



UNIVERSITY OF LEEDS

This is a repository copy of *Mixture Models for Spherical Data with Applications to Protein Bioinformatics*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/185512/>

Version: Accepted Version

---

**Book Section:**

Mardia, KV, Barber, S, Burdett, PM et al. (2 more authors) (2022) Mixture Models for Spherical Data with Applications to Protein Bioinformatics. In: SenGupta, A and Arnold, B, (eds.) Directional Statistics for Innovative Applications: A Bicentennial Tribute to Florence Nightingale. Forum for Interdisciplinary Mathematics . Springer Singapore , pp. 15-32. ISBN 978-9811910432

[https://doi.org/10.1007/978-981-19-1044-9\\_2](https://doi.org/10.1007/978-981-19-1044-9_2)

---

© 2022 The Author(s), under exclusive license to Springer Nature Singapore Pte Ltd. This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use (<https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at [https://doi.org/10.1007/978-981-19-1044-9\\_2](https://doi.org/10.1007/978-981-19-1044-9_2).

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Mixture models for spherical data with applications to protein bioinformatics

Kanti V. Mardia, Stuart Barber, Philippa M. Burdett, John T. Kent and Thomas Hamelryck

**Abstract** Finite mixture models are fitted to spherical data. Kent distributions are used for the components of the mixture because they allow considerable flexibility. Previous work on such mixtures has used an approximate maximum likelihood estimator for the parameters of a single component. However, the approximation causes problems when using the EM algorithm to estimate the parameters in a mixture model. Hence the exact maximum likelihood estimator is used here for the individual components.

This paper is motivated by a challenging prize problem in structural bioinformatics of how proteins fold. It is known that hydrogen bonds play a key role in the folding of a protein. We explore this hydrogen bond geometry using a data set describing bonds between two amino acids in proteins. An appropriate coordinate system to represent the hydrogen bond geometry is proposed, with each bond represented as a point on a sphere. We fit mixtures of Kent distributions to different subsets of the hydrogen bond data to gain insight into how the secondary structure elements bond together, since the distribution of hydrogen bonds depends on which secondary structure elements are involved.

---

Kanti V. Mardia

School of Mathematics, University of Leeds, LS2 9JT, UK, e-mail: K.V.Mardia@leeds.ac.uk;  
Department of Statistics, University of Oxford, OX1 3LB, UK, e-mail: mardia@stats.ox.ac.uk

Stuart Barber

School of Mathematics, University of Leeds, LS2 9JT, UK, e-mail: S.Barber@leeds.ac.uk

Philippa M. Burdett

School of Mathematics, University of Leeds, LS2 9JT, UK.

John T. Kent

School of Mathematics, University of Leeds, LS2 9JT, UK, e-mail: J.T.Kent@leeds.ac.uk

Thomas Hamelryck

Department of Biology/Department of Computer Science, University of Copenhagen, Denmark,  
e-mail: thamelry@binf.ku.dk

## 1 Introduction

Spherical data arise in many applications, such as palaeomagnetism (Hospers, 1955) optical crystallography (Vistelius, 1966, p.92) and stellar directions (Jupp, 1995); more examples are given in Fisher et al. (1987) and Mardia and Jupp (2000). Pewsey and García-Portugués (2020) have given a recent review with many other references. In simple cases, these can be described as a sample from an appropriate spherical distribution such as the Kent distribution (Kent, 1982). However, many data sets are not so straightforward and are better described as a mixture of components, such as the rock fracture data considered by Peel et al. (2001). Often, we must estimate which observations come from each component and a natural way to fit such a model is to use the EM algorithm.

Kent (1982) developed a bivariate normal approximation for the density of the Kent distribution. This corresponding approximate likelihood was used to construct a simple approximate maximum likelihood estimate (MLE) of the parameters in place of the more computationally intensive exact maximum likelihood estimator. Peel et al. (2001) used the approximate likelihood in an EM algorithm to fit a mixture of Kent distributions. Unfortunately, with this approximation the resulting EM algorithm is not necessarily monotone (the EM algorithm must be monotone when the exact likelihood is used). The main problem is the use of the bivariate normal approximation for the normalizing constant instead of the exact value given below in (5). In this paper the exact MLEs have been used to avoid the problem.

Our motivating example is the distribution of hydrogen bonds between secondary structures in proteins. Our data consist of four distances, three angles, and two categorical variables. Driven by these measurements and previous work by Paulsen (2009) a new ‘Euclidean-Latitude’ representation of the hydrogen bond is introduced, which reduces the original 12 variables to four directional variables, incorporating three spherical coordinates and a colatitude. We focus on two of the angles which capture the most important biological information about the hydrogen bond.


The joint distribution of these angles forms a ‘shell’ on the surface of the sphere, with the depth of the shell determined by the length of the hydrogen bond. This distribution is dependent on the secondary structures involved and the separation distance between the interacting pair. Paulsen et al. (2010) used bond lengths and PCA to develop a probabilistic model for the bonds. We propose a flexible model for the distribution of hydrogen bonds by conditioning on the hydrogen bond length and using the EM algorithm to fit a mixture distribution of an unknown number of Kent distribution components using the exact MLEs of the parameters.

### 1.1 Proteins and hydrogen bonding

For over forty years it has been known that a globular protein spontaneously folds into a compact structure determined by its amino acid sequence, but how the structure depends on the sequence is still unknown. A valuable tool in structure prediction

is the ability to simulate plausible candidate structures; examples are the methods described by Hamelryck et al. (2006) and Zhao et al. (2010). It is known that hydrogen bonds play a key role in folding a protein into its 3-D structure (Baker and Hubbard, 1984). Each hydrogen bond can be defined by four atomic coordinates in 3-D space in a representation equivalent to that of Kortemme et al. (2003). We use an alternative parameterisation based on directional statistics and focus on a probabilistic approach to modelling hydrogen bonds. Such a model has potential applications in protein structure prediction, simulation, validation and design. For example, the model could be used to evaluate the potential strength of a hydrogen bond. By evaluating the probability of the hydrogen bond network in a protein, potential errors or low quality protein structures could be identified.

