Published in final edited form as:

Magn Reson Imaging. 2014 July ; 32(6): 721–735. doi:10.1016/j.mri.2014.03.002.

Deconvolution filtering: Temporal smoothing revisited

Keith Bush^{a,*} and Josh Cisler^b

^aDepartment of Computer Science, University of Arkansas at Little Rock (UALR), Little Rock, AR 72204, USA

^bBrain Imaging Research Center, University of Arkansas for Medical Sciences (UAMS), Little Rock, AR 72205, USA

Abstract

Inferences made from analysis of BOLD data regarding neural processes are potentially confounded by multiple competing sources: cardiac and respiratory signals, thermal effects, scanner drift, and motion-induced signal intensity changes. To address this problem, we propose deconvolution filtering, a process of systematically deconvolving and reconvolving the BOLD signal via the hemodynamic response function such that the resultant signal is composed of maximally likely neural and neurovascular signals. To test the validity of this approach, we compared the accuracy of BOLD signal variants (i.e., unfiltered, deconvolution filtered, band-pass filtered, and optimized band-pass filtered BOLD signals) in identifying useful properties of highly confounded, simulated BOLD data: (1) reconstructing the true, unconfounded BOLD signal, (2) correlation with the true, unconfounded BOLD signal, and (3) reconstructing the true functional connectivity of a three-node neural system. We also tested this approach by detecting task activation in BOLD data recorded from healthy adolescent girls (control) during an emotion processing task.

Results for the estimation of functional connectivity of simulated BOLD data demonstrated that analysis (via standard estimation methods) using deconvolution filtered BOLD data achieved superior performance to analysis performed using unfiltered BOLD data and was statistically similar to well-tuned band-pass filtered BOLD data. Contrary to band-pass filtering, however, deconvolution filtering is built upon physiological arguments and has the potential, at low TR, to match the performance of an optimal band-pass filter. The results from task estimation on real BOLD data suggest that deconvolution filtering provides superior or equivalent detection of task activations relative to comparable analyses on unfiltered signals and also provides decreased variance over the estimate. In turn, these results suggest that standard preprocessing of the BOLD signal ignores significant sources of noise that can be effectively removed without damaging the underlying signal.

Keywords

Deconvolution; Filtering; fMRI; Imaging analysis; BOLD signal; Functional connectivity

1. Introduction

The need to account for confounding factors in fMRI is a well-known problem in fMRI analysis that uses BOLD signals. The essential problem is that a BOLD signal is multidetermined: it is a mixture of the neurovascular consequences of neural firing combined with confounding factors, such as head motion, scanner artifact and signal drift, neurophysiological variability, etc. Consequently, signal fluctuations cannot be solely attributed to neural causes, which limits inferences about neural processes in functional imaging studies. Accordingly, attempts must be made to account for and/or remove confounding (i.e., non-neural in origin) sources of variance from BOLD data in order to foster more precise inferences to inform cognitive and clinical neuroscience.

Attempts to address this problem date back to seminal papers by Friston and colleagues describing variations of the generalized linear modeling (GLM) approach to filtering fMRI BOLD signal [1–4], which models an observed BOLD signal, $\mathbf{X}$, as

$$\mathbf{X} = \mathbf{H} \cdot \mathbf{w}_s + \mathbf{D} \cdot \mathbf{w}_c + \eta \quad (1)$$

where $\mathbf{H}$ is a matrix that comprised a set of explanatory functions (i.e., kernel vectors); $\mathbf{w}_s$ is the linear weight vector of these explanatory kernel vectors; $\mathbf{D}$ is a matrix comprised of a set of temporally structured confounding processes, $\mathbf{w}_c$; and, the final term, η is white Gaussian noise. This framework is the bedrock on which nearly all subsequent fMRI analyses are based. Early debate in fMRI modeling focused on the magnitude and role of autocorrelated noise processes, as well as the structure and applicability of the HRF [5,6]. Numerous papers debated the pros and cons of voxel-wise temporal smoothing and filtering [5,7,8].

The consensus from this early work, formed at the turn of century, is that temporal smoothing, in general, damages the underlying signal in fMRI except in cases of appropriate experimental design combined with band-pass filtering [8]. Also, autocorrelation (i.e., cardiac and respiratory influences) is the dominant source of noise and should always be modeled in concert with Gaussian white noise (i.e., thermal and quantum effects) [6,9,10].

Early work on modeling and smoothing BOLD signals was predicated on the use of SPM [11] as a means of identifying statistically significant task-related activations. As the GLM and SPM approaches matured, and sophisticated understanding of the HRF functions became available, researchers migrated the focus of fMRI BOLD signal modeling efforts toward identification of causal relationships in neural processing, particularly the general problem of capturing the underlying temporal distribution of neural events. This is exemplified in dynamic causal modeling, where deconvolution of the observed signal into neural estimates is the basis of forming causal inferences. Indeed, researchers [12–14] have subsequently argued for the necessity of deconvolution of the BOLD signal into its mediating neural events (either implicitly or explicitly) in order to improve inferences about neural activity.

Deconvolution is an inversion of the observed BOLD signal into a temporal distribution of individual neural events. This inversion process has been well studied, and numerous recent algorithmic approaches to this problem may be found in the literature [15–18]. The majority of deconvolution algorithms (excluding Bayesian filtering approaches [16]) assume the GLM form of the BOLD signal [15,18] in which matrix $\mathbf{H}$ (see Eq. (1)) is a convolution (i.e., Toeplitz) matrix formed from temporal offsets of the canonical HRF. Thus, the solution of this system yields a maximum likelihood estimation of the underlying true BOLD signal given quasi-physiological constraints. What makes deconvolution a unique problem is the allowable form of the weight matrix, $\mathbf{w}_S$, such that neural activations exhibit positive values and temporal structure, e.g., clustering and sparsity [18].

Convolution of the appropriate HRF with the correct weight vector, $\mathbf{w}_S$, will generate the true BOLD signal, i.e., a BOLD signal containing only neural estimates combined with their neurovascular consequences detected by BOLD. This unique feature of deconvolution (i.e., removing BOLD fluctuations that are not feasibly caused by neural influences) allows for a potentially powerful filtering approach: deconvolution of an observed BOLD signal into its neural estimates followed by reconvolution with an HRF estimate to produce a BOLD signal that is filtered of non-neural sources of influence (confounds) and retains the true signal of interest (neural processes). Thus, reconvolution of the deconvolved signal weights (termed the encoding), under certain assumptions of the system's physiology, constitutes an optimal filter, which we term deconvolution filtering.

The goal of this work is to examine and understand the viability of deconvolution filtering in improving analysis of real-world fMRI data. To achieve this, we perform a three-part analysis. In the first part, we simulate resting state fMRI BOLD of a single neural node using a number of known confounds: normalization scale-effects, downsampling effects, autocorrelated noise processes, and Gaussian white noise processes. The use of a simulation allows us to precisely estimate the root mean squared error (RMSE) and correlation between the filtered signals and the true, unconfounded BOLD signals. In the second part of the analysis, we vary additional confounds such as HRF misspecification, noise process autocorrelation, Gaussian white noise processes, and connection topology in a three-node neural system; the use of a simulation in this case allows us to precisely estimate the true functional connectivity graph of the three-node system, thereby facilitating quantitative assessment of both the absolute and relative roles of the various confounds, both within a single filtering approach and across differing filtering approaches: band-pass linear filtering, optimized band-pass linear filtering, as well as deconvolution filtering using both a GLM [15] and a nonlinear method [17] that strictly enforces positive neural representations. We deem analysis of functional connectivity estimation error critical in assessing how BOLD signal filtering impacts real-world, commonly used analysis. In the third part, we examine statistical properties of functional activation results of real-world fMRI BOLD data as additional motivation for the use of deconvolution filtering in practice.

