

Research Master's programme: METHODOLOGY AND STATISTICS
OF BEHAVIORAL AND SOCIAL SCIENCES
Utrecht University, the Netherlands

MSc Thesis *Kees Mulder (3270882)*

TITLE: Extending Bayesian analysis of circular data to comparison of
multiple groups

May 2014

Supervisors:

Dr. *I.G. Klugkist*

Preferred journal of publication: BAYESIAN ANALYSIS

Word count: 4701

Extending Bayesian analysis of circular data to comparison of multiple groups

Kees Mulder * and Irene Klugkist †

Abstract. Circular data are data measured in angles and occur in a variety of scientific disciplines. Bayesian methods promise to allow for flexible analysis of circular data, for which few methods are available. Three existing MCMC methods (Gibbs, Metropolis-Hastings, and Rejection) for a single group of circular data were extended to be used in a between-subjects design, providing a novel procedure to compare groups of circular data. Investigating the performance of the methods by simulation study, all methods were found to overestimate the concentration parameter of the posterior, while coverage was reasonable. The rejection sampler performed best. In future research, the MCMC method may be extended to include covariates, or a within-subjects design.

Keywords: circular data, Bayesian inference, mcmc methods, gibbs, metropolis-hastings, rejection sampler

1 Introduction

Circular data are data measured in angles or orientations in two-dimensional space. For example, one may imagine directions on a compass ($0^\circ - 360^\circ$), times of the day ($0 - 24$ hours), or directions on a circumplex model, such as Leary's Circle (Leary 1957). Circular data are frequently encountered in many scientific disciplines, such as biology, social sciences, meteorology, astronomy, earth sciences, and medicine.

The analysis of circular data requires special directional statistical methods due to the periodicity of the sample space. For example, two angles of 10° and 350° differ by only 20° , while if treated linearly the distance between them would seem to be 340° . A similar mismatch occurs for the arithmetic mean of 10° and 350° , which is 180° , while their correct circular mean is 0° .

Three different approaches for analysis of circular data are discussed in the literature: the *intrinsic* approach, which uses the von Mises distribution (Von Mises 1918; Damien and Walker 1999); the *embedding* approach, which uses the Projected Normal distribution (Nunez-Antonio and Gutierrez-Pena 2005); and the *wrapping* approach, where distributions on the real line are wrapped around the circle (Ferrari 2009). Here, the scope is limited to the intrinsic approach.

Due to the difficulty of working with a circular sample space, few methods have

*Utrecht University, Department of Methodology and Statistics, Utrecht, The Netherlands, k.t.mulder@uu.nl

†Utrecht University, Department of Methodology and Statistics, Utrecht, The Netherlands, i.klugkist@uu.nl

been developed in the field of analysis of circular data. An overview of available frequentist methods for analysis of circular data can be found in [Fisher \(1995\)](#) and [Mardia and Jupp \(1999\)](#). Bayesian methods offer a promising new approach not only in the field of statistics at large, but also specifically in the analysis of circular data. Main advantages of the Bayesian approach are the flexibility of Markov chain Monte Carlo (MCMC) methods used in Bayesian analysis, the lack of asymptotic assumptions, and the possibility to incorporate knowledge from previous research. Some work has been done performing Bayesian estimation on circular data without utilising MCMC methods ([Dowe et al. 1996](#)), but such methods only perform point estimation without providing standard errors, while researchers are often interested in drawing inference.

In the case of directional statistics, MCMC methods may prove to be a flexible solution to the difficulty of drawing inference from circular data. A limited number of MCMC methods for circular data have been developed. Available methods generally employ the von Mises distribution, which is the natural analogue of the normal distribution on the circle. Early work by [Damien and Walker \(1999\)](#) provided a Gibbs sampler for a single group by adding latent variables to the model. Metropolis-Hastings algorithms have been developed for circular distributions in general ([Bhattacharya and Sengupta 2009](#)) and for the von Mises-Fisher distribution, which is the generalization of the von Mises distribution to the sphere ([Nunez-Antonio and Gutiérrez-Pena 2005](#)). Recent work has attempted to tune the parameters of a rejection sampling algorithm in order to obtain a computationally fast method to sample from the posterior of a von Mises distribution ([Forbes and Mardia 2014](#)). Although different in approach, these methods have in common that they draw from the posterior of the von Mises distribution given one group of circular data, which can be used to describe properties of a single sample. None of the methods may be used to compare groups.

In this paper existing MCMC methods will be extended to analyse data from between-subjects designs, where the research goal is to compare mean directions of multiple groups on a circular outcome. Many tests in between-subjects designs, such as ANOVA, assume equal variance across groups. Circular ANOVA methods that have been developed in a frequentist framework also carry this assumption ([Harrison and Kanji 1988](#); [Harrison et al. 1986](#)). A main aim of this paper is thus to extend available MCMC methods to between-subjects designs, so that the method samples multiple mean directions and a single measure of dispersion. Then, the performance of these methods will be assessed to decide which is the most commendable.

Section 2 provides the theoretical framework and notation for the von Mises distribution. Then, in Section 3, three MCMC methods are discussed and extended to between-subjects designs. These are compared by means of a simulation study in Section 4. Concluding remarks will be made in Section 5.

2 The intrinsic approach

The MCMC methods discussed in this paper all fall within the intrinsic approach, where it is assumed that the data follow the von Mises distribution. This section will discuss

basic properties of the von Mises distribution and provide a framework for the MCMC methods that will be discussed in Section 3. The first four sections will be restricted to the von Mises distribution for a single group, while Section 2.5 will introduce properties and notation to be used in the case with multiple groups.

