

---

# Learning Bijective Feature Maps for Linear ICA

---

Alexander Camuto<sup>\*,1,3</sup>

Matthew Willetts<sup>\*,1,3</sup>

Brooks Paige<sup>2,3</sup>

Chris Holmes<sup>1,3</sup>

Stephen Roberts<sup>1,3</sup>

<sup>1</sup>University of Oxford

<sup>2</sup>University College London

<sup>3</sup>Alan Turing Institute

## Abstract

Separating high-dimensional data like images into independent latent factors, i.e independent component analysis (ICA), remains an open research problem. As we show, existing probabilistic deep generative models (DGMs), which are tailor-made for image data, underperform on non-linear ICA tasks. To address this, we propose a DGM which combines bijective feature maps with a linear ICA model to learn interpretable latent structures for high-dimensional data. Given the complexities of jointly training such a hybrid model, we introduce novel theory that constrains linear ICA to lie close to the manifold of orthogonal rectangular matrices, the Stiefel manifold. By doing so we create models that converge quickly, are easy to train, and achieve better unsupervised latent factor discovery than flow-based models, linear ICA, and Variational Autoencoders on images.

## 1 Introduction

In linear Independent Component Analysis (ICA), data is modelled as having been created from a linear mixing of independent latent *sources* (Cardoso [1989a,b], 1997; Comon [1994]). The canonical problem is blind source separation; the aim is to estimate the original sources of a mixed set of signals by learning an *unmixing* matrix, which when multiplied with data recovers the values of these sources. While linear ICA is a powerful approach to unmix signals like sound (Everson & Roberts, 2001), it has not been as effectively developed for learning compact representations of high-dimensional data like images, where assuming linearity is limiting. Non-linear ICA methods, which assume non-linear mixing of latents, offer better performance on such data.

In particular, flow-based models have been proposed as a non-linear approach to *square* ICA, where we assume the dimensionality of our latent source space is the same as that of our data (Deco & Brauer [1995]; Dinh et al. [2015]). Flows parameterise a bijective mapping between data and a feature space of the same dimension and can be trained via maximum likelihood for a chosen base distribution in that space. While these are powerful generative models, for image data one typically wants fewer latent variables than the number of pixels in an image. In such situations, we wish to learn a non-square (dimensionality-reducing) ICA representation.

In this work, we highlight the fact that existing probabilistic deep generative models (DGMs), in particular Variational Autoencoders (VAEs), underperform on non-linear ICA tasks. As such there is a real need for a probabilistic DGM that can perform these tasks. To address this we propose a novel methodology for performing non-square non-linear ICA using a model, termed *Bijecta*, with two jointly trained parts: a highly-constrained non-square linear ICA model, operating on a feature space output by a bijective flow. The bijective flow is tasked with learning a representation for which linear ICA is a good model. It is as if we are *learning the data* for our ICA model.

We find that such a model fails to converge when trained naively with no constraints. To ensure convergence, we introduce novel theory for the parameterisation of decorrelating, non-square ICA matrices that lie close to the Stiefel manifold (Stiefel, 1935), the space of orthonormal rectangular matrices. We use this result to introduce a novel non-square linear ICA model that uses Johnson-Lindenstrauss projections (a family of randomly generated matrices). Using these projections, *Bijecta* successfully induces dimensionality reduction in flow-based models and scales non-square non-linear ICA methods to high-dimensional image data. Further we show that it is better able to learn independent latent factors than each of its constituent components in isolation and than VAEs. For a preliminary demonstration of the inability of VAEs and the ability of *Bijecta* to discover ICA sources see Fig 1.

---

\*Equal Contribution.

Proceedings of the 24<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).




Figure 1: Here we take a dSprites heart and, using a randomly sampled affine transformation, move it around a black background (a). The underlying sources of the dataset are *affine* transformations of the heart. In (b-c) images in the center correspond to the origin of the learnt source space. Images on either side correspond to linearly increasing values along one of the learnt latent sources whilst the other source remains fixed. Bijecta (c) has learned affine transformations as sources (white diagonals), whereas a VAE (with ICA-appropriate prior) (b) has learned non linear transforms (white curves). The VAE has not discovered the underlying latent sources.

## 2 Background

### 2.1 Independent Component Analysis

The goal of ICA is to learn a set of statistically independent sources that ‘explain’ our data. ICA is a highly diverse modelling paradigm with numerous variants: learning a mapping vs learning a model, linear vs non-linear, different loss functions, different generative models, and a wide array of methods of inference (Cardoso 1989a, Mackay 1996, Lee et al. 2000).

Here, we specify a generative model and find point-wise maximum likelihood estimates of model parameters in the manner of (Mackay, 1996; Cardoso, 1997). Concretely, we have a model with latent sources  $\mathbf{s} \in \mathcal{S} = \mathbb{R}^{d_s}$  generating data  $\mathbf{x} \in \mathcal{X} = \mathbb{R}^{d_x}$ , with  $d_s \leq d_x$ . The linear ICA generative model factorises as

$$p(\mathbf{x}, \mathbf{s}) = p(\mathbf{x}|\mathbf{s})p(\mathbf{s}), \quad p(\mathbf{s}) = \prod_{i=1}^{d_s} p(s_i),$$

where  $p(\mathbf{s})$  is a set of independent distributions appropriate for ICA. In linear ICA, where all mappings are simple matrix multiplications, the sources *cannot* be Gaussian distributions. Recall that we are mixing our sources to generate our data: A linear mixing of Gaussian random variables is itself Gaussian, so unmixing is impossible (Lawrence & Bishop, 2000). To be able to unmix, to break this symmetry, we can choose any heavy-tailed or light-tailed non-Gaussian distribution as our prior  $p(\mathbf{s})$  that gives us axis alignment and independence between sources. A common choice is the family of generalised Gaussian distributions,

$$p(s_i) = \text{GG}(s_i|\mu, \alpha, \rho)$$

$$p(s_i) = \frac{\rho}{2\alpha\Gamma(1/\rho)} \exp\left[\left(-\frac{|s_i - \mu|}{\alpha}\right)^\rho\right] \quad (1)$$

with mean  $\mu$ , scale  $\alpha$  and shape  $\rho$ . For  $\rho = 2$  we recover the Normal distribution, and for  $\rho = 1$  we have the (heavy-tailed) Laplace. As  $\rho \rightarrow \infty$  the distribution becomes increasingly sub-Gaussian, tending to a uniform distribution. As such, the generalised Gaussian is a flexible framework for specifying ICA-appropriate

priors as it allows for the specification of a sub or super Gaussian distribution by way of a single parameter:  $\rho$ .

