



## OPEN ACCESS

## EDITED BY

Luisa Ciobanu,  
CEA Saclay, France

## REVIEWED BY

Tomokazu Tsurugizawa,  
National Institute of Advanced  
Industrial Science and Technology  
(AIST), Japan  
Philippe Ciuciu,  
Commissariat à l'Energie Atomique et  
aux Energies Alternatives (CEA), France

## \*CORRESPONDENCE

Amanda F. Mejia  
✉ [mandy.mejia@gmail.com](mailto:mandy.mejia@gmail.com)

## SPECIALTY SECTION

This article was submitted to  
Brain Imaging Methods,  
a section of the journal  
Frontiers in Neuroscience

RECEIVED 22 November 2022

ACCEPTED 09 December 2022

PUBLISHED 04 January 2023

## CITATION

Parlak F, Pham DD, Spencer DA,  
Welsh RC and Mejia AF (2023) Sources  
of residual autocorrelation in  
multiband task fMRI and strategies for  
effective mitigation.  
*Front. Neurosci.* 16:1051424.  
doi: 10.3389/fnins.2022.1051424

## COPYRIGHT

© 2023 Parlak, Pham, Spencer, Welsh  
and Mejia. This is an open-access  
article distributed under the terms of  
the [Creative Commons Attribution  
License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution  
or reproduction in other forums is  
permitted, provided the original  
author(s) and the copyright owner(s)  
are credited and that the original  
publication in this journal is cited, in  
accordance with accepted academic  
practice. No use, distribution or  
reproduction is permitted which does  
not comply with these terms.

# Sources of residual autocorrelation in multiband task fMRI and strategies for effective mitigation

Fatma Parlak<sup>1</sup>, Damon D. Pham<sup>1</sup>, Daniel A. Spencer<sup>1</sup>,  
Robert C. Welsh<sup>2</sup> and Amanda F. Mejia<sup>1\*</sup>

<sup>1</sup>Department of Statistics, Indiana University, Bloomington, IN, United States, <sup>2</sup>Department of Psychiatry and Bio-behavioral Sciences, University of California, Los Angeles, Los Angeles, CA, United States

**Introduction:** Analysis of task fMRI studies is typically based on using ordinary least squares within a voxel- or vertex-wise linear regression framework known as the general linear model. This use produces estimates and standard errors of the regression coefficients representing amplitudes of task-induced activations. To produce valid statistical inferences, several key statistical assumptions must be met, including that of independent residuals. Since task fMRI residuals often exhibit temporal autocorrelation, it is common practice to perform “prewhitening” to mitigate that dependence. Prewhitening involves estimating the residual correlation structure and then applying a filter to induce residual temporal independence. While theoretically straightforward, a major challenge in prewhitening for fMRI data is accurately estimating the residual autocorrelation at each voxel or vertex of the brain. Assuming a global model for autocorrelation, which is the default in several standard fMRI software tools, may under- or over-whiten in certain areas and produce differential false positive control across the brain. The increasing popularity of multiband acquisitions with faster temporal resolution increases the challenge of effective prewhitening because more complex models are required to accurately capture the strength and structure of autocorrelation. These issues are becoming more critical now because of a trend toward subject-level analysis and inference. In group-average or group-difference analyses, the within-subject residual correlation structure is accounted for implicitly, so inadequate prewhitening is of little real consequence. For individual subject inference, however, accurate prewhitening is crucial to avoid inflated or spatially variable false positive rates.

**Methods:** In this paper, we first thoroughly examine the patterns, sources and strength of residual autocorrelation in multiband task fMRI data. Second, we evaluate the ability of different autoregressive (AR) model-based prewhitening strategies to effectively mitigate autocorrelation and control false positives. We consider two main factors: the choice of AR model order and the level of spatial regularization of AR model coefficients, ranging from local smoothing to global averaging. We also consider determining the AR model order optimally at every vertex, but we do not observe an additional benefit of this over the use of higher-order AR models (e.g. AR(6)). To overcome the computational challenge associated with spatially variable prewhitening, we developed a computationally efficient R implementation using parallelization and fast C++ backend code. This implementation is included in the open source R package *BayesfMRI*.

**Results:** We find that residual autocorrelation exhibits marked spatial variance across the cortex and is influenced by many factors including the task being performed, the specific acquisition protocol, mis-modeling of the hemodynamic response function, unmodeled noise due to subject head motion, and systematic individual differences. We also find that local regularization is much more effective than global averaging at mitigating autocorrelation. While increasing the AR model order is also helpful, it has a lesser effect than allowing AR coefficients to vary spatially. We find that prewhitening with an AR(6) model with local regularization is effective at reducing or even eliminating autocorrelation and controlling false positives.

**Conclusion:** Our analysis revealed dramatic spatial differences in autocorrelation across the cortex. This spatial topology is unique to each session, being influenced by the task being performed, the acquisition technique, various modeling choices, and individual differences. If not accounted for, these differences will result in differential false positive control and power across the cortex and across subjects.

#### KEYWORDS

temporal autocorrelation, prewhitening, false positive control, task fMRI analysis, multiband acquisition, surface-based analysis

## 1. Introduction

The general linear model (GLM) has long been a popular framework for the analysis of task functional magnetic resonance imaging (fMRI) data. In the GLM, a linear regression model is used to relate the observed blood oxygenation level dependent (BOLD) signal to the expected BOLD response due to each task or stimulus in the experiment, along with nuisance regressors, yielding an estimate of the activation across the brain due to each task (Friston et al., 1994). Hypothesis testing with multiplicity correction is used to determine areas of the brain that are significantly activated in response to each task or contrast. One well-known issue with the GLM approach is that BOLD data generally violates the ordinary least squares (OLS) assumption of residual independence (Lindquist, 2008; Monti, 2011). When this happens, standard errors associated with the model coefficients are biased, invalidating inference and generally giving rise to inflated false positive rates for areas of activation.

Violations of residual independence are most consequential for subject-level inference (the “first-level GLM”), since OLS-based group inference (the “second-level GLM”) has been found to be relatively robust to dependent errors in the first level (Mumford and Nichols, 2009). While group-level analysis has historically been the norm in fMRI studies, more recently subject-level analysis is gaining in relevance. This in part due to the rise of “highly sampled” datasets collecting lots of data on individual subjects (Choe et al., 2015; Laumann et al., 2015;

Braga and Buckner, 2017; Gordon et al., 2017), as well as growing interest in using fMRI data for biomarker discovery, clinical translation, and other contexts where robust and reliable subject-level measures are required. Unfortunately, fMRI data presents several challenges for proper statistical analysis (Monti, 2011), and subject-level task fMRI measures have been shown to be unreliable (Elliott et al., 2020). One important factor for reliable subject-level task fMRI analysis is dealing appropriately with temporal dependence to avoid inflated rates of false positives.

Generalized least squares (GLS) is a regression framework that accounts for dependent and/or heteroskedastic errors. Briefly, in a regression model  $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ ,  $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \mathbf{V})$ , assume that the residual covariance matrix  $\mathbf{V}$  is known. Then the GLS solution is  $\hat{\boldsymbol{\beta}}^{GLS} = (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}\mathbf{X}'\mathbf{W}\mathbf{y}$ , where  $\mathbf{W} = \mathbf{V}^{-1}$ , and  $Var(\hat{\boldsymbol{\beta}}^{GLS}) = (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1}$ . This is mathematically equivalent to pre-multiplying both sides of the regression equation by  $\mathbf{V}^{-1/2}$ , which induces independent and homoskedastic residuals and gives rise to an OLS solution of the same form as  $\hat{\boldsymbol{\beta}}^{GLS}$ . In the context of task fMRI analysis, such “prewhitening” is a common remedy to eliminate temporal dependence as it produces the best linear unbiased estimate (BLUE) of model coefficients (Bullmore et al., 1996). A key challenge in GLS analysis is determining the form of  $\mathbf{V}$ , which is not actually known in practice. In a conventional statistical analysis, an iterative approach to estimating  $\mathbf{V}$  is commonly used through iteratively reweighted least squares (IRLS). However, in fMRI analysis this is typically considered computationally prohibitive

and prone to overfitting, so  $V$  is often estimated in a single pass based on the OLS residuals (Woolrich et al., 2001). To constrain the estimation of  $V$ , a parametric form is typically assumed based on the temporal structure of the data, such as an autoregressive (AR) or AR moving average (ARMA) model. It is also common practice to regularize the AR or ARMA model parameters by smoothing or averaging across the brain or within tissue boundaries.

Prewhitening methods are implemented in the major software packages AFNI (Cox, 1996), FSL (Jenkinson et al., 2012) and SPM (Penny et al., 2011). Yet, many of these standard prewhitening techniques have received criticism for failing to effectively remove residual autocorrelation (Worsley et al., 2002; Eklund et al., 2012). These criticisms have pointed to two main sources of mismodeled residual autocorrelation: (1) use of overly parsimonious autocorrelation models that fail to fully capture the autocorrelation in the data—an issue of ever-increasing relevance with the rise of faster multi-band acquisitions—and/or (2) assuming the same degree of residual autocorrelation across the brain. Olszowy et al. (2019) performed a systematic comparison of prewhitening techniques implemented in SPM, AFNI and FSL using several task and rest datasets of varying repetition time (TR) between 0.645 and 3s. They found that, while some techniques clearly performed better than others, all failed to control false positives at the nominal level, especially for low TR data. The best performance was seen using AFNI, which assumes a first-order autoregressive moving average (ARMA) model with unsmoothed, spatially varying coefficients, and the FAST option in SPM, which employs a flexible but global model using a dictionary of covariance components (Corbin et al., 2018). Interestingly, these two methods represent opposite approaches: AFNI allows for spatially varying autocorrelation but uses a relatively restrictive ARMA(1,1) model, while SPM FAST uses a quite flexible temporal correlation model but imposes a restrictive global assumption. Neither AFNI, FSL, or SPM currently offers a prewhitening technique that provides for a flexible *and* spatially varying autocorrelation model. Therefore, the ability to fully account for residual autocorrelation remains a limitation of many first-level task fMRI analyses.

Several recent studies have considered the ability of higher-order autoregressive models to adequately capture residual autocorrelation in fast TR fMRI data (Bollmann et al., 2018; Chen et al., 2019; Luo et al., 2020). Bollmann et al. (2018) found that optimal AR model order and AR coefficient magnitude varied markedly across the brain in fast TR task fMRI data, and that physiological noise modeling reduced but by no means eliminated the spatial variability in residual autocorrelation or the need for a high AR model order. Woolrich et al. (2004) examined the spatial dependence of autocorrelation in a Bayesian setting by utilizing AR models with a range of orders from 0 to 3 which is insufficient subsecond TR scans. Luo et al. (2020) used resting-state fMRI data of varying sub-second TRs to examine false positives rates with assumed task paradigms. They found that the optimal AR model order varied spatially and

depended on TR, with faster TR requiring a higher AR model order. They found that a too-low or too-high AR model order resulted in inflated false positive rates, and that a global model order (even with spatially varying coefficients) performed worse than when model order was allowed to vary spatially. Their approach of allowing the AR model order to vary spatially also outperformed both SPM FAST and the ARMA(1,1) model used by AFNI, the two methods found to have the best performance by Olszowy et al. (2019).

These recent studies also highlight several challenges associated with the use of volumetric fMRI in prewhitening. Both Bollmann et al. (2018) and Luo et al. (2020) observed sharp differences in the strength of residual autocorrelation across tissue classes, with cerebral spinal fluid (CSF) exhibiting much stronger autocorrelation than gray matter, and white matter exhibiting relatively low autocorrelation. Because of this, they point out that the standard practice of spatial smoothing (of the data, of the AR model order, or of the AR coefficient estimates) may be problematic at tissue class boundaries: gray matter bordering CSF may have higher autocorrelation due to mixing with CSF signals, while gray matter bordering white matter may have decreased autocorrelation due to mixing with white matter signals. Indeed, Luo et al. (2020) found smoothing of the sample autocorrelations at 6 mm FWHM to result in inflated false positive rates. Yet some regularization of autocorrelation model parameters is believed to be necessary to avoid very noisy estimates (Worsley et al., 2002; Bollmann et al., 2018; Chen et al., 2019), and data smoothing is nearly universal practice in the massive univariate framework, given its ability to enhance signal-to-noise ratio (SNR) and increase power to detect activations. This presents a dilemma: smoothing across tissue classes can be detrimental for autocorrelation modeling, but regularization of autocorrelation coefficients is needed to avoid overly noisy estimates.

The use of cortical surface fMRI (cs-fMRI) could mitigate this dilemma in two ways. First, geodesic smoothing along the surface can increase SNR without blurring across tissue classes or neighboring sulcal folds. Second, by eliminating white matter and CSF, the spatial variability in residual autocorrelation is simplified, since the most dramatic spatial differences have been observed between tissue classes (Penny et al., 2003, 2007). An additional potential benefit of the use of cs-fMRI is the utility of spatial Bayesian models, which cs-fMRI is uniquely suited for (Mejia et al., 2020), to spatially regularize autocorrelation coefficients in a statistically principled way. Therefore, in this work we adopt cortical surface format fMRI.

In this work, we advance prewhitening methods for modern fMRI acquisitions in three ways. First, we thoroughly examine the spatial variability and influence of various factors on residual autocorrelation, including the task protocol, the acquisition technique, and systematic individual differences. We also examine the influence of potential model mis-specification for the GLM. Un-modeled neural activity is temporally correlated and may be absorbed into the model residuals, thus increasing

residual autocorrelation (Penny et al., 2003; Lindquist et al., 2009; Bollmann et al., 2018). Task-related neural activity is assumed to be captured through task regressors, which are typically constructed by convolving a hemodynamic response function (HRF) with a stick function representing the task paradigm. However, the shape and duration of the HRF are known to vary across the brain (Lindquist and Wager, 2007) as well as within and across individuals (Aguirre et al., 1998). Furthermore, sensitivity of the canonical HRF within and between subjects is discussed in Badillo et al. (2013). For example, the canonical HRF is found to be more suitable choice for modeling motor cortex regions while the choice of more flexible HRF modeling improved the sensitivity on parietal regions. Therefore, assuming a fixed “canonical” HRF may fail to adequately capture task-induced activity (Lindquist et al., 2009). More flexible models can capture differences in HRF height, width, and time to peak, including the model with temporal and dispersion derivatives (Friston et al., 1998b), the finite impulse response model (Glover, 1999), and the inverse logit model (Lindquist and Wager, 2007). Here, we consider the effect of including the temporal and/or dispersion derivatives of the HRF on autocorrelation. Another potential source of model misspecification is failure to account for noise resulting from head motion, scanner drift, and other sources. If such noise is not modeled, it will be reflected in the model residuals. Because such sources of noise tend to be temporally correlated, this will tend have the effect of increasing residual autocorrelation. Here, we consider the effect of including more or fewer head motion-based regressors on residual autocorrelation.

Second, we investigated the autocorrelation of GLM residuals by utilizing voxel specific AR models as suggested by Penny et al. (2003) and Luo et al. (2020). This investigation was conducted under two dependent conditions: (1) model order and (2) surface smoothing. We evaluate the effectiveness of different autoregressive (AR) model-based prewhitening strategies at reducing autocorrelation and controlling false positive rates. We consider AR model order varying from 1 to 6, as well as determining the AR model order optimally at each vertex. We also consider local regularization of AR model coefficients vs. global averaging. We find that local surface-based regularization of AR model coefficients is much more effective than a global prewhitening strategy at eliminating autocorrelation across the cortex.

Third, we overcome the major computational challenges associated with spatially-varying prewhitening. We have developed a computationally efficient implementation of the AR-based prewhitening techniques considered here. Using parallelization and backend code written in C++, we are able to perform spatially varying prewhitening very efficiently for surface-based analysis and “grayordinates” analysis more generally. This implementation is available in the open-source `BayesfMRI` R package (Mejia et al., 2022).

The remainder of this paper is organized as follows. In Section 2 we describe the data, the GLM approach, and the

methods for autocorrelation estimation, and prewhitening. We also describe a mixed effect modeling framework we use to assess the influence of several key factors on the strength and spatial variability of residual autocorrelation, including acquisition method, task protocol, modeling choices, and individual variability. In Section 3 we present results based on an analysis of several task and resting state fMRI studies from the Human Connectome Project, utilizing 40 subjects with test-retest data. In Section 4, we conclude with a discussion of these results and what they suggest for future research in prewhitening.

## 2. Materials and methods

### 2.1. Data collection and processing

The data used in this paper are from the Human Connectome Project (HCP) 1200-subject release (<http://humanconnectome.org>) collected at Washington University in St. Louis. The HCP includes task and resting-state fMRI data collected on a customized Siemens 3T Skyra scanner with a multiband factor of 8 to provide high spatial (2 mm isotropic voxels) and temporal ( $TR = 0.72s$ ) resolution. Functional MRI acquisition was performed using a simultaneous multi-slice gradient echo EPI sequence. Data were collected with a 32-channel head coil, with no in-plane acceleration and an echo-time of 33.1 ms (Van Essen et al., 2013). The fMRI data were processed according to the HCP minimal preprocessing pipelines including projection to the cortical surface, as described in Glasser et al. (2013). The resulting surface mesh for each hemisphere consists of 32,000 vertices. For all fMRI scans, we perform surface-based spatial smoothing using a 2-dimensional Gaussian kernel with 6mm full-width-at-half-maximum (FWHM). To reduce the computational burden of estimating the mixed-effects models described below, prior to smoothing, we resample to 6,000 vertices per hemisphere. Note that this level of resampling results in a much milder degree of interpolation than smoothing at 6 mm FWHM, and therefore results in a negligible loss of information when performed in combination with smoothing [see, e.g., Mejia et al. (2020) Figure A.1 in Appendix]. For both resampling and smoothing, we employ the Connectome Workbench (Marcus et al., 2011) via the `ciftiTools` R package (Pham et al., 2022).

