Hierarchical Topographic Factor Analysis

Jeremy R. Manning, Rajesh Ranganath, Waitsang Keung,
Nicholas B. Turk-Browne, Jonathan D. Cohen, Kenneth A. Norman, and David M. Blei
Princeton University

***Abstract*—Recent work has revealed that cognitive processes are often reflected in patterns of functional connectivity throughout the brain (for review see [16]). However, examining functional connectivity patterns using traditional methods carries a substantial computational burden (of computing time and memory). Here we present a technique, termed Hierarchical Topographic Factor Analysis, for efficiently discovering brain networks in large multi-subject neuroimaging datasets.**

I. Introduction

The most common approaches for analyzing functional Magnetic Resonance Imaging (fMRI) data involve relating the activity of individual voxels or multi-voxel spatial patterns of brain activity to the subject's cognitive state [4], [5], [11], [21]. In contrast, functional connectivity analyses estimate connections between brain regions by correlating the time series of activations (across images) of each pair of voxels [14]. This cutting-edge approach has already led to important new insights into how the brain's connections change during different experimental conditions [16].

Because the connectivity matrix grows with the square of the number of voxels, both filling in the entries of a connectivity matrix and storing it in memory can become intractable for fMRI images with tens of thousands of voxels. For example, the connectivity matrix for a 50,000 voxel image occupies approximately 5 GB of memory (assuming single precision floating point entries). Storing many such matrices in memory (e.g. to compare different subjects and/or experimental conditions) can be impractical on modern hardware (although see [18] for an alternative approach).

Topographic Factor Analysis (TFA) provides an alternative means of representing the brain's connectivity patterns that scales well to large images [9]. TFA casts each of N brain images as a linear combination of latent sources [Gaussian radial basis functions (RBFs)]. Each source can be interpreted as a node in a simplified representation of the brain's network. (The number of sources, K, is chosen by the practitioner.) Applying TFA to an fMRI dataset reveals the locations and sizes of the sources (i.e. the centers and widths of their RBFs), as well as the per-image source weights. In this way, the N by K matrix of source weights may be viewed as a low-dimensional embedding of the original dataset. Further, the covariance of the weight matrix can be taken as a proxy for the signs and strengths of the node-to-node connections.

One limitation of TFA is that it treats each subject's data as independent. This means that all hypothesis testing (e.g. on connection strengths that distinguish between experimental conditions) must be performed within-subject. Here we present Hierarchical TFA (HTFA), which extends the TFA model to incorporate data from multiple subjects. HTFA uses an entire multi-subject dataset to learn a *global template* of source centers, widths, and weights that describe how people's brains look and behave in general. Then we cast each individual subject's sources as a perturbation of that global template. This allows us to compare how a given source behaves in one subject's brain to how that same source behaves in another subject's brain. We can thereby test whether particular patterns of activation, or particular source-source interactions, vary reliably across subjects in different conditions of an experiment. The hierarchical implementation also has the advantage that, in theory, ambiguities in one subject's data may be resolved by examining the other subjects' data.

In the next section we formally define the HTFA model and describe how we fit the model's parameters to large fMRI datasets. We then use HTFA to discover function-dependent brain networks in three fMRI datasets.

978-1-4799-4149-0/14/$31.00

II. Methods

HTFA comes from a family of models, called *factor analysis* models, that also includes TFA [9], Topographic Latent Source Analysis (TLSA) [7], Principal Components Analysis (PCA) [12], Exploratory Factor Analysis (EFA) [15], and Independent Components Analysis (ICA) [3], [8], among others. If we have organized our collection of images (from a single subject) into an N by V 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:

$$\mathbf{Y} \approx \mathbf{WF}, \tag{1}$$

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 looks). The general idea is that different techniques place different constraints 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 computational tractability), or examine the factor images in $\mathbf{F}$ to gain insights into the dataset.

In standard approches such as PCA and ICA, the entries of $\mathbf{W}$ and $\mathbf{F}$ are real numbers. In PCA, each row of $\mathbf{F}$ is an eigenvector of the data covariance matrix, and the weights ($\mathbf{W}$) are chosen to minimize the reconstruction error (i.e. to make $\mathbf{WF}$ as close as possible to $\mathbf{Y}$). 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. activation patterns), of the same complexity as observations in the original dataset.

