

Searching for analogues, how long must we wait?

By H. M. VAN DEN DOOL, *Climate Analysis Center, NMC/NWS/NOAA, Washington DC 20233, USA*

(Manuscript received 24 February 1993; in final form 13 September 1993)

ABSTRACT

A three-way relationship is derived between the size of a library (M years) of historical atmospheric data, the distance between an arbitrarily picked state of the atmosphere and its nearest neighbor (or analogue), and the size of the spatial domain, as measured by the number of spatial degrees of freedom (N). It is found that it would take a library of order 10^{30} years to find 2 observed flows that match to within current observational error over a large area such as the Northern Hemisphere. Obviously, with only 10–100 years of data, the probability of finding natural analogues is very small, unless one is satisfied with analogy over small areas or in just 2 of 3 degrees of freedom as represented, for instance, by 2 or 3 leading empirical orthogonal modes. We further propose the notion that analogues can be constructed by combining a number of observed flow patterns. We have found at least one application where linearly constructed analogues are conclusively better at specifying US surface weather from concurrent 700 mb geopotential height than natural analogues are.

1. Introduction

It has long been assumed that if 2 observed atmospheric states are very close initially they will remain close for some time to follow. This intuitive assumption has two aspects that need to be distinguished, one diagnostic and one prognostic. The analogue diagnostic aspect is concerned with identifying atmospheric states that are so close that they can be called each others' analogue.

The identification of analogues is a classical theme in meteorology. With some emphasis we invite the reader not to associate the word analogue with forecasting exclusively. To our knowledge there appear to be at least 6 areas of interest where the concept of analogues is relevant or even applied in practice. They are (one reference per area added): (1) short-range forecasting through analogues (Shabbar and Knox, 1986), (2) specification of surface weather through analogues to a prognostic upper level map (Kruizinga and Murphy, 1983), (3) long-range forecasting through analogues (Barnett and Preisendorfer, 1978), (4) estimation of atmospheric predictability (Lorenz, 1969), (5) estimates of the dimension of

atmospheric phase space (Fraedrich, 1986) and (6) cluster analysis (Wallace et al., 1991).

A very major problem in meteorological practice is that there are, diagnostically, no very good analogues to begin with (Lorenz, 1969; Gutzler and Shukla, 1984; Shabbar and Knox, 1986; Ruosteenoja, 1988; Toth, 1991a). This is generally attributed to the short historical record of observations (order 10–100 years) which makes it highly unlikely to find two matching atmospheric states over the entire globe given the large variety of atmospheric flow.

As an illustration of the problem of finding a good analogue, a fresh example is given in Fig. 1. For each of the 61 days between 16 June 1990 and 15 August 1990 we searched a library to find the nearest neighbor in terms of root-mean-square difference in global 00GMT 500 mb geopotential height on a 2.5×2.5 degree grid. The library consists of 00GMT 500 mb geopotential height for the same dates in the non-matching years 1987, 88, 89, 91 and 92. For each base date one of the 305 (5×61) candidates is the nearest. The distance to the nearest neighbor (DNN) is plotted in Fig. 1 as a function of time. Typically the DNN is 80 to




Fig. 1. The RMS distance of the 500 mb geopotential height field on the base data (1–61; 16 June–15 August 1990) to the nearest neighbor in a 5-year library. Thinner dashed curve is the distance to the climate mean. Units gpm. The domain is global.

110 gpm, which is far too large to be called analogous. The lighter dashed curve in Fig. 1 represents the distance of the base to the climate mean, or the anomaly magnitude (AM). Indeed, with only a few years of data, and a large domain, it is hard to find an analogue to within a small distance (15 gpm say). On most days, we cannot find a neighbor to within the distance of the base to the climate mean (about 90 gpm on average; this is the space-time ensemble averaged standard deviation). Hence, the climate mean, rather than a real observed flow, is often the “best” analogue on this large domain. The base for which we find the smallest DNN (about 80 gpm) happens to be 2 August 1990 (case #48). Note also that DNN and AM are somewhat correlated, particularly in the first month. This is to be expected on finite libraries because as we approach the mean (small AM) the higher probability density tends to make DNN smaller too (Van den Dool, 1989; Toth, 1991b).

