

Some Fundamental Properties of a Multivariate von Mises Distribution

Kanti V. Mardia and Jochen Voss

17th February 2012

Abstract

In application areas like bioinformatics multivariate distributions on angles are encountered which show significant clustering. One approach to statistical modelling of such situations is to use mixtures of unimodal distributions. In the literature (Mardia et al., 2011), the multivariate von Mises distribution, also known as the multivariate sine distribution, has been suggested for components of such models, but work in the area has been hampered by the fact that no good criteria for the von Mises distribution to be unimodal were available. In this article we study the question about when a multivariate von Mises distribution is unimodal. We give sufficient criteria for this to be the case and show examples of distributions with multiple modes when these criteria are violated. In addition, we propose a method to generate samples from the von Mises distribution in the case of high concentration.

keywords: bioinformatics, directional distributions, mixture models, modes, simulation, sine distribution

1 Introduction

In biochemistry it is well known that the structure of macro-molecules such as proteins, DNA, and RNA can be described in terms of conformational angles. For proteins, these angles could be the dihedral and bond angles describing the conformation of the backbone together with additional angles for the configuration of the side chains (see *e.g.* Branden and Tooze, 1998). Data sets consist of the angles to describe each monomer in a macro-molecule, the number of angles required to give the conformation of a monomer determines the dimensionality of the problem. In non-coding RNA there can be 7 or 8 dihedral angles of importance per amino acid (Frellsen et al., 2009) and, if the side chains angles are included, many angles are required for amino acids in proteins (*e.g.* Harder et al., 2010). The resulting distributions on angles are multivariate, often highly structured, featuring various modes together with regions excluded by steric constraints (*e.g.* Mardia et al., 2011).

One way to approach the statistical modelling of such multimodal, multivariate distributions is to use mixture models with unimodal components. In the Euclidean space \mathbb{R}^p , an obvious choice for the components is to use normal distributions with appropriately chosen covariance matrices. For angular data, as considered in this article, the choice of component distribution is less clear, but a simple analogue of the multivariate normal distribution is the multivariate von Mises distribution (Mardia et al., 2008). This distribution is suggested for mixture modelling in Mardia et al. (2011). In order for a mixture model to be a useful description of a multimodal distribution, it is essential that the component distribution is unimodal. In case of the multivariate von Mises distribution, this constraint excludes some of the parameter

¹Department of Statistics, University of Leeds, Leeds, LS2 9JT, United Kingdom

²corresponding author

range. Previous work has been complicated by the problem that no characterisation of the parameter values corresponding to the unimodal case was available. To solve this problem, this article provides sufficient criteria for the multivariate von Mises distribution to be unimodal and we show examples of distributions with multiple modes (where these criteria are violated). It should be noted that univariate circular distributions are well established (see, for example, Mardia and Jupp, 2000) but understanding of multicircular distributions is still evolving.

The multivariate von Mises distribution, first introduced in Mardia et al. (2008) and also known as the multivariate sine distribution, is denoted by $\text{MVM}(\mu, \kappa, \Lambda)$. It is a distribution on the torus $\mathbb{T}^p = [0, 2\pi)^p$ and is given by the density (w.r.t. the uniform distribution on angles)

$$\varphi(\theta; \mu, \kappa, \Lambda) = \frac{1}{Z(\kappa, \Lambda)} \exp(\kappa^\top c(\theta) + \frac{1}{2} s(\theta)^\top \Lambda s(\theta)) \quad (1)$$

for all $\theta \in \mathbb{T}^p$. Here $Z(\kappa, \Lambda)$ is the normalisation constant and we use the abbreviations

$$c_i(\theta) = \cos(\theta_i - \mu_i), \quad s_i(\theta) = \sin(\theta_i - \mu_i)$$

for $i = 1, \dots, p$. The parameters of the distribution are the “mean” $\mu \in \mathbb{T}^p$, the “concentration parameter” $\kappa \in \mathbb{R}^p$ with $\kappa_i \geq 0$ for $i = 1, \dots, p$ and $\Lambda = (\lambda_{ij}) \in \mathbb{R}^{p \times p}$ with $\Lambda^\top = \Lambda$ and $\lambda_{ii} = 0$ for $i = 1, \dots, p$.

