

Scalable Group Level Probabilistic Sparse Factor Analysis

Jesper L. Hinrich, Søren F. V. Nielsen, Nicolai A. B. Riis,
Casper T. Eriksen, Jacob Frøsig, Marco D. F. Kristensen, Mikkel N. Schmidt,
Kristoffer H. Madsen and Morten Mørup ^{*}
Technical University of Denmark

December 15, 2016

Abstract

Many data-driven approaches exist to extract neural representations of functional magnetic resonance imaging (fMRI) data, but most of them lack a proper probabilistic formulation. We propose a group level scalable probabilistic sparse factor analysis (psFA) allowing spatially sparse maps, component pruning using automatic relevance determination (ARD) and subject specific heteroscedastic spatial noise modeling. For task-based and resting state fMRI, we show that the sparsity constraint gives rise to components similar to those obtained by group independent component analysis. The noise modeling shows that noise is reduced in areas typically associated with activation by the experimental design. The psFA model identifies sparse components and the probabilistic setting provides a natural way to handle parameter uncertainties. The variational Bayesian framework easily extends to more complex noise models than the presently considered.

1 Introduction

In functional magnetic resonance imaging (fMRI) large amounts of data are currently being generated due to high spatial resolution, an increase in the typical number of sessions acquired and a trend towards multi-center acquisition and data sharing. There is therefore a growing need for methods and algorithms that scale, while still keeping reasonable model assumptions.

A common problem in functional neuroimaging is finding a good latent representation of the data. One approach is to use data-driven methods, which allows the data to “speak for itself”. Two popular approaches in fMRI is to assume orthogonality or independence, giving rise to principal component analysis (PCA) [1] and independent component analysis (ICA) [2], respectively. For

^{*}This work was supported by the Lundbeck Foundation, grant no. R105-9813.

group level analysis group-ICA [3] and independent vector analysis (IVA) [4] are the most prominent data-driven methods for inferring components of neural activity. Group-ICA and IVA have the advantage that independent components often are sparse providing interpretable spatial activation maps, due to the typical assumption of sparse source distributions. This interpretability in fMRI may also be achieved by models optimizing for sparsity instead of independence [5], which in some cases achieves similar results despite optimizing different properties [6].

ICA models in general assume noise free data, which is commonly achieved by modeling additive homoscedastic Gaussian noise, or the noise is sought removed through a combination of (PCA) pre-whitening and standardization to unit voxel variance. A probabilistic ICA (pICA) approach with spatial and temporal noise modeling was suggested in [7], which included a noise estimation step in their framework. However, the noise is estimated in a separate step from the ICA components, which can be problematic in comparison with joint estimation [8]. While [7] does incorporate probabilistic elements into ICA, it is not probabilistic from a Bayesian modeling perspective. In contrast, joint estimation has shown promising results on fMRI (cf. [9, 10]).

We propose a probabilistic sparse factor analysis (psFA) model for group level analysis of fMRI data with heteroscedastic noise. Model inference is done using variational Bayes with a mean-field approximation and automatic relevance determination (ARD) [11] to promote sparsity on individual voxels of the spatial maps and prune components by learning their relevance in time. To overcome the large computational burden involved massive parallelization of the updates is exploited using a graphical processing unit (GPU). The model is first investigated on synthetic data, where the advantages of using a Bayesian approach to factor analysis (FA) and principal component analysis (PCA) is briefly assessed. The proposed model’s applicability to fMRI is tested on a motor task experiment [12] and on a resting state experiment [13] and contrasted with pICA (MELODIC[7]).

2 Methods and Data

2.1 Probabilistic Sparse Factor Analysis

We propose a group level probabilistic sparse factor analysis model (psFA), which is a combination of the probabilistic sparse PCA model proposed in [14], the group level PCA analysis proposed in [15] and the inclusion of heteroscedastic voxel noise, first proposed in the context of fMRI and variational inference in [16]. The generative model for a data array, $\mathbf{X} \in \mathbb{R}^{V \times T \times B}$, with V voxels, T timepoints and B subjects, can be written as,

$$\begin{aligned}
\alpha_{vd} &\sim \mathcal{G}(a_\alpha, b_{\alpha_{vd}}), \gamma_d \sim \mathcal{G}(a_\gamma, b_{\gamma_d}), \tau_v^{(b)} \sim \mathcal{G}(a_\tau, b_{\tau_v^{(b)}}) \\
\mathbf{a}_v &\sim \mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\alpha}_v)^{-1}), \mathbf{s}_t^{(b)} \sim \mathcal{N}(\mathbf{0}, \text{diag}(\boldsymbol{\gamma})^{-1}) \\
\mathbf{x}_t^{(b)} &\sim \mathcal{N}(\mathbf{A}\mathbf{s}_t^{(b)}, \text{diag}(\boldsymbol{\tau}^{(b)})^{-1}),
\end{aligned}$$

in which d indexes the latent space dimension D , v indexes voxels, t indexes time and b indexes subjects. The parameter α_{vd} is the precision on the spatial maps in the matrix \mathbf{A} , and acts as a sparsity pattern. The parameter γ_d is the precision on the d 'th component in the time-courses. Thus we have two 'forces' that can prune in the model, α to make the maps sparse and γ to prune away irrelevant components.