Proteins are polypeptide chains of amino acids joined end to end by peptide bonds; see for example Branden and Tooze (1999) and Mardia (2013). The sequence of these amino acids is known as the primary structure of the protein. There are twenty amino acids which all have a common structure consisting of a central carbon atom, denoted  $C_\alpha$ , a hydrogen atom, an amino group and a carboxyl group. A side chain, unique to each of the 20 amino acid types, projects from the central carbon atom. As the polypeptide chain folds it forms secondary structures which are classified into three types: helices, sheets and loops.



**Fig. 1** The three angles  $\chi$ ,  $\theta$ ,  $\psi$  and bond length  $\delta$  associated with hydrogen bonds. Note that the bonds C-O and H-N are assumed to be of constant length.

The geometry of a hydrogen bond is shown in Figure 1. A naive representation would be to use 12 variables (Euclidean coordinates of the four atoms) to represent one bond. However, the spatial orientation and location of the bond are not considered important, so we effectively lose two degrees of freedom for the fixed bond lengths C-O and H-N, three for rotation and a further three for translation. Thus we retain full information about the geometry of the hydrogen bond by using only four variables: the C-O-H angle  $\psi$ , the dihedral angle  $N_i, C_i, O_i, H_j$ ,  $\chi$  (where  $i$  and  $j$  are the residue numbers of the hydrogen bond acceptor and donor, respectively), the O-H distance  $\delta$  and the N-H-O angle  $\theta$ . We focus on a subset of the angles,  $\chi$  and  $\psi$ . These angles define a vector  $\mathbf{x} = [x_1, x_2, x_3]^T$  on the surface of the unit sphere using

$$x_1 = \sin \psi \sin \chi, \quad x_2 = \sin \psi \cos \chi, \quad x_3 = \cos \psi, \quad |\mathbf{x}| = 1, \quad (1)$$

so that  $0 \leq \psi \leq \pi$  and  $0 \leq \chi < 2\pi$  represent the colatitude and longitude of  $\mathbf{x}$ . Section 5.1 gives an analysis of the data using the angles  $\chi$  and  $\psi$ , together with the distance  $\delta$ , and a separation distance  $\Delta_L$  defined there.

## 2 Fitting a single Kent distribution

### 2.1 The Kent distribution

The Kent (or Fisher-Bingham 5 Parameter,  $FB_5$ ) distribution (Kent, 1982) is the key distribution in the analysis of directional data on  $S^2$ , the unit sphere in  $\mathbb{R}^3$ . It is the spherical analogue of the general multivariate normal distribution allowing for distributions of any elliptical shape, size or orientation on the surface of the sphere. We write  $\text{Kent}(\kappa, \beta, \Gamma)$  to denote the Kent distribution with density

$$f(\mathbf{x}) = \frac{1}{c(\kappa, \beta)} \exp \left\{ \kappa \boldsymbol{\gamma}_{(3)}^T \mathbf{x} + \beta [(\boldsymbol{\gamma}_{(1)}^T \mathbf{x})^2 - (\boldsymbol{\gamma}_{(2)}^T \mathbf{x})^2] \right\}, \quad |\mathbf{x}| = 1, \quad (2)$$

where  $\Gamma = [\boldsymbol{\gamma}_{(1)} \ \boldsymbol{\gamma}_{(2)} \ \boldsymbol{\gamma}_{(3)}]$  is a  $3 \times 3$  orientation matrix and  $c(\kappa, \beta)$  is a normalizing constant. Here,  $\kappa$  determines the concentration of the data,  $\beta$  describes the ovalness of the contours of constant probability,  $\boldsymbol{\gamma}_{(3)}$  is the population mean vector or pole of the distribution, and  $\boldsymbol{\gamma}_{(1)}, \boldsymbol{\gamma}_{(2)}$  denote the major and minor axes of the ellipse-like contours of constant probability.

An alternative angular parameterisation is obtained using Euler angles by noting that  $\Gamma$  can be written as a product of three simple rotations,

$$\Gamma = R_\phi R_\eta R_\omega,$$

where

$$\mathbf{R}_\phi = \begin{bmatrix} \cos \phi & -\sin \phi & 0 \\ \sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{R}_\eta = \begin{bmatrix} \cos \eta & 0 & \sin \eta \\ 0 & 1 & 0 \\ -\sin \eta & 0 & \cos \eta \end{bmatrix}, \quad \mathbf{R}_\omega = \begin{bmatrix} \cos \omega & -\sin \omega & 0 \\ \sin \omega & \cos \omega & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (3)$$

where  $0 \leq \phi, \omega < 2\pi$ ,  $0 \leq \eta \leq \pi$ . The first and third matrices represent rotations about the third coordinate axis (the north pole) and the second matrix represents a rotation about the second coordinate axis (i.e. along the prime meridian). In particular, the final column of  $\Gamma$  has elements

$$\gamma_{13} = \sin \eta \sin \phi, \quad \gamma_{23} = \sin \eta \cos \phi, \quad \gamma_{33} = \cos \eta, \quad (4)$$

so that  $\eta$  and  $\phi$  represent the colatitude and longitude of the mean direction.

Evaluation of the normalizing constant  $c(\kappa, \beta)$  is not trivial, but can be computed in a numerically stable way through a series expansion. From Kent (1982),

$$c(\kappa, \beta) = 2\pi \sum_{j=0}^{\infty} \frac{\Gamma(j + \frac{1}{2})}{\Gamma(j + 1)} \beta^{2j} (\frac{1}{2}\kappa)^{-2j - \frac{1}{2}} I_{2j + \frac{1}{2}}(\kappa); \quad (5)$$

to calculate the normalizing constant we need to evaluate the modified Bessel function  $I_{2j + \frac{1}{2}}(\kappa)$ . This can be calculated using a backwards recursion which is more stable than forward recursion. From Amos (1974),

$$\begin{aligned}
I_{\nu-1}(x) &= \frac{2\nu}{x} I_{\nu}(x) + I_{\nu+1}(x), \\
r_{\nu-1}(x) &= \frac{1}{2\nu/x + r_{\nu}(x)}, \quad r_{\nu}(x) = \frac{I_{\nu+1}(x)}{I_{\nu}(x)}, \\
I_{\nu}(x) &= I_{\alpha}(x) \prod_{k=1}^{\lfloor \nu \rfloor} r_{\nu-k}(x), \quad \alpha = \nu - \lfloor \nu \rfloor,
\end{aligned}$$

where  $\lfloor \nu \rfloor$  is the integer part of  $\nu$ . Therefore,

$$\begin{aligned}
I_{2j+\frac{1}{2}}(\kappa) &= I_{\frac{1}{2}}(\kappa) \prod_{k=1}^{2j} r_{2j+\frac{1}{2}-k}(x), \\
\text{where } I_{\frac{1}{2}}(\kappa) &= \sqrt{\frac{2}{\pi\kappa}} \sinh \kappa.
\end{aligned}$$

The truncation point of the series,  $m$ , can be altered for the desired level of accuracy. In practice, we have found  $m = 20$  to be sufficient.

