Journal of Statistical Planning and Inference 211 (2021) 315–325

Contents lists available at ScienceDirect

Journal of Statistical Planning and Inference

journal homepage: www.elsevier.com/locate/jspi

Bayesian tests for circular uniformity

Kees Tim Mulder *, Irene Klugkist

Utrecht University, The Netherlands

A R T I C L E I N F O

Article history:
Received 5 October 2018
Accepted 4 June 2020
Available online 17 July 2020

Keywords:
Circular statistics
Marginal likelihood
Jeffreys prior
Conjugate prior
Bayes factor
von Mises distribution

A B S T R A C T

Circular data are data measured in angles or directions, which occur in a wide variety of scientific fields. An often investigated hypothesis is that of circular uniformity, or isotropy. Frequentist methods for assessing the circular uniformity null hypothesis exist, but do not allow the user faced with an insignificant result to distinguish lack of power from support for the null hypothesis. Bayesian hypothesis tests, which solve this issue and several others, are developed here. They are easy to compute and perform well, which is shown in a simulation. Two alternative hypotheses are considered. One is based on the von Mises distribution and performs well against unimodal alternatives. Another is based on a kernel density, which acts as an omnibus test against all other scenarios. Assessing the performance of the tests using different priors, it is shown that they are powerful and allow more elaborate conclusions than classical tests of circular uniformity.

1. Introduction

Circular data are measured in angles or directions. They are frequently encountered in scientific fields as diverse as life sciences (Mardia, 2011), behavioral biology (Bulbert et al., 2015), cognitive psychology (Kaas and Van Mier, 2006), bioinformatics (Mardia et al., 2008), political sciences (Gill and Hangartner, 2010) and environmental sciences (Arnold and SenGupta, 2006). In psychology, circular data occur often in motor behavior research (Mechsner et al., 2001, 2007; Postma et al., 2008; Baayen et al., 2012), as well as in the application of circumplex models (Gurtman and Pincus, 2003; Gurtman, 2009; Leary, 1957). Circular data differ from linear data in the sense that circular data are measured in a periodical sample space. For example, an angle of 1° is quite close to an angle 359°, although linear intuition suggests otherwise.

A fundamental hypothesis of interest is that of circular uniformity. A test for circular uniformity can be used to assess a hypothesis of theoretical interest by itself, but can also be used as a preliminary assessment, because most tests performed in circular statistics are only valid if the data is non-uniform. Several methods for assessing circular uniformity exist in the frequentist framework. These will be reviewed in Section 2.

In the rest of this paper, Bayesian hypothesis tests will be added to this arsenal. In order to create a Bayesian test of circular uniformity, the Bayes factor will be employed, which is often hailed as the standard way of performing Bayesian hypothesis tests (Kass and Raftery, 1995; Jeffreys, 1961). A major advantage of this Bayesian method is that through specifying the alternative hypothesis and the associated prior, we can precisely quantify support for either the null hypothesis or the alternative hypothesis. Methods based on null hypothesis significance testing only signify whether or not the null hypothesis can be rejected, but never provide support in favor of the null hypothesis. In practice, failure to reject the null hypothesis in a frequentist test is often taken as evidence for the null. However, a failure to reject the null might just as well be caused by a lack of power, so that the evidence in the data is indifferent to circular uniformity. In

* Corresponding author.
E-mail address: keestimmulder@gmail.com (K.T. Mulder).

https://doi.org/10.1016/j.jspi.2020.06.002

contrast, the Bayesian hypothesis test developed here is able to provide support for the null hypothesis. This alleviates some of the well-described issues with null hypothesis significance testing, and is particularly useful for tests that are used as a preliminary assessment, such as circular uniformity tests.

To compute the Bayes factor, the so-called marginal likelihood must be obtained for each hypothesis. The marginal likelihood is the normalizing constant of the posterior, and the key ingredient of the Bayes factor. To obtain the marginal likelihood one must specify the prior distribution of the parameters in each hypothesis. The null hypothesis of circular uniformity has no parameters, so no prior is needed for it. The alternative hypothesis, however, requires specification of both a model for the data, would they be non-uniform, and second, a prior for the parameters in this model. This paper will investigate two alternative hypotheses, one based on the von Mises distribution and one based on a kernel density alternative. In addition, several priors will be investigated that can be used with each model.

The rest of the paper is structured as follows. A short review of frequentist tests of circular uniformity is provided in Section 2. The Bayesian circular uniformity test for a von Mises alternative is discussed in Section 3. The Bayesian circular uniformity test for a kernel density alternative, which functions as an omnibus test, is discussed in Section 4. The methods are applied to example datasets in Section 5. Section 6 provides a discussion.

2. Frequentist tests of circular uniformity

Here, we will shortly review frequentist tests of circular uniformity. Four commonly used tests are Kuiper's test (Kuiper, 1960), Rayleigh's test (Mardia and Jupp, 2000; Brazier, 1994), Rao's test of equal spacing (Rao, 1976) and Ajne's test (Ajne, 1968). Perhaps the most common of these is the Rayleigh test. Let $\boldsymbol{\theta} = (\theta_i = 1, \ldots, \theta_n)$ denote a set of data consisting of angles, and let the mean resultant length $\bar{R} = n^{-1}\sqrt{(\sum_{i=1}^n \cos\theta_i)^2 + (\sum_{i=1}^n \sin\theta_i)^2}$. Then the Rayleigh test statistic can be computed simply as $2n\bar{R}^2$, which has approximately a χ^2_2 distribution. It can be shown that the Rayleigh test is the most powerful test against von Mises alternatives, as well as Projected Normal (PN) alternatives (Bhattacharyya and Johnson, 1969). Although the Rayleigh test is consistent against unimodal alternatives, it is not consistent against alternatives that have resultant length $\rho = 0$, in particular distributions with antipodal symmetry (Mardia and Jupp, 2000).

Another test is Kuiper's test (Kuiper, 1960), which is based on the maximum difference between the theoretical and empirical distribution function. It is consistent against all alternatives to uniformity (Mardia and Jupp, 2000). A similar test uses Watson's U^2 statistic (Watson, 1961), which is instead based on the *mean* difference between the theoretical and empirical distribution function.

Several other tests for circular uniformity exist, among which Rao's equal spacing test (Rao, 1976), the range test (Laubscher and Rudolph, 1968), the Hodges–Ajne test (Hodges, 1955; Ajne, 1968), Ajne's A_n test (Ajne, 1968), and the Hermans–Rasson test (Hermans and Rasson, 1985). Somewhat more recently, a smooth test for circular uniformity was developed by Bogdan et al. (2002). A test specifically targeting multimodal alternatives was developed by Pycke (2010).

