

# Representation Learning for Maximization of MI, Nonlinear ICA and Nonlinear Subspaces with Robust Density Ratio Estimation

**Hiroaki Sasaki**

HSASAKI@FUN.AC.JP

*Department of Complexity and Intelligent Systems*

*Future University Hakodate*

*Hokkaido, Japan*

**Takashi Takenouchi**

T-TAKENOUCI@GRIPS.AC.JP

*National Graduate Institute for Policy Studies*

*RIKEN AIP*

*Tokyo, Japan*

**Editor:** Stefan Harmeling

## Abstract

Unsupervised representation learning is one of the most important problems in machine learning. A recent promising approach is contrastive learning: A feature representation of data is learned by solving a *pseudo* classification problem where class labels are automatically generated from unlabelled data. However, it is not straightforward to understand what representation contrastive learning yields through the classification problem. In addition, most of practical methods for contrastive learning are based on the maximum likelihood estimation, which is often vulnerable to the contamination by outliers. In order to promote the understanding to contrastive learning, this paper first theoretically shows a connection to maximization of mutual information (MI). Our result indicates that density ratio estimation is necessary and sufficient for maximization of MI under some conditions. Since popular objective functions for classification can be regarded as estimating density ratios, contrastive learning related to density ratio estimation can be interpreted as maximizing MI. Next, in terms of density ratio estimation, we establish new recovery conditions for the latent source components in nonlinear independent component analysis (ICA). In contrast with existing work, the established conditions include a novel insight for the dimensionality of data, which is clearly supported by numerical experiments. Furthermore, inspired by nonlinear ICA, we propose a novel framework to estimate a nonlinear subspace for lower-dimensional latent source components, and some theoretical conditions for the subspace estimation are established with density ratio estimation. Motivated by the theoretical results, we propose a practical method through outlier-robust density ratio estimation, which can be seen as performing maximization of MI, nonlinear ICA or nonlinear subspace estimation. Moreover, a sample-efficient nonlinear ICA method is also proposed based on a variational lower-bound of MI. Then, we theoretically investigate outlier-robustness of the proposed methods. Finally, we numerically demonstrate usefulness of the proposed methods in nonlinear ICA and through application to a downstream task for linear classification.

**Keywords:** representation learning, contrastive learning, density ratio estimation, maximization of MI, nonlinear ICA, nonlinear subspace estimation, outlier-robustness

## 1. Introduction

Unsupervised representation learning has been a long-term issue in machine learning (Raina et al., 2007; Bengio et al., 2013). In contrast with supervised learning, a feature representation of data is learned only from unlabelled data. Since the acquisition cost of unlabelled data is often low, the learned representation would be useful particularly when the number of data is very limited in a downstream task (e.g., supervised tasks). Successful examples include video classification (Wang and Gupta, 2015; Arandjelovic and Zisserman, 2017; Sun et al., 2019), motion capture (Tung et al., 2017), and natural language processing (Peters et al., 2018; Devlin et al., 2019). More examples can be found in recent review papers for representation learning (Jing and Tian, 2020; Liu et al., 2020).

Recently, a number of methods for unsupervised representation learning have been proposed, and most of them can be roughly divided into three categories based on variational autoencoder (VAE) (Kingma and Welling, 2014; Higgins et al., 2017; Khemakhem et al., 2020), generative adversarial network (GAN) (Goodfellow et al., 2014; Chen et al., 2016) and contrastive learning (Dosovitskiy et al., 2014; Doersch et al., 2015; Noroozi and Favaro, 2016). VAE learns a useful representation of data by maximizing a tractable lower-bound of the likelihood function. However, VAE makes a restrictive prior assumption on densities such as the Gaussian assumption. The GAN approach learns both a representation and generator of data, but it requires to solve a max-min optimization problem, which could be numerically unstable. On the other hand, contrastive learning does not make strong assumptions on densities, and simply solves a straightforward optimization problem.

In order to learn a feature representation of data, contrastive learning solves a *pseudo* classification problem where class labels are automatically generated from unlabelled data. For example, Arandjelovic and Zisserman (2017) make a dataset for binary classification from unlabelled video data: The positive labels are assigned to pairs of video frames and audio clips taken from the same video, while the negative labels are given to pairs from different videos. Then, representations of both video frames and audio clips are learned by solving the binary classification problem. Misra et al. (2016) also solve a binary classification problem for unsupervised representation learning where the positive data is temporally consecutive video frames in a video, but the negative one is temporally shuffled frames in the same video. This learning scheme based on classification problems is also called *self-supervised learning* (SSL). See a recent review article for SSL (Liu et al., 2020) and references therein. However, it is not straightforward to understand what representation contrastive learning yields by solving classification problems.

Contrastive learning has been justified as being closely related to a classical framework so called *maximization of mutual information* (MI) (Linsker, 1989; Bell and Sejnowski, 1995). This comes from an intuitive idea that MI must be high if a classifier can accurately distinguish positive samples drawn from the joint density (e.g., for image frames and audio clips) and negative ones drawn from the product of the marginal densities (Belghazi et al., 2018; Hjelm et al., 2019; Tschannen et al., 2019). However, this intuitive idea is rather superficial because it is not necessarily obvious why solving the classification problem leads to maximization of MI. A recent promising approach is based on maximization of tractable variational lower-bounds of MI and also related to contrastive learning (Nguyen et al., 2008; Belghazi et al., 2018; van den Oord et al., 2018), but the original purpose of maximizing the variational lower-bounds is *estimation* of MI, not necessarily maximization of MI. More rigorous justification would be preferable to understand how and when contrastive learning can be regarded as performing maximization of MI.

Contrastive learning has been employed in recent algorithms of *nonlinear independent component analysis* (ICA) (Hyvärinen and Morioka, 2016, 2017; Hyvärinen et al., 2019), which is a solid framework for unsupervised representation learning. Nonlinear ICA rigorously defines a generative model where the input data is assumed to be observed as a general nonlinear mixing of the latent source components. Then, the problem is to recover the source components from the observations of input data. Regarding the linear ICA where the mixing function is linear, the recovery conditions are well-established and the important condition is mutual independence of the source components (Comon, 1994). On the other hand, the problem of nonlinear ICA has been proved to be fundamentally ill-posed under the same independence condition because there exist an infinite number of decompositions of a random vector into mutually independent variables (Hyvärinen and Pajunen, 1999; Locatello et al., 2019). Very recently, novel recovery conditions for nonlinear ICA have been established (Sprekeler et al., 2014; Hyvärinen and Morioka, 2016, 2017; Hyvärinen et al., 2019). The main idea is to introduce additional data called the *complementary data* in this paper, and alternative condition to mutual independence is conditional independence of the source components given complementary data. An example of the complementary data is time segment labels obtained by dividing time series data into a number of time segments (Hyvärinen and Morioka, 2016), while Hyvärinen and Morioka (2017) employed the history of time-series data as complementary data. With the complementary data, contrastive learning based on the logistic regression has been performed for nonlinear ICA.

Most of practical methods in contrastive learning have been based on maximum likelihood estimation (MLE): In logistic regression, the conditional probability (i.e., posterior probability) is estimated by MLE, while maximizing a variational lower-bound of MI can be interpreted as performing MLE as we show later. MLE has a number of excellent properties such the asymptotic efficiency (Wasserman, 2006), but it is known not to be robust against outliers: If data is contaminated by outliers, MLE often suffers from a severe bias. Thus, estimation of existing methods for maximization of MI and nonlinear ICA might be magnified by outliers. This is a serious problem because outliers are observed in many practical situations. For instance, the contamination of outliers has been an important issue in functional MRI data (Poldrack, 2012) to which ICA methods have been applied (Monti et al., 2020).

This paper first shows that unsupervised representation learning on three frameworks can be performed through density ratio estimation. Previous work has already discussed that contrastive learning can be seen as density ratio estimation when popular objective functions are used for classification such as the cross entropy (Nguyen et al., 2008; Belghazi et al., 2018; van den Oord et al., 2018). Our primary contributions are to perform theoretical analysis on the three frameworks in terms of density ratio estimation. The first framework to be tackled is maximization of MI. We show that density ratio estimation is necessary and sufficient for maximization of MI under some conditions. This result supports the intuitive belief above between contrastive learning and maximization of MI, and clarifies when and how contrastive learning can be considered as performing maximization of MI. The second framework is nonlinear ICA, where we provide two new proofs for source recovery in terms of density ratio estimation. The key point is that the recovery conditions in both proofs include a novel insight, which has not been seen in previous work of nonlinear ICA (Hyvärinen et al., 2019): The dimensionality of complementary data is an important factor for source recovery. This insight is clearly supported by experimental results that the latent sources are more accurately recovered as the dimensionality of complementary data increases. Furthermore, one of the proofs can be seen as a generalization of Hyvärinen and Morioka (2017). The third

framework is *nonlinear subspace estimation* proposed in this paper. This framework is inspired by nonlinear ICA, and aimed at estimating a nonlinear subspace of lower-dimensional latent source components. To this end, we propose a novel generative model where data is generated as a nonlinear mixing of lower-dimensional source components and nuisance variables. Unlike nonlinear ICA, the source components are not necessarily assumed to be conditionally independent, and thus the proposed generative model is more general than nonlinear ICA in terms of the conditional independence. Moreover, as in nonlinear ICA, with density ratio estimation, we establish theoretical conditions that a nonlinear subspace of the lower-dimensional source components, which is separated from the nuisance variables, can be estimated.

Motivated by the density-ratio view, we propose an outlier-robust method for unsupervised representation learning. The proposed method performs density ratio estimation based on a robust alternative to the KL-divergence called the  $\gamma$ -divergence (Fujisawa and Eguchi, 2008), which has a favorable robustness property expressed as the *strong robustness* (Cichocki and Amari, 2010; Amari, 2016): The latent bias caused from outliers can be small even in the case of heavy contamination of outliers. Furthermore, we develop a new method of nonlinear ICA by applying an existing variational lower-bound of MI (Belghazi et al., 2018). Then, we theoretically investigate outlier-robustness of these methods. Finally, we numerically demonstrate that the proposed method based on the  $\gamma$ -divergence is very robust against outliers both in nonlinear ICA and a downstream task for linear classification, while the nonlinear ICA method based on the variational lower-bound is experimentally shown to perform better than existing ICA methods when the number of data is small.

This paper is organized as follows<sup>1</sup>: Section 2 formulates the problem of density ratio estimation for unsupervised representation learning, and reviews existing works for contrastive learning, variational MI estimation and nonlinear ICA. Section 3 theoretically shows that unsupervised representation learning on the three frameworks can be performed through density ratio estimation, and discusses theoretical contributions in each of the frameworks. Section 4 proposes two practical methods for unsupervised representation learning and theoretically analyzes them in terms of outlier-robustness. Section 5 numerically demonstrates usefulness of the proposed methods in nonlinear ICA and a downstream task for linear classification. Section 6 concludes this paper.

## 2. Problem formulation and background

This section first formulates the problem of unsupervised representation learning based on density ratio estimation, and then reviews existing works for contrastive learning, variational estimation of mutual information and nonlinear independent component analysis.

### 2.1 Problem formulation of density ratio estimation for representation learning

Suppose that we are given  $T$  pairs of two data samples drawn from the joint distribution with density  $p(\mathbf{x}, \mathbf{u})$ :

$$\mathcal{D} := \left\{ (\mathbf{x}(t)^\top, \mathbf{u}(t)^\top)^\top \mid \mathbf{x}(t) = (x_1(t), \dots, x_{D_x}(t))^\top, \mathbf{u}(t) = (u_1(t), \dots, u_{D_u}(t))^\top \right\}_{t=1}^T, \quad (1)$$

where  $\mathbf{x}(t)$  and  $\mathbf{u}(t)$  denote the  $t$ -th observations of  $\mathbf{x}$  and  $\mathbf{u}$ , respectively. Throughout the paper, we call  $\mathbf{x}$  and  $\mathbf{u}$  as *input* and *complementary* data, respectively. Our primary goal is to estimate

---

1. A very preliminary version of this paper was published in Sasaki et al. (2020).

Table 1: List of notations used in Section 2.

|  |  |
|--|--|
| $\mathbf{x} \in \mathbb{R}^{D_x}$  | Input data   |
| $\mathbf{u} \in \mathbb{R}^{D_u}$  | Complementary data   |
| $D_x$  | Dimensionality of input data   |
| $D_u$  | Dimensionality of complementary data   |
| $p(\mathbf{x}, \mathbf{u})$  | Joint probability density function of $\mathbf{x}$ and $\mathbf{u}$                      |
| $p(\mathbf{x})$  | Marginal probability density function of $\mathbf{x}$                                    |
| $p(\mathbf{u})$  | Marginal probability density function of $\mathbf{u}$                                    |
| $\perp$  | Statistical independence   |
| $\mathbf{u}^* \in \mathbb{R}^{D_u}$  | $\mathbf{u}^* \sim p(\mathbf{u})$ and $\mathbf{u}^* \perp \mathbf{x}$                    |
| $\mathbf{h}_x(\mathbf{x}) : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{d_x}$ | Representation function of $\mathbf{x}$  |
| $\mathbf{h}_u(\mathbf{u}) : \mathbb{R}^{D_u} \rightarrow \mathbb{R}^{d_u}$ | Representation function of $\mathbf{u}$  |
| $d_x$  | Dimensionality of representation function $\mathbf{h}_x(\mathbf{x})$                     |
| $d_u$  | Dimensionality of representation function $\mathbf{h}_u(\mathbf{u})$                     |
| $r(\mathbf{x}, \mathbf{u})$  | Model to approximate $\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$ |
| $\psi : \mathbb{R}^{d_x} \times \mathbb{R}^{d_u} \rightarrow \mathbb{R}$   | Function in $r(\mathbf{x}, \mathbf{u})$  |
| $a : \mathbb{R}^{d_x} \rightarrow \mathbb{R}$                              | Function in $r(\mathbf{x}, \mathbf{u})$  |
| $b : \mathbb{R}^{d_u} \rightarrow \mathbb{R}$                              | Function in $r(\mathbf{x}, \mathbf{u})$  |
| $J_{LR}$   | Cross entropy in logistic regression   |

$\mathbf{h}_x(\mathbf{x}) = (h_{x,1}(\mathbf{x}), \dots, h_{x,d_x}(\mathbf{x}))^\top$  called a *representation function* of  $\mathbf{x}$  such that the logarithmic ratio of the joint density to the product of the marginal densities  $p(\mathbf{x})$  and  $p(\mathbf{u})$ ,

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}, \quad (2)$$

is accurately approximated up to a constant under the following basic form of a model  $r(\mathbf{x}, \mathbf{u})$ :

$$r(\mathbf{x}, \mathbf{u}) := \psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x})) + b(\mathbf{h}_u(\mathbf{u})), \quad (3)$$

where  $\mathbf{h}_u(\mathbf{u}) = (h_{u,1}(\mathbf{u}), \dots, h_{u,d_u}(\mathbf{u}))^\top$  is a representation function of  $\mathbf{u}$ , and  $\psi(\cdot, \cdot)$ , and  $a(\cdot)$  and  $b(\cdot)$  are scalar functions. For nonlinear ICA,  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  is slightly modified: In Theorem 2,  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  is expressed by the sum of elementwise functions with respect to  $\mathbf{h}_x(\mathbf{x})$  (not  $\mathbf{h}_u(\mathbf{u})$ ) as  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) = \sum_{i=1}^{d_x} \psi_i(h_{x,i}(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$ . All of  $\mathbf{h}_u$ ,  $\psi$ ,  $a$  and  $b$  are also estimated from  $\mathcal{D}$ . Here, we assume that the dimensionalities of the representation functions are smaller than or equal to ones of input and complementary data, i.e.,  $d_u \leq D_u$  and  $d_x \leq D_x$ . Furthermore, in order to simplify the review of existing works and our theoretical results in Section 3, it is supposed that data samples is *not* contaminated by outliers, while we propose a practical method and theoretically investigate its robustness under the contamination of outliers in Section 4.

Function  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  in (3) takes a role of capturing the statistical dependencies between  $\mathbf{x}$  and  $\mathbf{u}$  with dimensionality reduction, and has been modeled previously such as  $\psi(\mathbf{h}_x, \mathbf{h}_u) = \mathbf{h}_x^\top \mathbf{h}_u$  (Bachman et al., 2019) and  $\psi(\mathbf{h}_x, \mathbf{h}_u) = \mathbf{h}_x^\top \mathbf{W} \mathbf{h}_u$  (van den Oord et al., 2018; Tian et al., 2019) where  $\mathbf{W}$  is a  $d_x$  by  $d_u$  matrix. More generally,  $\psi(\mathbf{h}_x, \mathbf{h}_u)$  can be modeled by a feedforward neural network (Arandjelovic and Zisserman, 2017). Scalar functions of  $a(\mathbf{h}_x(\mathbf{x}))$  and  $b(\mathbf{h}_u(\mathbf{u}))$  express any functions in the log-density ratio (2) that depend only on either  $\mathbf{x}$  or  $\mathbf{u}$ . Section 3 often

assumes that there exist functions  $\psi^*$ ,  $\mathbf{h}_x^*$ ,  $\mathbf{h}_u^*$ ,  $a^*$  and  $b^*$  such that the logarithmic density ratio is universally approximated at  $\psi = \psi^*$ ,  $\mathbf{h}_x = \mathbf{h}_x^*$ ,  $\mathbf{h}_u = \mathbf{h}_u^*$ ,  $a = a^*$  and  $b = b^*$  in (3) as follows:

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})).$$

This universal approximation assumption would be realistic when the log-density ratio is a continuous function and neural networks are employed for modelling  $\psi$ ,  $\mathbf{h}_x$ ,  $\mathbf{h}_u$ ,  $a$  and  $b$  (Hornik, 1991). Section 4 proposes practical methods to approximate  $\psi^*$ ,  $\mathbf{h}_x^*$ ,  $\mathbf{h}_u^*$ ,  $a^*$  and  $b^*$  from data samples.

The approach of using the complementary data is recently becoming more popular as in multi-view learning (Li et al., 2018). Furthermore, it might not be so expensive to obtain complementary data samples, but they are rather generated from a single unlabeled dataset in many practical situations. Let us list the following examples:

- Denoising autoencoder makes complementary data samples  $\mathbf{u}(t)$  by injecting noises to input data samples  $\mathbf{x}(t)$  (Vincent et al., 2010).
- $\mathbf{x}(t)$  and  $\mathbf{u}(t)$  can be drawn by dividing a single video data into image (i.e., a video frame) and sound data samples at each time index  $t$ , respectively (Arandjelovic and Zisserman, 2017).
- Regarding time series data  $\mathbf{x}(t)$  including video data, past input data is an example of the complementary data (e.g.,  $\mathbf{u}(t) = \mathbf{x}(t - 1)$ ) (Misra et al., 2016; Hyvärinen and Morioka, 2017).
- Suppose that  $\mathbf{x}(t)$  are image patches extracted from a large single image where the index  $t$  conveys some positional information of patches. Then, pairs of image patches at  $t$  and  $t' (\neq t)$  can be used as input and complementary data samples, that is,  $\mathbf{u}(t) = \mathbf{x}(t')$  (Noroozi and Favaro, 2016).
- Clustering labels to input data samples  $\mathbf{x}(t)$  can be used as the complementary data samples  $\mathbf{u}(t)$  (Caron et al., 2018).
- In Hjelm et al. (2019),  $\mathbf{x}(t)$  is an image, while  $\mathbf{u}(t)$  is a smaller image patch extracted from the single image  $\mathbf{x}(t)$ .

Density ratio estimation is useful particularly when neural networks are used for  $r(\mathbf{x}, \mathbf{u})$  because the density ratio is invariant under any invertible transformations or reparametrizations of  $\mathbf{x}$  and/or  $\mathbf{u}$ . This invariant property has been exploited by noise contrastive estimation as well (Gutmann and Hyvärinen, 2012). In addition to the practical usefulness, in this paper, we theoretically show that three frameworks for unsupervised representation learning can be performed by density ratio estimation.

## 2.2 Contrastive learning and variational estimation of mutual information

In order to estimate the representation functions, contrastive learning solves a classification problem where class labels are automatically generated from unlabelled data. A common setting is based on the following two datasets:

$$\mathcal{D}_+ := \{(\mathbf{x}(t), \mathbf{u}(t))\}_{t=1}^T \sim p(\mathbf{x}, \mathbf{u}) \quad \text{vs.} \quad \mathcal{D}_- := \{(\mathbf{x}(t), \mathbf{u}^*(t))\}_{t=1}^T \sim p(\mathbf{x})p(\mathbf{u}), \quad (4)$$

where  $\mathbf{u}^*(t)$  is a  $D_u$ -dimensional data vector sampled from the marginal density of  $\mathbf{u}$ , and thus independent to  $\mathbf{x}(t)$ . In practice,  $\mathbf{u}^*(t)$  can be generated by randomly shuffling  $\mathbf{u}(t)$  with respect to  $t$  under the i.i.d. assumption. Interestingly, the random shuffling has been heuristically used in a number of previous works (Misra et al., 2016; Lee et al., 2017). One of the most popular objective functions in contrastive learning is the following cross entropy for binary classification used in logistic regression:

$$J_{\text{LR}}(r) := -E_{\mathbf{x}\mathbf{u}} \left[ \log \frac{e^{r(\mathbf{X}, \mathbf{U})}}{1 + e^{r(\mathbf{X}, \mathbf{U})}} \right] - E_{\mathbf{x} \times \mathbf{u}} \left[ \log \frac{1}{1 + e^{r(\mathbf{X}, \mathbf{U})}} \right], \quad (5)$$

where  $E_{\mathbf{x}\mathbf{u}}$  and  $E_{\mathbf{x} \times \mathbf{u}}$  denote the expectations over  $p(\mathbf{x}, \mathbf{u})$  and  $p(\mathbf{x})p(\mathbf{u})$ , respectively. The representation functions  $\mathbf{h}_x(\mathbf{x})$  and  $\mathbf{h}_u(\mathbf{u})$  in  $r(\mathbf{x}, \mathbf{u})$  can be estimated by minimizing the empirical version of  $J_{\text{LR}}(r)$ . Logistic regression has been previously used to estimate a density ratio (Sugiyama et al., 2012), and the minimizer of  $J_{\text{LR}}(r)$  with respect to  $r$  is equal to

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$$

up to a constant. Intuitively, minimizing  $J_{\text{LR}}(r)$  enables to well-capture statistical dependencies between  $\mathbf{x}$  and  $\mathbf{u}$  by contrasting  $\mathcal{D}_+$  with  $\mathcal{D}_-$ , and thus to estimate data representations  $\mathbf{h}_x(\mathbf{x})$  and  $\mathbf{h}_u(\mathbf{u})$  having high mutual information. We theoretically justify this intuition, and clarifies when contrastive learning based on  $J_{\text{LR}}(r)$  can be considered to maximize mutual information.

*Infomax* is a classical framework for unsupervised representation learning, and advocates maximizing mutual information (MI) between input data and its representation (Linsker, 1989; Bell and Sejnowski, 1995). However, in practice, MI usually requires density estimation, which makes it hard to apply neural networks because of the notorious partition function problem. Recent promising approach employs complementary data and alternatively maximizes variational lower-bounds of MI between input and complementary data. For instance, the following lower bounds are often employed:

$$I(\mathcal{X}, \mathcal{U}) \geq E_{\mathbf{x}\mathbf{u}}[r(\mathbf{X}, \mathbf{U})] - E_{\mathbf{x} \times \mathbf{u}}[e^{r(\mathbf{X}, \mathbf{U})-1}] \quad (\text{Nguyen et al., 2008; Sugiyama et al., 2008}) \quad (6)$$

$$I(\mathcal{X}, \mathcal{U}) \geq E_{\mathbf{x}\mathbf{u}}[r(\mathbf{X}, \mathbf{U})] - \log E_{\mathbf{x} \times \mathbf{u}}[e^{r(\mathbf{X}, \mathbf{U})}] \quad (\text{Ruderman et al., 2012; Belghazi et al., 2018}), \quad (7)$$

where  $I(\mathcal{X}, \mathcal{U})$  denotes mutual information between  $\mathbf{x}$  and  $\mathbf{u}$ , and is defined by

$$I(\mathcal{X}, \mathcal{U}) := \int p(\mathbf{x}, \mathbf{u}) \log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} d\mathbf{x}d\mathbf{u}.$$

For other lower-bounds of MI, we refer to Poole et al. (2019). The key advantage is that these lower bounds enable us to employ neural networks without any special efforts, while it comes at the price for a statistical limitation that these lower-bounds may require an exponentially number of samples to accurately estimate MI (McAllester and Stratos, 2020, Theorem 3.1). As in logistic regression (5), the maximizers of both lower bounds in (6) and (7) with respect to  $r$  have been shown to be equal to

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$$

up to constants (Nguyen et al., 2008; Ruderman et al., 2012). Thus, these lower bounds can be used for density ratio estimation. The primary purpose of maximizing these lower-bounds is to estimate mutual information between  $\boldsymbol{x}$  and  $\boldsymbol{u}$ , yet it has been believed that maximizing these lower bounds leads to maximization of mutual information between  $\boldsymbol{h}_x(\boldsymbol{x})$  and  $\boldsymbol{h}_u(\boldsymbol{u})$ . Here, we provide a theoretically rigorous support to this belief.

### 2.3 Nonlinear independent component analysis

Nonlinear ICA is a solid framework for unsupervised representation learning, and assumes that data  $\boldsymbol{x} = (x_1, \dots, x_{D_x})^\top$  is generated as the following nonlinear mixing of the latent source  $\boldsymbol{s} = (s_1, \dots, s_{D_x})^\top$ :

$$\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{s}), \quad (8)$$

where  $\boldsymbol{f}(\boldsymbol{s}) = (f_1(\boldsymbol{s}), \dots, f_{D_x}(\boldsymbol{s}))^\top$ , and  $\boldsymbol{f}$  is an invertible function. The problem is to recover (or identify) the latent source components  $s_i$ . In the case of the linear mixing, i.e.,  $\boldsymbol{f}(\boldsymbol{s}) = \boldsymbol{A}\boldsymbol{s}$  with an invertible matrix  $\boldsymbol{A}$ , the latent source  $\boldsymbol{s}$  can be recovered up to the permutation (i.e., ordering) and scales of  $s_1, s_2, \dots, s_{D_x}$  when they are mutual independent and follow a nonGaussian density (Comon, 1994). However, the problem of nonlinear ICA has been proved to be seriously illposed under the same condition as the linear case because there exist an infinite number of decompositions of a random vector into mutually independent variables (Hyvärinen and Pajunen, 1999; Locatello et al., 2019). Nonetheless, a number of methods for nonlinear ICA have been previously proposed (Tan et al., 2001; Almeida, 2003; Blaschke et al., 2007), but most of them lack theoretical guarantees for source recovery and it is thus unclear to what extent these methods can recover the source components under the nonlinear mixing function  $\boldsymbol{f}$ .

Recently, novel recovery conditions for nonlinear ICA have been established (Sprekeler et al., 2014; Hyvärinen and Morioka, 2016, 2017; Hyvärinen et al., 2019). The key condition alternative to mutual independence is conditional independence of  $s_1, s_2, \dots, s_{D_x}$  given some complementary data  $\boldsymbol{u}$ . Time contrastive learning (TCL) employs time segment labels as complementary data samples, and assumes that the conditional density of the source  $\boldsymbol{s}$  given a time segment label is conditionally independent and belongs to an exponential family (Hyvärinen and Morioka, 2016). Then, the representation function  $\boldsymbol{h}_x(\boldsymbol{x})$  learned by the multinomial logistic regression has been shown to asymptotically correspond to elementwise nonlinear functions of the latent source components  $s_1, s_2, \dots, s_{D_x}$  up to a linear transformation. Hyvärinen et al. (2019) removed the exponential family assumption in TCL and has established recovery conditions in terms of more general conditional densities (nonexponential family). Then, it was proved that the representation function  $\boldsymbol{h}_x(\boldsymbol{x})$  learned by contrastive learning based on  $J_{\text{LR}}(r)$  (i.e., logistic regression) is asymptotically equal to the latent source components  $s_1, s_2, \dots, s_{D_x}$  up to their permutation and elementwise invertible functions. Details of the recovery conditions in Hyvärinen et al. (2019) are discussed in Section 3.2. Nonlinear ICA has been applied to causal analysis (Monti et al., 2020; Wu and Fukumizu, 2020) and transfer learning (Teshima et al., 2020).

