

Saddlepoint approximations for the normalizing constant of Fisher–Bingham distributions on products of spheres and Stiefel manifolds

BY A. KUME

*Institute of Mathematics, Statistics and Actuarial Science, University of Kent,
Canterbury CT2 7NF, U.K.
a.kume@kent.ac.uk*

S. P. PRESTON AND ANDREW T. A. WOOD

*School of Mathematical Sciences, University of Nottingham, University Park,
Nottingham NG7 2RD, U.K.
simon.preston@nottingham.ac.uk andrew.wood@nottingham.ac.uk*

SUMMARY

In an earlier paper Kume & Wood (2005) showed how the normalizing constant of the Fisher–Bingham distribution on a sphere can be approximated with high accuracy using a univariate saddlepoint density approximation. In this sequel, we extend the approach to a more general setting and derive saddlepoint approximations for the normalizing constants of multicomponent Fisher–Bingham distributions on Cartesian products of spheres, and Fisher–Bingham distributions on Stiefel manifolds. In each case, the approximation for the normalizing constant is essentially a multivariate saddlepoint density approximation for the joint distribution of a set of quadratic forms in normal variables. Both first-order and second-order saddlepoint approximations are considered. Computational algorithms, numerical results and theoretical properties of the approximations are presented. In the challenging high-dimensional settings considered in this paper the saddlepoint approximations perform very well in all examples considered.

Some key words: Directional data; Fisher matrix distribution; Kent distribution; Orientation statistics.

1. INTRODUCTION

The density of the Fisher–Bingham distribution on the unit sphere $\mathcal{S}_{d-1} = \{y \in \mathbb{R}^d : y^T y = 1\}$, with respect to Lebesgue measure on \mathcal{S}_{d-1} , is

$$f_0(y; \delta, G) = \mathcal{C}_0(\delta, G)^{-1} \exp(\delta^T y + y^T G y), \quad (1)$$

where the elements of $\delta \in \mathbb{R}^d$ and the $d \times d$ real symmetric matrix G are parameters and $\mathcal{C}_0(\delta, G)$ is the normalizing constant. The family of distributions (1) contains a number of subfamilies of importance in direction statistics and shape analysis, including the von Mises–Fisher distribution, the Bingham distribution, the complex Bingham distribution and the Kent distribution; see, for example, Bingham (1974), Kent (1982, 1994) and, for a general review, Mardia & Jupp (2000). Kume & Wood (2005) provided a convenient and accurate saddlepoint method for approximating the normalizing constant $\mathcal{C}_0(\delta, G)$.

The primary purpose of this paper is to extend the [Kume & Wood \(2005\)](#) saddlepoint approach to enable the accurate approximation of the normalizing constant of two types of generalization of (1): to Fisher–Bingham distributions on Cartesian products of spheres, and to Fisher–Bingham distributions on Stiefel manifolds. The common feature in the two cases is that we consider the most general exponential family in which the exponent consists of the sum of linear and quadratic terms. These distributions may be defined by conditioning on suitable quadratic constraints in a multivariate Gaussian distribution. Both first- and second-order saddlepoint approximations are considered; as in [Kume & Wood \(2005\)](#), it turns out that the second-order approximations typically have higher accuracy, often to an order of magnitude. In the present setting, this finding is perhaps rather surprising, due to the complexity of the second-order term.

As a particular case of the Fisher–Bingham distribution on a product-of-spheres, we briefly consider a multicomponent generalization of the [Kent \(1982\)](#) distribution. It turns out that for this distribution, the limiting relative errors of the saddlepoint approximations are zero as concentration goes to infinity, provided the distribution is unimodal.

Let S_{d-1}^p denote the Cartesian product of p copies of S_{d-1} , and consider $x = (x_1^T, \dots, x_p^T)^T \in S_{d-1}^p$, i.e., each x_i is a unit vector in S_{d-1} . We define the density, with respect to Lebesgue measure on S_{d-1}^p , of the Fisher–Bingham distribution on S_{d-1}^p by

$$\begin{aligned} f_1(x; a, B) &= C_1(a, B)^{-1} \exp(a^T x + x^T B x) \\ &= C_1(a, B)^{-1} \exp\left(\sum_{i=1}^p a_i^T x_i + \sum_{i,j=1}^p x_i^T B_{ij} x_j\right), \end{aligned} \tag{2}$$

where $a = (a_1^T, \dots, a_p^T)^T \in \mathbb{R}^{dp}$, $B = (B_{ij})_{i,j=1}^p$ is a block matrix with p^2 blocks B_{ij} of dimension $d \times d$, and $C_1(a, B)$ is the normalizing constant chosen so that the integral of $f_1(x; a, B)$ over S_{d-1}^p is 1. Without loss of generality we assume B is symmetric, so that $B_{ij} = B_{ji}^T$. [Mardia \(1975\)](#) considered (2) in the case $p = 2$; see also [Jupp & Mardia \(1980\)](#), [Mardia & Jupp \(2000\)](#), [Singh et al. \(2002\)](#) and [Mardia et al. \(2008\)](#).

The second case we consider is a generalization of (1) to a Stiefel manifold. For $d \geq q$, the Stiefel manifold $\mathcal{V}_{d,q}$ is defined by $\mathcal{V}_{d,q} = \{Y (d \times q) : Y^T Y = I_q\}$, where I_q is the $q \times q$ identity matrix. Write $X = (x_1, \dots, x_q)$ for the $d \times q$ matrix with columns x_1, \dots, x_q and define the vec operator by $\text{vec}(X) = (x_1^T, \dots, x_q^T)^T$. The density of the Fisher–Bingham distribution on $\mathcal{V}_{d,q}$, with respect to Lebesgue measure on $\mathcal{V}_{d,q}$, is

$$\begin{aligned} f_2(X; A, B) &= C_2(A, B)^{-1} \exp\{\text{tr}(A^T X) + \text{vec}(X)^T B \text{vec}(X)\} \\ &= C_2(A, B)^{-1} \exp\left(\sum_{i=1}^q a_i^T x_i + \sum_{i,j=1}^q x_i^T B_{ij} x_j\right), \end{aligned} \tag{3}$$

where $\text{tr}(\cdot)$ denotes trace, $A = (a_1, \dots, a_q)$ is a $d \times q$ matrix, $B = (B_{ij})_{i,j=1}^q$ has the same structure as in (2), and $C_2(A, B)$ is chosen so that the integral of $f_2(X; A, B)$ over $\mathcal{V}_{d,q}$ is 1. The difference between (2) and (3) is that in the latter we have $X^T X = I_q$, so that not only are the x_i unit vectors, but they also satisfy the orthogonality conditions $x_i^T x_j = 0$ if $i \neq j$. The Fisher matrix distribution, corresponding to the case where B in (3) is a matrix of zeros, has been considered by a number of authors, including [Downs \(1972\)](#), [Khatri & Mardia \(1977\)](#), [Jupp & Mardia \(1979\)](#), [Prentice \(1986\)](#) and [Wood \(1993\)](#). It turns out that in the Fisher matrix case the proposed saddlepoint approximations are available in closed form. The Bingham version of (3), corresponding

to the case where A is the matrix of zeros, has been studied by [Arnold & Jupp \(2013\)](#); see also [Hoff \(2009\)](#) for simulation from (3). In § 5 we consider a parametric model for rigid body motion data in which (3) arises as a marginal distribution.

