
Representation Learning of Compositional Data

Marta Avalos-Fernandez[†] Richard Nock^{‡§¶} Cheng Soon Ong^{‡§} Julien Rouar[†] Ke Sun^{‡*}

[†]Université de Bordeaux, [‡]Data61,

[§]the Australian National University and [¶]the University of Sydney

first.last@{u-bordeaux.fr,data61.csiro.au}

Abstract

We consider the problem of learning a low dimensional representation for compositional data. Compositional data consists of a collection of nonnegative data that sum to a constant value. Since the parts of the collection are statistically dependent, many standard tools cannot be directly applied. Instead, compositional data must be first transformed before analysis. Focusing on principal component analysis (PCA), we propose an approach that allows low dimensional representation learning directly from the original data. Our approach combines the benefits of the log-ratio transformation from compositional data analysis and exponential family PCA. A key tool in its derivation is a generalization of the scaled Bregman theorem, that relates the perspective transform of a Bregman divergence to the Bregman divergence of a perspective transform and a *remainder* conformal divergence. Our proposed approach includes a convenient surrogate (upper bound) loss of the exponential family PCA which has an easy to optimize form. We also derive the corresponding form for nonlinear autoencoders. Experiments on simulated data and microbiome data show the promise of our method.

1 Introduction

Compositional data analysis (CoDA) is a subfield of statistics introduced more than three decades ago [3, 2, 1, 29]. Compositional data consist of a collection of nonnegative measurements that sum to a constant value, typically, proportions that sum to 1. Because knowing the sum, one component can be determined from the sum of the remainder, the parts that make up the composition are mathematically and statistically dependent. This distinct structure complicates analysis and does not allow standard statistical analyses. Ignoring the underlying nature of the data studied might give rise to misleading conclusions.

Among others, [1] and [13] provided a framework to perform CoDA by mapping data from the constrained simplex space to the Euclidian space using nonlinear log-ratio transforms. In this paper, we focus on Principal Components Analysis (PCA), one of the main tools for exploratory analysis of compositional data. Just like in standard Euclidean data, it is particularly useful when the first few principal components explain enough variability to be considered as representative. Unfortunately, any operation of centering or scaling destroys the compositional nature of the data, which complicates a direct application of PCA.

Our motivation for studying CoDA comes from the recent explosion of microbiome studies [14, 15]. Indeed, spectacular advances in 16S rRNA gene sequencing of the bacterial component of the human microbial community (microbiota) have enabled researchers to investigate human health and disease, leading to new insights into the role of these microbial communities. The microbiota sequencing data are measured as read counts interpreted as a species' abundance in a microbial community. To make the microbial abundance comparable across samples, data are normalized to the relative abundances

* Authors in alphabetical order




Figure 1: 2D Visualization of the low dimensional representation A on the arms dataset.

of all bacteria observed. On the other hand, because high-throughput experiments produce large amounts of data, multivariate analysis is indispensable [27, 21]. There is a stress to understand the soundness of models [5].

In this paper, we propose to learn a low dimensional representation of CoDA from the *original* data. To account for the nonlinearity due to the compositional nature of the data, we start from exponential family PCA [12] that we augment with the compositional constraint and then simplify the loss to be optimized via a generalization of a recent result [25] on Bregman divergences, which may be of independent interest. We also propose a nonlinear autoencoder (AE) version to learn the low dimensional representation.

Let us examine a toy example to illustrate our approach. We generate the arms dataset in \mathfrak{S}^{19} by evenly interpolating the simplex center and each of the 20 vertices with 100 points, therefore yielding a matrix $X_{20 \times 2000}$. Figure 1 shows the 2D representation $A_{2 \times 2000}$ computed by five methods: the standard PCA; `clr`-PCA computes the standard PCA after performing `clr`; `CoDA-PCA` and `CoDA-AE` are our proposed methods; `t-SNE` [32] is a popular nonlinear dimensionality reduction method and is applied on X directly. In the PCA plot, the black segments indicate that the PCA reconstruction is outside of the simplex. PCA cannot be directly adapted to CoDA because the projection on the principal components may go beyond the convex hull of the vertices. It is clear that only `CoDA-PCA` and `CoDA-AE` uncover the true structure, where all the arms are clearly presented, and their connections are faithfully presented.

2 Compositional Data Analysis

We briefly review some definitions of CoDA. Compositional data are proportions: X is a compositional dataset if and only if $X \in \mathfrak{R}^{d \times m}$ such that $\forall i \in [1, \dots, m]$ the vector column \mathbf{x}_i of X is in the simplex $\mathfrak{S}^d = \{\mathbf{x} \in \mathfrak{R}^d : \forall j, x_j > 0; \sum_{j=1}^d x_j = \kappa\}$, where $\kappa > 0$ is a constant, classically 1. Here the superscript d does not denote the dimensionality as $\dim(\mathfrak{S}^d) = d - 1$. For a dataset X' which contains counts of strictly positive values, we reduce it to a compositional dataset by dividing out the totals, that is we compute the CoDA set X such that: $\mathbf{x}_i = \mathbf{x}'_i \frac{1}{\sum_{j=1}^d x'_{ji}}$ is the vector of proportions for individual i .

Using Bregman divergences makes explicit a dual affine [6] coordinate space which is in fact the log coordinates of Aitchison [1]. It is in this space that we have affine constraints, which are therefore non-linear in the "primal", ambient space. To manage this nonlinear structure, it has been proposed [4] to first apply a log-ratio transformation to transpose the data into real Euclidean space. For instance, the additive log-ratio transformation (`alr`) applies log-ratio between a component and a reference component; the centered log-ratio transformation (`clr`) scales each subject vector by its geometric mean; and the isometric log-ratio transformation (`ilr`) is associated with an orthogonal coordinate system in the simplex. Afterwards, standard PCA is performed.