This paper provides two new recovery proofs for nonlinear ICA, both of which include a novel insight: The dimensionality of complementary data is an important factor for source recovery. This insight is clearly supported by numerical experiments. Furthermore, we propose a novel generative model where data is generated as a nonlinear mixing of lower-dimensional latent source components and nuisance variables. The proposed generative model is more general than the one (8) in

Table 2: List of notations mainly used in Section 3. Wrt is an abbreviation of “with respect to”.

|   |   |
|---|---|
| $\mathbf{h}_x^*, \mathbf{h}_u^*$  | Optimal representation functions of $\mathbf{h}_x$ and $\mathbf{h}_u$   |
| $\psi^*, a^*, b^*$  | Optimal functions of $\psi, a$ and $b$  |
| $I(\mathcal{X}, \mathcal{U})$   | Mutual information (MI) between $\mathbf{x}$ and $\mathbf{u}$   |
| $I(\mathcal{Y}_x, \mathcal{Y}_u)$   | MI between $\mathbf{y}_x = \mathbf{h}_x(\mathbf{x})$ and $\mathbf{y}_u = \mathbf{h}_u(\mathbf{u})$  |
| $\mathbf{h}_x^\perp : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_x - d_x}$            | Function such that $(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_x^\perp(\mathbf{x}))$ is invertible.  |
| $\mathbf{h}_u^\perp : \mathbb{R}^{D_u} \rightarrow \mathbb{R}^{D_u - d_u}$            | Function such that $(\mathbf{h}_u(\mathbf{u}), \mathbf{h}_u^\perp(\mathbf{u}))$ is invertible.  |
| $\mathbf{s} \in \mathbb{R}^{d_x}$   | Vector of latent source components  |
| $\mathbf{f} : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_x}$                          | Nonlinear mixing function   |
| $q_i : \mathbb{R} \times \mathbb{R}^{D_u} \rightarrow \mathbb{R}$ (Proposition 2)     | Exponent of the conditional density $p(s_i   \mathbf{u})$   |
| $q_i', q_i''$   | First- and second-order derivatives of $q_i(t, \mathbf{u})$ wrt $t$   |
| $\nabla_{\mathbf{u}}$   | Gradient wrt $\mathbf{u}$   |
| $\mathbf{w} : \mathbb{R}^{D_x} \times \mathbb{R}^{D_u} \rightarrow \mathbb{R}^{2D_x}$ | $\mathbf{w}(\mathbf{v}, \mathbf{u}) := (q_1'(v_1, \mathbf{u}), \dots, q_{D_x}''(v_{D_x}, \mathbf{u}))^\top$   |
| $q_i : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ (Theorem 3)               | Exponent of the conditional density $p(s_i   \mathbf{u})$   |
| $\alpha_i^1, \alpha_i^2 : \mathbb{R} \rightarrow \mathbb{R}$                          | $\alpha_i^1(v) := \left. \frac{\partial^2 q_i(v, r)}{\partial v \partial r} \right _{r=\lambda_i(\mathbf{u}_1)}$ and $\alpha_i^2(v) := \left. \frac{\partial^2 q_i(v, r)}{\partial v \partial r} \right _{r=\lambda_i(\mathbf{u}_2)}$ |
| $\boldsymbol{\alpha} : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_x}$                 | $\boldsymbol{\alpha}(\mathbf{v}) := \left( \frac{\alpha_1^2(v_1)}{\alpha_1^1(v_1)}, \dots, \frac{\alpha_{D_x}^2(v_{D_x})}{\alpha_{D_x}^1(v_{D_x})} \right)^\top$  |
| $\mathbf{n} \in \mathbb{R}^{D_x - d_x}$   | Vector of nuisance variables  |

Table 3: Summary and comparison of theoretical conditions in maximization of MI, nonlinear ICA and nonlinear subspace estimation. Gen. model and source cond. are abbreviations of generative model and source condition, respectively. “-” means that the condition is not required. See the main text for more details.

|                  | Maximization of MI   | Nonlinear ICA   | Nonlinear subspace  |
|------------------|--|---|---|
| Form of $\psi^*$ | $\psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u}))$ | $\sum_{i=1}^{D_x} \psi_i^*(\mathbf{h}_{x,i}^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u}))$<br>or $\sum_{i=1}^{D_x} \psi_i^*(\mathbf{h}_{x,i}^*(\mathbf{x}), \mathbf{h}_{u,i}^*(\mathbf{u}))$ | $\psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u}))$                |
| Dim. assump.     | $d_x \leq D_x$ and $d_u \leq D_u$                                | $d_x = D_x$ and<br>$2D_x \leq D_u$ or $D_x \leq D_u$  | $d_x \leq D_x$ and $d_x \leq d_u$   |
| Gen. model       | -  | $\mathbf{x} = \mathbf{f}(\mathbf{s})$   | $\mathbf{x} = \mathbf{f}(\mathbf{s}, \mathbf{n})$                               |
| Source cond.     | -  | $s_i \perp s_j   \mathbf{u}$ ( $i \neq j$ )   | $\mathbf{s}, \mathbf{u} \perp \mathbf{n}$ and $\mathbf{s} \not\perp \mathbf{u}$ |

nonlinear ICA in the sense that the conditional independence of the latent source components is no longer assumed. Based on the proposed generative model, we establish theoretical conditions that complementary data enables us to automatically ignore the nuisance variables, and to estimate a nonlinear subspace related to only the lower-dimensional latent source components.

### 3. Three frameworks for unsupervised representation learning through density ratio estimation

This section shows that unsupervised representation learning on three frameworks can be performed by density ratio estimation, and discusses theoretical contributions on each of the frameworks. Ta-

bles 2 and 3 are a list of notations and summary of theoretical conditions in the three frameworks, respectively.

### 3.1 Maximization of mutual information

Maximization of mutual information (MI) (Barlow, 1961; Linsker, 1989; Bell and Sejnowski, 1995) is a classical yet recently retrieved framework for unsupervised representation learning combined with the recent development of deep neural networks (Hjelm et al., 2019; Tschannen et al., 2019). Density ratio estimation seems not to be strongly related, but our analysis implies that it is essential for maximization of MI.

Let us re-denote the representation functions of  $\mathbf{x}$  and  $\mathbf{u}$  by

$$\mathbf{y}_x := \mathbf{h}_x(\mathbf{x}) \quad \text{and} \quad \mathbf{y}_u := \mathbf{h}_u(\mathbf{u}),$$

respectively. The goal of maximization of MI is to find  $\mathbf{h}_x$  and  $\mathbf{h}_u$ , which maximize MI between  $\mathbf{y}_x$  and  $\mathbf{y}_u$  defined by

$$I(\mathcal{Y}_x, \mathcal{Y}_u) := \int p(\mathbf{y}_x, \mathbf{y}_u) \log \frac{p(\mathbf{y}_x, \mathbf{y}_u)}{p(\mathbf{y}_x)p(\mathbf{y}_u)} d\mathbf{y}_x d\mathbf{y}_u.$$

Data processing inequality shows that  $I(\mathcal{Y}_x, \mathcal{Y}_u)$  is a lower bound of MI between  $\mathbf{x}$  and  $\mathbf{u}$ , i.e.,

$$I(\mathcal{X}, \mathcal{U}) \geq I(\mathcal{Y}_x, \mathcal{Y}_u). \quad (9)$$

Inequality (9) indicates that any reparametrizations of  $\mathbf{x}$  and  $\mathbf{u}$  never exceed  $I(\mathcal{X}, \mathcal{U})$ , and  $I(\mathcal{Y}_x, \mathcal{Y}_u)$  is maximized at  $I(\mathcal{X}, \mathcal{U})$  if there exist such representation functions  $\mathbf{h}_x$  and  $\mathbf{h}_u$ .

The following theorem proved in Appendix A clarifies how and when contrastive learning and variational MI estimation can be regarded as performing maximization of MI, and establishes conditions on which  $I(\mathcal{Y}_x, \mathcal{Y}_u) = I(\mathcal{X}, \mathcal{U})$ , i.e.,  $I(\mathcal{Y}_x, \mathcal{Y}_u)$  is maximized:

**Theorem 1.** *We make the following assumptions:*

(A1)  $p(\mathbf{x}, \mathbf{u}) > 0$ ,  $p(\mathbf{x}) > 0$  and  $p(\mathbf{u}) > 0$ .

(A2) *There exist some functions  $\mathbf{h}_x^\perp : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_x - d_x}$  and  $\mathbf{h}_u^\perp : \mathbb{R}^{D_u} \rightarrow \mathbb{R}^{D_u - d_u}$  such that*

$$\begin{pmatrix} \mathbf{h}_x(\mathbf{x}) \\ \mathbf{h}_x^\perp(\mathbf{x}) \end{pmatrix} \in \mathbb{R}^{D_x} \quad \text{and} \quad \begin{pmatrix} \mathbf{h}_u(\mathbf{u}) \\ \mathbf{h}_u^\perp(\mathbf{u}) \end{pmatrix} \in \mathbb{R}^{D_u}$$

*are both invertible<sup>2</sup>.*

(A3) *There exist functions  $\psi^*$ ,  $\mathbf{h}_x^*$ ,  $\mathbf{h}_u^*$ ,  $a^*$  and  $b^*$  such that the following equation holds:*

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})). \quad (10)$$

*Then,  $I(\mathcal{Y}_x, \mathcal{Y}_u) = I(\mathcal{X}, \mathcal{U})$  at  $\mathbf{h}_x = \mathbf{h}_x^*$ , and  $\mathbf{h}_u = \mathbf{h}_u^*$ . Conversely, suppose that  $I(\mathcal{Y}_x, \mathcal{Y}_u) = I(\mathcal{X}, \mathcal{U})$  at  $\mathbf{h}_x = \mathbf{h}_x^*$  and  $\mathbf{h}_u = \mathbf{h}_u^*$  under Assumptions (A1-2). Then, there exist functions  $\psi^*$ ,  $a^*$  and  $b^*$  such that (10) holds.*

---

2. Invertibility means that there exist  $\mathbf{g}_x$  and  $\mathbf{g}_u$  such that  $\mathbf{x} = \mathbf{g}_x(\mathbf{y}_x, \mathbf{y}_x^\perp)$  and  $\mathbf{u} = \mathbf{g}_u(\mathbf{y}_u, \mathbf{y}_u^\perp)$  where  $\mathbf{y}_x^\perp := \mathbf{h}_x^\perp(\mathbf{x})$  and  $\mathbf{y}_u^\perp := \mathbf{h}_u^\perp(\mathbf{u})$ .

Theorem 1 implies that  $I(\mathcal{Y}_x, \mathcal{Y}_u)$  can be maximized through density ratio estimation, and thus motivates us to develop practical methods for maximization of MI through density ratio estimation. Eq.(10) is inspired by *sufficient dimension reduction* (Li, 1991; Cook, 1998; Fukumizu et al., 2004), which is a solid framework for *supervised* dimensionality reduction and whose goal is to find an informative lower-dimensional subspace to the output variable based on the conditional independence condition. As shown in the proof of Theorem 1, (10) can be rewritten as the following conditional independence conditions:

$$\mathbf{u} \perp \mathbf{x} \mid \mathbf{y}_x = \mathbf{h}_x^*(\mathbf{x}) \quad \text{and} \quad \mathbf{x} \perp \mathbf{u} \mid \mathbf{y}_u = \mathbf{h}_u^*(\mathbf{u}).$$

The conditional independence between  $\mathbf{u}$  and  $\mathbf{x}$  given  $\mathbf{y}_x = \mathbf{h}_x^*(\mathbf{x})$  implies that the lower-dimensional representation  $\mathbf{h}_x^*(\mathbf{x})$  has the same amount of information for  $\mathbf{u}$  as the original input data  $\mathbf{x}$ . The same implication holds the conditional independence between  $\mathbf{u}$  and  $\mathbf{x}$  given  $\mathbf{y}_u = \mathbf{h}_u^*(\mathbf{u})$  as well. Thus, accurately estimating the density ratio would yield representations of  $\mathbf{u}$  and  $\mathbf{x}$  possibly with minimum information loss.

An interesting point of Theorem 1 is that  $I(\mathcal{Y}_x, \mathcal{Y}_u) = I(\mathcal{X}, \mathcal{U})$  conversely implies the density ratio equation (10), and is useful for understanding when contrastive learning does not perform maximization of MI. In Section 2.2, we suppose that  $\mathbf{u}^*(t)$  are drawn from the marginal density  $p(\mathbf{u})$  in (4), but in practice,  $\mathbf{u}^*(t)$  are often taken from another dataset (Arandjelovic and Zisserman, 2017)<sup>3</sup> whose probability density is different from the marginal density  $p(\mathbf{u})$  and denoted by  $\tilde{p}(\mathbf{u})$ . Then, by assuming that the  $\mathbf{u}^*$  are independent to  $\mathbf{x}$ , contrastive learning related to density ratio estimation yields an estimate of

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})\tilde{p}(\mathbf{u})}$$

up to a constant, and (10) is never fulfilled. Thus, contrastive learning based on another marginal density  $\tilde{p}(\mathbf{u})$  might not be regarded as maximizing  $I(\mathcal{Y}_x, \mathcal{Y}_u)$  in general. On the other hand, as long as  $\mathbf{u}^*(t)$  are drawn from the marginal density  $p(\mathbf{u})$ , Theorem 1 would be a direct support to the belief that contrastive learning and variational MI estimation can be considered to perform maximization of MI because the popular objective functions are related to density ratio estimation as reviewed in Section 2.2.

Another simple practical point of Theorem 1 is the form of the right-hand side on (10): The right-hand side indicates that all terms have to be functions of  $\mathbf{h}_u^*$  and/or  $\mathbf{h}_x^*$ . Thus, even when  $a^*(\mathbf{h}_x^*(\mathbf{x}))$  is simply replaced with  $a(\mathbf{x})$  (i.e.,  $a$  is not a function of  $\mathbf{h}_x^*$ ), Theorem 1 does not hold.

### 3.2 Nonlinear ICA: A new insight for source recovery

Here, we perform two theoretical analyses for source recovery in nonlinear ICA with density ratio estimation (Proposition 2 and Theorem 3). These analyses shed light on a novel insight that the dimensionality of complementary data is an important factor for source recovery, which has not been revealed in previous work of nonlinear ICA. Furthermore, Theorem 3 can be regarded as a generalization of Theorem 1 in Hyvärinen and Morioka (2017).

3. For instance, in Arandjelovic and Zisserman (2017), the positive pairs of input and complimentary data samples are video frames and audio clips in the same video that overlap in time, while the negative pairs are randomly extracted from two different videos. Thus, the marginal density of the complimentary data (i.e., video frame or audio clip) can be different in the positive and negative data.

**Importance of the dimensionality of complementary data:** As reviewed in Section 2.3, nonlinear ICA assumes that input data  $\mathbf{x} = (x_1, x_2, \dots, x_{D_x})^\top$  is generated as a nonlinear mixing of the latent source components  $s_1, s_2, \dots, s_{D_x}$  as follows:

$$\mathbf{x} = \mathbf{f}(\mathbf{s}),$$

where  $\mathbf{s} = (s_1, s_2, \dots, s_{D_x})^\top$  and  $\mathbf{f} : \mathbb{R}^{D_x} \rightarrow \mathbb{R}^{D_x}$  is assumed to be invertible. Then, the goal is to recover the latent source components  $s_i$ . To this end, recent work of nonlinear ICA employs complementary data  $\mathbf{u}(t)$  in addition to input data  $\mathbf{x}(t)$ . For instance, by regarding  $\mathbf{x}(t)$  as time series data at the time index  $t$ , Hyvärinen and Morioka (2017) use past input data as  $\mathbf{u}(t)$  (e.g.,  $\mathbf{u}(t) = \mathbf{x}(t-1)$ ).

We first establish the following theorem showing that the latent source components can be recovered up to their permutation (i.e., ordering) and elementwise invertible functions, and that density ratio estimation plays an important role in nonlinear ICA as well:

**Proposition 2.** *Suppose that  $d_x = D_x$ . We further make the following assumptions:*

(B1) *The latent source components  $s_i$  are conditionally independent given  $\mathbf{u}$ . More specifically, the conditional density of  $\mathbf{s}$  given  $\mathbf{u}$  takes the following form:*

$$\log p(\mathbf{s}|\mathbf{u}) = \sum_{i=1}^{D_x} q_i(s_i, \mathbf{u}) - \log Z(\mathbf{u}),$$

where  $Z(\mathbf{u})$  denotes the partition function and  $q_i$  are differentiable functions.

(B2) *Input data  $\mathbf{x}$  is generated according to (8) where the mixing function  $\mathbf{f}$  is invertible.*

(B3) *Dimensionality of complementary data is twice larger than or twice as large as input data, i.e.,  $2D_x \leq D_u$ .*

(B4) *There exists a single point  $\mathbf{u}_1$  such that the rank of  $\nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u}) \in \mathbb{R}^{2D_x \times D_u}$  at  $\mathbf{u} = \mathbf{u}_1$  is  $2D_x$  for all  $\mathbf{v} = (v_1, v_2, \dots, v_{D_x})^\top$  where  $\nabla_{\mathbf{u}}$  denotes the differential operator with respect to  $\mathbf{u}$ , and*

$$\mathbf{w}(\mathbf{v}, \mathbf{u}) := (q'_1(v_1, \mathbf{u}), \dots, q'_{D_x}(v_{D_x}, \mathbf{u}), q''_1(v_1, \mathbf{u}), \dots, q''_{D_x}(v_{D_x}, \mathbf{u}))^\top \in \mathbb{R}^{2D_x},$$

with  $q'_i(t, \mathbf{u}) := \frac{\partial}{\partial t} q_i(t, \mathbf{u})$  and  $q''_i(t, \mathbf{u}) := \frac{\partial^2}{\partial t^2} q_i(t, \mathbf{u})$ .

(B5) *There exist functions  $\psi_i^*$ ,  $\mathbf{h}_x^*$ ,  $\mathbf{h}_u^*$ ,  $a^*$  and  $b^*$  such that the following equation holds:*

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \sum_{i=1}^{D_x} [\psi_i^*(h_{x,i}^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u}))] + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})), \quad (11)$$

where  $\psi_i^*$  are differentiable functions and  $\mathbf{h}_x^*(\mathbf{x}) := (h_{x,1}^*(\mathbf{x}), \dots, h_{x,D_x}^*(\mathbf{x}))^\top$  is invertible.

Then, under Assumptions (B1-5), the representation function  $\mathbf{h}_x^*(\mathbf{x})$  is equal to  $\mathbf{s}$  up to a permutation and elementwise invertible functions.

The proof is given in Appendix B. We essentially followed the proof of Theorem 1 in Hyvärinen et al. (2019) and derived the same conclusion. Here, the main difference is Assumptions (B3-4), which give a new insight for source recovery. Hyvärinen et al. (2019) adopted an alternative assumption to Assumptions (B3-4) called the *assumption of variability*, which assumes that there exist  $2D_x + 1$  points,  $\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{2D_x}$ , such that the following  $2D_x$  vectors are linearly independent:

$$\mathbf{w}(\mathbf{v}, \mathbf{u}_j) - \mathbf{w}(\mathbf{v}, \mathbf{u}_0) \text{ for } j = 1, \dots, 2D_x \text{ are linearly independent for all } \mathbf{v}.$$

The dimensionality assumption  $2D_x \leq D_u$  in Assumption (B3) would correspond to the  $2D_x$  vectors in the assumption of variability, but, in contrast, sheds light on a novel insight: The dimensionality of complementary data is an important factor for source recovery, which has not been revealed in previous work of nonlinear ICA. In fact, we numerically demonstrate that the accuracy of source recovery clearly depends on the dimensionality of complementary data in Section 5.1.2. Furthermore, Assumption (B3) is practically useful because it can be checked very easily. Regarding Assumption (B4), it implies that the complementary data  $\mathbf{u}$  is clearly dependent to  $\mathbf{s}$ : When the  $i$ -th element in  $\mathbf{u}$  is independent to  $\mathbf{s}$  and  $2D_x = D_u$  (i.e.,  $\nabla_{\mathbf{u}}\mathbf{w}(\mathbf{v}, \mathbf{u})$  is a square matrix), the  $i$ -th column in  $\nabla_{\mathbf{u}}\mathbf{w}(\mathbf{v}, \mathbf{u})$  is the zero vector, and thus the rank assumption in Assumption (B4) is never satisfied.

According to Hyvärinen et al. (2019), the assumption of variability implies that the underlying conditional density  $p(\mathbf{s}|\mathbf{u})$  has to be diverse and complex to recover the source components. As in Theorem 2 in Hyvärinen et al. (2019), we show that Assumption (B4) includes the same implication under the following exponential family:

$$\log p(\mathbf{s}|\mathbf{u}) = \sum_{i=1}^{D_x} \sum_{k=1}^K \lambda_{ik}(\mathbf{u}) q_{ik}(s_i) - \log Z(\mathbf{u}), \quad (12)$$

where  $\lambda_{ik}$  and  $q_{ik}$  are some scalar functions and  $K$  is a positive integer. Appendix C proves that

- When  $K = 1$ ,  $\text{rank}(\nabla_{\mathbf{u}}\mathbf{w}(\mathbf{v}, \mathbf{u})) \leq D_x$  where  $\text{rank}(\cdot)$  denotes the rank of a matrix.
- When  $K > 1$ ,  $\text{rank}(\nabla_{\mathbf{u}}\mathbf{w}(\mathbf{v}, \mathbf{u})) \leq 2D_x$

Thus, when  $K = 1$ , the  $2D_x$  rank assumption in Assumption (B4) is never fulfilled. This implies that in order to recover the source components, the conditional density has to be diverse and complex such as the exponential family (12) with a relatively large mixture number  $K$ .

**A milder dimensionality assumption:** The dimensionality assumption  $2D_x \leq D_u$  in Assumption (B3) might be strong because we need to have relatively high-dimensional complementary data. However, this assumption can be relaxed by restricting the underlying conditional density of  $\mathbf{s}$  given  $\mathbf{u}$ . The following theorem is based on a milder dimensionality assumption than Assumption (B3), and can be seen as a generalization of Theorem 1 in Hyvärinen and Morioka (2017):

**Theorem 3.** *Suppose that  $d_x = d_u = D_x$ . We make the following assumptions:*

(B'1) *The latent source components  $s_i$  are conditionally independent given  $\mathbf{u}$ , and the conditional density of  $\mathbf{s}$  given  $\mathbf{u}$  takes the following form:*

$$\log p(\mathbf{s}|\mathbf{u}) = \sum_{i=1}^{D_x} q_i(s_i, \lambda_i(\mathbf{u})) - \log Z(\mathbf{u}), \quad (13)$$

where  $q_i$  and  $\lambda_i$  are differentiable functions for  $i = 1, \dots, D_x$ .

(B'2) Input Data  $\mathbf{x}$  is generated according to (8) where the mixing function  $\mathbf{f}$  is invertible.

(B'3) Dimensionality of complementary data is larger than or equal to input data, i.e.,  $D_x \leq D_u$ .

(B'4) There exist two points,  $\mathbf{u}_1$  and  $\mathbf{u}_2$ , such that  $\alpha_i^1(v) := \left. \frac{\partial^2 q_i(v,r)}{\partial v \partial r} \right|_{r=\lambda_i(\mathbf{u}_1)} \neq 0$  and  $\alpha_i^2(v) := \left. \frac{\partial^2 q_i(v,r)}{\partial v \partial r} \right|_{r=\lambda_i(\mathbf{u}_2)} \neq 0$  for all  $i$  and  $v$ .

(B'5) There exist  $D_x$  points,  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{D_x}$ , such that the  $D_x$  vectors  $\boldsymbol{\alpha}(\mathbf{v}_1), \boldsymbol{\alpha}(\mathbf{v}_2), \dots, \boldsymbol{\alpha}(\mathbf{v}_{D_x})$  are linearly independent where with  $\mathbf{v} := (v_1, v_2, \dots, v_{D_x})^\top$ ,

$$\boldsymbol{\alpha}(\mathbf{v}) := \left( \frac{\alpha_1^2(v_1)}{\alpha_1^1(v_1)}, \frac{\alpha_2^2(v_2)}{\alpha_2^1(v_2)}, \dots, \frac{\alpha_{D_x}^2(v_{D_x})}{\alpha_{D_x}^1(v_{D_x})} \right)^\top \in \mathbb{R}^{D_x}.$$

(B'6) There exist functions  $\psi_i^*, \mathbf{h}_x^*, \mathbf{h}_u^*, a^*$  and  $b^*$  such that the following equation holds:

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \sum_{i=1}^{D_x} [\psi_i^*(h_{x,i}^*(\mathbf{x}), h_{u,i}^*(\mathbf{u}))] + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})), \quad (14)$$

where  $\psi_i^*$  are differentiable functions,  $\mathbf{h}_x^*(\mathbf{x}) := (h_{x,1}^*(\mathbf{x}), \dots, h_{x,D_x}^*(\mathbf{x}))^\top$  is invertible, and  $\mathbf{h}_u^*(\mathbf{u}) := (h_{u,1}^*(\mathbf{u}), \dots, h_{u,D_x}^*(\mathbf{u}))^\top$ .

Then, under Assumptions (B'1-6), the representation function  $\mathbf{h}_x^*(\mathbf{x})$  is equal to  $\mathbf{s}$  up to a permutation and elementwise invertible functions.

The proof is given in Appendix D. The key point in Theorem 3 is that the dimensionality assumption  $D_x \leq D_u$  in Assumption (B'3) is milder than  $2D_x \leq D_u$  in Assumption (B3). However, this milder assumption comes at a cost of making the conditional density  $p(\mathbf{s}|\mathbf{u})$  less general than Proposition 2 and of restricting the form of the right-hand side on (14): The conditional density  $p(\mathbf{s}|\mathbf{u})$  is restricted into a form of pairwise combinations of  $s_i$  and  $\lambda_i(\mathbf{u})$  for  $i = 1 \dots, D_x$  in (13), and the right-hand side on (14) also takes a pairwise form of  $h_{x,i}^*(\mathbf{x})$  and  $h_{u,i}^*(\mathbf{u})$ .

Similarly as Assumption (B4) in Proposition 2, Assumption (B'5) also implies that the underlying conditional density (13) is diverse and complex. Indeed, when  $p(\mathbf{s}|\mathbf{u})$  belongs to the exponential family (12) in  $K = 1$ , the ratio  $\frac{\alpha_i^2(v_i)}{\alpha_i^1(v_i)}$  is equal to a constant  $\frac{\lambda_{i1}(\mathbf{u}_2)}{\lambda_{i1}(\mathbf{u}_1)}$  for all  $i$ . Then,  $\boldsymbol{\alpha}(\mathbf{v})$  is a constant vector, and cannot be linearly independent over the  $D_x$  points. Thus, Assumption (B'5) is never satisfied under the exponential family (12) in  $K = 1$ .

Theorem 3 generalizes Theorem 1 in Hyvärinen and Morioka (2017): By regarding  $\mathbf{x}(t), t = 1, \dots, T$  as time series data at the time index  $t$ , Theorem 1 in Hyvärinen and Morioka (2017) is a special case of Theorem 3 where  $\mathbf{u}(t) = \mathbf{x}(t-1)$  (thus, implicitly suppose  $D_x = D_u$  and Assumption (B'3) is satisfied),  $\boldsymbol{\lambda}(\mathbf{u}) = \mathbf{u}$ , and  $\mathbf{h}_u = \mathbf{h}_x$  (e.g., the same neural architecture with weight sharing). This generalization is not straightforward because we had to derive a new lemma (Lemma 10), and thus the proof is substantially different. Furthermore, by this generalization, Theorem 3, again, reveals that the dimensionality of complementary data is an important factor, which has not been seen in Theorem 1 of Hyvärinen and Morioka (2017).

