

Whitening-Free Least-Squares Non-Gaussian Component Analysis

Hiroaki Shiino

HSHIINO@YAHOO-CORP.JP

Tokyo Institute of Technology, Tokyo 152-8550, Japan.

(Current Address: Yahoo Japan Corporation, Tokyo 102-8282, Japan.)

Hiroaki Sasaki

HSASAKI@IS.NAIST.JP

Nara Institute of Science and Technology, Nara 630-0192, Japan.

Gang Niu

GANG@MS.K.U-TOKYO.AC.JP

The University of Tokyo, Chiba 277-8561, Japan.

RIKEN Center for Advanced Intelligence Project, Tokyo 103-0027, Japan.

Masashi Sugiyama

SUGI@K.U-TOKYO.AC.JP

RIKEN Center for Advanced Intelligence Project, Tokyo 103-0027, Japan.

The University of Tokyo, Chiba 277-8561, Japan.

Editors: Yung-Kyun Noh and Min-Ling Zhang

Abstract

Non-Gaussian component analysis (NGCA) is an unsupervised linear dimension reduction method that extracts low-dimensional non-Gaussian “signals” from high-dimensional data contaminated with Gaussian noise. NGCA can be regarded as a generalization of *projection pursuit* (PP) and *independent component analysis* (ICA) to multi-dimensional and dependent non-Gaussian components. Indeed, seminal approaches to NGCA are based on PP and ICA. Recently, a novel NGCA approach called *least-squares NGCA* (LSNGCA) has been developed, which gives a solution analytically through least-squares estimation of *log-density gradients* and eigendecomposition. However, since *pre-whitening* of data is involved in LSNGCA, it performs unreliably when the data covariance matrix is ill-conditioned, which is often the case in high-dimensional data analysis. In this paper, we propose a *whitening-free* variant of LSNGCA and experimentally demonstrate its superiority.

Keywords: non-Gaussian component analysis, unsupervised dimension reduction

1. Introduction

Dimension reduction is a common technique in high-dimensional data analysis to mitigate the *curse of dimensionality* (Vapnik, 1998). Among various approaches to dimension reduction, we focus on unsupervised linear dimension reduction in this paper.

It is known that the distribution of randomly projected data is close to Gaussian (Diederichs et al., 2013). Based on this observation, *non-Gaussian component analysis* (NGCA) (Blanchard et al., 2006) tries to find a subspace that contains non-Gaussian sig-

Table 1: NGCA methods.

	MIPP	IMAK	SNGCA	LSNGCA	WF-LSNGCA (proposed)
Manual NGIF design	Need	No Need	Need	No Need	No Need
Computational efficiency	Efficient (iterative)	Inefficient (iterative)	Inefficient (iterative)	Efficient (analytic)	Efficient (analytic)
Pre-whitening	Need	Need	No Need	Need	No Need

nal components so that Gaussian noise components can be projected out. NGCA is formulated in an elegant *semi-parametric* framework and non-Gaussian components can be extracted without specifying their distributions. Mathematically, NGCA can be regarded as a generalization of *projection pursuit* (PP) (Friedman and Tukey, 1974) and *independent component analysis* (ICA) (Hyvärinen and Oja, 2000) to multi-dimensional and dependent non-Gaussian components.

The first NGCA algorithm is called *multi-index PP* (MIPP). PP algorithms such as FastICA (Hyvärinen and Oja, 2000) use a *non-Gaussian index function* (NGIF) to find either a super-Gaussian or sub-Gaussian component. MIPP uses a family of such NGIFs to find multiple non-Gaussian components and apply principal component analysis (PCA) to extract a non-Gaussian subspace. However, MIPP requires us to prepare appropriate NGIFs, which is not necessarily straightforward in practice. Furthermore, MIPP requires *pre-whitening* of data, which can be unreliable when the data covariance matrix is ill-conditioned.

To cope with these problems, MIPP has been extended in various ways. The method called *iterative metric adaptation for radial kernel functions* (IMAK) (Kawanabe et al., 2007) tries to avoid the manual design of NGIFs by *learning* the NGIFs from data in the form of radial kernel functions. However, this learning part is computationally highly expensive and pre-whitening is still necessary. *Sparse NGCA* (SNGCA) (Diederichs et al., 2010, 2013) tries to avoid pre-whitening by imposing an appropriate constraint so that the solution is independent of the data covariance matrix. However, SNGCA involves *semi-definite programming* which is computationally highly demanding, and NGIFs still need to be manually designed.

Recently, a novel approach to NGCA called *least-squares NGCA* (LSNGCA) has been proposed (Sasaki et al., 2016). Based on the *gradient* of the log-density function, LSNGCA constructs a vector that belongs to the non-Gaussian subspace from each sample. Then the method of *least-squares log-density gradients* (LSLDG) (Cox, 1985; Sasaki et al., 2014) is employed to directly estimate the log-density gradient without density estimation. Finally, the principal subspace of the set of vectors generated from all samples is extracted by eigendecomposition. LSNGCA is computationally efficient and no manual design of NGIFs is involved. However, it still requires pre-whitening of data.

The existing NGCA methods reviewed above are summarized in Table 1. In this paper, we propose a novel NGCA method that is computationally efficient, no manual design of NGIFs is involved, and no pre-whitening is necessary. Our proposed method is essentially an extension of LSNGCA so that the covariance of data is implicitly handled without explicit pre-whitening or explicit constraints. Through experiments, we demonstrate that

our proposed method, called *whitening-free* LSNGCA (WF-LSNGCA), performs very well even when the data covariance matrix is ill-conditioned.

2. Non-Gaussian Component Analysis

In this section, we formulate the problem of NGCA and review the MIPP and LSNGCA methods.

2.1. Problem Formulation

Suppose that we are given a set of d -dimensional i.i.d. samples of size n , $\{\mathbf{x}_i | \mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$, which are generated by the following model:

$$\mathbf{x}_i = \mathbf{A}\mathbf{s}_i + \mathbf{n}_i, \quad (1)$$

where $\mathbf{s}_i \in \mathbb{R}^m$ ($m \leq d$) is an m -dimensional signal vector independently generated from an unknown non-Gaussian distribution (we assume that m is known), $\mathbf{n}_i \in \mathbb{R}^d$ is a noise vector independently generated from a centered Gaussian distribution with an unknown covariance matrix \mathbf{Q} , and $\mathbf{A} \in \mathbb{R}^{d \times m}$ is an unknown mixing matrix of rank m . Under this data generative model, probability density function $p(\mathbf{x})$ that samples $\{\mathbf{x}_i\}_{i=1}^n$ follow can be expressed in the following semi-parametric form (Blanchard et al., 2006):

$$p(\mathbf{x}) = f(\mathbf{B}^\top \mathbf{x}) \phi_{\mathbf{Q}}(\mathbf{x}), \quad (2)$$

where f is an unknown smooth positive function on \mathbb{R}^m , $\mathbf{B} \in \mathbb{R}^{d \times m}$ is an unknown linear mapping, $\phi_{\mathbf{Q}}$ is the centered Gaussian density with the covariance matrix \mathbf{Q} , and $^\top$ denotes the transpose. We note that decomposition (2) is not unique; multiple combinations of \mathbf{B} and f can give the same probability density function. Nevertheless, the following m -dimensional subspace E , called the *non-Gaussian index space*, can be determined uniquely (Theis and Kawanabe, 2006):

$$E = \text{Null}(\mathbf{B}^\top)^\perp = \text{Range}(\mathbf{B}), \quad (3)$$

where $\text{Null}(\mathbf{B}^\top)$ denotes the null space of \mathbf{B}^\top , $^\perp$ denotes the orthogonal complement, and $\text{Range}(\mathbf{B})$ denotes the column space of \mathbf{B} .

The goal of NGCA is to estimate the non-Gaussian index space E from samples $\{\mathbf{x}_i\}_{i=1}^n$.

2.2. Multi-Index Projection Pursuit (MIPP)

MIPP (Blanchard et al., 2006) is the first algorithm of NGCA.

