Journal of Applied Statistics, Vol. 26, No. 4, 1999, 423–434

Change-point problems for the von Mises distribution

KAUSHIK GHOSH[1], **S. RAO JAMMALAMADAKA**[1] & **MANGALAM VASUDAVEN**[1,2], [1]*University of California, Santa Barbara, USA and* [2]*Curtin University of Technology, Perth, Australia*

ABSTRACT *A generalized likelihood ratio procedure and a Bayes procedure are considered for change-point problems for the mean direction of the von Mises distribution, both when the concentration parameter is known and when it is unknown. These tests are based on sample resultant lengths. Tables that list critical values of these test statistics are provided. These tests are shown to be valid even when the data come from other similar unimodal circular distributions. Some empirical studies of powers of these test procedures are also incorporated.*

1 Introduction

Suppose that we have a set of independent and identically distributed measurements on two-dimensional directions, say $\alpha_1, \alpha_2, \ldots, \alpha_n$. These measurements, called angular or circular data, can be represented as points on the circumference of a circle with unit radius. They may be wind directions, the vanishing angles at the horizon for a group of birds or the times of arrival at a hospital emergency room, where the 24-hour cycle is represented as a circle. Such data may have one or more peaks, or may show no preferred direction, corresponding to an isotropic distribution. We are interested in studying whether there has been a change in the preferred direction in the time-ordered data set, where the location of the change-point, if any, is unknown.

As an example, consider a meteorologist studying wind directions. Using previously gathered data, the meteorologist might be interested in knowing if there has been a change in the direction of the wind flow at some intermediate time point. Another interesting example is given in Lombard (1986), concerning the evaluation

Correspondence: K. Ghosh, Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106, USA.

0266-4763/99/040423-12 $7.00

of flares. Flares are launched upward attached to a projectile from a point O in a fixed direction. The quantity of interest is the latitude θ of the vector $\overrightarrow{OP}$, where P is the point at which the flare starts to burn. The variability of θ represents the stability of the flare-projectile assembly.

Formally, let $\alpha_1, \ldots, \alpha_n$ be angular measurements measured in a time-ordered or space-ordered sequence. Assume that, for some unknown (but fixed) k, $(1 \leqslant k \leqslant n)$, $\alpha_1, \ldots, \alpha_k \sim F_1$ and $\alpha_{k+1}, \ldots, \alpha_n \sim F_2 (\neq F_1)$. Here, k is called the change-point of the data. If $k = n$, then there are no observations from F_2, meaning that all the observations are from the same population so that there is no change-point. We are interested in testing for the presence of a change-point. Hence, we are testing H_0: $k = n$ vs H_1: $1 \leqslant k \leqslant n - 1$. For concreteness and simplicity, we assume that the two populations F_1 and F_2 follow the von Mises distribution (sometimes also known as the circular normal distribution) with a common concentration parameter κ, and that the mean directions are given by μ_1 and μ_2 respectively.

Change-point problems have evoked a considerable amount of interest in the statistical community, as a result of their application-oriented flavor. The reader is referred to Sen and Srivastava (1975), Horváth (1993), Gombay and Horváth (1990) and Page (1955) for a good exposition on the subject. All these authors (and many others) deal with the problem for data on the (real-) line. The interest in change-point problems for directional data is relatively new. Fisher (1993) presents some discussion on this topic, as do Lombard (1986) and Csörgő and Horváth (1996)—all of which deal with non-parametric methods. No work appears to have been done on parametric tests for the change-point problem on a circle.

Section 2 deals with the derivation of tests when κ is known and when κ is unknown. In both cases, we use the generalized likelihood ratio method to derive tests for H_0 vs H_1. We obtain the exact critical values of the test statistics, through simulation.

An alternative method, with a Bayesian flavor, assumes that the change-point is equally likely to be at any one of the intermediate points. Hence, using a discrete uniform prior over the possible change-point values, we obtain an alternate statistic. If we have further information about the possible point of change, we can incorporate that into an appropriate prior on k and derive the corresponding Bayes procedure.