From the form of the density it is obvious that whenever κ is “large” compared to Λ , the density will have exactly one maximum (where the vector $c(\theta)$ is approximately aligned with κ) and exactly one minimum (where $c(\theta)$ is approximately aligned with $-\kappa$). This effect is studied in section 2 where we give a sufficient criterion for the distribution to be unimodal. Conversely, for small κ the quadratic term $s^\top(\theta)\Lambda s(\theta)$ in the density φ dominates and one expects the occurrence of multimodal distributions. This situation is studied in section 3 where we show, by example, that a high number of modes is possible even in low dimensions. Finally, in section 4, we give an algorithm for generating samples of a $\text{MVM}(\mu, \kappa, \Lambda)$ distribution for the unimodal case. This will be required as part of any algorithm to sample from a mixture model with $\text{MVM}(\mu, \kappa, \Lambda)$ components.

2 High Concentration

In this section we derive a sufficient criterion for the $\text{MVM}(\mu, \kappa, \Lambda)$ to be unimodal. Since the exponential function \exp in the density (1) is strictly monotonically increasing and since the normalisation constant $Z(\kappa, \Lambda)$ does not depend on θ , it suffices to consider the extrema of

$$f(\theta) = \kappa^\top c(\theta) + \frac{1}{2} s(\theta)^\top \Lambda s(\theta) \quad (2)$$

instead. These can be found by setting the partial derivatives

$$\partial_i f(\theta) = -\kappa_i s_i(\theta) + c_i(\theta) \sum_{k=1}^p \lambda_{ik} s_k(\theta) \quad (3)$$

equal to 0: Since \mathbb{T}^p is a compact, closed manifold, all local extrema of f are located at $\theta \in \mathbb{T}^p$ with $\partial_i f(\theta) = 0$ for $i = 1, \dots, p$, *i.e.* at critical points of f .

To characterise the critical points of f , we consider the second derivatives

$$\partial_{ij} f(\theta) = -\left(\kappa_i c_i(\theta) + s_i(\theta) \sum_{k=1}^p \lambda_{ik} s_k(\theta) \right) \delta_{ij} + c_i(\theta) \lambda_{ij} c_j(\theta) \quad (4)$$

where δ_{ij} denotes the Kronecker delta. If the Hessian matrix $H_f(\theta) = (\partial_{ij} f(\theta))_{i,j}$ at a critical point θ is negative definite, θ is a local maximum of f and thus of $\varphi(\cdot; \mu, \kappa, \Lambda)$; if $H_f(\theta)$ is positive definite, θ is a local minimum; finally, if $H_f(\theta)$ has both positive and negative eigenvalues, the point θ is a saddle point.

For reference in the arguments below, we note that the biggest eigenvalue λ_{\max} of a symmetric matrix $A = (a_{ij}) \in \mathbb{R}^{p \times p}$ satisfies

$$\lambda_{\max} = \sup_{x \in \mathbb{R}^p, |x|=1} x^\top A x \geq \max_{i=1, \dots, p} e_i^\top A e_i = \max_{i=1, \dots, p} a_{ii}$$

where (e_1, \dots, e_p) denotes the standard basis in \mathbb{R}^p . In particular, if the Hessian matrix at a critical point θ has a positive diagonal element, it has at least one positive eigenvalue and thus θ cannot be a local maximum. Similarly, the smallest eigenvalue λ_{\min} satisfies $\lambda_{\min} \leq \min_{i=1, \dots, p} a_{ii}$ and if $H_f(\theta)$ has a negative diagonal element, θ cannot be a local minimum.

Proposition 2.1. Assume that the matrix

$$P = \text{diag}(\kappa_1, \dots, \kappa_p) - \Lambda$$

is positive definite. Then the global maximum of $\varphi = \varphi(\cdot; \mu, \kappa, \Lambda)$ is attained at $\theta = \mu$ and φ has no other (local) maxima.

proof. For $\theta = \mu$ we get $\nabla f(\mu) = 0$ and $H_f(\mu) = -P$; by assumption, this matrix is negative definite and thus $\theta = \mu$ is a *local* maximum. We now show that this is the only local, and thus the *global*, maximum of f .

Since P is positive, the smallest eigenvalue λ_{\min} of P satisfies $0 < \lambda_{\min} \leq \min_{i=1, \dots, p} P_{ii} = \min_{i=1, \dots, p} \kappa_i$ and thus we have $\kappa_i > 0$ for $i = 1, \dots, p$. From equation (3) we see that $\partial_i f(\theta) = 0$ implies $c_i \neq 0$ and consequently $\sum \lambda_{ik} s_k = \kappa_i s_i / c_i$. Substituting this into the expression for $\partial_{ij} f$ in (4) we find that the Hessian matrix H_f at a critical point has the elements

$$\partial_{ij} f(\theta) = -\kappa_i \left(x_i + \frac{s_i^2}{c_i} \right) \delta_{ij} + c_i \lambda_{ij} c_j = -\frac{\kappa_i}{c_i} \delta_{ij} + c_i \lambda_{ij} c_j$$

where we write c for $c(\theta)$ and s for $s(\theta)$ to improve readability. If a critical point θ has $c_i(\theta) < 0$ for an index $i \in \{1, \dots, p\}$, then $H_f(\theta)_{ii} = -\kappa_i / c_i > 0$ and thus θ cannot be a local maximum. Therefore we can assume $c_i(\theta) > 0$ for $i = 1, \dots, p$.

