




Contents lists available at ScienceDirect

NeuroImage

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

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

Jeremy R. Manning^{a,*}, Xia Zhu^b, Theodore L. Willke^b, Rajesh Ranganath^c, Kimberly Stachenfeld^d, Uri Hasson^c, David M. Blei^e, Kenneth A. Norman^c

^a Dartmouth College, Hanover, NH, United States

^b Intel Corporation, Hillsboro, OR, United States

^c Princeton University, Princeton, NJ, United States

^d DeepMind, London, UK

^e Columbia University, New York, NY, United States

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 Turk-Browne, 2013). 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 (Friston et al., 1994, 1996; Norman et al., 2006; Zarahn et al., 1997). In contrast, functional connectivity analyses correlate the time series of activities *across images* of pairs of voxels (Rubinov and Sporns, 2010). Functional connectivity analyses have already led to new insights into how the brain's correlational structure changes during different experimental conditions (Turk-Browne, 2013).

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 matrix occupies approximately 5 GB (storing only the upper triangle, using single precision floating point entries). Storing many such matrices in memory (e.g. to compare different subjects and/or experimental conditions) can exceed the limits of modern hardware. Further, many of the algorithms used to relate multivariate patterns of voxel activities in individual images to cognitive states or experimental conditions (e.g. Haxby et al., 2001; Norman et al., 2006) use superlinear time and memory with respect to the number of features; this computational expense makes it impractical to use the same techniques to examine correlational data (although see (Wang et al., 2015) for another promising approach using massive parallelization on a specialized computing cluster).

Previous work circumvents this issue by using functionally (Gonzalez-Castillo et al., 2015; Craddock et al., 2012; Power et al., 2011; Yeo et al., 2011) or anatomically (Betzel et al., 2017; Fox and Greicius, 2010; Shen et al., 2013) defined voxel clusters or regions of interest. These

* Corresponding author.

E-mail address: jeremy.r.manning@dartmouth.edu (J.R. Manning).

<https://doi.org/10.1016/j.neuroimage.2018.01.071>

Received 14 February 2017; Received in revised form 9 January 2018; Accepted 28 January 2018

Available online xxx

1053-8119/© 2018 Published by Elsevier Inc.

approaches segment the brain into discrete components, and then examine interactions or correlations between the activity patterns exhibited by those components (rather than attempting to examine every voxel-to-voxel interaction). In other words, these approaches encapsulate an intuition about how our brains work—specifically, that our brains are composed of a small number of network nodes that interact with each other.

One such class of methods entails pre-selecting a small number of regions of interest [ROIs; e.g. motor cortex (Biswal et al., 1995)] or a seed voxel [e.g. a single voxel within the posterior cingulate (Greicius et al., 2003)]. This procedure reduces the connectivity matrix from a $V \times V$ matrix (where V is the total number of voxels in each brain volume) to a much smaller $V_{ROI} \times V$ matrix (where V_{ROI} is the number of seed or ROI voxels). However, reducing the connectivity matrix in this way precludes finding connectivity patterns unrelated to the ROI or seed region. For example, if the analysis is limited to connectivity patterns between motor cortex and the rest of the brain, this precludes finding patterns of connectivity that do not involve the motor cortex (e.g. connectivity between prefrontal cortex and the hippocampus). Other work has examined full-brain connectivity patterns via ROI-to-ROI connectivity matrices (e.g. as defined using anatomical data; Honey et al., 2009). This method provides a tractable means of examining full-brain connectivity patterns by assuming that each ROI is perfectly uniform and discrete.

A related technique that does not require pre-selecting ROIs or seed regions is to compute the full voxel-to-voxel connectivity matrix, and then to threshold the connection strengths such that one only examines the most reliable connections (Cao and Worsley, 1999). This approach yields a sparse voxel-to-voxel connectivity matrix that may be efficiently manipulated (provided that it is sufficiently sparse). One drawback to this approach is that it is not always clear how to set the connectivity strength threshold; for example, setting too high a threshold will leave out potentially important patterns, whereas setting too low a threshold will not substantially reduce the computational burden (as compared with examining the original voxel-to-voxel connectivity matrix). However, regularization techniques such as sparse regression (e.g., Valdés-Sosa et al., 2005) show promise. Another potential drawback is that it is not clear that the strongest connections are necessarily the most informative; for example, a sub-threshold connection may still carry cognitively relevant information.

Other approaches have focused on reducing the dimensionality of the connectivity patterns (see Worsley et al., 2005) for a comparison between threshold-based and dimensionality-reduction-based approaches). For example, clustering-based approaches attempt to group together voxels into “factors” that exhibit similar activity patterns over time, across subjects (Cordes et al., 2002). One may then examine connectivity patterns between the factors rather than between the voxels [this approach was originally developed in the positron emission tomography literature using Principle Components Analysis (PCA; Friston et al., 1993)].

Here we propose *hierarchical topographic factor analysis* (HTFA), a Bayesian factor analysis model specifically designed to facilitate analyses of brain network dynamics. HTFA is focused on finding spherical nodes that can be used to efficiently explore, analyze, and understand full-brain network dynamics. Like other dimensionality reduction methods, HTFA provides a compact means of representing full brain connectivity patterns that scales well to large datasets (Anderson et al., 2016). But HTFA goes beyond these methods: (a) it provides a natural means of determining how many network nodes should be used to describe a given dataset; (b) it allows those nodes to be overlapping rather than forcing nodes to be fully distinct; and (c) it constrains the nodes to be in similar (but not necessarily identical) locations across people. Further, HTFA decomposes brain images into sums of spatial functions, which supports seamless mapping between images of different resolutions (e.g. different voxel sizes) and potentially between different recording modalities (e.g. fMRI, EEG, ECoG, MEG). We return to these latter points in the *Discussion*.

HTFA casts each subjects' brain images as linear combinations of latent factors [Gaussian radial basis functions (RBFs)]. Each RBF can be

interpreted as a spherical *node* in a simplified representation of the brain's networks. (The number of nodes, K , is determined from the data.) In this way, HTFA is a spatial model: these nodes reflect structures localized in 3D space whose activity patterns influence the observed voxel activities. This is conceptually different from approaches that interpret voxel activities directly: HTFA defines an explicit model for how voxels relate to each other according to their relative locations in space (also see Gershman et al., 2011). One advantage of constraining HTFA's nodes to be RBFs is that they are spatially compact. In this way, HTFA is qualitatively similar to spatial ICA (Calhoun et al., 2001). However, unlike spatial ICA, HTFA further constrains nodes to be in similar locations across subjects, providing a natural means of combining or comparing connectivity across individuals. [ICA with dual regression (Beckmann et al., 2009) also models hierarchical patterns across participants.] Note that in this paper we have constrained the RBF nodes to be isotropic, but in principle this assumption could be relaxed to support any parameterizable spatial function (e.g., non-isotropic Gaussians, mixtures of Gaussians, 3D wavelets, etc.). These spatial functions could also be constructed to incorporate detailed conductance models (e.g. informed by diffusion tensor imaging data). We have left these extensions to be explored in future work.

Applying HTFA to an fMRI dataset reveals the locations and sizes of these network nodes (i.e. the centers and widths of their RBFs), as well as the per-image node weights. If a given subject has contributed N images to the dataset, then the subject's N by K node weights matrix may be viewed as a low-dimensional embedding of their original data. Further, the pairwise correlations between columns of this weight matrix reflect the signs and strengths of the node-to-node connections (just as the pairwise correlations between voxel time series reflect the corresponding “connectivity” in voxel-based functional connectivity analyses).