### 2.2 Manifolds for the unmixing matrix $\mathbf{A}^+$

In linear ICA we want to find the linear mapping  $\mathbf{A}^+$  resulting in *maximally independent* sources. This is more onerous than merely finding decorrelated sources, as found by principal component analysis (PCA).

When learning a linear ICA model we typically have the mixing matrix  $\mathbf{A}$  as the (pseudo)inverse of the unmixing matrix  $\mathbf{A}^+$  and focus on the properties of  $\mathbf{A}^+$  to improve convergence.  $\mathbf{A}^+$  linearly maps from the data-space  $\mathcal{X}$  to the source space  $\mathcal{S}$ . It can be decomposed into two linear operations. First we *whiten* the data such that each component has unit variance and these components are mutually uncorrelated. We then apply an orthogonal transformation and a scaling operation (Hyvärinen et al. 2001, §6.34) to ‘rotate’ the whitened data into a set of coordinates where the sources are independent *and* decorrelated. Whitening on its own is not sufficient for ICA — two sources can be uncorrelated *and* dependent (see Appendix A).

Thus we can write the linear ICA unmixing matrix as

$$\mathbf{A}^+ = \Phi \mathbf{R} \mathbf{W} \quad (2)$$

where  $\mathbf{W} \in \mathbb{R}^{d_s \times d_x}$  is our whitening matrix,  $\mathbf{R} \in \mathbb{R}^{d_s \times d_s}$  is an orthogonal matrix and  $\Phi \in \mathbb{R}^{d_s}$  is a diagonal matrix. Matrices that factorise this way are known as the *decorrelating matrices* (Everson & Roberts, 1999): members of this family decorrelate through  $\mathbf{W}$ , and  $\Phi \mathbf{R}$  ensures that sources are statistically independent, not merely uncorrelated. The optimal ICA unmixing matrix is the decorrelating matrix that decorrelates *and* gives independence.

### 2.3 Flows

Flows are models that stack numerous invertible changes of variables. One specifies a simple base distribution and learns a sequence of (invertible) transforms to construct new distributions that assign high probability to observed data. Given a variable  $\mathbf{z} \in \mathcal{Z} = \mathbb{R}^{d_x}$ ,

we specify the distribution over data  $\mathbf{x}$  as

$$p(\mathbf{x}) = p(\mathbf{z}) \left| \det \frac{\partial f^{-1}}{\partial \mathbf{z}} \right|, \quad (3)$$

where  $f$  is a bijection from  $\mathcal{Z} \rightarrow \mathcal{X}$ , ie  $\mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_x}$ , and  $p(\mathbf{z})$  is the base distribution over the latent  $\mathbf{z}$  (Rezende & Mohamed, 2015; Papamakarios et al. 2019).

For more flexible distributions for  $\mathbf{x}$ , we specify  $\mathbf{x}$  through a series of composed functions, from our simple initial  $p$  into a more complex multi-modal distribution; for example for a series of  $K + 1$  mappings,  $\mathbf{z} = f_K \circ \dots \circ f_0(\mathbf{x})$ . By the properties of determinants under function composition

$$p(\mathbf{x}) = p(\mathbf{z}_K) \prod_{i=0}^K \left| \det \frac{\partial f_i^{-1}}{\partial \mathbf{z}_{i+1}} \right|, \quad (4)$$

where  $\mathbf{z}_{i+1}$  denotes the variable resulting from the transformation  $f_i(\mathbf{z}_i)$ ,  $p(\mathbf{z}_K)$  defines a density on the  $K^{\text{th}}$ , and the bottom most variable is our data ( $\mathbf{z}_0 = \mathbf{x}$ ).

Computing the determinant of the Jacobian ( $\det \frac{\partial f^{-1}}{\partial \mathbf{z}}$ ) in Eq. (3) can be prohibitively costly, especially when composing multiple functions as in Eq. (4). To address this, flows use *coupling layers* that enforce a lower triangular Jacobian such that the determinant of the Jacobian is simply the product of its diagonal elements. We use recently proposed coupling layers based on rational quadratic splines (RQS) to enforce this lower triangular structure (Durkan et al. 2019). They form highly flexible flows that typically require fewer composed mappings to achieve good performance relative to other coupling layers. See Appendix H for details.

### 3 Non-Square ICA using Flows

Variational Autoencoders seem like a natural fit for learning a compressed set of statistically independent latent variables (Kingma & Welling 2014; Rezende et al. 2014). It seems natural to train a VAE with an appropriate non-Gaussian prior, and expect that it would learn an appropriate ICA model. However, this is not the case. In Khemakhem et al. (2020) some experiments suggest that VAEs with ICA-appropriate priors are unsuited to performing non-linear ICA. In our experiments (§5) we further verify this line of inquiry and show that VAEs struggle to match their aggregate posteriors to non-Gaussian priors and thus are unable to discover independent latent sources.

Though source separation can be achieved by ‘disentangling’ methods such as the  $\beta$ -VAE (Higgins et al. 2017) and  $\beta$ -TCVAE (Chen et al. 2018), these methods require post-hoc penalisation of certain terms of the VAE objective, at times inducing improper priors (in the  $\beta$ -TCVAE in particular (Mathieu et al. 2019)). Further, precise tuning of this penalisation, a form of

soft supervision, is key to getting appropriate representations (Rolinek et al. 2019; Locatello et al. 2019). Stühmer et al. (2019) obtains a variety of non-linear ICA using VAEs with sets of Generalised Gaussian priors, but even then  $\beta$  penalisation is required to obtain ‘disentangled’ representations.

As such there is a need for probabilistic DGMs that can separate sources without added hyperparameter tuning and that can do so by matching ICA-appropriate priors. Our solution combines linear ICA with a dimensionality-preserving invertible flow  $f_\theta$ . The flow acts between our data space of dimensionality and the representation fed to the linear ICA generative model; learning a representation that is well fit by the simple, linear ICA model. As we demonstrate in experiments (§5), this hybrid model, which we call Bijecta, succeeds where VAEs fail: it can match non-Gaussian priors and is able to discover independent latent sources on image datasets.

