Published in final edited form as:

Magn Reson Imaging. 2015 December ; 33(10): 1314–1323. doi:10.1016/j.mri.2015.07.007.

Improving the Precision of fMRI BOLD Signal Deconvolution with Implications for Connectivity Analysis

Keith Bush^a, Josh Cisler^b, Jiang Bian^c, Gokce Hazaroglu^a, Onder Hazaroglu^a, and Clint Kilts^b

Keith Bush: kabush@ualr.edu; Josh Cisler: jcisler@uams.edu; Jiang Bian: jbian@uams.edu; Gokce Hazaroglu: gxhazaroglu@ualr.edu; Onder Hazaroglu: oxhazaroglu@ualr.edu; Clint Kilts: cdkilts@uams.edu

^aDepartment of Computer Science, University of Arkansas at Little Rock (UALR), 2801 S. University Ave., Little Rock, Arkansas, USA 72204

^bBrain Imaging Research Center, University of Arkansas for Medical Sciences (UAMS), 4301 W. Markham St., Little Rock, Arkansas, USA 72205

^cDivision of Biomedical Informatics, University of Arkansas for Medical Sciences (UAMS), 4301 W. Markham St., Little Rock, Arkansas, USA 72205

Abstract

An important, open problem in neuroimaging analyses is developing analytical methods that ensure precise inferences about neural activity underlying fMRI BOLD signal despite the known presence of confounds. Here, we develop and test a new meta-algorithm for conducting semi-blind (i.e., no knowledge of stimulus timings) deconvolution of the BOLD signal that estimates, via bootstrapping, both the underlying neural events driving BOLD as well as the confidence of these estimates. Our approach includes two improvements over the current best performing deconvolution approach; 1) we optimize the parametric form of the deconvolution feature space; and, 2) we pre-classify neural event estimates into two subgroups, either *known* or *unknown*, based on the confidence of the estimates prior to conducting neural event classification. This knows-what-it-knows approach significantly improves neural event classification over the current best performing algorithm, as tested in a detailed computer simulation of highly-confounded fMRI BOLD signal. We then implemented a massively parallelized version of the bootstrapping-based deconvolution algorithm and executed it on a high-performance computer to conduct large scale (i.e., voxelwise) estimation of the neural events for a group of 17 human subjects. We show that by restricting the computation of inter-regional correlation to include only those neural events estimated with high-confidence the method appeared to have higher sensitivity for identifying the default mode network compared to a standard BOLD signal correlation analysis when compared across subjects.

Corresponding Author: Keith Bush, Ph.D., Department of Computer Science, University of Arkansas at Little Rock, 2801 S. University Ave., Little Rock, AR 72204, kabush@ualr.edu, Phone: 501.569.8143, Fax: 501.569.8144.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

1. Introduction

Functional magnetic resonance imaging (fMRI) is the predominant methodology of contemporary neuroimaging and is most commonly performed using blood-oxygen-level dependent (BOLD) contrast [1]. BOLD imaging is founded on neuronal activity-dependent changes in local magnetic fields [2,3,4]. Neuronal activity is followed by a reliable influx of oxygenated hemoglobin molecules that alters the ratio between oxygenated and deoxygenated hemoglobin molecules in the local blood supply [1]; due to oxygen's role in masking the magnetic field of hemoglobin, this changing ratio alters the local magnetic field surrounding the neural activity and is captured as BOLD. This local change in magnetic field is the basis of the BOLD contrast mechanism. Differences in detected signal between two experimental states (e.g., viewing faces versus no stimulation) reflect differences in local ratios between oxygenated and deoxygenated hemoglobin, which in turn reflect differences in neural activity. While the exact physiological mechanisms mediating the BOLD contrast mechanisms are not clear, research has demonstrated that the BOLD signal is strongly correlated with local field potentials and is a valid, though indirect, measure of neural activity [1,4].

The mediating relationship between neural activation and BOLD contrast is the hemodynamic response function (HRF) [5,6], which is well-approximated by a double gamma kernel function, shown in Figure 1a. Via the HRF, BOLD signal evolves through time proportionally to the neural activation. HRFs are also assumed to be linearly additive; those occurring in close proximity produce a BOLD response that sums the individual HRF functions, as pictured in Figure 1b. Noise processes (both physiological and thermal) confound real-world fMRI BOLD signal acquisition [7,8], depicted in Figure 1c.

One limitation of the HRF is the need to assume parameters which govern its temporal characteristics (e.g., time-to-peak 6 s). In practice, HRF time-to-peak parameters vary substantially (± 2 s) across subjects and across different regions of the brain [9,10,11,12]. Such variations raise serious questions concerning the validity of inferences, particularly causal inferences [13], drawn directly from BOLD [14,15]. E.g., by pairing high-temporal fidelity intracerebral EEG recordings with fMRI recordings of the same brain regions, David et al. convincingly showed that region-wise variations in the HRF confounded the discovery of key causal relationships directly from BOLD data [14]; rather, it was shown that deconvolution (the technique of solving for the underlying neural activations of a BOLD signal) was found to be a necessary condition for identifying the true, directed brain organization. As a result of this work, there has been renewed interest in the development of improved deconvolution algorithms [16,17,18,19]. The purpose of this research is to demonstrate the ability of two well-known machine learning techniques to improve the classification performance of existing, semi-blind deconvolution algorithms on highly-confounded fMRI BOLD signal. It is also shown that improved deconvolution increases accuracy of detecting simulated BOLD functional connectivity and increased statistical power for detecting real BOLD functional connectivity among healthy adults.

In a recent survey of deconvolution efficacy [18], an explicit inverse model of fMRI BOLD signal proved to be the most robust deconvolution approach over both simulated and real-

world datasets. A key innovation of this model was the representation of discrete neural events (termed encodings) via a continuously differentiable function of neural activation. This approach facilitated gradient optimization while naturally bounding the encoding to the range of zero to one without regularization, a significant improvement over competing algorithms. In the original study, the functional mapping from neural activations onto encodings was modeled as the canonical logistic function. However, a continuum of logistic function shapes exist [20], approaching the step-function at its parametric extreme. Thus, it is an interesting, open question as to whether non-canonical variants of the logistic function may achieve better generalization performance in deconvolving confounded fMRI BOLD signal. We will term this approach *tuned-deconvolution* for the remainder of this manuscript.

