Multi-modal data fusion using source separation: Application to medical imaging

Tülay Adalı, *Fellow, IEEE,* Yuri Levin-Schwartz, and Vince D. Calhoun, *Fellow, IEEE*

Abstract—The Joint ICA (jICA) and the Transposed IVA (tIVA) models are two effective solutions based on blind source separation that enable fusion of data from multiple modalities in a symmetric and fully multivariate manner. In [1], their properties and the major issues in their implementation are discussed in detail. In this accompanying paper, we consider the application of these two models to fusion of multi-modal medical imaging data—functional magnetic resonance imaging (fMRI), structural MRI (sMRI), and electroencephalography (EEG) data collected from a group of healthy controls and patients with schizophrenia performing an auditory oddball task. We show how both models can be used to identify a set of components that report on differences between the two groups, *jointly*, for all the modalities used in the study. We discuss the importance of algorithm and order selection as well as trade-offs involved in the selection of one model over another. We note that for the selected dataset, especially given the limited number of subjects available for the study, jICA provides a more desirable solution, however the use of an ICA algorithm that uses flexible density matching provides advantages over the most widely used algorithm, Infomax, for the problem.

Index Terms—Keywords: data fusion, multimodality, fMRI, MRI, EEG, medical imaging, source separation, ICA, IVA.

I. Introduction

In the medical field, acquisition of multiple brain imaging modalities from the same participants has been a common practice for some time. Different modalities report on different aspects of the given problem, and contain common, or complementary, information as well as distinct. Hence, it is of particular interest to leverage the use of information that is common across different modalities to enable their full interaction for inference. Given the fact that data from multiple modalities has been available for some time, it is not a surprise that the field has been a particularly active one. While most of the earlier approaches, and a great number of those today, are essentially integrative in nature, where one modality is used to constrain another [2], [3], recently the emphasis has been shifting to methods that enable *true fusion*, where different modalities can influence each other, hence making full use of the "cross-information" that exists across modalities.

T. Adalı and Y. Levin-Schwartz are with the Department of CSEE, University of Maryland, Baltimore County, Baltimore, MD 21250, USA. V. D. Calhoun is with the University of New Mexico and the Mind Research Network, Albuquerque, NM 87106, USA. Email:{adali@umbc.edu}. This work was supported by the grants NSF-IIS 1017718, NSF-CCF 1117056, NIH 2R01EB000840, and NIH COBRE P20GM103472. The MATLAB codes of the ICA and IVA algorithms used in the examples are available at http://mlsp.umbc.edu/ resources.html and most are also incorporated into the toolboxes GIFT http://mialab.mrn.org/software/gift and FIT http://mialab.mrn.org/software/fit.

Methods based on latent variable analysis use either a matrix or tensor representation to decompose the available multi-modal data into factors/components to explain the relationships, and hence are both multivariate and data driven. While coupled or linked tensor decompositions [4], [5], dictionary learning [6] or nonnegative matrix factorization [7] enable such decompositions as well, their application to medical image fusion has been limited, see *e.g.*, [2], [8], [9]. Techniques based on blind source separation (BSS) on the other hand, have found wide use in multi-modal medical image fusion [2], [3], [10]. A reason for their attractiveness has been their ability to generate useful decompositions with few assumptions and, in general, without the need to introduce additional constraints for uniqueness.

The Joint ICA (jICA) model introduced in [11] has been implemented in the Fusion ICA Toolbox [12] and widely used in fusion of multiple medical imaging modalities including functional magnetic resonance imaging (fMRI) and electroencephalography (EEG) [3], [11], [13]–[17], genetic and fMRI [18], structural MRI (sMRI) and EEG [19], functional and structural cardiac imaging data (current density imaging and diffusion tensor imaging) [20], functional near-infrared spectroscopy and fMRI [21], diffusion tensor imaging and fMRI [22], EEG and electromyography (EMG) [23], positron emission tomography and sMRI [24] data, among many others.

Two other multivariate approaches, canonical correlation analysis (CCA) [25] and partial least squares (PLS) [26] also provide full interaction of two datasets, and enable fusion using second-order-statistics (SOS), as opposed to the higher-order statistics (HOS) considered in ICA. Both have been applied to fusion of multi-modal medical imaging data and examples of CCA include [27], [28] and of PLS [29], [30], where the focus has been primarily on association of two functional imaging modalities, EEG and fMRI. The extension of CCA to multiple datasets, multiset CCA (MCCA) [31], has been used to find associations across three modalities, fMRI, EEG, and sMRI [28], [32], and to perform fusion of concurrently acquired EEG and fMRI data [13], [33]. In addition, the latter problem has been also addressed using an extension of PLS to multiset data [30]. On the other hand, when jICA is applied to fusion of three modalities, generally, its extensions are considered as in [34], [35] since common covariations across three modalities is a serious constraint.

In the accompanying paper [1], we introduce the Transposed independent vector analysis (tIVA) model as an extension of the MCCA-based model proposed in [27], [28],

and study the properties of tIVA along with jICA that has been widely applied to fusion of multi-modal medical image data. Through its generalization of ICA to multiple datasets, IVA enables the incorporation of HOS into the MCCA-based model and, as demonstrated with multiple simulation examples in [1], can provide advantages over the jICA model that has been widely used. In this paper, we use fMRI, sMRI, and EEG data collected from healthy controls and patients with schizophrenia performing an auditory oddball (AOD) task, and demonstrate the application of the two models, jICA and tIVA, to this dataset. In the process, we revisit the main considerations such as algorithm and order choice, and discuss implementation details. We address the difficult problem of validation for the medical domain and hope to provide guidance in the application of these two models and others such as the MCCA+jICA model [36]. We show that for this given dataset, the jICA model might be preferable over tIVA, and in the case of jICA, use of a flexible ICA algorithm with adaptive density matching capability, like the entropy bound minimization (EBM) algorithm [37], is more preferable than the Infomax algorithm [38], which is typically used for jICA. Next, we introduce the data used in the study, and then, in Section III, discuss the details of the implementation, and present the results in Section IV. The paper concludes with a discussion of the results in Section V.

II. Data: Modalities and features used

We use medical imaging data from three modalities, fMRI, sMRI, and EEG, to compare the performance of the two models, jICA and tIVA, introduced in [1], and consider issues such as algorithm choice and order selection, again discussed in [1], for this selected dataset. The data used in the study was collected from 36 subjects—22 healthy controls and 14 schizophrenia patients. The fMRI and EEG data were collected separately while the subjects performed an auditory oddball (AOD) task that required them to press a button when they detect a particular infrequent sound among three kinds of auditory stimuli. Details of the task design and the participants are given in [39]. Both fMRI and EEG report on the brain function, but fMRI by measuring changes in the blood-oxygenation level, while EEG by measuring electrical field through the scalp. Since we work with multivariate features extracted from both—as explained next—associations become possible for these two very different types of functional measurement. Structural MRI, on the other hand, provides information about the tissue type of the brain: gray matter, white matter, and cerebrospinal fluid. Since structure underlies brain function, we study the effect of bringing this third modality into the fusion study. Next, we briefly introduce each modality, the preprocessing used for each dataset, and the multivariate features generated for the fusion of these data.

