Autoencoding Topographic Factors

ANTONIO MORETTI,^{1,*} ANDREW STIRN,^{1,*} GABRIEL MARKS,² and ITSIK PE'ER¹

ABSTRACT

Topographic factor models separate overlapping signals into latent spatial functions to identify correlation structure across observations. These methods require the underlying structure to be held fixed and are not robust to deviations commonly found across images. We present autoencoding topographic factors, a novel variational inference scheme, to decompose irregular observations on a lattice into a superposition of low-rank sources. By exploiting recent developments in variational autoencoders, we replace fixed sources with a nonlinear mapping that parameterizes an unnormalized distribution on the lattice. In doing so, we permit sources to drift dynamically, filtering residual differences in location across comparable areas of interest. This gives an implicit mapping to a unique latent representation while simultaneously forcing the posterior to model group variability in spatial structure. Simulation results and applications to functional imaging demonstrate the effectiveness of our method and its ability to outperform existing spatial factor models.

Keywords: biomedical image analysis, functional imaging, spatial statistics, variational inference.

1. INTRODUCTION

THE ANALYSIS OF BIOMEDICAL IMAGES has accelerated in recent years due to domain-specific methodologies developed for multiple application areas. Calcium imaging in neurons (Pnevmatikakis et al., 2016), transcriptome profiling from single cells (Svensson et al., 2017), and functional imaging of various biomarkers (Gershman et al., 2014; Manning et al., 2014a) are exciting examples. Latent variable models are the predominant method for visualizing and extracting structure in spatial data. These data are characterized by a location vector $\mathbf{x}_i \in \Omega \subseteq \mathbb{R}^d$ parameterizing each observation $\mathbf{y}(\mathbf{x}_i)$. Given a tensor $\mathcal{Y} \equiv \{\mathbf{y}(\mathbf{x}_1), \dots, \mathbf{y}(\mathbf{x}_m)\}_{n=1}^N$ of N realizations, each sequence of m correlated random variables $Y(\mathbf{x}_1), \dots, Y(\mathbf{x}_m)$, a fundamental challenge is to identify a subset of physical locations that define areas of interest. To this end, lattice-based models formalize an encoding of a latent probability distribution over $Y(\mathbf{x}_1), \dots, Y(\mathbf{x}_m)$ to quantify statistical dependencies based on distance. This representation is often used for Gaussian process regression or Kriging methods to predict covariance structure between hidden variables and observed features across a physical location in an ensemble (Svensson et al., 2017). For example, extracting relevant voxels from a collection of functional images to discover a latent hemodynamic response enables comparing baseline versus pathological populations (Worsley et al., 1996).

Techniques such as robust principal component analysis (Candès et al., 2011), independent component analysis, and dictionary learning (DL) are commonly applied to blind source separation problems; however,

¹Department of Computer Science, Columbia University, New York, New York.

²Department of Computer Science, Brown University, Providence, Rhode Island.

*These authors contributed equally to this work.

Joint ICML and IJCAI Workshop on Computational Biology, Stockholm, Sweden, 2018

they require an inherently linear demixing or deconvolution and may fail if there is no linear mixture that leads to independent outputs (Mukamel et al., 2009). Notably, these methods do not learn a distribution on the lattice that can be used to quantify uncertainty or to generate new data. Topographic factor models (Gershman et al., 2011, 2014; Manning et al., 2014a) are a family of Bayesian variational techniques for images that require underlying structure on the set of random variables to be held constant to produce a matrix factorization with spatially interpretable sources.

In this study, we develop autoencoding topographic factors (AETF), a novel Bayesian algorithm, to infer spatial dependencies by decomposing observations on a lattice into a weighted set of low-rank sources. We are particularly interested in a solution that generalizes to unseen data and that is robust to noncollocated regions of interest. The key insight of AETF is to leverage recent advances in variational inference (Gershman et al., 2014; Ranganath et al., 2014) and stochastic gradient variational Bayes (Kingma and Welling, 2013; Rezende et al., 2014) to learn a latent probability model that preserves group variability in spatial structure. Our contributions are to combine two paradigms where convolutional neural networks define the loading matrix, and the factor matrix itself maps data to source functions that transform across observations. This is achieved without hard coding hyperparameters that control an a-priori generative model. In doing so, we remove the propensity on initialization of domain-specific priors. Experiments on two simulated data sets and on functional imaging data show that our model returns a higher proportion of variance explained than existing topographic factor models.

This article is organized as follows: Section 2 reviews related methods for blind source separation. Section 3 defines the model we propose for spatial factor analysis. Section 4 provides experimental results. Section 5 concludes and discusses future work.

