

Simultaneous Estimation of Nongaussian Components and Their Correlation Structure

Hiroaki Sasaki

hsasaki@is.naist.jp

Graduate School of Information Science, Nara Institute of Science and Technology, Nara 630-0192, Japan

Michael U. Gutmann

michael.gutmann@ed.ac.uk

School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, U.K.

Hayaru Shouno

shouno@uec.ac.jp

Graduate School of Informatics and Engineering, University of Electro-Communications, Tokyo 182-8585, Japan

Aapo Hyvärinen

a.hyvarinen@ucl.ac.uk

Helsinki Institute for Information Technology, University of Helsinki, Helsinki 00560, Finland, and Gatsby Computational Neuroscience Unit, University College London, London W1T 4JG, U.K.

The statistical dependencies that independent component analysis (ICA) cannot remove often provide rich information beyond the linear independent components. It would thus be very useful to estimate the dependency structure from data. While such models have been proposed, they have usually concentrated on higher-order correlations such as energy (square) correlations. Yet linear correlations are a fundamental and informative form of dependency in many real data sets. Linear correlations are usually completely removed by ICA and related methods so they can only be analyzed by developing new methods that explicitly allow for linearly correlated components. In this article, we propose a probabilistic model of linear nongaussian components that are allowed to have both linear and energy correlations. The precision matrix of the linear components is assumed to be randomly generated by a higher-order process and explicitly parameterized by a parameter matrix. The estimation of the parameter matrix is shown to be particularly simple because using score-matching (Hyvärinen, 2005), the objective function is a quadratic form. Using simulations with artificial data, we demonstrate that the proposed method improves the identifiability of nongaussian components by

simultaneously learning their correlation structure. Applications on simulated complex cells with natural image input, as well as spectrograms of natural audio data, show that the method finds new kinds of dependencies between the components.

1 Introduction

Estimating latent nongaussian components is important in modern statistical data analysis and machine learning. A well-known method for that purpose is independent component analysis (ICA) (Comon, 1994; Hyvärinen & Oja, 2000), whose goal is to identify nongaussian components as statistically independent as possible. ICA has been applied in a wide range of fields, such as brain imaging analysis (Vigário, Särelä, Jousmäki, Hämäläinen, & Oja, 2000), image processing (Hyvärinen, Hurri, & Hoyer, 2009), pattern recognition (Bartlett, Movellan, & Sejnowski, 2002), and causal analysis (Shimizu, Hoyer, Hyvärinen, & Kerminen, 2006).

The components estimated by ICA, however, are often not independent at all. For natural images, for instance, the estimated components may have variance dependencies, that is, the squares of the components may be correlated, and the same holds for the wavelet coefficients (Hyvärinen & Hoyer, 2000; Karklin & Lewicki, 2005; Simoncelli, 1999). Inspired by this fact, extensions of ICA have been developed that take dependencies between the components into account: independent subspace analysis (ISA), which combines the techniques of multidimensional ICA (Cardoso, 1998; Theis, 2005) and the principle of invariant-feature subspaces (Kohonen, 1995, 1996), divides the components into predefined groups where the components in each group have variance dependencies (Hyvärinen & Hoyer, 2000). When applied to natural images, ISA produces a phase invariant pooling of the nongaussian components. Methods for learning topographic representation assume that nearby components have statistical dependencies, while far-away components are as statistically independent as possible (Hyvärinen, Hoyer, & Inki, 2001; Mairal, Jenatton, Obozinski, & Bach, 2011; Sasaki, Gutmann, Shouno, & Hyvärinen, 2013). The dependencies are used to order the components and arrange them on a grid, which provides a convenient visualization of properties of the data. Thus, the statistical dependencies that ICA cannot remove often contain rich information. However, a limitation of the work already cited is that it assumes prefixed dependency structures, which can be problematic because specifying a wrong dependency structure can hamper the identifiability of the nongaussian components (Sasaki et al., 2013).

This limitation can be removed by estimating dependency structures themselves from data. Two-layer models are suitable for this purpose if they further incorporate some higher-order process for nongaussian components. Karklin and Lewicki (2005) proposed a method to estimate a

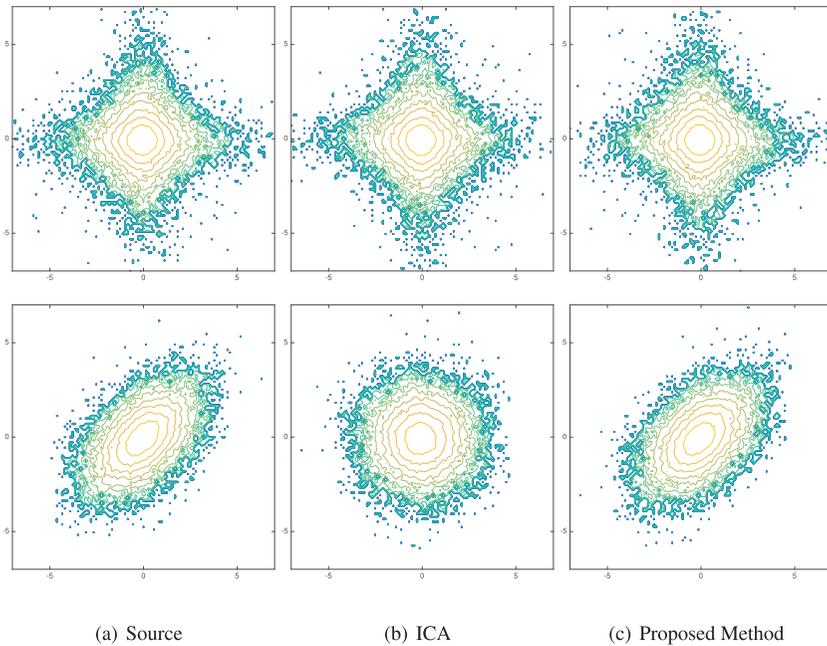


Figure 1: The logarithms of the nonparametric estimates of the two-dimensional probability density functions (a) for artificially generated source components, and (b, c) estimated source components by ICA (maximum likelihood estimation without the decorrelation constraint as in equation 3.8 and the proposed method, respectively). The top row is for independent nongaussian components and the bottom one for linearly correlated ones. The correlation coefficients in the second row are (a) 0.499, (b) 0.01, and (c) 0.472, respectively.

two-layer model whose first layer consists of ICA-like components and whose second layer represents density components. Osindero, Welling, and Hinton (2006) proposed a two-layer topographic model with nongaussian components in the first layer and weighted connections among them in the second layer. A related two-layer model was proposed by Köster and Hyvärinen (2010). However, most of the models focus on higher-order dependencies only and ignore linear correlations, usually even assuming that they are zero. But linear correlations between nongaussian components can be observed in real data for a couple of models (Coen-Cagli, Dayan, & Schwartz, 2012; Gómez-Herrero, Atienza, Egiazarian, & Cantero, 2008). It is thus probable that the underlying components that are forced to be linearly uncorrelated in the estimation would actually be correlated. For instance, Figure 1 illustrates that the components estimated by ICA are linearly uncorrelated even when the underlying source components are correlated.

Therefore, it would be meaningless to analyze linear correlations in components estimated by ordinary ICA methods; it is necessary to incorporate the linear correlations in the very definition of the model. Along these lines, the estimation of topographic representations has recently been improved by taking into account both linear and variance correlations (Sasaki et al., 2013). Linear and higher-order dependencies between the components have also been exploited in order to find correspondences between features in multiple data sets (Campi, Parkkonen, Hari, & Hyvärinen, 2013; Gutmann & Hyvärinen, 2011; Gutmann, Laparra, Hyvärinen, & Malo, 2014).

In this article, we propose a novel method to estimate latent nongaussian components and their dependency structure simultaneously. The dependency structure includes both linear and higher-order correlations and is parameterized by a single matrix. The off-diagonal elements of this dependency matrix represent the conditional dependencies much like the precision matrix does for gaussian Markov random fields. More generally, the dependency matrix defines a distance matrix that can be used for visualization via an undirected graph, multidimensional scaling, or some other suitable technique. The proposed method can be interpreted as a generalization of ICA and correlated topographic analysis (CTA) (Sasaki et al., 2013), where the dependency structures are assumed to be known. To develop the method, we begin with a new generative model for sources, which generalizes previous models of Hyvärinen et al. (2001), Karklin and Lewicki (2005), Köster and Hyvärinen (2010), and Osindero, Welling, and Hinton (2006) by capturing linear correlations between the components. The previous models generate nongaussian components without linear but with higher-order correlations (Hyvärinen et al., 2009). Divisive normalization models (Ballé, Laparra, & Simoncelli, 2015; Heeger, 1992; Schwartz & Simoncelli, 2001) are also closely related and focus on higher-order correlations.

Estimating two-layer models or Markov random fields for nongaussian components is often difficult because sophisticated parametric models generally have an intractable partition function so that the standard maximum likelihood estimation cannot be applied. To cope with this problem, several estimation methods have been proposed, such as contrastive divergence (Hinton, 2002), score matching (Hyvärinen, 2005), and noise contrastive estimation (Gutmann & Hyvärinen, 2012) and its extensions (Gutmann & Hirayama, 2011; Pihlaja, Gutmann, & Hyvärinen, 2010; see Gutmann & Hyvärinen, 2013a, for an introductory paper). Score matching has a particularly useful property for the proposed model: the objective function for the estimation of the dependency parameters takes a simple quadratic form and can be optimized by standard quadratic programming. Due to this computational simplification, here we use score matching and empirically show that our method estimates the dependency structure and improves the identifiability of the nongaussian components.

The article is organized as follows. In section 2, we begin with a novel probabilistic generative model for conditional precision matrices. Based on

the generative model, we derive an approximation of the marginal density for nongaussian components where the dependency structure is explicitly parameterized like a precision matrix. Section 3 deals with estimation of the model, presenting a powerful algorithm for identifying the nongaussian components and their dependency structure. In section 4, we perform numerical experiments and compare the proposed method with existing methods on artificial data. Section 5 demonstrates the applicability of the method to real data. Connections to past work and extensions of the proposed method are discussed in section 6. Section 7 concludes this article. A preliminary version of this article was presented at AISTATS 2014 (Sasaki, Gutmann, Shouno, & Hyvärinen, 2014).

2 Generative Model with Dependent Nongaussian Components _____

Here, we introduce a novel generative model with dependent nongaussian components. The probability density function (PDF) of the components and the data is shown to be only implicitly defined via an intractable integral. We derive an approximation of the PDF where the dependency structure of the components is explicitly parameterized and demonstrate the validity of the approximation using both analytical arguments and simulations.

2.1 The Generative Model. As in previous work related to ICA (Hyvärinen & Oja, 2000), we assume a linear mixing model for the data,

$$x = s_1 a_1 + s_2 a_2 + \dots + s_d a_d = \mathbf{A} s, \tag{2.1}$$

where $x = (x_1, x_2, \dots, x_d)^T$ denotes the d -dimensional data vector, $\mathbf{A} = (a_1, a_2, \dots, a_d)$ is the d by d mixing matrix formed by the basis vectors a_i , and $s = (s_1, s_2, \dots, s_d)^T$ is the d -dimensional vector consisting of the latent nongaussian components (the sources). The nongaussianity assumption about s is fundamental for identification of the mixing model (Comon, 1994).

We next construct a model for the components s_i which allows them to be statistically dependent, in contrast to ICA. We assume that s is generated from a gaussian density with precision matrix Λ whose elements λ_{ij} of Λ are random variables themselves, generated before s in a higher level of hierarchy as

$$\lambda_{ij} = \begin{cases} -u_{ij}, & i \neq j, \\ \sum_{k=1}^d u_{ik}, & i = j. \end{cases} \tag{2.2}$$

The u_{ij} are independent nonnegative random variables, and because precision matrices are symmetric, we set $u_{ij} = u_{ji}$. To ensure invertability, we also require that $u_{ii} > 0$. Nonzero u_{ij} produce positive correlations between s_i and s_j (given the remaining variables), and larger values of u_{ij} result in more strongly correlated variables. The motivation for the definition of the diagonal elements λ_{ii} is that it guarantees realizations of $\mathbf{\Lambda}$, which are invertible and positive definite ($\mathbf{\Lambda}$ is symmetric and strictly diagonally dominant, from which the stated properties follow; see theorem 6.1.10 in Horn & Johnson, 1985, and appendix E). Readers familiar with graph theory will recognize that $\mathbf{\Lambda}$ equals the Laplacian matrix of a weighted graph defined by the matrix \mathbf{U} with elements u_{ij} (Bollobás, 1998).

