

Title:

Task activations produce spurious but systematic inflation of task functional connectivity estimates

Authors:

Michael W. Cole¹, Takuya Ito^{1,2}, Douglas Schultz¹, Ravi Mill¹, Richard Chen^{1,2}, Carrisa Cocuzza^{1,2}

¹ Center for Molecular and Behavioral Neuroscience, Rutgers University, Newark, NJ, 07102, USA

² Behavioral and Neural Sciences PhD Program, Rutgers University, Newark, NJ 07102, USA

Corresponding author:

Michael W. Cole
Center for Molecular and Behavioral Neuroscience
197 University Ave, Suite 212
Newark, NJ 07102
973-353-3249
mwcole@mwcole.net

Keywords: functional connectivity, networks, fMRI, computational modeling, method validation

Acknowledgements:

The authors acknowledge support by the US National Institutes of Health under awards R01 AG055556 and R01 MH109520. Data were provided, in part, by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: D. Van Essen and K. Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Highlights:

- Computational model shows task inflation of functional connectivity estimates
- Hemodynamic response shape causes task activations to further inflate estimates
- Standard approach to remove task activations leaves many false positives
- Methods that flexibly fit hemodynamic response shape effectively correct inflation
- Correction of functional connectivity inflation verified with empirical fMRI data

Abstract

Task-state functional connections – such as those measured using functional MRI (fMRI) – are thought to coordinate distributed cognitive processes throughout the brain. Utilizing a neural mass computational model we found that the conversion of neural signals into fMRI hemodynamic responses substantially and inappropriately inflates task-state functional connectivity (FC) estimates (temporal correlations). Such activation-induced inflation of task FC estimates was postulated previously, but this phenomenon has not been conclusively established either theoretically or empirically, leading many task FC studies to simply ignore the issue. We found that activation-based task FC inflation was primarily driven by task-evoked fMRI activations introducing a similar hemodynamic response shape to underlying neural time series. This demonstrates that isolating task-state network changes from task-evoked activations is essential for ensuring discovery of unique functional network effects, independent of mechanistically-distinct activation effects. Standard approaches for fitting and removing task-evoked activations were unable to correct these inflated correlations. In contrast, methods that flexibly fit hemodynamic response shapes (especially finite impulse response-based regression) effectively corrected the inflated correlations. Results with empirical fMRI data confirmed the model's predictions, revealing activation-induced task FC inflation for both Pearson correlation and psychophysiological interaction approaches. These results demonstrate that removal of task activations using an approach that flexibly models hemodynamic response shape is an essential preprocessing step for valid estimation of task-state FC with fMRI.

INTRODUCTION

Converging evidence across a wide variety of neuroscientific methods applied across multiple species suggests cognition emerges from widespread neural interactions (Cole et al., 2013b; Gratton, 2013; Likhtik et al., 2005; M. Siegel et al., 2015). One of the most logically-direct and widely-used ways to characterize these cognitive brain network interactions involves estimating task-state functional connectivity (FC). Rather than conflate the method (e.g., correlation) with what the method seeks to estimate, we define FC as any interaction among neural units (e.g., neurons, local neural populations, brain regions). Notably, FC estimates often lack details about directionality or whether interactions are indirect (e.g., via a third region). Nonetheless, if an FC estimate is valid it should still reflect real (potentially indirect) neural interactions. Therefore, we used a neural mass computational model – in which we could know the ground-truth neural interactions – to test the efficacy of widely-used FC estimation approaches involving functional MRI (fMRI). We further focused on task-state FC given its close relationship with experimentally-controlled cognitive variables, in which, relative to resting-state FC alone or a single task-state FC condition, quantifying changes of task-state FC relative to a control condition (e.g., resting state or another task condition) provides improved experimental control of FC inferences. We also focused on Pearson correlation as an FC-estimation method given that it is among the most widely-used fMRI FC measures. Since Pearson correlation does not indicate whether an interaction is direct or indirect we focused primarily on testing for false

positives (among other potential issues) in cases where even indirect neural interactions would be impossible.

Many advances in characterizing task-state brain network dynamics have relied on fMRI (Cole et al., 2013b; 2014; Friston et al., 1997; Krienen et al., 2014; Rissman et al., 2004a). This is likely due to a unique combination of benefits when using fMRI to estimate task-related changes in the relationship among brain activity time series – i.e., task-state FC. For instance, fMRI provides for accurate spatial localization (e.g., relative to electroencephalography), which is essential for valid task-state FC inferences. Additionally, the non-invasive nature of fMRI facilitates the efficiency of cognitive task manipulations (e.g., relative to non-human animal studies), given that humans can be rapidly instructed to perform new tasks (Cole et al., 2013a). Finally, the whole-brain coverage of fMRI (e.g., relative to multi-unit recording or intracranial electroencephalography) allows for calculating comprehensive functional network graphs (Cole et al., 2013b; Power et al., 2011). This unique set of beneficial traits belies one potentially major issue – fMRI is an indirect measure of neural activity (Logothetis and Wandell, 2004). Here we utilize computational modeling of the relationship between neural and fMRI time series to test whether fMRI-based inferences of task-state FC are likely to be valid.

Several studies have proposed the possibility that cross-event average changes in fMRI activity amplitudes – as measured by standard fMRI general linear model (GLM) analyses – could induce false positive task-state FC estimates (Cole et al., 2014; Fair et al., 2007; Friston et al., 1997; Gratton et al., 2016; Norman-Haignere et al., 2012). The argument is that simultaneous increases in brain activity induced by task events could create spurious increases in FC estimates (e.g., correlations) given that this would not be induced by neural interactions but rather by the experimenter (the task timing; **Figure 1**). For instance, presenting a visual stimulus simultaneously with an auditory stimulus would increase activity simultaneously in primary visual and primary auditory cortices. This would increase the task-state FC estimate (e.g., Pearson correlation) between visual and auditory cortices, despite there being no true task-state FC – no change in neural interaction – between those regions simultaneous with stimulus onset. Note, however, that simply calculating a Pearson correlation using the post-stimulus time series would likely fix this problem when using direct neural signals (e.g., spike rates, local field potentials) since moment-to-moment fluctuations between non-interacting visual and auditory cortices are unlikely to be correlated after an initial task-state transient. In contrast, fMRI hemodynamic responses are known to be similar (but not identical) between regions (Handwerker et al., 2004; Ogawa et al., 1992), combining a delay from task event onset with a common shape that is largely independent from moment-to-moment fluctuations. It is this task-state introduction of a common (i.e., correlated) response shape that may have the biggest impact on inflating task-state FC estimates.




Figure 1 – Illustration of the possibility that task event timing acts as a confounder (third variable), producing false positives. Task-evoked activity (activity time-locked to task event onsets) has been proposed to induce false positives in task-state FC estimates. Despite many studies assuming task timing is indeed a confounder (Greenland et al., 1999) (and many other studies that have not), to our knowledge this has never been conclusively established. We sought to rigorously test this possibility here. **A)** A scenario with a true neural interaction between neural populations A and B induced by the task context. Even with this true positive, the task-evoked activity inflates the FC estimate relative to the ground truth interaction. **B)** A scenario with no true neural interaction between A and B. The task event timing nonetheless acts as a confounder to create an artefactual correlation between the neural populations. In this case, A and B both increase their activity in response to task events, but they do not interact either directly or indirectly. We test for this possibility, along with the possibility that the fMRI hemodynamic response function further inflates the effect of task timing as a confounder (due to the similarity of hemodynamic response functions associated with distinct neural populations).

Despite this hypothesized confound, to our knowledge such task-state FC false positives have not been systematically investigated in either simulations or real fMRI data. Adding doubt that this proposed confound exists, some studies do not acknowledge this potential confound (Bassett et al., 2011; 2013; Krienen et al., 2014; Shirer et al., 2012; Tomasi et al., 2013), suggesting many researchers do not consider it to be a problem. Indeed, these studies were justified in not worrying about this putative problem, given that it has not been conclusively established in the literature (it has only been *assumed* to be a problem by some researchers). The standard approach to correct for this confound is to fit an event-averaged general linear model (GLM) of the task events either simultaneously with task-state FC (as with psychophysiological interaction; PPI) (McLaren et al., 2012; O'Reilly et al., 2012) or calculate FC estimates using the residual time series of such a model (Al-Aidroos et al., 2012; Cole et al., 2013b; Gratton et al., 2016; Summerfield et al., 2006). Critically, without showing that task-state FC estimates are meaningfully altered by these preprocessing steps, it remains possible that the proposed confound does not exist or (even if it exists) that the proposed correction for the confound is ineffective.

One potentially confusing aspect of these task-regression correction approaches is that – despite removing some task-related variance – these approaches are not meant to eliminate all task-related variance from the time series. Rather, these approaches are designed to leave moment-to-moment (and event-to-event) task-related variance in the time series, but remove cross-event variance correlated with the task's timing. This allows for a distinction between two neural populations merely responding in a similar way to a task event (task co-activation), versus an ongoing interaction between those neural populations as evidenced by covariance among moment-to-

moment activity fluctuations. In the parlance of electroencephalography, these task-state FC approaches remove cross-event mean evoked responses (responses time-locked to task events) to isolate induced responses (responses influenced by task events but that vary in timing across multiple instances of those events) (Tallon-Baudry and Bertrand, 1999). Note that evoked (event-time-locked) responses that vary in amplitude across event instances (e.g., trials) remain in addition to induced responses (Truccolo et al., 2002a). Neural time series correlations that remain after removing the cross-event mean response are termed "noise correlations" in the non-human animal neurophysiology literature (Cafaro and Rieke, 2010; Cohen and Kohn, 2011). One goal of the present study is to determine whether only removing cross-event mean (evoked) responses is adequate for eliminating task-activation-driven FC inflation.

Demonstrating that the proposed confound is real and problematic would have a meaningful impact on our understanding of task FC, given that many task FC studies use fMRI and make no attempt to correct for the proposed confound. The impact of demonstrating this potential confound would be larger to the extent that the task-state FC estimate inflations are large or numerous. The impact would be even larger if the proposed confound were real and standard methods to correct for the confound were ineffective, given that this would implicate much of the task-state FC fMRI literature in potentially-false conclusions regarding task-state FC. Critically, however, it is difficult to conclude the exact impact of a methodological error such as this post-hoc, since it involves an inflation in false positives rather than a guarantee that all results are false (Eklund et al., 2016). Thus, it would be essential to correct for the confound in ongoing and future studies (and in reanalysis of previous studies when possible), with improved understanding of task-state FC effects being the best way to estimate the impact of the confound on prior studies.

We began by testing for the existence of the proposed task-state FC confound using a neural mass computational model that balances simplicity with biological interpretability. Briefly, the computational model is based on a standard firing rate model of neural populations (Hopfield, 1984). Note that unlike some firing rate models we analyze the input time series rather than the output (population spike rate) time series. This is due to the input into each neural population better reflecting local field potentials, which have been shown to better relate to a variety of signals used to investigate task-state FC such as fMRI (Logothetis et al., 2001). Once the task-state FC confound was identified in simulated fMRI data, we tested a variety of methods to correct for the confound. Once a confound-corrected method was identified, we tested its efficacy in real fMRI data. Critically, demonstrating that this confound-correcting method has a strong impact on results with real fMRI data would provide more conclusive evidence that the confound exists and that correcting it matters in practice.

METHODS

Neural mass model

We developed a neural mass model to simulate the large-scale activity and interaction patterns of sets of thousands of neurons. We sought to optimize the model simultaneously for simplicity and biological interpretability of its internal variables. We

expected simplicity to increase understanding/interpretability of results and computational tractability, with biological interpretability facilitating the relation of simulation results to neuroimaging results. The core of the model is a standard firing rate model. This increased the simplicity of the model compared to some alternatives, while remaining biologically plausible based on evidence that neural populations exhibit a sigmoid-like transfer function reflecting variability in the exact firing threshold across individual neurons (Hopfield, 1984).

We defined each node's output as:

$$u_i(t) = f(I_i + bias) \quad i = [1..n] \quad (1)$$

where u_i is the output activity (population spike rate) for unit i at time t , I_i is the input (population field potential) as defined below, and $bias$ is the bias (population resting potential, or excitability).

$$I_i(t) = \sum_{j=1}^n G w_{ij} u_j(t-1) + d_i + stim_i \quad (2)$$

where $I_i(t)$ is the input (population field potential) for unit i at time t , G is the global coupling parameter (a scalar influencing all connection strengths), w_{ij} is the synaptic weight from unit j to i , $u_j(t-1)$ is the output activity from unit j at the previous time step (τ_{ij} was set to 1 for simplicity), d_i is spontaneous activity (independent Gaussian random values across nodes), and $stim_i$ is task stimulation (if any).