The purpose of this paper is to present saddlepoint approximations to $\mathcal{C}_1(a, B)$ and $\mathcal{C}_2(A, B)$. Matlab code for performing the calculations is available in the Supplementary Material.

2. REPRESENTATION OF THE NORMALIZING CONSTANTS

First, a brief heuristic outline of our approach is given, with the aim of explaining how and why it works. The key theoretical result is then presented formally in Proposition 1 below.

Consider a $d \times q$ matrix X with real components and assume $d \geq q$. Provided $U = X^T X$ has full rank, we can write $X = YU^{1/2}$, where $U^{1/2}$ is the upper-triangular matrix square root of U . Multiplying on the right by the inverse of $U^{1/2}$ yields $Y = X(U^{1/2})^{-1}$, which implies $Y^T Y = I_q$ and therefore $Y \in \mathcal{V}_{d,q}$. If X has a matrix Gaussian density $f(X)$ and is such that $U = X^T X$ has full rank with probability 1, then, ignoring regions with probability zero,

$$f(X)dX = f(YU^{1/2})J(U, Y)(dU) d\mathcal{V}_{d,q}(Y), \tag{4}$$

where dU is the implied volume element on the space of positive-definite matrices, $d\mathcal{V}_{d,q}(Y)$ is the implied volume element on $\mathcal{V}_{d,q}$ and $J(U, Y)$ is the relevant Jacobian determinant.

The key points are now summarized without proof: (a) $J(U, Y) \equiv J(U)$ does not depend on Y , and $J(U)$ may be calculated explicitly using Jacobian theory presented in the Supplementary Material; (b) since $J(U, Y) \equiv J(U)$ and $f(X)$ is a Gaussian density, the conditional distribution of Y given $U = U^\dagger$ is of the form (3); (c) without loss of generality we may take $U^\dagger = I_q$, the $q \times q$ identity matrix; (d) integrating the right-hand side of both (3) and (4) over Y shows that the normalizing constant $\mathcal{C}_2(A, B)$ is equal to the marginal density of U at $U = I_q$; (e) the marginal density of U is equal to a known function of U multiplied by the joint density, evaluated at $U = I_q$, of a set of quadratic forms in Gaussian variables; (f) although this density is not known in closed form, its joint moment generating function is known, so that a saddlepoint density approximation may be used; (g) this way of expressing the normalizing constant can be extended from a single $Y \in \mathcal{V}_{d,q}$ to (Y_1, \dots, Y_p) in a Cartesian product of Stiefel manifolds, not necessarily of the same dimension.

[Kume & Wood \(2005\)](#) cover the case where $q = 1$, Y is a unit vector and U is a positive scalar. With general integer $q \geq 1$ considered here, points (a) and (d) are more challenging, and the calculations involved in (e) are more difficult to implement because a multivariate saddlepoint approximation is needed and the second-order term is far more complex.

Before stating Proposition 1, some definitions are needed. Let $\mathcal{V}_{d,q}^p = \mathcal{V}_{d_1,q_1} \times \dots \times \mathcal{V}_{d_p,q_p}$ denote a Cartesian product of p Stiefel manifolds, not necessarily of the same dimension. The Fisher–Bingham density with respect to Lebesgue measure on $\mathcal{V}_{d,q}^p$ is

$$f_3(X_1, \dots, X_p; A_1, \dots, A_p; B) = \mathcal{C}_3(A_1, \dots, A_p; B)^{-1} \exp \left\{ \sum_{i=1}^p \text{tr}(A_i^T X_i) + \sum_{i,j=1}^p \text{vec}(X_i)^T B_{ij} \text{vec}(X_j) \right\}, \tag{5}$$

where for $i, j = 1, \dots, p$, $X_i \in \mathcal{V}_{d_i,q_i}$, A_i is a $d_i \times q_i$ matrix and B_{ij} is a $(d_i q_i) \times (d_j q_j)$ matrix; and $\mathcal{C}_3(A_1, \dots, A_p; B)$ is chosen so that the integral of $f_3(X_1, \dots, X_p; A_1, \dots, A_p; B)$ over

$\mathcal{V}_{d,q}^p$ is unity. Both (2) and (3) are particular cases of (5). This density (5) was mentioned by Mardia (1975, Reply to Discussion) in the case $p = 2, d_1 = d_2, q_1 = q_2$.

Let $V^{-1} = (V^{ij})_{i,j=1}^p$ denote a symmetric matrix consisting of p^2 blocks, where V^{ij} is a $(d_i q_i) \times (d_j q_j)$ matrix given by $V^{ij} = -2B_{ij}$. We shall assume that V^{-1} and therefore V are positive definite; it will be seen in § 3.5 that there is no loss of generality in doing so. Now define μ via the equation $\mu = V\{\text{vec}(A_1)^T, \dots, \text{vec}(A_p)^T\}^T$. In this case, we have the following result. The vech operator, designed for symmetric matrices, stacks the lower triangular part of the matrix with the diagonal included, using the ordering implied by vec .

PROPOSITION 1. *Assume V is a positive-definite symmetric matrix of dimension $r = \sum_{i=1}^p d_i q_i$ and that $x = \{\text{vec}(X_1)^T, \dots, \text{vec}(X_p)^T\}^T \sim \mathcal{N}_r(\mu, V)$. Then the normalizing constant $C_3(A_1, \dots, A_p; B)$ in (5) may be written*

$$C_3(A_1, \dots, A_p; B) = g(s_1, \dots, s_p; \mu, V)(2\pi)^{r/2} |V|^{1/2} \exp\left(\frac{1}{2} \mu^T V^{-1} \mu\right) 2^{\sum_{i=1}^p q_i},$$

where $s_i = \text{vech}(I_{q_i})$ and $g(u_1, \dots, u_p; \mu, V)$ is the joint density of the $u_i = \text{vech}(X_i^T X_i)$.

Although the density $g(u_1, \dots, u_p; \mu, V)$ is not available in explicit form, the moment generating function of the distribution with density function g is available in convenient form. Consequently, g can be approximated using saddlepoint methods, given by (6) or (7). Although there is no problem in principle in developing saddlepoint approximations in the general setting of Proposition 1, we have chosen to focus on the particular cases (2) and (3) in the remainder of the paper, which are likely to be the cases of most practical interest.

Case 1. $C_1(a, B)$. This arises when $d_1 = \dots = d_p = d, q_1 = \dots = q_p = 1$. Then put $A_i = a_i$, where a_i is $d \times 1$, and $a = (a_1^T, \dots, a_p^T)^T$, in which case $C_1(a, B) = C_3(A_1, \dots, A_p; B)$. Moreover, in this case g is the joint density of $u_1 = x_1^T x_1, \dots, u_p = x_p^T x_p$, where $x = (x_1^T, \dots, x_p^T)^T \sim \mathcal{N}_{dp}(\mu, V)$, and the distribution with density (2) is obtained by conditioning $x \sim \mathcal{N}_{dp}(\mu, V)$ on $u_1 = \dots = u_p = 1$.