The results for the critical values of the test statistics, obtained through simulations, are presented in Section 3. They are provided as ‘nomograms’ from which one can read the 5% values. The authors may be contacted for the code, if other values are of interest. In Section 4, we analyze a real data set using our tests and, in Section 5, we compare the powers of the two procedures. Finally, in Section 6, we ascertain the model robustness of these procedures to the von Mises distributional assumption.

2 The tests

We say that a random angle α has a von Mises distribution with mean direction μ and concentration parameter κ (i.e. $vM(\mu, \kappa)$) if it has the probability density function

$$f(\alpha) = \frac{1}{2\pi I_0(\kappa)} \exp[\kappa \cos(\alpha - \mu)], \qquad -\pi \leqslant \mu, \alpha < \pi, \ \kappa > 0$$

The change-point in the mean direction at k implies that $\alpha_1, \ldots, \alpha_k \sim vM(\mu_1, \kappa)$ and that $\alpha_{k+1}, \ldots, \alpha_n \sim vM(\mu_2, \kappa)$.

2.1 κ known case

If $\boldsymbol{\theta} = (k, \mu_1, \mu_2)$ denotes the parameter vector for our problem, then the parameter space is $\boldsymbol{\Omega} = \{1, \ldots, n\} \times [-\pi, \pi) \times [-\pi, \pi)$. Technically, we might say that the hypothesis of no change corresponds to the change-point at n, so that, under H_0, the parameter space becomes $\omega = \boldsymbol{\Omega} \quad H_0 = \{n\} \times [-\pi, \pi) \times [-\pi, \pi)$. If the change in mean takes place at the point k, then the likelihood function is given by

$$L(\boldsymbol{\theta}) = \frac{1}{[2\pi I_0(\kappa)]^n} \exp\left\{\kappa\left[\sum_{i=1}^{k} \cos(\alpha_i - \mu_1) + \sum_{i=k+1}^{n} \cos(\alpha_i - \mu_2)\right]\right\} \tag{1}$$

Let $(R_{1k}, \bar{\alpha}_{1k})$, $(R_{2k}, \bar{\alpha}_{2k})$ and $(R, \bar{\alpha}_0)$ denote the length and direction of the resultant of the subsets $(\alpha_1, \ldots, \alpha_k)$, $(\alpha_{k+1}, \ldots, \alpha_n)$ and the combined sample $(\alpha_1, \ldots, \alpha_n)$ respectively. It can easily be verified that, under H_0, the maximum-likelihood estimate (MLE) of μ satisfies

$$\sum_{i=1}^{n} \sin(\alpha_i - \hat{\mu}_0) = 0\,. \tag{2}$$

Similarly, under H_1, for a given k, the MLEs of μ_1 and μ_2 satisfy

$$\sum_{i=1}^{k} \sin(\alpha_i - \hat{\mu}_{1k}) = 0 \tag{3}$$

and

$$\sum_{i=k+1}^{n} \sin(\alpha_i - \hat{\mu}_{2k}) = 0\,. \tag{4}$$

The solutions to equations (2)–(4) are given by $\hat{\mu}_0 = \bar{\alpha}_0$, $\hat{\mu}_{1k} = \bar{\alpha}_{1k}$ and $\hat{\mu}_{2k} = \bar{\alpha}_{2k}$ respectively.

Since

$$\begin{aligned}\sum_{i=1}^{n} \cos(\alpha_i - \bar{\alpha}_0) &= \sum_{i=1}^{n} \cos\alpha_i \cos\bar{\alpha}_0 + \sum_{i=1}^{n} \sin\alpha_i \sin\bar{\alpha}_0 \\ &= R\cos\bar{\alpha}_0 \cos\bar{\alpha}_0 + R\sin\bar{\alpha}_0 \sin\bar{\alpha}_0 \\ &= R\,,\end{aligned} \tag{5}$$

the likelihood ratio becomes

$$\begin{aligned}\Lambda_k &= \exp\left\{\kappa\left[\sum_{i=1}^{n} \cos(\alpha_i - \bar{\alpha}_0) - \sum_{i=1}^{k} \cos(\alpha_i - \bar{\alpha}_{1k}) - \sum_{i=k+1}^{n} \cos(\alpha_i - \bar{\alpha}_{2k})\right]\right\} \\ &= \exp[\kappa(R - R_{1k} - R_{2k})]\end{aligned}$$