Let us *whiten* the samples $\{\mathbf{x}_i\}_{i=1}^n$ so that their covariance matrix becomes identity:

$$\mathbf{y}_i := \mathbf{\Sigma}^{-\frac{1}{2}} \mathbf{x}_i,$$

where $\mathbf{\Sigma}$ is the covariance matrix of \mathbf{x} . In practice, $\mathbf{\Sigma}$ is replaced by the sample covariance matrix. Then, for an NGIF h , the following vector $\boldsymbol{\beta}(h)$ was shown to belong to the non-Gaussian index space E (Blanchard et al., 2006):

$$\boldsymbol{\beta}(h) := \mathbb{E}[\mathbf{y}h(\mathbf{y}) - \nabla_{\mathbf{y}}h(\mathbf{y})],$$

where $\nabla_{\mathbf{y}}$ denotes the differential operator w.r.t. \mathbf{y} and $\mathbb{E}[\cdot]$ denotes the expectation over $p(\mathbf{x})$. MIPP generates a set of such vectors from various NGIFs $\{h_l\}_{l=1}^L$:

$$\widehat{\boldsymbol{\beta}}_l := \frac{1}{n} \sum_{i=1}^n [\mathbf{y}_i h_l(\mathbf{y}_i) - \nabla_{\mathbf{y}} h_l(\mathbf{y}_i)], \quad (4)$$

where the expectation is estimated by the sample average. Then $\widehat{\boldsymbol{\beta}}_l$ is normalized as

$$\widehat{\boldsymbol{\beta}}_l \leftarrow \widehat{\boldsymbol{\beta}}_l / \sqrt{\frac{1}{n} \sum_{i=1}^n \|\mathbf{y}_i h_l(\mathbf{y}_i) - \nabla_{\mathbf{y}} h_l(\mathbf{y}_i)\|^2 - \|\widehat{\boldsymbol{\beta}}_l\|^2}, \quad (5)$$

by which $\|\widehat{\boldsymbol{\beta}}_l\|$ is proportional to its signal-to-noise ratio. Then vectors $\widehat{\boldsymbol{\beta}}_l$ with their norm less than a pre-specified threshold $\tau > 0$ are eliminated. Finally, PCA is applied to the remaining vectors $\widehat{\boldsymbol{\beta}}_l$ to obtain an estimate of the non-Gaussian index space E .

The behavior of MIPP strongly depends on the choice of NGIF h . To improve the performance, MIPP actively searches informative h as follows. First, the form of h is restricted to $h(\mathbf{y}) = s(\mathbf{w}^\top \mathbf{y})$, where $\mathbf{w} \in \mathbb{R}^d$ denotes a unit-norm vector and s is a smooth real function. Then, estimated vector $\widehat{\boldsymbol{\beta}}$ is written as

$$\widehat{\boldsymbol{\beta}} = \frac{1}{n} \sum_{i=1}^n \left(\mathbf{y}_i s(\mathbf{w}^\top \mathbf{y}_i) - s'(\mathbf{w}^\top \mathbf{y}_i) \mathbf{w} \right),$$

where s' is the derivative of s . This equation is actually equivalent to a single iteration of the PP algorithm called *FastICA* (Hyvärinen, 1999). Based on this fact, the parameter \mathbf{w} is optimized by iteratively applying the following update rule until convergence:

$$\mathbf{w} \leftarrow \frac{\sum_{i=1}^n (\mathbf{y}_i s(\mathbf{w}^\top \mathbf{y}_i) - s'(\mathbf{w}^\top \mathbf{y}_i) \mathbf{w})}{\left\| \sum_{i=1}^n (\mathbf{y}_i s(\mathbf{w}^\top \mathbf{y}_i) - s'(\mathbf{w}^\top \mathbf{y}_i) \mathbf{w}) \right\|}.$$

The superiority of MIPP has been investigated both theoretically and experimentally (Blanchard et al., 2006). However, MIPP has the weaknesses that NGIFs should be manually designed and pre-whitening is necessary.

2.3. Least-Squares Non-Gaussian Component Analysis (LSNGCA)

LSNGCA (Sasaki et al., 2016) is a recently proposed NGCA algorithm that does not require manual design of NGIFs (Table 1). Here the algorithm of LSNGCA is reviewed, which will be used for further developing a new method in the next section.

Derivation: For whitened samples $\{\mathbf{y}_i\}_{i=1}^n$, the semi-parametric form of NGCA given in Eq.(2) can be simplified as

$$p(\mathbf{y}) = \widetilde{f}(\widetilde{\mathbf{B}}^\top \mathbf{y}) \phi_{\mathbf{I}_d}(\mathbf{y}), \quad (6)$$

where \widetilde{f} is an unknown smooth positive function on \mathbb{R}^m and $\widetilde{\mathbf{B}} \in \mathbb{R}^{d \times m}$ is an unknown linear mapping. Under this simplified semi-parametric form, the non-Gaussian index space E can be represented as

$$E = \boldsymbol{\Sigma}^{-\frac{1}{2}} \text{Range}(\widetilde{\mathbf{B}}).$$

Taking the logarithm and differentiating the both sides of Eq.(6) w.r.t. \mathbf{y} yield

$$\nabla_{\mathbf{y}} \ln p(\mathbf{y}) + \mathbf{y} = \tilde{\mathbf{B}} \nabla_{\tilde{\mathbf{B}}^\top \mathbf{y}} \ln \tilde{f}(\tilde{\mathbf{B}}^\top \mathbf{y}), \quad (7)$$

where $\nabla_{\tilde{\mathbf{B}}^\top \mathbf{y}}$ denotes the differential operator w.r.t. $\tilde{\mathbf{B}}^\top \mathbf{y}$. This implies that $\mathbf{u}(\mathbf{y}) := \nabla_{\mathbf{y}} \ln p(\mathbf{y}) + \mathbf{y}$ belongs to the non-Gaussian index space E . Then applying eigendecomposition to $\sum_{i=1}^n \mathbf{u}(\mathbf{y}_i) \mathbf{u}(\mathbf{y}_i)^\top$ and extracting the m leading eigenvectors allow us to recover $\text{Range}(\tilde{\mathbf{B}})$. In LSNGCA, the method of *least-squares log-density gradients* (LSLDG) (Cox, 1985; Sasaki et al., 2014) is used to estimate the log-density gradient $\nabla_{\mathbf{y}} \ln p(\mathbf{y})$ included in $\mathbf{u}(\mathbf{y})$, which is briefly reviewed below.

LSLDG: Let ∂_j denote the differential operator w.r.t. the j -th element of \mathbf{y} . LSLDG fits a model $g^{(j)}(\mathbf{y})$ to $\partial_j \ln p(\mathbf{y})$, the j -th element of log-density gradient $\nabla_{\mathbf{y}} \ln p(\mathbf{y})$, under the squared loss:

$$\begin{aligned} J(g^{(j)}) &:= \mathbb{E}[(g^{(j)}(\mathbf{y}) - \partial_j \ln p(\mathbf{y}))^2] - \mathbb{E}[(\partial_j \ln p(\mathbf{y}))^2] \\ &= \mathbb{E}[g^{(j)}(\mathbf{y})^2] - 2\mathbb{E}[g^{(j)}(\mathbf{y})\partial_j \ln p(\mathbf{y})]. \end{aligned} \quad (8)$$

The second term in Eq.(8) yields

$$\begin{aligned} \mathbb{E}[g^{(j)}(\mathbf{y})\partial_j \ln p(\mathbf{y})] &= \int g^{(j)}(\mathbf{y})(\partial_j \ln p(\mathbf{y}))p(\mathbf{y})d\mathbf{y} = \int g^{(j)}(\mathbf{y})\partial_j p(\mathbf{y})d\mathbf{y} \\ &= - \int \partial_j g^{(j)}(\mathbf{y})p(\mathbf{y})d\mathbf{y} = -\mathbb{E}[\partial_j g^{(j)}(\mathbf{y})], \end{aligned}$$

where the second-last equation follows from *integration by parts* under the assumption $\lim_{|\mathbf{y}^{(j)}| \rightarrow \infty} g^{(j)}(\mathbf{y})p(\mathbf{y}) = 0$. Then sample approximation yields

$$J(g^{(j)}) = \mathbb{E}[g^{(j)}(\mathbf{y})^2 - 2\partial_j g^{(j)}(\mathbf{y})] \approx \frac{1}{n} \sum_{i=1}^n [g^{(j)}(\mathbf{y}_i)^2 + 2\partial_j g^{(j)}(\mathbf{y}_i)]. \quad (9)$$

