# The Continuum of Northern Hemisphere Teleconnection Patterns and a Description of the NAO Shift with the Use of Self-Organizing Maps

NATHANIEL C. JOHNSON

*Department of Meteorology, The Pennsylvania State University, University Park, Pennsylvania*

STEVEN B. FELDSTEIN

*Earth and Environment Systems Institute, The Pennsylvania State University, University Park, Pennsylvania*

BRUNO TREMBLAY

*Department of Atmospheric and Oceanic Sciences, McGill University, Montreal, Quebec, Canada*

ABSTRACT

In this study, the method of self-organizing maps (SOMs) is used with NCEP–NCAR reanalysis data to advance the continuum perspective of Northern Hemisphere teleconnection patterns and to shed light on the secular eastward shift of the North Atlantic Oscillation (NAO) that began in the late 1970s. A 20-pattern SOM analysis of daily, wintertime, Northern Hemisphere sea level pressure reveals a continuum of patterns that correspond closely with well-known teleconnection patterns. This analysis also reveals that interdecadal variability of the hemispheric sea level pressure field may be understood in terms of changes in the frequency distribution within the continuum of sea level pressure patterns described by the SOM. Based on the continuum perspective illustrated with the SOM, the above secular shift of the NAO may be understood as a change in dominance from westward-displaced, negative NAO-like patterns to eastward-displaced, positive NAO-like patterns, though westward- and eastward-displaced NAO-like patterns existed during all time periods and for both phases.

## 1. Introduction

Low-frequency variability in the Northern Hemisphere geopotential height field is often described in terms of recurrent, persistent teleconnection patterns (e.g., Wallace and Gutzler 1981; Barnston and Livezey 1987), where a teleconnection may be defined here as a "simultaneous significant temporal correlation" in the geopotential height field between two widely separated locations (Van den Dool 2007). The North Atlantic Oscillation (NAO) and the Pacific–North American (PNA) patterns are considered the two dominant teleconnection patterns in the Northern Hemisphere. Previous works have suggested at least two interpretations for Northern Hemisphere teleconnection patterns. Tra-ditionally, investigators have treated such patterns as discrete, recurrent regimes. This interpretation has resulted in their description as "standing oscillations with geographically fixed nodes and antinodes" (Wallace and Gutzler 1981). Alternatively, a few studies have suggested the existence of a continuum of teleconnection patterns rather than a small number of geographically fixed standing oscillations. In support of the continuum perspective, Kushnir and Wallace (1989) found that teleconnection patterns obtained in rotated empirical orthogonal function (EOF) analysis often occur in orthogonal pairs with comparable amplitudes. This finding suggests that, rather than interpreting the loading vectors as discrete modes of variability, we instead may interpret the loading vectors as basis functions that describe a continuum of patterns. Franzke and Feldstein (2005) elaborated on this perspective by investigating the dynamical processes associated with a continuum of teleconnection patterns defined by a set of nonorthogonal basis functions. In providing addi-

tional support for the continuum perspective, they found that most members of the continuum have variance, autocorrelation time scales, and dynamical characteristics that are intermediate between those of the constituent basis functions.

Although the continuum perspective provides an attractive framework for describing teleconnection patterns, such a framework also introduces constraints on the interpretation of patterns obtained by conventional means such as EOF analysis. For example, if we adopt the perspective suggested by Kushnir and Wallace (1989), then we must consider linear combinations of loading vectors to describe teleconnection patterns most accurately. A compact method of illustrating the continuum of patterns is the use of joint probability distribution functions (e.g., Kimoto and Ghil 1993; Franzke and Feldstein 2005). The drawbacks of this method, however, include the limitation to only two basis functions in a single plot, which means that the two-dimensional phase space may explain only a relatively small fraction of the total variance. In addition, physical interpretation becomes challenging when visualizing probability distributions in phase space rather than patterns in physical space. To circumvent these challenges, we present an alternative method called self-organizing maps (SOMs) for the purpose of visualizing the continuum of Northern Hemisphere teleconnection patterns. This clustering approach, which we describe more thoroughly in section 2 and in the appendix, provides a means of visualizing the distribution of large-scale circulation patterns, yet treats the data as a continuum (Hewitson and Crane 2002; Reusch et al. 2007).