3. Tests with a von Mises alternative

In this section, a Bayesian hypothesis test for circular uniformity against a von Mises alternative will be developed. The von Mises distribution is a natural distribution on the circle, given by

$$\mathcal{M}(\theta \mid \mu, \kappa) = [2\pi I_0(\kappa)]^{-1} \exp\{\kappa \cos(\theta - \mu)\}, \tag{1}$$

where $\theta \in [0, 2\pi)$ is an angular observation, $\mu \in [0, 2\pi)$ is the mean direction, $\kappa \in \mathbb{R}^+$ is a concentration parameter and $I_0(\cdot)$ is the modified Bessel function of the first kind and order zero.

The test will be based on the Bayes factor, which is the ratio of two marginal likelihoods, given by

$$BF_{10} = \frac{m_1(\boldsymbol{\theta})}{m_0(\boldsymbol{\theta})} = \frac{\int_{\boldsymbol{\phi}} p(\boldsymbol{\phi}, \boldsymbol{\theta} \mid H_1) d\boldsymbol{\phi}}{\int_{\boldsymbol{\phi}} p(\boldsymbol{\phi}, \boldsymbol{\theta} \mid H_0) d\boldsymbol{\phi}}, \tag{2}$$

where $\boldsymbol{\theta} = \theta_1, \ldots, \theta_n$ is a dataset of i.i.d. angles, and $\boldsymbol{\phi}$ is a vector of parameters belonging to the chosen model. For the von Mises distribution $\boldsymbol{\phi} = (\mu, \kappa)^T$.

Because the null hypothesis does not feature parameters and assigns equal probability to each data point, $m_0(\boldsymbol{\theta})$ depends only on the sample size. The circular uniform distribution is defined by $p(\theta) = (2\pi)^{-1}$, so the marginal likelihood for H_0 is obtained by

$$m_0(\boldsymbol{\theta}) = \prod_{i=1}^n p(\theta_i) = (2\pi)^{-n}. \tag{3}$$

The marginal likelihood of H_1 is given by

$$m_1(\boldsymbol{\theta}) = \int_{\boldsymbol{\phi}} p(\boldsymbol{\phi}, \boldsymbol{\theta} \mid H_1) d\boldsymbol{\phi} = \int_0^{\infty} \int_0^{2\pi} p(\mu, \kappa, \boldsymbol{\theta} \mid H_1) d\mu d\kappa, \tag{4}$$

where

$$p(\mu, \kappa, \boldsymbol{\theta} \mid H_1) \propto p(\mu, \kappa \mid H_1)\, p(\boldsymbol{\theta} \mid \mu, \kappa, H_1) \tag{5}$$

is the kernel of the posterior, where the prior $p(\mu, \kappa \mid H_1)$ must still be chosen, and $p(\boldsymbol{\theta} \mid \mu, \kappa, H_1)$ is the likelihood. The likelihood of the von Mises distribution is given by

$$p(\boldsymbol{\theta} \mid \mu, \kappa, H_1) = \prod_{i=1}^{n} \mathcal{M}(\theta_i \mid \mu, \kappa) = [2\pi I_0(\kappa)]^{-n} \exp\left\{R\kappa \cos(\bar{\theta} - \mu)\right\}, \tag{6}$$

where R is the resultant length and $\bar{\theta}$ is the mean direction.

For any prior that does not depend on μ, so that $p(\mu, \kappa) = p(\kappa)/2\pi$, the Bayes factor simplifies to

$$BF_{10} = (2\pi)^n \int_0^\infty \frac{p(\kappa)}{2\pi} \int_0^{2\pi} [2\pi I_0(\kappa)]^{-n} \exp\left\{R\kappa \cos(\bar{\theta} - \mu)\right\} d\mu d\kappa \tag{7}$$

$$= \int_0^\infty I_0(\kappa)^{-n} \frac{p(\kappa)}{2\pi} \int_0^{2\pi} \exp\left\{R\kappa \cos(\bar{\theta} - \mu)\right\} d\mu d\kappa \tag{8}$$

$$= \int_0^\infty p(\kappa) \frac{I_0(R\kappa)}{I_0(\kappa)^n} d\kappa, \tag{9}$$

where the last step uses the fact that $I_0(x) = [2\pi]^{-1} \int_0^{2\pi} \exp\{x \cos\theta\}\, d\theta$. Thus, computation of the Bayes factor requires only univariate integration.

3.1. Choosing priors

Choosing the prior for this hypothesis test is not trivial. In principle, the prior for $\{\mu, \kappa\}$ should capture our actual belief about the possible values of the parameters, given that the alternative hypothesis is true. Although researchers are free to determine their own prior for this test, we propose some general guidelines for the set of possible priors to be considered here.

First, it should be noted that choosing improper priors generally do not result in useful Bayesian hypothesis tests. Therefore, only proper priors will be considered here.

Second, if a test for circular uniformity is considered, the researcher will generally not already have an idea about the mean direction of the data if H_1 is true, because they are investigating whether there even is a preferred (mean) direction. Therefore, we suggest taking a circular uniform prior on μ. This is done by taking $p(\mu) = [2\pi]^{-1}$ and independent of κ, so that $p(\mu, \kappa) = p(\mu)p(\kappa) = p(\kappa)/2\pi$ and we can concern ourselves only with choosing the prior for κ.

Finally, a researcher that considers circular uniformity to be a reasonable hypothesis rarely expects strongly concentrated distributions, even if the alternative hypothesis were true. Therefore, we suggest setting a prior for κ that gives most of its probability to fairly low values of κ. Should the data follow a concentrated distribution anyway, the test will be powerful regardless.

In practice, whether these expectations are reasonable should be assessed by the researcher themselves. However, taking this approach allows us to build default methods that work well in most research scenarios in which the test would be applied. In the following sections, different choices for priors are considered, and for each the resulting test is assessed.

3.2. Priors based on the conjugate prior

A conjugate prior for the von Mises distribution was suggested by Guttorp and Lockhart (1988), and is given by

$$p(\mu, \kappa) \propto I_0(\kappa)^{-c} \exp\{R_0\kappa \cos(\mu - \mu_0)\}, \tag{10}$$

where μ_0, R_0, and c are the prior mean, prior resultant length, and prior 'sample size', respectively. As discussed previously, we would like to remove the necessity to choose a prior mean μ_0. This can be done by putting a circular uniform prior on μ_0 and integrating it out so that

$$p(\mu, \kappa) \propto \int_0^{2\pi} [2\pi]^{-1} I_0(\kappa)^{-c} \exp\{R_0\kappa \cos(\mu - \mu_0)\}\, d\mu_0 = \frac{I_0(R_0\kappa)}{I_0(\kappa)^c}, \tag{11}$$

which only depends on κ. Then, all that remains is choosing values for R_0 and c. It can easily be seen that imagining a single datapoint on the circle results in $R_0 = 1$ and $c = 1$, producing the constant prior on κ. Because the constant prior is improper and therefore invalid for hypothesis testing, we examine two valid alternatives.

