HAL
open science

A Riemannian approach to blind separation of t-distributed sources

Florent Bouchard, Arnaud Breloy, Guillaume Ginolhac, Alexandre Renaux

To cite this version:

Florent Bouchard, Arnaud Breloy, Guillaume Ginolhac, Alexandre Renaux. A Riemannian approach to blind separation of t-distributed sources. 2020 28th European Signal Processing Conference (EUSIPCO), Jan 2021, Amsterdam, Netherlands. 10.23919/eusipco47968.2020.9287783 . hal-02988356

HAL Id: hal-02988356

https://hal.univ-grenoble-alpes.fr/hal-02988356

Submitted on 4 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Riemannian approach to blind separation of t-distributed sources

Florent Bouchard*, Arnaud Breloy†, Guillaume Ginolhac*, Alexandre Renaux‡
*LISTIC, Univ. Savoie Mont Blanc, France, †LEME, Univ. Paris Nanterre, France, ‡ L2S, Univ. Paris Sud, France
Email: florent.bouchard@univ-smb.fr

***Abstract*—The blind source separation problem is considered, more specifically the approach based on non-stationarity and coloration. In both cases, sources are usually assumed to be Gaussian. In this paper, we extend previous works in order to handle sources drawn from the multivariate Student t-distribution. After studying the data model in this case, a new blind source separation criterion based on the log-likelihood of the considered distribution is proposed. To solve the resulting optimization problem, Riemannian optimization on the parameter manifold is leveraged. The performance of the proposed method is illustrated on simulated data.**

***Index Terms*—blind source separation, Student t-distribution, Riemannian optimization**

I. Introduction

Blind source separation (BSS) is a tremendous tool for signal processing and data analysis with applications in many engineering fields such as radar, communications, image processing and biomedical signal analysis; see [1] for a complete overview. In this work, we consider the determined linear instantaneous mixing model

$$\boldsymbol{x} = \boldsymbol{A}\boldsymbol{s}, \tag{1}$$

where $\boldsymbol{x} \in \mathbb{R}^n$ corresponds to the observations, $\boldsymbol{s} \in \mathbb{R}^n$ contains the sources and $\boldsymbol{A} \in \mathrm{GL}_n$ ($n \times n$ non-singular matrices) is the mixing matrix. Given some observations of $\boldsymbol{x}$, the goal is to estimate the sources $\boldsymbol{s}$ and the mixing process $\boldsymbol{A}$.

Up to now, two different types of methods have been proposed to perform blind source separation with model (1) depending on the data at hand. Historically, in the original formulation of the problem, which corresponds to independent component analysis (ICA), sources are assumed to be independent and identically distributed (i.i.d.) [2]–[4]. To ensure separability, sources have to be non-Gaussian and higher order statistics must be employed. On the other hand, if one wants to be able to separate Gaussian sources, the i.i.d. assumption must be removed, *i.e.*, some structure of the sources is exploited such as non-stationarity or coloration [5], [6].

In this paper, we are interested in the second approach, *i.e.*, when non-stationarity or coloration is exploited to perform blind source separation. In this case, one considers T realizations of K Gaussian random variables $\boldsymbol{x}_k = \boldsymbol{A}\boldsymbol{s}_k$ with covariance matrices $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T \in \mathcal{S}_n^{++}$ ($n \times n$ symmetric positive definite matrices), where the $\boldsymbol{s}_k$ are the random variables corresponding to the sources and whose covariance matrices are $\boldsymbol{\Lambda}_k \in \mathcal{D}_n^{++}$ ($n \times n$ diagonal positive definite matrices). In practice, the K different sets of observations $\{\boldsymbol{x}_k(t)\}$ can for instance correspond to different epochs (non-stationarity) or different frequencies (coloration) [1], [6], [7]. To solve the problem in the Gaussian case, one usually performs a two-steps procedure. First, the sample covariance matrices $\boldsymbol{C}_k$ of the observations $\{\boldsymbol{x}_k(t)\}$ are computed. Then, matrices $\boldsymbol{C}_k$ are jointly diagonalized with a matrix $\boldsymbol{B} \in \mathrm{GL}_n$, solution to the optimization problem

$$\underset{\boldsymbol{B}\in\mathrm{GL}_n}{\operatorname{argmin}} \quad \sum_k f(\boldsymbol{B}, \boldsymbol{C}_k), \tag{2}$$

where f is a diagonality criterion of $\boldsymbol{B}\boldsymbol{C}_k\boldsymbol{B}^T$. Estimates $\widehat{\boldsymbol{s}}$ and $\widehat{\boldsymbol{A}}$ of sources $\boldsymbol{s}$ and mixing matrix $\boldsymbol{A}$ are finally given by $\widehat{\boldsymbol{s}} = \boldsymbol{B}\boldsymbol{x}$ and $\widehat{\boldsymbol{A}} = \boldsymbol{B}^{-1}$.

