. To the best of our knowledge, this is currently not done.

Interestingly, the proposed approaches for circumventing the issues outlined in this article are often utilized when analyzing task fMRI using the general linear model (GLM). Here, it is common to perform motion regression and temporal filtering simultaneously by including both terms in the design matrix. If these steps are instead performed modularly (i.e., with low-pass filtering performed prior to fitting the GLM), the same issues described in this article will impact the results unless the design matrix used in the GLM has also been temporally filtered. Taking this step ensures that the frequencies removed in the low-pass filtering are not reintroduced when fitting the GLM. This is actually the default in most common fMRI software packages (e.g., SPM and FMRIB Software Library (FSL)). Hence, it is not clear if the issues raised in this article are problematic for most task fMRI studies, and therefore we focus on rs-fMRI, though we urge researchers to be cautious when performing modular analysis to ensure that artifacts are not reintroduced.

It should be noted that our framework provides a mechanism for investigating various preprocessing choices made in the literature. For example, Bright et al. (2017) recommend using prewhitening when performing motion regression. Prewhitening is equivalent to computing the generalized least squares solution

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{X}'\mathbf{V}^{-1}\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{V}^{-1}\mathbf{y}.$$

Thus, to compute the motion corrected residuals, one would use the projection matrix

$$\mathbf{P} = \mathbf{I} - \mathbf{X}\left(\mathbf{X}'\mathbf{V}^{-1}\mathbf{X}\right)^{-1}\mathbf{X}\mathbf{V}^{-1}.$$

Importantly, the residuals in this setting reside on the subspace orthogonal to the space spanned by $\mathbf{V}^{-1}\mathbf{X}$. Therefore, they will not necessarily be orthogonal to the space spanned by $\mathbf{X}$ (i.e., the space spanned by the motion parameters). Thus, the residuals will potentially still be correlated with motion if it is removed in this manner. This example illustrates the power of the geometrical framework to critically evaluate preprocessing choices made it the literature.

The complete preprocessing of rs-fMRI data can conveniently be separated into spatial and temporal preprocessing pipelines (Smith et al., 2013). The goal of spatial preprocessing is to remove spatial artifacts from the data. Here, the relevant steps include correction for spatial distortions caused by gradient nonlinearity, rigid-body correction for motion, correction for B_0 distortion, co-registration of structural and functional data, and normalization to a standard template (i.e., MNI space). These steps are performed prior to temporal preprocessing, which includes motion regression, scrubbing, removal of WM and CSF signal, global signal removal, and temporal filtering. In this work, we have focused entirely on potential problems with temporal preprocessing, and our geometrical framework is not explicitly designed for evaluating spatial preprocessing pipelines. Future work involves extending the framework, or alternatively developing a companion framework, for working in this setting.

It appears to be a trend for large data consortium (e.g., the Human Connectome Project) to make minimally preprocessed data available (Glasser et al., 2013). These data have only undergone spatial preprocessing, and should thus not be effected by the problems outlined in this article. However, we note that once researchers take these data sets and perform temporal preprocessing, they are susceptible to the problems outlined herein.

It is difficult to provide an exact estimate of the number of rs-fMRI papers that suffer from the issues raised in this work. This is in large part due to the wide variability in the reporting of preprocessing methods used in rs-fMRI studies (Waheed et al., 2016). In general, methods sections are often extremely terse when discussing preprocessing, and neither specify the order or the manner in which the different steps are performed. In addition, even if the order is specified it is not always clear whether proper orthogonalization has been performed. That said we suspect that the problems outlined in this article negatively impact a large number of papers published every year.

ACKNOWLEDGMENTS

The work presented in this article was supported in part by NIH grants R01 EB016061 and R01 EB026549 from the National Institute of Biomedical Imaging and Bioengineering.

ORCID

Martin A. Lindquist https://orcid.org/0000-0003-2289-0828

REFERENCES

Ashburner, J., & Friston, K. J. (2005). Unified segmentation. *NeuroImage*, *26*(3), 839–851.