First, the prior used by McVinish and Mengersen (2008) has $R_0 = 0$, $c = 1$, so that we obtain

$$p(\mu, \kappa) \propto I_0(\kappa)^{-1}. \tag{12}$$

This prior will be referred to as prior (12), and is displayed in Fig. 1, in red.

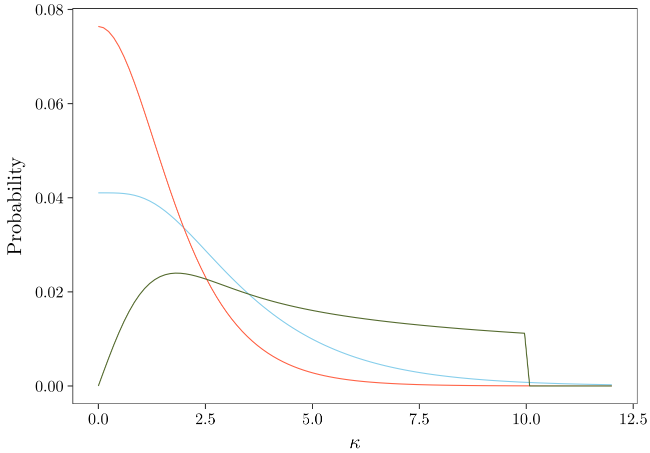

Fig. 1. Graphs of three different choices of priors for κ: Prior (12) (red) has $R_0 = 0, c = 1$, prior (13) (blue) has $R_0 = \sqrt{2}, c = 2$, and the Jeffreys prior (green) has $\kappa_u = 10$.

Second, the prior could be taken to be proportional to the likelihood of an imagined dataset $\{a, a + \pi/2\}$, with a any angle. This imagined dataset, somewhat arbitrarily, has two angles at 90° from one another. This results in $R_0 = \sqrt{2}, c = 2$, so we obtain

$$p(\mu, \kappa) \propto I_0\left(\sqrt{2}\kappa\right) I_0(\kappa)^{-2}. \tag{13}$$

This prior will be referred to as prior (13), and is displayed in Fig. 1, in blue. It can be seen that this prior has more mass at higher values of κ.

Denoting the normalizing constant of either prior by $g = 2\pi \int_0^\infty I_0(R_0\kappa)I_0(\kappa)^{-c} d\kappa$, the marginal likelihood for H_1 for these von Mises-based priors is

$$m_1(\boldsymbol{\theta}) = \int_0^\infty \int_0^{2\pi} p(\mu, \kappa) p(\boldsymbol{\theta} \mid \mu, \kappa) d\mu d\kappa \tag{14}$$

$$= g\,[2\pi]^{-n} \int_0^\infty \frac{I_0(R_0\kappa)}{I_0(\kappa)^c} I_0(\kappa)^{-n} \int_0^{2\pi} \exp\left\{R\kappa \cos\left(\bar{\theta} - \mu\right)\right\} d\mu d\kappa \tag{15}$$

$$= g\,[2\pi]^{-(n+1)} \int_0^\infty I_0(R_0\kappa) I_0(R\kappa) I_0(\kappa)^{-(n+c)} d\kappa. \tag{16}$$

The Bayes factor in favor of the alternative is then

$$BF_{10} = \frac{m_1(\boldsymbol{\theta})}{m_0(\boldsymbol{\theta})} = [2\pi]^n m_1(\boldsymbol{\theta}) = g\,[2\pi]^{-1} \int_0^\infty I_0(R_0\kappa) I_0(R\kappa) I_0(\kappa)^{-(n+c)} d\kappa. \tag{17}$$

This can be computed by univariate numerical integration. For computational stability, it can be beneficial to first compute the log of the product of Bessel functions inside the integral, using

$$I_0(R_0\kappa) I_0(R\kappa) I_0(\kappa)^{-(n+c)} = \exp\left\{\log I_0(R_0\kappa) + \log I_0(R\kappa) - (n+c)\log I_0(\kappa)\right\}.$$

3.3. *Jeffreys prior*

The Jeffreys prior is a common choice for non-informative priors, especially in low-dimensional parameter spaces as is the case here. The Jeffreys prior is proportional to the square root of the determinant of the Fisher Information Matrix $\mathcal{I}(\boldsymbol{\phi})$ for a single observation, so that for the von Mises distribution it is given by

$$p(\boldsymbol{\phi}) \propto \sqrt{\det\left[\mathcal{I}(\boldsymbol{\phi})\right]} = \sqrt{\kappa A(\kappa) A'(\kappa)}, \tag{18}$$

where $A(\kappa) = I_1(\kappa)/I_0(\kappa)$ and $A'(\kappa) = \frac{d}{d\kappa} A(\kappa)$.

An attractive property of this prior is that it has $p(\kappa = 0) = 0$. However, this prior is improper, which means it cannot be used directly in hypothesis testing. Therefore, we suggest to take a truncation of this prior from above at some value κ_u. A proper prior based on the Jeffreys prior is then given by

$$p(\mu, \kappa \mid \kappa_u) = \frac{I(\kappa < \kappa_u)\sqrt{\kappa A(\kappa) A'(\kappa)}}{2\pi \int_0^{\kappa_u} \sqrt{\kappa A(\kappa) A'(\kappa)} d\kappa}, \tag{19}$$

where $I(\cdot)$ is an indicator function. This prior with $\kappa_u = 10$ is shown in Fig. 1, in green.

To choose κ_u, it might be thought of as an upper bound for the values of κ for which we will be able to find support. If the data favors a value of κ higher than κ_u, the marginal likelihood of the alternative hypothesis H_1 will be underestimated, although H_1 will still be preferred. Conversely, it should be noted that even if the likelihood strongly suggests $\kappa < \kappa_u$, the resulting Bayes Factor will still depend on κ_u through the integral in the normalizing constant. The concern that a somewhat arbitrary choice must be made can be alleviated somewhat by performing a sensitivity analysis. In Section 3.4, it will be shown that the hypothesis test using this prior performs well even for some fixed values of κ_u.

The Bayes Factor is given by

$$BF_{10} = (2\pi)^n \int_0^\infty \int_0^{2\pi} p(\mu, \kappa) f(\boldsymbol{\theta} \mid \mu, \kappa) d\mu d\kappa \tag{20}$$

$$= \int_0^\infty p(\mu, \kappa) I_0(\kappa)^{-n} \int_0^{2\pi} \exp\left\{R\kappa \cos(\bar{\theta} - \mu)\right\} d\mu d\kappa \tag{21}$$

$$= 2\pi \left[\int_0^{\kappa_u} \sqrt{\kappa A(\kappa) A'(\kappa)} d\kappa\right]^{-1} \int_0^{\kappa_u} \sqrt{\kappa A(\kappa) A'(\kappa)} I_0(R\kappa) I_0(\kappa)^{-n} d\kappa. \tag{22}$$

3.4. Simulation

In order to assess the performance of the Bayesian hypothesis tests with a von Mises alternative and the three priors discussed previously, a simulation study was performed. One million datasets were sampled from the von Mises distribution with κ set to $\{0, 0.5, 1, 2, 5\}$, where $\kappa = 0$ was used three times more often as it represents H_0. Samples sizes were randomly selected from $\{2, \ldots, 15, 20, 30, \ldots, 190, 200\}$.