As a model of the log-density gradient, LSLDG uses a linear-in-parameter form:

$$g^{(j)}(\mathbf{y}) = \sum_{k=1}^b \theta_{k,j} \psi_{k,j}(\mathbf{y}) = \boldsymbol{\theta}_j^\top \boldsymbol{\psi}_j(\mathbf{y}), \quad (10)$$

where b denotes the number of basis functions, $\boldsymbol{\theta}_j := (\theta_{1,j}, \dots, \theta_{b,j})^\top$ is a parameter vector to be estimated, and $\boldsymbol{\psi}_j(\mathbf{y}) := (\psi_{1,j}(\mathbf{y}), \dots, \psi_{b,j}(\mathbf{y}))^\top$ is a basis function vector. The parameter vector $\boldsymbol{\theta}_j$ is learned by solving the following regularized empirical optimization problem:

$$\hat{\boldsymbol{\theta}}_j = \underset{\boldsymbol{\theta}_j}{\operatorname{argmin}} \left[\boldsymbol{\theta}_j^\top \hat{\mathbf{G}}_j \boldsymbol{\theta}_j + 2\boldsymbol{\theta}_j^\top \hat{\mathbf{h}}_j + \lambda_j \|\boldsymbol{\theta}_j\|^2 \right],$$

where $\lambda_j > 0$ is the regularization parameter,

$$\hat{\mathbf{G}}_j = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\psi}_j(\mathbf{y}_i) \boldsymbol{\psi}_j(\mathbf{y}_i)^\top, \quad \hat{\mathbf{h}}_j = \frac{1}{n} \sum_{i=1}^n \partial_j \boldsymbol{\psi}_j(\mathbf{y}_i).$$

Algorithm 1 Pseudo-code of WF-LSNGCA.

input Element-wise standardized data samples: $\{\mathbf{x}_i\}_{i=1}^n$.

- 1: Obtain an estimate $\hat{\mathbf{v}}(\mathbf{x})$ of $\mathbf{v}(\mathbf{x}) = \nabla_{\mathbf{x}} \ln p(\mathbf{x}) - \nabla_{\mathbf{x}}^2 \ln p(\mathbf{x})\mathbf{x}$ by the method described in Section 3.2.
 - 2: Apply eigendecomposition to $\sum_{i=1}^n \hat{\mathbf{v}}(\mathbf{x}_i)\hat{\mathbf{v}}(\mathbf{x}_i)^\top$ and extract the m leading eigenvectors as an orthonormal basis of non-Gaussian index space E .
-

This optimization problem can be analytically solved as

$$\hat{\boldsymbol{\theta}}_j = -\left(\hat{\mathbf{G}}_j + \lambda_j \mathbf{I}_b\right)^{-1} \hat{\mathbf{h}}_j,$$

where \mathbf{I}_b is the b -by- b identity matrix. Finally, an estimator of the log-density gradient $g^{(j)}(\mathbf{y})$ is obtained as

$$\hat{g}^{(j)}(\mathbf{y}) = \hat{\boldsymbol{\theta}}_j^\top \boldsymbol{\psi}_j(\mathbf{y}).$$

All tuning parameters such as the regularization parameter λ_j and parameters included in the basis function $\boldsymbol{\psi}_{k,j}(\mathbf{y})$ can be systematically chosen based on cross-validation w.r.t. Eq.(9).

3. Whitening-Free LSNGCA

In this section, we propose a novel NGCA algorithm that does not involve pre-whitening. A pseudo-code of the proposed method, which we call *whitening-free LSNGCA* (WF-LSNGCA), is summarized in Algorithm 1.

3.1. Derivation

Unlike LSNGCA which used the simplified semi-parametric form (6), we directly use the original semi-parametric form (2) without whitening. Taking the logarithm and differentiating the both sides of Eq.(2) w.r.t. \mathbf{x} yield

$$\nabla_{\mathbf{x}} \ln p(\mathbf{x}) + \mathbf{Q}^{-1}\mathbf{x} = \mathbf{B}\nabla_{\mathbf{B}^\top \mathbf{x}} \ln f(\mathbf{B}^\top \mathbf{x}), \quad (11)$$

where $\nabla_{\mathbf{x}}$ denotes the derivative w.r.t. \mathbf{x} and $\nabla_{\mathbf{B}^\top \mathbf{x}}$ denotes the derivative w.r.t. $\mathbf{B}^\top \mathbf{x}$. Further taking the derivative of Eq.(11) w.r.t. \mathbf{x} yields

$$\mathbf{Q}^{-1} = -\nabla_{\mathbf{x}}^2 \ln p(\mathbf{x}) + \mathbf{B}\nabla_{\mathbf{B}^\top \mathbf{x}}^2 \ln f(\mathbf{B}^\top \mathbf{x})\mathbf{B}^\top, \quad (12)$$

where $\nabla_{\mathbf{x}}^2$ denotes the second derivative w.r.t. \mathbf{x} . Substituting Eq.(12) back into Eq.(11) yields

$$\nabla_{\mathbf{x}} \ln p(\mathbf{x}) - \nabla_{\mathbf{x}}^2 \ln p(\mathbf{x})\mathbf{x} = \mathbf{B}\left(\nabla_{\mathbf{B}^\top \mathbf{x}} \ln f(\mathbf{B}^\top \mathbf{x}) - \nabla_{\mathbf{B}^\top \mathbf{x}}^2 \ln f(\mathbf{B}^\top \mathbf{x})\mathbf{B}^\top \mathbf{x}\right). \quad (13)$$

This implies that $\mathbf{v}(\mathbf{x}) := \nabla_{\mathbf{x}} \ln p(\mathbf{x}) - \nabla_{\mathbf{x}}^2 \ln p(\mathbf{x})\mathbf{x}$ belongs to the non-Gaussian index space E . Then we apply eigendecomposition to $\sum_{i=1}^n \mathbf{v}(\mathbf{x}_i)\mathbf{v}(\mathbf{x}_i)^\top$ and extract the m leading eigenvectors as an orthonormal basis of non-Gaussian index space E .

Now the remaining task is to approximate $\mathbf{v}(\mathbf{x})$ from data, which is discussed below.

3.2. Estimation of $\mathbf{v}(\mathbf{x})$

Let $v^{(j)}(\mathbf{x})$ be the j -th element of $\mathbf{v}(\mathbf{x})$:

$$v^{(j)}(\mathbf{x}) = \partial_j \ln p(\mathbf{x}) - (\nabla_{\mathbf{x}} \partial_j \ln p(\mathbf{x}))^\top \mathbf{x}.$$

To estimate $v^{(j)}(\mathbf{x})$, let us fit a model $w^{(j)}(\mathbf{x})$ to it under the squared loss:

$$\begin{aligned} R(w^{(j)}) &:= \mathbb{E}[(w^{(j)}(\mathbf{x}) - v^{(j)}(\mathbf{x}))^2] - \mathbb{E}[v^{(j)}(\mathbf{x})^2] \\ &= \mathbb{E}[w^{(j)}(\mathbf{x})^2] - 2\mathbb{E}[w^{(j)}(\mathbf{x})v^{(j)}(\mathbf{x})] \\ &= \mathbb{E}[w^{(j)}(\mathbf{x})^2] - 2\mathbb{E}[w^{(j)}(\mathbf{x})\partial_j \ln p(\mathbf{x})] + 2\mathbb{E}[w^{(j)}(\mathbf{x})(\nabla_{\mathbf{x}} \partial_j \ln p(\mathbf{x}))^\top \mathbf{x}]. \end{aligned} \quad (14)$$