Each subject underwent several task and resting-state fMRI protocols across two sessions. Each task and rest session was performed twice, using opposite phase encoding directions (LR and RL). For a subset of 45 participants, the entire imaging protocol was repeated. We analyze data from the 40 participants having a complete set of test and retest data for the protocols we analyze in this study. We analyze four task experiments, namely the emotion, gambling, motor, and relational tasks (Table 1). In the emotion processing task, developed by Hariri et al. (2002), participants are shown sets of faces or geometric shapes, and are

TABLE 1 Type and duration of each task protocol analyzed.

Task	Block or event	Frames per run	Run duration (mm:ss)
Emotion	Block	176	2:16
Gambling	Event	253	3:12
Motor	Block	284	3:34
Relational	Block	232	2:56

Adopted from Barch et al. (2013).

asked to determine which of two faces/shapes match a reference face/shape. Each face has an angry or fearful expression. A 3 s cue ("shape" or "face") precedes a block of 6 trials, lasting 18 s in total. Each run includes three blocks of each condition (shape or face).

In the gambling task, adopted from Delgado et al. (2000), participants play a game in which they are asked to guess the value of a mystery card to win or lose money. They indicate their guess for the value, which can range from 1 to 9, as being more or <5. Their response is evaluated by a program which predetermines whether the trial is a win, loss, or neutral event. In each run, there are two mostly win blocks (six win trials and two non-win trials), two mostly loss blocks (six loss trials and two non-loss trials), alternating with four 15 s fixation blocks. While this protocol can be considered a block design since the task protocol is comprised of short events rather than continuous blocks of stimulus, we analyze it as an event-related design.

In the motor task, developed by Buckner et al. (2011), participants are given a 3 s visual cue which instructs them to perform one of five motor actions: tap left or right fingers, wiggle left or right toes, or move tongue. Each action block lasts 12 s, and each run includes two blocks of each action as well as three 15 s fixation blocks.

In the relational task, developed by Smith Rachele and Kalina (2007), subjects undergo two conditions: relational processing and control matching. In the relational condition, one pair of objects is shown at the top of the screen and another pair is shown at the bottom of the screen. In this condition, participants are asked to determine the dimension (shape or texture) across which the pair displayed at the top differs. Next, they determine whether the bottom pair differs along the same dimension. During the matching condition, two objects are displayed at the top of the screen and one is shown at the bottom, and the word "shape" or "texture" appears in the middle of the screen. In this condition, participants are asked to determine whether the bottom object matches either of the top objects, based on the dimension displayed in the middle. Each condition is administered as blocks of trials of the same condition, with each block lasting 18 s total, with three blocks of each condition (relational, matching, and fixation block) in each run.

To quantify false positive rates, we also analyze resting-state fMRI data acquired for the same subjects and sessions, which we

analyze under a false task protocol. To emulate the duration of the task fMRI runs, we truncate the resting-state runs to have 284 volumes, the same number of volumes as the motor task, after dropping the first 15 rest volumes. The boxcar design consists of a single event with three "boxcars:" three periods of stimulus lasting 10 s each, with 10 s in between each consecutive stimulus. The first stimulus begins at the 20th second, or <28th volume. We use only three boxcars instead of extending the boxcars to the duration of the scan, in order to more closely resemble the number of stimuli in the HCP task scans. We used the same GLM model as with the task analysis, except we did not include any HRF derivatives since there is no true task-evoked signal to potentially mis-model.

## 2.2. Statistical analysis

Our analysis consists of three primary steps. First, for each subject, session, task, acquisition protocol, HRF modeling strategy, and motion regression strategy, we fit a vertex-wise general linear model (GLM) to estimate the amplitude of task-evoked activation assuming residual independence. Based on the fitted residuals, we estimate the degree of autocorrelation at every location in the brain. Second, we fit a series of mixed effects models to identify the effects of acquisition and modeling factors on residual autocorrelation across the brain, as well as systematic individual variability. Finally, we prewhiten the data using a range of strategies, varying the parametric model order, and spatial regularization level. We evaluate the ability of each prewhitening strategy to effectively mitigate autocorrelation and control false positives.

### 2.2.1. GLM estimation

We first fit a series of GLMs to each task fMRI dataset assuming residual independence in order to quantify residual autocorrelation and examine its patterns and sources. Let  $y_v$  be the BOLD response at vertex  $v$ , and let  $X$  be a design matrix containing an intercept, task-related regressors, and nuisance regressors. For each vertex  $v$ , the GLM proposed by Friston et al. (1994) can be written as:

$$y_v = X\beta_v + \epsilon_v, \quad \epsilon_v \sim MVN(\mathbf{0}, \Sigma_v). \quad (1)$$

$\Sigma_v$  ( $T \times T$ ) encodes the residual autocorrelation and variance, which may differ across the brain. If  $\Sigma_v \neq \sigma_v^2 \mathbf{I}$ , the OLS assumption of residual independence is violated, and a generalized least squares (GLS) approach is appropriate in place of OLS to improve estimation efficiency and to avoid invalid statistical inference. In some cases, such as in spatial Bayesian variants of the GLM where a single Bayesian linear model is fit to all vertices, spatially homogeneous variance may also be assumed, i.e.,  $\sigma_v^2 \equiv \sigma^2$ . GLS can be used

to satisfy this assumption by inducing unit variance across the brain. In GLS, OLS is first used to obtain an initial set of fitted residuals  $\hat{\epsilon}_v$ , which are utilized to estimate  $\Sigma_v$  (Kariya and Kurata, 2004). The GLS coefficient estimates are given by  $\hat{\beta}_v^{GLS} = (\mathbf{X}'\Sigma_v^{-1}\mathbf{X})^{-1}\mathbf{X}'\Sigma_v^{-1}\mathbf{y}$ , and their covariance is  $\text{Var}(\hat{\beta}_v^{GLS}) = (\mathbf{X}'\Sigma_v^{-1}\mathbf{X})^{-1}$ . Equivalently, prewhitening involves pre-multiplying both sides of the regression equation (1) by  $\mathbf{W}_v = \Sigma_v^{-1/2}$  to induce residual independence. Traditionally this process is repeated until convergence, but due to computational considerations a single iteration is often assumed to be sufficient for task fMRI analysis (Woolrich et al., 2001).

To avoid overly noisy estimates of  $\Sigma_v$ , restrictive parametric models (e.g., low-order autoregressive models) and/or aggressive regularization (e.g., averaging across all gray matter) are often used to estimate  $\Sigma_v$ . Recent work has suggested that these approaches generally fail to fully account for autocorrelation or control false positives as a nominal rate (Bollmann et al., 2018; Corbin et al., 2018; Chen et al., 2019; Olszowy et al., 2019; Luo et al., 2020). The challenge is how to produce a sufficiently efficient estimate of  $\Sigma_v$  while accurately representing the differential levels of autocorrelation across the brain. In Section 2.2.3, we consider various strategies for estimation of  $\mathbf{W}_v = \Sigma_v^{-1/2}$ . However, our first step is to examine the sources and patterns of residual autocorrelation, and therefore we do not impose any parametric model or regularization in estimating  $\Sigma_v$ . Instead, we use empirical autocorrelation function (ACF) at each vertex. Consider the timeseries of fitted OLS residuals  $\mathbf{e}_v = \mathbf{y}_v - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$  at vertex  $v$  for a particular fMRI dataset. The ACF of  $\mathbf{e}_v$  at lag  $u$  is defined as

$$\rho_{v,u} = \frac{\text{Cov}(e_{v,t}, e_{v,t+u})}{\sqrt{\text{Var}(e_{v,t})\text{Var}(e_{v,t+u})}},$$

for  $u = 0, \dots, T - 1$ , with lag-0 ACF  $\rho_{v,0} = 1$  (Venables and Ripley, 2013). We summarize the ACF  $\rho_{v,u}$  into a single metric of autocorrelation, the autocorrelation index (ACI) (Afyouni et al., 2019), which is given by

$$\tau(v) = \sum_{u=0}^{T-1} \rho_u^2(v).$$

We consider the effect of two potential sources of model misspecification on temporal autocorrelation: unmodeled neuronal activity *via* the task regressors in  $\mathbf{X}$ , and unmodeled head motion-induced noise *via* the nuisance regressors in  $\mathbf{X}$ . Since both neuronal activity and motion-induced noise exhibit temporal dependence, failing to adequately account for either may contribute to residual autocorrelation. The task regressors in  $\mathbf{X}$  are constructed by convolving a stimulus function representing the timing of the tasks or stimuli with a canonical HRF, which is typically modeled as a gamma function or a difference of two gamma functions (Worsley et al., 2002).

However, HRF onset and duration is known to vary across the brain and across individuals (Aguirre et al., 1998), so using a fixed HRF may fail to accurately capture the task-evoked BOLD signal (Loh et al., 2008; Lindquist et al., 2009). Therefore, we consider three models for the HRF: one assuming a fixed canonical HRF; one including the temporal derivative of the HRF to allow for differences in HRF onset timing; and one additionally including its dispersion derivative to allow for differences in HRF duration (Friston et al., 1998a; Lindquist et al., 2009). Regarding nuisance regressors, the inclusion of measures of head motion is a common practice to account for head motion-induced noise in the data. We therefore consider two sets of motion regressors: the six rigid body realignment parameters and their one-back differences (RP12) or those terms plus their squares (RP24). In all models, we include discrete cosine transform (DCT) bases to achieve high-pass filtering at 0.01 Hz, which is important to satisfy the stationarity assumption of AR-based prewhitening.

In sum, we estimate a vertex-wise GLM *via* OLS for each subject  $i = 1, \dots, 40$ , session  $j = 1, 2$ , task  $k = 1, 2, 3, 4$ , and phase encoding direction  $\in \{\text{LR}, \text{RL}\}$ . Each GLM is fit using the canonical HRF only, with its temporal derivative, and with its temporal and dispersion derivatives. Each GLM is also fit with 12 or 24 motion realignment parameters. In total, we fit 3,840 GLMs ( $40 \times 2 \times 4 \times 2 \times 3 \times 2$ ) before prewhitening. All models are fit using the `BayesfMRI` R package Mejia et al. (2022). In the next section, we describe the mixed effects modeling framework we use to disentangle the influence of each factor (e.g. subject effects vs. acquisition effects vs. modeling effects) on residual autocorrelation.

### 2.2.2. Examining sources of residual autocorrelation through mixed effects modeling

Let  $\tau_{ijk}^{hr\ell}(v)$  be the ACI at vertex  $v$  for subject  $i$ , session  $j$ , task  $k$ , phase encoding direction  $\ell$ , HRF modeling strategy  $h$  and motion regression strategy  $r$ . To determine the influences of population variability, spatial variability, and other factors on ACI, we fit a mixed effect model at each vertex. We include fixed effects for each task, for the interaction between task and HRF modeling strategy, and for the motion regression strategy. For each of these fixed effects, we also include a random effect to represent population heterogeneity. Finally, we include a fixed effect for phase encoding direction. The mixed effect model at vertex  $v$  is given by:

$$\begin{aligned} \tau_{ijk}^{hr\ell}(v) = & \underbrace{\{\alpha_k(v) + a_{k,i}(v)\}}_{\text{baseline effects}} + \underbrace{\{\beta_k(v) + b_{k,i}(v)\}}_{\text{HRF modeling effects}} h \\ & + \underbrace{\{\gamma(v) + g_i(v)\}}_{\text{motion regression effects}} r + \underbrace{\{\theta(v)\}}_{\text{acquisition effects}} \ell + \epsilon_{ijk}^{hr\ell}(v) \end{aligned} \quad (2)$$

$$\begin{aligned} \epsilon_{ijk}^{hr\ell}(\nu) &\stackrel{iid}{\sim} N(0, \sigma_\nu^2), \mathbf{b}_i(\nu) \\ &= \{a_{1,i}(\nu), \dots, a_{4,i}(\nu), b_{1,i}(\nu), \dots, b_{4,i}(\nu), g_i(\nu)\}' \stackrel{iid}{\sim} N(\mathbf{0}, \mathbf{G}(\nu)). \end{aligned}$$

The covariates  $h$ ,  $r$ , and  $\ell$  are constructed as dummy variables equalling zero for the “baseline” conditions (canonical HRF only, RP12 motion regression, and LR phase encoding direction acquisition) and equalling one for the alternative conditions (canonical HRF plus derivative(s), RP24 motion regression, and RL phase encoding direction acquisition). The model in equation (2) is estimated separately for  $h = 1$  representing the inclusion of HRF temporal derivatives or the HRF temporal and dispersion derivatives.

We perform model fitting using the `lmer` function from the `lme4` R package (version 1.1.-30) (Bates et al., 2015) to estimate each fixed effect [the  $\alpha_k(\nu)$ ,  $\beta_k(\nu)$ ,  $\gamma(\nu)$ , and  $\theta(\nu)$ ], the error variance, and the covariance of all the random effects [the  $a_{k,i}(\nu)$ ,  $b_{k,i}(\nu)$ , and  $g_i(\nu)$ ]. The fixed effect for task,  $\alpha_k(\nu)$ , represents the baseline autocorrelation for task protocol  $k$  for the model including the canonical HRF only ( $h = 0$ ). The corresponding random subject effect,  $a_{k,i}(\nu)$ , represents the difference in autocorrelation for subject  $i$ , vs. the average over subjects for task  $k$ . The fixed effect for HRF modeling,  $\beta_k(\nu)$ , represents the change in autocorrelation when the temporal derivative of the HRF is included ( $h = 1$ ) for task  $k$ , while  $b_{k,i}(\nu)$  represents the random variation in that change over subjects. We generally expect negative values for  $\beta_k(\nu)$ , representing a reduction in residual autocorrelation when the HRF derivative is included, since discrepancies between the true HRF and the canonical HRF tend to exhibit temporal dependence. The fixed effect for motion regression strategy,  $\gamma(\nu)$ , represents the change in autocorrelation associated with the use of RP24 ( $r = 1$ ) vs. RP12 ( $r = 0$ ) motion regression. The corresponding random effect,  $g_i(\nu)$ , represents random variation in that effect over subjects.  $\theta(\nu)$  represents the difference in autocorrelation when using phase encoding direction *RL* ( $\ell = 1$ ), compared with phase encoding direction *LR* ( $\ell = 0$ ).

Since the model includes multiple sessions from each subject, the random effects  $a_{k,i}(\nu)$ ,  $b_{k,i}(\nu)$ , and  $g_i(\nu)$  represent systematic effects that are consistently observed across sessions for subject  $i$ .  $\mathbf{G}(\nu)$  is the covariance matrix of the random effects vector  $\mathbf{b}_i(\nu)$ . It encodes population variance for each random effect, as well as the correlation between different random effects. For example,  $Cor\{a_{1,i}(\nu), a_{4,i}(\nu)\}$  represents the correspondence between the direction and strength of subject  $i$ 's deviation from the population mean autocorrelation for tasks 1 (emotion) and 4 (relational), using the canonical HRF only. A strong positive correlation here would suggest that the same subjects tend to exhibit stronger or weaker autocorrelation, possibly due to their having a longer or shorter HRF than the canonical HRF, regardless of task.  $Cor\{a_{1,i}(\nu), b_{1,i}(\nu)\}$  represents the correspondence between the direction and strength of subject  $i$ 's deviation from the population mean on

task 1 and the effect of including the HRF derivative for the same task. A strong anti-correlation here would suggest that including the HRF derivative has a bigger effect on subjects who exhibit stronger autocorrelation when using the canonical HRF—in short, that inclusion of the HRF derivative achieves the goal of accounting for some population variability in HRF timing.

### 2.2.3. Prewhitening strategies

To evaluate the effectiveness of different prewhitening strategies on mitigating residual autocorrelation, we use an AR model with varying model order and varying degrees of regularization to estimate the prewhitening matrix  $\mathbf{W}_\nu$ . Specifically, we vary AR model order from  $p = 1$  to  $p = 6$ . We also consider automatic selection of the optimal AR model order at each vertex using Akaike information criterion (AIC) (Sakamoto et al., 1986; Bozdogan, 1987) as proposed by Luo et al. (2020), with a maximum model order of 10. AIC is inversely related to the log-likelihood and includes a penalty term to account for differences in the number of parameters in different models. AIC is commonly used as a model selection criterion, where the model with the lowest AIC has the best fit, accounting for model complexity. The AR model coefficients and residual variance are estimated using the Yule-Walker equations (Brockwell and Davis, 2009). We consider both local and global regularization of the AR coefficients and white noise variance. Local regularization refers to surface smoothing with a 5 mm FWHM Gaussian kernel; global regularization refers to smoothing with an infinitely-wide Gaussian kernel, nearly equivalent to averaging the AR coefficients across the cortex. In the case of optimal AR model order selection, we impute a value of 0 for any AR coefficients above the selected model order prior to regularization.