Finally, we note a subtle difference of  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  in nonlinear ICA. Nonlinear ICA requires us to slightly modify  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  as follows: For Proposition 2,

$$\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) = \sum_{i=1}^{D_x} \psi_i(h_{x,i}(\mathbf{x}), h_{u,i}(\mathbf{u})), \quad (15)$$

while in Theorem 3,

$$\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) = \sum_{i=1}^{D_x} [\psi_i(h_{x,i}(\mathbf{x}), h_{u,i}(\mathbf{u}))]. \quad (16)$$

In contrast with maximization of mutual information (Theorem 1) as well as nonlinear subspace estimation (Theorem 4),  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  in (15) (or in (16)) is expressed as the sum of element-wise functions with respect to  $\mathbf{h}_x(\mathbf{x})$  (or  $\mathbf{h}_x(\mathbf{x})$  and  $\mathbf{h}_u(\mathbf{u})$ ). Thus, in practice, we need to slightly modify the form of  $\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))$  to (15) or (16) when performing nonlinear ICA.

### 3.3 Nonlinear subspace estimation with complementary data

Here, we first propose a new generative model where input data is generated as a nonlinear mixing of lower-dimensional latent source components and nuisance variables. Then, we establish some theoretical conditions to estimate a lower-dimensional nonlinear subspace of the latent source components only.

**A new generative model:** Let us consider the following novel generative model for input data  $\mathbf{x} \in \mathbb{R}^{D_x}$ :

$$\mathbf{x} = \mathbf{f}(\mathbf{s}, \mathbf{n}), \quad (17)$$

where  $\mathbf{f}$  is an invertible nonlinear mixing function,  $\mathbf{s} \in \mathbb{R}^{d_x}$  and  $\mathbf{n} \in \mathbb{R}^{D_x - d_x}$ . We further assume that

- Source components  $s_i$  are statistically independent to nuisance variables  $\mathbf{n}$ , i.e.,  $\mathbf{s} \perp \mathbf{n}$
- Complementary data is supposed to be available such that  $\mathbf{u} \perp \mathbf{n}$ , while  $\mathbf{u}$  depends on  $\mathbf{s}$ , i.e.,  $\mathbf{s} \not\perp \mathbf{u}$  where  $\not\perp$  denotes statistical dependence.

Based on the observations of  $\mathbf{x}$  and  $\mathbf{u}$ , the goal is to estimate a lower-dimensional nonlinear subspace of only  $\mathbf{s}$ , which is separated from  $\mathbf{n}$ . In contrast with nonlinear ICA, we do not necessarily assume that  $s_i$  are conditionally independent given  $\mathbf{u}$ . Thus, (17) is more general than the generative model (8) in nonlinear ICA in the sense that the conditional independence of  $s_i$  given  $\mathbf{u}$  is no longer assumed.

The new generative model (17) can be motivated by many practical situations. For example, time-series data such as brain signals (Dornhege et al., 2007) could be observed as a mixture of stationary and nonstationary components, and the nonstationarity could be important for a range of tasks such as change detection (Blythe et al., 2012). For this example, by using the past data in time series data as complementary data (e.g.,  $\mathbf{u}(t) = \mathbf{x}(t - 1)$ ), we may estimate a subspace of the useful nonstationary components. Another example is image data. The pixel values around the center of an image usually depends on surrounding pixels, while the pixels around the corners of the image are often almost independent to the other pixels (e.g., MNIST images). By using surrounding pixels as the complementary data, it would be very informative to estimate some lower-dimensional subspace for image data, which constitutes the fundamental part of image data.

**Estimating a lower-dimensional subspace of the source components:** First of all, it is important to understand whether we can estimate a lower-dimensional subspace of the latent sources  $\mathbf{s}$ , which is separated from nuisance variables  $\mathbf{n}$ . The following theorem gives conditions to estimate such a subspace as a (vector-valued) function of  $\mathbf{s}$  only:

**Theorem 4.** *Assume that*

(C1) *Input data  $\mathbf{x}$  is generated according to (17) where the mixing function  $\mathbf{f}$  is invertible,  $\mathbf{s} \perp \mathbf{n}$ ,  $\mathbf{s} \not\perp \mathbf{u}$  and  $\mathbf{u} \perp \mathbf{n}$ .*

(C2) *There exist functions  $\psi^*$ ,  $\mathbf{h}_x^*$ ,  $\mathbf{h}_u^*$ ,  $a^*$  and  $b^*$  such that the following equation holds:*

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})), \quad (18)$$

where  $\mathbf{h}_x^*$  is surjective.

(C3) *Dimensionality of  $\mathbf{h}_x^*(\mathbf{x})$  is smaller than or equal to  $\mathbf{h}_u^*(\mathbf{u})$ , i.e.,  $d_x \leq d_u$ .*

(C4) *There exists at least a single point  $\mathbf{u}_1$  such that the rank of  $\nabla_v \nabla_u \psi^*(\mathbf{v}, \mathbf{h}_u^*(\mathbf{u})) \in \mathbb{R}^{d_u \times d_x}$  is  $d_x$  at  $\mathbf{u} = \mathbf{u}_1$  for all  $\mathbf{v} \in \mathbb{R}^{d_x}$  where  $\mathbf{v} := \mathbf{h}_x^*(\mathbf{x})$ .*

Under Assumptions (C1-4), the representation function  $\mathbf{h}_x^*(\mathbf{x})$  is a nonlinear function of only  $\mathbf{s}$ .

The proof is given in Appendix E. The interesting point of Theorem 4 is that complementary data  $\mathbf{u}$  enables us to automatically ignore the nuisance variables  $\mathbf{n}$ , and Theorem 4 implies that a lower-dimensional nonlinear subspace only for  $\mathbf{s}$  can be estimated through density ratio estimation. Assumption (C3) is a necessary condition that the rank of  $\nabla_v \nabla_u \psi^*(\mathbf{v}, \mathbf{h}_u^*(\mathbf{u}))$  is  $d_x$ : By defining  $\mathbf{r} := \mathbf{h}_u^*(\mathbf{u})$ ,

$$\text{rank}(\nabla_v \nabla_u \psi^*(\mathbf{v}, \mathbf{h}_u^*(\mathbf{u}))) = \text{rank}(\mathbf{J}_r(\mathbf{u})^\top [\nabla_v \nabla_r \psi^*(\mathbf{v}, \mathbf{r})]) \leq \min(d_x, d_u),$$

where  $\mathbf{J}_r(\mathbf{u}) := \nabla_u \mathbf{r} \in \mathbb{R}^{d_u \times d_u}$  is the Jacobian of  $\mathbf{r}$  at  $\mathbf{u}$  and  $\nabla_v \nabla_r \psi^*(\mathbf{v}, \mathbf{r}) \in \mathbb{R}^{d_u \times d_x}$ . Thus, it must be  $d_x \leq d_u$  to satisfy Assumption (C4).

In order to understand the implication of Assumption (C4), we derive the following equation from (18) in Appendix F:

$$\mathbf{J}_v(\mathbf{s}) [\nabla_s \nabla_u \log p(\mathbf{s}|\mathbf{u})]_{\mathbf{u}=\mathbf{u}_1} = \nabla_v \nabla_u \psi^*(\mathbf{v}, \mathbf{u})|_{\mathbf{u}=\mathbf{u}_1}, \quad (19)$$

where  $\mathbf{J}_v(\mathbf{s}) := \nabla_s \mathbf{v} \in \mathbb{R}^{d_x \times d_x}$  is the Jacobian of  $\mathbf{v} = \mathbf{h}_x^*(\mathbf{f}(\mathbf{s}, \mathbf{n}))$  at  $\mathbf{s}$ . According to Lemma 11 and Assumption (C4), (19) ensures that the rank of  $\nabla_s \nabla_u \log p(\mathbf{s}|\mathbf{u})|_{\mathbf{u}=\mathbf{u}_1}$  is  $d_x$ . Next, as discussed in nonlinear ICA (Section 3.2), we investigate whether or not the rank of  $\nabla_s \nabla_u \log p(\mathbf{s}|\mathbf{u})|_{\mathbf{u}=\mathbf{u}_1}$  is  $d_x$  under the following exponential family:

$$\log p(\mathbf{s}|\mathbf{u}) = \sum_{j=1}^D \lambda_j(\mathbf{u}) q_j(\mathbf{s}) - \log Z(\mathbf{u}), \quad (20)$$

where  $\lambda_j$  and  $q_j$  are some scalar functions. Then, we have

$$\nabla_{\mathbf{s}} \nabla_{\mathbf{u}} \log p(\mathbf{s}|\mathbf{u}) = \sum_{j=1}^D \{\nabla_{\mathbf{u}} \lambda_j(\mathbf{u})\} \{\nabla_{\mathbf{s}} q_j(\mathbf{s})\}^{\top},$$

which indicates that by the rank factorization theorem,

$$\text{rank}(\nabla_{\mathbf{s}} \nabla_{\mathbf{u}} \log p(\mathbf{s}|\mathbf{u})) \leq \min(d_x, D)$$

where we used  $d_x \leq d_u \leq D_u$  from Assumption (C3). Thus, when  $D < d_x$ , Assumption (C4) is not fulfilled. As in nonlinear ICA, Assumption (C4) possibly implies that in order to estimate a nonlinear subspace of  $\mathbf{s}$ ,  $p(\mathbf{s}|\mathbf{u})$  has to be sufficiently complex such as the exponential family (20) with relatively large  $D (> d_x)$ .

Compared with nonlinear ICA in Section 3.2, Theorem 4 stands on a more general setting in the sense that the elements in  $\mathbf{s}$  are not necessarily assumed to be conditionally independent. Furthermore, the dimensionality condition  $d_x \leq d_u$  in Assumption (C3) is milder than Assumption (B3) in Proposition 2 and Assumption (B'3) in Theorem 3 because  $d_x$  and  $d_u$  are the dimensionalities of the representation functions  $\mathbf{h}_x^*$  and  $\mathbf{h}_u^*$  and assumed to be smaller than  $D_x$  and  $D_u$ , respectively. However, these milder conditions come at a price for losing the recover of each latent source component  $s_i$ : Nonlinear ICA recovers each component  $s_i$  up to a permutation and elementwise invertible function, while Theorem 4 only guarantees that a nonlinear function of  $\mathbf{s}$  can be estimated.

A similar generative model as (17) was proposed in multi-view learning (Gresele et al., 2019), which considers that two data,  $\mathbf{x}^{(1)}$  and  $\mathbf{x}^{(2)}$ , are generated as

$$\mathbf{x}^{(1)} = \mathbf{f}^{(1)}(\mathbf{s}) \quad \text{and} \quad \mathbf{x}^{(2)} = \mathbf{f}^{(2)}(\mathbf{g}(\mathbf{s}, \mathbf{n})), \quad (21)$$

where  $\mathbf{f}^{(1)}$  and  $\mathbf{f}^{(2)}$  are the mixing functions, and  $\mathbf{g}$  denotes an elementwise vector-valued function. The generative model (21) assumes that  $\mathbf{s} \perp \mathbf{n}$ ,  $s_i \perp s_j$  and  $n_i \perp n_j$  for  $i \neq j$  where  $n_i$  denotes the  $i$ -th element in  $\mathbf{n}$ . Then, Gresele et al. (2019) proved that each component  $s_i$  can be recovered up to elementwise invertible functions. In contrast, the assumptions in our generative model (17) can be more general because it is not assumed that the elements both in  $\mathbf{s}$  and  $\mathbf{n}$  are independent, i.e.,  $s_i \not\perp s_j$  and  $n_i \not\perp n_j$  for  $i \neq j$  in our generative model.

## 4. Estimation methods

Section 3 theoretically showed that density ratio estimation plays a key role in the three frameworks, and motivates us to develop practical methods for unsupervised representation learning through density ratio estimation. This section proposes two practical methods to estimate the log-density ratio based on neural networks. The first method employs the  $\gamma$ -cross entropy (Fujisawa and Eguchi, 2008), which is a robust variant of the standard cross entropy against outliers. For nonlinear ICA, the second one employs the Donsker-Varadhan variational estimation for mutual information (Ruderman et al., 2012; Belghazi et al., 2018). After describing the proposed methods, we investigate the outlier-robustness of the proposed methods. A list of notations mainly used in this section is given in Table 4.

Table 4: List of notations used in Section 4.

|  |   |
|--|---|
| $E_{\mathbf{x}\mathbf{u}}$                                 | Expectation over $p(\mathbf{x}, \mathbf{u})$  |
| $E_{\mathbf{x} \times \mathbf{u}}$                         | Expectation over $p(\mathbf{x})p(\mathbf{u})$   |
| $E_{\mathbf{x}}$   | Expectation over $p(\mathbf{x})$  |
| $E_{\mathbf{u}}$   | Expectation over $p(\mathbf{u})$  |
| $J_{\gamma}$   | $\gamma$ -cross entropy   |
| $\hat{J}_{\gamma}, \tilde{J}_{\gamma}$                     | Empirical approximations of $J_{\gamma}$  |
| $J_{\text{DV}}$  | Donsker-Varadhan variational lower-bound of MI  |
| $\hat{J}_{\text{DV}}, \tilde{J}_{\text{DV}}$               | Empirical approximations of $J_{\text{DV}}$   |
| $\mathbf{u}_p(t)$  | Random permutation of $\mathbf{u}(t)$ with respect to $t$   |
| $\bar{p}(\mathbf{x}, \mathbf{u})$                          | Contaminated joint probability density function of $\mathbf{x}$ and $\mathbf{u}$  |
| $\bar{p}(\mathbf{x})$                                      | Contaminated marginal probability density function of $\mathbf{x}$  |
| $\bar{p}(\mathbf{u})$                                      | Contaminated marginal probability density function of $\mathbf{u}$  |
| $\delta(\mathbf{x}, \mathbf{u})$                           | Joint probability density function for outliers   |
| $\delta(\mathbf{x}), \delta(\mathbf{u})$                   | Marginal probability density functions for outliers   |
| $\epsilon$   | Contamination ratio ( $0 \leq \epsilon \leq 1$ )  |
| $\bar{\mathbf{x}}, \bar{\mathbf{u}}$                       | Outlier points  |
| $\delta_{\bar{\mathbf{x}}}(\mathbf{x})$                    | Dirac delta function with a point mass at $\bar{\mathbf{x}}$  |
| $\delta_{\bar{\mathbf{u}}}(\mathbf{u})$                    | Dirac delta function with a point mass at $\bar{\mathbf{u}}$  |
| $r_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{u})$          | Model parametrized by $\boldsymbol{\theta}$   |
| $\mathbf{g}_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{u})$ | Gradient of $r_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{u})$ with respect to $\boldsymbol{\theta}$   |
| $\boldsymbol{\theta}^*$                                    | Parameters optimized over the (noncontaminated) densities   |
| $\boldsymbol{\theta}_{\epsilon}$                           | Parameters optimized over the contaminated densities  |
| $\text{IF}_{\text{DV}}$                                    | Influence function over $J_{\text{DV}}$   |
| $\mathbf{V}_{\text{DV}}$                                   | $\mathbf{V}_{\text{DV}} := E_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})^{\top}] - E_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})]E_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})]^{\top}$ |
| $\text{IF}_{\gamma}$                                       | Influence function over $J_{\gamma}$  |
| $\mathbf{V}_{\gamma}$                                      | $\mathbf{V}_{\gamma} := E_{\mathbf{x}\mathbf{u}}[\eta(\mathbf{X}, \mathbf{U})^{\frac{\gamma}{1+\gamma}} S(\mathbf{X}, \mathbf{U})^{\frac{1}{1+\gamma}} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})^{\top}]$  |
| $S(\mathbf{x}, \mathbf{u})$                                | $S(\mathbf{x}, \mathbf{u}) := \frac{1}{1+e^{(1+\gamma)r_{\boldsymbol{\theta}^*}(\mathbf{x}, \mathbf{u})}}$  |
| $\eta(\mathbf{x}, \mathbf{u})$                             | $\eta(\mathbf{x}, \mathbf{u}) := S(\mathbf{x}, \mathbf{u})(1 - S(\mathbf{x}, \mathbf{u}))$  |
| $\bar{E}_{\mathbf{x}\mathbf{u}}$                           | Expectation over $\bar{p}(\mathbf{x}, \mathbf{u})$  |
| $\bar{E}_{\mathbf{x} \times \mathbf{u}}$                   | Expectation over $\bar{p}(\mathbf{x})\bar{p}(\mathbf{u})$   |
| $\bar{J}_{\gamma}$   | $\gamma$ -cross entropy over $\bar{p}(\mathbf{x}, \mathbf{u})$ and $\bar{p}(\mathbf{x})\bar{p}(\mathbf{u})$   |

#### 4.1 Robust representation learning based on the $\gamma$ -cross entropy

The first method is based on the following  $\gamma$ -cross entropy for binary classification (Fujisawa and Eguchi, 2008; Hung et al., 2018). Let us recall the following two datasets:

$$\mathcal{D}_+ := \{(\mathbf{x}(t), \mathbf{u}(t))\}_{t=1}^T \sim p(\mathbf{x}, \mathbf{u}) \quad \text{vs.} \quad \mathcal{D}_- := \{(\mathbf{x}(t), \mathbf{u}^*(t))\}_{t=1}^T \sim p(\mathbf{x})p(\mathbf{u}).$$

By assigning class labels  $y = 0$  and  $y = 1$  to  $\mathcal{D}_+$  and  $\mathcal{D}_-$  respectively, the  $\gamma$ -cross entropy for posterior probability estimation can be formulated as

$$\begin{aligned} J_\gamma(f_+, f_-) &:= -\frac{1}{\gamma} \log \left[ \iint \sum_{y=0}^1 p(y, \mathbf{x}, \mathbf{u}) \left( \frac{\{f_+(\mathbf{x}, \mathbf{u})^{\gamma+1}\}^{1-y} \{f_-(\mathbf{x}, \mathbf{u})^{\gamma+1}\}^y}{f_+(\mathbf{x}, \mathbf{u})^{\gamma+1} + f_-(\mathbf{x}, \mathbf{u})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \right], \\ &= -\frac{1}{\gamma} \log \left[ p(y=0) E_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{f_+(\mathbf{X}, \mathbf{U})^{\gamma+1}}{f_+(\mathbf{X}, \mathbf{U})^{\gamma+1} + f_-(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] \right. \\ &\quad \left. + p(y=1) E_{\mathbf{x}\times\mathbf{u}} \left[ \left( \frac{f_-(\mathbf{X}, \mathbf{U})^{\gamma+1}}{f_+(\mathbf{X}, \mathbf{U})^{\gamma+1} + f_-(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] \right], \end{aligned}$$

where  $f_+$  and  $f_-$  are models for posterior probabilities,  $p(y=0)$  and  $p(y=1)$  are class probabilities, and we used the following relation based on the datasets  $\mathcal{D}_+$  and  $\mathcal{D}_-$ :

$$p(\mathbf{x}, \mathbf{u}|y=0) = p(\mathbf{x}, \mathbf{u}) \quad \text{and} \quad p(\mathbf{x}, \mathbf{u}|y=1) = p(\mathbf{x})p(\mathbf{u}).$$

By denoting  $\log \frac{f_+(\mathbf{x}, \mathbf{u})}{f_-(\mathbf{x}, \mathbf{u})}$  by  $r(\mathbf{x}, \mathbf{u})$  and assuming symmetric class probabilities (i.e.,  $p(y=0) = p(y=1) = \frac{1}{2}$ ), the  $\gamma$ -cross entropy can be written as

$$J_\gamma(r) := -\frac{1}{\gamma} \log \left[ E_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}}{1 + e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}} \right)^{\frac{\gamma}{\gamma+1}} \right] + E_{\mathbf{x}\times\mathbf{u}} \left[ \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}} \right)^{\frac{\gamma}{\gamma+1}} \right] \right], \quad (22)$$

where the term related to  $p(y=0)$  and  $p(y=1)$  is omitted because it is irrelevant to estimation of a model  $r$ . Since  $r(\mathbf{x}, \mathbf{u})$  is a model for the log-ratio of posterior probabilities,  $J_\gamma(r)$  is minimized at

$$\log \frac{p(y=0|\mathbf{x}, \mathbf{u})}{p(y=1|\mathbf{x}, \mathbf{u})} = \log \frac{p(\mathbf{x}, \mathbf{u}|y=0)p(y=0)}{p(\mathbf{x}, \mathbf{u}|y=1)p(y=1)} = \log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}.$$

Thus, the  $\gamma$ -cross entropy (22) can be used for density ratio estimation. As proven in Fujisawa and Eguchi (2008), the cross entropy in logistic regression can be obtained in the limit of the  $\gamma$ -cross entropy as follows:

$$\lim_{\gamma \rightarrow 0} J_\gamma(r) = J_{\text{LR}}(r),$$

indicating  $J_\gamma(r)$  can be regarded as a generalization of the cross entropy in logistic regression. The remarkable property of  $J_\gamma(r)$  is robustness against outliers. The positive parameter  $\gamma$  controls the

robustness, and a larger value of  $\gamma$  tends to be more robust to outliers. We theoretically characterize the robustness of the  $\gamma$ -cross entropy in the context of density ratio estimation later. On the other hand, as discussed in Fujisawa and Eguchi (2008), there seems to exist a trade-off between robustness and efficiency, indicating that outlier-robustness for the  $\gamma$ -cross entropy may come at a price of making estimation sample-inefficient. In fact, the experimental results in Section 5.1.1 partially demonstrate this trade-off: The proposed method based on the  $\gamma$ -cross entropy is the most outlier-robust method, but it is not the best in terms of sample-efficiency.

In practice, we empirically approximate  $J_\gamma(r)$  as

$$\hat{J}_\gamma(r) := -\frac{1}{\gamma} \log \left[ \frac{1}{T} \sum_{t=1}^T \left\{ \left( \frac{e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}(t))}}{1 + e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}(t))}} \right)^{\frac{\gamma}{\gamma+1}} + \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}_p(t))}} \right)^{\frac{\gamma}{\gamma+1}} \right\} \right],$$

where  $\mathbf{u}_p(t)$  denotes a random permutation of  $\mathbf{u}(t)$  with respect to  $t$ . Another empirical approximation is also possible as

$$\tilde{J}_\gamma(r) := -\frac{1}{\gamma} \log \left[ \frac{1}{T} \sum_{t=1}^T \left( \frac{e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}(t))}}{1 + e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}(t))}} \right)^{\frac{\gamma}{\gamma+1}} + \frac{1}{T^2} \sum_{t=1}^T \sum_{t'=1}^T \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}(t), \mathbf{u}(t'))}} \right)^{\frac{\gamma}{\gamma+1}} \right].$$

These empirical objective functions are minimized with a minibatch stochastic gradient method in Section 5.

## 4.2 Nonlinear ICA with the Donsker-Varadhan variational estimation

Our second method is based on variational estimation of mutual information (Ruderman et al., 2012; Belghazi et al., 2018), and employs the negative of the lower-bound in (7) as the objective function:

$$J_{\text{DV}}(r) := -E_{\mathbf{x}\mathbf{u}}[r(\mathbf{X}, \mathbf{U})] + \log(E_{\mathbf{x}\mathbf{u}}[e^{r(\mathbf{X}, \mathbf{U})}]). \quad (23)$$

As proved in Banerjee (2006) and Belghazi et al. (2018),  $J_{\text{DV}}(r)$  is minimized at  $\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$  up to some constant, and thus can be used for density ratio estimation.  $J_{\text{DV}}(r)$  has been already employed in the context of maximization of mutual information (Hjelm et al., 2019). Here, our contribution is to apply  $J_{\text{DV}}(r)$  to nonlinear ICA, and we numerically demonstrate its usefulness.

The objective function  $J_{\text{DV}}(r)$  has been independently derived in terms of density ratio estimation based on the KL-divergence (Sugiyama et al., 2008; Tsuboi et al., 2009). Since  $r(\mathbf{x}, \mathbf{u})$  is a model for  $\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$ , the following model  $p_m(\mathbf{x}, \mathbf{u})$  can be regarded as a density model for  $p(\mathbf{x}, \mathbf{u})$ :

$$p_m(\mathbf{x}, \mathbf{u}) := \frac{p(\mathbf{x})p(\mathbf{u})e^{r(\mathbf{x}, \mathbf{u})}}{E_{\mathbf{x}\mathbf{u}}[e^{r(\mathbf{X}, \mathbf{U})}]},$$

where the denominator ensures that  $p_m(\mathbf{x}, \mathbf{u})$  is normalized to be one. Then, the KL-divergence between  $p(\mathbf{x}, \mathbf{u})$  and  $p_m(\mathbf{x}, \mathbf{u})$  is given by

$$\begin{aligned} \text{KL}[p(\mathbf{x}, \mathbf{u}) || p_m(\mathbf{x}, \mathbf{u})] &= \iint p(\mathbf{x}, \mathbf{u}) \log \frac{p(\mathbf{x}, \mathbf{u})}{p_m(\mathbf{x}, \mathbf{u})} d\mathbf{x}d\mathbf{u} \\ &= I(\mathcal{X}, \mathcal{U}) - \underbrace{E_{\mathbf{x}\mathbf{u}}[r(\mathbf{X}, \mathbf{U})] + \log(E_{\mathbf{x}\mathbf{u}}[e^{r(\mathbf{X}, \mathbf{U})}])}_{=J_{\text{DV}}(r)}. \end{aligned} \quad (24)$$

Eq.(24) shows that minimizing  $\text{KL}[p(\mathbf{x}, \mathbf{u})||p_m(\mathbf{x}, \mathbf{u})]$  is equal to minimizing  $J_{\text{DV}}(r)$  because the mutual information  $I(\mathcal{X}, \mathcal{U})$  on the right-hand side is a constant with respect to  $r$ . As shown in Kanamori et al. (2010), the density ratio estimator based on the KL-divergence could be more accurate than the estimator based on the cross entropy  $J_{\text{LR}}$  (i.e., logistic regression) under the misspecified setting where the true density ratio is not necessarily included in the function class of a model  $r$ . In fact, we experimentally demonstrate that the nonlinear ICA method based on  $J_{\text{DV}}$  performs better than an existing method based on  $J_{\text{LR}}$  when the number of data samples is small.

In practice, we empirically approximate  $J_{\text{DV}}(r)$  as

$$\hat{J}_{\text{DV}}(r) := -\frac{1}{T} \sum_{t=1}^T r(\mathbf{x}(t), \mathbf{u}(t)) + \log \left( \frac{1}{T} \sum_{t=1}^T e^{r(\mathbf{x}(t), \mathbf{u}_p(t))} \right),$$

or

$$\tilde{J}_{\text{DV}}(r) := -\frac{1}{T} \sum_{t=1}^T r(\mathbf{x}(t), \mathbf{u}(t)) + \log \left( \frac{1}{T^2} \sum_{t=1}^T \sum_{t'=1}^T e^{r(\mathbf{x}(t), \mathbf{u}(t'))} \right).$$

Section 5 uses both approximations in numerical experiments with a mini-batch stochastic gradient method.