Finding the posterior $P(\boldsymbol{\theta}|\mathbf{X})$ is analytically intractable and an approximate solution is found through variational inference, as originally proposed for PCA by [17]. For the psFA model the mean-field approximation to the posterior,

$$\begin{aligned}
Q(\boldsymbol{\theta}|\mathbf{X}) &= \prod_{v=1}^V \mathcal{N}(\mathbf{a}_v | \boldsymbol{\mu}_{\mathbf{A}_v}, \boldsymbol{\Sigma}_{\mathbf{A}}^{(v)}) \prod_{b,t=1}^{B,T} \mathcal{N}(\mathbf{s}_t^{(b)} | \boldsymbol{\mu}_{\mathbf{S}_t}^{(b)}, \boldsymbol{\Sigma}_{\mathbf{S}}^{(b)}) \\
&\prod_{b,v=1}^{B,V} \mathcal{G}(\tau_v^{(b)} | \tilde{a}_\tau, \tilde{b}_{\tau_v^{(b)}}) \prod_{d=1}^D \mathcal{G}(\gamma_d | \tilde{a}_\gamma, \tilde{b}_{\gamma_d}) \prod_{d,v=1}^{D,v} \mathcal{G}(\alpha_{vd} | \tilde{a}_\alpha, \tilde{b}_{\alpha_{vd}}),
\end{aligned}$$

is used due to its similarity to the actual $P(\cdot)$ distributions yielding closed form solutions in the update rules. We use coordinate ascent variational inference, updating the moments of each variational distribution in a cyclic fashion conditioning on the other moments. The derived moments and a MATLAB implementation of the method are provided online¹.

The computational burden of the proposed model lies in calculating $\boldsymbol{\Sigma}_{\mathbf{A}}^{(v)}$ for each voxel v , which has $O(D^3V)$ time complexity due to inversion of V matrices of size $D \times D$. These inversions are embarrassingly parallel, and can be calculated quickly using GPUs, but comes at the price of having to keep the matrices in memory, requiring $O(D^2V)$ space.

If desired, subject specific mean values can be modeled such that $\mathbf{X}^{(b)} = \mathbf{A}\mathbf{s}_t^{(b)} + \boldsymbol{\mu}^{(b)}$, where $\boldsymbol{\mu}^{(b)} \sim \mathcal{N}(\mathbf{0}, \beta^{-1}\mathbf{I}_V)$. In practice, we have removed the empirical mean values prior to analysis.

2.2 Motor-task Data

We investigate the proposed model on a motor task experiment, which was previously acquired and analyzed in [12, 18]. The experiment consisted of $B = 29$ young and healthy adults, scanned while performing a block design motor task. The participants were visually cued by a blinking light, to indicate either right (green light) or left (red light) hand finger tapping. Each scanning session

¹<https://brainconnectivity.compute.dtu.dk/>

consisted of 10 task blocks, where each block consisted of a sequence of four tasks, i.e. “right/rest/left/rest”, with in total 240 images for each session. The data was pre-processed using standard techniques and parameter settings of the SPM8 software package². Each subject was realigned to the mean volume (rigid-body), normalized to a common Montreal Neurological Institute (MNI) template, and resliced to native 3 mm³ resolution. Afterwards spatial smoothing with a 3D Gaussian kernel (6mm FWHM) was applied, wavelet despiking to remove temporal outliers [19], voxel means subtracted, and data detrended via high-pass filtering with a 128 s cutoff. Finally we applied a rough grey-matter mask with 48799 voxels, and afterwards each subject was z-scored individually.

2.3 Resting State Data

We used the resting-state data³ from [13] and applied the following pre-processing steps to sessions 014-104⁴ using SPM12. All sessions were coregistered to the first image of the first functional session (session 014), and all sessions were then jointly corrected for motion artefacts using a rigid-body transformation towards the mean volume. A T1-weighted anatomical image from session 012 was coregistered to the functional space and grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF) was segmented using the standard tissue probability map from SPM. All functional sessions were then highpass filtered (1/128 Hz), nuisance regressed using motion parameters and eroded CSF and WM masks, and wavelet despiked [19]. Finally, all sessions were resliced (due to a change in the number of slices after session 027) to the first session and smoothed using a FWHM 5mm Gaussian kernel. The GM-mask was then resliced to the functional images and thresholded yielding a data matrix of size 69430 voxels \times 518 timepoints for each session. Due to memory limitations on the GPU we only considered the 25 first sessions of the data. For visualization purposes, we normalized the components from psFA and MELODIC to MNI space (2 mm³ resolution) using the deformation field estimated in the segmentation step.

