NeuroImage 82 (2013) 616–633

Contents lists available at ScienceDirect

NeuroImage

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

Resting state network estimation in individual subjects

Carl D. Hacker [a,*], Timothy O. Laumann [b], Nicholas P. Szrama [a], Antonello Baldassarre [b,e], Abraham Z. Snyder [b,c], Eric C. Leuthardt [a,b,1], Maurizio Corbetta [b,c,d,e,1]

[a] Department of Biomedical Engineering, Washington University, USA
[b] Department of Neurology and Neurosurgery, Washington University, USA
[c] Department of Radiology, Washington University, USA
[d] Department of Anatomy and Neurobiology, Washington University, USA
[e] Department of Neuroscience and Imaging, University of Chieti, "G. d'Annunzio" Italy

ARTICLE INFO

Article history:
Accepted 22 May 2013
Available online 2 June 2013

Keywords:
fMRI
Resting state network
Multilayer perceptron
Functional connectivity
Supervised classifier
Brain mapping

ABSTRACT

Resting state functional magnetic resonance imaging (fMRI) has been used to study brain networks associated with both normal and pathological cognitive functions. The objective of this work is to reliably compute resting state network (RSN) topography in single participants. We trained a supervised classifier (multi-layer perceptron; MLP) to associate blood oxygen level dependent (BOLD) correlation maps corresponding to pre-defined seeds with specific RSN identities. Hard classification of maps obtained from *a priori* seeds was highly reliable across new participants. Interestingly, continuous estimates of RSN membership retained substantial residual error. This result is consistent with the view that RSNs are hierarchically organized, and therefore not fully separable into spatially independent components. After training on *a priori* seed-based maps, we propagated voxel-wise correlation maps through the MLP to produce estimates of RSN membership throughout the brain. The MLP generated RSN topography estimates in individuals consistent with previous studies, even in brain regions not represented in the training data. This method could be used in future studies to relate RSN topography to other measures of functional brain organization (*e.g.*, task-evoked responses, stimulation mapping, and deficits associated with lesions) in individuals. The multi-layer perceptron was directly compared to two alternative voxel classification procedures, specifically, dual regression and linear discriminant analysis; the perceptron generated more spatially specific RSN maps than either alternative.

Introduction

Biswal et al. (1995) first described resting state functional magnetic resonance imaging (fMRI) in 1995. The literature in this field has since been growing exponentially (Snyder and Raichle, 2012). Most of this work has been directed towards describing the statistical properties of intrinsic blood oxygenation level dependent (BOLD) signal fluctuations in health and disease (Biswal et al., 2010; Fox and Greicius, 2010; Pievani et al., 2011; Zhang and Raichle, 2010). Spontaneous BOLD activity recapitulates, in the topographies of its temporal covariance structure, task-based fMRI responses to a wide variety of behavioral paradigms (Smith et al., 2009). These topographies currently are known as resting state networks (RSNs) or, equivalently, intrinsic connectivity networks (ICNs). RSNs have now been mapped over virtually all of the cerebral cortex as well as many subcortical structures including the cerebellum (Buckner et al., 2011; Choi et al., 2012; Lee et al., 2012; Power et al., 2011; Yeo et al., 2011). Critically, although RSN topographies differ across individuals (Mennes et al., 2010; Mueller et al., 2013), results generally have been reported at the group level. Effectively capturing individual differences in RSN organization would enhance the study of how intrinsic activity accounts for individual differences in human behavior and cognition.

Reliable RSN mapping in individuals has multiple applications, for example, in the study of the physiological basis of inter-individual differences in cognition, *e.g.*, (Cole et al., 2012; Koyama et al., 2011). Similarly, improved RSN mapping in individuals could be useful in the study of how focal lesions, *e.g.*, strokes, lead to performance deficits (Carter et al., 2010; Golestani et al., 2013; He et al., 2007); such studies are difficult at the group level because of lesion heterogeneity. Yet another application is to improve the delineation of "eloquent" cortex prior to neurosurgery, to potentially reduce iatrogenic deficits (Otten et al., 2012; Tie et al., 2013; Zhang et al., 2009). Pre-operative task-fMRI has been used for this purpose (Wurnig et al., 2013) but often fails because patients are unable to comply with task paradigms. Lastly, individual RSN mapping could enhance functional coregistration, *i.e.*, using RSN features to refine anatomical registration (Conroy et al., 2013; Sabuncu et al., 2010).

Two analytic strategies, seed-based correlation mapping (Biswal et al., 2010) and spatial independent component analysis (sICA)

* Corresponding author at: Washington University in St Louis, 660 S Euclid Ave, Neurosurgery, Campus Box 8057, Saint Louis MO 63110-1010, USA. Fax: +1 314 362 6110.
E-mail address: hackerc@wusm.wustl.edu (C.D. Hacker).
[1] Equally contributed to this work.

1053-8119/$ – see front matter
http://dx.doi.org/10.1016/j.neuroimage.2013.05.108

(Beckmann, 2012), have so far dominated the field of resting state fMRI. RSNs obtained by sICA are theoretically unbiased by prior assumptions. However, ICA has not been shown to be robust at the single subject level; results obtained by this technique generally are reported at the group level. Seed-based correlation mapping uses priors if it is limited to only a few seeds. However, systematically defining many seeds over the entire brain (Wig et al., 2013) and analyzing the results using graph theoretic tools (Power et al., 2011) or inner product based clustering (Lee et al., 2012; Yeo et al., 2011) effectively achieves independence from priors. Both ICA and systematic seed-based correlation mapping exemplify unsupervised learning. Therefore, RSNs obtained by these methods may differ not only in topography (*i.e.*, extent and shape), but also in topology (*i.e.*, number of distinct nodes making up a single RSN), depending on the granularity of the recovered components within the modular hierarchy of RSNs (Meunier et al., 2010). To illustrate, the default mode network (DMN) is a constellation of regions including the posterior cingulate–precuneus cortex (PCC), midline prefrontal cortex, lateral parietal cortex, superior frontal cortex and posterior cerebellum. The DMN may be recovered in its entirety (Fox et al., 2005) using highly supervised methods. However, unsupervised strategies variably recover the DMN in fragments (Kahn et al., 2008; Smith et al., 2009), or combined with fragments of other networks (Doucet et al., 2011; Lee et al., 2012; Yeo et al., 2011). Such inconsistencies stem from the fact that unsupervised learning procedures are not constrained to a particular topological scale; therefore, some *post-hoc* classification strategy (*e.g.*, template matching) must be used to establish RSN identity.