By definition, the `clr` transformation is

$$\mathbf{c}_{\text{clr}}(\mathbf{x}) = \log \left(\frac{\mathbf{x}}{g(\mathbf{x})} \right) = \mathbf{C}^{\text{clr}} \log(\mathbf{x}) = \log(\mathbf{x}) - \overline{\log(\mathbf{x})} \mathbf{1}_d, \quad (1)$$

where $g(\mathbf{x}) = (\prod_{j=1}^d x_j)^{1/d}$ the geometric mean of \mathbf{x} , $\bar{x} = \frac{1}{d} \sum_{j=1}^d x_j$ is the arithmetic mean of \mathbf{x} , $\mathbf{1}_d$ is the $1 \times d$ vector of all ones, and $\mathbf{C}^{\text{clr}} = \mathbf{I}_d - \frac{1}{d} \mathbf{1}_d \mathbf{1}_d^T$. The purpose of the log-ratio transformation

(centered or not) is to go back to \mathfrak{R}^d from \mathfrak{S}^d without losing information. Notice that $\log(x_j)$, $\overline{\log(\mathbf{x})} \in (-\infty, 0)$ so the compositional data is embedded in \mathfrak{R}^d under the clr transformation. The reverse operation: $\mathbf{x} = \mathbf{c}_{\text{KL}}^{-1}(\mathbf{x}') = \exp(\mathbf{x}') \prod_{j=1}^d \exp(\frac{1}{1-d}x'_j)$ embeds \mathfrak{R}^d into \mathfrak{S}^d . See the table below for a comparison of clr, alr and ilr, presented as different transformations $\mathbf{C} \log(\mathbf{x})$. They are equivalent up to linear transformations. Without loss of generality we focus on the clr.

clr	alr	ilr
$\mathbf{C}_{d \times d}^{\text{clr}} = \mathbf{I}_d - \frac{1}{d} \mathbf{1}_d \mathbf{1}_d^\top$	$\mathbf{C}_{(d-1) \times d}^{\text{alr}} = [\mathbf{I}_{d-1}, -\mathbf{1}_{d-1}]$	$\mathbf{C}_{(d-1) \times d}^{\text{ilr}} \in \{\mathbf{RC}^{\text{clr}} : \mathbf{RR}^\top = \mathbf{I}_{d-1}\}$

However interpreting the resulting coordinates is still challenging [13, 24]: alr transformation is no distance-preserving; clr leads to degenerate distributions and singular covariance matrices; ilr avoids the precedent drawbacks, but still, results from complicated nonlinear transformations are difficult to interpret. Currently, there seems to be no consensus about the best practices ([16] versus [31]) and, in all cases, log-transforming is not a remedy for all the difficulties arisen by CoDA [20].

3 Exponential Family Principal Component Analysis

Another way to apply dimension reduction is to perform a generalized PCA on crude count data. Based on the same ideas as the generalised linear model, [12] described a generalized PCA model for distributions from the exponential family. We first recall the standard PCA setting.

3.1 Principal Component Analysis

For simplicity suppose that the data matrix \mathbf{X} is already centered that can be easily achieved by appending to \mathbf{A} matrix a row of ones.

(Traditional PCA) We have a dataset $\mathbf{X} \in \mathfrak{R}^{d \times m}$ that we approximate as $\mathbf{X} \sim \mathbf{V}^\top \mathbf{A}$ by minimizing the following loss wrt the constraints $\mathbf{A} \in \mathfrak{R}^{\ell \times m}$, $\mathbf{V} \in \mathfrak{R}^{\ell \times d}$, $\mathbf{V}\mathbf{V}^\top = \mathbf{I}_\ell$:

$$\ell_{\text{PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \doteq \|\mathbf{X} - \mathbf{V}^\top \mathbf{A}\|_{\text{F}}^2. \quad (2)$$

Hence, observations are column-wise. $\mathbf{V} : \mathfrak{R}^d \rightarrow \mathfrak{R}^\ell$ is surjective with $\mathbf{V}^\top \mathbf{V}$ defining a rank- ℓ projection, assuming in general $\ell < d$. \mathbf{A} is the representation of data points. The goodness of fit of the representation is measured by the squared Frobenius norm. We summarise the different transformations and loss functions in Table 1.

Observe that instead of finding a linear representation \mathbf{A} and its corresponding linear loadings \mathbf{V} , we can consider nonlinear functions for encoding and decoding the latent representation. When the nonlinear encoder and decoder are implemented as feed-forward neural networks, we arrive at the autoencoder setting.

3.2 Bregman Divergence and φ -PCA

As mentioned in the introduction, compositional data do not live in Euclidean space. Count data are naturally linked to the Poisson distribution, and therefore we should consider an exponential family model for count data. From the Bayesian viewpoint, the PCA goal is to minimize a distance (L^2 for usual PCA) which is equivalent to a Bregman divergence minimization (or a Likelihood function maximization).

Definition 1 (Bregman Divergence) Let $\varphi : \mathfrak{R}^d \rightarrow \mathfrak{R}$ convex differentiable. The Bregman divergence D_φ with generator φ is

$$D_\varphi(\mathbf{x} \parallel \mathbf{x}') \doteq \varphi(\mathbf{x}) - \varphi(\mathbf{x}') - (\mathbf{x} - \mathbf{x}')^\top \nabla \varphi(\mathbf{x}'). \quad (3)$$

A Bregman divergence is just a truncation of the Taylor expansion of a function. It can therefore be defined for any differentiable function, not just the convex ones. If φ is *not* convex, we call D_φ a *Bregman distortion*, which is a signed dissimilarity. We denote by $\varphi^*(\mathbf{x}) \doteq \sup_{\mathbf{y}} \{\mathbf{x}^\top \mathbf{y} - \varphi(\mathbf{y})\}$ the convex conjugate of the generator φ [9].

Table 1: Summary of methods in this paper

Method	Original	Reconstruction	Distortion	Notes
PCA	\mathbf{X}	$\mathbf{V}^\top \mathbf{A}$	$\ \cdot - \cdot\ _{\mathbb{F}}^2$	classical PCA(2)
φ -PCA	\mathbf{X}	$\nabla^* \varphi(\mathbf{V}^\top \mathbf{A})$	$D_\varphi(\cdot \ \cdot)$	exponential family PCA (4)
clr-PCA	$\mathbf{c}_{\text{KL}}(\mathbf{X})$	$\mathbf{V}^\top \mathbf{A}$	$\ \cdot - \cdot\ _{\mathbb{F}}^2$	CoDA with clr (5)
gauged- φ -PCA	$\nabla \varphi(\tilde{\mathbf{X}})$	$\mathbf{V}^\top \mathbf{A}$	$D_{\varphi^*}(\cdot \ \cdot)$	General Bregman PCA (8)
CoDA PCA	$\mathbf{c}_{\text{KL}}(\mathbf{X})$	$\mathbf{V}^\top \mathbf{A}$	$D_{\text{exp}}(\cdot \ \cdot)$	(11) is a special case of (8)
S-CoDA-PCA	$\tilde{\mathbf{x}}_i$	$\nabla \check{\text{KL}}(\exp(\mathbf{V}^\top \mathbf{a}_i))$	inner product	upper bound (17)
CoDA AE	\mathbf{X}	$g_\theta \circ h_\Phi(\mathbf{X})$	$D_{\text{exp}}(\cdot \ \cdot)$	neural networks g_θ and h_Φ

