Constrained Tensor Decomposition via Guidance: Increased Inter and Intra-Group Reliability in fMRI Analyses

Peter B. Walker[1(✉)], Sean Gilpin[2], Sidney Fooshee[3], and Ian Davidson[2]

[1] Naval Medical Research Center, Silver Spring, MD, USA
peter.b.walker.mil@mail.mil
[2] University of California – Davis, Davis, CA, USA
sagilpin@davis.edu, davidson@cs.davis.edu
[3] United States Navy – Aerospace Experimental Psychology, Ann Arbor, MI, USA
sidney.fooshee@navy.mil

Abstract. Recently, Davidson and his colleagues introduced a promising new approach to analyzing functional Magnetic Resonance Imaging (fMRI) that suggested a more appropriate analytic approach is one that views the spatial and temporal activation as a multi-way tensor [1]. In this paper, we illustrate how the use of prior domain knowledge might be incorporated into the deconstruction of the tensor so as to increase analytical reliability. These results will be discussed in reference to implications towards military selection and classification.

Keywords: Tensor decomposition · Functional magnetic resonance imaging · Reliability

1 Introduction

In a series of papers, Davidson and his colleagues [1–4] recently introduced a promising new approach to analyzing functional connectivity that views the spatial and temporal activation as a multi-way tensor. The authors argued that such approaches may pave the way for new and innovative improvements across military domains. For example, the graph analytic approaches were applied to problems ranging from analyzing neuroimaging data to exploring the cause and effect relationships behind event networks.

Similarly, their approach to viewing multi-way data as a tensor, promises a number of different innovative opportunities across military domains. For example, many of the DoD services have begun to investigate the use of neuroimaging modalities for selection and classification [5, 6]. In addition, there has been tremendous interest in the use of using specific neuroimaging techniques to examine mental health issues including Traumatic Brain Injury (TBI) and Post-Traumatic Stress Disorder (PTSD) [7]. One such technique, functional Magnetic Resonance Imaging (fMRI), will be discussed in this paper.

D.D. Schmorrow and C.M. Fidopiastis (Eds.): AC 2015, LNAI 9183, pp. 361–369, 2015.
DOI: 10.1007/978-3-319-20816-9_34

fMRI provides the unique opportunity to visualize neural activity in the brain in real time. Perhaps one of the most promising applications of fMRI data has been on the analysis of functional connectivity [8]. The term functional connectivity has come to be understood as the temporal correlation of neuronal activation for spatially discrete locations. The excitement over this analytic approach is due, in part, to the applicability of these findings in both clinical and diagnostic settings [9].

The theoretical and practical advantages of several of the benchmark analytic approaches are well known [10]. For example, innovations in clustering such as spectral partitioning (SP) allows for the aggregation of several images to form a composite pattern with a single global minimum which can be approximated efficiently via a generalized eigenvalue problem and has been applied to fMRI clustering with promising results [11].

However, limitations in the applicability and interpretability of the results from fMRI following SP have been largely ignored. Specifically, SP and its specific use toward fMRI have been challenged with respect to test/retest reliability. For example, a review by Bennett and Miller [12] reported a mean Intra-class Correlation Coefficient (ICC) from 13 separate studies to be ICC = .50. However, the ICC appeared to vary widely depending on a number of different factors including scanner characteristics (magnet strength; signal-to-noise ratio), subject-specific (cognitive state), and/or task-specific (training effects). Depending on any number of these factors, the ICC ranged between .88 and .16.

The primary aim of this paper is to outline an approach to analyzing functional connectivity that overcomes many of the limitations of current data analytic approaches. Specifically, this paper demonstrates an increase in reliability over and above benchmark analytic approaches through the use of a technique we have termed *Constrained Tensor Decomposition* (hereafter, CTD). CTD attempts to transfer knowledge from existing domain knowledge (i.e., anatomical regions) to assist in the clustering/segmentation process thereby creating more stable fMRI images. In doing so, we show that we are more readily able to dissociate findings from different populations of scans. To accomplish this, we apply a form of spectral clustering to two separate fMRI scans. Unlike traditional clustering algorithms such as Independent Components Analysis, and/or Principle Components Analysis that attempt to segment a graph based on a single image, our approach incorporates knowledge from multiple graphs that might share the same set of nodes with the first graph, but have a different set of edges. Intuitively, the extra knowledge from the second graph may help to find a better partition than the one we can find with the first graph alone.

