

A folded model for compositional data analysis

Michail Tsagris and Connie Stewart

Department of Economics, University of Crete, Rethymnon, Greece,
mtsagris@uoc.gr

Department of Mathematics and Statistics, University of New
Brunswick, Saint John, Canada, connie.stewart@unb.ca

Abstract

A folded type model is developed for analyzing compositional data. The proposed model, which is based upon the α -transformation for compositional data, provides a new and flexible class of distributions for modeling data defined on the simplex sample space. Despite its rather seemingly complex structure, employment of the EM algorithm guarantees efficient parameter estimation. The model is validated through simulation studies and examples which illustrate that the proposed model performs better in terms of capturing the data structure, when compared to the popular logistic normal distribution.

Keywords: α -transformation, EM-algorithm, compositional data, folding transformation

1 Introduction

Compositional data are positive multivariate data which sum to the same constant, usually 1. In this case, their sample space is the standard simplex

$$\mathbb{S}^{D-1} = \left\{ (x_1, \dots, x_D)^T \mid x_i \geq 0, \sum_{i=1}^D x_i = 1 \right\}, \quad (1)$$

where D denotes the number of variables (better known as components).

Compositional data are met in many different scientific fields. In sedimentology, for example, samples were taken from an Arctic lake and their composition of water, clay and sand were the quantities of interest. Data from oceanography studies involving Foraminiferal (a marine plankton species) compositions at 30 different sea depths from oceanography were analyzed in Aitchison (2003, pg 399). Schnute & Haigh (2007) also analyzed marine compositional data through catch curve models for a quillback rockfish (*Sebastes maliger*) population. In hydrochemistry, Otero et al. (2005) used regression analysis to draw conclusions about anthropogenic and geological pollution sources of rivers in Spain. Stewart & Field (2011) modeled compositional diet estimates with an abundance of zero values obtained through quantitative fatty acid signature analysis. In another biological setting, Ghosh & Chakrabarti (2009) were interested was in the classification of immunohistochemical data. Other applications areas of compositional data analysis include archaeometry (Baxter et al., 2005), where the composition of ancient glasses, for instance, is of interest, and in economics (Fry, Fry, & McLaren, 2000), where the focus is on the percentage of the household expenditure allocated to different products. Compositional data are also met in political science (Katz, & King, 1999) for modeling

electoral data and in forensic science where the compositions of forensic glasses are compared and classified (Neocleous, Aitken, & Zadora, 2011). In demography, compositional data are met in multiple-decrement life tables and the mortality rates amongst age groups are modeled (Oeppen, 2008). In a study of the brain, Prados et al. (2010) evaluated the diffusion anisotropy from diffusion tensor imaging using new measures derived from compositional data distances. Some recent areas of application include bioinformatics and specifically microbiome data analysis (Xia et al., 2013; Chen & Li, 2016; Shi, Zhang & Li, 2016). These examples illustrate the breadth of compositional data analysis applications and consequently the need for parametric models defined on the simplex.

The Dirichlet distribution is a natural distribution for such data due to its support being the simplex space. However, it has long been recognized that this distribution is not statistically rich and flexible enough to capture many types of variabilities (especially curvature) of compositional data and, for this reason, a variety of transformations have been proposed that map the data outside of the simplex. In Aitchison (1982) the log-ratio transformation approach was developed, and later the so called isometric log-ratio transformation methodology which was first proposed in Aitchison (2003, pg 90) and examined in detail by Egozcue et al. (2003). More recently, Tsagris, Preston & Wood (2011) suggested the α -transformation which includes the isometric transformation as a special case. The α -transformation is a Box-Cox type transformation and has been successfully applied in regression analysis (Tsagris, 2015) and classification settings (Tsagris, Preston & Wood, 2016).

Regardless of the transformation chosen, the usual approach for modeling compositional data is to assume that the transformed data is multivariate normally distributed. While the α -transformation offers flexibility, a disadvantage of this transformation is that it maps the compositional data from the simplex (\mathbb{S}^{D-1}) to a subset of \mathbb{R}^{D-1} , and not \mathbb{R}^{D-1} itself on which the multivariate normal density is defined. An improvement to this method can be obtained by using a folded model procedure, similar to the approach used in Scealy & Welsh (2014) and to the folded normal distribution in \mathbb{R} employed in Leone, Nelson & Nottingham (1961) and Johnson (1962). The folded normal distribution in \mathbb{R} , for example, corresponds to taking the absolute value of a normal random variable and essentially “folding” the negative values to the positive side of the distribution. The model we propose here works in a similar fashion where values outside the simplex are mapped inside of it via a folding type transformation. An advantage of this approach over the aforementioned log-ratio methodology is that it allows one to fit any suitable multivariate distribution on \mathbb{S}^{D-1} through the parameter α .

The paper is structured as follows. In Section 2 we describe the folding approach and in Section 3 we introduce the α -folded multivariate normal model for compositional data, along with the EM algorithm for maximum likelihood estimation (MLE), an algorithm for generating data from the proposed folded distribution, inference for α and an approach for principal component analysis. Simulation studies are conducted in Section 4 to assess the estimation accuracy of the parameters and two data sets are analyzed to further assess the performance of our model. Finally, concluding remarks may be found in Section 5.

2 The α folding technique

2.1 The α -transformation

For a composition $\mathbf{x} \in \mathbb{S}^{D-1}$, the centered log-ratio transformation is defined in [2] as

$$w_i^0(\mathbf{x}) = \log \left(\frac{x_i}{\prod_{j=1}^D x_j^{1/D}} \right), \text{ for } i = 1, \dots, D. \quad (2)$$

The sample space of Equation (2) is the set

$$\mathbb{Q}_0^{D-1} = \left\{ (w_1^0, \dots, w_D^0)^T : \sum_{i=1}^D w_i^0 = 0 \right\} \quad (3)$$

which is a subset of \mathbb{R}^{D-1} . Note that the zero sum constraint in Equation (2) is an obvious drawback of this transformation as it can lead to singularity issues. In order to remove the redundant dimension imposed by this constraint, one can apply the isometric log-ratio transformation

$$\mathbf{z}_0(\mathbf{x}) = \mathbf{H}\mathbf{w}_0(\mathbf{x}), \quad (4)$$

where \mathbf{H} is the Helmert matrix (Lancaster, 1965) (an orthonormal $D \times D$ matrix) after deletion of the first row. The left multiplication by the Helmert sub-matrix maps the data onto \mathbb{R}^{D-1} thus, in effect, removing the zero sum constraint. The Helmert sub-matrix is a standard orthogonal matrix in shape analysis used to overcome singularity problems (Dryden & Mardia, 1998;; Le & Small, 1999).

Tsagris, Preston & Wood (2011) developed the α -transformation as a more general transformation than that in Equation (4). Let

$$\mathbf{u}_\alpha(\mathbf{x}) = \left(\frac{x_1^\alpha}{\sum_{j=1}^D x_j^\alpha}, \dots, \frac{x_D^\alpha}{\sum_{j=1}^D x_j^\alpha} \right)^T \quad (5)$$

denote the power transformation for compositional data as defined by Aitchison (2003). In a manner analogous to Equations (2-4), first define

$$\mathbf{w}_\alpha(\mathbf{x}) = \frac{D\mathbf{u}_\alpha - \mathbf{1}}{\alpha}. \quad (6)$$

The sample space of Equation (6) is then the set

$$\mathbb{Q}_\alpha^{D-1} = \left\{ (w_{1,\alpha}, \dots, w_{D,\alpha})^T : \frac{-1}{\alpha} \leq w_{i,\alpha} \leq \frac{D-1}{\alpha}, \sum_{i=1}^D w_{i,\alpha} = 0 \right\}. \quad (7)$$

Note that the inverse of Equation (6) is as follows

$$\mathbf{x} = \mathbf{w}_\alpha^{-1}(\mathbf{m}) = \frac{(1 + \alpha m_i)^{1/\alpha}}{\sum_{j=1}^D (1 + \alpha m_j)^{1/\alpha}}, \text{ for } i = 1, \dots, D \quad (8)$$

for $\mathbf{m} \in \mathbb{Q}_\alpha^{D-1}$. Finally, the α -transformation is defined as

$$\mathbf{z}_\alpha(\mathbf{x}) = \mathbf{H}\mathbf{w}_\alpha(\mathbf{x}). \quad (9)$$

The transformation in Equation (9) is a one-to-one transformation which maps data inside the simplex onto a subset of \mathbb{R}^{D-1} and vice versa for $\alpha \neq 0$. The corresponding sample space of Equation (9) is

$$\mathbb{A}_\alpha^{D-1} = \left\{ \mathbf{H}\mathbf{w}_\alpha \mid -\frac{1}{\alpha} \leq w_{i,\alpha} \leq \frac{D-1}{\alpha}, \sum_{i=1}^D w_{i,\alpha} = 0 \right\}. \quad (10)$$

For $\mathbf{y} = \mathbf{z}_\alpha(\mathbf{x})$, the inverse transformation from \mathbb{A}_α^{D-1} to \mathbb{S}^{D-1} is given by $\mathbf{z}_\alpha^{-1}(\mathbf{y}) = \mathbf{w}_\alpha^{-1}(\mathbf{H}^T \mathbf{y})$ where $\mathbf{w}_\alpha^{-1}(\cdot)$ is given in Equation (8).

Note that vectors in \mathbb{A}_α^{D-1} are not subject to the sum to zero constraint and that $\lim_{\alpha \rightarrow 0} \mathbb{A}_\alpha^{D-1} \rightarrow \mathbb{R}^{D-1}$.

For convenience purposes we allow α to lie within $[-1, 1]$. From Equations (5) and (6), when $\alpha = 1$, the simplex is linearly expanded as the values of the components are simply multiplied by a scalar and then centered. When $\alpha = -1$, the inverse of the values of the components are multiplied by a scalar and then centered.

If we assume that the α -transformed data (for any value of $\alpha \in [-1, 1]$) follow a multivariate normal distribution, then a way to choose the value of α is via maximum likelihood estimation. However, given that the multivariate normal distribution is defined on \mathbb{R}^{D-1} and that $\mathbb{A}_\alpha^{D-1} \subseteq \mathbb{R}^{D-1}$, this approach might neglect an important amount of volume (probability) of the multivariate normal. The same problem arose in Scealy & Welsh (2011b), who then developed the folded Kent distribution (Scealy & Welsh, 2014) and we propose a similar solution here.