2. RELATED WORK

Non-negative matrix factorization has been adapted to calcium imaging data to infer both location and spiking dynamics of neurons from fluorescence movies (Friedrich et al., 2015; Pnevmatikakis et al., 2016). To facilitate inference, significant domain-specific preprocessing must be done to constrain the region of interest to patches and model spatial background structure. A more general approach to factor analysis for space-time variation involves identifying a common set of shared factors whose time varying dynamics are modeled with autoregressive weights (Lopes et al., 2008). Our method differs in that we allow each observation to have a unique set of factors generated by a common nonlinear mapping from the input space. In the case of spatial dynamic factor analysis, Bayesian inference is performed with a reversible jump MCMC algorithm whose convergence is difficult to assess and must be tailored for the problem at hand.

Topographic factor analysis (TFA) for modeling functional magnetic resonance imaging via factor decomposition was proposed (Manning et al., 2014a) and fit using Black Box VI (Ranganath et al., 2014). Our approach uses similar spatial functions and a posterior, which is mean field in time, however (Manning et al., 2014a), hard codes the structure of the factors, which are shared across observations. These parameters are often prefit using a preprocessing algorithm. Previous models are unable to capture the variation among observations while fitting individual-specific factors. In this sense, we develop a robust and expressive posterior distribution that does not require hand tuning hyperparameters for the priors. In addition, the TFA model does not generalize to unseen data and suffers linear parameter growth with respect to the size of the data set.

3. AUTOENCODING TOPOGRAPHIC FACTORS

3.1 Standard lattice modeling

Following the convention of factor analysis, we assume that our data $\mathbf{Y} \in \mathbb{R}^{N \times V}$ can be decomposed into a set of unobserved weights and latent factors. We use N to denote the number of observations (images), K the number of sources, and V the number of lattice positions (voxels). We are discussing lattices in both 2D and 3D for our analysis. Each latent source is defined using a function that assigns a value to each point on the lattice (in voxel space) based on its location. For example, using the multivariate normal (MVN):

$$K(\mathbf{x}_i | \mu, \Sigma) = \exp \left\{ -\frac{1}{2} (\mathbf{x}_i - \mu)^T \Sigma^{-1} (\mathbf{x}_i - \mu) \right\}. \quad (1)$$

We posit each observation $\mathbf{y}_n \in \mathbb{R}^{1 \times V}$ has a low-rank approximation that is a product of factor loadings $\mathbf{w}_n \in \mathbb{R}^{1 \times K}$ and a factor matrix $\mathbf{F} \in \mathbb{R}^{K \times V}$. The generative distribution of our model factorizes using a Gaussian as follows:

$$P(\mathbf{Y}) = \prod_{n=1}^N P(\mathbf{y}_n). \quad (2)$$

$$P(\mathbf{y}_n) = \mathcal{N}(\mathbf{y}_n | \mathbf{w}_n \mathbf{F}, \sigma_y^2), \quad (3)$$

where σ_y^2 denotes the location or voxel noise. In a study by Manning (Manning et al., 2014b), radial basis source functions $f_k \in \mathbb{R}^V$ are used to generate basis images and to define $\mathbf{F}$, the source image matrix. In general, rows of $\mathbf{F}$ are computed by evaluating each of the K source functions at all V lattice points of the voxel space.

While it is common to focus on $\Sigma = \sigma \mathbf{I}$ or the MVN case in which Σ is full, a larger class of kernels are supported through the Matérn family of covariance functions. In this study, $K_\nu(\cdot)$ is the modified Bessel function of the second kind with order parameter ν , where ρ defines correlation length and $\lfloor \nu \rfloor$ describes the smoothness of the process. $\Gamma(\cdot)$ is the gamma function.

$$K(\mathbf{x}_i | \mu, \nu) = \frac{1}{\Gamma(\nu) 2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{\rho} \cdot \|\mathbf{x}_i - \mu\| \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{\rho} \|\mathbf{x}_i - \mu\| \right).$$

The above simplifies for half-integer values of ν and reduces to the rational quadratic function with $\nu, \rho > 0$ to express a scale mixture of squared exponentials:

$$K(\mathbf{x}_i | \mu, \nu, \rho) = \left(1 + \frac{\|\mathbf{x}_i - \mu\|^2}{2\nu\rho^2} \right)^{-\nu}. \quad (4)$$

Samples from the Gaussian process are $\lfloor \nu - 1 \rfloor$ times differentiable producing the radial basis function (RBF) case when $\nu \rightarrow \infty$. As with the above, the choice of distance metric can produce isotropy or anisotropy.