Using the notation

$$P(\theta) = \text{diag}\left(\frac{\kappa_1}{c_1(\theta)}, \dots, \frac{\kappa_p}{c_p(\theta)}\right) - \Lambda \quad (5)$$

we can equivalently re-write the condition $\nabla f(\theta) = 0$ as

$$P(\theta)s(\theta) = 0. \quad (6)$$

Since

$$P(\theta) = P + \text{diag}\left(\kappa_1 \left(\frac{1}{c_1(\theta)} - 1\right), \dots, \kappa_p \left(\frac{1}{c_p(\theta)} - 1\right)\right)$$

is the sum of two positive matrices, it is positive and in particular non-singular. Thus, the only solution of (6) is $s = 0$ which implies that the maximum at $\theta = \mu$ is the only critical point with $c_i \geq 0$ for $i = 1, \dots, p$. This completes the proof. \blacksquare

From the proof of proposition 2.1 we see that, if $\theta = \mu$ is the global and thus a local maximum of φ , the matrix P must be positive *semi*-definite, *i.e.* the positivity condition is almost equivalent to φ having the global maximum at μ . The following corollary gives a sufficient (but not necessary) condition for the statement to hold; the given condition is often easier to verify in practice. Coincidentally, this stronger condition allows to also identify the minima of the von Mises density φ .

Corollary 2.2. Assume

$$\kappa_i > \sum_{j=1}^p |\lambda_{ij}| \quad \text{for all } i = 1, \dots, p. \quad (7)$$

Then the global maximum of $\varphi = \varphi(\cdot; \mu, \kappa, \Lambda)$ is attained at $\theta = \mu$, the global minimum is at $\theta = (\mu_1 + \pi, \dots, \mu_p + \pi)$ and these two points are the only (local) extrema of φ .

proof. By the Gershgorin theorem (Horn and Johnson, 1985, Theorem 6.1.1), the eigenvalues of P are contained in the union of the closed discs $B(\kappa_i, r_i) \subseteq \mathbb{C}$ with radii $r_i = \sum_{j \neq i} |\lambda_{ij}|$ for $i = 1, \dots, p$. Since P is symmetric, its eigenvalues are real and since we have $\kappa_i > \sum_j |\lambda_{ij}| = r_i$, all eigenvalues of P are positive. Thus the condition of the proposition is satisfied and $\theta = 0$ is the global maximum of φ .

Similarly, the eigenvalues of the matrix $P(\theta)$ from (5) are contained in the union of the closed discs $B(\kappa_i/c_i(\theta), r_i) \subseteq \mathbb{R}$ with radii $r_i = \sum_{j \neq i} |\lambda_{ij}|$ for $i = 1, \dots, p$. Since we have

$$\left| \frac{\kappa_i}{c_i(\theta)} \right| \geq \kappa_i > \sum_j |\lambda_{ij}| = r_i,$$

none of the discs contain 0 and the matrix $P(\theta)$ cannot have 0 as an eigenvalue. This shows that all solutions of (6), *i.e.* the critical points of f , satisfy $s(\theta) = 0$ and thus $c(\theta) \in \{-1, 1\}^p$.

To classify the critical points, we consider the Hessian matrix $H_f = (\partial_{ij} f(\theta))_{ij}$. Invoking the Gershgorin theorem again, the eigenvalues of H_f are contained in the union of the closed discs with centres $\partial_{ii} f(\theta)$ and radii $\sum_{j \neq i} |\partial_{ij} f(\theta)|$ for $i = 1, \dots, p$. Using (4) we have

$$|\partial_{ii} f(\theta)| = |\kappa_i c_i(\theta)| = |\kappa_i| > \sum_j |\lambda_{ij}| = \sum_j |c_i(\theta) \lambda_{ij} c_j(\theta)| = \sum_{j \neq i} |\partial_{ij} f(\theta)|,$$

none of these discs contain 0 and thus the circles corresponding to i with $c_i = 1$ and with $c_i = -1$ respectively form two disjoint groups. We can conclude that for each i with $c_i = 1$ the matrix H_f has a negative eigenvalue and for each i with $c_i = -1$ the Hessian has a positive eigenvalue. Consequently, $\theta = \mu$ is the only local maximum of f , $\theta = (\mu_1 + \pi, \dots, \mu_p + \pi)$ is the local minimum of f and all other critical points are saddle points. \blacksquare