The next section provides a descriptive overview of the HTFA model. *Materials and methods* describes how we efficiently fit the model to a multi-subject fMRI dataset using *maximum a posteriori* inference. The *Supplemental materials* contain a complete formal (mathematical) description of the model. To validate our approach, we first generated a set of 50 synthetic datasets for which the underlying activity patterns, node locations and sizes, etc. were known, and we show that HTFA recovers these known patterns. Then we applied HTFA to two (real) fMRI datasets collected as participants listened to an audio recording of a story (Case Study 2) and watched an episode of a television show (Case Study 3). We show that the moment-by-moment patterns uncovered by HTFA may be used to decode which moments in the story or show the participants were experiencing.

Model description

Overview

HTFA is a member of a family of models, called *factor analysis* models, that includes Topographic Factor Analysis (TFA; Manning et al., 2014), Topographic Latent Source Analysis (TLSA; Gershman et al., 2011), Principal Components Analysis (PCA; Pearson, 1901), Exploratory Factor Analysis (EFA; Spearman, 1904), and Independent Components Analysis (ICA; Jutten and Herault, 1991; Comon et al., 1991), among others. If we have organized our collection of images (from a single subject) into an N by V data matrix \mathbf{Y} (where N is the number of images and V is the number of voxels), then factor analysis models decompose \mathbf{Y} as follows:

$$\mathbf{Y} \approx \mathbf{W}\mathbf{F}, \quad (1)$$

where \mathbf{W} is an N by K weight matrix (which describes how each of K factors are activated in each image), and \mathbf{F} is a K by V matrix of factor images (which describes how each factor maps onto the brain). Note that, in the general case, this decomposition is underspecified— in other words, there are infinitely many solutions for \mathbf{W} and \mathbf{F} that approximate the data equally well. What differentiates factor analysis models is the particular

constraints they place on what form \mathbf{W} and/or \mathbf{F} should take (i.e. by changing the function being optimized in order to settle on a specific choice of \mathbf{W} and \mathbf{F}). We may then use \mathbf{W} as a low-dimensional embedding of the original data (e.g. to facilitate interpretability or improve computational tractability), or we may choose to examine the factor images in \mathbf{F} to gain insights into the spatial structure of the data.

In related approaches such as PCA and ICA, the entries of \mathbf{W} and \mathbf{F} can, in principle, be any pattern of real numbers. In PCA, each row of \mathbf{F} is an eigenvector of the data covariance matrix, and \mathbf{W} is chosen to minimize the reconstruction error (i.e. to make \mathbf{WF} as close as possible to \mathbf{Y} in terms of mean squared error). In ICA, the goal is to minimize the statistical dependence between the rows of \mathbf{F} while also adjusting \mathbf{W} to minimize the reconstruction error. In this way, the factor images (the rows of \mathbf{F}) obtained using PCA and ICA are unstructured images (i.e. activity patterns) of the same complexity as individual observations in the original dataset: each factor is parameterized by V numbers (1 parameter per voxel).

In TFA (and TLSA), each row of \mathbf{F} is parameterized by the center parameter, μ , and the width parameter, λ , of an RBF. If an RBF has center μ and width λ , then its activity RBF($\mathbf{r}|\mu, \lambda$) at location \mathbf{r} is:

$$\text{RBF}(\mathbf{r}|\mu, \lambda) = \exp\left\{-\frac{\|\mathbf{r} - \mu\|^2}{\lambda}\right\}. \quad (2)$$

The factor images are filled in by evaluating each RBF, defined by the corresponding parameters for each factor, at the location of each voxel. In contrast to the factors obtained using PCA or ICA, TFA's more constrained factors may be represented much more compactly; each factor corresponds to the structure or group of structures in the brain over which the factor spreads its mass (which is governed by μ and λ). As highlighted above, TFA's factors may be conceptualized as nodes located in 3D space whose activity patterns influence the observed brain data.

HTFA works similarly to TFA, but places an additional constraint over the nodes to bias all of the subjects to exhibit similar nodes. Whereas TFA attempts to find the nodes that best explain an individual subject's data, HTFA also attempts to find the nodes that are common across a group of subjects (Fig. 1). This is important, because it allows the model to jointly consider data from multiple subjects.

HTFA (Fig. 2) handles multi-subject data by defining a *global template*, which describes in general where each RBF is placed, how wide it is, and how active its node tends to be. In addition to estimating how nodes look and behave in general (across subjects), HTFA also estimates each individual's *subject-specific template*, which describes each subject's particular instantiations of each RBF (i.e. that subject's RBF locations and widths) and the node activities (i.e. the activities of each of that subject's RBF nodes in each of that subject's images). These node activities, in turn, may be used to estimate each subject's full-brain functional connectivity patterns. Further, because the subject-specific templates are related to each other [hierarchically (Gelman and Hill, 2007), via the global template], a given node's RBF will tend to be located in about the same location, and be about as large, across all of the subject-specific templates. Because each subject has the same set of nodes (albeit in slightly different locations and with slightly different sizes) we can run analyses




Fig. 2. Graphical model for HTFA. Each variable in the model appears as a circle; hidden variables are unshaded and observed variables are shaded. The variables include $y_{n,s}$ (observed image n from subject s); $w_{n,k,s}$ (node k 's weight, or activity, in image n from subject s); $\mu_{k,s}$ (node k 's center coordinate for subject s); $\lambda_{k,s}$ (node k 's width for subject s); $\hat{\mu}_k$ (node k 's center in the global template); and $\hat{\lambda}_k$ (node k 's width in the global template). Hyperparameters (defined in the *Supplemental materials*) are denoted by dots. Arrows denote conditional dependence, originating at terms that appear on the right sides of conditionals and pointing towards terms that appear on the left sides. Rectangular plates denote repeated structure, where the number of copies is indicated within each plate: N_s (number of images from subject s); S (number of subjects), and K (number of nodes). For a comprehensive introduction to graphical models see Bishop (2006).

that relate node activity across subjects (as in the inter-subject functional connectivity analyses described below).

Materials and methods

Applying HTFA to multi-subject fMRI datasets

We use a *maximum a posteriori* (MAP) inference procedure to compute the most probable RBF nodes and node weights. The procedure has three basic steps: initialization (during which we set the prior over each node's parameters); fitting subject-specific node parameters for each subject (given the prior); and updating the global template (using the subject-specific parameters). When we carry out the full inference procedure,




Fig. 1. Hierarchical topographic factor analysis. **A. A brain image and its associated reconstruction.** The left sub-panel displays a single horizontal slice from a single subject; the right sub-panel displays its associated HTFA reconstruction, which we obtain by summing together the weighted images of the subject's RBF nodes. **B. Explaining data across subjects.** The left and middle sub-panels display example images from five subjects (left sub-panel), and their associated reconstructions (middle sub-panel). The right sub-panel displays the approximation of all of the single-subject images in the left sub-panel, obtained by setting the weights

of the global template's nodes to their average weights in the subject-specific reconstructions. The locations of the RBFs in the global template reflect commonalities across subjects, whereas the single-subject RBF locations reflect the associated subject-specific idiosyncrasies.

we first perform a cross-validation step (described in the next subsection) to determine the optimal number of nodes, K . We then randomly initialize K RBF nodes by drawing their parameters from the prior distribution (see *Supplemental materials*). Finally, we iterate between updating the subject-specific parameters (using the current global template as the prior) and the global template (using the latest estimates of the subject-specific parameters) until the largest change in any parameters value from the previous iteration to the current iteration is less than a pre-determined threshold value, ϵ . (We typically set ϵ to be the length of the longest voxel dimension.) Once the global centers and widths have converged, we run an additional step whereby we recompute the per-image node weights for each subject. We developed an efficient implementation of our algorithm for applying HTFA to large multi-subject fMRI datasets (Anderson et al., 2016); additional details may be found there. We also published a Python toolbox that enables users to apply HTFA to multisubject fMRI datasets (Capota et al., 2017).