Initial condition is set to a Gaussian random value (mean 0, standard deviation 1) as input for all units at time point 0 (independent Gaussian random values across nodes).

The sigmoid $f(x)$ (population threshold) in the node output equation above is defined as:

$$f(x) = \frac{1}{1 + e^{-x}} \quad (3)$$

We reduced the arbitrariness of model parameter selection using a principled parameter search. Parameters for the model were determined based on optimizing for task-state FC change relative to resting-state FC without fMRI simulation. Specifically, we optimized for the average correlation-based task-state FC (relative to the average resting-state FC) among all pairs of the first 50 nodes in the 300-node network described in the next section. Optimizing for only a portion of the entire network reduced the chance that the optimization overfit to the particular network structure. Rerunning the model with multiple initial random conditions (for the main analyses) also ensured overfitting was not an issue. Notably, we did not optimize for task-state FC false positives nor for fMRI-based FC, such that we could test for fMRI-based false positives

as a hypothesis independent of how the model parameters were chosen. The parameter search involved all permutations of the model parameters varied, with the following ranges: $G=1$ to 10, $bias=-15$ to 0, $d=1$ to 10, $stim=0.1$ to 1.0 (in 0.1 increments).

Settings used for the model: d_i was a Gaussian random value with mean 0 and standard deviation 3, G was set to 5, $bias$ was set to -5, $stim$ was set to 0.3, and all self-connections (diagonals in w) were set to 1. Setting the self-connection above 0 reflects the theoretical neurons within the modeled neural population having synaptic connections among each other, such that the same outputs sent to other units also affect the unit that sent it. The model was implemented in Python (version 2.7).

The model's network organization

The model's network included 300 nodes, selected to be in the same range as some recent estimates of the number of functional regions in human neocortex (Glasser et al., 2016; Van Essen et al., 2012). This 300-node network was given a functional network community structure, based on empirical evidence of such large-scale network structure in the human brain (Ito et al., 2017; Power et al., 2011; Spronk et al., 2017). Briefly, the construction and running of each "subject's" network went as follows: 1) Build structural and synaptic connectivity network architecture, 2) Apply all model parameters, running both a resting-state run and a task-state run, 3) Simulate fMRI data collection by converting each node's "input" time series to fMRI via convolving with a hemodynamic response function (HRF) and downsampling the resulting time series.

Network construction involved a series of steps, with the construction of the network model randomly initialized separately for each "subject". First, there was a 10% probability of any node in the network connecting to any other. Next, three structural communities were created by increasing the probability of connectivity within each set of 100 nodes to 50%. This was then converted to a synaptic connectivity matrix by adding a Gaussian random value to each structural connection (mean of 1, standard deviation of 0.001). The first structural community was then split into two "functional" communities by multiplying the synaptic weights among the first 50 nodes (and, separately, the second 50 nodes) by 1.2 and multiplying the connections to/from the first and second 50 nodes by -0.2. Next, all connections to/from the final 100 nodes and all other nodes were multiplied by 0, completely isolating the final community from the rest of the network. Finally, each node's synaptic connectivity was normalized such that all inputs summed to 1.0. Input weight normalization is thought to be a biologically realistic process (e.g., via each neuron regulating the number of channels at each synapse) (Barral and D Reyes, 2016).

Task stimulation amplitude targeted 25 nodes in the first and last network communities. Note that the setting of the bias to -5 was consistent with units starting out at a near-0 firing rate (given the sigmoid activation function that was used), modeling most neurons within a modeled population being at a sub-threshold resting potential. Modeling conversion to fMRI data involved convolution of variable HRFs with the input time series from each node. The HRF differed for each simulated subject and each region, though it differed more between subjects than between regions, consistent with empirical evidence (Handwerker et al., 2004). Specifically, the values for peak time (3 to

9 in increments of 0.5 s), undershoot time (3 to 17 in increments of 0.5 s), and undershoot ratio (0 to 1 in increments of 0.1) of a double-gamma HRF were varied randomly (uniform distribution) by subject. Then, each region had these three parameters varied from a given subject's selected values based on a Gaussian random distribution centered on 0 with a standard deviation of 1, with that value being the array index selecting from the set of allowed values for each parameter (as indicated in the previous sentence). Note that results were similar without HRF variability (i.e., with the same non-canonical HRF shape used for all subjects and all regions). HRF convolution was followed by sampling (selecting a single time point) of the convolved time series at a time to repetition (TR) of 0.785 seconds, in the range of multiband fMRI protocols (Chen et al., 2015).

The model was implemented with 24600 time steps per “run”, with each time step conceptualized as 50 ms, such that the total simulated time was conceptualized as 20.5 minutes in duration. Each run was implemented across 24 “subjects”, with a separate random seed used for each subject for the spontaneous activity. The first run consisted of a resting-state simulation with no task stimulation. The second run consisted of a task-state simulation, with 6 task “blocks” of 2.5 minutes of constant stimulation of the two sets of nodes indicated above. There was 30 seconds of non-stimulation before and after each task block. All FC analyses used the time points included in the 6 task blocks, ignoring the inter-block periods.

FC estimation

Estimates of time series association were calculated using either MATLAB (version R2014b) or R (version 2.15.1). Pearson correlation was calculated as:

$$r = \frac{cov_{XY}}{S_X S_Y} = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}$$

Where S is the time series standard deviation, cov is the time series covariance, X and Y are brain activity time series, n is the number of time points, and \bar{X} and \bar{Y} are the time series means. Most analyses also involved the Fisher's z-transform of the resulting Pearson correlation, which increases the dynamic range of correlation values to go beyond ± 1.0 . This is critical when investigating changes in functional connectivity, as forgoing the Fisher's z-transform would result in artificial restrictions in dynamics. The Fisher's z-transform:

$$Fz = atanh(r)$$

Psycho-physiological interaction (PPI) was estimated using simple linear regression, which is equivalent to:

$$\beta = \frac{cov_{XY}}{var_X} = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sum_{i=1}^n (X_i - \bar{X})^2}$$

Where var is the time series variance. The beta for each condition was estimated separately for each condition, consistent with generalized PPI (McLaren et al., 2012).

Empirical fMRI data collection

The empirical fMRI dataset was collected as part of the Washington University-Minnesota Consortium Human Connectome Project (HCP) (Van Essen et al., 2013). These data are publicly available, accessible at <https://www.humanconnectome.org>. Participants were recruited from Washington University (St. Louis, MO) and the surrounding area. All participants gave informed consent. The data used were selected by the HCP as the "100 unrelated" dataset, consisting of data from 100 participants with no family relations. Data from 25 subjects were not used because of excessive in-scanner movement (defined as over 50% of volumes in any run with mean framewise displacement > 0.25 mm) for these subjects, such that data from 75 subjects were included in the final analyses. Framewise displacement was calculated as described by Power et al. (2012), with a low-pass filter of 0.3 Hz applied as suggested by Siegel et al. (2016) for multiband fMRI data.

Whole-brain echo-planar imaging acquisitions were acquired with a 32 channel head coil on a modified 3T Siemens Skyra with TR = 720 ms, TE = 33.1 ms, flip angle = 52°, BW = 2290 Hz/Px, in-plane FOV = 208 × 180 mm, 72 slices, 2.0 mm isotropic voxels, with a multi-band acceleration factor of 8 (Ugurbil et al., 2013). Data were collected over two days. On each day 28 minutes of rest (eyes open with fixation) fMRI data across two runs were collected (56 minutes total), followed by 30 minutes of task fMRI data collection (60 minutes total). Each of the 7 tasks was completed over two consecutive fMRI runs. Resting-state data collection details for this dataset can be found elsewhere (Smith et al., 2013), as can task data details (Barch et al., 2013).

Empirical fMRI dataset analysis

The empirical dataset preprocessing consisted of standard functional connectivity preprocessing (typically performed with resting-state fMRI data), with several modifications given that analyses were also performed on task-state data. Resting-state and task-state data were preprocessed identically to facilitate comparisons between them. Spatial normalization to a template (MSM-sulc), motion correction, intensity normalization (normalized to a 4D whole brain mean of 10,000) were already implemented in a minimally-processed version of the empirical fMRI dataset described elsewhere (Glasser et al., 2013), so we began preprocessing with this version of the data. With the surface (rather than the volume) version of the minimally preprocessed data, we used custom scripts in MATLAB to additionally remove nuisance time series (motion, ventricle, and white matter signals, along with their derivatives) using linear regression, and remove the linear trend for each run. Note that the main results were broadly similar with and without whole brain (global) signal regression. Unlike standard resting-state functional connectivity preprocessing a low-pass temporal filter was not applied. This was due to the possible presence of task signals at higher frequencies than the relatively slow resting-state fluctuations.

Data were sampled from a set of 360 brain regions (rather than individual voxels/vertices) to make inferences at the region and systems levels. We used an independently-identified set of putative functional brain regions (Glasser et al., 2016) so as to reduce any potential circularity in analyses (Kriegeskorte et al., 2009). The use of this parcellation also reduces the chance of combining signals from multiple functional

regions as compared to anatomically-defined parcellations (Wig et al., 2011). These brain regions were identified using parcellation of a variety of data types, including resting-state functional connectivity, task activation, and myelin maps (Glasser et al., 2016). Data were summarized for each region by averaging signal in all vertices falling inside each region.

Preprocessing was carried out using Freesurfer (version 5.3.0-HCP), FSL (version 5.0.8), and custom code in MATLAB 2014b (Mathworks) for the 7-task dataset (using the minimally preprocessed version of the data (Glasser et al., 2013)). Further analysis was carried out with MATLAB and R.

We estimated FC using Pearson correlations and regressions between time series from all pairs of brain regions using MATLAB (version R2014b). For Pearson correlations, all computations used Fisher's z-transformed values. FC estimation was straightforward for resting-state data, as there were no additional steps after preprocessing prior to calculating these values. For task data there were additional steps related to task activation regression, as described in the following section.

FC differences were assessed using two-way t-tests paired by subject. Multiple comparisons were corrected for using false discovery rate (FDR) (Genovese et al., 2002). When comparing task-state FC to resting-state FC estimates the number of time points contributing to those estimates were matched. The beginning of the first resting-state fMRI run was used in all cases, due to the increased likelihood of subjects falling asleep later in the rest run (Tagliazucchi and Laufs, 2014).

Task-activation regression for task-state FC

Cross-event (trial and block) mean activations during task fMRI might unduly influence task-related changes in FC. This was rigorously tested using computational modeling, which informed our empirical fMRI data analysis. We sought to suppress or remove such influences with task regression techniques. This involved running standard fMRI general linear model (GLM) analysis, and calculating FC based on the residuals. Specifically, each region's task time series was modeled using a GLM, with a distinct model depending on the analysis (as described below). To improve removal of task-related activation variance, a separate regressor was included for each task condition (e.g., face stimuli vs. tool stimuli in the N-back task; 24 task conditions total). Note that regressing out task events using GLM primarily removes the cross-event signal means, retaining event-to-event and sub-event fluctuations in time series such that these sources of variability likely contribute the most to task FC estimates (Rissman et al., 2004a; Truccolo et al., 2002b). The residuals from this regression model were used for FC estimation, restricted to time points corresponding to the current task. A standard hemodynamic lag was included when determining task timing, by convolving the timing with a canonical HRF and selecting time points with a value above 0.

FC estimation was conducted along with no task regression, canonical HRF task regression, constrained basis set task regression, or finite impulse response (FIR) task regression. Other than the task regression step, all steps were identical in the no-task-regression case as when task regression was used. Canonical HRF task regression involved use of the SPM software function `spm_hrf.m` with the default parameters to create the HRF. This HRF was then convolved with each of the 24 task condition time series, then fit using ordinary least squares regression in MATLAB (function `regress.m`).

Constrained basis set task regression involved creation of a set number of basis set regressors (either 5 or 28) in the FLOBS interface in FSL software (version 5.0.8; default parameters) (Woolrich et al., 2004). Note that the first three basis function regressors are highly similar to the canonical, time, and dispersion derivatives often used together to model task activations in SPM software (Woolrich et al., 2004). These basis set functions were then convolved with each of the 24 task condition time series before fitting them to the brain region time series (identically to the canonical HRF approach).