Case 2. $C_2(A, B)$. Here we have $p = 1, d_1 = d, q_1 = q$. Then $C_2(A, B) = C_3(A, B)$, and here g is the joint density of $\text{vech}(X^T X)$, where $x = \text{vec}(X) \sim \mathcal{N}_{dq}(\mu, V)$. The distribution with density (3) is obtained by conditioning $x \sim \mathcal{N}_{dq}(\mu, V)$ on $X^T X = I_q$.

3. SADDLEPOINT APPROXIMATIONS

3.1. Review of multivariate saddlepoint density approximations

Later in this section we derive saddlepoint approximations for $C_1(a, B)$ and $C_2(A, B)$. First, general multivariate saddlepoint density approximations are briefly summarized; see, for example, Butler (2007) for further details.

Let $M(\theta)$ and $K(\theta) = \log M(\theta)$ denote, respectively, the moment generating function and cumulant generating function of a random m -vector x . If x has an absolutely continuous distribution on \mathbb{R}^m , then the first-order saddlepoint approximation, $\hat{f}_1(x)$, of the density $f(x)$ of x is

$$\hat{f}_1(x) = (2\pi)^{-m/2} |\hat{K}''|^{-1/2} \exp(\hat{K} - \hat{\theta}^T x), \tag{6}$$

where $\hat{\theta}$ is the unique solution to the saddlepoint equation $K'(\theta) \equiv \partial K(\theta)/\partial \theta = x$; K'' is the $d \times d$ matrix of second derivatives of K ; $|\cdot|$ denotes determinant; and, here and below, a circumflex means evaluation of a function of θ at $\theta = \hat{\theta}$.

Two versions of the second-order saddlepoint approximation are

$$\hat{f}_2(x) = \hat{f}_1(x)(1 + T), \quad \hat{f}_3(x) = \hat{f}_1(x) \exp(T), \tag{7}$$

where

$$T = \frac{1}{8}\hat{\rho}_4 - \frac{1}{24}(3\hat{\rho}_{13}^2 + 2\hat{\rho}_{23}^2) \tag{8}$$

and

$$\hat{\rho}_4 = \hat{K}_{ijkl}\hat{K}^{ij}\hat{K}^{kl}, \quad \hat{\rho}_{23}^2 = \hat{K}_{ijk}\hat{K}_{rst}\hat{K}^{ir}\hat{K}^{js}\hat{K}^{kt}, \quad \hat{\rho}_{13}^2 = \hat{K}_{ijk}\hat{K}_{rst}\hat{K}^{ij}\hat{K}^{kr}\hat{K}^{st}, \tag{9}$$

where \hat{K}^{ij} is component (i, j) of $(\hat{K}'')^{-1}$, and the summation convention is assumed in (9), i.e., when an index occurs as both a subscript and a superscript in a product of terms, then summation over that index is implied. The approximation \hat{f}_3 in (7) has the potential advantage that it is positive even when $T \leq -1$, although in the extensive numerical calculations performed for this paper, we never encountered an example where $T \leq -1$. Our general findings are that \hat{f}_3 and \hat{f}_2 are typically very close, but that \hat{f}_3 tends to be slightly more accurate.

3.2. Saddlepoint approximation for $\mathcal{C}_1(a, B)$

The joint moment generating function of $r_1 = x_1^T x_1, \dots, r_p = x_p^T x_p$ is

$$\begin{aligned} & \int \exp\left(\sum_{i=1}^p \theta_i x_i^T x_i\right) \phi_{dp}(x; \mu, V) dx_1, \dots, dx_p \\ &= \frac{1}{|C(\theta)|^{1/2}|V|^{1/2}} \exp\left\{\frac{1}{2}\mu^T V^{-1}C(\theta)^{-1}V^{-1}\mu - \frac{1}{2}\mu^T V^{-1}\mu\right\} \end{aligned}$$

where $\phi_m(x; \mu, V)$ is the m -dimensional multivariate Gaussian density of $\mathcal{N}_m(\mu, V)$,

$$C(\theta) = V^{-1} - 2 \sum_{i=1}^p \theta_i J_i, \tag{10}$$

where J_i is a block diagonal matrix with p^2 blocks, with the $d \times d$ identity matrix in the i th diagonal block and zeros elsewhere. Consequently, the cumulant generating function $K(\theta)$ is

$$K(\theta) = -\frac{1}{2} \log |C(\theta)| - \frac{1}{2} \log |V| + \frac{1}{2}\mu^T V^{-1}C(\theta)^{-1}V^{-1}\mu - \frac{1}{2}\mu^T V^{-1}\mu. \tag{11}$$

Then

$$K_i(\theta) = \frac{\partial K(\theta)}{\partial \theta_i} = \text{tr}\{C(\theta)^{-1}J_i\} + \mu^T V^{-1}C(\theta)^{-1}J_i C(\theta)^{-1}V^{-1}\mu,$$

where we have used the standard results

$$\frac{\partial}{\partial \theta_i} \log |C(\theta)| = \text{tr}\left\{C(\theta)^{-1} \frac{\partial C}{\partial \theta_i}(\theta)\right\}, \quad \frac{\partial}{\partial \theta_i} C(\theta)^{-1} = -C(\theta)^{-1} \left\{ \frac{\partial C}{\partial \theta_i}(\theta) \right\} C(\theta)^{-1}. \tag{12}$$

Higher derivatives of $K(\theta)$, which are needed to obtain T in (8) via (9), may be calculated using (12) repeatedly. Simplified expressions for these derivatives are given in the Supplementary Material; see (S1)–(S3).

Applying Proposition 1 and (11), and depending on which of the saddlepoint approximations in (6) and (7) is used, we obtain the following approximations for the normalizing constant $C_1(a, B)$:

$$\hat{C}_{1,k}(a, B) = 2^p (2\pi)^{p(d-1)/2} |\hat{K}''|^{-1/2} |\hat{C}|^{-1/2} \exp\left(\frac{1}{2} a^\top \hat{C}^{-1} a - \sum_{i=1}^p \hat{\theta}_i\right) R_k(T), \tag{13}$$

where T is defined in (8) and

$$R_k(T) = \begin{cases} 1, & k = 1, \\ 1 + T, & k = 2, \\ \exp(T), & k = 3. \end{cases} \tag{14}$$

3.3. Saddlepoint approximation for $C_2(A, B)$

Here we need the joint moment generating function of $\{x_i^\top x_j : 1 \leq i \leq j \leq q\}$, which is

$$\begin{aligned} & \int \exp\left(\sum_{i \leq j} \theta_{ij} x_i^\top x_j\right) \phi_{dq}(x; \mu, V) dx_1, \dots, dx_p \\ &= \frac{1}{|C(\theta)|^{1/2} |V|^{1/2}} \exp\left\{\frac{1}{2} \mu^\top V^{-1} C(\theta)^{-1} V^{-1} \mu - \frac{1}{2} \mu^\top V^{-1} \mu\right\}, \end{aligned}$$

i.e., of the same form as (14), but where now $\theta = \{\theta_{ij} : 1 \leq i \leq j \leq q\}$,

$$C(\theta) = V^{-1} - \sum_{1 \leq i \leq j \leq q} \theta_{ij} (J_{ij} + J_{ji}), \tag{15}$$

and J_{ij} is a $(dq) \times (dq)$ block matrix with q^2 blocks each of dimension $d \times d$, with block (i, j) equal to the $d \times d$ identity matrix I_d and all other blocks equal to the $d \times d$ matrix of zeros. The cumulant generating function $K(\theta)$ is given by the expression (11), but it should be noted that μ, V and $C(\theta)$ are defined differently.

