



Fast tensorial JADE

Joni Virta^{1,2}  | Niko Lietzén¹ | Pauliina Ilmonen¹ | Klaus Nordhausen³ 

¹Department of Mathematics and Systems Analysis, Aalto University School of Science

²Department of Mathematics and Statistics, University of Turku

³Institute of Statistics & Mathematical Methods in Economics, Vienna University of Technology

Correspondence

Klaus Nordhausen, Institute of Statistics & Mathematical Methods in Economics, Vienna University of Technology, Wiedner Hauptstrasse 8-10/105, A-1040 Wien, Austria
Email: Klaus.Nordhausen@tuwien.ac.at

Funding information

Academy of Finland, 321883; Austrian Science Fund, P31881-N32; CRoNoS COST Action, IC1408; Emil Aaltonen Foundation, 170156 N

Abstract

We propose a novel method for tensorial-independent component analysis. Our approach is based on TJADE and k -JADE, two recently proposed generalizations of the classical JADE algorithm. Our novel method achieves the consistency and the limiting distribution of TJADE under mild assumptions and at the same time offers notable improvement in computational speed. Detailed mathematical proofs of the statistical properties of our method are given and, as a special case, a conjecture on the properties of k -JADE is resolved. Simulations and timing comparisons demonstrate remarkable gain in speed. Moreover, the desired efficiency is obtained approximately for finite samples. The method is applied successfully to large-scale video data, for which neither TJADE nor k -JADE is feasible. Finally, an experimental procedure is proposed to select the values of a set of tuning parameters. Supplementary material including the R-code for running the examples and the proofs of the theoretical results is available online.

KEYWORDS

independent component analysis, joint diagonalization, Kronecker structure, limiting normality, tensorial-independent component analysis

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Scandinavian Journal of Statistics* published by John Wiley & Sons Ltd on behalf of The Board of the Foundation of the Scandinavian Journal of Statistics.

1 | TENSORIAL-INDEPENDENT COMPONENT ANALYSIS

In modern data analysis, an increasingly common and a natural assumption is that the covariance matrices are Kronecker-structured: the random vector $\mathbf{x} \in \mathcal{R}^{pq}$ has a covariance matrix $\mathbf{\Sigma} = \mathbf{\Sigma}_2 \otimes \mathbf{\Sigma}_1$, where $\mathbf{\Sigma}_1 \in \mathcal{R}^{p \times p}$, $\mathbf{\Sigma}_2 \in \mathcal{R}^{q \times q}$ are positive-definite and \otimes is the Kronecker product. In order for an estimator of $\mathbf{\Sigma}$ to be adequate, it should utilize this special structure (see, e.g., Leng & Pan, 2018; Roś, Bijma, de Munck, and de Gunst, 2016; Srivastava, von Rosen, & von Rosen, 2008; Werner, Jansson, & Stoica, 2008; Wiesel, 2012). In particular, if the Kronecker structure is ignored, the amount of parameters is inflated from $p(p+1)/2 + q(q+1)/2$ to $pq(pq+1)/2$.

Basic properties of the Kronecker product ensure that any zero-mean random vector \mathbf{x} , with the Kronecker covariance structure, admits a natural representation as a random matrix by means of the matrix location-scatter model,

$$\mathbf{x} = \mathbf{\Sigma}_1^{1/2} \mathbf{Z} \left(\mathbf{\Sigma}_2^{1/2} \right)',$$

where the original random vector is obtained from vectorization, $\text{vec}[\mathbf{X}] = \mathbf{x}$, and the latent random matrix \mathbf{Z} is standardized such that $\text{Cov}[\text{vec}[\mathbf{Z}]] = \mathbf{I}_{pq}$. In order to find interesting structures, the uncorrelatedness assumption for the elements of \mathbf{Z} is occasionally not enough. In Virta, Li, Nordhausen, and Oja (2017), the uncorrelatedness assumption was replaced by the assumption of full-statistical independence and the location-scatter model was extended to the matrix-valued-independent component model,

$$\mathbf{X} = \mathbf{\Omega}_1 \mathbf{Z} \mathbf{\Omega}_2', \quad (1)$$

where the invertible $\mathbf{\Omega}_1 \in \mathcal{R}^{p \times p}$, $\mathbf{\Omega}_2 \in \mathcal{R}^{q \times q}$ are unknown parameters and the latent tensor \mathbf{Z} is assumed to have mutually independent, marginally standardized components. The estimation procedure of a latent vector with independent components, using only the information provided by observed linear combinations of the components, is referred to as independent component analysis (ICA; see Comon & Jutten, 2010; Hyvärinen, Karhunen, & Oja, 2004; Ilmonen & Paindaveine, 2011; Miettinen, Taskinen, Nordhausen, & Oja, 2015; Nordhausen & Oja, 2018; Samarov & Tsybakov, 2004 for different approaches under classic multivariate settings). Since the location has no effect on the estimation of the so-called mixing matrices $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$, we can without loss of generality assume that \mathbf{Z} is centered. Under the additional assumption that at most one entire row of \mathbf{Z} has a multivariate normal distribution and at most one entire column of \mathbf{Z} has a multivariate normal distribution, it can be shown that the latent \mathbf{Z} is identifiable up to the order, joint scaling, and signs of its rows and columns (Virta et al., 2017). This pair of assumptions is less restrictive than its counterpart in standard vector-valued ICA, which allows a maximum of one normal component in the latent vector (Comon & Jutten, 2010). On the contrary, the assumptions of the matrix model allow \mathbf{Z} to contain up to $pq - \max(p, q)$ normal components, if the nonnormal elements are suitably located.

Vectorizing Equation (1) yields a Kronecker-structured standard independent component model

$$\mathbf{x} = (\mathbf{\Omega}_2 \otimes \mathbf{\Omega}_1) \mathbf{z},$$

where $\mathbf{x} = \text{vec}[\mathbf{X}] \in \mathcal{R}^{pq}$, $\mathbf{z} = \text{vec}[\mathbf{Z}] \in \mathcal{R}^{pq}$. Again, here any reasonable ICA approach should take the special form of the mixing matrix into account. To our knowledge, no structured

estimation methods have yet been developed for standard vector-valued ICA in this setting. However, the problem has been considered under the matrix representation of eq. (1) in Virta et al. (2017) and Virta, Li, Nordhausen, and Oja (2018), where two classical ICA procedures, fourth-order blind identification (FOBI; Cardoso, 1989) and joint diagonalization of eigenmatrices (JADE; Cardoso & Souloumiac, 1993), are extended to TFOBI and TJADE in order to solve Equation (1). Even though TJADE provides a definite improvement in efficiency with respect to TFOBI and the classical multivariate versions of the procedures, it is relatively expensive to compute. Consequently, the main goal of this article is to address this issue. We provide a computationally more powerful, alternative TJADE-based estimator, which retains the desirable statistical properties of TJADE. In particular, the estimator retains the consistency and the limiting distribution of TJADE.

We wish to point out that the statistical models for matrix data we are about to discuss are markedly different from the ones usually encountered when discussing array-valued data in engineering contexts. The relevant engineering literature is populated mostly with different tensor decompositions, the most popular ones being the Tucker decomposition and the CP-decomposition, which extend the singular value decomposition to higher order structures in different ways (see Cichocki et al., 2015, Kolda & Bader, 2009 for comprehensive reviews). A fundamental difference between the classical tensor decompositions and our current model is that the former rarely incorporate the concept of sample in the formulation. As such, tensor decompositions generally reduce also the dimension of the observation space, which is very unusual for classical statistical methods. These fundamentally different objectives also prevent any simple, meaningful comparisons, and therefore we have refrained from including comparisons to classical tensor decompositions in this work.

The rest of the article is structured as follows. In Section 2, we review the existing methodology on tensorial ICA and, in particular, TJADE on which our proposed extension will be based. In Section 3, we formulate k -TJADE in the case of matrix-valued observations and thoroughly establish its limiting behavior along with the necessary assumptions. The theoretical results are extended to a general tensor-valued model in Section 4 using the technique of matricization. Section 5 explores the finite-sample performance of the proposed method under both simulations and data examples on videos and handwritten digits. We also investigate the estimation of a set of tuning parameters of k -TJADE in the context of the latter example. In Section 6, we provide a short summary and list ideas for future work. Technical proofs are presented in the supplementary material.

2 | TENSORIAL JADE

We begin by briefly describing the theory behind TFOBI and TJADE. Everything in the following is formulated on the population level for a random matrix $\mathbf{X} \in \mathcal{R}^{p \times q}$. In practice, one would obtain a sample of matrices, $\mathbf{X}_1, \dots, \mathbf{X}_n$, from the distribution of \mathbf{X} and the expected values below should be replaced by sample means.

