MEAN FIELD IMPLEMENTATION OF BAYESIAN ICA

Pedro Højen-Sørensen[*], *Lars Kai Hansen*[†]

Informatics and Mathematical Modelling
Technical University of Denmark, B321
DK-2800 Lyngby, Denmark

Ole Winther[‡]

Center for Biological Sequence Analysis
Technical University of Denmark, B208
DK-2800 Lyngby, Denmark

ABSTRACT

In this contribution we review the mean field approach to Bayesian independent component analysis (ICA) recently developed by the authors [1, 2]. For the chosen setting of additive Gaussian noise on the measured signal and Maximum Likelihood II estimation of the mixing matrix and the noise, the expected sufficient statistics are obtained from the two first posterior moments of the sources. These can be effectively estimated using variational mean field theory and its linear response correction. We give an application to feature extraction in neuro-imaging using a binary (stimuli/no stimuli) source paradigm. Finally, we discuss the possibilities of extending the framework to convolutive mixtures, temporal and 'spatial' source prior correlations, identification of common sources in mixtures of different media and ICA for density estimation.

1. INTRODUCTION

Independent component analysis (ICA) is an active research field that has seen the development of many successful algorithms for blind separation of source signals, for a review see [3]. The algorithms have either been statistical based upon maximum likelihood and information maximization or have exploited higher-order moments information for performing the source separation.

Bayesian methods are formally straight forward extensions of maximum likelihood approaches, i.e. average over the uncertainty of the model rather than optimizing model parameters. Instead of just selecting point estimates according to the maximum of the posterior densities, the Bayesian paradigm takes into account the entire probability mass of the posterior densities which, in turn, leads to better estimates. However, the averages – which are often high dimensional – are analytically intractable and one thus has to resort to Monte Carlo integration [4].

* Current address at Nokia Mobile Phones, Algorithm Design, Denmark Pedro.Hoejen-Soerensen@nokia.com.

† lkhansen@imm.dtu.dk

‡ winther@cbs.dtu.dk

An interesting alternative that has attracted much recent interest in many areas of probabilistic modeling, e.g. graphical models, is mean field theory (MFT). MFT covers a range of approximate techniques where variational MFT (often called ensemble learning in neural computation) is the most widely used. Other more advanced methods are linear response theory and TAP for densely connected systems and structured mean field methods such as belief propagation and Kickuchi cluster variational approximations for sparsely connected systems, see [5] for a review. Common to all MFT is that a – hopefully unimportant – part of the correlations between the stochastic variables is ignored. The resulting approximation is either written in terms of a set of non-linear fix point (mean field) equations in the primal parameters, typically the posterior mean of the stochastic variables, or in terms of the dual set of parameters as in belief propagation. MFT also gives an estimate of the likelihood of the model, i.e. the probability of the observed data given the model.

A number of authors have previously applied mean field theory to ICA [6, 7]. We will briefly discuss the relation of our contributions [1, 2] to these. Attias [6] uses a mixture of Gaussian source prior in a maximum likelihood ML-II setting, i.e. only average over the primary variables, the sources, and estimating secondary level parameters; the mixing matrix and noise by maximizing the likelihood of the model. This gives rise to an Expectation Maximization (EM) type algorithm. In our contribution we also use ML-II but generalizes the analysis to arbitrary source priors and furthermore use linear response theory to give an improved estimate of covariances (needed for estimating secondary level parameters). Variational mean field trivially predicts the covariances to be $\langle S_i S_i \rangle = \langle S_i \rangle \langle S_j \rangle$ for $i \neq j$. This has turned out to be important in hard e.g. very noisy problems. Miskin and MacKay [7] also introduces priors for mixing matrix and noise. With their specific choices of priors they can give perform a full variational treatment within the Bayesian paradigm at the expense of getting a somewhat more involved system of mean field equations to solve. The fully Bayesian framework should be advantageous in cases where the distribution of secondary level parameters are *not* sharply peaked. However, these parameters which are typically assumed to be 'time independent' are often more well determined than the sources. Linear response theory should also be able to give improved second order statistics estimates for this framework.