Estimating the optimal number of nodes

When K (the number of nodes) equals V (the number of voxels), HTFA can exactly recover the data by setting the RBF widths to be very small and the per-image node activities equal to the voxel activities. Therefore setting $K = V$ represents one logical extreme whereby HTFA loses no information about the original data but also achieves no gains in how efficiently the data are represented. At the other extreme, when $K = 1$, HTFA achieves excellent compression but poor reconstruction accuracy (only a single activity value may be represented for each image). In practice, we will typically want to set K to some value between these extremes. Specifically, we want to choose the minimum K that is expected to explain the data up to a pre-defined level of precision, q .

Given a multi-subject dataset, we first select s_{training} subjects at random (from the full set) to participate in the cross validation procedure. For each of those subjects, we select (at random) a set of n_{training} images from each subject's data (here we set this number to be 70% of the subject's images). We fit HTFA to these randomly selected images to estimate the subject-specific node centers ($\mu_{1\dots K,1\dots s}$) and widths ($\lambda_{1\dots K,1\dots s}$). Next, of the remaining $n_{\text{test}} = N_s - n_{\text{training}}$ images for each subject, we select (at random) 70% of the voxels to estimate the per-image node weights, $\mathbf{w}_{1\dots N,1\dots K,1\dots s}$. Finally, we use the estimated centers, widths, and weights to reconstruct the voxel activities for the remaining 30% of the voxels in those n_{test} images. The mean squared error between the reconstructed and true (observed) voxel activities provides an error signal that we can use to optimize K . In particular, starting from a minimum value of $K = \delta_K$, we use the above procedure to compute the mean squared error for δ_K . We then increase K (in increments of δ_K) until the mean squared error is less than our pre-defined threshold, q . (In this paper we set $\delta_K = 100$ and $q = 0.25$.)

Inferring dynamic full-brain inter-subject functional connectivity patterns

Inference yields, for each subject, an N_s by K matrix, \mathbf{W}_s , of per-image node weights (i.e. activity patterns). We can estimate the functional connectivity between each pair of nodes by computing the correlation between the columns of \mathbf{W}_s . This approach is analogous to standard voxel-wise techniques for estimating functional connectivity (Biswal et al., 1995). Further, because the columns of $\mathbf{W}_{1\dots s}$ correspond to the same nodes across the different subjects, since all of the nodes are linked through the global template, the set of these weight matrices provide a convenient means of testing hypotheses related to the connectivity strengths.

In our analyses for Case Studies 2 and 3, we used *inter-subject functional connectivity* (ISFC; Simony et al., 2016) to isolate the time-varying correlational structure (functional connectivity patterns) that was specifically driven by the story participants listened to (Case Study 2) and television show participants watched (Case Study 3). We first applied

HTFA to the fMRI datasets to obtain a time series of node activities for every participant (where $K = 700$ for Case Study 2 and $K = 900$ for Case Study 3; see *Estimating the optimal number of nodes*). We obtained ISFC matrices for each of a series of overlapping temporal windows (*sliding windows*). To do so, we performed the following analysis for each sliding window (which contained a 90 s time series of node activities for each node and participant). For each participant, we computed the correlation between the activities of each node from that participant (during that sliding window) and the average activities of every node during the same window (where the average was taken across all of the other participants). The result, $C_{s,t}$ was a K by K correlation matrix for a single participant (s), during a single sliding window (t). We computed the ISFC matrix (across participants) during time t as:

$$\bar{C}_t = R \left(\frac{1}{2S} \sum_{s=1}^S Z(C_{s,t})^T + Z(C_{s,t}) \right), \quad (3)$$

where Z is the Fisher z -prime transformation (Zar, 2010):

$$Z(r) = \frac{\log(1+r) - \log(1-r)}{2} \quad (4)$$

and R is the inverse of Z :

$$R(z) = \frac{\exp(2z-1)}{\exp(2z+1)}. \quad (5)$$

For additional details and discussion of ISFC see Simony et al. (2016).

Decoding analysis


We asked whether the moment-by-moment HTFA-derived patterns we identified in Case Studies 2 and 3 were reliably preserved across participants, and (in Case Study 2) whether the degree of agreement across participants was modulated according to the cognitive salience of the stimuli participants experienced. For example, prior work has shown that different participants exhibit similar responses while experiencing richly structured stimuli (such as story listening), whereas participants exhibit less stereotyped responses while experiencing less structured stimuli (such as resting with their eyes open in the scanner; Simony et al., 2016).

To study these phenomena, we randomly divided the Case Study 2 participants into two groups, for each experimental condition: intact (i.e. participants who listened to the original story recording), paragraph-scrambled (i.e. participants who listened to an altered recording where the paragraphs occurred in a randomized order), word-scrambled (i.e. participants who listened to an altered recording where the words occurred in a randomized order), and rest (i.e. participants who rested with their eyes open in the scanner, without listening to any story). Participants within each condition experienced the same auditory stimuli, but the cognitive salience (i.e. how meaningful the stimuli were) varied systematically across these experimental conditions. (For the experiment presented in Case Study 3, every participant experienced the analog of the ‘‘intact’’ condition.)

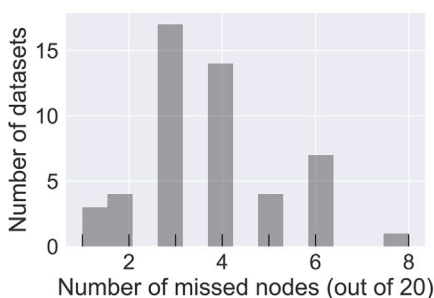
For each experimental condition, we computed the mean voxel, node, or factor activities within each sliding window (this resulted in either a V -dimensional or K -dimensional vector for each moment of the story). For each group of participants in turn, we compared these activity patterns (using Pearson correlations) to estimate the story times each pattern corresponded to. In the activity-based analyses shown in Fig. 4 (darker bars in each color group), we used these activity vectors to decode which moments of the story participants were listening to. Specifically, we asked, for each sliding window (t): what are the correlations between the first group's activity pattern at time t and the second group's activity patterns in every sliding window (this yielded one correlation value per sliding window). We used the best-matching pattern (i.e. the activity pattern with the strongest positive correlation) to estimate which story

A. Center recovery


True vs. recovered node centers for one synthetic dataset



Missed nodes across 50 synthetic datasets

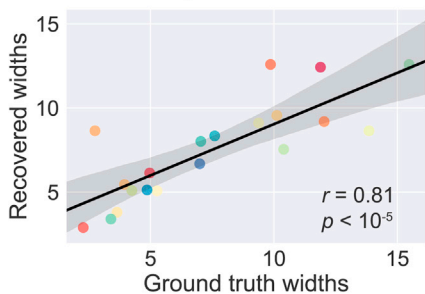


Distances between true/recovered node centers across 50 synthetic datasets




B. Width recovery

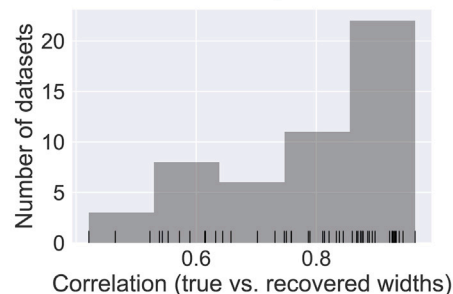
True vs. recovered node widths for one synthetic dataset