It is easy to see that the statements about the minimum in corollary 2.2 do not necessarily hold under the weaker assumption from proposition 2.1. For example, the matrix

$$\Lambda = \begin{pmatrix} 0 & -2 & 2 \\ -2 & 0 & 2 \\ 2 & 2 & 0 \end{pmatrix}$$

has eigenvalues $-4, 2$ and 2 . Thus, for $\kappa = (3, 3, 3)$ the matrix P is positive (the eigenvalues are 1, 1 and 7), and the assumption of proposition 2.1 is satisfied. On the other hand, the Hessian matrix of f at $\theta = (\mu_1 + \pi, \mu_2 + \pi, \mu_3 + \pi)$ is $H_f(\theta) = \text{diag}(\kappa_1, \kappa_2, \kappa_3) + \Lambda$ and, since this matrix is not positive semi-definite (the eigenvalues are $-1, 5$ and 5), the minimum of the distribution cannot be at $(\mu_1 + \pi, \mu_2 + \pi, \mu_3 + \pi)$.

3 Low Concentration

In this section we consider the case of “small” κ . In this case the structure of the extrema of a $\text{MVM}(\mu, \kappa, \Lambda)$ distribution is much more complicated than for the concentrated case. We illustrate some of the possible scenarios with the help of examples, starting with the boundary case $\kappa = (0, 0, \dots, 0)$ and then considering small but non-zero κ .

The following lemma shows that for $\kappa = 0$ the case of a single global maximum can never occur.

Lemma 3.1. For $\kappa = 0$, the following statements hold:

1. The density of the multivariate von Mises distribution $\text{MVM}(\mu, 0, \Lambda)$ takes its maximal value on the set $\{\frac{1}{2}\pi, \frac{3}{2}\pi\}^p \subseteq \mathbb{T}^p$, *i.e.*

$$\sup_{\theta \in \mathbb{T}^p} \varphi(\theta; \mu, 0, \Lambda) = \sup_{\theta \in \{\frac{1}{2}\pi, \frac{3}{2}\pi\}^p} \varphi(\theta; \mu, 0, \Lambda).$$

2. If θ is a maximum, then so is $(\theta_1 + \pi, \dots, \theta_p + \pi)$. In particular the number of isolated maxima of f is always even (and thus cannot be 1).

proof. Without loss of generality, we can assume $\mu = 0$. Let $\theta^* \in \mathbb{T}^p$ be a global maximum of $\varphi(\cdot; 0, 0, \Lambda)$. As in proposition 2.1, this is equivalent to θ^* being a maximum of the function f from equation (2). Since we assume $\kappa = 0$, the formula for f simplifies to

$$f(\theta) = \frac{1}{2} s(\theta)^\top \Lambda s(\theta) \quad (8)$$

and the partial derivatives of f are given by

$$\partial_i f(\theta) = c_i(\theta) \sum_{k=1}^p \lambda_{ik} s_k(\theta). \quad (9)$$

Let $i \in \{1, 2, \dots, p\}$. Since $\lambda_{ii} = 0$, the value $\sum_{k=1}^p \lambda_{ik} s_k(\theta)$ does not depend on θ_i and thus $\theta_i \mapsto \partial_i f(\theta)$ can only change sign at the points $\theta_i = \frac{1}{2}\pi, \frac{3}{2}\pi$. Consequently, $\theta_i \mapsto f(\theta)$ changes monotonically between the values $\theta_i = \frac{1}{2}\pi, \frac{3}{2}\pi$. Defining θ^+ and θ^- by $\theta_i^+ = \frac{1}{2}\pi$, $\theta_i^- = \frac{3}{2}\pi$, and $\theta_j^+ = \theta_j^- = \theta_j^*$ for $j \neq i$ this shows that one of the two inequalities $f(\theta^+) \geq f(\theta^*) \geq f(\theta^-)$ and $f(\theta^-) \geq f(\theta^*) \geq f(\theta^+)$ holds. Since θ^* is a global maximum of f , equality holds in the upper bound and thus either θ^+ or θ^- is also a global maximum. By repeating this procedure for $i = 1, 2, \dots, p$ we find a global maximum where each coordinate is in the set $\{\frac{1}{2}\pi, \frac{3}{2}\pi\}$. This completes the proof of the first statement.

