

Repeated decompositions reveal the stability of infomax decomposition of fMRI data

Jeng-Ren Duann¹, Tzyy-Ping Jung¹, Terrence J. Sejnowski^{1,2}, and Scott Makeig¹

¹Institute for Neural Computation,
University of California San Diego 0961,
La Jolla CA 92093-0961, USA

²Computational Neurobiology Lab.,
The Salk Institute for Biological Studies,
La Jolla CA 92037, USA

Abstract – In this study, we decomposed 12 fMRI data sets from six subjects each 101 times using the infomax algorithm. The first decomposition was taken as a reference decomposition; the others were used to form a component matrix of 100 by 100 components. Equivalence relations between components in this matrix, defined as maximum spatial correlations to the components of the reference decomposition, were found by the Hungarian sorting method and used to form 100 equivalence classes for each data set. We then tested the reproducibility of the matched components in the equivalence classes using uncertainty measures based on component distributions, time courses, and ROC curves. Infomax ICA rarely failed to derive nearly the same components in different decompositions. Very few components per data set were poorly reproduced, even using vector angle uncertainty measures stricter than correlation and detection theory measures.

I. INTRODUCTION

Independent component analysis (ICA) has been successfully applied in a variety of applications including biomedical time series (EEG/MEG, ECG, etc.) and images (fMRI, PET, and SPECT) [1]. For example, Makeig et al. [2] applied ICA to separate the multi-channel EEG and found that the basis of the event-related potential, averaged across similar experimental trials, can typically be dissected into several source components accounting for temporally and functionally distinct EEG processes. This method can better explain the variability of event-related EEG results and can give more complete and consistent results, within and between subjects, than standard methods that extract data features at individual scalp channels [3].

ICA is a relatively model-free blind source separation method that attempts to separate the underlying source contributions to the data without knowing in advance what the sources are or how they are mixed. Thus, there are no *a priori* suitable statistical models to evaluate the goodness-of-fit of the resultant independent components. For both EEG and fMRI data, the ‘true’ nature of the signal sources is not known, making it difficult to evaluate the performance of ICA precisely.

To study the stability of ICA decomposition of fMRI data, McKeown et al. [4][5] decomposed one fMRI data set twice, first with random initial weights and then again starting with the weights learned in the first decomposition. They then compared the components from the two decompositions using mutual information, selecting a mutual information threshold for detecting highly reproducible components, and suggested that these components might be of interest for further analysis. In a subsequent study,

Duann et al. [6] found that infomax ICA could consistently derive very similar components by matching independent components obtained in 10 different decompositions using a Hungarian sorting algorithm that maximizes the sum of the positive correlation coefficients of the matched components. In contrast to McKeown’s findings, matching components by mutual information did not give reasonable pair assignments. The rapid drop off in pairwise mutual information they reported was mainly due to poor identification of corresponding components from the different decompositions. On the other hand, Meinecke et al. [7][8] applied a resampling approach, first randomly selecting independent components from an ICA decomposition to create surrogate data, and then applying ICA to the surrogate data. This process was repeated several times to test whether a given component was found consistently.

In this study, we used infomax ICA to decompose 12 fMRI data sets collected from six subjects participating in a passive visual perception experiment involving infrequent 8-Hz flickering checkerboard pattern stimuli, then employed quantitative approaches to assess the reproducibility of the ICA decompositions. In the following sections, we first describe the fMRI data, how we performed the decompositions, and how we matched the components across different decompositions. We then define several similarity measures including ROC curve and uncertainty measures applied to both time courses and ROA maps of components obtained from repeated decompositions. We then present the resulting component equivalence sets. Finally, we discuss the implication of our results for fMRI data analysis.

II. MATERIAL AND METHODS

A. fMRI Data Sets

Twelve fMRI data sets, two sets each from six subjects were used for multiple ICA decompositions. The subjects passively viewed reversing checkerboard patterns presented for either 500 ms or 3000 ms every 30 s (see Duann et al. [9]). The BOLD data consisted of five 64 x 64 T2*-weighted brain slice images covering mainly the primary visual cortex (parallel to the calcarine sulcus) collected with an inter-scan interval (TR) of 0.5 s on a 3-T MR research scanner at the Taipei Veterans General Hospital, Taiwan. The study was approved by the local ethics committee of the hospital and all subjects provided written consent.