The FIR task regression approach involved fitting the cross-trial/cross-block mean response for each time point in a set window length that is time-locked to the trial/block onset for a given task condition. This allows the fit to be completely flexible with regard to the HRF response shape, so long as it is consistent across trials/blocks for that condition. Each of the 24 task conditions were fit with a series of regressors, one per time point. Each condition's window length matched the duration of the events, with an additional 18 s (25 regressors) added to account for the likely duration of the HRF.

There were some noteworthy issues with the timing of the HCP tasks that were especially relevant for FIR modeling. First, in some cases not all task events were the same duration for a given condition. For instance, some events were cut off at the end of a run. In such cases we cut off (or extended) the FIR window length to match the duration of the individual event. Simulations indicated that this properly removed cross-event variance, though of course fewer events contributed to the estimates of the later time points in the FIR window. Notably, this choice to model time points that had a smaller number of contributing events might differ from standard FIR approaches, which primarily seek to estimate robust activations rather than remove task-evoked activation variance. Second, the exact timing of the modeling was extremely important and was somewhat difficult given the many timing variants across the HCP tasks. The first version of our FIR analysis actually resulted in worse task-state FC inflation than all other methods, due to errors in the timing of the FIR regressors. This should be considered when using FIR regression, perhaps in terms of double-checking the regressor timings and comparing to canonical-HRF regression results (e.g., to make sure FIR reduces task-state FC more than canonical-HRF regression). For instance, we plotted the canonical-HRF and FIR timings and compared them carefully to make sure they lined up, since it was easier to set the timings for the canonical-HRF GLM (given that there was only one regressor per task condition).

RESULTS

Testing for false positives using a neural mass computational model

Many studies have reported task-state changes in FC using large-scale neuroimaging methods such as fMRI and EEG/MEG (Cole et al., 2014; Friston, 2011; 1994; Krienen et al., 2014; Mill et al., 2017). We sought to determine the efficacy of standard task-state FC estimation methods using a simple neural mass model with biologically interpretable parameters. We began with a standard firing rate neural mass

model, with a standard sigmoid transfer function relating network nodes on each time step (see Methods).

Unlike some firing rate neural mass models we focused on the “input” signal, given that it is analogous to a neural population’s local field potential. Variance in this input signal is thought to be the mechanistic basis of the blood-oxygen-level dependent (BOLD) signal underlying fMRI (Logothetis et al., 2001) as well as the electromagnetic fields underlying EEG/MEG signals. Thus, the present neural mass model is likely more compatible with simulating these methods than standard firing rate models, similar to others that have used input (synaptic activity) estimates for simulating fMRI data in the past (AdrianPonceAlvarez:2015fo; Schirner et al., 2016). Note, however, that each node is given self-connections to simulate the effect of neurons within each modeled neural mass connecting to one another (which, in addition to being nearby spatially, is what defines them as part of the same neural mass). This makes each node’s output feed back into its own input on the next time step, making the input and output signals similar in many circumstances.

We constructed a series of large-scale network communities, given the presence of such communities in many real-world networks (Girvan and Newman, 2002) including the human brain (Power et al., 2011; Yeo et al., 2011). We began by making three structural communities of 100 nodes each (**Figure 2A**; see Methods). Importantly, we removed all structural connections to/from the last community, allowing us to test for false positives in subsequent analyses (see upper-right corner of **Figure 2A**).

We took a principled approach to model parameter selection, optimizing all model parameters for high task-state FC (relative to resting-state FC) among the first 50 nodes of the model (see Methods). This was done by performing an exhaustive parameter search including a wide range of parameter values. Note that restricting the search to the first 50 nodes (and running the search without fMRI simulation) helped ensure independence of model parameter selection from the hypothesis that task-evoked fMRI activity could induce false positives within the “no connectivity” zone.




Figure 2 – The 300-node model, with fMRI simulation. The same model used in the 3-node simulation was used in a 300-node simulation. **A)** Three structural communities were constructed (100 nodes each), with the first community split into two communities via synaptic connectivity. The first and second structural communities had random connectivity (10% density), while the third community had no connections with the rest of the network. Connections to/from the third community acted as tests for false positives in subsequent simulations. **B)** We simulated fMRI by convolving the input time series of each unit with a hemodynamic response function (HRF) and downsampling (every 785 ms). Spontaneous activity without task stimulation was used to produce this FC matrix. T-tests vs. 0 were based on across-subject variance, with each “subject” being a random initialization of the synaptic connectivity matrix and spontaneous activity. Note the low false positive rate (0.81%). **C)** Two populations of 25 nodes (indicated by yellow stars) were stimulated simultaneously across 6 task blocks. Two completely unconnected communities were stimulated to test for false positives. Note the increase in false positives (41%). **D)** T-tests indicated an inflated false positive rate of 40% when comparing task FC to rest FC. Note that without fMRI simulation (i.e., no HRF or downsampling) the false positive rate was 1.4%.

We next used the neural mass model to simulate collection and analysis of resting-state FC with fMRI (**Figure 2B**). We began by simulating 20.5 minutes of resting-state data, driven by spontaneous activity only. This was repeated for a total of 24 “subjects” (random initializations of the synaptic connectivity matrix and spontaneous activity). Simulation of fMRI data collection involved the input (population field potential) time series convolved with an HRF and down sampled by a standard multiband fMRI TR of 0.785 s. Note that the HRF varied for each region and for each subject, with more variation between subjects than between regions, as is the case in empirical fMRI data (Handwerker et al., 2004). We used Pearson correlation for estimating FC to make the results more directly relevant to typical practice in fMRI FC studies (Zalesky et al., 2012). We found that the resting-state FC matrix was significantly similar to the large-scale structure of the synaptic connectivity matrix (mean $r=0.47$, $t(23)=266$, $p<0.00001$). Further, there were minimal false positives (0.8%) in the “no connectivity” zone at a t-test threshold of $p<0.01$. Note that there is no need to correct for multiple comparisons in the computational model analyses, given that we know the ground truth and the p-value indicates the expected percentage of false positives. Given that we use $p<0.01$,

we interpret false positives greater than 1% of the tested connections to be true false positives (with false positives at or below 1% likely being due to multiple comparisons).

We then simulated task-state FC by stimulating two sets of units, in the first and last functional communities (**Figure 2C**). Task stimulation consisted of a small constant input (0.3) across 50 nodes (25 for each of the two communities) in six task blocks. Each block lasted 2.5 minutes with 30 seconds of non-stimulation between each block, with a total duration of 20.5 minutes of simulated time. Only the on-stimulation times were analyzed for task FC to reduce the influence of on/off task transients. Task-state FC over-and-above resting-state FC was widespread (**Figure 2D**). This is consistent with the observation of task-state FC across a wide variety of brain systems and tasks in the fMRI literature (Cole et al., 2014; Krienen et al., 2014). However, it was apparent that a large number of false positives were present in the “no connectivity” zone: 42.58% false positives for task vs. rest FC ($p < 0.01$). This appeared to be driven primarily by the fMRI simulation, since the false positive rate was only 1.99% (task vs. rest, $p < 0.01$) with the same data prior to fMRI simulation.

Testing for correction of the false positive rate with fMRI

Such a high false positive rate is consistent with the possibility – hypothesized as part of several fMRI studies in the past – that task co-activation with fMRI could lead to false positives in task-state FC estimates (Cole et al., 2014; 2013b; Fair et al., 2007; Friston et al., 1997). Notably, however, these studies did not necessarily hypothesize that fMRI data were especially problematic in this respect, as we show here. We nonetheless treated the approaches those studies used to correct the false-positives problem as a guide here. These typically involve regressing out the task timing, which involves using the residuals of a GLM as the time series to compute task-state FC. This is very similar to simultaneous fitting of task-state FC and task activations when using PPI (Friston et al., 1997; McLaren et al., 2012; O'Reilly et al., 2012). Critically, however, since our simulations provided “ground truth” knowledge of the false positive rate we were able to validate the approaches and verify their efficacy for reducing the false positive rate.

There remains much confusion in the literature about the nature of these task-regression correction approaches. It is often assumed that these approaches will eliminate all (or a substantial portion of) task-related variance from the time series, but this is mistaken. Rather, these approaches are designed to leave moment-to-moment (and event-to-event) task-related variance in the time series, but remove cross-event variance correlated with the task's timing. This allows for task-state FC to remain (due to moment-to-moment and event-to-event task-related variance remaining), while removing the variance directly correlated with the task timing.

We began by using the most common approach for reducing false positives – fitting the “canonical” HRF shape to remove cross-event mean variance correlated with task timing (**Figure 3A**). This is the same HRF shape used in PPI (O'Reilly et al., 2012) and related approaches (Cole et al., 2014). Note that the variable HRF shapes used across regions and subjects were all based on the canonical HRF, such that the canonical HRF would be expected to fit the mean responses relatively well. We found

that task regression with this canonical HRF shape reduced the false positive rate somewhat but failed to bring it below the 1% specified by the p-value threshold ($p < 0.01$): 20.34% false positive rate (**Figure 3B**).




Figure 3 – Testing task-timing regression approaches to reduce false positive rate. While some researchers investigating task FC fMRI ignore this problem, there are several standard approaches for attempting to reduce potential false positives. Critically, the 300-node computational model can provide a ground-truth scenario for testing the validity of these approaches. Note that all approaches are designed to leave moment-to-moment (and event-to-event) task-related variance in the time series, but to remove cross-event variance related to the task's timing. Task vs. rest Pearson correlation differences (t-test $p < 0.01$ thresholded) are shown. **A)** The 4 tested approaches are illustrated. The canonical HRF shape is what is typically used to reduce false positives in the literature, as with PPI. To assess whether the HRF shape mattered a “wrong” HRF was also used. The finite impulse response (FIR) and constrained basis set approaches are flexible, allowing them to fit the actual HRF shape. **B)** The canonical HRF shape task regression. There was a reduction from the no-regression condition (42.58%) but the remaining high false positive rate (20.34%) demonstrates that task regression with the canonical HRF is helpful but fails to correct the problem. **C)** Task regression with the “flipped” HRF shape does a worse job at correcting the problem (25.43% false positive rate). **D)** Task regression with the FIR approach eliminates the problem, with the false positive rate just below the expected detection rate of 1% (given our $p < 0.01$ threshold). **E)** Task regression with a basis set of 5 regressors (accounting for 99.5% of the variance among 1000

plausible HRF shapes) was also successful in reducing the false positive rate (1.05%). **F**) False positive rates across six variants of the analyses. Since results were thresholded at $p < 0.01$, any values above 1% can be considered false positives.

We next used a worst-case scenario “wrong” HRF shape to determine if having an approximately-correct HRF shape (as with the canonical HRF) mattered for reducing the false positive rate. This “wrong” shape was created by moving the HRF undershoot in the canonical HRF to the beginning instead of the end of the HRF shape (**Figure 3A**). We found that using the wrong HRF shape did a worse job of reducing the false positive rate than the canonical HRF (**Figure 3C**): 25.43% false positive rate. This suggests that the relative accuracy of the HRF shape matters.

A standard approach for empirically determining the correct HRF shape for task regression is finite impulse response (FIR) modeling (Cordova et al., 2016; Fair et al., 2007; Norman-Haignere et al., 2012). This involves including a binary regressor for every time point in the task event/block (**Figure 3A**). This is sometimes referred to as “background connectivity” analysis when used with a block experimental design (as here) (Cordova et al., 2016; Norman-Haignere et al., 2012). As expected, we found that FIR modeling successfully reduced the false positive rate below the 1% specified by the p-value threshold ($p < 0.01$): 0.94% false positive rate.

The success of the FIR approach suggested that flexibly fitting each region's (for each subject's) HRF shape was critical for correcting the false positive rate. We next tested this hypothesis more fully by using an alternative approach that also flexibly fits HRF shapes, but with fewer regressors. This approach – the constrained basis set approach (Woolrich et al., 2004) – involves reducing many plausible HRF shapes (variants on the canonical HRF) to a select set of basis functions using singular value decomposition. We chose this approach for theoretical reasons: It can be conceptualized as a regularized approach that reduces variance in model fitting, moving away from the FIR model's high variance in the bias-variance trade-off. Unlike the extremely high bias implicit in assuming an exact HRF shape, however, the constrained basis set approach involves a restricted search for the HRF to a set of plausible HRF shapes. We used a standard approach to producing the basis set implemented in the FSL software package (Woolrich et al., 2004). This involved varying the parameters in a standard equation for producing HRF shapes to produce 1000 plausible HRF shapes. This set of HRF shapes were then reduced to five components that account for 99.5% of the variance using singular value decomposition. Note that the first three regressors were highly similar to the canonical, temporal derivative, and dispersion derivative regressors (respectively) commonly used with SPM software (Woolrich et al., 2004).

