Statistics and Probability Letters 108 (2016) 40–44

Contents lists available at ScienceDirect

Statistics and Probability Letters

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

A nonparametric test of stationarity for independent data

Jeffrey D. Hart

Texas A&M University, Department of Statistics, United States

ARTICLE INFO

Article history:
Received 11 September 2015
Accepted 11 September 2015
Available online 9 October 2015

Keywords:
Randomness statistics
Consistent test
Local alternatives

ABSTRACT

A nonparametric test of stationarity for independent data is investigated. The test is based on comparing kernel density estimates calculated from subsamples of the data. Asymptotic distribution theory is developed and results of a modest simulation study are presented.

1. Introduction

A common problem in statistics is trying to determine whether or not data have a common distribution. This problem arises in quality control, for example, where one wishes to be able to detect some change in a process. Suppose one observes independent random variables $X_1, \ldots, X_N$ and wants to test whether or not they are stationary, i.e., whether or not they have a common distribution. A number of methods exist for doing so. One class of methods falls under the heading "change-point detection". These methods seek to identify an abrupt change in a sequence of observations. The cusum test proposed by Page (1954) is a long-standing tool for detecting change-points. It is designed mainly for detecting level changes in the observed process. Methods for detecting more general types of change, such as in variation or skewness, have also been proposed. These include so-called randomness statistics (McDonald, 1991) and the stationarity tests of Kapetanios (2007), Busetti and Harvey (2010) and Lima and Neri (2013).

The test of stationarity proposed in this paper is an adaptation of the test of Zhan and Hart (2014). Their test is for a setting where one has a large number of small data sets and wishes to test whether all these data sets come from the same distribution. In the current setting, one may partition the N observations into, say, p smaller data sets of equal size, and then apply the test of Zhan and Hart (2014) to these p sets. The statistic of Zhan and Hart (2014) is analogous to one proposed by Lehmann (1951). Let F^i be the empirical distribution function (edf) for the ith small data set, $i = 1, \ldots, p$, and let F_N be the edf for all N observations. Then the statistic of Lehmann (1951) is

$$\sum_{i=1}^{p} \int (F^i(x) - F_N(x))^2 \, dF_N(x). \qquad (1)$$

Lehmann (1951) and McDonald (1991) obtain the asymptotic distribution of (1) in the respective cases (a) N/p tends to ∞ with p fixed, and (b) p tends to ∞ with N/p fixed.

The statistic of Zhan and Hart (2014) is an analog of (1) that compares kernel density estimates rather than edfs. A number of authors, including Eubank et al. (1994), Martínez-Camblor and de Uña Álvarez (2009) and Rayner et al. (2009), have made the case that goodness-of-fit tests based on density estimates are generally more powerful than ones based on edfs.

E-mail address: hart@stat.tamu.edu.

http://dx.doi.org/10.1016/j.spl.2015.09.024
0167-7152/

The simulation study in Section 4 (of the current paper) suggests that a similar conclusion may be valid when testing for stationarity.

The rest of the paper may be outlined as follows. The test statistic and its asymptotic distribution are described in Section 2. Power properties against both fixed and local alternatives are developed in Section 3. A brief simulation study illustrating the potential power gain over edf-based statistics is presented in Section 4. Finally, concluding remarks and proofs of theoretical results are given in Section 5 and Appendix, respectively.

2. The test statistic and its null distribution

Let $\boldsymbol{X} = (X_1, \ldots, X_N)$ be independent observations such that X_i has density f_i, $i = 1, \ldots, N$. We wish to test the null hypothesis

$$H_0 : f_1 \equiv f_2 \equiv \cdots \equiv f_N \tag{2}$$

against the negation of H_0. Merely for the sake of notational simplicity, suppose that $np = N$ for integers $n > 1$ and p. We assume that n is fixed. Now divide the data set into p groups, the ith of which is

$$\boldsymbol{X}_i = (X_{n(i-1)+1}, \ldots, X_{ni}), \quad i = 1, \ldots, p.$$

Having grouped the data in this way, we may now test H_0 using methodology like that proposed in Zhan and Hart (2014).

The test of Zhan and Hart (2014) is, effectively, a comparison of the p data distributions defined by the above grouping. In a completely arbitrary setting where the alternative hypothesis is true, such a comparison would not necessarily be very powerful. However, the main application envisioned here is that the data are observed chronologically, and changes in densities are such that the difference between f_i and f_j tends to be larger for large values of $|i-j|$ than for smaller values of $|i-j|$.

For a data set $\boldsymbol{Y} = (Y_1, \ldots, Y_k)$, define the following kernel density estimate:

$$\hat{f}_h(x|\boldsymbol{Y}) = \frac{1}{kb}\sum_{i=1}^{k}\phi\left(\frac{x-Y_i}{b}\right),$$

where ϕ is the standard normal density and b is a positive bandwidth. Then our test statistic is a properly standardized version of

$$\sum_{i=1}^{p}\int_{-\infty}^{\infty}(\hat{f}_b(x|\boldsymbol{X}_i) - \hat{f}_b(x|\boldsymbol{X}))^2\,dx. \tag{3}$$

To define the standardization, let

$$S_W = \frac{1}{pn(n-1)\sqrt{2}b}\sum_{i=1}^{p}\sum_{j=1}^{n}\sum_{l=1,l\neq j}^{n}\phi\left(\frac{X_{n(i-1)+j} - X_{n(i-1)+l}}{\sqrt{2}b}\right),$$

$$S_B = \frac{1}{p(p-1)n^2\sqrt{2}b}\sum_{i=1}^{p}\sum_{k=1,k\neq i}^{p}\sum_{j=1}^{n}\sum_{l=1}^{n}\phi\left(\frac{X_{n(i-1)+j} - X_{n(k-1)+l}}{\sqrt{2}b}\right).$$

The proposed test statistic is

$$T_b = \frac{\sqrt{p}(S_W - S_B)}{\hat{\sigma}},$$

where $\hat{\sigma}^2/p$ is an estimator (to be defined subsequently) of $\mathrm{Var}(S_W - S_B)$ assuming that H_0 is true.

The reader is referred to Zhan and Hart (2014) for an explanation of why $S_W - S_B$ is a centered version of (3). When H_0 is true, $E(S_W - S_B) = 0$, and under the "smooth" alternatives defined in Section 3 $E(S_W - S_B) > 0$. Furthermore, under hypothesis (2) and with n and b fixed, the results of Zhan and Hart (2014) imply that T_b converges in distribution to a standard normal random variable as $N \to \infty$. The W and B on S_W and S_B stand for *within* and *between*, reflecting the fact that in S_W and S_B the argument of ϕ depends upon two observations that come from the same and different groups, respectively. Our test rejects H_0 at level α if and only if $T_b \geq z_\alpha$, where z_α is the $1-\alpha$ percentile of the standard normal distribution.

To define the variance estimator $\hat{\sigma}^2$, let

$$h_1(\boldsymbol{X}_i) = \frac{1}{n(n-1)\sqrt{2}b}\sum_{j=1}^{n}\sum_{l=1,l\neq j}^{n}\phi\left(\frac{X_{n(i-1)+j} - X_{n(i-1)+l}}{\sqrt{2}b}\right),$$

$$h_2(\boldsymbol{X}_i, \boldsymbol{X}_k) = \frac{1}{n^2\sqrt{2}b}\sum_{j=1}^{n}\sum_{l=1}^{n}\phi\left(\frac{X_{n(i-1)+j} - X_{n(k-1)+l}}{\sqrt{2}b}\right)$$