The lack of test/retest reliability through the use of SP is illustrated in Fig. 1. This figure shows SP results for two resting state fMRI data sets of the same healthy young individual, acquired in short succession on an MRI machine. Barring a major medical event between the two scans, the spatial and temporal patterns of resting state activity should largely be the same. Yet, the spectral partitioning of the two data sets has little in common despite the algorithm finding the global minimum of its objective function. This suggests that the clustering resulting from spectral partitioning are strongly influenced by random “noise” (i.e., MRI scanner noise) that have little to do with the true similarity of the time course of brain activity.

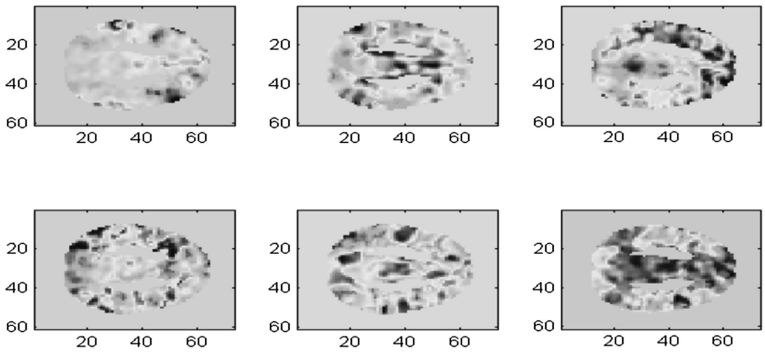

Fig. 1 Lack of test-retest reliability of spectral partitioning. A healthy young individual received two fMRI scans in rapid succession. The top and bottom rows show the three top-ranked spectral partitions of a corresponding slice from the first and second scans. Voxels more strongly associated with one partition, vs. the other, are shown in lighter vs. darker shades, respectively.

Our approach, CTD, offers significant improvements over previous attempts at measuring brain connectivity [13]. The aim of our work was the design of algorithms that can take event data in the form of activations over a three-dimensional spatial region x, y, z, over time t and simplify that data into a network [14]. This involves both aggregation so that the active cohesive regions (nodes) are identified and the formation of relationships (edges) between these regions. The edges and their weights can then be used to indicate properties such as information flow, excitation/inhibition or probabilistic relationships.

2 Constrained Tensor Decomposition

We view fMRI data as containing a complex interaction of signals and noise [14, 15] with a natural question being how to simplify the activity into the underlying cognitive network being used. Let $\boldsymbol{\chi}$ (*2D space x time*) be a three-mode tensor representing the fMRI data for a mid-level slice of the brain. The aim of CTD is to decompose (simplify) this tensor into f factors ($\widehat{\boldsymbol{\chi}} = \boldsymbol{\chi}_1 + \boldsymbol{\chi}_2 \ldots + \boldsymbol{\chi}_f$) using a PARAFAC model (though more complex decompositions could be used). Let factor i be defined by the outer product of three factor vectors ($\boldsymbol{\chi}_i = \mathbf{a_i} \circ \mathbf{b_i} \circ \mathbf{t_i}$) and for brevity we write ($\boldsymbol{\chi}_i = A \circ B \circ T$) with the factor vectors being stacked column-wise so that each factor matrix has f columns. Then $\boldsymbol{\chi}_i$ can potentially represent a region of the brain with the outer product of $\mathbf{a_i}$ and $\mathbf{b_i}$ being the active region and $\mathbf{t_i}$ their activations over time. However, with unconstrained tensor decomposition $\boldsymbol{\chi}_\mathbf{i}$ is typically not a spatially contiguous region nor does it necessarily match an anatomical region. To achieve this, guidance is introduced.