Consistent with our hypothesis, we found that the constrained basis set approach also reduced the false positive rate to the level expected with the p-value threshold ($p < 0.01$) (**Figure 3E**): 1.05% false positive rate. These results suggest that the constrained basis set approach was able to fit unknown HRF shapes in a manner similar to FIR modeling.

Checking for false negatives due to cross-event mean task regression

Given that the approaches involving more regression parameters did better, it is possible that the reduction in false positives was due to removing variance generally (rather than just the variance associated with false positives). This possibility predicts that the FIR and constrained basis set approaches would inflate false negatives along with reducing false positives. We tested this possibility using the pre-fMRI results as "ground truth" FC – Pearson correlations among the time series prior to HRF convolution and downsampling. We focused on the entire 300 x 300 FC matrix for these tests (rather than just the no-connectivity zone), since we needed some true positive effects to test for false negatives. We hypothesized that the FIR and basis set approaches reduced false positives without inflating false negatives.

We began by setting a baseline by comparing no-task-regression fMRI FC estimates to no-task-regression pre-fMRI FC estimates. This isolated the effect of the fMRI simulation on the FC results, given that fMRI simulation was the only difference between these two conditions. We found a 15.04% false negative rate (along with a 18.50% false positive rate) for no-task-regression fMRI FC relative to no-task-regression pre-fMRI FC. Note that the false negatives were likely due to both the HRF convolution (similar to a low-pass filter on the time series) and time series downsampling reducing the number of data points contributing to the analysis. Based on this, a 15.04% or lower false negative rate when using the FIR or basis set approach would indicate that these approaches did not increase the false negative rate.

As expected, the false negative rate for the FIR and basis set approaches were both below 15.04%: 13.94% for FIR and 13.15% for basis set. These results suggest that the FIR and basis set approaches removed variance that was inappropriately altering FC estimates, both in terms of false positives and false negatives. Note that, when using the entire FC matrix (rather than just the no-connectivity zone), the false positive rate dropped from 18.50% for no-task-regression to 0.60% for FIR and 0.71% for basis set approaches – smaller false positive rates than observed when focusing solely on the no-connectivity zone. Together these results suggest that the extra regression parameters included in the FIR and basis set approaches are unlikely to reduce false positives by also reducing true effects (and that they can actually increase detection of true effects).

Factors driving the false positive rate: Task-state FC inflation occurs to the extent that HRF shapes are similar across regions/voxels

We next sought to isolate major factors underlying the observed inflation in task-state FC estimation following fMRI simulation. To facilitate isolating the primary cause of the inflation, we first sought an extremely simple demonstration of the inflation effect. This involved creating two Gaussian random time series with very low correlation ($r = -0.01$), followed by adding a value of 2 for the second half of both time series (**Figure 4A**). This can be thought of as an increase in activity/excitability for both "nodes". Unlike the above simulations this manipulation was extremely simple, allowing isolation of causes of any observed inflation. Pre-"task" correlation was $r = -0.03$ and $r = 0.01$ during

“task”. We next convolved the exact same time series with a canonical HRF (**Figure 4B**). This did not substantially change the pre-“task” correlation ($r=-0.08$), but the “task” correlation was highly inflated ($r=0.98$). This demonstrates that the correlation inflation is mostly driven by an interaction between an increase in time series amplitude and HRF convolution.




Figure 4 – Highly simplified model (involving even fewer assumptions than the 300-node model): Convolution of simple Gaussian random time series demonstrates that fMRI task-state FC inflation is primarily driven by HRF convolution. **A)** Two Gaussian random time series were generated. At time point 500 a constant value of 2 was added to all subsequent time points to simulate a sustained “task” activity increase in both time series. The Pearson correlation was low both before ($r=-0.03$) and during ($r=0.01$) the “task” time period. Note that the entire time series went from a correlation of $r=-0.01$ prior to the constant value being added to $r=0.49$ after the constant value was added. This suggests initial transients driven by task activity increases can drive FC even in non-fMRI data, if the task period isn’t fully separated from the pre-task period (e.g., due to autocorrelated time series). **B)** The exact same time series as in panel A were convolved with a standard HRF. The correlation between the time series was now much higher ($r=0.98$) during the “task” time period. **C)** As expected, GLM-based task regression with the known HRF shape reduced the task-evoked correlation increase to a near-original level ($r=0.07$).

We next sought to test if task GLM regression could reduce the correlation inflation, as shown in the more complex 300-node model. This involved fitting a GLM with the known HRF shape for each time series, then using the residuals to compute the correlation (**Figure 4C**). As expected, the “task” correlation was substantially reduced ($r=0.07$). This further demonstrates the efficacy of task regression for reducing task-amplitude-induced correlation/FC inflation.

Note that the pre-HRF-convolved time series (**Figure 4A**) went from an original whole-time-series correlation of $r=-0.01$ to a whole-time-series correlation (i.e., not restricted to just the “task” portion of the time series) of $r=0.49$ when the constant value of 2 was added to the second half of each time series. This demonstrates the importance of isolating task from non-task time periods when calculating task-state FC, since transient rest-to-task activity transitions can drive correlation increases. Further, this suggests the potential utility of regressing out task transitions even from non-fMRI time series if one wants to avoid simple co-activations masquerading as FC network reconfigurations. Finally, this result suggests much of the HRF-induced correlation inflation effect is likely driven by the coincident activity increase being delayed in time by HRF convolution (see delay in task-driven activity increases in **Figure 4B** relative to **Figure 4A**). Consistent with this, if we delayed the task time window (the time points

used to calculate task FC) by 12.5 s, the correlation dropped from $r=0.98$ to $r=0.87$. Note that this drop was often higher with different random initial conditions, but still typically retaining the inflated correlation (mean across 10 random initial conditions: $r=0.37$, $t(9)=1.90$, $p=0.09$ vs. $r=0$). This suggests that inflation of task FC could be reduced by introducing a delay relative to task event onsets, but that mean task regression is likely a much more effective approach.

We next sought to determine the role of other factors that could reduce this correlation-inflation effect. In particular, we reasoned that real fMRI involves some variability in the HRF shape across regions/voxels (Handwerker et al., 2004), along with additional sources of noise (physiological and magnetic resonance noise). We simulated the effect of HRF variability by using our previously-chosen “wrong” HRF shape (see **Figure 3A**) for convolving with the second time series in **Figure 4A**. This reduced the post-convolution “task” correlation from $r=0.98$ to $r=0.30$. This reduced correlation reflects the low similarity between the “correct” HRF and “wrong” HRF shapes: $r=-0.66$. This demonstrates that the task-state FC inflation effect occurs to the extent that HRF shapes are similar across regions/voxels. Since HRF shapes are relatively similar across regions (varying more across subjects than regions) (Handwerker et al., 2004), we expect a positive-correlation inflation in real fMRI data. Nonetheless, we expect the inflation effect to be somewhat smaller in real fMRI data than the Figure 4 simulation given that HRF shapes were identical across regions in this simulation. Note that we included HRF variability in the 300-node simulations.

We next focused on likely effects of physiological and magnetic resonance noise on task-induced correlation inflation. This involved repeating the above analysis of the “task” time series correlation (see **Figure 4B**) with Gaussian random noise added to the “correct” HRF shape for one of the time series. The amplitude of the added HRF noise varied from 0.5 to 1 to 2. The correlation between the “task” time series decreased as HRF noise increased: $r=0.93$ to $r=0.80$ to $r=0.18$. This again demonstrates, this time in a more generalized sense, that the task-state FC inflation effect occurs to the extent that HRF shapes are similar across regions/voxels, with a variety of noise and HRF variability factors potentially reducing the effect. Notably, these sources of HRF variability and fMRI noise likely also reduce the ability to detect real task-state FC changes. Thus, these factors likely inflate false negatives while reducing the chance of false positives.

Testing the efficacy of false-positive-reduction approaches in empirical fMRI data

We next sought to test the ability of the FIR and constrained basis set approaches to reduce task-state FC false positives relative to other standard false-positive-reduction approaches. Unlike the computational model, we did not know the “ground truth” here, so we had to rely on any reduction in detected task-state FC as a proxy for false-positive reduction. Importantly, the FIR and constrained basis set approaches are unlikely to create false negatives given that they did not inflate the false negative rate in the computational model.

Using the Washington University-Minnesota Human Connectome Project (HCP) dataset (“100 unrelated”), we calculated cortex-wide FC across seven highly distinct

tasks in 100 healthy young adults. A set of 360 functionally-defined nodes were used (Glasser et al., 2016). Without task regression the percentage of connections that increased from resting-state FC to task-evoked FC (false discovery rate corrected for multiple comparisons) was 7.22% across the seven tasks (**Figure 5**). Only slightly reduced values were found for task regression with the canonical HRF approach (4.89%). Critically, there were substantial reductions in the percentage of task-state FC increases when using the FIR (2.48%) and constrained basis set (3.01%) approaches. These results suggest that the model likely showed a “worst case” scenario, but that false positives can nonetheless almost triple the rate of detected task-state FC changes when an effective task regression approach is not used.




Figure 5 – Analysis of empirical fMRI data suggests both false positives and false negatives in task-state FC estimates. **A)** The cross-7-task average rate of significant task-state FC increases from resting-state FC are shown (using Pearson correlation, FDR corrected for multiple comparisons, $p < 0.05$). To the extent that the FIR approach eliminates false positives, the percentages suggest a false positive rate of 65.5% without task-regression preprocessing, and 49.2% if canonical HRF or 17.3% if constrained basis set model approaches were used. There were 2.9 times more significant FC increases without task regression compared to when FIR task regression was used. **B)** The cross-7-task average rate of significant task-state FC decreases from resting-state FC are shown (using Pearson correlation, FDR corrected for multiple comparisons, $p < 0.05$). **C)** Example task vs. rest difference: Statistically significant N-back task-state FC vs. resting-state FC differences ($p < 0.05$, FDR corrected for multiple comparisons) are shown. No task regression preprocessing was used. Note the increases in task vs. rest FC within the

frontoparietal network, which is a network with high activation amplitudes during N-back task performance. **D**) Results when using the canonical HRF for task regression during preprocessing ($p < 0.05$, FDR corrected). **E**) Identical processing as in panel A, except that constrained-basis-set task regression was used during preprocessing. **F**) Results when using FIR modeling for task regression during preprocessing ($p < 0.05$, FDR corrected). Note, for example, that the within-frontoparietal-network task-rest FC increases are no longer present. This demonstrates that scientific inferences (here, whether FC increased or not) can be fundamentally altered by which preprocessing step is used prior to task FC analysis.

These effects were relatively consistent across the seven tasks performed by each participant, despite differences in timing, duration, and cognitive processes across the tasks. The percentage of connections with task-state FC increases from resting-state FC (false discovery rate corrected for multiple comparisons, $p < 0.05$) for each of the seven tasks without task-regression preprocessing was: 2.0% (emotion task), 4.3% (gambling task), 8.9% (language task), 4.6% (motor task), 7.9% (reasoning task), 14.7% (social task), and 7.7% (working memory task; **Figure 5C**). In contrast, for the FIR approach the rate of task-state FC increases were: 0.8% (emotion task), 2.2% (gambling task), 3.0% (language task), 1.9% (motor task), 3.5% (reasoning task), 2.1% (social task), and 3.1% (working memory task; **Figure 5D**). Thus, there were fewer task-state FC increases for every task when using the FIR approach, demonstrating consistency in this result.

There were also effects of the FIR approach on task-state FC decreases from resting-state FC (**Figure 5B**). Consistent with task-state FC being inflated positively without correction, the FIR approach identified a somewhat larger (but overall similar) number of task-state FC decreases from resting-state FC. This was apparent from the shift from 20.03% of FC decreases without task-regression preprocessing to 21.45% with FIR task-regression preprocessing. This was consistent with all but three of the 7 tasks individually, as most tasks showed more decreases for task-state FC relative to resting-state FC. No-task-regression decreases: 20.0% (emotion task), 25.0% (gambling task), 22.5% (language task), 7.6% (motor task), 20.1% (reasoning task), 15.4% (social task), and 28.5% (working memory task; **Figure 5C**). In contrast, for the FIR approach the rates of task-state FC decreases were: 19.7% (emotion task), 33.9% (gambling task), 14.3% (language task), 11.1% (motor task), 26.0% (reasoning task), 11.3% (social task), and 33.6% (working memory task; **Figure 5F**). Similar results were found when using the constrained basis set approach, though with even more task-state FC decreases (27.39% on average). It is unclear why there were more FC decreases with the basis set approach than the FIR approach (it will be important for future research to determine the cause of this difference). Focusing on the FIR approach (given its greater flexibility for fitting HRF shape), these results suggest that task timing regression results in a similar number of task-state FC decreases from resting-state FC compared to when no task regression is used.