FMRI data provides a measure of brain function on a millimeter spatial scale and a sub-second (and delayed) temporal scale. The data consists of repeatedly imaging the 3-D volume of the brain slice-by-slice, usually while the subject performs a particular task. A number of preprocessing steps important for fMRI: slice timing correction to account for the sequential acquisition of the slices, registration to correct for subject motion in the scanner, spatial filtering to reduce noise, and spatial normalization to be able to compare brains across different individuals and to use standardized atlases to identify particular brain regions. For fMRI data, we use the task-related spatial activity maps as calculated by the general linear regression approach using Statistical Parametric Mapping (SPM8) [40] as the spatial features for the fusion analysis.

SMRI analysis refers to the acquisition and processing of T1-, T2-, and/or proton density-weighted images. Multiple structural images are often collected to enable multispectral segmentation approaches. The primary outcome measure in a structural image may include a measure of a particular structure (*e.g.*, volume or surface area) or a description of the tissue type (*e.g.*, gray matter or white matter). There are many methods for preprocessing sMRI data which may include bias field correction for intensity changes caused by radio frequency (RF) or main magnetic field (B_0) inhomogeneities [41], spatial linear or non-linear [42] filtering and normalization. MRI images are typically segmented using a tissue classifier producing images showing the spatial distribution of gray matter, white matter, and cerebrospinal fluid. Both supervised and automated segmentation approaches have been developed for sMRI analysis [43]–[45] and each technique is optimized to detect specific features. We use probabilistically segmented gray matter images as the sMRI features for the fusion study.

EEG is a technique that measures brain function by recording and analyzing the scalp electrical activity generated by brain structures. Like MRI, it is a noninvasive procedure that can be applied repeatedly in patients, healthy adults, and children with virtually no risks or limitations. Local current flows are produced when brain cells are activated. It is believed that contributions are primarily driven by large synchronous populations of firing neurons. The recorded electrical signals are then amplified, digitized, and stored. Event-related potentials (ERPs) are small voltage fluctuations resulting from evoked neural activity and are one of many ways to process EEG data. These electrical changes are extracted from scalp recordings by computer averaging of the epochs—recording periods—of EEG time-locked to repeated occurrences of sensory, cognitive, or motor events. The spontaneous background EEG fluctuations, which are typically random relative to when the stimuli occurred, are averaged out, leaving the ERPs. These electrical signals reflect only the activity that is consistently associated with the stimulus processing in a time-locked way. The ERP, thus, reflects, with high temporal resolution, the patterns of neuronal activity evoked by a stimulus. Due to their high temporal resolution, ERPs provide unique and important timing information about the processing in the brain and are an ideal tool for studying the timing aspects of both normal and abnormal cognitive processes. We use ERPs from the midline central position (Cz) for the fusion study because it appeared to be the best single channel to detect both anterior and posterior sources for the given task.

The feature data from the three modalities is thus formed

into three matrices, a feature for each of 36 subjects by the number of voxels for fMRI and sMRI, $V_f = 60,261$, $V_s = 306,626$, respectively, and subjects by time points for ERP, 36 by $T = 451$.

III. EXPERIMENTAL SET-UP AND IMPLEMENTATION

When addressing the choice of a model and an algorithm for the given model, first, the available data determines certain choices. In our case, since the data comes from a total of 36 subjects, for the tIVA model, use of an algorithm that takes HOS into account would not be feasible as for reliable computation of HOS one would require higher number of samples, at least a couple of hundred. Thus, we compare the performance of jICA using two algorithms, Infomax [38] and EBM [37], and use only algorithms that take SOS into account for the tIVA model. We consider two implementations of MCCA, those using the MAXVAR and the GENVAR cost functions, and for tIVA, consider only IVA with multivariate Gaussian (IVA-G) [46]. The MAXVAR cost achieves MCCA by maximizing the largest eigenvalue of the source component vector autocovariance matrix and GENVAR by minimizing the product of the eigenvalues of the autocovariance matrix [1], [31]. When data from only two modalities is fused, then both MAXVAR and GENVAR reduce to CCA. In the subsequent discussions, we use jICA and tIVA when specifically referring to the model, and when specifying the algorithm, we use two versions for the tIVA model, tMCCA and tIVA such that tMCCA-MAXVAR and tMCCA-GENVAR refer to the use of MAXVAR and GENVAR costs, and tIVA-G to use of IVA-G. For jICA, we have jICA-Infomax and jICA-EBM for jICA using Infomax and EBM. Next, we discuss the main considerations in the implementation of these two models for the given dataset.

Selection of the most stable run: Other than CCA and MAX-VAR, all the algorithms are of iterative type, and hence, as in [47], [48], it is important to make sure that we evaluate the consistency of a given algorithm using multiple runs and identify a result that represents a consistent estimate to use as the final output of that algorithm. This is achieved by running the algorithm multiple times and identifying what we call the *best run*. We used the minimum spanning tree (MST) method introduced in [49] and implemented in the toolbox Group ICA of fMRI (GIFT) [50] for the task. The method generates a matrix for each pair of runs whose elements are given by 1 minus the correlation coefficient between the estimated components from those runs. The minimum cost of aligning the sources between runs is found, and a graph is constructed with nodes corresponding to runs and edges corresponding to the minimum cost of aligning the runs. An MST, with corresponding central node, is generated for this graph. Since the method was originally proposed for ICA, it can be directly applied to the spatially concatenated "joint" components for jICA. For the tIVA model, the components (spatial maps) are aligned based on the sign of the correlation between the corresponding profiles such that whenever the correlation is negative for a given modality, then both the estimated

component and corresponding profile for that component are multiplied by -1. We have used 10 runs for each iterative method and used average pairwise correlation between aligned components of each run to evaluate the consistency of a given algorithm, hence a modified version of the consistency compared to [49].

Selection of the components of interest: Since the data we use comes from a patient and a control group, we can use the fusion results to identify differences between the groups, *i.e.*, to identify biomarkers for schizophrenia. These correspond to components estimated directly as the result of the ICA or IVA decomposition, for fMRI, these are spatial activation maps, responses for the ERP, and structural maps for sMRI, reflecting increase/decrease in the case of fMRI and sMRI (in activation levels and gray matter volume respectively), and deviation from the expected ERP for the given task in the case of EEG data. To identify the components of interest, we perform a 2-sample t-test on the loadings (values) of the profiles for the two groups and identify those that pass a given significance threshold. The component corresponding to a profile that demonstrates significant difference, as measured by the t-test then yields the result of the fusion for the given modalities. This is demonstrated in the simulation example given in Section IV and shown Figure 4 in [1]. Hence, the sign of the t-statistic determines whether the activation of the component is higher for patient or controls. We display only components that have a positive sign—meaning greater activation for controls than patients—in the t-statistic corresponding to the two-sample t-test run on the corresponding column of the mixing matrix. If t-statistic is negative, we multiply the column of the mixing matrix and corresponding estimated source by -1. In the components that we display, which are all normalized through the Z-score—the activations (red, orange, yellow) correspond to an increase in controls over patients and deactivations (blue) correspond to a decrease in controls versus patients. We use a p-value of 0.05 as the threshold to determine the candidate component maps, and without correction to be able to generate a significant number of candidates as components of interest. We retain the most significant five components for fMRI, sMRI, and ERP that correspond to the profiles with p-values that exceed this significance level, *i.e.*, have p-values lower than 0.05. These components are then further evaluated using prior information about the properties of spatial and temporal components to identify the final components of interest. For fMRI, we expect smooth and focal regions of activity, and similarly for sMRI, smooth regions in the gray matter. The ERP components, on the other hand, would be expected to have a smooth response, typically with peaks that would coincide with those of the average ERP such as N1, N2, and P3 shown in Figure 1 along with the estimated ERP components. To complement this information, in a joint BSS framework, the results of a single modality ICA results are helpful as well, since they provide guidance with evaluation of the results. We performed single modality analysis using Infomax and EBM.