In TLSA (and TFA, which is a special case of TLSA), each row of $\mathbf{F}$ is parameterized by the center parameter, μ,

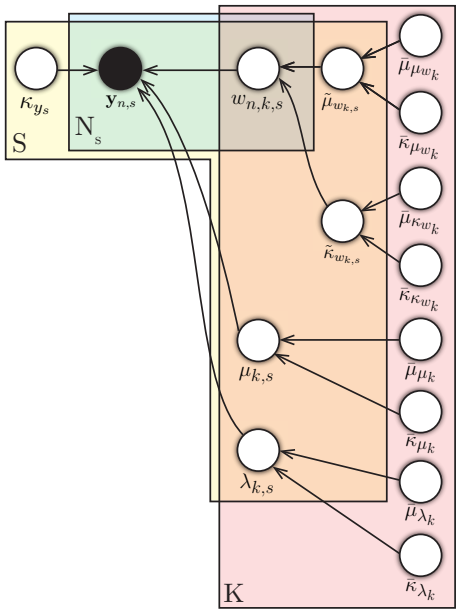

Fig. 1. **Graphical model for HTFA.** Each variable in the model appears as a circle; hidden variables are unshaded and observed variables are shaded. Arrows denote conditional dependence, originating at terms that appear on the right sides of conditionals and pointing towards terms that appear on the left sides. Rectangular plates denote repeated structure, where the number of copies is indicated within each plate (e.g. N_s, S, or K). For a comprehensive introduction to graphical models see [2]. Variables are defined in Algorithm 1.

and width parameter, λ, of an RBF. If an RBF has center μ and (log) width λ, then its activation $\mathrm{RBF}(\mathbf{r}|\mu,\lambda)$ at location $\mathbf{r}$ is given by:

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

The factor images, termed *sources*, 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, TLSA's more constrained factors are easily interpretable: 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). While constraining the factors in this way reduces reconstruction accuracy, it substantially improves interpretability. For example, PCA's factors are not constrained (i.e. each factor's mass need not be concentrated in one brain region), and therefore explain the data well but are difficult to interpret without additional processing. HTFA works similarly to TFA, but places an additional constraint over the factors to bias all of the subjects to exhibit similar factors. In this way, whereas TFA attempts to find the factors that best explain an individual subject's data, HTFA attempts to find the factors that are common across a group of subjects.

We formulate HTFA as a probabilistic model, which can be represented in graphical model notation (Fig. 1). The graphical model specifies the conditional dependencies in the model:

$$p(\mathbf{Y}_{1\ldots S}, \mathbf{\Omega}) = LG, \tag{3}$$

where $\mathbf{Y}_s$ is the N_s by V_s matrix of images from subject s,[1] $\mathbf{\Omega}$ is the set of hidden variables in the model

$$\begin{aligned}\mathbf{\Omega} = \{&\sigma^2_{y_{1\ldots S}}, w_{1\ldots N,1\ldots K,1\ldots S}, \mu_{1\ldots K,1\ldots S}, \lambda_{1\ldots K,1\ldots S},\\ &\tilde{\mu}_{w_{1\ldots K,1\ldots S}}, \tilde{\kappa}1_{1\ldots K,1\ldots S},\\ &\bar{\mu}_{\mu_{w_{1\ldots K}}}, \bar{\kappa}_{\mu_{w_{1\ldots K}}}, \bar{\mu}_{\kappa_{w_{1\ldots K}}}, \bar{\kappa}_{\kappa_{w_{1\ldots K}}}, \bar{\mu}_{\mu_{1\ldots K}}, \bar{\kappa}_{\mu_{1\ldots K}},\\ &\bar{\mu}_{\lambda_{1\ldots K}}, \bar{\kappa}_{\lambda_{1\ldots K}}\},\end{aligned} \tag{4}$$

the probability of the subject-specific template L is given by

$$\begin{aligned}L = &\prod_{s=1}^{S} p(\sigma^2_{y_s}) \prod_{n=1}^{N_s} p(\mathbf{y}_{n,s}|w_{n,1\ldots K,s}, \mu_{1\ldots K,s}, \lambda_{1\ldots K,s})\\ &\prod_{k=1}^{K} p(w_{n,k,s}|\tilde{\mu}_{w_{k,s}}, \tilde{\kappa}_{w_{k,s}})\\ &p(\tilde{\mu}_{w_{k,s}}|\bar{\mu}_{\mu_{w_k}}, \bar{\kappa}_{\mu_{w_k}}) p(\tilde{\kappa}_{w_{k,s}}|\bar{\mu}_{\kappa_{w_k}}, \bar{\kappa}_{\kappa_{w_k}})\\ &p(\mu_{k,s}|\bar{\mu}_{\mu_k}, \bar{\kappa}_{\mu_k}) p(\lambda_{k,s}|\bar{\mu}_{\lambda_k}, \bar{\kappa}_{\lambda_k}),\end{aligned} \tag{5}$$