3 Results

For all analysis the psFA (or pFA) model the following parameters are fixed; $a_\gamma, a_\alpha, a_\tau, b_{\gamma_d}, b_{\tau_v^{(b)}}, b_{\alpha_{v,d}} = 1e-6$. Variational inference was performed for the remaining parameters (except α for pFA), starting from an initial initial solution where the elements of \mathbf{A} were drawn from a $\mathcal{N}(0, 1)$ distribution and the subject specific time courses were then back reconstructed, $\mathbf{S}^{(b)} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{X}^{(b)}$. In our analysis we try to mitigate the effect of local minima by running the psFA analysis multiple times with random initializations. We note that this is not sufficient to avoid local minima, and this should be investigated further. In

²<http://fil.ion.ucl.ac.uk/spm>

³<https://openfmri.org/dataset/ds000031/>

⁴Some sessions did not contain resting state data and were thus discarded

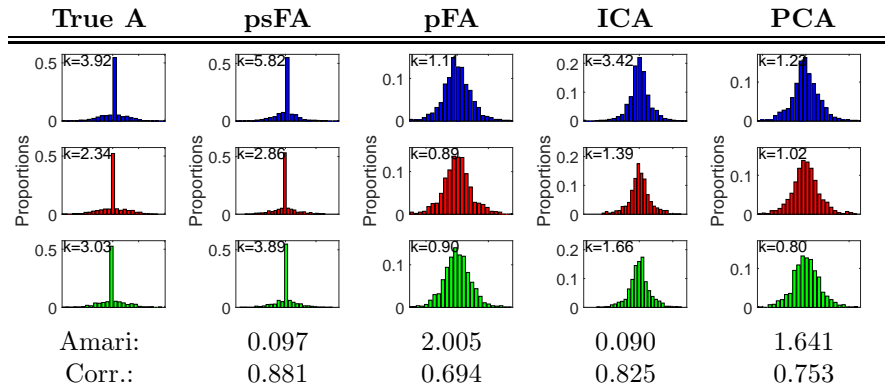


Figure 1: **Synthetic experiments:** Comparison between the true and estimated (first three) spatial maps \mathbf{A} . A histogram for each map is shown. Further, the Amari distance and average correlation between \mathbf{A}_{true} and \mathbf{A}_{est} is given.

all results in this section, only the run achieving the maximum lowerbound is further analysed.

3.1 Synthetic Experiments

We investigate the model in a synthetic setting with $B = 3$ subjects. For each subject, we generated three sources ($D = 3$), of length $T = 25$, from a normal distribution with zero mean and unit variance. These sources were then mapped to a higher dimensional space of size $V = 1000$ through a sparse matrix \mathbf{A} . The elements of \mathbf{A} were generated from a $\mathcal{N}(0, 1)$ and element-wise multiplied by a binary indicator from $\text{Uniform}(0, 1) > 0.5$. Heteroscedastic voxel- and subject-specific noise variance $\tau_v^{(b)-1}$ was drawn from $\mathcal{N}(0.009, 0.002)$. Noise drawn from $\mathcal{N}(0, \tau_v^{(b)-1})$ was then added to the corresponding voxel and subject, yielding a data set of size $V \times T \times B$.

The psFA (sparsity) and pFA (no sparsity) model were then run for 500 iterations with $D = 6$, with fifty random restarts, results are shown in Fig. 1. These methods are compared to regular PCA and infomax ICA⁵ on temporally concatenated data. The best performing methods are psFA and ICA, achieving high correlation and low Amari distance [20].

3.2 Motor-task Experiment

The psFA and pFA models were run for 1000 iterations with $D = 25$ components and we restarted the algorithm 5 times. The results obtained by psFA and pFA are contrasted to those found by MELODIC-ICA⁶ (pICA) with default settings, but using the same grey matter mask as described in section 2.2.

⁵<http://cogsys.imm.dtu.dk/toolbox/ica/>

⁶<http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/MELODIC>

The Pearson correlation between the estimated components and a set of reference maps was then calculated. The reference maps were, the default mode network (DMN) from [21] and eight anatomical regions from [22] which were: 1) Visual hOc1, hOc2, FG1, FG2; 2) Left, right sensorimotor, left and right motor cortex. For each model a visual and two motor components with highest absolute Pearson correlation to the reference maps are shown in Fig. 2. The components are sign corrected such that there is a positive correlation. The components found by psFA and pICA, Fig. 2a and 2c, have more well-defined spatial and temporal activation than those found by pFA, Fig. 2b. While pFA does capture the experimental design, the resulting spatial maps are more dense making them difficult to interpret. From the histograms, it is evident that both psFA and pICA enforce super-Gaussian distributions, which pFA does not.

The expectation of the log precision of the noise τ , averaged over subjects, is shown in Fig. 4a. As the estimated precision varies over voxels, this hints that the assumption of heteroscedastic noise is supported by the data. Furthermore, the regions of high precision are related to the experimental design, where a high signal to noise ratio is expected. The noise precision estimates by pFA are similar those of psFA and are therefore not shown.