The rest of the paper is organized as follows. In sections 2-4, we discuss the model, the estimation problem and mean field theory. We give an application to neuro-imaging in section 5 and wrap up with an outlook and conclusion in sections 6 and 7.

2. MODEL

In the simplest possible ICA-model the measured signal $\mathbf{X}_t$ at time t is assumed to an instantaneous linear mixing of the sources corrupted with additive white Gaussian noise $\mathbf{\Gamma}_t$

$$\mathbf{X}_t = \mathbf{A}\mathbf{S}_t + \mathbf{\Gamma}_t \;, \tag{1}$$

where $\mathbf{A}$ is a (time independent) mixing matrix and the noise is assumed to be without temporal correlations and with time independent covariance matrix $\mathbf{\Sigma}$. The likelihood for parameters and sources at time t is thus

$$P(\mathbf{X}_t|\mathbf{A},\mathbf{\Sigma},\mathbf{S}_t) = \mathcal{N}(\mathbf{X}_t|\mathbf{A}\mathbf{S}_t,\mathbf{\Sigma}) \;, \tag{2}$$

where $\mathcal{N}$ is shorthand for a multi-variate Gaussian distribution. The likelihood for the entire set of observations, $\mathbf{X} = \{\mathbf{X}_t\}$, is simply obtained as the product

$$P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma},\mathbf{S}) = \prod_{t=1}^{N} P(\mathbf{X}_t|\mathbf{A},\mathbf{\Sigma},\mathbf{S}_t) \;. \tag{3}$$

The aim of independent component analysis is to recover the unknown quantities: the sources $\mathbf{S}$, the mixing matrix $\mathbf{A}$ and the noise covariance $\mathbf{\Sigma}$ from the observed data using the further assumption of statistical independence of the sources $P(\mathbf{S}_t) = \prod_{m=1}^{M} P(S_{mt})$ and $P(\mathbf{S}) = \prod_t P(\mathbf{S}_t)$.

3. SUFFICIENT STATISTICS

Under a quadratic loss, the Bayes optimal estimate of the sources is the posterior mean $\langle\mathbf{S}\rangle$, where we have denoted an average over the posterior

$$P(\mathbf{S}|\mathbf{X},\mathbf{A},\mathbf{\Sigma}) = \frac{P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma},\mathbf{S})P(\mathbf{S})}{P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma})} \tag{4}$$

by angle brackets $\langle\cdot\rangle$. A valid criticism of the statistical approach is that it requires knowledge of $P(\mathbf{S})$ which is usually not available. However, in practice it turns out that the detailed form is not very important (but rather whether the prior is sub- or super-Gaussian). Secondly, within the Bayesian maximum likelihood II (MLII) approach, the likelihood

$$P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma}) = \int d\mathbf{S} P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma},\mathbf{S})\, P(\mathbf{S}) \tag{5}$$

can be used for model selection between different priors and data analysis models such as Bayesian PCA. We can thus let the data decide which prior is the appropriate for the data. In a fully Bayesian approach one should rather average over different possible priors. Here, we will limit the averaging to the first level of parameters, i.e. the sources.

Maximum likelihood II (MLII) estimation of $\mathbf{A}$ and $\mathbf{\Sigma}$, i.e. maximizing the likelihood $P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma})$ leads to [1]

$$\mathbf{A}_{\mathrm{MLII}} = \sum_t \mathbf{X}_t \langle\mathbf{S}_t\rangle^T \left(\sum_{t'} \langle\mathbf{S}_{t'}\mathbf{S}_{t'}^T\rangle\right)^{-1} \tag{6}$$

$$\mathbf{\Sigma}_{\mathrm{MLII}} = \frac{1}{N}\sum_t \langle(\mathbf{X}_t - \mathbf{A}\mathbf{S}_t)(\mathbf{X}_t - \mathbf{A}\mathbf{S}_t)^T\rangle \;. \tag{7}$$

The sufficient statistics of the model is thus the first and second moments of the sources. Further non-negativity constraints on $\mathbf{A}$ relevant for e.g. feature extraction in images (see section 5) and text mining can be enforced via Lagrange multipliers [1].