Fig. 2 shows the performance of $BF_{10} > 1$ as a decision criterion for all priors, as well as a plot of the obtained log Bayes factors. In general, all three tests perform well, and are particularly good at correctly classifying data generated under the null hypothesis. Prior (12) and prior (13) show very similar performance, although prior (13) is more prone to select H_0. The Jeffreys prior with $\kappa_u = 20$ is even more prone to select H_0. When data is almost uniform with $\kappa = 0.5$, the tests need a large sample size to select H_1 more than half of the time (around $n > 50$ for prior (12) and prior (13), and $n > 100$ for the Jeffreys prior with $\kappa_u = 20$).

Compared to the error rates of the Rayleigh test, the current test has better power in all situations but those with $\kappa = .5, n > 30$ and $\kappa = 0, n < 30$. For those cases, it can be seen in the plots on the right of Fig. 2 that the Bayes factors that are produced are somewhat indecisive, so they may not be taken as evidence in favor of either hypothesis at all. Also, it can be seen that if H_0 is true, $p(H_0 \mid \boldsymbol{\theta}) \to 1$ as $n \to \infty$, which is not the case for the Rayleigh test.

It can be seen that in some cases, such as in 2(a) with $\kappa = 0.5$, increasing the sample size from 1 to 10 actually decreases the probability of selecting H_1, even though it is the true hypothesis. This is a known property of some Bayesian hypothesis tests. It should be noted that in these cases, the Bayes factor shows indecision.

4. Tests with a kernel density alternative

If the von Mises alternative is insufficient, the correct alternative distribution to test against is often unknown. A pure Bayesian approach could be to formulate a set of possible models, and choose between this set of alternatives. However, this requires attempting to fit an infinite set of models which might be hard to do in practice.

Instead, it may be useful to fit a very flexible model as the alternative, which can mimic the true distribution well, so as to provide an omnibus test against many possible models. A kernel density fulfills this role, being able to approximate any density given enough data. Recent developments of kernel density methods for circular data have focused on kernel density bandwidth selection and kernel regression (Di Marzio et al., 2009; Oliveira et al., 2012; Di Marzio et al., 2013; Oliveira et al., 2014).

Here, we will build a test for circular uniformity which uses a von Mises kernel density as the alternative. The pdf of the kernel density based on a dataset $\boldsymbol{\Theta} = \Theta_1, \ldots, \Theta_n$ is given by

$$f(\theta \mid \boldsymbol{\Theta}, \kappa) = \frac{1}{n} \sum_{i=1}^{n} \mathcal{M}(\theta \mid \Theta_i, \kappa). \tag{23}$$

Our interest is to obtain a posterior for the bandwidth κ, which is the only free parameter. However, if the likelihood is specified as

$$f(\boldsymbol{\Theta} \mid \kappa) = \prod_{j=1}^{n} f(\theta_j \mid \boldsymbol{\Theta}, \kappa) = \prod_{j=1}^{n} \sum_{i=1}^{n} \mathcal{M}(\theta_j \mid \Theta_i, \kappa). \tag{24}$$

then $\mathcal{M}(\theta_j \mid \Theta_i, \kappa) \to \infty$ if $i = j, \kappa \to \infty$. Therefore, following Hall et al. (1987), we specify the likelihood in a manner reminiscent of leave-one-out cross-validation, by setting

$$f(\boldsymbol{\Theta} \mid \kappa) = \prod_{j=1}^{n} \sum_{i \neq j} \mathcal{M}(\theta_j \mid \Theta_i, \kappa). \tag{25}$$

