The Use of Canonical Correlation to Study Teleconnections

NEVILLE NICHOLLS

Bureau of Meteorology Research Centre, Melbourne, Australia

(Manuscript received 11 March 1986, in final form 7 August 1986)

ABSTRACT

Canonical correlation is proposed as an exploratory technique for studying teleconnections. It is suggested that the technique can elucidate the temporal signature (i.e., the seasonally varying nature) of the teleconnections and the lags between variables. Several teleconnections associated with the Southern Oscillation are subjected to canonical correlation as examples. The teleconnections studied are between Darwin pressure and Tahiti pressure, southeast Australian rainfall, and Willis Island air temperature. In each example the canonical correlation analysis confirms the teleconnections uncovered previously by other statistical techniques but also suggests the existence of other interesting features of these teleconnections.

1. Introduction

In recent decades there has been much research aimed at documenting and explaining the existence of teleconnections between climate fluctuations in widely separated areas. The strong and complex suite of teleconnections labeled the Southern Oscillation (SO) by Walker early this century has received particular attention. Walker and Bliss (1932) described the Southern Oscillation thus:

> When pressure is high in the Pacific Ocean it tends to be low in the Indian Ocean from Africa to Australia; these conditions are associated with low temperatures in both these areas, and rainfall varies in the opposite direction to pressure. Conditions are related differently in winter and summer.

Three aspects of the SO teleconnections of particular interest are their large-scale nature (they cover large areas of the globe), their seasonal variations, and the lags between the different variables in the teleconnection. A variety of statistical techniques have been used to elucidate the nature of these teleconnections. The techniques used most commonly are simple linear correlation (both simultaneous and lagged) and compositing.

Correlations between an SO index and seasonal means of different meteorological or oceanographic fields have been calculated by many investigators. Some studies (e.g., Egger et al., 1981) have examined only simultaneous relationships while others (e.g., Wright, 1977; Newell et al., 1982) have studied both simultaneous and lag relationships. This approach displays clearly the spatial nature of the teleconnections but produces an enormous set of results consisting of relationships at various lags for each season. The magnitude of this set of results complicates both the description and the understanding of the lagged relationships between the fields.

In the composite approach a set of years is selected according to some criterion, e.g., when an SO index exceeds a certain threshold, and the mean values of various fields over these years are calculated (e.g., Van Loon and Madden, 1981; Rasmusson and Carpenter, 1982; Pan and Oort, 1983). Wright et al. (1985) note that this approach is appropriate if the SO can be considered as a sequence of discrete "events" separated by periods within which the variations are of less interest. Compositing, unlike the correlation approach, does not assume that the relationships between the fields and the SO are linear. It is also simpler, usually, to display and interpret the lagged relationships through compositing than with the huge set of lagged correlations produced in the correlation approach. There is, however, a certain degree of arbitrariness in selecting the threshold to determine which events to be composited, and even in selecting the SO index to which a threshold will be applied in composite selection. Numerous SO indices exist (e.g., Darwin pressure or the difference in pressure between Tahiti and Darwin), all closely but not perfectly related, and different composites will result from selecting different indices or different thresholds.

Wright et al. (1985) performed an analysis that contained elements of both the composite and correlation approaches. Their approach was to calculate correlations and regressions between *seasonal* anomalies of meteorological fields and *12-month* average anomalies of an SO index. The 12-month average for the index was taken from April of one year through March of the next. This choice maximizes the year-to-year variability of the index and locates the months of strongest autocorrelation of the index (July to December) near the middle of the averaging period. Correlations between this 12-month average index and meteorological fields in six 3-month seasons were calculated. The first season was the December–January–February several months prior to the start of the SO index 12-month average and the last season was the March–April–May near the end of the 12-month average. This approach differed from most previous correlation studies in that it took account of the strong link between the teleconnections and the annual cycle. It differed from the composite approach by using all years, rather than just those in which a particular threshold was passed. The use of seasonal average anomalies may tend to smooth out some of the interesting behavior, such as the lags between the SO index and the various fields. This could be overcome, presumably, by using a shorter averaging period such as a month, although this would result in large numbers of complicated diagrams that may be difficult to interpret or present.