The criterion for a truly good analogue is a bit arbitrary at this point and depends on the application. For forecasts based on analogues to be

competitive to day-to-day numerical weather forecasting the initial error, i.e., the DNN, would have to be less than 15 gpm for the 500 mb height field, with simultaneous close matching in other variables and levels too. Clearly natural analogues are nowhere near such accurate initial matching.

Figs. 2, 3 are the same as Fig. 1, but now for progressively smaller areas, i.e., the extratropical Northern Hemisphere (20–90°N) and Eastern North America/Western Atlantic (30–50°N, 50–70°W) respectively. Over a small enough area one can find reasonably good analogues (DNN < 40 gpm) for almost every base.

The purpose of this manuscript is first of all to present a method that enables us to estimate how many years of data it would take to have a certain level of confidence in finding an analogue of a certain desired quality (read small DNN) over an area of a certain size. By size, we do not mean $X \text{ km}^2$, but rather a dynamical size best expressed by the number of spatial degrees of freedom. The method is outlined in Section 2, and the results are presented in Section 3. Essentially, the result is a three-way relationship between the library size




Fig. 2. As Fig. 1, but now for the extratropical Northern hemisphere (20–90 N) domain.




Fig. 3. As Fig. 1, but now for an area over Eastern North America and the West Atlantic (30–50 N, 50–70 W).

(M), the number of spatial degrees of freedom (N) and the desired accuracy of matching (ε). The outcome, no surprise, is that we need to wait very many years, order 10^{30} if we want an accurate match over a large area.

We proceed in Section 4 by discussing the best ways to use a short data set. With current libraries of 10–100 years we can find a good match in only 2 or 3 spatial degrees of freedom. We will also (re)interpret some earlier papers on analogues in the literature from this point of view. Clearly, the best way of spending 2 degrees of freedom depends on the intended application.

Finally, in Section 5 we present a possibility to shorten the waiting time. The idea developed here is to *construct an analogue* of a desired quality by combining a set of atmospheric states observed in the past. For at least one of the 6 applications listed above (#2: specification) this idea appears to have merit.

2. How long? Method

We can pose the problem of finding an analogue as follows: Given a library of M years, how large must M be to have a 95% probability to find a match in the library that resembles any randomly selected base case to within a desired tolerance? In the discussion below the base case is simply a 500 mb height map (today's or any other), but the reasoning can be readily applied to multilevel data, including boundary conditions.

The answer is derived as follows. The chance of two randomly chosen extratropical 500 mb height fields to be within tolerance ε depends on 3 factors. One is the natural variability of height, represented here, for simplicity, by a single standard deviation (sd). The second is the number of spatial degrees of freedom (N) that characterizes 500 mb height, while the third is the tolerance ε itself. In the calculation below, we will take $sd = 90$ gpm (see Fig. 1 or Fig. 2). N is at least 20 for the extratropical Northern Hemisphere (Van den Dool and Chervin, 1986; Wallace et al., 1991) but probably less than 35, so we will tabulate M for various values of N in that range, and for ε we take 15 and 7.5 gpm, the latter thought to be a low estimate of the current observational/analysis error in 500 mb height in densely observed areas.

Additional assumptions of varying degree of

importance have to be made as we go along. For instance, considering day-to-day autocorrelation, we assume that in any given year there are only 20 independent cases inside a 2-month-window in time centered around the date of the base case. Rather importantly we assume 500 mb height to behave like a multi-normal probability distribution with N dimensions in phase space (Toth, 1991b).