Assuming that the random matrix $\mathbf{X} \in \mathcal{R}^{p \times q}$ follows the model in Equation (1), the first step is to simultaneously standardize both, the rows and the columns of the matrix, using the left and right covariance matrices of \mathbf{X} ,

$$\Sigma_1[\mathbf{X}] = \frac{1}{q} \mathbb{E}[\mathbf{X}\mathbf{X}'] \quad \text{and} \quad \Sigma_2[\mathbf{X}] = \frac{1}{p} \mathbb{E}[\mathbf{X}'\mathbf{X}].$$

We denote the unique symmetric inverse square root of the positive-definite matrix \mathbf{S} by $\mathbf{S}^{-1/2}$. The standardized variable $\mathbf{X}^{\text{st}} = (\boldsymbol{\Sigma}_1^{-1/2}[\mathbf{X}])\mathbf{X}(\boldsymbol{\Sigma}_2^{-1/2}[\mathbf{X}])'$ satisfies $\mathbf{X}^{\text{st}} = \tau\mathbf{U}_1\mathbf{Z}\mathbf{U}_2'$ for some $\tau = \tau(\boldsymbol{\Omega}_1, \boldsymbol{\Omega}_2) > 0$ and some orthogonal matrices, $\mathbf{U}_1 \in \mathcal{R}^{p \times p}$, $\mathbf{U}_2 \in \mathcal{R}^{q \times q}$ (see Virta et al., 2017). The unknown constant of proportionality τ is a result of the joint scaling of the model in Equation (1) being left unfixed. After the standardization, solving the IC problem is reduced to the estimation of the orthogonal matrices $\mathbf{U}_1, \mathbf{U}_2$, a task commonly addressed in ICA using higher order cumulants. The TFOBI and TJADE procedures also utilize the higher order cumulants. In TFOBI, a Fisher consistent (under mild assumptions, see Section 3) estimator, $\mathbf{\Gamma}^{\text{F}}[\mathbf{X}]$, for the inverse of the matrix $\boldsymbol{\Omega}_1$ is constructed such that

$$\mathbf{\Gamma}^{\text{F}}[\mathbf{X}] = (\mathbf{V}^{\text{F}}[\mathbf{X}])' \boldsymbol{\Sigma}_1^{-1/2}[\mathbf{X}],$$

where the columns of $\mathbf{V}^{\text{F}}[\mathbf{X}]$ are the eigenvectors of the matrix

$$\mathbf{B}[\mathbf{X}] = \frac{1}{q} \mathbb{E}[\mathbf{X}\mathbf{X}'\mathbf{X}\mathbf{X}'] .$$

Thus, $\mathbf{\Gamma}^{\text{F}}[\mathbf{X}]$ provides a solution for the left-hand side of the IC model in Equation (1) (Virta et al., 2017).

The TJADE procedure utilizes a set of matrices, $C = \{ \mathbf{C}^{ij}[\mathbf{X}^{\text{st}}] : i, j \in \{1, \dots, p\} \}$, referred to as the set of cumulant matrices, such that

$$\mathbf{C}^{ij}[\mathbf{X}] = \frac{1}{q} \mathbb{E}[\mathbf{e}_i' \mathbf{X}\mathbf{X}' \mathbf{e}_j \mathbf{X}\mathbf{X}'] - \boldsymbol{\Sigma}_1[\mathbf{X}] (\delta_{ij} q \mathbf{I}_p + \mathbf{E}^{ij} + \mathbf{E}^{ji}) (\boldsymbol{\Sigma}_1[\mathbf{X}])', \tag{2}$$

where δ_{ij} is the Kronecker delta, $\mathbf{e}_k \in \mathcal{R}^p$, $k \in \{1, \dots, p\}$, are the standard basis vectors of \mathcal{R}^p , and $\mathbf{E}^{kl} = \mathbf{e}_k \mathbf{e}_l'$. The left covariance matrix of the standardized matrix satisfies $\boldsymbol{\Sigma}_1[\mathbf{X}^{\text{st}}] = \tau^2 \mathbf{I}_p$ and is included in Equation (2) solely for the estimation of the constant of proportionality τ . The authors in Virta et al. (2018) proved that the joint diagonalizer of C is under mild assumptions, see Section 3, equal to the orthogonal matrix \mathbf{U}_1 up to the order and signs of its columns. The joint diagonalizer of the set C is defined as any orthogonal matrix $\mathbf{V} \in \mathcal{R}^{p \times p}$ that minimizes

$$\tilde{g}(\mathbf{V}, \mathbf{X}^{\text{st}}) = \sum_{i,j=1}^p \left\| \text{off}(\mathbf{V}' \mathbf{C}^{ij}[\mathbf{X}^{\text{st}}] \mathbf{V}) \right\|_{\text{F}}^2, \tag{3}$$

where $\text{off}(\mathbf{S}) \in \mathcal{R}^{p \times p}$ is equal to $\mathbf{S} \in \mathcal{R}^{p \times p}$ with the diagonal elements set to zero.

The joint diagonalizer defines a coordinate system in which the linear transformations $\mathbf{C}^{ij}[\mathbf{X}^{\text{st}}]$, $i, j \in \{1, \dots, p\}$, have minimal sum of squared off-diagonal elements. There exists several algorithms for optimizing Equation (3), the most popular being the Jacobi-rotation technique, for details see Belouchrani, Abed-Meraim, Cardoso, and Moulines (1997) and Illner et al. (2015). After the estimation of the joint diagonalizer $\mathbf{V}^{\text{J}}[\mathbf{X}]$, an estimated inverse for the matrix $\boldsymbol{\Omega}_1$ is obtained as the TJADE-functional, $\mathbf{\Gamma}^{\text{J}}[\mathbf{X}] = (\mathbf{V}^{\text{J}}[\mathbf{X}])' \boldsymbol{\Sigma}_1^{-1/2}[\mathbf{X}]$.

The results of this article are derived only for the left-hand side of the matrix-valued model. The right-hand side of the matrix-valued model can be solved exactly as the left-hand side. One can simply take the transpose of \mathbf{X} and proceed as with the left-hand side. Only the sizes of the cumulant and transformation matrices change from $p \times p$ to $q \times q$. Moreover, matricization allows us to extend the estimators beyond the matrix-valued model to arbitrary-dimensional

tensor-valued IC models (see Virta et al., 2017, 2018). Matricization allows us to hide a considerable amount of the unpleasant notation related to tensor algebra (see Section 4). In total, it is sufficient to present the results in matrix form and only for the left-hand side of the model in Equation (1).

When the dimension q is equal to one, the Tjade procedure for the left-hand side of the model is equivalent to the standard JADE for vector-valued data. Extensive comparisons between JADE and Tjade are conducted in Virta, Taskinen, and Nordhausen (2016) and Virta et al. (2017, 2018) with the conclusion that the latter is uniformly superior to the former under the Kronecker-structured IC model. Moreover, the tensorial version is computationally significantly faster. Consider a tensor of r th order with all dimensions of size p . Standard JADE requires a single joint diagonalization of p^{2r} matrices that are of size $p^r \times p^r$, whereas Tjade requires r joint diagonalizations of p^2 matrices that are of size $p \times p$. In essence, adding dimensions to a tensor has a multiplicative effect on the number of operations the classic vectorial methods require and merely an additive effect on the tensorial methods. However, even with its considerable advantages over JADE, running the Tjade procedure is slow for large tensors.

To obtain a faster method, we approach the problem in the spirit of Miettinen, Nordhausen, Oja, and Taskinen (2013), where a faster version of JADE, k -JADE, is derived. The modification can be described very succinctly: instead of diagonalizing the entire set of cumulant matrices \mathbf{C}^{ij} , we diagonalize only a specific subset of them, chosen such that the desirable statistical properties of Tjade are still carried over to the extension. Since the subset of cumulant matrices can be chosen separately in each direction of the tensor, the k -JADE approach provides even more significant improvements in tensorial settings compared with its original use in improving JADE. Note that, similar ideas as in Miettinen et al. (2013) were used already in Cardoso (1999) to formulate shifted blocks for blind separation, where only those cumulant matrices of regular JADE with matching indices are diagonalized.

3 | TENSORIAL k -JADE

In this section, we propose a novel extension of the Tjade procedure. We formulate the extension, k -Tjade, such that it retains the following three key properties of Tjade. The first of these properties is the ability to solve the tensor-independent component model, manifesting either as Fisher consistency or consistency, depending on whether we are at the population or sample level, respectively. The second property is orthogonal equivariance under arbitrary data. ICA-estimators are customarily expected to have some form of equivariance, which makes the generalization of limiting properties more straightforward (Miettinen et al., 2015; Virta et al., 2017). The third desired property is the limiting distribution of Tjade, the one with the lowest limiting variance of the known tensorial ICA methods. Next, we establish these properties one-by-one. As mentioned in the previous section, all results derived for the left-hand side of the model also hold for the right-hand side, prompting us to consider only the former in the following. In the same spirit, even though the proposed k -Tjade method has, in the case of matrix data, two tuning parameters, k_1 for the rows and k_2 for the columns, the method still acts only on a single mode at a time, and as such we omit the subscripts and speak in the following only of the tuning parameter k . Moreover, the same idea is reflected in the name of the method which should (for matrix data) technically be (k_1, k_2) -Tjade. In order to keep the presentation more readable, we prefer to call the method simply k -Tjade.

We define matrix-independent component functionals, the extension of independent component functionals (Miettinen et al., 2015) to matricial ICA.

Definition 1. A $p \times p$ matrix-valued functional Γ is a matrix-independent component (IC) functional if

1. $\Gamma[\mathbf{X}] \equiv \Omega_1^{-1}$ for all $\mathbf{X} \in \mathcal{R}^{p \times q}$ that follow the matrix IC model of Equation (1),
2. $\Gamma[\mathbf{U}_1 \mathbf{X} \mathbf{U}_2'] \equiv \Gamma[\mathbf{X}] \mathbf{U}_1'$ for all $\mathbf{X} \in \mathcal{R}^{p \times q}$ and all orthogonal $\mathbf{U}_1 \in \mathcal{R}^{p \times p}$, $\mathbf{U}_2 \in \mathcal{R}^{q \times q}$,

where two matrices $\mathbf{A}, \mathbf{B} \in \mathcal{R}^{p \times p}$ satisfy $\mathbf{A} \equiv \mathbf{B}$ if $\mathbf{A} = c\mathbf{P}\mathbf{J}\mathbf{B}$ for some $c > 0$, some diagonal matrix $\mathbf{J} \in \mathcal{R}^{p \times p}$ with diagonal elements equal to ± 1 , and some permutation matrix $\mathbf{P} \in \mathcal{R}^{p \times p}$.