The second statement is a direct consequence of the fact that the function f from (8) is invariant under the map $\theta \mapsto \theta + (\pi, \dots, \pi)$. \blacksquare

Lemma 3.2. Let $\kappa = 0$ and $\Lambda \neq 0$. Then every global maximum θ of $\varphi(\theta; \mu, 0, \Lambda)$ satisfies $\|s(\theta)\|_\infty = 1$.

proof. Since the trace of a matrix equals the sum of its eigenvalues and since Λ is a non-zero matrix with zero trace, Λ must have a strictly positive eigenvalue λ . Let x be a corresponding eigenvalue with $\|x\|_\infty \leq 1$. Then we can find $\theta = (\theta_1, \dots, \theta_p)$ with $\sin(\theta_i) = x_i$ for $i = 1, 2, \dots, p$. This vector θ satisfies

$$f(\theta) = \frac{1}{2} s(\theta)^\top \Lambda s(\theta) = \frac{\lambda}{2} s(\theta)^\top s(\theta) > 0.$$

Consequently the maximal value of f is strictly positive.

Now let $\theta \in \mathbb{T}^p$ with $f(\theta) > 0$ and $\|s(\theta)\|_\infty < 1$, *i.e.* $|s_i(\theta)| < 1$ for all $i \in \{1, \dots, p\}$. Let $c = 1/\|s(\theta)\|_\infty > 1$ and $\tilde{s} = cs(\theta)$. Since $\|\tilde{s}\|_\infty = 1$, we can find $\tilde{\theta} \in \mathbb{T}^p$ with $\sin(\tilde{\theta}_i) = \tilde{s}_i$ for $i = 1, 2, \dots, p$. This point satisfies

$$f(\tilde{\theta}) = \frac{1}{2} \tilde{s}^\top \Lambda \tilde{s} = c^2 \frac{1}{2} s(\theta)^\top \Lambda s(\theta) > f(\theta).$$

Therefore, θ cannot have been a maximum of f . \blacksquare

Example 1 ($\kappa = 0$, two isolated modes). Consider a MVM($\mu, 0, \Lambda$) distribution with

$$\Lambda = \begin{pmatrix} 0 & 1.75 & 0.77 \\ 1.75 & 0 & 0.06 \\ 0.77 & 0.06 & 0 \end{pmatrix}.$$

Since $\kappa = 0$, lemma 3.2 applies and shows that all maxima of $\varphi(\cdot; \mu, 0, \Lambda)$ correspond to θ where $s(\theta)$ lies on the surface of the cube $Q = [-1, 1]^3$. Thus, we can find the local extrema of f by first finding the local extrema of $g(s) = \frac{1}{2} s^\top \Lambda s$ on the surface of Q and then identifying the corresponding values θ . To aid with finding the maxima of g , figure 1 shows a plot of g on the (unwrapped) surface of Q . In the figure, the top-most square corresponds to $s_3 = +1$, the centre square to $s_2 = -1$, the right-most square to $s_1 = +1$ and so on. One can see that the distribution has two modes, corresponding to $s(\theta) = (-1, -1, -1)$ and $s(\theta) = (+1, +1, +1)$.




Figure 1: *Visualisation of a von Mises density from example 1, as a function of $s(\theta)$ restricted to the surface of the cube $[0,1]^3$. The plot shows that the distribution has two modes.*

Example 2 ($\kappa = 0$, one extended mode). Consider a MVM($\mu, 0, \Lambda$) distribution with

$$\Lambda = \begin{pmatrix} 0 & -1 & 1 \\ -1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

We can find the modes of this distribution in the same way as we did in example 1, the corresponding plot of g is shown in figure 2. The figure shows that the density of this MVM($\mu, 0, \Lambda$) distribution has an extended maximum which forms a loop on the surface of the cube. Figure 2 shows the density as a function of $s(\theta)$. To give an idea of the distribution of the corresponding angles $\theta_1, \theta_2, \theta_3$ themselves, we show a scatter plot of a sample in figure 3. While this (much more conventional) diagram shows the distribution of the sample clearly, comparison with figure 2 makes it clear that the structure of the mode is difficult to understand from a scatter plot alone.

The case of small, non-zero κ can be seen as a perturbation of the case $\kappa = 0$. Such a perturbation would normally just shift the extrema of the density around, but the following example shows that such a perturbation can also break a spatially extended maximum into a set of isolated maxima, thus increasing the number of modes.

Example 3 ($\kappa > 0$, six isolated modes). The maximum of the von Mises distribution illustrated in figure 2 lives on a “ring” formed as the union of six lines in \mathbb{T}^3 , aligned with the grid $\{\frac{1}{2}\pi, \frac{3}{2}\pi\}^3$. Since the c_i are zero on the grid and take their maxima between the grid points, we would expect that adding a perturbation term $\kappa^\top c(\theta)$ with small κ will not only shift these lines, but will also collapse this extended maximum into a collection of isolated maxima which live on the shifted lines, at the point where the perturbation was maximal. The following, explicit example gives a von Mises distribution in \mathbb{T}^3 with six isolated maxima.