One advantage of the Bayesian approach compared to maximum a posteriori (MAP)

$$\mathbf{S}_{\mathrm{MAP}} = \operatorname*{argmax}_{\mathbf{S}} P(\mathbf{S}|\mathbf{X},\mathbf{A},\mathbf{\Sigma}) \tag{8}$$

is that we take correlation of sources into account, i.e. compare eq. (6) with the extremum of $P(\mathbf{S}|\mathbf{X},\mathbf{A},\mathbf{\Sigma})$. Another advantage is that we straightforwardly can deal with any kind of variables, e.g. discrete or non-negative. We will give an application to binary variables in section 5.

The clear disadvantage is that the sufficient statistics cannot be obtained analytically apart for the Gaussian case which is not relevant for instantaneous ICA mixtures. Mean field theory – as presented in the next section – represents a powerful alternative to the most general approach, Monte Carlo integration.

4. MEAN FIELD THEORY

In a standard variational mean field theoretic approach, the posterior distribution, $P(\mathbf{S}|\mathbf{X},\mathbf{A},\mathbf{\Sigma})$ is approximated in a family of product distributions $Q(\mathbf{S}) = \prod_{mt} Q(S_{mt})$. For the Gaussian likelihood $P(\mathbf{X}|\mathbf{A},\mathbf{\Sigma},\mathbf{S})$, we set [1]

$$Q(S_{mt}) \propto P(S_{mt}) e^{-\frac{1}{2}\lambda_{mt} S_{mt}^2 + \gamma_{mt} S_{mt}} \;. \tag{9}$$

We can simplify the notation by writing

$$P(\mathbf{X}_t|\mathbf{A},\mathbf{\Sigma},\mathbf{S}_t) = P(\mathbf{X}_t|\mathbf{J},\mathbf{h}_t,\mathbf{S}_t) \propto e^{-\frac{1}{2}\mathbf{S}_t^T\mathbf{J}\mathbf{S}_t + \mathbf{h}_t^T\mathbf{S}_t} \;,$$

where the interaction matrix $\mathbf{J}$ and the field $\mathbf{h}_t$ are given by

$$\begin{aligned} \mathbf{J} &= \mathbf{A}^T \mathbf{\Sigma}^{-1} \mathbf{A} & (10) \\ \mathbf{h}_t &= \mathbf{A}^T \mathbf{\Sigma}^{-1} \mathbf{X}_t \, . & (11) \end{aligned}$$

The starting point of the variational derivation of mean field equations is the Kullback-Leibler divergence between the product distribution $Q(\mathbf{S})$ and the true source posterior, i.e.

$$\begin{aligned} KL &= \int d\mathbf{S} Q(\mathbf{S}) \log \frac{Q(\mathbf{S})}{P(\mathbf{S}|\mathbf{X}, \mathbf{A}, \mathbf{\Sigma})} & (12) \\ &= \log P(\mathbf{X}|\mathbf{A}, \mathbf{\Sigma}) - \log P(\mathbf{X}|\mathbf{A}, \mathbf{\Sigma}, \mathrm{NMF}) \, , \end{aligned}$$

where $P(\mathbf{X}|\mathbf{A}, \mathbf{\Sigma}, \mathrm{NMF})$ is the naive mean field approximation to the likelihood (see [1] for an explicit expression). The Kullback-Leibler is zero when $P = Q$ and positive otherwise. The parameters of Q should consequently be chosen as to minimize KL. The saddle points define the mean field equations of which the two first yield

$$\begin{aligned} \boldsymbol{\gamma}_t &= \mathbf{h}_t - (\mathbf{J} - \mathrm{diag}(\mathbf{J})) \langle \mathbf{S}_t \rangle & (13) \\ \lambda_{mt} &= J_{mm} \, , & (14) \end{aligned}$$

where $\mathrm{diag}(\mathbf{J})$ is the diagonal matrix of $\mathbf{J}$. The remaining two equations $\frac{\partial KL}{\partial \gamma_{mt}} = 0$ and $\frac{\partial KL}{\partial \lambda_{mt}} = 0$ depend explicitly on the source prior, $P(S)$. The first define the variational mean $f(\gamma_{mt}, \lambda_{mt})$:

$$\begin{aligned} \langle S_{mt} \rangle^{\mathrm{VMF}} &= \frac{\partial}{\partial \gamma_{mt}} \log \int dS P(S) e^{-\frac{1}{2}\lambda_{mt} S^2 + \gamma_{mt} S} \\ &\equiv f(\gamma_{mt}, \lambda_{mt}) \, , & (15) \end{aligned}$$

where $\langle \cdot \rangle^{\mathrm{VMF}}$ denotes the variational mean field estimate of the posterior mean. Since the variational distribution is diagonal, the off-diagonal elements of the covariance matrix

$$\boldsymbol{\chi}_t \equiv \langle \mathbf{S}_t \mathbf{S}_t^T \rangle - \langle \mathbf{S}_t \rangle \langle \mathbf{S}_t^T \rangle \qquad (16)$$

estimates are zero. Using the condition $\frac{\partial KL}{\partial \lambda_{mt}} = 0$, we can write the variational MF estimate of $\boldsymbol{\chi}_t$ as

$$\chi^{\mathrm{VMF}}_{mm't} = \delta_{mm'} \frac{f(\gamma_{mt}, \lambda_{mt})}{\partial \gamma_{mt}} \, . \qquad (17)$$

Eqs. (13), (14) and (15) define the variational mean field equations. These can be solved by iteration together with equations for $\mathbf{A}$ and $\mathbf{\Sigma}$ (using eq. (17)) in an EM-fashion, see [1] for an algorithmic recipe.

We now turn to the derivation of the linear response (LR) estimate of $\boldsymbol{\chi}_t$. Notice that $\mathbf{h}_t$ acts as an external field from which all moments of the sources can be obtained exactly, i.e.

$$\begin{aligned} \langle S_{mt} \rangle &= \frac{\partial \log P(\mathbf{X}_t|\mathbf{J}, \mathbf{h}_t)}{\partial h_{mt}} & (18) \\ \chi_{mm't} &= \frac{\partial^2 \log P(\mathbf{X}_t|\mathbf{J}, \mathbf{h}_t)}{\partial h_{m't} \partial h_{mt}} = \frac{\partial \langle S_{mt} \rangle}{\partial h_{m't}} \, . & (19) \end{aligned}$$

If we in the second line instead use $\langle S_{mt} \rangle^{\mathrm{VMF}}$, we get the linear response estimate $\boldsymbol{\chi}^{\mathrm{LR}}_t \equiv \frac{\partial \langle \mathbf{S}_t \rangle^{\mathrm{VMF}}}{\partial \mathbf{h}_t^T}$ [1]:

$$\begin{aligned} \chi^{\mathrm{LR}}_{mm't} &= \frac{\partial f(\gamma_{mt}, \lambda_{mt})}{\partial \gamma_{mt}} \frac{\partial \gamma_{mt}}{\partial h_{m't}} \\ \frac{\partial \gamma_{mt}}{\partial h_{m't}} &= \delta_{mm'} - \sum_{m'', m'' \neq m} J_{mm''} \chi^{\mathrm{LR}}_{m''m't} \, . & (20) \end{aligned}$$

This equation is (by construction) linear in $\boldsymbol{\chi}^{\mathrm{LR}}_t$ and can be solved to give the linear response estimate of the covariances:

$$\boldsymbol{\chi}^{\mathrm{LR}}_t = (\mathbf{\Lambda}_t + \mathbf{J})^{-1} \, , \qquad (21)$$

where we have defined the diagonal matrix

$$\begin{aligned} \mathbf{\Lambda}_t &= \mathrm{diag}\,(\Lambda_{1t}, \ldots, \Lambda_{Mt}) & (22) \\ \Lambda_{mt} &\equiv \left(\frac{\partial f(\gamma_{mt}, \lambda_{mt})}{\partial \gamma_{mt}} \right)^{-1} - J_{mm} \, . & (23) \end{aligned}$$

Now we can simply use eq. (21) instead of the variational estimate eq. (17). In general we expect the LR estimate to be better than the variational [1]. Obtaining the LR correction gives an additional complexity of $\mathcal{O}(NM^3)$, where N denotes the number of time steps and M the number of sources. This poses no computational problem for typical applications. However, once convolutive models or priors with correlations – as discussed below – are considered, the increase in complexity can be prohibitively large.