Then, using the identities in (12), we obtain

$$\begin{aligned} K_{(rs)}(\theta) &= \frac{\partial}{\partial \theta_{rs}} K(\theta) \\ &= \frac{1}{2} \text{tr}\{C(\theta)^{-1} (J_{rs} + J_{sr})\} + \frac{1}{2} \mu^\top V^{-1} C(\theta)^{-1} (J_{rs} + J_{sr}) C(\theta)^{-1} V^{-1} \mu, \end{aligned}$$

and the saddlepoint equation in this case is $K_{(rs)}(\theta) = 0$ if $r < s$ and $K_{(rr)}(\theta) = 1$. Further applications of (12) yield higher derivatives that are needed to obtain T in (8) via (9); see (S4)–(S6) in the Supplementary Material.

The resulting saddlepoint approximations, $\hat{C}_{2,k}(A, B)$, with $k = 1, 2, 3$, of the normalizing constant $C_2(A, B)$ are

$$\hat{C}_{2,k}(A, B) = \frac{2^q (2\pi)^{dq/2 - q(q+1)/4}}{|\hat{K}''|^{1/2} |\hat{C}|^{1/2}} \exp\left\{\frac{1}{2} \text{vec}(A)^\top \hat{C}^{-1} \text{vec}(A) - \sum_{i=1}^q \hat{\theta}_{ii}\right\} R_k(T) \tag{16}$$

where $R_k(T)$ is defined in (14).

3.4. The Fisher matrix case

The Fisher matrix distribution is the special case of (3) in which $B = 0_{dq,dq}$, a matrix of zeros, and the density with respect to Lebesgue measure on $\mathcal{V}_{d,q}$ simplifies to

$$C_2(A, 0_{dq,dq})^{-1} \exp\{\text{tr}(A^T X)\}, \quad X \in \mathcal{V}_{d,q}. \tag{17}$$

Consider the isometry $Y = Q^T X R$ on $\mathcal{V}_{d,q}$, where Q and R are, respectively, $d \times d$ and $q \times q$ orthogonal matrices, obtained from the singular valued decomposition $A = Q\Omega R^T$, and $\Omega = (\omega_{ij})$ is $d \times q$, with $\omega_{ij} = 0$ if $i \neq j$ and $\omega_{ii} = \omega_i > 0$ for $i = 1, \dots, q$. Applying this isometry, the density (17) simplifies to

$$C_2(\Omega, 0_{dq,dq})^{-1} \exp\left(\sum_{i=1}^q \omega_i x_{ii}\right), \tag{18}$$

where $C_2(\Omega, 0_{dq,dq}) = C_2(A, 0_{dq,dq})$. For the underlying multivariate normal distribution in § 2, we take $\mu = (\omega_1 e_1^T, \dots, \omega_q e_q^T)^T$, where e_j is a d -vector with all components zero except for component j , which is 1. Then $C(\theta)$ in (15) is $C(\theta) = -\Theta^\dagger \otimes I_d$ where \otimes denotes the Kronecker product, $\Theta^\dagger = \Theta^\dagger(\theta) = (\theta_{ij}^\dagger)_{i,j=1}^q$, and $\theta_{ij}^\dagger = \theta_{ji}^\dagger = \theta_{ij}$ if $i < j$ and $\theta_{ij}^\dagger = 2\theta_{ii}$ if $i = j$. Then $C(\theta)^{-1} = \Phi \otimes I_d$, where $\Phi = (\phi_{rs})_{r,s=1}^q = -(\Theta^\dagger)^{-1}$, and consequently $C(\theta)^{-1} = (\phi_{rs} I_d)_{r,s=1}^q$. Therefore the saddlepoint equation is $K_{(r,s)} = \delta_{r,s}$, where $\delta_{r,s}$ is the Kronecker delta and

$$K_{(r,s)} = \frac{\partial K}{\partial \theta_{rs}} = \begin{cases} d\phi_{rs} + \sum_{u=1}^q \omega_u^2 \phi_{ur} \phi_{us}, & (r < s), \\ d\phi_{rr} + \omega_r^2 \phi_{rr}^2, & (r = s). \end{cases}$$

The unique solution is

$$\hat{\phi}_{rs} = 0 \quad (r < s), \quad \hat{\phi}_{rr} = \{-d + (d^2 + 4\omega_r^2)^{1/2}\} / (2\omega_r^2).$$

As $\hat{\Phi} = (\hat{\phi}_{rs})_{r,s=1}^q$ is diagonal, so is $\hat{\Theta}^\dagger$, and so $\hat{\theta}_{rs} = 0$ if $r < s$ and $\hat{\theta}_{rr} = -1/(2\hat{\phi}_{rr}) = -\omega_r^2 / \{(d^2 + 4\omega_r^2)^{1/2} - d\}$. Further calculations using (S4) in the Supplementary Material show that

$$\hat{K}_{(r_1,s_1)(r_2,s_2)} \begin{cases} 0, & (r_1 \neq r_2 \text{ or } s_1 \neq s_2), \\ d\hat{\phi}_{rr} \hat{\phi}_{ss} + \hat{\phi}_{rr} \hat{\phi}_{ss} (\omega_r^2 \hat{\phi}_{rr} + \omega_s^2 \hat{\phi}_{ss}), & (r_1 = r_2 < s_1 = s_2), \\ 2d\hat{\phi}_{rr}^2 + 4\omega_r^2 \hat{\phi}_{rr}^3, & (r_1 = r_2 = s_1 = s_2), \end{cases}$$

where always $r_i \leq s_i$ by convention.

If, as is the case here at the saddlepoint, $\hat{K}_{ij} = \hat{K}^{ij} = 0$ unless $i = j$, i.e., \hat{K}_{ij} and \hat{K}^{ij} are diagonal, then the formulae in (9) simplify to the following:

$$\hat{\rho}_4^2 = \sum_{i=1}^m \sum_{j=1}^m \hat{K}_{iijj} \hat{K}^{ii} \hat{K}^{jj}, \quad \hat{\rho}_{13}^2 = \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m \hat{K}_{iij} \hat{K}_{jkk} \hat{K}^{ii} \hat{K}^{jj} \hat{K}^{kk}, \tag{19}$$

$$\hat{\rho}_{23}^2 = \sum_{i=1}^m \sum_{j=1}^m \sum_{k=1}^m (\hat{K}_{ijk})^2 \hat{K}^{ii} \hat{K}^{jj} \hat{K}^{kk}. \tag{20}$$

In the case of the Stiefel manifold $\mathcal{V}_{d,q}$, i, j and k are of the form (r, s) where $1 \leq r \leq s \leq q$, and $m = q(q + 1)/2$. Convenient expressions for (19) and (20) in the Fisher matrix case are given in the Supplementary Material; see (S7)–(S12).