Several studies have used factor analysis or empirical orthogonal function (EOF) analysis to study the spatial nature of SO teleconnections (e.g., Barnett, 1977). Lanzante (1984) examined the joint associations between two fields, sea surface temperatures and 700 mb heights, by finding the eigenvalue solution of the mean product matrix of the cross correlations. This produced two sets of spatial patterns that explain part of the linear correlation between the two fields.

Extended or complex EOF analyses that take account of the temporal, as well as the spatial, nature of teleconnections have been used recently (e.g., Weare and Nasstrom, 1982; Barnett, 1983). These analyses appear to result in functions that are interpretable as dominant modes of space–time sequences of teleconnections. Such analysis techniques use all the data available, do not use arbitrary thresholds for selecting "events," and may detect propagating features, whereas conventional EOF analysis can detect standing oscillations only. Horel (1985) notes that complex EOF patterns may, however, be difficult to interpret, may result in a misleading reconstruction of data in some areas examined and that sudden transitions and noisy spikes may be emphasized unduly.

In an attempt to simplify the interpretation of teleconnections, some investigators have concentrated on presenting the relationships between SO indices and related variables at just a few "core" sites, rather than with global or ocean basin fields. Thus, Rasmusson and Carpenter (1982) presented composites for a few variables (e.g., Darwin pressure) at selected sites. After selecting the "events" to be composited, they calculated 3-month running means of the anomalies of Darwin pressure associated with these events, starting about 12 months prior to the event and continuing well after the completion of the event. In this way the cycle of Darwin pressure anomalies typically associated with the "event" could be displayed and compared with the changes in, say, Tahiti pressure anomalies. Again, the selection of the index and threshold used to select events was, to some extent, arbitrary, and the composite approach ignored all other years.

Cross-spectral analysis between just a few SO indices or variables has also been used to study teleconnections, especially their lagged relationships (e.g., Trenberth, 1976; Chen, 1982). Chen calculated the cross-spectra between monthly anomalies of pressure at Darwin and Tahiti (among others) to investigate the strength of the relationship and the lags between fluctuations at the two sites. This approach uses all years of data, unlike the composite approach, but may not be appropriate for teleconnections characterized by short-term irregular features and/or cyclo-stationarity such as phase locking of anomalies with the annual cycle (Barnett, 1983).

Canonical correlation is proposed here as yet another statistical technique that can assist in the study and depiction of teleconnections, especially their temporally varying nature. The technique proposed uses all the data without applying arbitrary thresholds (in effect it treats the teleconnection as a continuum in time rather than arising from several distinct events), copes with sinusoidal or irregular temporal patterns, and appears to be useful for determining lags between sites or variables.

The data for a canonical correlation analysis comprise a set of observations for each of two sets (**X** and **Y**) of variables. Canonical correlation finds a linear combination, called the first canonical variable, from each set such that the correlation between them is maximized. This correlation is the first canonical correlation. The first canonical variates are two series of values (or scores) of the linear combinations of the original sets of variables. There are as many scores in each series as there are observations for the original variables. The analysis continues by finding a second linear combination from each set, uncorrelated with the first pair, that produces the next highest correlation coefficient, and so on. The form of the canonical vectors is best depicted by the canonical structure patterns, which are the correlations of each of the original variables with the canonical variates.

Johnston (1980) provides a clear nonmathematical description of canonical correlation analysis and several simple examples of its use. A mathematical description, following that of Srivastava and Carter (1983) and Lloyd (1984), is provided in the Appendix. Glahn (1968) provides further discussion of canonical correlation, especially its relationship to discriminant analysis and multiple regression, along with an example of its possible application in meteorology.