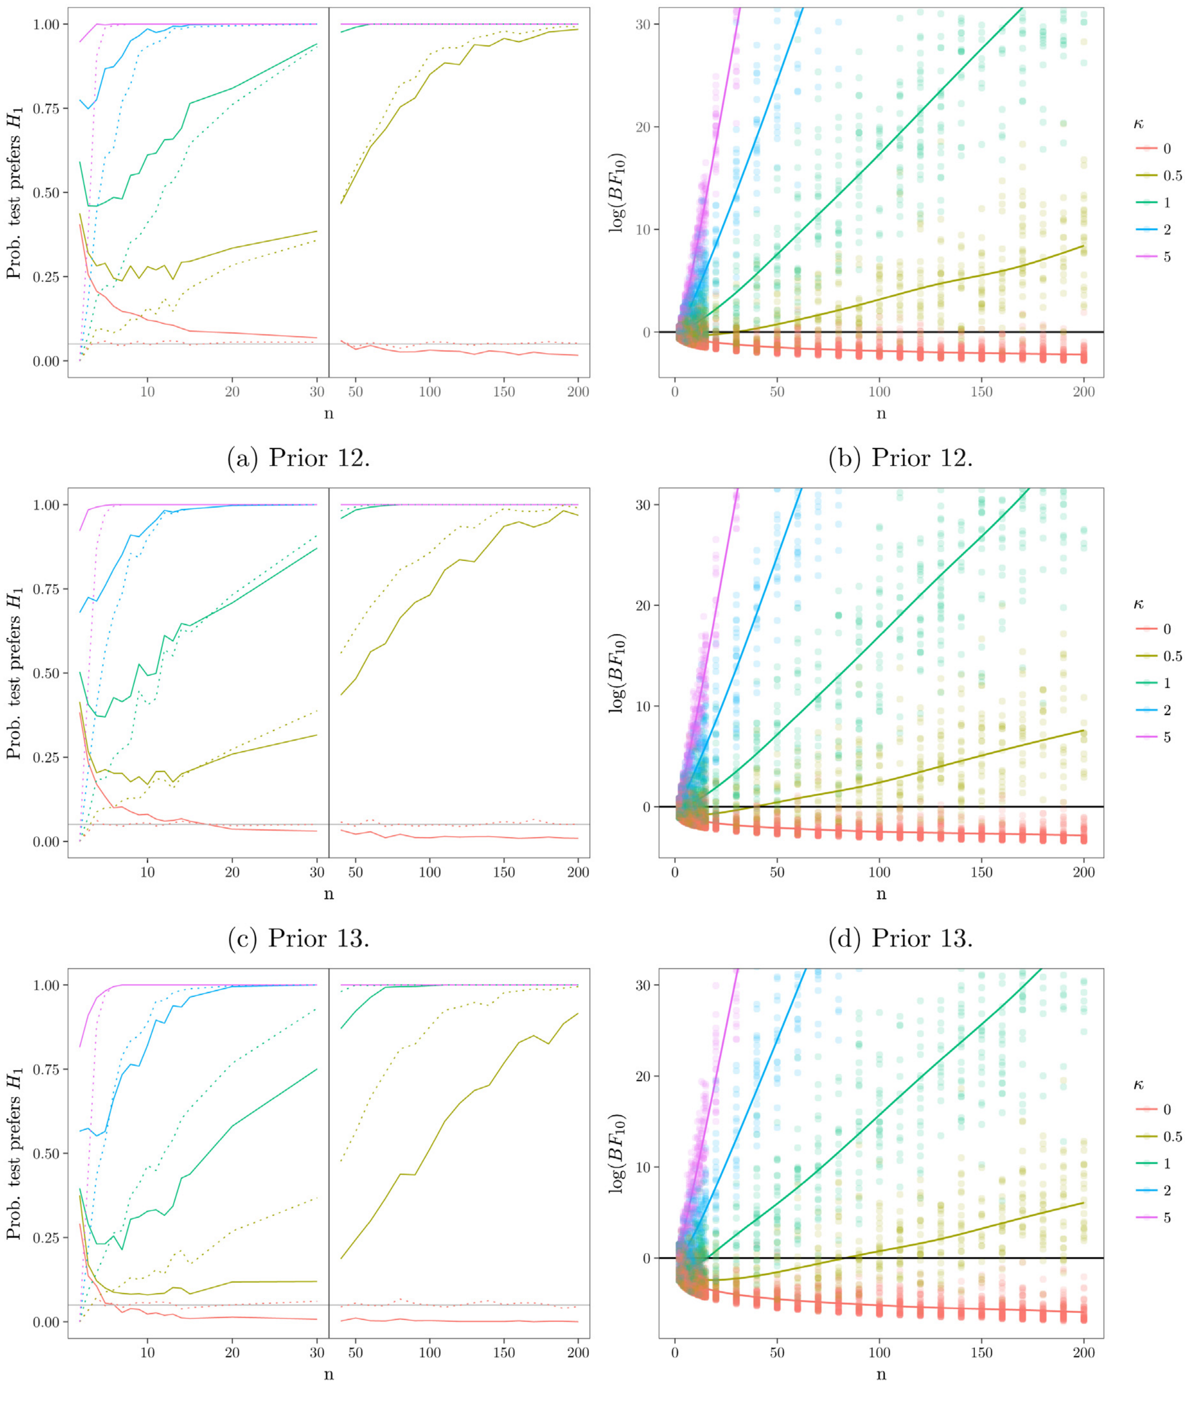

(a) Prior 12.

(b) Prior 12.

(c) Prior 13.

(d) Prior 13.

(e) Jeffreys prior with $\kappa_u = 20$.

(f) Jeffreys prior with $\kappa_u = 20$.

Fig. 2. Results of the simulation study for prior (12) (top), prior (13) (middle) and the Jeffreys prior (bottom) with $\kappa_u = 20$. The left plots show the proportion of simulations which obtained a Bayes factor in favor of the alternative hypothesis ($BF_{10} > 1$). Error rates for the Rayleigh test with $\alpha = .05$ are provided as dotted lines, with the nominal significance displayed as a gray line at .05. The right plots show a subsample of the log-Bayes factors obtained for different sample sizes n and κ, as well as a solid trendline computed from the full simulation showing the average log Bayes factor for each sample size.

This leads to the posterior

$$p(\kappa \mid \boldsymbol{\Theta}) \propto f(\boldsymbol{\Theta} \mid \kappa) p(\kappa) \tag{26}$$

where the prior for $p(\kappa)$ must still be set. Note that for this model, κ has a different interpretation than in the von Mises, so a different prior is in order. Specifically, in the von Mises model κ refers to the concentration of the full dataset, while in the kernel density model κ refers to the concentration around each separate data point. Therefore, higher values of κ should be considered likely a priori.

The Bayes factor is given by

$$BF_{10} = [2\pi]^n \int_0^\infty p(\kappa) \prod_{j=1}^n \sum_{i \neq j} [2\pi I_0(\kappa)]^{-1} \exp\left\{\kappa \cos(\theta_j - \Theta_i)\right\} d\kappa \tag{27}$$

$$= \int_0^\infty \frac{p(\kappa)}{I_0(\kappa)^n} \prod_{j=1}^n \sum_{i \neq j} \exp\left\{\kappa \cos(\theta_j - \Theta_i)\right\} d\kappa, \tag{28}$$

which is once again computed by univariate numerical integration.

For priors of the type discussed in Section 3.2, the Bayes factor for some R_0 and c can be written as

$$BF_{10} = \left[\int_0^\infty I_0(R_0\kappa) I_0(\kappa)^{-c} d\kappa\right]^{-1} \int_0^\infty I_0(R_0\kappa) I_0(\kappa)^{-(n+c)} \prod_{j=1}^n \sum_{i \neq j} \exp\left\{\kappa \cos(\theta_j - \Theta_i)\right\} d\kappa. \tag{29}$$

Another particularly good option for the kernel density model is the Jeffreys prior discussed in 3.3, as it allows tuning κ_u to accommodate reasonably high values for the concentration. For this prior, the Bayes factor can be written as

$$BF_{10} = \left[2\pi \int_0^{\kappa_u} \sqrt{\kappa A(\kappa) A'(\kappa)} d\kappa\right]^{-1} \int_0^{\kappa_u} \frac{\sqrt{\kappa A(\kappa) A'(\kappa)}}{I_0(\kappa)^{n+1}} \prod_{j=1}^n \sum_{i \neq j} \exp\left\{\kappa \cos(\theta_j - \Theta_i)\right\} d\kappa. \tag{30}$$

4.1. Simulation

We assess the performance of the kernel based circular uniformity test for antipodal von Mises. The antipodal von Mises is an antipodally symmetric mixture of two von Mises distributions, where data was obtained by drawing from the pdf

$$f(\theta \mid \mu, \kappa) = \frac{1}{2}\mathcal{M}(\theta \mid \mu, \kappa) + \frac{1}{2}\mathcal{M}(\theta \mid \mu + \pi, \kappa). \tag{31}$$

This alternative hypothesis is chosen to be especially hard for the von Mises based tests developed in Section 3. The setup in terms of sample sizes and chosen true values for κ is the same as in Section 3.4.

Results for data generated from the antipodal von Mises distribution are displayed in Fig. 3. It can be seen that the Rayleigh test performs abysmally, which is expected, because it is based on rejection of H_0 for large values of the resultant length, which for the antipodal von Mises is zero on average. Our method picks up the difference with reasonable power when data was generated with $\kappa \geq 2$. In order to detect non-uniformity for antipodal von Mises data with $\kappa = 1$, a very large sample is needed, but it must be noted that antipodal data with small κ is almost uniform. Evidence in favor of H_0 is collected slowly, but with larger sample sizes, H_0 is selected more and more.

5. Examples

In this section, the method will be applied to two real-world examples.