The second term in Eq.(14) yields

$$\begin{aligned} \mathbb{E}[w^{(j)}(\mathbf{x})\partial_j \ln p(\mathbf{x})] &= \int w^{(j)}(\mathbf{x})(\partial_j \ln p(\mathbf{x}))p(\mathbf{x})d\mathbf{x} = \int w^{(j)}(\mathbf{x})\partial_j p(\mathbf{x})d\mathbf{x} \\ &= - \int \partial_j w^{(j)}(\mathbf{x})p(\mathbf{x})d\mathbf{x} = -\mathbb{E}[\partial_j w^{(j)}(\mathbf{x})], \end{aligned}$$

where the second-last equation follows from *integration by parts* under the assumption $\lim_{|\mathbf{x}^{(j)}| \rightarrow \infty} w^{(j)}(\mathbf{x})p(\mathbf{x}) = 0$. $\partial_j \ln p(\mathbf{x})$ included in the third term in Eq.(14) may be replaced with the LSLDG estimator $\widehat{g}^{(j)}(\mathbf{x})$ reviewed in Section 2.3. Note that the LSLDG estimator is obtained with non-whitened data \mathbf{x} in this method. Then we have

$$\begin{aligned} R(w^{(j)}) &\approx \mathbb{E}[w^{(j)}(\mathbf{x})^2 + 2\partial_j w^{(j)}(\mathbf{x}) + 2w^{(j)}(\mathbf{x})(\nabla_{\mathbf{x}} \widehat{g}^{(j)}(\mathbf{x}))^\top \mathbf{x}] \\ &\approx \frac{1}{n} \sum_{i=1}^n [w^{(j)}(\mathbf{x}_i)^2 + 2\partial_j w^{(j)}(\mathbf{x}_i) + 2w^{(j)}(\mathbf{x}_i)(\nabla_{\mathbf{x}} \widehat{g}^{(j)}(\mathbf{x}_i))^\top \mathbf{x}_i]. \end{aligned} \quad (15)$$

Here, let us employ the following linear-in-parameter model as $w^{(j)}(\mathbf{x})$:

$$w^{(j)}(\mathbf{x}) := \sum_{k=1}^{b'} \alpha_{k,j} \varphi_{k,j}(\mathbf{x}) = \boldsymbol{\alpha}_j^\top \boldsymbol{\varphi}_j(\mathbf{x}), \quad (16)$$

where b' denotes the number of basis functions, $\boldsymbol{\alpha}_j := (\alpha_{1,j}, \dots, \alpha_{b',j})^\top$ is a parameter vector to be estimated, and $\boldsymbol{\varphi}_j(\mathbf{x}) := (\varphi_{1,j}(\mathbf{x}), \dots, \varphi_{b',j}(\mathbf{x}))^\top$ is a basis function vector. The parameter vector $\boldsymbol{\alpha}_j$ is learned by minimizing the following regularized empirical optimization problem:

$$\widehat{\boldsymbol{\alpha}}_j = \underset{\boldsymbol{\alpha}_j}{\operatorname{argmin}} \left[\boldsymbol{\alpha}_j^\top \widehat{\mathbf{S}}_j \boldsymbol{\alpha}_j + 2\boldsymbol{\alpha}_j^\top \widehat{\mathbf{t}}_j(\mathbf{x}) + \gamma_j \|\boldsymbol{\alpha}_j\|^2 \right],$$

where $\gamma_j > 0$ is the regularization parameter,

$$\begin{aligned} \widehat{\mathbf{S}}_j &= \frac{1}{n} \sum_{i=1}^n \boldsymbol{\varphi}_j(\mathbf{x}_i) \boldsymbol{\varphi}_j(\mathbf{x}_i)^\top, \\ \widehat{\mathbf{t}}_j &= \frac{1}{n} \sum_{i=1}^n \left(\partial_j \boldsymbol{\varphi}_j(\mathbf{x}_i) + \boldsymbol{\varphi}_j(\mathbf{x}_i) \left(\nabla_{\mathbf{x}} \widehat{g}^{(j)}(\mathbf{x}_i) \right)^\top \mathbf{x}_i \right). \end{aligned}$$

This optimization problem can be analytically solved as

$$\hat{\boldsymbol{\alpha}}_j = - \left(\hat{\mathbf{S}}_j + \gamma_j \mathbf{I}_{b'} \right)^{-1} \hat{\mathbf{t}}_j.$$

Finally, an estimator of $v^{(j)}(\mathbf{x})$ is obtained as

$$\hat{v}^{(j)}(\mathbf{x}) = \hat{\boldsymbol{\alpha}}_j^\top \boldsymbol{\varphi}_j(\mathbf{x}).$$

All tuning parameters such as the regularization parameter γ_j and parameters included in the basis function $\varphi_{k,j}(\mathbf{y})$ can be systematically chosen based on cross-validation w.r.t. Eq.(15).

3.3. Theoretical Analysis

Here, we investigate the convergence rate of WF-LSNGCA in a parametric setting.

Let $g^*(\mathbf{x})$ be the optimal estimate to $\nabla_{\mathbf{x}} \ln p(\mathbf{x})$ given by LSLDG based on the linear-in-parameter model $g(\mathbf{x})$, and let

$$\begin{aligned} \mathbf{S}_j^* &= \mathbb{E} \left[\boldsymbol{\varphi}_j(\mathbf{x}) \boldsymbol{\varphi}_j(\mathbf{x})^\top \right], \quad \mathbf{t}_j^* = \mathbb{E} \left[\partial_j \boldsymbol{\varphi}_j(\mathbf{x}) + \boldsymbol{\varphi}_j(\mathbf{x}) \left(\nabla_{\mathbf{x}} g^{*(j)}(\mathbf{x}) \right)^\top \mathbf{x} \right], \\ \boldsymbol{\alpha}_j^* &= \operatorname{argmin}_{\boldsymbol{\alpha}} \left\{ \boldsymbol{\alpha}^\top \mathbf{S}_j^* \boldsymbol{\alpha} + 2 \boldsymbol{\alpha}^\top \mathbf{t}_j^* + \gamma_j^* \boldsymbol{\alpha}^\top \boldsymbol{\alpha} \right\}, \quad w^{*(j)}(\mathbf{x}) = \boldsymbol{\alpha}_j^{*\top} \boldsymbol{\varphi}_j(\mathbf{x}), \end{aligned}$$

where $(\mathbf{S}_j^* + \gamma_j^* \mathbf{I}_{b'})$ must be strictly positive definite. In fact, \mathbf{S}_j^* should already be strictly positive definite, and thus $\gamma_j^* = 0$ is also allowed in our theoretical analysis.

We have the following theorem (its proof is given in Section 3.4):

Theorem 1 *As $n \rightarrow \infty$, for any \mathbf{x} ,*

$$\|\hat{\mathbf{v}}(\mathbf{x}) - \mathbf{w}^*(\mathbf{x})\|_2 = \mathcal{O}_p \left(n^{-1/2} \right),$$

provided that γ_j for all j converge in $\mathcal{O}(n^{-1/2})$ to γ_j^ , i.e., $\lim_{n \rightarrow \infty} n^{1/2} |\gamma_j - \gamma_j^*| < \infty$.*

Theorem 1 is based on the theory of perturbed optimizations (Bonnans and Cominetti, 1996; Bonnans and Shapiro, 1998) as well as the convergence of LSLDG shown in Sasaki et al. (2016). It guarantees that for any \mathbf{x} , the estimate $\hat{\mathbf{v}}(\mathbf{x})$ in WF-LSNGCA converges to the optimal estimate $\mathbf{w}^*(\mathbf{x})$ based on the linear-in-parameter model $w(\mathbf{x})$, and it achieves the optimal parametric convergence rate $\mathcal{O}_p(n^{-1/2})$. Note that Theorem 1 deals only with the estimation error, and the approximation error is not taken into account. Indeed, approximation errors exist in two places, from $\mathbf{w}^*(\mathbf{x})$ to $\mathbf{v}(\mathbf{x})$ in WF-LSNGCA itself and from $g^*(\mathbf{x})$ to $\nabla_{\mathbf{x}} \ln p(\mathbf{x})$ in the plug-in LSLDG estimator. Since the original LSNGCA also relies on LSLDG, it cannot avoid the approximation error introduced by LSLDG. For this reason, the convergence of WF-LSNGCA is expected to be as good as LSNGCA.