The first condition in Definition 1 requires that a matrix IC functional must be able to solve the left-hand side of the model in Equation (1) (Fisher consistency). The second condition essentially states that the functional cancels out any orthogonal transformations on the observed matrices (orthogonal equivariance). As a particularly useful consequence of the latter, the limiting distribution of a matrix IC functional under trivial mixing, $\Omega_1 = \mathbf{I}_p$, $\Omega_2 = \mathbf{I}_q$, instantly generalizes to any orthogonal mixing as well.

Let $\boldsymbol{\kappa} \in \mathcal{R}^p$ be the vector of the row means of the elementwise kurtoses, $\mathbb{E}[x_{kl}^4] - 3$, of the elements of \mathbf{Z} .

Assumption 1. At most one element of $\boldsymbol{\kappa}$ equals zero.

Assumption 2 (ν). The multiplicities of the elements of $\boldsymbol{\kappa}$ are at most ν .

The TFOBI functional Γ^F is a matrix IC functional in the sense of Definition 1 if Assumption 2(1) is satisfied and the TJADE functional Γ^J is a matrix IC functional if Assumption 1 is satisfied. Naturally, the column mean analogues of the assumptions are required to separate the right-hand side of the model in Equation (1).

The same comparison as was done between the normality assumptions in Section 1 holds analogously between Assumption 1 and Assumption 2(ν) and their vectorial counterparts, allowing maximally one zero-kurtosis component or only distinct kurtoses, respectively. The main implication is that in matrix ICA numerous latent components may have identical kurtoses as long as their row means (and column means when separating the right-hand side of the model) satisfy the necessary requirements. The assumptions also satisfy the following set of relations,

$$\text{Assumption 2(1)} \subset \dots \subset \text{Assumption 2(p)},$$

where \subset means “implies.” Moreover, we have also **Assumption 2(1) \subset Assumption 1**.

In order to speed up TJADE such that the properties in Definition 1 are retained, we proceed as in Miettinen et al. (2013), and instead of diagonalizing the set C , we diagonalize only those members of it which satisfy $|i - j| < k$, for some predefined value of the tuning parameter $k \in \{1, \dots, p\}$. This discarding can be motivated in two ways. First, all except the repeated index matrices, $\mathbf{C}^{11}, \dots, \mathbf{C}^{pp}$, vanish asymptotically. Every matrix $\mathbf{C}^{ij}[\mathbf{Z}] = \mathbf{0}$, $i \neq j$, implying that with increasing sample size, all the separation information is eventually contained in the p repeated index matrices. Second, by assuming that the values in $\boldsymbol{\kappa}$ are in decreasing order (this is guaranteed by using TFOBI as a preprocessing step, see the next paragraph) the i th row of \mathbf{Z} is the most difficult to separate from its immediate neighboring rows and the separation information between them is contained precisely in the matrices \mathbf{C}^{ij} and \mathbf{C}^{ji} where j is close to i .

Analogously to k -JADE, we use the TFOBI-algorithm to obtain an initial value for the functional. This ensures that even after the previous modification, the functional remains orthogonally equivariant. The following definition and theorem formalize our resulting novel method, called k -TJADE.

Definition 2. Fix $k \leq p$. The k -TJADE functional is

$$\mathbf{\Gamma}^k[\mathbf{X}] = (\mathbf{V}[\mathbf{X}^F])' \mathbf{\Gamma}^F[\mathbf{X}],$$

where $\mathbf{\Gamma}^F$ is the TFOBI functional, $\mathbf{X}^F = \mathbf{\Gamma}^F[\mathbf{X}] \mathbf{X} (\mathbf{\Gamma}^F[\mathbf{X}'])'$ is the TFOBI-solution for \mathbf{X} , and the orthogonal matrix $\mathbf{V}[\mathbf{X}^F] = (\mathbf{v}_1, \dots, \mathbf{v}_p)$ is the joint diagonalizer of the set of matrices $\mathcal{C}_k = \{\mathbf{C}^{ij}[\mathbf{X}^F] : |i - j| < k\}$.

Theorem 1. Let Assumptions 1 and 2(v) hold for some fixed v . Then the k -TJADE functional $\mathbf{\Gamma}^k$ is a matrix IC functional for all $k \geq v$.

Theorem 1 provides the Fisher consistency and the orthogonal equivariance for the k -TJADE functional. The assumptions that Theorem 1 requires are interesting since they provide an interpretation for the tuning parameter k —the parameter is the maximal number of allowed kurtosis mean multiplicities. The values $k = 1$ and $k = p$ correspond to the extreme cases where all the kurtosis means have to be distinct (as in TFOBI) and where no assumptions are made on the multiplicities of the nonzero kurtosis means (as in TJADE). Thus, k -TJADE can be seen as a middle ground between TFOBI and TJADE. As the assumptions, also the methods can be ordered according to the strictness of the assumptions they require,

$$\text{TFOBI} \geq 1\text{TJADE} \geq \dots \geq p - \text{TJADE} = \text{TJADE},$$

where “ \geq ” is read as “makes at least as many assumptions as” and “=” as “makes the same assumptions as”.

We next give some intuition behind the proof of Theorem 1 (given in the supplementary material). The proof relies on a specific interplay between the preliminary TFOBI-step and the set of cumulant matrices \mathcal{C} . Being an eigendecomposition-based method, TFOBI is able to estimate the matrix $\mathbf{\Omega}_1^{-1}$ only up to row blocks determined by the elements of $\boldsymbol{\kappa}$, such that the rows of the estimate corresponding to equal kurtosis values remain mixed by orthogonal matrices (c.f., eigenvectors corresponding to multiple eigenvalues are not uniquely defined but their span is). However, the proof also shows that the joint diagonalization of all cumulant matrices with $|i - j| < k$, for some given k , can separate mixed row blocks of size at most k (intuitively, increasing k means using more cumulant matrices which increases the amount of available information). Thus, by sequencing the two steps together and having $k \geq v$ ensures that any blocks left still mixed after the TFOBI-step will be unmixed in the joint diagonalization step.

Assumptions 1 and 2(v) not only provide Theorem 1 but are also sufficient to guarantee the two remaining desired properties of k -TJADE, consistency and the same limiting distribution as that of TJADE. These asymptotic properties are formalized in the following two theorems. Remarkably as a special case, when $q = 1$, the latter also proves a previously unsolved conjecture posed about the limiting behavior of vectorial k -JADE in Miettinen et al. (2015).

Theorem 2. Let Assumptions 1 and 2(v) hold for some fixed v and let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be an i.i.d. sequence from the matrix IC model in Equation (1) with identity mixing matrices, $\mathbf{\Omega}_1 = \mathbf{I}_p, \mathbf{\Omega}_2 = \mathbf{I}_q$. Assume that the population quantity \mathbf{X} has finite eighth moments. Then, for all $k \geq v$, there exists a consistent

sequence of k -TJADE-estimators (indexed by n). That is,

$$\hat{\mathbf{\Gamma}}^k \rightarrow_{\mathbb{P}} \mathbf{I}_p.$$

Theorem 3. Under the assumptions of Theorem 2, we have for all $k \geq \nu$ that,

$$\sqrt{n}(\hat{\mathbf{\Gamma}}^k - \mathbf{I}_p) = \sqrt{n}(\hat{\mathbf{\Gamma}}^J - \mathbf{I}_p) + o_p(1),$$

where $\hat{\mathbf{\Gamma}}^J$ is the TJADE-estimator. The notation $o_p(1)$ refers to a sequence of random matrices that converges in probability to the zero matrix.

The proofs of Theorems 2 and 3 are based on the fact that the kurtosis matrices of the latent \mathbf{Z} with nonmatching indices vanish asymptotically, $\mathbf{C}^{ij}[\mathbf{Z}] = \mathbf{0}$, $i \neq j$. The key point of the proof is then to show that this property transfers to the corresponding kurtosis matrices $\mathbf{C}^{ij}[\mathbf{X}^F]$, of the TFOBI-standardized observation, to such extent that the asymptotical contribution of the sample estimates of these matrices to the final k -TJADE estimate is negligible. Thus, regardless of the choice of k , the asymptotical behavior of the method is determined by the p repeated index matrices $\mathbf{C}^{11}, \dots, \mathbf{C}^{pp}$. Finally, the proof reveals that the rotation specified by the joint diagonalization dominates over the preliminary TFOBI-rotation and the effect of the latter also turns out to be asymptotically negligible.

Note that, $\hat{\mathbf{\Gamma}}^J = \hat{\mathbf{\Gamma}}^p$ does not generally hold as the latter estimator utilizes the preliminary TFOBI-step, while the former does not. The limiting distribution of $\sqrt{n}(\hat{\mathbf{\Gamma}}^J - \mathbf{I}_p)$ is in Virta et al. (2018) shown to be multivariate normal and closed form expressions for its limiting variances are also given therein. By the orthogonal equivariance of matrix-independent component functionals (property (ii) of Definition 1), the limiting results of Theorem 3 generalize to any orthogonal mixing matrices. Note that the original k -JADE in Miettinen et al. (2013) is affine equivariant (equivariant under all coordinate system changes, not just orthogonal). The problem of achieving affine equivariance in the context of tensorial ICA is discussed in Virta et al. (2018). There it was conjectured that tensorial ICA cannot be affine equivariant.

The two limiting theorems above show that, under suitable assumptions, k -TJADE indeed has all the desirable properties listed at the beginning of this section. This can be summarized by saying that k -TJADE makes a trade-off between assumptions and computational load: With the price of added assumptions, we obtain a method with the same limiting efficiency as TJADE, but with significantly lighter computational burden. As the claim about efficiency holds only asymptotically, we conduct a simulation study (in Section 5) to compare the finite-sample efficiency of the estimators.

