Journal of Neuroscience Methods 219 (2013) 205–219

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Computational Neuroscience

Reproducible paired sources from concurrent EEG-fMRI data using BICAR

Kevin S. Brown[a,b,c,d,*], Ryan Kasper[e,b], Barry Giesbrecht[e,b], Jean M. Carlson[a,b], Scott T. Grafton[e,b]

[a] Department of Physics, University of California, Santa Barbara, CA 93106, USA
[b] Institute for Collaborative Biotechnologies, University of California, Santa Barbara, CA 93106, USA
[c] Chemical and Biomolecular Engineering, University of Connecticut, Storrs, CT 06269, USA
[d] Department of Marine Sciences, University of Connecticut, Groton, CT 06340, USA
[e] Department of Psychological and Brain Sciences, University of California, Santa Barbara, CA 93106, USA

HIGHLIGHTS

- We introduce BICAR, a new algorithm for subject-level EEG-fMRI data fusion.
- BICAR ranks each joint source by a task-independent measure of reproducibility.
- We derive an analytical reproducibility cutoff below which components are discarded.
- We apply BICAR to human subjects performing a visual search task.
- Among the most reproducible sources are visual, motor, and attentional components.

ARTICLE INFO

Article history:
Received 11 March 2013
Received in revised form 22 July 2013
Accepted 22 July 2013

Keywords:
EEG
fMRI
Concurrent EEG-fMRI
Multimodal data fusion
Independent component analysis

ABSTRACT

We introduce BICAR, an algorithm for obtaining robust, reproducible pairs of temporal and spatial components at the individual subject level from concurrent electroencephalographic and functional magnetic resonance imaging data. BICAR assigns a task-independent measure of component quality, reproducibility, to each paired source. Under BICAR a reproducibility cutoff is derived that can be used to objectively discard spuriously paired EEG-fMRI components. BICAR is run on minimally processed data: fMRI images undergo the standard preprocessing steps (alignment, motion correction, etc.) and EEG data, after scanner artifact removal, are simply bandpass filtered. This minimal processing allows the secondary scoring of the same set of BICAR components for a variety of different endpoint analyses; in this manuscript we propose a general method for scoring components for task event synchronization (evoked response analysis), but scoring using many other criteria, for example frequency content, are possible. BICAR is applied to five subjects performing a visual search task, and among the most reproducible components we find biologically relevant paired sources involved in visual processing, motor planning, execution, and attention.

1. Introduction

Concurrent surface electroencephalography and whole-brain functional magnetic resonance imaging (EEG-fMRI) hold tremendous promise for obtaining non-invasive high spatiotemporal resolution measurements of human brain dynamics (Debener et al., 2006). Multiple studies support the idea that EEG sources and fMRI activations in healthy subjects can colocalize (Christmann et al., 2007; Whittingstall et al., 2007; Roberts et al., 2008; Esposito et al., 2009; Dubois et al., 2012; Logothetis et al., 2001). In addition, in certain patient populations, especially epileptics being considered for surgical therapy, EEG-fMRI is emerging as a promising alternative to invasive corticography (Rosenkranz and Lemieux, 2010). The most compelling reason for fusing EEG and fMRI is that they are complementary. Trying to infer the locations of cortical sources from EEG data alone is a highly underdetermined inverse problem which will not yield a unique solution without employing *ad hoc* constraints on solution norm or smoothness (Sekihara et al., 2001; Pascual-Marqui et al., 1994; Hämäläinen and Ilmoniemi, 1994). Conversely, even if major advances in pulse-sequence technology could vastly shorten MRI acquisition times (Lustig et al., 2007;

* Corresponding author at: 1080 Shennecossett Road, Groton, CT 06340-6040, USA. Tel.: +1 860 405 9136.
E-mail address: kevin.s.brown@uconn.edu (K.S. Brown).

0165-0270/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jneumeth.2013.07.012

Candès and Wakin, 2008; Jung et al., 2009; Jeromin et al., 2012), fMRI is fundamentally a slow metabolic/hemodynamic measurement that cannot yield the kind of time resolution obtainable from electrodynamic measurements.

Despite a general belief that EEG-fMRI is better than either measurement on its own, general methods for combining these data of vastly different spatial and temporal resolution are lacking (Chaudhary et al., 2011). Some existing methods make heavy trade-offs for or against spatial or temporal resolution, for example in the case of fMRI event (Grouiller et al., 2011; Mulert et al., 2002; Wu et al., 2010; Yuah et al., 2012) and EEG source (Babiloni et al., 2002; Liu et al., 2006; Strobel et al., 2008) weighting. Other methods build specialized solutions geared towards obtaining a solution useful for analysis in frequency space, as is the case for wavelets (Schultze-Kraft et al., 2011) or partial least squares (Martinez-Montes et al., 2004). A variety of other methods based on response functions (Sato et al., 2010), mutual information (Ostwald et al., 2010, 2012), or variational Bayesian approximations (Daunizeau et al., 2007) have been proposed, but typically time locking to a behavioral event of interest is essential for obtaining a solution.

Joint estimations aim to take a symmetric approach to integrating EEG-fMRI data sets; the most straightforward way to do this is with Independent Component Analysis (ICA) (Moosmann et al., 2008). EEG is modeled using temporal ICA (Makeig et al., 1996) and fMRI via spatial ICA (McKeown et al., 1998). These approaches have the advantage of requiring no source localization model, though one could be used if desired (Brookings et al., 2009; Brown et al., 2010). Joint estimations typically have the disadvantage of relying heavily on the task structure and/or rather arbitrary processing/reshaping of the data matrices to obtain conformable dimensions across subjects or modalities.

It is extremely challenging to establish reliable and unbiased methods to match ICA components systematically. This is a difficult problem even when considering a single imaging modality, e.g. selection of spatial fMRI sources alone. One solution is to match components at a population level by finding components that are consistent across subjects (Eichele et al., 2008). However, in order to be diagnostically important, a solution must be obtainable at the individual subject level. Another promising approach to the problem of unbiased component selection for single data modalities is given by the RAICAR algorithm (Yang et al., 2008), which uses a repeated estimations technique to assess the robustness of components. Each component is given a task-independent measure of component quality: the reproducibility. This value is calculated by performing stochastic ICA many times on the same dataset and then matching components across realizations. Components which repeatedly occur in many realizations are more reproducible and hence more trustworthy. However, RAICAR is limited to use on single modalities and therefore not suitable for joint EEG-fMRI data mining.

Motivated by the strategy of RAICAR, we recently developed BICAR: a general-purpose data-driven method for obtaining robust, reproducible *pairs* of temporal and spatial components from complementary spatiotemporal data matrices (Brown et al., 2012). In that paper the algorithm is applied to many different types of data, demonstrating the general applicability of BICAR for spatiotemporal data fusion. In this manuscript, the extension and application of BICAR to event-related analysis of EEG-fMRI data as a specific application is introduced. To apply BICAR to EEG-fMRI data, a number of significant modifications and improvements were necessary. BICAR follows RAICAR (Yang et al., 2008) in assigning an experiment-independent measure of component quality (reproducibility), but goes beyond RAICAR in several ways, most notably in the application to joint EEG-fMRI data. BICAR has five key advantages over existing methods applied to this problem:

Objectivity: In BICAR one can analytically calculate a data- and algorithm-dependent reproducibility cutoff below which a subject's components should be deemed insignificant; this cutoff is derived in this manuscript. This entirely removes the subjectivity that accompanies many ICA analyses, in which components are selected by eye using prior anatomical or task information.
Flexibility: Only scanner acquisition times and data sampling rates are used in BICAR; no information about the task or paradigm is required. BICAR is run on the preprocessed data, with no epoching, calculation of t-maps, etc. This allows the same set of BICAR components to be used for multiple endpoint analyses: evoked responses, frequency content, spatial priors, etc. This is a crucial point; while we consider event-related analysis in this manuscript, that choice is made *after* BICAR is run on the data. BICAR itself remains completely ignorant of the task events in the experiment.
Extensibility: BICAR is *not* a source localization method, but results from BICAR could be directly used in source localization studies: highly reproducible components could be used as basis functions in a joint inversion approach (Brookings et al., 2009; Brown et al., 2010).
Specificity: BICAR produces joint components at the individual subject level. This allows the discovery of both subject-common and subject-unique dynamics, which could be important for understanding behavioral variability (Miller et al., 2002, 2011). Methods like group ICA (Calhoun et al., 2001) which begin at the population level only reveal activity common to the entire population.
Symmetry: Unlike event or source weighting methods, BICAR treats the two data sources on identical footing by taking the best resolved dimension of each: time from EEG and space from fMRI. In addition, it is possible to adjust the relative weighting between the two data sources to accentuate temporal or spatial information.

In what follows, the BICAR algorithm is discussed in the context of application to EEG-fMRI data (see Section 2). Full details on the algorithm and its performance on a variety of test data are available elsewhere (Brown et al., 2012). Here we introduce for the first time a general method for scoring BICAR components for significant association with task events; this allows selection of BICAR components which are both reproducible and significantly associated with a set of task events of interest.

Not only can BICAR assign a reproducibility value to each joint source, but it is also possible to calculate a reproducibility cutoff below which BICAR components are deemed to be spuriously paired and should be ignored. This derivation is discussed in detail in this manuscript, and is an important advance. Simply ranking the components by reproducibility gives one an order for consideration (most reproducible first), but offers no guidance as to which joint components to neglect for further analysis. The reproducibility cutoff provides this guidance.

In this manuscript, BICAR is applied to a set of healthy participants with normal attention performing a visual search task. BICAR is able to find biologically relevant paired sources involved in visual processing, motor planning, execution, and attention, which are highly reproducible and present in multiple subjects. In these examples, we focus on finding components that are relevant for an event related experimental design for individual subjects and not to perform a large sample, group analysis. It should be stressed that this is only one application. Other scoring criteria, such as frequency analysis, can be applied to the same set of BICAR components after the algorithm is used. BICAR only needs to be repeated if the EEG or fMRI data itself changes, for example because of different choices in image preprocessing.

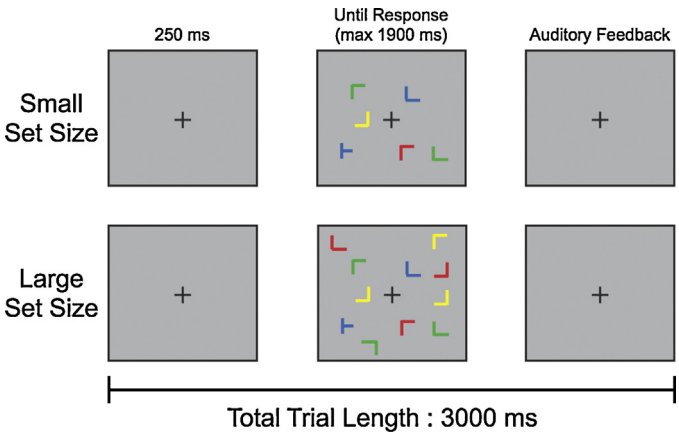

Fig. 1. Single target-search task used in this study. After presentation of a fixation cross, a field of "L"s is presented along with one rotated "T", in one of two possible orientations. Subjects search the visual field until they find the "T", reporting on its orientation via a button press. Search fields were either from a small item set (six items total) or a large item set (twelve items total).

2. Methods

2.1. Experimental design

2.1.1. Participants

Five study participants (four female, one male) performed a challenging visual search task during a two hour simultaneous EEG-fMRI experiment at the University of California, Santa Barbara Brain Imaging Center. All of them self-reported as right handed, and ranged in age from 20 to 30 years. All volunteers had normal or corrected-to-normal visual acuity. Prior to participation, volunteers provided written informed consent that had been approved by the Ethical Committee of the University of California, Santa Barbara.