In passing, note that various approximations to the normalizing constant have been developed over the years that improve over the bivariate normal approximation in (6), notably the saddlepoint approximation. See, e.g., Kume and Wood (2005) and Scaaly and Welsh (2014). However, the exact series used here (5) has the advantage that it can be computed to arbitrary accuracy.

## 2.2 High concentration bivariate normal approximation

Under high concentration, a bivariate normal approximation to the Kent distribution was given by Kent (1982). Let  $\mathbf{x} \sim \text{Kent}(\kappa, \beta, \mathbf{\Gamma})$ , and rotate to  $\mathbf{y} = \mathbf{\Gamma}^T \mathbf{x}$  so that  $\mathbf{y} \sim \text{Kent}(\kappa, \beta, \mathbf{I})$ , where  $\mathbf{I}$  denotes the identity matrix. The limiting distribution as  $\kappa \rightarrow \infty$  and  $\frac{\beta}{\kappa} \rightarrow d$ , where  $0 \leq d < \frac{1}{2}$ , can be described using the bivariate normal distribution for the standardized variables  $\mathbf{y}_1^* = \sqrt{\kappa} y_1$  and  $\mathbf{y}_2^* = \sqrt{\kappa} y_2$ . In particular

$$\mathbf{y}_1^* \sim N(0, (1 - 2d)^{-1}), \quad \mathbf{y}_2^* \sim N(0, (1 + 2d)^{-1}), \quad (6)$$

independently. The condition number of the corresponding covariance matrix of the bivariate normal distribution is defined by

$$\text{CN} = \frac{\kappa + 2\beta}{\kappa - 2\beta}, \quad (7)$$

and measures the extent of any anisotropy, where  $d = 0$  (i.e.  $\beta = 0$ ) for the isotropic case.

### 2.3 Maximum likelihood estimators for the Kent distribution

Consider  $n$  independent observations  $\mathbf{x}_1, \dots, \mathbf{x}_n$  from a Kent distribution  $\text{Kent}(\kappa, \beta, \Gamma)$ . The log-likelihood is given by

$$\begin{aligned} \ell(\Psi; \mathbf{x}) &= \sum_{j=1}^n \left\{ \kappa \boldsymbol{\gamma}_{(3)}^T \mathbf{x}_j + \beta [(\boldsymbol{\gamma}_{(1)}^T \mathbf{x}_j)^2 - (\boldsymbol{\gamma}_{(2)}^T \mathbf{x}_j)^2] \right\} - n \log c(\kappa, \beta) \\ &= n \kappa \boldsymbol{\gamma}_{(3)}^T \bar{\mathbf{x}} + n \beta \left( \boldsymbol{\gamma}_{(1)}^T \mathbf{S} \boldsymbol{\gamma}_{(1)} - \boldsymbol{\gamma}_{(2)}^T \mathbf{S} \boldsymbol{\gamma}_{(2)} \right) - n \log c(\kappa, \beta). \end{aligned} \quad (8)$$

Here  $\Psi = (\kappa, \beta, \eta, \phi, \omega)$  contains the parameters. The log-likelihood depends on the data through the sample mean vector and the sample second moment matrix

$$\bar{\mathbf{x}} = \frac{1}{n} \sum \mathbf{x}_j, \quad \mathbf{S} = \frac{1}{n} \sum \mathbf{x}_j \mathbf{x}_j^T.$$

The approximate MLE of Kent (1982) can be computed as follows.

- (a) The estimated orientation matrix  $\hat{\Gamma}$  is determined so that  $\hat{\Gamma}^T \bar{\mathbf{x}} \propto [0 \ 0 \ 1]^T$  and  $(\hat{\Gamma}^T \mathbf{S} \hat{\Gamma})_{12} = 0$ , where the subscript indicates element (1, 2) of a  $3 \times 3$  matrix. This estimate depends just on the first two sample moments and does not require any estimates of  $\kappa$  and  $\beta$ .
- (b) The normalizing constant is approximated using the bivariate normal approximation in (6). With this substitution the estimation of  $\kappa$  and  $\beta$  is straightforward.

Mixture models are discussed below in Section 3. As noted earlier, the use of the approximate MLE in mixture models can cause problems for the EM algorithm, because the approximate log-likelihood can sometimes fail to increase at each iteration. Hence for this paper we limit attention to the exact MLE for the single-component Kent distribution. The exact MLE needs to be found numerically, e.g. by using the function `optim` in `R`, and for this purpose an analytic form for the derivatives of  $\ell$  is very helpful. For the concentration parameters

$$\begin{aligned} \partial \ell / \partial \kappa &= n \boldsymbol{\gamma}_{(3)}^T \bar{\mathbf{x}} - n \partial \log c(\kappa, \beta) / \partial \kappa, \\ \partial \ell / \partial \beta &= n \left( \boldsymbol{\gamma}_{(1)}^T \mathbf{S} \boldsymbol{\gamma}_{(1)} - \boldsymbol{\gamma}_{(2)}^T \mathbf{S} \boldsymbol{\gamma}_{(2)} \right) - n \partial \log c(\kappa, \beta) / \partial \beta, \end{aligned}$$

where the derivatives of  $c(\kappa, \beta)$  are obtained by differentiating the series (5) termwise. For the orientation parameters, the derivatives of  $\Gamma$  are obtained by differentiating each of its building blocks. For example,

$$\partial \Gamma / \partial \phi = R'_\phi R_\eta R_\omega, \quad R'_\phi = \begin{bmatrix} -\sin \phi & -\cos \phi & 0 \\ \cos \phi & -\sin \phi & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

and similarly for  $\partial \Gamma / \partial \eta$  and  $\partial \Gamma / \partial \omega$ .