2.2 The folding transformation

Let $\mathbf{y} = \mathbf{z}_\alpha(\mathbf{x})$ in Equation (9) for $\mathbf{x} \in \mathbb{S}^{D-1}$ and some value of α , then the inverse transformation $\mathbf{z}_\alpha^{-1}(\mathbf{y})$ provides a transformation from $\mathbf{y} \in \mathbb{A}^{D-1}$ to \mathbb{S}^{D-1} . Now suppose we have a point $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1}$ (that is, outside of \mathbb{A}_α^{D-1}) and we want to map it inside the simplex \mathbb{S}^{D-1} . It can be shown (Tsagris, 2013) that the following transformation maps \mathbf{y} to \mathbb{S}^{D-1}

$$\mathbf{x} = \mathbf{w}_\alpha^{-1} \left(\frac{\mathbf{H}^T \mathbf{y}}{q_\alpha^{*2}(\mathbf{y})} \right), \quad (11)$$

where $\mathbf{w}_\alpha^{-1}(\cdot)$ is defined in Equation (8) and $q_\alpha^*(\mathbf{y}) = \alpha \min \{\mathbf{H}^T \mathbf{y}\}$.

Note that the inverse of Equation (11) is the $D - 1$ dimensional vector

$$\mathbf{y} = \frac{1}{w_\alpha^{*2}(\mathbf{x})} \mathbf{H}\mathbf{w}_\alpha(\mathbf{x}) = \frac{1}{w_\alpha^{*2}(\mathbf{x})} \mathbf{z}_\alpha(\mathbf{x}), \quad (12)$$

where $w_\alpha^*(\mathbf{x}) = \alpha \min \{\mathbf{w}_\alpha(\mathbf{x})\}$ and $\mathbf{w}_\alpha(\mathbf{x})$ is defined in Equation (6).

In summary, we propose the following folding transformation from $\mathbf{y} \in \mathbb{R}^{D-1}$ to \mathbb{S}^{D-1}

$$\mathbf{x} = \begin{cases} g_0^\alpha(\mathbf{y}) & \text{if } \mathbf{y} \in \mathbb{A}^{D-1} \\ g_1^\alpha(\mathbf{y}) & \text{if } \mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1} \end{cases} \quad (13)$$

where $g_0^\alpha(\mathbf{y}) = \mathbf{z}_\alpha^{-1}(\mathbf{y}) = \mathbf{w}_\alpha^{-1}(\mathbf{H}^T \mathbf{y})$ and $g_1^\alpha(\mathbf{y}) = \mathbf{w}_\alpha^{-1} \left(\frac{\mathbf{H}^T \mathbf{y}}{q_\alpha^{*2}(\mathbf{y})} \right)$.

3 The α -folded multivariate normal distribution

3.1 Definition

The distribution of \mathbf{X} in Equation (13) can be derived as a mixture distribution if we assume that \mathbf{Y} is multivariate normally distributed with parameters $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{\Sigma}_\alpha$, that is if $\mathbf{Y} \sim N^{D-1}(\boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha)$. Motivated by the folded normal distribution in Leone (1961), we define the distribution of \mathbf{X} on \mathbb{S}^{D-1} and refer to it as the α -folded multivariate normal distribution. The term α -folded is a combination of the α -transformation and the folding type approach that is applied.

The distribution of $\mathbf{X} \in \mathbb{S}^{D-1}$ can be written as

$$f(\mathbf{x}|\alpha, p) = pf_0(\mathbf{x}|\alpha) + (1-p)f_1(\mathbf{x}|\alpha), \quad (14)$$

where f_0 and f_1 refer to the densities of $g_0^\alpha(\mathbf{y})$ and $g_1^\alpha(\mathbf{y})$ respectively when $\mathbf{Y} \sim N^{D-1}(\boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha)$, and p is the probability that $Y \in \mathbb{A}^{D-1}$.

To specify the density $f(\mathbf{x}|\alpha, p)$, we require the relevant Jacobians of the transformations $g_0^\alpha(\mathbf{y})$ and $g_1^\alpha(\mathbf{y})$. These are given in the following two Lemmas with corresponding proofs in the Appendix.

Lemma 3.1 *The Jacobian of $g_0^\alpha(\mathbf{y}) = \mathbf{z}_\alpha^{-1}(\mathbf{y})$ is*

$$|J_\alpha^0| = D^{D-1+\frac{1}{2}} \prod_{i=1}^D \frac{x_i^{\alpha-1}}{\sum_{j=1}^D x_j^\alpha} \quad (15)$$

Lemma 3.2 *The Jacobian of $g_1^\alpha(\mathbf{y}) = \mathbf{w}_\alpha^{-1}\left(\frac{\mathbf{H}^T \mathbf{y}}{q_\alpha^{*2}(\mathbf{y})}\right)$ is*

$$|J_\alpha^1| = \left(\frac{1}{\alpha w_\alpha^*(\mathbf{x})}\right)^{2(D-1)}.$$

From the two lemmas, it follows that

$$\begin{aligned} f(\mathbf{x}|\alpha, p, \boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha) &= p \frac{|J_\alpha^0|}{|2\pi\boldsymbol{\Sigma}_\alpha|^{1/2}} \exp\left[-\frac{1}{2}(z_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha)^T \boldsymbol{\Sigma}_\alpha^{-1}(z_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha)\right] \\ &+ (1-p) \frac{|J_\alpha^1|}{|2\pi\boldsymbol{\Sigma}_\alpha|^{1/2}} \exp\left[-\frac{1}{2}\left(\frac{z_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha\right)^T \boldsymbol{\Sigma}_\alpha^{-1}\left(\frac{z_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha\right)\right], \end{aligned} \quad (16)$$

where $\mathbf{x} \in S^{D-1}$, $\alpha \in [-1, 1]$, $0 \leq p \leq 1$ and $z_\alpha(\mathbf{x})$ is given in Equation (9). We will refer to the density in Equation 16 as the α -folded multivariate normal distribution.

Although we have indicated that $\mathbf{x} \in S^{D-1}$, it is important to note that boundaries on the simplex (that is, zero values) are not allowed due to the log of $|J_\alpha^0|$ being undefined in such cases. This potential limitation also occurs with conventional transformations for compositional data analysis as well, including the isometric transformation in (Equation 4). Recent approaches to handle zero values include mapping the data onto the hyper-sphere (Scealy and Welsh, 2011b, Scealy and Welsh), assuming latent variables (Butler & Glasbey, 2008) or conditioning on the zero values (Stewart & Field 2011; Tsagris & Stewart, 2018).

A few special cases are worthy of mention. First, when $p = 1$, the second term in Equation (16) vanishes and we end up with the α -normal distribution [41]. When $\alpha = 1$, $|J_\alpha^0| = D^{D-1+\frac{1}{2}}$.

Finally, when $\alpha \rightarrow 0$, Equation (16) reduces to the multivariate logistic normal on \mathbb{S}^d ([3])

$$f(\mathbf{x}) = (2/\pi)^{-(D-1)/2} |\boldsymbol{\Sigma}|^{-1/2} \times \exp \left[-\frac{1}{2} (z_0(\mathbf{x}) - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}_0^{-1} (z_0(\mathbf{x}) - \boldsymbol{\mu}_0) \right] \prod_{i=1}^D x_i^{-1}, \quad (17)$$

where $z_0(\mathbf{x})$ is defined in Equation (4). The density function in Equation (17) is that of the logistic normal on \mathbb{S}^d . This leads us to the following Definition which generalizes Aitchison's (2003) Definition 6.2.

A D-part composition \mathbf{X} is said to follow the α -folded multivariate normal distribution if \mathbf{Y} follows a multivariate normal distribution in \mathbb{R}^d where

$$\mathbf{y} = \begin{cases} \mathbf{z}_\alpha(\mathbf{x}) & \text{with probability } p \\ \frac{\mathbf{z}_\alpha(\mathbf{x})}{w_\alpha^*(\mathbf{x})}, & \text{with probability } 1-p, \end{cases} \quad (18)$$

$w_\alpha^*(\mathbf{x}) = \alpha \min \{\mathbf{w}_\alpha(\mathbf{x})\}$ and $\mathbf{z}_\alpha(\mathbf{x})$ is defined in Equation (9).

The generic power transformation Equation (9) includes the isometric log-ratio transformation in Equation (4) and thus the same concepts can be used to build the model. This means that once the data are transformed via Equation 18, a multivariate normality test can be applied. However, numerical optimization is required for obtaining the maximum likelihood estimates of the mean vector and covariance matrix, as described in Section 3.3.

The α -folded multivariate normal distribution resolves the problem associated with the assumption of multivariate normality of the α -transformed data (Tsagris, Preston & Wood, 2011); the ignored probability left outside the simplex due to the sample space of the transformation being a subset of R^d . With this distribution, any ignored probability is folded back onto the simplex, and hence the density has the form of Equation (16).

3.2 Contour plots of the α -folded bivariate normal distribution

We now consider a visual illustration of the bivariate normal distribution with and without folding. In particular, we examine contours of the normal distribution in \mathbb{R}^2 and compare these to the contours of the folded model in Equation 16 plotted on the simplex, with $\alpha = 1$. We consider two settings in which $\boldsymbol{\mu} = (0.561, 0.547)^T$ in both cases, but use two different covariances matrices as follows

$$\boldsymbol{\Sigma}_1 = \begin{pmatrix} 0.5 & 0.25 \\ 0.25 & 0.35 \end{pmatrix} \text{ and } \boldsymbol{\Sigma}_2 = \begin{pmatrix} 2.5 & 1.25 \\ 1.25 & 1.75 \end{pmatrix}. \quad (19)$$

The mean vector was selected by applying the α -transformation in Equation (9) with $\alpha = 1$ to a sub-composition formed by the first three components of the Hongite data (artificial dataset used by Aitchison, 2003). The elements of the first covariance matrix ($\boldsymbol{\Sigma}_1$) were chosen so that the correlation was positive, whereas the second covariance matrix is such that $\boldsymbol{\Sigma}_2 = 5\boldsymbol{\Sigma}_1$.