PCA has been generalized to the exponential families in a way that makes fitting occur in the *natural parameter* space [12, 19] (and references therein). The optimization problem is non-convex. The algorithmic strategy proposed by [12] is to use an alternating sequence of convex minimizations under constraints. Alternatively, [19] proposed maximizing the deviance (as a generalized notion of variance) and [10] proposed maximizing the likelihood function via a variational algorithm and gradient descent.

We denote exponential family PCA as φ -PCA, where φ is the cumulant of the exponential family, which is strictly convex differentiable with convex conjugate φ^* , and uniquely determines the exponential family under mild conditions [7]. Note that for φ -PCA, \mathbf{X} is not necessarily in a vector space (e.g. $\mathbf{X} \notin \mathbb{R}^{d \times m}$).

(φ -PCA) We have a dataset $\mathbf{X}_{d \times m}$ that we approximate as $\mathbf{X} \sim \nabla \varphi^*(\mathbf{V}^\top \mathbf{A})$ with $\mathbf{A} \in \mathbb{R}^{\ell \times m}$, $\mathbf{V} \in \mathbb{R}^{\ell \times d}$, $\mathbf{V}\mathbf{V}^\top = \mathbf{I}_\ell$, through minimizing the Bregman loss

$$\ell_{\varphi\text{-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \doteq \sum_i D_\varphi(\mathbf{x}_i \| \nabla \varphi^*(\mathbf{V}^\top \mathbf{a}_i)) = D_\varphi(\mathbf{X} \| \nabla \varphi^*(\mathbf{V}^\top \mathbf{A})). \quad (4)$$

Vectors are column-vectors: $\mathbf{x}_i, \mathbf{a}_i$ are respectively column observation i in the ambient and principal spaces, respectively. This formulation has a major advantage that linear algebra may be used to fit \mathbf{A}, \mathbf{V} while \mathbf{X} may not lie in a vector space, see for example [12, 19] and references therein. We remark that because of the dual symmetry of Bregman divergences, we have $D_\varphi(\mathbf{X} \| \nabla \varphi^*(\mathbf{V}^\top \mathbf{A})) = D_{\varphi^*}(\mathbf{V}^\top \mathbf{A} \| \nabla \varphi(\mathbf{X}))$ [8]. Notice there exists a little "hole" in the φ -PCA definition, as \mathbf{X} is not necessarily easy to center when it is not in a vector space.

φ -PCA includes standard PCA as a special case as when $\varphi(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_{\mathbb{F}}^2$ and the corresponding Bregman divergence becomes $D_\varphi(\mathbf{x} \| \mathbf{x}') = \frac{1}{2} \|\mathbf{x} - \mathbf{x}'\|_{\mathbb{F}}^2$.

4 Exponential family PCA on Compositional Data

CoDA has found a workaround for the centering problem, *centered log-ratio coordinates*. From [3, Def. 4.6, Chap. 8] the associated loss is the standard PCA loss on clr transformed data:

$$\ell_{\text{clr-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \doteq \frac{1}{2} \|\mathbf{c}_{\text{KL}}(\mathbf{X}) - \mathbf{V}^\top \mathbf{A}\|_{\mathbb{F}}^2 = D_\varphi(\mathbf{c}_{\text{KL}}(\mathbf{X}) \| \nabla \varphi^*(\mathbf{V}^\top \mathbf{A})), \quad (5)$$

where $\mathbf{c}_{\text{KL}}(\mathbf{X})$ is the centered log-ratio transform defined in Equation (1) and $\varphi(\mathbf{x}) = \frac{1}{2} \|\mathbf{x}\|_{\mathbb{F}}^2$. Recall from the previous section that we could deal with crude count data by using exponential family PCA. However if we wish to perform PCA on the crude count data, while maintaining the clr transform, we need an additional normalization term, which requires us to obtain a gauged version of the Bregman divergence.

4.1 Scaled Bregman Theorem with Remainder

In this section we generalize the Scaled Bregman Theorem from [25, Theorem 1] to allow for a remainder term. We use it in this paper to deal with the perspective transform required for CoDA, but it may be of independent interest. Recall that φ is the generator of the Bregman distortion (Definition 1). We additionally define a perspective (or gauge) function g to deal with the fact that we

are considering data on the simplex. Whenever φ and g are differentiable, the following is immediate from [25, Theorem 1].

Theorem 2 (Scaled Bregman Theorem with Remainder) For any $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{X} \rightarrow \mathbb{R}_*$ ($\mathbb{R}_* = \mathbb{R} \setminus \{0\}$) that are both differentiable, denoting

$$\tilde{\mathbf{x}} \doteq \frac{\mathbf{x}}{g(\mathbf{x})} \quad \text{and} \quad \tilde{\varphi}(\mathbf{x}) \doteq g(\mathbf{x}) \cdot \varphi\left(\frac{\mathbf{x}}{g(\mathbf{x})}\right), \quad (6)$$

the following holds true:

$$g(\mathbf{x}) \cdot D_\varphi(\tilde{\mathbf{x}} \parallel \tilde{\mathbf{y}}) = D_{\tilde{\varphi}}(\mathbf{x} \parallel \mathbf{y}) + R_{\varphi,g}(\mathbf{x} \parallel \mathbf{y}), \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{X}, \quad (7)$$

where $R_{\varphi,g}(\mathbf{x} \parallel \mathbf{y}) \doteq \varphi^*(\nabla\varphi(\tilde{\mathbf{y}})) \cdot D_g(\mathbf{x} \parallel \mathbf{y})$ is called the remainder.