Our R/C++ implementation of prewhitening is available in the open-source `BayesfMRI` R package (Mejia et al., 2022), which is compatible with cortical surface and “grayordinates” neuroimaging file formats *via* the `ciftiTools` R package (Pham et al., 2022). After estimating the prewhitening matrix  $\mathbf{W}_\nu$  as described in Appendix 2, (Algorithm 1) the response and design matrix are pre-multiplied at each vertex by  $\mathbf{W}_\nu$ , changing the GLM in (1) to

$$\tilde{\mathbf{y}}_\nu = \tilde{\mathbf{X}}_\nu \boldsymbol{\beta}_\nu + \tilde{\boldsymbol{\epsilon}}_\nu, \quad \tilde{\boldsymbol{\epsilon}}_\nu \sim \text{MVN}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (3)$$

where  $\tilde{\mathbf{y}}_\nu = \mathbf{W}_\nu \mathbf{y}_\nu$ ,  $\tilde{\mathbf{X}}_\nu = \mathbf{W}_\nu \mathbf{X}_\nu$  and  $\tilde{\boldsymbol{\epsilon}}_\nu = \mathbf{W}_\nu \boldsymbol{\epsilon}_\nu$ . Note that the prewhitened design matrix  $\tilde{\mathbf{X}}_\nu$  may vary across vertices when using a local approach to prewhitening. This increases the computational burden associated with GLM coefficient estimation: with a common design matrix  $\tilde{\mathbf{X}}$ , the model coefficients for all vertices can be estimated with a single matrix multiplication step as  $(\tilde{\mathbf{X}}' \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}' \tilde{\mathbf{Y}}$ , where  $\tilde{\mathbf{Y}} = (\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_V)$ ; when the design matrix varies spatially, however, we must perform  $V$  matrix multiplications  $(\tilde{\mathbf{X}}'_\nu \tilde{\mathbf{X}}_\nu)^{-1} \tilde{\mathbf{X}}'_\nu \tilde{\mathbf{y}}_\nu$  to obtain the

coefficient estimate at each vertex  $v$ . Even more computationally burdensome is estimating  $\mathbf{W}_v$  at each location, which involves performing  $V$  different eigendecompositions. These obstacles are perhaps one reason that global prewhitening approaches are often preferred in a practical sense. To overcome these challenges, we have developed a highly computationally efficient implementation using parallelization and C++ backend code. This implementation typically completes in 1 min per scan for the task fMRI data we analyze here.

#### 2.2.4. Evaluation metrics

To evaluate the performance of each prewhitening strategy, we take two approaches. First, we directly assess the degree of residual autocorrelation still present in each task fMRI dataset after prewhitening. Using on the fitted GLS residuals at each vertex, we compute the autocorrelation index (ACI) as in Section 2.2.2. We also perform a Ljung-Box (LB) test at every vertex (Ljung and Box, 1978) to identify vertices exhibiting statistically significant levels of residual autocorrelation after prewhitening. We use the `Box.test` function in the `stats` R package, version 4.2.0. As in Corbin et al. (2018), we use the first 100 volumes of each session and consider up to 20 lags. We consider two approaches to determine the degrees of freedom (DOF) for the test: accounting for the intercept only, or accounting for the intercept and the AR model coefficients. We consider the intercept-only approach for maximum comparability with Corbin et al. (2018).<sup>1</sup> When accounting for the AR( $p$ ) model coefficients as well, the DOF for the LB test is  $20 - \lceil p * 100/T \rceil - 1$ , where  $T$  is the original number of volumes in the task fMRI session (given in Table 1). We scale the number of AR coefficients  $p$  by  $100/T$  to account for the fact that the AR model parameters were estimated across the whole duration of the scan, not just the 100 volumes used for the LB test. We determine vertices whose residuals exhibit significant autocorrelation based on those with  $p < 0.05$  after false discovery rate correction (Benjamini and Hochberg, 1995). We compute the proportion of vertices exhibiting significant autocorrelation before and after prewhitening with each technique to determine the ability of each prewhitening method to effectively eliminate autocorrelation.

Second, we quantify false positives using resting-state fMRI data, assuming a false boxcar task paradigm. For each resting-state fMRI dataset, we perform GLS using the estimated prewhitening matrix for each prewhitening strategy. We then perform a  $t$ -test at every vertex. We

<sup>1</sup> Though Corbin et al. (2018) did not state explicitly how DOF was determined in their analysis, their implementation used the Matlab Ljung-Box test function, where ignoring the model DOF is the default. Further, the higher-order SPM FAST models considered in Corbin et al. (2018) contain more than 20 DOF, which would result in negative DOF for the LB test if taken into account.

correct for multiple comparisons across all vertices with Bonferroni correction to control the family-wise error rate (FWER) at 0.05. While Bonferroni correction is typically considered overly conservative for whole-brain voxel-wise analysis involving potentially hundreds of thousands of tests, here we are performing fewer than 6,000 tests per hemisphere. In previous work we have found Bonferroni correction to have similar power as permutation testing for the cortical surface resampled to a similar resolution (Spencer et al., 2022). We visualize the spatial distribution of false positive vertices and quantify the false positive rate and FWER before and after prewhitening. We obtain 95% confidence intervals for the FWER using Agresti-Coull intervals for proportions (Agresti and Coull, 1998).

## 3. Results

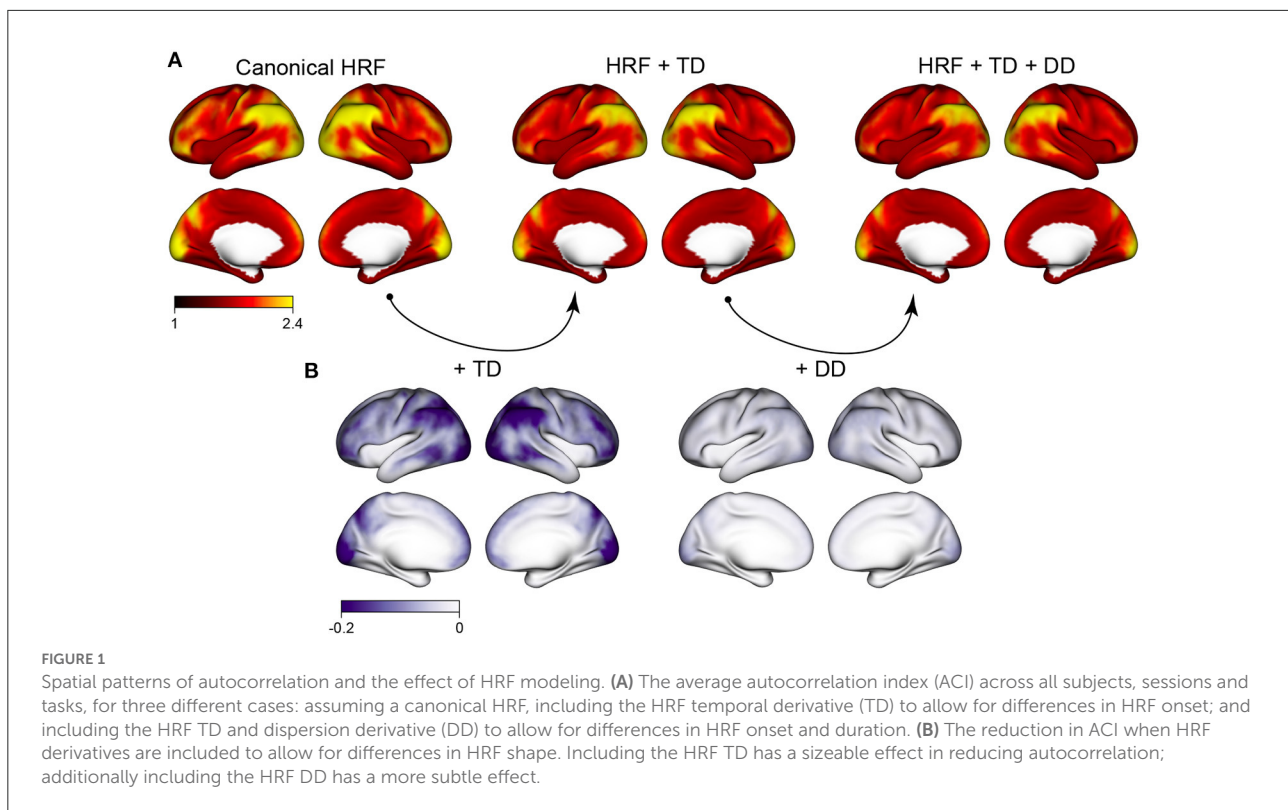
### 3.1. Overview

We first examine the spatial patterns and factors influencing autocorrelation in task fMRI prior to any prewhitening, using a random effects analysis of task fMRI data from the HCP retest dataset. This allows us to understand to what degree residual autocorrelation varies across the cortex in task fMRI studies, which in turn helps inform an effective approach to prewhitening. We find that residual autocorrelation varies markedly across the cortex, and the spatial topology is influenced by the task being performed, the phase encoding direction, and systematic inter-subject differences. Effective modeling choices (e.g., HRF flexibility, nuisance regression) can mitigate autocorrelation, but their effects are relatively modest and they do not eliminate spatial variability. We then assess the ability of different AR-based prewhitening strategies to effectively mitigate residual autocorrelation and control false positive rates. We consider low- and high-order AR models, as well as optimal determination of the AR model order at each vertex. We also consider two opposing approaches to spatial regularization of AR model parameters: local spatial smoothing and global averaging across the cortex. We find that higher-order AR models that allow for spatial variability in AR model parameters are able to effectively mitigate autocorrelation, while global averaging and very low-order AR models retain substantial levels of autocorrelation.

### 3.2. Spatial patterns of residual autocorrelation

Figure 1 shows the autocorrelation index (ACI) across the cortex, averaged over all subjects, sessions, runs, and tasks in the



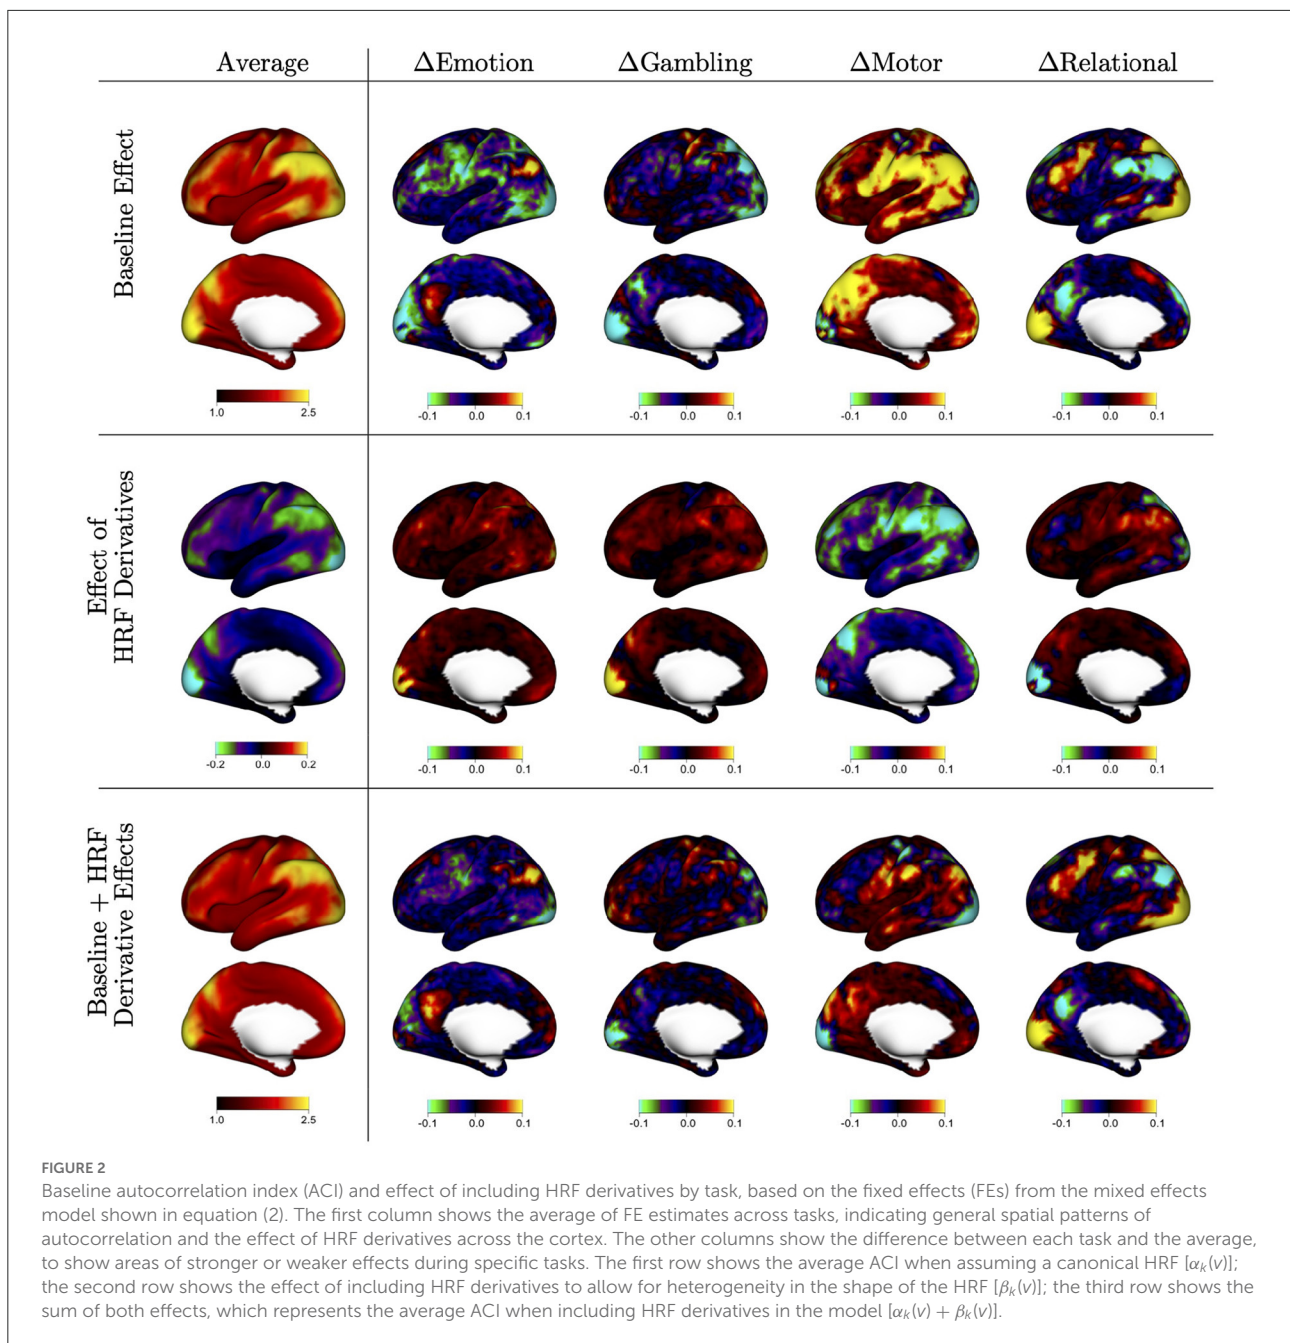


HCP retest study. The spatial topology of autocorrelation across the cortex is striking: higher autocorrelation is seen in frontal, parietal, and occipital areas, particularly the inferior parietal cortex and the occipital pole. Much lower autocorrelation is seen in the insula, the fusiform gyrus and the temporal pole, for example. The regions with highest autocorrelation tend to be near the edge of the brain, possibly reflecting in part effects of motion. Three different modeling strategies for the hemodynamic response function (HRF) are considered, ranging from rigid [canonical HRF] to more flexible (canonical HRF plus its temporal derivative (TD) and dispersion derivative (DD)). The degree of autocorrelation is reduced by the inclusion of HRF derivatives, illustrating that HRF mis-modeling is one source of autocorrelation that can be mitigated with more flexible HRF models. **Figure 1B** shows that inclusion of the TD serves to reduce autocorrelation most in the same areas where autocorrelation tends to be the highest, suggesting that the spatial patterns of autocorrelation may be due in part to heterogeneity in HRF onset and duration that varies across the cortex. Even after including HRF derivatives, however, there are still marked differences in the degree of residual autocorrelation across the cortex. Including the dispersion derivative has a more subtle effect, compared with just including the temporal derivative. For the remaining analyses, we therefore focus on inclusion of the HRF temporal derivative only.