For each combination of parameters we calculated the density of the normal and the folded normal for a grid of two-dimensional vectors and then plotted their contours. While for the unconstrained normal case, the density was simply calculated for all grid points, for the folded model, the transformation in Equation 13 was first applied (with $\alpha = 1$) and then the density in Equation 16 was calculated. The contour plots are presented in Figure 1.

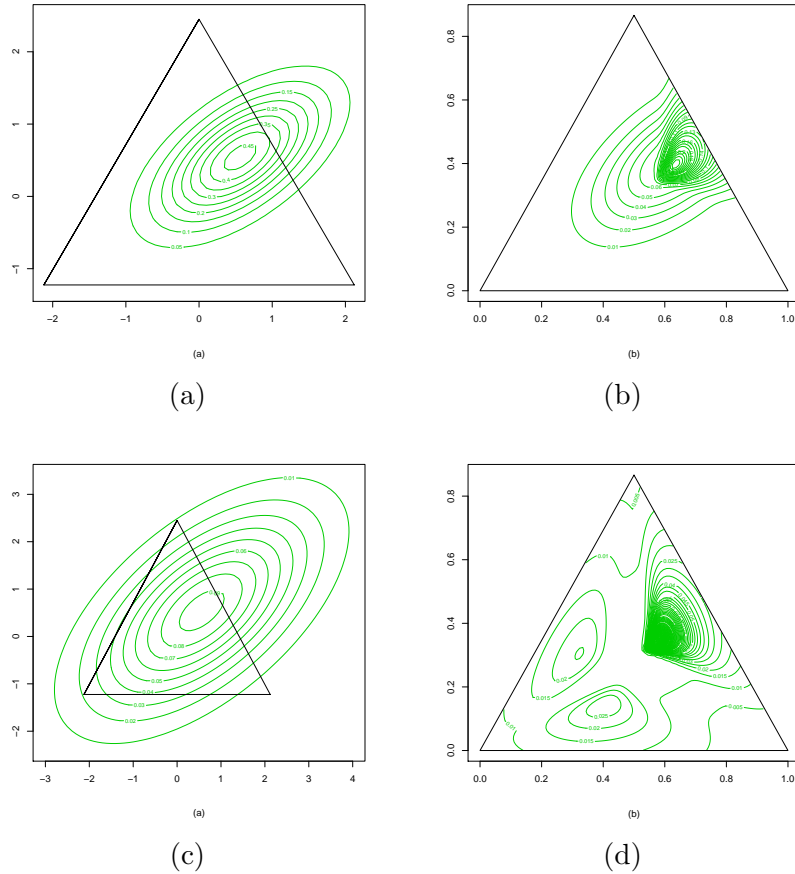


Figure 1: All contour plots refer to a normal distribution with mean $\boldsymbol{\mu} = (0.561, 0.547)^T$. The covariance matrix of the first row is $\boldsymbol{\Sigma}_1$ and of the second row is $\boldsymbol{\Sigma}_2 = 5\boldsymbol{\Sigma}_1$. The left plot is the normal in the real space and the triangle (the simplex after the α -transformation in Equation (9) with $\alpha = 1$) is for illustration purposes. The left column shows the contour plots of the multivariate normal in \mathbb{R}^2 and the right column shows the contours of the α -folded normal on the simplex.

Figures (a) and (c) depict the contours of the multivariate normal distribution (defined in \mathbb{R}^2) while Figures (b) and (d) show the contours of the α -folded normal (defined on the simplex). The first row corresponds to $\boldsymbol{\Sigma}_1$ in Equation (19) while the second row is derived from $\boldsymbol{\Sigma}_2$. The triangles in the second column (Figures (b) and (d)) display the simplex while the corresponding triangles in Figures (a) and (c) were obtained through the α transformation in Equation 9.

What is perhaps most evident from the contour plots is that points falling outside the triangle in Figures (a) and (c) result in modes inside the simplex in Figures (b) and (d) respectively. When there is a high probability of being left outside of two or more sides of the triangle (or faces of the pyramid or hyper-pyramid in higher dimensions), as in Figures 1(c) and 1(d), the contours of the folding model will have a somewhat peculiar shape due to a multi-modal distribution arising on the simplex. The multi-modality depends upon the allocation of the probability left outside the simplex along the components. If only one side of the simplex has probability left outside (as in Figure 1(a)) then the resulting distribution will be unimodal (see Figure 1(b)).

An estimate of the probability left outside each side of the triangle (or the simplex) may be obtained through simulation in a straightforward manner. To accomplish this for our example, we generated data from the multivariate normal distribution with the parameters previously specified and applied the inverse of the α -transformation (that is, $g_0^\alpha(\mathbf{y})$) in Equation 13. For the points left outside of the simplex, we simply calculated how many are outside of each edge of the triangle and divided by the total number of the simulated data points.

If we partition the missed probability into three parts, where each part refers to one of the three components, then we obtain the values (0.008, 0.018, 0.124) for the case when $\Sigma = \Sigma_1$ (see Figures 1(a) and 1(b)). In this case, most of the probability is left outside the third component and the total probability left outside of the simplex is therefore 0.15. The total probability left outside of the simplex when $\Sigma = \Sigma_2$ (see Figures 1(c) and 1(d)) is 0.557 and the allocation to the three components is different than in the previous example, namely (0.141, 0.138, 0.278). Since, in this case, all of the estimates are relatively high, multi-modality appears in Figure 1(d).

3.3 Maximum likelihood estimation

The estimation of the parameters of the α -folded model on \mathbb{S}^2 (that is, $D = 3$) is not too complicated mainly because there are not too many parameters (2 for the mean vector, 3 for the covariance matrix and two more including the value of α and the probability p) involved in the maximization of the log-likelihood. We can use the ‘‘simplex’’ algorithm of Nelder & Mead (1965) to maximize the logarithm of Equation (16), available via the command *optim* in R (R Core Team, 2015). This algorithm is generally robust and is derivative free (Venables & Ripley, 2002). However when moving to higher dimensions (roughly $D = 5$ or larger), the maximization is not straightforward.

For this reason, we use the E-M algorithm (McLachlan & Krishnan, 2007) to maximize the log-likelihood corresponding to Equation 16. Let \mathbf{X} denote the set of n compositional vectors in \mathbb{S}^{D-1} . Following Jung, Foskey & Marron (2011) who applied the E-M algorithm in the context of a univariate folded normal we propose the algorithm below.

E-M algorithm

- Step 1. For a fixed value of α , apply the α -transformation without the Helmert sub-matrix multiplication (that is, Equation 6) to the compositional data \mathbf{X} to obtain the matrix \mathbf{W}_α .
- Step 2. Calculate w_i^* for each vector \mathbf{w}_i^α , for $i = 1, \dots, n$.
- Step 3. Left multiply each \mathbf{w}_i^α by the Helmert sub-matrix \mathbf{H} to obtain \mathbf{y}_{1i}^α . Then multiply each \mathbf{y}_{1i}^α by $\frac{1}{w_i^{*2}}$ to obtain \mathbf{y}_{2i}^α (see Equation 18). The \mathbf{y}_{1i}^α and \mathbf{y}_{2i}^α are the transformed compositional data onto \mathbb{A}_α^{D-1} and $\mathbb{R}^{D-1} \setminus \mathbb{A}_\alpha^{D-1}$, for $i = 1, \dots, n$ respectively.
- Step 4. Choose initial values for the estimates of the parameters, for example

$$\hat{\boldsymbol{\mu}}_\alpha^0 = \frac{\sum_{i=1}^n \mathbf{y}_{1i}^\alpha}{n} \quad \& \quad \hat{\Sigma}_\alpha^0 = \frac{1}{n} \sum_{i=1}^n (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^0) (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^0)^T$$

and

$$\hat{t}_i^0 = \frac{f_0(\mathbf{y}_{1i}^\alpha)}{f_0(\mathbf{y}_{1i}^\alpha) + f_1(\mathbf{y}_{2i}^\alpha)} \quad \& \quad \hat{p}^0 = \frac{\sum_{i=1}^n \hat{t}_i^0}{n},$$

where \hat{t}_i is the estimated conditional expectation of the indicator function that indicates whether the i -th observation belongs to f_0 or f_1 .

Step 5. Update all the parameters each time, for $k \geq 1$

$$\begin{aligned} \hat{\boldsymbol{\mu}}_\alpha^k &= \frac{\sum_{i=1}^n \hat{t}_i^{k-1} \mathbf{y}_{1i}^\alpha + \sum_{i=1}^n (1 - \hat{t}_i^{k-1}) \mathbf{y}_{2i}^\alpha}{n} \\ \hat{\boldsymbol{\Sigma}}_\alpha^k &= \frac{1}{n} \left[\sum_{i=1}^n \hat{t}_i^{k-1} (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k) (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k)^T \right. \\ &\quad \left. + \sum_{i=1}^n (1 - \hat{t}_i^{k-1}) (\mathbf{y}_{2i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k) (\mathbf{y}_{2i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k)^T \right] \\ \hat{t}_i^k &= \frac{f_0(\mathbf{y}_{1i}^\alpha)}{f_0(\mathbf{y}_{1i}^\alpha) + f_1(\mathbf{y}_{2i}^\alpha)} \\ \text{and } \hat{p}^k &= \frac{\sum_{i=1}^n \hat{t}_i^k}{n} \end{aligned}$$

Step 6. Repeat Step 5 until the change between two successive log-likelihood

$$\ell_\alpha = \sum_{i=1}^n \log [p f_0(\mathbf{y}_{1i}^\alpha) + (1 - p) f_1(\mathbf{y}_{2i}^\alpha)] \quad (20)$$

values is less than a tolerance value.

The above described procedure should be repeated for a grid of values of α , ranging from -1 to 1 , choosing the value of α which maximizes the log-likelihood. A more efficient search for the best alpha is via Brent's algorithm (Brent, 2013). When $\alpha = 0$, the MLE estimates of the transformed data are obtained directly; no E-M algorithm is necessary as all the probability is retained with the simplex. The fact that (9) tends to (4) as $\alpha \rightarrow 0$ ensures the continuity of the log-likelihood at $\alpha = 0$.