In this manuscript, we show how the continuum perspective, as illustrated with a SOM, may shed light on a secular change in the NAO position that began in the late 1970s. Hilmer and Jung (2000) first pointed out that an eastward shift of the northern center of the NAO pattern coincided with a positive trend in the NAO index (Hurrell 1995). This eastward displacement resulted in a significant increase in the correlation between the NAO time series and Arctic sea ice export through Fram Strait. Hilmer and Jung (2000) illustrated the shift by comparing NAO patterns for two separate 20-yr periods, 1958–77 (P1) and 1978–97 (P2), where the patterns were obtained by separate linear regressions of sea level pressure (SLP) anomalies onto a winter NAO index, which was defined as the difference between the normalized SLP time series from the Azores and Iceland. Jung et al. (2003) described additional changes in NAO-related anomaly patterns that occurred during the shift, including changes in patterns

of North Atlantic cyclone activity, surface air temperature anomalies, and turbulent surface heat flux anomalies in the Labrador Sea.

A few studies have suggested that the eastward displacement of the NAO may be closely related to the above trend from the negative to the positive NAO phase. This suggestion implies that the positive and negative phases of the NAO have different spatial patterns. In support of this contention, based on simulations with a dry general circulation model (GCM), Peterson et al. (2003) observed a nonlinear dependence on the NAO index such that the spatial pattern of the NAO was displaced eastward (westward) for anomalously high (low) index values. From this result, the authors concluded that the eastward shift of the NAO may have resulted from the transition from a negative to a positive NAO state. Using a weakly nonlinear NAO model, Luo and Gong (2006) illustrated that the mechanism responsible for the shift may have been an increase in the mean westerly winds in the NAO region during P2. From an alternative approach, Cassou et al. (2004) used cluster analysis of the National Centers for Environmental Prediction–National Center for Atmospheric Research (NCEP–NCAR) reanalysis data to categorize SLP patterns over the North Atlantic. This nonlinear approach to classification revealed spatial asymmetries between the two NAO phases: most prominently, the positive phase of the NAO featured an eastward shift of approximately $30°$ longitude in the position of the Azores high. Similar to Peterson et al. (2003), Cassou et al. (2004) concluded that the apparent shift in the NAO position resulted from the recent predominance of an intrinsically more eastward displaced NAO. In an earlier coupled ocean–atmosphere GCM study, Ulbrich and Christoph (1999) suggested that an eastward shift of the NAO may continue to accompany a positive trend in the NAO index in response to increased greenhouse gas forcing.

In this study, we investigate how a SOM analysis may present the continuum of Northern Hemisphere teleconnection patterns and offer insight into the changes in the NAO pattern that coincided with the observed eastward shift. The following section provides a brief overview of the SOM methodology, with a more thorough discussion of the method found in the appendix. In section 3, we present the SOM analysis of Northern Hemisphere SLP patterns and evaluate how the analysis describes hemispheric SLP anomaly patterns for three periods: P1 (1958–77), P2 (1978–97), and P3 (1998–2005). Section 4 applies the SOM analysis to describe the eastward shift of the NAO. The article follows with the conclusions in section 5.

## 2. SOMs and the continuum of teleconnection patterns

If we adopt a continuum perspective of teleconnection patterns, then each dominant pattern actually comprises a continuum of similar patterns, each with a unique probability of occurrence. For example, we might consider the NAO pattern as a weighted mean of all NAO-like SLP patterns. If we describe each NAO-like pattern as an $n$-dimensional vector $\mathbf{m}$, where $n$ describes the number grid points in the spatial domain, then we may express the mean NAO pattern as

$$\overline{\mathbf{m}} = \int \mathbf{m} p(\mathbf{m}) \, d\mathbf{m}, \qquad (1)$$

where $p(\mathbf{m})$ is the probability density function of $\mathbf{m}$, $d\mathbf{m}$ is a shorthand notation for the $n$-dimensional volume differential in phase space, and the integral is taken over the complete phase space of $\mathbf{m}$. Because we cannot examine all members of the continuum, we may choose a finite number of *representative patterns* to describe this distribution of NAO-like patterns. If we choose $K$ representative patterns, where $K$ is sufficiently large and denote each of these patterns as $\mathbf{m}_c^*$, then we may choose each $\mathbf{m}_c^*$ in a way that allows us to express a discretized form of (1) by

$$\overline{\mathbf{m}} = \sum_{c=1}^{K} \mathbf{m}_c^* p(\mathbf{m}_c^*). \qquad (2)$$

We may interpret $p(\mathbf{m}_c^*)$ as the probability of occurrence of the NAO-like pattern $\mathbf{m}_c^*$. If we extend (2) from phase space to physical space, we may express the mean, two-dimensional NAO pattern as

$$\overline{m}(x, y) = \sum_{c=1}^{K} m_c^*(x, y) p(m_c^*), \qquad (3)$$

where $x$ and $y$ denote positions in the physical domain. The SOM method provides a means of describing a continuum of patterns, as in (1), by a discrete number of representative patterns such that (2) and (3) may hold approximately (the method for determining the $\mathbf{m}_c^*$ is presented in the appendix).