The validity of H_0 is questioned if Λ_k is 'sufficiently small'. However, since we do not know k, we have two options: one option is the frequentist approach, where we treat k also as a parameter as indicated and minimize the likelihood ratio Λ_k over k; the other option is to take a Bayesian approach and average Λ_k over a prior distribution on $\{1, \ldots, n\}$. In the first case, the likelihood ratio test (LRT) of H_0 vs H_1 would reject H_0 for small values of $\inf_k \Lambda_k$ or, equivalently, since κ is known, it would reject H_0 for large values of

$$\Lambda = \sup_{\kappa \in \{1, \ldots, n\}} (R_{1k} + R_{2k}) - R\,.$$

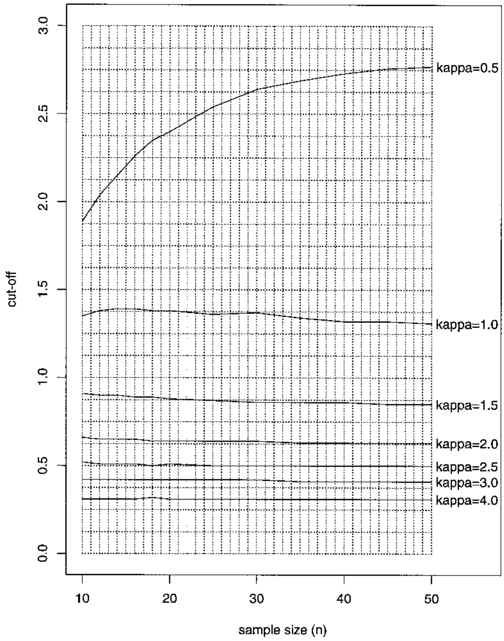

FIG. 1. 5% cut-offs of the avg statistic, κ known.

This leads to rejecting H_0 when

$$\Lambda > c \tag{6}$$

where the cut-off point c is determined based on the significance level of the test. In all further discussion, we refer to this test as the supremum or 'sup' test.

In contrast, for our Bayesian analog, assuming a uniform prior on the possible values of k, we reject H_0 whenever

$$\frac{1}{n}\sum_{k=1}^{n}(R_{1k}+R_{2k}) - R > c' \tag{7}$$

where c' is also determined based on the significance level. We refer to this test as the average or 'avg' test.

The joint distribution of (R_{1k}, R_{2k}) is known (see, for example, Mardia, 1972, p. 97) for any fixed k and is indeed, independent of κ, conditional on R. However, the joint distributions of $(R_{1k}, R_{2k})_{k=1}^{n}$, on which the exact theory of 'sup' and 'avg' depend, do not appear to have any reasonable analytic form, partly because they are not independent for different k. The only solution appears to be to simulate their cut-off points for various combinations of κ and n, which is what we do. The results are represented graphically in Figs 1 and 2.

2.2 *κ unknown case*

Here, we have a four-dimensional parameter $\theta = (k, \mu_1, \mu_2, \kappa)$ with the parameter space $\Omega = \{1, \ldots, n\} \times [-\pi, \pi) \times [-\pi, \pi) \times (0, \infty)$. Arguing as before, under H_0, the parameter space becomes $\omega = \Omega \;\; H_0 = \{n\} \times [-\pi, \pi) \times [-\pi, \pi) \times (0, \infty)$.

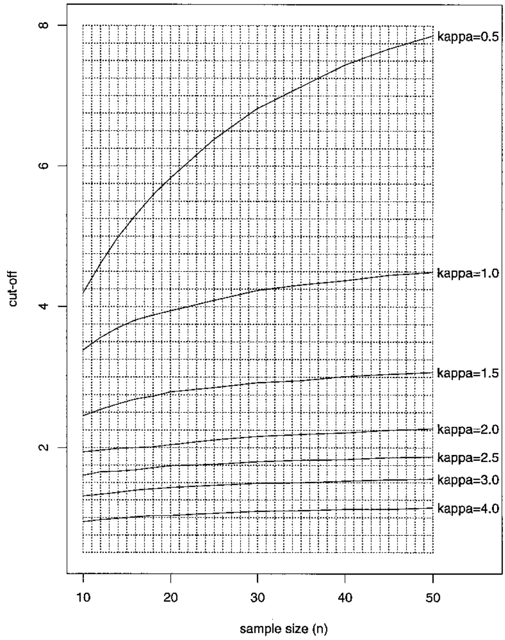

FIG. 2. 5% cut-offs of the sup statistic, κ known.