3.4 Generating data from the α -folded multivariate normal distribution

The algorithm below describes how to simulate a random vector from the α -folded model in Equation (16) when $\alpha \neq 0$. The case when $\alpha = 0$, is considered subsequently.

Step 1. Choose α , $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, where $\alpha \neq 0$.

Step 2. Generate a $D - 1$ by 1 vector \mathbf{Y} from a $N_{D-1}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

Step 3. As per Equation 13, determine whether $\mathbf{y} \in \mathbb{A}^{D-1}$. To do this, compute $\mathbf{w} = \mathbf{H}^T \mathbf{y}$. If $-\frac{1}{\alpha} \leq w_{i,\alpha} \leq \frac{D-1}{\alpha}$ for all components of \mathbf{w} and $\sum_{i=1}^D w_{i,\alpha} = 0$, then $\mathbf{y} \in \mathbb{A}^{D-1}$ and let $\mathbf{x} = \mathbf{z}_\alpha^{-1}(\mathbf{y})$. Otherwise, $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1}$ and let $\mathbf{x} = \mathbf{w}_\alpha^{-1} \left(\frac{\mathbf{H}^T \mathbf{y}}{q_\alpha^*(\mathbf{y})} \right)$ where $q^* = \alpha \min \{ \mathbf{H}^T \mathbf{y} \}$.

When $\alpha = 0$, a the following simplified algorithm is used:

Step 1. Choose α , $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

Step 2. Generate a $D - 1$ by 1 vector \mathbf{Y} from a $N_{D-1}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

Step 3. Compute $\mathbf{w} = \mathbf{H}^T \mathbf{y}$.

Step 4.

$$x_i = \frac{e^{w_i}}{\sum_{j=1}^D e^{w_j}}, \text{ for } i = 1, \dots, D. \quad (21)$$

Equation (21) is the inverse of the centered log-ratio transformation in Equation (2) and is the mechanism behind Aitchison's idea that generates compositional data.

3.5 Inference for α

In the previous two subsections, simplifications arose if $\alpha = 0$ and it may, consequently, be worthwhile to test whether the simpler multivariate normal in Equation (17) (corresponding to $\alpha = 0$) is appropriate for the data at hand.

Consider the hypothesis test: $H_0 : \alpha = 0$ versus $H_1 : \alpha \neq 0$. While one option is to use a log-likelihood ratio test, depending on the alternative hypothesis (a. $\alpha \neq 0$ & $p < 1$, b. $\alpha \neq 0$ & $p = 1$, c. $\alpha = 1$ & $p < 1$ or d. $\alpha = 1, p = 1$), the degrees of freedom will vary. We have not encountered case d. so far in our data analyses but, in this case, would recommend using a Dirichlet model. In fact, if we generate data from a Dirichlet distribution, case d. is expected to arise from the MLE of Equation (16). An alternative to the log likelihood ratio test is to use a parametric bootstrap such as the hypothesis testing procedure described below.

Step 1. For a given dataset, estimate the value of α obtained via the E-M algorithm. This is the observed test statistic denoted by α_{obs} .

Step 2. Apply the α -transformation in Equation (9) with $\alpha = \alpha_{obs}$ to the data. The data are now mapped onto set in Equation (10).

Step 3. Apply the inverse of the isometric transformation with $\alpha = 0$ in Equation (4) to the data in Step 2. to form a new sample of compositions acquired with $\alpha = 0$. That is, the data has been transformed under the null hypothesis.

Step 4. Re-sample B times from this new composition and each time estimate the value of α_b for $b = 1, \dots, B$.

The p -value is then given by (Davison & Hinkley, 1997)

$$p - \text{value} = \frac{\sum_{b=1}^B \mathbf{I}\{b : \alpha_b \geq \alpha_{obs}\} + 1}{B + 1}, \quad (22)$$

where \mathbf{I} is the indicator function.

One might argue that the value of α itself is not a pivotal statistic, in fact it is not even standardized, so a second bootstrap should be performed to obtain the standard error of the

estimate for each bootstrap sample. In order to avoid this extra computational burden, the parametric bootstrap hypothesis testing could alternatively be carried out using the log-likelihood ratio test statistic in Steps 1 and 4 above.

Inference for α could also be achieved via the construction of bootstrap confidence intervals. For this approach, simply re-sample the observations (compositional vectors) from the compositional dataset and find the value of α for which the log-likelihood derived from Equation (16) is maximized for each bootstrap sample. By repeating this procedure many times, we can empirically estimate the distribution of α , including the standard error of $\hat{\alpha}$. A variety of confidence intervals may be formed based on this distribution (see Davison & Hinkley, 1997 and R package *boot*). The percentile method, for example, simply uses the 2.5% lower and upper quantiles of the bootstrap distribution as confidence limits.

A less computationally intensive approach to obtain confidence intervals is based upon the second derivative of the profile log-likelihood of α , that is, the observed Fisher’s information measure. Assuming asymptotic normality of the estimator, the inverse of the observed information serves as an estimate of the standard error of the maximum likelihood estimator (Cox & Hinkley, 1979).

3.6 Principal component analysis

Principal component analysis for compositional data has been described by Aitchison (1983). The centred log-ratio transformation (2) is applied to the compositional data and standard eigen analysis is applied to the covariance matrix (which has at least one zero eigenvalue).

If $\alpha \neq 0$, we suggest an analogous approach in which the estimated covariance matrix is mapped to \mathbb{Q}_α^{D-1} space (Equation 7), using the Helmert sub-matrix

$$\hat{\Sigma}_\alpha^* = \mathbf{H}^T \hat{\Sigma}_\alpha \mathbf{H},$$

If $\alpha = 0$, the covariance is mapped to \mathbb{Q}_0^{D-1} (Equation (3)).

Step 3 of the algorithm presented in 3.4 is used to back-transform the principal components onto the simplex.

4 Simulation study and data analysis examples

4.1 Simulation study

Recovery of the mean, covariance and probability left outside the simplex

For this part of our simulation study, two values of α were chosen, namely -0.5 and 0.5 , and 4-dimensional data ($D = 5$) from a multivariate normal distribution with two different mean vectors and a variety of different covariance parameters were generated. Specifically, we used mean vectors

$$\boldsymbol{\mu}_{-0.5} = (1.715, 0.914, 0.115, 0.167) \quad \text{and} \quad \boldsymbol{\mu}_{0.5} = (-0.566, -0.979, -0.648, -0.651),$$

and covariance matrices

$$\boldsymbol{\Sigma} = \kappa \begin{pmatrix} 0.149 & -0.458 & 0.002 & -0.005 \\ -0.458 & 1.523 & 0.000 & 0.007 \\ 0.002 & 0.000 & 0.037 & -0.047 \\ -0.005 & 0.007 & -0.047 & 0.061 \end{pmatrix} \quad (23)$$

where $\kappa = 0.5, 1, 2, 3, 5, 7, 10$. Note that the value of k changes the probability that a point is left outside of the simplex. The “true” values of p for each value of k were computed through Monte-Carlo simulations and $1 - p$ (denoted as “Probability”) for each k value is presented in Table 1.

For each combination of α , $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, seven sample sizes, namely $n = (50, 100, 200, 300, 500, 750, 1000)$, were considered. Results are based on 1000 simulated data sets (for each n) and, for each simulated sample, estimates of p , $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ were calculated, and their distance from the true parameters were estimated. All computations took place in R 3.2.3 R (R Core Team, 2015) using a desktop computer with Intel Core i5 at 3.5 GHz processor and 32GB RAM.

For the probability left outside of the simplex (that is, $1-p$), the absolute difference between the estimated probability and the true probability is computed. For the mean vector, the Euclidean distance was calculated to measure the discrepancy between the estimated vector and true vector whereas for the covariance matrix, the following metric [15] was calculated

$$d(\hat{\boldsymbol{\Sigma}}, \boldsymbol{\Sigma}) = \sqrt{\sum_{i=1}^{D-1} \left[\log \Lambda_i(\hat{\boldsymbol{\Sigma}}\boldsymbol{\Sigma}^{-1}) \right]^2},$$

where $\Lambda_i(\mathbf{A})$ denotes the i -th eigenvalue of the matrix \mathbf{A} .


Note that while we could have used the Kullback-Leibler divergence of the fitted multivariate normal from the true multivariate normal to evaluate the overall performance of our estimation method, we would then have had no individual information regarding the accuracy of our procedure in terms of estimating the probability left outside the simplex, the mean vector and the covariance matrix. The results of the simulation studies are presented in Table 1 and Figure 2.

	κ	0.5	1	2	3	5	7	10
Estimated	$\alpha = -0.5$	0.0281	0.0925	0.1997	0.2800	0.3900	0.4630	0.5377
Probability	$\alpha = 0.5$	0.0223	0.1048	0.1849	0.3060	0.4217	0.4750	0.5356

Table 1: Probability left outside the simplex when $\alpha = -0.5$ and $\alpha = 0.5$, calculated via Monte Carlo with 50,000,000 iterations.

From Figure (2) we observe that when the probability left outside the simplex grows larger (k is larger), a larger sample size is required in order to get better estimates, for both the probability and the mean vector. The covariance matrix seems to be unaffected by the probability left outside the simplex.

Simulation study when $\alpha = -0.5$



Simulation study when $\alpha = 0.5$




Figure 2: All graphs contain the mean distance from each set of the parameters. The first column refers to the Euclidean distance of the estimated mean vector from the true mean vector. The second column refers to the mean distance between the estimated and the true covariance matrix. The third column refers to the mean absolute distance between the estimated probability and the true probability inside the simplex.

Estimation of the computational time required by the maximum likelihood estimation

Using only $\alpha = 0.5$, we generated data as before with increasing sample sizes and, for each sample size, recorded the time (in seconds) required to estimate the true value of α . As expected, the computational cost is mostly affected by p . For large sample sizes the computational burden is similar regardless of the probability of being outside of the simplex.