We structure the manuscript into four parts. First, we describe a computational model of the BOLD signal's generation and observation for unifying performance results across filtering approaches. Second, we describe the deconvolution filtering framework, as well as two different deconvolution algorithms that will be used inside the framework during

experimentation; we also briefly review two common band-pass filtering approaches and an optimized band-pass filtering approach that we use as benchmark comparisons on which to judge the performance of deconvolution filtering. Third, we describe and execute a set of experiments on both simulated and real BOLD that examine the performance of deconvolution filtering under numerous signal confounds. Finally, we combine the experimental results into a single assessment of the practical application of deconvolution filtering, and we offer guidance on fruitful directions of future research.

2. Methods

Our methods are composed of (1) a comprehensive parametric computational model of an fMRI BOLD signal's generation, (2) a novel algorithm for filtering a BOLD signal via its deconvolved representation, (3) band-pass filters based on Fourier analysis, and (4) an optimal Fourier-based band-pass filter; we describe these components below.

2.1. A generative model of fMRI BOLD data

Following prior simulation models [19,20], we construct a model of resting state neural activity using standard assumptions about the nature of the observed BOLD signal. We model fMRI BOLD data in four distinct processes: (1) a functional network generation process; (2) a neural generation process that captures temporal structure of neural events induced either by external, unmodeled processes or communication between brain regions; (3) a theoretical BOLD generation process that maps neural events onto an ideal set of BOLD signals (using either the canonical HRF or the balloon model); and, (4) an observation process that maps theoretically ideal BOLD signals onto low-frequency, noise-corrupted BOLD observations that represent real-world data. Each of these processes is described in detail below.

2.1.1. Functional network generation—We describe the relationships in space and time between V distinct brain regions as a functional network, $\mathcal{F} = \{\mathbf{C}, \mathbf{L}, \rho\}$, comprised of a connectivity model, $\mathbf{C}$, a communication lag model, $\mathbf{L}$, and an external activity model, ρ . The connectivity model, $\mathbf{C}$, is a matrix of conditional probabilities, $\mathbf{C}(i, j) \in [0, 1]$, $i, j \in 1, \dots, V$, determining the probability that a neural event in brain region j will be induced by a neural event in brain region i . The lag model, $\mathbf{L}$, is a real-valued matrix, $\mathbf{L}(i, j) \in [0, L_{\max}]$, $i, j \in 1, \dots, V$ dictating the time delay (in seconds) required for brain region i to influence region j . The bound $L_{\max}$ is the anatomically constrained maximum time required for a neural signal to traverse the brain. The external activity model, ρ , is a real-valued vector, $\rho(i) \in [0, 1]$, $i \in 1, \dots, V$, which determines the probability of neural activity intrinsically generated by brain region i . Note, this intrinsic activity is intended to model neural events induced by external, “unmodeled” relationships between the regions of interest and the remainder of the brain.

2.1.2. Neural event generation process—The model simulates neural activity in a two-step process. The intrinsic neural activity is first sampled. Then functional neural activity (arising from interaction among the modeled brain regions) is sampled.

Each brain region, $i, i \in 1, \dots, V$, models neural activity as a binary-valued sequence, stored as a column vector, $\mathbf{e}_i$ (termed the encoding), of length N , $\mathbf{e}_i(n) \in \{0, 1\}$, $n \in 1, \dots, N$. Each element of this vector indicates whether a neural event occurred or did not occur at step n . This sequence may be mapped to a time-series by assuming an underlying maximum frequency of neural event generation, ν_g , such that elements of $\mathbf{e}_i$ may be correspondingly indexed as $\mathbf{e}_i(t)$, $t \in 1/\nu_g, \dots, N/\nu_g$. The individual elements of $\mathbf{e}_i$ may be populated with values generated from a number of deterministic or random processes; we choose to populate this vector via a thresholded uniformly random distribution, such that

$$\mathbf{e}_i(n) = \begin{cases} 1 & \text{if } \mathcal{U}(0, 1) < \rho(i), \\ 0 & \text{otherwise.} \end{cases}$$

From this description of the intrinsic neural activity the model calculates functional neural activity as follows. For each simulation step, n , for each brain region, i , if $\mathbf{e}_i(n) = 1$ then for each brain region, $j, j \in 1, \dots, V$, the existence of a neural event in region j induced by region i is calculated as,

$$\mathbf{e}_i(n + n_{lag}, j) = \begin{cases} 1 & \text{if } \mathcal{U}(0, 1) \leq \mathbf{C}(i, j), \\ \mathbf{e}_i(n + n_{lag}, j) & \text{otherwise,} \end{cases} \quad (2)$$

where n_{lag} is the number of simulation steps of lag required for the message from region i to travel to region j , given by the function,

$$n_{lag} = \left\lceil \frac{\mathbf{L}(i, j)}{\nu_g^{-1}} \right\rceil + n_{rem}, \quad (3)$$

where,

$$n_{rem} = \begin{cases} 1 & \text{if } \mathcal{U}(0, 1) \leq \frac{\mathbf{L}(i, j)}{\nu_g^{-1}} - \left\lfloor \frac{\mathbf{L}(i, j)}{\nu_g^{-1}} \right\rfloor, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

To summarize, a functional neural event is randomly sampled from the probability distribution that a neural event in one brain region will induce a future neural event in a connected neighbor. The simulation step during which this neural event will be induced is calculated as the current simulation step plus the ratio of the neural communication time between two brain regions and the simulation step-size, Eq. (3). In cases where the lag between regions is not wholly divisible by the simulation step-size, the simulation index corresponding to the arrival of the induced neural event is sampled from a weighted distribution of the nearest preceding and succeeding simulation indices. For example, if the true arrival time is temporally closer to the simulation time corresponding to the succeeding simulation index, then sampling this index is more likely and vice versa.

2.1.3. Mapping of neural events to a BOLD signal—The model generates a BOLD signal through a two-step process, identical to that described in Bush and Cisler [17]. In the first step, the model generates the true (i.e., theoretical) BOLD signal of the neural event encoding in either of two ways: via canonical HRF or via a more biophysically sophisticated hemodynamic model based on Buxton's balloon model [16,21,22]. The result of this step is a true BOLD signal (i.e., a theoretically ideal signal). The second step is an observation modeling process by which we introduce a wide range of confounding factors.

We will specify, briefly, the mathematical steps involved in this process in order to introduce relevant notation and parameters that will be addressed in the experiments. See Bush and Cisler [17] for a more detailed explanation. Here we derive our methods for the canonical HRF,¹ which assumes a kernel vector, $\mathbf{k}$, of length K . We convolve this kernel with the encoding corresponding to the i th brain region to form the true BOLD signal, $\mathbf{x}_i$, at sample, n :

$$\mathbf{x}_i(n) = \sum_{k=1}^K \mathbf{k}(k) \mathbf{e}_i(n - (k-1)), i \in 1, \dots, V, n \in 1, \dots, N$$

Each vector $\mathbf{x}_i$ is then mapped individually, via the following processing steps, to form an observation vector, $\bar{\mathbf{y}}_i^{\text{scan}}$.

Latent neural activity: to model latent neural events that occur prior to observation, we decompose the N elements of the true BOLD signal, $\mathbf{x}_i$, into two sequences of length $K-1$ and M , respectively, such that $N = (K-1) + M$. Thus, $\mathbf{x}_{trans}(n_{trans}) = \mathbf{x}_i(n_{trans})$, $n_{trans} \in 1, \dots, K-1$, which we term the transient signal (which is discarded), and, $\mathbf{x}_i^{obs}(n_{obs}) = \mathbf{x}_i(n_{obs} + (K-1))$, $n_{obs} \in 1, \dots, M$, which we term the observed BOLD signal.