Another property of the model is that the components s_i are supergaussian: Hyvärinen et al. (2001) showed that the marginal PDFs of the gaussian variables with random variances have heavier tails than a gaussian PDF. Furthermore, since the conditional variances are dependent on each other in model 2.2, higher-order correlations are likely to exist (Hyvärinen et al., 2001; Sasaki et al., 2013).

We thus assume that the conditional PDF of \mathbf{s} given $\mathbf{\Lambda}$ equals $p(\mathbf{s}|\mathbf{\Lambda})$,

$$p(\mathbf{s}|\mathbf{\Lambda}) = \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2}\mathbf{s}^\top \mathbf{\Lambda} \mathbf{s}\right) \quad (2.3)$$

$$= \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2} \sum_{i=1}^d \left\{ \lambda_{ii} s_i^2 + \sum_{j \neq i} \lambda_{ij} s_i s_j \right\}\right), \quad (2.4)$$

which we can write in terms of the u_{ij} as $p(\mathbf{s}|\mathbf{U})$,

$$p(\mathbf{s}|\mathbf{U}) = \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp\left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{j>i} (s_i - s_j)^2 u_{ij} \right\}\right), \quad (2.5)$$

as proved in appendix A. While not explicitly visible from the notation, the determinant $|\mathbf{\Lambda}|$ is a function of the u_{ij} .

The specification of the PDF of the u_{ij} completes the model, but this is a complex issue that we postpone to section 2.2. Denoting the PDF of the u_{ij} generally by p_u , the PDF of the sources equals $p_s(\mathbf{s})$,

$$p_s(\mathbf{s}) = \int_0^\infty p(\mathbf{s}|\mathbf{U}) p_u(\mathbf{U}) d\mathbf{U}, \quad (2.6)$$

and the PDF of \mathbf{x} follows from the standard formula for linear transformations of random variables,

$$p_x(\mathbf{x}) = p_s(\mathbf{W}\mathbf{x}) |\mathbf{W}|, \quad (2.7)$$

with $\mathbf{W} = \mathbf{A}^{-1}$. However, since the determinant $|\mathbf{\Lambda}|$ in $p(\mathbf{s}|\mathbf{U})$ depends on \mathbf{U} , solving the multidimensional integral in equation 2.6 is practically impossible for any choice of p_u . While the integral can be estimated using Monte Carlo methods, it could be time-consuming. Therefore, we consider next an analytical approximation of the determinant, which allows us to find an approximation of p_s that holds qualitatively for any p_u . Once we have an approximation, say, \tilde{p}_s , we can use equation 2.7 to obtain a tractable approximation \tilde{p}_x of the PDF of \mathbf{x} ,

$$\tilde{p}_x(\mathbf{x}) = \tilde{p}_s(\mathbf{W}\mathbf{x})|\mathbf{W}|. \tag{2.8}$$

2.2 Approximating the Density of the Dependent Non-Gaussian Components. In order to derive an approximation \tilde{p}_s of p_s , we approximate the determinant of $\mathbf{\Lambda}$ via a product over the u_{ii} ,

$$|\mathbf{\Lambda}| \approx \prod_{i=1}^d u_{ii}. \tag{2.9}$$

This is the only approximation we need to obtain the tractable \tilde{p}_s . Another meaning of this approximation is to give a lower bound of $p(\mathbf{s}|\mathbf{U})$; consequently, \tilde{p}_s is a lower bound of p_s , which is proved by using Ostrowski's inequality in appendix B. In other words, \tilde{p}_s is an unnormalized model defined up to a multiplicative factor not depending on \mathbf{s} . This is not an insurmountable problem but needs to be taken into account when performing the estimation (see, e.g., Gutmann & Hyvärinen, 2013a).

Inserting the approximation 2.9 and the independence assumption of the u_{ij} into equation 2.6 yields the following approximation of $p_s(\mathbf{s})$:

$$\begin{aligned} p_s(\mathbf{s}) &\approx \tilde{p}_s(\mathbf{s}) \\ &\propto \int_0^\infty \left[\prod_{i=1}^d \sqrt{u_{ii}} \right] \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{j>i} (s_i - s_j)^2 u_{ij} \right\} \right) p_u(\mathbf{U}) d\mathbf{U} \end{aligned} \tag{2.10}$$

$$\begin{aligned} &\propto \int_0^\infty \left[\prod_{i=1}^d \sqrt{u_{ii}} \right] \left[\prod_{i=1}^d \exp \left(-\frac{1}{2} s_i^2 u_{ii} \right) \prod_{j>i} \exp \left(-\frac{1}{2} (s_i - s_j)^2 u_{ij} \right) \right] \\ &\quad \times \left[\prod_{i=1}^d \prod_{j>i} p_{ij}(u_{ij}) \right] d\mathbf{U}, \end{aligned} \tag{2.11}$$

where the product over the p_{ij} in the last line is the PDF p_u of the u_{ij} due to their independence. The expression for \tilde{p}_s can be simplified by grouping

together terms featuring u_{ii} and u_{ij} only,

$$\begin{aligned} \tilde{p}_s(\mathbf{s}) &\propto \prod_{i=1}^d \left[\int_0^\infty \sqrt{u_{ii}} \exp\left(-\frac{1}{2}s_i^2 u_{ii}\right) p_{ii}(u_{ii}) du_{ii} \right] \\ &\quad \times \prod_{j>i} \left[\int_0^\infty \exp\left(-\frac{1}{2}(s_i - s_j)^2 u_{ij}\right) p_{ij}(u_{ij}) du_{ij} \right] \end{aligned} \quad (2.12)$$

$$\propto \prod_{i=1}^d g_{ii}(s_i^2) \prod_{j>i} g_{ij}((s_i - s_j)^2), \quad (2.13)$$

where we have introduced the nonnegative functions $g_{ii}(v)$ and $g_{ij}(v)$, defined for $v \geq 0$,

$$g_{ii}(v) \propto \int_0^\infty \sqrt{u_{ii}} \exp\left(-\frac{v}{2} u_{ii}\right) p_{ii}(u_{ii}) du_{ii}, \quad (2.14)$$

$$g_{ij}(v) \propto \int_0^\infty \exp\left(-\frac{v}{2} u_{ij}\right) p_{ij}(u_{ij}) du_{ij} \text{ for } i \neq j. \quad (2.15)$$

The proportionality sign is used because \tilde{p}_s is defined only up to the partition function. The approximation of the determinant thus allowed us to transform the multidimensional integral in equation 2.6 into a product of functions defined by one-dimensional integrals. The one-dimensional integrals can be easily solved numerically for arbitrary p_{ij} or analytically for particular choices of them. We also note that the g_{ij} are related to the Laplace transform of the p_{ij} .

Different PDFs p_{ij} yield different functions g_{ij} . But appendix C shows that unless $g_{ij}(v)$ is a constant, the different $\log g_{ij}(v)$ are monotonically decreasing convex functions for any choice of p_{ij} . We thus focus on the following particular class of functions:

$$\log g_{ij}(v) = -m_{ij}\sqrt{v} + \text{const}. \quad (2.16)$$

The m_{ij} are free parameters, which will be estimated from the data. They are symmetric, $m_{ij} = m_{ji}$. Further, we require that $m_{ii} > 0$ so that \tilde{p}_s depends on all s_i . For m_{ij} , $i \neq j$, we require only nonnegativity: if $m_{ij} = 0$, then $g_{ij}((s_i - s_j)^2) = \text{const}$, which happens when the variable u_{ij} is deterministically zero.

The particular choice, equation 2.16, is motivated by its simplicity, but we show in appendix D that it corresponds to choosing an inverse-gamma density for the u_{ij} :

$$p_{ij}(u_{ij}) = \begin{cases} \frac{m_{ii}^2}{2} u_{ii}^{-2} \exp\left(-\frac{m_{ii}^2}{2u_{ii}}\right), & i = j, \\ \frac{m_{ij}}{\sqrt{2\pi}} u_{ij}^{-3/2} \exp\left(-\frac{m_{ij}^2}{2u_{ij}}\right), & i \neq j. \end{cases} \tag{2.17}$$

The parameters m_{ij} determine the mode of the p_{ij} (i.e., the point at which p_{ij} is maximum): the mode is $m_{ii}^2/2$ for $i = j$ and $2m_{ij}^2/3$ otherwise.

Denoting by \mathbf{M} the matrix formed by the m_{ij} and its upper-triangular part by \mathbf{m} , $\mathbf{m} = (m_{11}, \dots, m_{1d}, m_{22}, \dots, m_{2d}, m_{33}, \dots, m_{3d}, \dots, m_{dd})$, we obtain the approximation $\tilde{p}_s(\mathbf{s}) = \tilde{p}_s(\mathbf{s}; \mathbf{m})$,

$$\tilde{p}_s(\mathbf{s}; \mathbf{m}) \propto \prod_{i=1}^d \exp\left(-m_{ii}|s_i| - \sum_{j>i} m_{ij}|s_i - s_j|\right), \tag{2.18}$$

which we will use in the following sections. The terms $|s_i|$ in equation 2.18 are related to modeling the s_i as supergaussian, and the terms $|s_i - s_j|$ capture statistical dependencies between the components. The dependencies can be read out from the dependency matrix \mathbf{M} : if $m_{ij} = 0$ for some $j \neq i$, s_i is independent from the s_j conditioned on the other variables. Furthermore, larger m_{ij} imply stronger conditional (positive) dependencies between s_i and s_j .

The model, equation 2.18 generalizes PDFs used in previous work in the following ways:

1. \tilde{p} approaches the Laplacian factorizable PDF when $m_{ij} \rightarrow 0$ for all $j \neq i$. Laplacian factorizable PDFs are often assumed for supergaussian components in ICA.
2. \tilde{p} approaches the topographic PDF in CTA (Sasaki et al., 2013) when $m_{ii} \rightarrow 1$, $m_{i,i+1} \rightarrow 1$ for all i and $m_{ij} \rightarrow 0$ otherwise. The topographic PDF was derived using a different generative model and resorting to two quite heuristic approximations. This is contrast to this article, where a single relatively well-justified approximation was used. The new derivation is not only more elegant, but allows for further extensions of the model as well, which is done in appendix E.
3. \tilde{p} was heuristically proposed in our preliminary conference paper (Sasaki et al., 2014) as a simple extension of CTA. In this article, we have derived \tilde{p} from a novel generative model for random precision matrices.

2.3 Numerical Validation of the Approximation. Here, we investigate the validity of the approximative PDF $\tilde{p}_s(\mathbf{s}; \mathbf{m})$ in equation 2.18 using numerical simulations.

We generated a large number of samples for dependent nongaussian components according to the generative model in section 2.1, fitted a non-parametric density to the sample, and compared it with our approximation in equation 2.18. The dimension of \mathbf{s} was $d = 2$, and the size of the sample was $T = 10^6$. The u_{ij} were drawn from the inverse-gamma density in equation 2.17 with $m_{11} = m_{22} = 1$. We performed the comparison for multiple m_{12} between 0 and 1. Using the generated sample, the density p_s in equation 2.6 was estimated as a normalized histogram. The approximative PDF $\tilde{p}_s(\mathbf{s}, \mathbf{m})$ was normalized using numerical integration and evaluated with the same m_{ij} used to generate the sample.