### 4.3 Theoretical analysis for outlier-robustness

Here, we investigate outlier-robustness of estimators based on  $J_\gamma$  and  $J_{\text{DV}}$ . To this end, we consider the following two scenarios of contamination by outliers:

- *Contamination model 1:* Given input data  $\mathbf{x}$ , complementary data  $\mathbf{u}$  is conditionally contaminated by outliers as follows:

$$\bar{p}(\mathbf{u}|\mathbf{x}) := (1 - \epsilon)p(\mathbf{u}|\mathbf{x}) + \epsilon\delta(\mathbf{u}|\mathbf{x}), \quad (25)$$

where  $\bar{p}(\mathbf{u}|\mathbf{x})$  is the contaminated conditional density of  $\mathbf{u}$  given  $\mathbf{x}$ ,  $\delta(\mathbf{u}|\mathbf{x})$  denotes the conditional density for outliers, and  $0 \leq \epsilon < 1$  denotes the contamination ratio. By assuming that input data  $\mathbf{x}$  is *noncontaminated*, (25) leads to the following contaminated joint and marginal densities:

$$\begin{aligned} \bar{p}(\mathbf{x}, \mathbf{u}) &:= (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta(\mathbf{u}|\mathbf{x})p(\mathbf{x}), \\ \bar{p}(\mathbf{x}) &:= p(\mathbf{x}) \quad \text{and} \quad \bar{p}(\mathbf{u}) := (1 - \epsilon)p(\mathbf{u}) + \epsilon \int \delta(\mathbf{u}|\mathbf{x})p(\mathbf{x})d\mathbf{x}, \end{aligned}$$

where  $\bar{p}(\mathbf{x}, \mathbf{u})$  is the contaminated joint density, and  $\bar{p}(\mathbf{x})$  and  $\bar{p}(\mathbf{u})$  denote the contaminated marginal densities of  $\mathbf{x}$  and  $\mathbf{u}$ , respectively.

- *Contamination model 2:* Input data  $\mathbf{x}$  and complementary data  $\mathbf{u}$  are jointly contaminated by outliers:

$$\bar{p}(\mathbf{x}, \mathbf{u}) := (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta(\mathbf{x}, \mathbf{u}), \quad (26)$$

Table 5: Comparison of contamination models in influence function analysis (Propositions 5 and 6).  $\delta_{\bar{x}}(\mathbf{x})$  and  $\delta_{\bar{u}}(\mathbf{u})$  denote the Delta functions having point masses at  $\bar{x}$  and  $\bar{u}$ , respectively.

| Contamination model 1   | Contamination model 2  |
|---|--|
| $\bar{p}(\mathbf{x}, \mathbf{u}) = (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta_{\bar{u}}(\mathbf{u})p(\mathbf{x})$ | $\bar{p}(\mathbf{x}, \mathbf{u}) = (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta_{\bar{x}}(\mathbf{x})\delta_{\bar{u}}(\mathbf{u})$ |
| $\bar{p}(\mathbf{x}) = p(\mathbf{x})$   | $\bar{p}(\mathbf{x}) = (1 - \epsilon)p(\mathbf{x}) + \epsilon\delta_{\bar{x}}(\mathbf{x})$   |
| $\bar{p}(\mathbf{u}) = (1 - \epsilon)p(\mathbf{u}) + \epsilon\delta_{\bar{u}}(\mathbf{u})$                                      | $\bar{p}(\mathbf{u}) = (1 - \epsilon)p(\mathbf{u}) + \epsilon\delta_{\bar{u}}(\mathbf{u})$   |

where  $\delta(\mathbf{x}, \mathbf{u})$  denotes the joint density for outliers. Eq.(26) leads to the following contaminated marginal densities:

$$\bar{p}(\mathbf{x}) := (1 - \epsilon)p(\mathbf{x}) + \epsilon\delta(\mathbf{x}) \quad \text{and} \quad \bar{p}(\mathbf{u}) := (1 - \epsilon)p(\mathbf{u}) + \epsilon\delta(\mathbf{u}),$$

where

$$\delta(\mathbf{x}) := \int \delta(\mathbf{x}, \mathbf{u})d\mathbf{u} \quad \text{and} \quad \delta(\mathbf{u}) := \int \delta(\mathbf{x}, \mathbf{u})d\mathbf{x}.$$

The fundamental difference between these contamination scenarios is whether or not input data  $\mathbf{x}$  is contaminated by outliers. Next, we perform two analyses for outlier-robustness: One analysis is performed under the condition that the contamination ratio  $\epsilon$  is small, while another one does *not* necessarily assume that  $\epsilon$  is small and thus heavy contamination of outliers is also within the scope of the analysis.

#### 4.3.1 INFLUENCE FUNCTION ANALYSIS

Here, we perform influence function analysis (Hampel et al., 2011), which is an established tool in robust statistics and assumes that the densities for outliers are given as follows:

$$\delta(\mathbf{x}, \mathbf{u}) = \delta_{\bar{x}}(\mathbf{x})\delta_{\bar{u}}(\mathbf{u}), \quad \delta(\mathbf{x}) = \delta_{\bar{x}}(\mathbf{x}) \quad \text{and} \quad \delta(\mathbf{u}) = \delta_{\bar{u}}(\mathbf{u}), \quad (27)$$

where  $\delta_{\bar{x}}(\mathbf{x})$  and  $\delta_{\bar{u}}(\mathbf{u})$  are the Dirac delta functions having point masses at  $\bar{x}$  and  $\bar{u}$ , respectively. Eq.(27) implies that  $\delta(\mathbf{u}|\mathbf{x}) = \delta_{\bar{u}}(\mathbf{u})$  and  $\int \delta(\mathbf{u}|\mathbf{x})p(\mathbf{x})d\mathbf{x} = \delta_{\bar{u}}(\mathbf{u})$  in the contamination model 1. Table 5 is a comparison of contamination models 1 and 2 in this analysis.

In this analysis, we suppose that a model  $r_{\theta}(\mathbf{x}, \mathbf{u})$  is parametrized by  $\theta$  and has the following properties:

- $r_{\theta}(\mathbf{x}, \mathbf{u})$  and its gradient  $\mathbf{g}_{\theta}(\mathbf{x}, \mathbf{u})$  are continuous where  $\mathbf{g}_{\theta}(\mathbf{x}, \mathbf{u}) := \nabla_{\theta}r_{\theta}(\mathbf{x}, \mathbf{u})$ .
- $|r_{\theta}(\mathbf{x}, \mathbf{u})| \rightarrow \infty$  and  $\|\mathbf{g}_{\theta}(\mathbf{x}, \mathbf{u})\| \rightarrow \infty$  as  $\|\mathbf{x}\| \rightarrow \infty$  and/or  $\|\mathbf{u}\| \rightarrow \infty$ .

These properties are simply used to discuss the robustness of the estimators based on the influence function, and satisfied by standard models such neural networks with an unbounded activation function (e.g., the softplus function). Based on  $r_{\theta}(\mathbf{x}, \mathbf{u})$ , we define two estimators as

$$\theta^* := \underset{\theta}{\operatorname{argmin}} J(r_{\theta}) \quad \text{and} \quad \theta_{\epsilon} := \underset{\theta}{\operatorname{argmin}} \bar{J}(r_{\theta}),$$

where  $J$  is some objective function over the *noncontaminated* densities (e.g.,  $J_\gamma$ ), while  $\bar{J}$  is the contaminated version of  $J$  computed over the contaminated densities  $\bar{p}(\mathbf{x}, \mathbf{u})$ ,  $\bar{p}(\mathbf{x})$  and  $\bar{p}(\mathbf{u})$  instead of  $p(\mathbf{x}, \mathbf{u})$ ,  $p(\mathbf{x})$  and  $p(\mathbf{u})$ . Then, we define the *influence function* as

$$\text{IF}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) = \lim_{\epsilon \rightarrow 0} \frac{\boldsymbol{\theta}^* - \boldsymbol{\theta}_\epsilon}{\epsilon}. \quad (28)$$

Eq.(28) indicates that  $\|\text{IF}(\bar{\mathbf{x}}, \bar{\mathbf{u}})\|$  measures how the estimator  $\boldsymbol{\theta}^*$  is influenced by small contamination of outliers  $\bar{\mathbf{x}}$  and  $\bar{\mathbf{u}}$ , and a larger  $\|\text{IF}(\bar{\mathbf{x}}, \bar{\mathbf{u}})\|$  implies stronger influence by outliers. Based on the influence function, some robust properties of estimators can be characterized: Estimator  $\boldsymbol{\theta}^*$  is called *B-robust* if  $\sup_{\bar{\mathbf{x}}, \bar{\mathbf{u}}} \|\text{IF}(\bar{\mathbf{x}}, \bar{\mathbf{u}})\| < \infty$  (Hampel et al., 2011). B-robustness guarantees that the influence from outliers is limited. A more favorable property is the *redescending property* (Hampel et al., 2011), which defined as  $\lim_{\|\bar{\mathbf{x}}\|, \|\bar{\mathbf{u}}\| \rightarrow \infty} \|\text{IF}(\bar{\mathbf{x}}, \bar{\mathbf{u}})\| = 0$ . The redescending property ensures that  $\boldsymbol{\theta}^*$  has almost no influence from even strongly deviated data  $\bar{\mathbf{x}}$  and/or  $\bar{\mathbf{u}}$ .

First, the following proposition shows the forms of the influence functions for  $J_{\text{DV}}$ :

**Proposition 5.** *Assume that  $r_\theta(\mathbf{x}, \mathbf{u}) = \log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$  at  $\boldsymbol{\theta} = \boldsymbol{\theta}^*$ . For the contamination model 1, the influence function of the estimator  $\boldsymbol{\theta}^*$  based on  $J_{\text{DV}}$  is given by*

$$\text{IF}_{\text{DV}}(\bar{\mathbf{u}}) = \mathbf{V}_{\text{DV}}^{-1} E_{\mathbf{x}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}}) - e^{r_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}})], \quad (29)$$

where  $\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{x}, \mathbf{u}) = \nabla_{\boldsymbol{\theta}} r_\theta(\mathbf{x}, \mathbf{u})|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*}$  and

$$\mathbf{V}_{\text{DV}} := -E_{\mathbf{xu}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})^\top] + E_{\mathbf{xu}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})]E_{\mathbf{xu}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})]^\top. \quad (30)$$

Regarding the contamination model 2, the influence function is given by

$$\begin{aligned} \text{IF}_{\text{DV}}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) = \mathbf{V}_{\text{DV}}^{-1} & \left[ \mathbf{g}_{\boldsymbol{\theta}^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) + E_{\mathbf{xu}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})] \right. \\ & \left. - E_{\mathbf{x}}[e^{r_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}})] - E_{\mathbf{u}}[e^{r_{\boldsymbol{\theta}^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\boldsymbol{\theta}^*}(\bar{\mathbf{x}}, \mathbf{U})] \right], \end{aligned} \quad (31)$$

where  $E_{\mathbf{x}}$  and  $E_{\mathbf{u}}$  denote the expectations over  $p(\mathbf{x})$  and  $p(\mathbf{u})$ , respectively.

The proof is deferred in Appendix G. The key to interpret the influence functions in Proposition 5 is unboundedness of  $r_{\boldsymbol{\theta}^*}(\mathbf{x}, \mathbf{u})$  and its gradient  $\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{x}, \mathbf{u})$ . For example, the right-hand side of (29) indicates that  $\|\text{IF}_{\text{DV}}(\bar{\mathbf{u}})\|$  can be unbounded when  $r_{\boldsymbol{\theta}^*}(\mathbf{x}, \bar{\mathbf{u}}) \rightarrow \infty$  and  $\|\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{x}, \bar{\mathbf{u}})\| \rightarrow \infty$  as  $\|\bar{\mathbf{u}}\| \rightarrow \infty$ . This means that the estimator based on  $J_{\text{DV}}$  is not B-robust, and can be strongly influenced by outliers  $\bar{\mathbf{u}}$ . The same discussion is applicable to the influence function (31) of the contamination model 2. Thus, these results imply that estimation based on  $J_{\text{DV}}$  can be sensitive to outliers.

Next, we show the influence functions for the  $\gamma$ -cross entropy  $J_\gamma$  in the following proposition:

**Proposition 6.** *Suppose the same assumption holds as Proposition 5. Regarding the contamination model 1, the influence function of the estimator  $\boldsymbol{\theta}^*$  based on  $J_\gamma$  is obtained as follows:*

$$\text{IF}_\gamma(\bar{\mathbf{u}}) = V_\gamma^{-1} E_{\mathbf{x}}[\{S(\mathbf{X}, \bar{\mathbf{u}})^{\frac{1}{1+\gamma}} - (1 - S(\mathbf{X}, \bar{\mathbf{u}}))^{\frac{1}{1+\gamma}}\} \eta(\mathbf{X}, \bar{\mathbf{u}})^{\frac{\gamma}{1+\gamma}} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \bar{\mathbf{u}})], \quad (32)$$

where

$$V_\gamma := E_{\mathbf{xu}}[\eta(\mathbf{X}, \mathbf{U})^{\frac{\gamma}{1+\gamma}} S(\mathbf{X}, \mathbf{U})^{\frac{1}{1+\gamma}} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U}) \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})^\top], \quad (33)$$

$$S(\mathbf{x}, \mathbf{u}) := \frac{1}{1 + e^{(1+\gamma)r_{\theta^*}(\mathbf{x}, \mathbf{u})}} \quad \text{and} \quad \eta(\mathbf{x}, \mathbf{u}) := S(\mathbf{x}, \mathbf{u})(1 - S(\mathbf{x}, \mathbf{u})). \quad (34)$$

Regarding the contamination model 2, the influence function is given by

$$\begin{aligned} \text{IF}_\gamma(\bar{\mathbf{x}}, \bar{\mathbf{u}}) &= V_\gamma^{-1} \left[ \eta(\bar{\mathbf{x}}, \bar{\mathbf{u}})^{\frac{\gamma}{1+\gamma}} S(\bar{\mathbf{x}}, \bar{\mathbf{u}})^{\frac{1}{1+\gamma}} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) \right. \\ &\quad - E_{\mathbf{x}}[\eta(\mathbf{X}, \bar{\mathbf{u}})^{\frac{\gamma}{1+\gamma}} (1 - S(\mathbf{X}, \bar{\mathbf{u}}))^{\frac{1}{1+\gamma}} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] \\ &\quad - E_{\mathbf{u}}[\eta(\bar{\mathbf{x}}, \mathbf{U})^{\frac{\gamma}{1+\gamma}} (1 - S(\bar{\mathbf{x}}, \mathbf{U}))^{\frac{1}{1+\gamma}} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})] \\ &\quad \left. + E_{\mathbf{x} \times \mathbf{u}}[\eta(\mathbf{X}, \mathbf{U})^{\frac{\gamma}{1+\gamma}} (1 - S(\mathbf{X}, \mathbf{U}))^{\frac{1}{1+\gamma}} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] \right]. \quad (35) \end{aligned}$$

The proof of Proposition 6 is not given because it is almost the same as Proposition 5. In contrast to  $J_{\text{DV}}$ , the influence functions based on the  $\gamma$ -cross entropy  $J_\gamma$  include the weight functions  $\eta(\mathbf{x}, \mathbf{u})$  and  $S(\mathbf{x}, \mathbf{u})$ , and  $\|\text{IF}_\gamma(\bar{\mathbf{u}})\|$  and  $\|\text{IF}_\gamma(\bar{\mathbf{x}}, \bar{\mathbf{u}})\|$  can be bounded even when  $\|\mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}})\| \rightarrow \infty$  as  $\|\bar{\mathbf{x}}\| \rightarrow \infty$  and/or  $\|\bar{\mathbf{u}}\| \rightarrow \infty$ . The following corollary characterizes the robustness of the  $\gamma$ -cross entropy and shows that in stark contrast with  $J_{\text{DV}}$ , the estimator based on  $J_\gamma$  is B-robust for both contamination models, and has the redescending property for the contamination model 1 under some conditions:

**Corollary 7.** *Assume that*

$$\sup_{\mathbf{x}, \mathbf{u}} \|\eta(\mathbf{x}, \mathbf{u})^{\frac{\gamma}{1+\gamma}} \mathbf{g}_{\theta^*}(\mathbf{x}, \mathbf{u})\| < \infty. \quad (36)$$

*Then, the estimator  $\theta^*$  based on  $J_\gamma$  is B-robust both for the contamination models 1 and 2. Another assumption is made as follows:*

$$\lim_{\|\bar{\mathbf{u}}\| \rightarrow \infty} \|\eta(\mathbf{x}, \bar{\mathbf{u}})^{\frac{\gamma}{1+\gamma}} \mathbf{g}_{\theta^*}(\mathbf{x}, \bar{\mathbf{u}})\| = 0 \quad \text{for all } \mathbf{x}. \quad (37)$$

*Then, the estimator  $\theta^*$  based on  $J_\gamma$  has the redescending property for the contamination model 1.*

The results are an almost direct consequence of Proposition 6, and thus the detailed calculation is omitted. Assumptions (36) and (37) would be mild: By definition (34), the function  $\eta(\mathbf{x}, \mathbf{u})$  exponentially converges to zero when  $r_{\theta^*}(\mathbf{x}, \mathbf{u}) \rightarrow \infty$  as  $\|\mathbf{x}\| \rightarrow \infty$  and/or  $\|\mathbf{u}\| \rightarrow \infty$ . Thus, B-robustness would hold under both contamination models for standard models (e.g., neural networks). On the other hand, the redescending property is shown only for the contamination model 1. This comes from the fact that the contamination of outliers is more complicated in the contamination model 2: Both  $\mathbf{x}$  and  $\mathbf{u}$  are contaminated by outliers in the contamination model 2, while  $\mathbf{u}$  is only contaminated in the contamination model 1. However, we numerically demonstrate that the proposed method based on the  $\gamma$ -cross entropy is robust against outliers even for the contamination model 2.

Proposition 5 also reveals the outlier weakness of the cross entropy  $J_{\text{LR}}$  used in logistic regression. We recall that the  $\gamma$ -cross entropy  $J_\gamma$  approaches  $J_{\text{LR}}$  as  $\gamma \rightarrow 0$ . Then, it is obvious that Assumptions (36) and (37) in Corollary 7 are never satisfied in the limit of  $\gamma \rightarrow 0$  when  $\|\mathbf{g}_{\theta^*}(\mathbf{x}, \mathbf{u})\|$

is an unbounded function (e.g., neural networks). This implies that the estimator based on  $J_{LR}$  is neither B-robust nor redescending, and estimation based on  $J_{LR}$  can be hampered by outliers.

In short, the analysis based on the influence function implies that it is a promising approach to use the  $\gamma$ -cross entropy for density ratio estimation in the presence of outliers even with neural networks, while estimation based on  $J_{DV}$  and  $J_{LR}$  can be strongly influenced by outliers particularly when  $r_{\theta}(\mathbf{x}, \mathbf{u})$  is modelled by using unbounded functions such as neural networks.

#### 4.3.2 STRONG ROBUSTNESS OF THE $\gamma$ -CROSS ENTROPY

Here, we investigate the robustness of the  $\gamma$ -cross entropy even when the contamination ratio  $\epsilon$  is *not* necessarily assumed to be small. Unlike influence function analysis, the contaminated densities,  $\delta(\mathbf{x}, \mathbf{u})$ ,  $\delta(\mathbf{x})$  and  $\delta(\mathbf{u})$ , are not Dirac delta functions, but rather general probability density functions. To clarify the notations, we denote the  $\gamma$ -cross entropy over the contaminated densities as

$$\bar{J}_{\gamma}(r) := -\frac{1}{\gamma} \log \left[ \bar{E}_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}}{1 + e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}} \right)^{\frac{\gamma}{\gamma+1}} \right] + \bar{E}_{\mathbf{x} \times \mathbf{u}} \left[ \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{X}, \mathbf{U})}} \right)^{\frac{\gamma}{\gamma+1}} \right] \right],$$

where  $\bar{E}_{\mathbf{x}\mathbf{u}}$  and  $\bar{E}_{\mathbf{x} \times \mathbf{u}}$  denote the expectations over  $\bar{p}(\mathbf{x}, \mathbf{u})$  and  $\bar{p}(\mathbf{x})\bar{p}(\mathbf{u})$  in contamination models, respectively.

The remarkable property of the  $\gamma$ -cross entropy is *strong robustness* (Fujisawa and Eguchi, 2008): The latent bias caused from outliers can be small even in the case of heavy contamination (i.e., nonsmall  $\epsilon$ ). To prove the strong robustness in the context of density ratio estimation, we recall that  $r(\mathbf{x}, \mathbf{u})$  is defined by

$$r(\mathbf{x}, \mathbf{u}) = \psi(\mathbf{h}_{\mathbf{x}}(\mathbf{x}), \mathbf{h}_{\mathbf{u}}(\mathbf{u})) + a(\mathbf{h}_{\mathbf{x}}(\mathbf{x})) + b(\mathbf{h}_{\mathbf{u}}(\mathbf{u})).$$

The following proposition proved in Appendix H focuses on the contamination model 1 and establishes certain conditions for strong robustness:

**Proposition 8.** *Let us define a constant  $\nu_1$  as*

$$\nu_1 := \iint p(\mathbf{x})\delta(\mathbf{u}|\mathbf{x}) \left( \frac{e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} + \iint p(\mathbf{x})\delta(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u}.$$

*We suppose the contamination model 1 and make the following assumptions:*

(D1) *Assume that*

$$\iint p(\mathbf{x})\delta(\mathbf{u}|\mathbf{x}) \left( \frac{e^{(\gamma+1)\{\psi(\mathbf{h}_{\mathbf{x}}(\mathbf{x}), \mathbf{h}_{\mathbf{u}}(\mathbf{u})) + a(\mathbf{h}_{\mathbf{x}}(\mathbf{x}))\}/2}}{e^{-(\gamma+1)b(\mathbf{h}_{\mathbf{u}}(\mathbf{u}))} + e^{(\gamma+1)\{\psi(\mathbf{h}_{\mathbf{x}}(\mathbf{x}), \mathbf{h}_{\mathbf{u}}(\mathbf{u})) + a(\mathbf{h}_{\mathbf{x}}(\mathbf{x}))\}}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{u}d\mathbf{x} < \infty$$

and

$$\iint p(\mathbf{x})\delta(\mathbf{u}) \left( \frac{e^{-(\gamma+1)b(\mathbf{h}_{\mathbf{u}}(\mathbf{u}))/2}}{e^{-(\gamma+1)b(\mathbf{h}_{\mathbf{u}}(\mathbf{u}))} + e^{(\gamma+1)\{\psi(\mathbf{h}_{\mathbf{x}}(\mathbf{x}), \mathbf{h}_{\mathbf{u}}(\mathbf{u})) + a(\mathbf{h}_{\mathbf{x}}(\mathbf{x}))\}}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{u}d\mathbf{x} < \infty.$$

(D2) *The following integrals are sufficiently small:*

$$\iint p(\mathbf{x})\delta(\mathbf{u}|\mathbf{x}) e^{\gamma\{\psi(\mathbf{h}_{\mathbf{x}}(\mathbf{x}), \mathbf{h}_{\mathbf{u}}(\mathbf{u})) + a(\mathbf{h}_{\mathbf{x}}(\mathbf{x}))\}} d\mathbf{u}d\mathbf{x} \text{ and } \int \delta(\mathbf{u}) e^{-\gamma b(\mathbf{h}_{\mathbf{u}}(\mathbf{u}))} d\mathbf{u}.$$

Then, it holds that  $\nu_1$  is sufficient small and

$$\bar{J}_\gamma(r) = J_\gamma(r) - \frac{1}{\gamma} \log(1 - \epsilon) + O\left(\frac{\epsilon}{1 - \epsilon} \nu_1\right). \quad (38)$$

The proof can be found in Appendix H. Eq.(38) in Proposition 8 indicates that for the contamination model 1, minimization of  $\bar{J}_\gamma(r)$  is approximately equal to minimization of  $J_\gamma(r)$  over *noncontaminated* densities under Assumptions (D1-2) even when the contamination ratio  $\epsilon$  is not necessarily small (e.g.,  $\epsilon = 0.2$  as in the numerical experiments of Fujisawa and Eguchi (2008, Section 8)). This means that we could perform density ratio estimation in many practical situations almost as if outliers did not exist.

Assumption (D1) implies that  $\delta(\mathbf{u}|\mathbf{x})$  and  $\delta(\mathbf{u})$  are localized on the domain of  $\mathbf{u}$  and have a fast decay on their tails as often seen in densities for outliers: For example, when both  $\delta(\mathbf{u})$  and  $\delta(\mathbf{u}|\mathbf{x})$  have finite supports<sup>4</sup>, Assumption (D1) is fulfilled. Assumption (D2) means that outliers drawn from  $\delta(\mathbf{u}|\mathbf{x})$  and  $\delta(\mathbf{u})$  exist on the respective tails of  $e^{\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x}))}$  and  $e^{-b(\mathbf{h}_u(\mathbf{u}))}$ , which can be roughly regarded as models of  $p(\mathbf{u}|\mathbf{x})$  and  $p(\mathbf{u})$  respectively because  $r(\mathbf{x}, \mathbf{u})$  is modelling  $\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \log p(\mathbf{u}|\mathbf{x}) - \log p(\mathbf{u})$ . Thus, when  $r(\mathbf{x}, \mathbf{u})$  is near to  $\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$ , Assumption (D2) implies that the outliers exist on the tails of  $p(\mathbf{x}, \mathbf{u})$  and  $p(\mathbf{u})$ , and indeed reflects typical contamination by outliers. Overall, Assumptions (D1-2) would be practically reasonable.

For the contamination model 2, the strong robustness does not hold, but the  $\gamma$ -cross entropy still has a favorable property which we call *semistrong robustness*. The detail of the semistrong robustness is given in Appendix I. In addition, we experimentally demonstrate that the proposed method based on the  $\gamma$ -cross entropy is robust against outliers under the contamination model 2.

## 5. Numerical Experiments

This section numerically investigates how the proposed methods work in nonlinear ICA and a downstream task for linear classification, and empirically supports the implication of the theoretical analysis for nonlinear ICA in Section 3.2.

### 5.1 Nonlinear ICA on artificial data

Here, we demonstrate the numerical performance of the proposed methods in nonlinear ICA and investigate how the dimensionality of complementary data affects the source recovery as implied in Section 3.2. All experiments suppose  $d_x = D_x$  according to Proposition 2 and Theorem 3.

#### 5.1.1 SAMPLE EFFICIENCY AND OUTLIER ROBUSTNESS

We followed the experimental setting of a nonlinear ICA method called *permutation contrastive learning* (PCL) in Hyvärinen and Morioka (2017), which is intended for temporally dependent data  $\mathbf{x}(t)$ ,  $t = 1, \dots, T$  and supposes the complementary data samples are  $\mathbf{u}(t) = \mathbf{x}(t - 1)$  (i.e., past input data). First, as in the autoregressive process, the ten-dimensional temporally dependent  $T$  sources were generated from the contaminated density model as

$$\bar{p}(\mathbf{s}(t)|\mathbf{s}(t - 1)) = (1 - \epsilon)p(\mathbf{s}(t)|\mathbf{s}(t - 1)) + \epsilon\delta(\mathbf{s}(t)|\mathbf{s}(t - 1)), \quad (39)$$

---

4. The support of  $\delta(\mathbf{u})$  is defined by  $\{\mathbf{u} \mid \delta(\mathbf{u}) \neq 0\}$ .

where  $\epsilon$  denotes the outlier ratio, and with  $\rho = 0.7$  and  $D_x = 10$ ,

$$p(\mathbf{s}(t)|\mathbf{s}(t-1)) \propto \prod_{i=1}^{D_x} \exp\left(-\sqrt{2}|s_i(t) - \rho s_i(t-1)|/(1 - \rho^2)\right) \quad (40)$$

$$\delta(\mathbf{s}(t)|\mathbf{s}(t-1)) \propto \prod_{i=1}^{D_x} \exp\left(-\frac{\{s_i(t) + \rho s_i(t-1)\}^2}{2}\right). \quad (41)$$

Then, input data  $\mathbf{x}$  was generated according to  $\mathbf{x} = \mathbf{f}(\mathbf{s})$  in (8) where  $\mathbf{f}$  was modelled by a multi-layer feedforward neural network with the leaky ReLU activation function and random weights. The numbers of all hidden and output units were the same as the dimensionality of data (i.e.,  $D_x$ ). Since complementary data  $\mathbf{u}(t)$  is past input data  $\mathbf{x}(t-1)$ , the contamination process corresponds to the contamination model 2 in Section 4.3 where both  $\mathbf{x}(t)$  and  $\mathbf{u}(t)$  are contaminated by outliers. As preprocessing, we performed whitening based on the  $\gamma$ -cross entropy (Chen et al., 2013).