2.1.2. Task and procedure

The orientation task we used is one that is commonly used in visual search tasks (Chun and Jiang, 1998; Johnson et al., 2007; Giesbrecht et al., 2013). The search field consisted of rotated colored 'L' distractors and a single rotated 'T' target (see Fig. 1). The item location for each array was randomly sampled from a grid of sixteen possible locations with the constraint that across all trials the target was equally likely to be in any of these locations. Each trial could be either a small set size (six items) or a large set size (twelve items), as well as either 'repeated' (identical array of search items to one seen previously) or 'novel' (random layout of search items). For repeated displays, the distractor locations and identities were held constant, but the target was rotated 90° to the left or 90° to the right (equal numbers of trials). In order to facilitate the contextual cueing effect participants were instructed to adopt a passive search strategy (Lleras and Von Mühlenen, 2004). For the analysis in this manuscript, we aggregate across all target presentations and call them all "stimuli".

The search task procedure was as follows. Each trial started with a central fixation to be maintained at all times. After 250 ms of fixation, a search array was displayed until either the subject responded or 1900 ms elapsed. The subject had to indicate whether the target 'T' was rotated to the left or right by pressing one of two buttons as quickly as possible. Responses were made with the right hand using a button box. After response, auditory feedback was given to indicate whether the answer was correct (high pitched beep) or incorrect (buzzer). In this article we discard all incorrect responses when studying the "response" category.

Each trial lasted a total of three seconds. A total of eight blocks of trials were used for each subject, with each block consisting of eight small and eight large displays repeated twice for a total of 256 experimental trials. In addition, approximately 85 blank trials (one third of 256) were included in each block to introduce temporal jitter into the paradigm. During the blank trials, no target/distractor field was presented; the fixation cross simply remained on the screen. Importantly, the orientation of the target (left or right with respect to an upright 'T') was completely counterbalanced and randomized so that in each block, every display (regardless of whether it was 'novel' or 'repeated') was presented twice, once with the target rotated left and once with the target rotated right. The 'repeated' and 'novel' aspects refer only to configuration of the target and distractor encountered previously. Across blocks, new displays consisting of unique target/distractor spatial configurations were created for the 'novel' condition and the original configurations from the 'repeated' condition were again used twice per block, once for each orientation. This counterbalancing and randomization procedure removes any motor biases associated with a particular display and ensures that the task (orientation discrimination) is orthogonal to the spatial configuration manipulation ('novel' vs. 'repeated').

In order to avoid removing task-relevant signal from concurrently recorded EEG during MR artifact correction, the search array presentations were not explicitly tied to a TR onset (Allen et al., 2000). Instead, the initial search array was presented following the first TR after a random delay. This delay and the three-second trial length avoided regular synchronization with the two second TR.

2.1.3. Data collection and preprocessing

All data was collected at the UCSB Brain Imaging Center.

Electroencephalographic data was acquired simultaneously with fMRI data using an MR-compatible 64 channel EEG system (www.brainproducts.com). Eight five-minute sessions of EEG-fMRI data were collected for each subject. The data was acquired at 1000 Hz and re-referenced offline to an average of electrodes TP9 and TP10 (near mastoids). MR gradient switching artifacts were removed via BrainVision Analyzer software version 2.0 (www.brainproducts.com); the correction process creates an artifact template for each TR based on scanner triggers, and this template is subtracted from the raw EEG data (Allen et al., 2000). Following this correction, the data was downsampled to 250 Hz.

The ballistocardiogram artifact was removed via Niazy's OBS method (Niazy et al., 2005), using the FMRIB plugin for EEGLab (sccn.ucsb.edu/eeglab). After BCG correction the ECG electrode was discarded. Each electrode in each imaging session was band pass filtered between 0.1 and 30 Hz, downsampled to 62.5 Hz, and linearly detrended before concatenation. This resulted in a two-dimensional array of size number of electrodes (n_E) by number of EEG samples (t_E).

A 3T TIM TRIO Siemens Magnetom scanner with a 12-channel phased-array head coil was used for MRI data collection. The functional data were acquired using a $T2^*$-weighted gradient-echo sequence with a repetition time (TR) of 2 s, echo time (TE) of 30 ms, and flip angle (FA) of 90°, resulting in 56 contiguous slices at 3 mm × 3 mm × 3.5 mm voxel resolution. For anatomical data, a T1-weighted MPRAGE sequence with 1 mm isometric voxel resolution was used.

fMRI data were preprocessed using SPM 5.0 (www.fil.ion.ucl.ac.uk/spm) (Friston et al., 1995). fMRI image volumes were slice time corrected, motion corrected, unwarped, spatially normalized to the Montreal 152 Average T1 atlas, and resliced to 4 mm × 4 mm × 4 mm voxel sizes. No smoothing was performed on the fMRI images. Each voxel in each imaging session was linearly detrended and the sessions were concatenated; these steps occurred after all SPM preprocessing steps. The high resolution T1 anatomical MRI was segmented into gray and white matter, warped to the 152 T1 atlas, and resliced to 4 mm × 4 mm × 4 mm

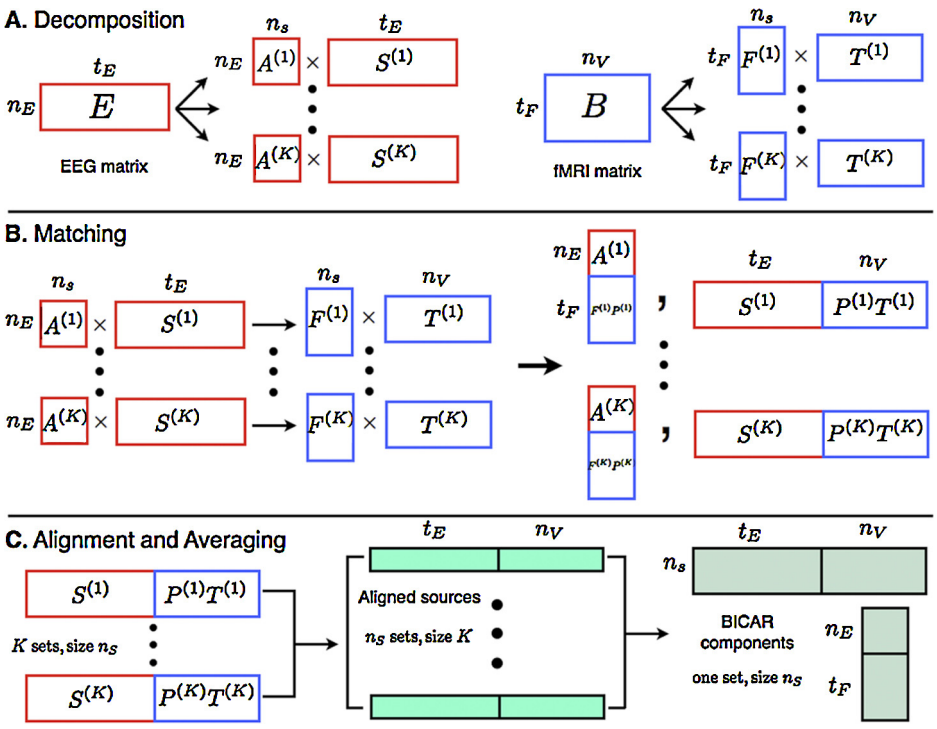

Fig. 2. Schematic for the BICAR algorithm. BICAR proceeds first by K-fold FastICA decomposition of the EEG and fMRI data matrices E and B; n_S is the number of extracted sources. Then sources are paired using a transfer function between the EEG and fMRI dynamics. The cross-realization correlation matrices between the paired sources are searched and aligned via absolute correlation; this results in n_S groups of paired sources, each of which consists of one source from each of the K ICA realizations. Components are then averaged with weighting to obtain one set of n_S paired BICAR sources and associated mixing matrices. See Section 2 for details on these steps.

voxel sizes. The combined gray and white matter images were used as masks for each subject to remove non-brain voxels before session concatenation and reshaping into a two-dimensional array of the number of scans (t_F) by the number of voxels (n_V). This process was replicated for all five subjects.

For two of the five subjects, one of the eight sessions had to be discarded as there were errors in the time stamps needed to make a correspondence between the EEG record and fMRI volume acquisition times within the five-minute session. For these two subjects, the other seven sessions were used. For the remaining three subjects all eight recording sessions were usable.

2.2. BICAR

2.2.1. Algorithm

The steps of the BICAR algorithm (Brown et al., 2012) are discussed here in order to highlight particulars of applying it to EEG-fMRI data. Fig. 2 gives an algorithm schematic. In Fig. 2A K-fold FastICA (Hyvärinen and Oja, 1997) decomposition is performed on both the EEG data matrix (size $n_E \times t_E$) and the fMRI data matrix (size $t_F \times n_V$). By K-fold FastICA we mean the following. Each data matrix (E and B in Fig. 2) is unmixed via FastICA with a different random initial condition; we refer to a single pair of FastICA compositions as a *realization*. The initial unmixing matrix for each decomposition is a matrix of Gaussian distributed random numbers with mean zero and unit variance of the appropriate size ($n_E \times n_S$ for E, $t_F \times n_S$ for B). This matrix is orthogonalized in at the start of the FastICA algorithm. Temporal ICA is used on the EEG matrix and spatial ICA on the fMRI data matrix; the fMRI data matrix has been transposed in Fig. 2A so that time (scans) is the row dimension, as is appropriate for spatial ICA. All initial conditions – both for different values of K and for spatial and temporal ICA within a single realization – are independent. $K = 30$ was used for this study. Both applications of ICA in each realization extract n_S sources; for all analyses in this paper, $n_S = 63$, representing a full-rank decomposition of the EEG data matrix E. Hence in what follows $n_S = n_E = 63$ in Fig. 2, though this is not generally required (see below for additional information on this point). Each ICA decomposition for each dataset produces both n_S sources and n_S mixing elements. For the EEG data, the sources are time series and the mixing elements are typically called scalp maps. For the fMRI data, the sources are image volumes and the mixing elements are time series.

After decomposition, EEG sources are matched to the fMRI sources (Fig. 2B). This matching step is performed by pairing transformed, decimated EEG sources (time series) with fMRI mixing elements (also time series) via absolute Pearson correlation. Sources are paired one-to-one. One-to-many matching is also possible, but our previous work has considered only one-to-one source matching (Brown et al., 2012). In order to carry out the matching step, a transfer function between the EEG and fMRI dynamics needs to be specified. We have shown elsewhere via simulation that BICAR is relatively robust to transfer function misspecification (Brown et al., 2012); in addition, it is possible to optimize the transfer function from within BICAR itself, a subject for a future study. For the transfer function a simple parameterized function of the type often assumed for the hemodynamic response function (HRF) in fMRI research was employed. This is a single non-delayed gamma of the form

$$h(t, \{\alpha, \tau\}) = \left(\frac{t}{\tau}\right)^{\alpha} \frac{e^{-t/\tau}}{\tau\Gamma(\alpha+1)}, \quad (1)$$

which basically acts as a delayed low pass filter. The parameter values used were $\alpha = 8.6$, $\tau = 0.55$, resulting in a peak location of 4.73 s. After transformation by the HRF and before the correlation calculations, the transformed EEG sources were decimated to establish temporal correspondence with the fMRI scans. Correlation matching proceeds as follows and occurs independently in each realization. The source pair in each realization with the largest absolute transformed EEG source/fMRI loading correlation

are associated and removed from the pool of potential matches. This process is then repeated, in this case 62 additional times, to yield $n_S = 63$ paired sources in each of the $K = 30$ ICA realizations.

The paired sources are then aligned across realizations (Fig. 2C). This process is rather complicated and the details for both RAICAR (Yang et al., 2008) and BICAR (Brown et al., 2012) can be found elsewhere. The goal can be understood as follows. The ordering of the n_S paired sources in each of the K realizations is arbitrary. Hence, even if an *identical* set of paired sources was found in each realization, we still must construct a consistent ordering in each realization so that similar sources can be grouped. Therefore, we use a realization-realization cross-correlation technique in order to compute an optimal ordering. We then construct each of the n_S BICAR sources as a weighted average of K sources, one instance from each of the ordered K realizations, with the weighting factors representing the averaged pairwise reproducibility of sources across the K estimates. For example, in the case of identical sources in each realization (rarely achieved in practice), all the weights would be unity.