Of course, the use of (1), (2), and (3) may apply more broadly to other teleconnection patterns and to other meteorological fields. In practical applications, we must rely on a set of discrete data samples to determine the representative patterns and the corresponding probability or frequency of occurrence. The SOM algorithm determines the representative patterns such that each $\mathbf{m}_c^*$ approximates the mean of the nearby data samples in the phase space, thus approximately defining a group centroid. In addition, the SOM method captures the

distribution of the data in that the algorithm tends to minimize the average differences between the samples in the dataset, that is, the daily SLP fields and the nearest representative patterns in phase space. If we describe each SLP field in the dataset as an $n$-dimensional vector $\mathbf{z}$, then we may express the difference between $\mathbf{z}$ and the best-matching representative pattern $\mathbf{m}_c^*$ in terms of Euclidean distance, $\|\mathbf{z} - \mathbf{m}_c^*\|$.[1] Thus, we may determine the best-matching representative pattern $\mathbf{m}_c^*$ for a particular sample $\mathbf{z}$ with

$$\|\mathbf{z} - \mathbf{m}_c^*\| = \min_i \{ \|\mathbf{z} - \mathbf{m}_c^*\| \}, \qquad (4)$$

where $i \in \{1, \ldots K\}$, $K$ is the number of representative patterns, and the subscript $c$ denotes the pattern for which the Euclidean distance $\|\mathbf{z} - \mathbf{m}_c^*\|$ attains a minimum. The self-organizing map algorithm tends to minimize the average distance between each sample in the dataset and the best-matching representative pattern. We describe this optimization as a minimization of

$$E = \frac{1}{N} \left( \sum_{t=1}^{N} \|\mathbf{z}_t - \mathbf{m}_c^*\| \right), \qquad (5)$$

where $E$ is termed the *average quantization error* (Kohonen 2001) and $N$ describes the number of samples in the dataset.[2]

The use of (4) and (5) indicates that each sample in the dataset becomes associated with a similar representative pattern in phase space. Therefore, we may classify each sample to a group defined by the approximate centroid, $\mathbf{m}_c^*$, as in cluster analysis. In addition, we may determine the frequency of occurrence of $\mathbf{m}_c^*$, or $f(\mathbf{m}_c^*)$, by calculating the ratio of samples associated with $\mathbf{m}_c^*$ to $N$, the total number of samples in the dataset. We then may use $f(\mathbf{m}_c^*)$ as an estimate for $p(\mathbf{m}_c^*)$ in (2) and (3).

The previous discussion indicates that, given a specified number of representative patterns, the SOM analysis tends to maximize the similarity between the representative patterns and the daily SLP fields in the dataset. In addition, the SOM patterns approximate the

---

[1] If the gridpoint density within the SLP fields is variable, then each entry in $\mathbf{z}$ must be weighted in proportion to the square root of the gridbox area to achieve the correct distance calculation. In this study, where the area of each grid box is proportional to the cosine of latitude, we may express $\mathbf{z}$ as $\mathbf{z} = \mathbf{A}\mathbf{z}_{uw}$, where $\mathbf{z}_{uw}$ is the unweighted SLP field and $\mathbf{A}$ is an $n \times n$ diagonal matrix containing the square root of cosine of latitude at the $n$ grid points.

[2] The SOM algorithm does not explicitly minimize this error criterion. The optimization described here happens to hold true approximately, which suggests a similarity between the SOM method and k-means clustering. A precise optimization criterion for the SOM is difficult to define (Kohonen 2001).

mean of all samples within a cluster so that we can view each SOM pattern as a composite of relatively similar daily SLP fields within a local cluster. These characteristics of SOM analysis provide reasonable assurance that *each representative pattern* corresponds with a *physical pattern* to some degree. This characteristic provides an *advantage of SOM analysis over EOF analysis*, as EOFs *may* or *may not represent* physical patterns. The degree of similarity between the representative patterns and the daily SLP fields depends on the number of representative patterns: as we increase the number of representative patterns, the daily and representative patterns attain a greater degree of similarity (correspondingly, the average quantization error decreases). A significant advantage of SOM analysis over other methods is that the SOM assumes that the data are continuous, so the method allows the representative patterns to span the phase space even when data are not highly clustered (see the appendix for two examples with artificial, unclustered data).