2.1 Von Mises distribution

The von Mises distribution is a symmetric unimodal distribution, which is given by

$$\text{VM}(\theta|\mu, \kappa) = \{2\pi I_0(\kappa)\}^{-1} \exp\{\kappa \cos(\theta - \mu)\}, \quad 0 \leq \theta < 2\pi, \kappa \geq 0$$

where θ represents the data, μ represents the mean direction, κ is the concentration parameter, and $I_0(\cdot)$ is the modified Bessel function of order 0 (Abramowitz and Stegun 1972). A higher κ represents less variation, and thus more concentrated data. Let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ be a sample of angular measurements θ_i of size n .

Each angle in the dataset may be viewed as a vector of length 1 in direction θ_i . As illustrated in Figure 1, the summation of these vectors results in a vector in direction $\bar{\theta}$ of length R . $\bar{\theta}$ is an unbiased estimator of μ , while R is called the resultant length and may be obtained from

$$R = \sqrt{\left(\sum_{i=1}^n \cos \theta_i\right)^2 + \left(\sum_{i=1}^n \sin \theta_i\right)^2},$$

which increases with concentration and sample size. The mean resultant length can be computed as $\bar{R} = R/n$, which is a metric of concentration independent of the sample size.

2.2 Prior distribution

Guttorp and Lockhart (1988) present a conjugate prior for the von Mises distribution. It is given up to a constant of proportionality by

$$p(\mu, \kappa) \propto I_0(\kappa)^{-c} \exp\{R_0 \kappa \cos(\mu - \mu_0)\},$$

which represents c observations with prior mean direction μ_0 and prior resultant length R_0 . In all methods applied in this paper, this conjugate prior will be used.

2.3 Posterior distribution

To obtain the posterior distribution, the data and the prior are combined to obtain the posterior mean μ_n ¹ and the posterior resultant length R_n by

$$C_n = R_0 \cos \mu_0 + \sum_{i=1}^n \cos \theta_i, \quad S_n = R_0 \sin \mu_0 + \sum_{i=1}^n \sin \theta_i,$$

¹In R (R Development Core Team 2013), calculation of μ_n is readily available in `atan2(S_n, C_n)`.




Figure 1: Illustration of the mean direction and resultant length of $\theta = \{56^\circ, 77^\circ, 344^\circ\}$. The summation of the vectors results in a vector of length R in direction $\bar{\theta}$.

$$\mu_n = \begin{cases} \tan^{-1}(S_n/C_n) & \text{if } C_n > 0, S_n > 0 \\ \tan^{-1}(S_n/C_n) + \pi & \text{if } C_n < 0 \\ \tan^{-1}(S_n/C_n) + 2\pi & \text{if } C_n > 0, S_n < 0 \end{cases}$$

and

$$R_n = \sqrt{C_n^2 + S_n^2}.$$

Then, the joint posterior distribution is given up to a constant of proportionality, by

$$f(\mu, \kappa | \theta) \propto \{I_0(\kappa)\}^{-m} \exp\{R_n \kappa \cos(\mu - \mu_n)\},$$

where $m = n + c$. This distribution is not of closed form due to the Bessel function.

2.4 Conditional distributions

The MCMC methods presented in Section 3 are based upon the conditional posterior distributions $f(\mu | \kappa, \theta)$ and $f(\kappa | \mu, \theta)$. The conditional posterior distribution of μ , up to a constant of proportionality, is given by

$$f(\mu | \kappa, \theta) \propto \exp\{R_n \kappa \cos(\mu - \mu_n)\},$$

which is the kernel of a von Mises distribution with mean direction μ_n and concentration parameter $R_n \kappa$. Several straightforward methods to sample data from the von Mises distribution are available (Best and Fisher 1979; Fisher 1995).

The conditional distribution of $f(\kappa | \mu, \theta)$, is given by

$$f(\kappa | \mu, \theta) \propto \{I_0(\kappa)\}^{-m} \exp\{R_n \kappa \cos(\mu - \mu_n)\}.$$

However, it is not straightforward to sample from this conditional distribution, so that special methods are required. In Section 3, three methods that can sample the concentration parameter will be discussed.

2.5 Notation for multiple groups

Here, basic notation and properties will be defined that will be used to extend the methods discussed in Section 3 to multiple groups. Denote the groups by $j = 1, \dots, J$. Then, for group j , the posterior mean is denoted by μ_{nj} and the posterior resultant length by R_{nj} . The sample size of group j is denoted by n_j , which will be combined with the prior property c_j to obtain $m_j = n_j + c_j$. Finally, let

$$R_t = \sum_{j=1}^J R_{nj} \quad \text{and} \quad m_t = \sum_{j=1}^J m_j.$$

Utilising this notation, the posterior for multiple groups with a common κ is given by

$$f(\boldsymbol{\mu}, \kappa | \boldsymbol{\theta}) \propto \{I_0(\kappa)\}^{-m_t} \exp \left[\kappa \sum_{j=1}^J R_{nj} \cos(\mu_j - \mu_{nj}) \right],$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_J)$ denotes the mean directions of the groups.

3 Available methods and their extensions

In this section, three MCMC methods will be presented and extended to be able to sample from the posterior of a von Mises distribution in a between-subjects design with $J \geq 1$ independent groups with common but unknown κ . Importantly, all three methods use the conjugate prior as described in Section 2.2.

3.1 A Gibbs sampler using latent variables

In one of the earliest attempts at sampling the concentration parameter of the von Mises distribution, [Damien and Walker \(1999\)](#) provide a Gibbs sampler that only requires sampling of uniform random variates. It is an application of the procedure of adding latent variables to a posterior distribution in order to be able to apply Gibbs samplers in situations where this may not have been feasible originally ([Damien et al. 1999](#)).