Order selection and other processing steps: Another important consideration is the best way to balance the contributions of each modality to the fusion results, which is especially important for the jICA model. Since the dimensionality of the datasets are significantly different, for the jICA model, to balance the contributions of each, ERP data is repeated 100 times—simple interpolation of ERP also generated the same result—and the sMRI and fMRI data are used as is. After this step in jICA, and for the original data for tIVA, the data from each modality is normalized to have zero mean and unit variance. For order selection, using the data that is balanced in contribution as described, we first determined the order for jICA using MDL/BIC criterion with the sampling approach to generate i.i.d. samples [51], which yielded the following orders: 14 for all three modalities, 9 for fMRI-ERP, 32 for fMRI-sMRI, and 32 for sMRI-ERP. To enable comparisons across models and two vs three modality results, a single order is chosen for all the results, which is 15. This decision was based on the guidance using the information theoretic criterion as well as evaluation of the stability and quality of the results as discussed in Section III.C in [1]. For this problem, the evaluation is based on the statistical significance of the estimated profiles and the interpretability of the estimated components as in [28], [52]. It was observed that the results were quite similar in terms of t-statistics and the estimation of components that are most significant for orders 10 and 20, whereas the performance started to degrade in terms of t-statistics and the estimated areas started to significantly change at orders 5 and 25. Hence, the final dimensions of the three datasets used for jICA were $\mathbf{X}_f \in \mathbb{R}^{N \times V_f}$, $\mathbf{X}_s \in \mathbb{R}^{N \times V_s}$, and $\mathbf{X}_e \in \mathbb{R}^{N \times T_e}$ for fMRI, sMRI, and ERP respectively $N = 15$, $V_f = 60,261$, $V_s = 306,626$, and $T = 45,100$.

For the tIVA model, since the number of samples—subjects in this case—is much smaller than the number of voxels or time points, using a data-driven technique such as pairwise CCA does not generate useful results for determining the order of common subspace across modalities. Using the procedure in [53], the dimensions of signal and common subspace (n, r) for pairwise combinations are computed as $(17, 3)$ for fMRI-ERP, $(17, 4)$ for fMRI-sMRI, and $(12, 1)$ for sMRI-ERP. Again, the estimation was stable for orders 10 and 20 indicating that 15 was a good target order for tIVA as well, and would allow direct comparison with the jICA results. Thus, for the tIVA model, all three datasets were of dimension 36×15.

IV. RESULTS

First, it is important to note that certain components, or more specifically, certain spatial areas for fMRI and sMRI and temporal characteristics for ERP, were consistently detected across models and algorithms, an aspect that increases the overall confidence in the final results. In particular, for the jICA model, the components for the two algorithms, Infomax and EBM were very similar with one important difference. For EBM, most of the components that passed our initial screening—*i.e.*, had a p-value in terms of group difference smaller than 0.05—were also deemed interpretable,
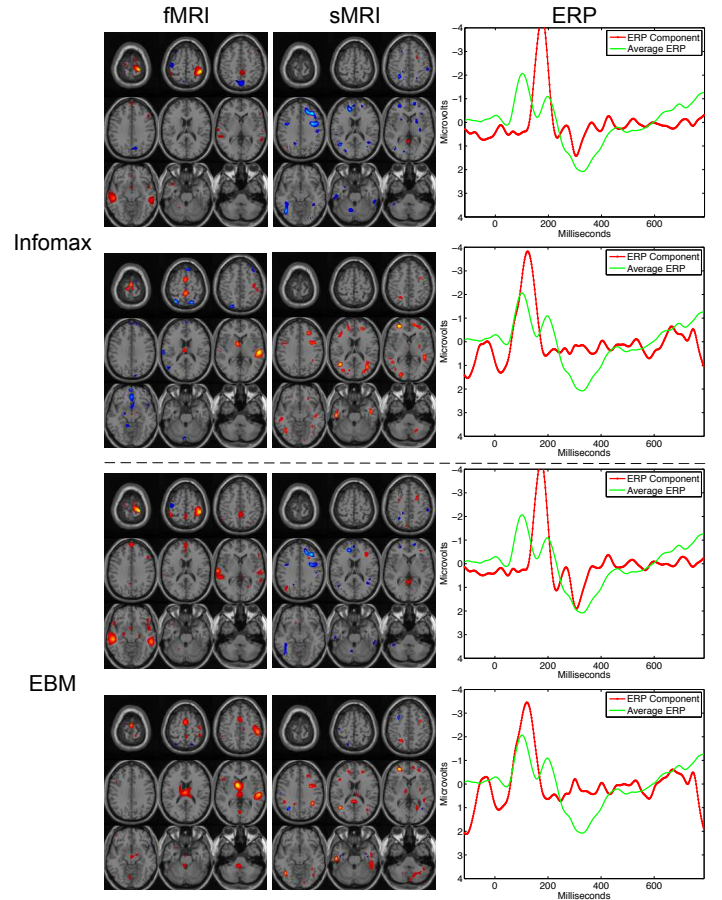

Fig. 1. Components generated for fusion of all three datasets using jICA. The significance and average consistency values for the components estimated using jICA-Infomax are (.0046, 1.00) for the top component and (.023, 1.00) for the second one. The values for jICA-EBM are (.0021, .96) for the first component and (.0048, .96) for the second.

i.e., physically meaningful and hence useful, while a good number of candidates generated by Infomax that exceeded the set threshold were eventually eliminated as they did not correspond to physically meaningful components. Among the modalities, ERP appeared to be most informative while sMRI the least if we were to use the group difference p-values as the main criterion for this given dataset. However, in all cases, all three modalities proved to be useful as we explain next. More importantly, when available, the results point to advantages of using all available modalities jointly in a fully multivariate framework as we propose. We performed the analyses using all three modalities as well as pairs of modalities—both with the jICA and tIVA models—and compared those against results from fusion of all three modalities. In general, while similar components were estimated, the three way results showed smaller areas, hence were more specific in terms of the detected activation differences for fMRI and gray matter differences for sMRI.

We first discuss the jICA results and then compare the tIVA results with those. It is important to reiterate that the components we show correspond to areas in fMRI and sMRI, and temporal characteristics in ERP that are fully linked
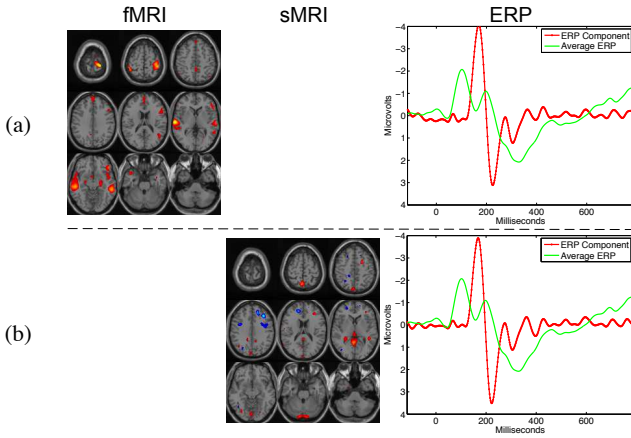
Fig. 2. Components generated using jICA-EBM with two modalities at a time. The significance and average consistency values for the two components are (a) (.0015, 1.00) for the fMRI-ERP pair, and (b) (.0021, .85) for sMRI-ERP.