### 3 Fitting Kent mixtures using the EM algorithm

#### 3.1 The mixture distribution

Consider a mixture of  $g$  Kent components with density

$$f(\mathbf{x}; \Psi) = \sum_{i=1}^g p_i f(\mathbf{x}; \Psi_i) \quad (9)$$

Here  $\mathbf{p} = (p_1, \dots, p_g)^T$ , where  $p_i > 0$ ,  $i = 1, \dots, g$  and  $\sum p_i = 1$  is a vector of *mixing proportions*. Also,  $f(\mathbf{x}; \Psi_i)$  is the Kent density for the  $i$ th component,  $i = 1, \dots, g$  with parameter vector  $\Psi_i$ , as defined below (8). In many cases it is useful to let one of the components denote the uniform distribution on the sphere with density  $f_0(\mathbf{x}) = 1/(4\pi)$ . Of course for the uniform component there are no parameters to estimate. This is the same model used by Peel et al. (2001).

Let  $\mathbf{Z} = (z_1, \dots, z_n)^T$  be an  $n \times g$  matrix membership matrix where  $z_{ji} = 1$  if observation  $j$  belongs to the  $i$ th component and 0 otherwise. In general the membership matrix is unobserved and represents a set of missing or latent variables. However, if the membership matrix is available, then a complete-data vector  $\mathbf{w}_j$  for observation  $j$  can be defined by  $\mathbf{w}_j = (\mathbf{x}_j^T, \mathbf{z}_j^T)^T$ . The complete-data log-likelihood function for  $\Psi$  and  $\mathbf{p}$  is given by

$$\ell_c(\Psi, \mathbf{p}) = \sum_{i=1}^g \sum_{j=1}^n z_{ji} \{\log p_i + \log f(\mathbf{x}_j; \Psi_i)\}. \quad (10)$$

When the membership matrix is not observed, then it is possible to compute the probability that each observation belongs to each group. Given parameter values  $\Psi_i$ ,  $i = 1, \dots, g$  and group probabilities  $p_i$ ,  $i = 1, \dots, g$ , the membership probabilities for observation  $j$  are given by

$$\tau_{ji} = P(Z_{ji} = 1 | \mathbf{x}_j) = \frac{p_i f(\mathbf{x}_j; \Psi_i)}{\sum_{h=1}^g p_h f(\mathbf{x}_j; \Psi_h)}. \quad i = 1, \dots, g. \quad (11)$$

#### 3.2 EM algorithm

The EM algorithm (Dempster et al., 1977) can be used to iteratively compute the exact maximum likelihood estimator for a mixture of  $g$  Kent components. It consists of alternating expectation (E) steps and maximisation (M) steps. It requires a knowledge of the number of groups  $g$  and initial estimates for the parameters  $\Psi_i^{(0)}$ ,  $p_i^{(0)}$ ,  $i = 1, \dots, g$ . Let  $\nu \geq 0$  index the current iteration.



- (a)(E-step). Given current values for the parameters  $\Psi_i^{(v)}$ ,  $p_i^{(v)}$ ,  $i = 1, \dots, g$ , estimate the membership probabilities using (11) to get estimates  $\tau_{ji}^{(v)}$ , say.
- (b)(M-step). Given the the estimated membership probabilities  $\tau_{ji}^{(v)}$ , compute weighted sample moments

$$\bar{\mathbf{x}}_i^{(v)} = \frac{\sum_{j=1}^n \tau_{ji}^{(v)} \mathbf{x}_j}{\sum_{j=1}^n \tau_{ji}^{(v)}}, \quad \mathbf{S}_i^{(v)} = \frac{\sum_{j=1}^n \tau_{ji}^{(v)} \mathbf{x}_j \mathbf{x}_j^T}{\sum_{j=1}^n \tau_{ji}^{(v)}}. \quad (12)$$

Using these weighted moments in (8), update the maximum likelihood estimates for the parameters  $\Psi_i^{(v+1)}$ ,  $i = 1, \dots, g$ . In addition update the estimates of the mixture proportions to

$$p_i^{(v+1)} = \sum_{j=1}^n \tau_{ji}^{(v)} / n, \quad i = 1, \dots, g. \quad (13)$$

Then cycle between these two steps until convergence. For more details on the EM algorithm, see e.g. Little and Rubin (2002) and McLachlan and Krishnan (2008).

To select the number  $g$  of components in the model, we choose that value which minimises the AIC,

$$2k^* - 2\hat{\ell}_g.$$

Here  $g$  is the number of components in the model, and  $\hat{\ell}_g$  denotes the maximized log-likelihood for the  $g$ -component mixture. If all the components are Kent distributions, then the total number of parameters is  $k^* = 5g + g - 1 = 6g - 1$  (5 parameters for each Kent distribution and  $g - 1$  parameters for the mixing probabilities). If one of the components is a uniform distribution, then the total number of parameters is  $k^* = 4g + g - 1 = 5g - 1$ .

## 4 Simulation study

To assess the performance of our EM algorithm for fitting the Kent mixture model given by (9), we consider several simulation experiments. In section 4.1 we consider how classification rate improves as components move further apart on the surface of the sphere, assuming the number of components is known. In section 4.2, we use a range of conditions to examine the accuracy of the method when the number of components must be estimated by minimizing the AIC.

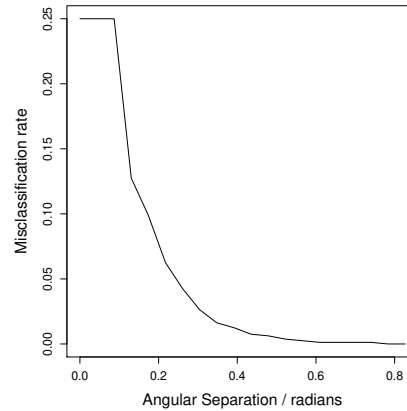
#### 4.1 Angular separation of components

Consider an observation  $j$ . A set of classification probabilities  $\hat{\tau}_{ji}$ ,  $i = 1, \dots, g$  can be *hardened* by assigning observation  $j$  to the group  $i$  with the largest classification probability. In any simulation experiment, the labels of all the observations are known. Hence for each true group  $i$ , the number misclassified observations can be counted. The results can be summarized either as a misclassification rate for each group  $i$  or as an overall misclassification rate.

The misclassification rate gives a useful indication of how distinct the groups are from one another. Clearly the angular separation of the components will affect the misclassification rate, where the angular separation between two unit vectors  $\mathbf{l}_1$  and  $\mathbf{l}_2$  is measured by  $\text{acos}(\mathbf{l}_1^T \mathbf{l}_2)$ .