Another simple and well-known strategy for improving classification performance is to select predictive features using domain knowledge [21]. One such form of domain knowledge, highly relevant to the issue of deconvolution, is knowledge of the presence of confounds in BOLD signal, including autoregressive and thermal noise as well as sample-rate and normalization effects. Such confounds introduce errors into the encoding estimates. Thus, we propose a new approach to deconvolution in which both the encoding and its confidence are simultaneously estimated via the bootstrapping method [22], a sample-based approach. As part of this approach we propose ignoring deconvolved encodings for which there exists little confidence in the estimate. The consequence of this approach is that it introduces a class label for each neural activation, either *known* or *unknown*. This label is in addition to the canonical label of *active* or *inactive*, i.e., whether or not a neural activation occurred, which is calculated during classification. A goal of interest for the deconvolution community, therefore, is identifying the level of classification performance achievable through confidence-based classification. We term this approach *resampled-deconvolution* for the remainder of the manuscript.

The primary value of this research stems from an improved ability to identify specific timings of neural activations. These results may directly impact ongoing research into effective connectivity [23] estimation algorithms, which strive to address fundamental questions of directed brain organization and function [24,25,26,27]. There also exists an immediate impact of this work on undirected measures of brain organization, such as functional connectivity [23]. Given concerns over variation in the HRF, described above, there is reason to examine functional connectivity at the level of neural activation rather than that of BOLD signal. In support of this, prior work demonstrated that variability in HRF has little impact on the performance of deconvolution algorithms [18], suggesting that estimates of functional connectivity conducted among neural event estimates are more robust to variations in HRF shape compared to functional connectivity conducted from the BOLD signal. Moreover, resampled-deconvolution provides the means of assessing confidence in the presence or absence of neural activations. This provides a unique measure of functional connectivity via *strong interactions* in which we create functional connectivity maps with increasingly higher requirements for the confidence of *known* neural activations. These maps, which we can compute for whole-brain fMRI BOLD signal recordings, allow us to examine and classify brain regions for which there are strong interactions.

The remainder of this manuscript details methodologies, experiments, and analysis by which *tuned-deconvolution* and *resampled-deconvolution* variants were constructed and tested. We then describe the application of resampled-deconvolution to two case studies of real-world, human fMRI BOLD signal (ROI and voxelwise) and provide evidence of both the validity of the approach as well as the utility of strong interaction functional connectivity maps as a neuroimaging analytical tool.

2. Materials and Methods

2.1 Simulated fMRI BOLD Signal

As in prior deconvolution studies [18] we employ a sophisticated parametric model of fMRI BOLD signal to simulate the experiments. The parameters of this model are detailed in the Appendix.

2.2 Base Deconvolution

We use the Bu13 algorithm [18] as the base deconvolution algorithm of the proposed meta-algorithm, which has been benchmarked against all known competing algorithms and shown to be robust to real-world confounds. For clarity in this work, we briefly describe the algorithm's design. The Bu13 algorithm models neural events as a vector of continuous values restricted to the range (0,1). To achieve this, it models the measured BOLD signal, $\mathbf{y}$, as a vector, $\tilde{\mathbf{y}}$, of length T , given by:

$$\tilde{\mathbf{y}} = z(\mathbf{F}\mathbf{h}) \quad [1]$$

where $\mathbf{F}$ is a feature matrix of size $T \times K$, $\mathbf{h}$ is the HRF kernel column-vector of length K , and $z(\cdot)$ is the normalization mapping. The feature matrix, $\mathbf{F}$, is a modification of the Toeplitz matrix such that

$$\mathbf{F}(i, k) = \begin{cases} \tilde{\mathbf{e}}(i-k) & : i-k > -K, i \in 2-K, \dots, M, k \in 1, \dots, K \\ 0 & : \text{otherwise} \end{cases}, \quad [2]$$

where $\tilde{\mathbf{e}}$ is the encoding (a vector of length $M + (K - 1)$, $\tilde{\mathbf{e}}(t) \in (0,1)$, $t \in 2 - K, \dots, M$). Each element of this vector 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 the desired range of $\tilde{\mathbf{e}}(t)$ the deconvolution algorithm assumes that neural events are driven by an unobserved time-series of real-valued neural activations, $\mathbf{a}$, $\mathbf{a}(t) \in \mathbb{R}$, $t \in 2 - K, \dots, T$, that are temporally independent and whose values determine the neural event encoding via the logistic function,

$$\tilde{\mathbf{e}}(t) = \frac{1}{1 + \exp(-\beta \cdot \mathbf{a}(t))}, \quad [3]$$

where $\beta = 1$. Using this model, we deconvolve the BOLD observations by optimizing neural activations, $\mathbf{a}$, such that they minimize the cost function, $\mathbf{J}$, given by

$$\mathbf{J} = \frac{1}{2}(\tilde{\mathbf{y}}(1:M) - \mathbf{y})^2. \quad [4]$$

2.3 Transfer function optimization

An important technical feature of the Bu13 algorithm is its representation of neural events using the canonical logistic function (Eqn 3) in which $\beta = 1$. This function is a model of the causal relationship between neural events and neural activation in fMRI BOLD signal. It has two useful properties: 1) it is continuously differentiable, which enables the application of gradient descent methods to optimize latent parameters; and, 2) its range is restricted to (0,1). Thus, this function mimics known biological constraints of neural events without additional regularization.

While the canonical logistic function proved valuable, there exists a continuum of logistic function shapes [20]. As a demonstration, multiple logistic shapes are presented in Figure 2, varying by their β parameter. As this parameter grows large, the logistic function exhibits an increasingly non-linear shape, approximating the step-function. Translated to the context of neuroimaging, the β parameter controls the logistic function's temporal precision. As β grows, optimized encodings become more binary in nature (i.e., the distribution of encoding values concentrate more at the extreme values 0 and 1). This is both more representative of the biological assumption as well as a more challenging constraint to satisfy. We hypothesize that there exists an optimal value of the β parameter that best generalizes this tradeoff.

2.4 Encoding variance analysis

From a machine learning perspective, deconvolution could be viewed as feature extraction incorporating physiological domain knowledge. We assess a deconvolution algorithm's performance by its ability to accurately reconstruct true neural events from BOLD signal. This is a binary classification problem. For each time-point of BOLD signal, there either is or is not an underlying neural event. Likewise, at each time-point, deconvolution produces a feature that either predicts or does not predict a neural event.