Physiological variation: We model the impacts of temporally correlated noise sources (e.g., cardiac and respiratory) as a first-order autoregressive, AR(1), signal [6,10], η_{phys} , of length M , which is added to $\mathbf{x}_i^{obs}$ to form $\mathbf{y}_i^{obs}$. The noise signal, η_{phys} , is built from zero-mean white Gaussian noise and correlation coefficient, γ . The scale of this process is dictated by a signal-to-noise parameter, SNR_{phys} (the lower the value the greater confound of the process).

Downsampling: we compute the downsampling rate, d , using the function $d = v_g/v_o$ (note, $\text{TR} = 1/v_o$) and then downsample the observed BOLD signal,

$$\mathbf{y}_i^{dcmp}(p) = \mathbf{y}_i^{obs}(d \cdot p), p \in 1, \dots, P, \quad (5)$$

to form the downsampled observed BOLD signal, $\mathbf{y}_i^{dcmp}$, of length, $P = M/d$.

¹We used the canonical HRF as provided in the commonly used SPM software package.

Measurement noise: we model thermal and quantum effects on the observation via a zero-mean white Gaussian noise process, η^{scan} , of length P that is added to $\mathbf{y}_i^{dmp}$, forming, $\mathbf{y}_i^{scan}$. The scale of this process is dictated by SNR_{scan} .

Normalization: we compute the normalized observed vector, $\bar{\mathbf{y}}_i^{scan}$, as

$$\bar{\mathbf{y}}_i^{scan} = z(\mathbf{y}_i^{scan}), \quad (6)$$

where $z(\cdot)$ is the normalization function (e.g., percent signal change).

2.2. Deconvolution filtering

We describe the deconvolution filtering method, formally, in Algorithm 1. Given an observation, $\bar{\mathbf{y}}_i^{scan}$, and a deconvolution mapping, h (parameterized by set, Θ , and HRF kernel, $\mathbf{k}$), deconvolution filtering is a two-step process: (1) deconvolve the BOLD signal into a maximally likely neural event encoding, $\tilde{\mathbf{e}}_i$, and (2) reconvolve an approximate theoretical BOLD signal, $\mathbf{y}_i^{fltr}$, via $\mathbf{k}$. Note, parameter Θ_H describes the number of time points that the deconvolution algorithm attempts to estimate prior to the first observed datapoint, $\bar{\mathbf{y}}_i^{scan}(1)$, where $K \geq 1$ and $\Theta_H \leq K - 1$.

Algorithm 1

Deconvolution Filtering

parameters: $h, \Theta, \mathbf{k}$

data: $\bar{\mathbf{y}}_i^{scan}(p) \in \mathbb{R}, p \in 1, \dots, P, i \in 1, \dots, V$

result: $\mathbf{y}_i^{fltr}(\hat{p}), p \in 1 - \Theta_H, \dots, 1, \dots, P$

Begin

for $i \in 1, \dots, V$

$\tilde{\mathbf{e}}_i(\hat{p}) = h_{\Theta, \mathbf{k}}(\bar{\mathbf{y}}_i^{scan}), \tilde{\mathbf{e}}_i(\hat{p}) \in [0, 1]$

$\hat{K}(\hat{p}) = \begin{cases} \min(K, \hat{p} - (1 - \Theta_H)) & \text{if } \hat{p} < 1 \\ K & \text{else} \end{cases}$

$\mathbf{y}_i^{fltr}(\hat{p}) = \sum_{k=1}^{\hat{K}(\hat{p})} \mathbf{k}(k) \tilde{\mathbf{e}}_i(\hat{p} - (k - 1))$

end-for

end

The strength of deconvolution filtering rests on the ability of the underlying deconvolution algorithm, h , to generate high-quality neural encodings. In prior work [17] we compared the performance of various deconvolution algorithms [15–18] in decoding neural events from both simulated and real BOLD data. Using the results of this comparison, we select, and briefly summarize, the two best performing algorithms from this study that we use for h .

2.2.1. Regularized linear regression deconvolution—We introduce the regularized linear regression (ridge regression) approach of Gaudes et al. [15], which we abbreviate as Gau11 for the remainder of this work. This approach assumes that the observed BOLD signal is modeled by the multiplication of a convolution (i.e., Toeplitz) matrix, $\mathbf{A}$, of size $P \times P$ and a vector of neural events, $\mathbf{e}_i$, combined with a vector of i.i.d. Gaussian random noise, η_i ,

$$\bar{\mathbf{y}}_i^{scan} = \mathbf{A}\mathbf{e}_i + \eta_i.$$

Due to its form, inversion of matrix $\mathbf{A}$ is poorly-determined, such that direct inversion methods, i.e., $\tilde{\mathbf{e}}_i = \mathbf{A}^{-1}\bar{\mathbf{y}}_i^{scan}$, will have difficulty producing a good encoding estimate. To overcome this weakness the cost function of the regression necessarily includes a regularization term, λ , such that

$$J_i = \|\mathbf{A}\tilde{\mathbf{e}}_i - \bar{\mathbf{y}}_i^{scan}\| + \|\lambda(\bar{\mathbf{y}}_i^{scan})\tilde{\mathbf{e}}_i\|,$$

which enforces an intended structure on the estimated encoding, $\tilde{\mathbf{e}}_i$. The Gau11 algorithm defines λ as a function of the statistical properties of the observations [23].

2.2.2. Nonlinear deconvolution—Bush and Cisler [17] noted that deconvolution algorithms should explicitly optimize representations of neural event encodings that are strictly positive. This hypothesis led to a reformulation of the general linear modeling approach of deconvolution to a nonlinear one, which we abbreviate as Bu13 for the remainder of this work. This approach models the BOLD signal, $\bar{\mathbf{y}}_i^{scan}$, as:

$$\begin{aligned} \hat{\mathbf{y}}_i(\hat{p}) &= z(\mathbf{F}_i(\hat{p}, \cdot) \mathbf{k}), \\ \bar{\mathbf{y}}_i^{scan}(p) &= \hat{\mathbf{y}}_i(p) + \eta_i^{Bu13}(p), \end{aligned} \quad (7)$$

where $\hat{\mathbf{y}}_i(p)$ is the approximate BOLD signal, $z(\cdot)$ is the normalization operation, $\mathbf{k}$ is the HRF kernel, η_i^{Bu13} is a vector of the algorithm's residual error, and $\mathbf{F}_i$ is a modified Toeplitz matrix of size $P + (K - 1) \times K$, given by

$$\mathbf{F}_i(\hat{p}, k) = \begin{cases} \tilde{\mathbf{e}}_i(\hat{p} - k) & \hat{p} - k > -K, \quad \hat{p} \in 2 - K, \dots, 1, \dots, P, \\ 0 & \text{otherwise,} \end{cases} \quad (8)$$

where the change between indices p and $\hat{p}$ reflects that the Bu13 algorithm estimates latent neural events that can influence the observed BOLD signal but occur prior to the start of observation; also, $\tilde{\mathbf{e}}_i \in (0, 1)$.

Each element of $\tilde{\mathbf{e}}_i$ represents the magnitude of neural activity (a value 0.5 equates to mean neural activity, whereas 0 and 1.0 represent minimum and maximum neural activity, respectively). To achieve this desired range the deconvolution algorithm assumes that neural events are driven by an unobserved time-series of real-valued neural activations, $\mathbf{a}_i(p) \in \mathbb{R}$,

that are temporally independent and whose values determine the neural event encoding via the logistic function,

$$\tilde{\mathbf{e}}_i(\hat{p}) = \frac{1}{1 + \exp(-\mathbf{a}_i(\hat{p}))}. \quad (9)$$

Using this model, we deconvolve the BOLD observations by optimizing neural activations, $\mathbf{a}_i$, such that they minimize the cost function, $\mathbf{J}_i^{Bu13}$, given by

$$\begin{aligned} \mathbf{J}_i^{Bu13}(p) &= \frac{1}{2} (\tilde{\mathbf{y}}_i(p) - \bar{\mathbf{y}}_i^{scan}(p))^2, \\ &= \frac{1}{2} (\eta_i^{Bu13}(p))^2. \end{aligned} \quad (10)$$

See Bush and Cisler [17] for details on the methods used to minimize Eq. (10).