True vs. recovered node widths across 50 synthetic datasets

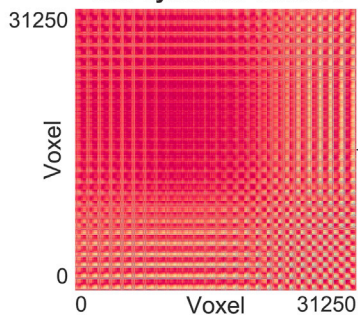


Correlations between true/recovered node widths across 50 synthetic datasets

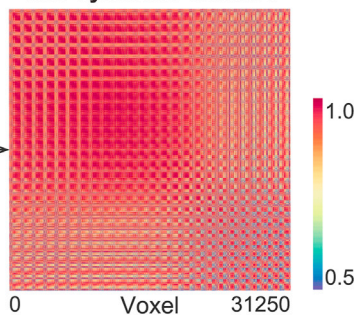


C. Across-voxel correlation recovery

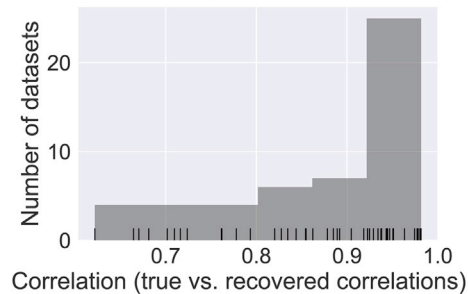
True voxel correlations for one synthetic dataset



Recovered voxel correlations for one synthetic dataset




Correlations between true/recovered voxel correlation matrices across 50 synthetic datasets




D. Across-image correlation recovery

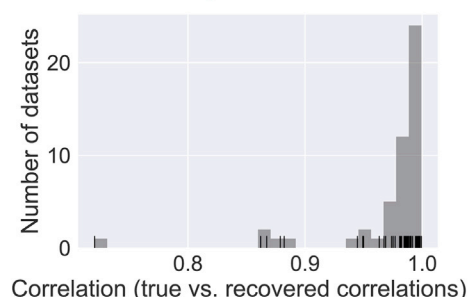
True image correlations for one synthetic dataset



Recovered image correlations for one synthetic dataset

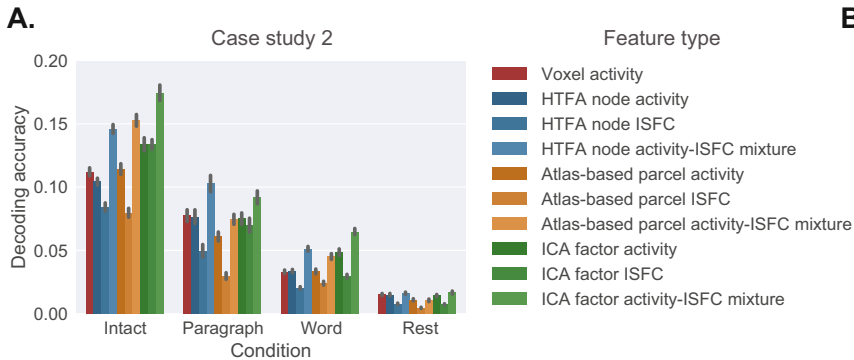


Correlations between true/recovered image correlation matrices across 50 synthetic datasets



continued on next page

Fig. 3. Recovered structure from 50 synthetic datasets. A. Center recovery. *Left.* Each color denotes a different node's RBF in the global template of one synthetic dataset. The filled circles indicate the true locations of the RBF centers, and the open circles indicate the RBF centers recovered by HTFA. To facilitate visual comparison between the true and recovered locations, we have drawn a line between each recovered RBF center and the closest matching RBF center in the synthetic template. One “missed” RBF that was not assigned any nodes by HTFA (as determined by using this matching technique) is denoted by a black \times . *Middle.* Number of missed nodes across 50 synthetic datasets. *Right.* Average distances between the true and recovered (global) node centers, across 50 datasets. The averages include distances between “missed” nodes and the nearest recovered node. **B. Width recovery.** *Left.* Each dot denotes an RBF's true and recovered (global) widths. The correlation reported in the panel is between the true and recovered RBF widths. To enable us to make comparisons in this panel, we have assigned each recovered node to the node in the original template with the closest RBF center (in terms of Euclidean distance). The colors match those in Panel A (left). *Middle.* This panel is in the same format as the left panel, but combines results across 50 synthetic datasets. The dot colors denote which dataset they came from (one color per dataset). *Right.* Correlations between true and recovered (global) widths across 50 synthetic datasets. **C. Across-voxel correlation matrices.** The left and middle panels display the true and recovered voxel-to-voxel correlation matrices for one synthetic dataset. The right panel displays a summary of the correlations between the true vs. recovered across-voxel correlation matrices across 50 synthetic datasets. **D. Across-image correlation matrices.** The left and middle panels display the true and recovered image-to-image correlation matrices for one synthetic dataset. The right panel displays a summary of the correlations between the true vs. recovered across-image correlation matrices across 50 synthetic datasets.



displays the decoding performance for Case Study 2 and panel B. displays the decoding performance for Case Study 3. All error bars denote bootstrap-estimated 95% confidence intervals, estimated using 5000 iterations.

time sliding window t corresponded to.

We used a similar approach to examine moment-by-moment ISFC patterns for each of the two groups of participants (i.e. for each condition, we obtained one ISFC pattern for each sliding window, for each of the two groups). For the ISFC analysis shown in Fig. 4 (medium-shaded bars in each color group), we reshaped these ISFC patterns into vectors, and used the same correlation-based technique to label each group's sliding windows according to how well they matched the ISFC patterns in the other group's sliding windows.

Finally, we carried out a mixed activity-based and ISFC-based decoding analysis by combining the estimates of the two above decoders. Specifically, for each sliding window (from one group of partic-

ipants), we computed the correlations between the activity patterns from each sliding window from the other group, and the correlations between the ISFC patterns from each sliding window of the other group. We used the average of these two correlations to label each group's timepoints in the “mixed” decoding analysis shown in Fig. 4 (lightly shaded bars in each color group). In Fig. 5 we extended this analysis by changing the relative weights of the activity-based and ISFC-based decoders using a mixing parameter, ϕ , where $\phi = 0$ corresponds to activity-based decoding; $\phi = 1$ corresponds to ISFC-based decoding; and $0 < \phi < 1$ reflects a weighted mixture of activity- and ISFC-based decoding. (Using this notation, we fixed $\phi = 0.5$ for all mixture analyses in Fig. 4.)