We applied the following methods to data samples:

- Permutation contrastive learning (PCL) (Hyvärinen and Morioka, 2017): A nonlinear ICA method for temporally dependent sources based on logistic regression whose objective function is an empirical version of  $J_{LR}$ .
- DV-PCL: A proposed nonlinear ICA method for temporally dependent sources based on the Donsker-Varadhan variational estimation  $\hat{J}_{DV}$  (Section 4.2).
- Robust PCL (RPCL): A proposed nonlinear ICA method for temporally dependent sources based on the  $\gamma$ -cross entropy  $\hat{J}_\gamma$  (Section 4.1). The value of  $\gamma$  was increased from 0.0 to 5.0 during minibatch stochastic gradient at every 100 epoch where  $\gamma = 0.0$  means to perform the standard logistic regression.

We used a same model  $r(\mathbf{x}, \mathbf{u})$  for all three methods: The representation function  $\mathbf{h}_x$  was modelled by a multi-layer feedforward neural network where the number of hidden units was  $4D_x$ , but the final layer was  $D_x$ . The number of layers was the same as  $\mathbf{f}(\mathbf{s})$  in the data generative model (8). The activation functions in hidden layers were the max-out function (Goodfellow et al., 2013) with two groups, while the final layer has no activation function. Regarding the representation function  $\mathbf{h}_u$ , the same architecture as  $\mathbf{h}_x$  was employed with parameter sharing (i.e.,  $\mathbf{h}_u = \mathbf{h}_x$ ). Following Hyvärinen and Morioka (2017),  $r(\mathbf{x}, \mathbf{u})$  was modelled by

$$r(\mathbf{x}, \mathbf{u}) = \sum_{i=1}^{D_x} |a_{i,1}h_{x,i}(\mathbf{x}) + a_{i,2}h_{u,i}(\mathbf{u}) + b_i| - (\bar{a}_i h_{x,i}(\mathbf{x}) + \bar{b}_i)^2 + c, \quad (42)$$

where  $a_{i,1}, a_{i,2}, b_i, \bar{a}_i, \bar{b}_i, c$  are parameters to be estimated from data. All parameters were estimated by the Adam optimizer for 1,600 epochs with mini-batch size 256 and learning rate 0.001. We also applied the  $\ell_2$  regularization to the weight parameters with the regularization parameter  $10^{-4}$ .

With the test sources  $s_i^{\text{te}}(t)$  without outliers drawn from  $p(\mathbf{s}(t)|\mathbf{s}(t-1))$ , the performance was measured by the mean absolute correlation as

$$\frac{1}{T_{\text{te}} D_x} \sum_{i=1}^{D_x} \left| \sum_{t=1}^{T_{\text{te}}} (\hat{h}_{x,i}(\mathbf{x}^{\text{te}}(t)) - \mu_i^{\hat{h}})(s_{\pi(i)}^{\text{te}}(t) - \mu_{\pi(i)}^{\text{te}}) \right|,$$

where  $\hat{h}_{x,i}$  denotes the learned representation function,  $T_{\text{te}} = 50,000$  denotes the number of test samples, and  $\mu_i^{\hat{h}}$  and  $\mu_i^{\text{te}}$  are the sample means of  $\hat{h}_{x,i}(\mathbf{x}^{\text{te}}(t))$  and  $s_i^{\text{te}}(t)$ , respectively. Since estimation of nonlinear ICA is indeterminate with respect to permutation (i.e., ordering) of the source components  $s_i$ , the permutation indices  $\pi(i) \in \{1, 2, \dots, D_x\}$  were determined so that the absolute mean correlation was maximized.

Fig.1(a-c) investigate the sample efficiency of each method when  $\epsilon = 0$  (i.e., no outliers) and the numbers of layers are one, two and three, respectively. When the number of samples is large, all methods well-recover the source components and work similarly. However, when the number of samples gets smaller, DV-PCL tends to perform better than PCL and RPCL. This presumably because the objective function  $J_{\text{DV}}$  in DV-PCL is based on the KL-divergence, and the density ratio estimator based on the KL-divergence would be accurate as implied in Kanamori et al. (2010). The same tendency can be seen when the numbers of layers are four (Fig.1(d)) and five (Fig.1(e)).

Fig.2 shows outlier-robustness of each nonlinear ICA method in  $T = 100,000$ . When the number of layers is one, all methods perform well (Fig.2(a)). This would be because we perform an outlier-robust whitening based on the  $\gamma$ -cross entropy (Chen et al., 2013) in preprocessing. However, when the number of layer is larger than one, the performance of DV-PCL and PCL quickly decreases as the contamination ratio  $\epsilon$  is increased (Fig.2(b-e)). On the other hand, RPCL keeps high correlation values on a wide range of  $\epsilon$ , and thus is very robust against outliers. These results are consistent with the theoretical results in Section 4.3 that estimators based on  $J_{\text{DV}}$  (DV-PCL) and  $J_{\text{LR}}$  (PCL) can be seriously hampered by outliers, while the  $\gamma$ -cross entropy (RPCL) is promising in the presence of outliers.

Overall, the proposed methods, DV-PCL and RPCL, have respective advantages over PCL. DV-PCL seems to be more useful in the limited number of samples, while RPCL is very robust against outliers particularly when the contamination ratio  $\epsilon$  is large.

### 5.1.2 IMPORTANCE OF THE DIMENSIONALITY OF COMPLEMENTARY DATA

Next, we investigate how the dimensionality of complementary data  $\mathbf{u}$  affects the source recovery in nonlinear ICA as implied in Proposition 2 and Theorem 3.

We followed the recovery conditions in Theorem 3. In order for the conditional density  $p(\mathbf{s}|\mathbf{u})$  to be differentiable, the conditionally independent sources were first generated from

$$p(\mathbf{s}|\mathbf{u}) \propto \prod_{i=1}^{D_x} \exp\left(-\log(\cosh(s_i - \mathbf{w}_i^{\text{s}\top} \mathbf{u}))\right), \quad (43)$$

where  $\mathbf{w}_i^{\text{s}}$ ,  $i = 1 \dots, D_x$  are  $D_u$ -dimensional vectors randomly determined from the independent uniform density on  $[-1, 1]^{D_u}$ . The function  $\log(\cosh(\cdot))$  is a smooth approximation of the absolute function  $|\cdot|$ , and thus  $p(\mathbf{s}|\mathbf{u})$  can be regarded as a smoothed version of the Laplace density (40) in Section 5.1.1. This smooth approximation has been previously used in linear ICA as well (Hyvärinen, 1999). In this experiment, complementary data samples  $\mathbf{u}(t)$  were simply drawn from the independent uniform density on  $[0, 1]^{D_u}$ . The total number of samples is  $T = 100,000$ . Input data  $\mathbf{x}$  was generated according to (8) where the mixing function  $\mathbf{f}$  is modelled by a feedforward neural network with random connections.

Here, we used the nonlinear ICA methods based on the  $\gamma$ -cross entropy  $\hat{J}_\gamma$  (Section 4.1) and Donsker-Varadhan variational estimation  $\hat{J}_{\text{DV}}$  (Section 4.2). As in Section 5.1.1, the representation function  $\mathbf{h}_x$  was modelled by a feedforward neural network where the number of hidden units

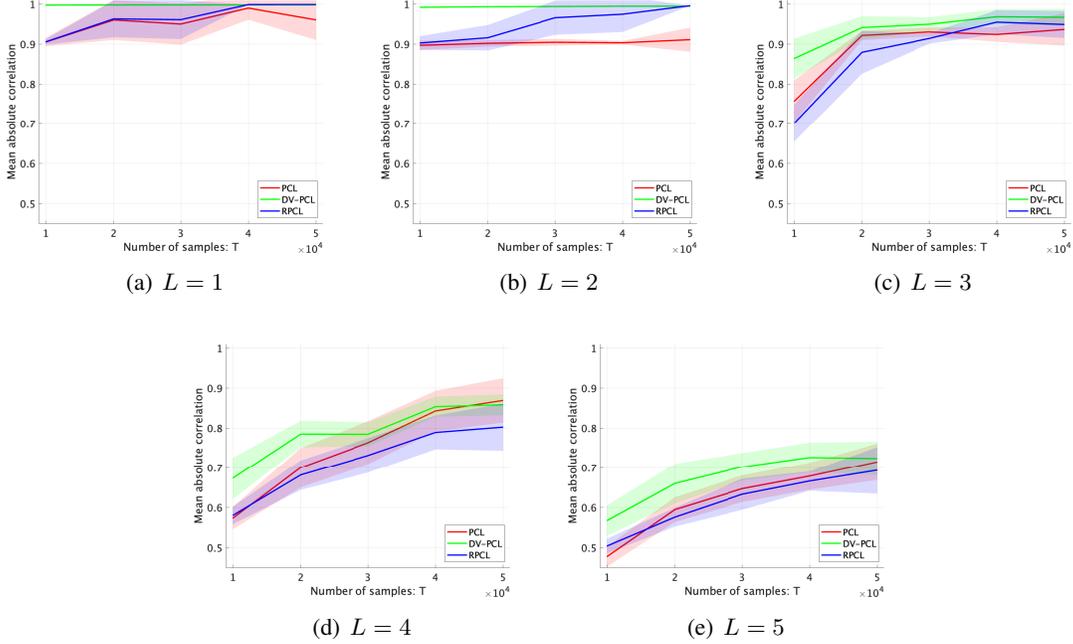


Figure 1: Sample efficiency of nonlinear ICA methods in the case of no outliers (i.e.,  $\epsilon = 0$ ). The averages of the mean absolute correlation were computed over 10 runs.  $L$  denotes the number of layers both in the nonlinear mixing function  $\mathbf{f}$  in (8) and representation function  $\mathbf{h}_x$ .

was  $4D_x$ , but the final layer was  $D_x$ . The number of layers was the same as  $\mathbf{f}$  in the data generative model.  $\mathbf{h}_u(\mathbf{u})$  was modelled by a one-layer neural network without the activation function as  $\mathbf{h}_u(\mathbf{u}) = \mathbf{W}_u \mathbf{u} + \mathbf{b}_u$  where  $\mathbf{W}_u \in \mathbb{R}^{D_x \times D_u}$  and  $\mathbf{b}_u \in \mathbb{R}^{D_x}$ . Since  $p(\mathbf{s}|\mathbf{u})$  in (43) is a smoother density,  $r(\mathbf{x}, \mathbf{u})$  was also modelled by a smoother function than (42) in Section 5.1.1 as follows:

$$r(\mathbf{x}, \mathbf{u}) = \sum_{i=1}^{D_x} \log(\cosh(a_{i,1} h_{x,i}(\mathbf{x}) + a_{i,2} h_{u,i}(\mathbf{u}) + b_i)) - (\bar{a}_i h_{x,i}(\mathbf{x}) + \bar{b}_i)^2 + c.$$

All parameters were optimized by Adam for 1,500 epochs with mini-batch size 256 and learning rate 0.001. The  $\ell_2$  regularization was also applied as done in Section 5.1.1. For the  $\gamma$ -cross entropy, the value of  $\gamma$  was increased from 0.1 to 3.0 during minibatch stochastic gradient at every 100 epoch. The performance was evaluated by the mean absolute correlation as in Section 5.1.1.

The two plots in the first row of Fig.3 are results for Donsker-Varadhan variational estimation, and clearly show that the performance for source recovery depends on the dimensionality of complementary data. When  $D_u < D_x$ , the mean correlation is small in all layers. However, when  $D_u$  is larger than or equal to  $D_x$ , the mean correlation gets significantly larger. This is consistent with implication of Theorem 3: In order to recover the source components, the dimensionality of complementary data is larger than or equal to input data (i.e.,  $D_x \leq D_u$  in Assumption (B'3)), which has not been revealed in previous work of nonlinear ICA. These empirical results clearly support our theoretical implications, and suggest to use fairly high-dimensional complementary data in

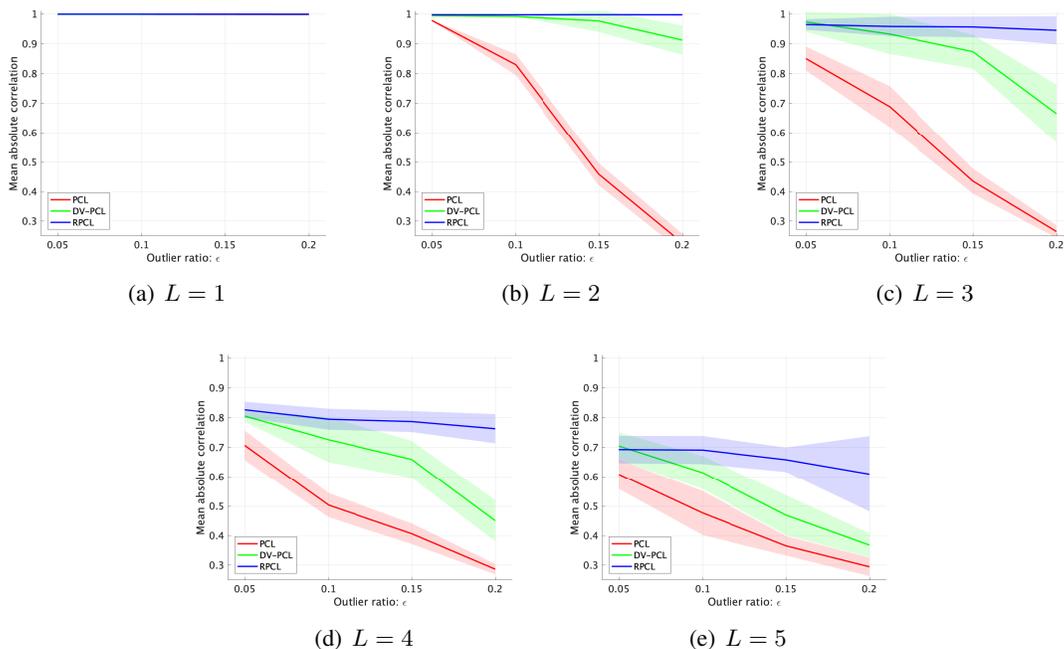


Figure 2: Outlier robustness of nonlinear ICA methods. The averages of the mean absolute correlation were computed over 10 runs.  $L$  denotes the number of layers both in the nonlinear mixing function  $f$  in (8) and representation function  $h_x$ .

practice. Similar results were observed for the  $\gamma$ -cross entropy as well (the second row of Fig.3). Furthermore, especially for Donsker-Varadhan variational estimation, another interesting point is that higher-dimensional complementary data often decreases the variance of the mean absolute correlation, and this implies that higher-dimensional complementary data takes a role of stabilizing estimation as well.

## 5.2 Evaluation on downstream linear classification

We finally demonstrate how the proposed method based on the  $\gamma$ -cross entropy works on benchmark datasets as done in the context of maximization of mutual information (Tschannen et al., 2019), and implicitly investigate Theorem 4 as well because the representation function  $h_x(x)$  outputs a lower-dimensional feature than input data. In order to evaluate methods for unsupervised representation learning, a number of protocols have been previously proposed: Multi-scale structural similarity (Wang et al., 2003), mutual information estimation and see more protocols in Hjelm et al. (2019). Here, we employ a linear classification protocol (van den Oord et al., 2018; Tian et al., 2019; Tschannen et al., 2019), which consists of two steps: First, a representation function  $h_x$  is learned with unlabelled data. Second, the learned representation function is fixed (i.e., not learned anymore), and the feature computed through the representation function is tested on a downstream linear classification task using labelled data. The classification accuracy is the measure for goodness of the representation function.

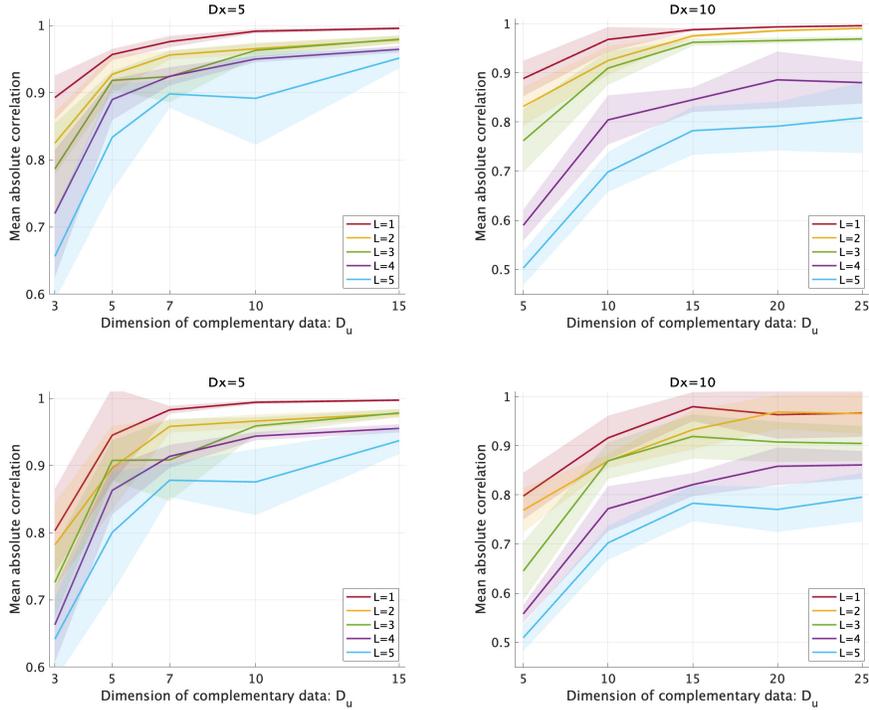


Figure 3: The averages of the mean absolute correlation against the dimensionality of complementary data over 10 runs. The left and right plots in the first (second) row are  $D_x = 5$  and  $D_x = 10$  for Donsker-Varadhan variational estimation ( $\gamma$ -cross entropy), respectively.  $L$  denotes the number of layers both in the nonlinear mixing function  $\mathbf{f}$  in (8) and representation function  $\mathbf{h}_x$ .

More specifically, we followed the experimental protocol in Tschannen et al. (2019)<sup>5</sup>, which has been used in the context of deep canonical correlation analysis as well (Andrew et al., 2013). We first divided a single image in half, and then used the upper and lower half images as input  $\mathbf{x}(t)$  and complementary data samples  $\mathbf{u}(t)$ , respectively (Left figures in Fig. 4). Based on these data samples  $\{\mathbf{x}(t), \mathbf{u}(t)\}_{t=1}^T$ , we applied five methods based on the following objective functions for representation learning:

- $\gamma$ -cross entropy ( $\gamma$ -CE)  $\tilde{J}_\gamma$ : In order for initialization, we first updated the parameters in  $r(\mathbf{x}, \mathbf{u})$  with  $\gamma = 0.1$  for ten epochs, and then used the updated parameters as the initial parameters for the  $\gamma$ -cross entropy with a larger  $\gamma$  value.
- Logistic regression (LR) in (5): LR can be seen as the limit of  $\gamma \rightarrow 0$  in the  $\gamma$ -cross entropy.
- Variational estimation of  $f$ -divergence ( $f$ -div) in (6) (Nguyen et al., 2008; Sugiyama et al., 2008).

5. We slightly modified the python codes available at [https://github.com/google-research/google-research/tree/master/mutual\\_information\\_representation\\_learning](https://github.com/google-research/google-research/tree/master/mutual_information_representation_learning).

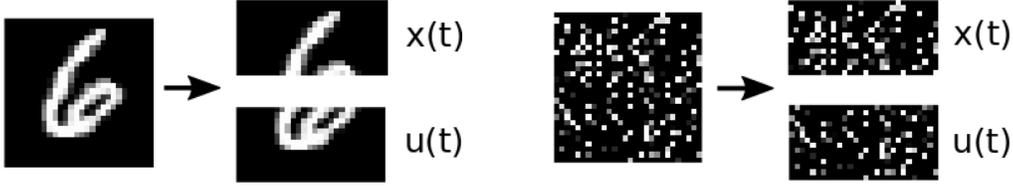


Figure 4: An example of input and complementary data. The half upper images are used for input data samples, while the half below images are complementary ones. The right images are randomly shuffled versions of left images used as outliers.

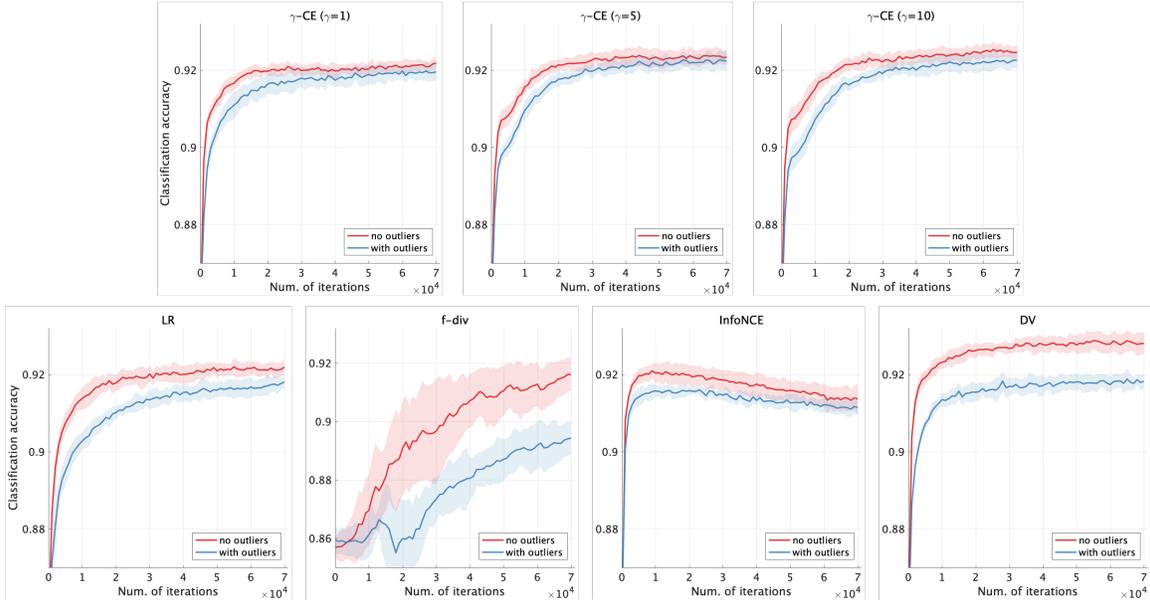


Figure 5: Averaged classification accuracy for MNIST over 10 runs. For the red lines, training is performed without outliers, while the blue lines are training with outliers. Note that the scale of the vertical axis only for  $f$ -div is different.

- *InfoNCE* (van den Oord et al., 2018):

$$\tilde{J}_{\text{NCE}}(r) := E \left[ \frac{1}{K} \sum_{i=1}^K \log \frac{e^{r(\mathbf{x}_i, \mathbf{u}_i)}}{\frac{1}{K} \sum_{j=1}^K e^{r(\mathbf{x}_i, \mathbf{u}_j)}} \right], \quad (44)$$

where the expectation is taken over  $\prod_{i=1}^K p(\mathbf{x}_i, \mathbf{u}_i)$ . In practice, we estimate (44) by averaging over multiple (mini-)batches of samples (Tschannen et al., 2019).

- *Donsker-Varadhan variational estimation of mutual information (DV)*  $\tilde{J}_{\text{DV}}$  (Ruderman et al., 2012; Belghazi et al., 2018): Unlike nonlinear ICA, this experiment excludes DV from the proposed methods because DV has been already used for maximization of MI (Hjelm et al., 2019).

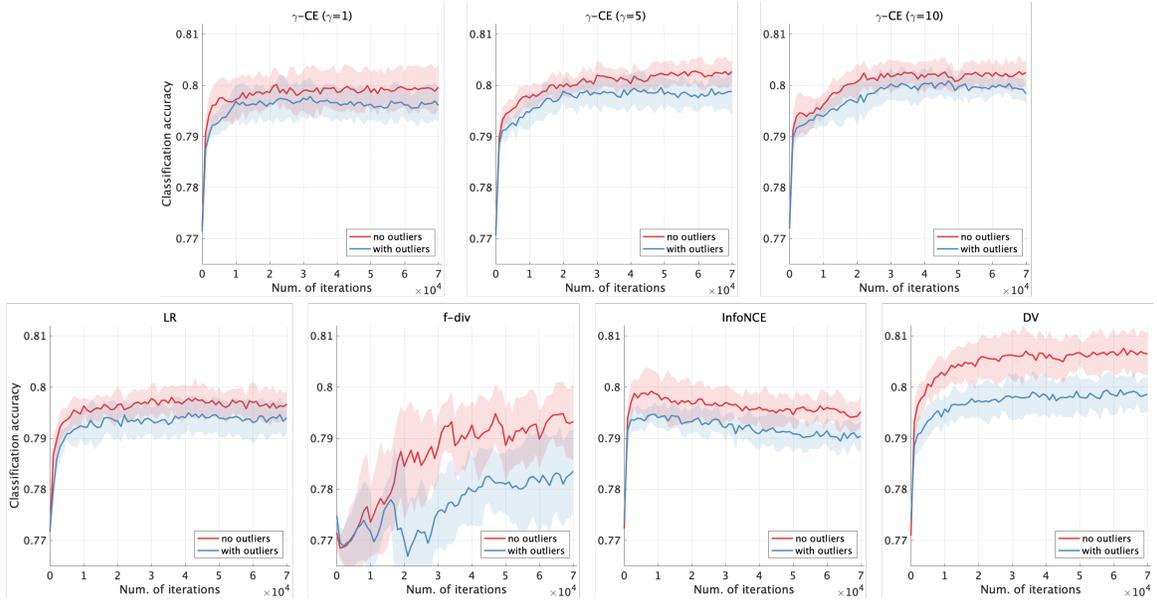


Figure 6: Averaged classification accuracy for Fashion-MNIST over 10 runs.

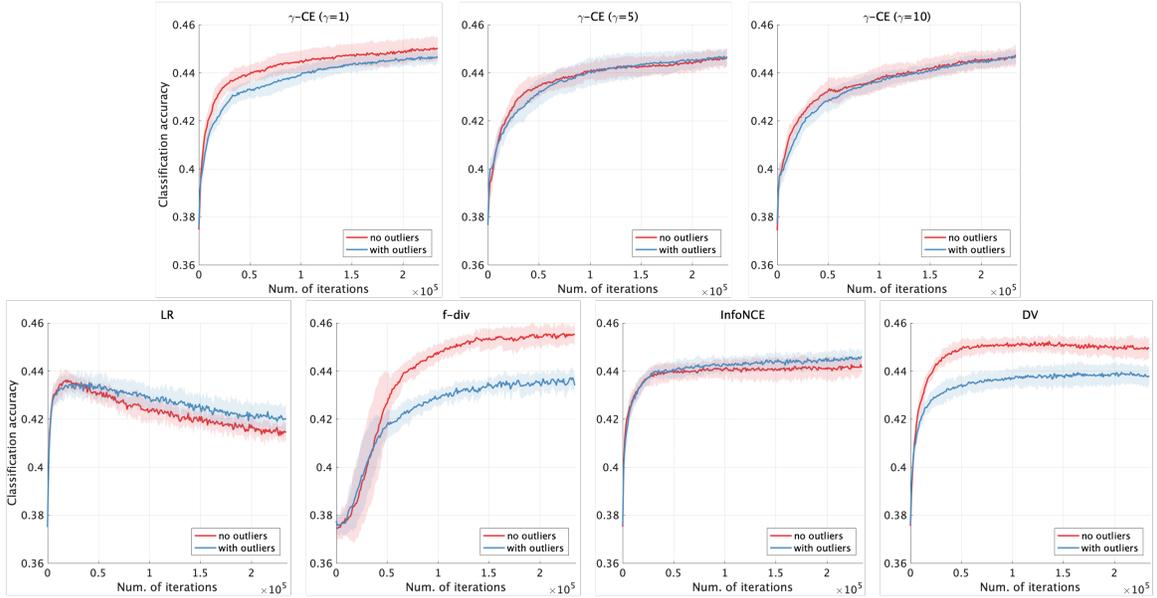


Figure 7: Averaged classification accuracy for CIFAR-10 over 10 runs.

