

A folded model for compositional data analysis

Michail Tsagris^{1*} and Connie Stewart²

Department of Economics, University of Crete and Department of Mathematics & Statistics, University of New Brunswick

Summary

A folded type model is developed for analyzing compositional data. The proposed model involves an extension of the α -transformation for compositional data and 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, and can be advantageous over a similar model without folding.

Key words: α -transformation, EM-algorithm, compositional data, folding transformation

1. Introduction

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

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

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 (Aitchison 2003) 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 different sea depths were analysed in Aitchison (2003). Schnute & Haigh (2007) analysed 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

*Author to whom correspondence should be addressed.

¹ Department of Economics, University of Crete, Rethymnon, Greece

² Department of Mathematics and Statistics, University of New Brunswick, Saint John, New Brunswick, Canada

Email: mtsagris@uoc.gr

anthropogenic and geological pollution sources of rivers in Spain. Stewart & Field (2011) modeled compositional diet estimates of fish predators obtained through quantitative fatty acid signature analysis. In another biological setting, Ghosh & Chakrabarti (2009) were interested in the classification of immunohistochemical data. Other application areas of compositional data analysis include archaeometry (Baxter et al. 2005), where the composition of ancient glasses is of interest, and economics (Fry, Fry & McLaren 2000), where the focus is on the percentage of the household expenditure allocated to different commodities. Compositional data are also met in political science (Katz & King 1999) for modeling electoral data and in forensic science where different types of forensic glasses are compared and classified according to their compositions (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 et al. 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 recognised that this distribution is not, statistically, rich and flexible enough to capture many types of variabilities (especially curvature) of compositional data. For this reason, a variety of transformations have been proposed that map the data outside 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) and examined in detail by Egozcue et al. (2003). The α -transformation, a Box-Cox type transformation, was suggested by Tsagris, Preston & Wood (2011) and includes the isometric transformation as a special case. The α -transformation has been successfully applied in regression analysis (Tsagris 2015) and classification settings (Tsagris 2014; Tsagris, Preston & Wood 2016).

Regardless of the transformation chosen, the usual approach for modeling compositional data is to assume that the transformed data are 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. (In the univariate case, for example, imagine fitting a normal, and not a truncated normal, to a limited space.) 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 support. The model we propose here works in a similar fashion; 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 a variety of multivariate distributions 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, an algorithm for generating data from the proposed folded distribution and inference for the transformation parameter α . Three data sets are analysed in Section 4 that illustrate both the advantages and limitations of the proposed model. Finally, simulation studies are provided in Section 5 to assess the estimation accuracy of the parameters and computational burden of the algorithm, and concluding remarks may be found in Section 6.

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 Aitchison (1983) as

$$\mathbf{w}_0(\mathbf{x}) = \left(\log \left(\frac{x_1}{\prod_{j=1}^D x_j^{1/D}} \right), \dots, \log \left(\frac{x_D}{\prod_{j=1}^D x_j^{1/D}} \right) \right). \quad (2)$$

The sample space of (2) is the set

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

which is a subset of \mathbb{R}^{D-1} . Note that the zero sum constraint in Equation (3) 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 $\mathbf{z}_0(\mathbf{x})$ is a $D - 1$ dimensional vector and \mathbf{H} is the $(D - 1) \times D$ Helmert (Lancaster 1965) sub-matrix. (This is the Helmert matrix after deletion of the first row. This sub-matrix

is a standard orthogonal matrix in shape analysis used to overcome singularity problems. Its structure and components are specified in the Appendix. For further information, see Dryden & Mardia (1998); Le & Small (1999).) The sample space of Equation (4) is \mathbb{R}^{D-1} because left multiplication by the Helmert sub-matrix maps the centred log-ratio transformed data onto \mathbb{R}^{D-1} , thus, in effect, removing the zero sum constraint.

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

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

denote the power transformation for compositional data as defined by Aitchison (2003), where α can take any value. (When zero values exist in the data, α can take only strictly positive values.) In a manner analogous to Equations (2-4), first define

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

The sample space of Equation (6) is

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

Note that the inverse of Equation (6) is

$$\mathbf{x} = \mathbf{w}_\alpha^{-1}(\mathbf{m}) = \left(\frac{(1 + \alpha m_1)^{1/\alpha}}{\sum_{j=1}^D (1 + \alpha m_j)^{1/\alpha}}, \dots, \frac{(1 + \alpha m_D)^{1/\alpha}}{\sum_{j=1}^D (1 + \alpha m_j)^{1/\alpha}} \right), \quad (8)$$

for $\mathbf{m} \in \mathbb{Q}_\alpha^{D-1}$. As $\alpha \rightarrow 0$ Equation (6) converges to Equation (2) and Equation (8) becomes

$$\mathbf{x} = \mathbf{w}_0^{-1}(\mathbf{m}) = \left(\frac{e^{m_1}}{\sum_{j=1}^D e^{m_j}}, \dots, \frac{e^{m_D}}{\sum_{j=1}^D e^{m_j}} \right). \quad (9)$$

Finally, the α -transformation is defined as

$$\mathbf{z}_\alpha(\mathbf{x}) = \mathbf{H}\mathbf{w}_\alpha(\mathbf{x}). \quad (10)$$

The transformation in Equation (10) 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 (10) is

$$\mathbb{A}_\alpha^{D-1} = \left\{ \mathbf{H}\mathbf{w}_\alpha(\mathbf{x}) \mid -\frac{1}{\alpha} \leq w_{i,\alpha} \leq \frac{D-1}{\alpha}, \sum_{i=1}^D w_{i,\alpha} = 0 \right\}. \quad (11)$$

For $\mathbf{y} = \mathbf{z}_\alpha(\mathbf{x})$, the inverse transformation from \mathbb{A}_α^{D-1} to \mathbb{S}^{D-1} is $\mathbf{z}_\alpha^{-1}(\mathbf{y}) = \mathbf{w}_\alpha^{-1}(\mathbf{H}^\top \mathbf{y})$ 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 (MLE). 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 Scaely & Welsh (2011b), who then developed the folded Kent distribution Scaely & Welsh (2014) and we propose a similar solution here.

2.2. The α -folding transformation

Let $\mathbf{y} = \mathbf{z}_\alpha(\mathbf{x})$ in Equation (10) for $\mathbf{x} \in \mathbb{S}^{D-1}$ and some value of α . The inverse transformation $\mathbf{z}_\alpha^{-1}(\mathbf{y})$ provides a transformation from $\mathbf{y} \in \mathbb{A}_\alpha^{D-1}$ to \mathbb{S}^{D-1} . Now suppose we have a point $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}_\alpha^{D-1}$ (that is, outside of \mathbb{A}_α^{D-1}) and we want to map it inside the simplex \mathbb{S}^{D-1} . The following transformation maps \mathbf{y} to \mathbb{S}^{D-1} (see Appendix for a proof)

$$\mathbf{x} = \mathbf{w}_\alpha^{-1} \left(\frac{\mathbf{H}^\top \mathbf{y}}{q_\alpha^{*2}(\mathbf{y})} \right), \quad (12)$$

where $\mathbf{w}_\alpha^{-1}(\cdot)$ is defined in Equation (8), $q_\alpha^*(\mathbf{y}) = \alpha \min \{\mathbf{H}^\top \mathbf{y}\}$ and $\min \{\mathbf{H}^\top \mathbf{y}\}$ refers to the minimum component value of the vector $\mathbf{H}^\top \mathbf{y}$. Note that for $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}_\alpha^{D-1}$ the inverse of Equation (12) is the $D - 1$ dimensional vector

$$\mathbf{y} = \frac{1}{w_\alpha^{*2}(\mathbf{x})} \mathbf{H}\mathbf{w}_\alpha(\mathbf{x}) = \frac{1}{w_\alpha^{*2}(\mathbf{x})} \mathbf{z}_\alpha(\mathbf{x}), \quad (13)$$

where $w_\alpha^*(\mathbf{x}) = \alpha \min \{\mathbf{w}_\alpha(\mathbf{x})\}$ and $\min \{\mathbf{w}_\alpha(\mathbf{x})\}$ is the minimum component of the vector $\mathbf{w}_\alpha(\mathbf{x})$ defined in Equation (6). The square term in the exponent of the denominator in Equation (12) was chosen for convenience purposes. The square could alternatively be

substituted by any other power, say β . However, this generalisation would make the model more complex and difficult to work with as there would be an additional parameter to estimate. In the Bayesian stance this would be easier to solve, yet computationally less efficient. In summary, we propose the following folding transformation from $\mathbf{y} \in \mathbb{R}^{D-1}$ to \mathbb{S}^{D-1}

$$\mathbf{x} = \begin{cases} g_0^\alpha(\mathbf{y}) & \text{if } \mathbf{y} \in \mathbb{A}_\alpha^{D-1} \\ g_1^\alpha(\mathbf{y}) & \text{if } \mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}_\alpha^{D-1}, \end{cases} \quad (14)$$

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

3. The α -folded multivariate normal distribution

3.1. Definition

The derivation of the α -folded multivariate normal distribution is carried out by first assuming that \mathbf{y} is multivariate normally distributed with parameters $\boldsymbol{\mu}_\alpha$ and $\boldsymbol{\Sigma}_\alpha$, and subsequently considering the distribution of \mathbf{x} , where \mathbf{x} is defined in Equation (14). The basis of this approach was used in Leone, Nelson & Nottingham (1961) to derive the univariate folded normal distribution.