Recovery of the true value of α

In the previous simulations, recall that the value of α was fixed. In order to have a better picture of our estimation procedure, we also assessed the accuracy of the true value of α for large sample sizes for insight into the asymptotic case. The mean vector was set to

$$\boldsymbol{\mu} = (1.715, 0.914, 0.115, 0.167)$$

and the covariance matrices were the same as in 23.

	κ						
Sample size	0.5	1	2	3	5	7	10
50	0.052	0.058	0.065	0.075	0.095	0.124	0.148
100	0.073	0.079	0.092	0.106	0.135	0.174	0.212
200	0.090	0.089	0.096	0.105	0.130	0.162	0.195
300	0.095	0.095	0.096	0.108	0.128	0.160	0.194
500	0.084	0.083	0.091	0.100	0.123	0.156	0.188
750	0.078	0.074	0.081	0.089	0.109	0.135	0.163
1000	0.082	0.078	0.086	0.095	0.117	0.140	0.168
2000	0.267	0.390	0.357	0.352	0.339	0.276	0.302
5000	0.638	0.933	0.838	0.841	0.816	0.675	0.722
10000	1.435	2.113	1.905	1.915	1.858	1.544	1.723

Table 2: Computational times (in seconds) required to estimate the value of α , averaged over 1000 repetitions.

For a range of values of α ranging from 0 up to 1 with a step of 0.1 we estimated these values for the different values of k using 4 sample sizes ($n = 1000, 5000, 10000, 20000$). For each combination of α , k and n we used 1000 repetitions.

Figure 3 shows the average bias of the α estimates in boxplots for each sample size. Each box corresponds to a value of k and is the average bias aggregated for all values of α . For example, Figure 3(a) refers to a sample size equal to 1000 and the first box contains information about the average biases of the 11 values of α . The range in the variances increases with the value of k as expected, since higher values of k indicate higher probability of being left outside of the simplex. Table 3 presents $1 - p$ for many combinations of values of α and k calculated using Monte Carlo simulation with 20,000,000 repetitions.

	κ						
α	0.5	1	2	3	5	7	10
0.0	0.000	0.000	0.000	0.000	0.000	0.000	0.000
0.1	0.000	0.000	0.000	0.000	0.000	0.000	0.002
0.2	0.000	0.000	0.001	0.007	0.034	0.071	0.128
0.3	0.000	0.004	0.040	0.091	0.187	0.316	0.348
0.4	0.006	0.047	0.156	0.245	0.367	0.445	0.522
0.5	0.043	0.149	0.306	0.402	0.516	0.583	0.648
0.6	0.132	0.284	0.448	0.536	0.632	0.687	0.741
0.7	0.258	0.423	0.571	0.644	0.722	0.768	0.812
0.8	0.398	0.551	0.673	0.731	0.794	0.830	0.866
0.9	0.535	0.661	0.757	0.802	0.851	0.880	0.907
1.0	0.66	0.756	0.827	0.861	0.898	0.918	0.937

Table 3: Probability left outside the simplex for many combinations of α and l values.

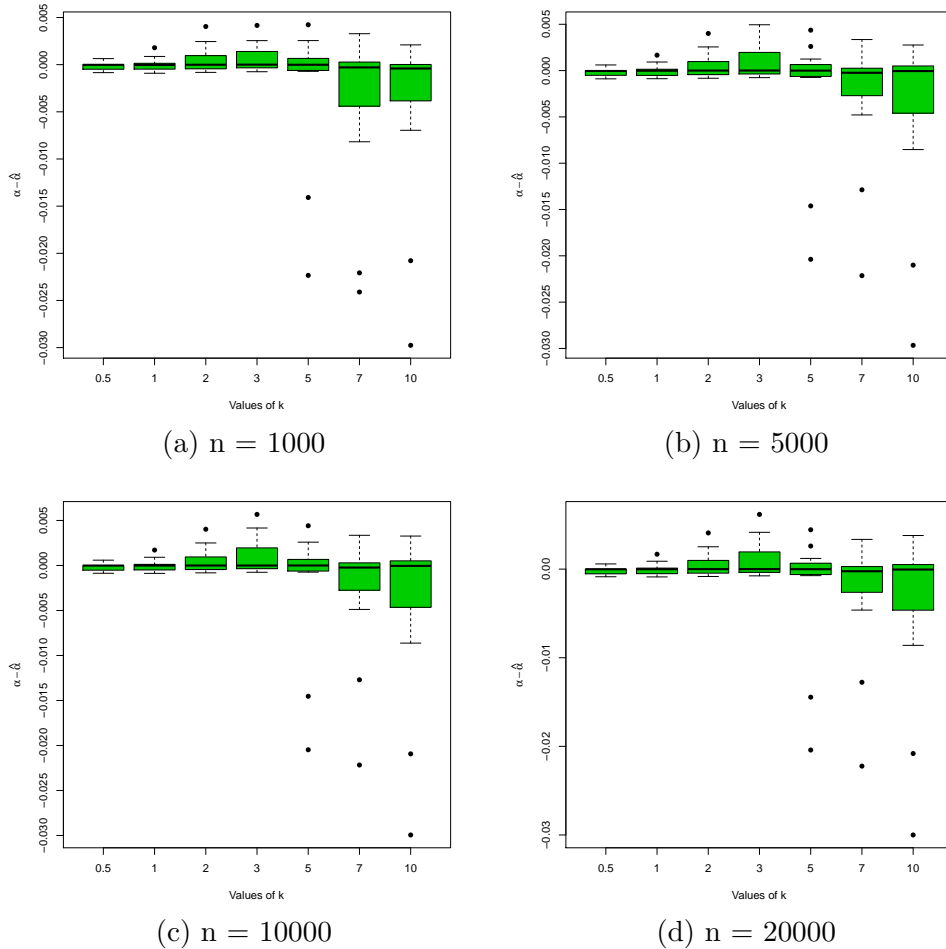


Figure 3: Box plots of the range of $\alpha - \hat{\alpha}$ as a function of k for 4 different sample sizes.

4.2 Example 1: Sharp's artificial data

We chose to analyze Sharp's artificial data (Sharp, 2006) because they are curved data and according to Aitchison (2003) the logistic normal (17) should produce a very good fit for curved data. Clearly a Dirichlet distribution would fail to capture the variability of such data and we would not expect a value of $\alpha = 1$ to do better. This dataset consists of 3 components and Sharp constructed them from Aitchison's Hongite data (Aitchison, 2003). [36] showed that the normalized geometric mean of the components (assuming a logistic normal distribution) fails to lie within the corpus of the data, whereas the spatial graph median does.

Figure 4(a) shows the profile log-likelihood of α and the maximum of the log-likelihood which occurs at $\alpha = 0.419$. The log-likelihood values at $\alpha = 0.419$ and $\alpha = 0$ are equal to 82.780 and 57.316 respectively. The log-likelihood ratio test of χ^2 at 2 degrees of freedom clearly rejects the logistic normal on the simplex (that $\alpha = 0$ is the optimal transformation) and this conclusion is in line with the confidence interval limits. Figure 4(b) shows the ternary diagram of the data and we can clearly see that both the simple and the normalized geometric mean fail to lie within the main bulk of the data. However, the Fréchet mean (or α -mean) (calculated at $\alpha = 0.419$)

achieves this goal. The three means in \mathbb{S}^2 are

$$\begin{aligned}\hat{\boldsymbol{\mu}}_0 &= (0.707, 0.241, 0.051) \quad (\text{Normalised geometric mean}) \\ \hat{\boldsymbol{\mu}}_1 &= (0.540, 0.275, 0.185) \quad (\text{Simple arithmetic mean}) \\ \hat{\boldsymbol{\mu}}_{0.419} &= (0.622, 0.272, 0.106) \quad (\alpha - \text{mean})\end{aligned}$$

and the estimated parameters assuming a multivariate normal in \mathbb{R}^2 with $\alpha = 0.419$ and $\alpha = 0$ are

$$\begin{aligned}\bar{\mathbf{y}}_{0.419} &= (0.679, 1.008) \quad \& \quad \hat{\boldsymbol{\Sigma}}_{0.419} = \begin{pmatrix} 0.256 & -0.684 \\ -0.684 & 1.932 \end{pmatrix} \quad \text{and} \\ \bar{\mathbf{y}}_0 &= (0.761, 1.701) \quad \& \quad \hat{\boldsymbol{\Sigma}}_0 = \begin{pmatrix} 0.588 & -1.629 \\ -1.629 & 5.986 \end{pmatrix}\end{aligned}$$

Note that in order to transform the α -mean (that is, $\bar{\mathbf{y}}_{0.419}$ in our example) inside the simplex, we apply Equation 13. The resulting compositional mean vector (α -mean) was termed Fréchet mean by (Tsagris, Preston & Wood, 2011).

In this example the probability left outside the simplex was calculated via Monte Carlo simulations (with 20,000,000 iterations) and was equal to 0.09301. The value of the mixing proportion in (16) was $\hat{p} = 0.95$ and $1 - \hat{p} = 0.05$. However, there were 2 out of 25 estimated values of the t_i 's less than 0.5, so without these (possible outliers) a proportion equal to 0.08 would have been obtained and this is closer to the Monte-Carlo estimated probability left outside the simplex.

The contour plots of the α -folded model with $\alpha = 0.419$ and $\alpha = 0$ appear in Figures 5(a) and 5(b) respectively. We generated 500 observations from each model and these are plotted in Figures 5(c) and 5(d). When $\alpha = 0.419$ the simulated data look more like the observed data, in contrast to the simulated data with $\alpha = 0$.

4.3 Example 2: Arctic lake data

Samples from an Arctic lake were collected and the composition of three minerals, sand silt and clay, was measured at different depths (Aitchison, 2003, pg 359). The estimated optimal value of α for this dataset was equal to 0.362 and Figure 6(b) shows the data along with the three α -means and the scores of the first principal component, evaluated at $\alpha = 0$ (geometric mean normalized to sum to 1) $\alpha = 0.362$ and $\alpha = 1$ (arithmetic mean).