### 3.3. Random effects analysis of residual autocorrelation

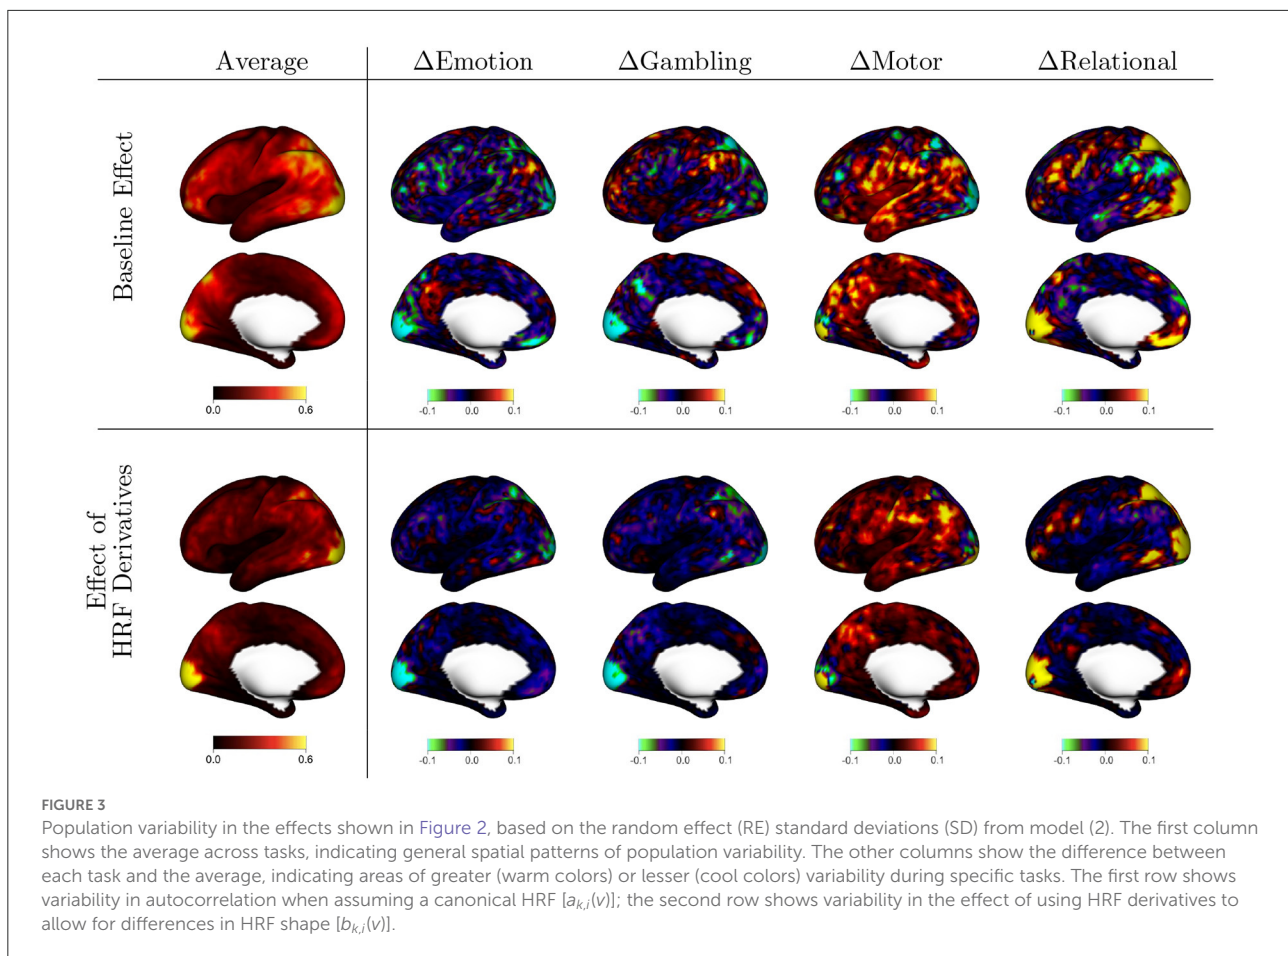
Here, we examine the sources of residual autocorrelation in task fMRI studies through a random effects analysis of repeated task fMRI scans from the HCP retest participants. We fit a series of general linear models (GLMs) to each task fMRI scan, varying the HRF modeling strategy and the number of motion regressors across GLMs. For each GLM, we quantify the autocorrelation index (ACI) of the model residuals at each surface vertex. We then fit the random effects model in (2) at each vertex to quantify the contribution of task-specific differences, HRF modeling strategy, phase encoding direction, and number of nuisance regressors on the residual autocorrelation index. The inclusion of random effects accounts for systematic between-subject variability and allows us to understand population heterogeneity in these effects.

**Figure 2** displays the fixed baseline effects associated with each task [ $\alpha_v^{(k)}$ ,  $k = 1, \dots, 4$  in model (2)], along with the effect of including the HRF derivative [ $\beta_v^{(k)}$ ,  $k = 1, \dots, 4$  in model (2)]. The first column displays the mean effect over tasks  $k = 1, \dots, 4$ ; the other columns display the difference between the task-specific effects and that mean effect. The baseline effects shown on the first row [ $\alpha_k(v)$  in model 2] represent the average ACI when assuming a canonical HRF, including 12



motion regressors, and using the LR phase encoding direction. The average pattern is very similar to the mean ACI in the dataset shown in Figure 1. The task-specific deviations show that ACI tends to be markedly higher or lower in certain regions depending on the task. For example, the motor task tends to have higher residual ACI in many areas, whereas the emotion task tends to have lower residual ACI. These results show that there are systematic task-related effects on autocorrelation that vary across the cortex.

The HRF derivative effects shown on the second row represent the change in average ACI when including the HRF derivative in each GLM [ $\beta_k(v)$  in model 2]. The mean effect shows that including HRF derivatives tends to decrease ACI, particularly in areas where ACI tends to be the highest, as observed in Figure 1. The task-specific deviations show that more flexible HRF modeling has the strongest effect for the motor task, mimicking the more severe autocorrelation seen in the motor task. The areas most affected by flexible HRF



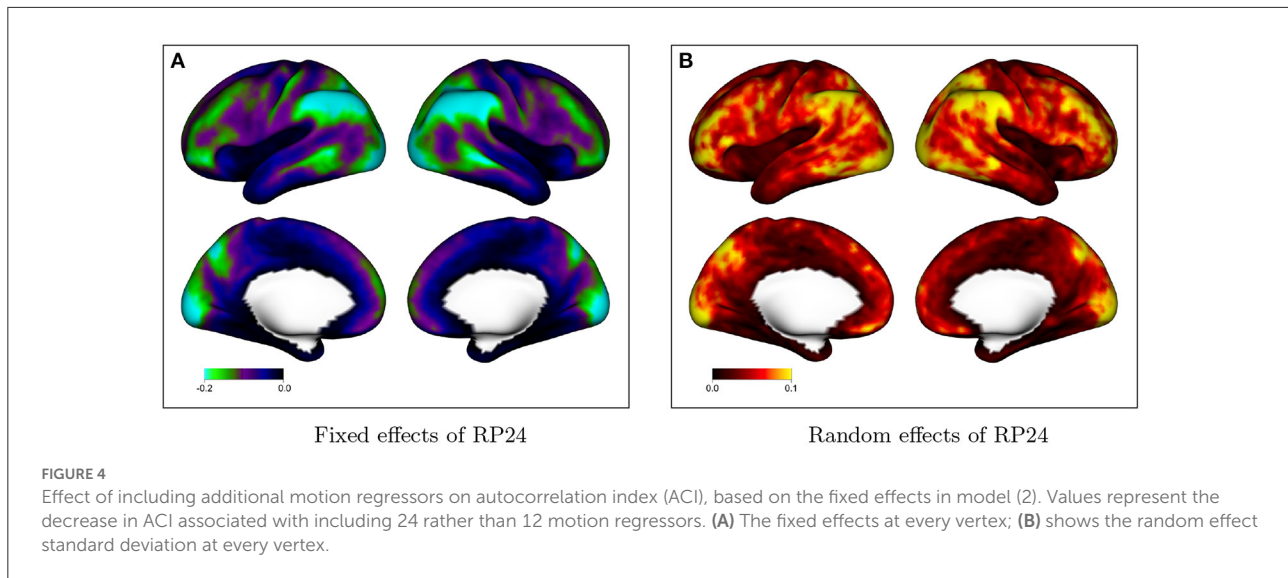
modeling for each task tend to somewhat mimic the spatial patterns unique to each task, but do not fully account for them.

The sum of both effects shown in the third row represent the average ACI when including HRF derivatives [ $\alpha_k(v) + \beta_k(v)$  in model 2]. The average image shows reduced autocorrelation on average, consistent with 1. The task-specific deviations show that task-related differences in residual autocorrelation are substantially reduced but not eliminated when using a more flexible HRF model. This illustrates that the differences in autocorrelation across the different task experiments may be largely due to differences in the HRF shape or onset that are not captured with the canonical HRF. This is particularly true for the motor task, which exhibits similar magnitude of autocorrelation compared to the other tasks once the HRF onset is allowed to vary.

Figure 3 displays the random effects associated with each fixed effect shown in Figure 2. These random effects represent reliable between-subject differences observed across repeated sessions from each subject. The scale of the values is standard deviation, so they share the same units as the fixed effects. The figure organization is the same as that of Figure 2, with the first column representing the mean over tasks and the remaining

columns representing differences between each task and that mean. The mean effects show that there is substantial population heterogeneity in the degree of residual autocorrelation, as well as in the reduction in residual autocorrelation achieved by more flexible HRF modeling. In general, the spatial patterns mimic that of the fixed effects: areas that tend to have higher autocorrelation on average also tend to exhibit greater systematic variability across subjects, and areas that benefit more from flexible HRF modeling on average also tend to exhibit the most population variability in the degree of reduction.

Figure 4 shows fixed and random effects of including additional motion regressors (24 vs. 12) in the GLM on the degree of residual autocorrelation across the cortex. In the GLM with 12 motion regressors, the six realignment parameters (RPs) plus their one-back differences are included as covariates; in the model with 24 motion regressors, their square of each term is also included. Figure 4A shows that on average, including additional motion regressors decreases residual autocorrelation. This illustrates that without adequate nuisance signal modeling, temporally correlated noise such as that arising from head motion will be at least partly absorbed into the residuals, which will consequently exhibit greater autocorrelation. The spatial



patterns mimic the spatial topology of baseline autocorrelation seen in Figures 1, 2, suggesting that thorough noise modeling helps to alleviate, without eliminating, the spatial heterogeneity in autocorrelation across the cortex. Figure 4B shows that there is population heterogeneity in the benefit of including additional motion regressors on autocorrelation, particularly in those areas with the most benefit on average.

Figure 5 takes a deeper look at population heterogeneity by examining the correlation between different random effects in the model. The lower triangle of the matrix is divided into several blocks, representing the different random components of the model. Blocks (A) and (B) show that there is moderate correlation between the different tasks in terms of the baseline effect and the effect of including HRF derivatives. This shows that for a particular subject, the spatial topology of autocorrelation and HRF shape are similar but not identical across different tasks. In block (C), the diagonal elements show a strong negative correlation between the baseline level of autocorrelation and the reduction in autocorrelation due to flexible HRF modeling. This shows that subjects having higher (or lower) baseline autocorrelation tend to benefit more (or less) from flexible HRF modeling. The strong but not perfect correlations suggest that accounting for between-subject differences in HRF shape reduces, but does not eliminate, systematic population variability in residual autocorrelation. Turning to the effect of motion, the strongly negative correlations in block (E) show that including additional motion regressors (24 instead of 12) reduces autocorrelation most for subjects who tend to exhibit stronger baseline autocorrelation. The moderate positive correlations in block (D) show that subjects who benefit more from more flexible HRF modeling also tend to benefit more from inclusion of additional motion

regressors, suggesting that these two approaches may be complementary in reducing residual autocorrelation.

Figure 6A displays the fixed effects associated with the utilizing an RL phase encoding direction acquisition, compared with an LR phase encoding direction. The effect of phase encoding direction on residual autocorrelation is clearly lateralized, with the RL phase encoding direction generally producing less autocorrelation on the right side of each hemisphere (the lateral cortex of the right hemisphere, and the medial cortex of the left hemisphere). This is likely due to lateralized distortions induced by the RL and LR phase encoding directions even after distortion correction, as shown in Figure 6B. The spatial distribution of residual autocorrelation is therefore sensitive to specific acquisition. This may result in an increased risk of false positives in certain areas, depending on the acquisition method. For example, using an LR phase encoding direction, residual autocorrelation is generally higher within the right lateral cortex. This will result in higher rates of false positives compared with the left lateral cortex if not accounted for with prewhitening techniques that account for such spatial discrepancies in residual autocorrelation.

We examine this further through a small simulation study from HCP subject 103,818, shown in Figure 7. We consider a strip of nine voxels overlapping with the edge of the brain of the subject, shown in Figure 7A. These include, sequentially, three voxels in CSF (red), two cortical gray matter voxels (yellow), and four white matter voxels. In addition, we include 26 background voxels on the left and 14 additional WM voxels on the right in order to absorb any edge effects. We generate autocorrelated timeseries for each voxel using an AR(3) model with white noise variance equal to 1. The AR coefficients are chosen to induce low ACI in white matter, moderate ACI in gray matter, high ACI in

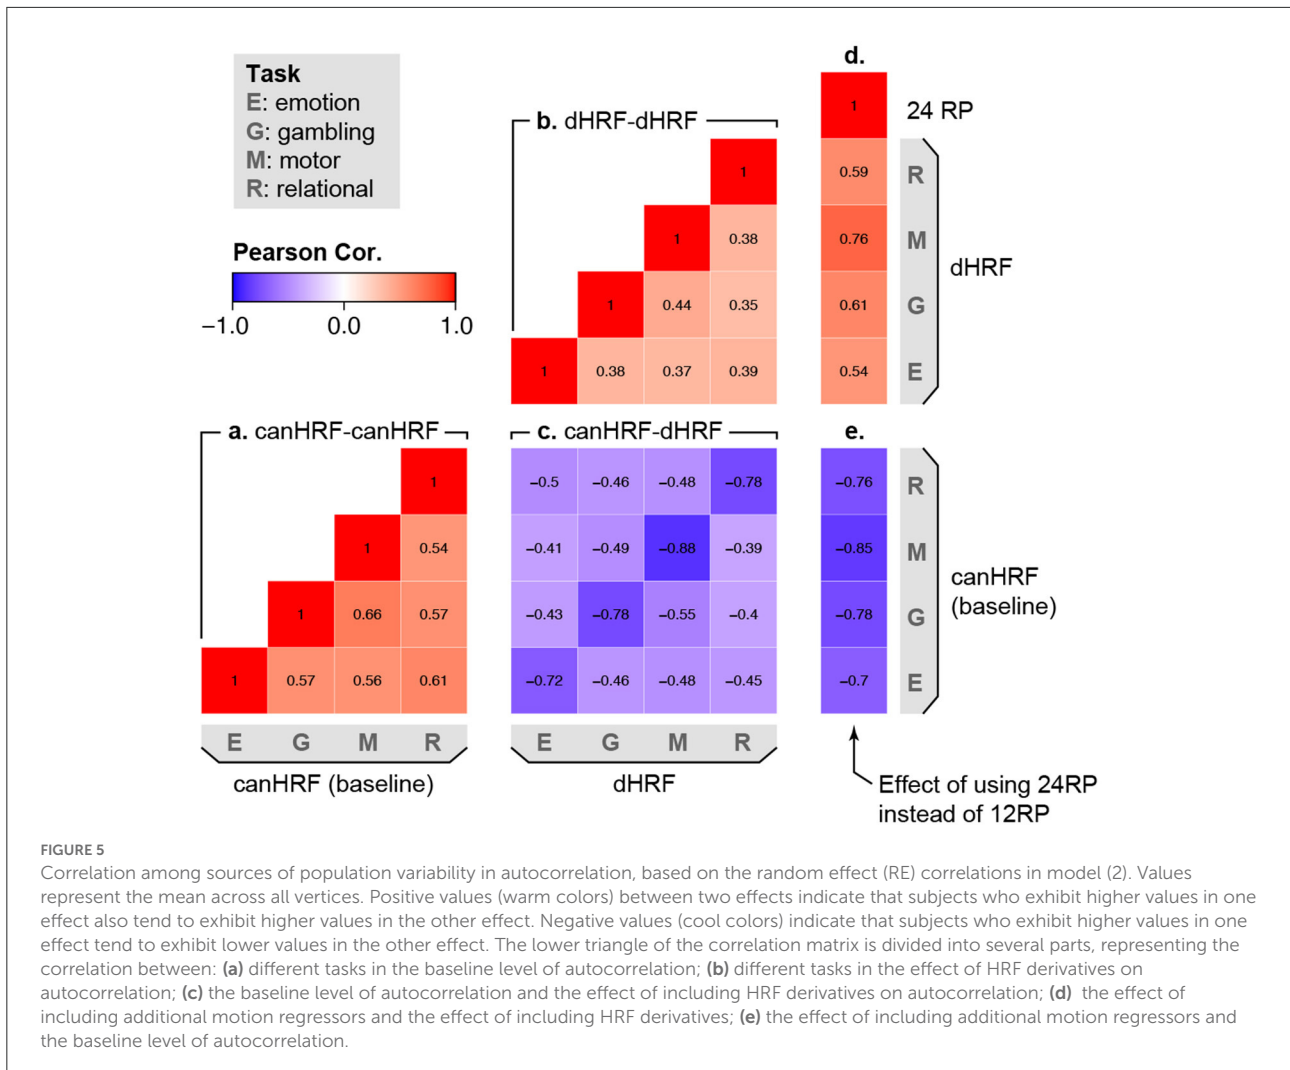


FIGURE 5

Correlation among sources of population variability in autocorrelation, based on the random effect (RE) correlations in model (2). Values represent the mean across all vertices. Positive values (warm colors) between two effects indicate that subjects who exhibit higher values in one effect also tend to exhibit higher values in the other effect. Negative values (cool colors) indicate that subjects who exhibit higher values in one effect tend to exhibit lower values in the other effect. The lower triangle of the correlation matrix is divided into several parts, representing the correlation between: (a) different tasks in the baseline level of autocorrelation; (b) different tasks in the effect of HRF derivatives on autocorrelation; (c) the baseline level of autocorrelation and the effect of including HRF derivatives on autocorrelation; (d) the effect of including additional motion regressors and the effect of including HRF derivatives; (e) the effect of including additional motion regressors and the baseline level of autocorrelation.