B. Infomax ICA

As noted by both McKeown et al. [10] and Duann et al. [9], the objective of applying ICA to fMRI data is to separate fMRI signals into spatially independent processes with distinct (though not necessarily independent) time courses each accounting for fluctuations in blood oxygenation during image acquisition produced by a single hypothesized source or process. ICA does not attempt to model the natures of the contributing sources before performing the separation – for this reason it is defined to be a

'blind source separation' method.

ICA derives an unmixing matrix, W , iteratively, here using the infomax ICA algorithm [11] in its natural gradient formulation [12] and optionally using the extended option of [13] that can also identify sub-Gaussian (e.g., spatially extensive) source distributions. The estimated source activations, U , can then be calculated as

$$U = W \times X$$

where X is the $T \times N$ matrix of observed time series data, T is the total number of time points, and N is the number of voxels submitted to ICA decomposition. U , the estimated source activation matrix, has the same dimension as input data matrix, X . W is a $T \times T$ ICA unmixing matrix that recovers the source activities from the original recordings. Its inverse, W^{-1} (the mixing matrix), gives the relative weight of the sources at each time point, thus specifying how the source activities contribute to the recorded data. When necessary to avoid numerical instability, principal component analysis (PCA) may be used *prior* to ICA decomposition to reduce the dimensionality (T) of the input time series.

PCA dimension reduction allows 'non-square' unmixing, in which fewer independent components than the number of time points in the input training data are extracted. PCA is of use when the number of acquired time points is large and the number of image slices (and voxels) is relatively small. In such cases, ICA decomposition performed without dimensionality reduction might not converge, since ICA needs a sufficient number of voxels (typically » $T \times T$) to learn the $T \times T$ weights of the unmixing matrix from the input data.

C. Image Analysis

Before image preprocessing, the first 10 ('dummy') scans of each fMRI data were first removed from the data. The data were then subjected to a slice-timing adjustment to minimize inhomogeneities arising from acquisition of the five image planes at slightly different time. We then removed voxels outside the brain volume to reduce the data size for ICA training and also to eliminate machine noise arising from susceptibility differences between air and brain tissues. The data were then transposed to form a matrix with dimension $T \times V$, where T is the number of scans ($T=600$ in this study) and V , the number of remaining voxels after preprocessing ($V \sim 4000$). The transposed matrix was then subjected to infomax ICA decomposition [14], producing an unmixing matrix that linearly separates the input fMRI time series into maximally spatially independent components (or source processes). Multiplying the data by this matrix produced the estimated component maps; inverting the unmixing matrix gave the component time courses. Since the number of voxels was not sufficient to derive a full (600×600) unmixing matrix, the dimensionality of the data was reduced to 100 by PCA *prior* to ICA training. The surrogate data were made by back-projecting the first 100 principal components to the original data space instead of using the principal components themselves. A unity matrix was used to initialize the weights of the unmixing matrix. The number of training steps was increased to 3000 to assure the full convergence of ICA training. For further detail, see Duann et al. [9].

D. Multiple ICA Decompositions

To evaluate the stability and consistency of ICA decomposition, each of the preprocessed fMRI data sets was decomposed by infomax ICA 101 times using the same initial conditions and training parameters. Variability in the output of the infomax algorithm [14] is induced by the random shuffling of the training data (i.e. voxel) order before each training step. Next, the activation matrices containing the component spatial maps were obtained by

$$\mathbf{u}^i = \mathbf{W}^i \times \mathbf{x} \quad (1)$$

where \mathbf{u}^i are the 'activation' map weights of the i^{th} decomposition, \mathbf{W}^i is the inverse matrix of the i^{th} decomposition of the training data set \mathbf{x} . These activation map distributions were transformed to z values by