Although the relative simplicity of the Gibbs sampler usually is appealing, it has been noted that this sampler shows high autocorrelation for more concentrated data, causing slow convergence ([Nunez-Antonio and Gutiérrez-Pena 2005](#), p. 990).

[Damien and Walker \(1999\)](#) add latent variables w, v, x , and $u = (u_1, u_2, \dots)$ to the joint posterior density $f(\boldsymbol{\mu}, \kappa | \boldsymbol{\theta})$, where u is an infinite set of latent variables. It is not necessary to sample an infinite number of values for u_k , as computing values for u_k

up to some sufficient k provides a good approximation of the correct solution. Let Z be the number of values of u_k that will be sampled, so that the set of sampled values is u_1, \dots, u_Z . For each analysis performed with this method, a value for Z must be chosen. This is a disadvantage of this method, because setting Z too high will prove computationally intensive, while setting Z too low produces biased results.

Another disadvantage is that this method requires setting starting values not only for μ and κ , but also for w . However, w does not have an intuitive interpretation, making the choice of a starting value somewhat arbitrary and possibly difficult.

Sampler for a single group

The posterior density of a single group, after inclusion of the latent variables, is given up to a constant of proportionality as

$$f(\mu, \kappa, w, v, u, x|\boldsymbol{\theta}) \propto e^{-R_n \kappa} I(v < e^{R_n \kappa \{1 + \cos(\mu - \mu_n)\}}, x < w^{m-1}) \times \left(e^{-w} \prod_{k=1}^{\infty} I(u_k < e^{-w \lambda_k \kappa^{2k}}) \right),$$

for which the marginal for (μ, κ) is $f(\mu, \kappa|\boldsymbol{\theta})$, as required. The Gibbs sampler works by drawing a value from the conditional distributions of x, v, μ, u_k, w and κ in sequential order, each conditional on the current other values. Further details, including the required conditional distributions, are found in [Damien and Walker \(1999\)](#) and will not be given here, as the Gibbs sampler for a single group is a special case of the Gibbs sampler described next with $J = 1$.

Sampler for multiple groups

This section will describe the adapted procedure to implement the Gibbs sampler for multiple groups, so that it will sample from the posterior density $f(\boldsymbol{\mu}, \kappa, w, v, u, x|\boldsymbol{\theta})$. It differs in two ways from the sampler provided in [Damien and Walker \(1999\)](#): first, means for multiple groups and a common κ are now sampled, and second, some steps were combined or simplified to facilitate implementation. Notably, the sampling of a set of values u_1, \dots, u_Z is rewritten to sample another set of values N_1, \dots, N_Z . The extended Gibbs sampler consists of the following 8 steps:

1. Set $\boldsymbol{\mu}, \kappa$, and w to their starting values.
2. Draw a random variate τ from $U(0, 1)$.
3. For each group j , draw a value for μ_j from $U(\mu_{nj} - \cos^{-1} g, \mu_{nj} + \cos^{-1} g)$, where

$$g = \max \left[-1, \frac{\ln \tau}{R_t \kappa} + \frac{\sum_{j=1}^J R_{nj} \{1 + \cos(\mu_j - \mu_{nj})\}}{R_t} - 1 \right].$$

4. Calculate $M = \tilde{w} + E$, where \tilde{w} is the current value of w and E is a random variate drawn from an exponential distribution with rate $I_0(\kappa) - 1$.
5. Draw a new value for w from $e^{-w}I(\tilde{w}r^{1/(m-1)} < w < M)$, where r is a uniform random variate from $U(0, 1)$.
6. Compute $N_k = \kappa(1 + F_k)^{1/(2k)}$, where F_k is an exponential r.v. with rate $\tilde{w}(k!)^{-2}(0.5\kappa)^{2k}$, and $k = 1, \dots, Z$. Set $N = \min N_k$. For advice on setting Z , see Section 3.1.
7. Draw a value for κ from $e^{-R_n\kappa}I(\max\{0, v_n\} < \kappa < N)$, where

$$v_n = \frac{\ln \tau}{\sum_{j=1}^J R_{nj} \{1 + \cos(\mu_j - \mu_{nj})\}} + \kappa.$$

8. Repeat steps 2 - 7 until a sufficient number of samples have been obtained.

Choosing Z

In step 6 of the procedure given above to draw from the conditional density of κ , a number of samples N_k are generated of which the smallest is retained. However, the number of N_k that should be sampled (here denoted by Z) was not discussed in [Damien and Walker \(1999\)](#). A small simulation study was performed to be able to give guidelines for setting Z when applying this algorithm.

For each combination of sample sizes $\{5, 30, 100\}$ and concentrations $\{0.1, 1, 4, 8\}$, 100 datasets with $J = 1$ and $J = 3$ were generated. The Gibbs sampler was then run for 10000 iterations with no lag and a burn-in of 500 on each dataset, with Z set to 85. In each iteration, the index number k of the selected (smallest) value for N_k was saved. This resulted in 100 chains (one for each dataset) of chosen index number k of 10000 iterations. Then, the overall maximum value of these chains was taken. The reasoning behind this is that if in all these iterations the chosen value never exceeds some number, setting Z to that number or slightly above it will ensure that Z is not too low to produce bias while still remaining somewhat computationally efficient.

From the results, given in [Table 1](#), it is apparent that a value for Z of about 25 should be sufficient in most applications, assuming the data is not strongly concentrated. The required Z decreases with higher sample sizes and less concentrated data.