In a classic experiment on pigeon homing (Schmidt-Koenig, 1963), the vanishing angles of homing pigeons were measured, with the initial question of whether the vanishing direction is circular uniform or follows some other circular distribution. Two datasets from this experiment are depicted in Fig. 4. In one experiment, also provided in Fisher (1995), fifteen homing pigeons were measured to have vanishing directions given by

$$\{85°, 135°, 135°, 140°, 145°, 150°, 150°, 150°, 160°, 285°, 200°, 210°, 220°, 225°, 270°\},$$

shown in Fig. 4(a). In another dataset, provided in Mardia and Jupp (2000), ten pigeons were measured to have vanishing directions

$$\{55°, 60°, 65°, 95°, 100°, 110°, 260°, 275°, 285°, 295°\},$$

shown in Fig. 4(b).

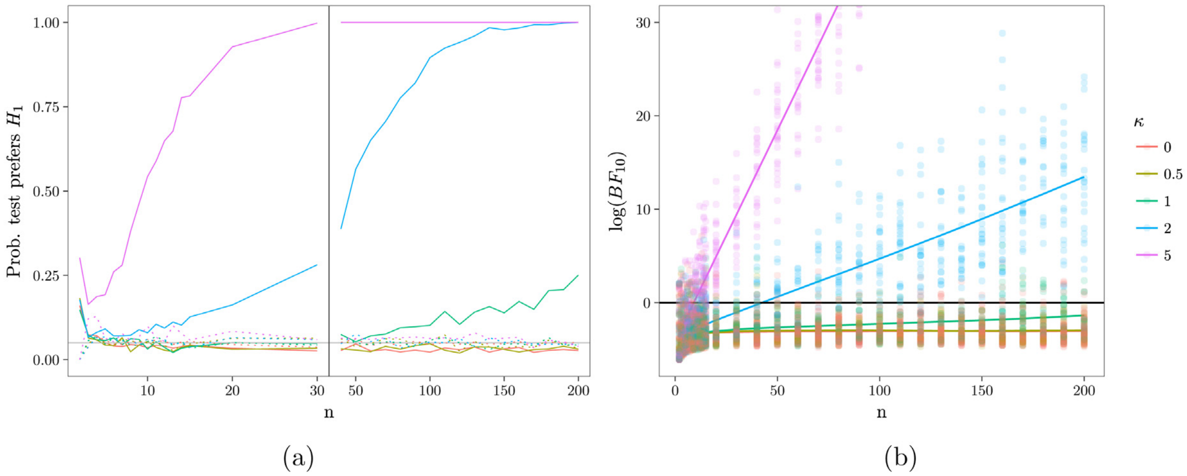

Fig. 3. Performance for data from the antipodal von Mises distribution with various true values for κ. The left plot shows the proportion of simulations which obtained a Bayes factor in favor of the alternative hypothesis ($BF_{10} > 1$). Error rates for the Rayleigh test with $\alpha = .05$ are provided as dotted lines, with that nominal significance displayed as a gray line at .05. The right plot shows a subsample of the log Bayes factors obtained for different sample sizes n and κ, as well as a solid trendline computed from the full simulation showing the average log Bayes factor for each sample size.

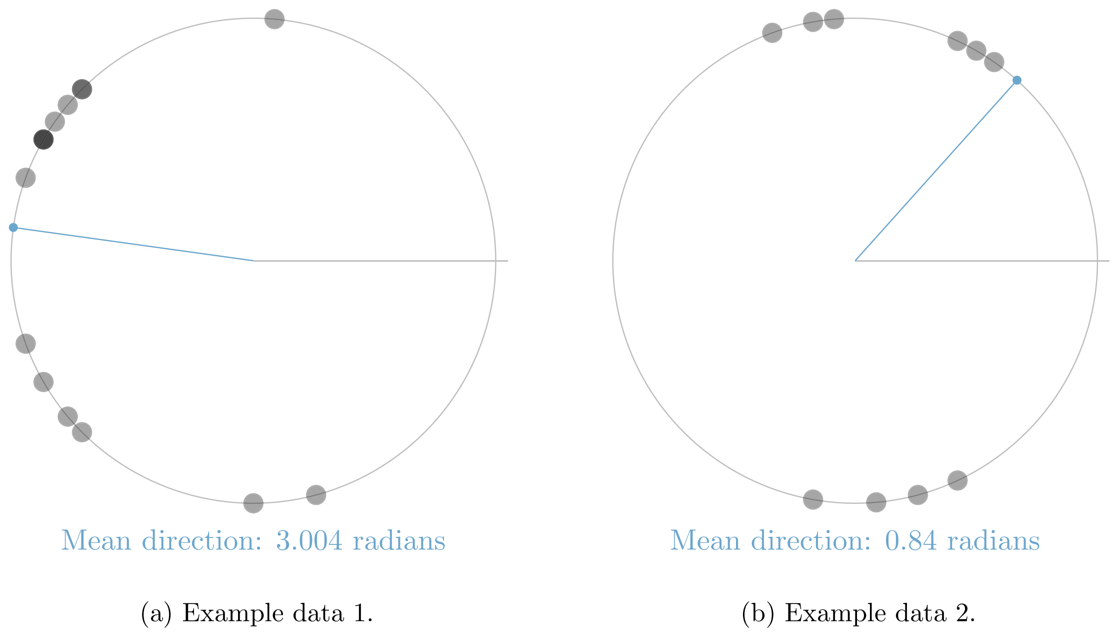

(a) Example data 1. (b) Example data 2.

Fig. 4. The two example datasets. For each subfigure, the blue line between the center and the circle depicts the mean direction, while the gray line depicts 0°.

5.1. Homing pigeon example 1

The results for the first dataset are given in Table 1. For this dataset, reasonable hypotheses are that the data are either circular uniform (which we call H_0), or that the data follow a symmetric unimodal distribution, where we pick the von Mises distribution (which we call $H_{\mathcal{M}}$ here). These two hypotheses are evaluated as discussed in Section 3.2, using prior (12) given by $p(\kappa) \propto I_0(\kappa)^{-1}$ because a low concentration is expected. Table 1 denotes the results of our hypothesis test, as well as the Rayleigh test for comparison. The log marginal likelihood of H_0 is -27.57, while the log marginal likelihood of $H_{\mathcal{M}}$ is -23.92, so $H_{\mathcal{M}}$ is most supported by the data. In fact, the Bayes factor in favor of $H_{\mathcal{M}}$ is 38.54, so that the posterior probability of $H_{\mathcal{M}}$ is 0.975, which constitutes strong support for this hypothesis.

Table 1
Results of example 1.

Bayes factor	$p(H_0 \mid \boldsymbol{\theta})$	$p(H_{\mathcal{M}} \mid \boldsymbol{\theta})$	Rayleigh statistic	Rayleigh p-value
38.542	0.025	0.975	0.637	0.001

Table 2
Results of example 2.