3.5. *Practical issues*

As a consequence of the constraint $y^T y = 1$ in (1), it follows that $C_0(\delta, G + \lambda I_d) = e^\lambda C_0(\delta, G)$; see, e.g., formula (2) of Kume & Wood (2005). The constraints in (2) and (3) lead to generalizations of this identity.

In (2) we have the constraints $x_i^T x_i = 1$ for $i = 1, \dots, p$, from which we conclude that $C_1(a, B)$ satisfies the following identity: for any diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_p)$,

$$C_1(a, B + \Lambda \otimes I_d) = \exp\left(\sum_{i=1}^p \lambda_i\right) C_1(a, B). \tag{21}$$

The corresponding identity for $C_2(A, B)$ is that for any symmetric matrix $\Lambda = (\lambda_{ij})_{i,j=1}^q$,

$$C_2(A, B + \Lambda \otimes I_d) = \exp\left(\sum_{i=1}^q \lambda_{ii}\right) C_2(A, B). \tag{22}$$

It is straightforward to check that the approximations (13) and (16) satisfy (21) and (22), respectively. A practical consequence is that we can always arrange for V to be positive definite when calculating (6) and (7).

Finally, we offer some brief comments on numerical implementation. One convenient and numerically robust way to solve the saddlepoint equation is to minimize $K(\theta) - \sum_{i=1}^p \theta_i$, or $K(\theta) - \sum_{i=1}^q \theta_{ii}$, over θ . An important point in practice is to ensure that we restrict attention to θ such that the cumulant generating function $K(\theta)$ is finite. A necessary and sufficient condition for θ to be feasible is that $C(\theta)$ in (10) or (15) is positive definite.

4. THEORETICAL RESULTS

Some results on the relative errors of the saddlepoint approximations of $C_1(a, B)$ and $C_2(A, B)$ are now presented; all proofs are given in the Supplementary Material. Propositions 2 and 3 describe relative accuracy in the neighbourhood of the relevant uniform distribution, while Propositions 4 and 5 describe limiting behaviour for, respectively, highly concentrated Fisher matrix distributions, and highly concentrated multicomponent Kent distributions, defined below.

PROPOSITION 2. *For any $(dp) \times 1$ vector a , any symmetric $(dp) \times (dp)$ matrix B , and any diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_p)$, the saddlepoint approximations in (13) with $k = 1, 2, 3$ satisfy*

$$\lim_{\psi \rightarrow 0} \frac{\hat{C}_{1,k}(\psi a, \psi B + \Lambda \otimes I_d)}{C_{1,k}(\psi a, \psi B + \Lambda \otimes I_d)} = \left\{ \frac{\Gamma(d/2)}{\hat{\Gamma}(d/2)} \right\}^p R_k\{-p/(6d)\}$$

where $R_k(T)$ is defined in (14), $\Gamma(\alpha) = \int_0^\infty z^{\alpha-1} e^{-z} dz$ is the gamma function and $\hat{\Gamma}(x) = (2\pi)^{1/2} x^{x-1/2} e^{-x}$ is Stirling's approximation to $\Gamma(x)$.

The above result is a generalization of Kume & Wood (2005, Proposition 2). The $\Lambda \otimes I_d$ term is included because of the degeneracy indicated in (21). Next we state the corresponding result for the uniform distribution on the Stiefel manifold $\mathcal{V}_{d,q}$. Define

$$\Gamma_m(\alpha) = \pi^{m(m-1)/4} \prod_{i=1}^m \Gamma\{\alpha - (i-1)/2\}, \quad \hat{\Gamma}_m(\alpha) = 2^{m/2} \alpha^{m\alpha} \left(\frac{\pi}{\alpha}\right)^{m(m+1)/4} e^{-m\alpha}, \tag{23}$$

where $\Gamma_m(\alpha)$ is the multivariate gamma function (Muirhead, 1982), and $\hat{\Gamma}_m(\alpha)$ is the large- α approximation to $\Gamma_m(\alpha)$ mentioned by Butler & Wood (2003).

PROPOSITION 3. For any $d \times q$ matrix A , any symmetric $(dq) \times (dq)$ matrix B and any symmetric $q \times q$ matrix Λ , the saddlepoint approximations in (16) with $k = 1, 2, 3$ satisfy

$$\lim_{\psi \rightarrow 0} \frac{\hat{C}_{2,k}(\psi A, \psi B + \Lambda \otimes I_d)}{C_2(\psi A, \psi B + \Lambda \otimes I_d)} = \frac{\Gamma_q(d/2)}{\hat{\Gamma}_q(d/2)} R_k(T)$$

where $T = -q(q^3 - 2q^2 + 17q - 8)/(48d)$, Γ_q is the multivariate gamma function, $\hat{\Gamma}_q$ is defined in (23) and $R_k(T)$ is defined in (14).

Our next result is a generalization of Kume & Wood (2005, Prop. 3).

PROPOSITION 4. For any $d \times q$ matrix A , the saddlepoint approximations in (16) with $k = 1, 2, 3$ satisfy $\lim_{\psi \rightarrow \infty} \hat{C}_{2,k}(\psi A, 0_{dq,dq})/C_2(\psi A, 0_{dq,dq}) = 1$.

For our final result, in (2) define

$$a = (\kappa_1 \mu_1^T, \dots, \kappa_p \mu_p^T)^T, \quad B_{ii'} = \begin{cases} \sum_{j,j'=1}^{d-1} a_{ij,i'j'} \mu_{ij} \mu_{i'j'}^T, & (i \neq i'), \\ \sum_{j=1}^{d-1} \beta_{ij} \mu_{ij} \mu_{ij}^T, & (i = i'), \end{cases} \quad (24)$$

where, without loss of generality, we assume $\kappa_i \geq 0$, $\sum_{j=1}^{d-1} \beta_{ij} = 0$ and $a_{ij,i'j'} = a_{i'j',ij}$ for all $i, i' = 1, \dots, p$ and $j, j' = 1, \dots, d - 1$; and for each i , $\{\mu_i, \mu_{i,1}, \dots, \mu_{i,d-1}\}$ is a set of orthonormal d -vectors. Also, define the $p(d - 1) \times p(d - 1)$ block matrix $\Sigma = (\Sigma_{ii'})_{i,i'=1}^p$ where each block is a $(d - 1) \times (d - 1)$ matrix,

$$\Sigma_{ii'} = \begin{cases} (\alpha_{ij,i'j'})_{j,j'=1}^{d-1}, & (i \neq i'), \\ \kappa_i I_{d-1} - 2 \text{diag}(\beta_{i1}, \dots, \beta_{i,d-1}), & (i = i'). \end{cases} \quad (25)$$

PROPOSITION 5. Consider the distribution (2) with vector a and block matrix $B = (B_{ii'})_{i,i'=1}^p$ defined by (24). Suppose that the matrix $\Sigma = (\Sigma_{ii'})_{i,i'=1}^p$ defined in (25) is positive definite. Then the saddlepoint approximations (13) with $k = 1, 2, 3$ satisfy $\lim_{\psi \rightarrow \infty} \hat{C}_{1,k}(\psi a, \psi B)/C_1(\psi a, \psi B) = 1$.