3.3 Resting state Experiment

The psFA model was run for 2000 iterations with $D = 50$ components and five random restarts. In Fig. 3, we show the component from psFA and pICA that has highest correlation with the DMN reference map (see sec. 3.2). Both components have super-Gaussian shape (cf. histograms and kurtosis), and seem to capture the posterior part of the DMN. However, it seems the psFA obtains a more sparse solution. Finally, we observe from the precision-noise maps in Fig. 4b, where it seems that highest certainty is found in areas contributing to DMN.

4 Summary

In this work we investigated a scalable sparse probabilistic extension of factor analysis (psFA) for fMRI. We found in two data sets; a motor-task experiment with 29 subjects and a resting-state data set with 25 sessions, that inducing this form of sparsity results in ICA-like components. The probabilistic approach enables joint modeling of noise and quantification of parameters and their uncertainties. This comes at a computational cost, but due to the model structure a lot of computations are trivially parallelizable which we have exploited in the implementation. The fast inference scheme is currently limited by the memory on the single GPU used, but the proposed model easily extends to multiple GPUs. Future directions should be to investigate more advanced noise models and prediction on previously unseen data.

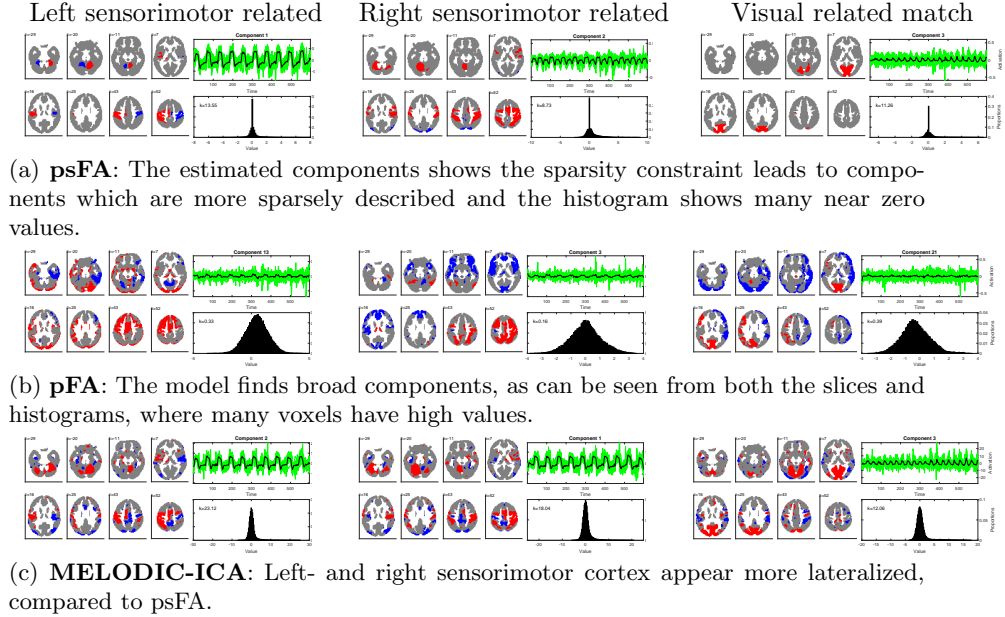


Figure 2: **Motor experiment** (psFA): The three estimated components for each model (psFA, pFA and MELODIC-ICA) which had highest correlation to sensorimotor (left and right) and visual related areas. For each component the following is shown; 1) eight slices with z-scored and thresholded (> 1) spatial activation (red: positive, blue: negative). 2) histogram (100 bins) of the spatial elements. We also report the empirical kurtosis k . 3) temporal activation of the component (black = mean over subjects, green = individual subjects).




Figure 3: **Resting state experiment**: For psFA (left) and MELODIC-ICA (right) we show the component with highest spatial correlation to the default mode network. We show; 1) Eight z-scored and thresholded (> 1) spatial activation slices (red: positive, blue: negative). 2) A histogram (100 bins) of the spatial elements. We also report the empirical kurtosis k .

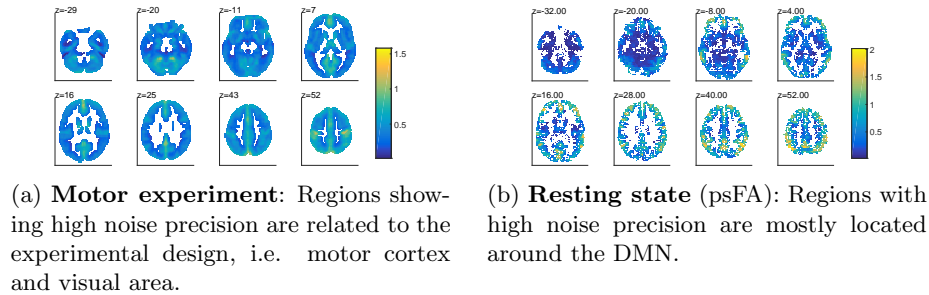


Figure 4: The expectation of log precision ($\langle \log(\tau^{(b)}) \rangle$), averaged over subjects/sessions.