and where the probability of the global template G is given by

$$\begin{aligned}G = &\prod_{k=1}^{K} p(\bar{\mu}_{\mu_{w_k}}) p(\bar{\kappa}_{\mu_{w_k}}) p(\bar{\mu}_{\kappa_{w_k}}) p(\bar{\kappa}_{\kappa_{w_k}}) p(\bar{\mu}_{\mu_k}) p(\bar{\kappa}_{\mu_k})\\ &p(\bar{\mu}_{\lambda_k}) p(\bar{\kappa}_{\lambda_k}).\end{aligned} \tag{6}$$

(Note that the hyperparameters have been omitted for notational simplicity.)

In the graphical model (Fig. 1), variables associated with the subject-specific template are found in the yellow plate. These include the subject-specific source centers, widths, and weights, as well as the observed images.[2] Variables associated with the global template are found outside of the yellow plate; these include the means and log precisions of the source center, width, and weight distributions. The subject-specific templates are conditioned on the global template, thereby associating data from different subjects. (This interaction between the subject-specific and global templates occurs in the orange region of the graphical model.)

Our goal is to infer a *posterior distribution* over the model's hidden variables (e.g. source centers, widths, and weights). The posterior distribution $p(\mathbf{\Omega}|\mathbf{Y}_{1\ldots S})$ tells us how likely each hidden variable is to be set to a particular value, given the data.

To begin to compute the posterior, we first write down HTFA's *generative process* (Alg. 1); this algorithm defines a recipe for producing the observations (in this case, an fMRI dataset). Each time we run the generative process, we obtain a single sample of each of the model's variables (e.g. the source centers, widths, and weights for each subject) and a synthetic dataset. When we fit HTFA to a real dataset, we "reverse" the generative process by computing the most probable settings of

[1] Note that both the number of images from subject s, N_s, and the number of voxels from subject s, V_s, may vary from subject to subject.

[2] We draw a distinction between the subject-specific *per-image* source weights, $w_{1\ldots N,1\ldots K,1\ldots S}$, and the subject-specific distributions from which they are drawn. Specifically, the per-image weights are drawn from a distribution whose mean is $\tilde{\mu}_{w_{1\ldots K,1\ldots S}}$ and whose log precision is $\tilde{\kappa}_{w_{1\ldots K,1\ldots S}}$. In other words, each source has an associated subject-specific mean that governs how much (or little) that source will be expressed across images, and a subject-specific log precision that governs how variable the source's weights will be across images.

each variable given the data. In effect, we tune the parameters of the model's generative process until the data it produces look like our observations. For example, the generative process posits that source k's center location is drawn (independently for each subject) from a Multivariate Gaussian distribution centered on source k's center location in the global template. The goal of posterior inference is to determine which particular subject-specific source centers, widths, and weights were most probably sampled from the global template, and what the global template looks like, given the observed brain images.

Algorithm 1: HTFA's generative process. Note that we parameterize the variances of the Gaussian distributions using log precision parameters (equal to the log of the inverse of the variance). This parameterization is equivalent to the usual one, and facilitates our approximate inference algorithm. We use $\mathrm{RBF}_s(\cdot)$ to denote a vector obtained by evaluating an RBF with the given parameters at the locations of each of subject s's voxels.