The likelihood function for the data is given by equation (1) with the new parameter vector $\boldsymbol{\theta}$. It can easily be verified that, under H_0, the MLEs of μ and κ satisfy equation (2) and

$$\frac{I_1(\hat{\kappa}_0)}{I_0(\hat{\kappa}_0)} = \frac{R}{n}. \tag{8}$$

Similarly, under H_1 for a given k, the MLEs satisfy equations (3), (4) and

$$\frac{I_1(\hat{\kappa}_k)}{I_0(\hat{\kappa}_k)} = \frac{R_{1k} + R_{2k}}{n}. \tag{9}$$

Again, the solutions to equations (3) and (4) are given by $\hat{\mu}_{1k} = \bar{\alpha}_{1k}$ and $\hat{\mu}_{2k} = \bar{\alpha}_{2k}$ respectively.

Hence, the likelihood ratio of the data (for a given k) is

$$\Lambda_k = \left[\frac{I_0(\hat{\kappa}_k)}{I_0(\hat{\kappa}_0)}\right]^n \exp\left\{\hat{\kappa}_0 \sum_{i=1}^{n} \cos(\alpha_i - \bar{\alpha}_0) - \hat{\kappa}_k \left[\sum_{i=1}^{k} \cos(\alpha_i - \bar{\alpha}_{1k}) + \sum_{i=k+1}^{n} \cos(\alpha_i - \bar{\alpha}_{2k})\right]\right\}$$

$$= \left[\frac{I_0(\hat{\kappa}_k)}{I_0(\hat{\kappa}_0)}\right]^n \exp[\hat{\kappa}_0 R - \hat{\kappa}_k (R_{1k} + R_{2k})] .$$

This gives

$$\begin{aligned} \lambda_k &= -\log \Lambda_k \\ &= n\{\log[I_0(\hat{\kappa}_0)] - \log[I_0(\hat{\kappa}_k)]\} + [\hat{\kappa}_k (R_{1k} + R_{2k}) - \hat{\kappa}_0 R] \\ &= n[\Psi(R_{1k} + R_{2k}) - \Psi(R)] \end{aligned}$$

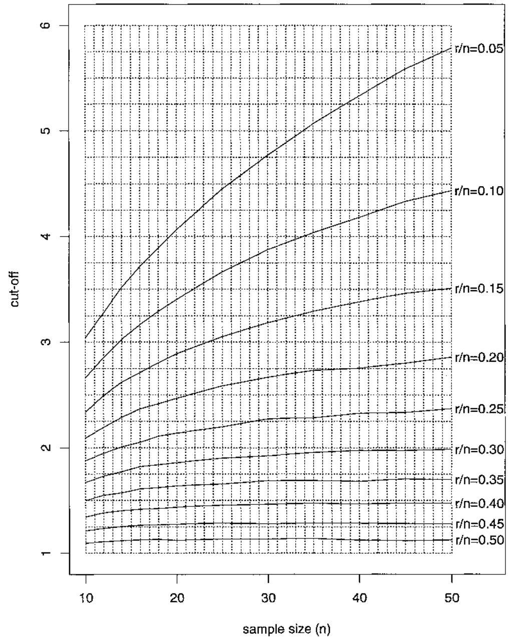

FIG. 3. 5% cut-offs of the avg statistic, κ unknown.

where

$$\Psi(t) = tA^{-1}(t) - \log\{I_0[A^{-1}(t)]\}$$

and $A(\cdot)$ is defined as

$$A(t) = \frac{I_1(t)}{I_0(t)}.$$

As before, since k is unknown, we employ the sup and avg methods to come up with two test criteria, i.e. $\sup_k \Psi(R_{1k} + R_{2k}) - \Psi(R)$ and

$$\frac{1}{n}\sum_{k=1}^{n} \Psi(R_{1k} + R_{2k}) - \Psi(R)$$

In both the cases, however, the null distribution of the test statistic depends on (the unknown) κ. To make our tests κ-free, we condition on the overall resultant length R. Making use of the monotonic nature of $\Psi(\cdot)$ (see Mardia, 1972, p. 134), the two (conditional) tests then become the same as the sup and avg tests given in equations (6) and (7) respectively. The only difference from the κ known case is that the test statistic values are conditional on R. The cut-offs of these two conditional tests are determined based on the significance level.