We evaluated the goodness of the approximation using three measures;

$$\text{ang}(p_s, \tilde{p}_s) = \frac{\sum_{l=1}^N \sum_{k=1}^N p_s(l, k) \tilde{p}_s(l, k)}{\sqrt{\sum_{l=1}^N \sum_{k=1}^N p_s(l, k)^2 \sum_{l=1}^N \sum_{k=1}^N \tilde{p}_s(l, k)^2}}, \quad (2.19)$$

$$\text{KL}(p_s, \tilde{p}_s) = \sum_{l=1}^N \sum_{k=1}^N \log \frac{p_s(l, k)}{\tilde{p}_s(l, k)} p_s(l, k), \quad (2.20)$$

$$\text{SQ}(p_s, \tilde{p}_s) = \sum_{l=1}^N \sum_{k=1}^N \{p_s(l, k) - \tilde{p}_s(l, k)\}^2 p_s(l, k), \quad (2.21)$$

where $p_s(l, k)$ and $\tilde{p}_s(l, k)$ denote the values of the two-dimensional normalized histogram and the normalized $\tilde{p}_s(\mathbf{s}; \mathbf{m})$ in bin (l, k) , respectively. Equation 2.19 is the cosine of the angle between p_s and \tilde{p}_s , and the larger the value, the better the approximation. Equations 2.20 and 2.21 are the KL divergence and the expected squared distance, respectively. For comparison, we additionally computed these distance measures for a Laplace factorizable PDF with the same mean and marginal variance of the generated sources \mathbf{s} and a gaussian density with the same mean and covariance matrix.

The logarithms of p_s and \tilde{p}_s for $m_{12} = 0.95$ are shown in Figures 2a and 2b, respectively. The two PDFs seem to have similar properties in terms of the heavy-tailed profiles and linear correlations. Figures 2c to 2e show that \tilde{p}_s approximates p_s better than the Laplace and gaussian densities for all m_{12} . This is due to the fact that our approximation captures both the heavy tails of p_s and its dependency structure. Thus, we conclude that our approximation \tilde{p}_s captures, at least qualitatively, the basic characteristics of the true density of the dependent nongaussian components.

3 Estimation of the Model

In this section, we first show how to estimate the linear mixing model, equation 2.1, and the dependency matrix \mathbf{M} formed by \mathbf{m} using score matching and then discuss important implementation details.

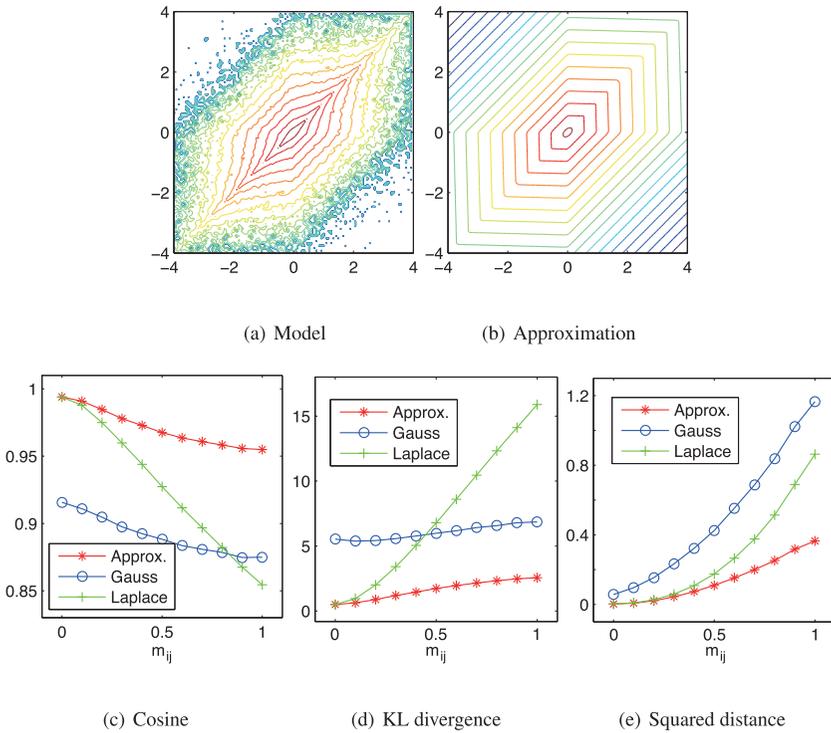


Figure 2: Numerical validation of the approximation in equation 2.18. (a) Contour plots of a simple nonparametric estimate of the density p_s . (b) Contour plot of the approximation \tilde{p}_s . (c–e) Goodness of the approximation as a function of m_{12} . “Approx.” is the approximation; “Gauss” and “Laplace” are standard models (multivariate gaussian density and Laplace factorizable density, respectively). Note that for panel c, larger is better, while for panels d and e, smaller is better.

3.1 Using Score Matching for the Estimation. Approximating p_s in equation 2.7 with \tilde{p}_s from equation 2.18 yields an approximative PDF for the data as

$$p_x(x) \approx \tilde{p}_x(x; \mathbf{W}, \mathbf{m}) \propto \prod_{i=1}^d \exp \left(-m_{ii} |\mathbf{w}_i^\top x| - \sum_{j>i} m_{ij} |\mathbf{w}_i^\top x - \mathbf{w}_j^\top x| \right) |\mathbf{W}|, \tag{3.1}$$

where \mathbf{w}_i denotes the i th row of \mathbf{W} . Both \mathbf{W} and \mathbf{m} are unknown parameters that we wish to estimate from a set of T observations $\{x_1, \dots, x_T\}$ of x . A conventional approach for the estimation would consist in maximizing

the likelihood. However, maximum likelihood estimation cannot be done here because $\tilde{p}_x(x; \mathbf{W}, \mathbf{m})$ is an unnormalized model, defined only up to a proportionality factor (partition function) that depends on \mathbf{m} . To cope with such estimation problems, several methods have been proposed (Gutmann & Hyvärinen, 2012; Hinton, 2002; Hyvärinen, 2005; see Gutmann & Hyvärinen, 2013a, for a review paper). One of the methods is score matching (Hyvärinen, 2005) whose objective function for models from the exponential family is a quadratic form (Hyvärinen, 2007). If \mathbf{W} is fixed, \tilde{p}_x in equation 3.1 belongs to the exponential family, so that estimation of \mathbf{m} given an estimate of \mathbf{W} would be straightforward with score matching. This motivated us to estimate the model in equation 3.1 by score matching and to optimize its objective function by alternating between \mathbf{W} and \mathbf{m} .

We next derive the score matching objective function $J(\mathbf{W}, \mathbf{m})$ for the joint estimation of \mathbf{W} and \mathbf{m} and then show how it is simplified when \mathbf{W} is considered fixed. By definition of score matching (Hyvärinen, 2005),

$$J(\mathbf{W}, \mathbf{m}) = \frac{1}{T} \sum_{t=1}^T \sum_{k=1}^d \frac{1}{2} \psi_k(x_t; \mathbf{W}, \mathbf{m})^2 + \partial_k \psi_k(x_t; \mathbf{W}, \mathbf{m}), \quad (3.2)$$

where ψ_k and $\partial_k \psi_k$ are the first- and second-order derivatives of $\log \tilde{p}_x$ with respect to the k th coordinate of x :

$$\psi_k(x; \mathbf{W}, \mathbf{m}) = \frac{\partial \log \tilde{p}_x(x; \mathbf{W}, \mathbf{m})}{\partial x_k}, \quad \partial_k \psi_k(x; \mathbf{W}, \mathbf{m}) = \frac{\partial^2 \log \tilde{p}_x(x; \mathbf{W}, \mathbf{m})}{\partial x_k^2}. \quad (3.3)$$

For our model in equation 3.1, we have

$$\begin{aligned} & \psi_k(x; \mathbf{W}, \mathbf{m}) \\ &= - \sum_{i=1}^d m_{ii} G'(\mathbf{w}_i^\top \mathbf{x}) w_{ik} - \sum_{j>i} m_{ij} G'(\mathbf{w}_i^\top \mathbf{x} - \mathbf{w}_j^\top \mathbf{x}) (w_{ik} - w_{jk}), \end{aligned} \quad (3.4)$$

$$\begin{aligned} & \partial_k \psi_k(x; \mathbf{W}, \mathbf{m}) \\ &= - \sum_{i=1}^d m_{ii} G''(\mathbf{w}_i^\top \mathbf{x}) w_{ik}^2 - \sum_{j>i} m_{ij} G''(\mathbf{w}_i^\top \mathbf{x} - \mathbf{w}_j^\top \mathbf{x}) (w_{ik} - w_{jk})^2, \end{aligned} \quad (3.5)$$

where the absolute value in equation 2.18 is approximated by $|u| \approx G(u) = \log \cosh(u)$ for numerical stability, $G'(u) = dG(u)/du = \tanh(u)$; and $G''(u) = \text{sech}^2(u)$.

Inspection of equations 3.4 and 3.5 shows that ψ_k and $\partial \psi_k$ are linear functions of the m_{ii} and m_{ij} . If we consider \mathbf{W} fixed and let $\mathbf{g}_k^w(x)$ be the column

vector formed by the terms in equation 3.4 that are multiplied by the elements of \mathbf{m} and let $\mathbf{h}_k^w(\mathbf{x})$ be the analogous vector for equation 3.5, we have $\psi_k(\mathbf{x}; \mathbf{W}, \mathbf{m}) = \mathbf{m}^\top \mathbf{g}_k^w(\mathbf{x})$ and $\partial_k \psi_k(\mathbf{x}; \mathbf{W}, \mathbf{m}) = \mathbf{m}^\top \mathbf{h}_k^w(\mathbf{x})$. The superscript w is used as a reminder that the vectors $\mathbf{g}_k^w(\mathbf{x})$ and $\mathbf{h}_k^w(\mathbf{x})$ depend on \mathbf{W} . For fixed \mathbf{W} , we can thus write $J(\mathbf{W}, \mathbf{m})$ in equation 3.2 as a quadratic form $J(\mathbf{m}|\mathbf{W})$,

$$J(\mathbf{m}|\mathbf{W}) = \frac{1}{2} \mathbf{m}^\top \left(\frac{1}{T} \sum_{t=1}^T \sum_{k=1}^d \mathbf{g}_k^w(\mathbf{x}_t) \mathbf{g}_k^w(\mathbf{x}_t)^\top \right) \mathbf{m} + \mathbf{m}^\top \left(\frac{1}{T} \sum_{t=1}^T \sum_{k=1}^d \mathbf{h}_k^w(\mathbf{x}_t) \right). \tag{3.6}$$

For fixed \mathbf{W} , an estimate of \mathbf{m} is obtained by minimizing $J(\mathbf{m}|\mathbf{W})$ under the constraint that the m_{ii} are positive and the other m_{ij} are nonnegative.

However, for estimation of \mathbf{W} , fixing \mathbf{m} does not lead to an objective that takes a simpler form than $J(\mathbf{W}, \mathbf{m})$ in equation 3.2. Therefore, we optimize $J(\mathbf{W}, \mathbf{m})$ by a simple gradient descent whose details follow.

3.2 Implementation Details. We estimate the parameters \mathbf{W} and \mathbf{m} of our statistical model, equation 3.1, for dependent nongaussian components by alternately minimizing $J(\mathbf{W}, \mathbf{m})$ in equation 3.2 with respect to \mathbf{W} and \mathbf{m} . We next discuss some important details in this optimization scheme.

In our discussion of \tilde{p}_s in equation 2.18, we noted that larger values of m_{ij} , $j > i$, indicate stronger (conditional) correlation between components s_i and s_j . In preliminary simulations, we observed that sometimes the estimated m_{ij} would take much larger values than the estimated m_{ii} , leading to estimated sources \hat{s}_i and \hat{s}_j , and hence estimated features $\hat{\mathbf{w}}_i$ and $\hat{\mathbf{w}}_j$, which were almost the same. In order to avoid this kind of degeneracy, we imposed the additional constraint that an m_{ii} had to be larger than the off-diagonal m_{ij} summed together. Having this additional constraint, we found that the strict positivity constraint on the m_{ii} could be relaxed to nonnegativity. In summary, we imposed the following constraints on \mathbf{m} :

$$(\forall (i, j) : i \leq j) \quad 0 \leq m_{ij}, (\forall i) \quad \sum_{j \neq i} m_{ij} \leq m_{ii}. \tag{3.7}$$

The constraints are linear, so that constrained minimization of $J(\mathbf{m}|\mathbf{W})$ can be done by standard methods from quadratic programming.