$$z_{j,k}^i = \frac{u_{j,k}^i - \frac{1}{K} \sum_k u_{j,k}^i}{\sqrt{\sum_k \left(u_{j,k}^i - \frac{1}{K} \sum_k u_{j,k}^i \right)^2}} \quad (2)$$

where i is the decomposition index, j the component index, and k the voxel index ($k = 1, 2, \dots, K$). $u_{j,k}^i$ and $z_{j,k}^i$ are the activation weights (or simply activations) and the z value of the k^{th} voxel contributed by the j^{th} component in the i^{th} decomposition. Assessments of how well components were matched across decompositions were performed on the z -transformed component maps with threshold 2.0 (except for the ROC curve measurements (see below) whereby the z -value threshold was subject to change).

E. Component Matching

Our first question was whether or not each component could be reliably found in the repeated decompositions. To this end, we had to find the best-matching component pairs for each pair of decompositions. Here, we used as a cost function the sum of positive cross correlation coefficients between components, which we found to yield the most satisfactory matching component pairs [6]. To match components, we used the first decomposition from each data set as a reference decomposition and then referred the other 100 decompositions to this reference, forming 100 decomposition pairs. For each pair of decompositions by Hungarian method, best matching component pairs were assigned by selecting without substitution the most similar pair of unselected components. Cross correlation coefficients between best matching component pairs proved to be quite high. Thus, the 100 (decompositions) \times 100 (components) were essentially separated into 100 component equivalence classes, each consisting of 100 equivalent components from each of the 100 decompositions. After the first component matching, the components of the reference decomposition were resorted according to the cross correlation coefficients. Then, the components of all other decompositions were also resorted according to the order of components in the reference decompositions after best matching assigned.

F. Evaluating the Stability of ICA Decomposition

To assess the stability of ICA decomposition, we used four different measures: (1) the cross correlation coefficient between component pair maps, (2) uncertainty in the component maps, (3) uncertainty in the component time courses, and (4) the area under the ROC curve.

To evaluate the strength of component pair cross correlations, we simply plotted the mean correlation coefficients across decomposition for each component pair with error bars representing their standard deviations.

Component uncertainty was calculated as the vector angle between the maps or time courses of equivalent component pairs:

$$U^i = \cos^{-1} \left(\frac{\mathbf{A}^i \cdot \mathbf{A}^0}{|\mathbf{A}^i| \cdot |\mathbf{A}^0|} \right) \quad (3)$$

Above, U^i is the uncertainty measure of the i^{th} component, \mathbf{A} is a vector representing the component spatial map or time course (0 indicates the reference decomposition), and $\mathbf{A} \cdot \mathbf{B}$ is the dot-product between vectors \mathbf{A} and \mathbf{B} . Since the component pairs were derived by the component matching process, the uncertainty measures were applied to the matched component pairs only. The statistical distribution of the uncertainty measures was estimated by the multiple decompositions. The mean and standard deviation of the uncertainty measures was plotted. The more reproducible a component, the lower the uncertainty measure. However, the uncertainty measures based on component spatial maps and component time courses may be different.

To compute ROC curves, we used the component maps of the first decomposition as the reference then varied the z -value threshold to construct component maps corresponding to the other

100 decompositions. If a voxel defined as active in one decomposition was also selected as active in the reference decomposition, we marked this voxel as a true positive. On the contrary, if a voxel was defined as active in the same decomposition but was marked as non-active in the reference decomposition, this voxel would be marked as false positive. The numbers of true positive voxels and false positive voxels were then divided by the total number of active voxels and non-active voxels in the reference decomposition to form true-positive and false-positive rates. Consequently, the ROC curve was plotted as the scatter plot of these two rates as a function of the z-value threshold. The area under this ROC curve was then calculated for each component. For highly reproducible components, this area should approach 1.