We note that the decoding test we used is more conservative than

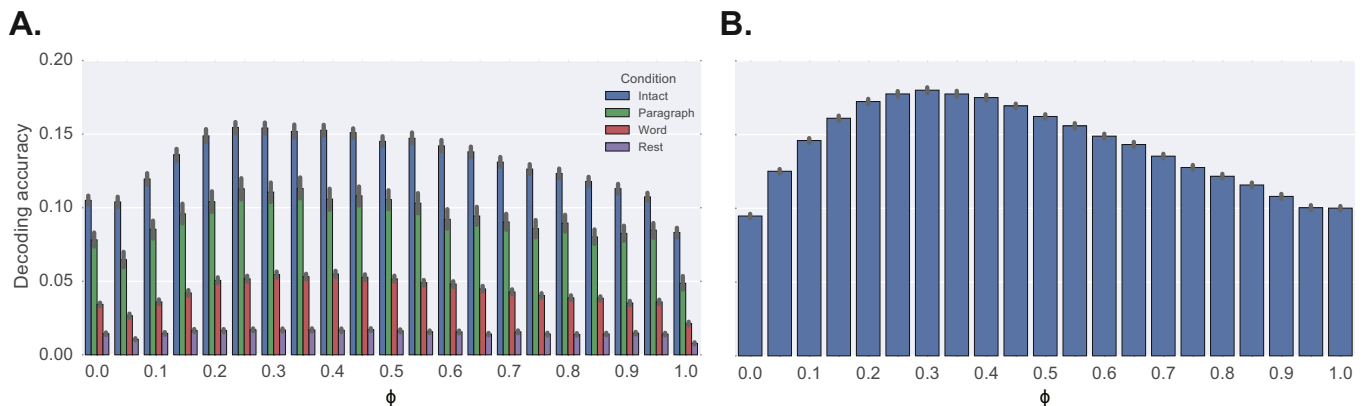


Fig. 5. Decoding accuracy as a function of mixing proportion (ϕ). The bar heights indicate the decoding accuracy achieved for each value of the mixing parameter, ϕ , where $\phi = 0$ corresponds to HTFA node activity-based decoding; $\phi = 1$ corresponds to ISFC-based decoding; and $0 < \phi < 1$ reflects a weighted mixture of activity- and ISFC-based decoding. Panel A. displays the decoding performance for Case Study 2 (colors denote the experimental condition) and panel B. displays the decoding performance for Case Study 3.

those used in some previously reported timepoint decoding studies (e.g. Haxby et al., 2011) because we count a timepoint label as incorrect if it is not an exact match, even if it overlaps substantially with the correct label. For example, if our decoder matches the 0–60 TR window from group 1 with the 1–61 TR window from group 2 (i.e. 88.5 of the 90 s are overlapping), our performance metric considers this to be a decoding failure, indistinguishable (performance-wise) from if the group 1 and group 2 windows had not overlapped at all; by contrast Haxby et al. (2011) used a more liberal procedure where they only compared the correct time window (e.g. 0–60 TR) to the exactly matching window or non-overlapping time windows (e.g. 61–120 TR), but not to partially overlapping windows (e.g. 1–61 TR). We chose to use this conservative test because our decoders attained over 99% accuracy for both voxel-based and HTFA-derived neural features (on data from the intact condition of the experiment in Case Study 2, as well as in Case Study 3) when we used the more liberal “standard” matching procedure. This made it difficult to achieve our goal of comparing and distinguishing between decoders trained on different neural features, hence our more conservative test.

Results

We applied HTFA to three types of data. In Case Study 1, we applied HTFA to synthetic data, and in Case Studies 2 and 3 we applied HTFA to fMRI data from human participants.

Case study 1: recovering known network dynamics from synthetic data

We generated 50 synthetic datasets, each comprising 200 brain images from each of 10 simulated subjects. Each image volume was a $50 \times 25 \times 25$ rectangular block of $31,250 \times 1 \times 2 \times 2$ mm voxels. To generate the voxel activities in each image, we randomly placed 20 RBF centers in a “template” brain volume (drawn randomly without replacement from the set of 31,250 voxel locations), and randomly assigned a positive width to each RBF [where $\log(\text{width}) \sim \mathcal{N}(\mu = 2\text{mm}, \sigma^2 = 0.5\text{mm})$]. We then generated each subject’s RBF centers and widths by adding a small amount of Gaussian noise to the corresponding parameters in the template [for each parameter dimension, noise $\sim \mathcal{N}(\mu = 0, \sigma^2 = 0.5\text{mm})$]. Next, we set the time-varying activities of each simulated subject’s RBF nodes to exhibit a pre-defined sequence of activity patterns. Specifically, the node activations were defined by a symmetric Toeplitz matrix whose first row was the sequence (1, 2, ..., 19, 20, 20, 19, ..., 2, 1), tiled 5 times to yield an activation for each node in each of the 200 images (Fig. 3D, left panel).

Our primary goal in examining the synthetic dataset was to assess whether the patterns (i.e. the RBF nodes and node weights) revealed by applying HTFA (with $K = 20$) to the dataset corresponded to the known patterns in the data. We first compared the global RBF centers and widths recovered by the model to the RBF centers and widths in the original template image. We found that the nodes’ RBF centers recovered by applying HTFA to the synthetic data closely matched the template’s RBF centers (Fig. 3A), and that the recovered RBF widths were correlated with the template RBF widths (Fig. 3B).

We next asked whether HTFA accurately recovered the correlational patterns in the synthetic data. We found that both the mean voxel-to-voxel correlation matrices (Fig. 3C) and the mean across-image correlation matrix (Fig. 3D) recovered by HTFA were strongly correlated with the true patterns we embedded into the synthetic data. (Recovery statistics for one example dataset, along with summaries of how well each pattern was recovered across all 50 datasets, are reported in the figure.) Taken together, these analyses show that HTFA is able to accurately infer (from synthetic data) the locations and sizes of the underlying RBF nodes, the correlation patterns within the images, and the correlations across images. We next turn to a series of analogous analyses on fMRI data.

Case study 2: full-brain network dynamics are modulated by story listening

We examined fMRI data collected as participants listened to an audio recording of a story (*intact* condition; 36 participants), listened to time-scrambled recordings of the same story (18 participants in the *paragraph-scrambled* condition listened to the paragraphs in a randomized order and 25 in the *word-scrambled* condition listened to the words in a randomized order), or lay resting with their eyes open in the scanner (*rest* condition; 36 participants). We sought to demonstrate how HTFA may be used to efficiently discover and examine dynamic functional connectivity patterns in (real) multi-subject fMRI datasets. This story listening dataset was collected as part of a separate study, where the full imaging parameters, image preprocessing methods, and experimental details may be found (Simony et al., 2016). The dataset is available at <http://arks.princeton.edu/ark:/88435/dsp015d86p269k>.

In contrast to the synthetic data we examined in Case Study 1, in real datasets there are no “ground truth” parameter values to compare to the recovered estimates. Instead, we sought to explore how well the patterns HTFA discovered could be used to decode which specific moments in the story participants were listening to. We also sought to explore whether (and how) decoding performance varied with the properties of the stimuli the participants experienced (see *Materials and methods* for a description of the decoding procedure). We first used a cross-validation procedure to determine the optimal number of nodes for efficiently representing the data while still capturing the relevant structure (see *Materials and methods*). We used all of the data from all of the experimental conditions (*intact*, *paragraph-scrambled*, *word-scrambled*, and *rest*) in this procedure— in other words, all of the experimental conditions were effectively lumped together into a single dataset. The analysis indicated that $K = 700$ nodes was optimal.

The fitted model provided estimates for each participants’ node locations, widths, and time-varying activities. Further, because the global model connects the subject-specific models, we were able to compare different participants’ activity patterns, even if their underlying nodes were not in exactly the same locations. We reasoned that, on one hand, this aspect of HTFA might improve our ability to capture cognitively relevant patterns (relative to examining the voxel activities directly). On the other hand, representing the brain images in the lower-dimensional space captured by HTFA necessarily results in information loss relative to the original voxel activity data. Effectively, HTFA blurs out high spatial frequency details from the images, where the precise amount of blurring depends on how large the nodes are and how many nodes there are overall.

