NeuroImage 152 (2017) 563–574

Contents lists available at ScienceDirect

NeuroImage

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

Gaussian process based independent analysis for temporal source separation in fMRI

Ditte Høvenhoff Hald[a,*], Ricardo Henao[a,b], Ole Winther[a]

[a] DTU Compute B321, Technical University of Denmark, DK-2800 Lyngby, Denmark
[b] Electrical and Computer Engineering, Duke University, Durham, NC 27708, USA

ARTICLE INFO

Keywords:
Gaussian processes
FMRI
Source separation
Independent component analysis
Convolutive mixing
Bayesian inference

ABSTRACT

Functional Magnetic Resonance Imaging (fMRI) gives us a unique insight into the processes of the brain, and opens up for analyzing the functional activation patterns of the underlying sources. Task-inferred supervised learning with restrictive assumptions in the regression set-up, restricts the exploratory nature of the analysis. Fully unsupervised independent component analysis (ICA) algorithms, on the other hand, can struggle to detect clear classifiable components on single-subject data. We attribute this shortcoming to inadequate modeling of the fMRI source signals by failing to incorporate its temporal nature. fMRI source signals, biological stimuli and non-stimuli-related artifacts are all smooth over a time-scale compatible with the sampling time (TR). We therefore propose Gaussian process ICA (GPICA), which facilitates temporal dependency by the use of Gaussian process source priors.

On two fMRI data sets with different sampling frequency, we show that the GPICA-inferred temporal components and associated spatial maps allow for a more definite interpretation than standard temporal ICA methods. The temporal structures of the sources are controlled by the covariance of the Gaussian process, specified by a kernel function with an interpretable and controllable temporal length scale parameter. We propose a hierarchical model specification, considering both instantaneous and convolutive mixing, and we infer source spatial maps, temporal patterns and temporal length scale parameters by Markov Chain Monte Carlo. A companion implementation made as a plug-in for SPM can be downloaded from https://github.com/dittehald/GPICA.

Introduction

The hemodynamic response (HDR) of the brain and the emerged blood-oxygen-level-dependent (BOLD) contrast image captured by functional magnetic resonance imaging (fMRI) is a window into many aspects of the brain. Independent component analysis (ICA) is one of the most common ways to extract signal constituents from fMRI data. These sub-parts of the fMRI signal arise from different independent processes related to the stimuli, non-stimuli effects (such as heart beat) and artifacts (such as head movement), (Duann et al., 2002; McKeown et al., 2003). Proper source separation will thus allow both identification of stimuli-related signals and artifact removal. These independent processes are expected to vary smoothly over time, on a time scale that is often comparable with the fMRI acquisition frequency. It therefore seems obvious to use this information in the modeling of the sources. This is, however, computationally demanding, and ICA models based on (non-Gaussian) independent identically distributed (i.i.d.) sources are the most commonly used in practice, e.g, Infomax and FastICA (Bell and Sejnowski, 1995; Hyvärinen, 1999).

The fMRI data, X, is assumed to be a linear combination of stimuli, non-stimuli and artifact signals plus additive noise. Each signal can be thought of as consisting of a source, Z, and its dispersion in space, W, by $X = WZ + \epsilon$, with ϵ being the noise contribution. The estimation of W and Z is identifiable (up to permutation and sign symmetries) for some choices of prior distributions on the model parameters (Kagan et al., 1973; Henao and Winther, 2011). The non-identifiable case is i.i.d. Gaussian priors on Z and W, corresponding to probabilistic PCA and standard factor analysis. In blind source separation, ICA is the most widely used generative model to solve the problem, and the most notably identifiable model specification is based upon non-Gaussian i.i.d. priors or temporally correlated Gaussian priors. The non-Gaussian i.i.d. priors are either used explicitly as in InfoMax (Bell and Sejnowski, 1995)[1], or implicitly as in FastICA (Hyvärinen, 1999). The Molgedey-Schuster algorithm (icaMS) (Molgedey and Schuster,

* Corresponding author.
E-mail address: ditha@dtu.dk (D.H. Hald).

[1] MacKay (2003) provides an alternative derivation of Infomax using a Bayesian formulation similar in spirit to the model specification proposed here.

http://dx.doi.org/10.1016/j.neuroimage.2017.02.070
Received 20 June 2016; Accepted 24 February 2017
Available online 27 February 2017
1053-8119/

1994) is a notable example of a model that implicitly uses a temporally correlated source prior.

In this paper, we work with temporally smooth sources. We accomplish this by modeling each source as a temporal Gaussian process. Gaussian processes (GP) are a generalization of multivariate Gaussian distributions into a stochastic process described by a mean function and a covariance (kernel) function (for a introduction to Gaussian processes, see, e.g., Rasmussen and Williams (2006)). In Gaussian processes, the covariance function's shape and parameters determine the length-scales on which the signal is correlated. Consequently, in our algorithm, Gaussian Process based ICA (GPICA[2]), different sources will differ both in terms of their spatial and temporal patterns, as in standard ICA, and additionally, by the length-scale of their characteristic temporal correlation. As a part of the algorithm we infer W and Z and for each source the length-scale parameter of its kernel function.

A prerequisite for this model to make sense is that the sampling time (TR, the inverse of the sample frequency) is on the same length-scale as temporal correlation in the signal components we would like to recover. Typically, the fMRI acquisition frequency is around 0.5-1 Hz (that is TR $\approx$ 1–2 secs.). We consider two data sets with different sampling rates: a fast acquisition visual paradigm dataset with TR = 1/3 secs. and a motor paradigm dataset with a more standard TR = 2.49 sec. In both cases, we show that GPICA can recover interpretable time-scale signals, compatible with the TR used.