The SOM method most significantly distinguishes itself from other clustering methods through its topological ordering of the representative patterns. The patterns become organized on a regular, two-dimensional grid such that the SOM assigns similar representative patterns to nearby locations on the grid and dissimilar patterns to widely separated locations on the grid. Ultimately, the SOM method allows us to visualize the full set of representatives that describe patterns in physical space on a two-dimensional display that describes the complete phase space of the system.

## 3. The SOM of Northern Hemisphere SLP patterns

To illustrate the continuum of Northern Hemisphere SLP anomaly patterns, we use the SOM method to classify daily SLP anomaly data for the 48-yr period from 1958 to 2005. We analyze wintertime (December–March) Northern Hemisphere SLP data poleward of 20°N, derived from the NCEP–NCAR reanalysis dataset (Kalnay et al. 1996). We have removed the seasonal cycle at each grid point and refer to the resulting quantities as anomalies. Because the SLP field is on a 2.5° latitude–longitude grid, we weight the data by the square root of cosine of latitude in all subsequent analyses to account for increasing gridpoint density with latitude.

As described in section 2 and more thoroughly in the appendix, the SOM method reduces a potentially large set of data samples to a much smaller collection of representative patterns on a spatially ordered, two-dimensional grid. The user chooses the size of this two-

dimensional array, and we chose various array sizes to assess the sensitivity of the analysis to the size of the array. Ultimately, we chose the 4 × 5 SOM to classify the full 48-yr SLP dataset because that size represents a balance between economy and the level of resolved detail in the representative patterns. We obtain qualitatively similar results with both smaller and larger SOMs, so the broad interpretation that follows is rather insensitive to the choice of array size. After determining these representative patterns, we classify each daily SLP field in the 48-yr dataset to one of the 20 patterns using the criterion of minimum Euclidean distance, as described in section 2. Therefore, each day is associated with one particular SOM pattern. Each SLP anomaly map in the SOM (Fig. 1) approximately equals the mean of all observations within the local cluster defined by the SOM pattern; the pattern correlations between the SOM patterns and the corresponding composites of clustered observations range from 0.92 to 0.97 for this SOM. Thus, we validate the characterization of these SOM patterns as approximate group centroids, as discussed in section 2.

Because of the topology-preserving properties of the SOM, the most similar patterns neighbor each other whereas the most dissimilar patterns are located at opposite corners of the array. We provide a more quantitative description of the between-map similarity in Fig. 2, which presents the SOM distortion surface created by a Sammon mapping scheme (Sammon 1969). The between-circle distances in this two-dimensional plot approximate the actual Euclidean distances in the larger-dimensional SOM space. Figure 2 reveals that the patterns in the corners are relatively similar and that the greatest pattern differences generally lie between the second and third rows.