After this process of weighted averaging one BICAR decomposition has the same shape as two regular ICA decompositions (one temporal, one spatial) concatenated; the BICAR source matrix will be of dimensions $n_S \times t_E + n_V$ (number of extracted sources × number of EEG samples + number of fMRI voxels) and the BICAR mixing matrix of dimensions $n_E + t_F \times n_S$ (number of EEG electrodes + number of fMRI scans × number of extracted sources). In practice, all the matrices are kept separate for computational efficiency. However, it is useful to think this way because alignment is driven by both temporal (from EEG sources) and spatial (from fMRI sources) information; highly similar groups of sources are similar in both time and space. Once BICAR has been run, the temporal and spatial parts of the BICAR sources are separated and related to anatomical and other experimental information.

Before source averaging and reproducibility calculation, one must deal with a sign problem. In both the EEG and fMRI cases, ICA produces a set of sources and mixing elements that reconstruct the corresponding data matrix. However, one can easily change the sign of any number of sources along with the signs of the corresponding columns of the mixing matrix and the data matrix will remain invariant. Since we have used absolute Pearson correlation for similarity calculations, sources in one realization may be aligned with their sign-reversed version in another. Hence a procedure is employed to fix the signs of each aligned component; this is accomplished by arbitrarily assigning one source in each group to have the "canonical" sign, and then computing signed correlations of that source with all others in the group. Any source having a negative correlation with the canonical source is multiplied by minus one; the corresponding mixing column also obtains a sign change.

Once alignment and sign canonicalization are finished the weighted averages are constructed. For each of the n_S aligned BICAR sources, we compute the $K \times K$ set of estimate-estimate absolute cross-correlations. There are $K(K-1)/2$ unique such correlations, accounting for the symmetry of the cross-correlations and ignoring autocorrelation. We define the reproducibility of a BICAR source and its corresponding mixing elements as the average of these unique correlations. The weight for each of the K estimates in constructing the BICAR source is given by the average of the pairwise absolute cross-correlations of that estimate with all other estimates. These definitions place reproducibility in [0, 1]. Contributions to the reproducibility come from both time (EEG sources) and space (fMRI sources). Matching, sign canonicalization, source averaging, and reproducibility are either not present in, or quite different from, the RAICAR algorithm (Yang et al., 2008; Brown et al., 2012).

The choice of 63 sources – a full-rank decomposition of the EEG data – is motivated by previous work (Brown et al., 2012). By using BICAR on synthetic data of known source content, we investigated the result of source overextraction from the temporal data. Specifically, ten sources were extracted from a mixture containing only five true sources. In all cases, BICAR source reproducibility was correlated with BICAR source similarity to a true source, and the spurious sources migrated to the tail of the reproducibility distribution. This was true even in low signal-to-noise scenarios; in fact, in cases of high signal-to-noise the shape of the BICAR source reproducibility spectrum (see Fig. 3 for an example) could be used to estimate the number of true sources in the data. In these low noise cases, the sorted reproducibility values drop sharply between the last true source and the first spurious source (sources five and six in the synthetic example). Given that we do not know the true number of sources in the real neuroimaging data, this previous result gives us confidence that overextraction is safer than underextraction, in which we might not capture all the true neural sources.

2.2.2. Component display

Figs. 4–6 show BICAR components from the subjects in this study. Common display conventions for the BICAR components are as follows. In every case, for sources which are significantly task-associated, the temporal portion of the BICAR source is shown as a stacked plot of individual trials (as a raster) and trial average (as a line plot). No information about the temporal source is shown for nonsynchronized components. For both synchronized and nonsynchronized components, representative slices are shown for the spatial sources. More specific information follows:

Temporal: Time windows of extracted signals for the stimulus and response categories were (−0.24 s, 0.64 s) and (−0.72 s, 0.24 s), respectively, where the event of interest (search field presentation or button press) occurred at 0 s. For components which are synchronized to both categories, rasters and event average plots are shown separately for both event types (stimuli and responses). All extracted windows were baselined using the pre-event interval. Each raster has been transformed with a nonlinear (hyperbolic tangent) contrast adjustment for easier viewing; this transformation is for visualization only and was not used in calculation of event synchronization or the trial means. Shaded gray areas in the trial average plots are envelopes representing two standard errors of the mean, and the event onset ($t = 0$) is marked with a vertical dotted line in each signal average plot. The y-axis units in the signal average plot are arbitrary; each signal has been standardized before display. A colormap in which green is more positive and red more negative has been employed, but the overall sign of the component is essentially arbitrary – it is jointly carried by the source and the mixing elements (not shown).

Spatial: Each spatial source is shown as a series of representative slices, superimposed on that subject's anatomy. In each figure showing BICAR components (Figs. 4–6), the set of slices is always the same. They are numbers 35, 38, 41, 44, and 47. This is noted on the first BICAR source in each figure (top left); the numbering is suppressed elsewhere due to space constraints. BICAR spatial sources have been standardized before display. All spatial sources have been smoothed with an isotropic Gaussian filter with a half-width of 4 mm (1 voxel); as detailed in Experimental Design, no smoothing was performed on the fMRI images themselves. In addition, voxels with absolute intensity less than a chosen threshold of 0.25 have been suppressed. The same red/green colormap used for the temporal rasters has been used for the spatial sources, but again the sign of the component is essentially arbitrary – it is jointly carried by the source and the mixing elements (not shown).

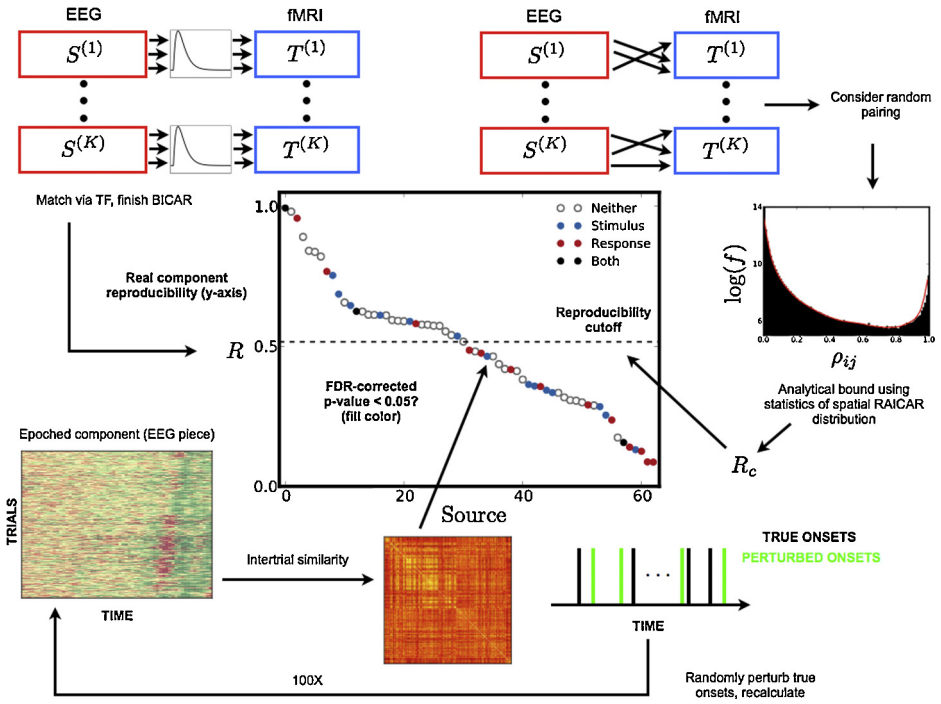

Fig. 3. Sample reproducibility spectrum with schematics for scoring. BICAR itself, via physical matching, gives each resulting component a reproducibility value (the *y*-axis in the central plot). The dotted line shows a sample reproducibility cutoff R_c, whose derivation is given in Section 3 and Appendix A. Filling the symbols in the central plot indicates significant task-association with one or both categories (all Stimuli, all Responses, or Both), with category indicated by color. Significance for task association is determined by comparing intertrial similarity when epoching to the real event locations versus a distribution of scrambled locations, in which each real trigger is moved in time by a uniform random number in (−500 ms, 500 ms). See Section 3 for additional details on these calculations.

2.3. Ranking components for task-relevance

BICAR run on each subject produces n_S joint components, each with an associated reproducibility value $R \in [0, 1]$. The components can then be secondarily scored on other criteria of interest. In this study only secondary scoring for phase synchronization to task events (i.e., evoked response analysis) is considered. Scores for task synchronization are obtained as follows. The EEG portion of the BICAR component is epoched separately with respect to each task event category of interest; for the analysis in this paper there are two categories of interest, all stimuli (search field presentations) and correct responses. The N_i trials from each category i are used to compute average inter-trial similarities for each component as follows. Computing the $N_i \times N_i$ correlation matrix $\mathbf{C}(i)$, we compute a synchronization index for component c and category i

$$s_c(i) = \frac{1}{N_i(N_i - 1)} \sum_{k \neq l} C_{kl}^2(i) \tag{2}$$

which sums the squares of the unique trial-to-trial cross correlations. The sum is restricted to off-diagonal terms, and the prefactor corrects for double counting the upper triangle and its counterpart reflected across the diagonal.

In this study each component has two values for s, corresponding to the categories chosen for analysis.[1] To determine whether these synchronization values are statistically significant, s in Eq. (2) is recomputed for sets of permuted stimuli in each category. For each category (all stimuli, all correct responses) a set of permuted task events is created in which the true trigger locations are moved in time by uniform random numbers in (−500 ms, 500 ms). These jitter times will be experiment-dependent, but the scoring method is not. These permuted trigger sets preserve the number of stimuli but onsets are shifted with respect to the true values in the experiment. For each permuted task $s_c(i)$ is recomputed for all components c. This process is replicated 100 times to obtain a category-specific distribution for $s_c(i)$. This distribution is then used to standardize the observed true s values, and the resulting standardized scores converted to p-values. A false discovery rate (FDR) correction (Benjamini and Hochberg, 1995) is applied to account for multiple hypothesis testing, and the calculated p-values compared to the FDR-corrected cutoffs for significance. In general, at the end of this process each component will be significantly associated with zero or more categories: for the subsequent analysis in this manuscript, that means stimuli, responses, both, or neither. The p-value threshold was 0.05 (corrected).

It should be emphasized here that this secondary scoring has a lot to do with how the experiment was performed and what quantities are of biological interest; for example, evoked versus spontaneous activity, correct or incorrect responses, etc. However, reproducibility is assessed in a *completely* task-independent manner; many different secondary analyses of interest can be applied to BICAR components, and none of them affect the BICAR decomposition.

3. Results

3.1. Reproducibility cutoff

BICAR ranks each paired component by reproducibility but by itself does not generate a significance level below which an individual subject's components are deemed to be spuriously paired during the matching step. Such an approximate cutoff can be analytically derived, and it is dependent on both the data and BICAR's

[1] There are two categories of interest in this study, but this process generalizes to any number of categories.