One potential problem with canonical correlation is that it requires a large sample-to-variable ratio to ensure stable canonical vectors. In the examples discussed later this ratio was not allowed to fall below about 2 and was usually well above this. The stability of the canonical vectors was tested by multiple replications omitting segments of the data (cross-validation). Classical confirmatory tests were also used to examine the significance of the canonical correlations. Given the a posteriori nature of canonical correlation, however, other approaches to significance testing, such as replication on further samples or cross-validation (e.g., Stone 1974) seem more relevant.

The canonical correlation method and the data used to study, as examples, three teleconnections (Darwin pressure with, in turn, Tahiti pressure, Willis Island air temperature, and southeast Australian rainfall) are described in section 2. The results are presented and discussed in section 3 and the conclusions presented in section 4.

2. Method and data

The canonical correlation analysis employed here aims to examine the temporal nature of the teleconnections between variables at two sites. In each case the **X** variables are bimonthly averages of Darwin pressures. The first variable in **X** is the September–October mean Darwin pressure. The second variable is the November–December mean Darwin pressure. Then follow, for Darwin pressure, the means for January–February, March–April, May–June, July–August, September–October, November–December (all for the following year), then January–February, March–April, and May–June (for the subsequent year). Bimonthly averages, rather than monthly averages, were used to restrict the number of variables to ensure stability of the canonical vectors. The **Y** set of variables for the three examples used means for exactly the same set of bimonths but for, in turn, Tahiti pressure, Willis Island air temperature, and southeast Australian rainfall.

The particular set of 22 months used for the 11 bimonthly variables was chosen to span a typical SO life cycle of about a year or so and to locate the period of strongest autocorrelations of SO indices (July to December) in the middle of the set.

Once the canonical vectors had been calculated, the canonical structure patterns were determined by correlating the first canonical variates with the monthly means of the two original variables (Darwin pressure and either Tahiti pressure, southeast Australian rainfall, or Willis Island air temperature) over 3 yr starting in the January prior to the first bimonth (September–October) used in the analysis. Thus, the canonical vectors were calculated using the 11 bimonths (to ensure a healthy sample-to-variable ratio), but the pattern of the teleconnections, as revealed by this analysis, is illustrated by correlating monthly means of each of the original variables with the first canonical variates. This provides the opportunity for a more detailed description of the teleconnection pattern in time.

The three examples of the use of canonical correlation analysis for the study of teleconnections were chosen because previous work had demonstrated the existence of strong teleconnections between the variables used. Darwin pressure was selected as the **X** variable set in each case because it is used frequently as a single variable index of the SO. The Australian Bureau of Meteorology's National Climate Centre provided Darwin monthly mean pressures from 1882–1985. The same organization also provided monthly rainfall data for southeast Australian towns (from 1882) and Willis Island air temperature (from 1939).

The relationship between Tahiti and Darwin pressures is perhaps the best known SO teleconnection so the choice of Tahiti as an example seemed obvious. Tahiti monthly mean pressures from 1882 were obtained from P. D. Jones of the Climatic Research Unit, University of East Anglia. These data had been corrected for some years by Jones on the basis of comparisons with other nearby stations. In some years prior to 1935 one or more months of Tahiti pressure data are missing. Such years have been omitted for the canonical correlation analysis between Tahiti and Darwin pressures. Eighty-three complete years of data remained.

Quayle (1929) proposed that southeast Australian rainfall was related to Darwin pressures in some seasons. Nicholls and Woodcock (1981) confirmed this relationship on independent data. The teleconnection between Darwin pressure and southeast Australian rainfall was examined using the canonical correlation technique to study the exact seasonal nature of the relationship. An index of southeast Australian rainfall was calculated for each month from 1882 by averaging rainfall received at the ten towns used by Quayle. Figure 1 in Nicholls and Woodcock (1981) indicates the locations of the stations.

Sea surface temperatures in the Coral Sea are related to the SO, and there is a possibility that changes in sea temperatures here lead changes in some other SO indices such as Darwin pressure (e.g., Nicholls, 1984; Wright, 1984), but a long sea temperature data series was not available so the monthly mean of the daily minimum air temperatures at Willis Island (17°S, 150°E) was used as a proxy. Wright (1984) demonstrated that relationships between air temperature and the SO are similar to relationships between sea surface temperature and the SO in this region. Data from before 1939 were not available for Willis Island.