A closer look at Figure 6(b) reveals that the α -mean evaluated at $\alpha = 0.362$ lies within the bulk of the data in contrast to the α -mean evaluated at $\alpha = 0$ as well as the arithmetic mean ($\alpha = 1$). The improved fit with $\alpha = 0.362$ (over $\alpha = 0$) can also be seen from the centers of the contour plots in Figures 7(a) and 7(b). There are two observations which seem to be highly influential.

We removed the two seemingly influential observations and re-implemented the analysis. The optimal value of α is now 0.443 and the contour plots of the α -folded model at $\alpha = 0.443$ and $\alpha = 0$ are presented in Figure 7(c) and 7(d). We again observe that a value of α other than 0 appears to provide a better fit to the data, leading to the same conclusion as before. That is, even though the data seem to be curved in the simplex, a log-ratio transformation is not the most suitable transformation.

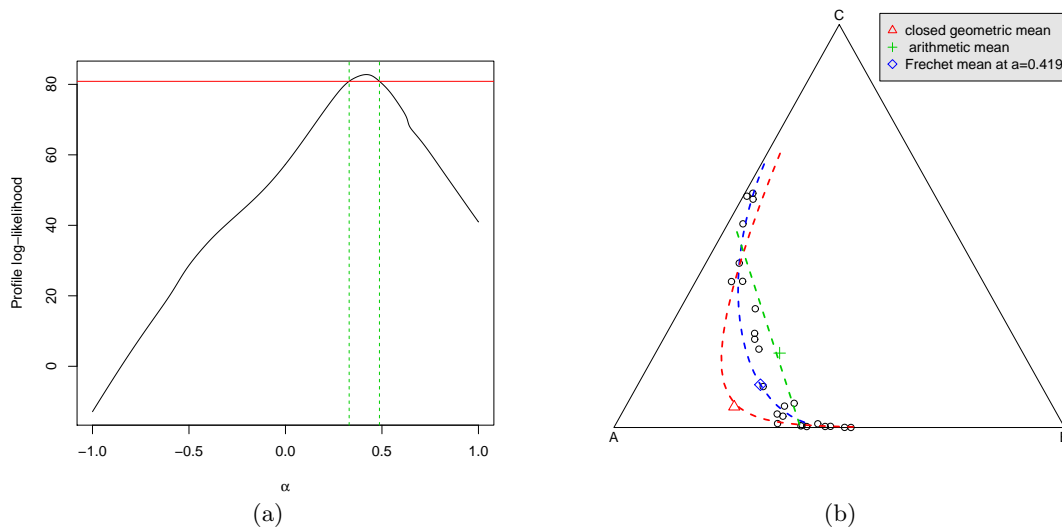


Figure 4: (a) Profile log-likelihood of α using Sharp’s data. The red and green lines indicate the 95% confidence interval of the true value of the parameter. (b) Ternary plot of the data along with three α -means evaluated at $\alpha = 0$ (geometric mean normalised to sum to 1), $\alpha = 0.419$ and $\alpha = 1$ (arithmetic mean). The lines correspond to the scores of the first principal component for each value of α .

5 Conclusions

In this paper we developed a novel parametric model, with nice properties, for compositional data analysis. Tsagris, Preston & Wood (2011) introduced the α -transformation and corresponding multivariate normal distribution. A drawback of that model is that it does not take into account the probability left outside the simplex. This is similar to the Box-Cox transformation, where the support of the transformed data is not the whole of \mathbb{R}^{D-1} as we have already mentioned. Another possible solution would be to use truncation [12]. However, in the examples presented the α -folded model appeared to fit the data adequately and, interestingly, an improved fit was obtained when $\alpha \neq 0$. For more examples where the log-ratio transformation fails to capture the variability of the data see Tsagris, Preston & Wood (2011), baxter (2006) and Sharp (2006).

The use of another multivariate model, such as the multivariate skew normal distribution (Azzalini & Valle, 1996) has been suggested. The difficulty, however, with this distribution is that more parameters need to be estimated, thus making the estimation procedure more difficult because the log-likelihood has many local maxima. This of course does not exclude the possibility of using this model. Another, perhaps simpler, alternative model is the multivariate t distribution. Bayesian analysis and regression modeling are two suggested research directions.

Similar to the Box-Cox transformation, the logarithm of the Jacobian of the α -transformation contains the term $(\alpha - 1) \sum_{j=1}^n \sum_{i=1}^D \log x_{ij}$. Thus, zero values are not allowed. Note that this issue also arises with the logistic normal. However, it is possible to generalize most of the analyzes suggested for the logistic normal distribution using our proposed folded model, including extensions that allow zeros.

As is standard practice in log-ratio transformation analysis, if one is willing to exclude from




Figure 5: Ternary plots of Sharp's data along with the contours of (a) the α -folded model with $\alpha = 0.419$ and (c) of the multivariate normal distribution applied to the α -transformed data with $\alpha = 0$. Plots (b) and (d) are 500 simulated data from the α -folded model with $\alpha = 0.419$ and from the multivariate normal distribution applied to the α -transformed data with $\alpha = 0$ respectively.

the sample space the boundary of the simplex, which corresponds to observations that have one or more components equal to zero, then the α -transformation (9) and its inverse are well defined for all $\alpha \in \mathbb{R}$, and the α -folded model provides a new approach for the analysis of compositional data with the potential to provide an improved fit over traditional models.

Proof of Lemma 3.1

Let us begin by deriving the Jacobian determinant of (5) at first. The map (5) is degenerate due to the constraints $\sum_{i=1}^D x_i = 1$ and $\sum_{i=1}^D u_i = 1$. In order to make (9) non-degenerate we consider the version of (5) as follows

$$u_a \{(x_i)\} = \frac{x_i^\alpha}{\sum_{j=1}^d x_j^\alpha + \left(1 - \sum_{j=1}^d x_j\right)^\alpha} \quad i = 1, \dots, d. \quad (\text{A.1})$$

The (A.1) is presented to highlight that in fact we have $d = D - 1$ and not D variables.

Let us start by proving the Jacobian of (5) or (A.1). We denote $S(\alpha) = \sum_{j=1}^D x_j^\alpha$, where $x_D = 1 - \sum_{j=1}^d x_j$. The diagonal and the non-diagonal elements of the Jacobian matrix are as

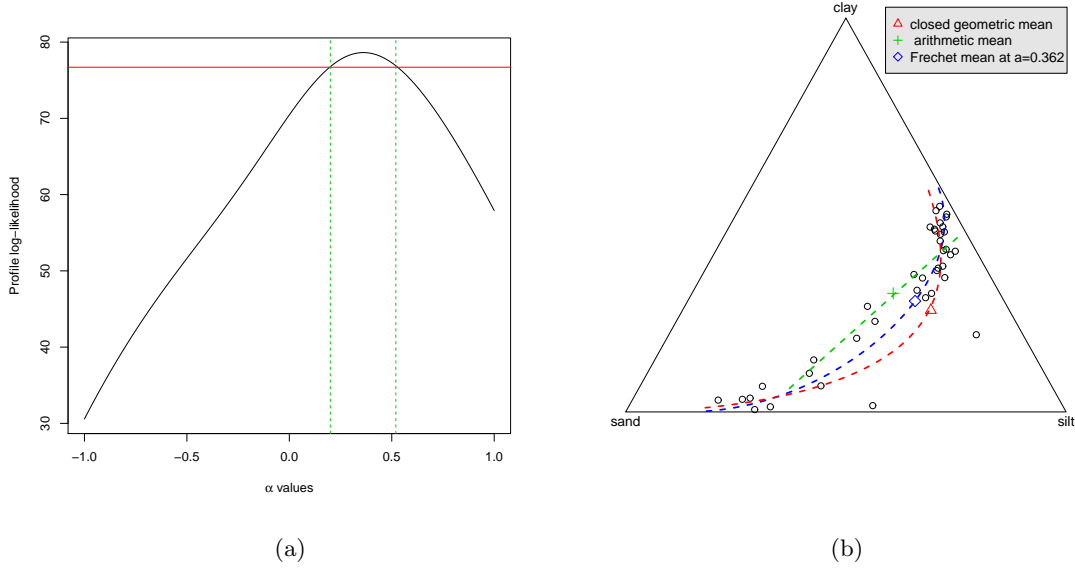


Figure 6: (a) Profile log-likelihood of α using the Arctic lake data. The red and green lines indicate the 95% confidence interval of the true value of the parameter. (b) Ternary plot of the data along with three α -means evaluated at $\alpha = 0$ (geometric mean normalised to sum to 1), $\alpha = 0.362$ and $\alpha = 1$ (arithmetic mean). The lines correspond to the scores of the first principal component for each value of α .

follows.

$$\frac{du_i}{dx_j} = \begin{cases} \frac{\alpha x_i^{\alpha-1} S(\alpha) - x_i^\alpha (\alpha x_i^{\alpha-1} - \alpha x_D^{\alpha-1})}{S^2(\alpha)} & i = j \\ \frac{x_i^\alpha (\alpha x_j^{\alpha-1} - \alpha x_D^{\alpha-1})}{S^2(\alpha)} & (i \neq j) \quad i \neq j \end{cases}$$

The Jacobian takes the following form [25]:

$$|\mathbf{J}| = |\mathbf{A} - \mathbf{B}\mathbf{C}^T| S^{-2d}(\alpha) = |\mathbf{A}| (1 - \mathbf{C}^T \mathbf{A}^{-1} \mathbf{B}) S^{-2d}(\alpha),$$

where \mathbf{A} is a diagonal $d \times d$ matrix with elements $\alpha x_i^{\alpha-1} S(\alpha)$ and \mathbf{B} and \mathbf{C} are defined as

$$\mathbf{B} = (x_1^\alpha, \dots, x_d^\alpha)^T \quad \text{and} \quad \mathbf{C} = \alpha (x_1^{\alpha-1} - x_D^{\alpha-1}, \dots, x_d^{\alpha-1} - x_D^{\alpha-1})^T.$$