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *360*(1457), 1001–1013.

Behzadi, Y., Restom, K., Liau, J., & Liu, T. T. (2007). A component based noise correction method (CompCor) for bold and perfusion based fMRI. *NeuroImage*, *37*(1), 90–101.

Birn, R. M., Diamond, J. B., Smith, M. A., & Bandettini, P. A. (2006). Separating respiratory-variation-related fluctuations from neuronal-activity-related fluctuations in fMRI. *NeuroImage*, *31*(4), 1536–1548.

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

Bright, M. G., Tench, C. R., & Murphy, K. (2017). Potential pitfalls when denoising resting state fMRI data using nuisance regression. *NeuroImage*, *154*, 159–168.

Caballero-Gaudes, C., & Reynolds, R. C. (2017). Methods for cleaning the bold fMRI signal. *NeuroImage*, *154*, 128–149.

Carp, J. (2012). The secret lives of experiments: Methods reporting in the fMRI literature. *NeuroImage*, *63*(1), 289–300.

Chang, C., Cunningham, J. P., & Glover, G. H. (2009). Influence of heart rate on the bold signal: The cardiac response function. *NeuroImage*, *44* (3), 857–869.

Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, *44*(2), 435–465.

Churchill, N. W., Spring, R., Afshin-Pour, B., Dong, F., & Strother, S. C. (2015). An automated, adaptive framework for optimizing preprocessing pipelines in task-based functional MRI. *PLoS One*, *10*(7), e0131520.

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

Cox, R. W. (1996). AFNI: Software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research*, *29*(3), 162–173.

De Luca, M., Beckmann, C., De Stefano, N., Matthews, P., & Smith, S. M. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *NeuroImage*, *29*(4), 1359–1367.

Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., … Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, *18*(11), 1664–1671.

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

Frackowiak, R. S. (2004). *Human brain function*. London: Elsevier.

Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., … WU-Minn HCP Consortium. (2013). The minimal preprocessing pipelines for the human connectome project. *NeuroImage*, *80*, 105–124.

Glover, G. H., Li, T.-Q., & Ress, D. (2000). Image-based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. *Magnetic Resonance in Medicine*, *44*(1), 162–167.

Hallquist, M. N., Hwang, K., & Luna, B. (2013). The nuisance of nuisance regression: Spectral misspecification in a common approach to resting-state fMRI preprocessing reintroduces noise and obscures functional connectivity. *NeuroImage*, *82*, 208–225.

Jones, T. B., Bandettini, P. A., & Birn, R. M. (2008). Integration of motion correction and physiological noise regression in fMRI. *NeuroImage*, *42* (2), 582–590.

Kay, K., Rokem, A., Winawer, J., Dougherty, R., & Wandell, B. (2013). GLMdenoise: A fast, automated technique for denoising task-based fMRI data. *Frontiers in Neuroscience*, *7*, 247.

Landman, B. A., Huang, A. J., Gifford, A., Vikram, D. S., Lim, I. A. L., Farrell, J. A., … Smith, S. A. (2011). Multi-parametric neuroimaging reproducibility: A 3-T resource study. *NeuroImage*, *54*(4), 2854–2866.

Lemieux, L., Salek-Haddadi, A., Lund, T. E., Laufs, H., & Carmichael, D. (2007). Modelling large motion events in fMRI studies of patients with epilepsy. *Magnetic Resonance Imaging*, *25*(6), 894–901.

Lindquist, M. A. (2008). The statistical analysis of fMRI data. *Statistical Science*, *23*(4), 439–464.

Murphy, K., Birn, R. M., & Bandettini, P. A. (2013). Resting-state fMRI confounds and cleanup. *NeuroImage*, *80*, 349–359.