Simulation results for the 5% cut-off points of the two conditional tests are given in Figs 3 and 4.

3 Simulation results

The tests proposed in the previous sections have no simple known distributional form. Thus, to obtain their cut-off values, we resorted to large-scale Monte Carlo

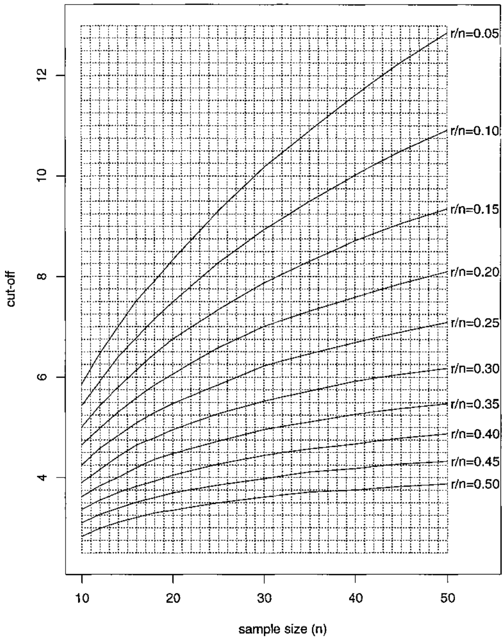

FIG. 4. 5% cut-offs of the sup statistic, κ unknown.

simulations. All the codes were written in the C language, with calls to the IMSL/C/STAT library for the random number generators. In particular, we made extensive use of the routine imsls_f_random_von_mises to generate all the von Mises random deviates.

For the κ known case, we sampled from a von Mises distribution with center zero and concentration κ (i.e. $vM(0, \kappa)$). We considered $\kappa = 0.5(0.5)3(1)4$ and $n = 10(2)20(5)50$. For each (n, κ) combination, we carried out 100 000 simulations to obtain upper 5% points of the two tests. The results appear in Figs 1 and 2. These figures, called 'nomograms', list the cut-offs at a specific κ value along each line.

For the κ unknown case, we sampled from a conditional von Mises distribution, the conditioning event being the given length r of the resultant R. Each (r, n) combination results in a different distribution and we considered $n = 10(2)20(5)50$ and $r/n = 0.05(0.05)0.95$. Since the conditional sampling discards many of the random numbers for not meeting the conditioning criterion, the sampling procedure was much slower compared with the unconditional case. Apart from using imsls_f_random_von_mises, we used imsls_f_random_binomial to draw the conditional samples. The results of 100 000 simulations appear as nomograms in Figs 3 and 4.

Casual examination of the nomograms suggest that high κ values for the unconditional case (and, correspondingly, high r/n values for the conditional case) make the test statistics (and hence, the cut-offs) free of n. This is reflected by the almost horizontal lines in the corresponding figures. In particular, for the avg statistic in the κ known case, the $100(1-\alpha)\%$ cut-offs are approximated very well by

$$\frac{1}{2\kappa}\chi^2_{1;\alpha}$$

when $\kappa \geqslant 2$.

This may be explained by the fact that, for large κ and any fixed k, $2\kappa(n - R)$ has an approximate χ^2_{n-1} distribution. Hence, for each k, we have

$$2\kappa(R_{1k} + R_{2k} - R) \sim \chi^2_1$$

When we take the average of this over k, we obtain an average of dependent χ^2_1 and the result is approximated quite well by a χ^2_1 random variable. There does not seem to be any easy explanation for the sup statistic.

A similar phenomenon (albeit less pronounced) occurs for the κ unknown case. Then, for high r/n, the cut-off for avg does not change greatly with changing n (r/n measures the unknown concentration in this case).

4 An example

To see how well our tests work, we used the following data from Schmidt-Koenig (1958).