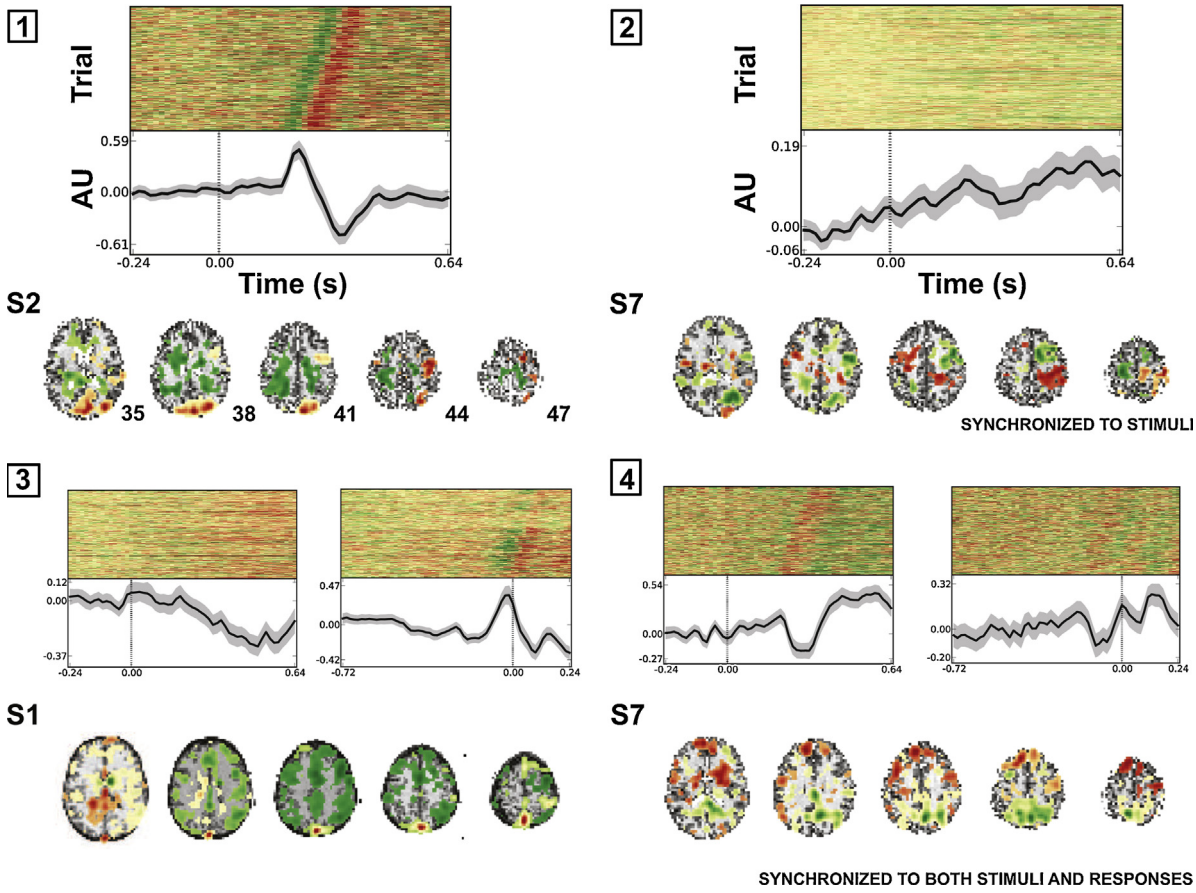

Fig. 4. Most reproducible stimulus-synchronized BICAR components. This figure and Figs. 5 and 6 show the four most reproducible BICAR components from each of the five subjects in this study. The portion of those twenty components that are significantly synchronized to either stimuli alone (blue box) or both stimuli and responses (black box) are shown here. The colors of the enclosing boxes have been chosen to match the point colors in Fig. 8 and give the task event association. All signals and slices are shown as described in Section 2. Slice numbering in the spatial sources is shown for the first source but suppressed thereafter; all spatial sources have the same displayed slices. Component numbering (in boxes) is arbitrary and used for easy reference in the text. Subject number (**S1**, **S2**, etc.) is included to the left of each component, and is the same as in Fig. 8.

parameters. More details on this derivation are given in Appendix A, but here we introduce the ideas behind, and summarize the results of, those calculations.

Consider a "worst-case" *random, nondegenerate* component pairing: if two individually reproducible components – one from EEG and one from fMRI – were paired spuriously, one could obtain an artificially highly reproducible BICAR component that does not represent a true joint component. Suppose there is a component in the EEG data which has reproducibility R^* close to unity; then the reproducibility R of a BICAR component whose spatial portion is randomly paired with this temporal source is

$$R = wR^* + (1-w)\left[\frac{2}{K(K-1)}\sum_{i,j;j>i} r_{ij}\right]. \tag{3}$$

Here, w is the weight given to the EEG data — all calculations in this paper use $w = 1/2$, but in deriving the cutoff we simply require $w \in [0, 1]$. r_{ij} is the absolute correlation between two spatial sources in different realizations. All BICAR sources consist of one component from each of K realizations, hence there will be $K(K-1)/2$ unique correlations to sum.

Calculating R in Eq. (3) therefore depends critically on what we will call the (spatial) RAICAR distribution, after Yang et al. (2008). Fig. 7A shows a sample RAICAR distribution from real fMRI data. Note the logarithmic scale on the y axis; most cross-realization absolute correlations are small. Considering random, nondegenerate pairing means that to calculate R one draws $K(K-1)/2$ random r_{ij}s from the RAICAR distribution and computes their mean. This then means that the reproducibility of random pairs depends on the data and algorithm parameters via the statistics of the RAICAR distribution.

The r_{ij} are certainly identically distributed, and if not completely independent very nearly so. For even moderate K the sum in Eq. (3) can have hundreds of terms. The central limit theorem thus gives the distribution for the term in brackets as:

$$\frac{2}{K(K-1)}\sum r_{ij} \sim N\left(\langle r_{ij}\rangle, \frac{2\sigma^2(r_{ij})}{K(K-1)}\right), \tag{4}$$

where $N(\mu, \sigma^2)$ is a normally distributed random number with mean μ and variance σ^2, and $\langle r_{ij}\rangle$ and $\sigma^2(r_{ij})$ are the mean and variance of the spatial RAICAR distribution, respectively. Since the sum is normally distributed, the spurious reproducibility R is also normally distributed as

$$R \sim N\left(wR^* + (1-w)\langle r_{ij}\rangle, (1-w)^2\frac{2\sigma^2(r_{ij})}{K(K-1)}\right), \tag{5}$$

a normal variate with

$$\langle R\rangle = wR^* + (1-w)\langle r_{ij}\rangle \tag{6}$$

$$\sigma(R) = \frac{\sqrt{2}(1-w)\sigma(r_{ij})}{\sqrt{K(K-1)}}. \tag{7}$$

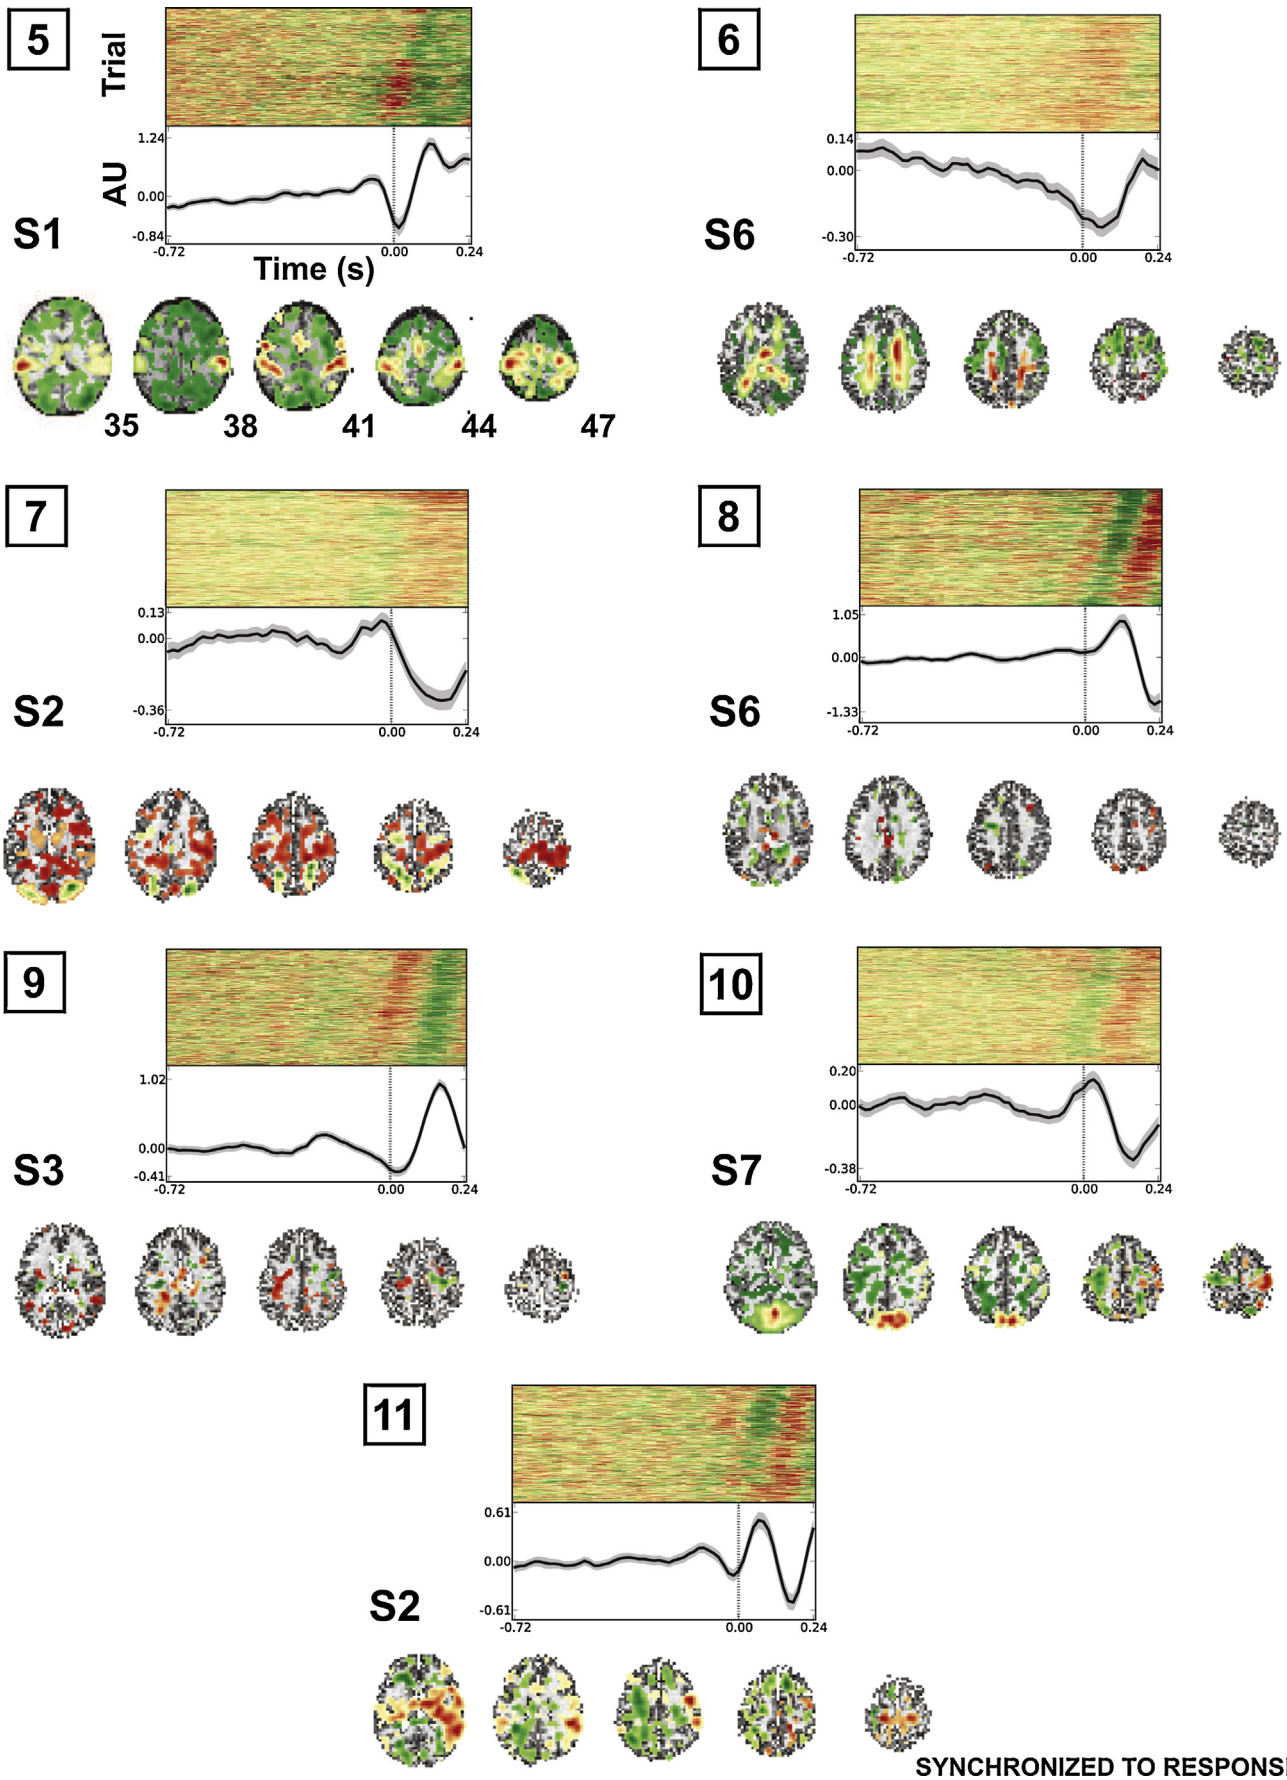

Fig. 5. Most reproducible response-synchronized BICAR components. The portion of the four most reproducible BICAR components from each of the five subjects which are significantly synchronized to responses. Component numbering continues that begun in Fig. 4. Slices, signals, and subject number are shown as described in Section 2 and the caption of Fig. 4.