The chance (α) that two arbitrarily chosen states are to within ε at a single point in physical space is given by

$$\alpha = \int_{-\varepsilon/sd'}^{\varepsilon/sd'} P(y) dy, \quad (1)$$

where $P(y)$ is the standardized Gaussian probability density function of height differences (zero mean, standard deviation $sd' = sd \cdot 2^{1/2}$). Consulting standard tables we estimate $\alpha = 0.08$ (0.04) for $\varepsilon = 15$ (7.5) gpm. This means, for instance, that there is an 8% chance that 2 arbitrarily chosen maps agree to within 15 gpm at the prechosen point of interest. Under those circumstances, one does not have to wait very long in order to be practically sure of finding an analogue. (We leave aside here whether or not it is of any practical significance to have an analogue at one point.)

Eq. (1) is "global" in the sense that we have implicitly integrated over all possible base cases. Therefore α does not depend on, for instance, the distance of a particular base case to the climate mean, a complication we want to avoid here. (Close to the mean it is often easier to find good analogues (see Fig. 1 for evidence), simply because the probability density is locally high there, compared to the extremes of the distribution where the points in phase space (i.e., observed flows) are farther apart.)

We are equating one spatial degree of freedom to one point in physical space where heights vary according to a Gaussian process, and look upon N degrees of freedom as N such processes going on independently from each other at N points in physical space. The chance that two arbitrarily chosen states are within ε at all N points is then given by α^N . Here, we have assumed that all dimensions have the same standard deviation, similar to Wallace et al. (1991) in their discussion of the linear paradigm. (Identification of N points in physical space where independent Gaussian

processes with identical standard deviation go on is not trivial and may be impossible. Some readers may prefer to think in terms of N empirical orthogonal functions (EOF). Likewise it is non-trivial to calculate equal variance EOFs.)

The chance (c) of finding an analogue can hence be written

$$c = 1 - (1 - \alpha^N)^{20 \times M}. \quad (2)$$

Demanding $c > 0.95$, we can solve for M from

$$M > \frac{\ln(0.05)}{20 \times \ln(1 - \alpha^N)} \approx \frac{\ln(0.05)}{20 \times (\alpha^N)} \quad (3)$$

(where the number 20 is the number of independent cases in time in the 2-month window). The final assumption underlying (3) is that we have a constant climate during M years.

With (3), we can answer the question how big a library it takes to have a 95% chance to can find a match to the base that agrees to within ε (α implied) at all N points simultaneously. Note that demanding two states to be within ε at all N points is a more strict criterion than requiring the root-mean-square-difference (RMSD) of 2 states to be less than ε . Therefore DNN in Figs. 1–2 does not

precisely have the meaning of ε , and they are numerically different, DNN being the smaller of the two.

3. How long? Result

The left hand of Table 1 shows values of M for 2 entries of α and 4 entries of N . Typically $M = 10^{30}$ years, obviously with an uncertainty of several orders of magnitude given that many of our assumptions are not quite valid. That uncertainty hardly impacts the conclusion that M is very very large. This conclusion was also reached by Lorenz (1969), although he gave no number for M . It appears that Ruosteenoja (1988) obtained $M = 10^{29}$ years (not explicitly mentioned in his paper, but kindly supplied) by extrapolating a transformed frequency distribution of RMSD's very far into the low values (15 gpm say). The lowest RMSD ever observed in winter is only a poor 65 gpm for the extratropical Northern Hemisphere (Ruosteenoja, 1988).

Eq. (3) provides a threeway relationship between α and M (i.e., the smallest M required), and N , for given probability level $c = 0.95$ (which we keep constant). Fig. 4 shows plots of α versus M




Fig. 4. The chance α as a function of L for 4 values of N . On the right-hand, α has been converted to ε (units gpm).