#### 3.1 A Linear ICA base distribution for flows

Our aim here is to develop a non-square ICA method that is both end-to-end differentiable *and* computationally efficient, such that it can be trained jointly with a flow via stochastic gradient descent. We begin by choosing our base ICA source distribution to be a set of independent generalised Gaussian distributions, Eq (1) with  $\mu = 0$ ,  $\alpha = 1$  and  $\rho$  varying per experiment; and the ICA model’s likelihood to be a Gaussian.

$$p(s_i) = \text{GG}(s_i | \mu = 0, \alpha = 1, \rho), \text{ for } i \in \{1, \dots, d_s\},$$

$$p(\mathbf{z} | \mathbf{s}) = \mathcal{N}(\mathbf{x} | \mathbf{A}\mathbf{s}, \Sigma_\theta),$$

where  $\mathbf{A} \in \mathbb{R}^{d_x \times d_s}$  is our (unknown) ICA mixing matrix, which acts on the sources to produce a linear mixture; and  $\Sigma_\theta$  is a learnt or fixed diagonal covariance. This linear mixing of sources yields an intermediate representation  $\mathbf{z}$  that is then mapped to the data by a flow. Our model has three sets of variables: the observed data  $\mathbf{x}$ , the flow representation  $\mathbf{z} = f^{-1}(\mathbf{x})$ , and ICA latent sources  $\mathbf{s}$ . It can be factorised as

$$p_\theta(\mathbf{x}, \mathbf{s}) = p_\theta(\mathbf{x} | \mathbf{s}) p(\mathbf{s}) = p(\mathbf{z} | \mathbf{s}) p(\mathbf{s}) \left| \det \frac{\partial f_\theta^{-1}}{\partial \mathbf{z}} \right| \quad (5)$$

While it is simple to train a flow by maximum likelihood method when we have a simple base distribution in  $\mathcal{Z}$ , here to obtain a maximum likelihood objective we would have to marginalise out  $\mathbf{s}$  to obtain the evidence in  $\mathcal{Z}$ ; a computationally intractable procedure:

$$p(\mathbf{z}; \mathbf{A}, \Sigma_\theta) = \int d\mathbf{s} p(\mathbf{z} | \mathbf{s}; \mathbf{A}, \Sigma_\theta) p(\mathbf{s}). \quad (6)$$

A contemporary approach is to use amortised variational inference for the linear ICA part of our model. This means we introduce an approximate amortised posterior for  $\mathbf{s}$  and perform importance sampling on




Figure 2: The generative model (a) and variational posterior (b), as defined in Eq (8).

Eq (6), taking gradients through our samples using the reparameterisation trick (Kingma & Welling 2014; Rezende et al. 2014). Amortised stochastic variational inference offers numerous benefits: it scales training to large datasets by using stochastic gradient descent, our trained model can be applied to new data with a simple forward pass, and we are free to choose the functional & probabilistic form of our approximate posterior. Further our ICA model is end-to-end differentiable, making it optimal for jointly training with a flow.

We choose a linear mapping in our posterior, with  $q_\phi(\mathbf{s}|\mathbf{z}) = \text{Laplace}(\mathbf{s}|\mathbf{A}^+\mathbf{z}, \mathbf{b}_\phi)$ , where we have introduced variational parameters  $\phi = \{\mathbf{A}^+, \mathbf{b}_\phi\}$  corresponding to an unmixing matrix and a diagonal diversity. Using samples from this posterior we can define a lower bound  $\mathcal{L}$  on the evidence in  $\mathcal{Z}$

$$\begin{aligned} \log p(\mathbf{z}; \mathbf{A}, \Sigma_\theta) &\geq \mathcal{L}(\mathbf{z}; \phi, \mathbf{A}, \Sigma_\theta) \\ &= \mathbb{E}_{\mathbf{s} \sim q} [\log p(\mathbf{z}|\mathbf{s}) - \text{KL}(q_\phi(\mathbf{s}|\mathbf{z})||p(\mathbf{s}))] \end{aligned} \quad (7)$$

Using the change of variables equation, Eq (3), and the lower bound on the evidence for ICA in (7) for  $\mathcal{Z}$ , we can obtain a variational lower bound on the evidence for our data  $\mathbf{x}$  as the sum of the ICA model’s ELBO (acting on  $\mathbf{z}$ ) and the log determinant of the flow:

$$\begin{aligned} \log p_\theta(\mathbf{x}; \mathbf{A}, \Sigma_\theta) &\geq \mathcal{L}(\mathbf{x}; \theta, \phi, \mathbf{A}, \Sigma_\theta) \\ &= \mathcal{L}(\mathbf{z}; \phi, \mathbf{A}, \Sigma_\theta) + \log \left| \det \frac{\partial f_\theta^{-1}}{\partial \mathbf{z}} \right| \end{aligned} \quad (8)$$

As such our model is akin to a flow model, but with an additional latent variable  $\mathbf{s}$ ; the base distribution  $p(\mathbf{z})$  of the flow is defined through marginalizing out the linear mixing of the sources. We refer to a model with  $n$  non-linear splines mapping from  $\mathcal{X}$  to  $\mathcal{Z}$  as an  $n$ -layer Bijecta model.

In the case of non-square ICA, where our ICA model is not perfectly invertible, errors when reconstructing a mapping from  $\mathcal{S}$  to  $\mathcal{Z}$  may amplify when mapping back to  $\mathcal{X}$ . To mitigate this we add an additional regularisation term in our loss that penalises the  $L_1$  error of each point when reconstructed into  $\mathcal{X}$ . This penalisation can be weighted according to the importance of high-fidelity reconstructions for a given application.

We attempted to train Bijecta with unconstrained mixing and unmixing matrices, but found that jointly train-

ing a linear model with a powerful flow was not trivial and models failed to converge when naively optimising Eq (8). We found it crucial to appropriately constrain the unmixing matrix to get models to converge. We detail these constraints in the next section.


## 4 Whitening in $\mathbf{A}^+$ , without SVD

What are good choices for the mixing and unmixing matrices? Recall in Sec 2.2 we discussed various traditional approaches to constraining the unmixing matrix. For our flow-based model, design choices as to the parameterisation of  $\mathbf{A}^+$  stabilise and accelerate training. As before, the mixing matrix  $\mathbf{A}$  is unconstrained during optimisation. However, without the constraints on  $\mathbf{A}^+$  we describe in this section, we found that joint training of a flow with linear ICA did not converge.