Previous research on this topic have mainly focused on two aspects: the choice of the diagonality criterion f in (2) and the conception of iterative methods to solve (2) given a criterion f. Concerning the diagonality criterion f, most articles employ either the least squares one proposed in [4] or the one based on the log-likelihood introduced in [6]. Recently, different criteria based on several divergences and distances on $\mathcal{S}_n^{++}$ have been proposed in [8], [9]. Many algorithms to solve (2) for various diagonality criteria f have been developed; see *e.g.*, [4], [6], [9]–[12]. In particular, methods based on Riemannian optimization can be found in [9], [10], [12].

Works based on non-stationarity and coloration consider that sources are Gaussian. In this paper, our main contribution is to extend this approach to handle sources drawn from the multivariate Student t-distribution [13]. Since it has a heavier tail than the Gaussian distribution, it is more adapted to account for outliers. To do so, we derive a new blind source separation criterion based on the log-likelihood of the data model. To find a minimizer of this criterion, we propose to exploit Riemannian optimization on the parameter manifold, which allows to find estimates $\widehat{\boldsymbol{A}} \in \mathrm{GL}_n$ and $\{\widehat{\boldsymbol{\Lambda}}_k\} \in (\mathcal{D}_n^{++})^K$ of the true parameters $\boldsymbol{A}$ and $\{\boldsymbol{\Lambda}_k\}$.

This paper contains four sections including this introduction. In section II, the data model is defined along with the manifold holding the parameters of the considered distribution. To handle the diagonal scaling ambiguity of blind source separation, an original constraint is proposed. Finally, the Riemannian geometry of the resulting manifold is studied. In section III, a Riemannian optimization framework is developed in order to be able to solve optimization problems on the parameter manifold of interest. A blind source separation criterion based on the log-likelihood function of the considered

distribution is then derived, thus yielding an original blind source separation method. In section IV, the performance of the proposed method are illustrated on simulated data. As expected, when sources are drawn from the multivariate Student t-distribution, the proposed method outperforms the state of the art method [6].

II. Model

To design a new blind source separation method adapted to the multivariate Student t-distribution, we first need to clearly establish the model of the data along with the parameters they depend on. In section II-A, the probability density function of the distribution of the data and the manifold holding the parameters of interest are defined. In section II-B, as the approach chosen to solve the blind source separation problem is based on Riemannian optimization, the Riemannian geometry of the parameter manifold is studied.

A. Data distribution and parameter manifold

We consider K sets of T observations $\{\boldsymbol{x}_k(t)\}$ of the random variables $\boldsymbol{x}_k$ in $\mathbb{R}^n$ following the elliptical distibutions with density generator function $g : \mathbb{R}^+ \to \mathbb{R}^+$ and covariance matrices $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$, where $\boldsymbol{A} \in \mathrm{GL}_n$ and $\boldsymbol{\Lambda}_k \in \mathcal{D}_n^{++}$. The probability density function f of $\{\boldsymbol{x}_k(t)\}$ is

$$f(\{\boldsymbol{x}_k(t)\}|\boldsymbol{A}, \{\boldsymbol{\Lambda}_k\}) = \prod_k f_d(\{\boldsymbol{x}_k(t)\}|\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T), \quad (3)$$

where f_d is the probability density function of the multivariate Student t-distribution with $d \in \mathbb{N}^*$ degrees of freedom. Given $\{\boldsymbol{x}(t)\}$ with covariance matrix $\boldsymbol{C}$, we have, up to a factor,

$$f_d(\{\boldsymbol{x}(t)\}|\boldsymbol{C}) = \prod_t \det(\boldsymbol{C})^{-1/2}\left(1 + \frac{1}{d}\boldsymbol{x}(t)^T\boldsymbol{C}^{-1}\boldsymbol{x}(t)\right)^{-(d+n)/2}. \quad (4)$$

The set of parameters $(\boldsymbol{A}, \boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_K)$ of the random variables $\boldsymbol{x}_k$ lies in the manifold $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. The parameters $(\boldsymbol{A}, \boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_K)$ of the $\boldsymbol{x}_k$ are not unique: given any permutation matrix $\boldsymbol{P}$ and non-singular diagonal matrix $\boldsymbol{\Sigma}$, we have

$$f(\{\boldsymbol{x}_k(t)\}|\boldsymbol{A}, \{\boldsymbol{\Lambda}_k\}) = f(\{\boldsymbol{x}_k(t)\}|\boldsymbol{A}\boldsymbol{P}\boldsymbol{\Sigma}, \{\boldsymbol{\Sigma}^{-1}\boldsymbol{P}^T\boldsymbol{\Lambda}_k\boldsymbol{P}\boldsymbol{\Sigma}^{-1}\}).$$

To deal with the diagonal scaling ambiguity resulting from the invariance through the action of $\boldsymbol{\Sigma}$, an additional constraint can be imposed. In this work, we propose a new constraint, related to the one introduced in [11]. It is given by

$$\sum_k \boldsymbol{\Lambda}_k = \boldsymbol{I}_n, \quad (5)$$

where $\boldsymbol{I}_n$ is the identity matrix. As this constraint is smooth with respect to $(\boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_K)$, it can be integrated in a smooth manifold. The permutation ambiguity resulting from the invariance through the action of $\boldsymbol{P}$ is more tricky as it cannot be handled in a manifold. However, contrary to the diagonal scaling ambiguity, it is not an issue from a numerical point of view as it does not yield algorithms to converge to degenerate solutions. As most works on blind source separation, we choose to ignore it. It follows that the manifold holding the parameters of the distribution of interest is chosen as

$$\mathcal{M} = \{(\boldsymbol{A}, \{\boldsymbol{\Lambda}_k\}) \in \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K : \textstyle\sum_k \boldsymbol{\Lambda}_k = \boldsymbol{I}_n\}, \quad (6)$$

which is a submanifold of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$.