for $k = 1$ **to** K **do**
  Pick mean of template weight distribution for source k: $\bar{\mu}_{w_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\mu}_w}, \hat{\kappa}_{\bar{\mu}_w}\right)$;
  Pick precision of template weight distribution for source k: $\bar{\kappa}_{w_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\kappa}_w}, \hat{\kappa}_{\bar{\kappa}_w}\right)$;
  Pick mean of template center distribution for source k (repeat for each dimension; not shown in graphical model): $\bar{\mu}_{\mu_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\mu}_\mu}, \hat{\kappa}_{\bar{\mu}_\mu}\right)$;
  Pick precision of template center distribution for source k (repeat for each dimension; not shown in graphical model): $\bar{\kappa}_{\mu_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\kappa}_\mu}, \hat{\kappa}_{\bar{\kappa}_\mu}\right)$;
  Pick mean of template width distribution for source k: $\bar{\mu}_{\lambda_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\mu}_\lambda}, \hat{\kappa}_{\bar{\mu}_\lambda}\right)$;
  Pick precision of template width distribution for source k: $\bar{\mu}_{\lambda_k} \sim \mathcal{N}\left(\hat{\mu}_{\bar{\mu}_\lambda}, \hat{\kappa}_{\bar{\mu}_\lambda}\right)$;
  for $s = 1$ **to** S **do**
    Pick mean of subject-specific weight distribution for source k: $\tilde{\mu}_{w_{k,s}} \sim \mathcal{N}\left(\bar{\mu}_{\mu_{w_k}}, \bar{\kappa}_{\mu_{w_k}}\right)$;
    Pick precision of subject-specific weight distribution for source k: $\tilde{\kappa}_{w_{k,s}} \sim \mathcal{N}\left(\bar{\mu}_{\kappa_{w_k}}, \bar{\kappa}_{\kappa_{w_k}}\right)$;
    Pick center of subject-specific source k (repeat for each dimension; not shown in graphical model): $\mu_{k,s} \sim \mathcal{N}\left(\bar{\mu}_{\mu_k}, \bar{\kappa}_{\mu_k}\right)$;
    Pick width of subject-specific source k: $\lambda_{k,s} \sim \mathcal{N}\left(\bar{\mu}_{\lambda_k}, \bar{\kappa}_{\lambda_k}\right)$;
    for $n = 1$ **to** N_s **do**
      Pick weight of source k in image n of subject s: $w_{n,k,s} \sim \mathcal{N}\left(\tilde{\mu}_{w_{k,s}}, \tilde{\kappa}_{w_{k,s}}\right)$;
    end
  end
end
for $s = 1$ **to** S **do**
  Pick subject-specific noise parameter: $\kappa_{y_s} \sim \mathcal{N}\left(\hat{\mu}_{\sigma_y^2}, \hat{\kappa}_{\sigma_y^2}\right)$;
  for $n = 1$ **to** N_s **do**
    Pick image n for subject s: $\mathbf{y}_{n,s} \sim \mathcal{N}\left(\sum_{k=1}^{K} w_{n,k,s}\mathrm{RBF}_s(\mu_{k,s}, \lambda_{k,s}), \kappa_{y_s}\mathbf{I}^V\right)$;
  end
end

In theory, we could compute the posterior using Bayes' rule (e.g. [6]):

$$p(\mathbf{\Omega}|\mathbf{Y}_{1...S}) = \frac{p(\mathbf{Y}_{1...S}|\mathbf{\Omega})p(\mathbf{\Omega})}{\int_{\mathbf{\Omega}} p(\mathbf{\Omega}, \mathbf{Y}_{1...S})d\mathbf{\Omega}}. \quad (7)$$

However, computing the denominator is intractable, as it would require (numerically) integrating over all possible settings of the hidden variables in the model. Therefore, rather than computing the posterior exactly, we instead approximate it using Black Box Variational Inference [13] (the full details of this technique are beyond the scope of this document due to space constraints, but are equivalent to those in [9]). Running this inference procedure yields an approximation of the posterior, $q(\mathbf{\Omega})$, which we use as a proxy for the true posterior $p(\mathbf{\Omega}|\mathbf{Y}_{1...S})$.

III. Results

Applying HTFA to a multi-subject fMRI dataset reveals the locations and widths of K sources that are reflected in general across subjects (in a global template) and in individual subjects (whose sources are modeled as perturbations of the global template). We also obtain, for each subject, a set of K source weights for each of that subject's images. We use the K by K covariance matrix of each subject's N_s by K weight matrix[3] to estimate the connection strengths between the sources. Because HTFA associates sources across subjects (via the global template), this gives us a means of comparing connectivity across subjects. In other words, source k in subject x will "mean" the same thing as source k in subject $y \neq x$. Specifically, the same source number will be located in about the same location, and exhibit about the same (average) weights, across different subjects. We can also use this technique to infer functional connectivity by examining how the source weights covary within each experimental condition.

We sought to evaluate our approach by applying HTFA to three fMRI datasets. For each dataset we computed, for each subject, the source-to-source connectivity matrices for each condition of the experiment. We then asked whether any connections varied their strengths between conditions in a similar way across subjects. Our approach was intended to assess whether the connectivity matrices identified by HTFA were stable (i.e. similar across subjects).