If $\mathbf{y} \sim N_{D-1}(\boldsymbol{\mu}_\alpha, \boldsymbol{\Sigma}_\alpha)$, then the density of \mathbf{y} is given by

$$f_{\mathbf{y}}(\mathbf{y}) = \frac{1}{|2\pi\boldsymbol{\Sigma}_\alpha|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{y} - \boldsymbol{\mu}_\alpha)^\top \boldsymbol{\Sigma}_\alpha^{-1}(\mathbf{y} - \boldsymbol{\mu}_\alpha)\right], \quad (15)$$

and we can apply standard statistical techniques for finding the distribution of a transformation of a random vector to derive the distribution of \mathbf{x} . Recognise that in our case we require the distribution of two transformations, namely $g_0^\alpha(\mathbf{y})$ and $g_1^\alpha(\mathbf{y})$ and that the inverse of $g_0^\alpha(\mathbf{y})$ is $\mathbf{z}_\alpha(\mathbf{x})$ and that of $g_1^\alpha(\mathbf{y})$ is $\mathbf{z}_\alpha(\mathbf{x})/w_\alpha^{*2}(\mathbf{x})$. The required Jacobians of the transformations can be derived and are given in the following two Lemmas, with corresponding proofs in the Appendix.

Lemma 1. *The Jacobian of $\mathbf{z}_\alpha(\mathbf{x})$ is*

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

Lemma 2. *The Jacobian of $\mathbf{z}_\alpha(\mathbf{x})/w_\alpha^{*2}(\mathbf{x})$ is*

$$|J_\alpha^1| = D^{D-1+\frac{1}{2}} \prod_{i=1}^D \frac{x_i^{\alpha-1}}{\sum_{j=1}^D x_j^\alpha} \times \left(\frac{1}{\alpha w_\alpha^*(\mathbf{x})} \right)^{2(D-1)}.$$

The densities based on the transformations g_0^α and g_1^α are then given in Equations (16) and (17), respectively, as follows

$$f_{\mathbf{x}_0}(\mathbf{x}|\alpha) = \frac{|J_\alpha^0|}{|2\pi\Sigma_\alpha|^{1/2}} \exp \left[-\frac{1}{2} (\mathbf{z}_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha)^\top \Sigma_\alpha^{-1} (\mathbf{z}_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha) \right] \quad (16)$$

$$f_{\mathbf{x}_1}(\mathbf{x}|\alpha) = \frac{|J_\alpha^1|}{|2\pi\Sigma_\alpha|^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{\mathbf{z}_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha \right)^\top \Sigma_\alpha^{-1} \left(\frac{\mathbf{z}_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha \right) \right]. \quad (17)$$

If we let p denote the probability that $\mathbf{y} \in \mathbb{A}^{D-1}$, the distribution of $\mathbf{x} \in \mathbb{S}^{D-1}$ can be written as

$$f_{\mathbf{x}}(\mathbf{x}|\alpha, p) = p f_{\mathbf{x}_0}(\mathbf{x}|\alpha) + (1-p) f_{\mathbf{x}_1}(\mathbf{x}|\alpha). \quad (18)$$

In words, $1-p$ is the probability that \mathbf{y} needs to be folded into the simplex through the transformation g_1^α in Equation (14). Substituting the density functions in Equations (16) and (17) into Equation (18), it follows that

$$\begin{aligned} f_{\mathbf{x}}(\mathbf{x}|\alpha, p, \boldsymbol{\mu}_\alpha, \Sigma_\alpha) = & \\ p \frac{|J_\alpha^0|}{|2\pi\Sigma_\alpha|^{1/2}} \exp \left[-\frac{1}{2} (z_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha)^\top \Sigma_\alpha^{-1} (z_\alpha(\mathbf{x}) - \boldsymbol{\mu}_\alpha) \right] & \\ + (1-p) \frac{|J_\alpha^1|}{|2\pi\Sigma_\alpha|^{1/2}} \exp \left[-\frac{1}{2} \left(\frac{z_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha \right)^\top \Sigma_\alpha^{-1} \left(\frac{z_\alpha(\mathbf{x})}{w_\alpha^{*2}(\mathbf{x})} - \boldsymbol{\mu}_\alpha \right) \right] & \end{aligned} \quad (19)$$

where $\mathbf{x} \in \mathbb{S}^{D-1}$, $\alpha \in [-1, 1]$, $0 \leq p \leq 1$. We will refer to the density in Equation (19) as the α -folded multivariate normal distribution.

Although we have indicated that $\mathbf{x} \in \mathbb{S}^{D-1}$, it is important to note that boundaries on the simplex (that is, zero values) are not allowed due to the product in Lemma 1. 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 & Welsh 2011a), assuming latent variables (Butler & Glasbey 2008) or conditioning on the zero values (Stewart & Field 2011; Tsagris & Stewart 2018).

When $p = 1$, the second term in Equation (19) vanishes and we end up with the α -normal distribution (Tsagris, Preston & Wood 2011). When $\alpha = 1$, $|J_\alpha^0| = D^{D-1+\frac{1}{2}}$. Finally, when $\alpha \rightarrow 0$, Equation (19) reduces to the multivariate logistic normal on \mathbb{S}^{D-1} (Aitchison

2003).

$$f_{\mathbf{x}}(\mathbf{x}) = \frac{1}{|2\pi\boldsymbol{\Sigma}|^{1/2}} \exp \left[-\frac{1}{2} (z_0(\mathbf{x}) - \boldsymbol{\mu}_0)^\top \boldsymbol{\Sigma}_0^{-1} (z_0(\mathbf{x}) - \boldsymbol{\mu}_0) \right] \prod_{i=1}^D x_i^{-1}, \quad (20)$$

where $z_0(\mathbf{x})$ is defined in Equation (4).

We can mimic Aitchison's Definition 6.2 of the additive logistic normal distribution Aitchison (2003) and similarly define the α -folded multivariate normal distribution by a transformation to multivariate normality. Specifically, a composition $\mathbf{x} \in \mathbb{S}^{D-1}$ is said to follow the α -folded multivariate normal distribution if \mathbf{y} follows a multivariate normal distribution in \mathbb{R}^{D-1} where

$$\mathbf{y} = \begin{cases} \mathbf{z}_\alpha(\mathbf{x}) & \text{with probability } p \\ \frac{\mathbf{z}_\alpha(\mathbf{x})}{w_\alpha^*(\mathbf{x})}, & \text{with probability } 1 - p, \end{cases} \quad (21)$$

$w_\alpha^*(\mathbf{x}) = \alpha \min \{w_\alpha(\mathbf{x})\}$ and $\mathbf{z}_\alpha(\mathbf{x})$ is defined in Equation (10).

Unlike in the case of log-ratio transformations to normality, numerical optimisation is required for obtaining the maximum likelihood estimates of the parameters in the α -folded multivariate normal distribution and the procedure is fully described in Subsection 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 α -transformation being a subset of \mathbb{R}^{D-1} . With the proposed distribution, any ignored probability is folded back onto the simplex, and hence the density has the form of Equation (19).

3.2. Contour plots of the α -folded bivariate normal distribution

To help visualise the α -folded multivariate normal distribution, we consider the two-dimensional case and compare the contours 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 (19) plotted on the simplex, with $\alpha = 1$. We consider two settings in which $\boldsymbol{\mu} = (0.561, 0.547)^\top$ in both cases, but the covariances matrices differ as follows

$$\boldsymbol{\Sigma}_1 = \begin{pmatrix} 0.5 & 0.25 \\ 0.25 & 0.35 \end{pmatrix} \text{ and } \boldsymbol{\Sigma}_2 = \begin{pmatrix} 2.5 & 1.25 \\ 1.25 & 1.75 \end{pmatrix}. \quad (22)$$

The mean vector was selected by applying the α -transformation in Equation (10) with $\alpha = 1$ to a sub-composition formed by the first three components of the Hongite data (an

artificial data set used by Aitchison (2003)). The elements of the first covariance matrix (Σ_1) were chosen so that the correlation was positive, whereas the second covariance matrix is such that $\Sigma_2 = 5\Sigma_1$.

While for the unconstrained normal case, the density was simply calculated for all grid points, for the folded model, the transformation in Equation (14) was first applied (with $\alpha = 1$) and then the density in Equation (19) was calculated.