III. RESULTS

Fig. 1 shows the average correlation coefficients of matched component pairings between the reference decomposition and all other 100 decompositions. Error bars give the standard deviation of the correlation coefficients across the 100 decompositions. Figs. 2 and 3 show the uncertainty measures based on the angles between the vectors of component maps and time courses, respectively. These measures may be considered more stringent measures of component similarity than the correlation coefficients shown in Fig. 1. Fig. 4 shows the average area under the ROC curves of each equivalence component or equivalence class across 100 decompositions (the error bar indicating their standard deviation).




Fig. 1 Average correlation coefficients between component spatial maps of matched components across decompositions. Error bars show the standard deviations of the correlation coefficients.




Fig. 2 Mean uncertainty in component maps for equivalent components in 100 decompositions. Components are ranked by their contribution to the data, from largest (1) to smallest (100).




Fig. 3 Mean uncertainty in component time course for equivalent components in 100 decompositions. Component ranking as in Fig. 2.




Fig. 4 Mean area under ROC curves of each component across 100 decompositions. The bars indicate the areas, averaged across decompositions, under the ROC curve of components. Error bars displayed on top show represent the standard deviations.

IV. DISCUSSION

Correctly determining best-matching component pairs is the key to the evaluation of the stability of ICA decompositions using either resampled [7] [8] or repeated decompositions [5] [6]. Here, we have shown that component matching by maximizing the sum of positive correlation coefficients over component pairs gives satisfactory results. Within equivalence classes, the best-matched components show extremely high correlations of their component maps (> 0.85 for all data sets). There were always more than 90 (of 100) component classes having mean pairwise correlation coefficients > 0.95 (with very small standard deviations). Our results further confirm the preliminary results of Duann et al. [6] that finding matching components by maximizing the summed positive correlations across component pairs from a decomposition provides more accurate component matching than using mutual information.

Although correlation coefficients can provide accurate component pair assignments, they are not sensitive enough to match component maps that contain many small groups of active voxels. Any spatially overlapping voxels will give high correlation coefficients without necessarily contributing to overall component reproducibility. This is the reason why the last few components, even if they do not seem to have similar component maps, still have high correlation coefficients, as shown in Fig. 1. Here we also used two uncertainty measures, the vector angles between component

maps and time courses. These uncertainty measures can be considered to be more sensitive than the correlation coefficient for detecting matching components. As shown in Figs. 2 and 3, these uncertainty measures revealed more differences between matched components within the same equivalence classes than did the correlation measure (Fig. 1).

It is worth noting that the component map angles suggested a lower uncertainty (range: from 0 to 0.45 rad) compared to the time course angles (range: from 0 to 0.9 rad). That is, the component maps were more consistent than the time courses. However, for the largest 50 components, the multiple ICA decompositions gave almost identical time courses (uncertainty $< \sim 0.1$ rad in Figs. 2 and 3). These uncertainty measures can be used not only to measure the consistency of different ICA decompositions but also to help select components of interest.

The result of ROC curve analysis also reveals that ICA decomposition can stably decompose fMRI data set into highly reproducible independent components. For all components, the average areas under the ROC curves were all close to 1. This means the 100 decompositions almost perfectly reproduced the same components with the same brain areas as highlighted in the reference decomposition.

This result was dependent on the choice of the reference decomposition and the z-value threshold. Preliminary testing using different z-value thresholds to construct the reference component ROA maps showed that z-value threshold of 2.0 gave the best ROC curve area for all the decompositions of one data set (data not shown). Thus, we employed z-value threshold of 2.0 for computing the component maps in this study.

We have recently developed an ICA component browser application, FMRLAB, that allows for computing and browsing through properties of the independent components extracted from any kind of functional neuroimaging data. It also allows experimenters to select components of interest for further analysis or of no interest for rejection. FMRLAB runs on several platforms under the MATLAB (The MathWorks, Inc., MA) environment, and is freely available for non-commercial use via our web site <http://www.sccn.ucsd.edu/fmrlab> [15].

V. CONCLUSION

Assessing the consistency or stability of ICA decomposition of fMRI data is an important issue for interpreting its results. Here, we performed multiple decompositions of the same fMRI data sets into spatially independent components to measure the stability of the infomax ICA algorithm in the form used by our group and available via our website [2]. For these data, infomax ICA never failed to converge to a stable solution. Moreover, measures of component maps and time courses showed that the repeated ICA decompositions were extremely stable. However, these promising results need to be further tested using fMRI data sets from experimental paradigms involving e.g. more complex cognitive processes.