$p(H_0 \mid \boldsymbol{\theta})$	$p(H_{\mathcal{M}} \mid \boldsymbol{\theta})$	$p(H_k \mid \boldsymbol{\theta})$	Rayleigh statistic	Rayleigh p-value
0.034	0.012	0.954	0.223	0.620

5.2. Homing pigeon example 2

For the second dataset, the hypothesis that the data is bimodal is also reasonable, although we might not want to assume antipodal symmetry. To demonstrate the flexibility of the Bayesian approach, we evaluate three hypotheses jointly. The hypotheses are circular uniformity (H_0), the von Mises distribution ($H_{\mathcal{M}}$), and the kernel density alternative (H_k) described in Section 4. For the von Mises distribution, the same conjugate prior is used as before, $p(\kappa) \propto I_0(\kappa)^{-1}$. For the kernel density alternative, higher concentrations are more plausible than for the von Mises hypothesis, because dispersion in the final kernel density model is not exclusively determined by κ, but also by the spread of the data. Therefore, we pick the Jeffreys prior here, truncated above at $\kappa_u = 40$. In a small sensitivity analysis for the truncation value (not reported further), the Bayes factor was robust to truncation values above 20, although setting the value extremely high will influence the marginal likelihood, and the inference as a result.

In order to compare the relative probability of each hypothesis, posterior model probabilities were computed. When choosing between a set of p models, we can compute the posterior model probability of model i, assuming equal prior model probabilities, as

$$p(H_i \mid \boldsymbol{\theta}) = \frac{m_i(\boldsymbol{\theta})}{\sum_{j=1}^{p} m_j(\boldsymbol{\theta})}, \tag{32}$$

where $m_a(\boldsymbol{\theta})$ denotes the marginal likelihood of model H_a. This will provide the relative probabilities of the models that are assessed.

Results are displayed in Table 2. The Rayleigh test is not significant ($p = 0.62$), suggesting no departure from uniformity. In contrast, our comparison of hypotheses shows a preference for the kernel density alternative, giving it a posterior model probability of 0.954. This can be seen as evidence that the data generation distribution is likely neither the circular uniform distribution nor the von Mises distribution. Rather, the correct model was likely not included in the set of models that were assessed, which should motivate the researcher to further investigate possible models. This result is easy to interpret and understand, and provides a more complete picture than the usual frequentist test.

6. Discussion

Bayesian hypothesis tests for assessing circular uniformity were developed in this paper. The Bayesian approach provides three major advantages for this type of hypothesis. First, the hypothesis of circular uniformity is precisely the type of hypothesis which might be true in reality, so that we would want to choose H_0 if the data supports it. The available frequentist tests do not support this, as an insignificant p-value does not allow us to draw conclusion on whether H_0 is true. Second, the Bayesian hypothesis test allows us to quantify the strength of the evidence, either in an odds ratio in the Bayes factor, or in an intuitive probability in the posterior model probability, which is more informative than the simple dichotomous decisions provided by null hypothesis tests. Third, the Bayesian framework allows us to add additional hypotheses to the comparison quite easily. In example 2 in Section 5.2, this is used by having the kernel density alternative effectively act as a 'none of the above' category, motivating the researcher to search for a model that fits the data better.

Among the most central critiques of the Bayesian method (and Bayesian testing in particular) lies the difficulty in choosing priors, as this seemingly requires us to know in advance what distribution the data may have should the alternative hypothesis be true. Moreover, the conclusions drawn in Bayesian hypothesis tests are often highly dependent on seemingly arbitrary quantities, most notably the parameters of the prior distribution. However, when choosing a frequentist test for circular uniformity, one is faced with a plethora of tests (see Section 2) which are each most powerful against different alternatives. This choice closely mirrors the choice of the prior in the alternative hypothesis of a Bayesian hypothesis test. For example, this can be seen in Landler et al. (2018), where different tests are recommended for different expected alternative distributions. In either case, we must use our expectations of the distribution of the data, should the alternative hypothesis be true. Furthermore, in Section 3.1 it was shown how the selection of priors can be dealt with to circumvent the concerns about their influence on the results.

Beyond circular uniformity, previously Bayesian analyses of circular models have been investigated from several viewpoints. Bayesian model assessment has been investigated for wrapped models (Ravindran and Ghosh, 2011), projected normal models (Nuñez-Antonio et al., 2015) and semiparametric intrinsic models (Bhattacharya and SenGupta, 2009; George and Ghosh, 2006). However, the only previously discussed Bayesian test for circular uniformity the authors are aware of is in McVinish and Mengersen (2008), where the alternative hypothesis is a Dirichlet process mixture of triangular distributions. Compared to that work, our focus is on adding parametric alternatives, simplifying computation, assessing performance of the Bayes factor, developing accessible computational tools and comparison of this method to frequentist methods, both conceptually and in a simulation study. Computation involved in evaluating the marginal likelihood of our models has been reduced to simple univariate numerical integration, which makes running these tests more straightforward and markedly faster. Also, the tools used in this paper are easily available from R through the package BayesCircIsotropy, available on GitHub.

Assessing the performance of the method, it was shown that the test is often powerful in selecting the correct model, both for the data from the null hypothesis as well as the alternative. The main difficulty in practice is selecting a prior, as is often the case in Bayesian analyses. In general, choosing a prior with larger variance will allow us to find support for a larger set of true models, but required sample size to find this support will increase. As is shown in the simulation, some default options perform quite well in common research settings. In practice, it is often advisable to perform a prior sensitivity analysis.

Although the philosophical underpinnings of Bayesian hypothesis testing are not the focus of this paper, we will shortly connect the current work with the ongoing discussion. The Bayesian framework is sometimes touted as inductive, which would suggest Bayesian model comparison is sufficient to draw scientific conclusions from data. Recently, Gelman and Shalizi (2013) refute this claim outright and advocate model checking, as models are usually wrong. We generally follow the view of Morey et al. (2013) and note that tools developed here are useful to give preference between models, but do not necessarily provide inductive evidence in favor of the model assessed, such as the von Mises model. The kernel density alternative presented in Section 4 functions as a form of model checking, circumventing the step of deciding on test statistics to be used in a posterior predictive check, or deciding on a specific alternative hypothesis to test against.

Finally, the approach of this paper is to apply Bayesian hypothesis testing to basic circular data analyses. Future work might attempt to obtain easily computable marginal likelihoods for more complex models. In circular data analysis, model selection is an important avenue that requires more attention.

Acknowledgments

This work was supported by a Vidi grant awarded to I. Klugkist from NWO, The Netherlands, the Dutch Organization for Scientific Research (NWO 452-12-010).

References

Ajne, B., 1968. A simple test for uniformity of a circular distribution. Biometrika 55 (2), 343–354.

Arnold, B.C., SenGupta, A., 2006. Recent advances in the analyses of directional data in ecological and environmental sciences. Environ. Ecol. Stat. 13 (3), 253–256.