For all methods, we used the same model  $r(\mathbf{x}, \mathbf{u})$  as

$$r(\mathbf{x}, \mathbf{u}) = \psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x})) + b(\mathbf{h}_u(\mathbf{u})).$$

The representation functions of  $\mathbf{h}_x$  and  $\mathbf{h}_u$  were modelled by the same neural architecture without parameter sharing: The first two hidden layers were convolution layers with the ReLU activation function and the third layer, which is the output layer, was a fully connected layer without the activa-

Table 6: Averages of classification accuracy for the last 5,000 iterations over 10 runs. Since we recorded classification accuracy at every 1,000 iterations, each method includes the 50 points of the accuracy over 10 runs. The numbers in parentheses indicate standard deviations. The best and comparable methods judged by the t-test at the significance level 1% are described in boldface.

| $\gamma$ -CE ( $\gamma = 5$ )      | $\gamma$ -CE ( $\gamma = 10$ ) | LR           | f-div.              | InfoNCE             | DV                  |
|------------------------------------|--------------------------------|--------------|---------------------|---------------------|---------------------|
| MNIST                              |                                |              |                     |                     |                     |
| 0.924(0.002)                       | 0.925(0.002)                   | 0.921(0.002) | 0.915(0.006)        | 0.914(0.003)        | <b>0.928(0.003)</b> |
| MNIST <i>with outliers</i>         |                                |              |                     |                     |                     |
| <b>0.922(0.002)</b>                | <b>0.922(0.002)</b>            | 0.917(0.002) | 0.893(0.006)        | 0.911(0.002)        | 0.918(0.002)        |
| Fashion-MNIST                      |                                |              |                     |                     |                     |
| 0.802(0.003)                       | 0.802(0.003)                   | 0.796(0.003) | 0.794(0.006)        | 0.795(0.003)        | <b>0.807(0.004)</b> |
| Fashion-MNIST <i>with outliers</i> |                                |              |                     |                     |                     |
| <b>0.799(0.003)</b>                | <b>0.799(0.002)</b>            | 0.794(0.003) | 0.783(0.008)        | 0.791(0.003)        | <b>0.799(0.003)</b> |
| CIFAR-10                           |                                |              |                     |                     |                     |
| 0.446(0.003)                       | 0.447(0.004)                   | 0.414(0.004) | <b>0.455(0.004)</b> | 0.442(0.004)        | 0.449(0.004)        |
| CIFAR-10 <i>with outliers</i>      |                                |              |                     |                     |                     |
| <b>0.446(0.003)</b>                | <b>0.446(0.003)</b>            | 0.420(0.005) | 0.436(0.004)        | <b>0.445(0.003)</b> | 0.438(0.004)        |

tion function. Before the output layer, we sequentially applied layer normalization (Ba et al., 2016) and average pooling. The output dimensions were fixed at  $d_x = d_u = 200$ .  $a(\mathbf{h}_x)$  and  $b(\mathbf{h}_u)$  were modelled by a one-layer feedforward network without activation functions. By following van den Oord et al. (2018) and Tian et al. (2019), we set  $\psi(\mathbf{h}_x, \mathbf{h}_u) = \mathbf{h}_x^\top \mathbf{W} \mathbf{h}_u$  where  $\mathbf{W}$  is a  $d_x$  by  $d_u$  matrix and learned from data. We optimized all parameters by the Adam optimizer (Kingma and Ba, 2015). After estimating the representation function,  $\mathbf{h}_x(\mathbf{x})$  was fixed and not learned anymore on the evaluation phase. As evaluation, we only trained a linear classifier to the learned features  $\{\mathbf{h}_x(\mathbf{x}(t))\}_{t=1}^T$  based on multinomial logistic regression. Classification accuracy on test data was used as the evaluation metric.

As datasets, we used the following three classification datasets with ten classes<sup>6</sup>:

- *MNIST* ( $D_x = 392, D_u = 392$  and  $T = 60,000$ )
- *Fashion-MNIST* ( $D_x = 392, D_u = 392$  and  $T = 60,000$ )
- *CIFAR10* ( $D_x = 1,536, D_u = 1,536$  and  $T = 50,000$ )

For MNIST and Fashion-MNIST, we used the learning rate  $10^{-4}$  and updated the parameters for 300 epochs, while the learning rate for CIFAR10 is  $10^{-5}$  and the number of epochs is 1200. In order to demonstrate the robustness to outliers, the training data was contaminated by randomly shuffled images with respect to pixels (Right figures in Fig. 4). The contamination ratio of outliers was fixed at 0.3.

6. All datasets were downloaded through the tensorflow library.

Classification accuracy of MNIST on test data samples over iterations is plotted in Fig. 5. In the case of no outliers (i.e., red lines), DV performs the best, while the classification accuracy of the  $\gamma$ -CE and LR is fairly good. For InfoNCE, the classification accuracy around the 70,000-th iteration is less than DV, LR and  $\gamma$ -CE, while  $f$ -div shows high variance of the classification accuracy over iterations. When data is contaminated by outliers (i.e., blue lines),  $f$ -div, LR and DV show clear performance degeneration. On the other hand, the performance of  $\gamma$ -CE is not so influenced by the contamination of outliers. Interestingly, estimation based on InfoNCE is also not strongly hampered by outliers, but the classification accuracy is worse than  $\gamma$ -CE around the 70000-th iteration. The same tendency of the robustness of the  $\gamma$ -cross entropy was observed both for Fashion-MNIST and CIFAR-10 (Figs. 6 and 7). Table 6 quantitatively indicates that the  $\gamma$ -cross entropy yields the best classification accuracy when data is contaminated by outliers, and shows a fairly good permanence even without outliers. Thus, the proposed method based on the  $\gamma$ -cross entropy can be robust against outliers as implied in our theoretical analysis and fairly works well even when data is not contaminated by outliers.

Finally, let us note that these results are not trivial because the connection between mutual information and quality of data representation has been experimentally demonstrated to be rather loose (Tian et al., 2019). Thus, it was unclear that robust estimation always leads to robust representations of data. Nonetheless, the proposed method based on the  $\gamma$ -cross entropy performed the best in the presence of outliers. This means that the proposed robust method for unsupervised representation learning is promising.

## 6. Conclusion

This paper theoretically showed that density ratio estimation plays a key role in three frameworks for unsupervised representation learning: Maximization of mutual information and nonlinear independent component analysis as well as nonlinear subspace estimation, which is a novel framework proposed in this paper. Furthermore, we made theoretical contributions in each of the three frameworks: We showed that density ratio estimation is necessary and sufficient in maximization of mutual information, while our analysis revealed a novel insight for source recover in nonlinear ICA that the dimensionality of complementary data is an important factor for source recovery, which was clearly supported by numerical experiments. In addition, the proposed generative model in nonlinear subspace estimation is more general than nonlinear ICA in the sense that the latent source components are no longer assumed to be conditionally independent, and theoretical conditions to estimate a nonlinear subspace of the latent source components were given. Motivated by the theoretical results, we developed a nonlinear ICA method by applying an variation lower-bound of mutual information, and proposed an outlier-robust method for unsupervised representation learning through density ratio estimation. We also theoretically investigated their outlier-robustness. The usefulness of the proposed methods were demonstrated through numerical experiments of nonlinear ICA and linear classification in a downstream task.

## Acknowledgments

The authors would like to thank Dr. Hiroshi Morioka for sharing his python codes of permutation contrastive learning with us. Takashi Takenouchi was partially supported by JSPS KAKENHI Grant Number 20K03753 and 19H04071.

## Appendix A. Proof of Theorem 1

**Proof** We first recall that  $\mathbf{y}_x := \mathbf{h}_x(x)$  and  $\mathbf{y}_u := \mathbf{h}_u(u)$ , and denote  $\mathbf{h}_x^\perp(x)$  and  $\mathbf{h}_u^\perp(u)$  by  $\mathbf{y}_x^\perp$  and  $\mathbf{y}_u^\perp$ , respectively. Based on Assumption (A2),  $(\mathbf{y}_x, \mathbf{y}_x^\perp) := (\mathbf{h}_x(x), \mathbf{h}_x^\perp(x))$  and  $(\mathbf{y}_u, \mathbf{y}_u^\perp) := (\mathbf{h}_u(u), \mathbf{h}_u^\perp(u))$  are both invertible in the sense that there exist functions  $\mathbf{g}_x$  and  $\mathbf{g}_u$  such that  $x = \mathbf{g}_x(\mathbf{y}_x, \mathbf{y}_x^\perp)$  and  $u = \mathbf{g}_u(\mathbf{y}_u, \mathbf{y}_u^\perp)$ . Based on the invertible functions, we first establish the following lemma, which decomposes  $I(\mathcal{X}, \mathcal{U})$  into three terms:

**Lemma 9.** *Assumptions (A1-2) hold. Then,  $I(\mathcal{X}, \mathcal{U})$  can be decomposed as follows:*

$$I(\mathcal{X}, \mathcal{U}) = I(\mathcal{Y}_x, \mathcal{Y}_u) + E_{Y_u}[I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u)] + E_{Y_x}[I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x)], \quad (45)$$

where  $E_{Y_u}$  and  $E_{Y_x}$  denote the expectations over  $p(\mathbf{y}_u)$  and  $p(\mathbf{y}_x)$  respectively,

$$\begin{aligned} I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) &= \iint p(\mathbf{y}_x, \mathbf{y}_u^\perp | \mathbf{y}_u) \log \frac{p(\mathbf{y}_x, \mathbf{y}_u^\perp | \mathbf{y}_u)}{p(\mathbf{y}_x | \mathbf{y}_u) p(\mathbf{y}_u^\perp | \mathbf{y}_u)} d\mathbf{y}_x d\mathbf{y}_u^\perp \\ I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) &:= \iint p(\mathbf{y}_x^\perp, \mathbf{u} | \mathbf{y}_x) \log \frac{p(\mathbf{y}_x^\perp, \mathbf{u} | \mathbf{y}_x)}{p(\mathbf{y}_x^\perp | \mathbf{y}_x) p(\mathbf{u} | \mathbf{y}_x)} d\mathbf{y}_x^\perp d\mathbf{u}. \end{aligned}$$

The proof is given in Appendix A.1. The first term on the right-hand side of (45) is mutual information between the representation functions of  $x$  and  $u$ , while  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x)$  and  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u)$  measure the conditional independence and are equal to 0 when  $\mathbf{u} \perp \mathbf{y}_x^\perp | \mathbf{y}_x$  and  $\mathbf{y}_x \perp \mathbf{y}_u^\perp | \mathbf{y}_u$ . It follows from (45) that  $I(\mathcal{X}, \mathcal{U}) = I(\mathcal{Y}_x, \mathcal{Y}_u)$  if and only if  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) = 0$  and  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0$ .

Next, we fix  $\mathbf{h}_x$  and  $\mathbf{h}_u$  at  $\mathbf{h}_x^*$  and  $\mathbf{h}_u^*$  respectively, and substitute them as  $\mathbf{y}_x = \mathbf{h}_x^*(x)$  and  $\mathbf{y}_u = \mathbf{h}_u^*(u)$ . Then, it is shown that (10) implies both  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) = 0$  and  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0$ . We first multiply  $p(u)$  to both sides and rewrite (10) as

$$p(\mathbf{u} | x) = p(\mathbf{u} | \mathbf{y}_x), \quad (46)$$

where with the partition function  $Z(\mathbf{y}_x) := \int \exp(\psi^*(\mathbf{y}_x, \mathbf{h}_u^*(u)) + a^*(\mathbf{y}_x) + b^*(\mathbf{h}_u^*(u))) p(u) du$ ,

$$p(\mathbf{u} | \mathbf{y}_x) := \frac{\exp(\psi^*(\mathbf{y}_x, \mathbf{h}_u^*(u)) + a^*(\mathbf{y}_x) + b^*(\mathbf{h}_u^*(u))) p(u)}{Z(\mathbf{y}_x)}.$$

By the invertible assumption for  $(\mathbf{y}_x, \mathbf{y}_x^\perp)$ , (46) can be equivalently expressed as the following conditional independence:

$$\mathbf{u} \perp \mathbf{y}_x^\perp | \mathbf{y}_x, \quad (47)$$

implying that

$$I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0.$$

Similarly, we can derive the following equation from (10):

$$p(\mathbf{x}|\mathbf{u}) = p(\mathbf{x}|\mathbf{y}_u), \quad (48)$$

where with the partition function  $Z(\mathbf{y}_u) := \int \exp(\psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{y}_u) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{y}_u)) p(\mathbf{x}) d\mathbf{x}$ ,

$$p(\mathbf{x}|\mathbf{y}_u) := \frac{\exp(\psi^*(\mathbf{h}_x^*(\mathbf{x}), \mathbf{y}_u) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{y}_u)) p(\mathbf{x})}{Z(\mathbf{y}_u)}.$$

Again, (48) implies the conditional independence as

$$\mathbf{x} \perp \mathbf{y}_u^\perp | \mathbf{y}_u. \quad (49)$$

Data processing inequality and conditional independence (49) lead to

$$I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) \leq I(\mathcal{X} | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) \stackrel{(49)}{=} 0.$$

Both  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) = 0$  and  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0$  enable us to ensure  $I(\mathcal{X}, \mathcal{U}) = I(\mathcal{Y}_x, \mathcal{Y}_u)$ .

Conversely, we suppose that  $I(\mathcal{X}, \mathcal{U}) = I(\mathcal{Y}_x, \mathcal{Y}_u)$  at  $\mathbf{h}_x = \mathbf{h}_x^*$  and  $\mathbf{h}_u = \mathbf{h}_u^*$ . Then, it follows from (45) that  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) = 0$  and  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0$ .  $I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x) = 0$  means the conditional independence (47), which can be equivalently expressed as (46). On the other hand, under the invertibility of  $(\mathbf{y}_u, \mathbf{y}_u^\perp)$ ,  $I(\mathcal{Y}_x | \mathcal{Y}_u, \mathcal{Y}_u^\perp | \mathcal{Y}_u) = 0$  implies

$$\mathbf{y}_x \perp \mathbf{y}_u^\perp | \mathbf{y}_u \quad \text{or equivalently} \quad p(\mathbf{y}_x | \mathbf{u}) = p(\mathbf{y}_x | \mathbf{y}_u). \quad (50)$$

A simple calculation yields

$$p(\mathbf{x}, \mathbf{u}) = p(\mathbf{u}|\mathbf{x})p(\mathbf{x}) \stackrel{(46)}{=} p(\mathbf{u}|\mathbf{y}_x)p(\mathbf{x}) \stackrel{(*)}{=} \frac{p(\mathbf{y}_x|\mathbf{u})p(\mathbf{u})}{p(\mathbf{y}_x)} p(\mathbf{x}) \stackrel{(50)}{=} \frac{p(\mathbf{y}_x|\mathbf{y}_u)}{p(\mathbf{y}_x)} p(\mathbf{x})p(\mathbf{u}), \quad (51)$$

where we applied Bayes' theorem on  $(*)$ . By dividing both sides of (51) by  $p(\mathbf{x})p(\mathbf{u})$ , we have

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} = \log p(\mathbf{y}_x, \mathbf{y}_u) - \log p(\mathbf{y}_x) - \log p(\mathbf{y}_u).$$

The existence of the functions is obvious by denoting  $\log p(\mathbf{y}_x, \mathbf{y}_u)$ ,  $-\log p(\mathbf{y}_x)$  and  $-\log p(\mathbf{y}_u)$  by  $\psi^*(\mathbf{y}_x, \mathbf{y}_u)$ ,  $a^*(\mathbf{y}_x)$  and  $b^*(\mathbf{y}_u)$  respectively. The proof is completed.  $\blacksquare$

### A.1 Proof of Lemma 9

**Proof** We first decompose  $I(\mathcal{X}, \mathcal{U})$  based on the invertibility assumption of  $(\mathbf{y}_x, \mathbf{y}_x^\perp)$  as follows:

$$\begin{aligned} I(\mathcal{X}, \mathcal{U}) &= \iiint p(\mathbf{y}_x, \mathbf{y}_x^\perp, \mathbf{u}) \log \frac{p(\mathbf{y}_x, \mathbf{y}_x^\perp, \mathbf{u})}{p(\mathbf{y}_x, \mathbf{y}_x^\perp)p(\mathbf{u})} d\mathbf{y}_x d\mathbf{y}_x^\perp d\mathbf{u} \\ &= \iint p(\mathbf{y}_x, \mathbf{u}) \log \frac{p(\mathbf{y}_x, \mathbf{u})}{p(\mathbf{y}_x)p(\mathbf{u})} d\mathbf{y}_x d\mathbf{u} \\ &\quad + \int p(\mathbf{y}_x) \left[ \iint p(\mathbf{y}_x^\perp, \mathbf{u} | \mathbf{y}_x) \log \frac{p(\mathbf{y}_x^\perp, \mathbf{u} | \mathbf{y}_x)}{p(\mathbf{y}_x^\perp | \mathbf{y}_x)p(\mathbf{u} | \mathbf{y}_x)} d\mathbf{u} d\mathbf{y}_x^\perp \right] d\mathbf{y}_x \\ &= I(\mathcal{Y}_x, \mathcal{U}) + E_{\mathcal{Y}_x} [I(\mathcal{Y}_x^\perp | \mathcal{Y}_x, \mathcal{U} | \mathcal{Y}_x)], \end{aligned} \quad (52)$$

where we used the following relation:

$$\frac{p(\mathbf{y}_x, \mathbf{y}_x^\perp, \mathbf{u})}{p(\mathbf{y}_x, \mathbf{y}_x^\perp)p(\mathbf{u})} = \frac{p(\mathbf{y}_x^\perp, \mathbf{u}|\mathbf{y}_x)p(\mathbf{y}_x)}{p(\mathbf{y}_x^\perp|\mathbf{y}_x)p(\mathbf{y}_x)p(\mathbf{u})} \cdot \frac{p(\mathbf{y}_x, \mathbf{u})}{p(\mathbf{y}_x, \mathbf{u})} = \frac{p(\mathbf{y}_x, \mathbf{u})}{p(\mathbf{y}_x)p(\mathbf{u})} \cdot \frac{p(\mathbf{y}_x^\perp, \mathbf{u}|\mathbf{y}_x)}{p(\mathbf{y}_x^\perp|\mathbf{y}_x)p(\mathbf{u}|\mathbf{y}_x)}.$$

Applying the same decomposition as (52) to  $I(\mathcal{Y}_x, \mathcal{U})$  under the change of variables by the invertible  $(\mathbf{y}_u, \mathbf{y}_u^\perp)$  yields

$$\begin{aligned} I(\mathcal{Y}_x, \mathcal{U}) &= \iiint p(\mathbf{y}_x, \mathbf{y}_u, \mathbf{y}_u^\perp) \log \frac{p(\mathbf{y}_x, \mathbf{y}_u, \mathbf{y}_u^\perp)}{p(\mathbf{y}_x)p(\mathbf{y}_u, \mathbf{y}_u^\perp)} d\mathbf{y}_x d\mathbf{y}_u d\mathbf{y}_u^\perp \\ &= \iint p(\mathbf{y}_x, \mathbf{y}_u) \log \frac{p(\mathbf{y}_x, \mathbf{y}_u)}{p(\mathbf{y}_x)p(\mathbf{y}_u)} d\mathbf{y}_x d\mathbf{y}_u \\ &\quad + \int p(\mathbf{y}_u) \left[ \iint p(\mathbf{y}_x, \mathbf{y}_u^\perp|\mathbf{y}_u) \log \frac{p(\mathbf{y}_x, \mathbf{y}_u^\perp|\mathbf{y}_u)}{p(\mathbf{y}_u^\perp|\mathbf{y}_u)p(\mathbf{y}_x|\mathbf{y}_u)} d\mathbf{y}_x d\mathbf{y}_u^\perp \right] d\mathbf{y}_u \\ &= I(\mathcal{Y}_x, \mathcal{Y}_u) + E_{\mathcal{Y}_u}[I(\mathcal{Y}_x|\mathcal{Y}_u, \mathcal{Y}_u^\perp|\mathcal{Y}_u)], \end{aligned} \quad (53)$$

where we used

$$\frac{p(\mathbf{y}_x, \mathbf{y}_u, \mathbf{y}_u^\perp)}{p(\mathbf{y}_x)p(\mathbf{y}_u, \mathbf{y}_u^\perp)} = \frac{p(\mathbf{y}_x, \mathbf{y}_u^\perp|\mathbf{y}_u)p(\mathbf{y}_u)}{p(\mathbf{y}_x)p(\mathbf{y}_u^\perp|\mathbf{y}_u)p(\mathbf{y}_u)} \cdot \frac{p(\mathbf{y}_x, \mathbf{y}_u)}{p(\mathbf{y}_x, \mathbf{y}_u)} = \frac{p(\mathbf{y}_x, \mathbf{y}_u)}{p(\mathbf{y}_x)p(\mathbf{y}_u)} \cdot \frac{p(\mathbf{y}_x, \mathbf{y}_u^\perp|\mathbf{y}_u)}{p(\mathbf{y}_u^\perp|\mathbf{y}_u)p(\mathbf{y}_x|\mathbf{y}_u)}$$

Substituting (53) into (52) completes the proof.  $\blacksquare$

## Appendix B. Proof of Proposition 2

**Proof** Let us denote the inverse of the mixing function  $\mathbf{f}$  in the generative model (8) as  $\mathbf{g} := \mathbf{f}^{-1}$ . From the conditional independence of  $\mathbf{s}$  given  $\mathbf{u}$  in Assumption (B1), we obtain the conditional density of  $\mathbf{x}$  given  $\mathbf{u}$  under the change of variables by  $\mathbf{s} = \mathbf{g}(\mathbf{x})$  as follows:

$$\log p(\mathbf{x}|\mathbf{u}) = \sum_{i=1}^{D_x} q_i(g_i(\mathbf{x}), \mathbf{u}) + \log |\det \mathbf{J}_g(\mathbf{x})| - \log Z(\mathbf{u}), \quad (54)$$

where  $\mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), g_2(\mathbf{x}), \dots, g_{D_x}(\mathbf{x}))^\top$  and  $\mathbf{J}_g(\mathbf{x})$  is the Jacobian of  $\mathbf{g}$  at  $\mathbf{x}$ . Similarly, by applying the same change of variables for the marginal density of  $\mathbf{s}$ , the marginal density of  $\mathbf{x}$  is given by

$$\log p(\mathbf{x}) = \log p(\mathbf{g}(\mathbf{x})) + \log |\det \mathbf{J}_g(\mathbf{x})|. \quad (55)$$

On the other hand, the log-density ratio equation (11) yields

$$\log p(\mathbf{x}|\mathbf{u}) - \log p(\mathbf{x}) = \sum_{i=1}^{D_x} \psi_i^*(h_{x,i}^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})). \quad (56)$$

Substitution of (54) and (55) into the left-hand side of (56) cancels out the Jacobian term and gives the following equation:

$$\sum_{i=1}^{D_x} q_i(g_i(\mathbf{x}), \mathbf{u}) - \log Z(\mathbf{u}) - \log p(\mathbf{g}(\mathbf{x})) = \sum_{i=1}^{D_x} \psi_i^*(h_{x,i}^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + a^*(\mathbf{h}_x^*(\mathbf{x})) + b^*(\mathbf{h}_u^*(\mathbf{u})). \quad (57)$$

Then, we compute the gradient of the both sides on (57) with respect to  $\mathbf{u}$  as

$$\sum_{i=1}^{D_x} \nabla_{\mathbf{u}} q_i(g_i(\mathbf{x}), \mathbf{u}) - \nabla_{\mathbf{u}} \log Z(\mathbf{u}) = \sum_{i=1}^{D_x} \nabla_{\mathbf{u}} \psi_i^*(h_{x,i}^*(\mathbf{x}), \mathbf{h}_u^*(\mathbf{u})) + \nabla_{\mathbf{u}} b^*(\mathbf{h}_u^*(\mathbf{u})).$$

By using the notations of  $\mathbf{z} := \mathbf{h}_x^*(\mathbf{x})$  and  $\mathbf{v}(\mathbf{z}) := \mathbf{g}((\mathbf{h}_x^*)^{-1}(\mathbf{z}))$  where  $(\mathbf{h}_x^*)^{-1}$  denotes the inverse of  $\mathbf{h}_x^*$ ,

$$\sum_{i=1}^{D_x} \nabla_{\mathbf{u}} q_i(v_i(\mathbf{z}), \mathbf{u}) - \nabla_{\mathbf{u}} \log Z(\mathbf{u}) = \sum_{i=1}^{D_x} \nabla_{\mathbf{u}} \psi_i^*(z_i, \mathbf{h}_u^*(\mathbf{u})) + \nabla_{\mathbf{u}} b^*(\mathbf{h}_u^*(\mathbf{u})), \quad (58)$$

where  $z_i$  and  $v_i$  are the  $i$ -th elements in  $\mathbf{z}$  and  $\mathbf{v}$ , respectively.