Recall Eq (2) — linear ICA methods carry out whitening  $\mathbf{W}$ , performing dimensionality reduction projecting from a  $d_x$ -dimensional space to a  $d_s$ -dimensional space, and the remaining rotation and scaling operations are square. When training with a flow the powerful splines we are learning can fulfill the role of the square matrices  $\mathbf{R}$  and  $\Phi$ , but doing this ahead of the whitening itself. Put another way, the outputs from the flow can be learnt such that they are simply a whitening operation away from being effective ICA representations in  $\mathcal{S}$ . Thus, to minimise the complexities of jointly training a powerful flow with a small linear model, we can simply set  $\mathbf{A}^+ = \mathbf{W}$ , such that the unmixing matrix projects from  $d_z$  to  $d_s$  and is decorrelating. Statistical independence will come from the presence of the KL term in Eq (7): the flow will learn to give  $\mathbf{z}$  representations that, when whitened, are good ICA representations in  $\mathcal{S}$ . See Fig 3 for a visual illustration of this process and a comparison with the steps involved in linear ICA.

In previous linear ICA methods, the whitening procedure  $\mathbf{W}$  has been derived in some data-aware way. A common choice is to whiten via the Singular Value Decomposition (SVD) of the data matrix, where  $\mathbf{W} = \Sigma \mathbf{U}^T$ ,  $\Sigma$  is the rectangular diagonal matrix of singular values of  $\mathbf{X}$ , and the columns of  $\mathbf{U}$  are the left-singular vectors. Computing the SVD of the whole dataset is expensive for large datasets; for us, in the context of Bijecta, we would be re-calculating the SVD of the representations  $\mathbf{Z} = f^{-1}(\mathbf{X})$  of the entire dataset after every training step. One route around this would be online calculation of the whitening matrix (Cardoso & Laheld 1996; Hyvärinen et al. 2001). This introduces an extra optimisation process that also has to be tuned, and would interact with the training of the flow.

To tackle these shortcomings of existing whitening methods, we propose a new method for linear non-square ICA that uses Johnson–Lindenstrauss (JL) transforms (also known as sketching) (Woodruff 2014),



(a) Linear ICA

(b) Bijeta

Figure 3: (a) Sequence of actions that are performed by the elements of  $\mathbf{A}^+$ , the unmixing matrix of linear ICA.  $\mathbf{W}$  whitens the correlated data and  $\Phi\mathbf{R}$  then ensures that the whitened (decorrelated) data is also independent. (b) Sequence of actions that are performed by the elements of Bijeta.  $f^{-1}$  maps data to a representation for which the whitening matrix is the ICA matrix.  $\mathbf{W}$  now whitens  $f^{-1}(\mathbf{x})$  and the result is *also* statistically independent.

which not only works effectively as a linear ICA method, but also works in conjunction with a flow model. These JL transforms have favourable properties for ICA, as we demonstrate in theoretical results. Further, this method samples part of the whitening matrix at initialisation and leaves it fixed for the remainder of training, requiring *no hyper-parameter tuning* and making it extremely computationally efficient. This method is novel and efficient when used as a whitening method within linear ICA, and when combined with a flow as in Bijeta is a powerful method for non-linear ICA as we demonstrate in experiments.

#### 4.1 Approximately-Stiefel matrices

We have set  $\mathbf{A}^+ = \mathbf{W}$ , the whitening matrix.  $\mathbf{W}$  has two aims in non-square ICA. The first is dimensionality reduction, projecting from a  $d_x$ -dimensional space to a  $d_s$ -dimensional space. The second is to decorrelate the data it transforms, meaning that the resulting projection will have unit variance and mutually uncorrelated components. More formally we wish for  $\mathbf{W}$  of dimensionality  $d_s \times d_x$  to be decorrelating.

The set of orthogonal decorrelating rectangular matrices lie on the Stiefel Manifold (Stiefel, 1935) denoted  $\mathcal{V}$ . For matrices with  $r$  rows and  $c$  columns, a matrix  $\mathbf{G} \in \mathcal{V}(r, c)$  iff  $\mathbf{G}\mathbf{G}^* = \mathbf{I}$  ( $\mathbf{G}^*$  the conjugate transpose of  $\mathbf{G}$ ). Constraining the optimisation of  $\mathbf{W}$  to this manifold can be computationally expensive and complex (Bakir et al., 2004; Harandi & Fernando, 2016; Siegel, 2019) and instead we choose for  $\mathbf{W}$  to be *approximately* Stiefel, that is to lie close to  $\mathcal{V}(d_s, d_x)$ . This is justified by the following theorem, proved in Appendix A:

**Theorem 1.** *Let  $\mathbf{G}$  be a rectangular matrix and  $\tilde{\mathbf{G}}$  be its projection onto  $\mathcal{V}(r, c)$ . As the Frobenius norm  $\|\mathbf{G} - \tilde{\mathbf{G}}\| \rightarrow 0$  we have that  $\|\mathbf{G}\mathbf{X}\mathbf{X}^T\mathbf{G}^T - \Psi\| \rightarrow 0$ , where  $\mathbf{G}\mathbf{X}\mathbf{X}^T\mathbf{G}^T$  is the cross-correlation of the projection of data  $\mathbf{X}$  by  $\mathbf{G}$ , and  $\Psi$  is some diagonal matrix.*

Simply put, this shows that as a matrix  $\mathbf{G}$  approaches

the Stiefel manifold  $\mathcal{V}(r, c)$  the off-diagonal elements of the cross-correlation matrix of the projection  $\mathbf{G}\mathbf{X}$  are ever smaller, so  $\mathbf{G}$  is ever more decorrelating. Given these properties we want our whitening matrix to lie close to the Stiefel manifold.