The mixing model, equation 1.1, has a scale indeterminacy because dividing a feature by some number while multiplying the corresponding source by the same amount does not change the value of \mathbf{x} . While this scale indeterminacy is a well-known phenomenon in ICA, the situation is here more complicated because we have terms of the form $m_{ii}|\mathbf{w}_i^\top \mathbf{x}|$ and $m_{ij}|\mathbf{w}_i^\top \mathbf{x} - \mathbf{w}_j^\top \mathbf{x}|$ in the model PDF instead of the simpler $|\mathbf{w}_i^\top \mathbf{x}|$ terms found in ICA. While in ICA, the scale indeterminacy is not a problem for

maximum likelihood estimation, we have found that for our model, it was necessary to explicitly resolve the indeterminacy by imposing a unit norm constraint on the \mathbf{w}_i (for whitened data).

For the optimization of $J(\mathbf{W}, \mathbf{m})$, we have to choose some initial values for \mathbf{W} and \mathbf{m} . We initialized \mathbf{W} using a maximum-likelihood-based ICA algorithm, additionally imposing the norm constraint on the \mathbf{w}_i . In more detail, we initialized \mathbf{W} as $\widehat{\mathbf{W}}$:

$$\widehat{\mathbf{W}} = \arg \min_{1 \leq i \leq d, \|\mathbf{w}_i\|=1} J_0(\mathbf{W}), \quad J_0(\mathbf{W}) = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^d G(\mathbf{w}_i^\top \mathbf{x}_t) - \log |\det \mathbf{W}|. \quad (3.8)$$

Given the initial value $\widehat{\mathbf{W}}$, we obtained an initial value for \mathbf{m} by minimizing $J(\mathbf{m}|\widehat{\mathbf{W}})$ in equation 3.6 under the constraints in equation 3.7.

While $J(\mathbf{m}|\mathbf{W})$ can be minimized by quadratic programming, minimization of $J_0(\mathbf{W})$ and $J(\mathbf{W}, \mathbf{m})$ for fixed \mathbf{m} has to be done by less powerful methods. We used a simple gradient descent algorithm where the step-size μ_k at each iteration k was chosen adaptively by trying out $2\mu_{k-1}$ and $1/2\mu_{k-1}$ in addition to μ_{k-1} , and selecting the one that yielded the smallest objective.

Algorithm 1 summarizes our approach to estimate the model, equation 3.1, where it is assumed that the data have already been preprocessed by whitening and, optionally, dimension reduction by principal component analysis (PCA). In the proposed method, good initialization is important because the objective function has local optima, which can produce spurious correlations in the estimated \mathbf{s} . Therefore, we first perform ICA to give a reasonable initialization for both \mathbf{M} and \mathbf{W} . One weakness of the proposed method is that optimization for high-dimensional and large data can be slow compared with ICA because we alternately repeat steps 1 and 2. To alleviate this problem, in section 5, we perform dimensionality reduction by PCA and estimate \mathbf{W} and \mathbf{M} based on a randomly chosen subset of data samples at every repeat of steps 1 and 2.

4 Simulations on Artificial Data

In this section, using artificial data, we evaluate how well the proposed method identifies the sources and their dependency structure. We compare the proposed method to ICA and CTA.

4.1 Methods. For our evaluation, we used data generated according to the model in section 2.1. We considered both data with independent components and data with components that had statistical dependencies within certain blocks. The interest of using independent components (sources) in the evaluation is to check that the model does not impose dependencies among the estimates when the underlying sources are truly independent.

Algorithm 1: Estimation of the Mixing Matrix \mathbf{A} and Dependency Matrix \mathbf{M} .

Input: Data $\{x_1, x_2, \dots, x_T\}$ which have been whitened.

- Initialization (ICA): Compute $\widehat{\mathbf{W}}$ as in equation 3.8. Fixing $\mathbf{W} = \widehat{\mathbf{W}}$, compute $\widehat{\mathbf{m}}$ by minimizing $J(\mathbf{m}|\mathbf{W})$ in equation 3.6 under the constraints in equation 3.7, using a standard solver for quadratic programs.

- Repeat steps 1 and 2 until some conventional convergence criterion is met:

Step 1: Fixing $\mathbf{m} = \widehat{\mathbf{m}}$, update $\widehat{\mathbf{W}}$ by taking one gradient step to minimize $J(\mathbf{W}, \mathbf{m})$ in equation 3.2 under the unit norm constraint on the rows w_i of \mathbf{W} .

Step 2: Fixing $\mathbf{W} = \widehat{\mathbf{W}}$, recompute $\widehat{\mathbf{m}}$ by minimizing $J(\mathbf{m}|\mathbf{W})$ in equation 3.6 under the constraints in equation 3.7, using a standard solver for quadratic programs.

Output: Mixing matrix $\widehat{\mathbf{A}} = \widehat{\mathbf{W}}^{-1}$, dependency matrix $\widehat{\mathbf{M}}$ formed by $\widehat{\mathbf{m}}$.

For the independent sources, the u_{ii} in equation 2.2 were sampled from an inverse-gamma density with shape parameter $k_{ii} = 2$ and scale parameter $(m'_{ii})^2 = 1$. The other elements u_{ij} , $i \neq j$ were set to zero. Thus, $\mathbf{\Lambda}$ was a diagonal matrix, with $\lambda_{ii} = u_{ii}$, and the generated sources s_i were statistically independent of each other. For the block-dependent sources, $\lambda_{ii} = \sum_{k=1}^d u_{ik}$ and $\lambda_{ij} = -u_{ij}$ as in equation 2.2. The u_{ii} were for all i from an inverse-gamma density with shape parameter $k_{ii} = 2$ and scale parameter $(m'_{ii})^2 = 1$. The variables u_{12} , u_{13} , and u_{23} were sampled from an inverse-gamma density with shape parameter $k_{ii} = 2$ and scale parameter $(m'_{ij})^2 = 1/3$, while the remaining u_{ij} were set to zero. With this setup, the s_1 , s_2 , and s_3 are statistically dependent, while the other sources are conditionally independent. This dependency structure creates a block structure in the linear and energy correlation matrices where any pairs of (s_1, s_2, s_3) show relatively stronger dependencies than the other pairs (see Figures 4c and 4d). The energy correlation matrix is the correlation matrix of the squared random variables whose (i, j) th element is given by

$$\frac{E\{s_i^2 s_j^2\} - E\{s_i^2\}E\{s_j^2\}}{\sqrt{E\{(s_i^2 - E\{s_i^2\})^2\}E\{(s_j^2 - E\{s_j^2\})^2\}}}, \quad (4.1)$$

where E denotes the expectation and $E\{s_i\} = 0$. After generating the sources, each component was standardized so that it has the zero mean and unit variance.

The observed data were generated from the mixing model, equation 2.1, where the elements in \mathbf{A} were sampled from the standard normal density. The data dimension was $d = 10$, and the total number of observations was $T = 20,000$. The preprocessing was whitening based on PCA. The performance matrix $\mathbf{P} = \mathbf{W}\mathbf{A}$ was used to visualize and evaluate the results. If \mathbf{P} is close to a permutation matrix, the sources are well identified.

To measure the goodness of the estimated dependency matrix, we used the scale parameters of the inverse-gamma densities employed to generate the sources in this simulation. For independent sources, we first constructed a reference matrix \mathbf{M}' by setting the diagonals and off-diagonals to $m'_{ii} = 1$ and zeros, respectively. To enable a comparison to the reference matrix, we normalize $\widehat{\mathbf{M}}$ by dividing \widehat{m}_{ij} by $\sqrt{\widehat{m}_{ii}\widehat{m}_{jj}}$ so that the diagonals are all ones and denote the normalized $\widehat{\mathbf{M}}$ by $\widehat{\mathbf{M}}'$. Finally, the goodness was measured by

$$\text{Error}_{\mathbf{M}'} = \|\mathbf{M}' - \widehat{\mathbf{M}}'\|_{\text{Fro}}, \quad (4.2)$$

where $\|\cdot\|_{\text{Fro}}$ denotes the Frobenius norm. The reason of the error definition, equation 4.2, is that since the different shape parameter k_{ii} from the ones in the inverse-gamma density, equation 2.17, was used for numerical stability, we could not know the exact \mathbf{M} and therefore had to focus on the relative values of the elements in $\widehat{\mathbf{M}}$. For block sources, we constructed the reference matrix by setting the diagonals in \mathbf{M}' to 1 and the off-diagonals to $m'_{ij}/\sqrt{\sum_{j=1}^d m'_{ij} \sum_{i=1}^d m'_{ji}}$ inside the block and to zeros outside the block. As a result, \mathbf{M}' becomes a diagonally dominant matrix with a block structure.

For comparison, we performed ICA and CTA (Sasaki et al., 2013) on the same data.¹ The ICA method was the same as the method used to initialize \mathbf{W} in algorithm 1, with the unit norm constraint on the rows \mathbf{w}_i of \mathbf{W} . For all methods, to avoid local optima, we performed 10 runs with different initialization of \mathbf{W} and chose the run with the best value of each objective function. For ICA and CTA, after estimating \mathbf{W} , we estimated their dependency matrices by minimizing equation 3.6 with the same constraints, equation 3.7.

¹The Matlab package for CTA is available at the first author's web page: <https://sites.google.com/site/hworksites/home/software/cta>.

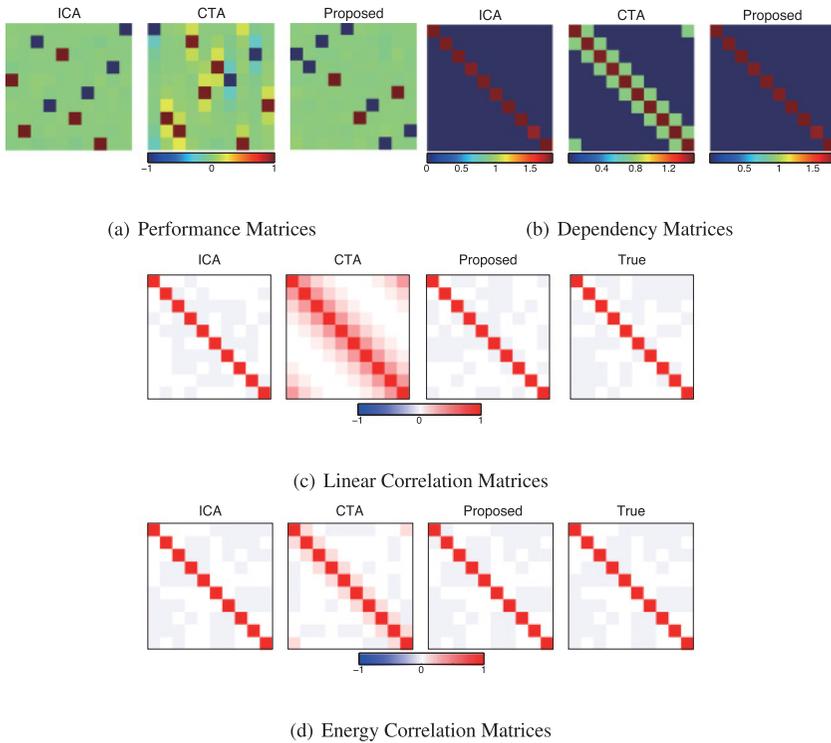


Figure 3: Simulation results for independent sources. From panels b to d, the permutation indeterminacy of ICA and the proposed method was compensated so that the largest element in each row of the performance matrix is on the diagonal. The true linear and energy correlation matrices of the sources are presented in the right-most figures of panels c and d.