Fig. 3. Components generated using EBM using single modality. The significance and average consistency values for components are given by (a) (.039, .78), (b) (.042, .85), (c) (.010, 1.00), and (d) (.045, 1.00).
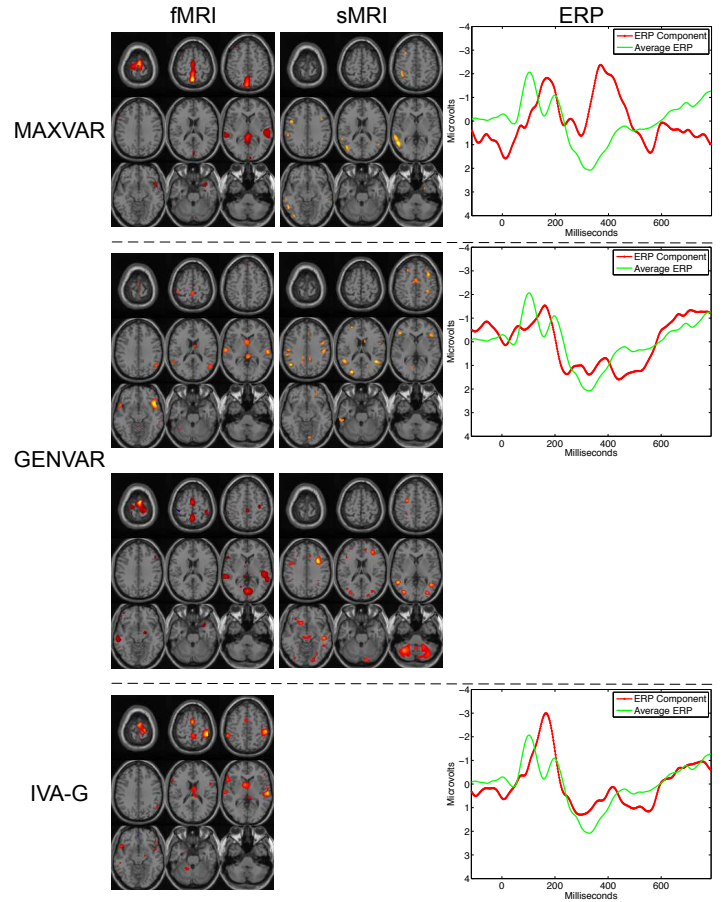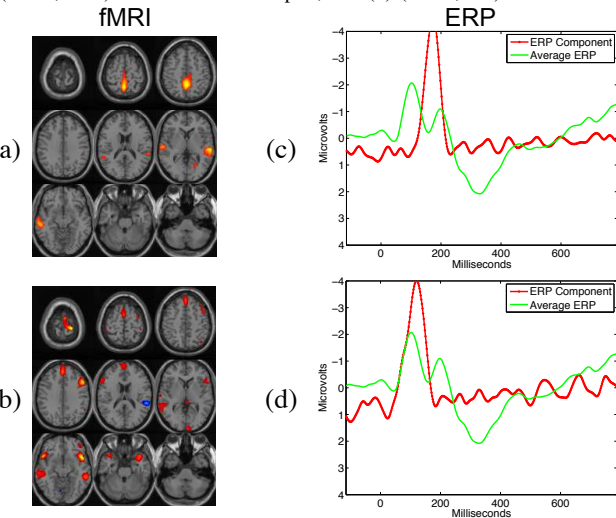

Fig. 4. Components generated for fusion of all three datasets using tIVA. The significance and average consistency values for the components estimated using MAXVAR are (.0069, 1.00), and (.0066, 1.00) for the top component and (.0093, 1.00) for the second one estimated using GENVAR. For IVA-G, we have a single component that shows significant group difference with (.013, .75).

across the modalities such that they jointly inform on the differences across patients with schizophrenia compared with the healthy controls. Hence, they link all three modalities that are essentially of different characteristics. In the figures that we show, red indicates areas where there is more activation or gray matter volume for controls, healthy participants, compared with patients as explained in Section III. All spatial maps are thresholded at a Z-score of 3.5, and as mentioned before, all p-values are not corrected for multiple comparisons unless specified otherwise.

jICA results: In Figure 1, we show the jICA results for Infomax and EBM, where we also give the values for the significance of the estimated components and their consistency. First, we note the similarity of the two components across algorithms, and then the higher significance of the components in terms of group difference generated by EBM compared to Infomax. The most significant component estimated by EBM even passes the threshold in terms of statistical significance using the conservative Bonferroni correction. This is further

supported by the fact that EBM resulted in only two significant components while for Infomax, there were two additional components that were not physically meaningful even though they were statistically significant. On the other hand, the consistency of EBM is lower compared to that of Infomax, which is expected as Infomax uses a fixed nonlinearity while EBM uses an adaptive approach for density matching. In Figure 1, the top component identifies changes in the motor execution region (anterior) associated with N2 (also similar to the N2/P3 complex), and the second component, differences in motor planning regions and motor (anterior to motor execution) as well as some thalamus associated with the N1 peak. These regions and the ERP deviation correspond to areas that have been also previously noted as affected in schizophrenia.

As we noted, we also performed pairwise fusion analyses. For the ERP component, use of three modalities rather than two yielded more physically interpretable results for jICA. We show an example in Figure 2 where the ERP component appears very similar for the two combinations, fMRI-ERP and sMRI-ERP, and corresponds to the top component shown for Infomax and EBM in Figure 1. In this example, we note that bringing in the third modality into the analysis provides
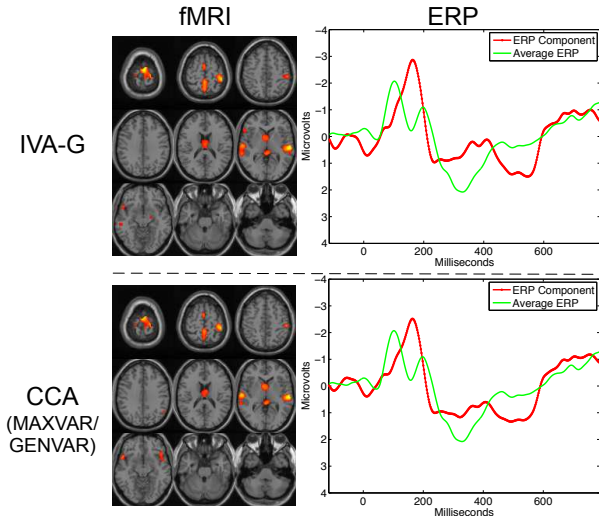
Fig. 5. Components generated using tIVA with two modalities at a time. Note that both MAXVAR and GENVAR reduce to CCA for two datasets. The significance and average consistency for components are given by (.0065, .92) for the component estimated using IVA-G and (.0066, 1) using CCA.

additional advantages, yields an ERP that is easier to interpret physically, and also yields more consistent differences in gray matter for sMRI. In addition, in Figure 3, we show four components estimated using EBM with a single modality at each time. The results of the single modality analysis hence can provide the guidance for the interpretation of the fusion results, and, as expected, similar components result in most cases. With multiple modalities, however, significance of the results increases and connections across modalities in terms of the observed group—healthy and patients—differences are established.