In the case of CTD, the objective function is complemented by adding constraints (guidance) as shown in Eq. 1. The addition of guidance helps rule out solutions that are

non-actionable by restricting them to be consistent with known domain knowledge and expectations.[1] Specifically, we identified 116 anatomical regions within a prototype scan. These groupings of voxels were defined by a matrix/mask with all given matrices being $Q_1 \ldots Q_m$. Examples of the Q matrices used in this work can be seen in Fig. 2 with one matrix for each anatomical region. However, most anatomical networks will only consist of a small number of nodes so we wish to use only a subset of Q matrices in the decomposition. We can encode which structures/nodes/matrices are to be used with a vector $\mathbf{w}$ with one component/entry per factor. We can represent how closely the discovered factors match these matrices (allowing for some deviation) which is upper-bounded by which is proportional to the number of voxels outside the given group. The formulation for this process is presented in Eq. 1.

$$\begin{aligned} \underset{w,A,B,T}{\arg\min} \quad & \|\chi - \sum_i w_i, a_i \circ b_i \circ t_i\| F + \|w\| \\ \text{subject to} \quad & w_1 a_1 b_1^T - Q_1\| \leq \in_1 \\ & \vdots \\ & w_m \|a_m b_m^T - Q_m\| \leq \in_m \end{aligned} \tag{1}$$

The output of this computation will be the smallest set of groups/nodes (defined by $\mathbf{a_i} \circ \mathbf{b_i}$) that best summarize the fMRI scan with $\mathbf{t_i}$ stating the activations over time. The penalty term $\|\mathbf{w}\|$ introduces a sparsity constraint to enforce that the simplest structure is discovered. Furthermore, by rank ordering the factors by their entry in $\mathbf{w}$ we can determine the most important nodes in the network. Figure 2 shows an example of our work where the Q matrices are just of the cores of a pre-defined anatomical network. The top four most important nodes as per the $\mathbf{w}$ matrix and the corresponding Q matrices are shown.

To illustrate the applicability of CTD over and above the benchmark analytic techniques for network segmentation, we applied CTD to a series of scans from resting-state fMRI. Typically, clustering from resting-state fMRI results in the segmentation of a particular set of voxels known as the default mode network (DMN) [16–18]. Further, this clustering, while reliably found in healthy populations, has has been shown to be less clearly differentiated elderly populations, especially those with Alzheimer's Disease.

For the present work, we identified an exemplar scan whose clustering clearly indicated the DMN as one of its clusters. We then applied CTD, using the exemplar scan to demarcate regions of interest, to partition groups of scans including young and elderly individuals. Using this technique, we will illustrate a boost in intra-individual reliability (clustering across scans), by showing greater differentiation between different populations.

[1] Unconstrained tensor decomposition results have been shown to be relatively poor for fMRI data since, many spatially adjacent voxels in the same structure are not always active in the same factor. Pre-processing the tensor by applying wavelets can alleviate this and could complement our work, though in practice this pre-processing was time intensive and yielded only marginally better results.

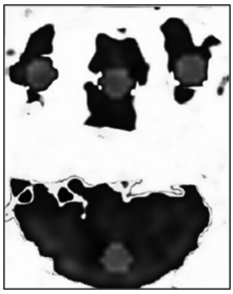

Fig. 2. Q matrices identified (a priori) as the top four nodes of the DMN amongst a possible 116 anatomical regions.

3 Empirical Results

Here, we used CTD to boost the reliability of intra-individual fMRI clustering and provide a new approach for assessing inter-individual clustering commonalities at a population level [3]. CTD uses an efficient generalized eigenvalue formulation to find a 2-category clustering of voxels such that voxels within the one cluster have highly similar time series. But at the same time, CTD attempts to satisfy user-provided guidance about which voxels should and should not cluster together. In this case, guidance took the form of the identification of 116 known anatomical regions used to constrain CTD to cluster a target scan. CTD provides a cluster quality metric that quantifies how difficult it was to provide a high-quality clustering of the target that is similar to that of the exemplar; we use this metric as a *clustering metric* between the groups. In our experiments, we summarize an entire population of fMRI scans in terms of clustering distance from an exemplar; however, we note that the clustering distance between any pair of scans can be quantified by applying CTD, thus providing a fine-grained assessment of the clustering quality between populations.

