Accepted Manuscript

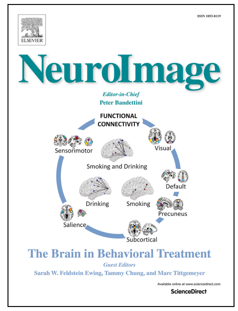

A probabilistic approach to discovering dynamic full-brain functional connectivity patterns

Jeremy R. Manning, Xia Zhu, Theodore L. Willke, Rajesh Ranganath, Kimberly Stachenfeld, Uri Hasson, David M. Blei, Kenneth A. Norman

PII: S1053-8119(18)30071-5

DOI: 10.1016/j.neuroimage.2018.01.071

Reference: YNIMG 14684

To appear in: *NeuroImage*

Received Date: 14 February 2017

Revised Date: 9 January 2018

Accepted Date: 28 January 2018

A probabilistic approach to discovering dynamic full-brain functional connectivity patterns

Jeremy R. Manning[1,*], Xia Zhu[2], Theodore L. Willke[2], Rajesh Ranganath[3], Kimberly Stachenfeld[4], Uri Hasson[3], David M. Blei[5], and Kenneth A. Norman[3]
[1]Dartmouth College, Hanover, NH
[2]Intel Corporation, Hillsboro, OR
[3]Princeton University, Princeton, NJ
[4]DeepMind, London, UK
[5]Columbia University, New York, NY
*Address correspondence to jeremy.r.manning@dartmouth.edu

February 8, 2018

Abstract

Recent research shows that the covariance structure of functional magnetic resonance imaging (fMRI) data – commonly described as *functional connectivity* – can change as a function of the participant's cognitive state (for review see [35]). Here we present a Bayesian hierarchical matrix factorization model, termed *hierarchical topographic factor analysis* (HTFA), for efficiently discovering full-brain networks in large multi-subject neuroimaging datasets. HTFA approximates each subject's network by first re-representing each brain image in terms of the activities of a set of localized *nodes*, and then computing the covariance of the activity time series of these nodes. The number of nodes, along with their locations, sizes, and activities (over time) are learned from the data. Because the number of nodes is typically substantially smaller than the number of fMRI voxels, HTFA can be orders of magnitude more efficient than traditional voxel-based functional connectivity approaches. In one case study, we show that HTFA recovers the known connectivity patterns underlying a collection of synthetic datasets. In a second case study, we illustrate how HTFA may be used to discover dynamic full-brain activity and connectivity patterns in real fMRI data, collected as participants listened to a story. In a third case study, we carried out a similar series of analyses on fMRI data collected as participants viewed an episode of a television show. In these latter case studies, we found that the HTFA-derived activity and connectivity patterns can be used to reliably decode which moments in the story or show the participants were experiencing. Further, we found that these two classes of patterns contained partially non-overlapping information, such that decoders trained on combinations of activity-based and dynamic connectivity-based features performed better than decoders trained on activity or connectivity patterns alone. We replicated this latter result with two additional (previously developed) methods for efficiently characterizing full-brain activity and connectivity patterns.

Introduction

The most common approaches for analyzing functional Magnetic Resonance Imaging (fMRI) data involve relating, in individual images, the activity of individual voxels or multi-voxel spatial patterns of brain activity to the subject's cognitive state [14, 15, 27, 41]. In contrast, functional connectivity analyses correlate the time series of activities *across images* of pairs of voxels [30]. Functional connectivity analyses have already led to new insights into how the brain's correlational structure changes during different experimental conditions [35].

The size of the full-brain functional connectivity matrix grows with the square of the number of voxels. Because of this rate of growth, filling in its entries and storing it in memory can become intractable for fMRI images with tens of thousands of voxels. For example, for a series of 50,000 voxel images, each connectivity