Table 1. *Minimum size of the library needed to find an analogue*

$\alpha \backslash N$	20	25	30	35	2	3
0.042	10^{27}	10^{34}	10^{41}	10^{48}	94	2344
0.084	10^{21}	10^{26}	10^{32}	10^{37}	23	293
0.168					7	44

Given is M in years calculated by eq. (3) as a function of spatial degrees of freedom ($N = 20, \dots, 35$) and observational error (for $\varepsilon = 15$ gpm, $\alpha = 0.08$; for $\varepsilon = 7.5$, $\alpha = 0.04$). For $N = 2$ and 3 we added $\alpha = 0.17$ ($\varepsilon = 30$ gpm) which approximately corresponds to the RMSD in Van den Dool's (1989) study on limited area analogues.

for $N = 2, 20, 35$, and 200. For simplicity we have rescaled the horizontal axis by $M = 10^L$, so as to plot against linear L . Along the vertical axis we have linear α (left) and ε (right), the latter being calculated from standard tables for the Gaussian distribution using $sd = 90$. Fig. 4 shows how the quality of analogues (i.e., the tolerance ε) depends on M (or L) and on N . For all N , α and ε decrease as M increases.

How does Fig. 4 compare to Figs. 1–3? For instance in Fig. 2 (Northern Hemisphere) DNN is plotted for $M = 5$ (round off to $L = 1$!). We think $sd = 90$ gpm and $N = 20$ apply here. According to Fig. 4 one cannot, under these circumstances count on finding analogues better than about $\varepsilon = 160$ gpm at all points simultaneously, which is very poor indeed. As shown in Fig. 2 the DNN is about 80 gpm. A RMSD of 80 gpm implies discrepancies in excess of 80 gpm at some of the N points. We suspect ε and DNN are proportional and since it is arbitrary to demand a $c = 95\%$ chance (the $\ln(0.05)$ factor in (3)), one could tune c downward to make DNN and ε numerically comparable.

4. Use of short libraries

Current libraries are only 10–100 years. We can therefore accommodate only a few degrees of freedom. How few? On the right in Table 1 we added the M calculated from (3) for $N = 2$ and 3 and added an entry for the DNN (30 gpm) reported in Van den Dool (1989) who sought analogues for 500 mb heights on a circular area over North

America (radius 900 km). In Fig. 3 we showed similarly small DNN, but now for summer. From Table 1, it appears we can match slightly over 2 degrees of freedom with current libraries at a tolerance of 30 gpm. What are the most profitable degrees of freedom to be matched? While Van den Dool (1989) matched height fields over small areas one could also match two remote centers of action, or the coefficients of two leading empirical orthogonal functions (EOF) over some domain. The advantage of small areas is that for short lead times, time-tendencies, to first order, are determined locally, hence making a discussion of time evolution aspects feasible (Van den Dool, 1991). With EOF's this forecast aspect is not clear.

It depends obviously on the application, how one wants to spend 2 degrees of freedom. Kruizinga and Murphy (1983) and Vislocky and Young (1989) have sought analogues on limited areas with the purpose of specifying surface weather given an upper level prognostic chart. This has been, by far, the most successful practical application of analogues. Many good analogues can be found on a small area (i.e., N is small) with current libraries which allows an ensemble average as well as a probability specification of surface weather, i.e., proper representation of uncertainty. The success of analogues in specification is remarkable also because other competitive methods for specification, i.e., linear multiple regression (see Klein and Bloom (1987)), are available.

The third application discussed in the literature is long-range (month/season) forecasting, where analogues are sought on *time-mean* flow, the Southern Oscillation Index, sea-surface temperatures in certain areas, etc., with the purpose of making a forecast for a subsequent time-mean (Barnett and Preisendorfer, 1978; Harnack et al., 1985; Shabbar and Knox, 1986; Van den Dool, 1987; Livezey and Barnston, 1988; Barnston and Livezey, 1989). This has not been a great success, although for lack of truly superior competitors analogues may still be used in long-range weather forecasting. No great success can be expected because although time-averaging (in combination with other filters) lowers N , it does not lower it enough, at least not in the publications referred to. Also, there is no firm dynamical basis for the hope that similar time-means are followed by similar subsequent developments; in fact it is easy to show that this need not be the case. The true challenge is

to identify a small number of degrees of freedom that are relevant to the forecast in the area of interest.