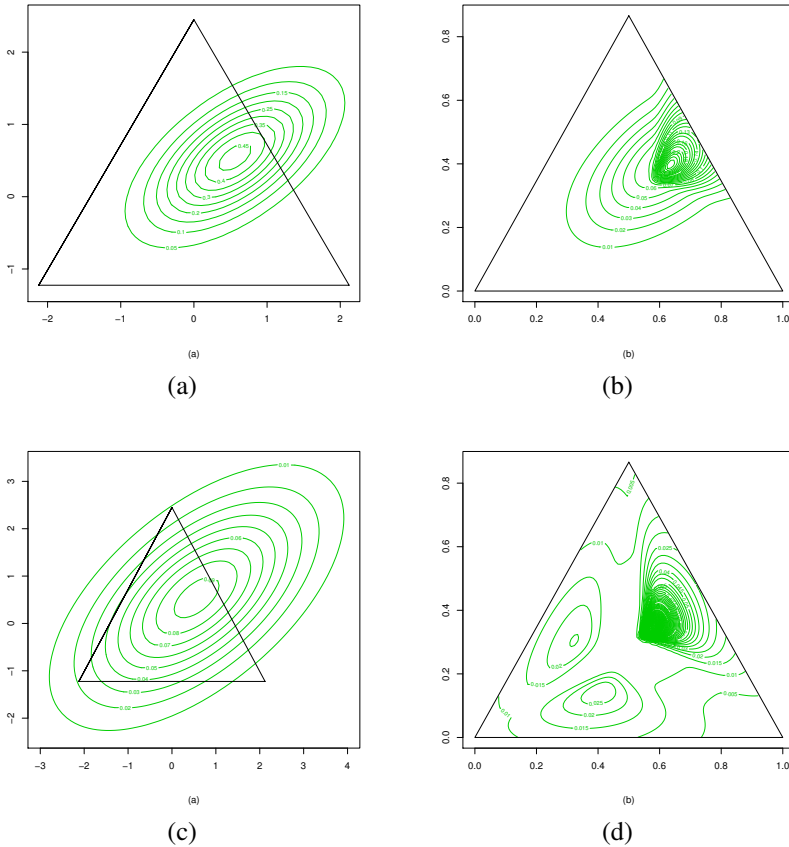


Figure 1. All contour plots refer to a normal distribution with mean $\mu = (0.561, 0.547)^\top$. The covariance matrix of the first row is Σ_1 and of the second row is $\Sigma_2 = 5\Sigma_1$. The left plot is the normal distribution in real space and the triangle (the simplex after the α -transformation in (10) 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 1(a) and 1(c) depict the contours of the multivariate normal distribution (defined in \mathbb{R}^2) while Figures 1(b) and 1(d) show the contours of the α -folded normal (defined on the simplex). The first row corresponds to Σ_1 in Equation (22) while the second row is derived from Σ_2 . The triangles in the second column (Figures 1(b) and 1(d)) display the simplex while

the corresponding triangles in Figures (a) and (c) were obtained through the α -transformation in Equation (10).

What is perhaps most evident from the contour plots is that points falling outside the triangle in Figures 1(a) and 1(c) result in modes inside the simplex in Figures 1(b) and 1(d) respectively. When there is a high probability of being left outside of two or more sides of the triangle (or faces of a 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 (14). 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 is not too complicated mainly because there are not too many parameters involved in the maximisation of the log-likelihood. In particular, when $D = 3$, there are 2 parameters for the mean vector, 3 for the covariance matrix as well as one for each of α and the probability p , and we can use the *simplex* algorithm of Nelder & Mead (1965) to maximise the logarithm of Equation (19), available via the command `optim` in `R` (R Core Team 2018). This algorithm is generally robust and is derivative free (Venables & Ripley 2002). However when moving to higher dimensions (roughly $D = 5$ or larger), the maximisation is not straightforward.

For this reason, we use the EM algorithm (McLachlan & Krishnan 2007) to maximise the log-likelihood corresponding to Equation (19). The corresponding code may be found in the R package *Compositional* (Tsagris, Athineou & Alenazi 2019). Following Jung, Foskey & Marron (2011) who applied the EM algorithm in the context of a univariate folded normal, we propose the algorithm below to estimate the unknown parameters $(\alpha, p, \boldsymbol{\mu}_\alpha$ and $\boldsymbol{\Sigma}_\alpha)$ from a sample of n compositional vectors \mathbf{X} in \mathbb{S}^{D-1} .

EM 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 $\mathbf{W}_\alpha(\mathbf{x})$.

Step 2. Calculate $w_{\alpha i}^*$ for each vector $\mathbf{w}_{\alpha i}$, $i = 1, \dots, n$, where $w_{\alpha i}^* = \alpha \min \{\mathbf{w}_{\alpha i}\}$.

Step 3. Left multiply each $\mathbf{w}_{\alpha i}$ by the Helmert sub-matrix \mathbf{H} to obtain \mathbf{y}_{0i}^α . Then multiply each \mathbf{y}_{0i}^α by $1/w_{\alpha i}^{*2}$ to obtain \mathbf{y}_{1i}^α (see Equation (21)). The \mathbf{y}_{0i}^α and \mathbf{y}_{1i}^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. Let $\mathbf{x}_{0i}^\alpha = g_0^\alpha(\mathbf{y}_{0i}^\alpha)$ and $\mathbf{x}_{1i}^\alpha = g_1^\alpha(\mathbf{y}_{1i}^\alpha)$.

Step 4. Choose initial values for the estimates of the parameters, for example

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

and

$$\hat{t}_i^0 = \frac{f_{\mathbf{x}_0}(\mathbf{x}_{0i}^\alpha)}{f_{\mathbf{x}_0}(\mathbf{x}_{0i}^\alpha) + f_{\mathbf{x}_1}(\mathbf{x}_{1i}^\alpha)} \quad \& \quad \hat{p}^0 = \frac{\sum_{i=1}^n \hat{t}_i^0}{n},$$

where \hat{t}_i is the estimated conditional expectation of the indicator function that indicates whether the i -th observation belongs to $f_{\mathbf{x}_0}$ or $f_{\mathbf{x}_1}$.

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

$$\hat{\boldsymbol{\mu}}_\alpha^k = \frac{\sum_{i=1}^n \hat{t}_i^{k-1} \mathbf{y}_{0i}^\alpha + \sum_{i=1}^n (1 - \hat{t}_i^{k-1}) \mathbf{y}_{1i}^\alpha}{n},$$

$$\widehat{\Sigma}_\alpha^k = \frac{1}{n} \left[\sum_{i=1}^n \hat{t}_i^{k-1} (\mathbf{y}_{0i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k) (\mathbf{y}_{0i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k)^\top + \sum_{i=1}^n (1 - \hat{t}_i^{k-1}) (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k) (\mathbf{y}_{1i}^\alpha - \hat{\boldsymbol{\mu}}_\alpha^k)^\top \right],$$

$$\hat{t}_i^k = \frac{f_{\mathbf{x}_0}(\mathbf{x}_{0i}^\alpha)}{f_{\mathbf{x}_0}(\mathbf{x}_{0i}^\alpha) + f_{\mathbf{x}_1}(\mathbf{x}_{1i}^\alpha)}$$

and

$$\hat{p}^k = \frac{\sum_{i=1}^n \hat{t}_i^k}{n}.$$

Step 6. Repeat Step 5 until the change in the log-likelihood $\ell_\alpha = \sum_{i=1}^n \log [p f_{\mathbf{x}_0}(\mathbf{x}_{0i}^\alpha) + (1-p) f_{\mathbf{x}_1}(\mathbf{x}_{1i}^\alpha)]$ between two successive values is less than a tolerance value, where $f_{\mathbf{x}_0}$ and $f_{\mathbf{x}_1}$ are given in Equations (16) and (17) respectively.

The above described procedure should be repeated for a grid of values of α , for example $(-1, -0.9, \dots, 0.9, 1)$, and the value of α which maximises the log-likelihood is chosen as its estimate. A more efficient search for the best α is via Brent's algorithm (Brent 2013). When $\alpha = 0$, the MLE estimates of the transformed data are obtained directly; no EM algorithm is necessary as all the probability is retained with the simplex. The fact that Equation (10) tends to Equation (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 (19) 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 (14), determine whether $\mathbf{y} \in \mathbb{A}^{D-1}$. To do this, compute $\mathbf{w} = \mathbf{H}^\top \mathbf{y}$. If $-1/\alpha \leq w_{i,\alpha} \leq (D-1)/\alpha$ 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 $\mathbf{x} = \mathbf{w}_\alpha^{-1}(\mathbf{H}^\top \mathbf{y} / q_\alpha^*(\mathbf{y}))$ where $q_\alpha^* = \alpha \min \{\mathbf{H}^\top \mathbf{y}\}$.

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

Step 1. Choose $\boldsymbol{\mu}$ 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 $\mathbf{w} = \mathbf{H}^\top \mathbf{y}$.

Step 4. Using Equation (9), compute $\mathbf{x} = \mathbf{w}_0^{-1}(\mathbf{w})$.

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 (20) (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 (19). 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 compositional data set, estimate the value of α obtained via the EM algorithm. This is the observed test statistic denoted by α_{obs} .

Step 2. Apply the α -transformation in Equation (10) with $\alpha = \alpha_{obs}$ to the data and map them onto set \mathbb{A}_α^{D-1} (11).