##### 4.1.1 Johnson-Lindenstrauss projections

By Theorem 1 we know that we want our whitening matrix to be close to  $\mathcal{V}(d_s, d_x)$ . How might we enforce this closeness? By the definition of the Stiefel manifold, we can intuit that a matrix  $\mathbf{G}$  will lie close to this manifold if  $\mathbf{G}\mathbf{G}^T \approx \mathbf{I}$ . We formalise this as:

**Theorem 2.** *Let  $\mathbf{G} \in \mathbb{R}^{d_s \times d_x}$  and let  $\tilde{\mathbf{G}}$  be its projection onto  $\mathcal{V}(d_s, d_x)$ . As the Frobenius norm  $\|\mathbf{G}\mathbf{G}^T - \mathbf{I}\| \rightarrow 0$ , we also have  $\|\tilde{\mathbf{G}} - \mathbf{G}\| \rightarrow 0$ .*

The proof for this is presented in Appendix D. Using this theorem, we now propose an alternative to SVD-based whitening. Instead of having  $\mathbf{W} = \Sigma^{-1}\mathbf{U}^T$  be the result of SVD on the data matrix, we define our whitening matrix as a data-independent Johnson-Lindenstrauss transform. We must ensure that  $\mathbf{W}$ , our rectangular matrix, is approximately orthogonal, lying close to the manifold  $\mathcal{V}(d_s, d_x)$ . More formally by Theorem 2 our goal is to construct a rectangular matrix  $\mathbf{W}$  such that  $\mathbf{W}\mathbf{W}^T \approx \mathbf{I}$ .

We construct approximately orthogonal matrices for  $\mathbf{W}$  by way of Johnson-Lindenstrauss (JL) Projections (Johnson & Lindenstrauss, 1984). A JL projection  $\mathbf{W}$  for  $\mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_s}$  is sampled at initialisation from a simple binary distribution (Achlioptas (2003):

$$W_{i,j} = \begin{cases} +1/\sqrt{d_s}, & \text{with probability } \frac{1}{2} \\ -1/\sqrt{d_s}, & \text{with probability } \frac{1}{2} \end{cases} \quad (9)$$

This distribution satisfies  $\mathbb{E}[\mathbf{W}\mathbf{W}^T] = \mathbf{I}$ , and such a draw has  $\mathbf{W}\mathbf{W}^T \approx \mathbf{I}$ . We choose to fix  $\mathbf{W}$  after initialisation such that  $\mathbf{A}^+ = \mathbf{W}$  never updates, greatly simplifying optimisation.




Figure 4: Here we run linear ICA on a pair of images (a) that are mixed linearly (mix =  $w_1 * \text{image}_1 + w_2 * \text{image}_2$ ) (b) to form a dataset with 512 points. In both cases  $w_1$  and  $w_2$  are sampled from a uniform distribution. We plot the mixing matrix  $\mathbf{A}$  for our JL-Cayley model with a quasi-uniform GG prior with  $\rho = 10$  (c) and for FastICA (Hyvärinen & Oja, 1997) as a benchmark.  $\mathbf{A}$  should recover the source images, which occurs for both models.




Figure 5: Here we demonstrate that Bijeta is capable of unmixing non-linearly mixed sources, better than VAEs with ICA-appropriate priors. We take a dSprites heart and, using a randomly sampled affine transformation, move it around a 32 by 32 background (a). With 2-D GG priors with  $\rho = 10$  for a convolutional VAE (b) and for Bijeta (c) we plot the generations resulting from traversing the 2-D latent-source space in a square around the origin. We sketch the learnt axis of movement of the sprite with white lines. In (b) the VAE does not ascribe consistent meaning to its latent dimensions. It has failed to discover consistent independent latent sources: it has a sudden change in the learnt axes of movement along the second dimension, as seen by the kink in the white vertical lines. In (c) Bijeta is able to learn a simple affine transformation along each latent dimension, consistently spanning the space. In Fig B.2 we show the posterior distributions of both these models and show that Bijeta is better able to match the GG prior than the VAE, supporting our findings here.

## 5 Experiments

Here we show that our approach outperforms VAEs and flows with ICA-priors at discovering ICA sources in image data. But first, as a sanity check, we show that a linear ICA model using JL projections to whiten can successfully unmix linearly mixed sources in Fig 4. For details on how to implement such a linear ICA model, see Appendix E. We take a pair of images from dSprites and create linear mixtures of them. We see that linear ICA with JL projections can successfully discover the true sources, the images used to create the mixtures, in the columns of  $\mathbf{A}$ .

**Affine Data** Given that we have established that our novel theory for decorrelating matrices can produce standalone linear ICA models, we now want to ascertain that our hybrid model performs well in non-linear mixing situations. To do so we create a dataset consisting of a subset of dSprites where we have a light-blue heart randomly uniformly placed on a black field. The true latent sources behind these randomly sampled affine transformations are simply the coordinates of the

heart. First, in Fig B.2 we demonstrate that linear ICA models are unable to uncover the true latent sources. As expected non-linear mixing regimes motivate the use of flexible non-linear models.

We now demonstrate that Bijeta can uncover the latent sources underpinning these affine transformations, whereas VAEs with ICA-appropriate priors fail to do so. For details of VAE architecture, see Appendix I. These VAEs are able to learn to reconstruct data well, but the learnt latent space does not correspond to the underlying statistically independent sources (see Figs 1 and B.2). In fact for VAEs the effect of the latent variables is not consistent throughout the latent space, as seen in Fig 5. For Bijeta, the learnt latent space corresponds to the underlying statistically independent sources (see Figs 1 and B.2), and the meaning of the latent variables is consistent in Fig 5. Further in Fig 5 the model seems able to extrapolate outside the training domain: it generates images where the heart is partially rendered at the edges of the frame, even removing the heart entirely at times, even though such images are not in the training set.