5. FEATURE EXTRACTION IN NEUROIMAGING

The low signal-to-noise ratio and the many possible sources of variability makes recordings from non-invasive functional neuroimaging techniques a most challenging data analysis problem. In this section we apply the proposed mean field algorithm for noisy ICA with adaptive binary sources for exploratory analysis of functional magnetic resonance imaging (fMRI) data. The number of hidden sources is determined using the Bayesian information criterion (BIC) [8] in which the naive mean field energy is used as an approximation to the likelihood.

The typical fMRI activation study consists of maintaining a healthy volunteer in controlled mental states; typically a baseline control state and an active task state. The temporal course of the baseline/activation paradigm is denoted the *reference function* whereas the regional hemodynamic response to focal neuronal activation is denoted the BOLD signal (blood oxygenation level dependent (BOLD) contrast.

Bandettini et al. [9] analyzed the correlation between a binary reference function and the BOLD signal. Lange and Zeger [10] discuss a parameterized hemodynamic response adapted by a least squares procedure whereas multivariate strategies have been pursued in [11] and [12]. Several explorative strategies have been proposed for finding spatio-temporal activation patterns without explicit reference to the activation paradigm. McKeown et al. [13] used the independent component analysis algorithm of [14] to identify *spatially* independent patterns and found several types of activations including components with “transient task related” response, i.e., responses that could not simply be accounted for by the paradigm; such components would typically not be identified by a simple principal component analysis (PCA) since they would tend not to be orthogonal to the task. Other authors have argued for identifying *temporally* independent patterns, e.g. [15]; a comparative study of temporal and spatial ICA approaches for analyzing fMRI were carried out in [16]. While previous ICA approaches succeeded in finding task related components, the ICA schemes applied were not well-matched to the binary on/off character of the stimulus (see also [16] for further discussion of this point). Hence, in this section we will focus on the more appropriate binary source distribution assumption. Furthermore, it is interesting to note that when analyzing the signals for spatially independent components, the binary source assumption corresponds to another popular approach for exploratory analysis of fMRI, namely wave form clustering. The binary spatial component can be viewed as the *binary association* of pixels with the corresponding time course.

In this particular application we have restricted ourself to consider binary $\{0, 1\}$–sources parameterized by

$$p(s_{i,t}) = \mu_i^{s_{i,t}} (1 - \mu_i)^{1 - s_{i,t}} \;, \qquad (24)$$

where μ_i is the mean of the binary sources which we also estimate by MLII from the data. The details of the following experiment can be found in [2]. In figure 1 we show in panel (a) the log marginal likelihood as function of the number of components for *temporal* ICA. A generative ICA model with a single latent component is optimal in this case, and in panel (b) we show the inferred posterior mean of the source and the on-off binary reference function (with values 0 and 0.5 for clarity). The ICA time series shows a few scattered activations and a large contiguous activation beginning approximately 3 seconds (10 scans) after stimulus onset. This is consistent with typical hemodynamic delays found in primary visual cortex [17]. The spatial pattern associated to the inferred source is presented in panels (c) and (d). In (c) we show the 2.5 % most positive (white) and negative (black) activation “hot spots” superimposed on an anatomical background which has the same spatial resolution as the data. In panel (d) we provide a quantitative representation of the spatial pattern. The spatial pattern is dominated by a large cluster of pixels in the primary visual areas. We note two relatively weak, but bilateral negative activations that could be auditive regions that are processing audible scanner noise when the subject is not attending to visual input, as suggested in [16] for the same data set.