Step 3. Transform the data under the null hypothesis by applying the inverse of the isometric transformation with $\alpha = 0$ in Equation (4) to the data in Step 2.

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

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

$$p\text{-value} = \frac{\sum_{b=1}^B \mathbf{1}\{b : \alpha_b \geq \alpha_{obs}\} + 1}{B + 1}, \quad (23)$$

where $\mathbf{1}$ 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 could 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 data set and find the value of α for which the log-likelihood derived from Equation (19) is maximised for each bootstrap sample. By repeating this procedure many times, we can empirically estimate the distribution of $\hat{\alpha}$, and construct a variety of confidence intervals based on this distribution (see Davison & Hinkley (1997)).

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 variance of the maximum likelihood estimator (Cox & Hinkley 1979).

4. Data analysis examples

4.1. Example 1: Sharp's data set I

Our first example makes use of Sharp's (Sharp 2006) first 25×3 artificial data set, (termed "3a" by Sharp), made up from Aitchison's Hongite data (Aitchison 2003). We chose to analyze Sharp's artificial data (Sharp 2006) because they are curved data and according to Aitchison (2003) the logistic normal in Equation (20) should produce a very good fit for curved data. 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. Sharp (2006) showed that, surprisingly enough, 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 2(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 based on a χ^2 distribution with 2 degrees of freedom clearly rejects the logistic normal on the simplex (that is, that $\alpha = 0$ is the optimal transformation) and this conclusion is in line with the confidence interval limits. For this example, p was estimated to be approximately 0.95 so the probability of a point needing to be folded into the simplex was small (about 0.05) and the optimal value of α without folding, as proposed by Tsagris, Preston & Wood (2011), was equal to 0.428.

Figure 2(b) shows the ternary diagram of the data. In order to obtain the displayed α -mean, termed the Fréchet mean by Tsagris, Preston & Wood (2011), the EM algorithm estimate of $\mu_{0.419}$, say $\bar{y}_{0.419}$ is transformed inside the simplex using Equation (14). When $\alpha = 0$, the closed geometric mean can be similarly obtained through the inverse of the ilr transformation applied to \bar{y}_0 . That is, $\mathbf{w}_0^{-1}(\mathbf{H}^\top \bar{y}_0)$ is computed where \mathbf{w}_0 is given in

Equation (2) and \mathbf{w}_0^{-1} in Equation (8). The closed geometric mean, arithmetic mean and the Fréchet α -mean in \mathbb{S}^2 , are respectively

$$\begin{aligned}\hat{\boldsymbol{\mu}}_0 &= (0.707, 0.241, 0.051)^\top \text{ (Normalized geometric mean)} \\ \hat{\boldsymbol{\mu}}_1 &= (0.540, 0.275, 0.185)^\top \text{ (Simple arithmetic mean)} \\ \hat{\boldsymbol{\mu}}_{0.419} &= (0.622, 0.272, 0.106)^\top \text{ (Fréchet mean using the proposed folded model)}\end{aligned}$$

Note that the Fréchet α -mean with $\alpha = 0.428$, corresponding to Tsagris, Preston & Wood (2011), is $\hat{\boldsymbol{\mu}}_{0.428} = (0.619, 0.271, 0.110)^\top$ which, as expected, is very similar to $\hat{\boldsymbol{\mu}}_{0.419}$ for this example. We can clearly see that both the simple and the closed geometric mean fail to represent the data. However, the Fréchet α -mean calculated at $\alpha = 0.419$ achieves this goal.

The curved dashed lines correspond to the scores of the first principal component for each value of α . Principal component analysis for compositional data has been described by Aitchison (1983). The centered log-ratio transformation in Equation (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 as follows

$$\widehat{\boldsymbol{\Sigma}}_\alpha^* = \mathbf{H}^\top \widehat{\boldsymbol{\Sigma}}_\alpha \mathbf{H}.$$

If $\alpha = 0$, the covariance is mapped to \mathbb{Q}_0^{D-1} (Equation (3)). Step 3 of the algorithm presented in Subsection 3.4 is used to back-transform the principal components onto the simplex.

The contour plots of the α -folded model with $\alpha = 0.419$ and $\alpha = 0$ appear in Figures 3(a) and 3(c) respectively. We generated 500 observations from each model and these are plotted in Figures 3(b) and 3(d). When $\alpha = 0.419$ the simulated data look more like the observed data, in contrast to the simulated data with $\alpha = 0$.

4.2. Example 2: Sharp's data set II

We analysed a second 25×3 artificial data set taken from Aitchison's Hongite data (Aitchison 2003), termed "3c" by Sharp (2006). Figure 4 presents the relevant graphical information about the estimation of the fitted α -folded model, namely the profile log-likelihood of α , the ternary diagram of the data (showing the curvature of the data), and simulated data from both the fitted α -folded model as well as the multivariate normal distribution applied to the α -transformed data with $\alpha = 0$.

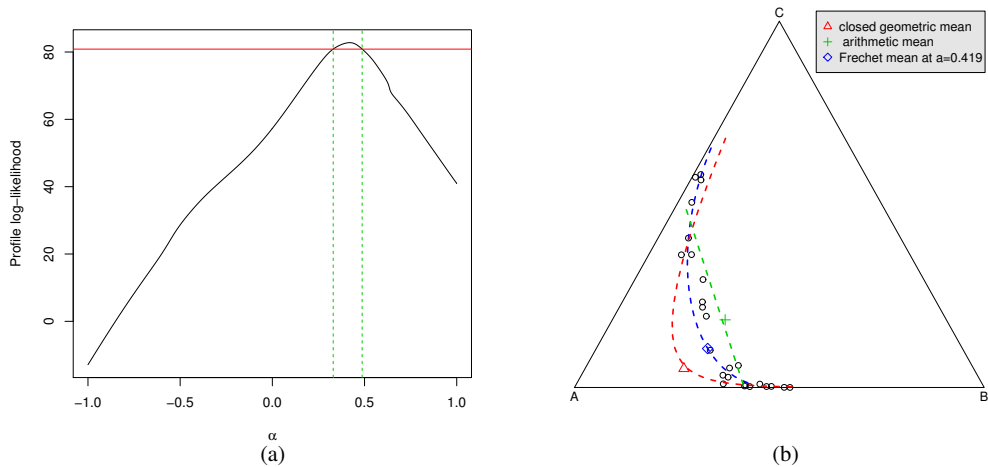


Figure 2. **Sharp's data set I** (a) Profile log-likelihood of α . The vertical dashed lines indicate the 95% confidence interval of the true value of the parameter. (b) Ternary plot of the data along with three α -means evaluated at $\alpha = 0$ (geometric mean normalised to sum to 1), $\alpha = 0.419$ and $\alpha = 1$ (arithmetic mean). The lines correspond to the scores of the first principal component for each value of α .

This is an example where the α -folded model has failed to capture the structure of the data. However, its competing model (the multivariate normal distribution applied to the isometric log-ratio transformed data) is clearly an even less attractive model for these data. We attribute this failure to the construction of the data. Sharp (2006) created these data in a manner that forced the log-ratio transformation to fail even though the data exhibit curvature. Curvature in the data is an argument favoring the log-ratio transformation (Aitchison 2003).

In terms of the potential benefit of folding for this example, the estimated probability left outside of the simplex was only 0.02 ($\hat{p} = 0.98$) and the optimal α using the model in Tsagris, Preston & Wood (2011) without folding was equal to 1, compared to 0.774 with folding. The Fréchet mean evaluated at $\alpha = 0.774$ was equal to $(0.553, 0.270, 0.177)^\top$ while the Fréchet mean evaluated at $\alpha = 1$ was similar and equal to $(0.540, 0.276, 0.184)^\top$.

4.3. Example 3: Labor force data

To illustrate the need for folding using our proposed model in a real-life setting, we consider a third data set in the area of economics. The data is accessible via the R package *rob-Compositions* (Templ, Hron & Filzmoser 2011) and comes from the United Nations' statistics division website (<https://unstats.un.org/unsd/demographic/products/indwm/default.htm>) which provides statistics and indicators on women and men for various 'fields of concern'. In particular, we analyze compositional labour force data for 124 countries and 6 categories determined by status in employment (either employee, employer or 'own-account workers




Figure 3. **Sharp's data set I** 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.

and contributing family workers') and gender (either male or female). For example, the first country listed in the data set is Algeria where, in 2004, the percentage of female employees was 49.8%, of male employees was 61.9%, of female employers was 1.3%, of male employers was 5.8%, of female owners was 48.9% and, finally, of male owners was 32.0%.

To help visualise the data in a manner that allows for data on men and women to be compared, ternary plots of the 124 3-part female subcompositions and of the 3-part male subcompositions are presented in Figure 5. Although the two plots are similar, for many countries, there appears to be a higher percentage of male employers than female employers.