References

- [1] A H Andersen, D M Gash, and M J Avison. Principal component analysis of the dynamic response measured by fmri: a generalized linear systems framework. *Magnetic Resonance Imaging*, 17(6):795–815, 1999.
- [2] M J McKeown, S Makeig, G G Brown, T-P Jung, S S Kindermann, A J Bell, and T J Sejnowski. Analysis of fmri data by blind separation into independent spatial components. Technical report, DTIC Document, 1997.
- [3] V D Calhoun, T Adali, G D Pearlson, and J J Pekar. A method for making group inferences from functional MRI data using independent component analysis. *Human brain mapping*, 14(3):140–151, November 2001.
- [4] T Adali, M Anderson, and Geng-Shen Fu. Diversity in independent component and vector analyses: Identifiability, algorithms, and applications in medical imaging. *Signal Processing Magazine, IEEE*, 31(3):18–33, May 2014.
- [5] I Daubechies, E Roussos, S Takerkart, M Benharrosh, C Golden, K D’ardenne, W Richter, JD Cohen, and J Haxby. Independent component analysis for brain fmri does not select for independence. *Proceedings of the National Academy of Sciences*, 106(26):10415–10422, 2009.
- [6] V D Calhoun, V K Potluru, R Phlypo, R F Silva, B A Pearlmutter, A Caprihan, S M Plis, and T Adali. Independent component analysis for brain fmri does indeed select for maximal independence. *PloS one*, 8(8):e73309, 2013.
- [7] C F Beckmann and S M Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE transactions on medical imaging*, 23(2):137–152, 2004.
- [8] C Chen and L-M Liu. Joint estimation of model parameters and outlier effects in time series. *Journal of the American Statistical Association*, 88(421):284–297, 1993.

- [9] R E Roge, K H Madsen, M Schmidt, and M Morup. Unsupervised segmentation of task activated regions in fMRI. In *Machine Learning for Signal Processing (MLSP), 2015 IEEE 25th International Workshop on*, pages 1–6. ieeexplore.ieee.org, 2015.
- [10] J L Hinrich, S Bardenfleth, R Roge, N Churchill, K H Madsen, and M Morup. Archetypal analysis for modeling Multi-Subject fMRI data. *IEEE journal of selected topics in signal processing*, PP(99):1–1, 2016.
- [11] D. J. MacKay. Probable networks and plausible predictions—a review of practical bayesian methods for supervised neural networks. *Network: Computation in Neural Systems*, 1995.
- [12] P M Rasmussen, L K Hansen, K H Madsen, N W Churchill, and S C Strother. Model sparsity and brain pattern interpretation of classification models in neuroimaging. *Pattern Recognition*, 45(6):2085–2100, 2012.
- [13] Poldrack, R. A. *et al.* Long-term neural and physiological phenotyping of a single human. *Nature communications*, 6:8885, 9 December 2015.
- [14] Yue Guan and Jennifer G Dy. Sparse probabilistic principal component analysis. In *AISTATS*, pages 185–192, 2009.
- [15] J L Hinrich, S F V Nielsen, K H Madsen, and M Mørup. Variational group-pca for intrinsic dimensionality determination in fmri data. *The 6th International Workshop on Pattern Recognition in Neuroimaging (PRNI 2016)*, 2016.
- [16] F B Nielsen. Variational approach to factor analysis and related models, 2004.
- [17] C M Bishop. Variational principal components. In *Proceedings of the 1999 the 9th International Conference on 'Artificial Neural Networks (ICANN99)*, pages 509–514. IEEE, 1999.
- [18] P M Rasmussen, T J Abrahamsen, K H Madsen, and L K Hansen. Nonlinear denoising and analysis of neuroimages with kernel principal component analysis and pre-image estimation. *NeuroImage*, 60(3):1807–1818, 2012.
- [19] A X Patel, P Kundu, M Rubinov, S Jones, P E Vértes, K D Ersche, J Suckling, and E T Bullmore. A wavelet method for modeling and despiking motion artifacts from resting-state fmri time series. *Neuroimage*, 95:287–304, 2014.
- [20] F R Bach and M I Jordan. Kernel independent component analysis. *Journal of machine learning research*, 3(Jul):1–48, 2002.
- [21] A R Franco, A Pritchard, V D Calhoun, and A R Mayer. Interrater and intermethod reliability of default mode network selection. *Human brain mapping*, 30(7):2293–2303, 2009.

- [22] S B Eickhoff, K E Stephan, H Mohlberg, C Grefkes, G R Fink, K Amunts, and Karl Zilles. A new spm toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *Neuroimage*, 25(4):1325–1335, 2005.

Appendix

The derived details for probabilistic sparse factor analysis model for group level analysis are given. The model was first proposed by the authors and applied to functional magnetic resonance imaging[?].