Theorem 1 is basically a theoretical guarantee that is similar to *Part One in the proof of Theorem 1* in Sasaki et al. (2016). Hence, based on Theorem 1, we can go along the line of *Part Two in the proof of Theorem 1* in Sasaki et al. (2016) and obtain the following corollary.

Corollary 2 For eigendecomposition, define matrices $\widehat{\mathbf{\Gamma}} = \frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{v}}(\mathbf{x}_i) \widehat{\mathbf{v}}(\mathbf{x}_i)^\top$ and $\mathbf{\Gamma}^* = \mathbb{E}[\mathbf{w}^*(\mathbf{x}) \mathbf{w}^*(\mathbf{x})^\top]$. Given the estimated subspace \widehat{E} based on n samples and the optimal estimated subspace E^* based on infinite data, denote by $\widehat{\mathbf{E}} \in \mathbb{R}^{d \times m}$ the matrix form of an arbitrary orthonormal basis of \widehat{E} and by $\mathbf{E}^* \in \mathbb{R}^{d \times m}$ that of E^* . Define the distance between subspaces as

$$\mathcal{D}(\widehat{E}, E^*) = \inf_{\widehat{\mathbf{E}}, \mathbf{E}^*} \|\widehat{\mathbf{E}} - \mathbf{E}^*\|_{\text{Fro}},$$

where $\|\cdot\|_{\text{Fro}}$ stands for the Frobenius norm. Then, as $n \rightarrow \infty$,

$$\mathcal{D}(\widehat{E}, E^*) = \mathcal{O}_p\left(n^{-1/2}\right),$$

provided that γ_j for all j converge in $\mathcal{O}(n^{-1/2})$ to γ_j^* and $\boldsymbol{\varphi}_j(\mathbf{x})$ are well-chosen basis functions such that the first m eigenvalues of $\mathbf{\Gamma}^*$ are neither 0 nor $+\infty$.

3.4. Proof of Theorem 1

Step 1. First of all, we establish the growth condition (see *Definition 6.1* in [Bonnans and Shapiro \(1998\)](#)). Denote the expected and empirical objective functions by

$$\begin{aligned} R_j^*(\boldsymbol{\alpha}) &= \boldsymbol{\alpha}^\top \mathbf{S}_j^* \boldsymbol{\alpha} + 2\boldsymbol{\alpha}^\top \mathbf{t}_j^* + \gamma_j^* \boldsymbol{\alpha}^\top \boldsymbol{\alpha}, \\ \widehat{R}_j(\boldsymbol{\alpha}) &= \boldsymbol{\alpha}^\top \widehat{\mathbf{S}}_j \boldsymbol{\alpha} + 2\boldsymbol{\alpha}^\top \widehat{\mathbf{t}}_j + \gamma_j \boldsymbol{\alpha}^\top \boldsymbol{\alpha}. \end{aligned}$$

Then $\boldsymbol{\alpha}_j^* = \operatorname{argmin}_{\boldsymbol{\alpha}} R_j^*(\boldsymbol{\alpha})$, $\widehat{\boldsymbol{\alpha}}_j = \operatorname{argmin}_{\boldsymbol{\alpha}} \widehat{R}_j(\boldsymbol{\alpha})$, and we have

Lemma 3 Let ϵ_j be the smallest eigenvalue of $(\mathbf{S}_j^* + \gamma_j^* \mathbf{I}_{b'})$, then the following second-order growth condition holds

$$R_j^*(\boldsymbol{\alpha}) \geq R_j^*(\boldsymbol{\alpha}_j^*) + \epsilon_j \|\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*\|_2^2.$$

Proof $R_j^*(\boldsymbol{\alpha})$ must be strongly convex with parameter at least $2\epsilon_j$. Hence,

$$\begin{aligned} R_j^*(\boldsymbol{\alpha}) &\geq R_j^*(\boldsymbol{\alpha}_j^*) + (\nabla R_j^*(\boldsymbol{\alpha}_j^*))^\top (\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*) + (\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*)^\top (\mathbf{S}_j^* + \gamma_j^* \mathbf{I}_{b'}) (\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*) \\ &\geq R_j^*(\boldsymbol{\alpha}_j^*) + \epsilon_j \|\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*\|_2^2, \end{aligned}$$

where we used the optimality condition $\nabla R_j^*(\boldsymbol{\alpha}_j^*) = \mathbf{0}$. ■

Step 2. Second, we study the stability (with respect to perturbation) of $R_j^*(\boldsymbol{\alpha})$ at $\boldsymbol{\alpha}_j^*$. Let

$$\mathbf{u} = \{\mathbf{u}_S \in \mathcal{S}_+^{b'}, \mathbf{u}_t \in \mathbb{R}^{b'}, u_\gamma \in \mathbb{R}\}$$

be a set of perturbation parameters, where $\mathcal{S}_+^{b'} \subset \mathbb{R}^{b' \times b'}$ is the cone of b' -by- b' symmetric positive semi-definite matrices. Define our perturbed objective function by

$$R_j(\boldsymbol{\alpha}, \mathbf{u}) = \boldsymbol{\alpha}^\top (\mathbf{S}_j^* + \mathbf{u}_S) \boldsymbol{\alpha} + 2\boldsymbol{\alpha}^\top (\mathbf{t}_j^* + \mathbf{u}_t) + (\gamma_j^* + u_\gamma) \boldsymbol{\alpha}^\top \boldsymbol{\alpha}.$$

It is clear that $R_j^*(\boldsymbol{\alpha}) = R_j(\boldsymbol{\alpha}, \mathbf{0})$, and then the stability of $R_j^*(\boldsymbol{\alpha})$ at $\boldsymbol{\alpha}_j^*$ can be characterized as follows.

Lemma 4 *The difference function $R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha})$ is Lipschitz continuous in $\boldsymbol{\alpha}$ modulus*

$$\omega(\mathbf{u}) = \mathcal{O}(\|\mathbf{u}_S\|_{\text{Fro}} + \|\mathbf{u}_t\|_2 + |u_\gamma|)$$

on a sufficiently small neighborhood of $\boldsymbol{\alpha}_j^$.*

Proof The difference function is

$$R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha}) = \boldsymbol{\alpha}^\top \mathbf{u}_S \boldsymbol{\alpha} + 2\boldsymbol{\alpha}^\top \mathbf{u}_t + u_\gamma \boldsymbol{\alpha}^\top \boldsymbol{\alpha},$$

with a partial gradient

$$\frac{\partial}{\partial \boldsymbol{\alpha}} (R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha})) = 2\mathbf{u}_S \boldsymbol{\alpha} + 2\mathbf{u}_t + 2u_\gamma \boldsymbol{\alpha}.$$

Notice that due to the ℓ_2 -regularization in $R_j^*(\boldsymbol{\alpha})$, $\exists M > 0$ such that $\|\boldsymbol{\alpha}_j^*\|_2 \leq M$. Now given a δ -ball of $\boldsymbol{\alpha}_j^*$, i.e., $B_\delta(\boldsymbol{\alpha}_j^*) = \{\boldsymbol{\alpha} \mid \|\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*\|_2 \leq \delta\}$, it is easy to see that $\forall \boldsymbol{\alpha} \in B_\delta(\boldsymbol{\alpha}_j^*)$,

$$\|\boldsymbol{\alpha}\|_2 \leq \|\boldsymbol{\alpha} - \boldsymbol{\alpha}_j^*\|_2 + \|\boldsymbol{\alpha}_j^*\|_2 \leq \delta + M,$$

and consequently

$$\left\| \frac{\partial}{\partial \boldsymbol{\alpha}} (R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha})) \right\|_2 \leq 2(\delta + M)(\|\mathbf{u}_S\|_{\text{Fro}} + |u_\gamma|) + 2\|\mathbf{u}_t\|_2.$$