In Table I, we show the mutual information (MI) values of the estimated components for the fusion results as well as single modality results. The values are calculated for the best run, and after normalizing the estimated component using *Z*-scoring to allow for comparisons across different pairs. As expected, in most cases, EBM yields lower MI values as it uses an adaptive density estimation procedure and hence can better optimize the criterion, *i.e.*, can better maximize independence among the estimated components. It is also worth noting that the MI values are higher for the fusion results than those for a single modality. This is also expected as jICA treats the combination of multiple modalities as a single underlying source, which typically will have a more complicated and richer distribution as it is formed by samples from multiple modalities.

tIVA results: In Figure 4, we show the results of the fusion of three modalities using the tIVA model, with tMCCA-MAXVAR, tMCCA-GENVAR, and tIVA-G. First, note that since the profile matrices are not constrained to be the same for all modalities and each modality has a separate profile matrix, we can now estimate components that are not significant in all three modalities, and in the case of tIVA-G, the only component that is estimated is an fMRI-ERP pair. We see correspondences to the most significant—shown on

top for Infomax and EBM in Figure 1—component detected by jICA in these results. The component shown on top for GENVAR and the only significant component detected using IVA-G show similar difference in fMRI and an affected ERP component with characteristics similar to the N2/P3 complex. But in the case of tMCCA-GENVAR, the areas with functional change—areas of decreased activity and decreased gray matter in patients—are split among two components, one showing stronger association with sMRI. Since tMCCA-GENVAR and tIVA-G make use of the same cost function with difference in constraint for the demixing matrix, which is orthogonal for MCCA, it is not surprising that the results are more similar, and in fact the combined areas of change for fMRI for the two GENVAR components bear much resemblance to the fMRI result in the single component estimated using IVA-G, which also yields a better estimate of the N2/P3-like complex. The MAXVAR cost only takes the maximum eigenvalue into account and hence estimates only a single significant component that links all three modalities while the GENVAR/IVA-G cost, which takes all eigenvalues of the source component vector matrix into account, provides a more balanced representation of the underlying common components, *i.e.*, group differences, among the datasets. In addition, IVA-G is the least constrained approach among those and is the one that that yields a single component that is significant, and the sMRI difference is not part of this result since, as a modality, sMRI has the weakest connection to the other two modalities, as also noted in the simulation results in [1].

Order selection: In Section III, we noted that the suggested orders for the signal subspace using [51] for the jICA model are 14 for all three modalities, 9 for fMRI-ERP, 32 for fMRI-sMRI, and 32 for sMRI-ERP. In general, these values follow a similar trend to the number of significant components determined using the modality combinations for jICA, in that the combination of fMRI-sMRI and sMRI-ERP resulted in more significant components than using fMRI-ERP and all three modalities. For the tIVA model on the other hand, the joint order selection scheme of [53] indicated a similar number of joint components for the two combinations, fMRI-ERP and fMRI-sMRI as three and four, and suggested a smaller number, one, for the sMRI-ERP combination. This finding also correlates with the number of significant components detected using methods that use the tIVA model, where between 2 and 4 significant components were found for fMRI-ERP, between 4 and 5 significant components were found for fMRI-sMRI, and between 2 and 3 components were found using sMRI-ERP. The consistent detection of the smallest order with the sMRI and ERP combination is also expected due to the more significant difference in the nature of these two modalities. Hence, order selection, and especially that of common and modality specific subspaces, provides another promising way to perform exploratory analysis of multi-modal data.

V. DISCUSSION

In this paper, we study the application of the jICA and the tIVA models to fusion of data from two groups that

	fMRI-sMRI-ERP	fMRI-sMRI	fMRI-ERP	sMRI-ERP
Infomax	0.59 ± 0.05	0.43 ± 0.02	0.35 ± 0.15	0.56 ± 0.04
EBM	0.30 ± 0.06	0.42 ± 0.06	0.22 ± 0.09	0.26 ± 0.05

	fMRI	sMRI	ERP
Infomax	0.05 ± 0.01	0.05 ± 0.01	0.33 ± 0.05
EBM	0.03 ± 0.01	0.09 ± 0.03	0.28 ± 0.06

TABLE I

AVERAGE PAIRWISE MUTUAL INFORMATION AMONG THE ESTIMATED COMPONENTS FOR JICA AND SINGLE MODALITY ICA RESULTS

perform an AOD task. We performed both three-way fusion among the modalities as well as two-way and single modality decompositions using the two models, and studied their relationships. As demonstrated with numerous examples for fMRI [54], the maximization of independence is a useful tool for brain imaging as enables us to identify non-systematically overlapping components, and also separates the ERP components into interpretable waveforms in most cases [11], [28]. In a joint analysis, the estimated components identify regions that are co-varying together, *e.g.*, the functional and structural regions in the case of fMRI and sMRI, providing important information not available from either modality alone. Such an approach also does not require the regions for the different modalities to be in the same location, an important point as structural and functional information is not co-localized in many cases

We show that both with tIVA and jICA, fusion of all three modalities together leads to meaningful areas that show group differences and increased focality, and these areas are what would be expected to be implicated in schizophrenia. In the three-way analysis, the areas that demonstrate change in fMRI and sMRI are, in general, smaller and the ERP characteristics are easier to interpret. However, the link between the two functional modalities fMRI and ERP appears to be the strongest and the sensitivity for the components showing group difference, in general, decreased when sMRI was added to the analysis. Thus, even though using all the available modalities in such a multivariate analysis framework is attractive for multiple reasons, there may be penalties associated as well, which in this case, manifests as an overall decrease in the sensitivity of the detected group differences with sMRI. For the given dataset, since the number of samples is only 36, the relatively small sample size limits the performance of tIVA, and having a strong constraint as jICA imposes, seems to lead to results that are easier to interpret. In addition, the stronger relation between the ICA results from each dataset separately and the jICA results, leads to greater confidence in jICA results as a whole. Also, use of an algorithm like Infomax that emphasizes sparse distributions leads to more focal estimates for the spatial maps and narrower responses for the ERP. However, in terms of the physical interpretability of the components we estimate as well as their overall significance values, EBM leads to more desirable performance than the Infomax algorithm, which has been the algorithm of choice in almost all jICA implementations. When interpreting the significance values,

however, some caution might be required. As suggested by the simulation results given in [1], the statistical significance values might be over-estimated with jICA and under-estimated with tIVA if the connection with a given dataset is weak, and the estimation performance with tIVA would be expected to be more reliable using the tIVA approach, especially when there are sufficient number of samples.

When interpreting the results with real data, it is important to remember the main points demonstrated by the simulation results. When the structure of the connection across the modalities is strong, both models perform similarly, but tIVA outperforms jICA when the connection between any of the multiple modalities decreases, and the estimation of the component with the weaker link will be significantly impacted for jICA. On the other hand, jICA is more robust to noise if the underlying structure is strong, since it assumes that the modalities share the same mixing matrix, *i.e.*, profiles. This strong assumption enables jICA to average out the effects of noise. Based on the overall results, we would expect the component estimations for fMRI-EEG to be reliable for both models, possible more for jICA, whereas for fMRI-sMRI and sMRI-EEG, the results obtained by the tIVA model might be more reliable. In addition, judging the quality of results using only the t-values for group differences might be somewhat misleading. Through its strong assumptions, jICA might lead to increased sensitivity while its estimation performance in terms of recovery of the underlying components might be deteriorating as shown for the simulation example in Figure 6 in [1]. Number of samples is another important factor affecting the performance, especially for the tIVA model. Simulation results demonstrate that the tIVA model becomes more attractive when sample size—number of subjects—increases. For the example with the real data, since we only have 36 subjects, the situation favors the jICA model. It is also worth noting that when the sample size is low, the limitations of the models or algorithms we have noted due to the assumptions on independence or orthogonality are of less concern as such statistical properties are better satisfied for larger sample sizes.