4.2 Results. We first report the performance on a single data set and then the average performance on 100 randomly generated data sets. The results for the independent sources from the single data set are presented in Figure 3. For ICA and the proposed method, the performance matrices are close to permutation matrices (see Figure 3a), the estimated dependency matrices resemble a diagonal matrix (see Figure 3b), as they should be, and the correlation matrices are almost diagonal (see Figures 3c and 3d). For CTA, the performance matrix includes more cross-talk, the dependency matrix is tridiagonal, and the linear and energy correlation matrices are clearly different from a diagonal matrix. These unsatisfactory results for CTA come from the fact that the dependency structure of CTA is prefixed, and thus CTA forcibly imposes linear correlations among the estimated neighboring components even though the original components are linearly

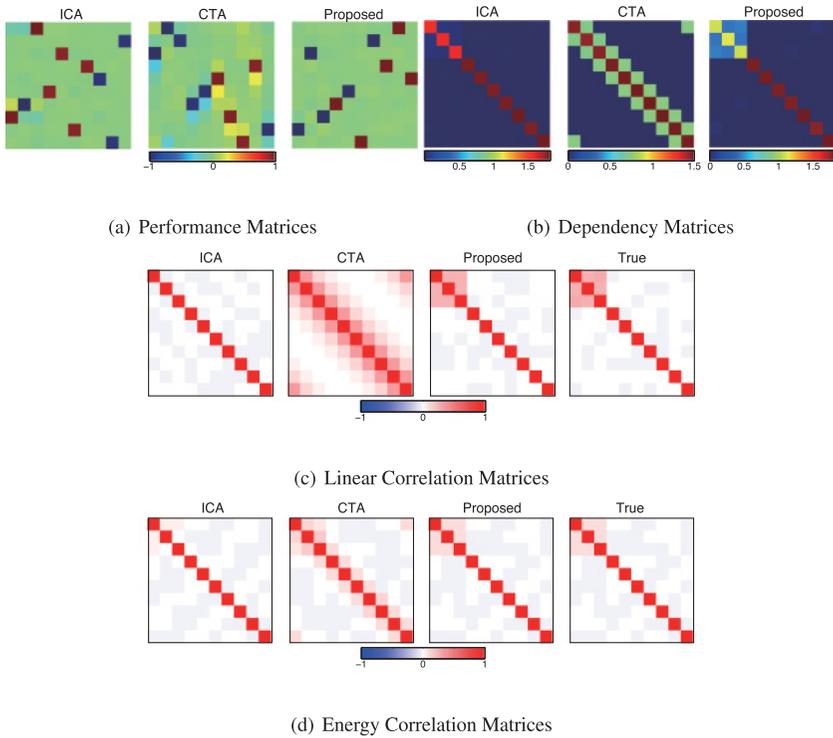


Figure 4: Simulation results for block-dependent sources. From panels b to d, the permutation indeterminacy of ICA and the proposed method was compensated so that the largest element in each row of the performance matrix is on the diagonal.

uncorrelated. This drawback has been already reported by Sasaki et al. (2013). In contrast, the proposed method learned automatically that the sources are independent and solved the identifiability issue of CTA.

The results for block-dependent sources are shown in Figure 4. The proposed method separates the sources, that is, estimates the linear components, with good accuracy, while ICA has more errors. The performance matrix for CTA includes again a lot of cross-talk (see Figure 4a). Regarding \mathbf{M} and the correlation matrices, we compensated for the permutation indeterminacy of the sources for both ICA and the proposed method so that the largest element in each row of the performance matrix is on the diagonal. Figure 4b shows that the proposed method yields a dependency matrix with a clearly visible block structure in the upper left corner, while ICA and CTA do not. In addition, the linear and energy correlation matrices have the block structure for the proposed method, while ICA and CTA do not

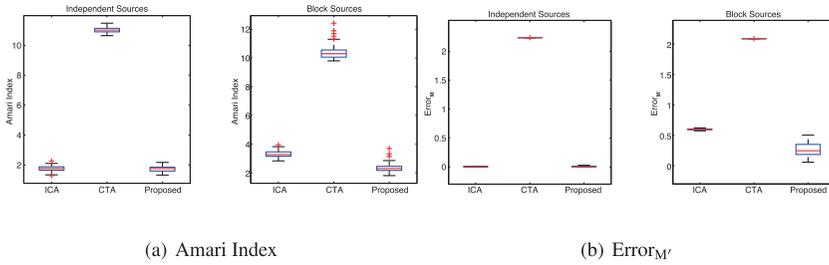


Figure 5: Estimation error for 100 runs for both independent and block sources, summarized in terms of Amari index and $Error_M$. In the comparisons, ICA and CTA were used to estimate \mathbf{W} , and their dependency matrices were estimated by minimizing the objective in equation 3.6 with the same constraints, equation 3.7.

produce it (see Figures 4c and 4d). Thus, only the proposed method was able to correctly identify the dependency structure of the latent sources.

We further analyzed the average performance for 100 randomly generated data sets. Figure 5a shows the distribution of the Amari index (AI) (Amari, Cichocki, & Yang, 1996) for the independent and block-dependent sources. AI is an established measure to assess the performance of blind source separation algorithms, and a smaller value means better performance. For independent sources, the proposed method has almost the same performance as ICA, while the performance of CTA is much worse. For block-dependent sources, the performance of the proposed method is better than ICA and CTA. This result means that when the original sources s_i are linearly correlated, ICA is not the best method in terms of identifiability of the mixing matrix and that taking into account the dependency structure for linear correlations improves the identifiability.

For the goodness of the estimated dependency matrix, we plot the distribution of $Error_M$ in Figure 5b. In this analysis, the permutation indeterminacy was compensated as done in Figure 4. The plot confirms the qualitative findings in Figures 3 and 4: the proposed method is able to capture the dependency structure of the sources better than ICA and CTA.

5 Application to Real Data

This section demonstrates the applicability of the proposed method on two kinds of real data: speech data and natural image data.

Basic ICA and related methods based on energy (square) correlation work well on raw speech and image data due to the symmetry of the data PDFs. The symmetry implies in particular that positive and negative values of linear features are to some extent equivalent, as implicitly assumed in the computation of Fourier spectra or complex cell outputs

in models of early (mammalian) vision, which is compatible with energy correlations.

However, on higher levels of feature extraction, such symmetry can no longer be found, and energy correlations cannot be expected to be meaningful. Our goal here is to apply our new method on such higher-level features, where linear correlations are likely to be important. In particular, we use speech spectrograms and outputs of complex cells simulating computations in the visual cortex, respectively.

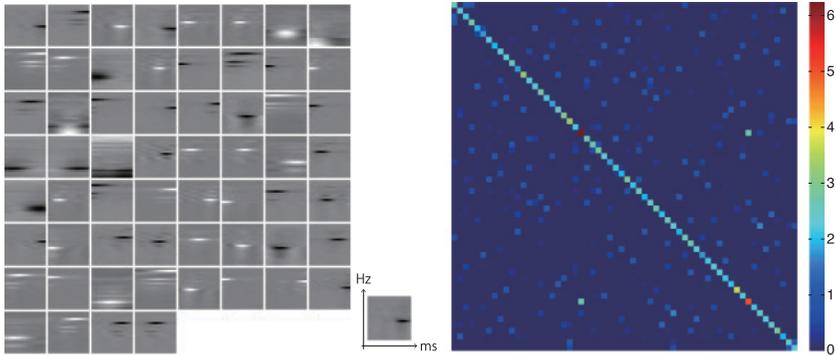
5.1 Speech Data. Previously, sparse coding (Olshausen & Field, 1996) and ICA-related methods have been applied to audio data to investigate the basic properties of cells in the primary auditory cortex (A1) (Klein, König, & Körding, 2003; Terashima & Hosoya, 2009; Terashima, Hosoya, Tani, Ichinohe, & Okada, 2013). More recently, topographic ICA (TICA) (Hyvärinen et al., 2001) was employed to analyze spectrogram data, and feature maps were learned that are similar to the tonotopic maps in A1 (Terashima & Okada, 2012). However, in TICA, the dependency structure is influenced by higher-order correlations only, and it is fixed to nearby components beforehand. Furthermore, using energy correlations for spectrograms may not be well justified. Here, we lift these restrictions and learn the dependency structure from the data by taking both linear and higher-order correlations between the latent sources into account.

Following Terashima and Okada (2012), we used human narratives data (International Phonetic Association, 1999). The data were downsampled to 8 kHz, and the spectrograms were computed by using the NSL toolbox.² After resizing the vertical (spectral) size of the spectrograms from 128 to 20, short spectrograms were randomly extracted with the horizontal (temporal) size equal to 20. The vectorized spectrogram patches were our $T = 100,000$ input data points $\{x_1, x_2, \dots, x_T\}$.

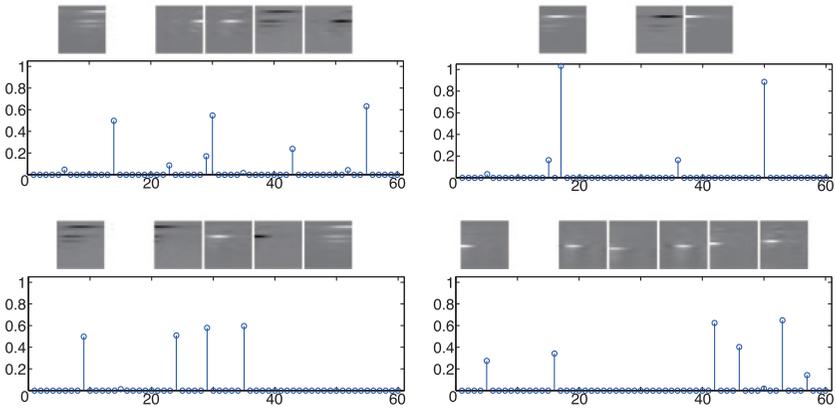
As preprocessing, we removed the DC component of each x_t and then rescaled each x_t to unit norm. Finally, whitening and dimensionality reduction were performed simultaneously by PCA. We retained $d = 60$ dimensions. To reduce the computational cost, in this experiment, at every repeat of steps 1 and 2 in algorithm 1, we randomly selected 30,000 data points from $T = 100,000$ data points to be used for estimation.

The estimated basis vectors a_i will be visualized in the original domain as spectrograms. For the estimated dependency matrix, we apply a multidimensional scaling (MDS) method to $\hat{\mathbf{M}}$ to visualize the dependency structure on the two-dimensional plane. To employ MDS, we constructed a distance matrix from $\hat{\mathbf{M}}$ similar to that done in Hurri and Hyvärinen (2003). We first normalized each element \hat{m}_{ij} by $\sqrt{\hat{m}_{ii}\hat{m}_{jj}}$ to make the diagonals ones, then computed the square root of each element in the normalized

²Available at <http://www.isr.umd.edu/Labs/NSL/Downloads.html>.



(a) Estimated basis vectors (left) and dependency matrix (right)



(b) For a selected i , the m_{ij} as stem plots, and the corresponding basis vectors with strong m_{ij} (the left-most being \hat{a}_i and the others \hat{a}_j)

Figure 6: Results for speech (spectrogram) data.

matrix, and finally subtracted each element from one. The purpose of MDS is to project the points in a high-dimensional space to the two-dimensional plane so that the distance in the high-dimensional space is preserved as much as possible in the two-dimensional space. Thus, applying MDS should yield a representation where the dependent features (basis vectors) are close to each other.

The estimated basis vectors \hat{a}_i and dependency matrix \hat{M} are presented in Figure 6a. Most of the basis vectors show vertically (spectrally) and horizontally (temporally) localized patterns with single or multiple peaks.



Figure 7: Speech data. Visualization of the estimated dependency structure between features (basis vectors) by MDS. In the figure, features with stronger m_{ij} should be closer to each other in this visualization. The positions of some features that are too close or too far away were magnified in order to show all features in a reasonable scale.

These properties have been also found in previous work (Terashima & Okada, 2012). But unlike previous work, we also estimated the dependency structure from the data. As shown in the right panel of Figure 6a, the off-diagonal elements of the dependency matrix are sparse: most of the elements are zero. Figure 6b shows that basis vectors with similar peak frequencies tend to have strong (conditional) dependencies. The visualization of $\hat{\mathbf{M}}$ further globally supports this observation (see Figure 7).

Compared with previous work, the properties of nearby features (basis vectors) in Figure 7 seem to be more consistent: nearby features tend to have similar peak positions along with the spectral (vertical) axis, while the peak positions on the temporal (horizontal) axis are more random. However, Terashima and Okada (2012) found that nearby features often show different peak positions on the spectral axis, and the estimated features on the topographic map are locally disordered. These results may reflect that linear correlations in sound spectrogram data should be important

dependencies and that the proposed method captured the structure in the data better than TICA.