We can abstract Theorem 2 by saying that for any φ, g differentiable, we have

$$\text{perspective-Bregman}(\varphi, g) = \text{Bregman}(\text{perspective}(\varphi)) + \text{conformal-Bregman}(g, \varphi),$$

where ‘‘perspective(φ)’’ is $\tilde{\varphi}$ in (6), and conformal divergences are defined and analyzed in [26]. General classes of perspective transforms of convex functions are introduced in [22, 23]. The notion of perspective transform of a Bregman divergence was introduced in [25]. In [25, Theorem 1], conditions are assumed that make $R_{\varphi,g}(\mathbf{x} \parallel \mathbf{y}) = 0$, resulting in the scaled Bregman theorem. Notice that D_φ is a Bregman distortion but not necessarily a Bregman divergence if φ is not convex. For reasons explained in [25], we call g a *gauge*. In the following we assume that φ is separable, so that we can use both notations $\nabla\varphi$ and φ' to denote the gradient and derivatives involving φ .

By Theorem 2, as long as $g(\mathbf{x})$ is homogeneous of degree one, $D_\varphi(\tilde{\mathbf{x}} \parallel \tilde{\mathbf{y}})$ and $\frac{1}{g(\mathbf{x})} [D_{\tilde{\varphi}}(\mathbf{x} \parallel \mathbf{y}) + R_{\varphi,g}(\mathbf{x} \parallel \mathbf{y})]$ are both invariant to re-scaling of \mathbf{x} and \mathbf{y} and can therefore be used to deal with compositional data. A general formulation of g satisfying this condition can be $g(\mathbf{x}) = \prod_{j=1}^d x_j^{w_j}$, where $\forall j, w_j \geq 0$ and $\sum_{j=1}^d w_j = 1$. In this paper, we focus on the special case $\forall j, w_j = \frac{1}{d}$ so that $D_\varphi(\tilde{\mathbf{x}} \parallel \tilde{\mathbf{y}})$ can be expressed in terms of the widely used clr transformation. Setting \mathbf{w} to be a one-hot vector $(1, 0, \dots, 0)$ can express $D_\varphi(\tilde{\mathbf{x}} \parallel \tilde{\mathbf{y}})$ with the alr. This latter case will be omitted here.

4.2 Exponential Family CoDA

We are now in a position to derive the exponential family version of the loss in (5). Let $\tilde{\mathbf{X}}$ denote the matrix of the column vectors $\tilde{\mathbf{x}}_i$. It turns out that in the same way as (2) is an approximation of (4), the loss in (5) is an approximation of the gauged loss:

$$\ell_{\text{gauged-}\varphi\text{-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \doteq D_{\varphi^*}(\mathbf{V}^\top \mathbf{A} \parallel \nabla\varphi(\tilde{\mathbf{X}})) = D_\varphi(\tilde{\mathbf{X}} \parallel \nabla\varphi^*(\mathbf{V}^\top \mathbf{A})). \quad (8)$$

Note that the above expression is in terms of the normalised matrix $\tilde{\mathbf{X}}$. To unpack it in terms of the original data \mathbf{X} , we apply Theorem 2. In the CoDA case, $\varphi^*(z) \doteq \exp z$, the convex dual of $\varphi(z) \doteq z \log z - z$. Indeed, after remarking that $\nabla\varphi(\tilde{\mathbf{X}}) = \mathbf{c}_{\text{KL}}(\mathbf{X})$, it follows

$$\begin{aligned} \ell_{\text{gauged-KL-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) &= D_{\text{exp}}(\mathbf{V}^\top \mathbf{A} \parallel \mathbf{c}_{\text{KL}}(\mathbf{X})) = D_{\text{KL}}(\tilde{\mathbf{X}} \parallel \exp(\mathbf{V}^\top \mathbf{A})) \\ &= \mathbf{1}^\top \exp(\mathbf{V}^\top \mathbf{A}) \mathbf{1} - \text{trace}(\tilde{\mathbf{X}}^\top \mathbf{V}^\top \mathbf{A}) + \text{constant}. \end{aligned} \quad (9)$$

In other words, the CoDA PCA is in fact fitting natural parameters from centered log-ratios being natural coordinates as well. From (9) we observe that both of them live in the same space. Therefore $\mathbf{V}^\top \mathbf{A}$ is centered in the same way as $\mathbf{c}_{\text{KL}}(\mathbf{X})$, and so

$$\forall \mathbf{1}_d \in \ker(\mathbf{A}^\top) \Leftrightarrow \mathbf{A}^\top \mathbf{V} \mathbf{1} = \mathbf{0}_m. \quad (10)$$

Remark that a centering assumption is also explicit in [3, Chapter 8, Eq. 8.1].

Hence, we can define the CoDA PCA problem as follows.

(CoDA PCA) We have a dataset $\mathbf{X} \in (\mathfrak{S}^d)^m$ that we approximate as $\mathbf{c}_{\text{KL}}(\mathbf{X}) \sim \mathbf{V}^\top \mathbf{A}$ by minimizing the following loss wrt the constraints $\mathbf{A} \in \mathfrak{R}^{\ell \times m}$, $\mathbf{V} \in \mathfrak{R}^{\ell \times d}$, $\mathbf{V}\mathbf{V}^\top = \mathbf{I}_\ell$, $\mathbf{A}^\top \mathbf{V}\mathbf{1} = \mathbf{0}$:

$$\ell_{\text{CoDA-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) = D_{\text{exp}}(\mathbf{V}^\top \mathbf{A} \parallel \mathbf{c}_{\text{KL}}(\mathbf{X})). \quad (11)$$

Regarding $\mathbf{c}_{\text{KL}}(\mathbf{x})$, $\tilde{\mathbf{x}}$ and \mathbf{x} as different coordinate systems of \mathfrak{S}^d , we use the Fisher information metric (FIM) [6], whose formulation is well studied on the \mathbf{x} coordinates, to define the corresponding *pullback metric* \mathcal{G} under the $\mathbf{c}_{\text{KL}}(\mathbf{x})$ and $\tilde{\mathbf{x}}$ coordinates, meaning that these metrics correspond to the same underlying geometry of \mathfrak{S}^d . We have the following proposition (proof omitted; see [30] for similar derivations).