Note that the maximal kurtosis multiplicity ν is unknown in practice, which makes the choosing of k such that $k \geq \nu$ a nontrivial task. However, the simulations of Section 5 reveal that not a lot of separation efficiency is necessarily lost when using a slightly too small value of k , making the problem less dire. To further alleviate the issue, in the final part of Section 5, we propose a procedure for estimating ν and apply it to handwritten digit data.

4 | A NOTE ON TENSORIAL ICA

In this section, we formulate the general tensorial IC model and discuss how it can be reduced to the matricial IC model. We begin with a short review of the basic concepts of multilinear algebra.

An r th order tensor $\mathcal{X} = (x_{i_1 \dots i_r}) \in \mathcal{R}^{p_1 \times \dots \times p_r}$ is an r -dimensional array containing a total of $\rho = \prod_{m=1}^r p_m$ elements and has a total of r modes or ways, that is, directions from which we can view it. For example, a matrix ($r = 2$) can be viewed either through its columns or through its rows. Two complementary ways of dividing a tensor into a disjoint collection of smaller tensors are called the m -mode vectors and the m -mode faces. In the former, we choose an index $m = 1, \dots, r$ and have the values of the indices $\{1, \dots, r\} \setminus \{m\}$ fixed and let the m th index vary over its range. Each fixed combination of the $r - 1$ indices then yields a single p_m -dimensional vector and the collection of all ρ/p_m such vectors is called the set of m -mode vectors of \mathcal{X} . On the other hand, if we fix the value of the m th index and let the others vary over their ranges, we get a total of p_m tensors of order $r - 1$ and size $p_1 \times \dots \times p_{m-1} \times p_{m+1} \times \dots \times p_r$, called the m -mode faces of \mathcal{X} . Illustrations of both, m -mode vectors and m -mode faces, $m = 1, 2, 3$, in the case of a three-dimensional tensor, are shown in Figures 1 and 2.

Two algebraic operations can now be defined in terms of the m -mode vectors. Let $\mathcal{X} \in \mathcal{R}^{p_1 \times \dots \times p_r}$ and $\mathbf{A}_m = (a_{i_m j_m}^{(m)}) \in \mathcal{R}^{q_m \times p_m}$ for some fixed $m = 1, \dots, r$. The m -mode multiplication $\mathcal{X} \times_m \mathbf{A}_m \in \mathcal{R}^{p_1 \times \dots \times p_{m-1} \times q_m \times p_{m+1} \times \dots \times p_r}$ of \mathcal{X} from the m th mode by \mathbf{A}_m is defined elementwise as,

$$(\mathcal{X} \times_m \mathbf{A}_m)_{i_1, \dots, i_r} = \sum_{j_m=1}^{p_m} x_{i_1 \dots i_{m-1} j_m i_{m+1} \dots i_r} a_{i_m j_m}^{(m)}.$$

The m -mode multiplication is easily understood: $\mathcal{X} \times_m \mathbf{A}_m$ is obtained by multiplying each m -mode vector of \mathcal{X} from the left by \mathbf{A}_m and collecting the resulting vectors back into an




FIGURE 1 The 1-mode, 2-mode, and 3-mode vectors of a three-dimensional tensor. Elina Vartiainen®




FIGURE 2 The 1-mode, 2-mode and 3-mode faces of a three-dimensional tensor. Elina Vartiainen®

r -dimensional tensor in the same order. The m -mode multiplications $\times_1, \dots, \times_r$ from distinct modes are commutative and we use the shorthand $\mathcal{X} \times_{m=1}^r \mathbf{A}_m = \mathcal{X} \times_1 \mathbf{A}_1 \dots \times_r \mathbf{A}_r$. In the case of matrices, $\mathcal{X} = \mathbf{X} \in \mathcal{R}^{p_1 \times p_2}$ ($r = 2$), the simultaneous multiplication from both modes simply gives $\mathbf{X} \times_{m=1}^2 \mathbf{A}_m = \mathbf{A}_1 \mathbf{X} \mathbf{A}_2'$.

To show the connection between the matricial and tensorial IC methods, we still need the concept of m -mode matricization. For a fixed mode $m = 1 \dots, r$, the m -mode matricization $\mathbf{X}^{(m)} \in \mathcal{R}^{p_m \times \rho/p_m}$ of \mathcal{X} is obtained by taking the m -mode vectors of \mathcal{X} and collecting them horizontally into a wide matrix. The arrangement order of the m -mode vectors has no effect on our further developments and we choose to use the cyclical ordering as in De Lathauwer, De Moor, and Vandewalle (2000). In this case, the relationship,

$$(\mathcal{X} \times_m \mathbf{A}_m)^{(m)} = \mathbf{A}_m \mathbf{X}^{(m)} (\mathbf{A}_{m+1} \otimes \dots \otimes \mathbf{A}_r \otimes \mathbf{A}_1 \otimes \dots \otimes \mathbf{A}_{m-1})' \tag{4}$$

holds. For more comprehensive introduction to multilinear algebra, see Cichocki, Zdunek, Phan, and Amari (2009) and De Lathauwer et al. (2000).

We are now sufficiently equipped to examine the connection between the matricial ICA and the tensorial ICA. The zero-mean r th-order random tensor $\mathcal{X} \in \mathcal{R}^{p_1 \times \dots \times p_r}$ is said to follow the tensorial independent component model if

$$\mathcal{X} = \mathcal{L} \times_{m=1}^r \mathbf{\Omega}_m, \tag{5}$$

where the random tensor $\mathcal{L} \in \mathcal{R}^{p_1 \times \dots \times p_r}$ has mutually independent, marginally standardized components and $\mathbf{\Omega}_m \in \mathcal{R}^{p_m \times p_m}$, $m = 1, \dots, r$, are invertible (Virta et al., 2017). We further assume that for each mode $m = 1, \dots, r$ at most one m -mode face of \mathcal{L} consists solely of normally distributed components. The objective of tensorial ICA is to estimate \mathcal{L} given a random sample of observed tensors from the distribution of \mathcal{X} .

We fix the mode $m = 1, \dots, r$ and consider the m -mode matricization of the model in Equation (5). It now follows from Equation (4) that

$$\mathbf{X}^{(m)} = \mathbf{\Omega}_m \mathbf{Z}^{(m)} (\mathbf{\Omega}_{m+1} \otimes \dots \otimes \mathbf{\Omega}_r \otimes \mathbf{\Omega}_1 \otimes \dots \otimes \mathbf{\Omega}_{m-1})'. \tag{6}$$

As Kronecker products of invertible matrices are themselves invertible, a comparison with Equation (1) now reveals that Equation (6) is in the form of a matrix IC model with $\mathbf{\Omega}_m$ replacing $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_{m+1} \otimes \dots \otimes \mathbf{\Omega}_r \otimes \mathbf{\Omega}_1 \otimes \dots \otimes \mathbf{\Omega}_{m-1}$, taking the role of $\mathbf{\Omega}_2$. In addition, $\mathbf{Z}^{(m)}$ satisfies all the assumptions of the matrix IC model. To be precise, it is possible that multiple columns of $\mathbf{Z}^{(m)}$ have multivariate normal distributions as a consequence of the matricization. However, as the m -mode matricization leaves the structure of the m th mode intact, the normality assumption on \mathcal{L} guarantees that at most one row of $\mathbf{Z}^{(m)}$ has a multivariate normal distribution, which is then sufficient for the identifiability of $\mathbf{\Omega}_m$, the current parameter of interest. Thus, using Equation (6), we can estimate $\mathbf{\Omega}_m$ exactly as $\mathbf{\Omega}_1$ in the matricial case. In the tensorial case, the kurtosis means in the vector $\boldsymbol{\kappa}$ in Assumption 1 and in Assumption 2(v) are computed over the rows of $\mathbf{Z}^{(m)}$, or equivalently, over the m -mode faces of \mathcal{L} . All our theoretical results, such as orthogonal equivariance (a Kronecker product of orthogonal matrices is itself orthogonal) and the asymptotic variances, hold fully under the tensorial IC model, as long as the relevant kurtosis assumptions are satisfied. Finally, we remind that for an r th order input tensor \mathcal{X} , the k -TJADE has a total of r tuning parameters, k_1, \dots, k_r , with the underlying idea that k_m should be chosen to be equal to or larger than the maximal kurtosis multiplicity in the corresponding mode (see Theorem 3).

5 | SIMULATIONS AND EXAMPLES

In this section, we illustrate the finite sample properties of the tensor-valued procedures TFOBI, TJADE, and the novel k -TJADE, which are implemented in the R-package *tensorBSS* (Virta, Nordhausen, Oja, & Li, 2016)). For comparison, we also consider the classical vector-valued versions of these estimators, denoted by VFOBI, VJADE, and k -VJADE, as implemented in the JADE package (Miettinen, Nordhausen, & Taskinen, 2017). In the classical procedures, the tensor-valued observations are first vectorized and the algorithms are then applied to the resulting data matrices. R-code for running the examples and obtaining the used datasets is available in the supplementary material.

Next, we consider a collection of observed i.i.d. tensors $\{\mathcal{X}_j\}_{j \in \{1, \dots, n\}}$, generated from the tensorial-independent component model, such that $\mathcal{X}_j \in \mathcal{R}^{p_1 \times \dots \times p_r}$, for every j . Let $\Omega_1, \dots, \Omega_r$ be the theoretical mixing matrices and let $\hat{\Gamma}_1, \dots, \hat{\Gamma}_r$ be the corresponding unmixing estimates produced by one of the tensor-valued procedures. We denote the Kronecker products of the matrices as $\hat{\Gamma} = \hat{\Gamma}_1 \otimes \dots \otimes \hat{\Gamma}_r$ and $\Omega = \Omega_1 \otimes \dots \otimes \Omega_r$. The vector-valued procedures produce a single unmixing estimate, denoted also by $\hat{\Gamma}$. Note that the estimates produced by the vectorial and the tensorial methods are comparable, since in both cases the matrix $\hat{\Gamma}$ estimates the inverse of the compound matrix Ω .