B. Riemannian geometry of the parameter manifold

In the following, $\theta = (\boldsymbol{A}, \{\boldsymbol{\Lambda}_k\})$, $\xi = (\boldsymbol{\xi}_{\boldsymbol{A}}, \{\boldsymbol{\xi}_k\})$ and $\eta = (\boldsymbol{\eta}_{\boldsymbol{A}}, \{\boldsymbol{\eta}_k\})$. Since GL_n and $\mathcal{D}_n^{++}$ are open in $\mathbb{R}^{n\times n}$ and $\mathcal{D}_n$ ($n \times n$ diagonal matrices), respectively, the tangent space $T_\theta\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ can be identified as $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$. From a statistical point of view, one ideally wants to choose the Fisher information metric associated with (3) as the Riemannian metric of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. However, as illustrated in [14] for Gaussian sources, the Fisher information metric appears complicated to deal with in this case and a different Riemannian metric is prefered. In this work, we choose

$$\langle \xi, \eta\rangle_\theta = \mathrm{tr}((\boldsymbol{\xi}_{\boldsymbol{A}}\boldsymbol{A}^{-1})^T\boldsymbol{\eta}_{\boldsymbol{A}}\boldsymbol{A}^{-1}) + \sum_k \mathrm{tr}(\boldsymbol{\Lambda}_k^{-2}\boldsymbol{\xi}_k\boldsymbol{\eta}_k). \quad (7)$$

The component that depends on $\boldsymbol{A}$ corresponds to the so-called right-invariant metric on GL_n [12], [15]. The component that depends on $\boldsymbol{\Lambda}_k$ correponds to the popular affine-invariant metric on $\mathcal{D}_n^{++}$ [16]. The main interest of the chosen metric on $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ is that it is invariant with respect to the actions of diagonal and permutation matrices $\boldsymbol{\Sigma}$ and $\boldsymbol{P}$ described in the previous section.

The manifold $\mathcal{M}$ defined in (6) that holds the parameters of interest is a submanifold of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. Thus, the tangent space $T_\theta\mathcal{M}$ at $\theta \in \mathcal{M}$ is a subspace of $T_\theta\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. Furthermore, $\mathcal{M}$ can inherit the Riemannian metric (7) of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. Given $\theta \in \mathcal{M}$, the tangent space $T_\theta\mathcal{M}$ along with the orthogonal projection map $P_\theta : \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K \to T_\theta\mathcal{M}$ according to (7) are given in Proposition 1.

Proposition 1. *The tangent space $T_\theta\mathcal{M}$ at $\theta \in \mathcal{M}$ is*

$$T_\theta\mathcal{M} = \{\xi \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K : \textstyle\sum_k \boldsymbol{\xi}_k = \boldsymbol{0}\}.$$

The orthogonal projection map $P_\theta : \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K \to T_\theta\mathcal{M}$ according to (7) *is given by*

$$P_\theta(\xi) = (\boldsymbol{\xi}_{\boldsymbol{A}}, \{\boldsymbol{\xi}_k - (\textstyle\sum_\ell \boldsymbol{\Lambda}_\ell^2)^{-1}\sum_\ell \boldsymbol{\xi}_\ell\boldsymbol{\Lambda}_k^2\}).$$

Proof. The manifold $\mathcal{M}$ of interest is defined as $\mathcal{M} = \{\theta \in \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K : \varphi(\theta) = \boldsymbol{0}\}$, where $\varphi : \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K \to \mathbb{R}^{n\times n}$ is the smooth mapping such that $\varphi(\theta) = \sum_k \boldsymbol{\Lambda}_k - \boldsymbol{I}_n$. It follows that the tangent space $T_\theta\mathcal{M}$ at $\theta \in \mathcal{M}$ is $T_\theta\mathcal{M} = \{\xi \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K : \mathrm{D}\,\varphi(\theta)[\xi] = \boldsymbol{0}\}$, where $\mathrm{D}\,\varphi(\theta)[\xi]$ is the derivative of φ at θ in direction ξ. Noticing that $\mathrm{D}\,\varphi(\theta)[\xi] = \sum_k \boldsymbol{\xi}_k$ yields the result.