This says that the gradient $\frac{\partial}{\partial \boldsymbol{\alpha}} (R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha}))$ has a bounded norm of order $\mathcal{O}(\|\mathbf{u}_S\|_{\text{Fro}} + \|\mathbf{u}_t\|_2 + |u_\gamma|)$, and proves that the difference function $R_j(\boldsymbol{\alpha}, \mathbf{u}) - R_j^*(\boldsymbol{\alpha})$ is Lipschitz continuous on the ball $B_\delta(\boldsymbol{\alpha}_j^*)$, with a Lipschitz constant of the same order. \blacksquare

Step 3. Lemma 3 ensures the unperturbed objective $R_j^*(\boldsymbol{\alpha})$ grows quickly when $\boldsymbol{\alpha}$ leaves $\boldsymbol{\alpha}_j^*$; Lemma 4 ensures the perturbed objective $R_j(\boldsymbol{\alpha}, \mathbf{u})$ changes slowly for $\boldsymbol{\alpha}$ around $\boldsymbol{\alpha}_j^*$, where the slowness is compared with the perturbation \mathbf{u} it suffers. Based on Lemma 3, Lemma 4, and *Proposition 6.1* in [Bonnans and Shapiro \(1998\)](#),

$$\|\widehat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_j^*\|_2 \leq \frac{\omega(\mathbf{u})}{\epsilon_j} = \mathcal{O}(\|\mathbf{u}_S\|_{\text{Fro}} + \|\mathbf{u}_t\|_2 + |u_\gamma|),$$

since $\widehat{\boldsymbol{\alpha}}_j$ is the exact solution to $\widehat{R}_j(\boldsymbol{\alpha}) = R_j(\boldsymbol{\alpha}, \mathbf{u})$ given $\mathbf{u}_S = \widehat{\mathbf{S}}_j - \mathbf{S}_j^*$, $\mathbf{u}_t = \widehat{\mathbf{t}}_j - \mathbf{t}_j^*$, and $u_\gamma = \gamma_j - \gamma_j^*$.

According to the *central limit theorem* (CLT), $\|\mathbf{u}_S\|_{\text{Fro}} = \mathcal{O}_p(n^{-1/2})$. Consider $\widehat{\mathbf{t}}_j - \mathbf{t}_j^*$:

$$\begin{aligned} \widehat{\mathbf{t}}_j - \mathbf{t}_j^* &= \frac{1}{n} \sum_{i=1}^n \partial_j \boldsymbol{\varphi}_j(\mathbf{x}_i) - \mathbb{E} [\partial_j \boldsymbol{\varphi}_j(\mathbf{x})] + \frac{1}{n} \sum_{i=1}^n \boldsymbol{\varphi}_j(\mathbf{x}_i) \left(\nabla_{\mathbf{x}} \widehat{g}^{(j)}(\mathbf{x}_i) \right)^\top \mathbf{x}_i \\ &\quad - \mathbb{E} \left[\boldsymbol{\varphi}_j(\mathbf{x}) \left(\nabla_{\mathbf{x}} g^{*(j)}(\mathbf{x}) \right)^\top \mathbf{x} \right]. \end{aligned}$$

The first half is clearly $\mathcal{O}_p(n^{-1/2})$ due to CLT. For the second half, the estimate $\widehat{g}^{(j)}(\mathbf{x})$ given by LSLDG converges to $g^{*(j)}(\mathbf{x})$ for any \mathbf{x} in $\mathcal{O}_p(n^{-1/2})$ according to *Part One in the*

proof of Theorem 1 in Sasaki et al. (2016), and $\nabla_{\mathbf{x}}\widehat{g}^{(j)}(\mathbf{x})$ converges to $\nabla_{\mathbf{x}}g^{*(j)}(\mathbf{x})$ in the same order because the basis functions in $\psi_j(\mathbf{x})$ are all derivatives of Gaussian functions. Consequently,

$$\frac{1}{n} \sum_{i=1}^n \varphi_j(\mathbf{x}_i) \left(\nabla_{\mathbf{x}}\widehat{g}^{(j)}(\mathbf{x}_i) \right)^\top \mathbf{x}_i - \frac{1}{n} \sum_{i=1}^n \varphi_j(\mathbf{x}_i) \left(\nabla_{\mathbf{x}}g^{*(j)}(\mathbf{x}_i) \right)^\top \mathbf{x}_i = \mathcal{O}_p(n^{-1/2}),$$

since $\nabla_{\mathbf{x}}\widehat{g}^{(j)}(\mathbf{x})$ converges to $\nabla_{\mathbf{x}}g^{*(j)}(\mathbf{x})$ for any \mathbf{x} in $\mathcal{O}_p(n^{-1/2})$, and

$$\frac{1}{n} \sum_{i=1}^n \varphi_j(\mathbf{x}_i) \left(\nabla_{\mathbf{x}}g^{*(j)}(\mathbf{x}_i) \right)^\top \mathbf{x}_i - \mathbb{E} \left[\varphi_j(\mathbf{x}) \left(\nabla_{\mathbf{x}}g^{*(j)}(\mathbf{x}) \right)^\top \mathbf{x} \right] = \mathcal{O}_p(n^{-1/2})$$

due to CLT, which proves $\|\mathbf{u}_t\|_2 = \mathcal{O}_p(n^{-1/2})$. Furthermore, we have already assumed that $|u_\gamma| = \mathcal{O}(n^{-1/2})$. Hence, as $n \rightarrow \infty$,

$$\|\widehat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_j^*\|_2 = \mathcal{O}_p\left(n^{-1/2}\right).$$

Step 4. Finally, for any \mathbf{x} , the gap of $\widehat{v}^{(j)}(\mathbf{x})$ and $w^{*(j)}(\mathbf{x})$ is bounded by

$$|\widehat{v}^{(j)}(\mathbf{x}) - w^{*(j)}(\mathbf{x})| \leq \|\widehat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_j^*\|_2 \cdot \|\varphi_j(\mathbf{x})\|_2,$$

where the *Cauchy-Schwarz inequality* is used. Since the basis functions in $\varphi_j(\mathbf{x})$ are again all derivatives of Gaussian functions, $\|\varphi_j(\mathbf{x})\|_2$ must be bounded uniformly, and then

$$|\widehat{v}^{(j)}(\mathbf{x}) - w^{*(j)}(\mathbf{x})| \leq \mathcal{O}(\|\widehat{\boldsymbol{\alpha}}_j - \boldsymbol{\alpha}_j^*\|_2) = \mathcal{O}_p\left(n^{-1/2}\right).$$

Applying the same argument for all $j = 1, \dots, d$ completes the proof. ■

4. Experiments

In this section, we experimentally investigate the performance of MIPP, LSNGCA, and WF-LSNGCA.

4.1. Configurations of NGCA Algorithms

4.1.1. MIPP

We use the MATLAB script which was used in the original MIPP paper (Blanchard et al., 2006)¹. In this script, NGIFs of the form $s_m^k(z)$ ($m = 1, \dots, 1000, k = 1, \dots, 4$) are used:

$$\begin{aligned} s_m^1(z) &= z^3 \exp\left(-\frac{z^2}{2\sigma_m^2}\right), & s_m^2(z) &= \tanh(a_m z), \\ s_m^3(z) &= \sin(b_m z), & s_m^4(z) &= \cos(b_m z), \end{aligned}$$

where σ_m , a_m , and b_m are scalars chosen at the regular intervals from $\sigma_m \in [0.5, 5]$, $a_m \in [0.05, 5]$, and $b_m \in [0.05, 4]$. The cut-off threshold τ is set at 1.6 and the number of FastICA iterations is set at 10 (see Section 2.2).