We are interested in the posterior distribution, involving integrating over a set of possible values for the latent variables:

$$P(\mathbf{W}, \mathbf{F} | \mathbf{Y}) = \frac{P(\mathbf{Y}, \mathbf{W}, \mathbf{F})}{P(\mathbf{Y})}. \quad (5)$$

$$P(\mathbf{Y}) = \int \int P(\mathbf{Y}, \mathbf{W}, \mathbf{F}) d\mathbf{W} d\mathbf{F}. \quad (6)$$

The problem is in general intractable to compute. To perform variational inference, a mean field distribution is defined in which each variable is independent:

$$Q(\mathbf{W}, \mathbf{M}, \Lambda) = \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}(w_{n,k} | \mathbf{m}_{w_{n,k}}, \Lambda_{w_{n,k}}) \mathcal{N}(c_{n,k} | \mathbf{m}_{c_{n,k}}, \Lambda_{c_{n,k}}) \mathcal{N}(s_{n,k} | \mathbf{m}_{s_{n,k}}, \Lambda_{s_{n,k}}). \quad (7)$$

We introduce notation for the set $\phi_k \in \phi$ to denote hyperparameters where c, s, w denote centers, width scales, and weights, respectively:

$$\phi_k = \{\mathbf{m}_{c,k}, \Lambda_{c,k}, \mathbf{m}_{s,k}, \Lambda_{s,k}, \mathbf{m}_{w,k}, \Lambda_{w,k}\}. \quad (8)$$

These allow drawing corresponding latent random variables for centers, width scales, and weights for the k th latent source:

$$\mathbf{Z}_k = \{z_{c,k}, z_{s,k}, z_{w,k}\}, \quad (9)$$

where $z_{\xi,k} \sim \mathcal{N}(\mathbf{m}_{\xi,k}, \Lambda_{\xi,k}^2)$ for $\xi \in \{c, s, w\}$. Note that in the isotropic case $\phi \in \mathbb{R}^{K(D+5)}$ and $\mathbf{Z} \in \mathbb{R}^{K(D+2)}$ where D is the dimensionality of the lattice.

Across all ξ, k one can define $\mathbf{m}_\phi = (\mathbf{m}_{\xi,k})_{\forall \xi,k}$ and $\Sigma_\phi = \Lambda_\phi \Lambda_\phi^T$ for $\Lambda_\phi = (\Lambda_{\xi,k})_{\forall \xi,k}$, thus the parameters $\mathbf{m}_\phi, \Sigma_\phi$ denote the means and covariances used to draw $\mathbf{Z}$. $\mathbf{Z}$ then defines $\mathbf{F}$, by f_k being a Gaussian function with parameters $z_{c,k}$ and $z_{s,k}$.

3.2. Autoencoding topographic factors

The idea of AETF is to replace the fixed latent sources by defining a function that parameterizes $\mathbf{Z}$ using the output of a probabilistic encoder. The encoder creates an implicit mapping from each $\mathbf{y}_n \in \mathbf{Y}$ across a set of observations to a unique factor representation, while requiring that ϕ encodes the group variability in spatial structure.

Formally, the variational inference framework states the ELBO for the marginal log likelihood $\mathcal{L}(\mathbf{Y}) \leq \log p(\mathbf{Y})$ with respect to the variational approximation $q_\phi(\mathbf{Z}|\mathbf{Y})$:

$$\begin{aligned} \mathcal{L}(\mathbf{Y}) &= \mathbb{E}_{q(\mathbf{z}|\mathbf{y})}[\log p_\theta(\mathbf{Y}, \mathbf{Z})] - \mathbb{E}_{q(\mathbf{z}|\mathbf{y})}[\log q_\phi(\mathbf{Z}|\mathbf{Y})] \\ &= \mathbb{E}_{q(\mathbf{z}|\mathbf{y})}[\log p_\theta(\mathbf{Y}|\mathbf{Z})] - D_{KL}(q_\phi(\mathbf{Z}|\mathbf{Y})||p(\mathbf{Z})) \end{aligned} \quad (10)$$

We wish to compute the expectation in Equation (10) numerically and differentiate with respect to ϕ . We now rewrite Equation (3) as

$$P(\mathbf{y}_n) = \mathcal{N}(\mathbf{y}_n | \mathbf{w}_n(\mathbf{y}_n) \mathbf{F}(\mathbf{y}_n), \sigma_y^2), \quad (11)$$

and decompose $\mathbf{F}$ as

$$\mathbf{F}(\mathbf{y}_n) = \begin{pmatrix} f_1(\mathbf{y}_n) \\ \vdots \\ f_K(\mathbf{y}_n) \end{pmatrix}, \quad (12)$$

where $f_k(\mathbf{y}_n)$ is the lattice values of a Gaussian function parameterized by $z_{c,k}(\mathbf{y}_n)$ and $z_{s,k}(\mathbf{y}_n)$. $z_{\xi,k}(\mathbf{y}_n)$ itself is a latent variable drawn from a normal distribution $z_{\xi,k}(\mathbf{y}_n) \sim \mathcal{N}(\mathbf{m}_{k,\xi}, \phi(\mathbf{y}_n), \Lambda_{k,\xi}, \phi(\mathbf{y}_n))$ whose parameters are the encoder output.