3.2 A Metropolis-Hastings sampler

Another approach is to employ the Metropolis-Hastings (MH) method ([Metropolis et al. 1953](#); [Hastings 1970](#)) to sample from the posterior of a von Mises distribution. Usually, MH algorithms are slower and encounter more autocorrelation and convergence problems than Gibbs samplers. However, considering the complicated nature of adding latent variables in the Gibbs sampler described above, an MH method may be advantageous. Another advantage is that the algorithm is reasonably straightforward. On

Table 1: Maximum k that was picked out as the smallest value after 10000 iterations of the Gibbs sampler applied to 100 datasets for different sample sizes (n), concentration (κ) and number of groups (J).

κ	$J = 1$			$J = 3$		
	Sample size			Sample size		
	5	30	100	5	30	100
0.1	11	5	5	7	5	5
1	14	6	6	8	5	6
4	25	10	8	14	8	8
8	25	12	11	14	9	7

the other hand, this method depends on a proper choice for the proposal density, which may limit its use.

Sampler for a single group

To apply the sampler for a single group, samples are needed for a single μ and κ . The conditional distribution of the mean direction μ is known and is easy to sample from using a Gibbs step. The conditional distribution of κ is known but difficult to sample from, which will be solved by applying an MH step.

For the MH step, two main ingredients are required: the posterior from which samples are required, and a proposal density from which it is straightforward to sample. The conditional posterior $f(\kappa|\mu, \theta)$ is given in Section 2.4. As a proposal, which must be non-negative, the χ^2 -distribution will be used. The full algorithm will not be presented here as it is a special case with $J = 1$ of the sampler described next.

Sampler for multiple groups

The MH sampler for multiple groups may employ the posterior $f(\mu, \kappa|\theta)$ as given in Section 2.5. However, in order to prevent underflow issues, the natural logarithm of the posterior is used, which is

$$\ln f(\mu, \kappa|\theta) = -m_t [I_0(\kappa)] + \kappa \sum_{i=1}^n R_{nj} \cos(\mu_j - \mu_{nj}).$$

Let κ_{cur} be the current value of κ , and $\chi^2(x|h)$ be the chi-square distribution with h degrees of freedom. Then, the MH method is given by the following 7 steps:

1. Set κ_{cur} to its starting value.
2. For each group j , draw a value μ_j from $\text{VM}(\mu_j|\mu_{nj}, R_n \kappa_{cur})$.
3. Draw a candidate κ_{can} from $\chi^2(\kappa_{can}|\kappa_{cur})$.

4. Calculate the MH ratio as

$$a = \ln f(\kappa_{can}|\boldsymbol{\mu}, \boldsymbol{\theta}) + \ln \chi^2(\kappa_{cur}|\kappa_{can}) \\ - \ln f(\kappa_{cur}|\boldsymbol{\mu}, \boldsymbol{\theta}) - \ln \chi^2(\kappa_{can}|\kappa_{cur}),$$

where $\boldsymbol{\mu} = \mu_1, \dots, \mu_J$, a vector of current values of μ for each group.

5. Draw a value u from $U(0, 1)$.
6. If $a > \ln u$, set $\kappa_{cur} = \kappa_{can}$. Otherwise, remain at κ_{cur} .
7. Repeat step 2 - 6 until a sufficient number of samples have been obtained.

3.3 A rejection sampler

In a recent paper, [Forbes and Mardia \(2014\)](#) presented a promising new algorithm to sample from the conditional posterior $f(\kappa|\mu, \boldsymbol{\theta})$. The approach is largely focused on computational speed, and was motivated by the fact that plugging a Bessel function approximation into the von Mises posterior leads to a Gamma distribution.

First, the algorithm sets $\eta = n$ and computes $\beta_0 = -n^{-1} \sum_{i=1}^n \cos(\theta_i - \mu)$. These are then used to compute the approximately optimal parameters for a Gamma proposal, such that the probability of rejection is minimized. In the rejection step, a candidate for κ is then repeatedly drawn from this Gamma proposal density until it is accepted.

Samples of μ are drawn outside of the algorithm, which may be done easily as in step 2 of the MH procedure in Section 3.2. As with the MH method, only κ requires a starting value.

Sampler for a single group

In their paper, [Forbes and Mardia \(2014\)](#) describe the rejection sampler for a single group of data, using a constant prior. The conjugate prior that is preferred here can be added as described below.

Using the sample mean direction $\bar{\theta}$, it can be shown that

$$\beta_0 = -n^{-1} \sum_{i=1}^n \cos(\theta_i - \mu) = -\frac{R \cos(\mu - \bar{\theta})}{n}.$$

This relation to the resultant length means that μ_n, R_n , and m from the desired posterior can be plugged into the formula for β_0 , to obtain

$$\beta_n = -\frac{R_n \cos(\mu - \mu_n)}{m}.$$

Then, the rejection algorithm can be applied exactly as given in [Forbes and Mardia \(2014\)](#), using β_n instead of β_0 and $\eta = m$.

Sampler for multiple groups

As the sampling of means occurs outside of the main algorithm, it is straightforward to sample separate means for each group. However, the common κ depends on the sampled means through β_n . After computation of β_n , the rejection algorithm no longer uses the data θ or the current value of μ . The sampler will thus be extended to multiple groups by once again rewriting β_n .

Using R_{nj} and μ_{nj} , and m_t as before, let

$$\beta_t = -\frac{\sum_{j=1}^J R_{nj} \cos(\mu - \mu_{nj})}{m_t}.$$

Then, the rejection algorithm can be applied using β_t instead of β_0 and $\eta = m_t$.

4 Simulation study

In the previous section, three distinct methods to sample from the posterior of the von Mises distribution with multiple groups were shown. In this section, these three methods will be evaluated on their performance and efficiency.