3. Results and discussion

The results of the study are shown in Fig. 1, which plots, for each of the three examples, the correlation of the monthly means of the two original variables (e.g., Tahiti pressure and Darwin pressure) with their respective first canonical variate. The first canonical correlations for each example are listed in Table 1, along with the number of observations (years) in the sample. In each example the first canonical correlation was significantly different from zero at the 1% level or better (see Appendix for significance test used).

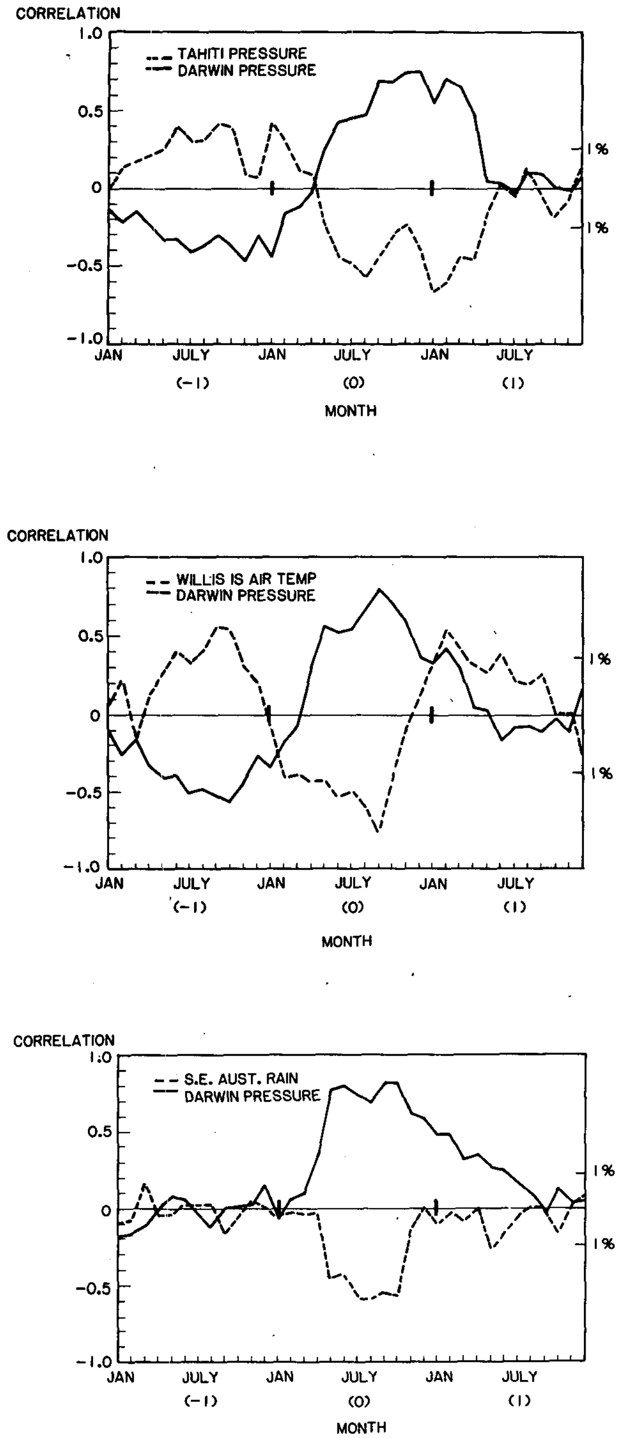

FIG. 1. Correlations of the first canonical variates with monthly means of the original variables. (a) Darwin pressure (solid line) and Tahiti pressure (dotted line). (b) Darwin pressure (solid line) and Willis Island mean monthly minimum air temperature (dotted line). (c) Darwin pressure (solid line) and southeast Australian rainfall (dotted line). See text for details of data used in analysis. The 1% significance levels for the correlations are indicated. The notation (0) on the abscissa denotes months in the middle of the three-year period for which correlations were calculated; (−1) and (+1) refer, respectively, to months in the first year and the last year.