The maps in Fig. 1 provide a visualization of the continuum of hemispheric, wintertime SLP patterns. Because of the relatively small size of the SOM, this analysis retains only the low-frequency patterns. For patterns of low-frequency variability, the daily SLP fields occupy similar regions in phase space for several consecutive days. Therefore, the SOM analysis identifies these patterns because they occupy regions of high data density in phase space. In contrast, the SOM analysis is less likely to capture the synoptic-scale baroclinic waves because their rapid propagation indicates that the daily SLP fields are more widely separated in phase space over the same period of several days. As a result, we would require a greater number of SOM patterns to resolve the synoptic-scale baroclinic waves. Furthermore, because lower-frequency patterns exhibit larger spatial scales (Kushnir and Wallace 1989), the SOM analysis in Fig. 1 is dominated by large spatial
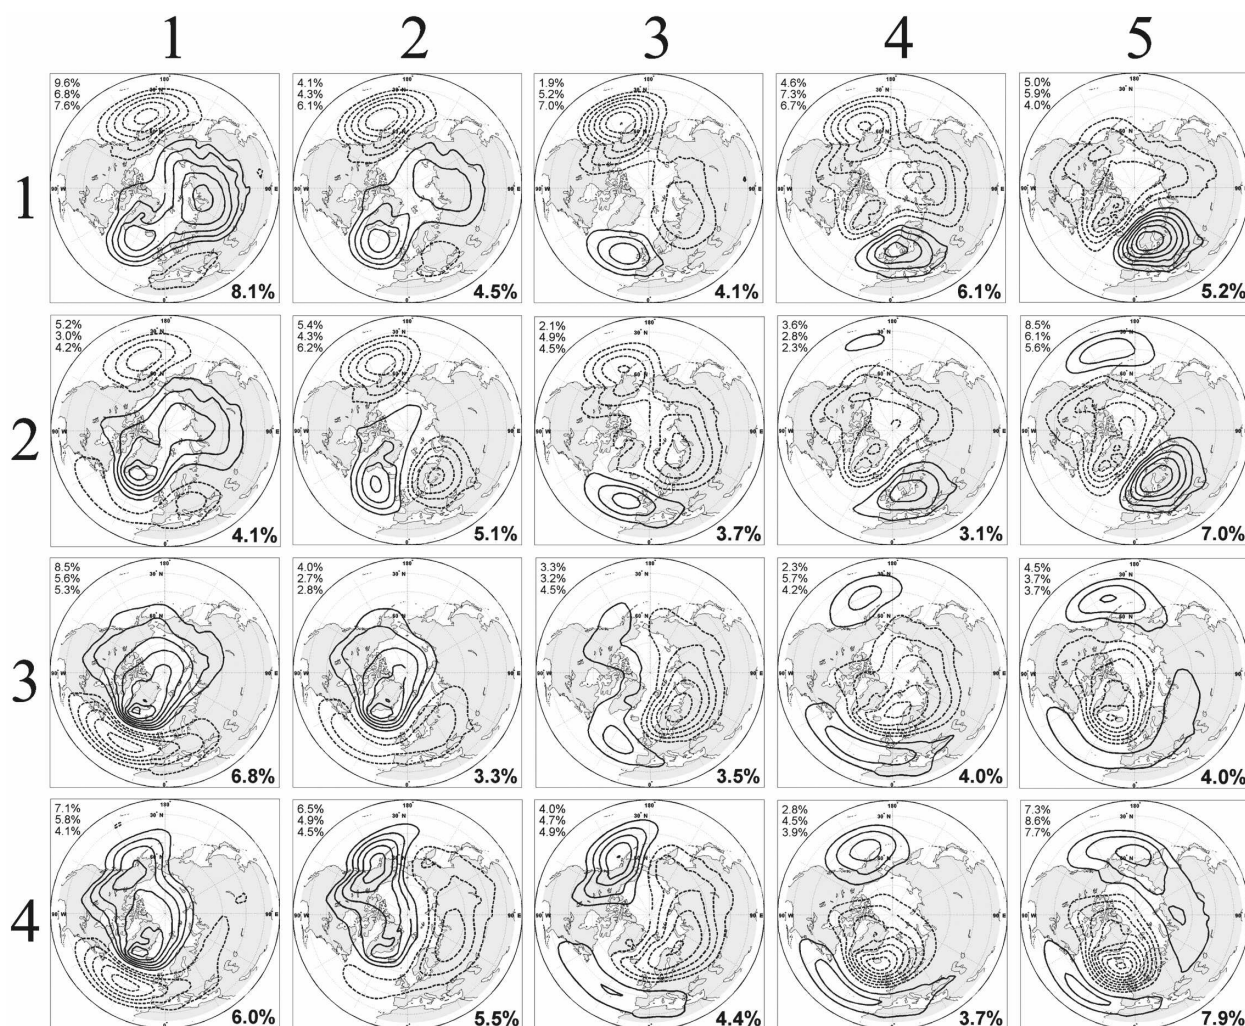
FIG. 1. The 4 × 5 SOM of SLP anomaly maps contoured at intervals of 2 hPa: solid (dashed) lines depict positive (negative) values, with the zero contour omitted. Percentages on the bottom right of each map describe the pattern frequency for the entire period; percentages on the top left of each map describe the pattern frequency for 1958–77 (P1: top), 1978–97 (P2: middle), and 1998–2005 (P3: bottom).
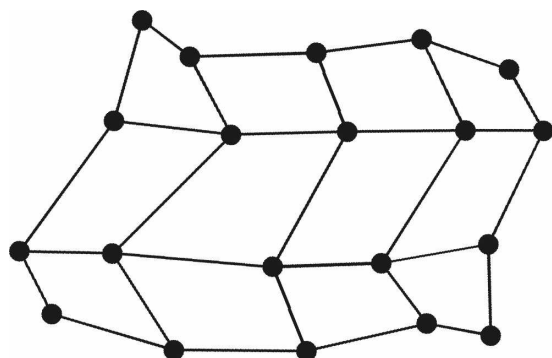