Example 1. In an experiment on pigeon-homing, the internal clocks of 10 birds were reset by 6 hours clockwise, while the clocks of nine birds were left unaltered. It is predicted from sun–azimuth compass theory that the mean direction of the vanishing angles in the experimental group should deviate by about 90° in the anti-clockwise direction with respect to the mean direction of the angles of the birds in the control group. The vanishing angles of the birds for this experiment are as follows, measured (in degrees) in the clockwise sense:

control group—75, 75, 80, 80, 80, 95, 130, 170, 210
experimental group—10, 50, 55, 55, 65, 90, 285, 285, 325, 355

Do the data support sun–azimuth compass theory?

To see how well our test procedures work, we consider the combined time-ordered sample of 19 observations. For this sample, we test for the presence of a change-point, if any. Since κ is unknown, we use the conditional tests. Calculations yield avg $= 1.745\,029$, sup $= 5.289\,31$ and $r/n = 0.48$. Consulting Figs 3 and 4, we see that both our tests reject the null hypothesis at the 1% level. This is not surprising, in view of the fact that Mardia (1972, p. 156) considers the same example and tests for the equality of mean directions, assuming that the two samples come from von Mises distributions with equal (but unknown) concentration. His test shows that the populations are indeed different.

5 Power comparisons

Using Monte Carlo simulations, a power comparison of the two statistics was made for the κ known case. All the results here are based on 1000 simulations each. We considered $n = 10$, 15 and 20 with change points at $k = 1$, $[n/4]$ and $[n/2]$ for each n, where $[x]$ denotes 'the greatest integer less than or equal to x'. We also chose $\kappa = 0.5$, 1, 2 and 4. Finally, the difference in the mean directions $\Delta = |\mu_1 - \mu_2|$ that we considered are $\pi/10(\pi/10)\pi$. Some of the graphs are presented in Figs 5–7. The results obtained may be summarized as follows.

(1) Both the tests are symmetric in the change-point k, in the sense that the power at k is approximately equal to the power at $n - k$, everything else being the same. This symmetry in our test procedures is to be expected.

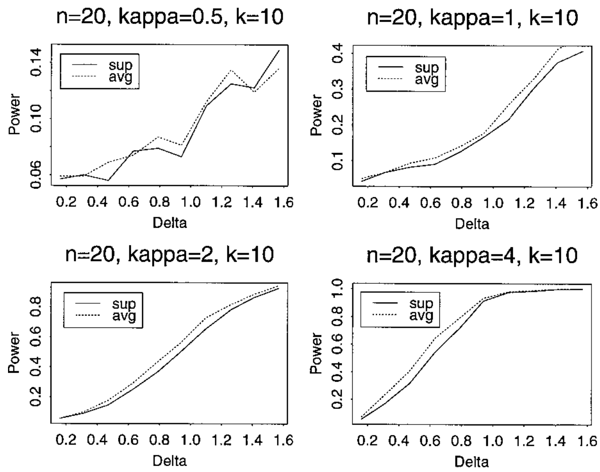

FIG. 5. Effect of κ on powers of the two tests.

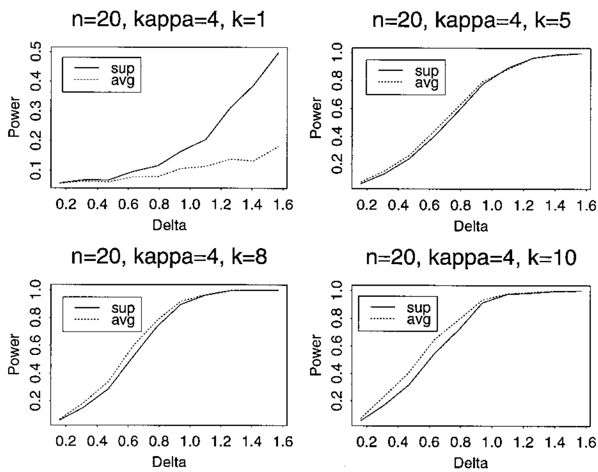

FIG. 6. Effect of k on powers of the two tests.