We next evaluated how well HTFA is able to capture cognitively relevant brain patterns. We performed a decoding analysis, using cross-validation to estimate (using other participants’ data) which parts of the story each HTFA-derived brain activity pattern corresponded to (see *Materials and methods*). We note that our primary goal was not to achieve perfect decoding accuracy, but rather to use decoding accuracy as a benchmark for assessing whether different neural features specifically capture cognitively relevant brain patterns.

Separately for each experimental condition, we divided participants into two groups. We then computed the average activity for each group, for each of 241 overlapping 90 s (60 TR) time windows. (The 90 s window length we used in our analyses followed Simony et al., 2016.) This yielded one activity pattern for each group of participants, for each time window. Next, for each time window, we correlated the group 1 activity patterns in that window with the group 2 activity patterns. Using these correlations, we labeled the group 1 timepoints using the group 2 timepoints with which they were most highly correlated; we then computed the proportion of correctly labeled group 1 time windows. (We also performed the symmetric analysis whereby we labeled the group 2 timepoints using the group 1 timepoints as a template.) We repeated this procedure 100 times (randomly re-assigning participants to the two groups each time) to obtain a distribution of decoding accuracies for each experimental condition. (There were 241 time windows, so chance

performance on this decoding test is $\frac{1}{241}$.)

As a baseline, we first used the participants' full-brain voxel activity patterns (44,415 voxels; $Y_{1...s}$) to decode story timepoints. These voxel-based decoders achieved reliably above-chance performance on data from all four experimental conditions ($t_s(99) > 22, p < 10^{-400}$; Fig. 4), with the best average performance in the intact condition (11.2% accuracy) and the worst average performance in the rest condition (1.5% accuracy). This shows that the original data (that we applied HTFA to) contained information about the story times that our correlation-based decoders could pick up on. Our finding that the decoders performed slightly (but reliably) above-chance during rest was unexpected. This may reflect reliable changes either in participants' attentional states or properties of the BOLD signal over the course of the resting state scans.

We compared the performance of the voxel-based decoders to the decoding performance achieved using decoders trained on the HTFA node weights – i.e. the inferred timepoint-by-timepoint activities of the 700 nodes derived for each participant ($W_{1...s}$). Like the voxel-based decoders, these node activity-based decoders achieved reliably above-chance performance on data from all four experimental conditions ($t_s(99) > 21, p < 10^{-40}$; Fig. 4A). The decoding performance of the node-based and voxel-based decoders were similar (left 2 bars of each group in Fig. 4A; decoding accuracies match within 1%) and neither performed reliably better. We also ran analogous decoding analyses using parcel activities derived from resting state fMRI data (Shen et al., 2013) and factor weights derived from ICA with dual regression (Beckmann et al., 2009). Each of these activity-based decoders performed similarly (maximum difference in decoding accuracy across all experimental conditions and methods: 2.2%); for detailed comparison statistics see Table S1.

Although the above analysis shows that HTFA node activities achieve decoding performance similar to voxel activities, the real strength of HTFA is in its enabling efficient computations that involve dynamic functional connectivity patterns (whereas these computations are in some cases intractable in the original voxel space). Following the logic of Simony et al. (2016), we reasoned that brain activities during story listening should capture two sources of information. First, some activity should reflect the story itself. Because every participant (within each condition) listened to the same stimulus, this story-driven activity should be similar across people. Second, some activity might reflect idiosyncratic thoughts or physiological processes specific to each individual, independent of the story. This non-story-driven activity should not be similar across people. To home in on the former (story-driven) contribution to functional connectivity, we used ISFC (Simony et al., 2016). This approach measures the correlations between brain regions of different individuals in each of several sliding windows (see Materials and methods for details). The resulting ISFC patterns are analogous to standard within-brain functional connectivity patterns (which reflect the correlational structure, across brain regions, within an individual's brain), but they should reflect stimulus-driven activity. Decoders trained and tested on these HTFA-derived ISFC patterns achieved reliably above-chance performance on data from all experimental conditions ($t_s(99) > 10, p < 10^{-15}$; Fig. 4A). Although these ISFC-based decoders were out-performed by the voxel-based and node activity-based decoders, it is important to note that the correlation-based features that the ISFC-based decoders utilize are fundamentally different than node activity patterns. We carried out analogous analyses using the parcel activity-based decoders and ICA-based factor activation-based decoders described above. The parcel-based analysis yielded similar results to what we observed using HTFA (whereby ISFC-based decoders performed reliably worse than activity-based decoders), and the ICA-based analyses yielded no significant differences between ISFC-based and activity-based decoders (Table S1).

Given the above results, we wondered whether the node activity-based decoders and ISFC-based decoders might be picking up on partially non-overlapping sources of cognitively relevant information. If

so, decoders trained on a mix of activity-based and ISFC-based features might outperform decoders trained on only a single class of features.

We therefore designed a fourth set of decoders whose predictions were a 50-50 mix of the node activity-based decoders and the ISFC-based decoders; see Materials and methods for additional details. These hybrid decoders reliably out-performed the other decoders we examined on all experimental conditions except the rest condition (intact, paragraph, and word conditions: $t_s(99) > 14, p < 10^{-8}$; Fig. 4A). This finding is consistent with the notion that activity-based and ISFC-based patterns contain different information about the story moments people were listening to. We replicated this finding that a blend of activity-based and ISFC-based decoders outperform solely activity-based or ISFC-based decoders using atlas-derived parcel activations and ICA factor weights (Fig. 4A, Table S1).

To follow up on this result, we set out to determine the optimal mixing proportion of the two types of HTFA-derived features. We defined a mixing parameter, ϕ , where $\phi = 0$ corresponds to activity-based decoding; $\phi = 1$ corresponds to ISFC-based decoding; and $0 < \phi < 1$ reflects a weighted mixture of activity- and ISFC-based decoding. We then re-ran the "mixture" decoding analysis in Fig. 4A using 21 linearly spaced values of ϕ ranging between 0 and 1, inclusive. As shown in Fig. 5A, the decoding accuracy (for the intact condition) peaked for $\phi = 0.25$, corresponding to a mix of 75% activity-based and 25% ISFC-based features.

These analyses indicate that HTFA enables the efficient implementation of more sophisticated decoders that would otherwise be computationally expensive. These decoders (that incorporate correlational information, or a mix of activity-based and correlation-based information) out-perform decoders trained only on raw voxel activity patterns. Further, HTFA yields similar decoding performance to other established methods (see Table S1 for detailed comparisons).

Case study 3: full-brain networks are modulated by movie viewing

As a secondary test of our results from Case Study 2, we ran an analogous set of analyses on data collected as 17 participants viewed an episode from the BBC television show *Sherlock* (Chen et al., 2017). (Experimental and imaging methods may be found in Chen et al., 2017.) As in Case Study 2, we used a cross validation procedure to determine the optimal number of network nodes for the dataset (see Materials and methods). This procedure revealed that $K = 900$ nodes were optimal. We then applied HTFA to the full dataset, which comprised (for each participant) 43,371 voxel activities over 1976 TRs. The dataset is available at <http://arks.princeton.edu/ark:/88435/dsp01nzz8062179>.