The canonical correlation analysis used here should reveal the presence of lags between the sets of original variables. Thus, Darwin and Tahiti pressures are known to be correlated strongly, and interannual fluctuations at the two stations tend to be out of phase. If, say, changes in Tahiti pressure tend to lead changes of the opposite sign in Darwin pressures, then such a lead might be revealed in the canonical structures in Fig. 1. In such a case one might expect that peaks (troughs) in the Tahiti canonical structure should occur earlier than corresponding troughs (peaks) in the Darwin canonical structure. The changes in sign in the Tahiti canonical structure should also lead changes in sign (of the opposite sense) in the Darwin structure. If, however, no leads or lags (other than a perfect out-of-phase relationship) exist between Tahiti and Darwin pressure, then the peaks in the canonical structure for Tahiti should align with the troughs in the Darwin structure, and vice versa.

a. Darwin and Tahiti pressure

The plot of correlations of Tahiti and Darwin pressure with their respective first canonical variate (Fig. 1a) reveals the familiar tendency for pressure anomalies at the two stations to be out of phase. The positive correlations at Darwin, and the negative ones at Tahiti, start about May of the middle year shown and last about 12 months. This behavior is very similar to that obtained from composites of SO events (e.g., Rasmusson and Carpenter, 1982). There is no evidence of any lag between the two sites, contrary to the lead of one month by Tahiti pressures found in some cross-spectral analyses. The figure suggests, by the correlations of reverse sign in the first year shown, that years with large Darwin and Tahiti pressure anomalies tend to be preceded by anomalies of the opposite sign. This tendency is well known and is also apparent, for instance, in the composites of Rasmusson and Carpenter (1982; their Fig. 15). The results in Fig. 1a thus match rather well the conclusions of other statistical studies of the SO relationship between pressures at Tahiti and Darwin.

One feature in Fig. 1a does not appear to have been noted before. The correlations of Tahiti pressure in October and November with its first canonical variate are much smaller in magnitude than those of the surrounding months, producing a double "trough" in the plot of correlations. This feature remained even when

TABLE 1. Canonical correlations and sample sizes.

Variables	First canonical correlation (sample size)
Darwin pressure–Tahiti pressure	0.80 (83)
Darwin pressure–Willis Island air temp.	0.94 (43)
Darwin pressure–SE Australian rainfall	0.72 (101)

the analysis was replicated numerous times, each time omitting a separate subset of data. The plot of correlations for Darwin pressure, on the other hand, did not show a two-peaked structure. East equatorial sea surface temperature anomalies associated with the SO show a double-peaked structure (Rasmusson and Carpenter, 1982) with a minimum about October and November. Perhaps the double-trough structure of the Tahiti correlation in Fig. 1a is related to the double-peaked structure in east equatorial Pacific sea surface temperature anomalies during SO events?

Time series of the first canonical variate for Darwin pressure and Tahiti pressure are shown in Fig. 2 for the period 1935–84. The major SO events, such as 1982, stand out clearly in this diagram. The strong positive correlation between the two is also clear, as expected from the strong canonical correlation (Table 1).

This example demonstrates how canonical correlation can be used to explore for, or confirm and provide more detail of, teleconnections. In particular, the technique has the freedom to uncover rather complicated temporal signatures not revealed by cross-spectral analysis and unlikely to be found without an exhaustive cross-correlation analysis.

b. *Darwin pressure and Willis Island air temperature*

The plot of correlations of Darwin pressure with the first canonical variate in this case (Fig. 1b) is very similar to that in the previous case. As before, the large positive correlations appear about May, last approximately 12 months, and are preceded by negative correlations. The Willis Island air temperature correlations are, in general, opposite in sign to the Darwin pressure correlations, as was found by Wright (1984). There is, however, a clear lag. Negative, though not significant, Willis Island correlations first appear in January, several months before the first appearance of positive Darwin correlations. The Willis Island negative correlations reach their greatest magnitude in September and then are quickly replaced by positive correlations, well before the positive Darwin correlations have dropped below the 1% significance level. This analysis suggests that Willis Island air temperature anomalies are related to the SO and tend to change sign some months before related changes in sign in the Darwin pressure anomalies. Nicholls (1984) noted a similar tendency for sea surface temperature in this region, using only two small datasets.