(2) As would be expected, the powers of both the statistics show an increasing trend as Δ increases, everything else being the same. The trend is not as pronounced for smaller values of the concentration parameter as in the case of larger values.

(3) For a high concentration parameter κ, the average statistic is more powerful than the sup statistic if the change-point is near the center (i.e. $k \approx n/2$). The maximum power reached in this case is almost 1 for both statistics. However, if the change point is near the end ($k \approx 1$), then the sup statistic becomes more powerful for high concentrations. The maximum power reached in this case is about 50% for the sup statistic compared with 20% for the avg statistic.

Hence, as intuition suggests, it is easier to detect a change-point in the middle than at the end. If one suspects change at the end points, then the sup statistic is recommended.

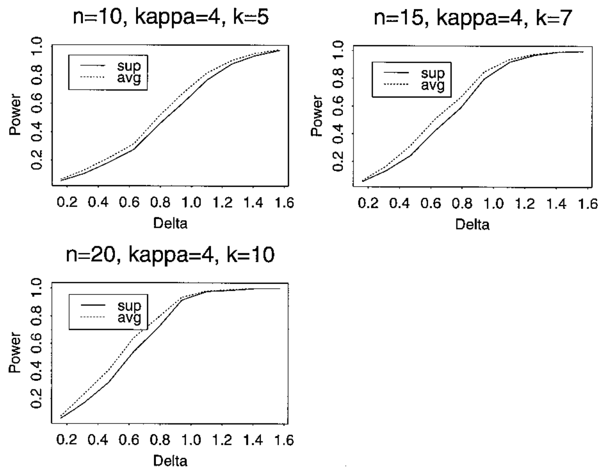

FIG. 7. Effect of n on powers of the two tests.

(4) For $k \approx n/4$, both the tests behave very similarly, having almost the same power curve, irrespective of n and κ.
(5) If there is a very low concentration in the data, then neither test is useful, since the maximum power obtained is as low as 15%, even for $n = 20$. Thus, trying to detect changes in the mean direction when the distributions are nearly uniform is futile.
(6) For a given κ, the power increases with n, as long as the k/n ratio remains constant. Hence, an increase in sample size means an increase in power.

Based on these findings, we would recommend the use of the sup statistic if we have some previous knowledge that the change-point occurred at the end and that the concentration is high ($\kappa > 3.5$). If we think that the change-point is in the middle and that the concentration is moderate to high, then we recommend using the avg statistic. If the concentration is low ($\kappa < 2$), then neither statistic is useful.

A power comparison was made using similar Monte Carlo techniques for the κ unknown case, and the results are a repetition of the κ known case, with r/n taking the role of κ. Hence, we do not give any separate pictures for this case.

6 Robustness

The simulation results we obtained depend entirely on the von Mises assumption for the distribution of the data. To investigate the effect of misspecification of the distribution, we carried out simulations with data from wrapped normal and wrapped Cauchy distributions. The cut-off points for the wrapped normal, wrapped Cauchy and von Mises distributions for $n = 20$ and $\kappa = 0.5(0.5)3(1)4$ data are given in Table 1. For each (n, κ) combination, we carried out 100 000 simulations. For low values of κ ($\kappa = 0.5$), the von Mises distribution gives results close to those with the wrapped Cauchy distribution. However, for higher κ, the von Mises results are closer to those with the wrapped normal distribution. Also, of the two tests, the avg statistic seems to be more robust than the sup statistic.

Suppose that we have angular data and we want to see if there is a change in the preferred mean direction at some intermediate (unknown) point, without making

TABLE 1. 5% (1%) cut-offs of the test statistics for the unconditional case, $n = 20$