The confounding factors are sample rate and noise. Currently, the neural event generation rate is significantly faster than the sample rate of fMRI technology. This causes individual neural events, when viewed through the lense of the HRF, to blur together to form an observed neural firing rate, the effects of which have been investigated previously [18]. As BOLD response is linearly additive with respect to the true, underlying neural events, the consequence of low sample rate is that the standard linear model of fMRI BOLD signal encodes in real-valued, rather than binary, numbers:

$$\mathbf{y} = \mathbf{F}\mathbf{e} + \boldsymbol{\eta}, \quad [5]$$

where $\mathbf{y}$ is the measured BOLD signal, $\mathbf{F}$ is a convolution (Toeplitz) matrix mapping the HRF kernel onto BOLD, $\mathbf{e}$, $\mathbf{e}(t) \in \mathbb{R}$, is the encoding vector, and $\boldsymbol{\eta}$, $\boldsymbol{\eta}(t) \in \mathbb{R}$, is a vector of

i.i.d. Gaussian noise attributable to thermal effects. Deconvolution algorithms, canonically, approximate the encoding, $\tilde{\mathbf{e}}$, by inversion, i.e.,

$$\mathbf{e} = \mathbf{F}^{-1}(\mathbf{y} - \boldsymbol{\eta}). \quad [6]$$

Thus, noise entrained in the BOLD signal impacts the resultant deconvolution feature. Applied statistics has provided us with bootstrapping [22], a sample-based method for approximating the distribution over model parameters from observations. We employ a variation of bootstrapping, termed residual-resampling [28,29,30,31,32], which samples from the distribution governing $\boldsymbol{\eta}$ (which can be estimated empirically) to induce a distribution governing the encoding, $\tilde{\mathbf{e}}$.

2.5 Resampled-Deconvolution

Residual-resampling, when formulated for use in the deconvolution of fMRI BOLD, we term resampled-deconvolution. The detailed mathematical description is given in the Appendix. We summarize the algorithm here. Given an observed BOLD signal $\tilde{\mathbf{y}}^{scan}$, a deconvolution mapping h (parameterized by set Θ), a deconvolution filtering [33] mapping g (parameterized by set Φ), a parameter N describing the number of surrogate samples to form, and an operator ψ which uniformly randomly reorders the elements of its vector argument; then resampled-deconvolution is a four-step process: 1) deconvolve the BOLD signal into a maximally likely neural event encoding, $\tilde{\mathbf{e}}$; 2) reconvolve an approximate theoretical BOLD signal, $\mathbf{y}^{fltr}$; 3) compute the residuals between the observed BOLD signal and the approximate theoretical BOLD signal, $\mathbf{r} = \tilde{\mathbf{y}}^{scan} - \mathbf{y}^{fltr}$; and, 4) iteratively (N times) deconvolve random surrogate BOLD signals, $\mathbf{y}_n^{surr}$. Each random surrogate BOLD signal is formed by creating a uniformly random ordering of the residual vector, $\mathbf{r}_n$, via the operator ψ and then linearly combining this vector with the approximate theoretical BOLD signal, $\mathbf{y}_n^{surr} = \mathbf{y}^{fltr} + \mathbf{r}_n$. As the residual $\mathbf{r}$ is assumed to be a vector of i.i.d Gaussian noise (see Eqn 2), each surrogate, $\mathbf{y}_n^{surr}$, is a sample from an approximation of the distribution which generated the observed BOLD, $\tilde{\mathbf{y}}^{scan}$. Deconvolving these surrogates, therefore, forms a distribution over neural event estimates which should contain the true, underlying encoding, $\mathbf{e}$ (see Eqn 5).

2.6 Classification Performance Analysis

From the set of surrogate encodings, $\tilde{\mathbf{e}}_n^{surr}$, we calculate the mean of $\tilde{\mathbf{e}}^{surr}$ and the 95% confidence intervals of the mean of $\tilde{\mathbf{e}}^{surr}$, yielding a distribution over each $\tilde{\mathbf{e}}^{surr}(t)$. We denote the minimum interval at time t as $\tilde{e}_{-}^{surr}(t)$ and the maximum interval as $\tilde{e}_{+}^{surr}(t)$. From these, we construct a labeling of each time-point, $L(t)$ such that

$$L(t) = \begin{cases} \text{known} & \text{if } \tilde{e}_{-}^{surr}(t) > \mu_{+} \text{ or } \tilde{e}_{+}^{surr}(t) < \mu_{-} \\ \text{unknown} & \text{otherwise} \end{cases}, \quad [7]$$

where the upper threshold, μ_{+} , and the lower threshold, μ_{-} , are defined on the range (0,1). These thresholds may be considered the 95% confidence intervals of the mean value of the

average (i.e., background) neural firing rate from which we detect higher or lower aggregate average activity when scanning with a conventional TR. Given this interpretation, the labeling process (Eqn 7) is a test of statistical significance. All time-points for which $L(t) = \textit{known}$ are statistically significantly unique from the background neural activity (either higher or lower, $p < 0.05$). We can then alter this test of significance by varying the uncertainty of the background neural activity symmetrically about the value 0.5, $\mu_+ = 0.5 + \delta$ and $\mu_- = 0.5 - \delta$, where $\delta \in [0.0, 0.5)$. This can be interpreted as increasing the uncertainty of the background neural activity. It should be clear that 1) the value of the thresholds as well as the properties of the confidence intervals dictate the number of *known* or *unknown* labels; 2) it is possible that some datasets may contain only *unknown* labels; and, 3) the number of *unknown* labels should increase as δ approaches 0.5.

We evaluated the effectiveness of confidence interval thresholding to improve neural event detection according to the area-under-the-curve (AUC) of the receiver operator characteristic (ROC) curve calculated from classification of neural event timings underlying the true BOLD signal. Performance was measured as follows. For each trial, the resampled-deconvolution algorithm was executed, and, using the resultant surrogate encodings, the maximum interval, $\tilde{e}_+^{surr}$, and minimum interval $\tilde{e}_-^{surr}$, of the encoding were calculated. These confidence intervals were linearly interpolated (from 1 Hz sample rate to 20 Hz sample rate) to form high-resolution vectors that represent the maximum and minimum confidence intervals for neural events existing at the simulation's generation rate, i.e., 20 Hz. Each time point of the encoding is labeled according to Eqn 7 for a given δ value. A neural event was classified, $B(t)$, $\forall t : L(t) = \textit{known}$, such that

$$B(t) = \begin{cases} \textit{active} & \text{if } \tilde{e}_n^{surr}(t) > \gamma \\ \textit{inactive} & \text{otherwise} \end{cases}, \quad [8]$$

where γ is the *ROC-threshold* parameter. By comparing the detected encoding against the true encoding (i.e., the simulated neural events), the specificity and sensitivity of the detected encoding was calculated for each threshold value, $\gamma \in [0, 1]$ sampled at intervals of 0.01 (simulation), 0.05 (1000 ROI atlas), and 0.1 (whole-brain voxelwise), generating a 101-point ROC curve. Note, one ROC curve was generated for each trial for each δ parameter, $\delta \in [0.0, 0.5)$ sampled at intervals of 0.01.