2.3. Spectral subtraction filtering

We employed a spectral subtraction fast Fourier transform(SS-FFT) filter as the baseline of performance. This filter applies the discrete FFT to the observed BOLD signal, which maps the signal into the frequency domain. It then zeroes out (via boxcar shape function) all transform components corresponding to frequencies outside of the range [0.01, 0.1] Hz and applies the inverse FFT transform to reconstruct a version of the observed BOLD signal. This frequency band aligns, approximately, with previously reported band-pass filtering studies [8,24]. This frequency band was also shown to be involved in resting state functional connectivity [25,26] and is often used as the spectral subtraction target when temporal smoothing fMRI; this filter represents a standardized approach found in the literature.

2.4. Fast Fourier transform filtering

We also employed a shaped band-pass fast Fourier transform (FFT) filter². This filter maps the BOLD signal into the frequency domain and selects frequency components via a Gaussian kernel centered at 0.05 Hz, selective for frequencies on the range 0.005–0.15 Hz. This frequency range also approximately aligns with previously reported band-pass ranges [8,24], which select for frequency components indicative of the canonical HRF function; however, the frequency range described was identified via cross-validation such that the resultant filter achieved good performance in our simulations (default model). After frequency domain selection, the signal was inverted back into the original BOLD signal space. Note the following: (1) in practice this tuning process is not possible as the performance of a filter cannot be directly measured against the true underlying BOLD signal; (2) the SS-FFT filter, described in Section 2.3, is analogous to this filter except for the shape function that selectively removes frequency components; and, (3) our FFT filter is intended to represent a reasonable, well-tuned filter that simulates the best possible FFT-based filter achievable in practice given an excellent simulation of resting state BOLD signal.

²We utilized the FouFilter (version 1.5, May 2007) algorithm developed by Tom O'Haver available at the Matlab File Exchange.

2.5. Optimal fast Fourier transform filtering

We also constructed an optimized band-pass FFT filter to facilitate our analysis of the effects of filtering on functional connectivity estimation. For each BOLD signal we generate a set of band-pass filters, each identical to the FFT filter described in Section 2.4 but varying in the range of frequencies for which they are selective. We then select the optimal filter, for a specific BOLD signal on a specific trial, as the filter yielding the lowest error according to the measures of filtering performance to be defined in Section 3.

The set of filters (from which the optimal is chosen) is constructed as follows. The lower frequency bound is defined to be 0.005 Hz. The upper frequency bound is sampled on the range [0.01,10.0] Hz at intervals of 0.01 Hz. Using a Gaussian kernel to select frequencies on the defined range (described above), we define the center of the kernel to be the midpoint of the frequency range. The upper bound of the frequency range 10.0 Hz was selected because the filter constructed from this bound reconstructs the original BOLD signal with error indistinguishable from zero at machine precision (i.e., does not filter the simulated BOLD signal).

We should note that this filter, like the FFT filter from Section 2.4, can never be constructed in practice and is used as a practical upper bound to the performance of FFT-based filtering. Selection of the optimal filter parameters in this manner requires knowledge of the true parameters generating the BOLD observations, which in practice are never known. The use of simulated BOLD signals permit employment of optimal filters as the strongest comparative measure against the proposed deconvolution filtering method.

The combination of this optimal filter, along with the SS-FFT and FFT filters, described above, represent a set of benchmarks on which to assess the performance of deconvolution filtering. It is assumed that, in practice, a temporal smoothing method would perform within the performance bounds achieved by SS-FFT (lower bound) and the optimal FFT filter (upper bound). Our performance analysis was based on measuring deconvolution filtering relative to these performance extremes.

3. Experiments

The goal of our experiments is to examine the impact of deconvolution filtering on the performance of standard fMRI analysis. To achieve this we measured the accuracy of BOLD signal variants (i.e., unfiltered, deconvolution filtered, band-pass filtered, and optimized band-pass filtered BOLD signals) in identifying useful properties of highly confounded, simulated BOLD data: (1) reconstructing the true, unconfounded BOLD signal, (2) correlation with the true, unconfounded BOLD signal, and (3) reconstructing the true functional connectivity of a three-node neural system. We also measure the degree to which deconvolution filtering impacts the magnitude of detected activations on functional connectivity analysis of real fMRI data. These experimental approaches are detailed in this section.

3.1. Simulation experiments

A convenient way to describe the simulated experiments is to describe the default configuration, followed by identifying each experiment's specific deviation from default. The default model's functional network is composed of three nodes³, $V = 3$. The connectivity matrix, $\mathbf{C}$, contains all non-zero entries. Throughout these experiments, we define regions i and j to be *connected* when $\mathbf{C}(i,j) = \text{logistic}(4) \approx 0.982$ and *unconnected* when $\mathbf{C}(i,j) = \text{logistic}(-4) \approx 0.018$. In the default network, $\mathbf{C}(1,2)$ and $\mathbf{C}(2,3)$ are connected and all other entries are unconnected, as depicted in Fig. 2(a). The lag matrix is defined to be $\mathbf{L}(i,j) = 0.050$ s, $i,j \in 1, \dots, V$, conforming to a similar network simulation approach [27]. The vector of the fraction of external neural activity for each node is $\rho(i) = 0.05$, $i \in 1, \dots, V$.

Consistent with prior work [17], for each node of the network, neural events are generated at $\nu_g = 20.0$ Hz for 200 s. Neural events in the default model are converted to a theoretical BOLD signal via the canonical HRF (time-to-peak 6 s), depicted in Fig. 1 (a). The default observation model is executed independently for each node in the network. The BOLD signal is observed at $\nu_o = 0.5$ Hz. The physiological signal-to-noise is $\text{SNR}_{\text{phys}} = 6.0$ and the autocorrelation coefficient of this process was $\gamma = 0.75$. The signal-to-noise ratio of the scanner noise process is $\text{SNR}_{\text{scan}} = 9.0$. The ratio of physiological to scanner noise as well as the autocorrelation coefficient are estimated from prior work that analyzed appropriatemagnitudes and ratios of these values [6,9,10]. The defaultmodel removes latent neural events (i.e., the first $K - 1$ observations) and each node's BOLD signal is normalized by percent signal change. An example observation is depicted in Fig. 1(b). Each trial of an experiment is an independent random sample of the external (prior) neural events, conditional neural events, and observation processes (where appropriate) for each region, i , and each simulation step, n . Thirty trials are generated for each experimental parameter configuration.

3.1.1. Connectivity structure—We examined four possible directed topologies of the three-node default model. The configurations are shown in Fig. 2. The default model corresponds to panel (a). Each panel of Fig. 2 is labeled with a shorthand denoting which nodes are *connected* and the direction of causality, e.g., the default network is denoted $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$.

3.1.2. Observation frequency—An important experimental parameter, which has been shown to impact deconvolution performance, is observation frequency [17]. Therefore, we vary this parameter from the default value of $\nu_o = 0.50$ Hz (equivalent to $\text{TR} = 2$ s, a standard value) through much higher observation frequencies, $\nu_o \in \{0.5, 1.0, 2.0, 4.0, 5.0, 10.0\}$ Hz, which may be achieved *in vivo* through improved scan technology [28] or alterations to the scan sequence. Note, for all subsequent experiments, the default connectivity configuration, $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$, is used.