4.1 Methods

All three methods were implemented in C++ within R ([R Development Core Team 2013](#)) via Rcpp ([Eddelbuettel and François 2011](#)). To illustrate the differences between the three methods, [Figure 2](#) shows example chains of the first 500 iterations for each of the three methods. It can be seen that the Gibbs sampler has large autocorrelation and slow convergence, that the MH algorithm can have low acceptance probability but converges fast, and that the rejection algorithm converges fast and mixes well.

The three sampling methods were applied to various scenarios, which differed in the following three properties. First, the samplers analyzed both a single group of data ($J = 1$) and three groups of data ($J = 3$). Second, sample sizes of 5, 30 and 100 were used. For $J = 3$, this sample size denotes the sample size per group (n_j), making the total sample size $3n_j$. Third, values for the concentration parameter κ were 0.1, 4 and 32. Because the multiple groups are assumed have equal κ , all three groups of data were sampled given the same true κ . These manipulations resulted in a 3x2x3x3 simulation study design, for a total of 54 cells. For $J = 1$, the true mean was set at 20° , while true means for $J = 3$ were set at 20° , 40° , and 60° .

For each cell, a thousand datasets were generated, each of which was analyzed with each sampler. Burn-in and lag (that is, how much the chain will be thinned) were set to appropriate values (see [Section 4.2](#)), after which the first 10000 retained iterations of both μ and κ were saved. Although all three methods allow inclusion of prior information, a non-informative prior was used throughout the simulation study by setting $\mu_0 = 0$, $R_0 = 0$, and $c = 0$. Each method requires a starting value for κ , which was set

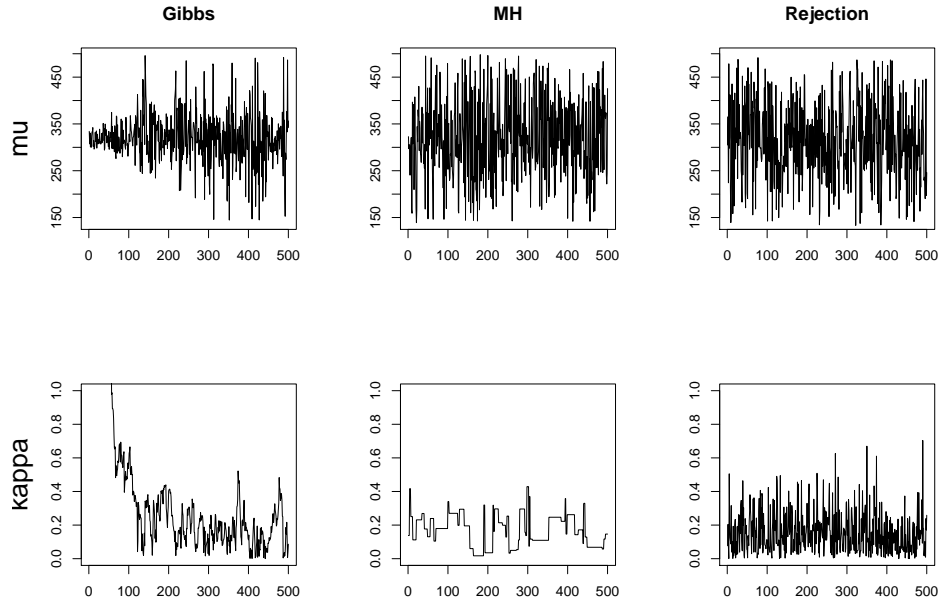


Figure 2: Example chains of μ and κ drawn in the first 500 iterations of each of the three methods, with no burn-in and without thinning the chain, where $J = 3$, true $\kappa = 0.1$, and $n_j = 30$.

to 2 in all cases. The Gibbs sampler required additional starting values for μ and w , which were set at 0 and 4 respectively, regardless of sample size or κ . Additionally, for the Gibbs sampler an appropriate Z must be chosen (see Section 3.1), which was set to 85 throughout this study.

4.2 Convergence

As convergence is achieved at a different number of iterations for each of the methods, several runs of each were assessed for each cell in order to assess convergence and required burn-in and lag, which were then set correctly for the simulation study.

The Gibbs sampler performed adequately for small samples with large dispersion. For example, a single group of 5 datapoints with true $\kappa = 0.1$ produced a reasonable sample from the posterior using a lag of 2, which means saving every other iteration. With larger sample sizes and more concentrated data, the autocorrelation increases quickly, requiring a lag of 250 for $\kappa = 4$ and $n_j = 100$ with $J = 3$. For values of κ above about 7, application of the Gibbs sampler simply becomes unfeasible, so results for the Gibbs sampling method with true $\kappa = 32$ are not reported.

The MH algorithm fared much better, converging quickly in all tested situations. However, applying MH methods requires reasonable acceptance rates, which can be computed by Q_{acc}/Q , where Q_{acc} is the number of accepted iterations, and Q is the total number of iterations. [Johnson and Albert \(1999\)](#) suggest an acceptance rate of about 50% to be ideal. A low acceptance rate may suggest a badly fitting proposal density, while a high acceptance rate (ie. close to 1) may suggest that the algorithm has yet to converge properly. As convergence was assessed separately and achieved quite quickly, only too low acceptance rates were of concern here. Acceptance rates were lower for larger sample sizes and less concentrated data. For example, for a single group of data with $n = 100$ and $\kappa = .1$ the average acceptance rate was .1, while for $n = 30, \kappa = 32$ the acceptance rate was as high as .71. One exception to this pattern was found: the smaller sample size for $n = 5, \kappa = 32$ should mean this cell would have an even higher acceptance rate than .71, but it was in fact lower at .55. The small sample size causes large differences in the concentration between datasets, so that this cell contains some extremely concentrated datasets. Acceptance rates of the MH algorithm were very low for these extreme concentrations.