Using the well-known reparameterization trick (Kingma and Welling, 2013; Rezende et al., 2014), we sample from $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ to compute the following as the Monte Carlo estimates of gradients have high variance:

$$\mathbf{Z}_c = \mu_c + \varepsilon \odot \sigma_c. \quad (13)$$

$$\mathbf{Z}_s = \mu_s + \varepsilon \odot \sigma_s. \quad (14)$$

One is now free to choose the weights $\mathbf{Z}_w \in \phi$ as variational parameters of the recognition model or parameters with the generative model: $\mathbf{Z}_w \in \phi$. Including the weights in ϕ gives the following:

$$\mathbf{Z}_w = \mu_w + \varepsilon \odot \sigma_w. \quad (15)$$

When $\mathbf{Z}_w \notin \phi$, we learn the weights as point estimates using the update rule:

$$\mathbf{W}^{i+1} \leftarrow \mathbf{W}^i \odot \mathbf{Y} \mathbf{F}(\mathbf{y}_n)^T \oslash \mathbf{W}^i \mathbf{F}(\mathbf{y}_n)^T. \quad (16)$$

Note that the problem is hard due to the nonconvexity in the source image matrix. With the parameters ϕ of the recognition model in hand, we have the full model specification. In contrast to standard autoencoder formalization, where the generative model involves a decoder whose parameters need to be inferred, AETF specifies the generative model. We thus compute the approximation $\mathbf{y} = \mathbf{W}(\mathbf{y}_n) \cdot \mathbf{F}(\mathbf{y}_n)$.

Standard autoencoders learn the respective encoder/decoder parameters θ, ϕ by maximizing the conditional log likelihood $\mathbb{E}_{q(\mathbf{z}|\mathbf{y}^i)}[\log p_\theta(\mathbf{y}^i|\mathbf{z})]$ by differentiating through $g \leftarrow \nabla_{\theta, \phi} \mathcal{L}^M(\theta, \phi \mathbf{Y}^M, \varepsilon)$ (Kingma and Welling, 2013; Rezende et al., 2014). AETF only needs to learn the encoder parameters ϕ , which is achieved by analogous maximization of the conditional log likelihood $\mathbb{E}_{q(\mathbf{z}|\mathbf{y}^i)}[\log p(\mathbf{y}^i|\mathbf{z})]$, differentiating through $g \leftarrow \nabla_{\phi} \mathcal{L}^M(\phi \mathbf{Y}^M, \varepsilon)$.

3.3. Implementation. The encoder takes as input an observation and outputs the parameters of the distributions over latent variables. Two recognition models are implemented, one with isotropic and another with full covariance source functions. The isotropic decoder receives as input the sampled latent space vector $\mathbf{Z} \in \mathbb{R}^{k(d+2)}$, including $\mathbf{Z}_c$, $\mathbf{Z}_s$, and $\mathbf{Z}_w$. Note that in the second case of a full covariance matrix $\mathbf{Z}_{ssk} = \Lambda \Lambda^T$, we learn parameters $\mathbf{Z}_{ssk} \in \mathbb{R}^{kd(d+1)/2}$. The spatial factorization constraints of our probability model are imposed within the decoder. Thus, unlike traditional variational autoencoders where both the encoder and decoder are neural networks, AETF uses a neural network only for the encoder. The

decoder uses the sampled latent space to reconstitute the input according to our imposed factorization and therefore is not parameterized by a neural network.

The encoder network can comprise any number of convolutional layers followed by any number of fully connected layers before the output layer. The convolutional layer executes an $L^{(1)} \otimes \dots \otimes L^{(D)}$ convolution along the number of lattice dimensions D (where L is specified for each layer) with k (the number of sources) output channels, a bias addition, a $\tanh$ nonlinearity, and max pooling with a $3^{(1)} \otimes \dots \otimes 3^{(D)}$ kernel and a stride of 1. Our fully connected layers begin operating on the flattened output of the last convolutional layer or the flattened image if a convolution layer is not used. Their output dimensions are specified ratiometrically according to their output-to-input dimensions. Like most autoencoders, our encoder seeks to compress information. Thus, we only consider output-to-input ratios for our fully connected layers that are all less than or equal to 1. These fully connected layers invoke an affine transformation followed by a $\tanh$ nonlinearity.