3.1.3. Misspecification of the HRF—We examined the performance impact when deconvolution algorithms are parameterized by an HRF kernel, $\mathbf{k}$, that is not the generating

³Except for direct comparisons between true BOLD signals and signal variants in which case $V = 1$.

kernel (i.e., HRF misspecification). To perform this experiment the model's mapping from neural events to a BOLD signal was altered to apply the balloon model rather than the canonical HRF. For each trial, the model sampled uniformly randomly from biologically plausible ranges of physiological parameters of the balloon model according to [16]. Note, the mean of the parameter ranges of the balloon model used in Havlicek et al. [16] generate an HRF with time-to-peak of 4 s. Therefore, for these experiments only, we modified the time-to-peak of the canonical HRF to be 4 s (instead of 6 s). We examine the impact of this confound with respect to observation frequency by repeating this experiment for $\nu_o \in \{0.5, 1.0, 2.0, 4.0, 5.0, 10.0\}$ Hz.

3.1.4. Noise magnitude—The default model uses $\text{SNR}_{\text{phys}} = 6.0$ and $\text{SNR}_{\text{scan}} = 9.0$. These values mimic an established ratio between physiological and scanner noise [9]. We acknowledge, however, that noise magnitudes vary depending on multiple factors. Thus, we attempted to establish robustness of the proposed methods over a wider range of noise. We executed simulations where the ratio of noise magnitudes, $\text{SNR}_{\text{phys}}/\text{SNR}_{\text{scan}}$, was maintained at 6/9, but we varied the total magnitude of noise, varying $\text{SNR}_{\text{phys}} \in 1, \dots, 8$.

3.1.5. Autocorrelation of physiological noise—The default model defines autocorrelation of the physiological noise process to be $\gamma = 0.75$, established experimentally in both Purdon et al. [6] and Lindquist [10]. To examine robustness of the proposed methods over a wider range of temporally structured noise, we examined the performance impact for simulations in which the autocorrelation coefficient was varied on range $\gamma \in \{0.1 - 0.9, 0.99\}$.

3.1.6. Lag variation—By default we uniformly assign causal lags between nodes to be the anatomical maximum values, L_{max} . To explore the impact of more plausible lag structures, we sampled the elements of the lag model, $\mathbf{L}$, uniformly randomly from an extreme range $[0.0, 0.100]$ s that subsumes anatomical variations. Further, we examined the performance impact of this confound with respect to observation frequency by repeating this experiment for $\nu_o \in \{0.5, 1.0, 2.0, 4.0, 5.0, 10.0\}$ Hz.

3.1.7. Connectivity strength—We also examined the relationship between connection strength and performance. We varied the meaning of *connected* to be *logistic (strength)*, $\text{strength} \in \{-4, \dots, 4\}$. We varied the two *connected* conditional probabilities independently (at incremental integer values), generating $9 \times 9 = 81$ unique networks.

3.2. Performance analysis

Signal reconstruction performance of signal variants was measured using RMSE and Pearson correlation. Functional connectivity performance was estimated as follows. For each trial, we computed Pearson correlation coefficients, $\phi(i, j)$, between all pairs of true BOLD signals, $\mathbf{x}_i, \mathbf{x}_j$; $i > j$, forming the true functional connectivity of the network, Φ_{true} , a vector of size $V^2/2$. Using the associated observed BOLD signals, $\bar{\mathbf{y}}_i^{\text{scan}}, \bar{\mathbf{y}}_j^{\text{scan}}; i > j$, we estimated the functional connectivity of the network for the BOLD signal variants: Φ_{obs} (unfiltered BOLD signal), Φ_{Bu13} (deconvolution filtering using Bu13 algorithm), Φ_{Gau11} (deconvolution filtering using Gau11 algorithm), $\Phi_{\text{ss-fft}}$ (baseline SS-FFT filter), and Φ_{opt}

(optimized band-pass FFT filter). See Fig. 1(a–c) for examples of the true BOLD signal, observed BOLD signal, and filtered signal variants that are used for performance analysis. Note, error and correlation calculations were made only for time points of the signal variants for which observed BOLD data exist (note, this omits some information generated by the Bu13 algorithm). We then computed the Euclidean distance, d , between each of these vectors and the true vector of coefficients, e.g., $d_{obs} = \|\Phi_{true} - \Phi_{obs}\|/2$ (smaller distance indicating better performance). This constitutes the measurement error. Finally, we plotted the distributions (over the 30 trials) of error according to their respective experiment and parameter configurations as a mean error, d as well as lower, a , and upper, b , confidence intervals such that $p(a \leq d \leq b) = 0.95$.

3.3. Real-world data

Nineteen healthy adolescent girls, aged 12–16 (mean ≈ 14.5), completed an emotion processing task during 3 T fMRI scanning. In this task, participants made button presses indicating gender decisions while viewing faces taken from the NimStim facial stimuli set [29]. The faces contained either neutral or fearful expressions, presented either overtly or covertly, in alternating blocks. Overt faces were presented for 500 ms, with a 1200 ms ISI displaying a blank screen with a fixation cross, in blocks of 8 presentations for a total block length of ~ 17 s. Covert face blocks used a similar design but were presented for 33 ms followed immediately by a neutral facial expression mask for 166 ms from the same actor depicted in the covert image, and the ISI was 1500 ms. Rest blocks were additionally included that displayed a blank screen with a fixation cross and lasted 10 s. The alternation between each block type (covert, overt, fear, neutral, rest) was additionally separated by 5000 ms of blank screen with a fixation cross. The task was presented in two runs, each lasting ~ 8 min, during which each block type was presented 5 times. Reaction times and response accuracy were recorded.

3.3.1. fMRI acquisition—A Philips 3 T Achieva X-series MRI system using an eight-channel head coil (Philips Healthcare, USA) was used to acquire imaging data. Anatomic images were acquired with a MPRAGE sequence (matrix = 192×192 , 160 sagittal slices, TR/TE/FA = min/min/ 8° , final resolution = $1 \times 1 \times 1$ mm³ resolution). Echo planar imaging sequences were used to collect the functional images using the following sequence parameters: TR/TE/FA = 2000 ms/30 ms/ 90° , FOV = 240×240 mm, matrix = 80×80 , 37 oblique slices (parallel to AC-PC plane to minimize OFC signal artifact), slice thickness = 3 mm, final resolution $3 \times 3 \times 3$ mm³.

3.3.2. Image preprocessing—Image preprocessing followed standard steps and was completed using AFNI [30] software. In the following order, images underwent despiking (3dDespike), slice timing correction (3dTshift), deobliquing (3dWarp), motion correction using rigid body alignment (3dvolreg), alignment to participant's normalized anatomical images (align_epi_anat.py), spatial smoothing using a 5 mm FWHM Gaussian filter (3dmerge), and rescaling into percent signal change (3dcalc and 3dTstat). Images were normalized using the MNI 452 template brain. The functional data remained in $3 \times 3 \times 3$ -mm resolution and were not resampled. To correct for respiratory and cardiovascular artifacts, fluctuations in white matter voxels and CSF were regressed out of time courses

from grey matter voxels following segmentation using FSL [31] and restricted maximum likelihood was used to account for autocorrelation (3dREML). In this step, singular value decomposition was applied separately to the white matter and CSF masks (3dmaskSVD) to define and extract the principal white matter and CSF timecourses to regress out of the grey matter voxels. This step was implemented directly after motion correction and normalization of the EPI images in the preprocessing stream. Additionally, to correct for residual motion artifacts in the signal that standard motion correction does not correct [32,33], Independent component analyses (ICA) were completed as a last preprocessing step separately on each participant and each run to identify and remove artifact components from the data [33–35]. In this step, ICA was implemented in Matlab using Group ICA of fMRI Toolbox (GIFT) [36] separately on each participant, in which the number of components to be solved is estimated from the data. The resulting components are then correlated with temporal confound templates (the SVD timecourses of white matter and CSF mentioned above), and components with timecourses correlating with these confound signals above a certain threshold (individually tuned to this particular data set) are then removed from the data using GIFT.