The so-called gain matrix is defined as $\hat{\mathbf{G}} = \hat{\Gamma}\Omega$. The unmixing is considered successful if the gain matrix is close to the identity matrix, up to the order and signs of its rows. We quantify this closeness using the performance measure called minimum distance (MD) index (Ilmonen, Nordhausen, Oja, & Ollila, 2010). The MD index is formulated as follows,

$$D(\hat{\mathbf{G}}) = \frac{1}{\sqrt{\rho - 1}} \inf_{\mathbf{C} \in C_0} \|\mathbf{C}\hat{\mathbf{G}} - \mathbf{I}_\rho\|_F, \quad (7)$$

where $\rho = \prod_{j=1}^r p_j$ and C_0 is the set of all matrices with exactly one nonzero element in each row and column. The range of the MD index is $[0, 1]$, where the value 0 corresponds to the case of the gain matrix being exactly a permuted and scaled identity matrix, that is, the estimate is perfectly accurate. In addition, the limiting distribution of the MD index can be obtained from the limiting distribution of the corresponding IC functional $\hat{\Gamma}$ (see Ilmonen et al., 2010, theorem 1; Virta et al., 2017, theorem 6).

In simulation studies, where multiple iterations are performed under identical conditions, the asymptotic value for the mean of the transformed MD index $n(\rho - 1)D(\hat{\mathbf{G}})^2$ can be obtained using the limiting variances of the applied IC functionals (see Virta et al., 2017 for further details). The convergence toward the theoretical limiting values given by Theorem 3 can then be demonstrated by visualizing the theoretical limits alongside $\ell_i^{-1} \sum_{j=1}^{\ell_i} n_i(\rho - 1)D(\hat{\mathbf{G}}_j)^2$, where ℓ_i is the number of iterations for the sample size n_i and $\hat{\mathbf{G}}_j$ is the estimated gain matrix for the corresponding j th iteration.

Before presenting the simulations, we discuss shortly the importance of the tuning parameter k , which in k -TJADE can be chosen separately for each mode. In the following, (k_1, \dots, k_r) -TJADE is used to refer to k -TJADE with the value k_m for the tuning parameter in the m th mode, $m \in \{1, \dots, r\}$. Based on Theorem 3, the computationally cheapest but still asymptotically optimal choice is the largest multiplicity of the kurtosis means in the current mode. This means that we have generally three choices for each mode: we may use the full TJADE if the mode is short; we can use k -TJADE for some small k if the mode is long; or we can choose not to separate the mode

at all if it is not expected to contain any relevant information, as might be the case, for example, for the color dimension of a sample of images.

5.1 | Finite-sample efficiency: Setting 1

We begin by demonstrating that the finite sample performance of k -TJADE is in line with the asymptotic results given in Theorem 3. In the first setting, we consider simulated collections of i.i.d. matrices of size 3×3 , $\mathbf{Z} := \{\mathbf{Z}_j\}_{j \in \{1, \dots, n\}}$. The components of every \mathbf{Z}_j are simulated independently from the following distributions,

$$\mathbf{Z}_j = \begin{pmatrix} \text{E} & \text{C} & \text{U} \\ \text{C} & \text{U} & \text{E} \\ \text{U} & \text{E} & \text{N} \end{pmatrix},$$

where C, E, U, and N denote independent replicates from the χ_1^2 , standard exponential, uniform, and normal distribution, respectively, all scaled to have zero mean and unit variance.

In this simulation setting, none of the theoretical row or column kurtosis means are zero. However, the theoretical mean kurtoses are the same for the first two rows and the first two columns. Thus, the preferable (k_1, k_2) -TJADE procedure here is 22-TJADE, $k_1 = k_2 = 2$. Note that in this setting, the requirements of TFOBI are not fulfilled.

When the observations are vectorized, the length nine vectors contain a single normally distributed component, two χ^2 -distributed elements, three elements from the exponential distribution, and three elements from the uniform distribution. The vector now contains one element (the normally distributed component) with theoretical kurtosis 0, making vectorial ICA viable. The most natural k -VJADE procedure here is 3-VJADE. The assumptions of VFOBI are violated.

The simulation was performed using 13 different sample sizes, $n_i = 2^{i-1} \times 1000$, $i \in \{1, \dots, 13\}$ with $\ell = 2000$ repetitions per sample size. To evaluate the equivariance properties and the effect of nonorthogonal mixing, we mixed the observations at each repetition using (a) identity mixing: $\mathbf{\Omega}_1 = \mathbf{\Omega}_2 = \mathbf{I}_p$, (b) orthogonal mixing: $\mathbf{\Omega}_1 = \mathbf{U}_1$ and $\mathbf{\Omega}_2 = \mathbf{U}_2$, where \mathbf{U}_1 and \mathbf{U}_2 are random orthogonal matrices uniformly sampled at each repetition with respect to the Haar measure, and (c) normal mixing, where $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$ were filled at each repetition with random elements from $N(0, 1)$.

To evaluate the effect of the mixing matrix and the limiting distributions, we present, in Figure 3, the transformed MD values for 3-VJADE, VJADE, 22-TJADE and TJADE and the corresponding theoretical limit values for the cases available. As our theory extends only the orthogonal mixing, the theoretical limit value of TJADE is missing from the third panel. The sample sizes n_i , $i \geq 8$, are omitted from Figure 3, since the sample size 64×10^3 is large enough for the convergence of both the vectorial and the tensorial methods.

Figure 3 clearly shows that for the vectorial methods, we have affine invariance and thus all of the three curves are identical. For the tensorial methods, we have only orthogonal invariance and hereby the curve corresponding to the normal mixing differs from the other two. Moreover, despite the limiting value of TJADE being unknown for normal mixing, based on the finite-sample curves, it seems that normal mixing does not change the relative separation accuracy of the tensorial methods over orthogonal mixing. Only the overall separation difficulty level is affected (the curves are higher in the plots). Furthermore, the benefit of applying the tensorial methods over the vectorized methods is impressive. Finally, Figure 3 also illustrates that all the methods

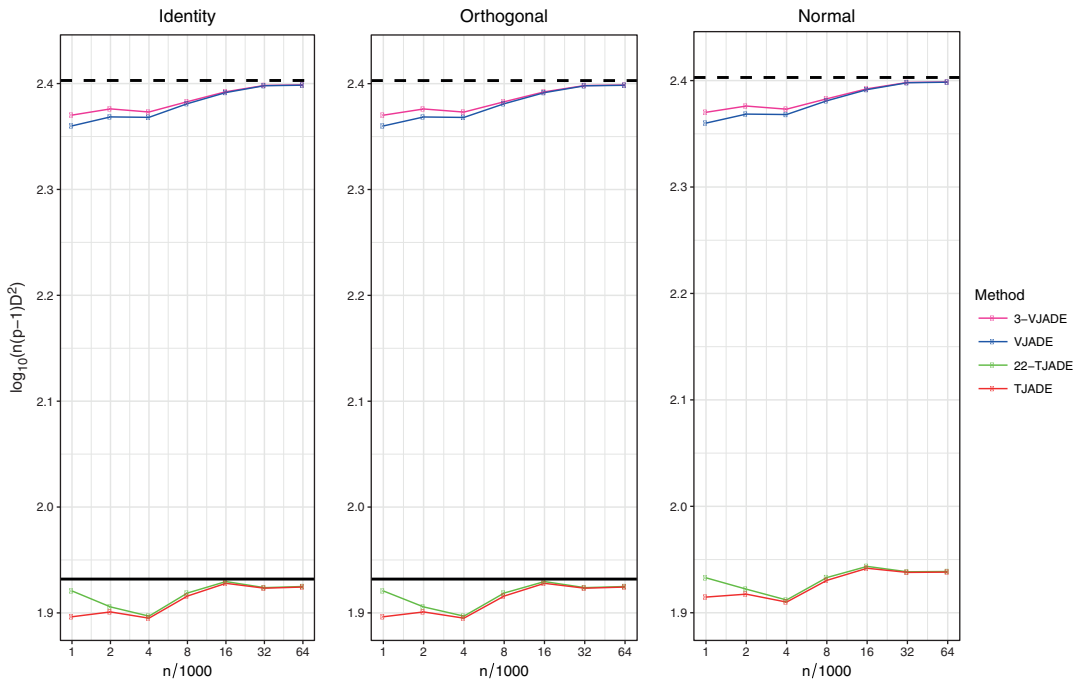


FIGURE 3 Means of the transformed minimum distance indices. The solid black line is the limiting value of TJADE (under orthogonal mixing) and the dashed black line is the limiting value of VJADE. The solid line is missing from the third panel as the limiting value of TJADE is not known under normal mixing [Colour figure can be viewed at wileyonlinelibrary.com]

converge to the theoretic values and that the k -TJADE versions have the same limiting values as TJADE.

We next, under the same setting as above, present the results for seven additional procedures that violate some of the required assumptions: VFOBI, TFOBI, 1-VJADE, 2-VJADE, 11-TJADE, 12-TJADE, and 21-TJADE. Note that TFOBI is the initial step in the k -TJADE procedure and hereby the comparison between k -TJADE and TFOBI illustrates the added benefit of the additional rotation after the TFOBI-solution. Likewise, the same holds between VFOBI and k -VJADE. The resulting mean values of the transformed MD indices are presented in Figure 4, where methods that have almost identical performance are presented using the same colors.

In Figure 4, TFOBI and VFOBI diverge at an exponential rate. Furthermore, the (k_1, k_2) -TJADE procedures that have either k_1 or k_2 less than the number of distinct kurtosis values, are not converging to anything reasonable, even at sample sizes greater than 4×10^6 . However, in this simulation setting, the k -VJADE procedure seems to allow slight deviations from the required assumptions. It seems that 2-VJADE converges to an asymptotic value that is at least close to that of VJADE (see Section 6 for further discussion).