As shown in Fig. 4B, all 10 classes of features we examined in Case Study 2 could also be used for this movie viewing dataset to reliably decode which timepoint participants were viewing (Fig. 4B; $t_s(99) > 78, p < 10^{-90}$). Following our approach in Case Study 2, we divided the experiment into 60 TR sliding windows. Therefore chance performance on this test is $\frac{1}{1917}$ ($1976 - 60 + 1 = 1917$ unique sliding windows).

We next carried out a series of analyses using HTFA, atlas-based parcels, and ICA factors to determine which aspects of the activity and connectivity patterns yielded the best decoding accuracy. As with Case Study 2, we found that mixed activity-ISFC decoders outperformed solely activity-based or ISFC-based decoders (for detailed performance statistics see Table S2).

We also carried out an exploratory analysis to estimate the optimal mixing proportions of HTFA node activity-based and ISFC-based features (Fig. 5B). The decoders with $\phi = 0.3$ (reflecting a mixture of 70% activity-based and 30% ISFC-based features) performed best. These analyses provide additional evidence that the activity and connectivity patterns contain partially non-overlapping information about the moments in the movie the participants were viewing.

Discussion

We proposed HTFA, a probabilistic approach to discovering and examining full-brain patterns of dynamic functional connectivity in multi-subject fMRI datasets. In Case Study 1, we used a synthetic dataset to demonstrate HTFA's ability to recover known patterns in synthetic data, and in Case Studies 2 and 3 we applied HTFA to real data, and showed that the resulting patterns could be used to decode story listening and movie viewing times. We also compared the performance of HTFA-derived decoders to decoders trained on features from other established methods, and we obtained similar decoding results.

Benefits and costs of our approach

HTFA was designed to provide an efficient means of examining dynamic full-brain connectivity patterns in multi-subject datasets. What would it have taken to study the same ISFC patterns we examined in Case Study 2 using voxel-based methods? After masking out non-gray matter voxels and warping every subject's data to a common brain space (see *Materials and methods*), each fMRI volume comprises 44,415 voxels. Therefore each voxelwise ISFC matrix contains roughly 986 million unique entries ($\frac{V^2-V}{2}$). Assuming single precision storage (4 bytes per number), the ISFC matrix for each timepoint would require roughly 3.9 GB of memory. By contrast, each HTFA-derived ISFC matrix for 700 nodes requires fewer than 250 thousand unique entries (1 MB). In general, for decoders that combine activity-based and ISFC-based features, the computational savings (in terms of representational compactness) gained by using HTFA-derived networks (of K nodes) rather than the original voxel data is given by $\frac{V^2+V}{K^2+K}$. For our Case Study 2, the HTFA-derived ISFC analyses in Fig. 4A were approximately 4000 times more computationally efficient than the corresponding voxel-based analyses would have been. Following this logic, we achieved a similar computational efficiency savings for Case Study 3; Fig. 4B. We note that these representational compactness benefits also apply to other dimensionality reduction approaches (e.g. see *Introduction*). Although HTFA and other dimensionality reduction approaches facilitate efficient computations of full-brain correlation patterns, we note that with sufficient computing power, optimized approaches have been developed for computing these matrices directly (Wang et al., 2015).

Another benefit of HTFA's separation from voxel space is its ability to naturally fill in missing observations (a property we exploit to determine the optimal number of nodes). Techniques like probabilistic PCA (Tipping and Bishop, 1999) can fill in missing voxel activities using the data covariance matrix, provided that we observe at least *some* activities from those missing voxels (in other images). However, suppose that *all* activities from a given voxel were missing—or more realistically, suppose that we wish to estimate what the activities would have been at any arbitrary point in space. Because PCA does not explicitly represent the voxels' spatial locations, neither PCA nor probabilistic PCA can accurately predict activity patterns at these never-observed voxels. HTFA, by contrast, naturally predicts the missing data by simply evaluating each node's RBF at the corresponding location in space.

These missing data examples also provide insights into other benefits of allowing nodes to exist in real space rather than considering only the set of voxel locations. For example, HTFA allows for different subjects' data to be sampled at different resolutions, or to contain different numbers of voxels. In principle, different subjects' data may even come from different recording modalities (e.g. one subject may contribute fMRI data and another may contribute EEG data). In this way, HTFA provides a common framework for describing neural data in general that transcends the specifics of the recording (modality, spatial or temporal resolution, etc.). For additional discussion of the benefits of spatial-based (rather than voxel-based) nodes see Gershman et al. (2011), and for an example of how similar models may be used to analyze EEG data see Gershman et al. (2014).

Note that using HTFA to examine connectivity patterns may not always out-perform voxel-based approaches. In particular, to the extent that the relevant patterns are high spatial frequency (at the level of single voxels), those patterns will be better described by voxel-based approaches than RBF nodes. (Representing brain images as sums of RBFs effectively blurs out the images in space, where the amount of blurring is inversely proportional to the number of nodes.) To address this issue, we developed an algorithm for estimating the minimum number of nodes required to reliably describe the data up to a desired level of precision (see *Materials and methods*).

Although we have focused in this manuscript on highlighting the strengths of HTFA in particular, it is worth noting the strengths of related approaches like the atlas-based segmentation and ICA with dual regression approaches that we compared with HTFA in Case Studies 2 and 3. For example, atlas-based segmentation methods obey anatomical boundaries well. This property is especially useful in applications aimed at discovering or exploring specific anatomical features or patterns. ICA-based approaches yield statistically independent features. This property is especially useful in applications that leverage or assume statistical independence between difference neural sources. HTFA, in turn, yields spatially compact features (which are useful for network visualizations and explorations) that partially overlap. This latter property may help detangle the contributions of spatially overlapping, but functionally distinct, activity patterns.

Concluding remarks

HTFA provides an efficient means of examining dynamic full-brain functional connectivity patterns, thereby making it easier to study how connectivity patterns relate to the cognitive processes they support. Further, HTFA's compact representations of connectivity facilitate studying connectivity with algorithms that are too expensive to use on the original data.

Acknowledgements

We acknowledge useful discussions with Jonathan Cohen, Janice Chen, Justin Hulbert, J. Benjamin Hutchinson, Talia Manning, Peter Ramadge, Erez Simony, and Nicholas Turk-Browne. This work was supported by the NSF/NIH Collaborative Research in Computational Neuroscience Program, grant number NSF IIS-1009542; NSF EPSCoR Award Number 1632738; and a grant from the Intel Corporation. The content is solely the responsibility of the authors and does not necessarily represent the official views of our supporting organizations. The implementation of HTFA that we used in this paper may be found in the BrainIAK Toolbox (<http://brainiak.org>); Capota et al. (2017).

Appendix A. Supplementary data

Supplementary data related to this article can be found at <https://doi.org/10.1016/j.neuroimage.2018.01.071>.

References

- Anderson, M.J., Capota, M., Turek, J.S., Zhu, X., Willke, T.L., Wang, Y., Chen, P.-H., Manning, J.R., Ramadge, P.J., Norman, K.A., 2016. Enabling factor analysis on thousand-subject neuroimaging datasets. In: *Proceedings of the IEEE International Conference on Big Data*, pp. 1242–1251.
- Beckmann, C.F., Mackay, C.E., Filippini, N., Smith, S.M., 2009. Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Neuroimage (Suppl. 1)*, S148.
- Betz, R.E., Medaglia, J.D., Papadopoulos, L., Baum, G., Gur, R., Roalf, D., Satterthwaite, T.D., Bassett, D.S., 2017. The modular organization of human anatomical brain networks: accounting for the cost of wiring. *Network Neuroscience* 1 (1), 42–68.
- Bishop, C., 2006. *Pattern Recognition and Machine Learning*. Springer.
- Biswal, B., Yetkin, F.Z., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* 34 (4), 537–541.