As long as there is a different GP for a each source (i.e., different length-scale), Kagan et al. (1973) theorem (Theorem 10.4.1) is fulfilled and will therefore lead to identifiable sources. And even if sources' GPs would be too close, a non-Gaussian prior on W can potentially drive the model to identifiable sources. The Gaussian process has the property that it reverts back to independent Gaussian variables (that is probabilistic PCA), if the temporal length-scale parameter in the covariance function is much shorter than TR. In that limit the kernel matrix approaches the identity matrix. We can in principle still get identifiable source recovery in this limit as long as the prior over the W matrix has a non-Gaussian distribution. We therefore use i.i.d. Student's t-distributions on the elements of W. This choice has the added benefit of promoting a more discriminative W because, relative to the Gaussian, Student's t has more shrinkage (heavier tails) towards zero for non-important parameters and less for important ones.

Gaussian processes are widely applied in probabilistic regression problems (Rasmussen and Williams, 2006). GPs have previously been employed in factor analysis (FA), where Yu et al. (2009) uses time-point FA tied together by a GP in order to perform dimension reduction and smoothing simultaneously. Luttinen and Ilin (2009) developed a GPFA algorithm for reconstructing missing data, with GPs attached to both factors and loadings. They use a variational Bayesian framework for learning the model, and factorize the posterior approximation either in time or space to reduce the computational complexity. Luttinen and Ilin (2009) are closely related to the GP work by Schmidt and Laurberg (2008), Schmidt (2009). Schmidt and Laurberg (2008) focuses on probabilistic non-negative matrix factorization (NMF), and models factor smoothness with a GP. In the follow-up work, Schmidt (2009) uses non-linear mapped GPs (Warped GPs similar to Snelson et al. (2004)) to perform function factorization. Bayesian inference is performed with Hamiltonian Markov Chain Monte Carlo. Preliminary work on source separation with GP sources has also been formulated by Park and Choi (2007), Park and Choi (2008). They use mutual information minimization (Park and Choi, 2007) or gradient-based optimization (Park and Choi, 2008) of the log-pseudo-likelihood to infer the mixing matrix. Their proposed model is closely related to our GPICA model. Our novel contributions are the hierarchical specification of the model, Markov Chain Monte Carlo (MCMC)-based inference, application to fMRI datasets and a companion plug-in for SPM to make the algorithm available for other researchers. The algorithmic workflow is illustrated in Fig. 1, and can be used as a point of reference through the model specifications together with the table of notation in Table A.1. We also propose a convolutive mixing matrix extension of the algorithm. This extension is related to the work by Olsson and Hansen (2006). The main difference in the model specification is that the temporal model is a linear state space model (AR model), that may be seen as a different parameterization of the GP temporal covariance (Hartikainen et al., 2010). Furthermore, the model as originally proposed is not hierarchical and inference is carried out with expectation-maximization (EM).

The proposed algorithm has relatively high complexity, mainly due to the number of MCMC samples needed to get sufficiently accurate estimates. However, our empirical results show that this is sometimes a price worth paying. For minimally pre-processed data, we can recover stimuli-related sources, even in cases where the amplitudes of these sources are only a small fraction of the total variance in the data (low signal-to-noise ratio regime). The careful time-consuming modeling and inference in GPICA are, without doubt, contributing factors to achieving meaningful sources.

The remainder of this paper is organized as follows: In Section 2, we introduce the GP-based ICA model with GP sources and both instantaneous and convolutive mixing. Appendix B gives the details of the Gibbs sampling-based inference scheme. In Section 3, we give a brief description of the two fMRI data sets used in Section 4. In Section 4.5 we describe an implementation as a plug-in to SPM. In Section 5 we interpret the empirical results. In Section 6 we provide an outlook and perspectives for the model.

Independent component analysis with Gaussian Process priors

In this Section, we will describe the independent component analysis (ICA) model with Gaussian process-based sources. We will specify hierarchical Bayesian priors for model parameters and describe a Markov Chain Monte Carlo (MCMC) framework for inference. We consider both instantaneous and convolutive versions of the mixing model.

Instantaneous independent component analysis

The basic ICA model is defined as

$$X_{ij} = \sum_{p=1}^{P} w_{ip} z_{pj} + \epsilon_{ij} \tag{1}$$

$$X = WZ + \epsilon \tag{2}$$

where X_{ij} is the data matrix with subscripts representing voxel and time sample, respectively. Each row in Z represents the time course of one source and W is the component map, or mixing matrix. P is the number of independent components in which the data is decomposed. P is much smaller than the number of voxels D and the number of time points N. ϵ_{ij}, the spatial-temporal noise contribution, is assumed to be i.i.d. Gaussian with zero mean and spatially varying noise variance Ψ_{ii}. We can thus write the likelihood for the jth time slice as

$$P(x_j|W, z_j, \Psi) = \mathcal{N}(x_j|Wz_j, \Psi) \tag{3}$$

and the joint distribution of the whole data set is

$$P(X|W, Z, \Psi) = \prod_{j=1}^{N} P(x_j \bigg| W, z_j, \Psi) \tag{4}$$

where Ψ is diagonal with elements Ψ_{ii}.

[2] Not to be confused with geometric post nonlinear ICA (gpICA) from Nguyen et al. (2007).