The normal space $N_\theta\mathcal{M}$ at $\theta \in \mathcal{M}$ is defined as $N_\theta\mathcal{M} = \{\eta \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K : \forall\xi \in T_\theta\mathcal{M}, \langle\xi, \eta\rangle_\theta = 0\}$. One can check that, in our case, it is given by

$$N_\theta\mathcal{M} = \{(\boldsymbol{0}, \{\boldsymbol{\Lambda}_k^2\boldsymbol{\eta}\}) : \boldsymbol{\eta} \in \mathcal{D}_n\}.$$

It follows that $P_\theta(\xi) = \xi - (\mathbf{0}, \{\mathbf{\Lambda}_k^2 \boldsymbol{\eta}\})$, where $\boldsymbol{\eta}$ is the diagonal matrix such that $P_\theta(\xi) \in T_\theta\mathcal{M}$, *i.e.*, $\sum_k (\boldsymbol{\xi_k} - \mathbf{\Lambda}_k^2 \boldsymbol{\eta}) = \mathbf{0}$. We thus obtain $\boldsymbol{\eta} = (\sum_\ell \mathbf{\Lambda}_\ell^2)^{-1} \sum_\ell \boldsymbol{\xi}_\ell$. ■

These elements of Riemannian geometry are sufficient to build the Riemannian optimization framework on $\mathcal{M}$ required for blind source separation in section III.

III. Blind source separation method

The blind source separation problem can be written as an optimization problem on the manifold $\mathcal{M}$. In section III-A, a Riemannian optimization framework on $\mathcal{M}$ is built. In section III-B, a blind source separation based on the likelihood of the multivariate Student t-distribution is proposed and the elements required for Riemannian optimization on $\mathcal{M}$ are provided.

A. Riemannian optimization on the parameter manifold

Let $f : \mathcal{M} \to \mathbb{R}$ be an objective function defined on $\mathcal{M}$. To be able to perform Riemannian optimization [17] of f on $\mathcal{M}$, it is needed to have the Riemannian gradient of f along with a retraction on $\mathcal{M}$, which, given $\theta \in \mathcal{M}$, maps the elements of $T_\theta\mathcal{M}$ back onto $\mathcal{M}$. Concerning the Riemannian gradient on $\mathcal{M}$, Propositon 2 shows that it can be obtained from the Euclidean gradient of f, which is defined on the ambient space $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$.

Proposition 2. *The Riemannian gradient* $\operatorname{grad}_{\mathcal{M}} f(\theta) \in T_\theta\mathcal{M}$ *of f at $\theta \in \mathcal{M}$ according to metric* (7) *is given by*

$$\operatorname{grad}_{\mathcal{M}} f(\theta) = P_\theta \left(\nabla f_{\boldsymbol{A}} \boldsymbol{A}^T \boldsymbol{A}, \{\mathbf{\Lambda}_k^2 \nabla f_k\} \right),$$

where $P_\theta : \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K \to T_\theta\mathcal{M}$ is defined in Proposition 1 and $\nabla f(\theta) = (\nabla f_{\boldsymbol{A}}, \{\nabla f_k\}) \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$ is the Euclidean gradient of f.

Proof. First recall that $\nabla f(\theta)$ is the only element of $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$ such that for all $\xi \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$,

$$\mathrm{D}\, f(\theta)[\xi] = \langle \nabla f(\theta), \xi \rangle^{\mathcal{E}},$$

where $\langle \cdot, \cdot \rangle^{\mathcal{E}}$ is the Euclidean metric on $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$, which is defined as

$$\langle \xi, \eta \rangle^{\mathcal{E}} = \operatorname{tr}(\boldsymbol{\xi}_{\boldsymbol{A}}^T \boldsymbol{\eta}_{\boldsymbol{A}}) + \sum_k \operatorname{tr}(\boldsymbol{\xi}_k \boldsymbol{\eta}_k).$$

Similarly, the Riemannian gradient of f on $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ according to the Riemannian metric (7) is the only element of $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$ such that, for all $\xi \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$,

$$\mathrm{D}\, f(\theta)[\xi] = \langle \operatorname{grad}_{\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K} f(\theta), \xi \rangle_\theta,$$

From $\langle \operatorname{grad}_{\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K} f(\theta), \xi \rangle_\theta = \langle \nabla f(\theta), \xi \rangle^{\mathcal{E}}$, one can check that

$$\operatorname{grad}_{\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K} f(\theta) = \left(\nabla f_{\boldsymbol{A}} \boldsymbol{A}^T \boldsymbol{A}, \{\mathbf{\Lambda}_k^2 \nabla f_k\} \right).$$

The rest of the proof follows from [17, Equation (3.37)], which indicates that $\operatorname{grad}_{\mathcal{M}} f(\theta)$ is obtained by projecting $\operatorname{grad}_{\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K} f(\theta)$ onto $T_\theta\mathcal{M}$. ■

The Riemannian gradient $\operatorname{grad}_{\mathcal{M}} f(\theta) \in T_\theta\mathcal{M}$ of f at $\theta \in \mathcal{M}$ allows to define a descent direction ξ of f in the tangent space $T_\theta\mathcal{M}$. In order to achieve a new point on the manifold $\mathcal{M}$ from ξ, a retraction R on $\mathcal{M}$ is needed. Given $\theta \in \mathcal{M}$, the retraction R_θ at θ is the mapping from $T_\theta\mathcal{M}$ onto $\mathcal{M}$ such that:

1) $R_\theta(0_\theta) = \theta$, where 0_θ is the zero element of $T_\theta\mathcal{M}$.
2) $\mathrm{D}\, R_\theta(0_\theta)[\xi] = \xi$.