Our final output layer varies according to the latent space parameter class. Those parameters that are means (μ_c , μ_s , and μ_e) have no restrictions on their values except the last one, which must be positive. We handle this exception in the decoder. Therefore, we are free to use a vanilla affine transform as the output layers for these parameters. Conversely, those parameters that are standard deviations (σ_c , σ_s , and σ_w) must be greater than or equal to zero. Thus, for those standard deviations that parameterize our latent space, we use an affine transformation followed by a custom nonlinearity we call PostReAct (Positive Real Activation in Equation (17)). This nonlinearity is a piece-wise combination of a shifted ReLU and a decaying exponential. In this manner, we benefit from ReLU's positive regimen that avoids vanishing gradients that are common with double-saturating activations while avoiding the potential of neuron death associated with ReLU's negative regimen.

$$\Psi(\lambda) = \begin{cases} \exp(\lambda) & , \lambda < 0 \\ \lambda + 1 & , \lambda \geq 0 \end{cases} \quad (17)$$

Our decoder has two responsibilities. First, it constructs the spatial factors using the $\mathbf{Z}_c$ and $\mathbf{Z}_s$ latent space. However, and as aforementioned, $\mathbf{Z}_s$ arrives at the encoder on the incorrect support. The RBF function assumes this number is positive. We convert $\mathbf{Z}_s$ to the correct support in two ways. First, we pass it through a PostReAct nonlinearity. Second, we square it in our isotropic implementation. Equation (18) captures this process that we use for each of our basis image calculations. In this study, $f_k(v)$ represents the value of the k th RBF source at voxel position v . Unlike traditional RBF functions, we add a 1 to the denominator to clamp the source's width in a continuously differentiable manner. Before this modification, sampled $\mathbf{Z}_s$, resulting in small source widths, produced exploding gradients for our optimizer. Once the decoder constructs the k basis images, it recombines them into a single image via a weighted summation that uses $\mathbf{Z}_w$.

$$f_k(v) = \exp\left(-\frac{\|\mathbf{Z}_{c,k} - v\|_2^2}{2 \cdot \Psi(\mathbf{Z}_{s,k})^2 + 1}\right). \quad (18)$$

We present two encoder network architectures in Figure 1. Our first uses only a 7×7 convolutional layer followed by the output layer. Our second uses—in order of appearance—a 7×7 convolutional layer, a 5×5 convolutional layer, a $1 : 1$ output-to-input fully connected layer, and a $4 : 3$ output-to-input fully connected layer followed by the output layer. We then permute these two architectures for differing numbers of latent sources. We note that k modifies the size of the network as it determines the number of output channels for each convolutional layer. Our implementation supports imposing a non-negative factorization in addition to one in which the weights are permitted to take negative values.

We modify the loss from Equation (10). Specifically, we introduce a β term in front of the regularizer as suggested in Liang et al. (2018). Furthermore, they suggest that β values less than 1 improve quality. The utilized per-sample loss function for AETF appears in Equation (19). In our experiments, we set β to zero such that our loss reduces to just the reconstruction error. In this study, n represents the n th sample and V is the cardinality of our voxel space such that subscript n, i corresponds to the i th voxel of the n th sample.

$$\mathcal{L}(\mathbf{Y}_n) = \frac{1}{V} \sum_{i=1}^V \left[(\hat{\mathbf{Y}}_{n,i} - \mathbf{Y}_{n,i})^2 \right] + \beta D_{KL}(q_\phi(\mathbf{Z}_i | \mathbf{Y}_i) || p(\mathbf{Z}_i)) \quad (19).$$

FIG. 1. Schematic illustrating encoder network architectures for the AETF framework. Shallow (a) and deep (b) architectures shown. AETF, autoencoding topographic factors.

4. EXPERIMENTS

Three results are presented, each of which illustrates a strength of the AETF model. We discuss (i) fitting in-model synthetic data, (ii) fitting noncollocated source functions to smooth, unmix, and localize spatial dependencies in random fields, and (iii) decomposing thousands of functional images into latent source functions and evaluating our ability to generalize on unseen data.

4.1. In-model data

We generate a synthetic data set using $k=2$ source functions over 1000 observations on a $20 \times 20 \times 20$ lattice. In our experiments, TFA was unable to handle larger lattice dimensions in $\mathbb{R}^3$. Unlike the generative process specified in TFA (Manning et al., 2014a), the position of each source function may shift across observations and is not restricted to be collocated on the lattice. This design choice is relevant given that the blood oxygen level dependency response is not static and often transforms dynamically as a time series. Figure 2a provides a schematic illustrating the position of two sources located near the vertices of the cube. Each source function is permitted to drift between one of two possible states, which are represented using the red and black colors. Figure 2b and c provides the variance explained on the training and testing sets, respectively, using $k=2$ components. TFA, independent component analysis (ICA), and principal component analysis (PCA) underperform relative to DL and AETF. Unlike DL, AETF is able to parameterize the transforming source functions while maintaining nearly all of the variances explained.