The distribution defined via (2) and (24) may be thought of as a multicomponent generalization of the Kent (1982) distribution when $p > 1$. When $p = 1$ and $d = 3$ we obtain FB_5 , the standard Kent distribution, and with $p = 1$ and general $d \geq 3$ we obtain the Kent distribution on \mathcal{S}_{d-1} . This distribution has $p(d - 1)\{3 + p(d - 1)\}/2$ free parameters: p parameters from the κ_i , $p(d - 2)$ from the β_{ij} , $p(p - 1)(d - 1)^2/2$ from the $a_{ij,i'j'}$ and $pd(d - 1)/2$ from the μ_i and μ_{ij} . This is the same number of free parameters as in the $p(d - 1)$ -dimensional multivariate normal distribution. The unit d -vectors μ_i and μ_{ij} determine the orientation of the distribution but do not appear in the normalizing constant $C_1(a, B)$. Also, from symmetry considerations, it follows that there is always a stationary point at $x = (\mu_1^T, \dots, \mu_p^T)^T$. Some further calculations show that a necessary and sufficient condition for this stationary point to be a local maximum, and therefore the global maximum of the density, is that Σ in (25) is positive definite. Moreover, if we define $\xi_{ij} = x_i^T \mu_{ij}$ for $i = 1, \dots, p$ and $j = 1, \dots, d - 1$, and replace κ_i , β_{ij} and $a_{ij,i'j'}$ by $\psi \kappa_i$, $\psi \beta_{ij}$ and $\psi a_{ij,i'j'}$ respectively, then, under the assumption that Σ in (25) is positive definite, it may

be shown with some further calculations that, as $\psi \rightarrow \infty$,

$$\psi^{1/2}(\xi_{11}, \dots, \xi_{1,d-1}, \xi_{2,1}, \dots, \xi_{p,d-1})^T \rightarrow \mathcal{N}_{p(d-1)}(0_{p(d-1)}, \Sigma) \tag{26}$$

in distribution.

Proposition 5 above generalizes Kume & Wood (2005, Proposition 3) in two senses: when $p = 1$, from the von Mises–Fisher distribution on \mathcal{S}_{d-1} to the unimodal Kent distribution on \mathcal{S}_{d-1} ; and to $p > 1$.

Limiting asymptotic normality of the type shown in (26) is not sufficient for limiting relative error to be zero, as is the case in Proposition 5 above. A counterexample is Proposition 4 of Kume & Wood (2005), in which a central limit theorem analogous to (26) holds but the limiting relative error is nonzero. In saddlepoint terminology, the reason that the limiting relative error is zero in the present setting is that the relevant tilted distribution, standardized to have mean zero and covariance matrix the identity, converges as $\psi \rightarrow \infty$ to a multivariate Gaussian distribution, for which the saddlepoint approximation is exact.

5. A PARAMETRIC MODEL FOR RIGID BODY MOTION DATA

5.1. The model

In an interesting recent paper, Oualkacha & Rivest (2012) develop statistical methods for the analysis of rigid body motion data, focusing on the question of how to define an average rigid body motion. A rigid body motion in \mathbb{R}^d relative to a fixed coordinate system is defined by a pair (X, y) where X is a $d \times d$ rotation matrix that represents the change in orientation of the object, while $y \in \mathbb{R}^d$ is a vector that represents the translation of the object. An issue not considered by Oualkacha & Rivest (2012) is the development of parametric models that incorporate dependence between X and y . One possibility, which so far as we are aware has not been considered before, is to consider an exponential family model of the form

$$f(X, y) \propto \exp \left\{ \text{tr}(\tilde{A}^T X) + x^T \tilde{B}x + y^T \tilde{C}x - (y - \mu)^T V^{-1}(y - \mu)/2 \right\}, \tag{27}$$

where $x = \text{vec}(X)$. Dependence between X and y is governed by the cross-product term $y^T \tilde{C}x$. Absorbing this cross-product term into the quadratic form in y by completing the square, and writing

$$A = \tilde{A} + \text{vec}_d^{-1}(\tilde{C}^T \mu), \quad B = \tilde{B} + \tilde{C}^T V \tilde{C}/2, \quad C = V \tilde{C},$$

where vec_d^{-1} , the inverse of the vec operator, maps a $d^2 \times 1$ vector argument onto a $d \times d$ matrix, we obtain

$$f(X, y) \propto \exp \left\{ \text{tr}(A^T X) + x^T Bx - (y - \mu - Cx)^T V^{-1}(y - \mu - Cx)/2 \right\}. \tag{28}$$

From (28), $y | X \sim N_d(\mu + Cx, V)$ and the marginal density function of X is given by $f_2(X; A, B)$ in (3). Moreover, from (27), $X | y \sim f_2(X; \tilde{A} + \text{vec}_d^{-1}(C^T y), \tilde{B})$ and the marginal density function of y is proportional to

$$C_2 \{ \tilde{A} + \text{vec}_d^{-1}(C^T y), \tilde{B} \} \exp \left\{ -(y - \mu)^T V^{-1}(y - \mu)/2 \right\},$$

where C_2 is the normalizing constant given in (3).

A convenient estimation approach is to estimate A and B using the marginal likelihood based on f_2 in (3), with $C_2(A, B)$ approximated using one of the saddlepoint approximations described

in this paper; and to estimate C , μ and V using multivariate linear regression based on the conditional likelihood $y | X \sim N_d(\mu + Cx, V)$.

5.2. A real data example

The model defined by (28) contains many parameters, some of which will be difficult to identify, particularly in high-concentration settings. It is therefore of interest to identify and test suitable submodels of (28). Our purpose here is to fit various submodels of (28) to the ankle-joint data from Rivest et al. (2008) and Oualkacha & Rivest (2012), which consist of a sample $\{(y_1, X_1), \dots, (y_{50}, X_{50})\}$. For simplicity, we treat the (y_i, X_i) as independent and identically distributed. We first consider the following models related to the Gaussian conditional likelihood $y | X$: $M_1, C = 0$; $M_2, \text{rank}(C) = 1$; M_3, C free. In each of the models M_1, M_2 and M_3 , μ and the symmetric matrix V are free. For M_2 , the unit-rank constraint can be imposed by setting $C = \beta\gamma^T$ where β is a d -vector and γ is a d^2 -vector, and the model can be estimated by iterating between the following three steps until convergence: fixing γ, Σ and estimating μ, β ; fixing μ, β, Σ and estimating γ ; and fixing μ, γ, β and estimating Σ . The Wilks statistic, i.e., twice the difference of maximized loglikelihoods, is used to compare nested models, and is referred to the χ^2 distribution with degrees of freedom given by the difference in the number of free parameters in each model. Fitting the models M_1, M_2, M_3 to the ankle-joint data from Rivest et al. (2008) and Oualkacha & Rivest (2012), the Wilks statistic for testing M_1 against M_2 is 335 on 11 degrees of freedom, and for testing M_2 against M_3 , it is 439 on 16 degrees for freedom. Consequently there is strong evidence for rejecting the simpler models M_1 and M_2 and preferring the more general model M_3 , suggesting that there is dependence between the y component and the X component that is not simply characterized in terms of a few parameters.