Proposition 3 contains a proper retraction on $\mathcal{M}$ obtained by projecting a second order approximation of the Riemannian exponential map of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ onto $\mathcal{M}$.

Proposition 3. *Given $\theta \in \mathcal{M}$, a proper retraction $R_\theta : T_\theta\mathcal{M} \to \mathcal{M}$ at θ is given, for all $\xi \in T_\theta\mathcal{M}$ by*

$$R_\theta(\xi) = \Gamma(\overline{R}_\theta(\xi)),$$

where $\overline{R}_\theta$ is a second order approximation of the Riemannian exponential map of $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ equiped with (7)*, which is given by*

$$\overline{R}_\theta(\xi) = \theta + \xi + \tfrac{1}{2}\left(\boldsymbol{\xi}_{\boldsymbol{A}} \boldsymbol{A}^{-1} \boldsymbol{\xi}_{\boldsymbol{A}} + \boldsymbol{\xi}_{\boldsymbol{A}} (\boldsymbol{A}^T \boldsymbol{A})^{-1} \boldsymbol{\xi}_{\boldsymbol{A}}^T \boldsymbol{A} - \boldsymbol{A}^{-T} \boldsymbol{\xi}_{\boldsymbol{A}}^T \boldsymbol{\xi}_{\boldsymbol{A}}, \{\mathbf{\Lambda}_k^{-1} \boldsymbol{\xi}_k^2\}\right),$$

and $\Gamma : \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K \to \mathcal{M}$ is the projection map defined, for all $\theta \in \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$, as

$$\Gamma(\theta) = (\boldsymbol{A}, \{(\textstyle\sum_\ell \Lambda_\ell)^{-1} \Lambda_k\}).$$

Proof. Since $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ equiped with metric (7) is a Riemannian product manifold, we know from [17] that the Riemannian exponential map is defined, for all $\theta \in \mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ and $\xi \in \mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$, as

$$\exp_\theta^{\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K}(\xi) = \left(\exp_{\boldsymbol{A}}^{\mathrm{GL}_n}(\boldsymbol{\xi}_{\boldsymbol{A}}), \{\exp_{\mathbf{\Lambda}_k}^{\mathcal{D}_n^{++}}(\boldsymbol{\xi}_k)\} \right),$$

where

$$\exp_{\boldsymbol{A}}^{\mathrm{GL}_n}(\boldsymbol{\xi}_{\boldsymbol{A}}) = \exp(\boldsymbol{\xi}_{\boldsymbol{A}} \boldsymbol{A}^{-1} - (\boldsymbol{\xi}_{\boldsymbol{A}} \boldsymbol{A}^{-1})^T) \exp(\boldsymbol{\xi}_{\boldsymbol{A}} \boldsymbol{A}^{-1})^T \boldsymbol{A}$$

is the Riemannian exponential map of GL_n equiped with the right-invariant metric [12], [15] and

$$\exp_{\mathbf{\Lambda}_k}^{\mathcal{D}_n^{++}}(\boldsymbol{\xi}_k) = \mathbf{\Lambda}_k \exp(\mathbf{\Lambda}_k^{-1} \boldsymbol{\xi}_k)$$

is the Riemannian exponential map of $\mathcal{D}_n^{++}$ equiped with the affine-invariant metric [16]. The matrix exponential admits the second order approximation

$$\exp(\boldsymbol{X}) = \boldsymbol{I}_n + \boldsymbol{X} + \frac{1}{2}\boldsymbol{X}^2 + o(\boldsymbol{X}^2).$$

By injecting this approximation in the formula of the Riemannian exponential map on $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$, one can show that the proposed operator $\overline{R}$ is a second order approximation of this Riemannian exponential map, which is also enough to prove that it is a retraction on $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$. Moreover, it is readily checked that Γ is a projection map from $\mathrm{GL}_n \times (\mathcal{D}_n^{++})^K$ onto $\mathcal{M}$. Therefore, from [18], we know that the proposed operator R is a proper retraction on $\mathcal{M}$. ■

The Riemannian gradient of $f : \mathcal{M} \to \mathbb{R}$ and the retraction

R on $\mathcal{M}$ are enough to define a Riemannian gradient descent algorithm on $\mathcal{M}$. Indeed, given iterate θ_i, the next iterate is obtained by

$$\theta_{i+1} = R_{\theta_i}(-t_i \operatorname{grad}_{\mathcal{M}} f(\theta_i)), \tag{8}$$

where t_i is the stepsize, which can for instance be computed through a linesearch; see [17] for details.