4.2. Gaussian random fields

Gaussian random fields (GRFs) are often used in image analysis to model stochastic processes on a lattice and to introduce noise. We illustrate how AETF recovers autocorrelation structure by filtering a sequence of GRFs simulated using spectral methods (Annika and Jrgen, 2011). The spectral density of a fixed covariance kernel is multiplied with a Fourier transformed white noise field before applying an inverse transform. This process introduces a nonsmooth signal in which spatial autocorrelations are not explicitly collocated across observations. Figure 3a provides a representative sample along with the inferred reconstruction in Figure 3b. We fit 10 source functions to 1000 observations on a cubic lattice. As a visualization, the planar cross section is provided in Figure 3. The surface is shifted above the image to illustrate the smoothness of the field along with contours presenting the location of the inferred spatial factors. Figure 3c provides the variance explained across models and fits. AETF outperforms TFA, ICA, DL, and PCA; the canonical method for Gaussian data. It is clear that TFA underperforms when the correlation structure is not held fixed.

4.2.1. Image noise. Spatial factor models learn a smooth statistical map in the presence of noise in which the desired signal extends over several lattice points. A good fit should be robust to variation between observations while preserving the correlation structure within the data. To achieve this, auto-encoding topographic factors learns a unique decomposition by simultaneously factorizing the observation matrix, inferring the position of spatial dependencies, and introducing flexibility for the location of factors across the lattice. This process is analogous to blurring residual differences in location between comparable areas of activation. When two observations are similar, this is captured in their latent spatial representations. For heterogeneous data, AETF parameterizes spatial dynamics.

4.2.2. Initializing TFA and hierarchical TFA. Heuristics are often suggested to initialize hyperparameters for TFA so that local optima in the source image matrix do not serve as an impediment for nonconvex optimization. There exist multiple values of parameters for the location and width of the sources that are equally likely to have generated an observation $\mathbf{y}_n$, due to the rotational invariance of $\mathbf{F}$. One proposed approach is to place hyperparameters a-priori in locations corresponding to high and low activation. Hotspot initialization (Manning et al., 2014a) refers to an iterative process in which the mean image is computed, the mean activation is subtracted, and the absolute value is taken of all of the remaining activations. The result is an energy landscape in which peaks correspond to extremum. These peaks are iteratively flattened as source centers are placed on these extremum. Values for $\mathbf{m}_{s_n,k}$ the mean of the distribution for source k 's width scale are then solved via Newton's method. Once preinitialized, the source centers and width scales frequently remain fixed. In our experiments, sources for TFA initialized using both

FIG. 2. Description of the first simulation: (a) schematic illustrating two source functions located near the vertices of the lattice. Each source transforms across observations drifting between one of two states (denoted with shades light and darker); (b) variance explained for different models using two components. AETF outperforms TFA on the train set; (c) TFA, PCA, and ICA underperform on the test set. ICA, independent components analysis; PCA, principal components analysis; TFA, topographic factor analysis.

FIG. 3. Summary of the AETF fit to the GRF simulation: (a) the cross section of a single observation and (b) the cross section of the AETF topographic reconstruction. The surface is shifted above the plane to illustrate the smoothness of the field along with contours presenting the location of the inferred spatial factors; (c) variance explained across models. AETF provides the highest R^2 . GRF, Gaussian random field.

FIG. 4. Summary of the sagittal cross-section NYU data: (a) variance explained for various models as a function of number of sources on the training data; (b) two source weights plotted across time frames illustrating strong subject-specific similarities. Dashed vertical lines denote unique subjects; (c) variance explained as a function of number of sources for test data. AETF returns a near perfect variance explained using $k=25$ source functions.

hotspot initialization and k -means outperformed experiments with no initialization. AETF outperformed both methods without being contingent on any such initialization to perform inference successfully.

4.3. NYU data set

We consider the problem of modeling functional images using the NYU test-retest data set (Shehzad et al., 2009). Data were obtained using a Siemens Allegra 3.0 Tesla scanner. Data consist of 26 participants each with 3 resting-state scans of 197 continuous EPI functional volumes. Each scan consists of 39 slices of a matrix 64×64 with an acquisition voxel size of $3 \times 3 \times 3$ mm. Scans 2 and 3 were conducted 45 minutes apart roughly 5–16 months after Scan 1.