To illustrate the robustness of our approach, CTD was compared directly to the results of clustering using one of the benchmark analytic approaches, Independent Components Analysis (ICA) [19, 20]. ICA has widely been used across several resting and functional imaging paradigms. Independent component analysis attempts to separate independent "sources" of data by identifying different factors. These factors then, can be used to assess different levels of functional connectivity across populations.

In contrast, CTD attempts to take imaging data in the form of activations over a three dimensional spatial region, x, y, z over time t and simplify that activation into a network. Previously, we illustrated how this can be accomplished so that active cohesive regions (nodes) can be identified along with the formation of relationships (edges) between those regions [2]. Here, we focus more on the use of CTD to boost intra-individual clustering across scans.

We used real resting-state fMRI scans of normal and demented elderly individuals to demonstrate the advantages of CTD over ICA. A particular set of voxels, referred to as the default mode network (DMN), is known to generate highly similar time series in fMRI scans of healthy individuals, and the tightness of this clustering is known to be

decreased in individuals with Alzheimer's disease [16]. Therefore, we analyzed the functional scans of 8 normal elderly and 8 demented participants. We used CTD to partition target scans including elderly and demented individuals based on constraints derived from the anatomical boundaries.

We present experimental results on fMRI data of eight healthy (Elderly), eight individuals with Mild Cognitive Impairment (MCI), and eight demented (Alzheimer's) individuals at rest. Each individual had been interviewed and measured using a series of cognitive tests to measure Episodic, Executive, Semantic and Spatial scores (whose range is from -2.5 to +2.5) and were categorized as Normal or having full-set Alzheimer's (Demented). When at rest state, an individual's brain activity exhibits the DMN in some form and in various degrees in all people. In normal individuals, the DMN is expected to be fully intact and well exhibited. However, for demented individuals the DMN may be only partially formed and the signal may be ery weak. These insights are well known and extensively published in the literature [16, 17] and we expect our method to be able to verify these results. The ability to determine the strength of the DMN from the scan is then akin to being able to predict the progression of Alzheimer's which we show is possible by predicting the cognitive scores (see below).

As can be clearly seen in Fig. 3 above, the use of SP was not sufficient in differentiating between the three populations of scans. However, through the use of CTD, we showed a differentiation in network quality between Elderly, MCI, and Demented. In other words, we show that the intra-individual reliability of CTD is greater than that of ICA. We also show that, as expected, the clustering metric (quality of cuts) between the healthy elderly individuals is less than that between the exemplar and demented elderly individuals, while there is no such significant difference for the analogous SP cluster quality metric. Finally, we show that this group difference between healthy and demented elders is robust across a range of constraint set sizes.

To provide convergence with the results presented above, we provided our algorithm with all 116 anatomical regions/masks of the brain encoded each in its own

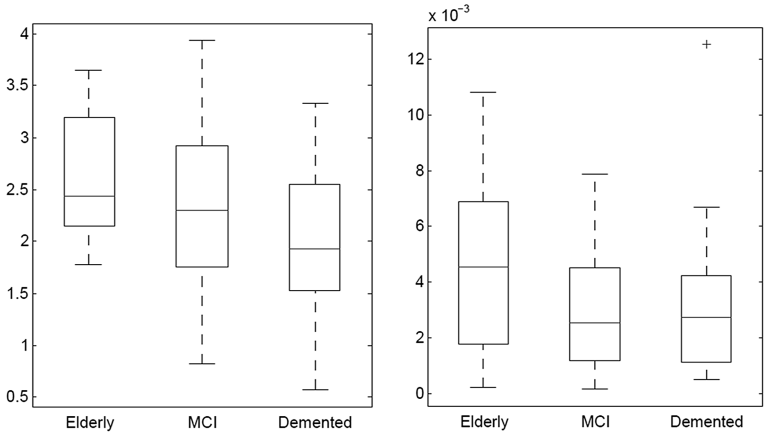

Fig. 3. SP and CTD partition cost differences. As can be seen, ICA (left) fails to differentiate clustering between different populations. However, the use of CTD (right) allows for clearer differentiation between Elderly and MCI and Demented.