To explore the misclassification rate, four Kent components each of size  $n = 200$ , and all with  $\kappa = 50$  and  $\beta = 10$  were generated. One component was located at the north pole, one at the south pole and two on the equator. The two equatorial components ( $E_1$  and  $E_2$ ) initially had the same mean direction. The two polar components, and  $E_1$  were kept stationary while  $E_2$  was rotated around the equator so the angular separation ranged from 0 to  $\pi$ . A 4-component Kent model (without a uniform component) was fitted to the data using the known values for the mean direction,  $\kappa$ ,  $\beta$ , and  $p_i$  as starting points.

The misclassification rate as the angular separation between  $E_1$  and  $E_2$  increased is shown in Figure 2. As expected, the misclassification decreases as the angular separation increases and the overlap between  $E_1$  and  $E_2$  decreases.



**Fig. 2** Misclassification rate against angular separation  $\text{acos}(E_1^T E_2)$ .

**Table 1** Case 1. True simulation parameters (top panel) and estimates of these parameters for the model with the lowest AIC (lower panel).

	Component				
	uniform	1	2	3	4
$\kappa$	-	5	10	20	20
$\beta$	-	2	4	9	2
$p_i$	0	0.25	0.25	0.25	0.25
$\gamma_{(3)}^T$	-	[0, 0, 1]	[0, 1, 0]	[0, -1, 0]	[0, 0, -1]
$\widehat{\kappa}$	-	6.94	10.70	18.22	22.16
$\widehat{\beta}$	-	3.14	5.00	8.55	2.82
$\widehat{p}_i$	0.06	0.21	0.24	0.24	0.24
$\widehat{\gamma}_{(3)}^T$	-	[0.0, 0.0, 1.0]	[0.0, 1.0, 0.0]	[0.0, -1.0, 0.0]	[0.0, 0.0, -1.0]

## 4.2 Model selection

We shall illustrate our method with data simulated under three conditions. In Case 1, four samples of equal size,  $n = 200$ , were generated from the Kent distribution with different values for  $\kappa, \beta$  and the mean vector  $\gamma_{(3)}$ . Case 2 looks at five samples of differing sizes and values for the parameters. In Case 3 the Case 2 data set is re-used with 100 spherical uniform random variates added. In all cases a uniform component is included in the mixture model fitted to the data.

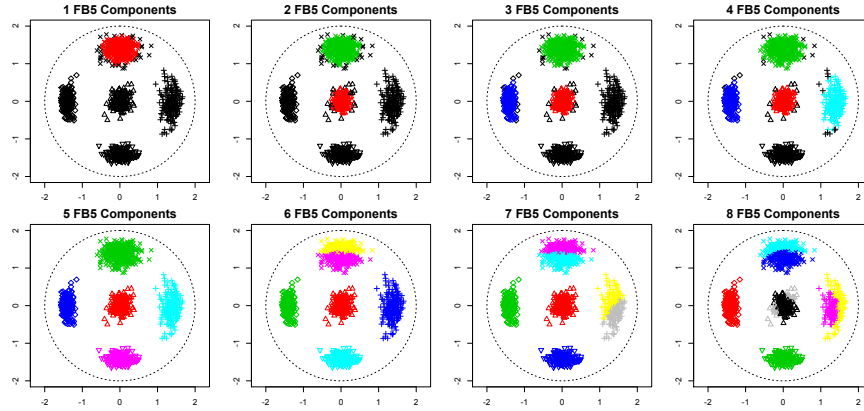
## 4.3 Case 1: Four equi-sized components

In this case samples of size  $n = 200$  were simulated from  $g = 4$  Kent distributions with parameters given in Table 1. Mixture models defined by (9) with 1, 2, ..., 7 Kent components, plus a uniform component, were fitted to the data set and the AIC for each model was computed.

The fitted model with the smallest AIC has the correct choice of four Kent components (plus a uniform component) and provides MLEs that are very close to the true parameters. In addition the misclassification rate is 2.5%, with eleven observations allocated to the uniform component. In this case, the mean directions of the components are well-separated.

#### 4.4 Case 2: Five unequally sized components

In this case  $n = 200$  samples were simulated from  $g = 5$  Kent distributions with parameters given in Table 2. Mixture models with  $1, \dots, 8$  Kent components, plus a uniform component, were fitted to the data set, and the AIC for each model was computed. The model with the smallest AIC has the correct choice of five Kent components. The parameter estimates for this model are shown in Table 2



**Fig. 3** Schmidt net projections of the data coloured according to each model. The plotting symbols indicate the known membership of the data. The top left panel shows the data with only one Kent component fitted, plus a uniform component. In each subsequent panel an additional Kent component has been added to the model. The fitted model with five Kent components (the correct number) is shown as the first plot in the bottom row. Points allocated to the uniform component are in black in each panel.

Figure 3 shows the Schmidt-net projections of the data coloured by their allocations for each of the models fitted. The plotting symbols show the known allocations. For Schmidt-net projection, see, for example, Mardia and Jupp (2000, p.161); this is an equal area projection of the unit sphere in  $\mathbb{R}^3$  onto a disc of radius two where

$$(y_1, y_2) = 2 \sin(\psi/2)(\cos \chi, \sin \chi),$$

where the polar coordinates  $(\psi, \chi)$  are defined in (1). From these plots it can also be seen that the model with five Kent components gives the best fit with no error, i.e. a misclassification rate of 0%. In particular, although a uniform component has been included in the fitted model, no points have been allocated to it.

With fewer than five components in the model, some of the true components are entirely allocated to the uniform component, and for those components that are fitted correctly, a few observations have also been allocated to be uniform — in these cases the fitted  $\kappa$  values are overly large making the fitted model more concentrated. For the model with more than five Kent components the clusters have been split into two

or more components. Again we see that the overly complex model results in some observations being erroneously assigned to the uniform component.

**Table 2** Case 2: True simulation parameters (top panel) and estimates (rounded to 2 d.p. to highlight the effect on  $\gamma_{(3)}$ ) for the selected model with the lowest AIC (lower panel).