Proposition 3 *The FIM that uniquely defines the geometry of $\mathbf{c} \in \{\mathbf{c}_{\text{KL}}(\mathbf{x}) : \mathbf{x} \in \mathfrak{S}^d\}$ is given by $\mathcal{G}_{ij}(\mathbf{c}) = \delta_{ij} \frac{\exp(c_i)}{\sum_{i=1}^d \exp(c_i)} - \frac{\exp(c_i + c_j)}{(\sum_{i=1}^d \exp(c_i))^2}$; the FIM under the coordinates $\tilde{\mathbf{x}}$ is given by $\mathcal{G}_{ij}(\tilde{\mathbf{x}}) = \delta_{ij} \frac{1}{\tilde{x}_i \sum_{i=1}^d \tilde{x}_i} - \frac{1}{(\sum_{i=1}^d \tilde{x}_i)^2}$, where $\delta_{ij} = 1$ if $i = j$ otherwise $\delta_{ij} = 0$.*

Intuitively, the metric \mathcal{G} measures the local distance $d\tilde{\mathbf{x}}^\top \mathcal{G}(\tilde{\mathbf{x}}) d\tilde{\mathbf{x}}$ of a tiny shift $d\tilde{\mathbf{x}}$. It is not everywhere identity as in a Euclidean space. Therefore the distance should not be measured by the Frobenious norm as in (5). In contrast, our loss $\ell_{\text{CoDA-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V})$ is based on the KL divergence which locally agrees with the FIM [6].

4.3 Relating CoDA PCA to φ -PCA

We now define and analyze a generalized perspective transform of the generator of KL divergence: let $\check{\text{KL}}(\mathbf{x}) \doteq g(\mathbf{x}) \cdot \sum_{j=1}^d \varphi(x_j/g(\mathbf{x}))$ where $\varphi(z) \doteq z \log(z) - z$ and $g(\mathbf{x}) \doteq (\prod_j x_j)^{1/d}$.

Lemma 4 *(Properties of $\check{\text{KL}}$) $\check{\text{KL}}$ satisfies the following properties:*

- (1) $\check{\text{KL}}$ is convex;
- (2) the general term of the Hessian \mathbf{H} of $\check{\text{KL}}$ is

$$\mathbf{H}_{ij} \doteq \mathbf{H}_{ij}(\check{\text{KL}}(\mathbf{x})) = \frac{1}{dx_j} \cdot \begin{cases} -u_{ji} & \text{if } j \neq i \\ \sum_{k \neq j} u_{kj} & \text{otherwise} \end{cases}, \quad (12)$$

where $u_{ab} \doteq 1 + x_a/x_b$. Furthermore,

$$\mathbf{z}^\top \mathbf{H} \mathbf{z} = \frac{1}{2d} \sum_{ij} (x_i + x_j) \cdot \left(\frac{z_i}{x_i} - \frac{z_j}{x_j} \right)^2, \forall \mathbf{z} \in \mathfrak{R}^d. \quad (13)$$

Hence, $\mathbf{z}^\top \mathbf{H} \mathbf{z} \geq 0, \forall \mathbf{x} \in \mathfrak{R}_{++}^d, \forall \mathbf{z} \in \mathfrak{R}^d$ and $\mathbf{z}^\top \mathbf{H} \mathbf{z} = 0$ only when $\mathbf{z} \propto \mathbf{x}$;

- (3) function $\check{\text{KL}} \circ \exp$ is 1-homogeneous on $\text{span}(\{\mathbf{1}\})^\perp$.

(Proof in SM, Section D) A consequence of Theorem 2 is the following Corollary.

Corollary 5 *For any \mathbf{A}, \mathbf{V} such that $\mathbf{A}^\top \mathbf{V}\mathbf{1} = \mathbf{0}$, we have*

$$\ell_{\text{CoDA-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \leq \tilde{\ell} \doteq \sum_i \frac{1}{g(\mathbf{x}_i)} \cdot D_{\check{\text{KL}}}(\mathbf{x}_i \parallel \exp(\mathbf{V}^\top \mathbf{a}_i)). \quad (14)$$

Hence, the CoDA PCA loss is upperbounded by a weighted generalized φ -PCA loss. Furthermore,

$$D_{\check{\text{KL}}}(\mathbf{x}_i \parallel \exp(\mathbf{V}^\top \mathbf{a}_i)) = \check{\text{KL}}(\mathbf{x}_i) - \mathbf{x}_i^\top \nabla \check{\text{KL}}(\exp(\mathbf{V}^\top \mathbf{a}_i)) \quad (15)$$

Proof Since g is concave (Example 1 in Supplement), $D_g(\mathbf{x} \parallel \mathbf{y}) = -D_{-g}(\mathbf{x} \parallel \mathbf{y}) \leq 0$, and (14) follows from Theorem 7 and the fact that $r_i \geq 0, \forall i$, which shows (14). (15) is a consequence of the analytical construct of Bregman divergences (Definition 1) and point (3) in Lemma 4 and the fact that $\mathbf{V}^\top \mathbf{a}_i \in \text{span}(\{\mathbf{1}\})^\perp$ by assumption. ■

Remark In [3, Chapter 8], CoDA PCA is presented as a (centered) regular PCA over data that been subject to *two* transforms via the centered log-ratio coordinates. What Corollary 5 shows is that we can solve the problem via a surrogate formulation using *non transformed data* but minimizing a loss which is that of a φ -PCA transformed *twice*: first taking a perspective transform of the KL generator ($\check{\text{KL}}$) and then having a weighted Bregman divergence minimization ($g^{-1}(\cdot)$). We remark that weights can also be folded in the arguments as we have:

$$\tilde{\ell} = \sum_i \text{KL}(\tilde{\mathbf{x}}_i) - \tilde{\mathbf{x}}_i^\top \nabla \check{\text{KL}}(\exp(\mathbf{V}^\top \mathbf{a}_i)) . \quad (16)$$

□

Furthermore, the leftmost argument in (16) plays no role in its minimization, and therefore we get the Surrogate CoDA PCA (s-CoDA-PCA) by replacing (11) with a simple inner product:

$$\ell_{\text{s-CoDA-PCA}}(\mathbf{X}; \mathbf{A}, \mathbf{V}) \doteq - \sum_i \tilde{\mathbf{x}}_i^\top \nabla \check{\text{KL}}(\exp(\mathbf{V}^\top \mathbf{a}_i)) . \quad (17)$$