The first dataset, which we refer to as the *Face/Scene Dataset*, contains data from 18 subjects who viewed images of male and female faces and indoor and outdoor scenes (the data were from a functional localizer task collected by [17]). Each testing session lasted 6 min 6 s, and was organized into 12 blocks (6 "face" blocks and 6 "scene" blocks). During each block, 12 stimuli were presented for 500 ms, every 1500 ms (with a 12 s fixation occurring between each block). As the images were displayed, the subjects made judgements (by pressing buttons) about whether the faces were male or female, and whether the scenes images were of indoor or outdoor locations.

Previous studies of this dataset [18], [19] had revealed that, whereas multivariate pattern classification analyses [11] typically identify only posterior regions (e.g. fusiform face area, parahippocampal place area) as distinguishing face and scene images, functional connectivity analyses reveal a more complete picture– specifically, frontal regions vary their connectivity with these posterior regions differently while viewing the two classes of stimuli. To examine these patterns using HTFA, we estimated face and scene connectivity matrices for each subject. We then used t-tests to compare the distributions of connectivity strengths during face vs. scene viewing. As shown in Figure 2A, HTFA (qualitatively) recovers the previously reported connectivity patterns reflected in the dataset.

We refer to the second fMRI dataset we examined as the *Tasks Dataset*. The Tasks Dataset (collected by W. Keung and

[3] We use the *maximum a posteriori* (MAP) per-image source weights for this computation, as in [9].

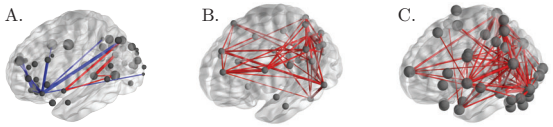

Fig. 2. **Function-dependent brain networks discovered using HTFA. A. Face/Scene Dataset.** Connections that were reliably stronger when subjects viewed images of faces are shown in blue, and connections that were reliably stronger when subjects viewed images of scenes are shown in red. Connections whose strengths did not reliably distinguish between the two conditions (after correcting for multiple comparisons using a $p = 0.05$ threshold [1]) are omitted. **B. Tasks Dataset.** Connections whose strengths reliably distinguished between *any* pair of tasks, after correcting for multiple comparisons, are shown in red. Connections that did not reliably vary their strengths across conditions are omitted. **C. Nouns Dataset.** Connections whose strengths reliably distinguished between *any* pair of categories, after correcting for multiple comparisons, are shown in red. Connections that did not reliably vary their strengths across categories are omitted. This figure was created using BrainNet viewer [20].

J. D. Cohen, *unpublished*) contains data from 18 subjects who performed a battery of 12 mini-tasks for 20 minutes (where each task lasted between 30 and 60 seconds and was sampled at least twice). The full set of tasks incorporated a number of classic paradigms, including the n-back memory task, Tower of London and scheduling tasks, and the Stroop task.

Our primary objective in examining the Tasks Dataset was to discover connections whose strengths varied reliably between *any* of the different task-types. We used HTFA to estimate, for each subject, a connectivity matrix for each of the 12 tasks. We then used repeated measures analyses of variance to detect pairs of nodes whose connection strengths varied from task to task in a similar way across subjects (Fig. 2B).

We also examined a third fMRI dataset, collected by [10], which we refer to as the *Nouns Dataset*. The dataset comprises data from 9 participants who each viewed 6 presentations of each of 60 line drawings, for a total of 360 viewings (each with an associated brain image). The drawing presentations were organized into six *epochs*, where all 60 drawings were presented in a random order during each epoch. The participants were instructed to think about the meaning of the word associated with each drawing as they viewed it. The drawings were selected from 12 categories: animals, body parts, buildings, building parts, clothing, furniture, insects, kitchen items, man made objects, tools, vegetables, and vehicles. We used a similar analysis to the one we carried out for the Tasks Dataset to identify connections whose strengths distinguished between categories (Fig. 2C).

IV. Concluding Remarks

We presented HTFA, a hierarchical model for efficiently discovering brain networks in large multi-subject fMRI datasets. We applied HTFA to three fMRI datasets as a proof of concept. HTFA automatically discovered stable networks (i.e. connections that, across subjects, varied in similar ways between experimental conditions) in all three datasets. In the first dataset (the Face/Scene Dataset), our results were similar to those obtained using a related approach, called *full correlation matrix analysis* (FCMA) that uses matrix algebra tricks and massive parallelization to exhaustively and efficiently compare the full (voxel-by-voxel) brain connectivity matrices across subjects between experimental conditions [18]. For the other datasets we examined (the Tasks Dataset and the Nouns Dataset), the fact that HTFA found stable task-specific and category-specific brain networks is promising. We view HTFA as a useful tool that will enable neuroscientists to perform efficient exploratory analyses of functional connectivity in fMRI datasets. A MATLAB toolbox implementing our model may be downloaded from `http://tinyurl.com/lynhm95`.