Probabilistic Sparse Factor Analysis

The model's likelihood function, for a data array of size $V \times T \times B$ with all relevant parameters collected in $\boldsymbol{\theta}$, can be written as,

$$\mathcal{L}(\mathbf{X}|\boldsymbol{\theta}) = \prod_{b=1}^B \prod_{t=1}^T \mathcal{N}\left(\mathbf{x}_t^{(b)} | \mathbf{A}\mathbf{s}_t^{(b)} + \boldsymbol{\mu}^{(b)}, \text{diag}\left(\boldsymbol{\tau}^{(b)}\right)^{-1}\right), \quad (1)$$

in which $\mathbf{x}_t^{(b)}$ is a vector of length V , \mathbf{A} is a matrix of size $V \times D$, where D is the size of the latent space, $\mathbf{s}_t^{(b)}$ is a vector of size D and $\boldsymbol{\tau}^{(b)}$ is a vector of length V . The distributions of the parameters in the model are,

$$\begin{aligned} P(\mathbf{A}|\boldsymbol{\alpha}) &= \prod_{v=1}^V \mathcal{N}\left(\mathbf{a}_v | \mathbf{0}, \text{diag}(\boldsymbol{\alpha}_v)^{-1}\right) \\ P(\mathbf{S}|\boldsymbol{\gamma}) &= \prod_{b=1}^B \prod_{t=1}^T \mathcal{N}\left(\mathbf{s}_t^{(b)} | \mathbf{0}, \text{diag}(\boldsymbol{\gamma})^{-1}\right) \\ P(\boldsymbol{\mu}) &= \prod_{b=1}^B \mathcal{N}\left(\boldsymbol{\mu}^{(b)} | \mathbf{0}, \beta^{-1}\mathbf{I}_V\right) \\ P(\boldsymbol{\tau}) &= \prod_{b=1}^B \prod_{v=1}^V \mathcal{G}\left(\tau_v^{(b)} | a_\tau, b_{\tau_v^{(b)}}\right) \\ P(\boldsymbol{\gamma}) &= \prod_{d=1}^D \mathcal{G}\left(\gamma_d | a_\gamma, b_{\gamma_d}\right) \\ P(\boldsymbol{\alpha}) &= \prod_{v=1}^V \prod_{d=1}^D \mathcal{G}\left(\alpha_{vd} | a_\alpha, b_{\alpha_{vd}}\right). \end{aligned}$$

Finding the parameters $\boldsymbol{\theta}$ from observed data \mathbf{X} can be done by inferring the posterior distribution,

$$P(\boldsymbol{\theta}|\mathbf{X}) = \mathcal{L}(\mathbf{X}|\boldsymbol{\theta})P(\boldsymbol{\theta})/P(\mathbf{X}).$$

Unfortunately an exact inference is unfeasible for all but the simplest problems. Therefore an approximate solution is sought. While there are numerous ways to tackle this, we use variational Bayesian (VB) inference as used by [17] for VB principal components analysis. We use a mean field approximation, and find the following variational distribution to approximate the posterior.

$$\begin{aligned}
Q(\mathbf{A}) &= \prod_{v=1}^V \mathcal{N}(\mathbf{a}_v | \boldsymbol{\mu}_{\mathbf{A}_v}, \boldsymbol{\Sigma}_{\mathbf{A}}^{(v)}) \\
Q(\mathbf{S}) &= \prod_{b,t=1}^{B,T} \mathcal{N}(\mathbf{s}_t^{(b)} | \boldsymbol{\mu}_{\mathbf{S}_t}^{(b)}, \boldsymbol{\Sigma}_{\mathbf{S}}^{(b)}) \\
Q(\boldsymbol{\mu}) &= \prod_{b=1}^B \mathcal{N}(\boldsymbol{\mu}^{(b)} | \boldsymbol{\mu}_{\boldsymbol{\mu}}^{(b)}, \boldsymbol{\Sigma}_{\boldsymbol{\mu}}^{(b)}) \\
Q(\boldsymbol{\tau}) &= \prod_{b,v=1}^{B,V} \mathcal{G}(\tau_v^{(b)} | \tilde{a}_{\tau}, \tilde{b}_{\tau_v^{(b)}}) \\
Q(\boldsymbol{\gamma}) &= \prod_{d=1}^D \mathcal{G}(\gamma_d | \tilde{a}_{\gamma}, \tilde{b}_{\gamma_d}) \\
Q(\boldsymbol{\alpha}) &= \prod_{d,v=1}^{D,v} \mathcal{G}(\alpha_{vd} | \tilde{a}_{\alpha}, \tilde{b}_{\alpha_{vd}})
\end{aligned}$$

The moments of a distribution are then found by conditioning them on all other distributions and using free-form optimization (see [17]).