5 Implementations

Both the CoDA-PCA in (11) and the s-CoDA-PCA in (17) can be equivalently written as the following unconstrained problems

$$\text{(CoDA-PCA)} \quad \underset{\mathbf{B}, \mathbf{U}}{\text{argmin}} \quad [\mathbf{1}_d^\top \exp(\mathbf{Y}) \mathbf{1}_m - \text{trace}(\tilde{\mathbf{X}}^\top \mathbf{Y})] , \quad (18)$$

$$\text{(s-CoDA-PCA)} \quad \underset{\mathbf{B}, \mathbf{U}}{\text{argmin}} \quad \text{trace} \left[\tilde{\mathbf{X}}^\top \left(\exp(-\mathbf{Y}) \circ \frac{\mathbf{1}_d \mathbf{1}_d^\top}{d} \exp(\mathbf{Y}) - \mathbf{Y} \right) \right] , \quad (19)$$

where “ \circ ” means element-wise product, $\exp(\cdot)$ is element-wise exponential, and $\mathbf{Y} = \mathbf{C} \mathbf{U}^\top \mathbf{B}$ with $\mathbf{C} \in \mathbb{R}^{d \times d}$, $\mathbf{U} \in \mathbb{R}^{\ell \times d}$, $\mathbf{B} \in \mathbb{R}^{\ell \times m}$. \mathbf{C} is a constant centering matrix satisfying $\text{rank}(\mathbf{C}) = d - 1$, $\mathbf{C}^\top \mathbf{1} = \mathbf{0}$, so that \mathbf{Y} ’s columns are automatically centered and $\mathbf{Y}^\top \mathbf{1} = \mathbf{B}^\top \mathbf{U} \mathbf{C}^\top \mathbf{1} = \mathbf{0}$. Any \mathbf{C} satisfying this condition corresponds to a valid re-parametrization of the feasible space, for example $\mathbf{C} = \mathbf{I}_d - \frac{1}{d} \mathbf{1}_d \mathbf{1}_d^\top$ or $\mathbf{C} = \mathbf{I}_d - \mathbf{I}_d^\#$ ($\mathbf{I}_d^\#$ circularly raises the diagonal entries of \mathbf{I}_d by 1 row). \mathbf{U} ’s rows form a nonorthogonal basis of \mathbb{R}^d . \mathbf{B} ’s columns are the sample coordinates in such a basis. After optimization, we take the QR decomposition $\mathbf{C}(\mathbf{U}^*)^\top = (\mathbf{V}^*)^\top \mathbf{T}^*$, where \mathbf{V}^* ’s rows are orthonormal. Therefore $\mathbf{C}(\mathbf{U}^*)^\top \mathbf{B}^* = (\mathbf{V}^*)^\top \mathbf{T}^* \mathbf{B}^*$ and $\mathbf{A}^* = \mathbf{T}^* \mathbf{B}^*$ is the corresponding coordinates. An optimal solution of the original constrained PCA problem is given by $(\mathbf{V}^*, \mathbf{A}^*)$.

Although the losses in (18) and (19) are non-convex, they are both *bi-convex*. Fixing \mathbf{U} , the loss is a strictly convex function of \mathbf{B} that is decomposed into a sum of per-sample convex functions of \mathbf{b}_i ; fixing \mathbf{B} , it is a strictly convex function of \mathbf{U} . These convex functions have the general form $f(\boldsymbol{\xi}) = \sum_i \exp(\boldsymbol{\alpha}_i^\top \boldsymbol{\xi} + \beta_i) + \boldsymbol{\zeta}^\top \boldsymbol{\xi}$. Its gradient and Hessian are both in simple closed form: $\nabla f = \sum_i \exp(\boldsymbol{\alpha}_i^\top \boldsymbol{\xi} + \beta_i) \boldsymbol{\alpha}_i + \boldsymbol{\zeta}$; $\nabla^2 f = \sum_i \exp(2\boldsymbol{\alpha}_i^\top \boldsymbol{\xi} + 2\beta_i) \boldsymbol{\alpha}_i \boldsymbol{\alpha}_i^\top$. One can apply an off-the-shelf convex optimizer, which in the simplest case can be the Newton method, to alternately minimize \mathbf{B} and \mathbf{U} until convergence. Our implementation simply uses L-BFGS [9] based on the gradient of the loss. In summary, we have the following result.

Proposition 6 *The CoDA-PCA and the s-CoDA-PCA are both equivalent to an unconstrained bi-convex optimization problem.*

As an alternative implementation, we assume a parametric mapping $\mathbf{b}_i = g_\Theta(\mathbf{x}_i)$ that is the ℓ -dimensional output of a feed-forward neural network with input $\mathbf{c}_{\text{KL}}(\mathbf{x}_i)$ and \mathbf{x}_i (or $\tilde{\mathbf{x}}_i$) and connection weights Θ . Then we minimize the cost function in (18) with respect to \mathbf{U} and Θ . If g_Θ is flexible enough, then the minimization recovers the CoDA-PCA projection. This approach could be favored as ① it learns an out-of-sample mapping $g_\Theta(\cdot)$ with a compact parametric structure that does not scale with the sample size m ; and ② it can be adapted to an online learning scenario. However, it requires tuning of the neural network architecture and the optimizer. In our experiments, the encoding map is modeled by a feed-forward neural network with two hidden layers of ELU [11] units, each of size 100. To distinguish between the two implementations, the method to directly optimize \mathbf{U} and \mathbf{B}