References

[1] Y. Benjamini and Y. Hochberg, "Controlling the False Discovery Rate: a practical and powerful approach to multiple testing," *Journal of Royal Statistical Society, Series B*, vol. 57, pp. 289–300, 1995.

[2] C. Bishop, *Pattern recognition and machine learning*. Springer, 2006.

[3] P. Comon, C. Jutten, and J. Herault, "Blind separation of sources, part ii: Problems statement," *Signal Processing*, vol. 24, no. 1, pp. 11 – 20, 1991.

[4] K. Friston, A. Holmes, K. Worsley, J.-P. Poline, C. Frith, and R. Frackowiak, "Statistical parameter maps in functional imaging: a general linear approach," *Human Brain Mapping*, vol. 2, pp. 189 – 210, 1995.

[5] K. Friston, C. Price, P. Fletcher, C. Moore, R. Frackowiak, and R. Dolan, "The trouble with cognitive subtraction," *NeuroImage*, vol. 4, pp. 97 – 104, 1996.

[6] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian data analysis, Third edition*. Chapman & Hall, 2013.

[7] S. Gershman, K. Norman, and D. Blei, "A topographic latent source model for fMRI data," *NeuroImage*, vol. 57, pp. 89 – 100, 2011.

[8] C. Jutten and J. Herault, "Blind separation of sources, part i: An adaptive algorithm based on neuromimetic architecture," *Signal Processing*, vol. 41, no. 1, pp. 1 – 10, 1991.

[9] J. R. Manning, R. Ranganath, K. A. Norman, and D. M. Blei, "Topographic factor analysis: a Bayesian model for inferring brain networks from neural data," *PLoS One*, 2014.

[10] T. Mitchell, S. Shinkareva, A. Carlson, K. Chang, V. Malave, R. Mason, and M. Just, "Predicting human brain activity associated with the meanings of nouns," *Science*, vol. 320, no. 5880, p. 1191, 2008.

[11] K. A. Norman, S. M. Polyn, G. J. Detre, and J. V. Haxby, "Beyond mind-reading: multi-voxel pattern analysis of fMRI data," *Trends in Cognitive Sciences*, vol. 10, no. 9, pp. 424–430, 2006.

[12] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosophical Magazine*, vol. 2, pp. 559–572, 1901.

[13] R. Ranganath, S. Gerrish, and D. M. Blei, "Black box variational inference," in *Artificial Intelligence and Statistics*, 2014.

[14] M. Rubinov and O. Sporns, "Complex network measures of brain connectivity: uses and interpretations," *NeuroImage*, vol. 52, pp. 1059 – 1069, 2010.

[15] C. Spearman, "General intelligence, objectively determined and measured," *Americal Journal of Psychology*, vol. 15, pp. 201 – 293, 1904.

[16] N. B. Turk-Browne, "Functional interactions as big data in the human brain," *Science*, vol. 342, pp. 580 – 584, 2013.

[17] N. B. Turk-Browne, M. G. Simon, and P. B. Sederberg, "Scene representations in parahippocampal cortex depend on temporal context," *Journal of Neuroscience*, vol. 32, no. 21, pp. 7202 – 7207, 2012.

[18] Y. Wang, J. D. Cohen, K. Li, and N. B. Turk-Browne, "Full correlation matrix analysis (FCMA): a high-performance toolbox and case study for unbiased functional connectivity in human brain imaging," *Submitted*, 2014.

[19] Y. Wang, K. Li, J. D. Cohen, and N. B. Turk-Browne, "What you see depends on how you look: Category-selective interactions with frontal cortex during object perception," *Submitted*, 2014.

[20] M. Xia, J. Wang, and Y. He, "BrainNet viewer: a network visualization tool for human brain connectomics," *PLoS One*, vol. 8, p. e68910, 2013.

[21] E. Zarahn, G. Aguirre, and M. D'Esposito, "A trial-based experimental design for fMRI," *NeuroImage*, vol. 6, pp. 122 – 138, 1997.