3.3.3. fMRI analysis—For the present purposes, we sought to define regions that were active during the task blocks versus the rest blocks. In the first-level analyses among each participant separately, we empirically defined the HRF by using 10 cubic splines, spaced evenly across a 19-s window, and corresponding to the onset of a task block (17 s) plus a TR length afterward to capture some of the delayed BOLD response to the block. This resulted in an HRF estimate (8 beta coefficients of percent signal change) for each of the four task conditions. We defined the HRF empirically, as opposed to using the canonical HRF, to avoid biasing the imaging results with the detection of regions that conform to the shape of the canonical HRF used by the deconvolution algorithms. Group-level analyses then consisted of voxel-wise one-sample *t*-tests, in which we contrasted area-under-the-curve estimates from the HRF for task vs rest (baseline) conditions. We controlled for multiple comparisons using cluster-level thresholding, in which a corrected alpha level of $p < .05$ was achieved through 17 contiguous voxels surviving an uncorrected $p < .005$, based on Monte Carlo simulations (AFNI's AlphaSim), which used 10,000 randomly generated data sets to determine the probability of observing clusters of differing sizes given an uncorrected *p* value [37]. A cluster size of 17 voxels with an uncorrected $p < .005$ was found in less than 5% of the randomly generated data sets, giving the corrected alpha level of .05.

These analyses identified three task-active regions in the left fusiform gyrus (cluster centroid in MNI coordinates: $X = -40$, $Y = -59$, $Z = -15$), dorsomedial PFC (cluster centroid: $X = -7$, $Y = 10$, $Z = 46$), and left motor cortex (cluster centroid: $X = -41$, $Y = -26$, $Z = 49$). The magnitude of the mean *F*-values (model fit values of the GLM from the task vs rest contrast) within each ROI then defined the *standard* baseline against which to compare the magnitude of effects from comparable analyses using the filtered data from the Bu13- and Gau11-based deconvolution filtering algorithms.

Following application of the Bu13- and Gau11-based deconvolution filtering algorithms to the timecourses of each voxel within the three ROIs, GLM-analyses were then conducted on these filtered timecourses that were identical to the GLM-analyses following *standard*

preprocessing. Mean F -values across the voxels within each ROI then defined the magnitude of task-activation effects following filtering to compare to the effects following *standard* preprocessing.

3.3.4. Comparisons between standard preprocessing and filtered data—We compared detection of task activation in two ways. First, we calculated the correlation across all 19 subjects of the magnitude of task activation (F -value from the GLM) between the standard preprocessing and the filtered approaches. This allows inferences regarding how similar results are between the filtered and non-filtered approaches. Second, we performed one-sample t -tests on the distribution of F -values from each of the three approaches. Given that filtering out noise should shrink the variance, it would be expected that accurately filtered data would result in higher t -values (i.e., shrinking variance in the context of a static mean would result in a higher t -value).

4. Results

4.1. Simulated data

Fig. 3 depicts a summary of the signal reconstruction accuracy achievable for the signal variants as a function of TR, $\log_2(\text{TR})$. This, and nearly all other performance results are plotted with respect to $\log_2(\text{TR})$ to reflect TR's role as a critical factor determining the efficacy of deconvolution algorithms [17]. Fig. 3(a) depicts a typical profile of signal reconstruction performance. Filtering performance is lowest at high TR (2 s) and improves as TR decreases. At high TR, SS-FFT and the deconvolution filtering algorithms struggle to significantly outperform unfiltered BOLD signal. In contrast, the filtering algorithms that were fit via knowledge of the underlying system (FFT and optimized FFT) perform well. The deconvolution algorithms are bracketed in performance by the SS-FFT and optimized FFT band-pass filtering algorithms (thus, subsequent analysis of filtering impact on functional connectivity omits FFT filter performance to reduce plot clutter, as it does not provide additional information). As TR decreases, the Bu13 deconvolution filter rapidly and continually improves performance whereas the Gau11 filter does not (much of the Gau11 filter error comes from the large reconstruction error at the beginning of the signal, see Fig. 1(c)). Unfiltered BOLD reconstruction error is independent of sample rate. Similar performance characteristics arise in correlation analysis, shown in Fig. 3(b); however, the confidence intervals of the mean performance are tighter.

Fig. 4 depicts a summary of the accuracy achievable using functional connectivity analysis across the four network topologies for each of the BOLD signal variants (unfiltered, deconvolution filtered, band-pass filtered, and optimized band-pass filtered). Fig. 4 exposes the key outcome of this work: filtering BOLD signals in all circumstances facilitates superior functional connectivity accuracy versus unfiltered BOLD signals. Moreover, deconvolution filtering (either Bu13 or Gau11) is statistically similar to baseline FFT filtering and is largely bounded by the optimized band-pass filter performance. Second, we observe that all filtering approaches benefit from faster TR, which stems from increased power to detect higher frequency components of the true signal structure in comparison to noise and this benefit is most pronounced up to TR = 0.5 s. Third, it is intuitive that analysis using unfiltered BOLD signals does not benefit from higher sample rates as the probability

of decorrelated scanner noise remains unchanged. Fourth, an interesting result of these experiments is that filtering has the greatest impact on functional connectivity structures that induce correlations between truly unconnected regions, e.g. $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$ and $\mathcal{F}_{1 \rightarrow 3}^{1 \rightarrow 2}$. This suggests that filtering removes a sufficient amount of the signal confound such that neural-level differences between signals are unmasked, e.g., intrinsic neural events and lag-effects.

Within the most challenging connectivity structure, $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$, the impact that other confounds have on functional connectivity estimation accuracy is summarized in Fig. 5. Panel (a) depicts the impact that HRF misspecification has on the performance of deconvolution filtering. Note that that band-pass filtering performances surpass those observed in Fig. 4(a), whereas the performance of the deconvolution filters remains more stable (in terms of error). The improvement in band-pass filtering arises from the use of a shorter time-to-peak in these experiments, 4 s versus 6 s. Fig. 5 suggests that deconvolution filters, in practice, are equivalent to well-fit band-pass filters over all: noise scales, panel (b); autocorrelation coefficient magnitudes, panel (c); and, simulations in which potential signal lag artifacts are removed through randomized lag connectivity, panel (d). The results also suggest that all filters improve functional connectivity estimation accuracy. This broad performance benefit of temporal filtering is an interesting result for fMRI network analysis.

Finally, Fig. 6 summarizes the robustness of the network correlation accuracy achieved by various filters with respect to a broad range of connection strengths within the $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$ network topology; we do this experiment to ensure that the results scale to the diversity of functional networks observed in practice. Highly correlated networks (in which both connections are strong) have the greatest total error (representing, possibly, the impact of chain-mediated connectivity), which is visible in the bottom-right corners (light gray pixels) of the image in panel (a). Optimized FFT filtering, shown in panel (e), is able to completely overcome these errors (all cells are dark gray). The temporal filters (despite the significant differences in the mathematical assumptions of each specific algorithm) produce similar performance distributions across signal intensity combinations. Thus, Fig. 6 suggests that temporal filtering (at least on small networks) is robust with respect to connection strength, which is important due to the wide range of task-mediated connection strengths that occur in real-world fMRI functional network analysis. Strongly divergent network connections are not necessary to achieve benefits; rather, the technique appears to be robust across intermediary network connection strengths as well.