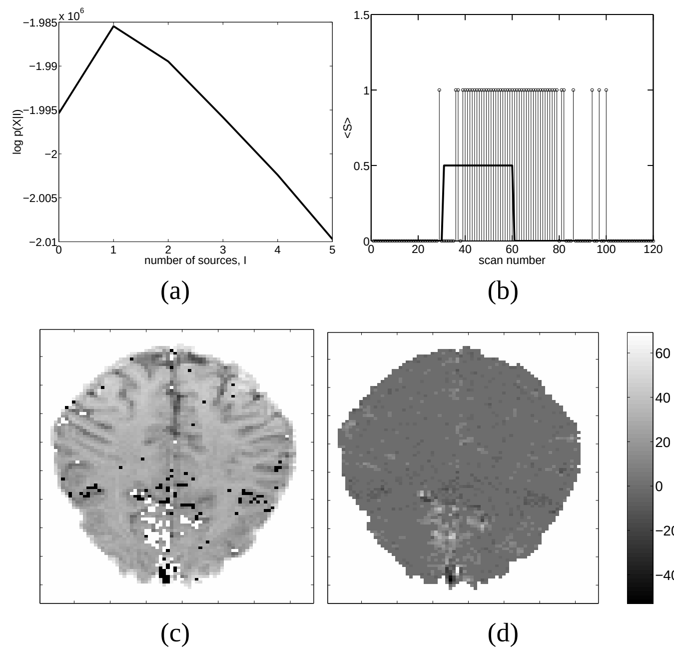

Fig. 1. Analysis of a fMRI data set using temporal binary ICA. (a) shows the log marginal likelihood as a function of the number of hidden sources. At $I = 0$ the log marginal likelihood is calculated for a Gaussian with common variance. (b) shows the experimental paradigm (in solid) and the inferred sources. (c) shows the .025 (black) and .975 (white) fractiles of the values in the eigenimage superimposed on an anatomical reference. (d) shows a quantitative representation of the spatial pattern retaining only the 0.1 and 0.9 fractiles of the eigenimage.

In figure 2 we present the results of searching for *spatially* independent components. In panel (a) we show the log marginal likelihood with nine being the most probable number of latent components. This is consistent with [13] who found a high number of interpretable components using spatial ICA. In panel (b) and (c) we show the binary images and the associated time series. The time series have been globally post-normalized such that the maximal value is 1. Notice that some of the binary images are mainly in the off-state, while others are mainly in the on-state. In particular, two components (4 and 7) have strong responses which are highly correlated with the stimulus. Furthermore, these two component time series are active only in the expected regions for visual stimulation. The weaker signals are confounds that are found globally (components 1,3,5,6,8,9) or locally (components 2).

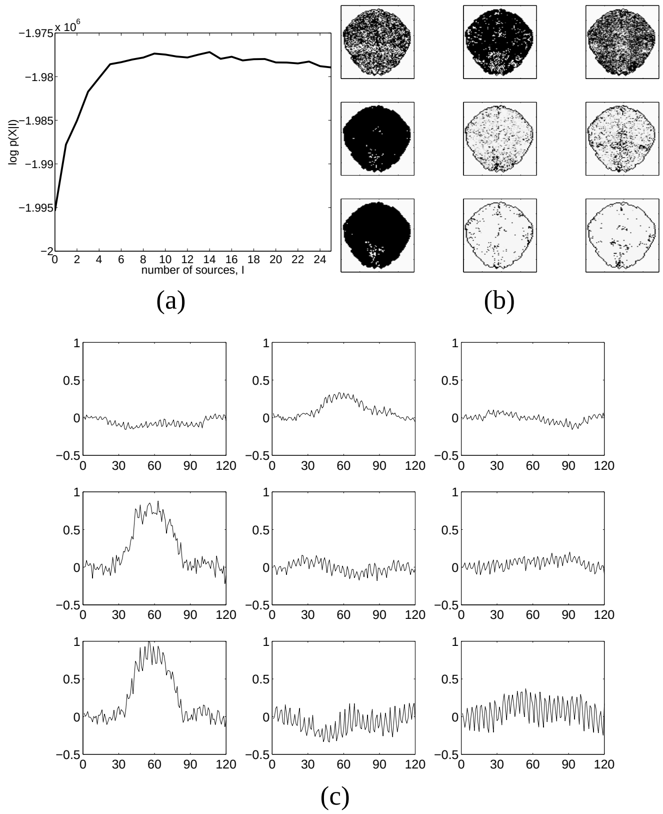

Fig. 2. Analysis of a fMRI data set using spatial binary ICA. (a) shows the log marginal likelihood as a function of the number of hidden sources. At $I = 0$ the log marginal likelihood is calculated for a Gaussian with common variance. The log marginal likelihood is the average of 30 random parameter initializations. (b) shows the posterior mean of the nine hidden binary images. (c) shows the corresponding responses associated to each of the nine hidden images.