CSF, and unit ACI (the minimum) in background voxels. Table 2 gives the AR coefficients and resulting ACI in each region.

To examine the effect of the distortions introduced through lateralized phase encoding on the ACI, we estimate the distortion map from the original temporal mean undistorted brain images of the subject. We distort the simulated timeseries then apply distortion correction using the Anima image processing toolbox.<sup>2</sup> Figure 7B shows the true ACI in each voxel, the ACI after distortion, and the ACI after distortion correction based on the mean over 7,000 randomly generated timeseries. The last row shows the bias between the ACI of the distortion-corrected data, as a proportion of the true ACI. We see inflated ACI in the gray matter voxel bordering CSF and diminished ACI in the gray matter voxel bordering white matter. This agrees with our findings in Figure 6 and supports the hypothesis that the LR and RL acquisitions result in changes in autocorrelation in gray

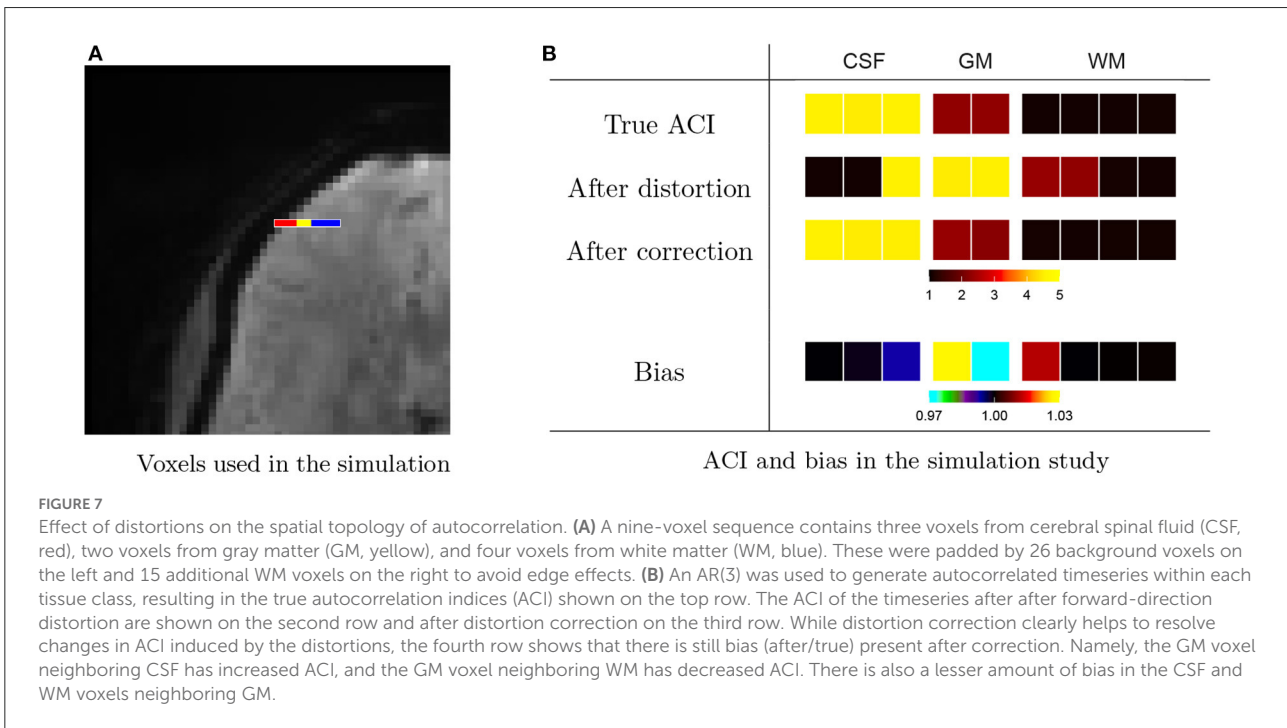
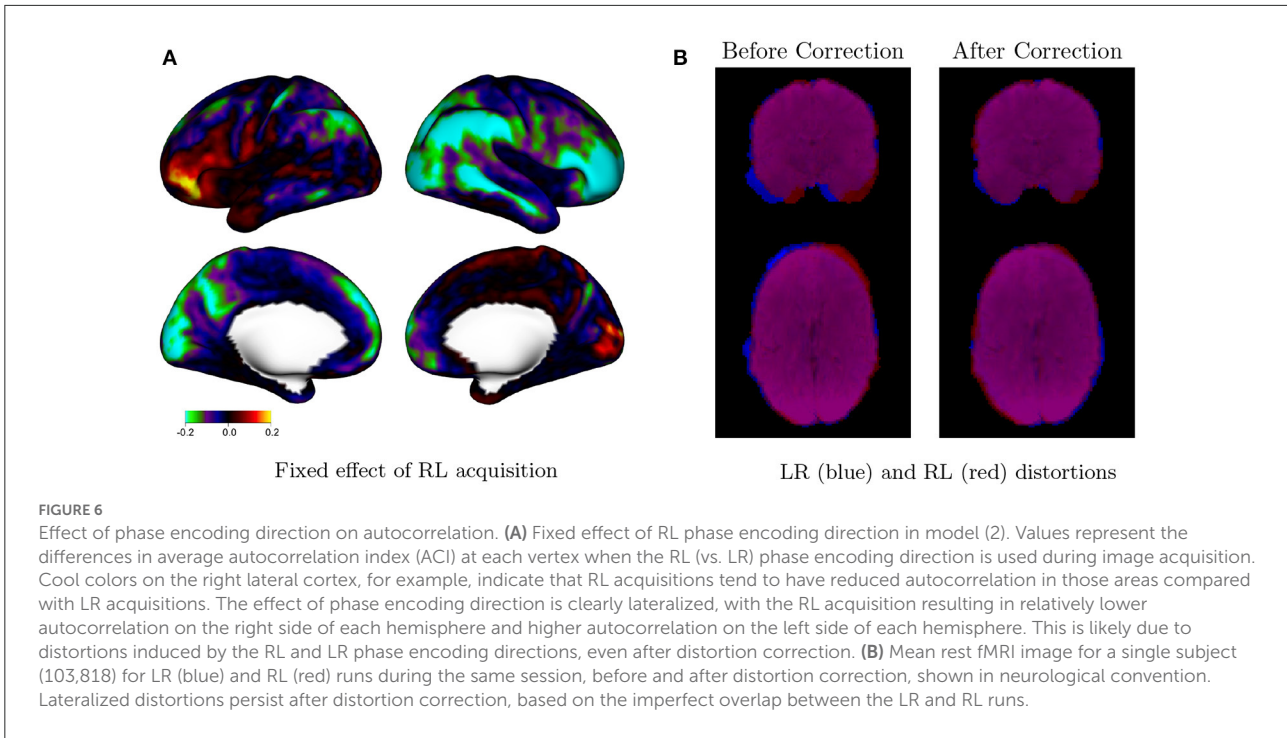
matter due to distortion-induced mixing of signals with white matter and CSF.

In sum, we observe marked spatial discrepancies in autocorrelation within cortical gray matter due to acquisition factors, modeling choices, task-related factors and individual differences. The following section evaluates the effect of different prewhitening strategies on mitigating autocorrelation, reducing spatial variability in autocorrelation, and controlling false positives.

### 3.4. The effect of prewhitening strategy on autocorrelation and false positives

Here, we apply several prewhitening strategies based on autoregressive (AR) modeling and evaluate their effect on both residual autocorrelation and false positive rates. Specifically, we consider AR model order ranging from 1 to 6, as well

<sup>2</sup> <https://github.com/Inria-Visages/Anima-Public>



as a spatially varying “optimal” model order based on Akaike information criterion (AIC). For each AR model order, we also consider two spatial regularization levels of the AR model coefficient estimates: “local” regularization is achieved by surface-based spatial smoothing the coefficient estimates

using a 5 mm full width at half maximum (FWHM) Gaussian kernel (Pham et al., 2022); “global” regularization is achieved by averaging the estimates across all cortical vertices. For optimal model order selection, all remaining coefficients after the AIC-based model order  $p^*$  (up to the maximum of 10) are set to

TABLE 2 Settings for the simulation study shown in Figure 7.

Tissue class	Number of voxels	AR(3) model	ACI
White matter	11 voxels	(0.1, 0.0, 0.0)	1.1
Gray matter	2 voxels	(0.425, 0.25, 0.1)	2.3
CSF	3 voxels	(0.5, 0.3, 0.1)	4.5
Background	11 voxels	White noise only	1

zero prior to regularization. All of the prewhitening methods considered are implemented in the open-source `BayesfMRI` R package, version 2.0 (Mejia et al., 2022; R Core Team, 2022).

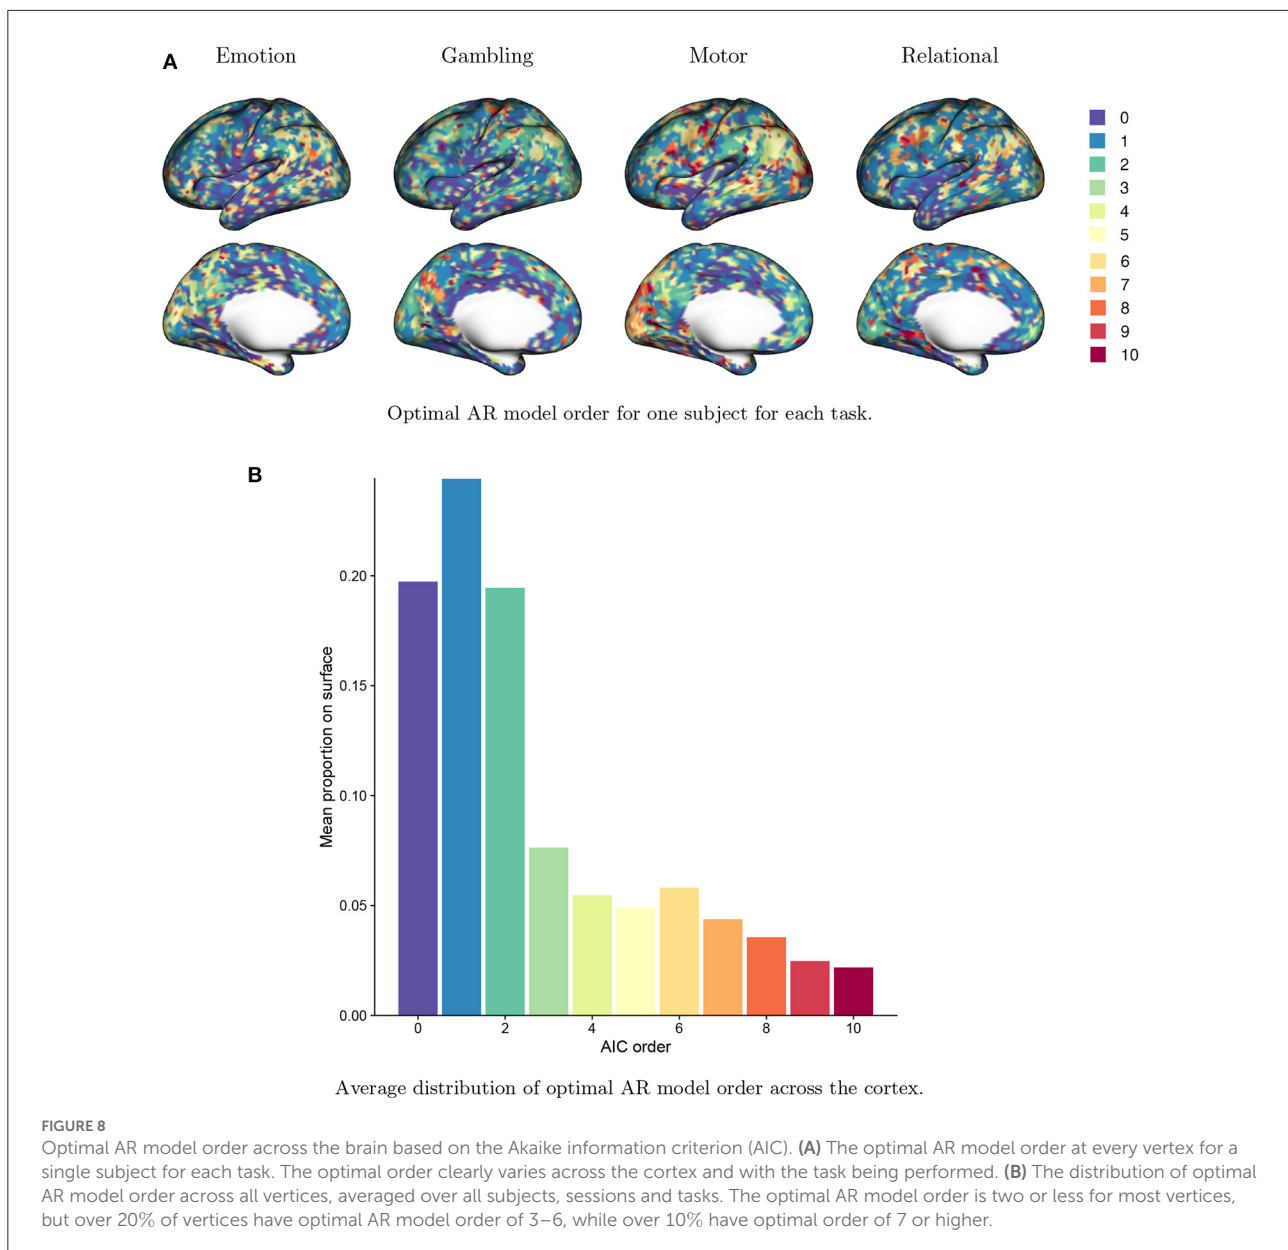
In this section, we include just 12 motion parameters (six rigid body realignment parameters and their one-back differences) in each GLM by default. This is because when combined with effective prewhitening, there does not appear to be a benefit of including the squares of the motion parameters and their one-back differences. In fact, it appears to be slightly detrimental, as shown in Figure A.1 in Appendix. This suggests that including these additional motion parameters no longer serves to reduce autocorrelation when effective prewhitening is performed, and the loss of degrees of freedom associated with the inclusion of superfluous covariates in the GLM actually worsens the performance of prewhitening.

Figure 8 displays the distribution of AIC-based AR model order,  $p^*$ , across the cortex. Figure 8A shows the value of  $p^*$  for a single run from a single subject across all four tasks. The spatial patterns mimic the general patterns of autocorrelation strength seen in Figure 1. Areas of higher autocorrelation generally require a higher AR model order (e.g., the inferior parietal cortex and the occipital pole), while areas of lower autocorrelation (e.g. the insula and the temporal pole) generally require a lower AR model order or no prewhitening at all ( $p^* = 0$ ). Differences across tasks can also be seen: for example, the motor task generally requires a higher AR model order, reflecting the stronger residual autocorrelation associated with the motor task as seen in Figure 2. The AIC-based model order is somewhat noisy, suggesting that some degree of regularization is needed. Figure 8B shows a histogram of AIC-based model order across all vertices. The proportions are averaged across all subjects, sessions and tasks. We see that on average, while the optimal AR model order is two or less for most vertices, over 30% of vertices have optimal AR model order of three or higher, while over 10% require an AR model order of seven or higher. This again underscores the important spatial differences in residual autocorrelation across the cortex and the need for prewhitening methods that account for those differences to avoid under- or over-whitening in a given area.

Figure 9 shows the effect of each prewhitening strategy on the degree of residual autocorrelation. Figure 9A displays the autocorrelation index (ACI) at each vertex averaged over all

subjects, sessions and tasks, prior to any prewhitening. Figure 9B displays the average ACI at each vertex after prewhitening with each strategy (varying AR model order, local vs. global regularization of the AR model coefficients). Figure 9C displays the mean and 95th quantile of ACI across the cortex by task, averaged over all subjects and sessions, after prewhitening with each strategy. Figures 9B, C show that there is a dramatic difference between local and global regularization in terms of reducing autocorrelation: local regularization reduces ACI more and mostly eliminates the spatial variability in ACI. The combination of higher AR model order [e.g., AR(6)] with local regularization is the most effective at reducing ACI. Notably, the use of higher model orders in combination with global regularization is not very effective at reducing autocorrelation in many areas of the cortex: even an AR(1) model with coefficients that are allowed to spatially vary appears to be more effective than an AR(6) model with global regularization. Interestingly, the use of AIC to select the optimal model order at each vertex does not appear to be advantageous over fitting an AR(6) model at every vertex. It is worth noting that an AR(6) model encompasses lower-order AR models, since the higher coefficients can equal zero. Local regularization of the AR model coefficients may have the effect of shrinking those higher coefficients closer to zero when that is appropriate. Therefore, fitting an AR(6) model at each vertex, combined with local regularization, may allow for less aggressive prewhitening in those areas that exhibit less autocorrelation.