Then

$$\mathbf{A}^{-1} \mathbf{B} = \begin{pmatrix} \frac{x_1^{1-\alpha}}{\alpha S(\alpha)} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \frac{x_d^{1-\alpha}}{\alpha S(\alpha)} \end{pmatrix} \begin{pmatrix} x_1^\alpha \\ \vdots \\ x_d^\alpha \end{pmatrix} = \begin{pmatrix} \frac{x_1}{\alpha S(\alpha)} \\ \vdots \\ \frac{x_d}{\alpha S(\alpha)} \end{pmatrix}.$$

Then the multiplication $\mathbf{C}^T \mathbf{A}^{-1} \mathbf{B}$ is

$$\begin{aligned} \mathbf{C}^T \mathbf{A}^{-1} \mathbf{B} &= \left(\alpha x_1^{\alpha-1} - \alpha x_D^{\alpha-1}, \dots, \alpha x_d^{\alpha-1} - \alpha x_D^{\alpha-1} \right) \begin{pmatrix} \frac{x_1}{\alpha S(\alpha)} \\ \vdots \\ \frac{x_d}{\alpha S(\alpha)} \end{pmatrix} \\ &= \frac{\sum_{i=1}^{D-1} x_i^\alpha}{S(\alpha)} - \frac{\alpha x_D^{\alpha-1}}{S(\alpha)} \left(\frac{x_1}{\alpha} + \dots + \frac{x_d}{\alpha} \right). \end{aligned}$$

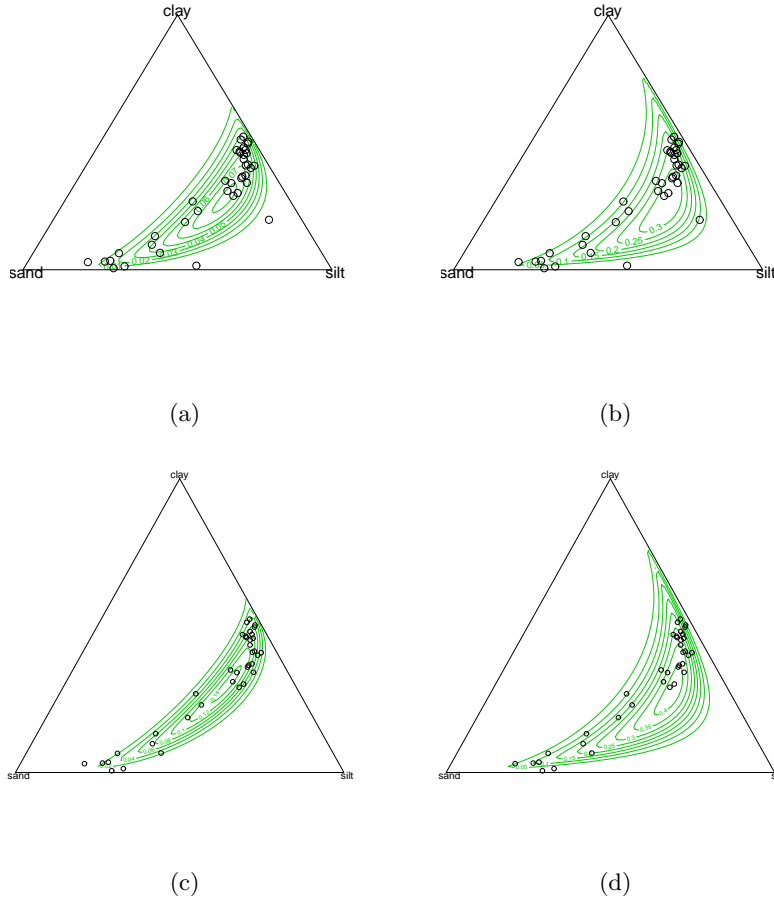


Figure 7: The contour plots refer to (a) the α -folded model with $\alpha = 0.362$ and (b) to the multivariate normal distribution applied to the α -transformed data with $\alpha = 0$. The second row presents the contour plots with two observations removed. Plots (c) and (d) correspond to $\alpha = 0.443$ and $\alpha = 0$ respectively.

So we end up with

$$\begin{aligned}
 1 - \mathbf{C}^T \mathbf{A}^{-1} \mathbf{B} &= \frac{S(\alpha) - (S(\alpha) - x_D^\alpha)}{S(\alpha)} + \frac{\alpha x_D^{\alpha-1}}{S(\alpha)} \sum_{i=1}^d \frac{x_i}{\alpha} = \frac{x_D^\alpha + \alpha x_D^{\alpha-1} \sum_{i=1}^d \frac{x_i}{\alpha}}{S(\alpha)} \\
 &= \frac{x_D^{\alpha-1} \left(x_D + \alpha \sum_{i=1}^d \frac{x_i}{\alpha} \right)}{S(\alpha)} = \frac{x_D^{\alpha-1}}{S(\alpha)}.
 \end{aligned}$$

Finally the Jacobian of (5) takes the following form

$$\begin{aligned}
 |\mathbf{J}| &= S^d(\alpha) \frac{\prod_{i=1}^d \alpha x_i^{\alpha-1} x_D^{\alpha-1}}{S^d(\alpha)^{-2d} S(\alpha)} = S^{-d-1}(\alpha) x_D^{\alpha-1} \prod_{i=1}^d \alpha x_i^{\alpha-1} \\
 &= \alpha^d \prod_{i=1}^d \frac{x_i^{\alpha-1}}{\sum_{j=1}^D x_j^\alpha}.
 \end{aligned}$$

The Jacobian of the α -transformation (9) without the left multiplication by the Helmert

sub-matrix \mathbf{H} is simply the Jacobian of (5) multiplied by $\frac{D^d}{\alpha^d}$

$$|\mathbf{J}| = D^{D-1} \prod_{i=1}^D \frac{x_i^{\alpha-1}}{\sum_{j=1}^D x_j^\alpha}$$

The multiplication by the Helmert sub-matrix adds an extra term to the Jacobian, which is \sqrt{D} and hence the Jacobian becomes.

$$|\mathbf{J}| = D^{D-1+1/2} \prod_{i=1}^D \frac{x_i^{\alpha-1}}{\sum_{j=1}^D x_j^\alpha}$$

Proof of Lemma 3.2

The proof presented below is about the case of $\alpha = 1$ for convenience purposes. The way to map a point \mathbf{x} from inside the simplex to a point \mathbf{y} outside the simplex, is given in Equation 12 and the component wise transformation can be expressed as

$$y_i = Z_i \left(\frac{1}{w^*} \right)^2 w_i + (1 - Z_i) \left(\frac{1}{w_D} \right)^2 w_i, \quad (\text{A.2})$$

where \mathbf{w} is defines in (6) and since $\alpha = 1$, we have exclude the superscript α .

$$w^* = |\min \{w_1, \dots, w_{D-1}\}| \quad \text{and} \quad Z_i = \left\{ \begin{array}{ll} 1 & \text{if } w^* \neq w_D \\ 0 & \text{if } w^* = w_D \end{array} \right\}, \quad \text{for } i = 1, \dots, D-1.$$

There are two cases to consider when calculating the Jacobian determinant of the transformation.

1. The first case is when $Z_i = 1$ and the transformation is

$$y_i = \left(\frac{1}{w^*} \right)^2 w_i.$$

There are two sub-cases to be specified.

- (a) $w^* = w_i$ where the derivatives are given by

$$\frac{\partial y_i}{\partial w_j} = \left\{ \begin{array}{ll} -\frac{1}{w_i^2} & i = j \\ 0 & i \neq j. \end{array} \right\}.$$

- (b) $w^* \neq w_i$ where the derivatives are given by

$$\frac{\partial y_i}{\partial w_i} = \left(\frac{1}{w^*} \right)^2 \quad \text{and} \quad \frac{\partial y_i}{\partial w_j} = \left\{ \begin{array}{ll} \frac{-2}{(w^*)^3} w_i & \text{if } w^* = w_j \\ 0 & \text{if } w^* \neq w_j \end{array} \right\}$$

The Jacobian matrix is

$$\begin{bmatrix} \frac{\partial y_1}{\partial w_1} & \cdots & \frac{\partial y_1}{\partial w_i} & \cdots & \frac{\partial y_1}{\partial w_{D-1}} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ \frac{\partial y_i}{\partial w_1} & \cdots & \frac{\partial y_i}{\partial w_i} & \cdots & \frac{\partial y_i}{\partial w_{D-1}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial y_{D-1}}{\partial w_1} & \cdots & \frac{\partial y_{D-1}}{\partial w_i} & \cdots & \frac{\partial y_{D-1}}{\partial w_{D-1}} \end{bmatrix} = \begin{bmatrix} \frac{1}{(w^*)^2} & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & 0 & \cdots & \vdots & \vdots \\ \vdots & \ddots & \frac{1}{(w^*)^2} & \frac{-2}{(w^*)^3} w_i & \vdots & \vdots \\ \vdots & \vdots & 0 & \ddots & 0 & \vdots \\ \vdots & 0 & \cdots & \ddots & \ddots & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & \frac{1}{(w^*)^2} \end{bmatrix}$$

and hence, the determinant is equal to

$$|J| = \left(\frac{1}{w^*} \right)^{2(D-1)}.$$

2. The second case is when $Z_i = 0$ and the transformation is

$$y_i = \frac{1}{w_D^2} w_i.$$

The derivatives are now given by

$$\frac{\partial y_i}{\partial w_j} = \begin{cases} \frac{1}{w_D^2} + \frac{2}{w_D^3} w_i & i = j \\ \frac{2}{w_D^3} w_i & i \neq j \end{cases}$$

Note that the sign for the derivative with respect to w_D is positive because

$$y_i = \frac{1}{w_D^2} w_i = \frac{w_i}{\left(-\sum_{j=1}^d w_j\right)^2} = \frac{w_i}{\left(\sum_{j=1}^d w_j\right)^2}, \text{ thus}$$

$$\frac{\partial y_i}{\partial w_D} = -2 \frac{w_i}{\left(\sum_{j=1}^d w_j\right)^3} = -2 \frac{w_i}{(-w_D)^3} = 2 \frac{w_i}{w_D^3}$$