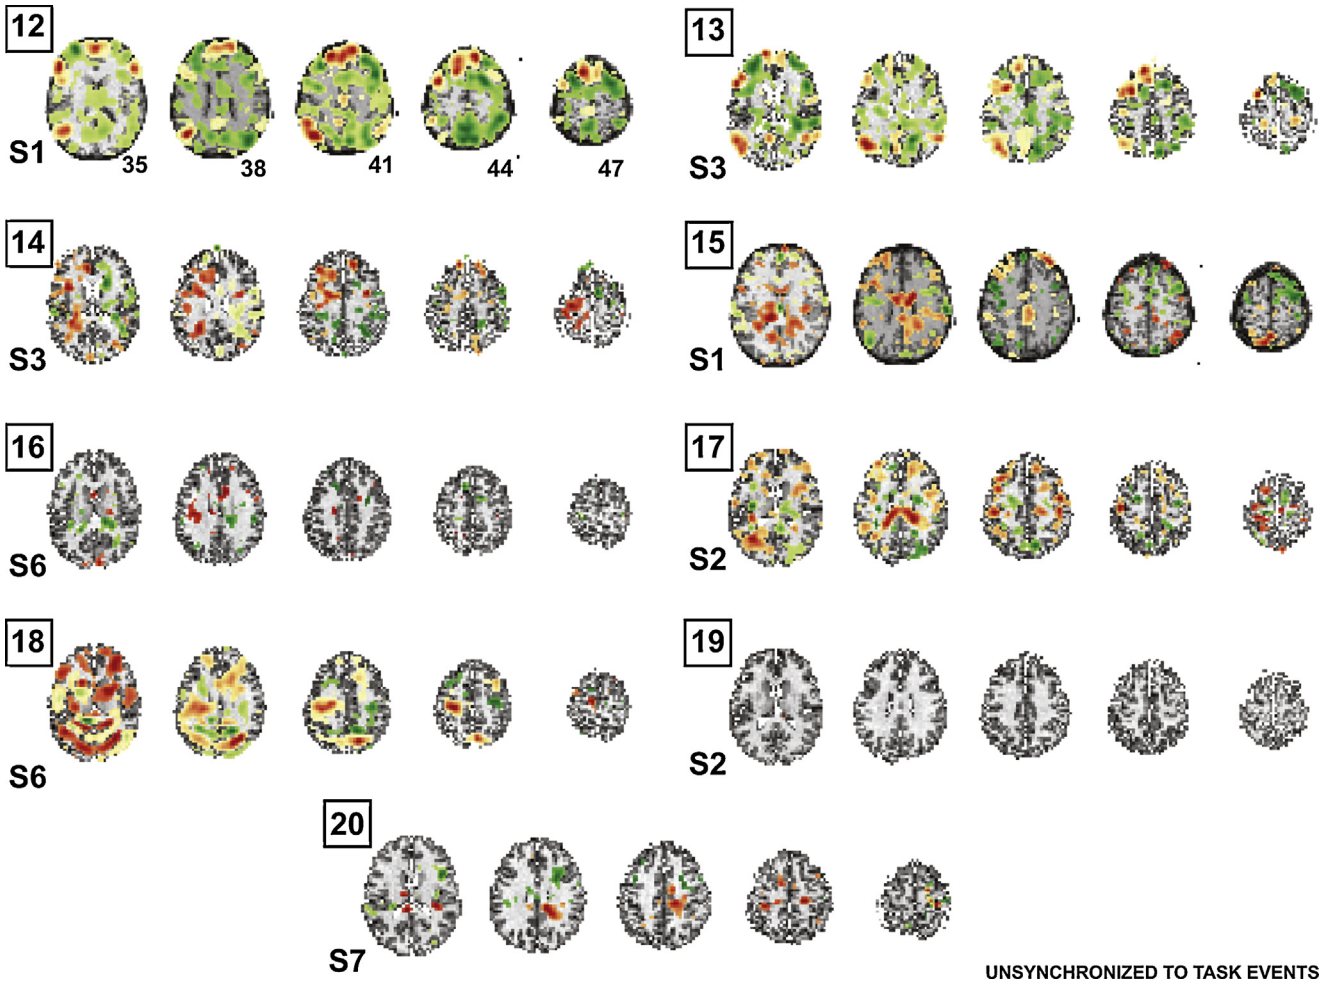

Fig. 6. Most reproducible nonsynchronized BICAR components. This figure shows the portion of the twenty BICAR components which were not found to be significantly task-associated. The component numbering here continues the numbering in Fig. 5 and is used for easy reference in the main text. Slice numbering is the same as in Figs. 4 and 5.

This spurious reproducibility distribution can be used to set a significance bound. Specifically, set the bound R_c as

$$R_c = \langle R \rangle + n_\sigma \sigma(R), \tag{8}$$

a desired n_σ standard deviations above the expected spurious reproducibility. Note that as $K \to \infty$, $R_c \to \langle R \rangle$. The reproducibility cutoffs calculated in this study use $w = 1/2$ and $n_\sigma = 2$. If we further assume the most reproducible component in the EEG data has $R^* = 1$, then

$$R_c = \frac{1 + \langle r_{ij} \rangle}{2} + \sqrt{\frac{2}{K(K-1)}}\sigma(r_{ij}) \tag{9}$$

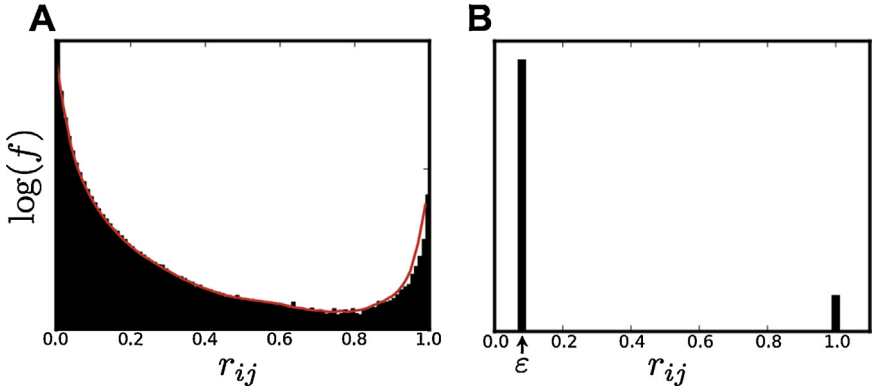

Fig. 7. Sample RAICAR distributions. (A) A sample real spatial RAICAR distribution for fMRI data, with a smoothed estimated of the histogram shown in red. Note the logarithmic scale; the vast majority of the cross-realization absolute correlations in real data are small. (B) Idealized RAICAR distribution used in the analytical calculation in Appendix A. For 63 extracted sources, the ratio of the density at ε to that at 1.0 is approximately 8.3. The scale in this figure is also logarithmic; the actual y values are unimportant in either case and both distributions have been rescaled for legibility.

One can go further than this if assumptions are made about the form of the RAICAR distribution; see Appendix A for details on these calculations. A few comments on what has led up to Eq. (9):

- Assuming that $R^* = 1$ is generally a slight overestimate but yields the most conservative bound.
- EEG (temporal) and fMRI (spatial) sources are assumed to be equally weighted for BICAR source reproducibility calculation ($w = 1/2$). Other choices, like weighting according to the number of samples in each, are possible, in which case Eq. (9) could be arrived at by combining Eqs. (6)–(8), with the chosen w.
- All cutoffs in this study use $K = 30$.
- This form for the cutoff has been calculated assuming nondegenerate (one-to-one) component pairing; degenerate (one-to-many) source matching would yield a different value of R_c.

From this point, the cutoff R_c can be determined using the data itself – via the spatial RAICAR distribution – and represents only a fraction of the total BICAR runtime. No simulation or perturbation testing is necessary to obtain this number. By making simple approximations to the true spatial RAICAR distribution, a closed-form, analytical expression for R_c can be obtained, which yields additional insight into the behavior of the cutoff as a function of the input data and algorithm parameters (see Appendix A). The cutoff is typically strongly dependent on n_S, the number of extracted sources, and weakly dependent on both K and the size n_V (number of voxels) of the spatial sources. The cutoffs computed in real EEG-fMRI data using the true, data-dependent spatial RAICAR distribution are quite close to those obtained using the analytical values with simplified spatial RAICAR distributions. For the algorithm parameters in this paper and typical EEG-fMRI data shapes, R_c is usually close to, but a bit larger than, one half.

3.2. Reproducibility spectra

Fig. 3 shows a sample BICAR reproducibility spectrum (center); the reproducibility spectrum is simply the reproducibility of each BICAR component, plotted in descending order. All BICAR decompositions in this paper have 63 components, so each reproducibility spectrum has 63 points. Along with the reproducibility values themselves, the spectral plots show the location of the reproducibility cutoff R_c (dotted, horizontal line, center of Fig. 3), and to which set of task events, if any, each source is significantly phase-locked (symbol color, center of Fig. 3). Refer to Section 3 and Appendix A for details about R_c and Section 2 for details about evoked response analysis.

Fig. 8 shows the BICAR reproducibility spectrum for each of the five subjects in this study. Refer to the schematic in Fig. 3 for help in interpreting the features of these graphs. All thresholds are drawn at their subject-specific levels, which are close to (but slightly larger than) 0.5. There are several features of note to these plots. First, all five subjects have many above-threshold BICAR components, but the number is quite variable. Subject **S7** has the fewest (13) while **S1** the most (31). It is also apparent that every subject save **S3** has a BICAR component with near perfect reproducibility, and all subjects have multiple components with $R > 0.75$. The number and type of task-synchronized components also varies widely among the five subjects; compare, for example, **S2** with **S6**. Despite this variability, all five subjects have one or more highly reproducible BICAR components significantly associated with either Stimuli (blue symbols), Responses (red symbols), or Both (black symbols). In general there are more response-locked BICAR components than any other category, which generally makes sense as the favorable signal-to-noise characteristics of the motor system likely make such components more easily extractable.

3.3. BICAR components

Ranking by reproducibility and employing the cutoff R_c provides a principled way of selecting BICAR components that avoids the use of prior information. In principle, the next step would be to group similar components across subjects, in which all above-threshold components for each subject are included. However, there are not enough subjects in this study to perform a proper group analysis, and there are far too many above-threshold components for each subject to show them all. There are two possible compromises that allow the display of an unbiased selection of BICAR results for this study. One would be to show the top N most reproducible components, irrespective of which subject they came from, and the other would be to show the top k components from each subject. The latter route was chosen, with $k = 4$, displayed in Figs. 4–6. Note that save for the presence of subject **S3**, these two strategies (the former using $N = 20$) would result in quite similar displays.