Muschelli, J., Nebel, M. B., Caffo, B. S., Barber, A. D., Pekar, J. J., & Mostofsky, S. H. (2014). Reduction of motion-related artifacts in resting state fMRI using aCompCor. *NeuroImage*, *96*, 22–35.

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

Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., & Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artifact in resting state fMRI. *NeuroImage*, *84*, 320–341.

Pruessmann, K. P., Weiger, M., Scheidegger, M. B., Boesiger, P. (1999). Sense: Sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, *42*(5), 952–962.

Pruim, R. H., Mennes, M., van Rooij, D., Llera, A., Buitelaar, J. K., & Beckmann, C. F. (2015). ICA-AROMA: A robust ICA-based strategy for removing motion artifacts from fMRI data. *NeuroImage*, *112*, 267–277.

Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., & Smith, S. M. (2014). Automatic denoising of functional MRI data: Combining independent component analysis and hierarchical fusion of classifiers. *NeuroImage*, *90*, 449–468.

Satterthwaite, T. D., Elliott, M. A., Gerraty, R. T., Ruparel, K., Loughead, J., Calkins, M. E., … Wolf, D. H. (2013). An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage*, *64*, 240–256.

Shen, X., Tokoglu, F., Papademetris, X., & Constable, R. T. (2013). Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *NeuroImage*, *82*, 403–415.

Shirer, W. R., Jiang, H., Price, C. M., Ng, B., & Greicius, M. D. (2015). Optimization of rs-fMRI pre-processing for enhanced signal-noise separation, test-retest reliability, and group discrimination. *NeuroImage*, *117*, 67–79.

Shmueli, K., van Gelderen, P., de Zwart, J. A., Horovitz, S. G., Fukunaga, M., Jansma, J. M., & Duyn, J. H. (2007). Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fMRI bold signal. *NeuroImage*, *38*(2), 306–320.

Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., … WU-Minn HCP Consortium. (2013). Resting-state fMRI in the human connectome project. *NeuroImage*, *80*, 144–168.

Stehling, M. K., Turner, R., & Mansfield, P. (1991). Echo-planar imaging: Magnetic resonance imaging in a fraction of a second. *Science*, *254* (5028), 43–50.

Van Den Heuvel, M. P., & Pol, H. E. H. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, *20*(8), 519–534.

Van Dijk, K. R., Sabuncu, M. R., & Buckner, R. L. (2012). The influence of head motion on intrinsic functional connectivity MRI. *NeuroImage*, *59* (1), 431–438.

Waheed, S. H., Mirbagheri, S., Agarwal, S., Kamali, A., Yahyavi-Firouz-Abadi, N., Chaudhry, A., … Sair, H. I. (2016). Reporting of resting-state functional magnetic resonance imaging preprocessing methodologies. *Brain Connectivity*, *6*(9), 663–668.

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

Wise, R. G., Ide, K., Poulin, M. J., & Tracey, I. (2004). Resting fluctuations in arterial carbon dioxide induce significant low frequency variations in bold signal. *NeuroImage*, *21*(4), 1652–1664.

Yan, C., & Zang, Y. (2010). DPARSF: A MATLAB toolbox for "pipeline" data analysis of resting-state fMRI. *Frontiers in Systems Neuroscience*, *4*, 13.

Yeo, B., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., … Buckner, R. L. (2011). The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology*, *106*(3), 1125–1165.

Zou, Q.-H., Zhu, C.-Z., Yang, Y., Zuo, X.-N., Long, X.-Y., Cao, Q.-J., … Zang, Y.-F. (2008). An improved approach to detection of amplitude of low-frequency fluctuation (ALFF) for resting-state fMRI: Fractional ALFF. *Journal of Neuroscience Methods*, *172*(1), 137–141.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

How to cite this article: Lindquist MA, Geuter S, Wager TD, Caffo BS. Modular preprocessing pipelines can reintroduce artifacts into fMRI data. *Hum Brain Mapp*. 2019;40: 2358–2376. https://doi.org/10.1002/hbm.24528