To be able to use more sophisticated Riemannian optimization algorithms such as conjugate gradient or BFGS, one further needs to define a vector transport operator $\mathcal{T}$, which allows to transport a tangent vector from one tangent space onto another [17]. Given $\theta \in \mathcal{M}$ and $\xi, \eta \in T_\theta\mathcal{M}$, the vector transport $\mathcal{T}(\theta, \xi, \eta)$ associated to the retraction R transports the tangent vector η in the tangent space $T_{R_\theta(\xi)}\mathcal{M}$ of $R_\theta(\xi)$. A generic solution [17] is to choose $\mathcal{T}(\theta, \xi, \eta)$ as

$$\mathcal{T}(\theta, \xi, \eta) = P_{R_\theta(\xi)}(\eta), \tag{9}$$

where, in our case, P is the projection map defined in Proposition 1 and R is the retraction defined in Proposition 3.

B. Application to the log-likelihood of the t-distribution

To perform the blind source separation of the K sets of T observations $\{\boldsymbol{x}_k(t)\}$ drawn from the distribution defined in Section II-A, the optimal solution is to employ the maximum likelihood estimator associated with the probability density function (3). It is equivalent to minimizing the negative log-likelihood of (3) on $\mathcal{M}$, which is given by

$$f(\theta) = \sum_k f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta)), \tag{10}$$

where $\psi_k : \mathcal{M} \to \mathcal{S}_n^{++}$ is such that $\psi_k(\theta) = \boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$ and f^{++} is the negative log-likelihood of the multivariate Student t-distribution with $d \in \mathbb{N}^*$ degrees of freedom. Given T observations $\{\boldsymbol{x}(t)\}$, the negative log-likelihood $f^{++}_{\boldsymbol{x}} : \mathcal{S}_n^{++} \to \mathbb{R}$ associated with (4) is defined, up to constants, as

$$f^{++}_{\boldsymbol{x}}(\boldsymbol{C}) = \frac{T}{2}\log\det(\boldsymbol{C}) + \frac{d+n}{2}\sum_t \log(d + \boldsymbol{x}(t)^T\boldsymbol{C}^{-1}\boldsymbol{x}(t)). \tag{11}$$

To be able to minimize (10) with a Riemannian optimization algorithm on $\mathcal{M}$ within the framework developed in secion III-A, it remains to provide the Euclidean gradient of f on $\mathbb{R}^{n\times n} \times (\mathcal{D}_n)^K$. This is achieved in Proposition 4.

Proposition 4. *The Euclidean gradient $\nabla f(\theta)$ of the blind source separation criterion f defined in* (10) *at $\theta \in \mathcal{M}$ is given by*

$$\nabla f(\theta) = \Big(2\textstyle\sum_k \nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta))\boldsymbol{A}\boldsymbol{\Lambda}_k, \{\operatorname{ddiag}(\boldsymbol{A}^T\nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta))\boldsymbol{A})\}\Big),$$

where $\operatorname{ddiag}(\cdot)$ *returns the diagonal part of its argument and $\nabla f^{++}_{\boldsymbol{x}}(\boldsymbol{C})$ is the Euclidean gradient of* (11) *at $\boldsymbol{C} \in \mathcal{S}_n^{++}$, which is*

$$\nabla f^{++}_{\boldsymbol{x}}(\boldsymbol{C}) = \boldsymbol{C}^{-1}\left[\frac{T}{2}\boldsymbol{C} - \frac{d+n}{2}\sum_t \frac{\boldsymbol{x}(t)\boldsymbol{x}(t)^T}{d + \boldsymbol{x}(t)^T\boldsymbol{C}^{-1}\boldsymbol{x}(t)}\right]\boldsymbol{C}^{-1}.$$

Proof. By definition of f,

$$\begin{aligned} \mathrm{D}\, f(\theta)[\xi] &= \textstyle\sum_k \mathrm{D}\, f_{\boldsymbol{x}_k}(\psi_k(\theta))[\mathrm{D}\,\psi_k(\theta)[\xi]] \\ &= \textstyle\sum_k \langle \nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta)), \mathrm{D}\,\psi_k(\theta)[\xi]\rangle^{\mathbb{R}^{n\times n}}, \end{aligned}$$

where $\mathrm{D}\,\psi_k(\theta)[\xi] = \boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{\xi}_{\boldsymbol{A}}^T + \boldsymbol{\xi}_{\boldsymbol{A}}\boldsymbol{\Lambda}_k\boldsymbol{A}^T + \boldsymbol{A}\boldsymbol{\xi}_k\boldsymbol{A}^T$ and $\langle\cdot,\cdot\rangle^{\mathbb{R}^{n\times n}}$ is the Euclidean inner product on $\mathbb{R}^{n\times n}$. Basic calculations yield

$$\begin{aligned} \mathrm{D}\, f(\theta)[\xi] &= \textstyle\sum_k 2\langle \nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta)), \boldsymbol{\xi}_{\boldsymbol{A}}\boldsymbol{\Lambda}_k\boldsymbol{A}^T\rangle^{\mathbb{R}^{n\times n}} \\ &+ \textstyle\sum_k \langle \nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta)), \boldsymbol{A}\boldsymbol{\xi}_k\boldsymbol{A}^T\rangle^{\mathbb{R}^{n\times n}} \\ &= \langle 2\textstyle\sum_k \nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta))\boldsymbol{A}\boldsymbol{\Lambda}_k, \boldsymbol{\xi}_{\boldsymbol{A}}\rangle^{\mathbb{R}^{n\times n}} \\ &+ \textstyle\sum_k \langle \operatorname{ddiag}(\boldsymbol{A}^T\nabla f^{++}_{\boldsymbol{x}_k}(\psi_k(\theta))\boldsymbol{A}), \boldsymbol{\xi}_k\rangle^{\mathbb{R}^{n\times n}}, \end{aligned}$$

where we used $\operatorname{tr}(\boldsymbol{X}\boldsymbol{\Sigma}) = \operatorname{tr}(\operatorname{ddiag}(\boldsymbol{X})\boldsymbol{\Sigma})$ for all $\boldsymbol{\Sigma} \in \mathcal{D}_n$. The result is finally obtained by identification. ■