	Component				
	1	2	3	4	5
$\kappa$	50	25	30	70	50
$\beta$	2	2	5	10	2
$p_i$	0.15	0.15	0.37	0.19	0.15
$\gamma_{(3)}$	$\begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$	$\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}$	$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$	$\begin{bmatrix} 0 \\ -1 \\ 0 \end{bmatrix}$	$\begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}$
$\hat{\kappa}$	50.07	24.97	29.54	68.22	52.31
$\hat{\beta}$	5.76	2.71	5.93	11.08	4.55
$\hat{p}_i$	0.15	0.15	0.37	0.19	0.15
$\hat{\gamma}_{(3)}$	$\begin{bmatrix} -0.02 \\ 0.01 \\ 1.00 \end{bmatrix}$	$\begin{bmatrix} 0.00 \\ 1.00 \\ 0.04 \end{bmatrix}$	$\begin{bmatrix} 1.00 \\ -0.01 \\ 0.00 \end{bmatrix}$	$\begin{bmatrix} 0.00 \\ -1.00 \\ 0.00 \end{bmatrix}$	$\begin{bmatrix} -1.00 \\ -0.02 \\ 0.00 \end{bmatrix}$

#### 4.5 Case 3: Case 2 with a uniform component added

To the previous data set (Case 2) of five Kent components, we added 100 random uniform variates. The algorithm was then applied with the same starting points. In this case, the AIC picks out the correct model with five Kent components (plus a uniform component), with a misclassification rate of 2.6%. Table 4 shows that over 75% of the misclassified points were from the uniform component and of these all can be shown to be close to a Kent component. Hence we conclude that our fitting procedure for the Kent mixtures performs well for components with little overlap.

## 5 Hydrogen bond data

### 5.1 Exploratory data analysis

We consider a data set of 11653 potential hydrogen bonds. It is important to note that the hydrogen bonds in this data set are not observed, but that the presence of a

**Table 3** Case 3. Estimates of the parameters (rounded to 2 d.p. to show the effect on  $\gamma_{(3)}$ ) when a uniform component is included. The true parameters are shown in Table 2.

	uniform	Component				
		1	2	3	4	5
$\widehat{\kappa}$	–	56.01	24.67	29.79	67.48	51.23
$\widehat{\beta}$	–	5.53	2.62	6.21	10.17	4.18
$p_i$	0.07	0.14	0.14	0.34	0.17	0.14
$\widehat{p}_i$	0.08	0.13	0.14	0.35	0.17	0.14
$\widehat{\gamma}_{(3)}$	–	$\begin{bmatrix} -0.02 \\ 0.01 \\ 1.00 \end{bmatrix}$	$\begin{bmatrix} -0.01 \\ 1.00 \\ 0.03 \end{bmatrix}$	$\begin{bmatrix} 1.00 \\ -0.01 \\ 0.00 \end{bmatrix}$	$\begin{bmatrix} 0.00 \\ -1.00 \\ 0.00 \end{bmatrix}$	$\begin{bmatrix} -1.00 \\ -0.02 \\ -0.01 \end{bmatrix}$

**Table 4** Case 3: The classifications of the data points from the fitted model compared to the known classifications.


Classification	True component					
	Uniform	Kent 1	Kent 2	Kent 3	Kent 4	Kent 5
Uniform	71	4	1	2	1	1
Kent 1	1	196	0	0	0	0
Kent 2	5	0	199	0	0	0
Kent 3	12	0	0	498	0	0
Kent 4	6	0	0	0	249	0
Kent 5	5	0	0	0	0	199

hydrogen bond is inferred between two non-adjacent residues within the protein chain where the oxygen-nitrogen bond has length less than  $3.5\text{\AA}$ . If several connections satisfied these criteria, then the one with the smallest oxygen-hydrogen distance,  $\delta$ , was selected, see Figure 1. We therefore assume that some of the alleged hydrogen bonds in the data are spurious, and could bias parameter estimates.

The data on each potential hydrogen bond consists of angles  $\chi$  and  $\psi$ , as defined in (1), separation distance ( $\Delta_L$ ), and secondary structure combination. Thus, our data are in the coordinate system  $(x_1, x_2, x_3)$  defined by  $(\chi, \psi)$  in (1). The variable  $\Delta_L$  is given by the number of amino acid residues separating the interacting pair where  $\Delta_L \in \{2, 3, 4, 5+\}$  as hydrogen bonds cannot form between adjacent amino acids.

The data also contain information on the secondary structure of the residues in the chain on the O-side and the N-side of the hydrogen bond (see Branden and Tooze (1999)). There are three types of secondary structure ( $\alpha$ -helix, loop and  $\beta$ -sheet), and hydrogen bonds cannot form between  $\beta$ -sheets and  $\alpha$ -helices, so there are seven possible secondary structure combinations. In this paper, we restrict ourselves to modelling the  $n = 3627$  potential hydrogen bonds formed between  $\alpha$ -helices.

Plots of the data can be made by projecting the data on to the sphere. Using the separation distance  $\Delta_L$  to break down the data further showed that this distance



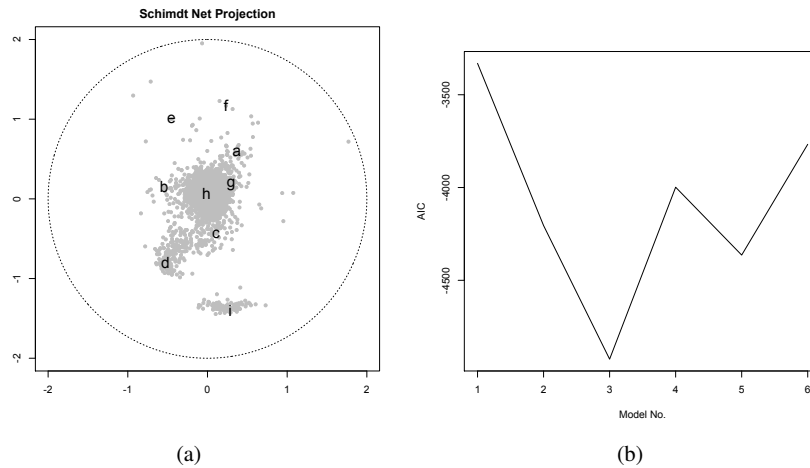
**Fig. 4** Plots showing the helix-helix data separated by  $\Delta_L$ , all shown from the same point of view. Panels (a)–(c) show data for  $\Delta_L = 2, 3,$  and  $4$  respectively.

affects the distribution of the hydrogen bonds. An example is shown in Figure 4 of the helix-helix data, with most of the data having  $\Delta_L = 4$ . From a biological point of view, it is to be expected that nearly all hydrogen bonds between helical amino acids have  $\Delta_L = 4$ , since this type of hydrogen bond provides the stability of the helix. The remaining hydrogen bonds can be regarded as outliers or examples of unusual hydrogen bond geometries as found in the relatively rare  $3_{10}$ - and  $\pi$ -helices".