The summary statistics for the reproducibilities of these 20 BICAR components are as follows: $\overline{R} = 0.864$, $\sigma(R) = 0.10$, median$(R) = 0.841$, max$(R) = 0.999$, and min$(R) = 0.689$. The components are divided into three groups: Fig. 4 shows components significantly synchronized to either stimuli or both stimuli and responses, Fig. 5 shows components significantly synchronized to responses only, and Fig. 6 shows those which are not significantly phase locked to either set of task events. The colors of the figure borders are picked to match the symbol coloring in Fig. 8. Methods give a full description of the conventions for display of BICAR components. For components which were significantly synchronized to either stimuli, responses, or both (enclosed in blue, red, and black boxes respectively), rasters and trial averages are shown for the temporal sources and slices for the spatial sources. For nonsynchronized components (gray boxes), only slices (the spatial portion of the BICAR sources) are shown.

It is clear that among the most reproducible BICAR components there are joint sources with biological significance. Component one (Fig. 4) has strong activity in dorsal occipital cortex with a peak location of roughly 250 ms after search field presentation. Note also component four (Fig. 4); while it synchronizes to both stimuli and responses, its alignment to stimuli is more impressive and its spatial source shows some overlap with component one. Areas involved in motor planning (finger selection) (Grafton et al., 1998) and execution (button press upon target selection) are both present in multiple subjects. For example, consider components three, five, and ten (Fig. 5). Among the nonsynchronized components (Fig. 6) numbers twelve and thirteen are clearly relevant to the task (though not synchronized to triggers), as they include portions of the attention network and the task is spatial search.

A subject may lack a particular component in this display, but that does not necessarily mean it was not found. For example, subject **S2** has a component with $R = 0.62$ which is markedly similar to component one in Fig. 4 (not shown). However, it is not among that subject's top four most reproducible BICAR components, and hence is not displayed in Fig. 4. Inspection of Fig. 8 reveals it is the eleventh most reproducible **S2** component, but it falls above the cutoff. Subject **S6** has no above-threshold stimulus-associated components, and in fact none at all even when ignoring R_c. It is clear that not every biologically relevant component is present in every subject; however we have too few subjects in this study to make any more general claims or to perform a proper group analysis of all the above-threshold BICAR components.

Reproducibility and trial-locking alone are not guaranteed to remove artifacts. Notice component six in which the dominant activity is localized to the white matter, and component seventeen

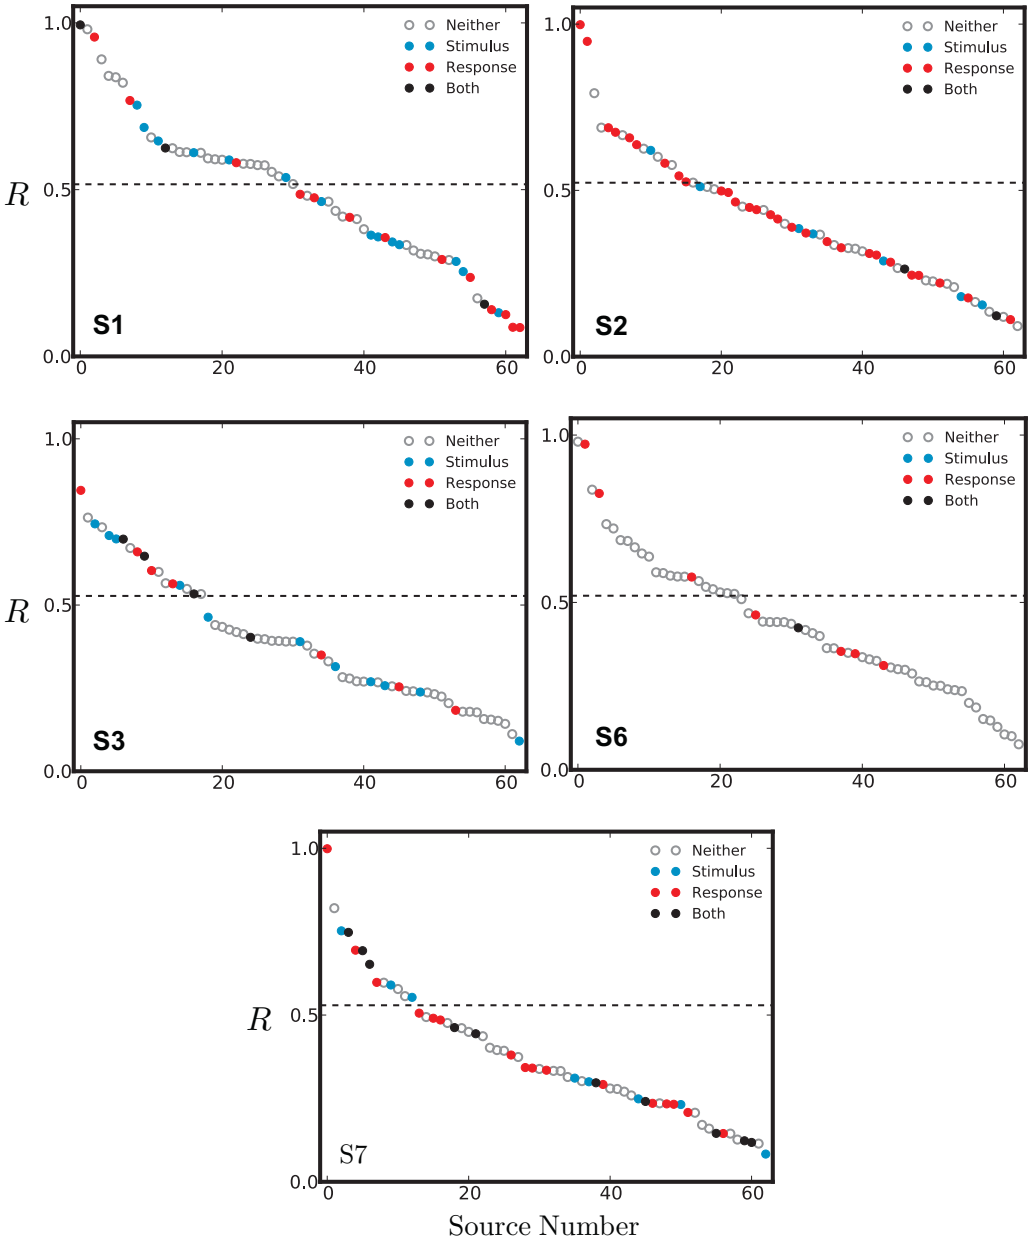

Fig. 8. Reproducibility spectra for the five subjects in this study. Fig. 3 gives a schematic guide to understanding the various features of reproducibility spectra and the calculations used to assess significance. Note that the number of highly reproducible components (those which sit above the subject-specific cutoff R_c, represented by a dotted horizontal line) varies among the five subjects, as do the number of significantly task-associated components. Compare, for example, subjects **S2** and **S6**.

that is at least partially a vascular artifact. Artifacts can also be reproducible! They can also be significantly phase-locked to task events, as for example susceptibility artifacts induced by a vertical head movement coincident with each button press or stimulus; while this particular artifact is not among the most reproducible components shown here, it is present as an above-threshold component in subject **S1** (not shown). Artifacts with stereotypical spatial or temporal patterns, in either EEG or fMRI, could be found by using a secondary scoring method that selects components based on similarity to temporal or spatial templates. Alternatively, an artifact correction (not trial-based rejection) method could be employed before running BICAR. This was not done in this study. Since many correction methods are ICA-based (Jung et al., 2000; Klados et al., 2011; Wallstrom et al., 2004; Urrestarazu et al., 2004; Vigário, 1997), if this route were chosen the resulting reduction in data dimensionality after correction would have to be taken into account when setting n_S in BICAR.

4. Discussion

BICAR, a new general algorithm for spatiotemporal data fusion was applied to concurrent EEG-fMRI data. The utility of the algorithm has been shown at the individual subject level of analysis using five participants who performed a spatial search task. A procedure was described to obtain a data-dependent significance cutoff for BICAR components, as well as a simple method to analyze BICAR components *post hoc* for evidence of statistically significant phase-locking to task events of interest. Some of the most reproducible BICAR components are clearly biologically relevant. Components representing visual processing, attention, motor planning, and execution were found in multiple subjects.

The response-locked components are more impressive than the stimulus-locked ones and easier to find; this likely has to do with the nature of the task. Since the real stimulus-associated cognitive event of interest is *not* associated directly with presentation of the

search field, but rather with localizing and discriminating the target item (Buschman and Miller, 2007), it can be difficult to locate. This is because localization and discrimination will occur at a variable time after search field presentation and before the motor response, making it poorly trial-aligned with either trigger. Nevertheless, for the purposes of this demonstration, there is a robust identification of movement and to a lesser degree, visual stimulus synchronized activity.

It bears repeating that *no* task information is used in BICAR itself. This means BICAR could be used in multiple ways, for example: (a) to reveal evoked activity, as in this manuscript, (b) to classify joint components based on frequency information, and (c) as an ICA-based pre-filter for removing non-reproducible activity from both the EEG and fMRI datasets, which could then be further analyzed using almost any existing EEG-fMRI method. Item (b) above is an obvious extension: trigger related or global (whole signal) time-frequency maps could be computed and components ranked according to power in frequency bands of interest (alpha, beta, mu, etc.). Yet another scoring method potentially relevant for the search task described in this manuscript would be to analyze the temporal mixing coefficients – scalp maps – for sterotypical patterns similar to those observed with the P300 event-related potential (Polich, 2007), and rank them in order of similarity to this pattern. Again, this would not require any reanalysis of the BICAR components, as the data itself has not changed, only the endpoint quantity of interest. Multiple secondary scoring methods could also be used. Scoring of components on any axes other than reproducibility will necessarily reflect the goals of the scientist and the experiment; reproducibility exists outside of these considerations.

There are several ways to improve BICAR, of which two stand out: optimization of the transfer function (TF) connecting the EEG and fMRI data from within BICAR itself, and integration of BICAR results into a population-level analysis. The first of these is important because, while we expect in general that the TF connecting the EEG and fMRI data will resemble an HRF, it is *not* an HRF. It connects dynamics to dynamics, in the form of transformed, decimated temporal sources to mixing matrix elements for spatial sources. BICAR needs a TF to run, and for this study a reasoned guess has been made about the likely shape of the TF and its parameters. The results even from this simple assumption are promising, but there are many reasons to believe a simple HRF is suboptimal for this problem. An algorithm for estimating the TF from within BICAR would benefit from calibration via simulations, with gains via simulation leading to direct improvements in real data applications, as TF estimates closer to the true value will generate better component matches and higher component reproducibility. Higher reproducibility components are of higher quality, as (a) low reproducibility is generated by averaging together groups of components with low similarity, and (b) our prior simulation studies have demonstrated that true (known) sources have higher reproducibility (Brown et al., 2012).

While the emphasis of the current work was to establish the reliability of the BICAR method at the individual subject level of analysis, additional methods will be needed to combine components across subjects for group level inference. The most effective method for placing BICAR results in a population-level analysis would likely involve data clustering. Clustering has been used to good effect with neuroimaging data, particularly for organizing ICA analyses of fMRI data (Himberg et al., 2004; Esposito et al., 2005), separating EEG data into microstates (Pascual-Marqui et al., 1995), and grouping voxel time series (Goutte et al., 1999). All above-threshold BICAR components across all study subjects could be clustered using a dissimilarity measure in both time (trigger-aligned averages of temporal sources) and space (image volumes). After clustering, group-average components could be constructed from the clusters, but reproducibility should again be taken into account here. Since the reproducibility of above-threshold BICAR components can vary by almost a factor of two (depending upon the value of R_c), the group averages should be reproducibility weighted. This allows the most reproducible components to carry the most weight in the group mean, which is as it should be, since they represent the most well-estimated activity.