Figure 10 displays the effect of prewhitening on the rate of vertices with statistically significant autocorrelation, based on performing a Ljung-Box (LB) test at every vertex (Ljung and Box, 1978). We correct for multiple comparisons by controlling the false discovery rate (FDR) at 0.05 using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). In Figures 10A, B, the value at each vertex represents the proportion of sessions that show significant autocorrelation across all subjects, sessions, and tasks. Figure 10C shows the proportion of significantly autocorrelated vertices in each session by task, averaged over all subjects, and sessions. For now we will focus on the solid lines, which represent the results of the LB test when we assume a single degree of freedom lost in all models. The patterns mimic those seen in Figures 9A, B: local regularization of AR prewhitening parameters is much more effective at reducing autocorrelation than global regularization, and even a parsimonious [e.g., AR(1)] AR model with local coefficient regularization is more effective than a high-order AR model with global regularization. Figure 10C shows that AR-based prewhitening with local regularization (the black lines) essentially eliminates statistically significant autocorrelation in all vertices, particularly when using an AR model order of 3 or higher. A globally regularized, high-order AR model approach is less effective, reducing the proportion of significantly autocorrelated vertices to 10–15%. Note that this is similar to the performance of the optimal 12-component



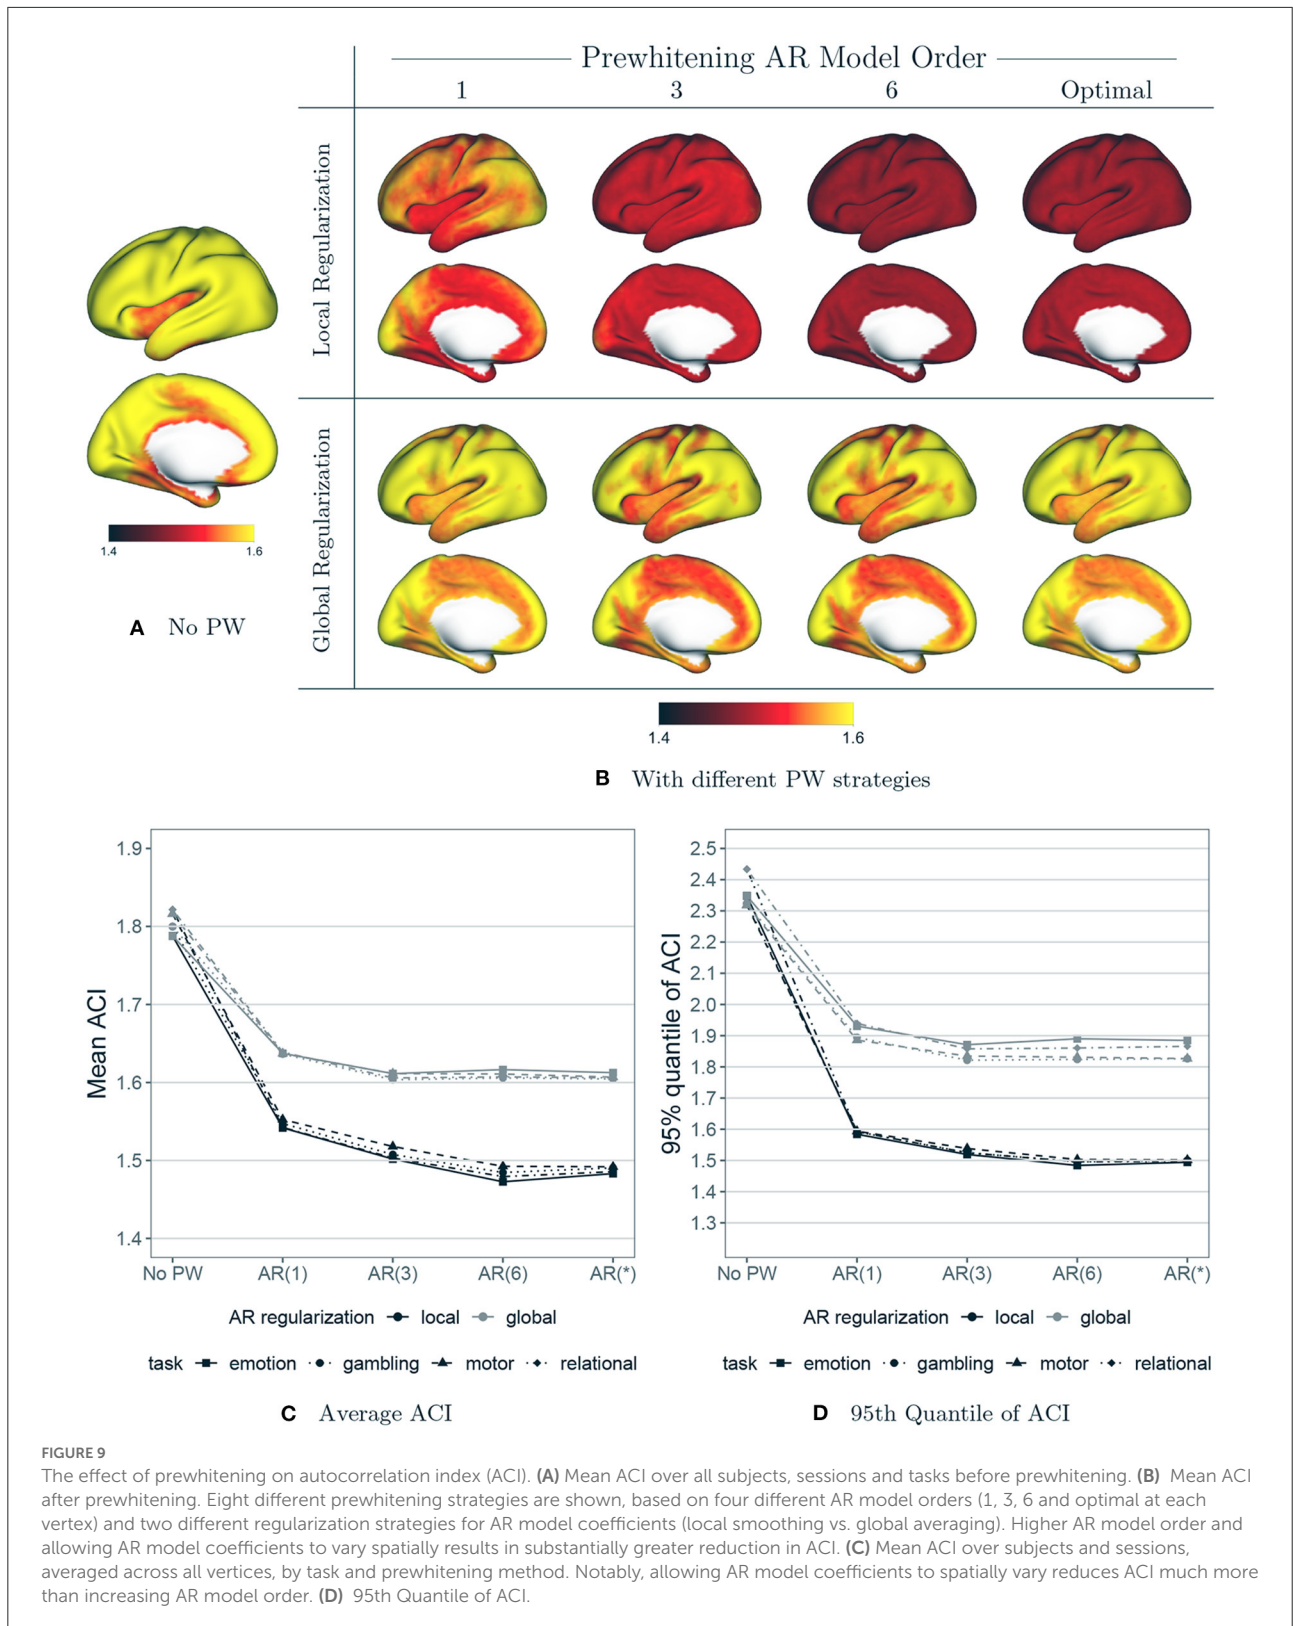
SPM FAST model for data with  $TR = 0.7$  (Corbin et al., 2018). The much greater reduction in autocorrelation with local regularization illustrates the need to consider spatial differences in autocorrelation for effective prewhitening.

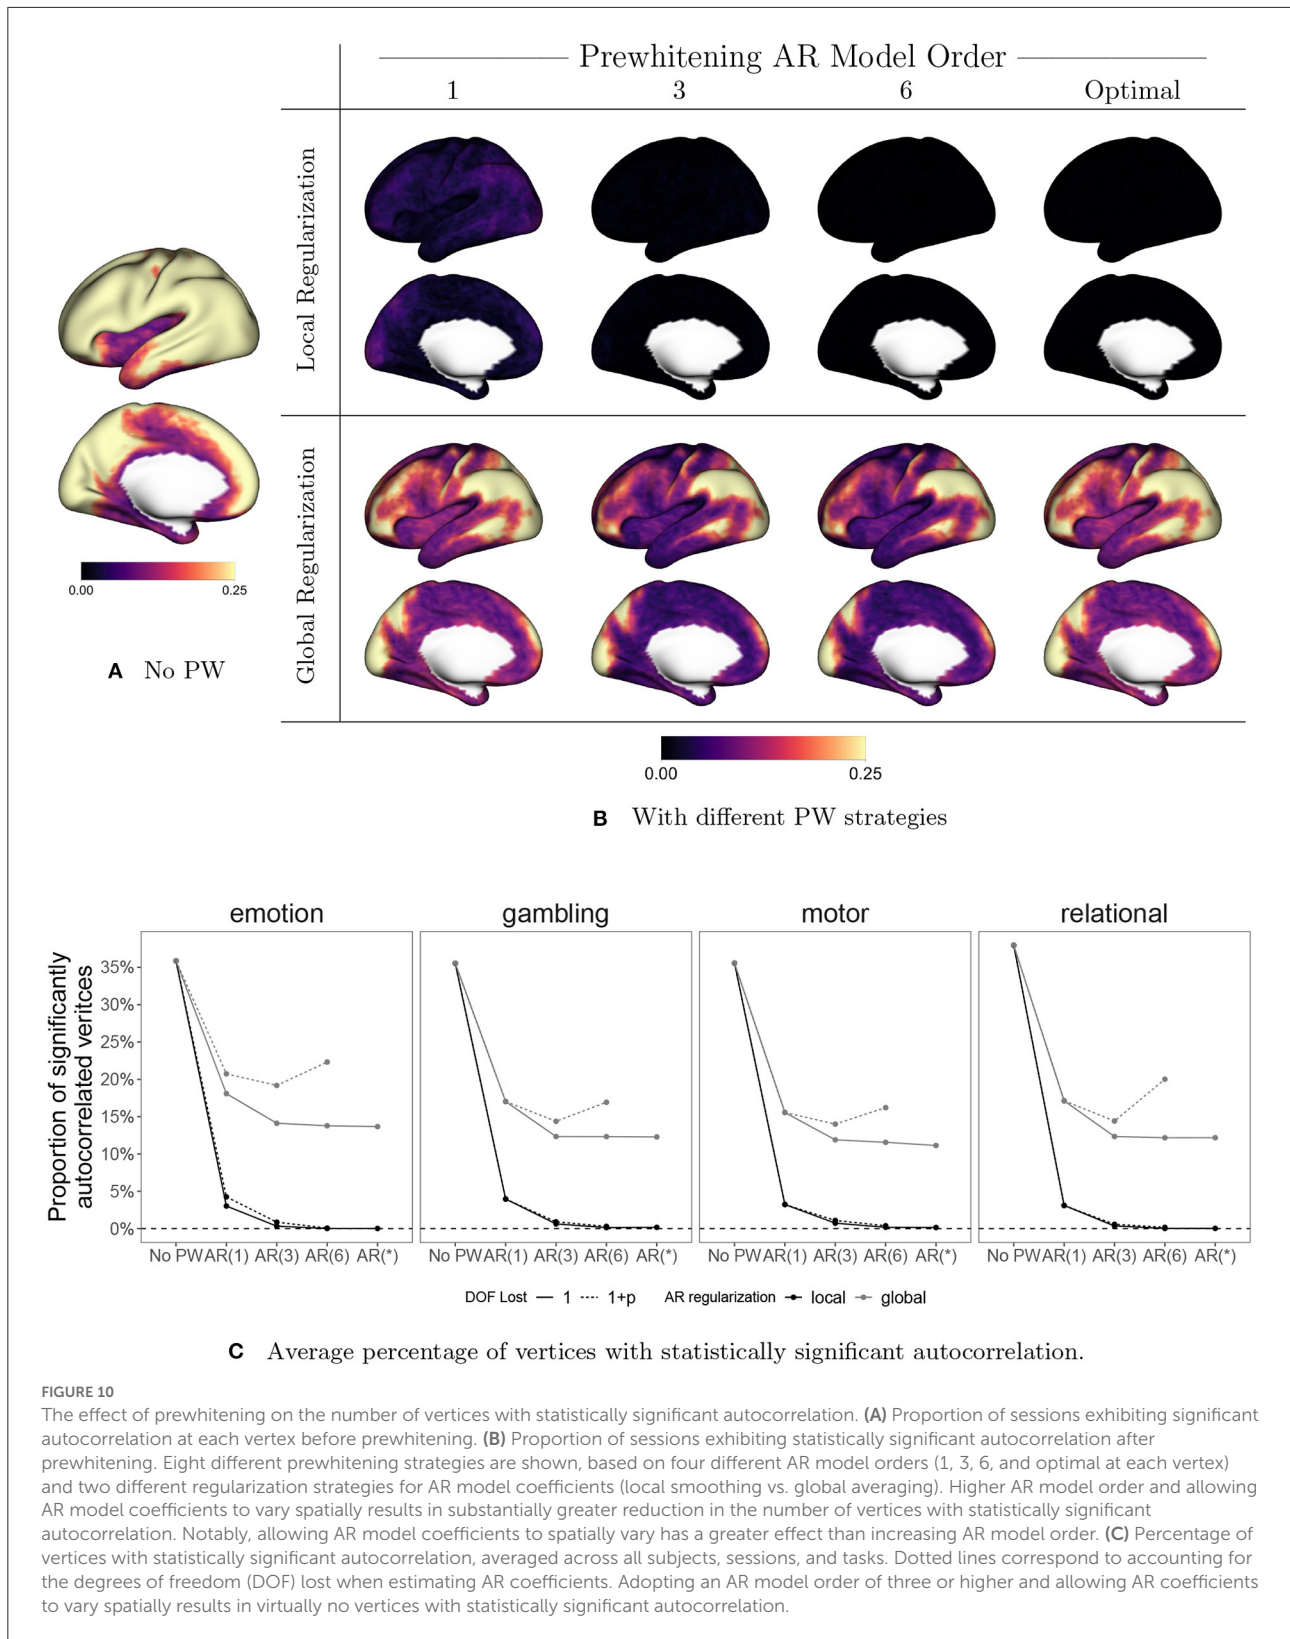
Note that in Figure 10C, we also consider the effect on the test result of accounting for the degrees of freedom (DOF) lost through the AR model fit ( $1 + p$ ) or just the intercept (1).<sup>3</sup> Though accounting for the AR fit in the total DOF is recommended, we account for the intercept only in