4.2. Real data

As can be seen in Figs. 7–9, there were strong correlations between the three approaches in mean model fit F -values for each of the three ROIs (r s ranging from .86 to .98). As can be seen in Fig. 10, the one-sample t -test for the model fit F -values was higher for data filtered via Bu13-based deconvolution versus Gau11-based deconvolution across all three ROIs (range 8%–18% higher). The one-sample t -test for the model fit F -values was also higher for the Bu13-based deconvolution filter versus the unfiltered data for two ROIs (fusiform gyrus = 5% higher; motor cortex = 25% higher), but the standard approach yielded a higher one-sample t -test in the dmPFC (2% higher).

5. Discussion

The BOLD contrast mechanism, upon which fMRI analysis is based, is a multidetermined signal. It reflects not only neural processes and the neurovascular and metabolic consequences of neural processes, but it also reflects noise sources such as cardiac and respiratory signals, scanner drift, and motion-induced signal intensity changes. Consequently, inferences regarding neural processes from analysis of BOLD data are potentially confounded by these competing sources of variance in the BOLD signal. Therefore, methods to remove the confounding influence of noise from BOLD data can improve the validity of neural inferences.

Here, we developed deconvolution filtering as a method toward this goal. Deconvolution filtering is based on the concept of (1) deconvolving an observed BOLD signal into the maximally likely neural event estimate, and (2) reconvolving the neural event estimate with an assumed HRF to reconstruct a BOLD signal. The filtering aspect of this approach comes from the strength of the deconvolution algorithm, which effectively only estimates neural events for which the corresponding timepoints in the BOLD signal could have plausibly been caused by a neural event; that is, nonneural sources of variance in the BOLD signal are filtered. Thus, the reconstructed BOLD signal is theoretically only composed of neural and neurovascular signals. To test this approach we applied deconvolution filtering to both simulated data (where ground truth is known) as well as on real task-activation fMRI data from healthy control adolescent girls during an emotion processing task.

In simulation, deconvolution filtering was shown to be significantly beneficial for removing signal confounds under a wide range of simulation conditions intended to account for real-world variations in scanner performance, intrinsic noise, and the neural system of interest. This conclusion holds for nearly all sample rates (the primary driver of filter performance) and was independent of the deconvolution algorithm used (either Bu13 or Gau11).

When comparing deconvolution filtering against band-pass frequency filtering, significant performance variations were observed under some simulation conditions. When attempting to reconstruct the true BOLD signal from a highly confounded BOLD signal with low sample rate (high TR), band-pass filtering was equivalent to deconvolution filtering. At higher sample rates, however, the performance of Bu13-based filtering demonstrated significantly better performance than Gau11-based filtering but similar performance to optimized band-pass filtering.

Within experiments designed to simulate functional connectivity analysis, deconvolution filtering performance was indistinguishable from band-pass performance; however, both filtering approaches showed significant improvement over analysis using unfiltered BOLD. Moreover, under highly confounded simulation parameters (those deemed most representative of real-world signal), deconvolution filtering often (particularly for sample rates $TR < 2$ s) exhibited performance that was not significantly different from optimized band-pass filtering in which the frequency band was selected such that the estimated functional connectivity best matched the true functional connectivity; we note that such an

optimized filter can never, in practice, be constructed because it requires knowledge of the true BOLD signal.

The results from analysis of the real fMRI data suggest two inferences. First, deconvolution filtering provided superior or equivalent detection of task activations relative to comparable analyses on non-deconvolution filtered data. Note, the regions used for this analysis were identified using the standard analysis methods (GLM analysis of unfiltered BOLD data); thus, this test is biased toward supporting the standard methodology. Nonetheless, deconvolution filtering (1) provided estimates of task activation highly correlated with the standard estimates, and (2) suggested decreased variance in these estimates. This latter result is exactly as hypothesized: a successful filter should retain the true signal while blocking noise signals. The removal of noise (but not signal) decreases the standard error of the sample estimates, consequently increasing the effect size of task activation estimates. These results are encouraging and suggest that deconvolution filtering removes noise variance without removing true signal variance, supporting this approach's application to real data.

In summary, results from both simulated and real data suggest the utility of deconvolution filtering in removing noise variance and improving estimates of true neural processes. Unlike band-pass filtering, deconvolution filtering achieves this performance using only physiological-based assumptions of the underlying signal. We hope continued testing and implementation of this approach will improve the overall quality of inferences in clinical and cognitive neuroscience. Additionally, future research will be necessary to test the deconvolution filtering approach on more varied sources of real fMRI data—for example, across different types of tasks across different types of populations (children, adults, clinical populations, etc). Nonetheless, the current data reported here suggest the utility of the approach, although more testing is needed before widespread application is recommended.

References

1. Friston K, Holmes A, Poline JB, Grasby P, Williams S, Frackowiak R. Analysis of fMRI time-series revisited. *NeuroImage*. 1995; 2:45–53. [PubMed: 9343589]
2. Friston K, Frith C, Turner R, Frackowiak R. Characterizing evoked hemodynamics with fMRI. *NeuroImage*. 1995; 2:157–65. [PubMed: 9343598]
3. Friston K, Frith C, Frackowiak R, Turner R. Characterizing dynamic brain responses with fMRI: a multivariate approach. *NeuroImage*. 1995; 2:168–72.
4. Worsley K, Friston K. Analysis of fmri time-series revisited—again. *NeuroImage*. 1995; 2:172–81.
5. Purdon P, Weisskoff R. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Hum Brain Mapp*. 1998; 6:239–49. [PubMed: 9704263]
6. Purdon P, Solo V, Weisskoff R, Brown E. Locally regularized spatiotemporal modeling and model comparison for functional MRI. *NeuroImage*. 2001; 14:912–23. [PubMed: 11554810]
7. Skudlarski P, Constable R, Gore J. Roc analysis of statistical methods used in functional MRI: individual subjects. *NeuroImage*. 1999; 9:311–29. [PubMed: 10075901]
8. Friston K, Josephs O, Zarahn E, Holmes A, Rouquette S, Poline JB. To smooth or not to smooth? Bias and efficiency in fMRI time-series analysis. *NeuroImage*. 2000; 12 (2):196–208. [PubMed: 10913325]
9. Kruger G, Glover G. Physiological noise in oxygenation-sensitive magnetic resonance imaging. *Magn Reson Med*. 2001; 46:631–7. [PubMed: 11590638]
10. Lindquist M. The statistical analysis of fMRI data. *Stat Sci*. 2008; 23(4):434–64.