Using the trapezoid rule to numerically integrate the AUC, we computed the distribution of AUCs achieved for each δ parameter. Each distribution constitutes the performance of the resampled-deconvolution algorithm at the respective δ parameter: random classification performance achieves AUC=0.5 and ideal classification performance achieves AUC=1.0.

Note, we assume a subordinate and unmodeled role of physiological noise in the temporal structure of the residuals, η (see Eqn 5), as it can largely be removed, via the regression of artifacts identified in the cerebrospinal fluid and white matter. Moreover, the effects that remnant physiological noise have on deconvolution, as modeled by an AR(1) process, have been systematically studied in prior work [18] and shown to have minimal performance impact for realistic simulation parameters. Minimal impacts on deconvolution, however, do

not guarantee similar impacts on the classification problem. For completeness, we include a study of the impact of autoregressive noise on classification performance in the Appendix.

2.7 Human data collection

Participants—We tested performance on real fMRI BOLD data collected among 17 healthy adults (9 female; mean age = 31.7; SD = 9.5) undergoing resting-state fMRI. Participants were included based on the absence of current mental health diagnoses, major medical conditions, or MRI contraindications (e.g., internal ferrous objects). The resting-state task was a 7.5 min scan during which participants were presented a fixation cross and told to lie still and not think about anything specific.

MRI acquisition—A 3T Achieva X-series small bore magnet using an 8-channel head coil (Philips Healthcare, USA) was used to acquire imaging data. Anatomic images were acquired with a MPRAGE sequence (matrix=192×192, 160 slices, TR/TE/FA=min/min/90°, final resolution=1×1×1 mm³ resolution). The echo planar imaging sequences used to collect the functional images were: TR/TE/FA=2000ms/30ms/90°, FOV=192×192mm, matrix=64×64, 34 oblique slices (parallel to AC-PC plane to minimize OFC signal artifact), slice thickness=3mm, final resolution 3×3×3 mm³.

Image preprocessing—Image preprocessing used AFNI software and followed standard steps. In order, images underwent despiking, slice timing correction, deobliquing, motion correction using rigid body alignment, alignment to participant's normalized anatomical images, spatial smoothing using a 6 mm FWHM Gaussian filter, temporal bandpass filtering (.01–.1 Hz), and scaling into percent signal change. Images were normalized using the ICBM 452 template brain. Additionally, to correct for respiratory and cardiovascular artifacts, fluctuations in white matter voxels and CSF were regressed out of time courses from grey matter (GM) voxels following segmentation using FSL [34] and using restricted maximum likelihood to account for autocorrelation. This step was implemented directly after motion correction and normalization of the EPI images in the preprocessing stream. Resulting images for each participant were inspected for artifacts and accuracy of normalization.

Data parcellation method—We conducted analyses with the human data using two methods: a voxelwise approach restricted within GM voxels, and an ROI approach using a whole-brain parcellation. For the GM voxelwise approach, each participant's high resolution T1 image was segmented into GM, white matter, and cerebrospinal fluid using FSL [34]. We then created a common group-level GM mask by selecting those GM voxels that were common among all participants (24,589 voxels). For the parcellation method, we used a previously published functional atlas [25,34], the 1000 ROI atlas, which used spatially constrained spectral clustering to reduce the voxel-wise timecourses into 883 unique regions-of-interest (ROIs). Given individual differences in brain coverage, there were 807 unique ROIs that were common across all 17 participants. We extracted the principal component timeseries of the voxels within each ROI through singular value decomposition (AFNI's 3dmaskSVD).

2.8 Parallelization

Resampled deconvolution of human data was made computationally feasible via high-performance implementation of the resampled-deconvolution algorithm. This implementation is based on massively parallelized variants of the deconvolution [18] and deconvolution filtering [33] algorithms, implemented via C++ threads [35] and Armadillo C++ linear algebra library [36]. Resampling was achieved first by voxel-wise deconvolution and deconvolution filtering of fMRI BOLD (2×807 tasks per subject for the 1000 ROI atlas and $2 \times 24,589$ tasks per subject for the GM voxelwise). Per subject encodings and filtered BOLD signals were then gathered to a central node where residuals were calculated and resampled ($N=30$) to form BOLD surrogates. Deconvolution was then performed surrogate-wise (30×807 tasks per subject for the 1000 ROI atlas and $30 \times 24,589$ tasks per subject for the GM voxelwise). Encodings of BOLD surrogates were then gathered to a central node for ROI/voxelwise calculation of statistical distributions. The parallelized code was on a Hewlett Packard ProLiant DL980 G7 Server (80 processors and 4 TB single-addressable memory). Benchmark performance for the experiments described in this work was 233.49 ± 58.62 ms (per task) executed concurrently on 40 threads.

3. Experiments

3.1 Tuned-deconvolution validation (simulated single voxel)

We sampled the transfer function shape parameter, β , from the set $\beta \in \{1, 4, 10, 20, 40, 60, 80\}$. We computed 30 random fMRI BOLD signals according to the default simulation parameters. We then deconvolved each BOLD signal using each of the shape parameters from the set of possible parameters and used the resulting encodings to classify the true, underlying simulated neural events. We defined the optimal β parameter to be the β parameter exhibiting the maximum of the lower 95% confidence interval of the mean performance.

3.2 Resampled-deconvolution validation (simulated single voxel)

We sampled 30 random fMRI BOLD signals from the default simulation parameter. Each BOLD signal was resampled-deconvolved using $N=30$ random surrogates and the default β parameter. We then used the resulting encoding distributions to classify the true, underlying simulated neural events as outline in Section 2.6.