1. <http://www.ms.k.u-tokyo.ac.jp/software.html>


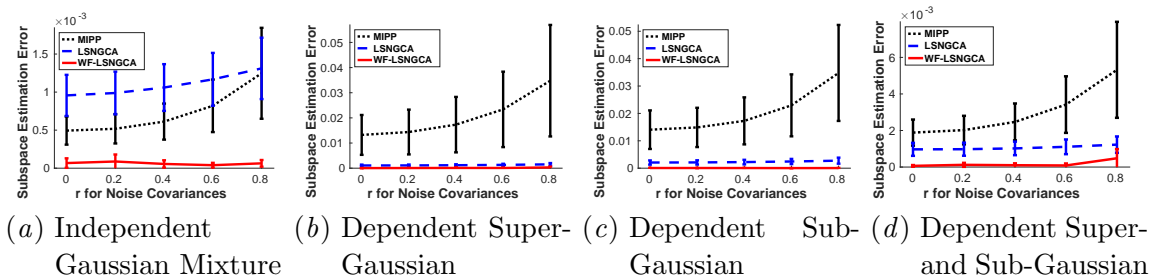


Figure 1: Distributions of non-Gaussian components.


 Figure 2: Averages and standard deviations of the subspace estimation error as the function of the condition-number controller r over 50 simulations on artificial datasets.

4.1.2. LSNGCA

Following [Sasaki et al. \(2014\)](#), the derivative of the Gaussian kernel is used as the basis function $\psi_{k,j}(\mathbf{y})$ in the linear-in-parameter model (10):

$$\psi_{k,j}(\mathbf{y}) = \partial_j \exp\left(-\frac{\|\mathbf{y} - \mathbf{c}_k\|^2}{2\sigma_j^2}\right),$$

where $\sigma_j > 0$ is the Gaussian bandwidth and \mathbf{c}_k is the Gaussian center randomly selected from the whitened data samples $\{\mathbf{y}_i\}_{i=1}^n$. The number of basis functions is set at $b = 100$. For model selection, 5-fold cross-validation is performed with respect to the hold-out error of Eq.(9) using 10 candidate values at the regular intervals in logarithmic scale for Gaussian bandwidth $\sigma_j \in [10^{-1}, 10^1]$ and regularization parameter $\lambda_j \in [10^{-5}, 10^1]$.

4.1.3. WF-LSNGCA

Similarly to LSNGCA, the derivative of the Gaussian kernel is used as the basis function $\varphi_{k,j}(\mathbf{x})$ in the linear-in-parameter model (16) and the number of basis functions is set as $b' = 100$. For model selection, 5-fold cross-validation is performed with respect to the hold-out error of Eq.(15) in the same way as LSNGCA.

4.2. Artificial Datasets

Let $\mathbf{x} = (s_1, s_2, n_3, \dots, n_{10})^\top$, where $\mathbf{s} := (s_1, s_2)^\top$ are the 2-dimensional non-Gaussian signal components and $\mathbf{n} := (n_3, \dots, n_{10})^\top$ are the 8-dimensional Gaussian noise compo-

nents. For the non-Gaussian signal components, we consider the following four distributions plotted in Figure 1:

(a) **Independent Gaussian Mixture:**

$$p(s_1, s_2) \propto \prod_{i=1}^2 \left(\exp\left(-\frac{(s_i - 3)^2}{2}\right) + \exp\left(-\frac{(s_i + 3)^2}{2}\right) \right).$$

(b) **Dependent super-Gaussian:**

$$p(\mathbf{s}) \propto \exp(-\|\mathbf{s}\|).$$

(c) **Dependent sub-Gaussian:**

$$p(\mathbf{s}) \text{ is the uniform distribution on } \{\mathbf{s} \in \mathbb{R}^2 \mid \|\mathbf{s}\| \leq 1\}.$$

(d) **Dependent super- and sub-Gaussian:**

$$p(s_1) \propto \exp(-|s_1|) \text{ and } p(s_2) \text{ is the uniform distribution on } [c, c + 1], \text{ where } c = 0 \text{ if } |s_1| \leq \log 2 \text{ and } c = -1 \text{ otherwise.}$$

For the Gaussian noise components, we include a certain parameter $r \geq 0$, which controls the condition number; the larger r is, the more ill-posed the data covariance matrix is. The detail is described in the supplementary material.

We generate $n = 2000$ samples for each case, and standardize each element of the data before applying NGCA algorithms. The performance of NGCA algorithms is measured by the following *subspace estimation error*:

$$\varepsilon(E, \hat{E}) := \frac{1}{2} \sum_{i=1}^2 \|\hat{\mathbf{e}}_i - \Pi_E \hat{\mathbf{e}}_i\|^2, \quad (17)$$

where E is the true non-Gaussian index space, \hat{E} is its estimate, Π_E is the orthogonal projection on E , and $\{\hat{\mathbf{e}}_i\}_{i=1}^2$ is an orthonormal basis in \hat{E} .

The averages and the standard derivations of the subspace estimation error over 50 runs for MIPP, LSNGCA, and WF-LSNGCA are depicted in Figure 2. The figure shows that, for all 4 cases, the error of MIPP and LSNGCA grow as r increases, while WF-LSNGCA performs much stably against the change in r . This shows that WF-LSNGCA is much more robust against ill-conditioning than MIPP and LSNGCA. Additionally, the figure shows that WF-LSNGCA performs well even when the covariance matrix is well-conditioned. This implies that the finite sample effect for pre-whitening is not ignorable even in well-conditioned cases.

In terms of the computation time, WF-LSNGCA is less efficient than LSNGCA and MIPP, but its computation time is still just a few times slower than LSNGCA, as seen in Figure 3. For this reason, the computational efficiency of WF-LSNGCA would still be acceptable in practice.

4.3. Benchmark Datasets

Finally, we evaluate the performance of NGCA methods using the *LIBSVM binary classification benchmark datasets*² (Chang and Lin, 2011). From each dataset, n points are

2. We preprocessed the LIBSVM binary classification benchmark datasets as follows:

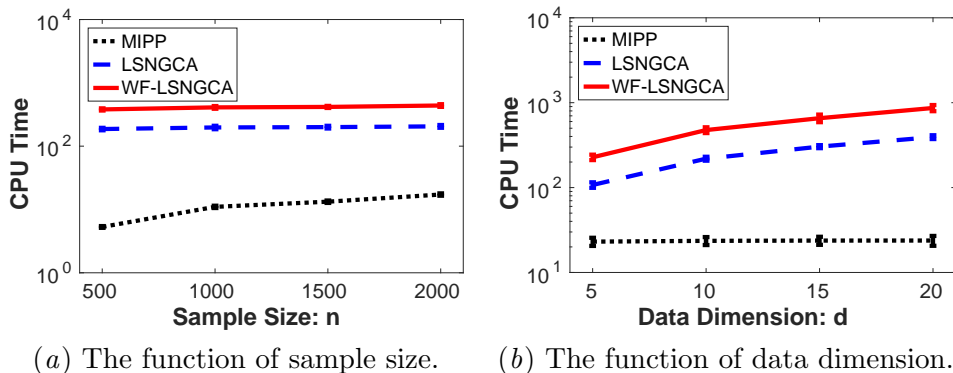


Figure 3: The average CPU time over 50 runs when the Gaussian mixture is used as non-Gaussian components and the condition-number controller $r = 0$. The vertical axis is in logarithmic scale.

selected as training (test) samples so that the number of positive and negative samples are equal, and datasets are standardized in each dimension. For an m -dimensional dataset, we append $(d - m)$ -dimensional noise dimensions following the standard Gaussian distribution so that all datasets have d dimensions. Then we use PCA, MIPP, LSNCGA, and WF-LSNGCA to obtain m -dimensional expressions, and apply the support vector machine (SVM) ³ to evaluate the test misclassification rate. As a baseline, we also evaluate the misclassification rate by the raw SVM without dimension reduction.

The averages and standard deviations of the misclassification rate over 50 runs for $d = 50, 100$ are summarized in Table 2. As can be seen in the table, the appended Gaussian noise dimensions have negative effects on each classification accuracy, and thus the baseline has relatively high misclassification rates. PCA has overall higher misclassification rates than the baseline since a lot of valuable information for each classification problem is lost. Among the NGCA algorithms, WF-LSNGCA overall compares favorably with the other methods. This means that it can find valuable low-dimensional expressions for each classification problem without harmful effects of a pre-whitening procedure.