Update rules

The found moments of the distributions are updated cyclically using Expectation-Maximization. In each iteration, after all distributions have been updated, the evidence lowerbound (ELBO) is calculated. After a number of iterations, when the relative change in ELBO is below a given threshold a set of local optimal

parameters is identified.

$$\begin{aligned}
\Sigma_{\mathbf{A}}^v &= \left(\text{diag} \langle \boldsymbol{\alpha}_v \rangle + \sum_{b=1}^B \langle \tau_v^{(b)} \rangle \langle \mathbf{S}^{(b)} \mathbf{S}^{(b)\top} \rangle \right)^{-1} \\
\boldsymbol{\mu}_{\mathbf{A}}^v &= \Sigma_{\mathbf{A}}^v \left(\sum_{b=1}^B \langle \tau_v^{(b)} \rangle \sum_{t=1}^T \langle \mathbf{s}_t^{(b)} \rangle (x_{tv}^{(b)} - \langle \boldsymbol{\mu}_v^{(b)} \rangle) \right) \\
\Sigma_{\mathbf{S}}^{(b)} &= \left(\text{diag} \langle \boldsymbol{\gamma} \rangle + \langle \mathbf{A}^\top \text{diag}(\boldsymbol{\tau}^{(b)}) \mathbf{A} \rangle \right)^{-1} \\
\boldsymbol{\mu}_{\mathbf{S}_t}^{(b)} &= \Sigma_{\mathbf{S}}^{(b)} \langle \mathbf{A}^\top \rangle \text{diag} \langle \boldsymbol{\tau}^{(b)} \rangle (\mathbf{x}_t^{(b)} - \langle \boldsymbol{\mu}^{(b)} \rangle) \\
\Sigma_{\boldsymbol{\mu}}^{(b)} &= \left(\beta \mathbf{I}_V + \text{diag} \langle \boldsymbol{\tau}^{(b)} \rangle \right)^{-1} \\
\boldsymbol{\mu}_{\boldsymbol{\mu}}^{(b)} &= \Sigma_{\boldsymbol{\mu}}^{(b)} \text{diag} \langle \boldsymbol{\tau}^{(b)} \rangle \sum_{t=1}^T (\mathbf{x}_t^{(b)} - \langle \mathbf{A} \rangle \langle \mathbf{s}_t^{(b)} \rangle) \\
\tilde{a}_\alpha &= a_\alpha + \frac{1}{2} \quad , \quad \tilde{b}_{\alpha_{vd}} = b_{\alpha_{vd}} + \langle a_{vd}^2 \rangle \\
\tilde{a}_\gamma &= a_\gamma + \frac{1}{2} \sum_{b=1}^B T^{(b)} \quad , \quad \tilde{b}_{\gamma_d} = b_{\gamma_d} + \frac{1}{2} \sum_{b=1}^B \text{trace}(\langle \mathbf{s}_d \mathbf{s}_d^\top \rangle) \\
\tilde{a}_{\tau^{(b)}} &= a_\tau + \frac{T^{(b)}}{2} \\
\tilde{b}_{\tau_v^{(b)}} &= b_{\tau_v^{(b)}} + \frac{1}{2} \left[\|\mathbf{x}_v^{(b)}\|_{\text{Fro}}^2 + T^{(b)} \langle \mu_v^{(b)2} \rangle \right. \\
&\quad - 2 \left(\langle \mathbf{a}_v \rangle \langle \mathbf{S}^{(b)} \rangle + \langle \mu_v^{(b)} \rangle \right) \mathbf{x}_v^\top \\
&\quad + 2 \langle \mathbf{a}_v \rangle \langle \mathbf{S}^{(b)} \rangle \mathbf{1}_{T^{(b)}} \langle \mu_v^{(b)} \rangle \\
&\quad \left. + \text{trace} \left(\langle \mathbf{a}_v^\top \mathbf{S}^{(b)} \mathbf{S}^{(b)\top} \mathbf{a}_v \rangle \right) \right]
\end{aligned}$$

Note $\langle \cdot \rangle$ is the expected value under the variational distributions. Further, using the properties of the trace operator, the expected value of the expression in $\Sigma_{\mathbf{S}}^{(b)}$ and $\tilde{b}_{\tau_v^{(b)}}$ are determined to be,

$$\begin{aligned}
\text{trace} \left(\langle \mathbf{a}_v^\top \mathbf{S}^{(b)} \mathbf{S}^{(b)\top} \mathbf{a}_v \rangle \right) &= \text{trace} \left(\langle \mathbf{S}^{(b)} \mathbf{S}^{(b)\top} \rangle \Sigma_{\mathbf{A}}^v \right) \\
&\quad + \langle \mathbf{a}_v^\top \rangle \langle \mathbf{S}^{(b)} \mathbf{S}^{(b)\top} \rangle \langle \mathbf{a}_v \rangle
\end{aligned}$$

and

$$\begin{aligned}
\langle \mathbf{A}^\top \text{diag}(\boldsymbol{\tau}^{(b)}) \mathbf{A} \rangle &= \left(\sum_{v=1}^V \Sigma_{\mathbf{A}}^v \langle \tau_v^{(b)} \rangle \right) \\
&\quad + \langle \mathbf{A}^\top \rangle \text{diag} \langle \boldsymbol{\tau}^{(b)} \rangle \langle \mathbf{A} \rangle .
\end{aligned}$$

Evidence Lowerbound (ELBO)

The evidence lowerbound is the sum of all the expression in this section. It can be divided into to categories; 1) the expected value of the P-distributions under the Q-distributions (i.e. substituting the moments of the P-distributions for the moments of the Q-distributions). 2) The entropy of the Q-distributions.