Slice timing correction, spatial normalization, smoothing, and noise stripping were performed using the *Nipype* interface to the FSL software library. The sequential dependency of the time series was not accommodated and each time frame was treated independently. An AETF model was trained using all three sessions reserving 20% for the testing set as a performance criteria to evaluate our fit. To test the significance of the lattice dimensions, models were fit to both sagittal cross-section data and full cubic volumes.

4.3.1. Sagittal cross sections in 2D. Sagittal cross-section data were fit to the 13 subjects using the first session. Figure 4a provides the variance explained for AETF, PCA, ICA, and DL as a function of number of sources on training data. TFA and hierarchical TFA (HTFA) implementations are not supported on the 2D lattice. The R^2 approaches 1 the number of sources K increases. Figure 4b plots the weight values for two randomly selected source functions across a subset of time frames. Dashed vertical lines distinguish subjects. Strong per-subject similarities are visible. Figure 4c provides the variance explained by AETF as K increases on the test data. Using $k=50$ source functions, 99% of variance is explained. We find that $K=25$ anisotropic sources are sufficient for high-quality reconstruction. Interestingly, AETF is able to converge without *any* preprocessing to preserve 89% of total variance on the raw NYU data. On the preprocessed data set, 25 source functions preserve 98% of total variance.

4.3.2. Functional imaging with 3D volumes. The three-session NYU data on the cubic lattice were modeled using AETF, TFA, and HTFA. The $64 \times 64 \times 40$ lattice was divided into eight $20 \times 20 \times 20$ cubic volumes. TFA and HTFA were unable to handle larger lattice dimensions on the full set of 7683 frames. We fit $k=10$ source functions to each cubic volume and average the cost across the total area. For TFA, one model was fit across subjects, whereas 39 subjects were fit using HTFA. Figure 7a displays a new frame evaluated using the trained model to illustrate the effect of applying the trained model on unseen data. The surface is plotted above the image to highlight the areas of activation above the corresponding factors on the mesh. Figure 7a and b provides the train and test R^2 , respectively. It is clear that AETF outperforms both methods. Unlike HTFA, the hierarchical covariance structure is inferred from the data and not specified a-priori.

FIG. 5. The distribution of the top 15 weighted source center means for subjects and time frames. Colors represent distinct source function center means. The outline of a single frame is provided as a reference. Each point is a parameter for an image in a collection of time series. The latent factors recovered by AETF exhibit strong spatial localization.

FIG. 7. Results for the cubic volume NYU data: (a) a cross section of a frame and the surface highlighting source intensities; R^2 values for training (b) and testing (c) for various models averaged across eight cubic volumes using $k = 10$ source functions. AETF consistently outperforms both TFA and HTFA. TFA, topographic factor analysis; HTFA, hierarchical topographic factor analysis.

5. DISCUSSION

In the context of functional imaging, a spatial model should be able to extract both global and individual characteristics. In examining how the model parameters for centers, widths, and weights varied across testing data, we find source centers are not only similar at the per-subject microscale but also marginally similar at the global macroscale. However, we see much more global variability with weight values. Compared with a similar factorization in Manning et al. (2014b) that constricts learning to globally shared sources and individual per-frame weights, our model naturally learns a similar representation. Namely, through a shared encoder mapping, source variability is less pronounced than weight variability.

AETF offers several advantages over unstructured blind source separation techniques. TFA, HTFA, DL, PCA, and ICA explicitly learn factor weights (loadings) for each observation. The number of trainable parameters is therefore linear with respect to N , the number of observations. AETF parameters φ are constant with respect to N . This paradigm reduces memory footprint for large N and allows AETF to handle unseen data. By design, the factor images learned by AETF possess lower complexity than the observed images.

AETF can accommodate any priors but is not contingent on an a-priori choice of generative model hyperparameters to converge. This is mitigated by choosing uniform priors for the generative model. In this way, AETF is not sensitive to preinitialization issues that plague TFA and HTFA. It is also possible to parameterize the priors of the generative model using a trainable decoder network. Unlike TFA and HTFA, source functions are allowed to transform across individual frames. This is advantageous for time series modeling. In our experiments, DL sometimes provided a comparable R^2 . AETF, however, returns a factorization along with spatially parameterized functions. AETF was written in TensorFlow. The source code and several visualizations are available online.

6. CONCLUSIONS

We have presented AETF, a novel variational inference scheme, for lattice-based measurements in which each observation is given a unique spatial decomposition. The proposed method is robust to high-dimensional data in which sources are not rigidly collocated, introduces nonlinearity, supports a family of kernels, and the ability to enforce a constrained or non-negative matrix factorization. AETF preserves a large proportion of variance even when factor positions shift dynamically across observations. Highlights include the ability to identify autocorrelation structure in a collection of random fields and the ability to scale to thousands of 3D functional images with a number of training parameters independent of data set size.