The rejection algorithm converged immediately and showed almost no autocorrelation. An acceptance rate can be computed by Q/Q_{can} , where Q is the total number of accepted candidates, which was chosen beforehand as the desired number of iterations, and Q_{can} is the total number of candidates, including those that were accepted. The algorithm rejected no more than 5% of the candidates in any case, which is even closer to 1 than the expected acceptance probabilities as given in [Forbes and Mardia \(2014\)](#).

4.3 Mode estimation for κ

Estimating κ as the mean or the median of the posterior sample may lead to biased results, as κ is non-negative and has a right-skewed distribution. For skewed distributions, the mode usually provides the least biased estimate. An estimate of the mode can be obtained by using the Highest Density Interval (HDI), which is the shortest interval containing a certain percentage of the data ([Venter 1967](#)). Here, the mode was estimated to be the midpoint of the 10% HDI.

4.4 Results

In [Table 2](#) and [3](#), results are displayed for a single group and three groups, respectively. As mentioned before, applying the Gibbs sampler to a situation with $\kappa = 32$ is unfeasible, and therefore these rows are left empty.

The column below posterior μ mean gives the average of the posterior mean of μ or $\{\mu_1, \mu_2, \mu_3\}$ of all replications. The coverage of the mean denotes the proportion of replications where 95% Central Credible Interval (CCI) contained the true μ . For $J = 3$, this coverage was averaged over the three means. The desired value of the coverage is .95. For the posterior κ , the estimated mode for each replication was saved, as well as the 95% HDI. The average of the mode over replications is provided in the column

Table 2: Average of posterior properties over 1000 replications for $J = 1$, with true $\mu = 20^\circ$, for different sample sizes (n) and concentration κ .

n	κ	Method	Posterior μ		Posterior κ ^a		Acc. ^b	MCT ^c
			Mean	Coverage	Mode	Coverage		
5	0.1	Gibbs	28.51	0.78	0.67	0.96	1	0.78
		MH	28.34	0.78	0.63	0.96	0.33	0.04
		Rejection	27.89	0.77	0.64	0.97	0.96	0.03
	4	Gibbs	19.04	0.94	5.11	0.96	1	7.60
		MH	19.01	0.90	9.83	0.93	0.63	0.04
		Rejection	19.02	0.91	7.65	0.95	1	0.04
	32	Gibbs	—	—	—	—	—	—
		MH	20.07	0.80	133.12	0.64	0.55	0.05
		Rejection	20.07	0.88	55.74	0.96	1	0.05
30	0.1	Gibbs	28.14	0.87	0.19	0.97	1	1.13
		MH	28.68	0.88	0.18	0.97	0.16	0.07
		Rejection	28.50	0.86	0.19	0.98	0.98	0.04
	4	Gibbs	19.96	0.94	4.25	0.92	1	9.91
		MH	19.96	0.94	4.26	0.94	0.36	0.08
		Rejection	19.96	0.94	4.17	0.96	1	0.04
	32	Gibbs	—	—	—	—	—	—
		MH	20.04	0.93	34.71	0.95	0.71	0.08
		Rejection	20.05	0.93	34.66	0.95	1	0.04
100	0.1	Gibbs	19.50	0.92	0.12	0.98	1	3.94
		MH	19.76	0.92	0.11	0.98	0.10	0.17
		Rejection	19.57	0.92	0.12	0.98	0.98	0.03
	4	Gibbs	20.02	0.95	4.10	0.94	1	59.03
		MH	20.02	0.95	4.11	0.95	0.21	0.17
		Rejection	20.02	0.95	4.04	0.96	1	0.03
	32	Gibbs	—	—	—	—	—	—
		MH	19.97	0.95	32.87	0.92	0.53	0.16
		Rejection	19.97	0.95	32.81	0.92	1	0.04

^aPosterior κ mode denotes the mode as described in section 4.3. Coverage denotes the proportion of replications in which the true κ fell within the 95 % HDI.

^bAcceptance ratio. For Gibbs sampling, this is always 1. For Metropolis-Hastings, Q_{acc}/Q is given. For the rejection method, this is Q/Q_{can} .

^cMean Computation Time of one replication in seconds.

posterior κ mode, followed by the posterior κ coverage, which denotes the proportion of replications for which the true value fell within the 95% HDI. The last two columns provide the acceptance rate and the mean computation time (MCT) per replication in seconds.

The true value of a parameter often influences the size of the bias. In order to investigate by what factor estimates are off, the relative bias can be calculated as *Bias/True value*. The relative bias may help facilitate interpretation of relative severity of bias for κ , in order to allow for more accurate comparisons between cells.

A single group

Posterior μ

All three methods provided similar results for the posterior mean, which was generally close to the true mean. Estimates were closer to the true value for increasing n and increasing κ . The worst case was found for $\kappa = 0.1, n = 5$, where the difference between the true μ (20°) and average posterior μ was as high as 8° . However, this difference is most likely due to sampling error of datasets instead of an issue with the MCMC methods. When $\kappa = 0.1$, the distribution of the sample mean direction $\bar{\theta}$ is close to the circular uniform distribution, so that the average over the sample mean directions shows some random variation. This is supported by the fact that the MCMC methods all show the same difference from the true value. In general, there seems to be no systematic bias in the estimation of the mean direction.

Coverage was generally adequate as well. Coverage for μ relies on correct procedures of sampling both μ and κ . For example, an upwards bias in κ results in a lower coverage for μ . Sampling methods for the distribution of μ are well-known and as such most likely correct, and thus deviations of the coverage from .95 are likely due to a deficient mechanism to sample κ , as the current value of κ is used in the distribution of μ .