To summarize this part of the simulation study: k -TJADE works as expected, when the theoretical conditions are met, and the convergence to the asymptotic value is relatively fast. Furthermore, even though k -TJADE is not affine invariant, its performance is better under all mixing scenarios, when compared with the affine invariant vectorial counterparts.

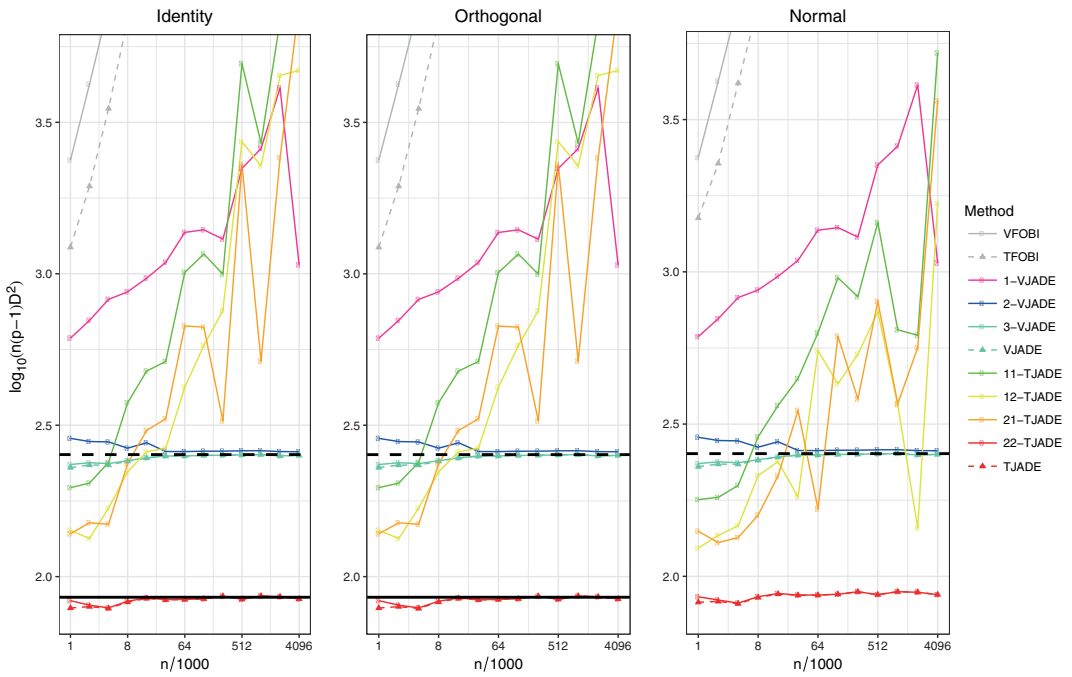


FIGURE 4 Means of the transformed minimum distance indices. The solid black line is the limiting value of TJADE (under orthogonal mixing) and the dashed black line is the limiting value of VJADE. The solid line is missing from the third panel as the limiting value of TJADE is not known under normal mixing. [Colour figure can be viewed at wileyonlinelibrary.com]

5.2 | Finite-sample efficiency: Setting 2

In our second simulation, we illustrate unmixing under a tensorial setting. We consider simulated collections of i.i.d. tensors of size $3 \times 3 \times 4$, $\mathcal{L} := \{\mathcal{L}_j\}_{j \in \{1, \dots, n\}}$. Let $\mathbf{z}_{jk}^{(3)}$ denote the k th 3-mode face of \mathcal{L}_j . The components of every \mathcal{L}_j are then simulated independently from the following distributions,

$$\mathbf{z}_{j1}^{(3)} = \begin{pmatrix} \text{E} & \text{N} & \text{N} \\ \text{N} & \text{U} & \text{N} \\ \text{N} & \text{N} & \text{E} \end{pmatrix}, \quad \mathbf{z}_{j2}^{(3)} = \begin{pmatrix} \text{E} & \text{N} & \text{N} \\ \text{N} & \text{U} & \text{N} \\ \text{N} & \text{N} & \text{E} \end{pmatrix}, \quad \mathbf{z}_{j3}^{(3)} = \begin{pmatrix} \text{E} & \text{N} & \text{N} \\ \text{N} & \text{U} & \text{N} \\ \text{N} & \text{N} & \text{E} \end{pmatrix}, \quad \mathbf{z}_{j4}^{(3)} = \begin{pmatrix} \text{N} & \text{U} & \text{E} \\ \text{E} & \text{E} & \text{E} \\ \text{E} & \text{E} & \text{N} \end{pmatrix},$$

where the different distributions are denoted as in Section 5.1.

All of the mean kurtoses over the different tensor faces are nonzero and none of the theoretical mean kurtoses are the same for the 1-mode faces. Moreover, two of them are the same for the 2-mode faces and three of them are the same for the 3-mode faces. Hereby, the preferable (k_1, k_2, k_3) -TJADE here is 123-TJADE, $k_1 = 1, k_2 = 2, k_3 = 3$. The vectorized versions of the observations contain several normal components and thus the assumptions for the vectorial methods are not satisfied here.

The simulation was performed using 11 different sample sizes, $n_i = 2^{i-1} \times 1000, i \in \{1, \dots, 11\}$ and the simulation was repeated 2000 times for each sample size. We considered the same three mixing scenarios as in Section 5.1 and generated the mixing matrices in the same




FIGURE 5 Means of the transformed minimum distance indices. The solid black line is the limiting value of TJADE (under orthogonal mixing) [Colour figure can be viewed at wileyonlinelibrary.com]

way, with the distinction that here we have three mixing matrices instead of two. We performed the unmixing using VFOBI, TFOBI, 1-VJADE, 2-VJADE, VJADE, TJADE and 11 different versions of k -TJADE. The resulting mean values of the transformed MD index, $\ell^{-1} \sum_{j=1}^{\ell} n_i(\rho - 1)D(\hat{G}_j)^2$, where $\ell = 2000$ and $\rho = 36$, are presented in Figure 5. The orthogonal mixing is omitted from Figure 5, since the tensorial methods are orthogonally invariant and the vectorial methods are affine invariant. The performance curves under the orthogonal mixing would be identical to those under the identity mixing, similarly as in Section 5.1.

The k -TJADE performs as expected for the values of k_1, k_2, k_3 that satisfy the assumptions and the convergence toward the theoretical limit is considerably fast. The vectorial methods fail completely in this example. Interestingly, k -TJADE has relatively nice performance when the elementwise deviation between (k_1, k_2, k_3) and $(1, 2, 3)$ is not too large. See Section 6 for further discussion.

5.3 | Finite-sample efficiency: Setting 3

In the third simulation setting, we have a collection of i.i.d. matrices $\mathbf{Z} := \{\mathbf{Z}_j\}_{j \in \{1, \dots, n\}}$, with components simulated independently from the following distributions,

$$\mathbf{Z}_j = \begin{pmatrix} G(0.9999) & G(1) & G(0.9) \\ G(0.9) & G(0.9998) & G(1) \\ G(0.9) & G(1) & G(1) \end{pmatrix},$$




FIGURE 6 Means of the transformed minimum distance indices. The solid black line is the limiting value of TJADE (under orthogonal mixing) and the dashed black line is the limiting value of VJADE [Colour figure can be viewed at wileyonlinelibrary.com]

where $G(\alpha)$ denotes the gamma distribution, with shape parameter α and rate parameter 1, scaled to have zero mean and unit variance. The theoretical row kurtosis means are approximately,

$$(6.222422 \ 6.222622 \ 6.222222) ,$$

and the theoretical column kurtosis means are approximately,

$$(6.444644 \ 6.000400 \ 6.222222) .$$

Hereby, in this simulation setting, the row kurtosis means are very close to each other. They start to differ at the fourth decimal. Also, the column kurtosis means are quite close to each other.

The simulation study was conducted exactly as in Section 5.1, with the distinction that the largest sample size of Section 5.1 was omitted here. Furthermore, we had more versions of k -TJADE in this study. The results are displayed in Figure 6.

Theoretically, the preferable (k_1, k_2) -TJADE would be 11-TJADE, $k_1 = 1, k_2 = 1$. However, with the current sample sizes, 11-TJADE, 12-TJADE, and 13-TJADE do not seem to converge (see Figure 6). Similarly, 21-TJADE and 31-TJADE exhibit erratic behavior with small sample sizes. Yet, with sample sizes over one million, the corresponding procedures are very close to the theoretical limiting value, given as the solid black line in Figure 6. It seems that with sample sizes of magnitude 10^6 , differences in column kurtosis means that are of magnitude 10^{-1} , are not a problem for the procedures. Conversely, the current sample sizes are not large enough to compensate the small differences in the row kurtosis means.

The curves corresponding 22-TJADE, 23-TJADE, 32-TJADE, and TJADE behave nicely in Figure 6, even with relatively small sample sizes. Especially, the good behavior of 22-TJADE is a little surprising. The corresponding procedure assumes that at most two of the column kurtosis means and at most two of the row kurtosis means are equal. We wish to emphasize that the assumptions required for the k -TJADE procedures are only sufficient conditions and it can be that in some special cases, for example, under some special distributional structures, some of the procedures behave well, even though the sufficient assumptions are not satisfied.

In Figure 6, the curves that correspond to VFOBI and TFOBI are again only visible in the top left corner. This is unsurprising since the required assumptions are not satisfied for VFOBI, and TFOBI cannot handle the close kurtosis means, even though they are distinct on the population level. With the current sample sizes, 1-VJADE does not seem to converge to the theoretical limiting value, given as the dashed black line. The vectorial 2-VJADE, 3-VJADE, and VJADE seem to function quite well in this setup. This could be explained by the fact that the vectorial methods have assumptions directly related to the theoretical kurtoses of the components, whereas the tensorial methods have assumptions related to means of the kurtoses. It is easier to detect small kurtosis differences, when compared with detecting differences in their means.