We next assessed the impact of the number of basis functions on the constrained basis set approach. For this analysis we increased the number of basis functions from 5 to 28 – to a number of regressors in the general range as the FIR approach (which included the number of TRs per block type + 25 regressors). Consistent with more basis functions accounting for more cross-event mean variance, we found that task-state FC increases dropped from 3.01% with 5 basis functions to 2.67% with 28 basis functions. Similarly, we found that task-state FC decreases

dropped from 27.39% with 5 basis functions to 26.66% with 28 basis functions. Together these results suggest that including more basis functions could be useful for reducing false positives when using the constrained basis set approach. More broadly, these results suggest FIR modeling may be more appropriate for reducing false positives in general, given that the fewest positive effects were identified using FIR modeling (even compared to the 28 basis function approach). This conclusion is further supported by the computational model results indicating that the false negative rate was not inflated by the additional regressors included with the FIR approach.

Having identified FIR as the preferred method, we next quantified the amount of likely task FC inflation by comparing the no-task-regression task FC estimates versus with FIR-based task-timing regression. **Figure 6** plots these statistically significant ($p < 0.05$, FDR corrected) differences for all seven tasks individually. The percentage of connections with significant ($p < 0.05$, FDR corrected) differences for each task were, respectively (increases/decreases): 10.9/7.2, 36.0/8.9, 27.9/25.0, 29.7/3.6, 41.5/13.2, 49.8/15.9, 28.3/8.8. These results demonstrate that task-timing regression matters in practice, as it significantly alters task-state FC estimates across a broad variety of brain regions across a broad variety of task manipulations.

Visualizing the relationship between task co-activation and task-state FC inflation

We next sought to visualize the correspondence between mean task-state co-activation patterns (as estimated using GLM analysis) and task-state FC inflation. This relationship was already established in previous sections, based on both a theoretical model and empirical results. Here we sought to further quantify and visualize this relationship to help further empirically establish its robustness.

First, we calculated task-state FC inflation as the difference between no-regression task FC and FIR-regressed task FC. We then visualized this difference for all connections for an example task – the "working memory" HCP task (**Figure 7A**). The working memory task was chosen as the example task due to there being more data per subject for that task than the others (increasing statistical power). This revealed that much of the task-state FC inflation was related to visual network connections, consistent with this being a task involving visual stimuli. Notably, not all connection changes were positive, suggesting the possibility that (among other possibilities) co-activations in the opposite direction (e.g., a positive activation for one region and a negative activation in the other) could lead to artificial FC reductions. We verified that this is a likely explanation for FC reductions by visualizing the FC inflation results alongside the actual activation pattern (**Figure 7A**). Specifically, it appeared that negative activation in default-mode network regions led to under-estimated FC with the positively-activated visual network regions. Note that we did not expect an exact correspondence between activations and the task-state FC inflations given that the activations were estimated using a standard GLM with a canonical HRF (for ease of interpretation).




Figure 7 – Visualizing the relationship between task co-activation and task FC inflation. **A)** Task-state FC inflation is shown (left) by subtracting the group-mean FIR-regressed task FC matrix from the group-mean non-regressed task FC matrix. An example task – the HCP "Working memory" task (which involves visual stimuli and button pressing) – is used for illustration (with no thresholding). The FC inflation values were summed (after taking the absolute value) by region to summarize the degree to which each region showed FC inflation. This was then compared with the task-evoked activation pattern (estimated using a standard GLM with a canonical HRF shape), showing a significant correspondence (Spearman rank $\rho=0.49$, $p<0.0001$). This provides a way to visualize the degree to which co-activation patterns are likely influencing task FC patterns. **B)** The group-mean task activation pattern was used to predict likely inflation of task-state FC estimates driven by co-activations. This involved multiplying each activation with all others in a pairwise manner, converting the activation vector into a co-activation matrix. There was a significant similarity between the co-activation matrix and the task-state FC inflation (Spearman rank $\rho = 0.60$, $p<0.0001$). This shows an alternative way to visualize the degree to which co-activation patterns are likely influencing task FC patterns.

We next sought to create a simple summary of the task-state FC inflation by region, so it could be compared directly to the task activation pattern. This involved summing the task-state inflation values by region (i.e., summing across all the columns in the task-state FC inflation matrix for each row), after taking the absolute value for

each number. This is visualized for the example task in **Figure 7A**. We found that this simple summary correlated significantly with the actual activation pattern for all seven tasks, all $p < 0.00001$ except for Task 3 (the "language" task; $p = 0.0003$). The Spearman rank correlation rho values for each task, respectively: 0.35, 0.31, 0.19, 0.47, 0.56, 0.56, 0.49. These results demonstrate the robustness of the association between task-evoked activations and task-state FC inflation.

To further illustrate the relationship between activation and task FC inflation, we next sought to create a simple prediction of task FC inflation based only on the co-activation pattern. Task-state co-activation inflation was conceptualized simply as the pairwise product of the task GLM estimates. Multiplying the activation values results in cases wherein large positive co-activations are expected to create the largest increases in task FC estimates. In contrast, co-activations in the opposite direction (e.g., a positive activation and a negative activation) are expected to cause task FC estimate decreases. This sort of prediction is visualized for the example task in **Figure 7B**, showing robust correspondence with the actual task-state FC inflation pattern (**Figure 7A**). This correspondence between the predicted and actual task-state FC inflation was statistically significant across all seven tasks (all $p < 0.00001$). The Spearman rank correlation rho values for each task was, respectively: 0.51, 0.74, 0.04, 0.28, 0.67, 0.60, 0.60. Note that the third ("language") task was still statistically significant despite having a small effect size, given the large N when comparing entire FC matrices (64,620). These results further demonstrate the robustness of the association between task-evoked activations and task-state FC inflation, this time by starting from the co-activation patterns to show how even complex patterns of FC can be driven by activation-based inflation. Note that we did not expect exact correspondence between the predicted and actual task-state FC inflations, given that (among other factors influencing FC inflation) HRF shape is known to vary across regions, which likely adds noise and reduces FC inflation (see above analysis with the highly simplified model).

Even in the no-task-regression case, task-state FC is primarily driven by moment-by-moment rather than cross-block mean variance

One potential concern with the task timing regression approach is that it removes the very cause of task-state FC of interest. The computational model already demonstrated that this is not the case, since removing cross-event mean activity did not induce false negatives. Nonetheless, there might be some concern that the task timing regression approach removes the primary source of task-state FC effects in empirical fMRI data. We tested this possibility by comparing the amount of task-state-FC-driving variance removed by task timing regression. We expected that most of the task-state-FC-driving variance would remain after this preprocessing step, consistent with the primary driver of task-state FC being moment-by-moment (rather than cross-event mean) fluctuations. Critically, however, removing this cross-mean variance would still be important, since the relatively small amount of cross-event mean variance was shown in previous sections to cause (false positive) statistically significant effects.

We tested this hypothesis by quantifying the change in between-region shared variance before versus after FIR task regression. We found that 89.39% of the shared

variance across all pairwise connections (across all 7 tasks) was preserved after FIR regression. This was computed after converting the r-values representing task-state FC to r-squared values (i.e., percent shared linear variance), then averaging across subjects, tasks, and connections. This revealed that mean shared variance during task went from $r^2=0.066$ without task regression to $r^2=0.059$ with FIR task regression on average. This small change indicated that 89.39% of the shared linear variance was preserved after FIR regression on average. This result confirms our hypothesis that, while critical for reducing the chance of a false positive for any single result, the FIR regression step removed only a small amount (less than 11%) of the variance driving task-state FC effects. This, in turn, demonstrates that task-state FC estimates (even when not performing task regression) are primarily driven by moment-by-moment variance rather than the cross-event mean variance removed by FIR regression.

Testing for generalization to task-to-task FC changes

The prior results demonstrate inflation of task-state FC, suggesting that task-to-task FC differences would also be altered. This result was not guaranteed, however, given the possibility that the task-state FC inflations reported above were subtle and therefore only detectable for large cognitive contrasts (such as between task and rest). We tested for cross-task alterations in the well-studied N-back task's 2-back vs. 0-back contrast (Barch et al., 2013). This is one of the seven tasks included in the prior analyses, with the 2-back and 0-back conditions estimated separately.

As expected, we found that results were similar with the cross-task FC comparison as the task-to-rest FC comparison. Specifically, the approaches that flexibly modeled the HRF shape (FIR and basis set approaches) produced fewer significant results than alternate approaches (**Figure 8**). Without task regression the percentage of connections with task-state FC changes (false discovery rate corrected for multiple comparisons, $p<0.05$) was 28.14% (**Figure 8A**). Only slightly reduced values were found for task regression with the canonical HRF approach (24.97%; **Figure 8B**). Consistent with the task-to-rest FC comparison results, there were substantial reductions in the percentage of task-state FC increases when using the constrained basis set (12.92%; **Figure 8C**) and FIR (2.89%; **Figure 8D**) approaches. In contrast with the task-to-rest FC comparison results, however, FIR regression reduced the number of significant results relative to the basis set approach (2.89% vs. 12.92%). Notably, the significant reduction of visual network FC with the dorsal attention network (from 2-back to 0-back) was present for three of the methods but went away with FIR regression – the method that most flexibly fits HRF shape and thus likely best reduces false positives. This demonstrates a large-scale conclusion that could have been reached erroneously if FIR regression was not used to remove task activations.




Figure 8 – Task-to-task FC comparison: 2-back vs. 0-back (N-back working memory task). **A)** 2-back vs. 0-back FC differences, with no task regression preprocessing ($p < 0.05$, FDR corrected for multiple comparisons). **B)** Identical to panel A, but with constrained basis set task regression preprocessing. **C)** Identical to panel A, but with canonical HRF task regression preprocessing. Note the visual similarity to the no-task-regression results. **D)** Identical to panel A, but with FIR task regression preprocessing.

These results suggest that the small FC differences between well-matched task conditions can be more sensitive than task-to-rest comparisons to the quality of GLM fit for the FC pattern that emerges. Based on the computational model results indicating that the fMRI data better reflect pre-fMRI (i.e., input/LFP) data when using the FIR approach, and the additional flexibility of the FIR approach (without inflated false negatives) relative to the basis set approach, we interpret the FIR results as likely being more accurate than the other approaches. Note, however, that (regardless of regression method) concluding a true change in FC occurred – rather than a change in unshared variance (e.g., noise) – would require additional tests such as unscaled covariance (Cole et al., 2016).

Testing an alternative FC estimation method: PPI

We next tested whether an alternate FC estimation method is similarly affected by fMRI-induced inflation. As we have shown previously, covariance is the common statistical measure underlying a variety of FC measures (Cole et al., 2016): Pearson correlation, Spearman rank correlation, multiple regression, and PPI are all forms of normalizing/modifying simple covariance. Specifically, Pearson correlation normalizes covariance by dividing by (a transform on) the standard deviations of the time series, while Spearman rank correlation is equivalent to calculating Pearson correlation on the rank orders of the time series values. PPI is the simple pairwise regression between time series along with some nuisance regressors (Cole et al., 2013b). Notably, simple pairwise regression (as used by PPI) is equivalent to the covariance divided by the variance of the source (in a source-target pair) time series. Finally, multiple regression is equivalent to the partial covariance (i.e., with variance from the linear best-fits of all other time series removed) divided by the variance of the source time series.