Fitting marginal models for X based on (3) is now considered. Let $B_0 = \hat{C}^T \hat{V}^{-1} \hat{C}$, where \hat{C} and \hat{V} are the estimates from the fitted model M_3 , and consider the models $M_4, A = 0, B = \lambda B_0$; M_5, A free, $B = 0$; M_6, A free, $B = \lambda B_0$, where λ is a free scalar parameter. Neither M_4 or M_5 is a submodel of the other, but each is a submodel of M_6 . For these models we find the Wilks statistic for testing M_4 against M_6 is 788 on 9 degrees of freedom, while for testing M_5 against M_6 it is 128 on one degree of freedom. Hence there is strong evidence in favour of M_6 over each of M_4 and M_5 , and in particular the Fisher matrix model M_5 for X provides an inferior fit.

The data of Rivest et al. (2008) are highly concentrated, yet M_6 is favoured strongly over the Fisher matrix model, M_5 , even though plots do not suggest any obvious inadequacy in the Gaussian tangent-space approximation. To understand this finding, consider the singular value decomposition $A = Q\Omega R^T$ and write $Y = Q^T X R = \exp(H)$, where $H = (h_{ij})$ is $d \times d$ skew-symmetric matrix. Under high concentration, $\exp(H) \approx I + H + H^2/2$ and $H \approx (Y - Y^T)/2$. Then

$$\exp\{\text{tr}(A^T X)\} \approx \exp \left[\text{tr} \left\{ \Omega \left(I + H + \frac{1}{2} H^2 \right) \right\} \right] \propto \exp \left\{ -\frac{1}{2} \sum_{i < j} h_{ij}^2 (\omega_i + \omega_j) \right\}.$$

It follows that, for large $\omega_1, \dots, \omega_d$, the elements h_{ij} , for $i < j$, have approximately independent Gaussian distributions with mean zero and variance $\sigma_{ij}^2 = (\omega_i + \omega_j)^{-1}$. Conversely, however, the h_{ij} being asymptotically independent and Gaussian does not necessarily lead to a well-matched Fisher matrix distribution for X : for the case $d = 3$, $\omega_1 = (\sigma_{12}^{-2} + \sigma_{13}^{-2} - \sigma_{23}^{-2})/2$, $\omega_2 = (\sigma_{12}^{-2} + \sigma_{23}^{-2} - \sigma_{13}^{-2})/2$, $\omega_3 = (\sigma_{13}^{-2} + \sigma_{23}^{-2} - \sigma_{12}^{-2})/2$, and these expressions impose constraints on the σ_{ij} in order that $\omega_i \geq 0$. If some of the σ_{ij}^2 are sufficiently different from the others, then the nonnegativity constraint $\omega_i \geq 0$ may not be achieved by all of the ω_i . For the data of Rivest et al.

Table 1. Numerical results showing the performance of the normalizing constant approximations for (i) a density with respect to Lebesgue measure on \mathcal{S}_1^2 given by (2) with $a = (\omega_1, 0, \omega_2, 0)^T$ and B having all elements zero except $b_{24} = \omega_3$; (ii) a density on \mathcal{S}_2^2 given by (2) with $a = (0, 0, \omega_1, 0, 0, \omega_2)^T$ and B having all elements zero except $b_{11} = -b_{22} = \omega_3$, $b_{44} = -b_{55} = -\omega_4$, $b_{14} = b_{25} = \omega_5$, $b_{15} = b_{24} = \omega_6$; and (iii) the density on $\mathcal{V}_{3,3}$ defined by (18). To investigate performance at various concentrations, for each of (i)–(iii) we considered three values of the parameter $\Omega = (\omega_i)$. The parameter values used were $\Omega_1 = (10, 5, 1)$, $\Omega_2 = (15, 5, 4, 3, 2, 1)$, $\Omega_3 = (30, 10, 6)$.

Density	Support	Ω	\tilde{Y}_1	\tilde{W}_1	\tilde{n}_1	\tilde{Y}_3	\tilde{W}_3	\tilde{n}_3
(i)	\mathcal{S}_1^2	Ω_1	-9.5×10^{-2}	4.9×10^{-3}	1.6×10^3	-6.6×10^{-3}	1.3×10^{-4}	6.1×10^4
		$\Omega_1/2$	-1.5×10^{-1}	4.2×10^{-3}	1.9×10^3	-1.2×10^{-2}	1.3×10^{-4}	5.9×10^4
		$\Omega_1/10$	-1.7×10^{-1}	3.7×10^{-4}	2.1×10^4	5.5×10^{-3}	5.8×10^{-5}	1.4×10^5
(ii)	\mathcal{S}_2^2	Ω_2	-1.2×10^{-2}	1.5×10^{-2}	8.5×10^2	-1.3×10^{-2}	3.2×10^{-2}	4.0×10^2
		$\Omega_2/3$	-9.2×10^{-2}	1.1×10^{-2}	1.2×10^3	-1.3×10^{-2}	1.8×10^{-3}	6.9×10^3
		$\Omega_2/15$	-1.1×10^{-1}	5.3×10^{-3}	2.4×10^3	3.0×10^{-4}	4.3×10^{-4}	2.9×10^4
(iii)	$\mathcal{V}_{3,3}$	Ω_3	-5.1×10^{-1}	1.1×10^{-1}	7.4×10^1	-5.7×10^{-2}	2.2×10^{-2}	3.6×10^2
		$\Omega_3/5$	-1.5	9.4×10^{-2}	8.3×10^1	-4.5×10^{-1}	1.8×10^{-2}	4.3×10^2
		$\Omega_3/30$	-1.6	4.5×10^{-4}	1.7×10^4	-5.0×10^{-1}	4.0×10^{-4}	2.0×10^4

(2008), the sample analogues of the σ_{ij}^2 are given by $\hat{\sigma}_{12}^2 = 0.0008$, $\hat{\sigma}_{13}^2 = 0.0352$, $\hat{\sigma}_{23}^2 = 0.0708$, which corresponds to a negative $\hat{\omega}_3$. This explains why the Fisher matrix model does not produce a good fit here, and it is interesting that the introduction of the one-parameter quadratic term in M_6 produces a much improved fit. More generally, further work is needed to identify useful subfamilies of (3) with nonzero B .

6. NUMERICAL ACCURACY

We present results on the numerical accuracy of the normalizing constant approximations for two particular Fisher–Bingham distributions on (i) \mathcal{S}_1^2 and (ii) \mathcal{S}_2^2 , and for (iii) the Fisher matrix distribution on $\mathcal{V}_{3,3}$. Table 1 shows results in terms of $Y_k = \log(C) - \log(\hat{C}_k)$, where C is the true normalizing constant and \hat{C}_k the corresponding saddlepoint approximation; a Wilks statistic, W_k , defined below; and the sample size, n_k , necessary for a likelihood ratio test to detect a difference between the true maximum likelihood estimate and an estimate that comes from maximizing an approximate likelihood based on \hat{C}_k .