5.4 | Timing comparison

The results from Sections 5.1 and 5.2 illustrate that the performances of TJADE and suitably chosen versions of k -TJADE are very similar. Next, we quantify the significant improvement in computational speed that k -TJADE provides when compared with TJADE. In the timing comparison, we consider a simulated collection of i.i.d. matrices of size $3 \times q$, $\mathbf{Z} := \{\mathbf{Z}_j\}_{j \in \{1, \dots, n\}}$, such that components of every \mathbf{Z}_j are simulated independently from the following distributions,

$$\mathbf{Z}_j = \begin{pmatrix} \chi_1^2 & \chi_4^2 & \cdots & \chi_{3(q-1)+1}^2 \\ \chi_2^2 & \chi_5^2 & \cdots & \chi_{3(q-1)+2}^2 \\ \chi_3^2 & \chi_6^2 & \cdots & \chi_{3(q-1)+3}^2 \end{pmatrix},$$

where χ_ν^2 denotes the χ^2 -squared distribution with ν degrees of freedom, and the width q of the matrix is the varying parameter in this simulation setting. We used parameter values $q = 5, 10, 15, 20, \dots, 50$ and the sample size $n = 1000$. We considered the same procedures as in Section 5.1: VFOBI, TFOBI, 1-VJADE, 2-VJADE, 3-VJADE, VJADE, 11-TJADE, 12-TJADE, 21-TJADE, 22-TJADE, TJADE and recorded the mean running times over a total of five iterations. The time it took R to vectorize the tensors was also considered as a part of the vectorized procedures. However, the time the vectorizing took, was negligible. We used two alternative stopping criteria for the methods that involve joint diagonalization, that is, for all the methods except for TFOBI and VFOBI. A single iteration was stopped, if either the converge tolerance of the Jacobi rotation based joint diagonalization was less than 10^{-6} or if the required tolerance was not satisfied after 100 iterations.

The average running times in minutes as a function of the dimension q are shown in Figure 7 and methods that have almost identical computing times, are presented using the same colors. Figure 7 clearly illustrates the superior computation speed of k -TJADE, when compared with either TJADE or any of the vectorized counterparts. The timing comparison was conducted on Ubuntu 16.04.4 LTS with Intel[®] Xeon[®] CPU E3-1230 v5 with 3.40 GHz and 64 GB.




FIGURE 7 The logarithms of the mean computation times (in minutes) as a function of the dimension q [Colour figure can be viewed at wileyonlinelibrary.com]

5.5 | Video example

We applied k -TJADE to the WaterSurface surveillance video (Li, Huang, Gu, & Tian, 2004) that is viewable at <http://pages.cs.wisc.edu/~jiaxu/projects/gosus/supplement/> and available to download as an .Rdata file at <https://wis.kuleuven.be/stat/robust/Programs/DO/do-video-data-rdata>. The video has already been used as an example for blind source separation (Virta & Nordhausen, 2017a, 2017b). Each frame of the video is of size $h \times w \times 3$ with the height $h = 128$, width $w = 160$, and a three-variate color channel (RGB). The total video consists of 633 such frames, making our data a sample of size $n = 633$ of random third-order tensors in $\mathcal{R}^{h \times w \times 3}$. The data constituting a single continuous surveillance video, the observations are naturally not independent and the assumptions of tensorial ICA are not fully satisfied. However, ICA is known to be robust against deviations from the independence assumption and applying it to sets of dependent data with success is common practice. We thus expect k -TJADE to successfully extract components of interest from our current data.

The video shows a beach scene with little to no motion until frame 480, when a man enters the scene from the left, staying in the picture for the remainder of the video. Figure 8 shows frames 50 and 550 of the video, illustrating moments before and after the man enters the scene. Our objective with the surveillance video is to find low-dimensional components that allow us to pinpoint the most obvious change point in the video, namely, the man's entrance. As change points are most likely captured in the data as outliers of some form, it is natural to seek the component of interest among those with high kurtosis values.

We proceed as follows: we run k -TJADE on the data with different choices of the tuning parameters, find the component with the highest absolute kurtosis, and plot its time course to



FIGURE 8 Frames 50 and 550 of the surveillance video data [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 9 The source with the highest absolute kurtosis found from the video data using $(1, 1, 0)$ -TJADE. The plot caption refers to the indices of the source in the extracted matrix \mathbf{Z}

visually assess whether it captures the change point. The component found with $(1, 1, 0)$ -TJADE is shown in Figure 9, where $k_3 = 0$ means that we do not unmix the supposedly uninformative color dimension at all. The time series is instantly seen to capture the desired time point as the spike coincides with the first frames the man spends in the scene. The running time of the method was 39 min on a Windows server with Intel[®] Xeon[®] CPU R5 2440 with 2.40 GHz and 64 GB. Applying $(2, 2, 0)$ -TJADE gave almost identical results with the increased running time of 1 hr and 54 min. However, the original TJADE proved to be very slow. The running time of the algorithm was over 5 days. Concluding, the example shows that k -TJADE can be used to reliably extract information from data where TJADE cannot be applied due to its extremely high-computational cost. In several real-world applications, for example, in crime scene investigation, waiting for days is not an option.

5.6 | Estimation of the maximal kurtosis multiplicity ν

In the previous video example, applying $(1, 1, 0)$ -TJADE already gave satisfactory results in the sense that it allowed us to find a latent component identifying the point of interest in time. However, in some cases, using the preferable value $k = \nu$ for the tuning parameter might be justified.

This will guarantee the successful estimation of all components, while simultaneously keeping computation times and the amount of noise in the estimation minimal. The results of Theorem 3 could possibly be used to formulate an asymptotic hypothesis test for the null hypothesis that $\nu \leq \nu_0$ for some given ν_0 in a given mode. Namely, if the null hypothesis holds, then the limiting distributions of the unmixing estimates $\hat{\Gamma}^{\nu_0}, \hat{\Gamma}^{\nu_0+1}, \dots, \hat{\Gamma}^p$ are identical, allowing us to pinpoint the true value of ν . In the following, we will pursue this idea from an empirical perspective and develop a heuristic procedure for the estimation of ν .

Consider a fixed mode of size p of a sample of data $\mathcal{X}_1, \dots, \mathcal{X}_n$ from the tensorial IC model in Equation (5) and let $\hat{\Gamma}^1, \hat{\Gamma}^2, \dots, \hat{\Gamma}^p$ be the unmixing matrices estimated from the corresponding data, respectively, to k -TJADE with the value of $k \in \{1, \dots, p\}$. Letting ν be the true maximal kurtosis multiplicity in the data, we thus expect $\hat{\Gamma}^\nu$ to separate the data well but $\hat{\Gamma}^{\nu-1}$ to leave some parts still unmixed (within the multiplicities). Analogously, we expect $m_{\nu-1} := D(\hat{\Gamma}^{\nu-1}(\hat{\Gamma}^\nu)^{-1})$, where $D(\cdot)$ is the MD index in Equation (7), to be large as the two matrices should have structures differing beyond row permutations, scaling, and sign changes (in the extreme case where $\hat{\Gamma}^{\nu-1} = \hat{\Gamma}^\nu$ we have $m_{\nu-1} = 0$). On the other hand, we expect all the following sequential MD indices, $m_{\nu+\ell}$, $\ell \in \{0, \dots, p - \nu - 1\}$, to be small as all values of $k \in \{\nu, \dots, p\}$ are sufficient for the separation and yield asymptotically equal solutions.

The true value of ν can now be located by plotting m_k versus $k \in \{1, \dots, p - 1\}$ and finding the value, starting from which the curve stays roughly constant. However, since using a single sequential MD index per value of k might leave us with a curve that is not smooth enough to accurately distinguish the true ν , we compute, for each $k \in \{1, \dots, p - 1\}$, the average of the forward sequential MD indices:

$$m_k^* := \frac{1}{p-k} \sum_{\ell=1}^{p-k} D\left(\hat{\Gamma}^k (\hat{\Gamma}^{k+\ell})^{-1}\right).$$

The plot of m_k^* versus k can be used to find the true value of ν similarly as with the plot of m_k versus k and has the interpretation of quantifying the efficiency loss one encounters when using a particular value of k , relative to the choice $k = p$. Note also that the choice of k can be done individually for each mode in the sense that the value of m_k^* in (k, k_0) -TJADE is invariant to the choice of k_0 , and similarly for higher order data.

To illustrate the proposed tool, we consider the handwritten digit dataset from the R-package *ElemStatLearn* (Halvorsen, 2015). The dataset consists of 7291 images of handwritten numerals 0–9 digitized to 16×16 matrices, with each element in $[-1, 1]$, describing the gray scale intensity of the corresponding pixel. The underlying objective of the dataset is to build a classifier that can correctly identify the digits that each of the matrices represent. For simplicity, we work on a subset of 400 randomly chosen images of the digits 1 and 7. These two digits were chosen as their visual similarity makes the classification task more difficult. In this kind of context, ICA is commonly used as a preprocessing step to reduce the data into a low-dimensional subspace, which hopefully contains all the classification information, simplifying the task for the subsequent classifier. In this spirit, we apply the proposed estimation strategy for ν to the sample and obtain the two plots shown in Figure 10. The plots should be interpreted similarly as the scree plot in PCA, where the aim is to find an “elbow” where the slope changes quickly. The two curves seem to imply that the correct parameter for k -TJADE is roughly $(4, 3)$. In Figure 11, we plot the components with the indices $(16, 16)$ and $(16, 15)$ of the $(4, 3)$ -TJADE solution. These components were chosen as k -TJADE orders, the rows/columns in descending order according to their mean kurtoses,




FIGURE 10 The plots of the sequential minimum distance index means m_k^* versus the tuning parameter value k . The left plot is for the rows and the right plot for the columns of the digit data




FIGURE 11 The scatter plot of two particular low-kurtosis-independent components found by (4, 3)-TJADE from the digit data [Colour figure can be viewed at wileyonlinelibrary.com]

and low kurtosis is often a sign of bimodality, making it natural to search for components with classifying ability in the lower right corner of \mathbf{Z} . Indeed, the chosen components reveal a clear separation of the digits. All ones are concentrated roughly in one spherical cluster, with a single outlying seven inside. The bulk of the sevens lies around the cluster of ones in a curved manner and the location of an individual seven is based on the angle of its stem. The analysis could further be continued from here by fitting a classifier, for example, a support vector machine, to the obtained components.