ACKNOWLEDGEMENT

This work was supported by the National Institute of Health USA, the Howard Hughes Medical Institute, and the Swartz Foundation (Old Field, NY)

REFERENCES

- [1] T.-P. Jung, S. Makeig, M. J. McKeown, A. J. Bell, T.-W. Lee, and T. J. Sejnowski, "Imaging brain dynamics using independent component analysis," *IEEE Proc.*, vol. 89, pp. 1107-1122, July 2001.
- [2] S. Makeig, T.-P. Jung, A. J. Bell, D. Ghahremani, and T. J. Sejnowski, "Blind separation of auditory event-related brain responses into independent components," *Proc. Natl. Acad. Sci. USA*, vol. 94, pp. 10979-10984, September 1997.
- [3] S. Makeig, M. Westerfield, T.-P. Jung, S. Enghoff, J. Twonson, E. Courchesne, and T. J. Sejnowski, "Dynamic brain sources of

- visual evoked responses," *Science*, vol. 295, pp. 690-694, January 2002.
- [4] M. J. McKeown, V. Varadarajan, S. Huettel, and G. McCarthy, "Deterministic and stochastic features of fMRI data: implications for analysis of event-related experiments," *J. Neurosci. Methods*, vol. 118, pp. 103-113, 2002.
- [5] M. J. McKeown, "Deterministic and stochastic features of fMRI data: Implications for data averaging," in *Exploratory Analysis and Data Modeling in Functional Neuroimaging*, F. T. Sommer and A. Wichert Eds. Cambridge: MIT Press, 2002, pp. 63.
- [6] J.-R. Duann, T.-P. Jung, S. Makeig, and T. J. Sejnowski, "Consistency of Infomax ICA decomposition of functional brain imaging data," pp. 289-294, April 2003 [4th International Symposium on Independent Component Analysis and Blind Signal Separation (ICA2003), Nara Japan].
- [7] F. Meinecke, A. Ziehe, M. Kawanabe, and K.-R. Muller, *Estimating the reliability of ICA projections*, vol. 14. Cambridge, MA: MIT Press, 2002.
- [8] F. Meinecke, A. Ziehe, M. Kawanabe, and K.-R. Muller, "A resampling approach to estimate the stability of one-dimensional or multidimensional independent components," *IEEE Trans. Biomedical Engr.*, vol. 49, pp. 1514-1525, 2002.
- [9] J.-R. Duann, T.-P. Jung, W.-J. Kuo, T.-C. Yeh, S. Makeig, J.-C. Hsieh, and T. J. Sejnowski, "Single-trial variability in event-related BOLD signal," *NeuroImage*, vol. 15, pp. 823-835, 2002.
- [10] M. J. McKeown, S. Makeig, G. G. Brown, T.-P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski, "Analysis of fMRI data by blind separation into independent spatial components" *Human Brain Mapp.* vol. 6, pp. 160-188.
- [11] A. J. Bell, and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural Comput.* vol. 7, pp. 1129-1159, 1995.
- [12] S.-I. Amari, "Natural gradient works efficiently in learning," *Neural Computation*, vol. 10, pp. 251-276, 1998.
- [13] T.-W. Lee, M. Girolami, and T. J. Sejnowski, "Independent component analysis using an extended Infomax algorithm for mixed sub-Gaussian and super-Gaussian sources," *Neural Computation*, vol. 11, pp. 609-633, 1999.
- [14] S. Makeig et al., World Wide Web publication, now available from <http://www.sccn.ucsd.edu/fmrlab>, 1997.
- [15] J.-R. Duann, T.-P. Jung, and S. Makeig, FMRLAB: Matlab software for independent component analysis of fMRI data. World Wide Web publication <http://www.sccn.ucsd.edu/fmrlab>, 2003.