One-to-one, nondegenerate matching of temporal and spatial sources was considered exclusively in this manuscript. In some situations, matching multiple temporal sources to a single spatial source may be appropriate (Yuah et al., 2012). Many-to-one degenerate matching is possible in BICAR but would require modifications to the existing algorithm. For one, the details of reproducibility cutoff calculation would change, though not its motivation and utility. In a situation where not all spatial components survive matching to be incorporated into BICAR components, to calculate the cutoff we would no longer be sampling from the full spatial RAICAR distribution (Fig. 7) but rather from a restricted distribution of only the cross-correlations among the subset of spatial components that are matched to a temporal component in any realization pair. This therefore means that we would have to calculate the cutoff after the matching step rather than after decomposition as done here, and the assumptions underlying our analytical calculations would no longer hold. However, even in the case of one-to-one matching, one could eschew the calculated cutoff we present here and instead create ensembles of random matches. While this would be much more computationally intensive, it would not be prohibitive, would be easily parallelizable, and would pose no additional theoretical difficulties. Hence one could always default to simulating the null model if a suitably accurate calculated cutoff could not be derived.

Another more subtle issue is the allowed degree of degeneracy: how few spatial components would we tolerate? It may be desirable to reward not only good match quality, in the sense of high correlation between transformed temporal sources and spatial loadings, but also to reward diversity of source selection. This would allow retention of more unique spatial components in the matches while sacrificing some total match quality, which would be particularly desirable if that match quality sacrifice was small compared to the number of additional components retained. A balance between goodness-of-fit and source inclusion could be obtained by matching using a multiobjective cost function that rewards high average correlation in the matches but penalizes the overall match set for including fewer unique spatial sources. However, the proper weighting of the goodness-of-fit and diversity terms would require substantial investigation. In addition, the use of a multiobjective cost function would raise the possibility of a Pareto front in the cost space (Messac and Mullur, 2007), in which one cannot simultaneously achieve both goals and must trade one for the other. Unlike for the reproducibility cutoff, the modifications to the matching algorithm that would allow for many-to-one degenerate matching would require substantial synthetic data studies in order to understand if and when source diversity is a problem and how to optimize the many-to-one matching process.

The results in this manuscript show that BICAR is a promising new data-driven method for mining concurrent EEG-fMRI data at the individual subject level. Future studies will be directed towards improving BICAR in the manner described above, as well as posing difficult experimental tests for BICAR's ability to separate sources of neural activity in both time and space.

Acknowledgments

This work was supported by the David and Lucile Packard Foundation and the Institute for Collaborative Biotechnologies through contract no. W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect

the position or the policy of the Government, and no official endorsement should be inferred.

Appendix A. Derivation of R_c

A.1. A workable example: setup

Calculating a subject and data-specific R_c in Eq. (8) requires calculation of the real spatial RAICAR distribution for each subject, which can be done anytime after the initial ICA step of BICAR is performed. In general, no further analytical expression can be obtained for R_c. However, insight into the dependence of the bound on algorithm parameters, particularly the number of extracted sources n_S, can be gained from considering a highly idealized spatial RAICAR distribution. Suppose decomposition of the spatial data produces one component that is perfectly reproduced in each of the K realizations, and all other components never appear in more than one of the K realizations. Then the spatial RAICAR distribution looks like that in Fig. 7B. The total number of r_{ij}s for n_S sources extracted in K realizations is

$$N_T = \left[\frac{n_S(n_S-1)}{2}\right]\left[\frac{K(K-1)}{2}\right], \tag{A.1}$$

$K(K-1)/2$ of which equal unity and $N_T-(K(K-1)/2)$ of which equal ε. The usage of a single nonzero number ε here is not strictly true; there will actually be a distribution of r_{ij}s with expected value ε. However, this will hardly matter, as will be shown below. Specifically,

- ε is small but not equal to zero. This is because, even for two sources with zero Pearson correlation, the expected value of their absolute correlation is *not* zero. The use of absolute Pearson correlation results in a small amount of skew to the right.
- The particular value of ε, and the spread around that value for many uncorrelated sources, depends on the size of the spatial source. In the EEG-fMRI case, this is the number of (independent) voxels.
- For sources the size of fMRI images ($O(10^4)$ voxels), ε is small ($O(10^{-2})$), and the error in using a single expected value in the calculation will be negligible.

A.2. Calculating ε

An approximate value for ε can be obtained as follows. Supposing the spatial sources are standardized, then the goal is to compute the expected value of

$$\varepsilon = \left\langle \left| \frac{1}{n}\sum_i x_i y_i \right| \right\rangle, \tag{A.2}$$

where i sums over the n voxels in the spatial source. Proceed by assuming x_i, $y_i \sim N(0, 1)$; this assumption of normality is certainly not true for the real sources, but will make little difference – only the mean and variance will matter. While the product of two Gaussian functions is also Gaussian, the distribution function of the product of two Gaussian random variates is distributed as Weisstein (2012)

$$P(z;\sigma_1,\sigma_2) = \frac{1}{\pi\sigma_1\sigma_2} K_0\left(\frac{|z|}{\sigma_1\sigma_2}\right), \tag{A.3}$$

where K is a modified Bessel function of the second kind. For the special case of K_0, K has the integral representation (Abramowitz and Stegun, 1972)

$$K_0(x) = \int_0^\infty \frac{\cos(xt)}{\sqrt{t^2+1}}\, dt. \tag{A.4}$$

$P(z;\sigma_1,\sigma_2)$ in Eq. (A.3) has a mean of zero and variance equal to $\sigma_1^2\sigma_2^2$.

It thus remains to compute

$$\varepsilon = \left\langle \left| \frac{1}{n}\sum_i z_i \right| \right\rangle, \tag{A.5}$$

with $z_i \sim P(z;\sigma_1,\sigma_2)$. The n in Eq. (A.5) will generally be quite large, as it is the number of brain voxels in the spatial sources. Thus the central limit theorem can be applied to this sum. Each term in the sum is iid. with finite mean and variance, hence as the number of terms in the sum grows large, it converges in distribution to a normal variate $N(\mu, \sigma^2/n)$, where μ and σ^2 here are the mean and variance of the distribution of each term. Since by assumption x_i, $y_i \sim N(0, 1)$, $\mu = 0$ and $\sigma^2 = (1*1)/n = 1/n$. Hence,

$$\varepsilon \approx \left\langle \left| N\left(0, \frac{1}{n}\right) \right| \right\rangle \tag{A.6}$$

The task of computing ε is now reduced to determining the distribution of $|z|$ when $z \sim N(\mu, \sigma^2)$. This is relatively straightforward using the cumulative density function. Suppose $f(z)$ is the distribution function for z, and $F(z)$ the cumulative density. Then

$$P(|x| < z) = P(x < z) - P(x < -z) \tag{A.7}$$

$$= F(z) - F(-z), \tag{A.8}$$

and the distribution function $g(|z|)$ for $|z|$ can be obtained via differentiation

$$g(|z|) = \frac{dP(|x| < z)}{dz} = f(z) - (-1)*f(-z) = f(z) + f(-z). \tag{A.9}$$

So far the fact that the distribution of z is known has not been used. For a normal density function, $f(-z) = f(z)$, and thus

$$g(|z|) = 2f(z). \tag{A.10}$$

For the special case of a normal $f(x)$, $g(x)$ is called the folded normal distribution (Leone et al., 1961), or for $\mu = 0$ a half-normal distribution. The folded normal for a normal distribution $N(\mu, \sigma^2)$ has mean $\langle x \rangle$ and variance σ_x^2

$$\langle x \rangle = \sigma\sqrt{\frac{2}{\pi}} e^{-\mu^2/2\sigma^2} + \mu\left[1 - 2\Phi\left(\frac{-\mu}{\sigma}\right)\right] \tag{A.11}$$

$$\sigma_x^2 = \mu^2 + \sigma^2 - \left(\sigma\sqrt{\frac{2}{\pi}} e^{-\mu^2/2\sigma^2} + \mu\left[1 - 2\Phi\left(\frac{-\mu}{\sigma}\right)\right]\right)^2, \tag{A.12}$$

where Φ is the CDF of a standard $N(0, 1)$ normal distribution. This simplifies considerably for a half-normal distribution to

$$\langle x \rangle = \sigma\sqrt{\frac{2}{\pi}} \tag{A.13}$$

$$\sigma_x^2 = \sigma^2\left(1 - \frac{2}{\pi}\right). \tag{A.14}$$

Hence by substituting $\sigma = \sqrt{1/n}$,

$$\varepsilon = \sqrt{\frac{2}{n\pi}} \tag{A.15}$$

$$\sigma_\varepsilon = \sqrt{\frac{1}{n}\left(1 - \frac{2}{\pi}\right)}. \tag{A.16}$$

Assuming a typical image size is 10^4 brain voxels, $\varepsilon = 0.008$, an extremely minor correction, and the width of the distribution is of the same order as the mean. For sources of size $n = 100$, $\varepsilon = 0.08$, which is not completely neglectable on a [0, 1] scale.

A.3. A workable example: calculation

Returning to the idealized RAICAR distribution in Fig. 7B, the mean and standard deviation can be directly computed. The probabilities for drawing a 1 and a ε from this distribution are, respectively,

$$p_1 = \frac{K(K-1)}{2N_T} = \frac{2}{n_S(n_S-1)} \tag{A.17}$$

$$p_\varepsilon = N_T - p_1 = \frac{(n_S-2)(n_S+1)}{n_S(n_S-1)}. \tag{A.18}$$

The mean and variance will therefore be, by definition,

$$\langle r_{ij}\rangle = p_1 + \varepsilon p_\varepsilon \tag{A.19}$$

$$\sigma^2(r_{ij}) = p_1 + \varepsilon^2 p_\varepsilon - \langle r_{ij}\rangle^2. \tag{A.20}$$

A bit of algebra and the use of the identity $p_1 + p_\varepsilon = 1$ yields

$$\langle r_{ij}\rangle = \frac{2}{n_S(n_S-1)}\left\{1 + \varepsilon\left(\frac{(n_S+1)(n_S-2)}{2}\right)\right\} \tag{A.21}$$

$$\sigma^2(r_{ij}) = (1-\varepsilon)^2\left[\frac{2(n_S-2)(n_S+1)}{n_S^2(n_S-1)^2}\right]. \tag{A.22}$$

To further simplify the cutoff, assume $R^* = 1$, $w = 1/2$, and $n_\sigma = 2$. Thus the cutoff in the example is

$$R_c(n_S, K, \varepsilon) = \frac{1}{2} + \frac{\langle r_{ij}\rangle}{2} + \sqrt{\frac{2}{K(K-1)}}\sigma(r_{ij}), \tag{A.23}$$

with $\langle r_{ij}\rangle$ and $\sigma^2(r_{ij})$ given by Eqs. (A.21) and (A.22). The notation for R_c has been modified to emphasize the explicit dependence on n_S, K, and ε. Notice that the cutoff will be strongly dependent on the number of extracted sources. For the parameters in this study,

$$R_c(63, 30, 0.008) = 0.506, \tag{A.24}$$

which is quite close to 1/2, a result satisfying to the intuition. However, for fewer ($n_S = 5$), smaller ($n = 10^2$) sources and fewer ($K = 10$) realizations,

$$R_c(5, 10, 0.08) = 0.675, \tag{A.25}$$

a substantially larger value.

The calculation of R_c for the idealized RAICAR distribution has clearly shown the strong dependence of the cutoff on both the input spatial data[2] and algorithm parameters. For this study, subject- and data-specific bounds were calculated using the empirical RAICAR distribution for that subject; they were within (0.51, 0.53), larger than the idealized bound but close to 1/2. For a large number of extracted sources, large K, and source vectors with thousands of elements, one could simply set $R_c = 1/2$. However, this would clearly be an underestimate in other situations, and this also ignores possibly important variations in the shape of the RAICAR distribution. Thus it is preferable to use the calculated R_c.

References

Abramowitz M, Stegun IA, editors. Handbook of mathematical functions with formulas, graphs, and mathematical tables. Dover; 1972.