IV. Numerical illustration

To simulate data, we generate $K = 30$ $n \times n$ (with $n = 10$) symmetric positive definite matrices $\boldsymbol{C}_k$ according to model

$$\boldsymbol{C}_k = \boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}, \tag{12}$$

where $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T$, with $\boldsymbol{U}$ and $\boldsymbol{V}$ random orthogonal matrices, and $\boldsymbol{\Sigma}$ random diagonal matrix whose minimal and maximal elements are $1/\sqrt{\alpha}$ and $\sqrt{\alpha}$, where $\alpha = 10$ is the condition number of $\boldsymbol{A}$. Matrices $\boldsymbol{\Lambda}_k \in \mathcal{D}_n^{++}$ correspond to the covariance matrices of the sources and have i.i.d. elements drawn from the chi squared distribution with expectation one. For $T \in \{15, 25, 50, 75, 100, 500, 1000\}$, 100 sets of T realizations $\{\boldsymbol{x}_k(t)\}$ randomly drawn from the multivariate Student t-distribution with $d = 3$ degrees of freedom.

In each case, the blind source separation of the simulated data is performed with two methods. For the first one, denoted *jadiag*, the sample covariance matrices $\hat{\boldsymbol{C}}_k = T^{-1}\sum_t \boldsymbol{x}_k(t)\boldsymbol{x}_k(t)^T$ are computed. These matrices are then jointly diagonalized with $\boldsymbol{B}$, minimizer of the log-likelihood criterion for Gaussian sources [6]. The estimated mixing matrix and source covariance matrices are given by $\hat{\boldsymbol{A}} = \boldsymbol{B}^{-1}$ and $\hat{\boldsymbol{\Lambda}}_k = \operatorname{ddiag}(\boldsymbol{B}\hat{\boldsymbol{C}}_k\boldsymbol{B}^T)$. The second method, denoted *MLE t-dist*, consists in optimizing the negative of the log-likelihood (10) on $\mathcal{M}$ for $d = 3$ with a Riemannian conjugate gradient alogrithm [17]. Optimization is performed with manopt toolbox [19].

To measure the performance of the two methods, three performance indexes are considered. The first one is the popular Moreau-Amari index [20], which measures how close $\boldsymbol{M} = \hat{\boldsymbol{A}}^{-1}\boldsymbol{A}$ is to a permuted diagonal matrix. It is defined

as

$$I_{\text{M-A}}(M) = \frac{1}{2n(n-1)} \sum_{p=1}^{n} \left(\frac{\sum_{q=1}^{n} |M_{pq}|}{\max\limits_{1 \le q \le n} |M_{pq}|} - 1 \right) + \frac{1}{2n(n-1)} \sum_{p=1}^{n} \left(\frac{\sum\limits_{q=1}^{n} |M_{qp}|}{\max\limits_{1 \le q \le n} |M_{qp}|} - 1 \right). \quad (13)$$

It is a measure between zero and one, with zero indicating a perfect recovery of the mixing process. The second one measures the distances in $\mathcal{S}_n^{++}$ between $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$ and $\hat{\boldsymbol{A}}\hat{\boldsymbol{\Lambda}}_k\hat{\boldsymbol{A}}^T$. It is defined as

$$I_{(\mathcal{S}_n^{++})^K} = \frac{1}{K} \sum_k \delta_R^2(\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T, \hat{\boldsymbol{A}}\hat{\boldsymbol{\Lambda}}_k\hat{\boldsymbol{A}}^T), \quad (14)$$

where δ_R is the usual Riemannian distance function on $\mathcal{S}_n^{++}$ [16]. The last performance index, denoted $I_{(s\mathcal{S}_n^{++})^K}$, is the same as $I_{(\mathcal{S}_n^{++})^K}$ except that we normalize the determinants of $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$ and $\hat{\boldsymbol{A}}\hat{\boldsymbol{\Lambda}}_k\hat{\boldsymbol{A}}^T$. It allows to only consider the structure of the estimated matrices without taking into account their power.

In Figure 1, we observe that the proposed method *MLE t-dist* features better performance than *jadiag* for all three performance indexes for every T. These results are expected since *MLE t-dist* corresponds to the optimal log-likelihood estimator according to these simulated data. Notice that the index $I_{(\mathcal{S}_n^{++})^K}$ of *jadiag* shows poor performance whereas $I_{(s\mathcal{S}_n^{++})^K}$ is more satisfying. It indicates that *jadiag* fails at retrieving the power of $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$ but captures their structure. On the other hand, *MLE t-dist* estimates both the power and the structure of $\boldsymbol{A}\boldsymbol{\Lambda}_k\boldsymbol{A}^T$ for these simulated data.