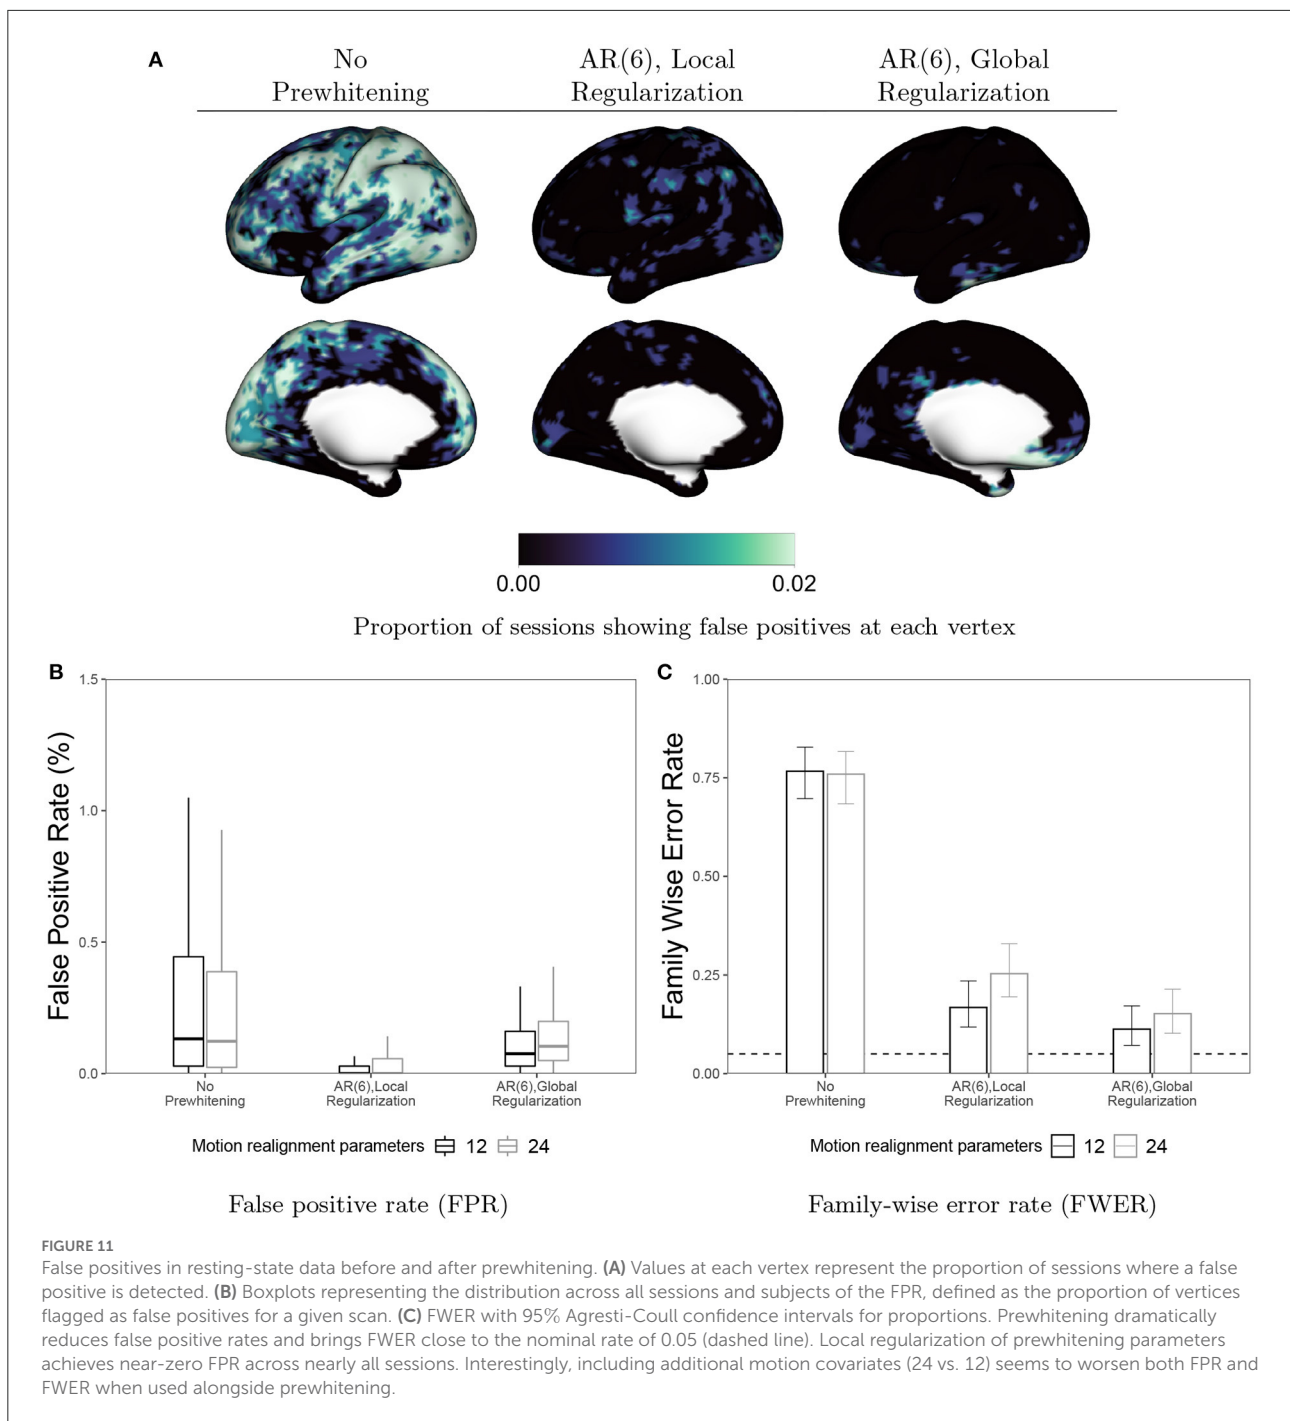
<sup>3</sup> Note that for the optimal AR model order approach [AR(\*)], we do not consider accounting for the DOF lost through model fitting when summarizing across vertices, since the DOF varies across vertices.

(A), (B) and the solid lines in (C). This is done in order to replicate the analysis of Corbin et al. (2018), which was based on the Matlab implementation of the Ljung-Box test, where ignoring the model DOF is the default. For more complex models involving many parameters, accounting for the model DOF generally results in apparently higher rates of autocorrelation, as seen in the U-shaped gray dashed lines in Figure 10C. This somewhat counterintuitive effect is simply a consequence of the loss in total degrees of freedom going from a more parsimonious model [e.g., AR(3)] to a more complex one involving more parameter estimates [e.g., AR(6)]. Accounting properly for the DOF lost helps to avoid overestimating the performance of more highly









parameterized models, which run the risk of overfitting to the data.

In Figure 11, we examine the effect of prewhitening on false positives in null (resting-state) data. Assuming a false on-off 10 s boxcar design, we fit a GLM and perform a *t*-test on the task coefficient at every vertex. We perform Bonferroni correction of the *p*-values to control the FWER at 0.05. Note that while Bonferroni correction is often considered overly

conservative for volumetric fMRI analyses involving hundreds of thousands of tests, here we have resampled the data to 6,000 vertices per hemisphere, so the number of tests being performed is an order of magnitude less. We have previously observed that Bonferroni correction is not more conservative than permutation testing in this data. Figure 11A displays the proportion of sessions showing a false detection at each vertex when no prewhitening is performed or when an AR(6)

model is used for prewhitening with local or global coefficient regularization. Figure 11B displays the false positive rate (FPR), the proportion of vertices labeled as active in each session, averaged across all subjects, sessions and tasks. Figure 11C displays the FWER, the proportion of sessions exhibiting a single false positive vertex. In Figures 11B, C we also consider the effect of including additional motion parameters (24) vs. the 12 included by default in our analyses. We observe that for both local and global regularization, the inclusion of additional motion parameters actually worsens the FPR and FWER. This is in line with the slight increase in ACI we observe when these parameters are included in combination with prewhitening (see Figure A.1 in Appendix). Taken together, these results suggest that the loss in degrees of freedom associated with including superfluous covariates in the GLM worsen the performance of prewhitening. This illustrates that overparameterized GLMs may actually result in inflated false positive rates, in addition to their well-known effect of reducing power to detect true effects. Comparing the FPR and FWER before and after prewhitening, we see that prewhitening drastically reduces the FPR within each session, and achieves FWER fairly close to the nominal rate of 0.05.

Interestingly, using global regularization achieves slightly lower FWER while achieving slightly worse but still very low FPR. This surprisingly strong performance of global regularization stands in contrast to its poor performance in our task fMRI-based analyses, displayed in Figures 9, 10. This may indicate some limitations of using resting-state fMRI as “null” data for evaluating false positive control in task fMRI. There are many features of task fMRI data that may not be reflected in resting-state fMRI data. For example, mismodeling of the task-induced HRF can induce residual autocorrelation, as shown in Figure 1. Inclusion of HRF derivatives only partly accounts for the task-related differences in autocorrelation, as shown in Figure 2.

## 4. Discussion

In this paper, we have made three primary advances in prewhitening in fMRI. First, we performed a comprehensive analysis to examine the spatial topology of autocorrelation across the cortex and identify the different factors driving autocorrelation. Second, we evaluated the efficacy of a range of AR-based prewhitening methods at eliminating autocorrelation and controlling false positives. We found that “local” prewhitening methods that account for spatial variability strongly outperform “global” methods where the same filter is applied to each voxel or vertex of the brain. Third, we developed a fast implementation of local prewhitening, available through the open-source *BayesfMRI* R package, that overcomes the computational challenges associated with performing prewhitening at thousands of locations.

### 4.1. Variable autocorrelation across the cortex results in spatially differential false positive control

Using a mixed effects modeling approach and test-retest data from the Human Connectome Project, we showed that autocorrelation varies markedly across the cortex and is influenced by task-related differences, modeling choices, acquisition factors, and population variability. As a result, the spatial topology of autocorrelation in each fMRI scan is unique. Given the spatial variability in autocorrelation, global prewhitening will result in differential levels of false positive control across the brain or cortex. And because the spatial topology of autocorrelation is unique to each fMRI scan, the topology of false positive control as well as power will likewise vary across fMRI scans, even within the same study. For example, one subject may be less likely to see a significant effect in a certain region compared with another subject in the same study, simply due to differences in autocorrelation in that region. These results illustrate the importance of prewhitening techniques that capture the spatial variability in autocorrelation, in order to avoid differential false positive rates across the cortex or across the brain.

Current prewhitening methods implemented in major fMRI software tools often use a global prewhitening approach. One likely reason for this is the computational efficiency of global prewhitening, since it requires a single  $T \times T$  matrix inversion, unlike local prewhitening which requires  $V$  such inversions. Likewise, the GLM coefficients can be estimated in a single matrix multiplication step with global prewhitening, whereas local prewhitening requires  $V$  multiplications. Another seeming advantage of global regularization of the prewhitening parameters is the low sampling variability in the estimates of those parameters, though this comes as the cost of large biases for specific locations. Local prewhitening can lead to noisier estimates of the prewhitening parameters, though smoothing can help combat this. While previous work based on volumetric fMRI found smoothing to be detrimental because of mixing signals across tissue classes (Luo et al., 2020), our use of surface-based smoothing largely avoids this limitation. Using cortical surface fMRI data also has the advantage of reduced dimensionality and the option to further reduce dimensionality through resampling without significant loss of spatial resolution. This lower dimensionality, combined with an implementation optimized for speed, makes our approach to local prewhitening quite feasible (1 min per run for the task fMRI we analyze here).

### 4.2. Implications for volumetric fMRI analyses

Our analysis focused on cortical surface-based analysis, but our findings have major implications for volumetric fMRI

analysis as well. The issue of spatially varying autocorrelation is actually more salient in volumetric fMRI, because autocorrelation is known to differ markedly across tissue classes, with CSF generally exhibiting higher autocorrelation and white matter exhibiting lower autocorrelation (Bollmann et al., 2018; Luo et al., 2020). A global prewhitening approach in volumetric fMRI may have more severe consequences than in surface-based analyses because of the more dramatic differences in autocorrelation across tissue classes. For example, to eliminate autocorrelation within CSF and thereby control false positive rates there, we may over-whiten within gray matter. Even if we target the gray matter, standard volumetric smoothing exacerbates differences in autocorrelation, increasing autocorrelation in voxels near CSF and decreasing autocorrelation in voxels near white matter. While these issues point to the importance of local prewhitening in volumetric fMRI analysis, the higher dimensionality of that data introduces new computational challenges. For example, our implementation of local prewhitening would take 10 min per run for a volumetric analysis involving 100,000 voxels or more, compared with 1 min per run for 12,000 surface vertices. Additionally, our local regularization approach of smoothing AR coefficients may not translate seamlessly to volumetric fMRI analysis, given the risks associated with smoothing across tissue classes. For volumetric fMRI analysis, it may be preferable to avoid AR coefficient smoothing or employ smoothing techniques that respect tissue class boundaries.

### 4.3. Low-order AR models perform surprisingly well when allowed to vary spatially

We were somewhat surprised by the fairly strong performance of AR(1) models when the AR coefficients were allowed to vary spatially through local regularization. This echoes the relatively strong performance of AFNI (Cox, 1996), which assumes a spatially varying ARMA(1,1) model with no smoothing, observed by Olszowy et al. (2019). In our analysis, locally regularized AR(1) prewhitening consistently outperformed globally regularized AR(6) prewhitening at reducing autocorrelation. We generally saw the best performance for local AR(6) prewhitening, but its improvement over local AR(1) was small compared to the difference between local and global regularization. This suggests that for volumetric fMRI analysis where the computational burden associated with voxel-specific prewhitening may be substantial, lower-order AR models may be worth consideration since they would speed up estimation of the prewhitening coefficients and matrix inversion.

### 4.4. FWER is not the whole picture

Because many multiplicity correction methods focus on controlling the family-wise error rate (FWER) or the probability of observing a single false positive voxel, vertex, or cluster, FWER control is often used as an evaluation metric for prewhitening methods. While this can be a useful metric in determining whether important modeling assumptions (e.g., independent residuals) have been satisfied, it is not the only important one. We advocate for two additional considerations: equal false positive control across the brain or cortex, and avoiding unnecessary loss of power. Global prewhitening that does not consider the spatial variance in autocorrelation may control the FWER, but will generally fail to achieve spatially homogeneous levels of false positive control. As a result, we may be much more likely to observe false positives in certain regions, which may also differ across subjects. In addition, global prewhitening will tend to over-whiten regions with low autocorrelation, which can lead to overly conservative inference, manifesting as a lack of power to detect effects. Just as we will be more likely to see false positives in certain regions, other regions will suffer disproportionately from a loss of power to detect real effects. Indeed, achieving nominal FWER without spatially homogeneous false positive control will almost surely come at the cost of unnecessary loss of power in many parts of the brain.

### 4.5. Acquisition-induced distortions change the spatial topology of autocorrelation and false positive control

One striking finding in our analysis was the effect of phase encoding direction on the spatial topology of residual autocorrelation. Comparing the RL and LR phase encoding directions, LR generally produced higher autocorrelation in the right lateral cortex and left medial cortex, while RL generally produced higher autocorrelation in the left lateral cortex and the right medial cortex. In other words, the choice of phase encoding direction generally had opposing effects on the lateral and medial cortices within each hemisphere, as well as across hemispheres. Why might this be? The LR and RL phase encoding directions are known to introduce lateralized distortions, even after distortion correction. These distortions are due to areas of varying magnetic susceptibility giving rise to signal stretching and signal pile-up based on the direction of the phase-encode direction for acquisitions such as echo-planar imaging (Jezzard and Clare, 1999). Those distortions cause a slight misalignment of the fMRI data on the structure of the brain. This has the result of mixing CSF signals with higher autocorrelation into some cortical areas, and mixing white matter signals with lower autocorrelation into others. For example, the higher autocorrelation on the right lateral cortex

with the LR acquisition may result from the introduction of CSF signals produced by a slight shift of gray matter voxels to the left; a similar shift to the left of the right medial cortex will introduce white matter signals, resulting in lower autocorrelation. This is somewhat analogous to the effect of volumetric smoothing on gray matter voxels bordering CSF and white matter observed by Luo et al. (2020) (though note that these effects of volumetric smoothing would be in addition to, and perhaps exacerbate, the effects of distortions). As a result, in the HCP there may be sizeable discrepancies in false positive control and power across and within each hemispheres before prewhitening or when using global prewhitening. While such lateralized effects are somewhat unique to HCP-style acquisitions that employ left-to-right or right-to-left phase encoding, other acquisitions can introduce different types of distortions that may also change the spatial topology of autocorrelation. Different acquisitions may therefore produce very different spatial distributions of false positive control, if not addressed through an effective local prewhitening strategy.

#### 4.6. Limitations and future directions

Our study has several limitations. First, our analysis was based on a single dataset, the Human Connectome Project, which includes data from healthy young adults collected using a certain multi-band acquisition protocol. Previous studies have found the baseline level of autocorrelation, as well as the efficacy of prewhitening, to differ across datasets of varying TR (Corbin et al., 2018; Olszowy et al., 2019). Future work should examine the generalizability of our findings, particularly the efficacy of our local prewhitening approach, to data collected with different acquisition protocols and in more diverse populations. Similarly, a valuable area of future work would be to assess the efficacy of local prewhitening strategies in volumetric task fMRI data, which presents unique challenges as described in Section 4.2.

Our study, as well as most prior studies on the efficacy of prewhitening in task fMRI analyses, focused on the ability of prewhitening techniques to effectively mitigate autocorrelation and control false positives. Here we discussed, but did not explicitly analyze, the possibility of a loss of power due to over-whitening. We argue that this is more likely with global prewhitening strategies that aim to achieve false positive control within a particular region or tissue class, at the risk of over-whitening in other areas. A valuable area of future work would be to examine the effect of different prewhitening techniques on power across the brain.

Here, we considered the effect of HRF modeling strategy on residual autocorrelation, and observed that the inclusion of temporal and dispersion derivatives of the HRF helped to alleviate it. However, we did not consider alternative, potentially advantageous HRF modeling strategies such as finite impulse response (FIR) (Ciuciu et al., 2003; Marrelec et al., 2003)

or inverse logit (Lindquist et al., 2009) models. For data with short TR such as the HCP, recently proposed techniques for blind deconvolution of the fMRI time series can be used to estimate the HRF associated with different regions of the brain in a particular dataset (Cherkaoui et al., 2021). When enough data is available to allow for accurate HRF estimation, these more flexible models may help account for additional spatial, within-subject and between-subject heterogeneity in HRF shape, onset and duration, particularly in more diverse populations (Woolrich et al., 2004). Thus, these more flexible HRF modeling strategies could further reduce autocorrelation and its spatial variance.

Finally, one limitation of our implementation of AR-based prewhitening implementation is that we did not account for potential bias in the prewhitening matrix due to using the fitted residuals as a proxy for the true residuals. Since the fitted residuals have a different dependence structure induced by the GLM, their covariance matrix is not equal to that of the true residuals. This bias will generally be worse in overparameterized GLMs, which may help explain why we observed a slightly detrimental effect of including all 24 motion regressors when prewhitening was also performed (see Figure A.1 in Appendix). A valuable topic of future work would be to develop prewhitening methods that formally model and adjust for this source of bias.