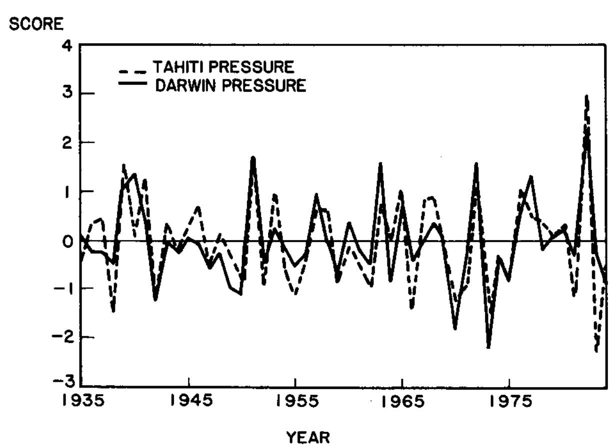

FIG. 2. Time series of the first canonical variates (i.e., scores of the first canonical vectors for each observation) for Darwin pressure (solid line) and Tahiti pressure (dotted line).

This example demonstrates how canonical correlation can be used to determine the lags between sites or variables known to be teleconnected. It provides an alternative to cross-spectral analysis and can cope with rather complicated temporal signatures, far removed from sinusoids.

c. *Darwin pressure and southeast Australian rainfall*

The plot of Darwin pressure correlations with the first Darwin canonical variate from the analysis with southeast Australian rainfall again shows strong positive correlations in the middle year plotted in Fig. 1c, as did the same plot for the other two examples. There are, however, substantial differences between the shape of this plot and those in the other two examples. Thus, in this case, there are only very weak correlations in the year before, and the positive correlations persist longer in the third year. The cause of this difference in behavior is not understood by the author, although it appears to suggest that the teleconnection between Darwin pressure and southeast Australian rainfall cannot be described as solely an SO teleconnection.

The plot of correlations of southeast Australian rainfall with its first canonical variate indicates clearly that this teleconnection with Darwin pressure only refers to rainfall in the period May–November. The correlations with summer (December–February) rainfall are very small and insignificant, even though there are strong, significant correlations at this time of year between Darwin pressure and *its* first canonical vector.

This example demonstrates how the canonical correlation analysis can be used to isolate the seasons contributing to a teleconnection.

4. Conclusions

This study proposed the use of canonical correlation as a technique for exploring teleconnections and provided three examples of its use. The canonical vectors appear to summarize the teleconnection pattern in a convenient form. The technique uses all data available without excluding any on the basis of a threshold, as is the case with composite analysis. Nor does the technique constrain the temporal signature of the teleconnections to be a simple sinusoid, but can reveal quite complicated signatures. The results seem to be easy to interpret and display. It is concluded that canonical correlation analysis can play a useful role in studying teleconnections, especially in determining the temporal signatures of the teleconnections and searching for lags between the variables involved in the teleconnection.

A potential problem with canonical correlation, especially for the study of teleconnections, is that low sample-to-variable ratios might result in unstable patterns. For the three examples discussed here the results were quite stable, with a sample-to-variable ratio of about 2 or better. Significance testing is also difficult since canonical correlation is essentially an a posteriori approach. The technique is proposed, therefore, only as an exploratory technique.