6. OUTLOOK

In this section we discuss a few extensions of the mean field framework to other problems in the ICA domain:

Convolutive mixtures. The instantaneous mixing model eq. (1) can be generalized to time-lagged models. Of special interested is an time-trans- lational invariant mixing matrix $\mathbf{A}_{t,t'} = \mathbf{A}_{t-t'}$ (with discrete time):

$$\mathbf{X}_t = \sum_{\Delta t \geq 0}^{\Delta t_{\max}} \mathbf{A}_{\Delta t} \mathbf{S}_{t-\Delta t} + \mathbf{\Gamma}_t \; , \qquad (25)$$

Formally this model fits within the mean field framework. However, the coupling of different time will imply that the coupling matrix $\mathbf{J}$ now is $MN \times MN$ rather than $M \times M$ as in the instantaneous mixing case. This makes linear response estimates difficult to obtain for e.g. speech analysis where the time series are long. Even though the matrix can be sparse, i.e. when $\Delta t_{\max}$ is small, it can be necessary for simplifying approaches. Another possibility is use a Fourier representation [18].

Temporally and 'spatially' correlated source priors. Another model that can be treated within the mean field approach is the one which relaxes the usual assumption of statistical independence of the sources, thus putting the basic assumption of ICA to the test in applications. The general source prior that achieves this is

$$P(\mathbf{S}) \propto \prod_{mt} \rho(S_{mt}) \exp\left(-\frac{1}{2}\sum_{t,t'} \mathbf{S}_t^T \mathbf{J}_{t-t'}^{(0)} \mathbf{S}_{t'}\right) \; ,$$

where $\rho(S)$ denotes the (typically non-Gaussian) single source constraint and $\mathbf{J}_{\Delta t}^{(0)}$ is the time-translational invariant Gaussian a priori source correlation. The source prior in this model thus includes both non-Gaussianity (from the single source constraint) and Gaussian inter correlations as well as temporal correlations (i.e. when $\mathbf{J}_{\Delta t}^{(0)} \neq 0$ for some $\Delta t \neq 0$).

The only additional complication comes from the ML-II estimate of $\mathbf{J}_{\Delta t}^{(0)}$ for which we also need to estimate the a priori source correlations $\langle \mathbf{S}_t \mathbf{S}_{t+\Delta t}^T \rangle_0$ coming from the normalization of $P(\mathbf{S})$. With respect to the computational complexity of the linear response correction, this model has same the problem as the convolutive mixture.

Both this approach and convolutive mixtures should be well suited for tackling hard underdetermined source separation problems, e.g. one microphone source separation, since contrary to instantaneous mixtures, temporal correlation are modeled. It is thus possible to extract more information using the more complicated model. We are currently testing these approaches in underdetermined source separation.

Mixtures of media. The mean field framework can straightforwardly be applied to recovering sources in mixtures of media, e.g. speech and images signals. This is a very interesting truly multimedia problem.

ICA for density estimation. Since the Bayesian approach gives an estimate of $P(\mathbf{X}|\mathbf{A}, \mathbf{\Sigma})$, Bayesian ICA can be used for density estimation.

7. CONCLUSION

We have presented an easily implementable, stable and flexible Bayesian mean field approach to independent component analysis (ICA). Standard variational mean field theory is used in combination with linear response theory. The latter provides improved estimates of second order statistics which are needed for ML-II estimation of the mixing and noise covariance matrix. With this approach we have been able to solve difficult source separation tasks for which the variational mean field theory failed [1].

In our opinion, the ICA field could benefit from a 'blind test' of the performance of algorithms. It is important to identify which types of algorithms work for specific problem domains and what problems are hard. Quantitative performance measures rather than often visual qualitative ones should be developed.