Next, we show from (58) that each element in  $\mathbf{v}(\mathbf{z})(= \mathbf{s})$  is a function of a distinct and single element in  $\mathbf{z}(= \mathbf{h}_x^*(\mathbf{x}))$ . To this end, we first take the partial derivative to both sides on (58) with respect to  $z_l$  as

$$\sum_{i=1}^{D_x} v_i^{(l)}(\mathbf{z}) \nabla_{\mathbf{u}} q_i'(v_i(\mathbf{z}), \mathbf{u}) = \frac{\partial}{\partial z_l} [\nabla_{\mathbf{u}} \psi_l^*(z_l, \mathbf{h}_u^*(\mathbf{u}))], \quad (59)$$

where  $q_i'(t, \mathbf{u}) := \frac{\partial q_i(t, \mathbf{u})}{\partial t}$  and  $v_i^{(l)}(\mathbf{z}) := \frac{\partial v_i(\mathbf{z})}{\partial z_l}$ . We further take the partial derivative of both sides on (59) with respect to  $z_m$  for  $m \neq l$  as

$$\sum_{i=1}^{D_x} \left[ v_i^{(l,m)}(\mathbf{z}) \nabla_{\mathbf{u}} q_i'(v_i(\mathbf{z}), \mathbf{u}) + v_i^{(l)}(\mathbf{z}) v_i^{(m)}(\mathbf{z}) \nabla_{\mathbf{u}} q_i''(v_i(\mathbf{z}), \mathbf{u}) \right] = \mathbf{0}, \quad (60)$$

where  $q_i''(t, \mathbf{u}) := \frac{\partial^2}{\partial t^2} q_i(t, \mathbf{u})$  and  $v_i^{(l,m)}(\mathbf{z}) := \frac{\partial^2}{\partial z_l \partial z_m} v_i(\mathbf{z})$ . In order to express (60) as a matrix form, we define the  $D_x(D_x - 1)/2$  by  $2D_x$  matrix consisting of a collection of  $v_i^{(l)}(\mathbf{z}) v_i^{(m)}(\mathbf{z})$  and  $v_i^{(l,m)}(\mathbf{z})$  with respect to  $i, l, m$  ( $l \neq m$ ) as

$$\mathbf{M}(\mathbf{v}) := \begin{pmatrix} v_1^{(1,2)} & \cdots & v_{D_x}^{(1,2)} & v_1^{(1)} v_1^{(2)} & \cdots & v_{D_x}^{(1)} v_{D_x}^{(2)} \\ v_1^{(1,3)} & \cdots & v_{D_x}^{(1,3)} & v_1^{(1)} v_1^{(3)} & \cdots & v_{D_x}^{(1)} v_{D_x}^{(3)} \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ v_1^{(D_x-1, D_x)} & \cdots & v_{D_x}^{(D_x-1, D_x)} & v_1^{(D_x-1)} v_1^{(D_x)} & \cdots & v_{D_x}^{(D_x-1)} v_{D_x}^{(D_x)} \end{pmatrix} \in \mathbb{R}^{\frac{D_x(D_x-1)}{2} \times 2D_x}.$$

We also define a collection of  $q_i'(v_i(\mathbf{z}), \mathbf{u})$  and  $q_i''(v_i(\mathbf{z}), \mathbf{u})$  by the following  $2D_x$ -dimensional vector:

$$\mathbf{w}(\mathbf{v}, \mathbf{u}) := (q_1'(v_1, \mathbf{u}), \dots, q_{D_x}'(v_{D_x}, \mathbf{u}), q_1''(v_1, \mathbf{u}), \dots, q_{D_x}''(v_{D_x}, \mathbf{u}))^\top \in \mathbb{R}^{2D_x}.$$

By fixing  $\mathbf{u}$  at  $\mathbf{u} = \mathbf{u}_1$ , (60) can be expressed by using  $\mathbf{M}(\mathbf{v})$  and  $\mathbf{w}(\mathbf{v}, \mathbf{u})$  as

$$\mathbf{M}(\mathbf{v}) \mathbf{W}(\mathbf{v}) = \mathbf{O}, \quad (61)$$

where  $\mathbf{W}(\mathbf{v}) := \nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u}) \Big|_{\mathbf{u}=\mathbf{u}_1} \in \mathbb{R}^{2D_x \times D_u}$  and  $\mathbf{O}$  denotes the null matrix. Since Assumptions (B3-4) ensure that the rank of  $\mathbf{W}(\mathbf{v})$  is  $2D_x$ ,  $\mathbf{W}(\mathbf{v})\mathbf{W}(\mathbf{v})^\top$  is a  $2D_x$  by  $2D_x$  invertible matrix. Thus, multiplying  $\mathbf{W}(\mathbf{v})^\top$  to (61) and taking the inverse of  $\mathbf{W}(\mathbf{v})\mathbf{W}(\mathbf{v})^\top$  on the right yields

$$\mathbf{M}(\mathbf{v}) = \mathbf{O}. \quad (62)$$

Eq.(62) indicates that  $v_i^{(l,m)}(\mathbf{z}) = 0$  and  $v_i^{(l)}(\mathbf{z})v_i^{(m)}(\mathbf{z}) = 0$  for all  $i, l, m = 1, \dots, D_x$  with  $l \neq m$ . This means that each element in  $\mathbf{v}(\mathbf{z})$  is a function of a distinct and single element in  $\mathbf{z}$  because  $\mathbf{v}(\mathbf{z})$  is an invertible function. Thus, the proof is completed.  $\blacksquare$

### Appendix C. Rank condition in the exponential family

We first express the exponential family (12) as

$$\log p(\mathbf{s}|\mathbf{u}) = \sum_{k=1}^K \boldsymbol{\lambda}_k(\mathbf{u})^\top \mathbf{q}_k(\mathbf{s}) - \log Z(\mathbf{u}),$$

where  $\boldsymbol{\lambda}_k(\mathbf{u}) := (\lambda_{1k}(\mathbf{u}), \dots, \lambda_{D_x k}(\mathbf{u}))^\top$  and  $\mathbf{q}_k(\mathbf{s}) := (q_{1k}(s_1), \dots, q_{D_x k}(s_{D_x}))$ . Then, a simple computation yields

$$\nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u}) = \sum_{k=1}^K \boldsymbol{\Lambda}_k \odot \mathbf{Q}_k, \quad (63)$$

where  $\odot$  denotes the Hadamard product (i.e., elementwise multiplication of two matrices)

$$\boldsymbol{\Lambda}_k := \begin{pmatrix} \nabla_{\mathbf{u}} \lambda_k(\mathbf{u}) \\ \nabla_{\mathbf{u}} \lambda_k(\mathbf{u}) \end{pmatrix},$$

and with  $q'_{ik}(t) := \frac{d}{dt} q_{ik}(t)$  and  $q''_{ik}(t) := \frac{d^2}{dt^2} q_{ik}(t)$  for  $i = 1, \dots, D_x$ ,

$$\mathbf{Q}_k = \begin{pmatrix} q'_{1k}(v_1) & q'_{1k}(v_1) & \cdots & q'_{1k}(v_1) \\ \vdots & \vdots & \cdots & \vdots \\ q'_{D_x k}(v_{D_x}) & q'_{D_x k}(v_{D_x}) & \cdots & q'_{D_x k}(v_{D_x}) \\ q''_{1k}(v_1) & q''_{1k}(v_1) & \cdots & q''_{1k}(v_1) \\ \vdots & \vdots & \cdots & \vdots \\ q''_{D_x k}(v_{D_x}) & q''_{D_x k}(v_{D_x}) & \cdots & q''_{D_x k}(v_{D_x}) \end{pmatrix}.$$

The rank of  $\boldsymbol{\Lambda}_k$  is at most  $D_x$  under Assumption (B3) because  $\boldsymbol{\Lambda}_k$  is a vertical concatenation of the two same matrices of  $\nabla_{\mathbf{u}} \lambda_k(\mathbf{u}) \in \mathbb{R}^{D_x \times D_u}$ . On the other hand,  $\mathbf{Q}_k$  is a rank-one matrix because the column vectors in  $\mathbf{Q}_k$  are all same. Thus, the rank factorization ensures that the following decompositions of  $\boldsymbol{\Lambda}_k$  and  $\mathbf{Q}_k$  exist:

$$\boldsymbol{\Lambda}_k = \sum_{i=1}^{D_x} \mathbf{a}_{ik} \mathbf{b}_{ik}^\top \quad \text{and} \quad \mathbf{Q}_k = \mathbf{c}_k \mathbf{d}_k^\top, \quad (64)$$

where  $\mathbf{a}_{ik}, \mathbf{c}_k \in \mathbb{R}^{2D_x}$  and  $\mathbf{b}_{ik}, \mathbf{d}_k \in \mathbb{R}^{D_u}$ . Substituting (64) into (63) yields

$$\nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u}) = \sum_{k=1}^K \sum_{i=1}^{D_x} (\mathbf{a}_{ik} \odot \mathbf{c}_k)(\mathbf{b}_{ik} \odot \mathbf{d}_k)^\top. \quad (65)$$

indicating that the rank of  $\nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u})$  is at most  $2D_x$  when  $K > 1$ . On the other hand, when  $K = 1$ , (65) is given by

$$\nabla_{\mathbf{u}} \mathbf{w}(\mathbf{v}, \mathbf{u}) = \sum_{i=1}^{D_x} (\mathbf{a}_{i1} \odot \mathbf{c}_1)(\mathbf{b}_{i1} \odot \mathbf{d}_1)^\top,$$

and thus the rank is at most  $D_x$ .

### Appendix D. Proof of Theorem 3

Our proof is based on the following lemmas:

**Lemma 10.** *Suppose that  $\mathbf{D}(\mathbf{v})$  is an  $n$  by  $n$  diagonal matrix whose diagonals  $d_i(\mathbf{v})$  is a function of  $\mathbf{v} \in \mathbb{R}^n$ , and  $\mathbf{A}$  and  $\mathbf{B}$  are  $n$  by  $n$  constant matrices. Furthermore, the following assumptions are made:*

- (1) *There exist  $n$  points  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$  such that  $\mathbf{d}(\mathbf{v}_1), \mathbf{d}(\mathbf{v}_2), \dots, \mathbf{d}(\mathbf{v}_n)$  are linearly independent where  $\mathbf{d}(\mathbf{v}) := (d_1(\mathbf{v}), d_2(\mathbf{v}), \dots, d_n(\mathbf{v}))^\top$  is the vector of the diagonal elements in  $\mathbf{D}(\mathbf{v})$ .*
- (2)  *$\mathbf{A}$  and  $\mathbf{B}$  are of full-rank.*

*Then, when  $\mathbf{A}\mathbf{D}(\mathbf{v})\mathbf{B}$  is a diagonal matrix at least at  $n$  points  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ , then both  $\mathbf{A}$  and  $\mathbf{B}$  are diagonal matrices multiplied by a permutation matrix.*

**Lemma 11** (Theorem 4.4.8 in Harville (2006)). *Suppose that  $\mathbf{A}$  is an  $m$  by  $n$  matrix of rank  $r$ . For any  $m$  by  $r$  matrix  $\mathbf{B}$  and  $r$  by  $n$  matrix  $\mathbf{T}$  such that  $\mathbf{A} = \mathbf{B}\mathbf{T}$ , both  $\mathbf{B}$  and  $\mathbf{T}$  have rank  $r$ .*

Lemma 10 is proved in Appendix D.2, while the reader may refer to the proof of Theorem 4.4.8 in Harville (2006) for Lemma 11.

#### D.1 Main proof

**Proof** We start by using the same notations and following the same line of the proof until (57) in Section B. Then, we obtain the following equation under Assumptions (B'1-2,6):

$$\begin{aligned} & \sum_{i=1}^{D_x} q_i(g_i(\mathbf{x}), \lambda_i(\mathbf{u})) - \log Z(\mathbf{u}) - \log p(\mathbf{g}(\mathbf{x})) \\ &= \sum_{i=1}^{D_x} \psi_i^*(h_{\mathbf{x},i}^*(\mathbf{x}), h_{\mathbf{u},i}^*(\mathbf{u})) + a^*(h_{\mathbf{x}}^*(\mathbf{x})) + b^*(h_{\mathbf{u}}^*(\mathbf{u})), \end{aligned} \quad (66)$$

where  $\mathbf{g}$  denotes the inverse of the mixing function  $\mathbf{f}$  and  $\mathbf{g}(\mathbf{x}) = (g_1(\mathbf{x}), g_2(\mathbf{x}), \dots, g_{D_x}(\mathbf{x}))^\top$ . Let us denote  $\mathbf{z} := \mathbf{h}_x^*(\mathbf{x})$ ,  $\mathbf{v}(\mathbf{z}) := \mathbf{g}((\mathbf{h}_x^*)^{-1}(\mathbf{z}))$ ,  $r_{u,i} := h_{u,i}^*(\mathbf{u})$  and  $r_{q,i} := \lambda_i(\mathbf{u})$ . Then, taking a cross-derivative with respect to  $v_l$  and  $r_{q,m}$  to both side on (66) yields

$$\sum_{i=1}^{D_x} \frac{\partial^2 q_i(v_i, r_{q,i})}{\partial v_l \partial r_{q,m}} = \sum_{i=1}^{D_x} \frac{\partial z_i}{\partial v_l} \frac{\partial^2 \psi_i^*(z_i, r_{u,i})}{\partial z_i \partial r_{u,i}} \sum_{j=1}^{D_x} \frac{\partial r_{u,i}}{\partial u_j} \frac{\partial u_j}{\partial r_{q,m}}. \quad (67)$$

To express (67) as a matrix form, we define the following  $D_x$  by  $D_x$  diagonal matrices:

$$\mathbf{D}_q(\mathbf{v}, \mathbf{r}_q) := \begin{pmatrix} \frac{\partial^2 q_1(v_1, r_{q,1})}{\partial v_1 \partial r_{q,1}} & 0 & \cdots & 0 \\ 0 & \frac{\partial^2 q_2(v_2, r_{q,2})}{\partial v_2 \partial r_{q,2}} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \frac{\partial^2 q_{D_x}(v_{D_x}, r_{q,D_x})}{\partial v_{D_x} \partial r_{q,D_x}} \end{pmatrix} \in \mathbb{R}^{D_x \times D_x}$$

$$\mathbf{D}_\psi(\mathbf{v}, \mathbf{r}_u) := \begin{pmatrix} \frac{\partial^2 \psi_1^*(v_1, r_{u,1})}{\partial v_1 \partial r_{u,1}} & 0 & \cdots & 0 \\ 0 & \frac{\partial^2 \psi_2^*(v_2, r_{u,2})}{\partial v_2 \partial r_{u,2}} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \frac{\partial^2 \psi_{D_x}^*(v_{D_x}, r_{u,D_x})}{\partial v_{D_x} \partial r_{u,D_x}} \end{pmatrix} \in \mathbb{R}^{D_x \times D_x},$$

where  $\mathbf{r}_q = (r_{q,1}, r_{q,2}, \dots, r_{q,D_x})^\top$  and  $\mathbf{r}_u = (r_{u,1}, r_{u,2}, \dots, r_{u,D_x})^\top$ . Then, (67) can be expressed as

$$\mathbf{D}_q(\mathbf{v}, \mathbf{r}_q) = \mathbf{J}_z(\mathbf{v}) \mathbf{D}_\psi(\mathbf{v}, \mathbf{r}_u) \mathbf{J}_{r_u}(\mathbf{u}) \mathbf{J}_u(\mathbf{r}_q), \quad (68)$$

where  $\mathbf{J}_z(\mathbf{v}) := \nabla_{\mathbf{v}} \mathbf{z} \in \mathbb{R}^{D_x \times D_x}$ ,  $\mathbf{J}_{r_u}(\mathbf{u}) := \nabla_{\mathbf{u}} \mathbf{r}_u \in \mathbb{R}^{D_x \times D_u}$  and  $\mathbf{J}_u(\mathbf{r}_q) := \nabla_{\mathbf{r}_q} \mathbf{u} \in \mathbb{R}^{D_u \times D_x}$  are the Jacobian matrices of  $\mathbf{z}$ ,  $\mathbf{r}_u$  and  $\mathbf{u}$ , respectively. We fix  $\mathbf{u}$  at  $\mathbf{u} = \mathbf{u}_1$  and express (68) as

$$\mathbf{D}_q^1(\mathbf{v}) = \mathbf{J}_z(\mathbf{v}) \mathbf{D}_\psi^1(\mathbf{v}) \mathbf{J}^1, \quad (69)$$

where  $\mathbf{D}_q^1(\mathbf{v}) := \mathbf{D}_q(\mathbf{v}, \mathbf{r}_q)$  at  $\mathbf{r}_q = \lambda(\mathbf{u}_1)$ ,  $\mathbf{D}_\psi^1(\mathbf{v}) := \mathbf{D}_\psi(\mathbf{v}, \mathbf{r}_u)$  at  $\mathbf{r}_u = \mathbf{h}_u^*(\mathbf{u}_1)$ , and  $\mathbf{J}^1 := \mathbf{J}_{r_u}(\mathbf{u}_1) \mathbf{J}_u(\mathbf{r}_q)$  at  $\mathbf{r}_q = \lambda(\mathbf{u}_1)$ . We can have a similar matrix expression as (69) at  $\mathbf{u} = \mathbf{u}_2$  as

$$\mathbf{D}_q^2(\mathbf{v}) = \mathbf{J}_z(\mathbf{v}) \mathbf{D}_\psi^2(\mathbf{v}) \mathbf{J}^2. \quad (70)$$

Here, we note that both  $\mathbf{J}^1 \in \mathbb{R}^{D_x \times D_x}$  and  $\mathbf{J}^2 \in \mathbb{R}^{D_x \times D_x}$  have at most rank  $D_x$  by Assumption (B'3).

Next, we confirm that  $\mathbf{J}^1$ ,  $\mathbf{J}^2$ ,  $\mathbf{D}_\psi^1(\mathbf{v})$  and  $\mathbf{D}_\psi^2(\mathbf{v})$  for all  $\mathbf{v}$  have rank  $D_x$ , and thus are invertible. To this end, we denote the  $i$ -th diagonals in  $\mathbf{D}_q^1(\mathbf{v})$  and  $\mathbf{D}_q^2(\mathbf{v})$  by

$$\alpha_i^1(v_i) := \left. \frac{\partial^2 q_i(v_i, r)}{\partial v_i \partial r} \right|_{r=\lambda_i(\mathbf{u}_1)} \quad \text{and} \quad \alpha_i^2(v_i) := \left. \frac{\partial^2 q_i(v_i, r)}{\partial v_i \partial r} \right|_{r=\lambda_i(\mathbf{u}_2)},$$

respectively. From Assumptions (B'4),  $\mathbf{D}_q^1(\mathbf{v})$  has nonzero diagonals  $\alpha_i^1(v_i)$  and rank  $D_x$  for all  $\mathbf{v}$ . Then, applying Lemma 11 to (69) ensures that  $\mathbf{D}_\psi^1(\mathbf{v}) \mathbf{J}^1$  have rank  $D_x$ . By applying Lemma 11

to  $D_\psi^1(\mathbf{v})\mathbf{J}^1$  again, it can be shown that both  $D_\psi^1(\mathbf{v})$  and  $\mathbf{J}^1$  also have rank  $D_x$ . Similarly, we can prove that both  $D_\psi^2(\mathbf{v})$  and  $\mathbf{J}^2$  have rank  $D_x$  under Assumption (B'4).

Finally, we show that  $\mathbf{J}_z(\mathbf{v})$  is the product of a permutation and diagonal matrices. Since  $D_\psi^1(\mathbf{v})$  and  $\mathbf{J}^1$  are invertible, from (69), we have

$$\mathbf{J}_z(\mathbf{v}) = D_q^1(\mathbf{v})[\mathbf{J}^1]^{-1}[D_\psi^1(\mathbf{v})]^{-1}. \quad (71)$$

Substituting (71) into (70), and multiplying  $\mathbf{J}^1$  and the inverses of  $D_q^1(\mathbf{v})$  and  $\mathbf{J}^2$  to both sides yields

$$\mathbf{J}^1 [D_q^1(\mathbf{v})]^{-1} D_q^2(\mathbf{v}) [\mathbf{J}^2]^{-1} = [D_\psi^1(\mathbf{v})]^{-1} D_\psi^2(\mathbf{v}). \quad (72)$$

Then, we compactly express all of diagonals in  $[D_q^1(\mathbf{v})]^{-1} D_q^2(\mathbf{v})$  as the following vector:

$$\boldsymbol{\alpha}(\mathbf{v}) := \left( \frac{\alpha_1^2(v_1)}{\alpha_1^1(v_1)}, \frac{\alpha_2^2(v_2)}{\alpha_2^1(v_2)}, \dots, \frac{\alpha_{D_x}^2(v_{D_x})}{\alpha_{D_x}^1(v_{D_x})} \right)^\top.$$

By Assumption (B'5),  $\boldsymbol{\alpha}(\mathbf{v}_1), \boldsymbol{\alpha}(\mathbf{v}_2), \dots, \boldsymbol{\alpha}(\mathbf{v}_{D_x})$  are linearly independent. Since  $\mathbf{J}^1$  and  $\mathbf{J}^2$  are proved to have rank  $D_x$  and thus constant full-rank matrices, applying Lemma 10 ensures that  $\mathbf{J}^1$  and  $\mathbf{J}^2$  are diagonal matrices multiplied by a permutation matrix. Thus, it follows from (71) that  $\mathbf{J}_z(\mathbf{v})$  is also a diagonal matrix multiplied by a permutation matrix. Thus, each  $v_i(\mathbf{z})(= s_i)$  corresponds to a single and distinct element in  $\mathbf{z}(= \mathbf{h}_x^*(\mathbf{x}))$ . The proof is completed.  $\blacksquare$

## D.2 Proof of Lemma 10

**Proof** Let us first denote the  $(l, m)$ -th elements of  $\mathbf{A}$  and  $\mathbf{B}$  by  $a_{lm}$  and  $b_{lm}$ . The  $(l, m)$ -th element in  $\mathbf{AD}(\mathbf{v})\mathbf{B}$  is given by

$$\sum_{i=1}^n d_i(\mathbf{v}) a_{li} b_{im}.$$

Since off-diagonal elements in  $\mathbf{AD}(\mathbf{v})\mathbf{B}$  are all zeros at  $\mathbf{v} = \mathbf{v}_j$ ,  $j = 1, \dots, n$ ,

$$\sum_{i=1}^n d_i(\mathbf{v}_j) a_{li} b_{im} = \mathbf{d}(\mathbf{v}_j)^\top \mathbf{c}^{lm} = 0 \quad (l \neq m), \quad (73)$$

where

$$\mathbf{c}^{lm} := (a_{l1}b_{1m}, a_{l2}b_{2m}, \dots, a_{ln}b_{nm})^\top.$$

Collecting  $\mathbf{d}(\mathbf{v})$  over the  $n$  points  $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$  based on (73) yields

$$\tilde{\mathbf{D}}^\top \mathbf{c}^{lm} = \mathbf{0}, \quad (74)$$

where

$$\tilde{\mathbf{D}} := [\mathbf{d}(\mathbf{v}_1), \mathbf{d}(\mathbf{v}_2), \dots, \mathbf{d}(\mathbf{v}_n)] \in \mathbb{R}^{n \times n}.$$

By Assumption (1),  $\tilde{\mathbf{D}}$  is of full-rank and invertible. Thus, from (74), we obtain  $\mathbf{c}^{lm} = \mathbf{0}$  for  $l, m = 1, \dots, n$  with  $l \neq m$ , indicating that

$$a_{li}b_{im} = 0 \quad (i, l, m = 1, \dots, n \text{ and } l \neq m). \quad (75)$$

Next, we show that both the  $i$ -th column vector in  $\mathbf{A}$  and the  $i$ -th row vector in  $\mathbf{B}$  have a single nonzero element, while the other elements are zeros. We first suppose that the  $i$ -th column vector in  $\mathbf{A}$  has at least two nonzeros elements such that  $a_{li} \neq 0$  and  $a_{l'i} \neq 0$  for  $l \neq l'$ .  $a_{li} \neq 0$  implies that  $b_{im} = 0$  except for  $m = l$ , but  $a_{l'i} \neq 0$  ensures  $b_{im} = 0$  at  $m = l$ . Thus, the  $i$ -th row vector in  $\mathbf{B}$  must be the zero vector. However, this contradicts to the assumption that  $\mathbf{B}$  is of full-rank. Therefore, the  $i$ -th column vector of  $\mathbf{A}$  must have the single nonzero element. Similarly, we can prove that the  $i$ -th row vector in  $\mathbf{B}$  also has a single nonzero element. Thus, by the full-rank assumption of  $\mathbf{A}$  and  $\mathbf{B}$ , all column and row vectors in  $\mathbf{A}$  and  $\mathbf{B}$  must have the single nonzero element at distinct positions. Thus,  $\mathbf{A}$  and  $\mathbf{B}$  are equal to diagonal matrices multiplied by a permutation matrix. ■

## Appendix E. Proof of Theorem 4

**Proof** Taking the gradient of both sides on (18) with respect to  $\mathbf{u}$  yields

$$\nabla_{\mathbf{u}} \log p(\mathbf{x}|\mathbf{u}) = \nabla_{\mathbf{u}} \psi^*(\mathbf{h}_{\mathbf{x}}^*(\mathbf{x}), \mathbf{h}_{\mathbf{u}}^*(\mathbf{u})) + \nabla_{\mathbf{u}} b^*(\mathbf{h}_{\mathbf{u}}^*(\mathbf{u})). \quad (76)$$

By the change of variables by  $\mathbf{x} = \mathbf{f}(\mathbf{s}, \mathbf{n})$  in the generative model (17), we re-express (76) as

$$\nabla_{\mathbf{u}} \log p(\mathbf{s}, \mathbf{n}|\mathbf{u}) = \nabla_{\mathbf{u}} \psi^*(\mathbf{h}_{\mathbf{x}}^*(\mathbf{f}(\mathbf{s}, \mathbf{n})), \mathbf{h}_{\mathbf{u}}^*(\mathbf{u})) + \nabla_{\mathbf{u}} b^*(\mathbf{h}_{\mathbf{u}}^*(\mathbf{u})), \quad (77)$$

where we note that the Jacobian term due to the change of variables is deleted by the differential operator  $\nabla_{\mathbf{u}}$ .

Next, with  $\mathbf{v} := \mathbf{h}_{\mathbf{x}}^*(\mathbf{f}(\mathbf{s}, \mathbf{n}))$ , we compute the gradient of the left-hand side on (77) with respect to  $\mathbf{n}$  as

$$\begin{aligned} \nabla_{\mathbf{n}} \nabla_{\mathbf{u}} \log p(\mathbf{s}, \mathbf{n}|\mathbf{u}) &= \nabla_{\mathbf{n}} [\nabla_{\mathbf{u}} \{\log p(\mathbf{s}|\mathbf{u}) + \log p(\mathbf{n})\}] \\ &= \nabla_{\mathbf{n}} [\nabla_{\mathbf{u}} \log p(\mathbf{s}|\mathbf{u})] \\ &= \mathbf{O}, \end{aligned} \quad (78)$$

where we used Assumption (C1) that  $\mathbf{s} \perp \mathbf{n}$ ,  $\mathbf{u} \perp \mathbf{n}$  and  $\mathbf{s} \not\perp \mathbf{u}$ . On the other hand, the gradient of the right-hand side on (77) is given by

$$\nabla_{\mathbf{n}} [\nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u}))] = [\nabla_{\mathbf{v}} \nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u}))] \mathbf{J}_{\mathbf{v}}(\mathbf{n}), \quad (79)$$

where  $\nabla_{\mathbf{v}} \nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u})) \in \mathbb{R}^{D_{\mathbf{u}} \times d_{\mathbf{x}}}$  and  $\mathbf{J}_{\mathbf{v}}(\mathbf{n}) := \nabla_{\mathbf{n}} \mathbf{v} \in \mathbb{R}^{d_{\mathbf{x}} \times (D_{\mathbf{x}} - d_{\mathbf{x}})}$  is the Jacobian matrix of  $\mathbf{v}$  at  $\mathbf{n}$ . We fix  $\mathbf{u}$  at  $\mathbf{u} = \mathbf{u}_1$  and equate (79) with (78) based on (77) as

$$\mathbf{M}(\mathbf{v}) \mathbf{J}_{\mathbf{v}}(\mathbf{n}) = \mathbf{O}, \quad (80)$$

where  $\mathbf{M}(\mathbf{v}) := \nabla_{\mathbf{v}} \nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u}))|_{\mathbf{u}=\mathbf{u}_1} \in \mathbb{R}^{D_{\mathbf{u}} \times d_{\mathbf{x}}}$ . By Assumption (C4), the rank of  $\mathbf{M}(\mathbf{v})$  is  $d_{\mathbf{x}}$ , and thus the inverse of  $\mathbf{M}(\mathbf{v})^{\top} \mathbf{M}(\mathbf{v})$  exists. Multiplying  $\mathbf{M}(\mathbf{v})^{\top}$  on the left and taking the inverse of  $\mathbf{M}(\mathbf{v})^{\top} \mathbf{M}(\mathbf{v})$  to (80) indicates  $\mathbf{J}_{\mathbf{v}}(\mathbf{n}) = \mathbf{O}$ . Recall that  $\mathbf{v} = \mathbf{h}_{\mathbf{x}}^* \circ \mathbf{f}(\mathbf{s}, \mathbf{n}) = \mathbf{h}_{\mathbf{x}}^*(\mathbf{x})$ . Since  $\mathbf{f}(\mathbf{s}, \mathbf{n})$  and  $\mathbf{h}_{\mathbf{x}}^*(\mathbf{x})$  are invertible and surjective respectively,  $\mathbf{J}_{\mathbf{v}}(\mathbf{n}) = \mathbf{O}$  ensures that  $\mathbf{v}$  is a (vector-valued) function of only  $\mathbf{s}$ . Thus, the proof is completed. ■

## Appendix F. Derivation of (19)

By Assumption (C1) that  $\mathbf{s} \perp \mathbf{n}$ ,  $\mathbf{u} \perp \mathbf{n}$  and  $\mathbf{s} \not\perp \mathbf{u}$ , (77) can be written as

$$\nabla_{\mathbf{u}} \log p(\mathbf{s}|\mathbf{u}) = \nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u})) + \nabla_{\mathbf{u}} b^*(\mathbf{h}_{\mathbf{u}}^*(\mathbf{u})),$$

where  $\mathbf{v} := \mathbf{h}_{\mathbf{x}}^*(\mathbf{f}(\mathbf{s}, \mathbf{n}))$ . Taking the gradient with respect to  $\mathbf{v}$  yields

$$\mathbf{J}_{\mathbf{v}}(\mathbf{s})[\nabla_{\mathbf{s}} \nabla_{\mathbf{u}} \log p(\mathbf{s}|\mathbf{u})] = \nabla_{\mathbf{v}} \nabla_{\mathbf{u}} \psi^*(\mathbf{v}, \mathbf{h}_{\mathbf{u}}^*(\mathbf{u})),$$

where  $\mathbf{J}_{\mathbf{v}}(\mathbf{s}) := \nabla_{\mathbf{s}} \mathbf{v}$  denotes the Jacobian of  $\mathbf{v}$  at  $\mathbf{s}$ . Thus, we obtain (19) by fixing  $\mathbf{u}$  at  $\mathbf{u} = \mathbf{u}_1$ .