The search for analogues has a very special meaning in meteorology because of its long and somewhat troublesome tradition in forecasting. However, searching for the nearest neighbor in phase space has taken on a new meaning recently with the advance of understanding deterministic chaos and dynamical systems (Fraedrich, 1986; Tsonis and Elsner, 1989). The estimation of the dimension of the phase space involves knowing the frequency distribution of distances between pairs, and in one estimate of the dimension, the limit for distance (essentially our ε) going to zero is taken. Accurate determination of this limit requires an enormous number of realizations (our M) unless the dimension (related to our N) is quite small (2 or so).

The ability to find high quality analogues is of relevance also to cluster analysis. This point was made rather explicit by Wallace et al. (1991), who used a cluster analysis method that starts off by finding the nearest neighbor in root-mean-square sense. In the absence of any true analogues, cluster analysis can hardly be expected to be a success. Indeed points are far apart in the multidimensional phase space with only 10–100 years of data. Perhaps cluster analysis should only be done on smaller regions (small N), so that clusters, if they exist, have a chance to be found.

5. Constructing an analogue

Is it really necessary to wait so long? Natural analogues being so impossible to find, we will now try to *construct an analogue* having greater similarity than the best natural analogue. In essence a construction is a mathematical combination of observed flows such that the combination is as close as possible to the base. Here, we will consider linear combinations only, thus limiting the further use of these constructed analogues to those applications where linearity is not a major obstacle. It has been shown (Klein, 1983) that linear regression equations are fairly successful in specifying US monthly surface weather from 700 mb height. Our specific task then is to show that constructed analogues can do a better job at specifying surface weather than natural analogues

(such as they are) can do, everything else being the same. (In doing so, we do not claim that constructed analogues are the best method to specify surface weather compared to all other methods imaginable.)

The data sets used are calendar means for 1950–91 for each of the 12 months. Upper level data is 700 mb height north of 20°N. Surface data is temperature and precipitation amount at 344 United States Climate Divisions.


For natural analogues the procedure is as follows. For any given base, say January 1967, we seek the nearest 10 other Januaries in terms of RMS distance of their height field to the January 1967 height field. The surface weather observed in these 10 nearest years is then averaged with equal weights and verified against the weather observed in January 1967. This is done for all bases exhaustively, and for all months.

We next proceed with construction as follows. For any given base, say January 1967, a linear combination is made of 41 January anomaly height fields (1950–91, 67 exclusive), so as to match the 1967 height field as closely as possible. This is done by classical least-squares minimization. (See appendix for some mathematical detail). For each base, we invert a 41×41 covariance matrix. The numerical stability of the resulting weights (for instance one would get $-0.08 \times \text{Jan. 1950} + 0.16 \times \text{Jan. 1951} \dots -0.34 \times \text{Jan. 1991}$) has been improved by artificially enhancing the diagonal elements of the covariance matrix using a method described by Meissner (1979). The same weights are then applied to put together a surface weather specification, which is then verified against observations in 1967. This is done for all bases and months. (More details about the construction process will be discussed elsewhere.)

Fig. 5 shows the verification for all months of the specification of US surface weather by the average of 10 natural analogues (the pluses), and by a single constructed analogue (the asterisks). The verification measure we use is the temporal correlation coefficient based on 42 entries (one for each month in the 1950–91 period). In Fig. 5 the correlation coefficient is averaged over all 344 Climate Divisions. The main message is that both for temperature and precipitation the constructed analogue scores about 10% higher in each and every month of the year. Not shown is that the similarity in terms of the height field is better than

skill of specifying US Temperature

method: circulation analogues



skill of specifying US Precipitation

method: circulation analogues






Fig. 5. The correlation between observed and specified temperature (top) and precipitation (bottom). Results are based on 42 years and averaged over all 344 Climate Divisions. Plusses are for 10 natural analogues averaged, the asterisk for the constructed analogue.



JAN 1950-91 CONSTRUCTED ANALOGUES



JAN 1950-91 NATURAL ANALOGUES

Fig. 6. Maps of the temporal correlation between observed and specified temperature over the US in January 1950-91. Upper map is for the constructed analogue, the lower for the average of 10 natural analogues. No shading for correlation less than 0.40, light stippling for correlation between 0.40 and 0.70, and heavy stippling for correlation in excess of 0.70.