## 5.2 Modelling the helix-helix data

The helix-helix data shows one large cluster containing most of the data, and several diffuse areas. Choosing starting points for the algorithm by conditioning on  $\Delta_L$  works for the larger clusters but cannot separate out the side clusters effectively. Also this method does not allow for choosing further starting points. Therefore the starting points for the mean directions were selected by plotting a Schmidt-net projection of the complete helix-helix data (Figure 5(a)) and identifying individual points by eye as possible choices. Initial values for  $\kappa, \beta$  and the mixing proportions,  $p_i$ , were also chosen by visual inspection.

Choosing the number  $g$  of Kent components in our mixture model is complicated by the fact that the fitted model found by the EM algorithm depends heavily on the initial selections for the mean directions. We conduct model selection in a stepwise manner. In the first step a mixture model of one Kent and one uniform component is fitted to the data using each starting point in turn and the one with the lowest AIC chosen as the preferred one-component model. At each subsequent step,  $g = 2, \dots, 9$ , the model with the lowest AIC in the previous step is used as the basis for the next model and then each remaining starting point added in turn to give a candidate model with  $g$  Kent components and one uniform component. The  $g$ -component model with the lowest AIC is retained. This process is repeated until the nine-component model using all starting points is fitted to the data. The final ‘optimal’ model is chosen by comparing the ‘best’ models from each step and choosing the one with the



**Fig. 5** Plots showing (a) Schmidt-net projection of the helix-helix data with eight different starting points  $a, \dots, i$  and (b) the AIC for each model fitted. The AIC is minimised by the model with three Kent components.

lowest AIC. Overall the AIC is minimized by a model consisting of the three Kent components labelled  $h$ ,  $d$ , and  $i$  in Figure 5(a). Component  $h$  was fitted first,  $d$  second and  $i$  third. The parameter estimates for this model are shown in the Table 5. Figure 6 shows the allocation of data points to model components achieved when this model is fitted to the data. The data are plotted with circles, squares, and triangles for components  $d$ ,  $h$ , and  $i$  respectively, while observations allocated to the uniform component are plotted with small points.

In this case, the mixture model has recovered the physically most plausible hydrogen bonds (those with  $\Delta_L = 4$ ) as mainly being in component  $h$ , despite  $\Delta_L$  not being used in the modelling process. Components  $d$  and  $i$  are mainly observations with  $\Delta_L$  being 3 and 2, respectively. This suggests that the mixture modelling has effectively separated the true hydrogen bonds from atypical and/or spurious ones in the data set. The adjusted Rand index (Hubert and Arabie, 1985) comparing the true  $\Delta_L$  and allocated component membership is 0.879, indicating a close agreement between the two classifications. Note that the condition numbers given by (7) for  $d$  and  $i$  are very high at 8.1 and 6.7 respectively.

## 6 Discussion


The ideas in this paper can be extended in various directions. A comprehensive strategy to fit mixtures of multivariate normal distributions has been developed by Fraley and Raftery (2002), with an implementation in R (Scrucca et al., 2016). One



**Table 5** Parameter estimates (top) and  $\Delta_L$  distances (bottom) for the three components fitted to the helix-helix data ( $h$ ,  $d$ , and  $i$ ).

	uniform	$h$	$d$	$i$
$\hat{\kappa}$	-	63.93	35.46	264.34
$\hat{\beta}$	-	1.74	13.89	98.07
$\hat{p}_i$	0.03	0.87	0.08	0.03
$\hat{\gamma}_1$	-	$\begin{bmatrix} -0.23 \\ 0.97 \\ 0.01 \end{bmatrix}$	$\begin{bmatrix} 0.33 \\ -0.93 \\ 0.14 \end{bmatrix}$	$\begin{bmatrix} 0.93 \\ -0.37 \\ 0.06 \end{bmatrix}$

$\Delta_L$	uniform	$h$	$d$	$i$	Total
2	8	0	0	94	102
3	11	10	244	0	263
4	43	3170	22	0	3235
5+	19	6	0	0	25
Total	81	3186	266	94	3627



**Fig. 6** The helix-helix data plotted according to the classifications given by the 'optimum' model. Points are indicated by  $\circ$ ,  $\square$ , and  $\triangle$  for the components  $d$ ,  $h$ , and  $i$  respectively. Points for the uniform component are indicated by small dots.

feature of this strategy is the specification of constraints on the covariance matrices.

If the covariance matrix of the  $i$ th component is written in terms of its spectral decomposition

$$\Sigma_i = c_i G_i \Lambda_i G_i^T,$$

where  $c_i > 0$  is a scaling constant,  $\Lambda_i$  is a diagonal matrix containing the scaled eigenvalues with  $\det(\Lambda_i) = 1$ , and  $G_i$  is an orthogonal matrix of eigenvectors, then various choices can be made, including the following.

- (a) Equality of scaling. The parameters  $c_i$  can either be constrained to be the same or be allowed to be different between components.
- (b) Equality of shape. The scaled eigenvalue matrices  $\Lambda_i$  can either be constrained to be the same or be allowed to be different between components.
- (c) Equality of orientation. The eigenvector matrices  $G_i$  can either be constrained to be the same or be allowed to be different between components.
- (d) Isotropy. If the scaled eigenvalue matrices  $\Lambda_i$  are constrained to be the same as one another, then they can be further constrained to equal the identity matrix (isotropy) or not (anisotropy).

It would be interesting to develop a similar strategy for directional mixtures, though some care is needed in adapting the above considerations to the directional setting.