Given such trade-offs in performance, it is natural to consider extensions of these models, and there are a number of them. These include the MCCA+jICA model [34], which performs MCCA using the tIVA model and then applies jICA to the reconstructed components post-MCCA. Since maximally independent components are desirable for interpretation but the assumption of common profiles across multiple datasets

might be too strong—especially when fusing more than two modalities—the goal here is to align the components that are *common* across the modalities prior to the jICA step. However, the additional MCCA step provides advantages when the link among modalities—as typically measured by correlation— is weak, and since in the setup considered for this paper we are primarily interested in components that are highly correlated across the modalities, for this dataset, the use of MCCA+jICA leads to identical results to those obtained using jICA alone. Another approach similar to [34] is used in [23] where the features are obtained using a multivariate latent variable preprocessing step for fusion of EEG and magnetoencephalography (MEG) data. Another approach, linked ICA [55] assumes a single mixing matrix for the modalities like jICA but estimates the weighting for each modality using a Bayesian approach. Finally, in parallel ICA [35], [56] multiple ICAs are performed while maximizing the correlation across profiles—selected columns of mixing matrices—through an alternating procedure between these two objectives. Reference [56] introduces the approach for two datasets with application to fusion of sMRI with genetic data and [35] extends it to three datasets, including fMRI in the analysis as the third modality. More comprehensive reviews of various data fusion methods based on blind source separation are given in [10], [57].

While extensions are always useful, with these two papers, the first one, [1], that discusses jICA and tIVA models in detail, and the current paper that examines their implementation for a specific dataset, our goal has been to fully understand the issues behind two fundamental BSS models for multi-modal data fusion and to provide guidance to the user. In the process, we noted issues such as the desirability of using a flexible ICA algorithm like EBM for the jICA model, the importance of order selection and its potential to enable exploratory analyses of multi-modal data, and the advantages of tIVA model over jICA when the connections among datasets are not all similar and when sufficient number of samples is available. Finally, we would like to note that diversity across and within datasets plays an important role in multi-modal data fusion enabling a fully multi-variate approach and BSS techniques like ICA and IVA are powerful solutions for the task.

References

[1] T. Adalı, Y. Levin-Schwartz, and V. D. Calhoun, "Multi-modal data fusion using source separation: Two effective models based on ICA and IVA and their properties," *Proc. IEEE*, vol. 103, no. 9, pp. 1478–1493, Sep. 2015.

[2] A. P. James and B. V. Dasarthy, "Medical image fusion: A survey of the state of the art," *Information Fusion*, vol. 19, pp. 4–19, 2014.

[3] V. D. Calhoun and T. Adalı, "Feature-based fusion of medical imaging data," *IEEE Trans. Info. Technology in Biomedicine*, vol. 13, no. 5, pp. 711–720, Sep. 2009.

[4] E. Acar, M. A. Rasmussen, F. Savorani, T. Naes, and R. B. a, "Understanding data fusion within the framework of coupled matrix and tensor factorizations," *Chemometrics and Intelligent Laboratory Systems*, vol. 129, pp. 53–63, 2013.

[5] I. V. Mechelen and A. K. Smilde, "A generic linked-mode decomposition model for data fusion," *Chemometrics and Intelligent Laboratory Systems*, vol. 104, pp. 83–94, 2010.

[6] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neural Computation*, vol. 15, no. 2, pp. 349–396, 2003.

[7] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing Systems*, vol. 13, 2000.

[8] M. Remazani, K. Marble, H. Trang, I. S. Johnsrude, and P. Abolmaesumi, "Joint sparse representation of brain activity patterns in multi-task fMRI data," *IEEE Trans. Med. Imaging*, vol. 34, no. 1, pp. 2–12, Jan. 2015.

[9] A. Anderson, P. K. Douglas, W. T. Kerr, V. S. Haynes, A. L. Yuille, J. Xie, Y. N. Wu, J. A. Brown, and M. S. Cohen, "Non-negative matrix factorization of multimodal MRI, fMRI and phenotypic data reveals differential changes in default mode subnetworks in ADHD," *NeuroImage*, vol. 102, no. 1, pp. 207–219, 2014.

[10] J. Sui, T. Adalı, Q. Yu, J. Chen, and V. Calhoun, "A review of multivariate methods for multimodal fusion of brain imaging data," *Journal of Neuroscience Methods*, vol. 204, pp. 68–81, 2012.

[11] V. D. Calhoun, T. Adalı, G. D. Pearlson, and K. A. Kiehl, "Neuronal chronometry of target detection: Fusion of hemodynamic and event-related potential data," *NeuroImage*, vol. 30, no. 2, pp. 544–553, April 2006.

[12] Fusion ICA toolbox: http://mialab.mrn.org/software/fit/. [Online]. Available: http://mialab.mrn.org/software/fit/

[13] T. Eichele, V. D. Calhoun, M. Moosmann, K. Specht, M. L. Jongsma, R. Q. Quiroga, H. Nordby, and K. Hugdahl, "Unmixing concurrent EEG-fMRI with parallel independent component analysis," *International Journal of Psychophysiology*, vol. 67, no. 3, pp. 222–234, 2008.

[14] B. Mijovic, K. Vanderperren, N. Novitskiy, B. Vanrumste, P. Stiers, B. V. Bergh, L. Lagae, S. Sunaert, J. Wagemans, S. V. Huffel, and M. D. Vos, "The "why" and "how" of JointICA: Results from a visual detection task," *NeuroImage*, vol. 60, pp. 1171–1185, 2012.

[15] B. Mijovic, M. D. Vos, K. Vanderperren, B. Machilsen, S. Sunaert, S. V. Huffel, and J. Wagemans, "The dynamics of contour integration: A simultaneous EEG–fMRI study," *NeuroImage*, vol. 88, pp. 10–21, 2014.

[16] J. Mangalathu-Arumana, S. Beardsley, and E. Liebenthal, "Within-subject joint independent component analysis of simultaneous fMRI/ERP in an auditory oddball paradigm," *NeuroImage*, vol. 60, pp. 2247–2257, 2012.

[17] W. Swinnen, B. Hunyadi, E. Acar, S. V. Huffel, and M. D. Vos, "Incorporating higher dimensionality in joint decomposition of eeg and fmri incorporating higher dimensionality in joint decomposition of eeg and fmri incorporating higher dimensionality in joint decomposition of eeg and fmri," in *Proc. European Signal Process. Conf. (EUSIPCO)*, Lisbon, Portugal, September 2014.

[18] V. D. Calhoun, J. Liu, and T. Adalı, "A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data," *NeuroImage*, vol. 45, no. 1, pp. S163–S172, March 2009.

[19] V. D. Calhoun, T. Adalı, N. R. Giuliani, J. J. Pekar, K. A. Kiehl, and G. D. Pearlson, "Method for multimodal analysis of independent source differences in schizophrenia: combining gray matter structural and auditory oddball functional data," *Human Brain Mapping*, vol. 27, pp. 47–62, 2006.