- Calhoun, V.D., Adali, T., Pearlson, G.D., Pekar, J.J., 2001. Spatial and temporal independent component analysis of functional MRI data containing a pair of task-related waveforms. *Hum. Brain Mapp.* 13 (1), 43–53.
- Cao, J., Worsley, K.J., 1999. The geometry of correlation fields, with an application to functional connectivity of the brain. *Ann. Appl. Probab.* 9 (4), 1021–1057.
- Capota, M., Turek, J., Chen, P.-H., Zhu, X., Manning, J.R., Sundaram, N., Keller, B., Wang, Y., Shin, Y.S., 2017. Brain Imaging Analysis Kit.
- Chen, J., Leong, Y.C., Norman, K.A., Hasson, U., 2017. Shared experience, shared memory: a common structure for brain activity during naturalistic recall shared experience, shared memory: a common structure for brain activity during naturalistic recall. *Nat. Neurosci.* 20 (1), 115–125.
- Comon, P., Jutten, C., Herault, J., 1991. Blind separation of sources, part II: problems statement. *Signal Process.* 24 (1), 11–20.
- Corde, D., Haughton, V.M., Carew, J.D., Arfanakis, K., Maravilla, K., 2002. Hierarchical clustering to measure connectivity in fMRI resting-state data. *Proc. Intl. Soc. Mag. Reson. Med.* 20 (4), 305–317.
- Craddock, R.C., James, G.A., Holtzheimer, P.E., Hu, X.P., Mayberg, H.S., 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum. Brain Mapp.* 33 (8), 1914–1928.
- Fox, M.D., Greicius, M., 2010. Clinical applications of resting state functional connectivity. *Front. Syst. Neurosci.* 4 (10), 1–13.
- Friston, K., Holmes, A., Worsley, K., Poline, J.-P., Frith, C., Frackowiak, R., 1994. Statistical parameter maps in functional imaging: a general linear approach. *Hum. Brain Mapp.* 2 (4), 189–210.
- Friston, K., Price, C., Fletcher, P., Moore, C., Frackowiak, R., Dolan, R., 1996. The trouble with cognitive subtraction. *Neuroimage* 4, 97–104.
- Friston, K.J., Frith, C.D., Liddle, P.F., Frackowiak, R.S.J., 1993. Functional connectivity: the principal-component analysis of large (PET) data sets. *J. Cerebr. Blood Flow Metabol.* 13 (1), 5–14.
- Gelman, A., Hill, J., 2007. *Data Analysis Using Regression and Multilevel/hierarchical Models*. Cambridge University Press.
- Gershman, S., Blei, D., Pereira, F., Norman, K., 2011. A topographic latent source model for fMRI data. *Neuroimage* 57, 89–100.
- Gershman, S.J., Blei, D.M., Norman, K.A., Sederberg, P.B., 2014. Decomposing spatiotemporal brain patterns into topographic latent sources. *Neuroimage* 98, 91–102.
- Gonzalez-Castillo, J., Hoy, C.W., Handwerker, D.A., Robinson, M.E., Buchanan, L.C., Saad, Z.S., Bandettini, P.A., 2015. Tracking ongoing cognition in individuals using brief, whole-brain functional connectivity patterns. *Proceedings of the National Academy of Science USA* 112 (28), 8762–8767.
- Greicius, M.D., Krasnow, B., Reiss, A.L., Menon, V., 2003. Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. *Proceedings of the National Academy of Science USA* 100 (1), 253–258.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., Pietrini, P., 2001. Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293, 2425–2430.
- Haxby, J.V., Guntupalli, J.S., Connolly, A.C., Halchenko, Y.O., Conroy, B.R., Gobbini, M.I., Hanke, M., Ramadge, P.J., 2011. A common, high-dimensional model of the representational space in human ventral temporal cortex. *Neuron* 72, 404–416.
- Honey, C.J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J.P., Meuli, R., Hagmann, P., 2009. Predicting human resting-state functional connectivity from structural connectivity. *Proceedings of the National Academy of Science USA* 106 (6), 2035–2040.
- Jutten, C., Herault, J., 1991. Blind separation of sources, part I: an adaptive algorithm based on neuromimetic architecture. *Signal Process.* 24 (1), 1–10.
- Manning, J.R., Ranganath, R., Norman, K.A., Blei, D.M., 2014. Topographic factor analysis: a Bayesian model for inferring brain networks from neural data. *PLoS One* 9 (5), e94914.
- Norman, K.A., Polyn, S.M., Detre, G.J., Haxby, J.V., 2006. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cognit. Sci.* 10 (9), 424–430.
- Pearson, K., 1901. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2, 559–572.
- Power, J.D., Cohen, A.L., Nelson, S.M., Wig, G.S., Barnes, K.A., Church, J.A., Vogel, A.C., Laumann, T.O., Miezin, F.M., Schlaggar, B.L., Peterson, S.E., 2011. Functional network organization of the human brain. *Neuron* 72, 665–678.
- Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. *Neuroimage* 52, 1059–1069.
- Shen, X., Tokoglu, F., Papademetris, X., Constable, R.T., 2013. Groupwise whole-brain parcellation from resting-state fMRI data for network node identification. *Neuroimage* 82, 403–415.
- Simony, E., Honey, C.J., Chen, J., Lositsky, O., Yeshurun, Y., Wiesel, A., Hasson, U., 2016. Dynamic reconfiguration of the default mode network during narrative comprehensions. *Nat. Commun.* 7 <https://doi.org/10.1038/ncomms12141> (12141).
- Spearman, C., 1904. General intelligence, objectively determined and measured. *American Journal of Psychology* 15, 201–292.
- Tipping, M.E., Bishop, C.M., 1999. Probabilistic principal component analysis. *J. Roy. Stat. Soc. B* 61 (3), 611–622.
- Turk-Browne, N.B., 2013. Functional interactions as big data in the human brain. *Science* 342, 580–584.
- Valdés-Sosa, P.A., Sánchez-Bornot, J.M., Lage-Castellanos, A., Vega-Hernández, M., Bosch-Bayard, J., Melie-García, L., Rodríguez, E.C., 2005. Estimating brain functional connectivity with sparse multivariate autoregression. *Philosophical Transactions of the Royal Society B* 360, 969–981.
- Wang, Y., Cohen, J.D., Li, K., Turk-Browne, N.B., 2015. Full correlation matrix analysis (FCMA): a high-performance toolbox and case study for unbiased functional connectivity in human brain imaging. *J. Neurosci. Meth.* 251, 108–119.
- Worsley, K.J., Chen, J.-I., Lerch, J., Evans, A.C., 2005. Comparing functional connectivity via thresholding correlations and singular value decomposition. *Philosophical Transactions of the Royal Society B* 360, 913–920.
- Yeo, B.T.T., Krienen, F.M., Sepulcre, J., Sabuncu, M.R., Lashkari, D., Hollinshead, M., Roffman, J.L., Smoller, J.W., Zolke, L., Polimieni, J.R., Fischl, B., Liu, H., Buckner, R.L., 2011. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of Neurophysiology* 106 (3), 1125–1165.
- Zar, J.H., 2010. *Biostatistical Analysis*. Prentice-Hall/Pearson.
- Zarahn, E., Aguirre, G., D'Esposito, M., 1997. A trial-based experimental design for fMRI. *Neuroimage* 6, 122–138.