The analysis of this data set again suggests that an α value other than zero may provide an improved fit. The optimal value of α using the model without folding in Tsagris, Preston & Wood (2011) was equal to 0.328, whereas its value using

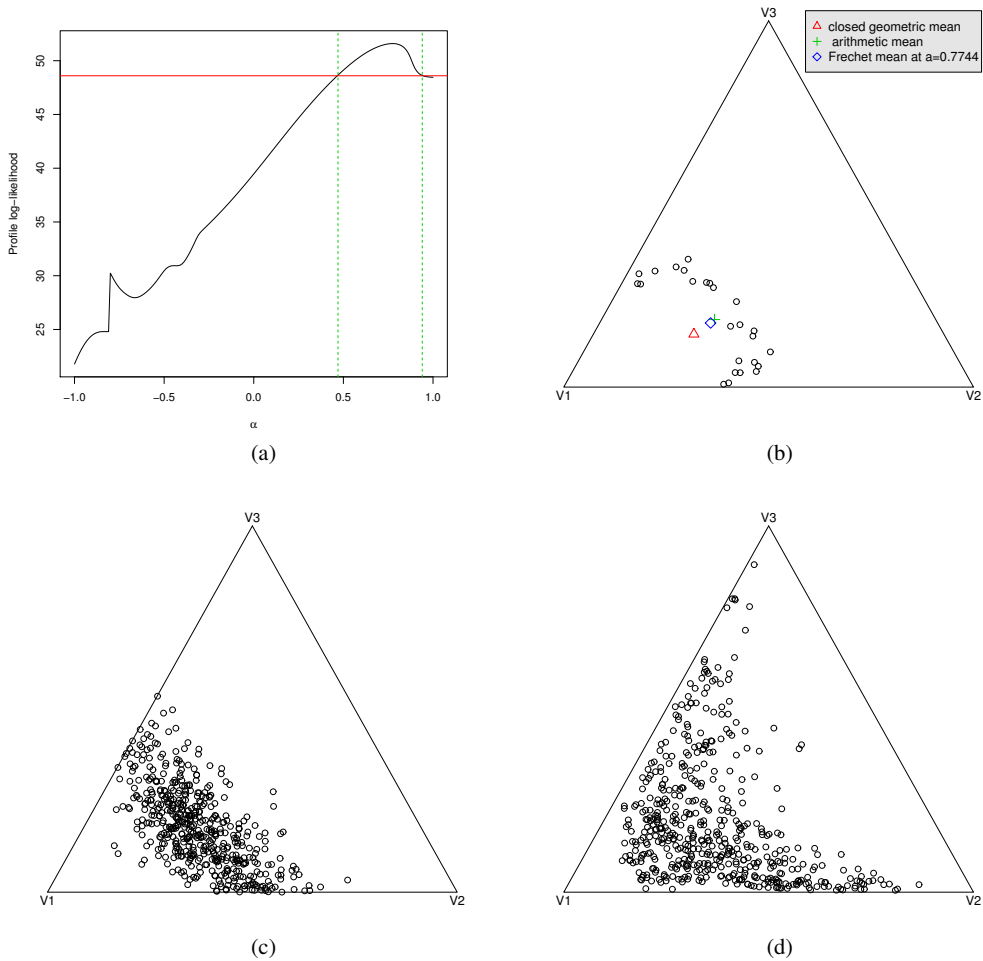


Figure 4. **Sharp's data set II** (a) Profile log-likelihood of α . The vertical dashed lines indicate the 95% confidence interval of the true value of the parameter. (b) Ternary plot of the data along with three Fréchet-means evaluated at $\alpha = 0$ (geometric mean normalised to sum to 1), $\alpha = 0.774$ and $\alpha = 1$ (arithmetic mean). Simulated data from the folded model with (c) $\alpha = 0.774$ and (d) $\alpha = 0$.

the proposed folded model was equal to 0.516. Furthermore, there appears to be a need for the folding transformation as the estimated probability of a point being left outside the simplex was large and equal to 0.756. The two Fréchet mean vectors were roughly equal ($\mu_{0.328} = (0.350, 0.366, 0.010, 0.024, 0.122, 0.127)^\top$ and $\mu_{0.516} = (0.348, 0.357, 0.011, 0.025, 0.129, 0.130)^\top$) and suggest that, on average, the percentage of female and male employees as well as owners are similar, but the percentage of male employers (while small) is roughly double that of female employers. In terms of the similarity

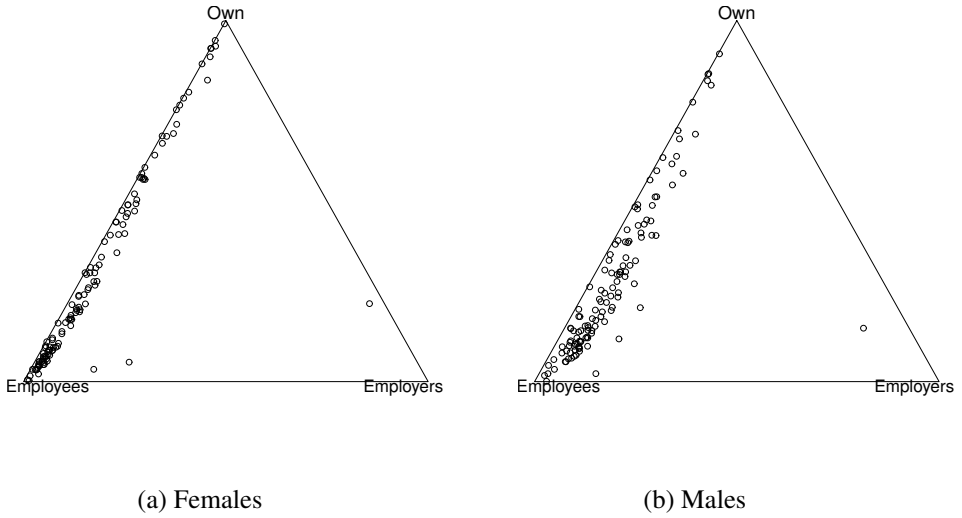


Figure 5. Ternary plots of subcompositions of the Labor Force data. The percentagewise allocation of females (a) and (b) males in the categories "Employees", "Employers" and "Own" (referring to own-account workers and contributing family workers) is shown.

between the two mean vectors, this is not surprising since as the number of components (or dimensionality) grows, the volume of the simplex becomes smaller and the mean vectors will be close for a wide range of values of α . The difference, however, between the two models can be observed in the estimated covariance matrices which are given below:

$$\widehat{\Sigma}_{0.328} = \begin{pmatrix} 0.071 & 0.086 & 0.058 & 0.249 & 0.150 \\ 0.086 & 0.538 & 0.340 & 0.557 & 0.356 \\ 0.058 & 0.340 & 0.255 & 0.367 & 0.233 \\ 0.249 & 0.557 & 0.367 & 1.854 & 1.242 \\ 0.150 & 0.356 & 0.233 & 1.242 & 0.906 \end{pmatrix},$$

$$\widehat{\Sigma}_{0.516} = \begin{pmatrix} 0.101 & 0.355 & 0.219 & 0.402 & 0.219 \\ 0.355 & 2.627 & 1.574 & 2.368 & 1.499 \\ 0.219 & 1.574 & 0.987 & 1.493 & 0.940 \\ 0.402 & 2.368 & 1.493 & 3.351 & 2.171 \\ 0.219 & 1.499 & 0.940 & 2.171 & 1.522 \end{pmatrix}.$$

The elements of $\widehat{\Sigma}_{0.516}$ (with folding) are larger than those of $\widehat{\Sigma}_{0.328}$ (without folding), ranging from 1.4 times up to 4.8 times larger and the generalised variance of $\widehat{\Sigma}_{0.516}$ is 7 times

that of the first covariance matrix. These differences are the due to \hat{p} being relatively small in this example.

5. Simulation studies

5.1. Estimation of μ_α , Σ_α and p

In this simulation study, we examined the accuracy of the EM algorithm in terms of estimating μ_α , Σ_α and p for increasing sample sizes. Specifically, 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

$$\mu_{-0.5} = (1.715, 0.914, 0.115, 0.167)^\top \quad \text{and} \quad \mu_{0.5} = (-0.566, -0.979, -0.648, -0.651)^\top,$$

and covariance matrices of this form

$$\Sigma = \kappa \begin{pmatrix} 0.149 & -0.458 & 0.002 & -0.005 \\ -0.458 & 1.523 & 0.000 & 0.007 \\ 0.002 & 0.000 & 0.037 & -0.047 \\ -0.005 & 0.007 & -0.047 & 0.061 \end{pmatrix}, \quad (24)$$

where $\kappa = \{0.5, 1, 2, 3, 5, 7, 10\}$. Note that the value of κ changes the probability that a point is left outside of the simplex. The *true* values of p for each value of κ were computed through Monte-Carlo simulations by generating many random vectors from a multivariate normal distribution with the parameters above and computing the proportion of vectors that belong to \mathbb{A}^{D-1} . We will refer to the probability that a vector is outside of \mathbb{A}^{D-1} (that is, $1 - p$) as the probability left outside the simplex since these vectors need to be folded into the simplex. This estimated probability, for each α and κ combination, is presented in Table 1.