[20] K. Magtibay, M. Beheshti, F. H. Foomany, K. Balasundaram, P. L. S. Masse, J. Asta, N. Zamiri, D. A. Jaffray, K. Nanthakumar, S. Krishnan, , and K. Umapathy, "Fusion of structural and functional cardiac magnetic resonance imaging data for studying ventricular fibrillation," in *Proc. IEEE Int. Conf. Eng. in Medicine and Biology Spciety*, 2014, pp. 5579–5582.

[21] Z. Yuan and J.-C. Ye, "Fusion of fNIRS and fMRI data: identifying when and where hemodynamic signals are changing in human brains," *Frontiers in Human Neuroscience*, Oct. 2013.

[22] J. Sui, G. Pearlson, A. Caprihan, T. Adalı, K. . A. Kiehl, J. Liu, J. Yamamoto, and V. D. Calhoun, "Discriminating schizophrenia and bipolar disorder by fusing fMRI and DTI in a multimodal CCA+ joint ICA model," *NeuroImage*, vol. 57, no. 3, pp. 839–855, 2011.

[23] X. Chen, Z. J. Wang, and M. J. McKeown, "A three-step multimodal analysis framework for modeling corticomuscular activity with application to Parkinson's disease," *IEEE J. Biomedical and Health Informatics*, vol. 18, no. 4, pp. 1232–1241, July 2014.

[24] K. Specht, R. Zahn, K. Willmes, S. Weis, C. Holtel, B. J. Krause, H. Herzog, and W. Huber, "Joint independent component analysis of structural and functional images reveals complex patterns of functional reorganisation in stroke aphasia," *NeuroImage*, vol. 47, pp. 2057–2063, 2009.

[25] H. Hotelling, "Relations between two sets of variates," *Biometrika*, vol. 28, no. 3/4, pp. 321–377, Dec. 1936.

[26] P. Geladi and B. R. Kowalski, "Partial least-squares regression: A tutorial," *Analytica Chimica Acta*, vol. 185, pp. 1–17, 1986.

[27] N. M. Correa, Y.-O. Li, T. Adalı, and V. D. Calhoun, "Canonical correlation analysis for feature-based fusion of biomedical imaging modalities and its application to detection of associative networks in schizophrenia," *IEEE J. Selected Topics in Signal Processing*, vol. 2, no. 6, pp. 998–1007, 2009.

[28] N. Correa, T. Adalı, and V. D. Calhoun, "Canonical correlation analysis for data fusion and group inferences: Examining applications of medical imaging data," *IEEE Signal Processing Magazine*, vol. 27, no. 4, pp. 39–50, July 2010.

[29] A. Krishnan, L. J. Williamd, A. R. McIntosh, and H. Abdi, "Partial least squares (pls) methods for neuroimaging: A tutorial and review," *NeuroImage*, vol. 56, pp. 455–475, 2011.

[30] E. Martinez-Montes, P. A. Valdés-Sosa, A. Miwakeichi, R. I. Goldman, and M. S. Cohen, "Concurrent EEG/fMRI analysis by multiway partial least squares," *NeuroImage*, vol. 22, pp. 1023–1034, 2004.

[31] J. R. Kettenring, "Canonical analysis of several sets of variables," *Biometrika*, vol. 58, no. 3, pp. 433–451, Dec. 1971.

[32] N. M. Correa, Y.-O. Li, T. Adalı, and V. D. Calhoun, "Fusion of fMRI, sMRI, and EEG data using canonical correlation analysis," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing (ICASSP)*, 2009.

[33] N. M. Correa, T. Eichele, T. Adalı, Y.-O. Li, and V. D. Calhoun, "Multi-set canonical correlation analysis for the fusion of concurrent single trial erp and functional mri," *NeuroImage*, vol. 50, no. 4, pp. 1438–1445, May 2010.

[34] J. Sui, H. He, G. D. Pearlson, T. Adalı, K. A. Kiehl, Q. Yu, V. P. Clark, E. Castro, T. White, B. A. Mueller, B. C. Ho, N. C. Andreasen, and V. D. Calhoun, "Three-way (N-way) fusion of brain imaging data based on mCCA+jICA and its application to discriminating schizophrenia," *NeuroImage*, vol. 66, pp. 119–132, Feb. 2013.

[35] V. M. Vergara, A. Ulloa, V. D. Calhoun, D. Boutte, J. Chen, and J. Liu, "A three-way parallel ICA approach to analyze links among genetics, brain structure and brain function," *NeuroImage*, vol. 98, pp. 386–394, Sep. 2014.

[36] J. Sui, T. Adalı, G. Pearlson, H. Yang, S. R. Sponheim, T. White, and V. D. Calhoun, "A CCA + ICA based model for multi-task brain imaging data fusion and its application to schizophrenia," *NeuroImage*, vol. 51, no. 1, pp. 123–134, May 2010.

[37] X.-L. Li and T. Adalı, "Independent component analysis by entropy bound minimization," *IEEE Trans. Signal Processing*, vol. 58, no. 10, pp. 5151–5164, Oct. 2010.

[38] A. Bell and T. Sejnowski, "An information maximization approach to blind separation and blind deconvolution," *Neural Computation*, vol. 7, no. 6, pp. 1129–1159, Nov. 1995.

[39] K. A. Kiehl, M. Stevens, K. R. Laurens, G. D. Pearlson, V. D. Calhoun, and P. F. Liddle, "An adaptive reflexive processing model of neurocognitive function: Supporting evidence from a large scale (n=100) fMRI study of an auditory oddball task," *NeuroImage*, vol. 25, pp. 899–915, 2005.

[40] K. Friston, "SPM8," http://www.fil.ion.ucl.ac.uk/spm/software/spm8/, April 2009.

[41] M. S. Cohen, R. M. DuBois, and M. M. Zeineh, "Rapid and effective correction of RF inhomogeneity for high field magnetic resonance imaging," *Human Brain Mapping*, vol. 10, pp. 204–211, 2000.

[42] G. Gerig, J. Martin, R. Kikinis, O. Kubler, M. E. Shenton, and F. Jolesz, "Unsupervised tissue type segmentation of 3D dual-echo MR head data," *Image and Vision Computing*, vol. 10, no. 6, 1992.

[43] Y. Zhang, M. Brady, and S. Smith, "Segmentation of brain mr images through a hidden markov random field model and the expectation-maximization algorithm," *IEEE Trans. Med. Imaging*, vol. 20, pp. 45–57, 2001.

[44] I. W. M. Wells, W. E. L. Grimson, R. Kikinis, and F. A. Jolesz, "Adaptive segmentation of MRI data," *IEEE Trans. Med. Imaging*, vol. 15, no. 4, pp. 429–442, 1996.

[45] J. C. Bezdek, L. O. Hall, and L. P. Clarke, "Review of MR image segmentation techniques using pattern recognition," *Med. Physics*, vol. 20, pp. 1033–1048, 1993.

[46] M. Anderson, X.-L. Li, and T. Adalı, "Joint blind source separation with multivariate Gaussian model: Algorithms and performance analysis," *IEEE Trans. Signal Processing*, vol. 60, no. 4, pp. 2049–2055, April 2012.

[47] Himberg, A. Hyvärinen, and A. Esposito, "Validating the independent components of neuroimaging time-series via clustering and visualization," *NeuroImage*, vol. 22, pp. 1214–1222, 2004.