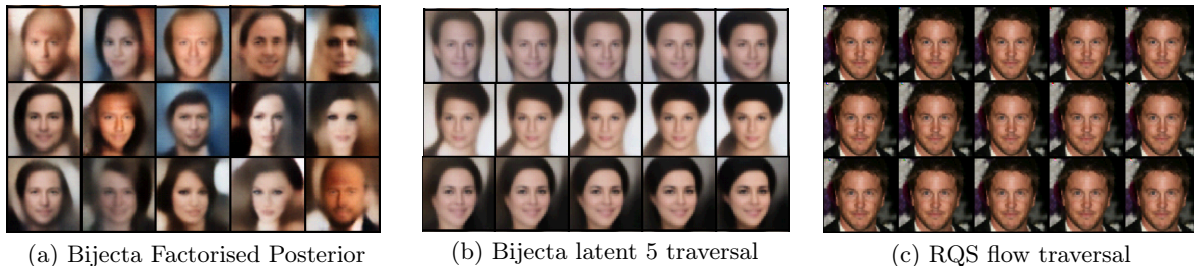


Figure 6: (a) shows decodings from an 8-layer Bijecta ( $d_s = 32$ ) trained on CelebA with a Laplace prior (GG  $\rho = 1$ ) where we sample from the factorised approximation to Bijecta’s posterior. See Fig J.6 for more such samples. (b) shows latent traversals for 3 different datapoints all along the same axis-aligned direction, for this same model. (c) shows traversals for a single embedded training datapoint from CelebA moving along 3 latent directions in an RQS flow with Laplace base distribution. Though we have selected 3 dimensions, all  $\mathcal{Z}$  dimensions had similar latent traversals. In (b-c) Images in the center correspond to the original latent space embedding, on either side we move up to 6 standard deviations away along this direction with other dimensions remaining fixed. The flow has not discovered axis-aligned transforms, whereas Bijecta has learned informative latent dimensions: here the dimension encodes hair thickness. Note that identity is maintained throughout and that the transform is consistent across different posterior samples. See Appendix J.3 for gallery of transforms for Bijecta.

**Natural Images** The previous experiments show that our model is capable of isolating independent sources on toy data. We complement this finding with experiments on a more complex natural image dataset, CelebA, and show that here too our model outperforms VAEs in learning factorisable representations.

An ersatz test of this can be done by synthesising images where we sample from a factorised approximation of Bijecta’s posterior. If the learned latent sources are actually independent, then the posterior over latent sources given the entire dataset should factorize into a product across dimensions, i.e.  $q(\mathbf{s}) = \prod_i q(s_i)$ . In this case, we can fit an approximation to the posterior by fitting  $d_s$  independent one-dimensional density estimates on  $q(s_i)$ . If the sources are not independent, then this factorized approximation to the posterior will be missing important correlations and dependencies. In Fig 6a samples from this factorised approximation look reasonable, suggesting that Bijecta has learnt representations that are statistically independent.

To quantify this source-separation numerically, we measure the total correlation (TC) of the aggregate posteriors of Bijecta ( $q(\mathbf{s}|\mathbf{z})$ ) and VAEs ( $q(\mathbf{z}|\mathbf{x})$ ) as Chen et al. (2018) do. Intuitively, the TC measures how well a distribution is approximated by the product of its marginals – or how much information is shared between variables due to dependence (Watanabe, 1960). It directly measures how well an ICA model has learnt decorrelated and independent latent representations (Eversson & Roberts, 2001). Formally, it is the KL divergence between a distribution  $r(\cdot)$  and a factorised representation of the distribution:  $\text{TC} = \text{KL}(r(\mathbf{s}) || \prod_i r(s_i))$ , where  $i$  indexes over (latent) dimensions.

In Table 1 we show that Bijecta learns an aggregate posterior with significantly lower TC values than both VAEs with Laplace priors, and  $\beta$ -TCVAEs – which in their training objective penalise the TC by a factor  $\beta$  (Chen et al., 2018). Our model has learnt a better ICA solution. We also include numerical results in Appendix J.1 showing that Bijecta outperforms linear ICA on a variety of natural image datasets.

Table 1: Total Correlation Results: We evaluate the source separation of different models on CelebA via the TC of the validation set embeddings in the 32-D latent space of: Laplace prior VAEs,  $\beta$ -TCVAEs ( $\beta = 15$ ), and Bijecta with a Laplace prior ( $\pm$  indicates the standard deviation over 2 runs). VAEs use the same architecture and training as Chen et al. (2018).

	Laplace-VAE	$\beta$ -TCVAE	Laplace-Bijecta
TC:	106.7 $\pm$ 0.9	55.7 $\pm$ 0.1	<b>13.1 <math>\pm</math> 0.4</b>

**Dimensionality reduction on flow models** To conclude, having shown that Bijecta outperforms VAEs on a variety of non-linear ICA tasks, we now contrast our model’s ability to automatically uncover sources relative to flow models with heavy-tailed base distributions. We do so by measuring the cumulative explained variance by the dimensions in  $\mathcal{Z}$  for both models. If a small number of dimensions explains most of  $\mathcal{Z}$ ’s variance then the model has learnt a bijection which only requires a small number of dimensions to be invertible. It has in effect learnt the generating sources underpinning the data.

In Fig 7 we show that Bijecta induces better-




Figure 7: Explained variance plots for the embedding in  $\mathcal{Z}$ , as measured by the sums of the eigenvalues of the covariance matrix of the embeddings, for both our Bijecta model and for an RQS model of equivalent size trained with a Laplace base distribution (GG distribution with  $\rho = 1$ ). For both Fashion-MNIST (left) and CIFAR 10 (right) datasets we see that the Bijecta model has learned a compressive flow, where most of the variance can be explained by only a few linear projections. The shaded region denotes the first 64 dimensions, corresponding to the size of the target source embedding  $\mathcal{S}$ .

compressed representations in  $\mathcal{Z}$  than non-compressive flows. We plot the eigenvalues of the covariance matrix on the output of the flow, i.e. on  $\text{Cov}(f(\mathbf{X}))$ , to see how much of the total variance in the learned feature space  $\mathcal{Z}$  can be explained in a few dimensions. In doing so we see that a flow trained jointly with a linear ICA model with  $d_s = 64$  effectively concentrates variation into a small number of intrinsic dimensions; this is in stark contrast with the RQS flows trained with only a Laplace base distribution. This demonstrates that our model is able to automatically detect relevant directions on a low dimensional manifold in  $\mathcal{Z}$ , and that the bijective component of our model is better able to isolate latent sources than a standard flow.

For a visual illustration of this source separation we show the difference in generated images resulting from smoothly varying along each dimension in  $\mathcal{S}$  for Bijecta models and in  $\mathcal{Z}$  for flows in Fig 6. Bijecta is clearly able to discover latent sources, whereby it learns axis-aligned transformations of CelebA faces, whereas a flow with equivalent computational budget and a heavy-tailed base distribution is not able to.

All flow-based baselines are trained using the objective in Eq (4), using Real-NVP style factoring-out (Durkan et al. 2019; Dinh et al. 2015), and are matched in size and neural network architectures to the flows of Bijecta models. See Appendix I for more details.