For each combination of α , μ and Σ , 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 , μ and Σ were calculated. All computations took place in R 3.2.3 (R Core Team, 2015) using a desktop computer with Intel Core i5 at 3.5 GHz processor and 32GB RAM. For various measures of distance (as described below), the mean distance between the 1000 estimates and the true parameters was calculated.

For the probability left outside of the simplex (that is, $1 - p$), the absolute difference between the estimated probability and the true probability was 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

(Förstner & Moonen, 2003) was calculated

$$d(\widehat{\Sigma}, \Sigma) = \sqrt{\sum_{i=1}^{D-1} \left[\log \Lambda_i(\widehat{\Sigma} \Sigma^{-1}) \right]^2}, \quad (25)$$

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 6.

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

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


From Figure 6 we observe that when the probability left outside the simplex grows larger (κ is larger), a larger sample size is required in order to get better estimates, for both the probability and the mean vector, whereas estimation of the covariance matrix seems to be unaffected by the probability left outside the simplex.

5.2. Estimation of α

We then examined the performance of our estimation algorithm in relation to α . We focused on the large sample case in order to gain insight into the asymptotic behavior of $\hat{\alpha}$. For this simulation study, we have generated data as in Subsection 5.1 with the mean vector set to $\boldsymbol{\mu} = (1.715, 0.914, 0.115, 0.167)^\top$ and the covariance matrices as in Equation (24). For values of α ranging from 0 up to 1 with a step of 0.1 we estimated these values for the different values of κ using 4 sample sizes ($n = \{1000, 5000, 10000, 20000\}$). For each combination of α , κ and n we used 1000 repetitions.

Figure 7 shows the average bias of the α estimates in boxplots for each sample size. Each box corresponds to a value of κ and is the average bias aggregated for all values of α . For example, Figure 7(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 . From the plots, as expected, the range in the variances increases with the value of κ , since higher values of κ correspond

Simulation study when $\alpha = -0.5$



Simulation study when $\alpha = 0.5$




Figure 6. 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.

to a higher probability of being left outside of the simplex. Table 2 presents $1 - p$ for many combinations of values of α and κ calculated using Monte Carlo simulation with 20,000,000 repetitions. Clearly, $1 - p$ increases as either κ or α increases.

Table 2. Probability left outside the simplex for many combinations of α and κ values.

α	κ						
	0.5	1	2	3	5	7	10
0.0	0.000	0.000	0.000	0.000	0.000	0.000	0.000
0.2	0.000	0.000	0.001	0.007	0.034	0.071	0.128
0.4	0.006	0.047	0.156	0.245	0.367	0.445	0.522
0.6	0.132	0.284	0.448	0.536	0.632	0.687	0.741
0.8	0.398	0.551	0.673	0.731	0.794	0.830	0.866
1.0	0.66	0.756	0.827	0.861	0.898	0.918	0.937




Figure 7. Box plots of the range of $\alpha - \hat{\alpha}$ as a function of κ for 4 different sample sizes.

5.3. Estimation of the computational cost

Using only $\alpha = 0.5$, we generated data as in Subsection 5.1 with increasing sample sizes and, for each sample size, recorded the time (in seconds) required to estimate the true value of α . The results are presented in Table 3. 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.

5.4. Estimation of data set parameters

Next, we evaluated the performance of the EM algorithm for parameter values observed in practice by using the estimated parameters from the four data sets analysed in Section 4 as

Table 3. Computational times (in seconds) required to estimate the value of α , averaged over 1000 repetitions.

Sample size	κ						
	0.5	1	2	3	5	7	10
100	0.073	0.079	0.092	0.106	0.135	0.174	0.212
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
1000	0.082	0.078	0.086	0.095	0.117	0.140	0.168
5000	0.638	0.933	0.838	0.841	0.816	0.675	0.722

the true parameters. To accomplish this, we generated compositional data using the proposed folded multivariate normal distribution (with the same number of dimensions and sample sizes as the real data) with estimated parameters $\hat{\alpha}$, \hat{p} , $\hat{\boldsymbol{\mu}}_\alpha$ and $\hat{\boldsymbol{\Sigma}}_\alpha$. For each data set, 1000 samples were simulated and we calculated the average distances of the estimated parameters from the true parameters, as described in Subsection 5.1. Table 4 shows the results of this Monte Carlo study. The large bias observed for α with the Sharp II data set is perhaps to be expected given the poor fit that was observed in Subsection 4.2. For Example 4 (*Labor Force*), the bias tends to be larger for \hat{p} , $\hat{\boldsymbol{\mu}}_\alpha$ and $\hat{\boldsymbol{\Sigma}}_\alpha$ compared to the other data sets.

Table 4. Estimated bias of the parameters in the real data sets using Monte Carlo.

Data set	Estimated bias					
	n	D	α	p	$\boldsymbol{\mu}_\alpha$	$\boldsymbol{\Sigma}_\alpha$
Sharp I	25	3	0.054	0.042	0.248	0.633
Sharp II	25	3	0.219	0.013	0.145	0.594
Labor Force	124	6	0.049	0.510	1.257	1.628

5.5. Comparison of the folded and simple α -transformations

We will now illustrate the effect of the folding transformation on the 3 aforementioned data sets. For each data set we generated data from the α -folded multivariate normal distribution using the estimated parameters as the true parameters, but with some modifications to induce various values of p . Specifically for the two Sharp's artificial data sets, we generated data from the estimated parameters but multiplied all the elements of the covariance matrix by κ , for various values of κ . Similarly, for the *Labor Force* data set we first multiplied all the elements of the mean vector by different values of λ and then generated data. The values of κ and λ affect the probability left outside the simplex. For every generated data set we estimated the parameters of the α -folded multivariate normal distribution and of the α -normal (Tsagris et al, 2011). In both cases we generated 1000 compositional vectors.

The accuracy of the mean vector was evaluated using the Euclidean distance, for the estimated covariance matrix the metric defined in (25) was used and the absolute difference between the true and the estimated value of α was computed. The results for the two Sharp's artificial data sets were similar and hence we only present the results for the first data set in Figure 8, whereas Table 5 contains the results for the *Labor Force* data set.

Figure 8 depicts the effect of the probability left outside the simplex. When this probability is considerably high, the α -normal fails to fit the data adequately and the estimates are highly biased. The results in Table 5 also suggest that the bias in the estimates decreases often substantially for the α -folded normal model as the probability left outside the simplex increases, but increases (though sometimes only slightly) for the α -model.

Table 5. Estimated differences between the estimated and the true parameters. For each of the three parameters and the probability left outside the simplex ($1 - p$), the smaller value is highlighted in bold. The α -folded normal is the model proposed in this paper, whereas the α -normal refers to the model proposed by Tsagris et al. (2011).

$1 - p$	Mean Vector		Covariance Matrix		α	
	α -folded normal	α -normal	α -folded normal	α -normal	α -folded normal	α -normal
0.246	1.277	1.619	1.515	2.016	0.017	0.024
0.586	0.729	1.668	0.939	2.014	-0.031	0.026
0.701	0.422	1.716	0.564	2.011	-0.035	0.029

6. Conclusions

In this paper we developed a novel parametric model, with nice properties, for compositional data analysis. The model is an extension of the model proposed by Tsagris, Preston & Wood (2011) in which the α -transformation and corresponding multivariate normal distribution was introduced. A drawback of their model is that it does not take into account the probability left outside the simplex, and this deficiency motivated the development of our proposed folded model. Simulation study results suggest that if the probability left outside the simplex is large (as it was for one of the real-life examples), parameter estimates using the model in Tsagris, Preston & Wood (2011) tend to be quite biased. While we dealt with the probability left outside the simplex through folding, another possible solution would be to use truncation (Dobigeon & Tournet 2007).

The proposed model is also an extension of the popular logistic normal distribution which corresponds to $\alpha = 0$. In the results that were presented, the α -folded model appeared to fit the data adequately when the logistic normal distribution did not. This is inline with other work in which the log-ratio transformation failed to capture the variability of the data. See, for example, (Tsagris, Preston & Wood 2011; Baxter & Freestone 2006; Sharp 2006).