Posterior κ

The mode of κ shows a systematic upward bias for all cells and all methods. The relative bias is worse for smaller κ and thus more dispersed data. The bias also decreases with increasing n and nearly disappears for $n = 100$.

Regardless of the observed bias, coverage for κ was generally acceptable, fluctuating around .95 for all methods. One exception was the MH algorithm with $n = 5, \kappa = 32$, which performed particularly badly with an average mode for κ of 133.12 and a coverage of .64. In this condition, many datasets have very high estimated concentrations, which the MH algorithm does not handle well. The rejection algorithm performed much better in this case, providing an adequate coverage of .96 even though the results were still strongly biased with the average estimate of κ at 55.74.




Figure 3: Coverages of κ for different sample sizes (n) and concentration (κ), all with $J = 3$.

Mean Computation Time

The final column denotes the computational time for the algorithms as implemented in C++ via Rcpp, which was averaged over all replications. As with previous results, the Gibbs sampling method performed worst by far. Its computational time increased as more lag was required, which occurred with higher sample sizes and concentration. The longest reported time was 59.03 seconds per replication. The computation time for the MH algorithm was independent of κ , but it did depend on n , becoming slower with increasing sample size, although the maximum time a replication took was only about .2 seconds. The rejection algorithm was very fast, as expected, and its computational time was independent of both sample size and κ , so that it never took more than .05 seconds for a replication.

Multiple groups

In Table 3, results for analysis of multiple groups of data are shown. For the posterior group means the observed pattern was similar to the single group case. Mean directions for cells with low concentration and small sample size showed unsystematic bias, which likely occurred due to sampling error of the datasets. The Gibbs sampler showed coverage for the posterior μ that was too low in all cells. For both the MH and rejection method, coverage of μ was generally adequate with the exception of $n_j = 5, \kappa = .1$ and $n_j = 30, \kappa = .1$, for which coverage was observed to be too low. As explained previously, this is likely due to a lacking mechanism to sample κ .

As with $J = 1$, a systematic bias was observed in κ for all methods, in particular for $n_j = 5$. For $n_j = 30$ and $n_j = 100$ much less bias was observed. Figure 3 shows the coverages of κ per method for different sample sizes and concentration, with $J = 3$.

Table 3: Average of posterior properties over 1000 replications for $J = 3$, with true means $\mu_1 = 20^\circ$, $\mu_2 = 40^\circ$, $\mu_3 = 60^\circ$, for different samples sizes per group (n_j) and concentration (κ).

n_j	κ	Method	Posterior μ				Posterior κ ^a		Acc. ^b	MCT ^c
			μ_1	μ_2	μ_3	Coverage	Mode	Coverage		
5	0.1	Gibbs	21.97	52.40	64.60	0.52	0.79	0.75	1	1.18
		MH	23.40	52.67	64.07	0.85	0.36	0.96	0.22	0.11
		Rejection	23.23	52.39	63.80	0.85	0.40	0.97	0.98	0.05
	4	Gibbs	20.94	39.34	59.95	0.69	5.39	0.89	1	20.13
		MH	20.95	39.33	59.96	0.94	4.69	0.96	0.47	0.13
		Rejection	20.95	39.34	59.96	0.94	4.57	0.97	1	0.06
	32	Gibbs	—	—	—	—	—	—	—	—
		MH	19.97	39.80	59.97	0.92	45.82	0.93	0.76	0.14
		Rejection	19.97	39.80	59.97	0.92	38.37	0.95	1	0.06
30	0.1	Gibbs	32.41	33.81	61.85	0.60	0.30	0.85	1	1.19
		MH	32.88	33.56	61.16	0.91	0.14	0.98	0.11	0.23
		Rejection	32.66	33.59	61.24	0.90	0.16	0.98	0.98	0.09
	4	Gibbs	20.10	39.95	60.27	0.73	4.19	0.93	1	20.25
		MH	20.10	39.95	60.27	0.95	4.12	0.96	0.22	0.24
		Rejection	20.10	39.95	60.27	0.95	4.04	0.97	1	0.14
	32	Gibbs	—	—	—	—	—	—	—	—
		MH	19.94	40.05	60.07	0.94	32.62	0.96	0.55	0.24
		Rejection	19.94	40.05	60.07	0.94	32.59	0.95	1	0.18
100	0.1	Gibbs	19.50	44.97	58.58	0.67	0.18	0.91	1	3.75
		MH	18.82	45.31	58.30	0.94	0.10	0.97	0.07	0.51
		Rejection	19.09	45.21	58.57	0.94	0.10	0.98	0.98	0.09
	4	Gibbs	19.93	40.01	60.04	0.74	4.05	0.94	1	99.10
		MH	19.93	40	60.04	0.95	4.03	0.95	0.13	0.52
		Rejection	19.93	40.01	60.04	0.95	3.97	0.96	1	0.10
	32	Gibbs	—	—	—	—	—	—	—	—
		MH	20.02	39.99	59.99	0.95	32.18	0.95	0.36	0.52
		Rejection	20.02	39.99	59.99	0.95	32.19	0.96	1	0.08

^aPosterior κ mode denotes the mode as described in section 4.3. Coverage denotes the proportion of replications in which the true κ fell within the 95 % HDI.

^bAcceptance ratio. For Gibbs sampling, this is always 1. For Metropolis-Hastings, Q_{acc}/Q is given. For the rejection method, this is Q/Q_{can} .

^cMean Computation Time of one replication in seconds.

Coverages for κ were adequate, although the Gibbs sampler once again performed badly. The MH and rejection sampler did perform well, although the rejection method seemed slightly more prone to coverages that are too high. These coverages in the range .95-1 indicate that the HDI would be chosen too wide so that the true value falls within the HDI more often than expected. Finally, computational time increased slightly with three groups for all methods.