We conclude with two remarks. First, experimenting reveals (not shown here) that also, for example, (2, 2)-TJADE finds similar separation of the two groups as is shown in Figure 11. There are two explanations for this: On one hand, locating the elbows in the scree plots in Figure 10 is visually difficult for a low sample size, and thus it could be that (2, 2) is the actual optimal choice. On the other hand, whatever the true value of ν is, using k larger than that is sure to guarantee that all independent components are estimated (asymptotically) correctly. However, it could be that the group separation structure is such a strong feature of the data that it can be found also with a suboptimal choice of k , and that using the optimal value of k just guarantees that also the 254 other (possibly less interesting) components are recovered successfully.

Second, the suggested procedure is somewhat excessive, as it requires the computation of the k -TJADE solution for a wide range of values of k . In addition, after having computed the solutions, one option would be to simply continue the analysis using the solution with the largest value of k and ignore the estimation altogether. However, when dealing with very low sample sizes, it could happen that using a too large k , although asymptotically sufficient for the estimation of all components, induces noise in the estimation. Still, the suggested procedure could prove useful in studies where a pilot/training dataset is used to determine the optimal value of k , which is then used for later datasets, saving computation time for all other datasets at the expense of the first. Finally, one option to save computation time would be to compute the values of m_k^* sequentially, increasing k one-by-one, and stopping when a significant drop in the plot is visible.

6 | DISCUSSION

We proposed a speed-up version of TJADE, the method with the lowest limiting variance among the currently studied methods of tensorial ICA. Under easily interpretable additional assumptions, the extension, k -TJADE, achieves the same limiting variance to TJADE, while simultaneously exhibiting significantly lower computational cost. A large part of this efficiency is preserved also for samples of finite size.


An interesting future research question is to derive the theoretic behavior of k -TJADE when Assumption 2(v) is violated. Based on our simulations, even when the value of k is chosen to be too small or too large, k -TJADE and k -JADE can still work. This can be seen as a safety net for the users of k -TJADE. The simulations suggest that the performance deteriorates the further down one goes from the optimal k .

ACKNOWLEDGEMENTS

The work of J.V. was supported by the Academy of Finland (grant 321883). The work of N.L. was supported by the Emil Aaltonen Foundation (grant 170156 N). The work of K.N. was supported by Austrian Science Fund (FWF) Grant number P31881-N32 and CRoNoS COST Action IC1408. All authors acknowledge the computational resources provided by the Aalto Science-IT project. The authors wish to express their gratitude to the two anonymous referees, whose insightful comments greatly helped improving the quality of the manuscript.

ORCID

Joni Virta  <https://orcid.org/0000-0002-2150-2769>

Klaus Nordhausen  <https://orcid.org/0000-0002-3758-8501>

REFERENCES

- Belouchrani, A., Abed-Meraim, K., Cardoso, J.-F., & Moulines, E. (1997). A blind source separation technique using second-order statistics. *IEEE Transactions on Signal Processing*, 45, 434–444.
- Cardoso, J.-F. (1989). *Source separation using higher order moments*. Paper presented at the International Conference on acoustics speech, IEEE. 2109–2112.
- Cardoso, J.-F., & Soudoumiac, A. (1993). *Blind beamforming for non-Gaussian signals*. In *IEE proceedings F - Radar and Signal Processing* (Vol. 140, pp. 362–370). IET Digital Library. <https://dx.doi.org/10.1049/ip-f-2.1993.0054>
- Cardoso, J.-F. (1999). High-order contrasts for independent component analysis. *Neural Computation*, 11, 157–192.

- Cichocki, A., Mandic, D., De Lathauwer, L., Zhou, G., Zhao, Q., Caiafa, C., & Phan, H. A. (2015). Tensor decompositions for signal processing applications: From two-way to multiway component analysis. *IEEE Signal Processing Magazine*, 32, 145–163.
- Cichocki, A., Zdunek, R., Phan, A. H., & Amari, S.-I. (2009). *Nonnegative matrix and tensor factorizations: Applications to exploratory multi-way data analysis and blind source separation*. Hoboken, NJ: John Wiley & Sons.
- Comon, P., & Jutten, C. (2010). *Handbook of blind source separation: Independent component analysis and applications*. London, UK: Academic Press.
- De Lathauwer, L., De Moor, B., & Vandewalle, J. (2000). A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications*, 21, 1253–1278.
- Halvorsen, K. B. (2015) ElemStatLearn: Data sets, functions and examples from the book: “The Elements of Statistical Learning, Data Mining, Inference, and Prediction” by Trevor Hastie, Robert Tibshirani and Jerome Friedman. R package version 2015.6.26. Retrieved from <https://CRAN.R-project.org/package=ElemStatLearn>.
- Hyvärinen, A., Karhunen, J., & Oja, E. (2004). *Independent component analysis*. Hoboken, NJ: John Wiley & Sons.
- Illner, K., Miettinen, J., Fuchs, C., Taskinen, S., Nordhausen, K., Oja, H., & Theis, F. J. (2015). Model selection using limiting distributions of second-order blind source separation algorithms. *Signal Processing*, 113, 95–103.
- Ilmonen, P., Nordhausen, K., Oja, H., & Ollila, E. (2010). A new performance index for ICA: Properties, computation and asymptotic analysis. *Lecture Notes in Comput Science*, 6365, 229–236.
- Ilmonen, P., & Paindaveine, D. (2011). Semiparametrically efficient inference based on signed ranks in symmetric independent component models. *Annals of Statistics*, 39, 2448–2476.
- Kolda, T. G., & Bader, B. W. (2009). Tensor decompositions and applications. *SIAM Review*, 51, 455–500.
- Leng, C., & Pan, G. (2018). Covariance estimation via sparse Kronecker structures. *Bernoulli*, 24, 3833–3863.
- Li, L., Huang, W., Gu, I. Y.-H., & Tian, Q. (2004). Statistical modeling of complex backgrounds for foreground object detection. *IEEE Transactions on Image Processing*, 13, 1459–1472.
- Miettinen, J., Nordhausen, K., Oja, H. & Taskinen, S. (2013) *Fast equivariant JADE*. Paper presented at the International Conference on Acoustics, Speech and signal processing, IEEE, 6153–6157.
- Miettinen, J., Nordhausen, K., & Taskinen, S. (2017). Blind source separation based on joint diagonalization in R: The packages JADE and BSSasyp. *Journal of Statistical Software*, 76, 1–31.
- Miettinen, J., Taskinen, S., Nordhausen, K., & Oja, H. (2015). Fourth moments and independent component analysis. *Statistical Science*, 30, 372–390.
- Nordhausen, K., & Oja, H. (2018). Independent component analysis: A statistical perspective. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10, e1440.
- Roś, B., Bijma, F., de Munck, J. C., & de Gunst, M. C. (2016). Existence and uniqueness of the maximum likelihood estimator for models with a Kronecker product covariance structure. *Journal of Multivariate Analysis*, 143, 345–361.
- Samarov, A., & Tsybakov, A. (2004). Nonparametric independent component analysis. *Bernoulli*, 10, 565–582.
- Srivastava, M. S., von Rosen, T., & von Rosen, D. (2008). Models with a Kronecker product covariance structure: Estimation and testing. *Mathematical Methods of Statistics*, 17, 357–370.
- Virta, J., Li, B., Nordhausen, K., & Oja, H. (2017). Independent component analysis for tensor-valued data. *Journal of Multivariate Analysis*, 162, 172–192.
- Virta, J., Li, B., Nordhausen, K., & Oja, H. (2018). JADE for tensor-valued observations. *Journal of Computational and Graphical Statistics*, 27, 628–637.
- Virta, J., & Nordhausen, K. (2017a). *Blind source separation for nonstationary tensor-valued time series*. Paper presented at 2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP), Tokyo, Japan, 1–6.
- Virta, J., & Nordhausen, K. (2017b). Blind source separation of tensor-valued time series. *Signal Processing*, 141, 204–216.
- Virta, J., Nordhausen, K., Oja, H., & Li, B. (2016). TensorBSS: Blind source separation methods for tensor-valued observations. R Package Version 0.3.3.
- Virta, J., Taskinen, S., & Nordhausen, K. (2016). *Applying fully tensorial ICA to fMRI data*. Paper presented at IEEE Signal Processing in Medicine and Biology Symposium (SPMB), Philadelphia, PA, 1–6.
- Werner, K., Jansson, M., & Stoica, P. (2008). On estimation of covariance matrices with Kronecker product structure. *IEEE Transactions on Signal Processing*, 56, 478–491.
- Wiesel, A. (2012). Geodesic convexity and covariance estimation. *IEEE Transactions on Signal Processing*, 60, 6182.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Virta J, Lietzén N, Ilmonen P, Nordhausen K. Fast tensorial JADE. *Scand J Statist.* 2021;48:164–187. <https://doi.org/10.1111/sjos.12445>