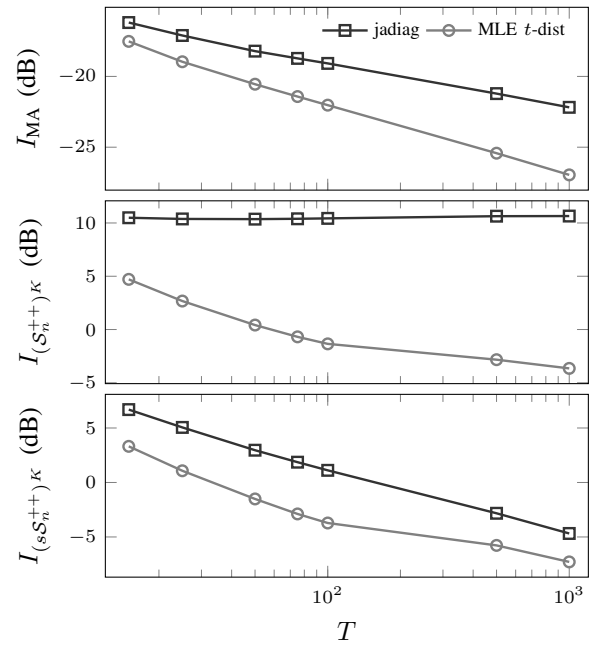

Fig. 1. Mean over 100 repetitions of the three performance indexes of the two considered blind source separation methods as functions of the number of samples T.

References

[1] P. Comon and C. Jutten. *Handbook of Blind Source Separation: Independent Component Analysis and Applications*. Academic Press, 1st edition, 2010.

[2] C. Jutten and J. Herault. Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture. *Signal processing*, 24(1):1–10, 1991.

[3] P. Comon. Independent component analysis, a new concept? *Signal processing*, 36(3):287–314, 1994.

[4] J.-F. Cardoso and A. Souloumiac. Blind beamforming for non Gaussian signals. *IEEE Proceedings-F*, 140(6):362–370, dec 1993.

[5] A. Belouchrani, K. Abed-Meraim, J.-F. Cardoso, and E. Moulines. A blind source separation technique using second-order statistics. *IEEE Transactions on signal processing*, 45(2):434–444, 1997.

[6] D. T. Pham and J.-F. Cardoso. Blind separation of instantaneous mixtures of nonstationary sources. *IEEE Transactions on Signal Processing*, 49(9):1837–1848, 2001.

[7] M. Congedo, C. Gouy-Pailler, and C. Jutten. On the blind source separation of human electroencephalogram by approximate joint diagonalization of second order statistics. *Clinical Neurophysiology*, 119(12):2677–2686, 2008.

[8] K. Alyani, M. Congedo, and M. Moakher. Diagonality measures of Hermitian positive-definite matrices with application to the approximate joint diagonalization problem. *Linear Algebra and its Applications*, 2016.

[9] F. Bouchard, J. Malick, and M. Congedo. Riemannian optimization and approximate joint diagonalization for blind source separation. *IEEE Transactions on Signal Processing*, 66(8):2041–2054, 2018.

[10] B. Afsari. Sensitivity analysis for the problem of matrix joint diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1148–1171, 2008.

[11] D.-T. Pham and M. Congedo. Least square joint diagonalization of matrices under an intrinsic scale constraint. In *Independent Component Analysis and Signal Separation*, pages 298–305. Springer, 2009.

[12] F. Bouchard, B. Afsari, J. Malick, and M. Congedo. Approximate joint diagonalization with Riemannian optimization on the general linear group. *SIAM Journal of Matrix Analysis and Applications*, 41(1):152–170, 2020.

[13] E. Ollila, D. E. Tyler, V. Koivunen, and H. V. Poor. Complex elliptically symmetric distributions: Survey, new results and applications. *IEEE Transactions on Signal Processing*, 60(11):5597–5625, 2012.

[14] F. Bouchard, A. Breloy, A. Renaux, and G. Ginolhac. Riemannian geometry and Cramér-Rao bound for blind separation of Gaussian sources. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP-2020)*. IEEE, accepted.

[15] B. Vandereycken, P.-A. Absil, and S. Vandewalle. A Riemannian geometry with complete geodesics for the set of positive semidefinite matrices of fixed rank. *IMA Journal of Numerical Analysis*, 33(2):481–514, 2012.

[16] R. Bhatia. *Positive definite matrices*. Princeton University Press, 2009.

[17] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, USA, 2008.

[18] P.-A. Absil and J. Malick. Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1):135–158, 2012.

[19] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.

[20] E. Moreau and O. Macchi. New self-adaptative algorithms for source separation based on contrast functions. In *Higher-Order Statistics, 1993., IEEE Signal Processing Workshop on*, pages 215–219. IEEE, 1993.