Baayen, C., Klugkist, I., Mechsner, F., 2012. A test of order-constrained hypotheses for circular data with applications to human movement science. J. Motor Behav. 44 (5), 351–363.

Bhattacharya, S., SenGupta, A., 2009. Bayesian inference for circular distributions with unknown normalising constants. J. Statist. Plann. Inference 139 (12), 4179–4192.

Bhattacharyya, G.K., Johnson, R.A., 1969. On Hodges's bivariate sign test and a test for uniformity of a circular distribution. Biometrika 56 (2), 446–449.

Bogdan, M., Bogdan, K., Futschik, A., 2002. A data driven smooth test for circular uniformity. Ann. Inst. Statist. Math. 54 (1), 29–44.

Brazier, K.T.S., 1994. Confidence intervals from the Rayleigh test. Mon. Not. R. Astron. Soc. 268 (3), 709–712.

Bulbert, M.W., Page, R.A., Bernal, X.E., 2015. Danger comes from all fronts: predator-dependent escape tactics of Túngara frogs. PLoS One 10 (4), e0120546.

Di Marzio, M., Panzera, A., Taylor, C.C., 2009. Local polynomial regression for circular predictors. Statist. Probab. Lett. 79 (19), 2066–2075.

Di Marzio, M., Panzera, A., Taylor, C.C., 2013. Non-parametric regression for circular responses. Scand. J. Stat. 40 (2), 238–255.

Fisher, N.I., 1995. Statistical Analysis of Circular Data. Cambridge University Press, Cambridge.

Gelman, A., Shalizi, C.R., 2013. Philosophy and the practice of Bayesian statistics. Br. J. Math. Stat. Psychol. 66 (1), 8–38.

George, B.J., Ghosh, K., 2006. A semiparametric Bayesian model for circular-linear regression. Comm. Statist. Simulation Comput. 35 (4), 911–923.

Gill, J., Hangartner, D., 2010. Circular data in political science and how to handle it. Political Anal. 18 (3), 316–336.

Gurtman, M.B., 2009. Exploring personality with the interpersonal circumplex. Soc. Personal. Psychol. Compass 3 (4), 601–619.

Gurtman, M.B., Pincus, A.L., 2003. The circumplex model: Methods and research applications. In: Handbook of Psychology. Wiley Online Library.

Guttorp, P., Lockhart, R.A., 1988. Finding the location of a signal: A Bayesian analysis. J. Amer. Statist. Assoc. 83 (402), 322–330.

Hall, P., Watson, G.S., Cabrera, J., 1987. Kernel density estimation with spherical data. Biometrika 74 (4), 751–762.

Hermans, M., Rasson, J.P., 1985. A new Sobolev test for uniformity on the circle. Biometrika 72 (3), 698–702.

Hodges, J.L., 1955. A bivariate sign test. Ann. Math. Stat. 26 (3), 523–527.

Jeffreys, H., 1961. Theory of Probability. Clarendon Press, Oxford.

Kaas, A.L., Van Mier, H.I., 2006. Haptic spatial matching in near peripersonal space. Exp. Brain Res. 170 (3), 403–413.

Kass, R.E., Raftery, A.E., 1995. Bayes factors. J. Amer. Statist. Assoc. 90 (430), 773–795.

Kuiper, N.H., 1960. Tests concerning random points on a circle. In: Indagationes Mathematicae (Proceedings), Vol. 63. Elsevier, pp. 38–47.

Landler, L., Ruxton, G.D., Malkemper, E., 2018. Circular data in biology: advice for effectively implementing statistical procedures. Behav. Ecol. Sociobiol. 72 (8), 128.

Laubscher, N.F., Rudolph, G.J., 1968. A Distribution Arising from Random Points on the Circumfence of a Circle. National Research Institute for Mathematical Sciences (South Africa).

Leary, T., 1957. Interpersonal Diagnosis of Personality. Ronald Press, New York.

Mardia, K.V., 2011. How new shape analysis and directional statistics are advancing modern life-sciences. In: Int. Statistical Inst.: Proc, 58th World Statistical Congress.
Mardia, K.V., Hughes, G., Taylor, C.C., Singh, H., 2008. A multivariate von Mises distribution with applications to bioinformatics. Canad. J. Statist. 36 (1), 99–109.
Mardia, K.V., Jupp, P.E., 2000. Directional Statistics, Vol. 494. John Wiley & Sons.
McVinish, R., Mengersen, K., 2008. Semiparametric Bayesian circular statistics. Comput. Statist. Data Anal. 52 (10), 4722–4730.
Mechsner, F., Kerzel, D., Knoblich, G., Prinz, W., 2001. Perceptual basis of bimanual coordination. Nature 414 (6859), 69–73.
Mechsner, F., Stenneken, P., Cole, J., Aschersleben, G., Prinz, W., 2007. Bimanual circling in deafferented patients: Evidence for a role of visual forward models. J. Neuropsychol. 1 (2), 259–282.
Morey, R.D., Romeijn, J.-W., Rouder, J.N., 2013. The humble Bayesian: model checking from a fully Bayesian perspective. Br. J. Math. Stat. Psychol. 66 (1), 68–75.
Nuñez-Antonio, G., Ausín, M.C., Wiper, M.P., 2015. Bayesian nonparametric models of circular variables based on Dirichlet process mixtures of normal distributions. J. Agric. Biol. Environ. Stat. 20 (1), 47–64.
Oliveira, M.a., Crujeiras, R., Rodrí guez Casal, A., 2014. NPCirc: An R package for nonparametric circular methods. J. Stat. Softw. 61 (1), 1–26. http://dx.doi.org/10.18637/jss.v061.i09, https://www.jstatsoft.org/index.php/jss/article/view/v061i09.
Oliveira, M., Crujeiras, R.M., Rodríguez-Casal, A., 2012. A plug-in rule for bandwidth selection in circular density estimation. Comput. Statist. Data Anal. 56 (12), 3898–3908.
Postma, A., Zuidhoek, S., Noordzij, M.L., Kappers, A.M.L., 2008. Keep an eye on your hands: on the role of visual mechanisms in processing of haptic space. Cogn. Process. 9 (1), 63–68.
Pycke, J.-R., 2010. Some tests for uniformity of circular distributions powerful against multimodal alternatives. Canad. J. Statist. 38 (1), 80–96.
Rao, J.S., 1976. Some tests based on arc-lengths for the circle. Sankhyā 329–338.
Ravindran, P., Ghosh, S.K., 2011. Bayesian analysis of circular data using wrapped distributions. J. Stat. Theory Pract. 5 (4), 547–561.
Schmidt-Koenig, K., 1963. On the role of the loft, the distance and site of release in pigeon homing (the "cross-loft experiment"). Biol. Bull. 125 (1), 154–164.
Watson, G.S., 1961. Goodness-of-fit tests on a circle. Biometrika 48 (1/2), 109–114.