5.2 Outputs of Complex Cells. We next apply our method to the outputs of simulated complex cells in the primary visual cortex when stimulated with natural image data. Previously, ICA, nonnegative sparse coding and CTA have been applied to these kinds of data, and some prominent features such as long contours and topographic maps have been learned (Hoyer & Hyvärinen, 2002; Hyvärinen, Gutmann, & Hoyer, 2005; Sasaki et al., 2013). However, these methods have either assumed that the features are independent or prefixed their dependency structure. Our method removes this restriction and learns the dependency structure from the data.

As in the previous work cited above, we computed the outputs of the simulated complex cells \mathbf{x} as

$$x'_k = \left(\sum_{x,y} W_k^o(x, y) I(x, y) \right)^2 + \left(\sum_{x,y} W_k^e(x, y) I(x, y) \right)^2,$$

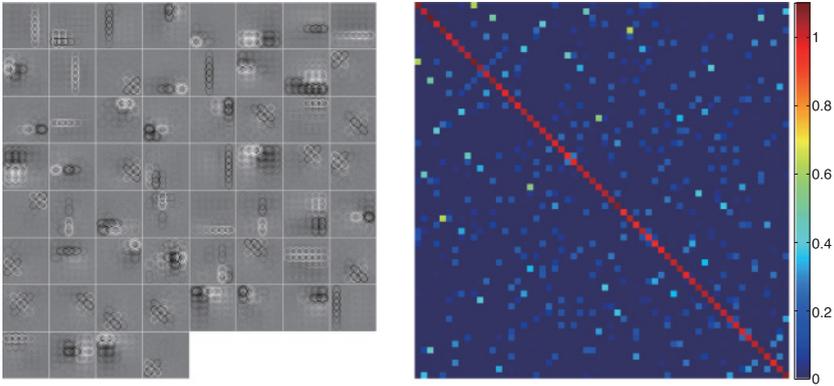
$$x_k = \log(x'_k + 1.0),$$

where $I(x, y)$ is a 24×24 natural image patch³ and $W_k^o(x, y)$ and $W_k^e(x, y)$ are odd- and even-symmetric Gabor functions with the same spatial positions, orientation, and frequency. The total number of outputs was $T = 100,000$. The complex cells were arranged on a 6 by 6 spatial grid, having four orientations each. In total, there were 144 cells. Since the simulated complex cells in this experiment are stimulated by natural images, we regard these data as real data. We performed the same preprocessing steps as in section 5.1; the dimensionality here was reduced to $d = 60$. As in the previous section, to reduce the computational cost, we randomly selected a subset of data points from the entire data points at every repeat of the two steps in algorithm 1. We visualized the basis vectors as in previous work (Hoyer & Hyvärinen, 2002; Hyvärinen et al., 2005): each basis vector is visualized by ellipses that have the orientation and spatial position preferences of the underlying complex cells.

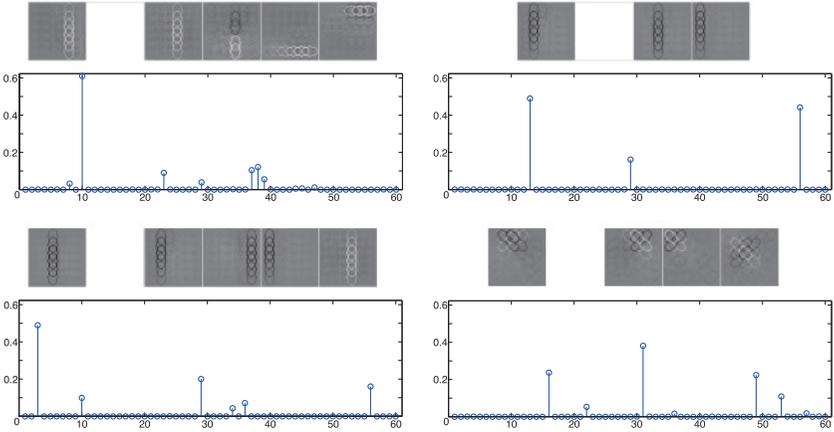
The estimated basis vectors $\hat{\mathbf{a}}_i$ and the dependency matrix $\hat{\mathbf{M}}$ are presented in Figure 8. Prominent features among the basis vectors are long contours, as also found in previous work (Hoyer & Hyvärinen, 2002; Hyvärinen et al., 2005).

Unlike in previous work, we also learned the dependencies among the features. As with the speech data, the off-diagonal elements of the

³To compute the complex cell outputs, we used the `contournet` package which is available at <http://www.cs.helsinki.fi/u/phoyer/software.html>.



(a) Estimated basis vectors (left) and dependency matrix (right)



(b) For a selected i , the m_{ij} as stem plots, and the corresponding basis vectors with stronger m_{ij} (the left-most being \hat{a}_i and the others \hat{a}_j)

Figure 8: Results for the simulated complex cell data.

dependency matrix are sparse (see Figure 8a, right), and similar features tend to have stronger dependencies (see Figure 8b).

Figure 9 visualizes the dependency structure by MDS as in section 5.1. This visualization supports the observation that the contour features tend to have stronger dependencies. In particular, it is often the case that contour features are closer to other contour features that are slightly shifted along



Figure 9: Complex cell outputs. Visualization of the estimated dependency structure between features (basis vectors) by MDS. In the graph, features with stronger m_{ij} should be closer to each other in this visualization.

their nonpreferred orientation. If put together, two such contour features would form either a broader contour of the same orientation or a slightly bent contour that is even longer. This property is in line with higher-level features learned using a three-layer model of natural images (Gutmann & Hyvärinen, 2013b).

We further investigate whether the proposed method estimated linearly correlated components on real data in contrast to previous energy-correlation-based methods (Karklin & Lewicki, 2005; Köster & Hyvärinen,

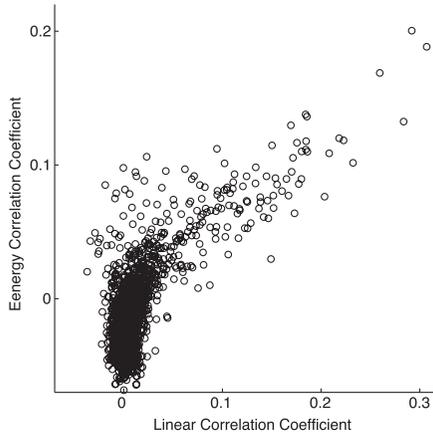


Figure 10: A scatter plot for linear and energy correlation coefficients from the estimated sources. Each point corresponds to one pair of estimated components (\hat{s}_i, \hat{s}_j) .

2010; Osindero et al., 2006). Figure 10 shows a scatter plot for linear and energy correlation coefficients for all the pairs of estimated components. As the sparsity of $\hat{\mathbf{M}}$ implies, most pairs have only weak statistical dependencies. However, some pairs show both strong linear and energy correlations. Thus, the proposed method did find linearly correlated components on real data as well.

6 Discussion

We discuss some connections to previous work and possible extensions of the proposed method.

6.1 Connection to Previous Work. We proposed a novel method to simultaneously estimate nongaussian components and their dependency (correlation) structure. So far, a number of methods to estimate nongaussian components have been proposed: ICA assumes that the latent nongaussian components are statistically independent, and ISA (Hyvärinen & Hoyer, 2000) and topographic methods (Hyvärinen et al., 2001; Mairal et al., 2011; Sasaki et al., 2013) prefix the dependency structure inside predefined groups of components, or overlapping neighborhoods of components. Methods to estimate tree-dependency structures have also been proposed (Bach & Jordan, 2003; Zoran & Weiss, 2009). In contrast to these methods, our method does not make any assumptions on the dependency structure to be estimated.

Some methods based on two-layer models also estimate the dependency structures of nongaussian components from data (Karklin & Lewicki, 2005; Köster & Hyvärinen, 2010; Osindero et al., 2006). Most of the methods mainly focus on high-order correlations, assuming that the components are linearly uncorrelated. The method proposed by Osindero et al. (2006) can estimate an overcomplete model as well, which necessarily makes the estimated components linearly correlated; however, it is unclear in what way such linear correlations reflect the dependency structure of the underlying sources. Our method models both linear and higher-order correlations explicitly. As a new theoretical approach, we have proposed in this article a generative model for random precision matrices. The proposed generative model generalizes a previous generative model for sources used in existing two-layer methods (Karklin & Lewicki, 2005; Osindero et al., 2006). The previous model corresponds to a special case in our model where the off-diagonal elements are deterministically zero.

Another line of related work is graphical models for latent factors (He, Qi, Kavukcuoglu, & Park, 2012). This work assumes that both the latent factors and the undirected graph are sparse. The main goal of that approach is to estimate a latent lower-dimensional representation of the input data where pair-wise dependencies between latent factors are represented by an undirected graph. The main difference to our work is that He et al. (2012) use a constant precision matrix instead of a stochastic one, and thus they do not really model nongaussian components. Instead, He et al. (2012) emphasize sparsity in the sense that the undirected graphs to be estimated have sparse edges. Our method can be easily extended to include the sparsity constraint while still keeping the objective function a simple quadratic form.

6.2 Extensions of the Proposed Method. One extension of the proposed method would be to incorporate prior information or an additional constraint on \mathbf{M} . Recently, a number of methods to estimate sparse gaussian graphical models have been proposed (Banerjee, El Ghaoui, & d'Aspremont, 2008; Friedman, Hastie & Tibshirani, 2008). In our method, the sparseness constraint on m_{ij} can be easily incorporated into the objective function $J(\mathbf{m}|\mathbf{W})$ in equation 3.8, yielding $J_\lambda(\mathbf{m}|\mathbf{W})$,

$$J_\lambda(\mathbf{m}|\mathbf{W}) = J(\mathbf{m}|\mathbf{W}) + \lambda \mathbf{m}^\top \mathbf{1}, \quad (6.1)$$

where λ is the regularization parameter and $\mathbf{1}$ is a vector of ones. The objective function $J_\lambda(\mathbf{m}|\mathbf{W})$ is still a quadratic form, and its minimization is not more difficult than the minimization of $J(\mathbf{m}|\mathbf{W})$. However, this extension involves selecting the parameter λ , which adds a complexity to the problem that needs to be solved. In practice, we may perform cross-validation (CV) under some criterion, but employing CV is computationally demanding,

especially in our alternate optimization. Therefore, selecting an appropriate value of λ is an issue for the future.

Another extension would be to estimate additional parameters modelling linear correlations. As shown in appendix E, the model for the sources in equation 2.18 can be generalized to

$$\tilde{p}_s(\mathbf{s}; \mathbf{m}, \Theta) \propto \prod_{i=1}^d \exp \left(-\theta_{ii} m_{ii} |s_i| - \sum_{j>i} m_{ij} \sqrt{\theta_{ii} s_i^2 + \theta_{jj} s_j^2 + 2\theta_{ij} s_i s_j} \right),$$

where the θ_{ii} are positive parameters and $\theta_{ij} \in \{-1, 0, 1\}$ for $i \neq j$. The PDF in equation 2.18 is a special case of the above PDF: it is obtained for $\theta_{ii} = 1$ and $\theta_{ij} = -1$ for $i \neq j$. Thus, by estimating all the θ_{ij} , we can estimate more complex dependency structures, generalizing our method. However, estimating θ_{ij} leads a complicated optimization problem because the objective function for score matching is no longer a quadratic form and the θ_{ij} are discrete variables. That is why we decided to use the simpler model in equation 2.18 and leave this interesting challenge for future work.

7 Conclusion

In this article, we proposed a method to simultaneously estimate nongaussian components and their dependency structure. The dependency structure is defined in terms of both linear and higher-order correlations, can be represented in a convenient form, and thus can be readily interpreted. When score matching is used, the estimation of the dependency structure is particularly simple because the objective function takes a quadratic form and can be minimized using standard methods from quadratic programming.

The proposed method generalizes previous methods: independent component and correlated topographic analysis both assume prefixed dependency structures for the nongaussian components, while our method flexibly estimates the dependency structure from the data themselves. Several existing methods based on two-layer models also aim at estimating the dependency structure, but they focus on higher-order correlations, ignoring linear correlations.