5 Discussion

This paper presented three different MCMC approaches for Bayesian estimation of the mean directions μ_j of multiple groups of circular data with common but unknown concentration κ . These approaches were based on existing knowledge on Bayesian analysis of circular data that which could be used for analysis of a single group of circular data. Additionally, a systematic investigation of the performance of the three approaches was performed.

Comparing the methods, clear differences became apparent. The Gibbs sampler encountered many problems, among which were bias, undesirable coverages, long computational time, and complexity in application. The MH method performed adequately, but it does not show desirable acceptance rates for large datasets with small concentration when using the current χ^2 proposal density. In addition, it did not seem to handle extremely concentrated data well. The rejection algorithm by [Forbes and Mardia \(2014\)](#) was found to be the most promising of the MCMC-methods available in the literature at present, due to fast computational speed, fast convergence and adequate coverage.

The model developed here is still limited in terms of scope; it provides a basic between-subjects design for multiple groups of circular data, but extensions of this model such as a between-within-design or the inclusion of covariates have yet to be developed. Although in the present study the rejection algorithm was the most advantageous, a general MH algorithm may prove more flexible for such extended models due to its more direct approach. It is expected that extending the MCMC methods provided here to more complex models will exacerbate any issues regarding acceptance rates in different ways, so it remains to be seen which method will perform best after such an extension.

This study is also limited to the assumption that the data follows the von Mises distribution. It is expected that a large chunk of the circular data encountered in practice will follow this distribution. Groundwork for a general method for any kind of circular distribution was provided by [Bhattacharya and Sengupta \(2009\)](#) and employs importance sampling. Because importance sampling relies on defining an additional density to approximate normalising constants, simpler methods such as the ones presented here are preferred where possible.

In sum, the intrinsic approach offers a promising and flexible approach to Bayesian analysis of circular data, and its extension to a model with J multiple groups is an important first step towards developing flexible modeling of circular data in between-subjects designs.

References

- Abramowitz, M. and Stegun, I. A. (1972). *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover Publications.
- Best, D. and Fisher, N. I. (1979). "Efficient simulation of the von Mises distribution." *Applied Statistics*, 152–157.
- Bhattacharya, S. and Sengupta, A. (2009). "Bayesian analysis of semiparametric linear-circular models." *Journal of agricultural, biological, and environmental statistics*, 14(1): 33–65.
- Damien, P., Wakefield, J., and Walker, S. (1999). "Gibbs sampling for Bayesian non-conjugate and hierarchical models by using auxiliary variables." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(2): 331–344.
- Damien, P. and Walker, S. (1999). "A full Bayesian analysis of circular data using the von Mises distribution." *Canadian Journal of Statistics*, 27(2): 291–298.
- Dowe, D. L., Oliver, J. J., Baxter, R. A., and Wallace, C. S. (1996). "Bayesian estimation of the von Mises concentration parameter." In *Maximum Entropy and Bayesian Methods*, 51–60. Springer.
- Eddelbuettel, D. and François, R. (2011). "Rcpp: Seamless R and C++ Integration." *Journal of Statistical Software*, 40(8): 1–18.
URL <http://www.jstatsoft.org/v40/i08/>
- Ferrari, C. (2009). "The wrapping approach for circular data Bayesian modeling."
- Fisher, N. I. (1995). *Statistical analysis of circular data*. Cambridge: Cambridge University Press.
- Forbes, P. G. and Mardia, K. V. (2014). "A Fast Algorithm for Sampling from the Posterior of a von Mises distribution." *arXiv preprint arXiv:1402.3569*.
- Guttorp, P. and Lockhart, R. A. (1988). "Finding the location of a signal: A Bayesian analysis." *Journal of the American Statistical Association*, 83(402): 322–330.
- Harrison, D. and Kanji, G. (1988). "The development of analysis of variance for circular data." *Journal of Applied Statistics*, 15(2): 197–223.
- Harrison, D., Kanji, G., and Gadsden, R. (1986). "Analysis of variance for circular data." *Journal of Applied Statistics*, 13(2): 123–138.
- Hastings, W. K. (1970). "Monte Carlo sampling methods using Markov chains and their applications." *Biometrika*, 57(1): 97–109.
- Johnson, V. E. and Albert, J. H. (1999). *Ordinal data modeling*. Springer.
- Leary, T. (1957). *Interpersonal diagnosis of personality*. New York: Ronald Press.
- Mardia, K. V. and Jupp, P. E. (1999). *Directional statistics*. New York: Wiley.

- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). "Equation of state calculations by fast computing machines." *The journal of chemical physics*, 21(6): 1087–1092.
- Nunez-Antonio, G. and Gutierrez-Pena, E. (2005). "A Bayesian analysis of directional data using the projected normal distribution." *Journal of Applied Statistics*, 32(10): 995–1001.
- Nunez-Antonio, G. and Gutiérrez-Pena, E. (2005). "A Bayesian analysis of directional data using the von Mises–Fisher distribution." *Communications in Statistics—Simulation and Computation*®, 34(4): 989–999.
- R Development Core Team (2013). "R: A language and environment for statistical computing." Technical report, ISBN 3-900051-07-0. R Foundation for Statistical Computing. Vienna, Austria, 2013. url: <http://www.R-project.org>.
- Venter, J. (1967). "On estimation of the mode." *The Annals of Mathematical Statistics*, 1446–1455.
- Von Mises, R. (1918). "Über die Ganzzahligkeit der Atomgewichte und verwandte Fragen." *Phys. Z*, 19: 490–500.