5. Conclusions

In this paper, we proposed a novel NGCA algorithm which is computationally efficient, no manual design of non-Gaussian index functions is required, and pre-whitening is not involved. Through experiments, we demonstrated that the effectiveness of the proposed method.

- *vehicle*: We convert original labels ‘1’ and ‘2’ to the positive label and original labels ‘3’ and ‘4’ to the negative label.
- *SUSY*: We convert original label ‘0’ to the negative label.
- *shuttle*: We use only the data labeled as ‘1’ and ‘4’ and regard them as positive and negative labels.
- *svmguide1*: We mix the original training and test datasets.

3. We used LIBSVM with MATLAB (Chang and Lin, 2011).

Table 2: Averages (and standard deviations in the parentheses) of the misclassification rates for the LIBSVM datasets over 50 runs. The best and comparable algorithms judged by the two-sample t-test at the significance level 5% are expressed as bold-face.

	$d = 50$				
Dataset [m, n]	No Dim. Red.	PCA	MIPP	LS NGCA	WF-LS NGCA
vehicle [18, 200]	0.340 (0.038)	0.404 (0.034)	0.328 (0.044)	0.324 (0.044)	0.286 (0.038)
svmguid3 [21, 200]	0.342 (0.035)	0.348 (0.037)	0.341 (0.041)	0.326 (0.039)	0.308 (0.036)
svmguid1 [3, 2000]	0.088 (0.007)	0.159 (0.012)	0.060 (0.005)	0.058 (0.006)	0.053 (0.008)
shuttle [9, 2000]	0.031 (0.004)	0.024 (0.007)	0.021 (0.004)	0.041 (0.015)	0.007 (0.002)
SUSY [18, 2000]	0.238 (0.010)	0.271 (0.012)	0.229 (0.010)	0.223 (0.012)	0.228 (0.010)
ijcnn1 [22, 2000]	0.102 (0.007)	0.273 (0.012)	0.084 (0.028)	0.061 (0.006)	0.057 (0.007)

	$d = 100$				
Dataset [m, n]	No Dim. Red.	PCA	MIPP	LS NGCA	WF-LS NGCA
vehicle [18, 200]	0.380 (0.033)	0.432 (0.033)	0.445 (0.036)	0.439 (0.045)	0.360 (0.051)
svmguid3 [21, 200]	0.363 (0.034)	0.367 (0.032)	0.443 (0.042)	0.427 (0.035)	0.343 (0.044)
svmguid1 [3, 2000]	0.102 (0.008)	0.175 (0.013)	0.087 (0.013)	0.088 (0.059)	0.067 (0.020)
shuttle [9, 2000]	0.038 (0.005)	0.069 (0.013)	0.065 (0.019)	0.209 (0.080)	0.017 (0.005)
SUSY [18, 2000]	0.250 (0.010)	0.280 (0.011)	0.234 (0.009)	0.228 (0.011)	0.226 (0.010)
ijcnn1 [22, 2000]	0.145 (0.008)	0.318 (0.013)	0.107 (0.021)	0.101 (0.010)	0.091 (0.020)

Acknowledgments

GN acknowledges support from JST CREST JPMJCR1403.

References

G. Blanchard, M. Kawanabe, M. Sugiyama, V. Spokoiny, and K. Müller. In search of non-Gaussian components of a high-dimensional distribution. *Journal of Machine Learning Research*, 7:247–282, 2006.

- F. Bonnans and R. Cominetti. Perturbed optimization in Banach spaces I: A general theory based on a weak directional constraint qualification; II: A theory based on a strong directional qualification condition; III: Semiinfinite optimization. *SIAM Journal on Control and Optimization*, 34:1151–1171, 1172–1189, and 1555–1567, 1996.
- F. Bonnans and A. Shapiro. Optimization problems with perturbations, a guided tour. *SIAM Review*, 40(2):228–264, 1998.
- C. Chang and C. Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2(3):27, 2011.
- D. D Cox. A penalty method for nonparametric estimation of the logarithmic derivative of a density function. *Annals of the Institute of Statistical Mathematics*, 37(1):271–288, 1985.
- E. Diederichs, A. Juditsky, V. Spokoiny, and C. Schütte. Sparse non-Gaussian component analysis. *IEEE Transactions on Information Theory*, 56(6):3033–3047, 2010.
- E. Diederichs, A. Juditsky, A. Nemirovski, and V. Spokoiny. Sparse non Gaussian component analysis by semidefinite programming. *Machine learning*, 91(2):211–238, 2013.
- J. H. Friedman and J. W. Tukey. A projection pursuit algorithm for exploratory data analysis. *IEEE Transactions on Computers*, C-23(9):881–890, 1974.
- A. Hyvärinen. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, 10(3):626–634, 1999.
- A. Hyvärinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13(4):411–430, 2000.
- M. Kawanabe, M. Sugiyama, G. Blanchard, and K. Müller. A new algorithm of non-Gaussian component analysis with radial kernel functions. *Annals of the Institute of Statistical Mathematics*, 59(1):57–75, 2007.
- H. Sasaki, A. Hyvärinen, and M. Sugiyama. Clustering via mode seeking by direct estimation of the gradient of a log-density. In *Machine Learning and Knowledge Discovery in Databases*, pages 19–34. Springer, 2014.
- H. Sasaki, G. Niu, and M. Sugiyama. Non-Gaussian component analysis with log-density gradient estimation. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 38, pages 1177–1185, 2016.
- F. J Theis and M. Kawanabe. Uniqueness of non-Gaussian subspace analysis. In *Independent Component Analysis and Blind Signal Separation*, pages 917–925. Springer, 2006.
- V. N. Vapnik. *Statistical Learning Theory*, volume 1. Wiley New York, 1998.

Supplementary Material to Whitening-Free Least-Squares Non-Gaussian Component Analysis

Detail of Artificial Datasets

Here, we describe the detail of the artificial datasets used in Section 4.2. The noise components \mathbf{n} are generated as follows:

1. The \mathbf{n} is sampled from the centered Gaussian distribution with covariance matrix $\text{diag}(10^{-2r}, 10^{-2r+4r/7}, 10^{-2r+8r/7}, \dots, 10^{2r})$, where $\text{diag}(\cdot)$ denotes the diagonal matrix.
2. The sampled \mathbf{n} is rotated as $\mathbf{n}'' \in \mathbb{R}^8$ by applying the following rotation matrix $\mathbf{R}^{(i,j)}$ for all $i, j = 3, \dots, 10$ such that $i < j$:

$$\begin{aligned} R_{i,i}^{(i,j)} &= \cos(\pi/4), & R_{i,j}^{(i,j)} &= -\sin(\pi/4), \\ R_{j,i}^{(i,j)} &= \sin(\pi/4), & R_{j,j}^{(i,j)} &= \cos(\pi/4), \\ R_{k,k}^{(i,j)} &= 1 \quad (k \neq i, k \neq j), & R_{k,l}^{(i,j)} &= 0 \quad (\text{otherwise}). \end{aligned}$$

3. The rotated \mathbf{n} is normalized for each dimension.

By this construction, increasing r corresponds to increasing the *condition number* of the data covariance matrix (see Figure 4). Thus, the larger r is, the more ill-posed the data covariance matrix is.

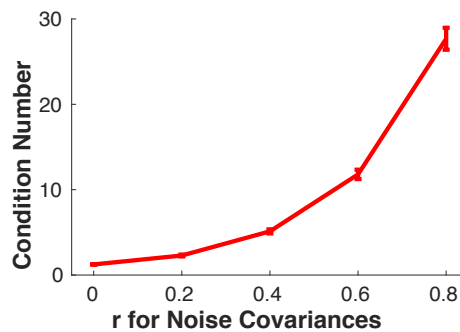


Figure 4: Condition number of the data covariance matrix as a function of experiment parameter r (with non-Gaussian components generated from the Gaussian mixture).