## 6 Related Work

One approach to extend ICA to non-linear settings is to have a non-linear mapping acting on the independent sources and data (Burel, 1992; Deco & Brauer, 1995; Yang et al. 1998; Valpola et al. 2003). In general, non-linear ICA models have been shown to be hard to train, having problems of unidentifiability: the model has numerous local minima it can reach under its training objective, each with potentially different learnt sources (Hyvärinen & Pajunen, 1999; Karhunen, 2001; Almeida, 2003; Hyvarinen et al. 2019). Some

non-linear ICA models have been specified with additional structure to reduce the space of potential solutions, such as putting priors on variables (Lappalainen & Honkela, 2000) or specifying the precise non-linear functions involved (Lee & Koehler, 1997; Taleb, 2002). Recent work shows that conditioning the source distributions on some always-observed side information, say time index, can be sufficient to induce identifiability in non-linear ICA (Khemakhem et al. 2020).

Modern flows were first proposed as an approach to non-linear square ICA (Dinh et al. 2015), but are also motivated by desires for more expressive priors and posteriors (Kingma et al. 2016; Papamakarios et al. 2019). Early approaches, known as symplectic maps (Deco & Brauer, 1995; Parra et al. 1995, 1996), were also proposed for use with ICA. Flows offer expressive dimensionality-preserving (and sometimes volume-preserving) bijective mappings (Dinh et al. 2017; Kingma & Dhariwal, 2018). Flows have been used to provide feature extraction for linear discriminative models (Nalisnick et al. 2019). Orthogonal transforms have been used in normalizing flows before, to improve the optimisation properties of Sylvester flows (Van Den Berg et al. 2018; Golinski et al. 2019). Researchers have also looked at constraining neural network weights to the Stiefel-manifold (Li et al. 2020).

## 7 Conclusion

We have developed a method for performing non-linear ICA large high-dimensional image datasets which combines state-of-the-art flow-based models and a novel theoretically grounded linear ICA method. This model succeeds where existing probabilistic deep generative models fail: its constituent flow is able to learn a representation, lying in a low dimensional manifold in  $\mathcal{Z}$ , under which sources are separable by linear unmixing. In source space  $\mathcal{S}$ , this model learns a low dimensional, explanatory set of statistically independent latent sources.



## Acknowledgments

This research was directly funded by the Alan Turing Institute under Engineering and Physical Sciences Research Council (EPSRC) grant EP/N510129/1. AC was supported by an EPSRC Studentship. MW was supported by EPSRC grant EP/G03706X/1. CH was supported by the Medical Research Council, the Engineering and Physical Sciences Research Council, Health Data Research UK, and the Li Ka Shing Foundation. SR gratefully acknowledges support from the UK Royal Academy of Engineering and the Oxford-Man Institute.

We thank Tomas Lazauskas, Jim Madge and Oscar Giles from the Alan Turing Institute’s Research Engineering team for their help and support.

## References

- Absil, P. A. & Malick, J. (2012). Projection-like retractions on matrix manifolds. *SIAM Journal on Optimization*, 22(1), 135–158.
- Achlioptas, D. (2003). Database-friendly random projections: Johnson-Lindenstrauss with binary coins. In *Journal of Computer and System Sciences*, volume 66 (pp. 671–687).
- Almeida, L. B. (2003). MISEP – Linear and Nonlinear ICA Based on Mutual Information. *Journal of Machine Learning Research*, 4, 1297–1318.
- Bakir, G. H., Gretton, A., Franz, M., & Schölkopf, B. (2004). Multivariate regression via Stiefel manifold constraints. In *Joint Pattern Recognition Symposium* (pp. 262–269).
- Bell, A. J. & Sejnowski, T. J. (1995). An information-maximisation approach to blind separation and blind deconvolution. *Neural Computation*, 7(6), 1004–1034.
- Burel, G. (1992). Blind separation of sources: A nonlinear neural algorithm. *Neural Networks*, 5(6), 937–947.
- Cardoso, J. F. (1989a). Blind identification of independent components with higher-order statistics. In *IEEE Workshop on Higher-Order Spectral Analysis*.
- Cardoso, J. F. (1989b). Source separation using higher order moments. In *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, volume 4 (pp. 2109–2112).
- Cardoso, J. F. (1997). Infomax and Maximum Likelihood for Blind Source Separation. *IEEE Letters on Signal Processing*, 4, 112–114.
- Cardoso, J. F. & Laheld, B. H. (1996). Equivariant adaptive source separation. *IEEE Transactions on Signal Processing*, 44(12), 3017–3030.
- Cayley, A. (1846). Sur quelques propriétés des déterminants gauches. *Journal für die reine und angewandte Mathematik*, 32, 119–123.
- Chen, R. T. Q., Li, X., Grosse, R., & Duvenaud, D. (2018). Isolating Sources of Disentanglement in Variational Autoencoders. In *NeurIPS*.
- Choudrey, R. (2000). *Variational Methods for Bayesian Independent Component Analysis*. PhD thesis, University of Oxford.
- Comon, P. (1994). Independent component analysis, A new concept? *Signal Processing*, 36(3), 287–314.
- Dasgupta, S. & Gupta, A. (2003). An Elementary Proof of a Theorem of Johnson and Lindenstrauss. *Random Structures and Algorithms*, 22(1), 60–65.
- Deco, G. & Brauer, W. (1995). Higher Order Statistical Decorrelation without Information Loss. In *NeurIPS*.
- Dinh, L., Krueger, D., & Bengio, Y. (2015). NICE: Non-linear Independent Components Estimation. In *ICLR*.
- Dinh, L., Sohl-Dickstein, J., & Bengio, S. (2017). Density estimation using Real NVP. In *ICLR*.
- Durkan, C., Bekasov, A., Murray, I., & Papamakarios, G. (2019). Neural Spline Flows. In *NeurIPS*.
- Everson, R. & Roberts, S. J. (1999). Independent Component Analysis: A Flexible Nonlinearity and Decorrelating Manifold Approach. *Neural Computation*, 11(8), 1957–83.
- Everson, R. & Roberts, S. J. (2001). *Independent Component Analysis*. Cambridge University Press.
- Golinski, A., Rainforth, T., & Lezcano-Casado, M. (2019). Improving Normalizing Flows via Better Orthogonal Parameterizations. In *ICML Workshop on Invertible Neural Networks and Normalizing Flows*.
- Harandi, M. & Fernando, B. (2016). Generalized Back-Propagation, Etude De Cas: Orthogonality.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., & Lerchner, A. (2017).  $\beta$ -VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *ICLR*.
- Hyvärinen, A., Karhunen, J., & Oja, E. (2001). *Independent Component Analysis*. John Wiley.
- Hyvärinen, A. & Oja, E. (1997). A fast fixed-point algorithm for independent component analysis. *Neural Computation*, 9(7), 1483–1492.
- Hyvärinen, A. & Pajunen, P. (1999). Nonlinear independent component analysis: Existence and uniqueness results. *Neural Networks*, 12(3), 429–439.
- Hyvarinen, A., Sasaki, H., & Turner, R. E. (2019). Nonlinear ICA Using Auxiliary Variables and Generalized Contrastive Learning. In *AISTATS*.