0.90 in almost all cases, natural analogues rarely exceed 0.70.

Fig. 6 gives the spatial distribution of the correlation for the two methods over the US in January. At each point the correlation is based on 42 specified versus 42 observed values. As in Fig. 5 the main message is that constructed analogues are considerably better, turning areas with almost no skill (New England, natural analogues) into 0.5 correlation areas upon using the constructed analogue. Areas in excess of 0.7 correlation cover a much larger part of the nation when construction is used.

6. Conclusion and discussion

We have shown that it is highly improbable that truly good multi-level multi-variate global natural analogues will be found during our lifetime. With a fair bit of assumptions one can derive a three-way relationship relating the size of the library (M), the number of spatial degrees of freedom (N) and the desired closeness of analogues. This allowed us to estimate that it would take order 10^{30} years to find analogues that match over the entire Northern Hemisphere 500 mb height field to within current observational error.

We have also shown conclusively that, everything else being the same, constructed analogues outperform natural analogues in specifying monthly surface weather over the US from monthly 700 mb height over the NH north of 20°N . This result is obtained because (a) constructed analogues are much closer matches to the base than any natural analogue will be in the next 1000+ years, and (b) the specification arena is one application where it is not terribly detrimental for the constructed analogue to be put together linearly.

It would be a particular (perhaps impossible) challenge to construct an analogue which leaves the forecast aspects intact. This would likely involve a non-linear combination, because in order to make both initial conditions and the time derivative similar one has to match also non-linear advection terms. Constructing an analogue is an outgrowth of the idea that antilogues or mirror images can be used (Harnack et al., 1985; Van den Dool, 1987; Barnston and Livezey, 1989; Van den Dool, 1991), for as long as tendencies behave

linearly. To the extent that methods of long-range forecasting of say seasonal averages is linear, constructed analogues should be competitive to other linear forecast methods.

7. Acknowledgement

The author had a series of most stimulating discussions on this article with Anthony Barnston and Zoltan Toth. Thanks are due also to David Rodenhuis, Åke Johansson, Ed Epstein, Klaus Fraedrich and Leonid Rukhovets, for their help and comments.

8. Appendix

The method of construction is presented here for the sake of brevity for 2 participating years only, but can be extended to any number of years. Given the height field in the base month, January 1967 for example, denoted by $Z(j_b)$, two Januaries in other years (j_1, j_2) are taken. The constructed analogue (expressed in terms of anomalies) is given by

$$Z_c(j_b) = \alpha_1 Z(j_1) + \alpha_2 Z(j_2),$$

where the alphas are to be determined by minimizing the distance between Z_0 and $Z(j_b)$. This is done by a least squares fit over the domain north of 20°N using standard matrix inversion as in multiple linear regression. We thus solve the matrix problem

$$\begin{pmatrix} \sum Z(j_1) Z(j_1) & \sum Z(j_1) Z(j_2) \\ \sum Z(j_2) Z(j_1) & \sum Z(j_2) Z(j_2) \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \end{pmatrix} = \begin{pmatrix} \sum Z(j_b) Z(j_1) \\ \sum Z(j_b) Z(j_2) \end{pmatrix},$$

where the summation is over space. In the construction one could combine for example January 1953 and 1971 to construct an analogue to January 1967. The surface weather is then constructed in the same way, using the alphas that were derived from the height only.

It is extremely important to note that anomalies in the above are taken relative to a climate average not including the base year, so as to avoid the error noted in Van den Dool (1987).