## 5. Conclusion

We performed a comprehensive investigation of the sources and patterns of residual autocorrelation across the cortical surface in multi-band task fMRI data. Our analysis revealed dramatic spatial differences in autocorrelation across the cortex. This spatial topology is unique to each session, being influenced by the task being performed, the acquisition technique, various modeling choices, and individual differences. If not accounted for, these differences will result in differential false positive control and power across the cortex and across subjects. We evaluated the efficacy of different prewhitening methods to mitigate autocorrelation and control false positives. Our findings demonstrate that allowing the prewhitening filter to vary spatially is crucial to effectively reducing autocorrelation and its spatial variability across the cortex. Our computationally efficient implementation of “local” prewhitening is available in the open-source `BayesfMRI` R package.

## Data availability statement

Publicly available datasets were analyzed in this study. This data can be found at: <http://humanconnectome.org>.

## Author contributions

FP and AM conceptualized and designed the study. AM, FP, and DP performed the statistical analysis. DP, DS, and AM developed the methods and software for the prewhitening techniques evaluated in the paper. FP and AM wrote the first draft of the manuscript, while RW and DP wrote sections. All authors reviewed and revised the manuscript. All authors contributed to the article and approved the submitted version.

## Funding

This work was supported by the National Institute of Biomedical Imaging and Bioengineering at the National Institutes of Health (NIH) under grant R01EB027119, by the National Institute of Neurological Disorders and Stroke at the NIH under the Grant No. R25NS117281, and by the National

Science Foundation (NSF) under grant IIS-2023985 from the Division of Information and Intelligent Systems.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

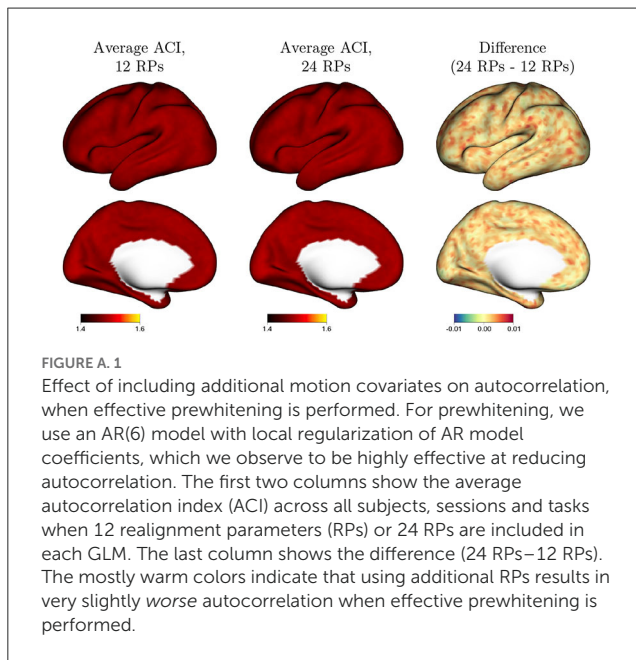
## References

- Afyouni, S., Smith, S. M., and Nichols, T. E. (2019). Effective degrees of freedom of the Pearson's correlation coefficient under autocorrelation. *NeuroImage* 199, 609–625. doi: 10.1016/j.neuroimage.2019.05.011
- Agresti, A., and Coull, B. A. (1998). Approximate is better than “exact” for interval estimation of binomial proportions. *Am. Stat.* 52, 119–126. doi: 10.1080/00031305.1998.10480550
- Aguirre, G. K., Zarahn, E., and D'Esposito, M. (1998). The variability of human, BOLD hemodynamic responses. *NeuroImage* 8, 360–369. doi: 10.1006/nimg.1998.0369
- Badillo, S., Vincent, T., and Ciuciu, P. (2013). Group-level impacts of within- and between-subject hemodynamic variability in fmri. *NeuroImage* 82, 433–448. doi: 10.1016/j.neuroimage.2013.05.100
- Barch, D. M., Burgess, G. C., Harms, M. P., Petersen, S. E., Schlaggar, B. L., Corbetta, M., et al. (2013). Function in the human connectome: task-fMRI and individual differences in behavior. *NeuroImage* 80, 169–189. doi: 10.1016/j.neuroimage.2013.05.033
- Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* 67, 1–48. doi: 10.18637/jss.v067.i01
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* 57, 289–300. doi: 10.1111/j.2517-6161.1995.tb02031.x
- Bollmann, S., Puckett, A. M., Cunningham, R., and Barth, M. (2018). Serial correlations in single-subject fMRI with sub-second TR. *NeuroImage* 166, 152–166. doi: 10.1016/j.neuroimage.2017.10.043
- Bozdogan, H. (1987). Model selection and Akaike's information criterion (AIC): The general theory and its analytical extensions. *Psychometrika* 52, 345–370.
- Braga, R. M., and Buckner, R. L. (2017). Parallel interdigitated distributed networks within the individual estimated by intrinsic functional connectivity. *Neuron* 95, 457–471. doi: 10.1016/j.neuron.2017.06.038
- Brockwell, P. J., and Davis, R. A. (2009). *Time Series: Theory and Methods*. Springer Science and Business Media.
- Buckner, R. L., Krienen, F. M., Castellanos, A., Diaz, J. C., and Yeo, B. T. T. (2011). The organization of the human cerebellum estimated by intrinsic functional connectivity. *J. Neurophysiol.* 106, 2322–2345. doi: 10.1152/jn.00339.2011
- Bullmore, E., Brammer, M., Williams, S. C. R., Rabe-Hesketh, S., Janot, N., David, A., et al. (1996). Statistical methods of estimation and inference for functional MR image analysis. *Mag. Reson. Med.* 35, 261–277. doi: 10.1002/mrm.1910350219
- Chen, J. E., Polimeni, J. R., Bollmann, S., and Glover, G. H. (2019). On the analysis of rapidly sampled fMRI data. *NeuroImage* 188, 807–820. doi: 10.1016/j.neuroimage.2019.02.008
- Cherkaoui, H., Moreau, T., Halimi, A., Leroy, C., and Ciuciu, P. (2021). Multivariate semi-blind deconvolution of fMRI time series. *NeuroImage* 241, 118418. doi: 10.1016/j.neuroimage.2021.118418
- Choe, A. S., Jones, C. K., Joel, S. E., Muschelli, J., Belegu, V., Caffo, B. S., et al. (2015). Reproducibility and temporal structure in weekly resting-state fMRI over a period of 3.5 years. *PLoS ONE* 10, e0140134. doi: 10.1371/journal.pone.0140134
- Ciuciu, P., Poline, J.-B., Marrelec, G., Idier, J., Pallier, C., and Benali, H. (2003). Unsupervised robust nonparametric estimation of the hemodynamic response function for any fmri experiment. *IEEE Trans. Med. Imaging* 22, 1235–1251. doi: 10.1109/TMI.2003.817759
- Corbin, N., Todd, N., Friston, K. J., and Callaghan, M. F. (2018). Accurate modeling of temporal correlations in rapidly sampled fMRI time series. *Hum. Brain Map.* 39, 3884–3897. doi: 10.1002/hbm.24218
- Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput. Biomed. Res.* 29, 162–173. doi: 10.1006/cbmr.1996.0014
- Delgado, M., Nystrom, L., Fissell, C., Noll, D., and Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *J. Neurophysiol.* 84, 3072–3077. doi: 10.1152/jn.2000.84.6.3072
- Eklund, A., Andersson, M., Josephson, C., Johansson, M., and Knutsson, H. (2012). Does parametric fMRI analysis with SPM yield valid results? *NeuroImage* 61, 565–578. doi: 10.1016/j.neuroimage.2012.03.093
- Elliott, M. L., Knodt, A. R., Ireland, D., Morris, M. L., Poulton, R., Ramrakha, S., et al. (2020). What is the test-retest reliability of common task-functional MRI measures? New empirical evidence and a meta-analysis. *Psychol. Sci.* 31, 792–806. doi: 10.1177/0956797620916786
- Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M., and Turner, R. (1998a). Event-related fMRI: characterizing differential responses. *NeuroImage* 7, 30–40. doi: 10.1006/nimg.1997.0306
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J.-P., Frith, C. D., and Frackowiak, R. S. (1994). Statistical parametric maps in functional imaging: a general linear approach. *Hum. Brain Map.* 2, 189–210. doi: 10.1002/hbm.460020402
- Friston, K. J., Josephs, O., Rees, G., and Turner, R. (1998b). Nonlinear event-related responses in fMRI. *Mag. Reson. Med.* 39, 41–52. doi: 10.1002/mrm.1910390109

- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., et al. (2013). The minimal preprocessing pipelines for the human connectome project. *Neuroimage* 80, 105–124. doi: 10.1016/j.neuroimage.2013.04.127
- Glover, G. H. (1999). Deconvolution of impulse response in event-related bold fMRI. *Neuroimage* 9, 416–429. doi: 10.1006/nimg.1998.0419
- Gordon, E. M., Laumann, T. O., Gilmore, A. W., Newbold, D. J., Greene, D. J., Berg, J. J., et al. (2017). Precision functional mapping of individual human brains. *Neuron* 95, 791–807. doi: 10.1016/j.neuron.2017.07.011
- Hariri, A. R., Tessitore, A., Mattay, V. S., Fera, F., and Weinberger, D. R. (2002). The amygdala response to emotional stimuli: a comparison of faces and scenes. *NeuroImage* 17, 317–323. doi: 10.1006/nimg.2002.1179
- Jenkinson, M., Beckmann, C. F., Behrens, T. E., Woolrich, M. W., and Smith, S. M. (2012). FSL. *Neuroimage* 62, 782–790. doi: 10.1016/j.neuroimage.2011.09.015
- Jezzard, P., and Clare, S. (1999). Sources of distortion in functional MRI data. *Hum. Brain Map.* 8, 80–85. doi: 10.1002/(SICI)1097-0193(1999)8:2/3<80::AID-HBM2>3.0.CO;2-C
- Kariya, T., and Kurata, H. (2004). *Generalized Least Squares*. John Wiley & Sons. doi: 10.1002/0470866993
- Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M.-Y., et al. (2015). Functional system and areal organization of a highly sampled individual human brain. *Neuron* 87, 657–670. doi: 10.1016/j.neuron.2015.06.037
- Lindquist, M. A. (2008). The statistical analysis of fMRI data. *Stat. Sci.* 23, 439–464. doi: 10.1214/09-STS282
- Lindquist, M. A., Loh, J. M., Atlas, L. Y., and Wager, T. D. (2009). Modeling the hemodynamic response function in fMRI: efficiency, bias and mis-modeling. *Neuroimage* 45, S187–S198. doi: 10.1016/j.neuroimage.2008.10.065
- Lindquist, M. A., and Wager, T. D. (2007). Validity and power in hemodynamic response modeling: a comparison study and a new approach. *Hum. Brain Map.* 28, 764–784. doi: 10.1002/hbm.20310
- Ljung, G. M., and Box, G. E. (1978). On a measure of lack of fit in time series models. *Biometrika* 65, 297–303. doi: 10.1093/biomet/65.2.297
- Loh, J. M., Lindquist, M. A., and Wager, T. D. (2008). Residual analysis for detecting mis-modeling in fMRI. *Stat. Sin.* 18, 1421–1448.
- Luo, Q., Misaki, M., Mulyana, B., Wong, C.-K., and Bodurka, J. (2020). Improved autoregressive model for correction of noise serial correlation in fast fMRI. *Magn. Reson. Med.* 84, 1293–1305. doi: 10.1002/mrm.28203
- Marcus, D., Harwell, J., Olsen, T., Hodge, M., Glasser, M., Prior, F., et al. (2011). Informatics and data mining tools and strategies for the human connectome project. *Front. Neuroinform.* 5, 4. doi: 10.3389/fninf.2011.00004
- Marrelec, G., Benali, H., Ciuciu, P., Péligrini-Issac, M., and Poline, J.-B. (2003). Robust bayesian estimation of the hemodynamic response function in event-related bold fmri using basic physiological information. *Hum. Brain Map.* 19, 1–17. doi: 10.1002/hbm.10100
- Mejia, A. F., Spencer, D., Pham, D., Bolin, D., Ryan, S., and Yue, Y. R. (2022). *BayesfMRI: Bayesian Methods for Functional MRI. R Package Version 0.2.0*.
- Mejia, A. F., Yue, Y., Bolin, D., Lindgren, F., and Lindquist, M. A. (2020). A Bayesian generallinear modeling approach to cortical surface fMRI data analysis. *J. Am. Stat. Assoc.* 115, 501–520. doi: 10.1080/01621459.2019.1611582
- Monti, M. M. (2011). Statistical analysis of fMRI time-series: a critical review of the GLM approach. *Front. Hum. Neurosci.* 5, 28. doi: 10.3389/fnhum.2011.00028
- Mumford, J. A., and Nichols, T. (2009). Simple group fMRI modeling and inference. *Neuroimage* 47, 1469–1475. doi: 10.1016/j.neuroimage.2009.05.034
- Olshowy, W., Aston, J., Rua, C., and Williams, G. B. (2019). Accurate autocorrelation modeling substantially improves fMRI reliability. *Nat. Commun.* 10, 1220. doi: 10.1038/s41467-019-09230-w
- Penny, W., Flandin, G., and Trujillo-Barreto, N. (2007). Bayesian comparison of spatially regularized general linear models. *Hum. Brain Map.* 28, 275–293. doi: 10.1002/hbm.20327
- Penny, W., Kiebel, S., and Friston, K. (2003). Variational bayesian inference for fmri time series. *NeuroImage* 19, 727–741. doi: 10.1016/S1053-8119(03)00071-5
- Penny, W. D., Friston, K. J., Ashburner, J. T., Kiebel, S. J., and Nichols, T. E. (2011). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier.
- Pham, D., Muschelli, J., and Mejia, A. (2022). ciftitools: a package for reading, writing, visualizing, and manipulating CIFTI files in R. *NeuroImage* 250, 118877. doi: 10.1016/j.neuroimage.2022.118877
- R Core Team (2022). *R: A Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.
- Sakamoto, Y., Ishiguro, M., and Kitagawa, G. (1986). *Akaike Information Criterion Statistics*. Dordrecht: D. Reidel.
- Smith Rachele, K. K., and Kalina, C. (2007). Localizing the rostralateral prefrontal cortex at the individual level. *NeuroImage* 36, 1387–1396. doi: 10.1016/j.neuroimage.2007.04.032
- Spencer, D., Yue, Y. R., Bolin, D., Ryan, S., and Mejia, A. F. (2022). Spatial Bayesian GLM on the cortical surface produces reliable task activations in individuals and groups. *NeuroImage* 249, 118908. doi: 10.1016/j.neuroimage.2022.118908
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., et al. (2013). The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79. doi: 10.1016/j.neuroimage.2013.05.041
- Venables, W. N., and Ripley, B. D. (2013). *Modern Applied Statistics With S-PLUS*. Springer Science and Business Media.
- Woolrich, M. W., Jenkinson, M., Brady, J. M., and Smith, S. M. (2004). Fully bayesian spatio-temporal modeling of fmri data. *IEEE Trans. Med. Imaging* 23, 213–231. doi: 10.1109/TMI.2003.823065
- Woolrich, M. W., Ripley, B. D., Brady, M., and Smith, S. M. (2001). Temporal autocorrelation in univariate linear modeling of fMRI data. *Neuroimage* 14, 1370–1386. doi: 10.1006/nimg.2001.0931
- Worsley, K. J., Liao, C. H., Aston, J., Petre, V., Duncan, G., Morales, F., et al. (2002). A general statistical analysis for fMRI data. *Neuroimage* 15, 1–15. doi: 10.1006/nimg.2001.0933



## Appendix



## Prewhitening Algorithm

**Result:** Prewhitened data values for the response and design across all vertices  $v = 1, \dots, V$

**for**  $v \leftarrow 1$  to  $V$  **do**

Find estimates  $\hat{\beta}_v$  by fitting the GLM at vertex  $v$ ;

Find AR coefficients  $\phi_{v,1}, \dots, \phi_{v,p}$  and white noise variance  $s_v$  after global or local regularization for the time series  $\mathbf{y}_v - \mathbf{X}_v \hat{\beta}_v$  using the Levinson-Durbin recursion (Brockwell and Davis, 2009);

Create the symmetric temporal precision matrix  $\mathbf{Q}_v^{-1}$  as a sparse band matrix taking the value 1 on the diagonal and  $-\phi_{v,q}$  for the  $q$ th location off the diagonal;

Set  $\mathbf{S}_v^{-1} = ((1/s_v)\mathbf{I}_T)\mathbf{Q}_v^{-1}$ ;

Find the eigenvectors and eigenvalues of  $\mathbf{S}_v^{-1}$ , where  $\mathbf{U}_v$  is the matrix of eigenvectors and  $\mathbf{D}_v$  is the diagonal matrix of corresponding eigenvalues;

Compute the prewhitening matrix  $\mathbf{W}_v = \mathbf{U}_v \mathbf{D}_v \mathbf{U}_v'$ ;

Set off-diagonal values  $W_{ij} = 0$  at row  $i$  and column  $j$  within  $\mathbf{W}_v$  if  $|i - j| > p$ ;

**end**

Create  $\mathbf{W}$  as a block-diagonal matrix for all prewhitening matrices  $\mathbf{W}_v$ ;

Output prewhitened response  $\tilde{\mathbf{Y}} = \mathbf{W}\mathbf{Y}$  and

block-diagonal prewhitened design matrix  $\tilde{\mathbf{X}}$  with elements  $\mathbf{W}_v \mathbf{X}_v$ ;

**Algorithm 1.** Prewhitening time series data in order to reduce temporal dependence in residuals. Steps in *italics* are performed using C++ for computational efficiency.