3.3 Human resting-state functional connectivity analysis

We focused analyses on the well-established default mode network [37,38] and selected a posterior cingulate cortex (PCC) seed for the voxelwise analyses (6mm spherical ROI; MNI center-of-mass coordinates = $-2, -53, 24$) and 807 ROI analyses (1000 atlas number = 226; MNI coordinates = $0, -49, 19$). For each subject, we calculated the pearson r correlation between the mean PCC timecourse and either every other GM voxel or every other mean ROI timecourse for the GM and ROI approaches, respectively. All seedmaps then underwent r -to- z transformation. The seedmaps resulting from the resampled-deconvolution timecourses were generated across various levels of thresholding, $\delta \in [-.1, .45]$ (negative values violate the interpretation of δ as a scaling of the *unknown* confidence interval, however, it is a useful analytical tool for ensuring a 1.0 fraction of neural events). One-

sample t-tests were conducted to compare the distribution of z-scores against zero separately for each of the seedmaps for the GM voxelwise and ROI approaches, resulting in group-level statistical maps indicating the strength of relationships between the PCC and each ROI or voxel. For the voxelwise approach, we set a corrected $p < .05$ using cluster-level thresholding, defined as 48 contiguous voxels surviving an uncorrected $p < .01$. For the ROI approach, we used uncorrected $p < .01$.

4. Results

4.1 Tuned-deconvolution validation (simulated single voxel)

Classification performance as a function of the logistic function's shape parameter is presented in Figure 3. The optimized parameter, $\beta = 60$ (selection described in Section 3.1) improves classification performance by a small but significant amount (2.18%) over the Bu13 algorithm ($\beta = 1$).

4.2 Resampled-deconvolution validation (simulated single voxel)

Resampled-deconvolution results are summarized in Figure 4. Figure 4a depicts classification performance as a function of δ . As this parameter defines the width of the probability range for which encodings values are labeled *unknown*, increases in this parameter should correlate with extreme-valued encoding probabilities; and, therefore, should also correlate with improved classification accuracy. Indeed, as the δ reaches a value of 0.48, classification performance reaches a mean maximum, improving from mean .637 (AUC of ROC) to mean .723 (a 13.5% increase) over the default value, $\delta = 0$. Figure 4b depicts the fraction of samples classified as *known* as a function of δ . As the δ increases, the width of the *unknown* portion of the probability space grows and the *known* portion shrinks; this change is reflected by consistent decrease in the fraction of *known* observations as the δ approaches .5. The key feature represented by Figures 4a and 4b is that there exists a trade-off between classification accuracy and the fraction of neural events that may be classified. Even when $\delta = 0$, 8.8% of the observed samples are removed from the classification problem because they are labeled *unknown*. These samples would include encodings for which the mean encoding estimate has a confidence interval containing .5. Thus, even the default resampled-deconvolution algorithm produces a small (1.62%) but significant ($p < .05$) improvement over tuned-deconvolution.

To examine performance of resampled-deconvolution more deeply, an example result from the experimental trials is depicted, graphically, in Figure 5. Figure 5a depicts the distribution of estimated encodings (mean and 95% confidence intervals) overlaid with the estimated encoding (i.e., without resampling) and the true neural events generated by the simulation. Figure 5b depicts the distribution of the estimate BOLD signal (mean and 95% confidence intervals) overlaid with the observed BOLD signal generated via the simulation. The first observation to be made of this case study is that for both encoding and BOLD, resampling decreases the influence of noise on the estimation. The second observation, visible in Figure 5a, is that resampling greatly improves the robustness of the estimation at extremes of the neural event sequence. When neural event generation is relatively high- or low-frequency, the resampled encoding captures the local trend more robustly than deconvolution without

resampling. There is a cost to this robustness, however. In some instances, deconvolution without resampling is more responsive to local neural events, which is most clearly illustrated at simulation steps 2550–2650 in Figure 5a; in this case the resampled encoding filters out a precise detection of a short burst of neural events that was captured by the base deconvolution algorithm. The benefits of resampling, however, are clearly illustrated at simulation steps 1250–1500 and 1500–1700 during which resampled-deconvolution identifies the respective local trends in neural event frequency better (and more smoothly) than the base deconvolution algorithm.

4.3 Resampled-deconvolution parameter optimization (simulated single voxel)

Rejection of *unknown* neural event estimates requires an assumption of the threshold for certainty. An important trade-off exists between the amount of data kept for classification and the precision of this data. Experiments studying resampled-deconvolution on simulated, but highly-confounded, fMRI BOLD signal (see Figure 4a) suggest the existence of a critical threshold at which classification performance benefits are not realized (and may be diminished) by increasing the δ parameter. In real-world application of resampled-deconvolution, however, it is not possible to measure the classification performance as the true neural events are not *known*. It is possible to measure the fraction of samples classified as *known*, which, in simulation, is highly correlated with classification performance. Based on the simulation results depicted in Figure 4b, a best estimate for real-world data would be to utilize a δ such that the fraction of *known* samples be equal to or greater than .33. In the simulations presented, which attempt to mimic real-world confounds as closely as possible, achieving this fraction requires a δ of .425 or smaller. However, this parameter depends on the temporal distribution of the neural events underlying BOLD signal. Temporal distributions applied in past simulations [13,16,18,19] differ substantially; identification of the true, underlying distribution of resting state neural events is not well-studied.

4.4 Human resting-state functional connectivity analysis

Results from the human resting-state data are depicted in Figures 6–9. Compared to the seedmaps generated from the preprocessed BOLD timeseries (top image in Figure 6 and Figure 8), the comparable seedmaps from the resampled-deconvolution timeseries detected overlapping patterns of functional connectivity specifically within the key nodes canonically implicated in the default mode: medial PFC, parahippocampus, and temporal lobes. Further, for the GM voxelwise approach, the seedmaps resulting from the resampled timecourses detected bilateral parahippocampus and temporal lobe clusters of functional connectivity at $\delta = -.10, 0.0$, and 0.20 , whereas the seedmap for standard BOLD timecourses detected unilateral temporal lobe and parahippocampal activation. For the ROI approach, the detected patterns of functional connectivity tended to demonstrate increased statistical significance among the resampled datasets compared to the standard BOLD seedmaps. For both approaches, as the δ approaches .30, the loss of data appears to begin to reduce the enhanced statistical power of the resampled-deconvolution seedmaps. Examining the relationship between δ and the fraction of neural events classified as *known*, depicted in Figures 7 and 9, we see that this fraction approaches .33 as δ increases to .25, as predicted by simulation.