Given that covariance underlies two common task-state FC methods used with fMRI – Pearson correlation and PPI – we expected that PPI would be similarly affected by task co-activations as compared to what we found with Pearson correlations. We tested this by calculating PPI using either no task regression, canonical-HRF regression (as used with standard PPI), or FIR regression. PPI involves a task-regression step that assumes the canonical HRF, such that comparison to the canonical-HRF condition will be the most relevant to existing PPI approaches. Note that we used a modified version of generalized PPI (McLaren et al., 2012), wherein the "psychological" variables are block-level boxcar regressors (Cole et al., 2013b). This aids with interpretation (and comparison to the Pearson correlation results), since the interaction term in the PPI calculation is not influenced by the chosen HRF shape. Also note that, unlike the original PPI approach (Friston et al., 1997), generalized PPI is calculated for each task condition separately (rather than using condition contrasts only) with contrasts calculated as subtraction of PPI estimates (McLaren et al., 2012). Another difference from typical PPI approaches was that the task activation regression occurred prior to (rather than simultaneous with) FC estimation. We did this primarily to make the PPI approach (slightly) more conservative, with as much variance as could be accounted for by the task regressors being taken out prior to PPI estimation. Thus, if anything, the approach used here should reduce the chance of false positives relative to typical PPI approaches.

We began by comparing no-regression to FIR-regression with PPI. As with Pearson correlation, we found that all seven tasks involved statistically significant ($p < 0.05$, FDR corrected) changes in FC estimates. The percentage of connections with significant ($p < 0.05$, FDR corrected) differences for each task were, respectively (increases/decreases): 8.6/6.3, 28.6/7.8, 25.0/22.9, 24.7/4.1, 35.0/12.1, 43.1/15.3, 23.4/8.2. These results demonstrate that task-timing regression matters for PPI analyses, as it significantly alters PPI estimates across a broad variety of brain regions across a broad variety of task manipulations.

We next tested the extent to which PPI results – which exclusively assume the standard HRF shape – likely include task-evoked activation-based FC inflation. This was quantified by comparing PPI calculated using canonical-HRF task regression

versus PPI calculated using FIR task regression. Consistent with the Pearson correlation results, canonical-HRF regression resulted in significantly distinct PPI estimates relative to when FIR regression was used. The percentage of connections with significant ($p < 0.05$, FDR corrected) differences for each task were, respectively (increases/decreases): 3.3/4.1, 25.7/1.3, 30.8/9.9, 6.4/0.4, 29.8/0.8, 36.8/10.6, 23.2/1.1. Notably, the percentage of changed connections tended to be smaller here than the no-regression vs. FIR regression case. This suggests that the canonical-HRF regression typically used with PPI likely helps reduce activation-induced FC inflation. However, given that a large number of significant differences remained when comparing canonical-HRF with FIR regression, the typical PPI approach appears to not be as effective as FIR regression.

DISCUSSION

We found strong evidence that task-evoked activations lead to spurious but systematic changes in fMRI-based task FC estimates. This was noted as a possibility in previous publications (Al-Aidroos et al., 2012; Cole et al., 2013b; Fair et al., 2007; Friston et al., 1997; Gratton et al., 2016) but, to our knowledge, has never been established either theoretically (using computational modeling) or empirically. Further, this hypothesized issue with task FC has typically been described generally, without reference to it being particularly problematic for fMRI analyses (though more fMRI researchers seem to have worried about it). We began by modeling the hypothesized effect using a neural mass computational model. Notably, we did not force the model to show activation-induced FC inflation, but discovered that it emerged simply from modeling fMRI task activations. Regression methods that flexibly fit hemodynamic response shape – FIR and basis set GLM approaches – were found to eliminate activation-induced FC inflation (without increasing false negatives), whereas alternative methods did not. Consistent with these theoretical results, we found that FIR and basis set approaches significantly reduced task FC estimates in empirical fMRI data. We found that the FIR approach reduced task FC estimates the most, consistent with its unique ability to flexibly fit any possible HRF shape, suggesting this as the preferred approach.

Why event-averaged task activation variance should be removed prior to estimating task FC

Our extensive computational and empirical investigation of activation-induced FC inflation suggests several reasons why event-averaged task activations should be removed prior to estimating task FC. For instance, we found that FC changes and activation amplitude changes are statistically and mechanistically distinct, such that they have meaningfully distinct implications for neuroscientific theory. Specifically, event-averaged task-evoked activations involve consistent cross-event activity amplitudes, while task-state FC involves synchronous moment-by-moment changes in activity (potentially with highly variable activity amplitudes) indicative of direct or indirect neural

interactions. This distinction can also be thought about in terms of task-evoked activation being enhanced by low variance (amplitude consistency) contrasting with task-state FC potentially being enhanced by high variance (moment-to-moment covariance). Thus, even if one finds event-averaged task-evoked activation patterns of interest, they should be investigated separately from task-state FC due to the mechanistic distinction between them. Indeed, there are already sub-fields to investigate task-evoked activation patterns separately from task-state FC – multivariate pattern analysis (Norman et al., 2006) and standard GLM analysis (Poline and Brett, 2012) – again supporting the conclusion that such effects should be isolated from task-state FC estimates.

Another reason to remove task activations prior to estimating task FC is that allowing task-evoked activations to inflate task-state FC estimates leaves open the possibility that new task-state FC effects are simply relabeling previously-discovered task-evoked activation effects as “connectivity”. This suggests that some previously-discovered effects that either did not remove any task activation variance (for example: Bassett et al., 2013; 2011; Krienen et al., 2014; Shirer et al., 2012; Tomasi et al., 2013), or that used a suboptimal approach for removing task activation variance (for example (including our own work): Banks et al., 2007; Cole et al., 2014; Iidaka et al., 2001; Lanius et al., 2004; Schultz and Cole, 2016), could have been driven to some extent by task activation changes. Notably, a handful of studies have already used FIR GLM to remove task activation variance prior to estimating task FC (Al-Aidroos et al., 2012; Cordova et al., 2016; Fair et al., 2007; Gratton et al., 2016; Norman-Haignere et al., 2012; Sadaghiani et al., 2015; Summerfield et al., 2006), suggesting these studies did not suffer from the task FC inflation effect identified here. Some have labeled this FIR-based removal of task activation variance followed by task FC estimation “background connectivity” (Al-Aidroos et al., 2012; Cordova et al., 2016; Griffis et al., 2015; Norman-Haignere et al., 2012). The present results suggest “background connectivity” and related approaches are effective in reducing (and likely even eliminating) task FC false positives driven by fMRI task activations.

A skeptic might argue that one could reverse this argument, with task FC being the real effect and task activations being secondary. The computational model analyses here demonstrate this is incorrect, since there are cases in which no true task FC is possible yet task FC is spuriously detected due to task co-activation (see the “no connectivity zone” in Figure 2). Further, it is clear that task activation is the first-order effect (simple change in cross-event mean amplitude), whereas task FC is a second-order effect building on covariance in moment-to-moment activation amplitudes. It is customary in science and statistics to account for simpler, first-order effects prior to interpreting second-order effects, such as interpreting ANOVA interactions only after accounting for main effects. Thus, an effect that can be explained as either a task activation or a task FC change would be preferentially interpreted as the simpler of the two – a task activation.

Another concern of a skeptic might be that removing task activation variance would remove the very task FC effects s/he is interested in. Both the model and the empirical results demonstrate that this is highly unlikely. First, we found that FIR task regression did not increase the rate of false negatives (for fMRI vs. pre-fMRI FC) in the computational model. Indeed, FIR task regression *reduced* the rate of false negatives,

suggesting FIR task regression might even increase the number of detected true task FC effects (rather than simply reducing false positives). Second, we found that the event-averaged task activation variance removed was only a small percentage (~10%) of the shared variance in the empirical fMRI data, suggesting that the bulk of the effects without activation regression was already driven by moment-by-moment variance independent from event-averaged activations. This suggests that even those who interpreted task FC in terms of event-averaged co-activation were actually observing primarily correlations of moment-by-moment fluctuations. Notably, despite most of the variance being driven by moment-by-moment fluctuations, we found that event-averaged activations alter task FC estimates substantially enough that many false conclusions are obtained without first removing event-averaged task activation variance. These two findings – that task activations were both a small portion of the overall variance and made a meaningful difference to results – can be reconciled by considering that a relatively small percentage of false positives among thousands of functional connections would nonetheless produce a large number of false inferences.

Limitations and opportunities for further research

As with most studies, many possible analyses related to the core research question were not included here, providing opportunities for future research. For instance, it could be informative to use a neuron-level computational model to further verify the results obtained using the neural mass model (Brette et al., 2007; Goodman, 2008). However, our neural mass model was intentionally kept simple and abstract, with the expectation that this abstraction will increase the probability that results will generalize to many different possible computational models (including highly realistic neuron-level models). The key idea is that abstraction to neuron-like units reduces the number of assumptions by identifying effects that are general enough to emerge from properties present in a variety of neuron-like interactions. Despite the plausibility of this expectation it is of course important to test this prediction using more detailed neuron-level modeling.

There were several aspects of the computational model results that did not completely agree with the empirical fMRI results. First, we empirically observed more task-state FC decreases from resting state, whereas the computational model results showed more task-state FC increases from resting state. This likely reflects our use of task-state FC increases from resting state (among the first 50 nodes) to select the model parameters. Notably, in the model we saw task-state FC decreases between the first 50 and second 50 nodes, due to there being inhibitory connections between those two network communities. This could suggest that more inhibitory connectivity should have been included in the model in order to match the empirical results. Alternatively, we could have selected model parameters based on maximal decreases in task-state FC relative to resting state. This may have resulted, for instance, in a higher bias parameter, equivalent to a larger amount of spontaneous activity leading to larger resting-state correlations. This issue is related to improving understanding of why the empirical results showed that most functional connections are lower during task relative to rest.

Another aspect of the computational model results that was not in complete agreement with the empirical fMRI results was the observation that FIR task regression reduced task-state FC estimates substantially more than basis set task regression. In the model the basis set approach involved only 1.05% false positives, very similar to the 0.94% false positives with the FIR approach. While the results were similar for task-state FC vs. resting-state FC (2.49% detected effects with FIR vs. 3.01% with basis set), our task-to-task FC comparison indicated substantially fewer detected FC differences when using FIR (FIR: 2.89%, basis set: 12.92%). Given the much more flexible fitting of HRF shape with FIR, it is likely that FIR task regression better fit and removed the task-evoked activations than the basis set approach. It is likely that the extra flexibility of FIR over fit the task-evoked time series, likely removing additional noise but also some covariance of interest. However, the computational model results suggest that, if anything, this extra flexibility likely reduced (rather than increased) false negatives, potentially by removing more noise than covariance signals. It will nonetheless be important for future research to quantify the degree to which FIR model overfitting results in inflation of false negatives in empirical results.

We were able to use the computational model to conclusively show that co-activations can induce spurious fMRI task FC by creating a "no connectivity zone" wherein no true task FC can be possible. Ideally, however, we would have had this sort of scenario in the empirical fMRI dataset. Instead, the empirical fMRI analyses supported the plausibility of task FC being inflated, with detected increases and decreases in task FC once event-averaged task activation variance was removed. This leaves open the possibility (however small) that removing task activation variance removed some true task FC effects. It will be important for future studies to investigate this possibility. Notably, however, the computational modeling results demonstrated that false negatives were not increased (and were in fact decreased) when task activation variance was removed. This suggests that, if anything, removing event-averaged task activation variance in turn increases the number of true task FC effects detected (rather than decreasing them).

We focused primarily on Pearson correlation-based task FC. It will be important for future research to test the generality of our conclusions to all task FC approaches. We showed that the results at least generalize to PPI analyses, suggesting the findings will likely generalize further. Indeed, the generalization to PPI suggests the task FC inflation effect is driven primarily by a change in covariance – the quantity underlying a variety of association measures used for task FC analysis (such as Pearson correlation and PPI) (Cole et al., 2016). This is consistent with the "highly simplified" model results (Figure 4), which shows that the underlying task FC inflation is driven by similarity in the hemodynamic response function. Such clear similarity – which was induced by convolution with a similar-shaped HRF – suggests a variety of association measures will be inflated by fMRI task co-activation, consistent with this effect generalizing to many task FC measures.

It will be important for future research to investigate alternative approaches to correcting the task FC inflation seen here. For instance, one promising approach is blind deconvolution (Havlicek et al., 2011), which flexibly removes HRF shape from entire time series. This could, in theory, correct the inflation by estimating the true neural time series separated from the HRF. Such a result would be consistent with our finding that

task FC was only minimally inflated in the pre-fMRI (i.e., truly neural) time series in the computational model results. Another method that we expect to be effective in reducing or eliminating task activation-based inflation of fMRI task FC is the "beta series" task FC approach (Rissman et al., 2004a). In this approach, a separate GLM parameter estimate is fit to each task event (with an assumed HRF shape), with Pearson correlation of the parameter estimates (across voxels or regions) estimating task FC. In theory, this approach estimates task FC based on event-by-event (e.g., trial-by-trial or block-by-block) covariance, rather than the moment-by-moment covariance that is typically used. This approach's use of an assumed HRF shape may result in false negatives (due to poor fit to activations in some cases), but appears unlikely to suffer from the same task FC false positives characterized here, given that beta series correlations do not include the moment-by-moment variance that is altered by HRF shape similarity between time series. This suggests that studies that used beta series correlations are unlikely to have been influenced by the false positives characterized here (Cisler et al., 2014; Gazzaley et al., 2004; for example: Nee and Brown, 2012; Rissman et al., 2004b; Zanto et al., 2011), though future research will be important for verifying this.