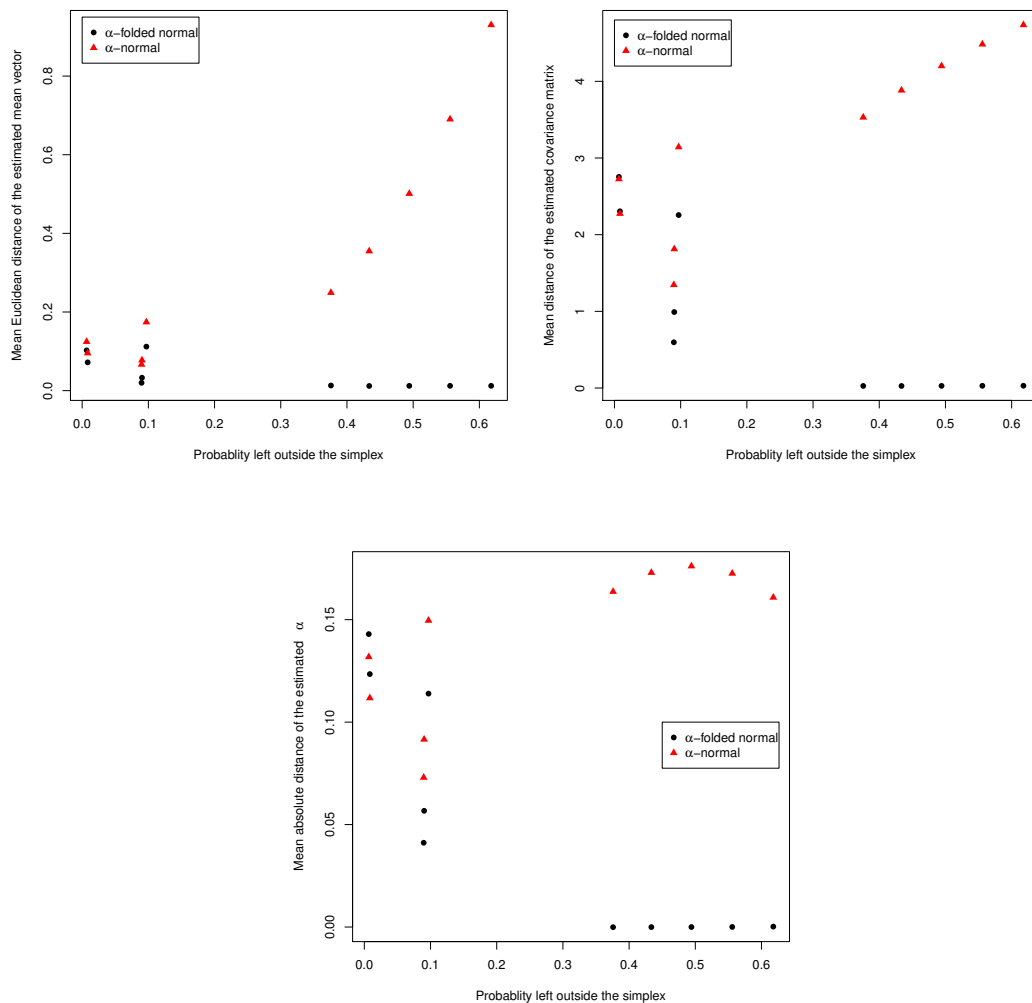


Figure 8. Mean estimated distances (for the mean vector, covariance matrix and values of α) between the true and estimated parameters of the α -folded multivariate normal distribution and of the multivariate normal distribution after applying the α -transformation for a range of probabilities left outside the simplex.

The use of a multivariate model other than the multivariate normal distribution, such as the multivariate skew normal distribution (Azzalini & Valle 1996) has also been suggested. The challenge, 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. Another, perhaps simpler, alternative model is the multivariate t distribution. Bayesian analysis and regression modeling are two suggested research directions.

As previously mentioned and similar to the Box-Cox transformation, zero values are not compatible with our proposed model. Note that the zero issue also arises with the logistic normal distribution. However, it is possible to generalise most of the analyzes suggested for the logistic normal distribution using our proposed folded model, including extensions that allow zeros. Furthermore, results (not shown) suggest that our estimation algorithm is most reliable for dimensions less than 30 and that for dimensions between 20 and 30, a large sample size is required.

As with the isometric log-ratio transformation, the α -transformed data have no clear interpretation. The same issue occurs with the additive log-ratio transformation $\left(\{\log(x_i/x_D)\}_{i=1,\dots,D-1}\right)$ (Aitchison 2003) and while the mean vector and covariance matrix are interpretable, they depend on the denominator component. Hence among models for compositional data, there is a trade-off between interpretability and better fit. Presumably the application would dictate which is of greater importance.

As is standard practice in log-ratio transformation analysis, if one is willing to exclude from the sample space the boundary of the simplex, which includes observations that have one or more components equal to zero, then the folded α -transformation (10) and its inverse are well defined for all $\alpha \in \mathbb{R}$, and the corresponding α -folded model provides a new approach for the analysis of compositional data with the potential to provide an improved fit over traditional models.

Appendix

The Helmert sub-matrix

The Helmert matrix is a $D \times D$ orthogonal matrix. The Helmert sub-matrix has the first row omitted, hence is a $(D - 1) \times D$ matrix, the structure of which is presented below.

$$\mathbf{H} = \begin{pmatrix} \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} & 0 & \dots & \dots & 0 \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & -\frac{2}{\sqrt{6}} & 0 & \dots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \frac{1}{\sqrt{i(i+1)}} & \dots & \dots & \frac{1}{\sqrt{i(i+1)}} & -\frac{i}{\sqrt{i(i+1)}} & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & 0 \\ \frac{1}{\sqrt{(D-1)D}} & \dots & \dots & \dots & \frac{1}{\sqrt{(D-1)D}} & -\frac{D-1}{\sqrt{(D-1)D}} \end{pmatrix} \quad (26)$$

Explanation of the folding transformation

We will show that if $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1}$ then Equation (12) transforms \mathbf{y} from $\mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1}$ to \mathbb{S}^{D-1} . We will consider only the case that $\alpha > 0$ as the case for $\alpha < 0$ is similar.

Suppose $\mathbf{y} \in \mathbb{R}^{D-1} \setminus \mathbb{A}^{D-1}$ and let $\mathbf{w} = \mathbf{H}^\top \mathbf{y}$. Then $\mathbf{w} \in \mathbb{Q}_0^{D-1}$ from Equation (4) and properties of \mathbf{H} . Also, for $\alpha \neq 0$, $\mathbf{w} \notin \mathbb{Q}_\alpha^{D-1}$ and this implies that there exists a $w_i < -1/\alpha$ (in which case, $\min(\mathbf{w}) < -1/\alpha$) or a $w_i > (D-1)/\alpha$ (in which case, $\max(\mathbf{w}) > (D-1)/\alpha$). Note, however, that we need only to consider the case in which $\min(\mathbf{w}) < -1/\alpha$ because if $\min(\mathbf{w}) \not< -1/\alpha$, we can show that $\max(\mathbf{w}) \leq (D-1)/\alpha$. Specifically, assume that $\min(\mathbf{w}) \geq -1/\alpha \Rightarrow w_i \geq -1/\alpha \forall i = 1, \dots, D$. If we let $w_D = \max(\mathbf{w})$, we have that

$$-\sum_{i=1}^{D-1} \frac{1}{\alpha} \leq \sum_{i=1}^{D-1} w_i \Rightarrow -\frac{D-1}{\alpha} \leq \sum_{i=1}^{D-1} w_i = -w_D \Rightarrow w_D \leq \frac{D-1}{\alpha}, \quad (27)$$

since $\mathbf{w} \in \mathbb{Q}_0$ and $\sum_i^D w_i = 0$.

Assume $\min(\mathbf{w}) < -1/\alpha$ and define

$$\mathbf{w}'^\alpha = \frac{\mathbf{w}}{\alpha |\min(\mathbf{w})|}. \quad (28)$$

Since $\sum_i^D w_i = 0$, at least one component must be negative (if not, all components are zero) and if we divide by the absolute value of the smallest component, it is straightforward to show that $\min(\mathbf{w}'^\alpha) = -1/\alpha$. Note that if Equation (8) is applied to \mathbf{w}'^α , \mathbf{y} is transformed to the boundary of the simplex (that is, one component of the resulting vector will be zero).

Let $q^* = \alpha \min(\mathbf{w})$. To transform \mathbf{w} inside \mathbb{Q}_α^{D-1} (rather than on the boundary), consider

$$\mathbf{w}^\alpha = \frac{\mathbf{w}}{q^{*2}}.$$

We need to show that $\mathbf{w}^\alpha \in \mathbb{Q}_\alpha^{D-1}$ since in this case, we can apply Equation (8) to transform \mathbf{w}^α to the simplex. Clearly $\sum_{i=1}^D \mathbf{w}_i^\alpha = 0$ so we need to show that $-1/\alpha < \mathbf{w}_i^\alpha < (D-1)/\alpha$. Since $\min(\mathbf{w}) < -1/\alpha \Rightarrow q^* < -1 \Rightarrow q^{*2} > 1$. Therefore $\min(\mathbf{w}/q^{*2}) > -1/\alpha$. To show that $\max(\mathbf{w}) < (D-1)/\alpha$, we can simply argue as we did in Equation (27) that $\min(\mathbf{w}/q^{*2}) > -1/\alpha \Rightarrow \max(\mathbf{w}/w^{*2}) < (D-1)/\alpha$. Therefore, $\mathbf{w}^\alpha \in \mathbb{Q}_\alpha^{D-1}$ and can be transformed inside the simplex via Equation (8).