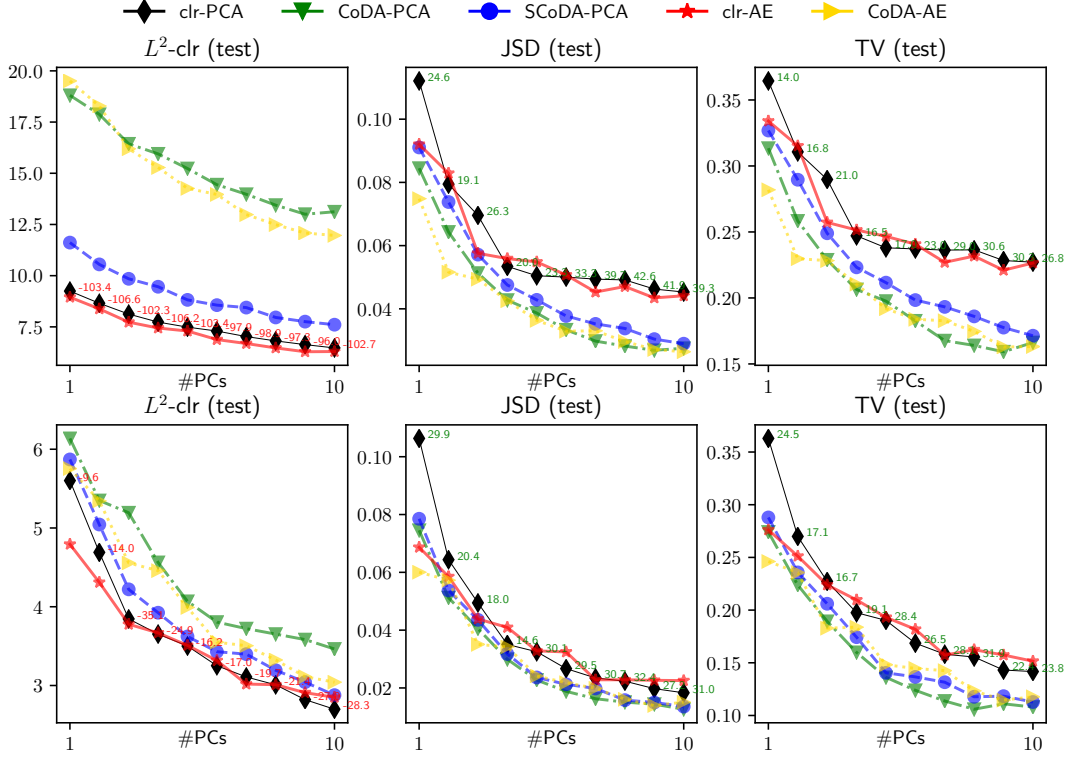


Figure 2: Testing errors (y-axis) against the number of principal components (x-axis) based on three different distance measures (from left to right) on the Atlas data (first row) and the diet swap data (second row). The numbers along the `clr-PCA` curves show the percentage of improvement (green) or disimprovement (red), comparing `CoDA-PCA` against `clr-PCA`.

without assuming the neural network mapping is called non-parametric CODA-PCA, and the latter parametric version is simply called CODA-PCA.

The above implementation resembles an auto-encoder structure: $\mathbf{x}_i \xrightarrow{\Theta} \mathbf{b}_i \xrightarrow{U} \mathbf{y}_i$, where the decoder is simply a linear mapping $\mathbf{y}_i = \mathbf{C}U^T\mathbf{b}_i$. In the general case, we apply a non-linear decoder $\mathbf{b}_i \xrightarrow{\Phi} \mathbf{y}_i$ in the form $\mathbf{y}_i = \mathbf{C}h_{\Phi}(\mathbf{b}_i)$, where $h_{\Phi}(\cdot)$ is a neural network with parameters Φ and d output dimensions. At the same time, we add a small random noise to the encoder input so as to avoid overfitting. In this way we obtain a denoising CODA-AUTOENCODER. In contrast to the CODA-PCA, the CODA-AUTOENCODER can only be trained by gradient-based optimizers.

In practice, the input matrix \mathbf{X} may contain zeros that lie on the boundary of \mathfrak{S}^d . In this case $\mathbf{c}_{\text{KL}}(\mathbf{x})$ and $\tilde{\mathbf{x}}$ are undefined. A simple way to tackle the zero entries is to replace them with a small positive number $\epsilon > 0$. Alternatively, one can redefine the gauge as $g(\mathbf{x}) = \prod_{j: x_j > 0} (x_j)^{1/\rho}$, where $\rho = |\{j : x_j > 0\}|$ so that $g(\mathbf{x})$ is always positive and $\tilde{\mathbf{x}}$ is well defined on $\mathfrak{S}^d \cup \partial\mathfrak{S}^d$.

6 Experiments

We compare the following methods: `clr-PCA` means PCA applied on the centered log-ratio coordinates; `CoDA-PCA` is the proposed CODA-PCA in (11); `SCoDA-PCA` is the proposed s-CODA-PCA in (17); `clr-AE` is an autoencoder with L^2 loss applied on the `clr` transformation; `CoDA-AE` is the proposed CODA-AUTOENCODER in subsection 5. Both `clr-AE` and `CoDA-AE` use exactly the same structure with one hidden layer of 100 ELU [11] units in their decoders.

The baselines are assessed based on an array of measures including (L^2 -clr) the L^2 -distance $\|\mathbf{c}_{\text{KL}}(\mathbf{x}) - \mathbf{c}_{\text{KL}}(\mathbf{x}')\|_F$ between the input data $\mathbf{x} \in \mathfrak{S}^d$ and the reconstruction $\mathbf{x}' \in \mathfrak{S}^d$ in the `clr` space; (JSD) the Jensen-Shannon divergence $\frac{1}{2}\text{KL}(\mathbf{x} : \frac{\mathbf{x} + \mathbf{x}'}{2}) + \frac{1}{2}\text{KL}(\mathbf{x}' : \frac{\mathbf{x} + \mathbf{x}'}{2})$; (TV) the total variation distance

$\frac{1}{2} \sum_{i=1}^d |x_i - x'_i|$. These measurements are all invariant to scaling or permutation of x and x' . See the supplementary material for more baselines and performance indicators.

We consider the following datasets available in the microbiome R package [18], each of which is randomly split into a training set (90%) and a testing set (10%). *The HITChip Atlas dataset* [17] contains 130 genus-level taxonomic groups that cover the majority of the known bacterial diversity of the human intestine. The data come from 1006 western adults from 15 western countries (Europe and the United States). Sample sets were analysed with three different DNA extraction methods. *The two-week diet swap study* between western (USA) and traditional (rural Africa) diets was reported in [28]. In this study, a two-week food exchange was performed in subjects from the same populations, where African Americans were fed a high-fibre, low-fat African-style diet and rural Africans a high-fat, low-fibre western-style diet. The group diet was indicated by HE (home environment days), DI (dietary intervention days) and ED (initial and final endoscopy days). Each subject served as his/her own control, given the known wide individual variation in colonic microbiota composition.

Fig. 2 shows the typical testing results. We observe that on most performance indicators CoDA-PCA and CoDA-AE show a much smaller testing error as compared to clr-PCA and clr-AE, respectively. The only exception is on L^2 -clr, where clr-PCA and clr-AE appear to be favored against our CoDA variants. This is because L^2 -clr is exactly the cost function of those two methods. We found that CoDA-AE is more robust against overfitting as compared to clr-AE. The performance of SCoDA-PCA is close to CoDA-PCA on most of the indicators and is better than CoDA-PCA on L^2 -clr.