## Appendix G. Proof of Proposition 5

**Proof** Firstly, we consider the contamination model 2. By definition, the estimator  $\boldsymbol{\theta}^*$  satisfies the following equation:

$$\nabla_{\boldsymbol{\theta}} \left\{ E_{\mathbf{x}\mathbf{u}}[r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})] - \log E_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} = \mathbf{0}, \quad (81)$$

while  $\boldsymbol{\theta}_{\epsilon}$  associated with the contaminated densities fulfills

$$\nabla_{\boldsymbol{\theta}} \left\{ \bar{E}_{\mathbf{x}\mathbf{u}}[r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})] - \log \bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}_{\epsilon}} = \mathbf{0}, \quad (82)$$

where  $\bar{E}_{\mathbf{x}\mathbf{u}}$  and  $\bar{E}_{\mathbf{x}\times\mathbf{u}}$  denote the expectations over  $\bar{p}(\mathbf{x}, \mathbf{u})$  and  $\bar{p}(\mathbf{x})\bar{p}(\mathbf{u})$ , respectively. Applying the Taylor series expansion of (82) around  $\boldsymbol{\theta}^*$  yields

$$\mathbf{0} = \nabla_{\boldsymbol{\theta}} \left\{ \bar{E}_{\mathbf{x}\mathbf{u}}[r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})] - \log \bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} + \bar{\mathbf{V}}_{\text{DV}}(\boldsymbol{\theta}_{\epsilon} - \boldsymbol{\theta}^*) + O(\|\boldsymbol{\theta}_{\epsilon} - \boldsymbol{\theta}^*\|^2). \quad (83)$$

where  $\bar{\mathbf{V}}_{\text{DV}}$  is the Hessian matrix at  $\boldsymbol{\theta} = \boldsymbol{\theta}^*$  and defined by

$$\bar{\mathbf{V}}_{\text{DV}} := \nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}} \left\{ \bar{E}_{\mathbf{x}\mathbf{u}}[r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})] - \log \bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*}.$$

Let us denote  $\nabla_{\boldsymbol{\theta}} r_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{u})$  by  $\mathbf{g}_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{u})$ . We recall that the contaminated densities in the contamination model 2 are defined as

$$\begin{aligned} \bar{p}(\mathbf{x}, \mathbf{u}) &= (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta_{\bar{\mathbf{x}}}(\mathbf{x})\delta_{\bar{\mathbf{u}}}(\mathbf{u}) \\ \bar{p}(\mathbf{x}) &= (1 - \epsilon)p(\mathbf{x}) + \epsilon\delta_{\bar{\mathbf{x}}}(\mathbf{x}) \\ \bar{p}(\mathbf{u}) &= (1 - \epsilon)p(\mathbf{u}) + \epsilon\delta_{\bar{\mathbf{u}}}(\mathbf{u}). \end{aligned}$$

Then, with sufficiently small  $\epsilon$ , the first term on the right-hand side of (83) is given by

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} \left\{ \bar{E}_{\mathbf{x}\mathbf{u}}[r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})] - \log \bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\boldsymbol{\theta}=\boldsymbol{\theta}^*} &= \bar{E}_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})] - \frac{\bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})]}{\bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})}]} \\ &\simeq \bar{E}_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})] - \bar{E}_{\mathbf{x}\times\mathbf{u}}[e^{r_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\boldsymbol{\theta}^*}(\mathbf{X}, \mathbf{U})], \end{aligned} \quad (84)$$

where  $\simeq$  denotes the equality up to terms of  $\epsilon^2$ , and we applied the following relation:

$$\bar{E}_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})}] \simeq E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})}] + \epsilon \left( E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})}] + E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})}] - 2E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})}] \right) = 1,$$

where we used

$$E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})}] = E_{\mathbf{X} \times \mathbf{U}} \left[ \frac{p(\mathbf{X}, \mathbf{U})}{p(\mathbf{X})p(\mathbf{U})} \right] = 1, \quad (85)$$

$$E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})}] = \int p(\mathbf{x}|\bar{\mathbf{u}})d\mathbf{x} = 1 \quad \text{and} \quad E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})}] = \int p(\mathbf{u}|\bar{\mathbf{x}})d\mathbf{u} = 1, \quad (86)$$

based on the assumption that  $r_{\theta^*}(\mathbf{x}, \mathbf{u}) = \log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})}$ . Then, we substitute

$$\bar{E}_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] = (1 - \epsilon)E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] + \epsilon \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}),$$

and

$$\begin{aligned} \bar{E}_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] &\simeq (1 - \epsilon)E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] + \epsilon \left( E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] \right. \\ &\quad \left. + E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})] - E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] \right), \end{aligned}$$

into (84), and have

$$\begin{aligned} &\nabla_{\theta} \left\{ \bar{E}_{\mathbf{XU}}[r_{\theta}(\mathbf{X}, \mathbf{U})] - \log \bar{E}_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\theta=\theta^*} \\ &\simeq (1 - \epsilon) \left( E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] \right) \\ &\quad + \epsilon \left( \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) + E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] - E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})] \right) \\ &= \epsilon \left( \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) + E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] - E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})] \right), \end{aligned} \quad (87)$$

where we applied  $E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] = E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})]$ , which is proved from (81) and (85) as follows:

$$\begin{aligned} E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] &= E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - \frac{E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})} \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})]}{E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta^*}(\mathbf{X}, \mathbf{U})}]} \\ &= \nabla_{\theta} \left\{ E_{\mathbf{XU}}[r_{\theta}(\mathbf{X}, \mathbf{U})] - \log E_{\mathbf{X} \times \mathbf{U}}[e^{r_{\theta}(\mathbf{X}, \mathbf{U})}] \right\} \Big|_{\theta=\theta^*} \\ &= \mathbf{0}. \end{aligned}$$

Substituting (87) into (83) yields

$$\begin{aligned} \mathbf{0} &\simeq \epsilon \left\{ \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) + E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{X}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] - E_{\mathbf{U}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})] \right\} \\ &\quad + \bar{\mathbf{V}}_{\text{DV}}(\theta_{\epsilon} - \theta^*) + O(\|\theta_{\epsilon} - \theta^*\|^2) \end{aligned} \quad (88)$$

In the limit of  $\epsilon \rightarrow 0$ ,  $O(\|\theta_{\epsilon} - \theta^*\|^2)$  in (88) quickly converges to zero and can be negligible, and  $\bar{\mathbf{V}}_{\text{DV}}$  approaches

$$-E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U}) \mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})^{\top}] + E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] E_{\mathbf{XU}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})]^{\top} =: \mathbf{V}_{\text{DV}},$$

where we applied (85). By dividing both sides on (88) by  $\epsilon$  and taking the limit of  $\epsilon \rightarrow 0$ , we obtain

$$\begin{aligned} & \mathbf{V}_{\text{DV}} \text{IF}_{\text{DV}}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) \\ &= \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) + E_{\mathbf{x}\mathbf{u}}[\mathbf{g}_{\theta^*}(\mathbf{X}, \mathbf{U})] - E_{\mathbf{x}}[e^{r_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})} \mathbf{g}_{\theta^*}(\mathbf{X}, \bar{\mathbf{u}})] - E_{\mathbf{u}}[e^{r_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})} \mathbf{g}_{\theta^*}(\bar{\mathbf{x}}, \mathbf{U})]. \end{aligned}$$

Thus, applying the inverse of  $\mathbf{V}_{\text{DV}}$  yields (31).

For the contamination model 1, we recall that the contaminated densities are given as follows:

$$\begin{aligned} \bar{p}(\mathbf{x}, \mathbf{u}) &= (1 - \epsilon)p(\mathbf{x}, \mathbf{u}) + \epsilon\delta_{\bar{\mathbf{u}}}(\mathbf{u})p(\mathbf{x}) \\ \bar{p}(\mathbf{x}) &= p(\mathbf{x}) \\ \bar{p}(\mathbf{u}) &= (1 - \epsilon)p(\mathbf{u}) + \epsilon\delta_{\bar{\mathbf{u}}}(\mathbf{u}). \end{aligned}$$

By following the same line of the proof, we reach the influence function (29) for the contamination model 1.  $\blacksquare$

## Appendix H. Proof of Proposition 8

### H.1 Main proof

We employ the following lemma, which is proved in Appendix H.2:

**Lemma 12.** *Constant  $\nu_1$  is sufficiently small under Assumptions (D1-2).*

Thus, we next focus on proving that (38) holds when  $\nu_1$  is sufficiently small. Denoting  $e^{r(\mathbf{x}, \mathbf{u})}$  by  $\varphi(\mathbf{x}, \mathbf{u})$ , for the contamination model 1, the expectations over  $\bar{p}(\mathbf{x}, \mathbf{u})$  and  $\bar{p}(\mathbf{x})\bar{p}(\mathbf{u})$  can be expressed as

$$\begin{aligned} \bar{E}_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{\varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] &= (1 - \epsilon) E_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{\varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] \\ &\quad + \epsilon \iint p(\mathbf{x}) \delta(\mathbf{u}|\mathbf{x}) \left( \frac{\varphi(\mathbf{x}, \mathbf{u})^{\gamma+1}}{1 + \varphi(\mathbf{x}, \mathbf{u})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x} d\mathbf{u} \\ \bar{E}_{\mathbf{x} \times \mathbf{u}} \left[ \left( \frac{1}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] &= (1 - \epsilon) E_{\mathbf{x} \times \mathbf{u}} \left[ \left( \frac{1}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] \\ &\quad + \epsilon \iint p(\mathbf{x}) \delta(\mathbf{u}) \left( \frac{1}{1 + \varphi(\mathbf{x}, \mathbf{u})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x} d\mathbf{u}. \end{aligned}$$

Substituting these expectations into  $\bar{J}_\gamma(\varphi)$  yields

$$\begin{aligned} \bar{J}_\gamma(\varphi) &= -\frac{1}{\gamma} \log \left[ E_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{\varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] + E_{\mathbf{x} \times \mathbf{u}} \left[ \left( \frac{1}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] + \frac{\epsilon}{1 - \epsilon} \nu_1 \right] \\ &\quad - \frac{1}{\gamma} \log(1 - \epsilon). \end{aligned}$$

By applying  $\log(z + \nu_1) = \log(z) + O(\nu_1)$  with a sufficiently small  $\nu_1$ ,

$$\begin{aligned} \bar{J}_\gamma(\varphi) &= -\frac{1}{\gamma} \log \left[ E_{\mathbf{x}\mathbf{u}} \left[ \left( \frac{\varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] + E_{\mathbf{x}\times\mathbf{u}} \left[ \left( \frac{1}{1 + \varphi(\mathbf{X}, \mathbf{U})^{\gamma+1}} \right)^{\frac{\gamma}{\gamma+1}} \right] \right] \\ &\quad - \frac{1}{\gamma} \log(1 - \epsilon) + O\left(\frac{\epsilon}{1 - \epsilon} \nu_1\right) \\ &= J_\gamma(\varphi) - \frac{1}{\gamma} \log(1 - \epsilon) + O\left(\frac{\epsilon}{1 - \epsilon} \nu_1\right). \end{aligned}$$

This completes the proof.

## H.2 Proof of Lemma 12

We recall that  $\nu_1$  is defined by

$$\nu_1 = \underbrace{\iint p(\mathbf{x}) \delta(\mathbf{u}|\mathbf{x}) \left( \frac{e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u}}_{\text{(A)}} + \underbrace{\iint p(\mathbf{x}) \delta(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u}}_{\text{(B)}}.$$

Here, we complete the proof by showing below that two terms (A) and (B) on the right-hand side are sufficiently small under Assumptions (D1-2).

**Term (A):** We apply the Cauchy–Schwartz inequality to Term (A) as

$$\begin{aligned} &\iint \delta(\mathbf{u}|\mathbf{x}) p(\mathbf{x}) \left( \frac{e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \\ &\leq C_A \left( \iint \delta(\mathbf{u}|\mathbf{x}) p(\mathbf{x}) e^{\gamma\{\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x}))\}} d\mathbf{x}d\mathbf{u} \right)^{\frac{1}{2}}, \end{aligned} \quad (89)$$

where

$$C_A := \left( \iint \delta(\mathbf{u}|\mathbf{x}) p(\mathbf{x}) \left( \frac{e^{(\gamma+1)\{\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x}))\}}/2}{e^{-(\gamma+1)b(\mathbf{h}_u(\mathbf{u}))} + e^{(\gamma+1)\{\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x}))\}}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \right)^{\frac{1}{2}}.$$

By Assumption (D1), the constant  $C_A$  is finite. Thus, Assumption (D2) ensures that Term (A) is sufficiently small.

**Term (B):** Applying the Cauchy–Schwartz inequality yields the upper bound of Term (B) as

$$\iint \delta(\mathbf{u}) p(\mathbf{x}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \leq C_B \left( \int \delta(\mathbf{u}) e^{-\gamma b(\mathbf{h}_u(\mathbf{u}))} d\mathbf{u} \right)^{\frac{1}{2}}, \quad (90)$$

where

$$C_B := \left( \iint \delta(\mathbf{u}) p(\mathbf{x}) \left( \frac{e^{-(\gamma+1)b(\mathbf{h}_u(\mathbf{u}))/2}}{e^{-(\gamma+1)b(\mathbf{h}_u(\mathbf{u}))} + e^{(\gamma+1)\{\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u})) + a(\mathbf{h}_x(\mathbf{x}))\}}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \right)^{\frac{1}{2}}.$$

By (90) and Assumption (D2), Term (B) is also sufficiently small.

## Appendix I. Semistrong robustness for the contamination model 2

For the contamination model 2, the following proposition is established:

**Proposition 13.** *Let us define a constant  $\nu_2$  as*

$$\begin{aligned} \nu_2 := & \iint \delta(\mathbf{x}, \mathbf{u}) \left( \frac{e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} + (1 - \epsilon) \iint p(\mathbf{x})\delta(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} \\ & + (1 - \epsilon) \iint \delta(\mathbf{x})p(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} + \epsilon \iint \delta(\mathbf{x})\delta(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u}. \end{aligned}$$

We make the following assumptions:

(D'1) Assume that

$$\begin{aligned} & \iint \delta(\mathbf{x}, \mathbf{u}) \left( \frac{e^{(\gamma+1)\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))/2}}{e^{-(\gamma+1)\{a(\mathbf{h}_x(\mathbf{x})+b(\mathbf{h}_u(\mathbf{u}))\}} + e^{(\gamma+1)\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} < \infty, \\ & \iint \{p(\mathbf{x})\delta(\mathbf{u}) + \delta(\mathbf{x})p(\mathbf{u}) + \delta(\mathbf{x})\delta(\mathbf{u})\} \\ & \quad \times \left( \frac{e^{-(\gamma+1)\{a(\mathbf{h}_x(\mathbf{x})+b(\mathbf{h}_u(\mathbf{u}))\}}/2}}{e^{-(\gamma+1)\{a(\mathbf{h}_x(\mathbf{x})+b(\mathbf{h}_u(\mathbf{u}))\}} + e^{(\gamma+1)\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))}} \right)^{\frac{2\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u} < \infty \end{aligned}$$

(D'2) The following integrals are sufficiently small.

$$\iint \delta(\mathbf{x}, \mathbf{u}) e^{\gamma\psi(\mathbf{h}_x(\mathbf{x}), \mathbf{h}_u(\mathbf{u}))} d\mathbf{u}d\mathbf{x}, \quad \int \delta(\mathbf{x}) e^{-\gamma a(\mathbf{h}_x(\mathbf{x}))} d\mathbf{u}d\mathbf{x} \quad \text{and} \quad \int \delta(\mathbf{u}) e^{-\gamma b(\mathbf{h}_u(\mathbf{u}))} d\mathbf{u}.$$

Then, it holds that  $\nu_2$  is sufficiently small and

$$\bar{J}_\gamma(r) = -\frac{1}{\gamma} \log [\exp(-\gamma J_\gamma(r)) - \epsilon I(r)] - \frac{1}{\gamma} \log(1 - \epsilon) + O\left(\frac{\epsilon}{1 - \epsilon} \nu_2\right), \quad (91)$$

where

$$I(r) := \iint p(\mathbf{x})p(\mathbf{u}) \left( \frac{1}{1 + e^{(\gamma+1)r(\mathbf{x}, \mathbf{u})}} \right)^{\frac{\gamma}{\gamma+1}} d\mathbf{x}d\mathbf{u}.$$

The proof is essentially the same as Proposition 8, and thus is omitted. Proposition 13 indicates when  $\nu_2$  is sufficiently small, minimization of  $\bar{J}_\gamma(r)$  is almost equal to minimization of

$$-\frac{1}{\gamma} \log \left[ e^{-\gamma J_\gamma(r)} - \epsilon I(r) \right]. \quad (92)$$

Unlike the contamination model 1 (Proposition 8), there exists an extra term  $\epsilon I(r)$  inside the logarithm in (92). However, this extra term may not cause any bias in terms of representation learning because the minimizer of (92) is given by

$$\log \frac{p(\mathbf{x}, \mathbf{u})}{p(\mathbf{x})p(\mathbf{u})} - \log(1 - \epsilon). \quad (93)$$

Since our goal of representation learning is to estimate the log-density ratio up to a constant, the minimizer (93) means that robust representation learning would be possible even for the contamination model 2 under the condition that  $\nu_2$  is sufficiently small. In addition, again, the contamination ratio  $\epsilon$  is not necessarily assumed to be small.

Assumptions (D'1-2) have the same implications as Assumptions (D1-2) in Proposition 8 (See Section 4.3.2), and thus would reflect the typical contamination of outliers. Thus, the condition that  $\nu_2$  is sufficiently small would be fairly reasonable. Indeed, Section 5 experimentally demonstrates that our method for representation learning is very robust against outliers.

## References

- L. B. Almeida. MISEP—linear and nonlinear ICA based on mutual information. *Journal of Machine Learning Research*, 4:1297–1318, 2003.
- S. Amari. *Information Geometry and Its Applications*. Springer, 2016.
- G. Andrew, R. Arora, J. Bilmes, and K. Livescu. Deep canonical correlation analysis. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1247–1255. PMLR, 2013.
- R. Arandjelovic and A. Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision (CVPR)*, pages 609–617, 2017.
- J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- P. Bachman, R. D. Hjelm, and W. Buchwalter. Learning representations by maximizing mutual information across views. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 15535–15545, 2019.
- A. Banerjee. On Bayesian bounds. In *Proceedings of the 23rd International Conference on Machine Learning (ICML)*, pages 81–88, 2006.
- H. B. Barlow. Possible principles underlying the transformations of sensory messages. *Sensory Communicatio*, 1961.
- M. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, D. Hjelm, and A. Courville. Mutual information neural estimation. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 530–539, 2018.
- A. J. Bell and T. J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural Computation*, 7(6):1129–1159, 1995.
- Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- T. Blaschke, T. Zito, and L. Wiskott. Independent slow feature analysis and nonlinear blind source separation. *Neural Computation*, 19(4):994–1021, 2007.

- D. A. Blythe, P. Von Bunau, F. C. Meinecke, and K.-R. Müller. Feature extraction for change-point detection using stationary subspace analysis. *IEEE Transactions on Neural Networks and Learning Systems*, 23(4):631–643, 2012.
- M. Caron, P. Bojanowski, A. Joulin, and M. Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 132–149, 2018.
- P. Chen, H. Hung, O. Komori, S.-Y. Huang, and S. Eguchi. Robust independent component analysis via minimum  $\gamma$ -divergence estimation. *IEEE Journal of Selected Topics in Signal Processing*, 7(4):614–624, 2013.
- X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2172–2180, 2016.
- A. Cichocki and S. Amari. Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities. *Entropy*, 12(6):1532–1568, 2010.
- P. Comon. Independent component analysis, a new concept? *Signal Processing*, 36(3):287–314, 1994.
- R. D. Cook. *Regression Graphics: Ideas for Studying Regressions Through Graphics*. John Wiley & Sons, 1998.
- J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4171–4186, 2019.
- C. Doersch, A. Gupta, and A. A. Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 1422–1430, 2015.
- G. Dornhege, J. d. R. Millán, T. Hinterberger, D. McFarland, and K.-R. Müller. *Toward brain-computer interfacing*. MIT press Cambridge, 2007.
- A. Dosovitskiy, J. T. Springenberg, M. Riedmiller, and T. Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 766–774, 2014.
- H. Fujisawa and S. Eguchi. Robust parameter estimation with a small bias against heavy contamination. *Journal of Multivariate Analysis*, 99(9):2053–2081, 2008.
- K. Fukumizu, F. R. Bach, and M. I. Jordan. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research*, 5:73–99, 2004.
- I. Goodfellow, D. Warde-Farley, M. Mirza, A. Courville, and Y. Bengio. Maxout networks. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, Proceedings of Machine Learning Research, pages 1319–1327. PMLR, 2013.

- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680, 2014.
- L. Gresele, P. Rubenstein, A. Mehrjou, F. Locatello, and B. Schölkopf. The incomplete Rosetta Stone problem: Identifiability results for multi-view nonlinear ICA. In *Proceedings of the 35th International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 296–313, 2019.
- M. Gutmann and A. Hyvärinen. Noise-contrastive estimation of unnormalized statistical models, with applications to natural image statistics. *Journal of Machine Learning Research*, 13:307–361, 2012.
- F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, and W. A. Stahel. *Robust statistics: the approach based on influence functions*. John Wiley & Sons, 2011.
- D. A. Harville. *Matrix Algebra From a Statistician’s Perspective*. Springer Science & Business Media, 2006.
- I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner.  $\beta$ -VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations (ICLR)*, 2017.
- R. D. Hjelm, A. Fedorov, S. Lavoie-Marchildon, K. Grewal, P. Bachman, A. Trischler, and Y. Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations (ICLR)*, 2019.
- K. Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2): 251–257, 1991.
- H. Hung, Z.-Y. Jou, and S.-Y. Huang. Robust mislabel logistic regression without modeling mislabel probabilities. *Biometrics*, 74(1):145–154, 2018.
- 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 H. Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ICA. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3765–3773, 2016.
- A. Hyvärinen and H. Morioka. Nonlinear ICA of temporally dependent stationary sources. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 460–469. PMLR, 2017.
- A. Hyvärinen and P. Pajunen. Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3):429–439, 1999.
- A. Hyvärinen, H. Sasaki, and R. E. Turner. Nonlinear ICA using auxiliary variables and generalized contrastive learning. In *Proceedings of the 22th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 89, pages 859–868, 2019.

- L. Jing and Y. Tian. Self-supervised visual feature learning with deep neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):4037–4058, 2020.
- T. Kanamori, T. Suzuki, and M. Sugiyama. Theoretical analysis of density ratio estimation. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 93(4): 787–798, 2010.
- I. Khemakhem, D. P. Kingma, R. P. Monti, and A. Hyvärinen. Variational autoencoders and non-linear ICA: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108, pages 2207–2217. PMLR, 2020.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, pages 1–15, 2015.
- D. P. Kingma and M. Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations (ICLR)*, 2014.
- H.-Y. Lee, J.-B. Huang, M. Singh, and M.-H. Yang. Unsupervised representation learning by sorting sequences. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 667–676, Oct 2017.
- K. Li. Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86(414):316–327, 1991.
- Y. Li, M. Yang, and Z. Zhang. A survey of multi-view representation learning. *IEEE Transactions on Knowledge and Data Engineering*, 31(10):1863–1883, 2018.
- R. Linsker. An application of the principle of maximum information preservation to linear systems. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 186–194, 1989.
- X. Liu, F. Zhang, Z. Hou, Z. Wang, L. Mian, J. Zhang, and J. Tang. Self-supervised learning: Generative or contrastive. *arXiv preprint arXiv:2006.08218*, 2020.
- F. Locatello, S. Bauer, M. Lucic, G. Raetsch, S. Gelly, B. Schölkopf, and O. Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97 of *Proceedings of Machine Learning Research*, pages 4114–4124. PMLR, 2019.
- D. McAllester and K. Stratos. Formal limitations on the measurement of mutual information. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108 of *Proceedings of Machine Learning Research*, pages 875–884, 2020.
- I. Misra, C. L. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 527–544. Springer, 2016.
- R. P. Monti, K. Zhang, and A. Hyvärinen. Causal discovery with general non-linear relationships using non-linear ICA. In *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 186–195, 2020.

- X. Nguyen, M. J. Wainwright, and M. I. Jordan. Estimating divergence functionals and the likelihood ratio by penalized convex risk minimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1089–1096, 2008.
- M. Noroozi and P. Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European Conference on Computer Vision (ECCV)*, pages 69–84. Springer, 2016.
- M. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2227–2237, 2018.
- R. A. Poldrack. The future of fMRI in cognitive neuroscience. *Neuroimage*, 62(2):1216–1220, 2012.
- B. Poole, S. Ozair, A. van den Oord, A. Alemi, and G. Tucker. On variational bounds of mutual information. In *Proceedings of the 36th International conference on Machine learning (ICML)*, pages 5171–5180, 2019.
- R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th International Conference on Machine Learning (ICML)*, pages 759–766, 2007.
- A. Ruderman, M. D. Reid, D. García-García, and J. Petterson. Tighter variational representations of  $f$ -divergences via restriction to probability measures. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pages 1155–1162, 2012.
- H. Sasaki, T. Takenouchi, R. Monti, and A. Hyvärinen. Robust contrastive learning and nonlinear ICA in the presence of outliers. In *Proceedings of the International Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 124, pages 659–668. PMLR, 2020.
- H. Sprekeler, T. Zito, and L. Wiskott. An extension of slow feature analysis for nonlinear blind source separation. *Journal of Machine Learning Research*, 15:921–947, 2014.
- M. Sugiyama, S. Nakajima, H. Kashima, P. V. Buenau, and M. Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1433–1440, 2008.
- M. Sugiyama, T. Suzuki, and T. Kanamori. *Density Ratio Estimation in Machine Learning*. Cambridge University Press, 2012.
- C. Sun, F. Baradel, K. Murphy, and C. Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019.
- Y. Tan, J. Wang, and J. M. Zurada. Nonlinear blind source separation using a radial basis function network. *IEEE Transactions on Neural Networks*, 12(1):124–134, 2001.
- T. Teshima, I. Sato, and M. Sugiyama. Few-shot domain adaptation by causal mechanism transfer. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, volume 119, pages 9458–9469, 2020.

- Y. Tian, D. Krishnan, and P. Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019.
- M. Tschannen, J. Djolonga, P. K. Rubenstein, S. Gelly, and M. Lucic. On mutual information maximization for representation learning. In *Proceedings of the 7th International Conference on Learning Representations (ICLR)*, 2019.
- Y. Tsuboi, H. Kashima, S. Hido, S. Bickel, and M. Sugiyama. Direct density ratio estimation for large-scale covariate shift adaptation. *Journal of Information Processing*, 17:138–155, 2009.
- H.-Y. Tung, H.-W. Tung, E. Yumer, and K. Fragkiadaki. Self-supervised learning of motion capture. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pages 5236–5246, 2017.
- A. van den Oord, Y. Li, and O. Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research*, 11:3371–3408, 2010.
- X. Wang and A. Gupta. Unsupervised learning of visual representations using videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2794–2802, 2015.
- Z. Wang, E. Simoncelli, and A. Bovik. Multiscale structural similarity for image quality assessment. In *Proceedings of the 37th Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402, 2003.
- L. Wasserman. *All of nonparametric statistics*. Springer, 2006.
- P. Wu and K. Fukumizu. Causal mosaic: Cause-effect inference via nonlinear ICA and ensemble method. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 108, pages 1157–1167, 2020.