Q matrix and performed a tensor decomposition and selected the top four factors (according to the w vector) to determine which nodes/structures were active in the brain during the scan. Table 1 (columns 2 and 3) shows the fraction of the various nodes/parts of the DMN found for Normal and Demented people in the top four factors.

Here, we clearly see that the DMN is completely discovered in seven out of the eight Elderly individuals with the eighth individual's prefrontal region being ranked fifth according to w. However, for the Demented group, the DMN is found in its entirety in only two of the eight patients. In other words, our approach clearly differentiated the signal strength between Elderly and Demented in terms of identifying specific brain regions associated with the DMN.

Figure 4 (top) shows the actual network reconstructed from the top four factors for the majority of the Elderly group. However, as can be seen in Fig. 4 (bottom), CTD was not able to recover the DMN for the majority of participants in the Demented group. This finding suggests, consistent with previous findings, that dementia has resulted in network activation that is less clearly organized and therefore does not preserve the DMN.

Table 1. Mean Pearson correlation for demented individuals for top four nodes of a network discovered for each individual. The nodes/structures denoted by * are part of the DMN.

	Elderly	Demented
Prefrontal Region	88 %	50 %
Posterior Cingulate	100 %	63 %
Inferior Parietal	100 %	38 %
Medial Temporal	100 %	25 %

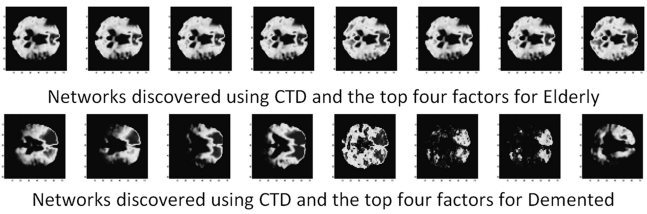

Fig. 4. Networks reconstructed using the top four factors from Elderly (top) and Demented individuals.

4 Discussion

The primary aim of this paper was to introduce a new method for image segmentation of fMRI data. Specifically, we showed how the use of Constrained Tensor Decomposition could be used to better differentiate three different populations of scans. Here, we showed how this data-driven approach can utilize previous domain knowledge to increase the overall validity of the results.

The widespread use of neuroimaging data has warranted the need for the development of new and innovative techniques that analyze the data with a specific focus towards incorporating domain knowledge. Here, we showed how the use of CTD can be used by incorporating domain knowledge on an a priori basis to drive the data towards answering specific questions. Using this type of technique, we believe new opportunities in areas such as military selection and classification will develop.

Acknowledgments. This research is supported by Office of Naval Research grant NAVY 00014-09-1-0712. The opinions of the authors do not necessarily reflect those of the United States Navy or the University of California - Davis.

Peter B Walker and Sidney Fooshee are military service members. This work was prepared as part of their official duties. Title 17 U.S.C. 101 defines U.S. Government work as a work prepared by a military service member or employee of the U.S. Government as part of that person's official duties.