8. REFERENCES

[1] P.A.d.F.R. Højen-Sørensen, O. Winther, and L. K. Hansen, "Mean field approaches to independent component analysis," Accepted for publication in *Neural Computation*, Available from `http://www.cbs.dtu.dk/winther/`, 2001.

[2] P.A.d.F.R. Højen-Sørensen, O. Winther, and L. K. Hansen, "Analysis of functional neuroimages using ica adaptive binary sources," Submitted to *Neurocomputing*, 2001.

[3] T.-W. Lee, *Independent Component Analysis*, Kluwer Academic Publishers, Boston, 1998.

[4] A. Belouchrani and J.-F. Cardoso, "Maximum likelihood source separation by the expectation-maximization technique: Deterministic and stochastic implementation," in *Proc. NOLTA*, 1995, pp. 49–53.

[5] M. Opper and D. Saad, Eds., *Advanced Mean Field Mathods – Theory and Practice*. MIT Press, 2001.

[6] H. Attias, "Independent factor analysis," *Neural Computation*, vol. 11, pp. 803–51, 1999.

[7] J. W. Miskin and D. J. C. MacKay, "Ensemble learning for blind image separation and deconvolution," in *Advances in Independent Component Analysis*, M. Girolami, Ed. Springer-Verlag Scientific Publishers, 2000.

[8] G. Schwarz, "Estimating the dimension of a model," *Annals of Statistics*, vol. 6, pp. 461–464, 1978.

[9] P. A. Bandettini, A. Jesmanowicz, E. C. Wong, and J. S. Hyde, "Processing strategies for time-course data sets in functional MRI of the human brain," *Magnetic Resonance in Medicine*, vol. 30, pp. 161–173, 1993.

[10] N. Lange and S. L. Zeger, "Non-linear Fourier time series analysis for human brain mapping by functional magnetic resonance imaging," *Journal of the Royal Statistical Society - Series C Applied Statistics*, vol. 46, no. 1–30, 1997.

[11] K. J. Worsley, J. B. Poline, K. J. Friston, and A. C. Evans, "Characterizing the response of PET and fMRI data using multivariate linear models (MLM)," *NeuroImage*, vol. 6, pp. 305–319, 1997.

[12] L. K. Hansen, J. Larsen, F. Å Nielsen, S. C. Strother, E. Rostrup, R. Savoy, N. Lange, J. J. Sidtis, and C. Svarer, "Generalizable patterns in neuroimaging: How many principal components?," *NeuroImage*, vol. 9, pp. 534–544, 1999.

[13] M. McKeown, T.-P. Jung, S. Makeig, G. Brown, S. S. Kindermann, T.-W. Lee, and T. J. Sejnowski, "Spatially independent activity patterns in functional MRI data during the Stroop color-naming task," in *Proc. Natl. Acad. Sci. USA*, February 1998, vol. 95, pp. 803–810.

[14] A. J. Bell and T. J. Sejnowski, "An information–maximization approach to blind separation and blind deconvolution," *Neural Computation*, vol. 7, no. 6, pp. 1129–1159, 1995.

[15] L. K. Hansen, "Blind separation of noisy image mixtures," in *Advances in Independent Components Analysis*, M. Girolami, Ed., pp. 165–187. Springer-Verlag, Berlin, 2000.

[16] K. S. Petersen, L. K. Hansen, T. Kolenda, E. Rostrup, and S. Strother, "On the independent components in functional neuroimages," in *Proc. Int. Workshop on Independent Component Analysis and Blind Signal Separation (ICA2000)*, P. Pajunen and J. Karhunen, Eds., Helsinki, Finland, June 2000, pp. 615–620.

[17] P. A. Bandettini and E. C. Wong, "Echo–planar magnetic resonance imaging of human brain activation," in *Echo–Planar Imaging*, F. Schmitt, M. K. Stehling, R. Turner, and Mansfield P., Eds. Springer-Verlag, New York, 1998.

[18] H. Attias and C. Schreiner, "Blind source separation and deconvolution: The dynamic component analysis algorithm," *Neural Computation*, vol. 11, pp. 803–852, 1998.