5. Discussion

We have described two algorithmic extensions to the Bu13 deconvolution algorithm. The first, tuned-deconvolution, is a technical refinement of the feature space used by the Bu13 algorithm and is specific to it. The second extension, resampled-deconvolution, is a general extension of deconvolution that leverages variance estimation principles first developed in the applied statistics literature as bootstrapping. Both extensions provided statistically significant neural event classification performance improvements over the base algorithm, which is the current performance leader for deconvolution of realistically confounded BOLD data. Tuned-deconvolution's results provide evidence that the encoding representation used by the Bu13 algorithm is significant but not the performance limiting factor. Resampled-deconvolution, rather, produced uncommonly large performance gains (as much as 13.5%) over the existing algorithm in simulation.

A limitation of resampled-deconvolution is that it removes *unknown* features from the classification problem. This introduction of an *unknown* label into the classification process has consequences for the application of resampled-deconvolution in practice: it is not suitable for domains in which temporally contiguous features need be present. The removal of *unknown* features can be severe; at the highest performance levels in simulation (.723 AUC) nearly 90% of the observations are classified as *unknown*. Simulation results also suggest a risk of variable classification performance when more than 67% of the observations are classified as *unknown*. Moreover, real-world utility of the estimated encodings decreased significantly at a similar fraction of *unknown* classification. Together, these results suggest that 67% be a rule-of-thumb for the highest allowable fraction of *unknown* classification when applying resampled-deconvolution in practice. Further research is necessary to better estimate the threshold.

One such domain where the authors believe resampled-deconvolution has potential is estimation of effective connectivity. Identifying temporally precise neural events is a critical precursor to identifying causal relationships among brain regions. Recent research has dampened the prospects of causal studies conducted directly on BOLD signal due to large variations in the HRF function, which mediates the relationship between neural events and observed BOLD. The push to deconvolve BOLD prior to causal identification provides an opportunity for algorithms that improve deconvolution performance. The authors propose this as a hypothetically interesting environment for application of resampled-deconvolution as the temporal delays between brain regions are anatomically constrained and (potentially) well-defined. Specifically, neural activation pairs (causal and caused) for which both labels are *known* may be extracted from the dataset and used for inference. Indeed, under such circumstances, resampled-deconvolution provides a tradeoff between improved data quality and decreased data quantity; tradeoffs for which the machine learning community is well-acquainted. Appropriate modeling and algorithm choices can be tailored to the properties of the dataset. It should be noted, that while resampled-deconvolution can improve the performance of deconvolution, it does not solve the fundamental parameter estimation problems that plague many deconvolution algorithms, e.g., intra- and inter-subject variability of the HRF function. However, it also does not preclude the use of an algorithm that blindly estimates these parameters from data. Blind deconvolution is an important area

of deconvolution research to which resampled-deconvolution may provide performance improvements.

Results from the human data conceptually replicate the simulation results. We observed that increasing the δ parameter generally yielded comparable or greater functional connectivity with the seed region for the resampled-deconvolution seedmaps relative to the original preprocessed BOLD timeseries seedmaps. This increase appeared to peak at a δ parameter on the range .15–.20, and began to decrease at a δ parameter of .30. The increased statistical power to detect functional connectivity demonstrates that the resampled-deconvolution methodology reduced noise that was confounding, and consequently attenuating, the group-level correlation estimates between the timecourses. Of note, there was specificity with respect to increased functional connectivity. We did not observe increased functional connectivity with the PCC seed regions across the brain; instead, we specifically observed increased functional connectivity between regions canonically correlated with the seed regions. These areas where the PCC seed region should be most highly correlated are the same regions where we saw the increased functional connectivity among the resampled-deconvolution seedmaps, providing strong support for the validity of the methodology and its applicability to resting-state data. It should be noted that this method is not appropriate for task-based data, due its potential removal of temporally contiguous signals; however, these simulation and real data results do suggests its viability for resting-state data.

Acknowledgments

This work was supported in part by National Institutes of Health grants R21MH097784-01 and R01DA036360-01 as well as by the National Science Foundation grants CRI CNS-0855248 and MRI ACI-1429160.

References

1. Ogawa S, Lee T, Kay A, Tank D. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proc Natl Acad Sci USA*. 1990; 87(24):9868–72. [PubMed: 2124706]
2. Logothetis N, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. *Nature*. 2001; 412(6843):150–7. [PubMed: 11449264]
3. Logothetis N. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci*. 2003; 23(10):3963–71. [PubMed: 12764080]
4. Logothetis N. What we can do and what we cannot do with fMRI. *Nature*. 2008; 453(7197):869–78. [PubMed: 18548064]
5. 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]
6. Friston K. Functional and effective connectivity in neuroimaging: A Synthesis. *Human Brain Mapping*. 1994; 2:56–78.
7. Purdon P, Weisskoff R. Effect of temporal autocorrelation due to physiological noise and stimulus paradigm on voxel-level false-positive rates in fMRI. *Human Brain Mapping*. 1998; 6:239–49. [PubMed: 9704263]
8. Lindquist M. The statistical analysis of fMRI data. *Statistical Science*. 2008; 23(4):434–64.
9. Saad Z, Ropella K, Cox R, DeYoe E. Analysis and use fMRI response delays. *Human Brain Mapping*. 2001; 13(2):74–93. [PubMed: 11346887]
10. Huettel S, McCarthy G. Regional differences in the refractory period of the hemodynamic response: an event-related fMRI study. *NeuroImage*. 2001; 14(5):967–76. [PubMed: 11697929]