The results, in particular Figure 4b, motivate an explicitly hierarchical AETF across individuals, as well as a temporally correlated AETF. A natural extension is to explore the method of normalizing flows (Rezende and Mohamed, 2015; Kingma et al., 2016) as an alternative to defining factors by specifying kernels for source functions. We expect that the approximate posterior would remain simple to compute while each source is permitted to undergo a sequence of transformations giving rise to complex and expressive spatial dependencies.

AUTHOR DISCLOSURE STATEMENT

The authors declare that no competing financial interests exist.

REFERENCES

- Annika, L., and Jrgen, P. 2011. Fast simulation of gaussian random fields. *Monte Carlo Methods Appl.* 17, 195–214.
- Candès, E.J., Li, X., Ma, Y., et al. 2011. Robust principal component analysis? *J. ACM* 58, 11:1–11:37.
- Gershman, S., Blei, D.M., Pereira, F., et al. 2011. A topographic latent source model for fmri data. *Neuroimage* 57, 89–100.
- Gershman, S., Blei, D.M., Norman, K.A., et al. 2014. Decomposing spatiotemporal brain patterns into topographic latent sources. *Neuroimage* 98, 91–102.

- Friedrich, J., Soudry, D., Mu, Y., et al. 2015. Fast constrained non-negative matrix factorization for whole-brain calcium imaging data. In *NIPS Workshop on Statistical Methods for Understanding Neural Systems*.
- Kingma, D.P., and Ba, J. 2015. Adam: A method for stochastic optimization. In 3rd International Conference for Learning Representations, San Diego.
- Kingma, D.P., and Welling, M. 2013. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013.
- Kingma, D.P., Salimans, T., Józefowicz, R., et al. 2016. Improving variational autoencoders with inverse autoregressive flow. In *NIPS*, 4736–4744.
- Lee, D.D., and Seung, S.H. 2001. Algorithms for non-negative matrix factorization, 556–562. In Leen, T.K., Dietterich, T.G., and Tresp, V., eds. *Advances in Neural Information Processing Systems*, Vol. 13. MIT Press.
- Liang, D., Krishnan, R., Hoffman, M., et al. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of The Web Conference (WWW)*, 2018, 2018.
- Lopes, H.F., Salazar, E., and Gamerman, D. 2008. Spatial dynamic factor analysis. *Bayesian Anal.* 3, 759–792.
- Manning, J.R., Ranganath, R., Keung, W., et al. 2014a. Hierarchical topographic factor analysis, 1–4. In International Workshop on Pattern Recognition in Neuroimaging, PRNI 2014, Tübingen, Germany, June 4–6, 2014. DOI: 10.1109/PRNI.2014.6858530
- Manning, J.R., Ranganath, R., Norman, K.A., et al. 2014b. Topographic factor analysis: A Bayesian Model for inferring brain networks from neural data. *PLoS ONE* 9, e94914.
- Mukamel, E.A., Nimmerjahn, A., and Schnitzer, M.J. 2009. Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron* 63, 747–760.
- Pnevmatikakis, E.A., Soudry, D., Gao, Y., et al. 2016. Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* 89, 285–299.
- Ranganath, R., Gerrish, S., and Blei, D.M. 2014. Black box variational inference, 814–822. In Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, AISTATS, Reykjavik, Iceland, April 22–25, 2014.
- Rezende, D.J., and Mohamed, S. 2015. Variational inference with normalizing flows, 1530–1538. In Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, July 6–11, 2015.
- Rezende, D.J., Mohamed, S., and Wierstra, D. 2014. Stochastic backpropagation and approximate inference in deep generative models, 1278–1286. In Xing, EP., and Jebara, T., eds. *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research (PMLR)*, Beijing, China, June 22–24, 2014.
- Shehzad, Z., Kelly, A.M., Clare, R., et al. 2009. The resting brain: Unconstrained yet reliable. *Cereb Cortex.* 19, 2209–2229.
- Svensson, V., Teichmann, S.A., and Stegle, O. 2017. SpatialDE—Identification of spatially variable genes. *bioRxiv*, 143321, May 2017. DOI: 10.1101/143321.
- Worsley, K.J., Marrett, S., Neelin, P., et al. 1996. A unified statistical approach for determining significant signals in location and scale space images of cerebral activation.

Address correspondence to:

Antonio Moretti
Department of Computer Science
Columbia University
New York, NY

E-mail: amoretti@cs.columbia.edu