The Jacobian matrix in this case can be written as

$$\begin{bmatrix} \frac{\partial y_1}{\partial w_1} & \cdots & \frac{\partial y_1}{\partial w_i} & \cdots & \frac{\partial y_1}{\partial w_{D-1}} \\ \vdots & \ddots & \vdots & & \vdots \\ \frac{\partial y_i}{\partial w_1} & \cdots & \frac{\partial y_i}{\partial w_i} & \cdots & \frac{\partial y_i}{\partial w_{D-1}} \\ \vdots & & \vdots & \ddots & \vdots \\ \frac{\partial y_{D-1}}{\partial w_1} & \cdots & \frac{\partial y_{D-1}}{\partial w_i} & \cdots & \frac{\partial y_{D-1}}{\partial w_{D-1}} \end{bmatrix} = \begin{bmatrix} \frac{1}{w_D^2} + \frac{2w_1}{w_D^3} & +\frac{2w_1}{w_D^3} & \cdots & \cdots & +\frac{2w_1}{w_D^3} \\ \vdots & \ddots & \cdots & \vdots & \vdots \\ \frac{2w_i}{w_D^3} & \vdots & \frac{1}{w_D^2} + \frac{2w_i}{w_D^3} & \vdots & +\frac{2w_i}{w_D^3} \\ \vdots & \vdots & \cdots & \ddots & \vdots \\ \frac{2w_d}{w_D^3} & \cdots & \cdots & \frac{2w_d}{w_D^3} & \frac{1}{w_D^2} + \frac{2w_{D-1}}{w_D^3} \end{bmatrix}$$

The determinant of such matrices is given by (author?) [25]

$$|J| = |\mathbf{A} + \mathbf{B}\mathbf{C}^T| = |\mathbf{A}| (1 + \mathbf{C}^T \mathbf{A}^{-1} \mathbf{B}) \quad (\text{A.3})$$

where

$$\mathbf{A} = \text{diag} \left(\frac{1}{w_D^2}, \dots, \frac{1}{w_D^2} \right), \quad \mathbf{B} = 2(w_1, \dots, w_{D-1})^T \quad \text{and} \quad \mathbf{C} = (w_D^{-3}, \dots, w_D^{-3})^T.$$

$$\begin{aligned} |J| &= \left(\frac{1}{w_D^2} \right)^d \left\{ 1 + 2(w_D^{-3}, \dots, w_D^{-3}) \begin{bmatrix} w_D^2 & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & w_D^2 \end{bmatrix} \begin{pmatrix} w_1 \\ \vdots \\ w_d \end{pmatrix} \right\} \\ &= \left(\frac{1}{w_D^2} \right)^d \left[1 + 2(w_D^{-1}, \dots, w_D^{-1}) \begin{pmatrix} w_1 \\ \vdots \\ w_d \end{pmatrix} \right] \\ &= \left(\frac{1}{w_D^2} \right)^{D-1} \left[1 + 2 \frac{\sum_{j=1}^d w_j}{w_D} \right] = \left(\frac{1}{w_D^2} \right)^{D-1} \left(1 + 2 \frac{\sum_{j=1}^d w_j}{-\sum_{j=1}^d w_j} \right). \end{aligned}$$

Finally, (A.3) becomes

$$|J| = \left| \left(\frac{1}{w_D^2} \right)^d (1 - 2) \right| = \left(\frac{1}{w_D} \right)^{2(D-1)}.$$

References

- [1] Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society. Series B*, 44, 139–177.
- [2] Aitchison, J. (1983). Principal component analysis of compositional data. *Biometrika*, 70, 57–65.
- [3] Aitchison, J. (2003). *The statistical analysis of compositional data*. New Jersey: Reprinted by The Blackburn Press.
- [4] Azzalini, A. & Valle, A. (1996). The multivariate skew-normal distribution. *Biometrika*, 83, 715–726.
- [5] Baxter, M., Beardah, C., Cool, H., & Jackson, C. (2005). Compositional data analysis of some alkaline glasses. *Mathematical geology*, 37, 183–196.
- [6] Baxter, M. & Freestone, I. (2006). Log-ratio compositional data analysis in archaeometry. *Archaeometry*, 48, 511–531.
- [7] Brent, R. P. (2013). *Algorithms for minimization without derivatives*. Courier Corporation.
- [8] Butler, A. & Glasbey, C. (2008). A latent gaussian model for compositional data with zeros. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 57, 505–520.
- [9] Chen, E. Z. & Li, H. (2016). A two-part mixed-effects model for analyzing longitudinal microbiome compositional data. *Bioinformatics*, 32, 2611–2617.
- [10] Cox, D. & Hinkley, D. (1979). *Theoretical statistics*. London: Chapman & Hall/CRC.
- [11] Davison, A. & Hinkley, D. (1997). *Bootstrap methods and their application*. Cambridge: Cambridge University Press.
- [12] Dobigeon, N. & Tournet, J.-Y. (2007). Truncated multivariate gaussian distribution on a simplex. Technical report 2007a, University of Toulouse.
- [13] Dryden, I. & Mardia, K. (1998). *Statistical Shape Analysis*. John Wiley & Sons.
- [14] Egozcue, J., Pawlowsky-Glahn, V., Mateu-Figueras, G., & Barceló-Vidal, C. (2003). Isometric logratio transformations for compositional data analysis. *Mathematical Geology*, 35, 279–300.
- [15] Förstner, W. & Moonen, B. (2003). A metric for covariance matrices. In *Geodesy-The Challenge of the 3rd Millennium*, 299–309. Springer.
- [16] Fry, J., Fry, T., & McLaren, K. (2000). Compositional data analysis and zeros in micro data. *Applied Economics*, 32, 953–959.
- [17] Ghosh, D. & Chakrabarti, R. (2009). Joint variable selection and classification with immunohistochemical data. *Biomarker insights*, 4, 103–110

- [18] Heath, T. L. (1921). *A history of Greek mathematics*. Clarendon.
- [19] Johnson, N. (1962). The folded normal distribution: Accuracy of estimation by maximum likelihood. *Technometrics*, 4, 249–256.
- [20] Jung, S., Foskey, M., & Marron, J. S. (2011). Principal arc analysis on direct product manifolds. *The Annals of Applied Statistics*, 5, 578–603.
- [21] Katz, J. & King, G. (1999). A statistical model for multiparty electoral data. *American Political Science Review*, 93, 15–32.
- [22] Lancaster, H. (1965). The Helmert matrices. *American Mathematical Monthly*, 72, 4–12.
- [23] Le, H. & Small, C. (1999). Multidimensional scaling of simplex shapes. *Pattern Recognition*, 32, 1601–1613.
- [24] Leone, F., Nelson, L., & Nottingham, R. (1961). The folded normal distribution. *Technometrics*, 3, 543–550.
- [25] Mardia, K., Kent, J., & Bibby, J. (1979). *Multivariate Analysis*. London: Academic Press.
- [26] McLachlan, G. & Krishnan, T. (2007). *The EM algorithm & extensions*. John Wiley & Sons.
- [27] Nelder, J. & Mead, R. (1965). A simplex algorithm for function minimization. *Computer Journal*, 7, 308–313.
- [28] Neocleous, T., Aitken, C., & Zadora, G. (2011). Transformations for compositional data with zeros with an application to forensic evidence evaluation. *Chemometrics & Intelligent Laboratory Systems*, 109, 77–85.
- [29] Oeppen, J. (2008). Coherent forecasting of multiple-decrement life tables: a test using japanese cause of death data. In *Proceedings of the 3rd Compositional Data Analysis Workshop, Girona, Spain*.
- [30] Otero, N., Tolosana-Delgado, R., Soler, A., Pawlowsky-Glahn, V., & Canals, A. (2005). Relative vs. absolute statistical analysis of compositions: A comparative study of surface waters of a mediterranean river. *Water research*, 39, 1404–1414.
- [31] Prados, F., Boada, I., Prats-Galino, A., Martín-Fernández, J., Feixas, M., Blasco, G., Puig, J., & Pedraza, S. (2010). Analysis of new diffusion tensor imaging anisotropy measures in the three-phase plot. *Journal of Magnetic Resonance Imaging*, 31, 1435–1444.
- [32] R Core Team (2015). *R: A Language & Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- [33] Scealy, J. & Welsh, A. (2011). Regression for compositional data by using distributions defined on the hypersphere. *Journal of the Royal Statistical Society. Series B*, 73, 351–375.
- [34] Scealy, J. & Welsh, A. (2014). Fitting Kent models to compositional data with small concentration. *Statistics & Computing*, 24, 165–179.

- [35] Schnute, J. & Haigh, R. (2007). Compositional analysis of catch curve data, with an application to sebastes maliger. *ICES Journal of Marine Science*, 64, 218–233.
- [36] Sharp, W. (2006). The graph median—a stable alternative measure of central tendency for compositional data sets. *Mathematical geology*, 38, 221–229.
- [37] Shi, P., Zhang, A., Li, H., et al. (2016). Regression analysis for microbiome compositional data. *The Annals of Applied Statistics*, 10, 1019–1040.
- [38] Stewart, C. & Field, C. (2011). Managing the essential zeros in quantitative fatty acid signature analysis. *Journal of Agricultural, Biological, & Environmental Statistics*, 16, 45–69.
- [39] Tsagris M. (2013). Novel methods for the statistical analysis of compositional data. *Unpublished PhD Thesis, University of Nottingham*.
- [40] Tsagris, M. (2015). Regression analysis with compositional data containing zero values. *Chilean Journal of Statistics*, 6, 47–57.
- [41] Tsagris, M., Preston, S., & Wood, A. (2011). A data-based power transformation for compositional data. In *Proceedings of the 4th Compositional Data Analysis Workshop, Girona, Spain*.
- [42] Tsagris, M., Preston, S., & Wood, A. T. (2016). Improved classification for compositional data using the α -transformation. *Journal of Classification*, 33, 243–261.
- [43] Tsagris, M. & Stewart C. (2018). A Dirichlet regression model for compositional data with zeros *Lobachevskii Journal of Mathematics*, To appear.
- [44] Venables, W. & Ripley, B. (2002). *Modern applied statistics with S*. Springer Verlag.
- [45] Xia, F., Chen, J., Fung, W. K., & Li, H. (2013). A logistic normal multinomial regression model for microbiome compositional data analysis. *Biometrics*, 69, 1053–1063.