The source codes to reproduce our experimental results are available online².

7 Conclusion

We propose an approach for learning a low dimensional representation directly on raw count data, which is compositional in nature. Our proposed algorithm generalizes PCA in two ways, first by going to the exponential family via the Bregman divergence, and second by converting the normalization of data to a change in the Bregman divergence. The key theorem used for transforming the Bregman divergence generalizes a recent result, and may be of independent interest.

Acknowledgements

The authors gratefully thank Perrine Soret, Frank Nielsen, Xinhua Zhang, and the anonymous NIPS reviewers, for their helpful and constructive feedback. This work was done while MAF was visiting Data61, CSIRO in Canberra, Australia.

References

- [1] J. Aitchison. The statistical analysis of compositional data (with discussion). *Journal of the Royal Statistical Society B*, 44(2):139–177, 1982.
- [2] J. Aitchison. Principal component analysis of compositional data. *Biometrika*, 70(1):57–65, 1983.
- [3] J. Aitchison. *The Statistical Analysis of Compositional Data*. Chapman and Hall, New York, 1986.
- [4] J. Aitchison. Principles of compositional data analysis. *Multivariate Analysis and its Applications*, 24:73–81, 1994.
- [5] J. Aitchison and J.-J. Egozcue. Compositional data analysis: Where are we and where should we be heading? *Mathematical Geology Journal*, 37:829–850, 2005.
- [6] S.-I. Amari. *Information Geometry and Its Applications*. Springer-Verlag, Berlin, 2016.
- [7] O. Barndorff-Nielsen. *Information and Exponential Families in Statistical Theory*. Wiley Publishers, 1978.

²<https://bitbucket.org/RichardNock/coda>

- [8] J.-D. Boissonnat, F. Nielsen, and R. Nock. Bregman Voronoi diagrams. *Discrete Comput. Geom.*, 44(2):281–307, 2010.
- [9] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [10] J. Chiquet, M. Mariadassous, and Stéphane Robin. Variational inference for probabilistic Poisson PCA. *Annals of Applied Statistics (to appear)*, 2018.
- [11] D.-A. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by exponential linear units (ELUs). In *4th ICLR*, 2016.
- [12] M. Collins, S. DasGupta, and R. Schapire. A generalization of principal components analysis to the exponential family. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *NIPS*15*, 2002.
- [13] J.-J. Egozcue, V. Pawlowsky-Glahn, G. Mateu-Figueras, and C. Barceló-Vidal. Isometric logratio transformations for compositional data analysis. *Mathematical Geology Journal*, 35:279–300, 2003.
- [14] G. Reid G. B. Gloor. Compositional analysis: a valid approach to analyze microbiome high-throughput sequencing data. *Can J Microbiol.*, 12:1–12, 2016.
- [15] G. B. Gloor, J. R. Wu, V. Pawlowsky-Glahn, and J. J. Egozcue. It’s all relative: analyzing microbiome data as compositions. *Ann Epidemiol.*, 26:322–9, 2016.
- [16] M. Greenacre. *Compositional Data Analysis in Practice*. Chapman and Hall, New York, 2018.
- [17] L. Lahti, J. Salojärvi, A. Salonen, M. Scheffer, and W. M. de Vos. Tipping elements in the human intestinal ecosystem. *Nat. Commun.*, 5:4344, 2014.
- [18] L. Lahti, S. Sudarshan, T. Blake, and J. Salojarvi. Microbiome R package. version 1.1.10012, 2017.
- [19] A. J. Landgraf. *Generalized Principal Component Analysis: Dimensionality Reduction through the Projection of Natural Parameters*. PhD thesis, Ohio State University, 2015.
- [20] D. Lovell, W. Müller, J. Taylor, A. Zwart, and C. Helliwell. Caution! compositions! can constraints on omics data lead analyses astray? Technical Report EP10994, CSIRO, 2010.
- [21] A.F. Andersson L.W. Hugerth. Analysing microbial community composition through amplicon sequencing: From sampling to hypothesis testing. *Frontiers in Microbiology*, 8:1561, 2017.
- [22] P. Maréchal. On a functional operation generating convex functions, part 1: duality. *J. of Optimization Theory and Applications*, 126:175–189, 2005.
- [23] P. Maréchal. On a functional operation generating convex functions, part 2: algebraic properties. *J. of Optimization Theory and Applications*, 126:375–366, 2005.
- [24] J. A. Martín-Fernández, V. Pawlowsky-Glahn, J. J. Egozcue, and R. Tolosona-Delgado. The statistical analysis of compositional data (with discussion). *Math Geosci*, 50:273–298, 2018.
- [25] R. Nock, A.-K. Menon, and C.-S. Ong. A scaled Bregman theorem with applications. In *NIPS*29*, pages 19–27, 2016.
- [26] R. Nock, F. Nielsen, and S.-I. Amari. On conformal divergences and their population minimizers. *IEEE Trans. IT*, 62:1–12, 2016.
- [27] V. Shankar O. Paliy. Application of multivariate statistical techniques in microbial ecology. *Molecular ecology*, 25:1032–57, 2016.
- [28] S. J. D. O’Keefe, J. V. Li, L. Lahti, J. Ou, F. Carbonero, K. Mohammed, J. M. Posma, J. Kinross, E. Wahl, E. Ruder, K. Vippera, V. Naidoo, L. Mtshali, S. Tims, P. G. B. Puylaert, J. DeLany, A. Krasinskas, A. C. Benefiel, H. O. Kaseb, K. Newton, J. K. Nicholson, W. M. de Vos, H. R. Gaskins, and E. G. Zoetendal. Fat, fiber and cancer risk in African Americans and rural Africans. *Nat. Commun.*, 6:6342, 2015.

- [29] V. Pawłowsky-Glahn and A. Buccianti. *Compositional Data Analysis, theory and applications*. Wiley, 2011.
- [30] Ke Sun and Stéphane Marchand-Maillet. An information geometry of statistical manifold learning. In *31st ICML*, pages 1–9, 2014.
- [31] R. Tolosona-Delgado V. Pawłowsky-Glahn, J. J. Egozcue. Lecture notes on compositional data analysis. Technical report, Girona University, 2007.
- [32] L. van der Maaten and G. E. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.