On the basis of the three examples it seems that

(i) there is no lag between SO-related fluctuations in Tahiti pressure anomalies and Darwin pressure anomalies of the opposite sign;

(ii) Tahiti pressures related to the SO appear to exhibit a double-trough structure with October and November pressures poorly related to the SO, reminiscent of the double-peaked structure of SO-related east equatorial Pacific sea surface temperature anomalies;

(iii) air temperatures at Willis Island are strongly related to the SO and changes in this variable appear to lead, by several months, changes in Darwin pressure; and

(iv) a Darwin pressure signature similar, but not identical to the SO signatures found in the other two examples, is related to southeast Australian rainfall, but only to rainfall in the period May–November.

Acknowledgments. I wish to thank T. P. Barnett, R. D. Rosen, and an anonymous reviewer for their suggestions on earlier versions of this paper. Their comments were of considerable help in the preparation of the final version. P. D. Jones provided the Tahiti pressure data.

APPENDIX

Canonical Correlation

The following description of canonical correlation follows the expositions in Srivastava and Carter (1983) and Lloyd (1984).

Consider two column vector random variables $\mathbf{X} = (X_1, X_2, \cdots, X_p)'$ and $\mathbf{Y} = (Y_1, Y_2, \cdots, Y_q)'$ with $q \leqslant p$. Let the covariance matrices of $\mathbf{X}$ and $\mathbf{Y}$ be

$$\mathbf{V}_{11} = E(\mathbf{X}\mathbf{X}'), \quad \mathbf{V}_{22} = E(\mathbf{Y}\mathbf{Y}')$$

respectively, and let the cross-covariance matrix between $\mathbf{X}$ and $\mathbf{Y}$ be

$$\mathbf{V}_{12} = E(\mathbf{X}\mathbf{Y}').$$

Canonical correlation attempts to simplify the relationships between $\mathbf{X}$ and $\mathbf{Y}$, given by the covariances in the matrix $\mathbf{V}_{12}$, by considering correlations between linear combinations of the two sets of random variables.

The first canonical correlation, say ρ, is the maximum correlation between a linear combination of $\mathbf{X}$, say $U = \mathbf{a}_1'\mathbf{X}$, and a linear combination of $\mathbf{Y}$, say $W = \mathbf{b}_1'\mathbf{Y}$; that is,

$$\rho_1 = \max_{\mathbf{a}_1\mathbf{b}_1} \mathbf{a}_1' \mathbf{V}_{12} \mathbf{b}_1 / (\mathbf{a}_1' \mathbf{V}_{11} \mathbf{a}_1')^{1/2} (\mathbf{b}_1' \mathbf{V}_{22} \mathbf{b}_1')^{1/2}.$$

Let

$$A = \mathbf{V}_{11}^{-1/2} \mathbf{V}_{12} \mathbf{V}_{22}^{-1/2}.$$

Then ρ_1^2 is the maximum eigenvalue of AA', $\mathbf{a}$ is the first eigenvector of AA' and $\mathbf{b}_1$ the first eigenvector of $A'A$. The first canonical vectors are $\mathbf{a}_1$ and $\mathbf{b}_1$ and U and W are the first canonical variates. Other canonical correlations and canonical vectors correspond to other eigenvalues and eigenvectors of AA' and $A'A$. The results of this study are displayed in Fig. 1 as the correlations of the canonical variates U and W with the original variables $\mathbf{X}$ and $\mathbf{Y}$, respectively.

The Statistical Analysis System (SAS, 1982) procedure CANCORR was used to calculate the canonical correlations and to test their significance using classical significance tests. A series of hypotheses, that each canonical correlation and all smaller canonical correlations are zero in the population, was tested. An F statistic based on Rao's approximation (Rao 1973, p. 556) was used.

REFERENCES

Barnett, T. P., 1977: The principal time and space scales of the Pacific tradewind field. *J. Atmos. Sci.*, **34**, 221–236.

——, 1983: Interaction of the monsoon and Pacific trade wind system at interannual time scales. Part I: The equatorial zone. *Mon. Wea. Rev.*, **111**, 756–773.