We experimentally demonstrated on artificial data that the proposed method improves identifiability of latent nongaussian components over existing methods by learning their dependency structure, and application to the outputs of simulated complex cells and the spectrograms of natural audio data unveiled new kinds of relationships among the latent nongaussian components.

Appendix A: Calculations for Equation 2.5 _____

With equation 2.2, the conditional PDF, equation 2.4, can be written as

$$p(\mathbf{s}|\mathbf{U}) = \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ \lambda_{ii} s_i^2 + \sum_{j \neq i} \lambda_{ij} s_i s_j \right\} \right), \tag{A.1}$$

$$= \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 \sum_{k=1}^d u_{ik} - \sum_{j \neq i} s_i s_j u_{ij} \right\} \right), \tag{A.2}$$

$$= \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{k \neq i} s_i^2 u_{ik} - \sum_{j \neq i} s_i s_j u_{ij} \right\} \right). \tag{A.3}$$

Since the u_{ij} are symmetric, we have further

$$p(\mathbf{s}|\mathbf{U}) = \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{k>i} (s_i^2 + s_k^2) u_{ik} - 2 \sum_{j>i} s_i s_j u_{ij} \right\} \right), \tag{A.4}$$

$$= \frac{|\mathbf{\Lambda}|^{1/2}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{j>i} (s_i - s_j)^2 u_{ij} \right\} \right), \tag{A.5}$$

which is equation 2.5.

Appendix B: Approximation of the Determinant in Equation 2.9 _____

We prove that equation 2.9 approximates the determinant with a lower bound and that the resulting approximation \tilde{p}_s of the source density p_s is an unnormalized model.

By definition in equation 2.2, $\mathbf{\Lambda}$ is a strictly diagonally dominant matrix. Ostrowski's inequality (Ostrowski, 1937) thus yields

$$|\mathbf{\Lambda}| \geq \prod_{i=1}^d \left(\lambda_{ii} - \sum_{\substack{j=1 \\ j \neq i}}^d \lambda_{ij} \right), \tag{B.1}$$

and by definition of the λ_{ij} in equation 2.2, we obtain

$$|\mathbf{\Lambda}| \geq \prod_{i=1}^d \left(\sum_{j=1}^d u_{ij} - \sum_{\substack{j=1 \\ j \neq i}}^d u_{ij} \right) = \prod_{i=1}^d u_{ii}, \quad (\text{B.2})$$

which shows that the approximation, equation 2.9, corresponds to a lower bound of the determinant.

Applying the approximation of the determinant to $p(\mathbf{s}|\mathbf{U})$ in equation 2.5 gives $\tilde{p}(\mathbf{s}|\mathbf{U})$,

$$\tilde{p}(\mathbf{s}|\mathbf{U}) = \frac{\prod_{i=1}^d \sqrt{u_{ii}}}{(2\pi)^{d/2}} \exp \left(-\frac{1}{2} \sum_{i=1}^d \left\{ s_i^2 u_{ii} + \sum_{j>i} (s_i - s_j)^2 u_{ij} \right\} \right), \quad (\text{B.3})$$

which is a lower bound for $p(\mathbf{s}|\mathbf{U})$ due to equation B.2:

$$\tilde{p}(\mathbf{s}|\mathbf{U}) \leq p(\mathbf{s}|\mathbf{U}). \quad (\text{B.4})$$

Using the approximation $\tilde{p}(\mathbf{s}|\mathbf{U})$ instead of $p(\mathbf{s}|\mathbf{U})$ in equation 2.6 yields an approximation \tilde{p}_s of the PDF of the dependent nongaussian components, which inherits the lower bound property,

$$\tilde{p}_s(\mathbf{s}) = \int_0^\infty \tilde{p}(\mathbf{s}|\mathbf{U}) p_u(\mathbf{U}) d\mathbf{U} \leq \int_0^\infty p(\mathbf{s}|\mathbf{U}) p_u(\mathbf{U}) d\mathbf{U} = p_s(\mathbf{s}). \quad (\text{B.5})$$

Since p_s integrates to one, we have

$$\int \tilde{p}_s(\mathbf{s}) d\mathbf{s} \leq 1, \quad (\text{B.6})$$

which means that \tilde{p}_s is an unnormalized model.

Appendix C: Properties of the Nonlinearities in Equation 2.15

We here investigate properties of the nonlinearities $g_{ij}(v)$ in equation 2.15. It is shown that unless constant, the $\log g_{ij}(v)$ are monotonically decreasing convex functions.

For the analysis, it is helpful to introduce the functions $\phi_{ij}(u)$, $u \geq 0$,

$$\phi_{ij}(u) = \begin{cases} \sqrt{u} p_{ii}(u) & \text{if } i = j \\ p_{ij}(u) & \text{otherwise,} \end{cases} \quad (\text{C.1})$$

which allow us to rewrite the $g_{ij}(v)$ as

$$g_{ij}(v) = \int_0^\infty \exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u) du, \quad v \geq 0, \tag{C.2}$$

for all (i, j) . The functions ϕ_{ij} are nonnegative, but they are not PDFs if $i = j$ because they do not integrate to one. We note that g_{ij} , $i \neq j$, is constant if p_{ij} corresponds to a Dirac measure concentrated at zero, that is, to a u_{ij} that is deterministically zero. In what follows, we assume that the p_{ij} are PDFs with nonzero variance.

The derivative of $\log g_{ij}(v)$ is

$$\frac{\partial \log g_{ij}(v)}{\partial v} = \frac{\frac{\partial g_{ij}(v)}{\partial v}}{g_{ij}(v)} \tag{C.3}$$

$$= \frac{-\frac{1}{2} \int_0^\infty u \exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u) du}{\int_0^\infty \exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u) du}, \tag{C.4}$$

which is negative for all $v > 0$. Importantly, the derivative is related to the expected value μ_{ijv} of the PDF $\gamma_{ijv}(u)$,

$$\gamma_{ijv}(u) = \frac{\exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u)}{\int_0^\infty \exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u) du}, \tag{C.5}$$

that is,

$$\frac{\partial \log g_{ij}(v)}{\partial v} = -\frac{1}{2} \mu_{ijv}. \tag{C.6}$$

The second derivative equals

$$\frac{\partial^2 \log g_{ij}(v)}{\partial v^2} = \frac{\partial^2 g_{ij}(v)}{\partial v^2} \frac{1}{g_{ij}(v)} - \left(\frac{\partial g_{ij}(v)}{\partial v}\right)^2 \frac{1}{g_{ij}(v)^2} \tag{C.7}$$

$$= \frac{\partial^2 g_{ij}(v)}{\partial v^2} \frac{1}{g_{ij}(v)} - \left(\frac{\partial \log g_{ij}(v)}{\partial v}\right)^2 \tag{C.8}$$

$$= \frac{\partial^2 g_{ij}(v)}{\partial v^2} \frac{1}{g_{ij}(v)} - \frac{1}{4} \mu_{ijv}^2, \tag{C.9}$$

which can be written in terms of the variance of the PDF γ_{ijv} . We have

$$\frac{\partial^2 g_{ij}(v)}{\partial v^2} = \frac{1}{4} \int_0^\infty u^2 \exp\left(-\frac{1}{2}vu\right) \phi_{ij}(u) du, \tag{C.10}$$

so that

$$\frac{\partial^2 g_{ij}(v)}{\partial v^2} \frac{1}{g_{ij}(v)} = \frac{1}{4} \int_0^\infty u^2 \gamma_{ijv}(u) du, \quad (\text{C.11})$$

and

$$\frac{\partial^2 \log g_{ij}(v)}{\partial v^2} = \frac{1}{4} \left(\int_0^\infty u^2 \gamma_{ijv}(u) du - \mu_{ijv}^2 \right). \quad (\text{C.12})$$

The term in the parentheses is the variance of the PDF γ_{ijv} and hence positive since the PDF is not degenerate. In conclusion, for nonconstant functions g_{ij} , we have

$$\frac{\partial \log g_{ij}(v)}{\partial v} < 0, \quad (\text{C.13})$$

$$\frac{\partial^2 \log g_{ij}(v)}{\partial v^2} > 0, \quad (\text{C.14})$$

for all $v > 0$, which makes the $\log g_{ij}$ monotonically decreasing convex functions.

Appendix D: Inverse-Gamma Densities and the Nonlinearities in Equation 2.16

We here show that the functions g_{ij} in equation 2.16 are obtained from inverse-gamma distributed u_{ij} with the PDFs in equation 2.17. Inserting the PDF p_{ii} from equation 2.17 into equation 2.15 yields $g_{ii}(s_i^2)$,

$$g_{ii}(s_i^2) = \int_0^\infty \sqrt{u_{ii}} \exp\left(-\frac{1}{2}s_i^2 u_{ii}\right) p_{ii}(u_{ii}) du_{ii} \quad (\text{D.1})$$

$$= \int_0^\infty \sqrt{u_{ii}} \exp\left(-\frac{1}{2}s_i^2 u_{ii}\right) \frac{m_{ii}^2}{2} u_{ii}^{-2} \exp\left(-\frac{m_{ii}^2}{2u_{ii}}\right) du_{ii} \quad (\text{D.2})$$

$$= \frac{m_{ii}^2}{2} \int_0^\infty u_{ii}^{-3/2} \exp\left(-\frac{1}{2}\left[s_i^2 u_{ii} + \frac{m_{ii}^2}{u_{ii}}\right]\right) du_{ii}, \quad (\text{D.3})$$

and the PDF p_{ij} yields $g_{ij}((s_i - s_j)^2)$:

$$g_{ij}((s_i - s_j)^2) = \int_0^\infty \exp\left(-\frac{1}{2}(s_i - s_j)^2 u_{ij}\right) p_{ij}(u_{ij}) du_{ij} \quad (\text{D.4})$$

$$= \int_0^\infty \exp\left(-\frac{1}{2}(s_i - s_j)^2 u_{ij}\right) \frac{m_{ij}}{\sqrt{2\pi}} u_{ij}^{-3/2} \exp\left(-\frac{m_{ij}^2}{2u_{ij}}\right) du_{ij} \quad (\text{D.5})$$

$$= \frac{m_{ij}}{\sqrt{2\pi}} \int_0^\infty u_{ij}^{-3/2} \exp\left(-\frac{1}{2} \left[(s_i - s_j)^2 u_{ij} + \frac{m_{ij}^2}{u_{ij}} \right]\right) du_{ij}. \tag{D.6}$$

It can be seen that both g_{ii} and g_{ij} are defined in terms of the integral

$$\int_0^\infty x^{-3/2} \exp\left(-\frac{1}{2} \left[\beta^2 x + \frac{\alpha^2}{x} \right]\right) dx,$$

where α corresponds to $m_{ii} > 0$ or $m_{ij} > 0$, and β to $|s_i|$ or $|s_i - s_j|$. The integral can be solved in closed form (Sasaki et al., 2013, equation 47)

$$\int_0^\infty x^{-3/2} \exp\left(-\frac{1}{2} \left[\beta^2 x + \frac{\alpha^2}{x} \right]\right) dx = \left(\frac{2\pi}{\alpha^2}\right)^{1/2} \exp(-|\alpha\beta|), \tag{D.7}$$

so that

$$g_{ii}(s_i^2) \propto \exp(-m_{ii}|s_i|) \tag{D.8}$$

$$g_{ij}((s_i - s_j)^2) \propto \exp(-m_{ij}|s_i - s_j|), \tag{D.9}$$

or, with $v > 0$,

$$\log g_{ii}(v) = -m_{ii}\sqrt{v} + \text{const}, \tag{D.10}$$

$$\log g_{ij}(v) = -m_{ij}\sqrt{v} + \text{const}, \tag{D.11}$$

as claimed in the main text.