- Johnson, W. B. & Lindenstrauss, J. (1984). Extensions of Lipschitz mappings into a Hilbert space. *Contemporary mathematics*, 26(1), 189–206.
- Karhunen, J. (2001). Nonlinear Independent Component Analysis. In R. Everson & S. J. Roberts (Eds.), *ICA: Principles and Practice* (pp. 113–134). Cambridge University Press.
- Khemakhem, I., Kingma, D. P., Monti, R. P., & Hyvärinen, A. (2020). Variational Autoencoders and Nonlinear ICA: A Unifying Framework. In *AISTATS*.
- Kingma, D. P. & Dhariwal, P. (2018). Glow: Generative flow with invertible 1x1 convolutions. *NeurIPS*.
- Kingma, D. P. & Lei Ba, J. (2015). Adam: A Method for Stochastic Optimisation. In *ICLR*.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., & Welling, M. (2016). Improved Variational Inference with Inverse Autoregressive Flow. In *NeurIPS*.
- Kingma, D. P. & Welling, M. (2014). Auto-encoding Variational Bayes. In *ICLR*.
- Lappalainen, H. & Honkela, A. (2000). Bayesian Non-Linear Independent Component Analysis by Multi-Layer Perceptrons. In M. Girolami (Ed.), *Advances in Independent Component Analysis* (pp. 93–121). Springer.
- Lawrence, N. D. & Bishop, C. M. (2000). *Variational Bayesian Independent Component Analysis*. Technical report, University of Cambridge.
- Lee, T.-W., Girolami, N., Bell, A. J., & Sejnowski, T. J. (2000). A Unifying Information-Theoretic Framework for Independent Component Analysis. *Computers & Mathematics with Applications*, 39(11), 1–21.
- Lee, T. W. & Koehler, B. U. (1997). Blind source separation of nonlinear mixing models. *Neural Networks for Signal Processing - Proceedings of the IEEE Workshop*, (pp. 406–415).
- Lee, T.-W. & Sejnowski, T. J. (1997). Independent Component Analysis for Mixed Sub-Gaussian and Super-Gaussian Sources. *Joint Symposium on Neural Computation*, (pp. 6–13).
- Li, J., Li, F., & Todorovic, S. (2020). Efficient Riemannian Optimization on the Stiefel Manifold via the Cayley Transform. In *ICLR*.
- Locatello, F., Bauer, S., Lucie, M., Rätsch, G., Gelly, S., Schölkopf, B., & Bachem, O. (2019). Challenging common assumptions in the unsupervised learning of disentangled representations. In *ICML*, volume 2019-June.
- Mackay, D. J. C. (1996). *Maximum Likelihood and Covariant Algorithms for Independent Component Analysis*. Technical report, University of Cambridge.
- Mathieu, E., Rainforth, T., Siddharth, N., & Teh, Y. W. (2019). Disentangling Disentanglement in Variational Autoencoders. In *ICML*.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., & Lakshminarayanan, B. (2019). Detecting Out-of-Distribution Inputs to Deep Generative Models Using Typicality.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. (2019). *Normalizing Flows for Probabilistic Modeling and Inference*. Technical report, DeepMind, London, UK.
- Parra, L., Deco, G., & Miesbach, S. (1995). Redundancy reduction with information-preserving nonlinear maps. *Network: Computation in Neural Systems*, 6(1), 61–72.
- Parra, L., Deco, G., & Miesbach, S. (1996). Statistical Independence and Novelty Detection with Information Preserving Nonlinear Maps. *Neural Computation*, 8(2), 260–269.
- Rezende, D. J. & Mohamed, S. (2015). Variational Inference with Normalizing Flows. In *ICML*.
- Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *ICML*.
- Rolinek, M., Zietlow, D., & Martius, G. (2019). Variational autoencoders pursue pca directions (by accident). In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2019-June (pp. 12398–12407).
- Roweis, S. & Ghahramani, Z. (1999). A unifying review of linear gaussian models. *Neural Computation*, 11(2), 305–345.
- Siegel, J. W. (2019). Accelerated Optimization With Orthogonality Constraints.
- Stiefel, E. (1935). Richtungsfelder und Fernparallelismus in n-dimensionalen Mannigfaltigkeiten. *Commentarii mathematici Helvetici*, 8, 305–353.
- Stühmer, J., Turner, R. E., & Nowozin, S. (2019). Independent Subspace Analysis for Unsupervised Learning of Disentangled Representations.
- Taleb, A. (2002). A generic framework for blind source separation in structured nonlinear models. *IEEE Transactions on Signal Processing*, 50(8), 1819–1830.
- Valpola, H., Oja, E., Ilin, A., Honkela, A., & Karhunen, J. (2003). Nonlinear blind source separation by variational Bayesian learning. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, E86-A(3), 532–541.
- Van Den Berg, R., Hasenclever, L., Tomczak, J. M., & Welling, M. (2018). Sylvester normalizing flows for variational inference. In *UAI*, volume 1 (pp. 393–402).

- Watanabe, S. (1960). Information Theoretical Analysis of Multivariate Correlation. *IBM Journal of Research and Development*, 4(1), 66–82.
- Woodruff, D. P. (2014). Sketching as a Tool for Numerical Linear Algebra. *Foundations and Trends in Theoretical Computer Science*, 10(2), 1–157.
- Yang, H. H., Amari, S. I., & Cichocki, A. (1998). Information-theoretic approach to blind separation of sources in non-linear mixture. *Signal Processing*, 64(3), 291–300.