11. Miezin F, Maccotta L, Ollinger J, Petersen S, Buckner R. Characterizing the Hemodynamic Response: Effect of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing. *NeuroImage*. 2000; 11:735–59. [PubMed: 10860799]
12. Lee A, Glover G, Meyer C. Discrimination of Large Venous Vessels in Time-Course Spiral Blood-Oxygen-Level-Dependent Magnetic Resonance Functional Neuroimaging. *Magn Reson Med*. 1995; 33:745–54. [PubMed: 7651109]
13. Smith S, Miller K, Salimi-Khorshidi G, Webster M, Beckmann C, Nichols E, Ramsey J, Woolrich M. Network modelling methods for fMRI. *Neuroimage*. 2011; 54(2):875–91. [PubMed: 20817103]
14. David O, Guillemain I, Saillet S, Reyt S, Deransart C, Segebarth C, Depaulis A. Identifying neural drivers with functional MRI: An electrophysiological validation. *PLOS Biology*. 2008; 6(12): 2683–97. [PubMed: 19108604]
15. Stephan K, Roebroeck A. A short history of causal modeling of fMRI data. *NeuroImage*. 2012; 62:856–63. [PubMed: 22248576]
16. 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]
17. Gaudes C, Petridou N, Dyrden I, Bai L, Francis S, Gowland P. Detection and characterization of single-trial fmri bold responses: Paradigm free mapping. *Human Brain Mapping*. 2011; 32(9): 1400–18. [PubMed: 20963818]
18. Bush K, Cislér J. Decoding neural events from fMRI BOLD signal: A comparison of existing approaches and development of a new algorithm. *Magn Reson Imaging*. 2013; 29(3):353–64.
19. 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]
20. Duch W, Jankowski N. Taxonomy of neural transfer functions. *Neural Networks*. 2000; 3:477–82.
21. Guyon I, Elisseeff A. An introduction to variable and feature selection. *Journal of Machine Learning Research*. 2003; 3:1157–82.
22. Efron B. Bootstrap methods: Another look at the jackknife. *Annals of Statistics*. 1979; 7:1–2.
23. Friston, K.; Ashburner, J.; Kiebel Nichols, T.; Penny, W. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Elsevier; London: Academic Press; 2006.
24. McKeown M, Jung T, Makeig S, Brown G, Kindermann S, Lee T, Sejnowski T. Spatially independent activity patterns in functional MRI data during the stroop color-naming task. *Proc Natl Acad Sci*. 1998; 95(3):803–810. [PubMed: 9448244]
25. Craddock R, James G, Holtzheimer P, Xiaoping P, Mayberg H. A Whole brain fMRI Atlas generated via spatially constrained clustering. *Human Brain Mapping*. 2012; 33(8):1914–28. [PubMed: 21769991]
26. Bullmore E, Sporns O. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews*. 2009; 10:186–98.
27. Dayan, P.; Abbott, LF. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. Cambridge, MA: MIT Press; 2001. p. 460
28. Wu C. Jackknife, Bootstrap and other resampling methods in regression analysis. *The Annals of Statistics*. 1986; 14(4):1261–95.
29. Mammen E. Bootstrap and wild bootstrap for high dimensional linear models. *Annals of Statistics*. 1993; 21:255–85.
30. Davison, A.; Hinkley, D. *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge, UK: Cambridge Univ Press; 1997. *Bootstrap methods and their application*; p. 594
31. Lahiri S. Theoretical comparisons of block bootstrap methods. *Annals of Statistics*. 1999; 27:386–404.
32. Cameron A, Gelbach J, Miller D. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*. 2008; 90:414–27.
33. Bush K, Cislér J. Deconvolution Filtering: Temporal Smoothing Revisited. *Mag Reson Imaging*. 2014; 32(6):721–35.
34. Smith S, Jenkinson M, Woolrich M, Beckmann C, Behrens T, Johansen-Berg H, Bannister P, De Luca M, Drobnjak I, Flitney D, Niazy R, Saunders J, Vickers J, Zhang Y, De Stefano N, Brady J,

- Matthews P. Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage. 2004; 23(S1):S208–19. [PubMed: 15501092]
35. Stallman, R., et al. Using the GNU Compiler Collection. GNU Press; Boston, MA: 2013.
36. Sanderson, C. Technical Report. National ICT Australia (NICTA); St. Lucia, Queensland, Australia: 2010. Armadillo: An Open Source C++ Linear Algebra Library for Fast Prototyping and Computationally Intensive Experiments.
37. Fox M, Snyder A, Vincent J, Corbetta M, Van Essen D, Raichle M. The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proc Natl Acad Sci USA. 2005; 102(27):9673–78. [PubMed: 15976020]
38. Raichle M, MacLeod A, Snyder A, Powers W, Gusnard D, Shulman G. A default mode of brain function. Proc Natl Acad Sci USA. 2001; 98(2):676–82. [PubMed: 11209064]

Appendix

fMRI BOLD Signal Simulation Generation and Observation Parameters

The default fMRI BOLD signal simulation [33], used in this work, employed parameters described in Table 1.

Resampled-Deconvolution

Given an observed BOLD signal $\mathbf{y}^{scan}$, of length P , a deconvolution mapping h (parameterized by set, Θ , where Θ_H describes the number of time-points that the deconvolution mapping attempts to estimate prior to the first observed datapoint, $\mathbf{y}^{scan}(1)$), a deconvolution filtering mapping g (parameterized by set, Φ), a parameter N describing the number of surrogate samples to form, and an operator ψ which uniformly randomly orders the elements of the argument vector.

algorithm: Resampled-Deconvolution

parameters: $h, g, \Theta, \Phi, N, \psi$

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

result: $\tilde{\mathbf{e}}_n^{sur}, \tilde{\mathbf{e}}_n^{sur}(p) \in [0,1], p \in 1 - \Theta_H, \dots, 1, \dots, P$

begin:

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

$\mathbf{y}^{fltr} = g_{\Phi}(\tilde{\mathbf{e}}), \mathbf{y}^{fltr}(p) \in \mathbb{R}$

$\mathbf{r} = \mathbf{y}^{scan} - \mathbf{y}^{fltr}$

for $n \in 1, \dots, N$

$\mathbf{r}_n = \psi(\mathbf{r})$

$\mathbf{y}_n^{sur} = \mathbf{y}^{fltr} + \mathbf{r}_n$

$\tilde{\mathbf{e}}_n^{sur} = h_{\Theta}(\mathbf{y}_n^{sur})$

end-for

end

Classification Performance Impact of Autoregressive Noise

We explore the impact that autoregressive noise could have on the classification of *known* And *unknown* labels by varying the AR_{phys} coefficient of the BOLD signal simulation on the range [0,1]. This parameter controls the strength of temporal structure in the physiological (autoregressive) noise component of the observation model of the simulation. Low values of this parameter would suggest that autoregressive filtering (via modeling and removal of CSF artifact) is successful. A value close to one would suggest that this filtering process is unsuccessful. Holding all other simulation parameters constant we conducted the following experiment. We performed 30 random BOLD signal simulations for each value of AR_{phys} on the range [0,1] sampled at intervals of 0.05 and performed resampled-deconvolution on these data. We then fit¹ the classification results for each level of background neural event confidence, δ , with respect to the relevant AR_{phys} coefficient via linear models (all of which had significantly negative slopes, $p < .05$). We then fit this relationship with a linear model. The resulting model has a statistically significant slope $m = -0.0126$ ($p < .05$) and explains 72.14% of the variance of the performance. We present this model in Figure 10 (raw performance plotted as a blue line; model plotted as a red line). Converting this relationship into a classification outcome, the model estimates that autoregressive noise, in the worst case, reduce the performance of *known/unknown* classification by .013 (which equates to a 2.0% reduction in performance in simulation). The increased variance at large δ values is consistent with results of AUC (see Fig 4), likely stemming from decreasing fraction of *known* events.