It will also be important for future research to investigate why the pre-fMRI simulation had some inflated task-state FC estimates. The inflation was quite small (a 1.99% false positive rate with a $p < 0.01$ threshold), especially relative to the no-regression fMRI results (42.58% false positive rate), but it was nonetheless higher than expected by chance (1%, given the $p < 0.01$ threshold). This likely reflects the small amount of coincident timing induced by the simultaneous stimulation across neural units, suggesting regression-based removal of non-fMRI data (Headley and Weinberger, 2013; Karamzadeh et al., 2010; Mill et al., 2017) could also be useful for reducing false positives (though the model results suggest this problem will likely be substantially smaller for non-fMRI relative to fMRI or other BOLD-based (Ferrari and Quaresima, 2012) data). Supporting this possibility, investigations of task-state FC with multi-unit recording in animal models (i.e., not involving the BOLD signal) have tended to remove cross-event average evoked responses prior to estimating correlations among neural time series (termed "noise correlations") in the interest of reducing false positives (Cafaro:2010im; Cohen and Kohn, 2011). It will be important for future research to investigate use of non-parametric approaches popular with spike timing cross-correlations – such as shuffling trial-by-trial events to estimate the contribution of confounding stimulus-evoked covariance (Brody, 1999; Grün, 2009) – with fMRI task-state FC estimation.

One remaining issue for the FIR GLM regression approach is that it relies on the particular set of regressors specified, when there might be additional task events unaccounted for. For instance, block onset and offset events with prominent fMRI activation responses have been identified (Dosenbach et al., 2006; Griffis et al., 2015; Visscher et al., 2003), such that a standard FIR model of an event-related task design would fail to remove fMRI activation variance from these prominent events. The variance from these events would likely inflate task FC estimates. One solution would be to model these block onset and offset events separately so as to remove this variance prior to task FC estimation, as has been done recently (Griffis et al., 2015). Another solution that was successfully applied here is to design task blocks of a given

condition to have identical trial timings, then model all blocks with a single long set of regressors (such that all consistent within-block events would be modeled, including block onset and offsets) (Al-Aidroos et al., 2012).

Similar issues arise from rare events with large fMRI activation responses such as error trials (Menon et al., 2001; Neta et al., 2015) or learning-induced changes in activations, which are typically not accounted for separately in GLM models. Such events might also inflate task FC estimates, though they could also be included in an FIR GLM to reduce this effect. It will be important for future studies to consider these various scenarios and determine whether they can meaningfully alter task FC estimates. Given that most of the task FC inflation effect is caused by the HRF shape, another possibility would be to utilize blind deconvolution (Havlicek et al., 2011) to reduce this effect no matter its source (even those unknown to the experimenter). Another possibility is that the task-activation false positives arise solely from the experimental manipulation (task timing) acting as a confounding third variable, implying that internally-generated activation events (such as error trials or learning-related activation changes) reflect the brain dynamics of interest and therefore do not need to be removed. It will be important for future studies to investigate this issue, however, given the ambiguity (regarding false positives) of situations like error trials and task learning being an interaction between experimenter-induced task timing and internal processes.

Conclusion

We identified strong evidence that fMRI-based task FC estimates are consistently and spuriously altered by task activations. This was shown across a neural mass computational model, a highly simplified model, and empirical fMRI data involving seven highly distinct tasks. The models and empirical fMRI data analyses converged in suggesting that methods that remove event-averaged task activation variance – when flexibly taking HRF shape into account (especially FIR GLM) – are able to correct for activation-induced task FC inflation. These results suggest prior task FC fMRI studies that did not use FIR GLM as a preprocessing step likely contain false positives. It will therefore be important to reanalyze data when possible, and begin using FIR GLM as a preprocessing step for task FC analyses moving forward. It might be tempting to retain event-averaged task activation variance in future task FC analyses given that the issue is not as problematic for non-fMRI data. However, the observation of inflated false positives in the "no connectivity zone" (1.99% with $p < 0.01$) for the pre-fMRI simulation data suggests this is a fundamental problem for task FC analysis, such that task activation regression should be used with non-fMRI data as well. Moving forward, it will be important to develop a deeper understanding of why event-averaged task activation causes false positives even for non-fMRI data, as well as identifying alternative approaches to removing event-averaged task activations in both fMRI and non-fMRI data.

REFERENCES

- Al-Aidroos, N., Said, C.P., Turk-Browne, N.B., 2012. Top-down attention switches coupling between low-level and high-level areas of human visual cortex. *Proceedings of the National Academy of Sciences* 109, 14675–14680. doi:10.1073/pnas.1202095109
- Banks, S.J., Eddy, K.T., Angstadt, M., Nathan, P.J., Phan, K.L., 2007. Amygdala–frontal connectivity during emotion regulation. *Social Cognitive and Affective Neuroscience* 2, 303–312. doi:10.1093/scan/nsm029
- Barch, D.M., Burgess, G.C., Harms, M.P., Petersen, S.E., Schlaggar, B.L., Corbetta, M., Glasser, M.F., Curtiss, S., Dixit, S., Feldt, C., Nolan, D., Bryant, E., Hartley, T., Footer, O., Bjork, J.M., Poldrack, R., Smith, S., Johansen-Berg, H., Snyder, A.Z., Van Essen, D.C., Consortium, F.T.W.-M.H., 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
- Barral, J., D Reyes, A., 2016. Synaptic scaling rule preserves excitatory-inhibitory balance and salient neuronal network dynamics. *Nat Neurosci*. doi:10.1038/nn.4415
- Bassett, D.S., Wymbs, N.F., Porter, M.A., Mucha, P.J., Carlson, J.M., Grafton, S.T., 2011. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences* 108, 7641–7646. doi:10.1073/pnas.1018985108
- Bassett, D.S., Wymbs, N.F., Rombach, M.P., Porter, M.A., Mucha, P.J., Grafton, S.T., 2013. Task-Based Core-Periphery Organization of Human Brain Dynamics. *PLoS Comput Biol* 9, e1003171. doi:10.1371/journal.pcbi.1003171.s012
- Brette, R., Rudolph, M., Carnevale, T., Hines, M., Beeman, D., Bower, J.M., Diesmann, M., Morrison, A., Goodman, P.H., Harris, F.C., Zirpe, M., Natschläger, T., Pecevski, D., Ermentrout, B., Djurfeldt, M., Lansner, A., Rochel, O., Vieville, T., Muller, E., Davison, A.P., Boustani, El, S., Destexhe, A., 2007. Simulation of networks of spiking neurons: A review of tools and strategies. *J Comput Neurosci* 23, 349–398. doi:10.1007/s10827-007-0038-6
- Brody, C., 1999. Correlations without synchrony. *Neural Comput* 11, 1537–1551.
- Cafaro, J., Rieke, F., 2010. Noise correlations improve response fidelity and stimulus encoding. *Nature* 468, 964–967. doi:10.1038/nature09570
- Chen, L., Vu, A.T., Xu, J., Moeller, S., Ugurbil, K., Yacoub, E., Feinberg, D.A., 2015. Evaluation of highly accelerated simultaneous multi-slice EPI for fMRI. *NeuroImage* 104, 452–459. doi:10.1016/j.neuroimage.2014.10.027
- Cisler, J.M., Bush, K., Steele, J.S., 2014. A comparison of statistical methods for detecting context-modulated functional connectivity in fMRI. *NeuroImage* 84, 1042–1052. doi:10.1016/j.neuroimage.2013.09.018
- Cohen, M.R., Kohn, A., 2011. Measuring and interpreting neuronal correlations. *Nat Neurosci* 14, 811–819. doi:10.1038/nn.2842

- Cole, M.W., Bassett, D.S., Power, J.D., Braver, T.S., Petersen, S.E., 2014. Intrinsic and task-evoked network architectures of the human brain. *Neuron* 83, 238–251. doi:10.1016/j.neuron.2014.05.014
- Cole, M.W., Laurent, P., Stocco, A., 2013a. Rapid instructed task learning: a new window into the human brain's unique capacity for flexible cognitive control. *Cogn Affect Behav Neurosci* 13, 1–22. doi:10.3758/s13415-012-0125-7
- Cole, M.W., Reynolds, J.R., Power, J.D., Repovs, G., Anticevic, A., Braver, T.S., 2013b. Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat Neurosci* 16, 1348–1355. doi:10.1038/nn.3470
- Cole, M.W., Yang, G.J., Murray, J.D., Repovs, G., Anticevic, A., 2016. Functional connectivity change as shared signal dynamics. *J Neurosci Methods* 259, 22–39. doi:10.1016/j.jneumeth.2015.11.011
- Cordova, N.I., Tompary, A., Turk-Browne, N.B., 2016. Attentional modulation of background connectivity between ventral visual cortex and the medial temporal lobe. *Neurobiology of Learning and Memory* 1–8. doi:10.1016/j.nlm.2016.06.011
- Dosenbach, N.U.F., Visscher, K.M., Palmer, E.D., Miezin, F.M., Wenger, K.K., Kang, H.C., Burgund, E.D., Grimes, A.L., Schlaggar, B.L., Petersen, S.E., 2006. A core system for the implementation of task sets. *Neuron* 50, 799–812. doi:10.1016/j.neuron.2006.04.031
- Eklund, A., Nichols, T.E., Knutsson, H., 2016. Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates. *Proc Natl Acad Sci USA* 201602413–6. doi:10.1073/pnas.1602413113
- Fair, D., Schlaggar, B., Cohen, A., Miezin, F., Dosenbach, N., Wenger, K., Fox, M., Snyder, A., Raichle, M., Petersen, S., 2007. A method for using blocked and event-related fMRI data to study “resting state” functional connectivity. *NeuroImage* 35, 396–405.
- Ferrari, M., Quaresima, V., 2012. A brief review on the history of human functional near-infrared spectroscopy (fNIRS) development and fields of application. *NeuroImage* 63, 921–935. doi:10.1016/j.neuroimage.2012.03.049
- Friston, K.J., 2011. Functional and Effective Connectivity: A Review. *Brain Connectivity* 1, 13–36. doi:10.1089/brain.2011.0008
- Friston, K.J., 1994. Functional and effective connectivity in neuroimaging: a synthesis. *Human brain mapping* 2, 56–78.
- Friston, K.J., Buechel, C., Fink, G.R., Morris, J., Rolls, E., Dolan, R.J., 1997. Psychophysiological and modulatory interactions in neuroimaging. *NeuroImage* 6, 218–229. doi:10.1006/nimg.1997.0291
- Gazzaley, A., Rissman, J., Desposito, M., 2004. Functional connectivity during working memory maintenance. *Cogn Affect Behav Neurosci* 4, 580–599.
- Genovese, C., Lazar, N., Nichols, T., 2002. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *NeuroImage* 15, 870–878.
- Girvan, M., Newman, M.E.J., 2002. Community structure in social and biological networks. *Proc Natl Acad Sci USA* 99, 7821–7826. doi:10.1073/pnas.122653799
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature*. doi:10.1038/nature18933

- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., WU-Minn HCP Consortium, 2013. The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* 80, 105–124. doi:10.1016/j.neuroimage.2013.04.127
- Goodman, D., 2008. Brian: a simulator for spiking neural networks in Python. *Front. Neuroinform.* 2. doi:10.3389/neuro.11.005.2008
- Gratton, C., 2013. The effect of theta-burst TMS on cognitive control networks measured with resting state fMRI 1–14. doi:10.3389/fnsys.2013.00124/abstract
- Gratton, C., Laumann, T.O., Gordon, E.M., Adeyemo, B., Petersen, S.E., 2016. Evidence for Two Independent Factors that Modify Brain Networks to Meet Task Goals. *CellReports* 17, 1276–1288. doi:10.1016/j.celrep.2016.10.002
- Greenland, S., Robins, J.M., Pearl, J., 1999. Confounding and collapsibility in causal inference. *Statist. Sci.* 14. doi:10.2307/2676645
- Griffis, J.C., Elkhetafi, A.S., Burge, W.K., Chen, R.H., Visscher, K.M., 2015. Retinotopic patterns of background connectivity between V1 and fronto-parietal cortex are modulated by task demands. *Front. Hum. Neurosci.* 9, 14675–14. doi:10.3389/fnhum.2015.00338
- Grün, S., 2009. Data-Driven Significance Estimation for Precise Spike Correlation. *J Neurophysiol* 101, 1126–1140. doi:10.1152/jn.00093.2008
- Handwerker, D., Ollinger, J., D'Esposito, M., 2004. Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *NeuroImage* 21, 1639–1651.
- Havlicek, M., Friston, K.J., Jan, J., Brazdil, M., Calhoun, V.D., 2011. Dynamic modeling of neuronal responses in fMRI using cubature Kalman filtering. *NeuroImage* 56, 2109–2128. doi:10.1016/j.neuroimage.2011.03.005
- Headley, D.B., Weinberger, N.M., 2013. Fear Conditioning Enhances Gamma Oscillations and Their Entrainment of Neurons Representing the Conditioned Stimulus. *J Neurosci* 33, 5705–5717. doi:10.1523/JNEUROSCI.4915-12.2013
- Hopfield, J.J., 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *Proc Natl Acad Sci USA* 81, 3088–3092.
- Iidaka, T., Omori, M., Murata, T., Kosaka, H., Yonekura, Y., Okada, T., Sadato, N., 2001. Neural interaction of the amygdala with the prefrontal and temporal cortices in the processing of facial expressions as revealed by fMRI. *Journal of Cognitive Neuroscience* 13, 1035–1047. doi:10.1162/089892901753294338
- Ito, T., Kulkarni, K.R., Schultz, D.H., Mill, R.D., Chen, R.H., Solomyak, L.I., Cole, M.W., 2017. Cognitive task information is transferred between brain regions via resting-state network topology. *Nature Communications* 8, 1027. doi:10.1038/s41467-017-01000-w
- Karamzadeh, N., Medvedev, A., Azari, A., Gandjbakhche, A., Najafizadeh, L., 2010. Capturing dynamic patterns of task-based functional connectivity with EEG. *NeuroImage* 66, 1–7. doi:10.1016/j.neuroimage.2012.10.032
- Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S.F., Baker, C.I., 2009. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci* 12, 535–540. doi:10.1038/nn.2303

- Krienen, F.M., Yeo, B.T.T., Buckner, R.L., 2014. Reconfigurable task-dependent functional coupling modes cluster around a core functional architecture. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369, 20130526–20130526. doi:10.1098/rstb.2013.0526
- Lanius, R.A., Williamson, P.C., Densmore, M., Boksman, K., Neufeld, R.W., Gati, J.S., Menon, R.S., 2004. The nature of traumatic memories: a 4-T fMRI functional connectivity analysis. *Am J Psychiatry* 161, 36–44. doi:10.1176/appi.ajp.161.1.36
- Likhtik, E., Pelletier, J.G., Paz, R., Paré, D., 2005. Prefrontal control of the amygdala. *Journal of Neuroscience* 25, 7429–7437. doi:10.1523/JNEUROSCI.2314-05.2005
- Logothetis, N.K., Pauls, J., Augath, M., Trinath, T., Oeltermann, A., 2001. Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157. doi:10.1038/35084005
- Logothetis, N.K., Wandell, B.A., 2004. Interpreting the BOLD signal. *Annu. Rev. Physiol.* 66, 735–769. doi:10.1146/annurev.physiol.66.082602.092845
- McLaren, D.G., Ries, M.L., Xu, G., Johnson, S.C., 2012. A generalized form of context-dependent psychophysiological interactions (gPPI): a comparison to standard approaches. *NeuroImage* 61, 1277–1286. doi:10.1016/j.neuroimage.2012.03.068
- Menon, V., Adelman, N., White, C., Glover, G., Reiss, A., 2001. Error-related brain activation during a Go/NoGo response inhibition task. *Hum Brain Mapp* 12, 131–143.
- Mill, R.D., Bagic, A., Bostan, A., Schneider, W., Cole, M.W., 2017. Empirical validation of directed functional connectivity. *NeuroImage* 146, 275–287. doi:10.1016/j.neuroimage.2016.11.037
- Nee, D.E., Brown, J.W., 2012. Dissociable Frontal-Striatal and Frontal-Parietal Networks Involved in Updating Hierarchical Contexts in Working Memory. *Cerebral Cortex*. doi:10.1093/cercor/bhs194
- Neta, M., Miezin, F.M., Nelson, S.M., Dubis, J.W., Dosenbach, N.U.F., Schlaggar, B.L., Petersen, S.E., 2015. Spatial and temporal characteristics of error-related activity in the human brain. *Journal of Neuroscience* 35, 253–266. doi:10.1523/JNEUROSCI.1313-14.2015
- Norman, K., Polyn, S., Detre, G., Haxby, J., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci (Regul Ed)* 10, 424–430. doi:10.1016/j.tics.2006.07.005
- Norman-Haignere, S.V., McCarthy, G., Chun, M.M., Turk-Browne, N.B., 2012. Category-Selective Background Connectivity in Ventral Visual Cortex. *Cerebral Cortex* 22, 391–402. doi:10.1093/cercor/bhr118
- O'Reilly, J.X., Woolrich, M.W., Behrens, T.E.J., Smith, S.M., Johansen-Berg, H., 2012. Tools of the trade: psychophysiological interactions and functional connectivity. *Social Cognitive and Affective Neuroscience* 7, 604–609. doi:10.1093/scan/nss055
- Ogawa, S., Tank, D.W., Menon, R., Ellermann, J.M., Kim, S.G., Merkle, H., Ugurbil, K., 1992. Intrinsic signal changes accompanying sensory stimulation: functional brain mapping with magnetic resonance imaging. *Proc Natl Acad Sci USA* 89, 5951–5955.
- Poline, J.-B., Brett, M., 2012. The General Linear Model and fMRI: Does love last forever? *NeuroImage* 1–18. doi:10.1016/j.neuroimage.2012.01.133

- Power, J.D., Barnes, K.A., Snyder, A.Z., Schlaggar, B.L., Petersen, S.E., 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *NeuroImage* 59, 2142–2154. doi:10.1016/j.neuroimage.2011.10.018
- Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., Petersen, S.E., 2011. Functional network organization of the human brain. *Neuron* 72, 665–678. doi:10.1016/j.neuron.2011.09.006
- Rissman, J., Gazzaley, A., D'Esposito, M., 2004a. Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage* 23, 752–763. doi:10.1016/j.neuroimage.2004.06.035
- Rissman, J., Gazzaley, A., D'Esposito, M., 2004b. Measuring functional connectivity during distinct stages of a cognitive task. *NeuroImage* 23, 752–763. doi:10.1016/j.neuroimage.2004.06.035
- Sadaghiani, S., Poline, J.-B., Kleinschmidt, A., D'Esposito, M., 2015. Ongoing dynamics in large-scale functional connectivity predict perception. *Proc Natl Acad Sci USA* 112, 8463–8468. doi:10.1073/pnas.1420687112
- Schirner, M., McIntosh, A.R., Jirsa, V., Deco, G., Ritter, P., 2016. Bridging multiple scales in the human brain using computational modelling 1–31. doi:10.1101/085548
- Schultz, D.H., Cole, M.W., 2016. Higher Intelligence Is Associated with Less Task-Related Brain Network Reconfiguration. *Journal of Neuroscience* 36, 8551–8561. doi:10.1523/JNEUROSCI.0358-16.2016
- Shirer, W.R., Ryali, S., Rykhlevskaia, E., Menon, V., Greicius, M.D., 2012. Decoding subject-driven cognitive states with whole-brain connectivity patterns. *Cerebral Cortex* 22, 158–165. doi:10.1093/cercor/bhr099
- Siegel, J.S., Mitra, A., Laumann, T.O., Seitzman, B.A., Raichle, M., Corbetta, M., Snyder, A.Z., 2016. Data Quality Influences Observed Links Between Functional Connectivity and Behavior. *Cereb Cortex* 1–11. doi:10.1093/cercor/bhw253
- Siegel, M., Buschman, T.J., Miller, E.K., 2015. Cortical information flow during flexible sensorimotor decisions. *Science* 348, 1352–1355. doi:10.1126/science.aab0551
- Smith, S.M., Vidaurre, D., Beckmann, C.F., Glasser, M.F., Jenkinson, M., Miller, K.L., Nichols, T.E., Robinson, E.C., Salimi-Khorshidi, G., Woolrich, M.W., Barch, D.M., Ugurbil, K., Van Essen, D.C., 2013. Functional connectomics from resting-state fMRI. *Trends Cogn Sci (Regul Ed)* 17, 666–682. doi:10.1016/j.tics.2013.09.016
- Sprong, M., Ji, J.L., Kulkarni, K., Repovs, G., bioRxiv, A.A., 2017, 2017. Mapping the human brain's cortical-subcortical functional network organization. *bioRxiv*. doi:10.1101/206292
- Summerfield, C., Greene, M., Wager, T., Egner, T., Hirsch, J., Mangels, J., 2006. Neocortical Connectivity during Episodic Memory Formation. *Plos Biol* 4, e128. doi:10.1371/journal.pbio.0040128
- Tagliazucchi, E., Laufs, H., 2014. Decoding Wakefulness Levels from Typical fMRI Resting-State Data Reveals Reliable Drifts between Wakefulness and Sleep. *Neuron* 82, 695–708. doi:10.1016/j.neuron.2014.03.020
- Tallon-Baudry, C., Bertrand, O., 1999. Oscillatory gamma activity in humans and its role in object representation. *Trends Cogn Sci (Regul Ed)* 3, 151–162.

- Tomasi, D., Wang, R., Wang, G.J., Volkow, N.D., 2013. Functional Connectivity and Brain Activation: A Synergistic Approach. *Cerebral Cortex*. doi:10.1093/cercor/bht119
- Truccolo, W.A., Ding, M., Knuth, K.H., Nakamura, R., Bressler, S.L., 2002a. Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clin Neurophysiol* 113, 206–226.
- Truccolo, W.A., Ding, M., Knuth, K.H., Nakamura, R., Bressler, S.L., 2002b. Trial-to-trial variability of cortical evoked responses: implications for the analysis of functional connectivity. *Clin Neurophysiol* 113, 206–226.
- Ugurbil, K., Xu, J., Auerbach, E.J., Moeller, S., Vu, A., Duarte-Carvajalino, J.M., Lenglet, C., Wu, X., Schmitter, S., Van de Moortele, P.F., 2013. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *NeuroImage*.
- Van Essen, D.C., Glasser, M.F., Dierker, D.L., Harwell, J., Coalson, T., 2012. Parcellations and Hemispheric Asymmetries of Human Cerebral Cortex Analyzed on Surface-Based Atlases. *Cerebral Cortex* 22, 2241–2262. doi:10.1093/cercor/bhr291
- Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K., Consortium, F.T.W.-M.H., 2013. The WU-Minn Human Connectome Project: An overview. *NeuroImage* 1–18. doi:10.1016/j.neuroimage.2013.05.041
- Visscher, K., Miezin, F., Kelly, J., Buckner, R., Donaldson, D., McAvoy, M., Bhalodia, V., Petersen, S., 2003. Mixed blocked/event-related designs separate transient and sustained activity in fMRI. *NeuroImage* 19, 1694–1708.
- Wig, G.S., Schlaggar, B.L., Petersen, S.E., 2011. Concepts and principles in the analysis of brain networks. *Ann N Y Acad Sci* 1224, 126–146. doi:10.1111/j.1749-6632.2010.05947.x
- Woolrich, M.W., Behrens, T.E.J., Smith, S.M., 2004. Constrained linear basis sets for HRF modelling using Variational Bayes. *NeuroImage* 21, 1748–1761. doi:10.1016/j.neuroimage.2003.12.024
- Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zöllei, L., Polimeni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *J Neurophysiol* 106, 1125–1165. doi:10.1152/jn.00338.2011
- Zalesky, A., Fornito, A., Bullmore, E., 2012. On the use of correlation as a measure of network connectivity. *NeuroImage* 60, 2096–2106. doi:10.1016/j.neuroimage.2012.02.001
- Zanto, T.P., Rubens, M.T., Thangavel, A., Gazzaley, A., 2011. Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. *Nat Neurosci* 14, 656–661. doi:10.1038/nn.2773