Chen, W. Y., 1982: Assessment of Southern Oscillation sea-level pressure indices. *Mon. Wea. Rev.*, **110**, 800–807.

Egger, J., G. Meyers and P. B. Wright, 1981: Pressure, wind, and cloudiness in the tropical Pacific related to the Southern Oscillation. *Mon. Wea. Rev.*, **109**, 1139–1149.

Glahn, H. R., 1968: Canonical correlation and its relationship to discriminant analysis and multiple regression. *J. Atmos. Sci.*, **25**, 23–31.

Horel, J. D., 1984: Complex principal component analysis: Theory and examples. *J. Climate Appl. Meteor.*, **23**, 1660–1673.

Johnston, R. J., 1980: *Multivariate Statistical Analysis in Geography.* Longman, 280 pp.

Lanzante, J. R., 1984: A rotated eigenanalysis of the correlation between 700 mb heights and sea surface temperatures in the Pacific and Atlantic. *Mon. Wea. Rev.*, **112**, 2270–2280.

Lloyd, E., 1984: *Handbook of Applicable Mathematics. Vol. VI, Statistics, Part B.* Wiley, 942 pp.

Newell, R. E., R. Selkirk and W. Ebisuzaki, 1982: The Southern Oscillation: Sea surface temperature and wind relationships in a 100 year data set. *J. Climatol.*, **2**, 357–373.

Nicholls, N., 1984: The Southern Oscillation and Indonesian sea surface temperature. *Mon. Wea. Rev.*, **112**, 424–432.

——, and F. Woodcock, 1981: Verification of an empirical long-range weather forecasting technique. *Quart. J. Roy. Meteor. Soc.*, **107**, 973–976.

Pan, Y. H., and A. H. Oort, 1983: Global climate variations associated with sea surface temperature anomalies in the eastern equatorial Pacific Ocean for the 1958–73 period. *Mon. Wea. Rev.*, **111**, 1244–1258.

Quayle, E. T., 1929: Long-range rainfall forecasting from tropical (Darwin) pressure. *Proc. Roy. Soc. Victoria*, **41**, 160–164.

Rao, C. R., 1973: *Linear Statistical Inference and Its Applications.* Wiley, 625 pp.

Rasmusson, E. M., and T. H. Carpenter, 1982: Variations in tropical sea surface temperature and surface wind fields associated with the Southern Oscillation/El Niño. *Mon. Wea. Rev.,* **110,** 354–384.

SAS, 1982: *SAS User's Guide: Statistics 1982 Edition.* SAS Institute, 584 pp.

Srivastava, M. S., and E. M. Carter, 1983: *An Introduction to Applied Multivariate Statistics.* North-Holland, 394 pp.

Stone, M., 1974: Cross validatory choice and assessments of physical predictions. *J. Roy. Statis. Soc.,* **B36,** 111–147.

Trenberth, K. E., 1976: Spatial and temporal variations of the Southern Oscillation. *Quart. J. Roy. Meteor. Soc.,* **102,** 639–644.

Van Loon, H., and R. A. Madden, 1981: The Southern Oscillation. Part I: Global associations with pressure and temperature in northern winter. *Mon. Wea. Rev.,* **109,** 1150–1162.

Walker, G. T., and E. W. Bliss, 1932: World weather IV. *Mem. R. Meteor. Soc.,* **4,** 53–84.

Weare, B. C., and J. S. Nasstrom, 1982: Examples of extended empirical orthogonal function analyses. *Mon. Wea. Rev.,* **110,** 481–485.

Wright, P. B., 1977: The Southern Oscillation—patterns and mechanisms of the teleconnections and the persistence. HIG-77-13, Hawaii Institute of Geophysics, 107 pp.

——, 1984: Relationship of SST near northern Australia to the Southern Oscillation. *Trop. Ocean–Atmos. Newslett.,* **27,** 17–18.

——, T. P. Mitchell and J. M. Wallace, 1985: Relationships between surface observations over the global oceans and the Southern Oscillation. NOAA Data Rep. ERL PMEL-12, 61 pp.