The present work is fundamentally different in that the objective is not to discover RSNs nor to study their functional relevance, but rather to map the topography of known RSNs in individuals. To this end, we trained a multi-layer perceptron (MLP) to estimate RSN memberships of brain loci on the basis of BOLD correlation maps. A perceptron is a feed-forward artificial neural network, originally modeled on the human visual system, trained to associate weighted sums of input features with pre-defined output classes (Rosenblatt, 1958). After training, the MLP decision boundaries are fixed; thus, subsequent results are guaranteed to represent the same entity (at the same topological scale) across individuals or populations. Perhaps the best-known application of perceptrons is to recognize (classify) handwritten digits (LeCun et al., 1989). This application has obvious utility in automatic routing of letters at the post office. To distinguish between supervised *vs*. unsupervised learning, consider discovering the characters used to represent numbers in the decimal system by analysis of a large sample of addressed letters. This is very different from training a perceptron to read (classify) known numerals, *e.g.*, zip codes on addressed letters. Analogously, RSN discovery, using group sICA or any other unsupervised method, is very different from preparing a trained MLP to map known RSNs in new subjects.

In the above example, each character must represent one and only one numeral. However, we do not assume that every brain region belongs to a single RSN. We allow each locus in the brain to belong to any RSN to a variable degree. Accordingly, RSN membership estimation represents regression rather than classification. However, classification and regression are closely related mathematically. MLP outputs, which approximate posterior probabilities of class membership (Ruck et al., 1990), can be converted to hard classifications by identifying the output class of greatest magnitude. We report both continuous RSN estimates and hard classifications ("winner-take-all" maps). MLP performance was characterized by residual error for the former and receiver operating characteristic (ROC) analysis for the latter (Section Quantitation of MLP performance).

Our methodology represents a solution to an engineering problem, namely, mapping RSNs in individuals. However, MLP training performance offers valuable information about the structure and separability of resting state networks. Differential performance across RSNs may provide insight into their relative inter-subject variability and complexity. MLP performance also provides an objective measure of data quality that can be used to study the effects of varying acquisition and preprocessing methodologies. We demonstrate this concept by determining the quantity of BOLD data required to reliably compute RSN topography in individual subjects. Similarly, we empirically determine the optimal ROI size for generation of correlation map training data. As a final result, two alternative strategies for extending group-level RSN topographies to individuals (linear discriminant analysis and dual regression) are compared to the MLP. This comparison shows that the MLP provides superior RSN mapping specificity.

Methods

The Methods section is organized as follows: We first describe the fMRI datasets (Section Participants) and neuroimaging methods (Section Neuroimaging methods). We next describe the task-fMRI meta-analyses (Section Meta-analysis of task fMRI and generation of training data) used to isolate seed ROIs. These seeds were used to generate the MLP training data. MLP-specific methodology is divided into design (Section MLP design) and application (Section Application of method to individuals). The design phase (Section MLP design) used correlation maps corresponding to seed ROIs with categorical RSN labels to train (Section MLP training), evaluate (Section Quantitation of MLP performance), and optimize (Section Architecture selection and Section Performance optimization by simulated annealing) the MLP. Application of the trained perceptron to individuals generated voxel-wise estimates of RSN membership throughout the brain (Section Application of method to individuals). MLP results then were compared to dual regression (DR) and linear discriminant analysis (LDA) (Section Comparison of the MLP to linear discriminant analysis and dual regression).

Participants

Perceptron training, optimization and validation used datasets previously acquired at the Neuroimaging Laboratories (NIL) at the Washington University School of Medicine. A second, large validation dataset was obtained from the Harvard-MGH Brain Genomics Superstruct Project (Yeo et al., 2011). All patients were young adults screened to exclude neurological impairment and psychotropic medications. Demographic information and acquisition parameters are given in Table 1.

Neuroimaging methods

MRI acquisition

Imaging was performed with a 3 T Allegra (NIL) or Tim Trio (Harvard-MGH) scanner. Functional images were acquired using a BOLD contrast sensitive gradient echo echo-planar sequence [parameters listed in Table 1] during which participants were instructed to fixate on a visual cross-hair, remain still and not fall asleep. Anatomical imaging included one sagittal T1-weighted magnetization prepared rapid gradient echo (MP–RAGE) scan (T1W) and one T2-weighted scan (T2W).

fMRI preprocessing

Initial fMRI preprocessing followed conventional practice (Shulman et al., 2010). Briefly, this included compensation for slice-dependent time shifts, elimination of systematic odd-even slice intensity differences due to interleaved acquisition (Supplemental text) and rigid body correction of head movement within and across runs. Atlas transformation was achieved by composition of affine transforms connecting the fMRI volumes with the T2W and T1W structural images. Head movement correction was included with the atlas transformation in a single resampling that generated volumetric timeseries in $(3\ mm)^3$ atlas space. Additional preprocessing in preparation for correlation mapping included spatial smoothing (6 mm full width at half