Let $\eta > 0$ and $\varepsilon = \sin(\eta)$. Define

$$\kappa = \begin{pmatrix} \varepsilon \\ \varepsilon \\ \varepsilon \end{pmatrix}, \quad \Lambda = \begin{pmatrix} 0 & -1 & 1 \\ -1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

We will show that, for small enough η , the function $f(\theta) = \kappa^\top c(\theta) + \frac{1}{2}s(\theta)^\top \Lambda s(\theta)$ has local maxima at the six points $\theta^1, \dots, \theta^6 \in \mathbb{T}^3$ given by the following table.

ℓ	θ_1^ℓ	θ_2^ℓ	θ_3^ℓ	$s_1(\theta^\ell)$	$s_2(\theta^\ell)$	$s_3(\theta^\ell)$	$c_1(\theta^\ell)$	$c_2(\theta^\ell)$	$c_3(\theta^\ell)$
1	0	$\frac{3}{2}\pi + \eta$	$\frac{3}{2}\pi + \eta$	0	$-\sqrt{1-\varepsilon^2}$	$-\sqrt{1-\varepsilon^2}$	1	ε	ε
2	$\frac{1}{2}\pi - \eta$	$\frac{3}{2}\pi + \eta$	0	$\sqrt{1-\varepsilon^2}$	$-\sqrt{1-\varepsilon^2}$	0	ε	ε	1
3	$\frac{1}{2}\pi - \eta$	0	$\frac{1}{2}\pi - \eta$	$\sqrt{1-\varepsilon^2}$	0	$\sqrt{1-\varepsilon^2}$	ε	1	ε
4	0	$\frac{1}{2}\pi - \eta$	$\frac{1}{2}\pi - \eta$	0	$\sqrt{1-\varepsilon^2}$	$\sqrt{1-\varepsilon^2}$	1	ε	ε
5	$\frac{3}{2}\pi + \eta$	$\frac{1}{2}\pi - \eta$	0	$-\sqrt{1-\varepsilon^2}$	$\sqrt{1-\varepsilon^2}$	0	ε	ε	1
6	$\frac{3}{2}\pi + \eta$	0	$\frac{3}{2}\pi + \eta$	$-\sqrt{1-\varepsilon^2}$	0	$-\sqrt{1-\varepsilon^2}$	ε	1	ε

For the convenience of the reader, the table also gives the vectors $s(\theta^\ell)$ and $c(\theta^\ell)$ for $\ell = 1, \dots, 6$. By substituting these values into the formula for $\partial_i f$ from equation (3), it is easy to check that $\nabla f(\theta^\ell) = 0$ for $\ell = 1, \dots, 6$ and thus all six points are critical points of f .

Substituting the values from the table into the formulas for $\partial_{ij} f$ from (4), we can compute the value of the Hessian matrix $H_\ell = H_f(\theta^\ell)$ for $\ell = 1, \dots, 6$. The results are as follows:

$$H_1 = H_4 = \begin{pmatrix} -\varepsilon & -\varepsilon & \varepsilon \\ -\varepsilon & -1 & \varepsilon^2 \\ \varepsilon & \varepsilon^2 & -1 \end{pmatrix}, \quad H_2 = H_5 = \begin{pmatrix} -1 & -\varepsilon^2 & \varepsilon \\ -\varepsilon^2 & -1 & \varepsilon \\ \varepsilon & \varepsilon & -\varepsilon \end{pmatrix}, \quad H_3 = H_6 = \begin{pmatrix} -1 & -\varepsilon & \varepsilon^2 \\ -\varepsilon & -\varepsilon & \varepsilon \\ \varepsilon^2 & \varepsilon & -1 \end{pmatrix}.$$

It can be checked that each of these matrices has eigenvalues $\lambda_1 = -1 + \mathcal{O}(\varepsilon^2)$, $\lambda_2 = -1 + \mathcal{O}(\varepsilon^2)$, and $\lambda_3 = -\varepsilon + \mathcal{O}(\varepsilon^2)$. Thus, for small enough $\varepsilon > 0$, all six points are local maxima as required.