	Wrapped normal		Wrapped Cauchy		von Mises	
	avg.	sup.	avg.	sup.	avg.	sup.
0.5	3.01 (4.01)	6.77 (8.41)	2.42 (3.46)	5.85 (7.56)	2.40 (3.42)	5.82 (7.56)
1.0	1.37 (2.10)	3.90 (5.33)	1.39 (2.19)	4.02 (5.65)	1.37 (2.12)	3.93 (5.44)
1.5	0.91 (1.39)	2.77 (3.79)	0.85 (1.35)	2.92 (4.07)	0.88 (1.36)	2.78 (3.83)
2.0	0.67 (1.04)	2.07 (2.87)	0.58 (0.95)	2.11 (3.41)	0.65 (1.00)	2.05 (2.93)
2.5	0.53 (0.82)	1.66 (2.29)	0.44 (0.72)	1.96 (2.83)	0.51 (0.79)	1.74 (2.30)
3.0	0.43 (0.66)	1.37 (1.88)	0.35 (0.57)	1.86 (2.40)	0.42 (0.66)	1.43 (1.95)
4.0	0.32 (0.49)	1.02 (1.39)	0.25 (0.42)	1.64 (2.00)	0.31 (0.48)	1.04 (1.48)

any assumptions about these two distributions. In other words, we wish to determine how justified the tests proposed here are in the absence of any parametric model. Although we depended heavily on the exact form of the parent distribution (von Mises) to derive these tests, we believe that these tests are justified when the underlying models are unimodal, for which the resultant length and direction are relevant measures. Since a good measure of the distance between two angles x and y in directional data is given by $1 - \cos(x - y)$ corresponding to $(x - y)^2$ for linear data, $(n - R)$ is a measure of dispersion of a set of n angles. If there is a change-point at k, then the overall dispersion might be large, while the dispersions of the two individual groups are expected to be small. Hence, a statistic to detect change-point is

$$(n - R) - [(k - R_{1k}) + (n - k - R_{2k})] = R_{1k} + R_{2k} - R$$

This is the same idea as used in the approximate analysis of variance for circular data (see Harrison *et al.*, 1986; Rao & Sengupta, 1970; Watson, 1966). Since k is unknown, we average over k or take the supremum over k, as we have done to obtain our tests.

Of course, the test statistic cut-offs cannot be obtained as before by simulation; we might have to resort to bootstrapping for them.

7 Concluding remarks

This paper discusses some specific change-point problems encountered in directional data analysis. In particular, we assume that the change-point is discrete and that the observations are all independent. A more practical scenario would be when the two assumptions do not hold—but the analysis becomes much more difficult in such cases. We plan to investigate further these cases in a later work.

Acknowledgements

The authors are grateful to the referee for valuable comments, the incorporation of which considerably improved the presentation of the paper. S. Rao Jammalamadaka was supported in part by NSF grant DMS-9803600.

REFERENCES

CSÖRGŐ, M. & HORVÁTH, L. (1996) A note on the change-point problem for angular data, *Statistics and Probability Letters*, 27, pp. 61–65.

FISHER, N. I. (1993) *The Statistical Analysis of Circular Data* (Cambridge, Cambridge University Press).

GOMBAY, E. & HORVÁTH, L. (1990) Asymptotic distributions of maximum likelihood tests for the change in the mean, *Biometrika*, 77, pp. 411–414.

HARRISON, D., KANJI, G. K. & GADSDEN, R. J. (1986) Analysis of variance for circular data, *Journal of Applied Statistics*, 13, pp. 197–223.

HORVÁTH, L. (1993) The maximum likelihood method for testing changes in the parameters of normal observations, *The Annals of Statistics*, 21, pp. 671–680.

LOMBARD, F. (1986) The change-point problem for angular data: a nonparametric approach, *Technometrics*, 28, pp. 391–397.

MARDIA, K. V. (1972) *Statistics of Directional Data*, 2nd Edn (London, Academic Press).

PAGE, E. S. (1955) A test for a change in a parameter occurring at an unknown point, *Biometrika*, 43, pp. 523–526.

RAO, J. S. & SENGUPTA, S. (1970) An optimum hierarchical sampling procedure for cross-bedding data, *The Journal of Geology*, 78, pp. 533–544.

SCHMIDT-KOENIG, K. (1958) Experimentelle Einflussnahme auf die 24-Stunden-Periodik bei Brieftauben und deren Auswirkungen unter besonderer Berucksichtigung des Heimfindevermogens, *Z. Tierpsychol*, 15, pp. 301–331.

SEN, A. & SRIVASTAVA, M. S. (1975) On tests for detecting change in mean, *Annals of Statistics*, 3, pp. 98–108.

WATSON, G. S. (1966) The statistics of orientation data, *Journal of Geology*, 74, pp. 786–797.