References

1. Walker, P.B., Davidson, I.: Exploring new methodologies for the analysis of functional magnetic resonance imaging (fMRI) following closed-head injuries. In: Schmorrow, D.D., Fidopiastis, C.M. (eds.) FAC 2011. LNCS, vol. 6780, pp. 120–128. Springer, Heidelberg (2011)
2. Davidson, I., Gilpin, S., Carmichael, O., Walker, P.: Guided network discovery via constrained tensor analysis of fMRI data. In: KDD 2013 Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 194–202. ACM, New York (2013)
3. Wang, X., Qian, B., Davidson, I.: On constrained spectral clustering and its applications. Data Min. Knowl. Disc. **28**(1), 1–30 (2014)
4. Phillips, H.L., Walker, P.B., Kennedy, C.H., Carmichael, O., Davidson, I.N.: Guided learning algorithms: an application of constrained spectral partitioning to functional magnetic resonance imaging (fMRI). In: Schmorrow, D.D., Fidopiastis, C.M. (eds.) AC 2013. LNCS, vol. 8027, pp. 709–716. Springer, Heidelberg (2013)
5. Olson, T.M., Walker, P.B., Phillips IV, H.L.: Assessment and selection of aviators in the US military. In: O'Connor, P.E., Cohn, J.V. (eds.) Human Performance Enhancement in High-Risk Environments: Insights, Developments, and Future Directions from Military Research: Insights, Developments, and Future Directions from Military Research, pp. 37–57. ABC-CLIO, Santa Barbara (2009)
6. Gordon, H.W., Leighty, R.: Importance of specialized cognitive function in the selection of military pilots. J. Appl. Psychol. **73**(1), 38 (1988)
7. Graner, J., Oakes, T. R., French, L. M., Riedy, G.: Functional MRI in the investigation of blast-related traumatic brain injury. Front. Neurol. **4** (2013)
8. Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J.: Functional connectivity: the principal-component analysis of large (PET) data sets. J. Cerebr. Blood Flow Metab. **13**, 5 (1993)
9. Hoge, C.W., Castro, C.A., Messer, S.C., McGurk, D., Cotting, D.I., Koffman, R.L.: Combat duty in Iraq and Afghanistan, mental health problems, and barriers to care. N. Engl. J. Med. **351**(1), 13–22 (2004)

10. Tan, P.N., Steinbach, M., Kumar, V.: Introduction to Data Mining, 1st edn. Addison-Wesley Longman Publishing Co. Inc., Redwood City (2005)
11. Liu, Y., Wang, K., Yu, C., He, Y., Zhou, Y., Liang, M., Jiang, T.: Regional homogeneity, functional connectivity and imaging markers of Alzheimer's disease: a review of resting-state fMRI studies. Neuropsychologia **46**(6), 1648–1656 (2008)
12. Bennett, C.M., Miller, M.B.: How reliable are the results from functional magnetic resonance imaging? Ann. N. Y. Acad. Sci. **1191**(1), 133–155 (2010)
13. Rubinov, M., Sporns, O.: Complex network measures of brain connectivity: uses and interpretations. Neuroimage **52**(3), 1059–1069 (2010)
14. Davidson, I., Gilpin, S., Walker, P.B.: Behavioral event data and their analysis. J. Knowl. Disc. Data Min. **25**, 635–653 (2012). Springer, Special Issue on Human Behavioral Modeling
15. Genovese, C.R., Lazar, N.A., Nichols, T.: Thresholding of statistical maps in functional neuroimaging using the false discovery rate. Neuroimage **15**(4), 870–878 (2002)
16. Greicius, M.: Resting-state functional connectivity in neuropsychiatric disorders. Curr. Opin. Neurol. **21**(4), 424–430 (2008)
17. Greicius, M.D., Srivastava, G., Reiss, A.L., Menon, V.: Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. Proc. Natl. Acad. Sci. U.S.A. **101**(13), 4637–4642 (2004)
18. Raichle, M.E., Snyder, A.Z.: A default mode of brain function: a brief history of an evolving idea. Neuroimage **37**(4), 1083–1090 (2007)
19. Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J.: A method for making group inferences from functional MRI data using independent component analysis. Hum. Brain Mapp. **14**(3), 140–151 (2001)
20. Ma, L., Wang, B., Chen, X., Xiong, J.: Detecting functional connectivity in the resting brain: a comparison between ICA and CCA. Magn. Reson. Imaging **25**(1), 47–56 (2007)