Figure 2: *Visualisation of a von Mises density from example 2. One can see that the distribution has an extended maximum which loops around the cube in a “zig-zag belt”.*




Figure 3: *Scatter plot of 1000 samples from the distribution from example 2. To make the structure of the maximum more visible, the matrix Λ was multiplied by 10, i.e. the plotted sample is from a $\text{MVM}(0, 0, 10\Lambda)$ distribution. The regions where the scatter plots have higher intensity are not isolated modes of the distribution but are artefacts caused by the projection of \mathbb{T}^3 onto \mathbb{T}^2 where straight segments of the extended maximum are seen “head-on”.*

4 Sampling

In this section we discuss a simple method to generate samples from a MVM(μ, κ, Λ) distribution, using the rejection sampling algorithm (Robert and Casella, 2004, Corollary 2.17). The method is restricted to small or moderate p , but works well for the case of high concentration. We assume that the matrix

$$P = \text{diag}(\kappa_1, \dots, \kappa_p) - \Lambda$$

is positive definite.

Without loss of generality we can assume $\mu = 0$, the general case is then obtained by a simple shift. We denote the smallest eigenvalue of P by $\lambda_{\min} > 0$. The proposed algorithm uses independent angles $\theta_1, \theta_2, \dots, \theta_p$ as proposals, distributed with density

$$g(\theta) = \prod_{i=1}^p \frac{\exp\left(\frac{\lambda_{\min}}{4} \cos(2\theta)\right)}{2\pi I_0\left(\frac{\lambda_{\min}}{4}\right)}.$$

This is the independent product of one-dimensional von Mises distributions, modified by replacing the angle θ by 2θ . Since we can efficiently generate samples $\tilde{\theta}_i$ from a one-dimensional von Mises distribution VM($0, \lambda_{\min}/4$) (*e.g.* Best and Fisher, 1979), we can obtain samples from the density g by taking $\theta_i = \tilde{\theta}/2$ with probability $1/2$ and $\theta_i = \tilde{\theta}/2 + \pi$ else.

The target density is the density of the multivariate von Mises distribution MVM($0, \kappa, \Lambda$), *i.e.* it is proportional to

$$f(\theta) = \exp\left(\kappa^\top c(\theta) + \frac{1}{2}s(\theta)^\top \Lambda s(\theta)\right).$$

Using the inequalities $\cos(\theta) + \sin(\theta)^2/2 \leq 1$ and $s(\theta)^\top P s(\theta) \geq \lambda_{\min} s(\theta)^\top s(\theta)$, we find

$$\begin{aligned} f(\theta) &= \exp\left(\kappa^\top c(\theta) + \frac{1}{2}s(\theta)^\top \Lambda s(\theta)\right) \\ &= \exp\left(\sum_{i=1}^p \kappa_i (c_i(\theta) + \frac{1}{2}s_i(\theta)^2) - \frac{1}{2}s(\theta)^\top P s(\theta)\right) \\ &\leq \exp\left(\sum_{i=1}^p \kappa_i - \frac{\lambda_{\min}}{2}s(\theta)^\top s(\theta)\right). \end{aligned}$$

Finally, since $\cos(2x) = 1 - 2\sin(x)^2$, we can rewrite this expression as

$$\begin{aligned} f(\theta) &\leq \exp\left(-\frac{p\lambda_{\min}}{4} + \sum_{i=1}^p \kappa_i\right) \cdot \exp\left(\frac{\lambda_{\min}}{4} \sum_{i=1}^p c_i(\theta)\right) \\ &= \exp\left(-\frac{p\lambda_{\min}}{4} + \sum_{i=1}^p \kappa_i\right) \cdot (2\pi I_0\left(\frac{\lambda_{\min}}{4}\right))^p \cdot g(\theta) \\ &=: Cg(\theta). \end{aligned}$$

Thus we have found a constant C with $f \leq Cg$ and the rejection sampling algorithm can be applied.

In the rejection sampling algorithm, a proposal θ is accepted with probability $f(\theta)/Cg(\theta)$, *i.e.* with probability

$$\begin{aligned} p(\theta) &= \frac{\exp\left(\kappa^\top c(\theta) + \frac{1}{2}s(\theta)^\top \Lambda s(\theta)\right)}{\exp\left(\sum_{i=1}^p \kappa_i - \frac{\lambda_{\min}}{2}s(\theta)^\top s(\theta)\right)} \\ &= \exp\left(\sum_{i=1}^p \kappa_i (c_i - 1) + \frac{1}{2}s^\top (\Lambda + \lambda_{\min} I) s\right) \end{aligned}$$

where I is the $p \times p$ identity matrix. Thus, the following algorithm can be used to generate samples of a MVM(μ, κ, Λ) distribution when P is positive:

1. Generate random variables

$$\begin{aligned}\tilde{\theta}_1, \dots, \tilde{\theta}_p &\sim \text{VM}(0, \lambda_{\min}/4) \\ \delta_1, \dots, \delta_n &\text{ with } P(\delta_i = 0) = P(\delta_i = \pi) = 1/2 \\ U &\sim \mathcal{U}([0, 1]),\end{aligned}$$

all independent of each other.