[48] S. Ma, N. Correa, X.-L. Li, T. Eichele, V. Calhoun, and T. Adalı, "Automatic identification of functional clusters in fMRI data using spatial dependence," *IEEE Trans. Biomedical Eng.*, vol. 58, no. 12, pp. 3406–3417, 2011.

[49] W. Du, S. Ma, G.-S. Fu, V. D. Calhoun, and T. Adalı, "A novel approach for assessing reliability of ICA for fMRI analysis," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Florence, Italy, May 2014.

[50] Group ICA of fMRI toolbox: http://mialab.mrn.org/software/gift/. [Online]. Available: http://mialab.mrn.org/software/gift

[51] Y.-O. Li, T. Adalı, and V. D. Calhoun, "Estimating the number of independent components for fMRI data," *Human Brain Mapping*, vol. 28, no. 11, pp. 1251–1266, Nov. 2007.

[52] M. Ramezani, P. Abolmaesumi, K. Marble, H. Trang, and I. Johnsrude, "Fusion analysis of functional mri data for classification of individuals based on patterns of activation," *Brain Imaging and Behavior*, pp. 1–13, 2014.

[53] Y. Song, P. Schreier, and N. J. Roseveare, "Determining the number of correlated signals between two data sets using PCA-CCA when sample support is extremely small," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Brisbane, Australia, April 2015.

[54] V. D. Calhoun and T. Adalı, "Multisubject independent component analysis of fMRI: A decade of intrinsic networks, default mode, and neurodiagnostic discovery," *IEEE Reviews in Biomedical Engineering*, vol. 5, pp. 60–73, Aug. 2012.

[55] A. R. Groves, C. F. Beckmann, S. M. Smith, and M. W. Woolrich, "Linked independent component analysis for multimodal data fusion," *NeuroImage*, vol. 54, no. 3, pp. 2198–2217, 2011. [Online]. Available: http://dx.doi.org/10.1016/j.neuroimage.2010.09.073

[56] J. Chen, V. D. Calhoun, G. D. Pearlson, N. Perrone-Bizzozero, J. Sui, J. A. Turner, J. R. Bustillo, S. Ehrlich, S. R. Sponheim, J. M. Canive, B.-C. Ho, and J. Liu, "Guided exploration of genomic risk for gray matter abnormalities in schizophrenia using parallel independent component analysis with reference," *NeuroImage*, vol. 83, pp. 384–396, Dec. 2013.

[57] D. Lahat, T. Adalı, and C. Jutten, "Multimodal data fusion: A methodological overview methods, challenges and prospectives," *Proc. IEEE*, vol. 103, no. 9, pp. 1449–1477, Sep. 2015.

Tülay Adalı (S'89–M'93–SM'98–F'09) received the Ph.D. degree in Electrical Engineering from North Carolina State University, Raleigh, NC, USA, in 1992 and joined the faculty at the University of Maryland Baltimore County (UMBC), Baltimore, MD, USA, the same year. She is currently a Distinguished University Professor in the Department of Computer Science and Electrical Engineering at UMBC. Prof. Adalı assisted in the organization of a number of international conferences and workshops including the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), the IEEE International Workshop on Neural Networks for Signal Processing (NNSP), and the IEEE International Workshop on Machine Learning for Signal Processing (MLSP). She was the General Co-Chair, NNSP (2001–2003); Technical Chair, MLSP (2004–2008); Program Co-Chair, MLSP (2008, 2009, and 2014), 2009 International Conference on Independent Component Analysis and Source Separation; Publicity Chair, ICASSP (2000 and 2005); and Publications Co-Chair, ICASSP 2008. She is Technical Program Co-Chair for ICASSP 2017 And Special Sessions Co-Chair for ICASSP 2018.

Prof. Adalı chaired the IEEE Signal Processing Society (SPS) MLSP Technical Committee (2003–2005, 2011–2013), served on the SPS Conference Board (1998–2006), and the Bio Imaging and Signal Processing Technical Committee (2004–2007). She was an Associate Editor for IEEE Transactions on Signal Processing (2003–2006), IEEE Transactions on Biomedical Engineering (2007–2013), IEEE Journal of Selected Areas in Signal Processing (2010-2013), and Elsevier Signal Processing Journal (2007–2010). She is currently serving on the Editorial Boards of the Proceedings of the IEEE and Journal of Signal Processing Systems for Signal, Image, and Video Technology, and is a member of the IEEE Signal Processing Theory and Methods Technical Committee.

Prof. Adalı is a Fellow of the IEEE and the AIMBE, a Fulbright Scholar, recipient of a 2010 IEEE Signal Processing Society Best Paper Award, 2013 University System of Maryland Regents' Award for Research, and an NSF CAREER Award. She was an IEEE Signal Processing Society Distinguished Lecturer for 2012 and 2013. Her research interests are in the areas of statistical signal processing, machine learning for signal processing, and biomedical data analysis.

PUBLISHED: PROC. IEEE, VOL. 103, NO. 9, PP. 1494–1506, SEP. 2015.

10

Yuri Levin-Schwartz received the B.S. degree in physics and the B.A. degree in mathematics, both from Brandeis University, Waltham, Massachusetts, in 2012. He is currently a Ph.D. candidate at the Machine Learning for Signal Processing Lab, University of Maryland, Baltimore County. His research interests include blind source separation, machine learning, statistical signal processing, and biomedical data analysis.

Vince D. Calhoun (S'91–M'02–SM'06–'F13) received a bachelor's degree in Electrical Engineering from the University of Kansas, Lawrence, Kansas, in 1991, master's degrees in Biomedical Engineering and Information Systems from Johns Hopkins University, Baltimore, in 1993 and 1996, respectively, and the Ph.D. degree in electrical engineering from the University of Maryland Baltimore County, Baltimore, in 2002. He worked as a research engineer in the psychiatric neuroimaging laboratory at Johns Hopkins from 1993 until 2002. He then served as the director of medical image analysis at the Olin Neuropsychiatry Research Center and as an associate professor at Yale University.

Dr. Calhoun is currently Executive Science Officer and Director of Image Analysis and MR Research at the Mind Research Network and is a Distinguished Professor in the Departments of Electrical and Computer Engineering (primary), Biology, Computer Science, Neurosciences, and Psychiatry at the University of New Mexico. He is the author of more than 400 full journal articles and over 500 technical reports, abstracts and conference proceedings. Much of his career has been spent on the development of data driven approaches for the analysis of brain imaging data. He has won over $85 million in NSF and NIH grants on the incorporation of prior information into ICA for functional magnetic resonance imaging, data fusion of multimodal imaging and genetics data, and the identification of biomarkers for disease, and leads a P20 COBRE center grant on multimodal imaging of schizophrenia, bipolar disorder, and major depression.

Dr. Calhoun is a fellow of the IEEE, The Association for the Advancement of Science, The American Institute of Biomedical and Medical Engineers, and the International Society of Magnetic Resonance in Medicine. He is also a member and regularly attends the Organization for Human Brain Mapping, the International Society for Magnetic Resonance in Medicine, the International Congress on Schizophrenia Research, and the American College of Neuropsychopharmacology. He is also a regular grant reviewer for NIH and NSF. He has organized workshops and special sessions at multiple conferences. He is currently chair of the IEEE Machine Learning for Signal Processing (MLSP) technical committee. He is a reviewer for many journals is on the editorial board of the Brain Connectivity and Neuroimage journals and serves as Associate Editor for Journal of Neuroscience Methods and several other journals.