REFERENCES

- Barnett, T. P. and Preisendorfer, R. W. 1978. Multifield analog prediction of short-term climate fluctuations using a climate state vector. *J. Atmos. Sci.* **35**, 1771–1787.
- Barnston, A. G. and Livezey, R. L. 1989. An operational multifield analog/anti-analog prediction system for United States seasonal temperatures. Part II: Spring, summer, fall and intermediate 3-month period experiments. *J. Climate* **2**, 513–541.
- Gutzler, D. S. and Shukla, J. 1984. Analogs in the winter-time 500 mb height field. *J. Atmos. Sci.* **41**, 177–189.
- Fraedrich, K. 1986. Estimating the dimensions of weather and climate attractors. *J. Atmos. Sci.* **43**, 419–432.
- Harnack, R., Cammarata, M., Dixon, K., Lanzante, J. and Harnack, J. 1985. Summary of US seasonal temperature forecast experiments. *Proc. of the 9th Conf. on Probability and Statistics in Atmospheric Sciences*. Virginia Beach, 175–179. American Meteorological Society, Boston, USA.
- Klein, W. H. 1983. Objective specification of monthly mean surface temperature from 700 mb heights in winter. *Mon. Wea. Rev.* **111**, 674–691.
- Klein, W. H. and Bloom, H. J. 1987. Specification of monthly precipitation over the United States from surrounding 700 mb height fields in winter. *Mon. Wea. Rev.* **115**, 2118–2132.
- Kruizinga, S. and Murphy, A. H. 1983. Use of an analogue procedure to formulate objectively probabilistic temperature forecasts in The Netherlands. *Mon. Wea. Rev.* **111**, 2244–2254.
- Livezey, R. E. and Barnston, A. G. 1988. An operational multifield analog/anti-analog prediction system for United States seasonal temperatures, Part I: System design and winter experiments. *J. Geoph. Res.* **93**, 10953–10974.
- Lorenz, E. N. 1969. Atmospheric predictability as revealed by naturally occurring analogues. *J. Atmos. Sci.* **26**, 636–646.
- Meissner, B. N. 1979. Ridge regression. Time extrapolation applied to Hawaiian rainfall normals. *J. Appl. Meteorol.* **18**, 904–913.
- Ruosteenoja, K. 1988. Factors affecting the occurrence and lifetime of 500 mb height analogues: A study based on a large amount of data. *Mon. Wea. Rev.* **116**, 368–376.
- Shabbar, A. and Knox, J. L. 1986. Monthly prediction by the analogue method. Proceedings of the 1st WMO workshop on the diagnosis and prediction of monthly and seasonal atmospheric variations over the globe. *Long-range forecasting Res. Rep. Ser. 6 Vol. II*, Tech. doc. WMO/TD 87, 672–681. World Meteorological Organization, Geneva, Switzerland.
- Toth, Z. 1991a. Estimation of atmospheric predictability by circulation analogs. *Mon. Wea. Rev.* **119**, 65–72.
- Toth, Z. 1991b. Circulation patterns in phase space: A multinormal distribution? *Mon. Wea. Rev.* **119**, 1501–1511.
- Tsonis, A. A. and Elsner, J. B. 1989. Chaos, strange attractors and weather. *Bull. Amer. Met. Soc.* **70**, 14–23.
- Van den Dool, H. M. 1987. A bias in skill in forecasts based on analogues and antilogues. *J. Climate Appl. Meteor.* **26**, 1278–1281.
- Van den Dool, H. M. 1989. A new look at weather forecasting through analogues. *Mon. Wea. Rev.* **117**, 2230–2247.
- Van den Dool, H. M. 1991. Mirror images of atmospheric flow. *Mon. Wea. Rev.* **119**, 2095–2106.
- Van den Dool, H. M. and Chervin, R. M. 1986. A comparison of month-to-month persistence of anomalies in a General Circulation Model and in the Earth's Atmosphere. *J. Atmos. Sci.* **42**, 1454–1466.
- Vislocky, R. L. and Young, G. S. 1989. The use of perfect prog forecasts to improve model output statistics forecasts of precipitation probability. *Weather Forecasting* **4**, 202–209.
- Wallace, J. M., Cheng, X. and Sun, D. 1991. Does low-frequency atmospheric variability exhibit regime-like behavior? *Tellus* **43AB**, 16–26.