Appendix E: A General Generative Model for Precision Matrices _____

We here extend the generative model for precision matrices, equation 2.2, to the more general form

$$\Lambda = \Omega \circ \Theta, \tag{E.1}$$

where \circ denotes the Hadamard product or element-wise multiplication. The matrix Ω is symmetric and its elements are nonnegative random variables, Θ is a d by d deterministic symmetric matrix whose diagonal elements θ_{ii} are positive and whose off-diagonal elements $\theta_{ij}, i \neq j$ take values in $\{-1, 0, 1\}$. The matrix Θ determines whether two components are positively, not at all, or negatively conditionally correlated. The matrix Ω scales the variances and correlations randomly and is defined as

$$\omega_{ij} = \begin{cases} \sum_{k=1}^d u_{ik}, & i = j, \\ u_{ij}, & i \neq j, \end{cases} \quad (\text{E.2})$$

where the $u_{ij} = u_{ji}$ are symmetric nonnegative random variables, with $u_{ii} > 0$. The model, equation E.1, generalizes the model, equation 2.2, of Λ in the main text, which is recovered for $\theta_{ii} = 1$ and $\theta_{ij} = -1$.

From the discussion of equation 2.2, it follows that Ω is always a strictly diagonally dominant matrix. We now show that Λ defined in equation E.1 has the same property.

Proposition 1: For Ω as in equation E.2, if $\theta_{ii} \geq |\theta_{ij}|$ for all i and j , Λ is a strictly diagonally dominant matrix.

Proof: It is sufficient to prove that for all i , $\lambda_{ii} > \sum_{j \neq i} |\lambda_{ij}|$. We compute

$$\begin{aligned} \lambda_{ii} - \sum_{\substack{j=1 \\ j \neq i}}^d |\lambda_{ij}| &= \theta_{ii} \omega_{ii} - \sum_{\substack{j=1 \\ j \neq i}}^d |\theta_{ij}| \omega_{ij}, \\ &= \theta_{ii} \sum_{j=1}^d u_{ij} - \sum_{\substack{j=1 \\ j \neq i}}^d |\theta_{ij}| u_{ij}, \\ &= \theta_{ii} u_{ii} + \sum_{\substack{j=1 \\ j \neq i}}^d (\theta_{ii} - |\theta_{ij}|) u_{ij}. \end{aligned}$$

Since $\theta_{ii} > 0$, $u_{ii} > 0$, and $u_{ij} \geq 0$, the proposition follows. \square

Since Λ is symmetric and strictly diagonally dominant, it is also invertible and positive definite (Horn & Johnson, 1985, theorem 6.1.10).

Following the same procedure as in section 2.2, we can derive the following approximation of the PDF of the dependent nongaussian components,

$$\tilde{p}_s(\mathbf{s}; \mathbf{m}, \Theta) \propto \prod_{i=1}^d \exp \left(-\theta_{ii} m_{ii} |s_i| - \sum_{j>i} m_{ij} \sqrt{\theta_{ii} s_i^2 + \theta_{jj} s_j^2 + 2\theta_{ij} s_i s_j} \right). \quad (\text{E.3})$$

For $\theta_{ii} = 1$ and $\theta_{ij} = -1$, we recover the model in equation 2.18. However, estimation of the parameters is much more difficult because equation E.3 no longer belongs to the exponentially family and the θ_{ij} are discrete variables.

Acknowledgments

H. Sasaki carried out most of the work when he was a JSPS fellow at the University of Electro-Communications. The work was partly done when M. U. Gutmann was with the Department of Mathematics and Statistics, University of Helsinki, and supported by the Finnish Centre of Excellence in Computational Inference Research COIN. H. Shouno was partly supported by MEXT/KAKENHI JSPS KAKENHI grant 26120515. A. Hyvärinen was supported by the Academy of Finland (Centre of Excellence in Inverse problems).

References

- Amari, S., Cichocki, A., & Yang, H. (1996). A new learning algorithm for blind signal separation. In D. Touretzky, M. Mozer, & M. Hasselmo (Eds.), *Advances in neural information processing systems, 8* (pp. 757–763). Cambridge, MA: MIT Press.
- Bach, F., & Jordan, M. (2003). Beyond independent components: Trees and clusters. *Journal of Machine Learning Research, 4*, 1205–1233.
- Ballé, J., Laparra, V., & Simoncelli, E. P. (2015). *Density modeling of images using a generalized normalization transformation*. arXiv:1511.06281
- Banerjee, O., El Ghaoui, L., & d'Aspremont, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. *Journal of Machine Learning Research, 9*, 485–516.
- Bartlett, M., Movellan, J., & Sejnowski, T. (2002). Face recognition by independent component analysis. *IEEE Transactions on Neural Networks, 13*(6), 1450–1464.
- Bollobás, B. (1998). *Modern graph theory*. New York: Springer.
- Campi, C., Parkkonen, L., Hari, R., & Hyvärinen, A. (2013). Non-linear canonical correlation for joint analysis of MEG signals from two subjects. *Frontiers in Neuroscience, 7*(107).
- Cardoso, J. (1998). Multidimensional independent component analysis. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing* (vol. 4, pp. 1941–1944). Piscataway, NJ: IEEE.
- Coen-Cagli, R., Dayan, P., & Schwartz, O. (2012). Cortical surround interactions and perceptual salience via natural scene statistics. *PLoS Computational Biology, 8*(3), e1002405.
- Comon, P. (1994). Independent component analysis, a new concept? *Signal Processing, 36*(3), 287–314.
- Friedman, J., Hastie, T., & Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics, 9*(3), 432–441.
- Gómez-Herrero, G., Atienza, M., Egiazarian, K., & Cantero, J. (2008). Measuring directional coupling between EEG sources. *NeuroImage, 43*(3), 497–508.
- Gutmann, M., & Hirayama, J. (2011). Bregman divergence as general framework to estimate unnormalized statistical models. In *Proc. Conf. on Uncertainty in Artificial Intelligence* (pp. 283–290). Corvallis, OR: AUAI Press.

- Gutmann, M., & Hyvärinen, A. (2011). Extracting coactivated features from multiple data sets. In *Proc. Int. Conf. on Artificial Neural Networks* (vol. 6791, pp. 323–330). Berlin: Springer.
- Gutmann, M., & Hyvärinen, A. (2012). Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13, 307–361.
- Gutmann, M., & Hyvärinen, A. (2013a). Estimation of unnormalized statistical models without numerical integration. In *Proc Workshop on Information Theoretic Methods in Science and Engineering*. <http://www.me.inf.kyushu-u.ac.jp/witmse2013/proceeding.html>
- Gutmann, M., & Hyvärinen, A. (2013b). A three-layer model of natural image statistics. *Journal of Physiology-Paris*, 107(5), 369–398.
- Gutmann, M., Laparra, V., Hyvärinen, A., & Malo, J. (2014). Spatio-chromatic adaptation via higher-order canonical correlation analysis of natural images. *PLoS ONE*, 9(2), e86481.
- He, Y., Qi, Y., Kavukcuoglu, K., & Park, H. (2012). Learning the dependency structure of latent factors. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 25 (pp. 2366–2374). Red Hook, NY: Curran.
- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9(2), 181–197.
- Hinton, G. (2002). Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8), 1771–1800.
- Horn, R., & Johnson, C. (1985). *Matrix analysis*. Cambridge: Cambridge University Press.
- Hoyer, P., & Hyvärinen, A. (2002). A multi-layer sparse coding network learns contour coding from natural images. *Vision Research*, 42(12), 1593–1605.
- Hurri, J., & Hyvärinen, A. (2003). Temporal and spatiotemporal coherence in simple-cell responses: A generative model of natural image sequences. *Network: Computation in Neural Systems*, 14(3), 527–551.
- Hyvärinen, A. (2005). Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6, 695–709.
- Hyvärinen, A. (2007). Some extensions of score matching. *Computational Statistics and Data Analysis*, 51(5), 2499–2512.
- Hyvärinen, A., Gutmann, M., & Hoyer, P. (2005). Statistical model of natural stimuli predicts edge-like pooling of spatial frequency channels in V2. *BMC Neuroscience*, 6, 12.
- Hyvärinen, A., & Hoyer, P. (2000). Emergence of phase- and shift-invariant features by decomposition of natural images into independent feature subspaces. *Neural Computation*, 12(7), 1705–1720.
- Hyvärinen, A., Hoyer, P., & Inki, M. (2001). Topographic independent component analysis. *Neural Computation*, 13(7), 1527–1558.
- Hyvärinen, A., Hurri, J., & Hoyer, P. (2009). *Natural image statistics*. New York: Springer.
- Hyvärinen, A., & Oja, E. (2000). Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4–5), 411–430.

- International Phonetic Association. (1999). *Handbook of the international phonetic association: A guide to the use of the international phonetic alphabet*. Cambridge: Cambridge University Press.
- Karklin, Y., & Lewicki, M. (2005). A hierarchical Bayesian model for learning nonlinear statistical regularities in nonstationary natural signals. *Neural Computation*, 17(2), 397–423.
- Klein, D. J., König, P., & Körding, K. P. (2003). Sparse spectrotemporal coding of sounds. *EURASIP Journal on Applied Signal Processing*, 2003, 659–667.
- Kohonen, T. (1995). *Self-organizing maps*. Berlin: Springer-Verlag.
- Kohonen, T. (1996). Emergence of invariant-feature detectors in the adaptive-subspace self-organizing map. *Biological Cybernetics*, 75(4), 281–291.
- Köster, U., & Hyvärinen, A. (2010). A two-layer model of natural stimuli estimated with score matching. *Neural Computation*, 22(9), 2308–2333.
- Mairal, J., Jenatton, R., Obozinski, G., & Bach, F. (2011). Convex and network flow optimization for structured sparsity. *Journal of Machine Learning Research*, 12, 2681–2720.
- Olshausen, B., & Field, D. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- Osindero, S., Welling, M., & Hinton, G. (2006). Topographic product models applied to natural scene statistics. *Neural Computation*, 18(2), 381–414.
- Ostrowski, A. (1937). Sur la détermination des bornes inférieures pour une classe des déterminants. *Bull. Sci. Math*, 61(2), 19–32.
- Pihlaja, M., Gutmann, M., & Hyvärinen, A. (2010). A family of computationally efficient and simple estimators for unnormalized statistical models. In *Proc. Conf. on Uncertainty in Artificial Intelligence* (pp. 442–449). Corvallis, OR: AUAI Press.
- Sasaki, H., Gutmann, M., Shouno, H., & Hyvärinen, A. (2013). Correlated topographic analysis: Estimating an ordering of correlated components. *Machine Learning*, 92(2–3), 285–317.
- Sasaki, H., Gutmann, M., Shouno, H., & Hyvärinen, A. (2014). Estimating dependency structures for non-Gaussian components with linear and energy correlations. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, JMLR: W&CP (vol. 33, pp. 868–876).
- Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nature Neuroscience*, 4(8), 819–825.
- Shimizu, S., Hoyer, P., Hyvärinen, A., & Kerminen, A. (2006). A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7, 2003–2030.
- Simoncelli, E. (1999). Modeling the joint statistics of images in the wavelet domain. In *Proc. SPIE, 44th Annual Meeting* (vol. 3813, pp. 188–195). Bellingham, WA: SPIE.
- Terashima, H., & Hosoya, H. (2009). Sparse codes of harmonic natural sounds and their modulatory interactions. *Network: Computation in Neural Systems*, 20(4), 253–267.
- Terashima, H., Hosoya, H., Tani, T., Ichinohe, N., & Okada, M. (2013). Sparse coding of harmonic vocalization in monkey auditory cortex. *Neurocomputing*, 103, 14–21.

- Terashima, H., & Okada, M. (2012). The topographic unsupervised learning of natural sounds in the auditory cortex. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems*, 35 (pp. 2321–2329). Red Hook, NY: Curran.
- Theis, F. (2005). Blind signal separation into groups of dependent signals using joint block diagonalization. In *IEEE International Symposium on Circuits and Systems, 2005* (pp. 5878–5881). Piscataway, NJ: IEEE.
- Vigário, R., Särelä, J., Jousmäki, V., Hämäläinen, M., & Oja, E. (2000). Independent component approach to the analysis of EEG and MEG recordings. *IEEE Transactions on Biomedical Engineering*, 47(5), 589–593.
- Zoran, D., & Weiss, Y. (2009). The “tree-dependent components” of natural images are edge filters. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems*, 22 (pp. 2340–2348). Red Hook, NY: Curran.

Received March 7, 2017; accepted May 26, 2017.