2. Let $s_i = \sin(\theta_i)$ and $c_i = \cos(\theta_i)$ for $i = 1, 2, \dots, p$.
3. If the condition

$$U \leq \exp\left(\sum_{i=1}^p \kappa_i (c_i(\theta) - 1) + \frac{1}{2} s(\theta)^\top (\Lambda + \lambda_{\min} I) s(\theta)\right)$$

is satisfied, output $\theta = (\theta_1 + \mu_1, \theta_2 + \mu_2, \dots, \theta_p + \mu_p)$ (*i.e.* the proposal is accepted).

4. Return to step 1.

We note that the algorithm still works when the eigenvalue λ_{\min} is replaced by a lower bound $0 < \hat{\lambda}_{\min} \leq \lambda_{\min}$ for the eigenvalues of P . This allows to apply the algorithm in situations where the eigenvalues of P are not exactly known.

The efficiency of this algorithm is determined by its acceptance rate: If Z is the normalisation constant which makes $\frac{1}{Z}f$ a probability density, then each proposal is accepted with probability Z/C . From Mardia et al. (2011, equation (3)) we know that, for high concentration, we have

$$Z \approx (2\pi)^{p/2} |P|^{-1/2} \exp\left(\sum_{i=1}^p \kappa_i\right)$$

where $|P|$ is the determinant of the matrix P . From Abramowitz and Stegun (1964, formula 9.7.1) we know

$$\sqrt{2\pi\kappa} e^{-\kappa} I_0(\kappa) \longrightarrow 1$$

as $\kappa \rightarrow \infty$. Consequently, the asymptotic acceptance probability for high concentration is

$$\begin{aligned}\frac{Z}{C} &\approx \frac{(2\pi)^{p/2} |P|^{-1/2} \exp(\sum_{i=1}^p \kappa_i)}{\exp(-p\lambda_{\min}/4 + \sum_{i=1}^p \kappa_i) \cdot (2\pi)^{p/2} (4/\lambda_{\min})^{p/2} \exp(p\lambda_{\min}/4)} \\ &= \frac{1}{2^p} \cdot \sqrt{\frac{\lambda_{\min}^p}{|P|}}.\end{aligned}\tag{10}$$

The proposed algorithm will be efficient if this probability is not too small. Considering the first factor on the right-hand side of (10), we see that the method only can be expected to perform well for sufficiently small values of p . The factor $1/2^p$ is expected, since the proposal distribution has 2^p modes, whereas the target distribution has only one. Since the determinant $|P|$ equals the product of all p eigenvalues of P (the smallest of which is λ_{\min}), the second factor on the right-hand side of (10) is big, if the eigenvalues of P are all of the same magnitude, *i.e.* if the mode of the distribution is approximately rotationally symmetric.

Acknowledgements. The authors wish to thank John Kent for many helpful discussions.

References

M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. Dover Publications, 1964.

- D. J. Best and N. I. Fisher. Efficient simulation of the von Mises distribution. *Journal of the Royal Statistical Society. Series C*, 28(2):152–157, 1979. URL <http://www.jstor.org/stable/2346732>.
- C. I. Branden and J. Tooze. *Introduction to Protein Structure*. Garland, second edition, 1998.
- J. Frellsen, I. Moltke, T. M., K. Mardia, J. Ferkinghoff-Borg, and T. Hamelryck. A probabilistic model of RNA conformational space. *PLoS Comput. Biol.*, 5(6), 2009. doi: 10.1371/journal.pcbi.1000406.
- T. Harder, W. Boomsma, M. Paluszewski, J. Frellsen, K. Johansson, and T. Hamelryck. Beyond rotamers: a generative, probabilistic model of side chains in proteins. *BMC Bioinformatics*, 11(306), 2010. doi: 10.1186/1471-2105-11-306.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.
- K. V. Mardia and P. E. Jupp. *Directional Statistics*. Wiley, 2000.
- K. V. Mardia, G. Hughes, C. C. Taylor, and H. Singh. A multivariate von Mises distribution with applications to bioinformatics. *The Canadian Journal of Statistics*, 36(1):99–109, 2008.
- K. V. Mardia, J. T. Kent, Z. Zhang, C. Taylor, and T. Hamelryck. Mixtures of concentrated multivariate sine distributions with applications to bioinformatics. Submitted, 2011.
- C. P. Robert and G. Casella. *Monte Carlo statistical methods*. Springer Texts in Statistics. Springer, second edition, 2004.