For each P-distribution the expected value the corresponding Q-distribution is given below,

$$\begin{aligned}
\langle \log P(\mathbf{A}|\boldsymbol{\alpha}) \rangle &= \sum_{v=1}^V -\frac{1}{2} \log(2\pi) + \frac{1}{2} \langle \log \alpha_{vd} \rangle - \frac{1}{2} \langle \alpha_{vd} \rangle \langle a_{vd}^2 \rangle \\
\langle \log P(\mathbf{S}|\boldsymbol{\gamma}) \rangle &= \sum_{b=1}^B \sum_{t=1}^{T^{(b)}} -\frac{D}{2} \log(2\pi) + \left(\frac{1}{2} \sum_{d=1}^D \langle \log \gamma_d \rangle \right) \\
&\quad - \frac{1}{2} \text{trace} \left(\text{diag} \langle \boldsymbol{\gamma} \rangle \langle \mathbf{s}_t^{(b)} \mathbf{s}_t^{(b)\top} \rangle \right) \\
\langle \log P(\boldsymbol{\mu}) \rangle &= \sum_{b=1}^B -\frac{V}{2} \log(2\pi) + \frac{V}{2} \log(\beta) - \frac{1}{2} \beta \langle \boldsymbol{\mu}_{\boldsymbol{\mu}}^{(b)\top} \boldsymbol{\mu}_{\boldsymbol{\mu}}^{(b)} \rangle \\
\langle \log P(\boldsymbol{\alpha}) \rangle &= \sum_{d=1}^D \sum_{v=1}^V -\log(\Gamma(a_{\alpha})) + a_{\alpha} \log(b_{\alpha_{vd}}) \\
&\quad + (a_{\alpha} - 1) \langle \log \alpha_{vd} \rangle - b_{\alpha_{vd}} \langle \alpha_{vd} \rangle \\
\langle \log P(\boldsymbol{\gamma}) \rangle &= \sum_{d=1}^D -\log(\Gamma(a_{\gamma})) + a_{\gamma} \log(b_{\gamma_d}) \\
&\quad + (a_{\gamma} - 1) \langle \log \gamma_d \rangle - b_{\gamma_d} \langle \gamma_d \rangle \\
\langle \log P(\boldsymbol{\tau}) \rangle &= \sum_{b=1}^B \sum_{v=1}^V -\log(\Gamma(a_{\tau})) + a_{\tau} \log(b_{\tau_v^{(b)}}) \\
&\quad + (a_{\tau} - 1) \langle \log \tau_v^{(b)} \rangle - b_{\tau_v^{(b)}} \langle \tau_v^{(b)} \rangle
\end{aligned}$$

Finally, the entropy for each Q-distribution is,

$$\begin{aligned}
- \langle \log Q(\mathbf{A}) \rangle &= \sum_{v=1}^V \left[\frac{1}{2} \log |\Sigma_{\mathbf{A}}^v| + \frac{D}{2} (1 + \log(2\pi)) \right] \\
- \langle \log Q(\mathbf{S}) \rangle &= \sum_{b=1}^B \sum_{t=1}^{T^{(b)}} \left[\frac{1}{2} \log |\Sigma_{\mathbf{S}}^{(b)}| + \frac{D}{2} (1 + \log(2\pi)) \right] \\
- \langle \log Q(\boldsymbol{\mu}) \rangle &= \sum_{b=1}^B \left[\frac{1}{2} \log |\Sigma_{\boldsymbol{\mu}}^{(b)}| + \frac{V}{2} (1 + \log(2\pi)) \right] \\
- \langle \log Q(\boldsymbol{\alpha}) \rangle &= \sum_{d=1}^D \sum_{v=1}^V \log(\Gamma(\tilde{a}_{\alpha_{vd}})) - (\tilde{a}_{\alpha} - 1)\psi(\tilde{a}_{\alpha}) \\
&\quad - \log(\tilde{b}_{\alpha_{vd}}) + \tilde{a}_{\alpha} \\
- \langle \log Q(\boldsymbol{\gamma}) \rangle &= \sum_{d=1}^D \log(\Gamma(\tilde{a}_{\gamma})) - (\tilde{a}_{\gamma} - 1)\psi(\tilde{a}_{\gamma}) - \log(\tilde{b}_{\gamma_d}) + \tilde{a}_{\gamma} \\
- \langle \log Q(\boldsymbol{\tau}) \rangle &= \sum_{b=1}^B \sum_{v=1}^V \log(\Gamma(\tilde{a}_{\tau^{(b)}})) - (\tilde{a}_{\tau^{(b)}} - 1)\psi(\tilde{a}_{\tau}) \\
&\quad - \log(\tilde{b}_{\tau^{(b)}}) + \tilde{a}_{\tau^{(b)}}
\end{aligned}$$

Implementation

A MATLAB implementation is available⁷. The implementation is limited to the use of a single GPU card, as well as analysis with $T^{(i)} = T^{(j)}, \forall i, j$. These limitations were deemed acceptable for the work in [?], as subjects with differing timesteps are not widespread in the field.

⁷<https://brainconnectivity.compute.dtu.dk/>