The loglikelihood corresponding to (2) or (3) is $\ell(\Omega) = -\log C(\Omega) + \Omega^T t$, where $\Omega = \{a^T, \text{vec}(B)^T\}^T$, $t = \{E(x)^T, E \text{vec}(xx^T)\}^T$ in the case of (2), and $\Omega = \{\text{vec}(A)^T, \text{vec}(B)^T\}^T$, $t = [E \text{vec}(X)^T, E \text{vec}\{\text{vec}(X)\text{vec}(X^T)\}]^T$ in the case of (3), where E denotes expectation with respect to the true density. An approximate loglikelihood, based on \hat{C}_k in place of C , is

$$\hat{\ell}_k(\Omega) = -\log \hat{C}_k(\Omega) + \Omega^T t. \tag{29}$$

Denote $\Omega_{\text{MLE}} = \text{argmax} \ell(\Omega)$ and $\hat{\Omega}_{\text{MLE},k} = \text{argmax} \hat{\ell}_k(\Omega)$, and define $W_k = 2|\ell(\Omega_{\text{MLE}}) - \ell(\hat{\Omega}_{\text{MLE},k})|$. Then the sample size necessary for a difference between Ω_{MLE} and $\hat{\Omega}_{\text{MLE},k}$ to be detected at the 5% level by a large-sample likelihood ratio test is $n_k = \chi_p^2(0.95)/W_k$, where $\chi_p^2(0.95)$ is the 95% quantile of a chi-squared distribution with degrees of freedom equal to p , the number of parameters in Ω .

In the absence of a convenient method for calculating $\mathcal{C}(\Omega)$ and t , we estimated them using 10^7 Monte Carlo samples. We substituted the estimate of t into (29) and used an optimization routine to calculate a Monte Carlo estimate, $\hat{\Omega}_{\text{MLE},k}^*$, of $\hat{\Omega}_{\text{MLE},k}$, then used a further 10^7 Monte Carlo samples to estimate $\mathcal{C}(\hat{\Omega}_{\text{MLE},k}^*)$, $l(\hat{\Omega}_{\text{MLE},k}^*)$ and hence W_k . To generate the data for each row of Table 1, we repeated the foregoing procedure 100 times, each run yielding a Monte Carlo estimate of Y_k and W_k , over which we averaged to obtain the \tilde{Y}_k and \tilde{W}_k shown in Table 1. The relative standard errors, not shown, are negligible for \tilde{Y}_k and in all cases less than ten percent for \tilde{W}_k . We estimated the critical sample size for the large-sample test by $\tilde{n}_k = \chi_p^2(0.95)/\tilde{W}_k$.

The results for the first-order approximation, $k = 1$, indicate good accuracy, but the second-order, $k = 3$, approximation is typically a major improvement, in some cases having values of \tilde{Y}_k and \tilde{W}_k an order of magnitude smaller than for $k = 1$. Further results given in the Supplementary Material show there is little to choose between $k = 2$ and $k = 3$, although the latter is slightly superior. The results show that in most cases considered, which span a range of concentrations, a large sample size would be necessary to detect a significant difference, at the 95% level, between the exact and approximate maximum likelihood estimates. The dimensions of the underlying saddlepoint approximation for distributions (i)–(iii) are 2, 2, 6, respectively. Even in the more challenging higher-dimensional setting, performance is strong. However, the computational cost of the second-order approximations grows rapidly with p ; typical timings to compute a second-order approximation of \mathcal{C}_1 for the product of $p = 2, 4, 6, 8, 10$ circles are 0.2, 1, 5, 21, 70 seconds.

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online contains formulae for the cumulant derivatives in the product-of-spheres and Stiefel manifold cases, formulae for $\hat{\rho}_{13}^2$, $\hat{\rho}_{23}^2$ and $\hat{\rho}_4$ in the matrix Fisher case, proofs of Propositions 1–5, numerical results for the $k = 2$ second-order approximations, and Matlab code for performing calculations in this paper.

ACKNOWLEDGEMENT

We are grateful to the editor and three reviewers for constructive comments, and to Professor Louis-Paul Rivest for sending us the rigid body motion dataset and for permission to use it.

REFERENCES

- ARNOLD, R. & JUPP, P. E. (2013). Statistics of orthogonal axial frames. *Biometrika* **100**, 571–86.
- BINGHAM, C. (1974). An antipodally symmetric distribution on the sphere. *Ann. Statist.* **2**, 1201–5.
- BUTLER, R. W. (2007). *Saddlepoint Approximations*. Cambridge: Cambridge University Press.
- BUTLER, R. W. & WOOD, A. T. A. (2003). Laplace approximations for Bessel functions of matrix argument. *J. Comp. Appl. Math.* **155**, 359–82.
- DOWNES, T. D. (1972). Orientation statistics. *Biometrika* **59**, 665–76.
- HOFF, P. (2009). Simulation of the matrix Bingham–von Mises–Fisher distribution, with applications to multivariate and relational data. *J. Comp. Graph. Statist.* **18**, 438–56.
- JUPP, P. E. & MARDIA, K. V. (1979). Maximum likelihood estimators for the matrix von Mises–Fisher and Bingham distributions. *Ann. Statist.* **7**, 599–606.
- JUPP, P. E. & MARDIA, K. V. (1980). A general correlation coefficient for directional data and related regression problems. *Biometrika* **67**, 163–73.
- KENT, J. T. (1982). The Fisher–Bingham distribution on the sphere. *J. R. Statist. Soc. B* **44**, 285–99.
- KENT, J. T. (1994). The complex Bingham distribution and shape analysis. *J. R. Statist. Soc. B* **56**, 285–99.
- KHATRI, C. G. & MARDIA, K. V. (1977). The von Mises–Fisher matrix distribution in orientation statistics. *J. R. Statist. Soc. B* **39**, 95–106.

- KUME, A. & WOOD, A. T. A. (2005). Saddlepoint approximations for the Bingham and Fisher–Bingham normalising constants. *Biometrika* **92**, 465–76.
- MARDIA, K. V. (1975). Statistics of directional data (with Discussion). *J. R. Statist. Soc. B* **37**, 349–93.
- MARDIA, K. V., HUGHES, G., TAYLOR, C. C. & SINGH, H. (2008). A multivariate von Mises distribution with applications to bioinformatics. *Can. J. Statist.* **36**, 99–109.
- MARDIA, K. V. & JUPP, P. E. (2000). *Directional Statistics*. Chichester: John Wiley.
- MUIRHEAD, R. J. (1982). *Aspects of Multivariate Statistical Theory*. New York: John Wiley.
- OUALKACHA, K. & RIVEST, L.-P. (2012). On the estimation of an average rigid body motion. *Biometrika* **99**, 585–98.
- PRENTICE, M. J. (1986). Orientation statistics without parametric assumptions. *J. R. Statist. Soc. B* **48**, 214–22.
- RIVEST, L.-P., BAILLARGEON, S. & PIERRYNOWSKI, M. R. (2008). A directional model for the estimation of the rotation axes of the ankle joint. *J. Am. Statist. Assoc.* **103**, 1060–9.
- SINGH, H., HNIZDO, V. & DEMCHUCK, E. (2002). Probabilistic model for two dependent circular variables. *Biometrika* **89**, 719–23.
- WOOD, A. T. A. (1993). Estimation of the concentration parameters of the Fisher matrix distribution on $SO(3)$ and the Bingham distribution on S_q , $q \geq 2$. *Aust. J. Statist.* **35**, 69–79.

[Received December 2011. Revised February 2013]