Allen PJ, Josephs O, Turner R. A method for removing imaging artifact from continuous EEG recorded during functional MRI. Neuroimage 2000;12:230–9.

Babiloni F, Babiloni C, Carducci F, Del Gratta C, Rossini P, Cincotti F. Cortical source estimate of combined high resolution EEG and fMRI data related to voluntary movements. Methods of Information in Medicine 2002;41:443–50.

Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. Journal of the Royal Statistical Society Series B Statistical Methodology 1995;57:289–300.

Brookings T, Ortigue S, Grafton S, Carlson J. Using ICA and realistic BOLD models to obtain joint EEG/fMRI solutions to the problem of source localization. Neuroimage 2009;44:411–20.

Brown KS, Grafton ST, Carlson JM. BICAR: a new algorithm for multiresolution spatiotemporal data fusion. PLoS ONE 2012;7:e50268.

Brown KS, Ortigue S, Grafton ST, Carlson JM. Improving human brain mapping via joint inversion of brain electrodynamics and the BOLD signal. Neuroimage 2010;49(February (3)):2401–15.

Buschman TJ, Miller EK. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science 2007;315:1860–2.

Calhoun VD, Adali T, Pearlson GD, Pekar JJ. A method for making group inferences from functional MRI data using independent component analysis. Human Brain Mapping 2001;14:140–51.

Candès EJ, Wakin MB. An introduction to compressive sampling. IEEE Signal Processing Magazine 2008:21.

Chaudhary UJ, Duncan JS, Lemieux L. Mapping hemodynamic correlates of seizures using fMRI: a review. Human Brain Mapping 2011, Epub.

Christmann C, Koeppe C, Braus DF, Ruf M, Flor H. A simultaneous EEG-fMRI study of painful electric stimulation. Neuroimage 2007;34:1428–37.

Chun MM, Jiang YH. Contextual cueing: implicit learning and memory of visual context guides spatial attention. Cognitive Psychology 1998;36:28–71.

Daunizeau J, Grova C, Marrelec G, Mattout J, Jbabdi S, Pélégrini-Issac M, et al. Symmetrical event-related EEG/fMRI information fusion in a variational Bayesian framework. Neuroimage 2007;36:69–87.

Debener S, Ullsperger M, Siegel M, Engel AK. Single-trial EEG-fMRI reveals the dynamics of cognitive function. Trends in Cognitive Sciences 2006;10(12):558–63.

Dubois C, Otzenberger H, Gounot D, Sock R, Metz-Lutz MN. Visemic processing in audiovisual discrimination of natural speech: a simultaneous fMRI-EEG study. Neuropsychologia 2012;50:1316–26.

Eichele T, Calhoun V, Moosmann M, Specht K. Unmixing concurrent EEG-fMRI with parallel independent component analysis. International Journal of Psychophysiology 2008;67:222–34.

Esposito F, Aragri A, Piccoli T, Tedeschi G, Goebel R, DiSalle F. Distributed analysis of simultaneous EEG-fMRI time-series: modeling and interpretation issues. Magnetic Resonance Imaging 2009;27:1120–30.

Esposito F, Scarabino T, Hyvarinen A, Himberg J, Formisano E, Comani S, et al. Independent component analysis of fMRI group studies by self-organizing clustering. Neuroimage 2005;25:193–205.

Friston KJ, Holmes AP, Worsley KJ, Pohline J-P, Frith CD, Frackowiak RSJ. Statistical parametric maps in functional imaging: a general linear approach. Human Brain Mapping 1995;2:189–210.

Giesbrecht B, Sy JL, Guerin SA. Both memory and attention systems contribute to visual search for targets cued by implicitly learned context. Vision Research 2013;85:80–9.

Goutte C, Toft P, Rostrup E, Nielsen F, Hansen L. On clustering fMRI time series. Neuroimage 1999;9:298–310.

Grafton ST, Fagg AH, Arbib MA. Dorsal premotor cortex and conditional movement selection: a PET functional mapping study. Journal of Neurophysiology 1998;79:1092–7.

Grouiller F, Thornton RC, Groening K, Spinelli L, Duncan JS, Schaller K, et al. With or without spikes: localization of focal epileptic activity by simultaneous electroencephalography and functional magnetic resonance imaging. Brain 2011;134:2867–86.

Hämäläinen M, Ilmoniemi R. Interpreting magnetic fields of the brain: minimum norm estimates. Medical and Biological Engineering and Computing 1994;32:35–42.

Himberg J, Hyvarinen A, Esposito F. Validating the independent components of neuroimaging time series via clustering and visualization. Neuroimage 2004;22:1214–22.

Hyvärinen A, Oja E. A fast fixed-point algorithm for independent component analysis. Neural Computation 1997;9(7):1483–92.

Jeromin O, Pattichis MS, Calhoun VD. Optimal compressed sensing reconstructions of fMRI using deterministic and stochastic sampling geometries. BioMedical Engineering OnLine 2012;11:25.

Johnson JS, Woodman GF, Braun E, Luck SJ. Implicit memory influences the allocation of spatial attention in visual cortex. Psychonomic Bulletin & Review 2007;14:834–9.

Jung H, Sung K, Nayak KS, Kim EY, Ye JC. k-t FOCUSS: a general compressed sensing framework for high resolution dynamic MRI. Magnetic Resonance in Medicine 2009;61:103–16.

Jung T-P, Makeig S, Humphries C, Lee T-W, McKeown MJ, Iragui V, et al. Removing electroencephalographic artifacts by blind source separation. Psychophysiology 2000;37:163–78.

Klados MA, Papadelis C, Braun C, Bamidis PD. REG-ICA: a hybrid methodology combining Blind Source Separation and regression techniques for the rejection of ocular artifacts. Biomedical Signal Processing and Control 2011;6:291–300.

Leone FC, Nottingham RB, Nelson LS. The folded normal distribution. Technometrics 1961;3:543–50.

Liu Z, Kecman F, He B. Effects of fMRI-EEG mismatches in cortical current density estimation. Clinical Neurophysiology 2006;117:1610–22.

Lleras A, Von Mühlenen A. Spatial context and top-down strategies in visual search. Spatial 2004;17:465–82.

[2] Of course, there is dependence on the input temporal data through R^* as well, but we will simply use the most conservative value of $R^* = 1$.

Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. Nature 2001;412:150–7.
Lustig M, Donoho D, Pauly JM. Sparse MRI: the application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine 2007;58:1182–95.
Makeig S, Bell A, Jung T, Sejnowski T. Independent component analysis of electroencephalographic data. In: Advances in neural information processing systems 8; 1996. p. 7.
Martinez-Montes E, Valdés-Sosa P, Miwakeichi F, Goldman RI, Cohen MS. Concurrent EEG/fMRI analysis by multiway Partial Least Squares. Neuroimage 2004;22:1023–34.
McKeown MJ, Makeig S, Brown GG, Jung T-P, Kindermann SS, Bell AJ, et al. Analysis of fMRI data by blind separation into independent spatial components. Human Brain Mapping 1998;6:160–88.
Messac A, Mullur AA. Multiobjective optimization: concepts and methods. In: Optimization of structural and mechanical systems. World Scientific Publishing Co; 2007. p. 121–48.
Miller MB, Donovan CL, Aminoff EM, Bennett CM, Mayer R. Individual differences in cognitive style and strategy predict similarities in the patterns of brain activity between individuals. Neuroimage 2011;59(1):83–93.
Miller MB, Van Horn J, Wolford GL, Handy TC, Valsangkar-Smyth M, Inati S, et al. Extensive individual differences in brain activations during episodic retrieval are reliable over time. Journal of Cognitive Neuroscience 2002;14(8):1200–14.
Moosmann M, Eichele T, Nordby H, Hugdahl K. Joint independent component analysis for simultaneous EEG-fMRI: principle and simulation. International Journal of Psychophysiology 2008;67:212–21.
Mulert C, Jager L, Pogarell O, Bussfeld P, Schmitt R, Juckel G, et al. Simultaneous ERP and event-related fMRI: focus on the time course of brain activity in target detection. Methods and Findings in Experimental and Clinical Pharmacology 2002;24(Suppl D):17–20.
Niazy R, Beckmann C, Iannetti G, Brady J. Removal of FMRI environment artifacts from EEG data using optimal basis sets. Neuroimage 2005;28:720–37.
Ostwald D, Porcaro C, Bagshaw AP. An information theoretic approach to EEG-fMRI integration of visually evoked responses. Neuroimage 2010;49:498–516.
Ostwald D, Porcaro C, Mayhew SD, Bagshaw AP. EEG-fMRI based information theoretic characterization of the human perceptual decision system. PLoS ONE 2012;7:e33896.
Pascual-Marqui R, Michel C, Lehmann D. Low resolution electromagnetic tomography: a new method for localizing electrical activity in the brain. International Journal of Psychophysiology 1994;18(1):49–65.
Pascual-Marqui R, Michel C, Lehmann D. Segmentation of brain electrical activity into microstates: model estimation and validation. Bio-Medical Engineering 1995;42(7):658–65.
Polich J. Updating P300: an integrative theory of P3a and P3b. Clinical Neurophysiology 2007;118:2128–48.
Roberts K, Papadaki A, Goncalves C, Tighe M, Atherton D, Shenoy R, et al. Contact heat evoked potentials using simultaneous EEG and fMRI and their correlation with evoked pain. BMC Anesthesiology 2008;8:8.
Rosenkranz K, Lemieux L. Present and future of simultaneous EEG-fMRI. MAGMA 2010;23:309–16.
Sato JR, Rondioni C, Sturzbecher M, de Araujo DB, Amaro E. From EEG to BOLD: brain mapping and estimating transfer functions in simultaneous EEG-fMRI acquisitions. Neuroimage 2010;50:1416–26.
Schultze-Kraft M, Becker R, Breakspear M, Ritter P. Exploiting the potential of three dimensional spatial wavelet analysis to explore the nesting of temporal oscillations and spatial variance in simultaneous EEG-fMRI data. Progress in Biophysics and Molecular Biology 2011;105:67–79.
Sekihara K, Nagarajan S, Poeppel D, Miyashita Y. Reconstructing spatio-temporal activities of neural sources using an MEG vector beamformer technique. IEEE Transactions on Biomedical Engineering 2001;48(7):760–71.
Strobel A, Debener S, Sorger B, Peters JC, Kranczioch C, Hoechstetter K, et al. Novelty and target detection processing during an auditory novelty oddball: a simultaneous event-related potential and functional magnetic resonance imaging study. Neuroimage 2008;40:869–83.
Urrestarazu E, Iriarte J, Alegre M, Valencia M, Viteri C, Artieda J. Independent component analysis removing artifacts in ictal recordings. Epsilepsia 2004;45:1071–8.
Vigário R. Extraction of ocular artifacts from EEG using independent component analysis. Electroencephalography and Clinical Neurophysiology 1997;103:395–404.
Wallstrom GL, Kass RE, Miller A, Cohn JF, Nathan AF. Automatic correction of ocular artifacts in the EEG: a comparison of regression-based and component-based methods. International Journal of Psychophysiology 2004;53:105–19.
Weisstein EW. Normal product distribution, from *MathWorld – A Wolfram Web Resource*; 2012.
Whittingstall K, Stroink G, Schmidt M. Evaluating the spatial relationship of event-related potential and functional MRI sources in the primary visual cortex. Human Brain Mapping 2007;28:134–42.
Wu L, Eichele T, Calhoun VD. Reactivity of hemodynamic responses and functional connectivity to different states of alpha synchrony: a concurrent EEG-fMRI study. Neuroimage 2010;52:1252–60.
Yang Z, Laconte S, Weng X, Hu X. Ranking and averaging independent component analysis by reproducibility (RAICAR). Human Brain Mapping 2008;29(6):711–25.
Yuah H, Zotev V, Phillips R, Drevets WC, Bodurka J. Spatiotemporal dynamics of the brain at rest – exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. Neuroimage 2012;60:2062–72.