11. Friston K, Frith C, Liddle P, Frackowiak R. Comparing functional (PET) images: the assessment of significant change. *J Cereb Blood Flow Metab.* 1991; 11:690–9. [PubMed: 2050758]
12. David O, Guillemain I, Saillet S, Reyt S, Deransart C, Segebarth C, et al. Identifying neural drivers with functional MRI: an electrophysiological validation. *PLoS Biol.* 2008; 6(12):2683–97. [PubMed: 19108604]
13. Friston K. Dynamic causal modeling and granger causality comments on: the identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. *NeuroImage.* 2011; 58(2):303–5. [PubMed: 19770049]
14. Valdes-Sosa P, Roebroeck A, Daunizeau J, Friston K. Effective connectivity: influence, causality and biophysical modeling. *NeuroImage.* 2011; 58(2):339–61. [PubMed: 21477655]
15. Gaudes C, Petridou N, Dyrden I, Bai L, Francis S, Gowland P. Detection and characterization of single-trial fMRI bold responses: paradigm free mapping. *Hum Brain Mapp.* 2011; 32(9):1400–18. [PubMed: 20963818]
16. Havlicek M, Friston K, Jan J, Brazdil M, Calhoun V. Dynamic modeling of neuronal responses in fMRI using cubature Kalman filtering. *NeuroImage.* 2011; 56(4):2109–28. [PubMed: 21396454]
17. Bush K, Cisler J. Decoding neural events from fMRI BOLD signal: a comparison of existing approaches and development of a new algorithm. *Magn Reson Imaging.* 2013; 31(6):976–89. [PubMed: 23602664]
18. Hernandez-Garcia L, Ulfarsson M. Neuronal event detection in fMRI time series using iterative deconvolution techniques. *Magn Reson Imaging.* 2011; 29(3):353–64. [PubMed: 21232893]
19. Roebroeck A, Formisano E, Goebel R. Mapping directed influence over the brain using granger causality and fMRI. *NeuroImage.* 2005; 25(1):230–42. [PubMed: 15734358]
20. Schippers M, Renken R, Keysers C. The effect of intra- and inter-subject variability of hemodynamic responses on group level granger causality analyses. *NeuroImage.* 2011; 57(1):22–36. [PubMed: 21316469]
21. Buxton R, Wong E, Frank L. Dynamics of blood flow and oxygenation changes during brain activation: the balloon model. *Magn Reson Med.* 1998; 39(6):855–64. [PubMed: 9621908]
22. Friston K, Mechelli A, Turner R, Price C. Nonlinear responses in fMRI: the balloon model, volterra kernels, and other hemodynamics. *NeuroImage.* 2000; 12(4):466–77. [PubMed: 10988040]
23. Seln Y, Abrahamsson R, Stoica P. Automatic robust adaptive beamforming via ridge regression. *Signal Process.* 2008; 88(1):33–49.
24. Biswal B, Yetkin FZ, Haughton V, Hyde J. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn Reson Med.* 1995; 34:537–41. [PubMed: 8524021]
25. Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, et al. Frequencies contributing to functional connectivity in the cerebral cortex in “resting-state” data. *Am J Neuroradiol.* 2001; 22:1326–33. [PubMed: 11498421]
26. Fox MD, Snyder AZ, Vincent JL, Corbetta M, Essen DCV, Raichle ME. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci.* 2005; 102:9673–8. [PubMed: 15976020]
27. Smith S, Miller K, Salimi-Khorshidi G, Webster M, Beckmann C, Nichols T, et al. Network modelling methods for fMRI. *NeuroImage.* 2011; 54(2):875–91. [PubMed: 20817103]
28. Feinberg D, Moeller S, Smith S, Auerbach E, Ramanna S, Glasser M, et al. Multiplexed echo planar imaging for sub-second whole brain fMRI and fast diffusion imaging. *PLoS ONE.* 2010; 5(12):e15710. [PubMed: 21187930]
29. Tottenham N, Tanaka JW, Leon AC, McCarry T, Nurse M, Hare TA, et al. The nimstim set of facial expressions: judgments from untrained research participants. *Psychiatry Res.* 2009; 168(3):242–9. [PubMed: 19564050]
30. Cox R. AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res.* 1996; 23(3):162–73. [PubMed: 8812068]
31. Smith SM, Jenkinson M, Woolrich MW, Beckmann CF, Behrens TE, Johansen-Berg H, et al. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage.* 2004; 23(Supplement 1)(0):S208–19. [PubMed: 15501092]

32. Friston K, Williams S, Howard R, Frackowiak R, Turner R. Movement-related effects in fMRI time-series. *Magn Reson Med*. 1996; 35:346–55. [PubMed: 8699946]
33. Tohka J, Foerde K, Aron A, Tom S, Toga A, Poldrack R. Automatic independent component labeling for artifact removal in fMRI. *NeuroImage*. 2008; 39:1227–45. [PubMed: 18042495]
34. Kelly RE, Alexopoulos G, Wang Z, Gunning F, Murphy C, Morimoto S. Visual inspection of independent components: defining a procedure for artifact removal from fMRI data. *J Neurosci Methods*. 2010; 189(2):233–45. [PubMed: 20381530]
35. Zeng W, Qiu A, Chodkowski B, Pekar JJ. Spatial and temporal reproducibility-based ranking of the independent components of BOLD fMRI data. *NeuroImage*. 2009; 46(4):1041–54. [PubMed: 19286465]
36. Calhoun VD, Adali T, Pearlson GD, Pekar JJ. A method for making group inferences from functional MRI data using independent component analysis. *Hum Brain Mapp*. 2001; 14:140–51. [PubMed: 11559959]
37. Forman SD, Cohen JD, Fitzgerald M, Eddy WF, Mintun MA, Noll DC. Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): use of a cluster-size threshold. *Magn Reson Med*. 1995; 33 (5):636–47. [PubMed: 7596267]

Fig. 1. Deconvolution filtering for network node i : (a) the true BOLD signal, $\mathbf{x}_i$, temporally aligned with timings of neural events (vertical lines); (b) observed BOLD signal, $\bar{y}_i^{scan}(p)$; and, (c) filtered BOLD signals, $y_i^{ftr}(\hat{p})$, reconvolved from the Gau11 and Bu13 approximate encodings as well as the baseline band-pass SS-FFT filter applied directly to the observed BOLD signal.

Fig. 2.

Graphical depiction of the four network connectivities utilized in this research. The default network, $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$, is depicted in the top-left. Solid arrowheads indicate *connected* network nodes, providing the direction of the conditional probability. Open arrowheads indicate the presence of external neural events.

Fig. 3.

Comparison of reconstruction performance of the true, unconfounded BOLD signal achieved by variants of the observed BOLD signal over a wide range of observation sample rates. Comparison is provided in terms of (a) RMSE and (b) Pearson correlation. Note, existing fMRI technology permits TR values on the range of 1–2 s, thus, $\log_2(\text{TR})$ will be on the range 0–1 for this technology (the right-hand side of both figures).

Fig. 4. Comparison of functional connectivity errors achieved by BOLD signal variants observed from the four network topologies utilized in simulation over a wide range of observation sample rates.

Fig. 5.

Comparison of functional connectivity errors achieved by various filtering methods when applied to BOLD signals of the default network, $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$, under confounds (a) misspecification of the HRF; (b) variation in the signal-to-noise ratio of scanner noise, SNR_{scan} (default model $\text{SNR}_{\text{scan}} = 6.0$); (c) variation in the autocorrelation coefficient, AR_{phys} , of the physiological noise process (default model $\text{AR}_{\text{phys}} = 0.75$); and (d) application of a randomized lag model, $\mathbf{L}$, to the default network.

Fig. 6.

Comparison of correlation errors achieved by various filtering methods when applied to BOLD signals of the default network, $\mathcal{F}_{2 \rightarrow 3}^{1 \rightarrow 2}$, as the conditional probabilities of *connected* network nodes, computed as $\text{logistic}(\text{strength})$, were redefined by varying the *strength* parameter on the range $\text{strength} \in [-4, 4]$, sampled at intervals of 1. Note, without filtering (a) correlation error is proportional to the connectedness of the network. This is also weakly true for the SS-FFT based band-pass filtering (b) but is largely overcome by both deconvolution filters (c, d) and is completely overcome by optimized FFT band-pass filtering.

Fig. 7. Mean model fit F -values for the fusiform gyrus between standard preprocessing (unfiltered BOLD), Gau11 filtered BOLD, and Bu13 filtered BOLD.

Fig. 8. Mean model fit F -values for the motor cortex between standard preprocessing (unfiltered BOLD signals), Gau11 filtered BOLD signals, and Bu13 filtered BOLD signals.

Fig. 9. Mean model fit F -values for the dorsomedial prefrontal cortex between standard preprocessing (unfiltered BOLD signals), Gau11 filtered BOLD signals, and Bu13 filtered BOLD signals.

Fig. 10.

Comparison of one-sample t -test statistics for the model fit F -values across regions of interest and between standard preprocessing (unfiltered BOLD signals), Gau11 filtered BOLD signals, and Bu13 filtered BOLD signals.