- (a) Equality of scaling. This constraint makes partial sense in the directional setting. Multiplying  $\kappa_i$  and  $\beta_i$  by a constant under high concentration corresponds to rescaling the projection of the distribution to a tangent plane, but its meaning is less clear under low concentration when the Kent distribution is close to a uniform distribution.
- (b) Equality of shape. This constraint still makes sense in the directional setting, where it just states that  $\kappa_i = \kappa$  and  $\beta_i = \beta$  do not depend on  $i$ .
- (c) Equality of orientation. This constraint does not extend easily to the directional setting. For the Kent distribution the major and minor axes lie in the tangent space at the mean direction. Since different mean directions have different tangent spaces, it is not possible to compare the major and minor axes for different mean directions. There is a partial exception to this statement if all the mean directions lie reasonably close together, when parallel transport can be used to map tangent spaces to one another.
- (d) Isotropy. This constraint still makes sense in the directional setting, where isotropy just corresponds to  $\beta_i = \beta = 0$ .

Another possibility in the directional setting is to consider mixtures of distributions other than Kent distributions. For example, Bingham distributions could be used as building blocks to model axial (i.e. antipodally symmetric) data.

This paper has focused on the main chain of a protein. However, the side chains of a protein are also very important. See, e.g., Branden and Tooze (1999); Harder et al. (2010); Kim et al. (2016); Mardia (2013); Mardia et al. (2012).

### Acknowledgments:

PMB was funded by a DTG award from the UK Engineering and Physical Sciences Research Council. The authors wish to thank G.J. McLachlan for helpful discussions and the referees for helpful comments. KVM thanks the Leverhulme Trust for an Emeritus Fellowship grant.

### References

- Amos, D. E. (1974). Computation of modified Bessel functions and their ratios, *Mathematics of Computation* **28**: 239–251.
- Baker, E. N. and Hubbard, R. E. (1984). Hydrogen bonding in globular proteins, *Progress in Biophysics and Molecular Biology* **44**: 97 – 179.
- Branden, C. and Tooze, J. (1999). *Introduction to Protein Structure*, second edn, Garland, New York.
- Dempster, A. P., Laird, N. M. and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm, *J. Roy. Statist. Soc., B* **39**: 1–38.
- Fisher, N. I., Lewis, T. and Embleton, B. J. (1987). *Statistical Analysis of Spherical Data*, Cambridge University Press, Cambridge.
- Fraley, C. and Raftery, A. E. (2002). Model-based clustering, discriminant analysis and density estimation, *Journal of the American Statistical Association* **97**: 611–631.
- Hamelryck, T., Kent, J. T. and Krogh, A. (2006). Sampling realistic protein conformations using local structural bias, *PLOS Computational Biology* **2**: e131.
- Harder, T., Boomsma, W., Paluszewski, M., Frellsen, J., Johansson, K. E. and Hamelryck, T. (2010). Beyond rotamers: a generative, probabilistic model of side chains in proteins, *BMC Bioinformatics* **11**: 306.  
**URL:** <https://doi.org/10.1186/1471-2105-11-306>
- Hospers, J. (1955). Rock magnetism and solar wandering, *J. Geology* **63**: 59–74.
- Hubert, L. and Arabie, P. (1985). Comparing partitions, *Journal of Classification* **2**(1): 193–218.
- Jupp, P. E. (1995). Some applications of directional data to astronomy, in E.-M. Titt, T. Kollo and H. Niemi (eds), *New Trends in Probability and Statistics. Vol 3: Multivariate Statistics and Matrices in Statistics*, VSP, Utrecht, pp. 123–133.
- Kent, J. T. (1982). The Fisher-Bingham distribution on the sphere, *J. Roy. Statist. Soc., B* **44**: 71–80.
- Kim, S., SenGupta, A. and Arnold, B. C. (2016). A multivariate circular distribution with applications to the protein structure prediction problem, *Journal of Multivariate Analysis* **146**: 374–382.
- Kortemme, T., Morozov, A. V. and Baker, D. (2003). An orientation-dependent hydrogen bonding potential improves prediction of specificity and structure for proteins and protein–protein complexes, *J. Mol. Biol.* **326**: 1239–1259.

- Kume, A. and Wood, A. T. A. (2005). Saddlepoint approximations for the Bingham and Fisher-Bingham normalising constants, *Biometrika* **92**: 465–476.
- Little, R. J. A. and Rubin, D. B. (2002). *Statistical Analysis with Missing Data*, second edn, John Wiley & Sons, New York.
- Mardia, K. V. (2013). Statistical approaches to three key challenges in protein structural bioinformatics, *J. Roy. Stat. Soc. C* pp. 487–514.
- Mardia, K. V. and Jupp, P. (2000). *Directional Statistics*, Wiley and Sons, New York.
- Mardia, K. V., Kent, J. T., Zhang, Z., Taylor, C. C. and Hamelryck, T. (2012). Mixtures of concentrated multivariate sine distributions with applications to bioinformatics, *Journal of Applied Statistics* **39**: 2475–2492.
- McLachlan, G. J. and Krishnan, T. (2008). *The EM algorithm and Extensions*, second edn, John Wiley & Sons, New York.
- Paulsen, J. (2009). *Probabilistic models of hydrogen bonds in proteins*, Master's thesis, University of Copenhagen.
- Paulsen, J., Paluszewski, M., Mardia, K. V. and Hamelryck, T. (2010). A probabilistic model of hydrogen bond geometry in proteins, in A. Gusnanto, K. V. Mardia, C. J. Fallaize and J. Voss (eds), *High-throughput Sequencing, Proteins and Statistics*, Leeds Univ. Press, Leeds, UK.
- Peel, D., Whiten, W. J. and McLachlan, G. J. (2001). Fitting mixtures of Kent distributions to aid in joint set identification, *J. Amer. Statist. Assoc.* **96**: 56–63.
- Pewsey, A. and García-Portugués, E. (2020). Recent advances in directional statistics, *arXiv preprint arXiv:2005.06889*.
- Scealy, J. and Welsh, A. H. (2014). Fitting Kent models to compositional data with small concentration, *Statistics and Computing* **24**(2): 165–179.
- Scrucca, L., Fop, M., Murphy, T. B. and Raftery, A. E. (2016). mclust 5: clustering, classification and density estimation using Gaussian finite mixture models, *The R Journal* **8**(1): 289–317.
- URL:** <https://doi.org/10.32614/RJ-2016-021>
- Vistelius, A. B. (1966). *Structural Diagrams*, Pergamon, London.
- Zhao, F., Peng, J. and Xu, J. (2010). Fragment-free approach to protein folding using conditional neural fields, *Bioinformatics* **26**: i310–i317.