Proof of Lemma 1

Let us begin by deriving the Jacobian determinant of Equation (5). The map in Equation (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

Equation (10) non-degenerate we consider the version of Equation (5) as follows

$$\mathbf{u}_a \{(x_i)\} = \frac{x_i^\alpha}{\sum_{j=1}^{D-1} x_j^\alpha + \left(1 - \sum_{j=1}^{D-1} x_j\right)^\alpha} \quad i = 1, \dots, D-1. \quad (29)$$

Equation (29) is presented to highlight that in fact we have $D-1$ and not D components (variables). Let us derive the Jacobian of Equation (5) or Equation (29). We denote $S(\alpha) = \sum_{j=1}^D x_j^\alpha$, where $x_D = 1 - \sum_{j=1}^{D-1} x_j$. The diagonal and the non-diagonal elements of the Jacobian matrix are as 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 (Mardia, Kent & Bibby 1979):

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

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

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

Then

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

and

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

So we end up with

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

Finally the Jacobian of Equation (5) takes the following form

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

The Jacobian of the α -transformation (10) without the left multiplication by the Helmert sub-matrix \mathbf{H} is simply the Jacobian of Equation (5) multiplied by D^{D-1}/α^{D-1}

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

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 2

We will prove Lemma 3.2 for the case that $\alpha = 1$ for convenience purposes. Furthermore, we will only prove the additional term in the Jacobian resulting from the division of $w_\alpha^{*2}(\mathbf{x})$, since this is the only difference between Lemma 3.1 and Lemma 3.2. To accomplish this with $\alpha = 1$, and using simplified notation, we need to find the Jacobian of

$$\mathbf{y} = \left(\frac{1}{w^*(\mathbf{x})} \right)^2 \mathbf{w}(\mathbf{x}) = \left(\frac{1}{w^*} \right)^2 \mathbf{w}.$$

The component wise transformation can be expressed as

$$y_i = Z_i \left(\frac{1}{w^*} \right)^2 w_i + (1 - Z_i) \left(\frac{1}{w_D} \right)^2 w_i, \quad (30)$$

where w_i refers to the i -th component of $\mathbf{w}_\alpha(\mathbf{x})$ defined in Equation (6) with $\alpha = 1$. Also,

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

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 $y_i = (1/w^*)^2 w_i$. There are two sub-cases to be specified.

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

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

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

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

The Jacobian matrix is

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

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 $y_i = w_i/w_D^2$. 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

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

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

The Jacobian matrix in this case can be written as

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

The determinant of such matrices is given by (Mardia, Kent, & Bibby, 1979)

$$|J| = \left| \mathbf{A} + \mathbf{B}\mathbf{C}^\top \right| = |\mathbf{A}| (1 + \mathbf{C}^\top \mathbf{A}^{-1} \mathbf{B}) \quad (31)$$

where

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

Therefore,

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

Finally, (31) becomes

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

References

- AITCHISON, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society. Series B* **44**, 139–177.
- AITCHISON, J. (1983). Principal component analysis of compositional data. *Biometrika* **70**, 57–65.
- AITCHISON, J. (2003). *The statistical analysis of compositional data*. New Jersey: Reprinted by The Blackburn Press.
- AZZALINI, A. & VALLE, A. (1996). The multivariate skew-normal distribution. *Biometrika* **83**, 715–726.
- BAXTER, M., BEARDAH, C., COOL, H. & JACKSON, C. (2005). Compositional data analysis of some alkaline glasses. *Mathematical geology* **37**, 183–196.
- BAXTER, M. & FREESTONE, I. (2006). Log-ratio compositional data analysis in archaeometry. *Archaeometry* **48**, 511–531.
- BRENT, R.P. (2013). *Algorithms for minimization without derivatives*. Courier Corporation.
- 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.
- CHEN, E.Z. & LI, H. (2016). A two-part mixed-effects model for analyzing longitudinal microbiome compositional data. *Bioinformatics* **32**, 2611–2617.
- COX, D. & HINKLEY, D. (1979). *Theoretical Statistics*. London: Chapman & Hall/CRC.
- DAVISON, A. & HINKLEY, D. (1997). *Bootstrap methods and their application*. Cambridge: Cambridge University Press.
- DOBIGEON, N. & TOURNERET, J.Y. (2007). Truncated Multivariate Gaussian Distribution on a Simplex. *Technical report*, University of Toulouse.
- DRYDEN, I. & MARDIA, K. (1998). *Statistical Shape Analysis*. Chichester: John Wiley & Sons.
- EGOZCUE, J., PAWLOWSKY-GLAHN, V., MATEU-FIGUERAS, G. & BARCELÓ-VIDAL, C. (2003). Isometric logratio transformations for compositional data analysis. *Mathematical Geology* **35**, 279–300.
- FRY, J., FRY, T. & McLAREN, K. (2000). Compositional data analysis and zeros in micro data. *Applied Economics* **32**, 953–959.
- GHOSH, D. & CHAKRABARTI, R. (2009). Joint Variable Selection and Classification with Immunohistochemical Data. *Biomarker insights* **4**, 103.
- JOHNSON, N. (1962). The folded normal distribution: Accuracy of estimation by maximum likelihood. *Technometrics* **4**, 249–256.

- JUNG, S., FOSKEY, M. & MARRON, J.S. (2011). Principal arc analysis on direct product manifolds. *The Annals of Applied Statistics* **5**, 578–603.
- KATZ, J. & KING, G. (1999). A statistical model for multiparty electoral data. *American Political Science Review*, 15–32.
- LANCASTER, H. (1965). The Helmert matrices. *American Mathematical Monthly* **72**, 4–12.
- LE, H. & SMALL, C. (1999). Multidimensional scaling of simplex shapes. *Pattern Recognition* **32**, 1601–1613.
- LEONE, F., NELSON, L. & NOTTINGHAM, R. (1961). The folded normal distribution. *Technometrics* **3**, 543–550.
- MARDIA, K., KENT, J. & BIBBY, J. (1979). *Multivariate Analysis*. London: Academic Press.
- MCLACHLAN, G. & KRISHNAN, T. (2007). *The EM algorithm and extensions*. New Jersey: John Wiley & Sons.
- NELDER, J. & MEAD, R. (1965). A simplex algorithm for function minimization. *Computer Journal* **7**, 308–313.
- NEOCLEOUS, T., AITKEN, C. & ZADORA, G. (2011). Transformations for compositional data with zeros with an application to forensic evidence evaluation. *Chemometrics and Intelligent Laboratory Systems* **109**, 77–85.
- OEPPEIN, 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*.
- 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.
- 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.
- R CORE TEAM (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- SCEALY, J. & WELSH, A. (2011a). Properties of a square root transformation regression model. In *Proceedings of the 4th Compositional Data Analysis Workshop, Girona, Spain*.
- SCEALY, J. & WELSH, A. (2011b). Regression for compositional data by using distributions defined on the hypersphere. *Journal of the Royal Statistical Society. Series B* **73**, 351–375.
- SCEALY, J. & WELSH, A. (2014). Fitting Kent models to compositional data with small concentration. *Statistics and Computing* **24**, 165–179.
- 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.
- SHARP, W. (2006). The graph median—A stable alternative measure of central tendency for compositional data sets. *Mathematical geology* **38**, 221–229.
- SHI, P., ZHANG, A., LI, H. et al. (2016). Regression analysis for microbiome compositional data. *The Annals of Applied Statistics* **10**, 1019–1040.
- STEWART, C. & FIELD, C. (2011). Managing the Essential Zeros in Quantitative Fatty Acid Signature Analysis. *Journal of Agricultural, Biological, and Environmental Statistics* **16**, 45–69.
- TEMPL, M., HRON, K. & FILZMOSER, P. (2011). robCompositions: an R-package for robust statistical analysis of compositional data. In V. Pawlowsky-Glahn and A. Buccianti, editors, *Compositional Data Analysis. Theory and Applications*. Wiley Online Library, pp. 341–355.
- TSAGRIS, M. (2014). The k-NN algorithm for compositional data: a revised approach with and without zero values present. *Journal of Data Science* **12**, 519–534.
- TSAGRIS, M. (2015). Regression analysis with compositional data containing zero values. *Chilean Journal of Statistics* **6**, 47–57.

- TSAGRIS, M., ATHINEOU, G. & ALENAZI, A. (2019). *Compositional: Compositional Data Analysis*. URL <https://CRAN.R-project.org/package=Compositional>. R package version 3.5.
- TSAGRIS, M., PRESTON, S. & WOOD, A. (2011). A data-based power transformation for compositional data. In *Proceedings of the 4rth Compositional Data Analysis Workshop, Girona, Spain*.
- TSAGRIS, M., PRESTON, S. & WOOD, A.T. (2016). Improved classification for compositional data using the α -transformation. *Journal of Classification* **33**, 243–261.
- TSAGRIS, M. & STEWART, C. (2018). A Dirichlet regression model for compositional data with zeros. *Lobachevskii Journal of Mathematics* **39**, 398–412.
- VENABLES, W. & RIPLEY, B. (2002). *Modern applied statistics with S*. Springer Verlag.
- 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.

2020

A folded model for compositional data analysis

Tsagris, Michail

Wiley

<https://doi.org/10.1111/anzs.12289>

This is the peer reviewed version of the following article: [A folded model for compositional data analysis. Australian & New Zealand Journal of Statistics 62, 2 p249-277 (2020)], which has been published in final form at <https://doi.org/10.1111/anzs.12289>

Downloaded from UNB Scholar