¹We use iteratively reweighted least squares to construct all linear models via Matlab's `robustfit` function (with default parameters).

Figure 1.

Theoretical BOLD signal: (a) canonical hemodynamic response to a single neural event; (b) simulated BOLD signal formed via the linearly additive composition of HRFs in response to a sequence of neural events; (c) simulated observed fMRI BOLD signal confounded by autocorrelated physiological noise, white thermal noise, normalization, and low-fidelity observations. Note, neural events are plotted as vertical bars for temporal reference; they are not to scale.

Figure 2.

Graphical depiction of the logistic transfer function for various shape parameters, β . As the β parameter increases the logistic function converges to the step function

Figure 3.

Performance of deconvolved encodings in classifying the true, underlying simulated neural events as a function of the deconvolution algorithm's shape parameter plotted for shape values $\beta \in \{1, 4, 10, 20, 40, 60, 80\}$. Mean classification performance (solid bold line) is plotted within the 95% confidence interval (solid thin lines). The optimized parameter, $\beta=60$ (selection described in Section 3.1) improves classification performance by 2.18% ($p < 0.05$) over the base Bu13 algorithm ($\beta=1$).

Figure 4.

Resampled-deconvolution performance summary: **(a)** classification performance (measured as AUC of ROC) as a function of δ presented as mean AUC (and 95% confidence interval); **(b)** fraction of *known* labels as a function of δ presented as mean fraction (solid bold line) and the 95% confidence interval (solid thin lines).

Figure 5.

Example of the resampled-deconvolution feature space: **(a)** resampled-deconvolution mean encoding (and 95% confidence intervals) plotted (blue) overlaid with the base deconvolution algorithm's encoding (red) and the true neural events (cyan) of a trial of the default simulation; **(b)** Estimated mean BOLD signal $\pm 2\sigma$ generated by reconvolving the resampled-deconvolution encodings with the canonical HRF (blue) plotted overlaid with the observed BOLD signal (red). The trial is the same as in panel (a), but the x-axis has been compressed to aid in the visualization of neural event details in panel (a). The y-axis has been z-scored.

Figure 6.

Comparison of functional connectivity seedmaps using the posterior cingulate cortex as the seed region within the 1000 ROI functional atlas. The comparison is calculated between the original BOLD signal's functional connectivity and various resampled-deconvolution functional connectivities generated as a function of the uncertainty of the background neural activity, δ . High t-values indicate greater connectivity for the resampled deconvolution approach versus the standard BOLD correlation.

Figure 7.

Fraction of *known* labels as a function of δ estimated for the fMRI BOLD signals of 17 human participants parcellated using the 1000 ROI functional atlas. The data are presented as mean fraction (solid bold line) and the 95% confidence interval (solid thin lines). At $\delta = -0.1$, all neural events are classified as *known* (fraction = 1.0). Results for $\delta = 0.5$ are not included in the plot as no data met this criteria, thus zero mean and zero variance. The vertical dashed line represents the boundary of δ as a scaling of the *unknown* confidence interval (Eqn. 7).

Figure 8.

Comparison of functional connectivity seedmaps using the posterior cingulate cortex as the seed region between the resampled-deconvolution approaches and the standard BOLD correlation, as a function of δ for GM voxelwise analysis. High t-values indicate greater connectivity for the resampled deconvolution approach versus the standard BOLD correlation.

Figure 9.

Fraction of *known* labels as a function of the δ estimated for the fMRI BOLD signals of 17 human participants. The data are presented as mean fraction (solid bold line) and the 95% confidence interval (solid thin lines). At $\delta = -0.1$, all neural events are classified as *known* (fraction = 1.0). Results for a $\delta = 0.5$ are not included in the plot as no data met this criteria, thus zero mean and zero variance. The vertical dashed line represents the boundary of δ as a scaling of the *unknown* confidence interval (Eqn. 7).

Figure 10.

Slopes of the relationship between performance of *known/unknown* classification (measured as AUC of ROC) and the temporal correlation, AR_{phys} , of the BOLD signal's noise, plotted as a function of background neural event uncertainty, δ . Raw slopes are plotted in blue. Linear model of the slopes plotted as a red line. Note, slopes are not plotted for $\delta=.50$ as this extreme constraint causes missing samples (i.e., 0.0 fraction of neural events).

Table 1

Parameter	Value	Description
TS	200 s	Simulation length.
N	1	The number of simulated regions-of-interest (ROI).
C	[0]	Connectivity model, a matrix (size $N \times N$) of the conditional probabilities. Each element, $C(i,j)$ determines the probability of a neural event in region j caused by a neural event in region i .
L	[0]	Lag model, a matrix (size $N \times N$) of communication times between simulated ROIs. Each element, $L(i,j)$ determines the temporal delay of a neural event in region j occurring due to a neural event in region i .
ρ	[0.05]	External activity model, a vector containing the latent probability for each ROI of the model to fire based on unmodeled, external influences.
FG	20 Hz	Simulated generation frequency. $1/FG$ represents the smallest unit of time modeled and is assumed to be the span of time in which a single neural event occurs.
FO	1 Hz	Simulated observation frequency; equivalent to $1/TR$.
HRF_d	6 s	Time-to-peak shape parameter of the canonical hemodynamic response function as required by the SPM software package.
SN_{phys}	6	Signal-to-noise ratio of the auto correlated noise process used to simulated cardiac and respiratory signals.
AR_{phys}	0.75	Correlation coefficient of the first-order autoregressive, AR(1), model.
SN_{scan}	9	Signal-to-noise ratio of the uncorrelated thermal noise process.
$transient$	3	The number of HRF function time-lengths to be prepended and then pruned from the simulated BOLD signal, i.e., the BOLD signal is simulated for $transient \cdot HRF_d + TS$ seconds; then the transient time period is cut from the observed BOLD. This parameter ensures that the first BOLD observation (and all subsequent observations) contains no simulation bias due to initial conditions.