FIG. 2. Distortion surface created by a Sammon mapping scheme: the black circles, which describe the positions of the SOM patterns, have the same general orientation as in Fig. 1.

scales. If we were to increase the size of the SOM, the average quantization error would decrease, and we would obtain the possibility that the SOM analysis may, in fact, capture the synoptic-scale baroclinic waves along the major storm tracks.

The percentages at the bottom right of each map depict the pattern frequency, which we define as the percentage of observations classified to that particular pattern, for the entire 48-yr period. For the purpose of describing the NAO shift discussed in the introduction and in the following section, we subdivide the data into three periods: 1958–77 (P1), 1978–97 (P2), and P3 (1998–2005). The designations, P1 and P2, are identical to those of previous works (Hilmer and Jung 2000; Jung

et al. 2003; Peterson et al. 2003), whereas P3 corresponds with a more recent period during which the NAO index had transitioned to more neutral values than those of P2 (Overland and Wang 2005). The percentages at the top left of each map describe the pattern frequency for P1 (top), for P2 (middle), and for P3 (bottom).

Inspection of Fig. 1 reveals several of the classic patterns of hemispheric variability that correspond closely with well-known teleconnection patterns. For example, the patterns on the bottom left of the array resemble the canonical negative phase of the NAO, whereas the patterns at the bottom right depict the positive NAO polarity. The occurrence of several NAO-like SOM patterns represents an approximation of the full continuum of NAO-like patterns. In the top (bottom) row, we also observe patterns with an anomalously strong (weak) Aleutian low. A calculation of composite SLP anomaly fields (not shown) based on the PNA index of the Climate Prediction Center (CPC) also reveals patterns dominated by an anomaly centered near the Aleutians. Following the methods presented in sections 3 and 4, we find that the positive (negative) phase PNA composite corresponds to a linear combination of the SOM patterns found primarily in the top (bottom) two rows of Fig. 1. The patterns on the top right side resemble a more zonally oriented dipole reminiscent of the Eurasian patterns of Barnston and Livezey (1987). In the following two subsections, we evaluate the ability of these 20 patterns derived from daily data to capture the observed decadal SLP variability.

## a. The P1 and P2 periods

The change in pattern frequencies from P1 to P2 illustrated in Fig. 1 reveals a general increase for the patterns on the bottom right and top middle and a decrease for the patterns on the left. This figure also reveals that both eastward and westward displaced NAO-like patterns exist for both phases. Based on the pattern frequencies, the negative NAO-like patterns occurred more frequently than the positive NAO-like patterns during P1, and the dominant negative NAO-like patterns generally were more westward displaced [notice the anomalously frequent patterns at (2, 1), (3, 1) and (4, 1)]. Although the positive NAO phase was less prevalent during P1, the more dominant positive NAO-like patterns generally displayed a westward displacement as well, as evidenced by the anomalously frequent patterns at (2, 4), (2, 5), and (3, 5). In contrast, during P2 the positive NAO-like patterns occurred more frequently than the negative NAO-like patterns. In addition, the dominant NAO-like patterns tended to be

more easterly displaced, as supported by the anomalously frequent patterns at (3, 4), (4, 4), and (4, 5). Therefore, the changes in pattern frequency displayed in Fig. 1 are consistent both with the positive trend and the eastward shift of the NAO. We note, however, that all NAO-like patterns within the continuum existed during each period.

To address more thoroughly the sufficiency of this SOM analysis in capturing the aggregate change in SLP patterns between P1 and P2, we evaluate how well the application of (3) to the SOM patterns approximates the aggregate anomalies for each period. Specifically, we evaluate

$$\bar{z}_i(x, y) = \sum_{c=1}^{20} m_c^*(x, y) f_i(m_c^*), \quad i = 1, 2, \quad (6)$$

where $\bar{z}_i(x, y)$ denotes the composite Northern Hemisphere SLP anomaly pattern for period $i$, $m_c^*(x, y)$ denotes SOM pattern $c$ in Fig. 1, and $f_i(m_c^*)$ is the pattern frequency of pattern $m_c^*$ for period $i$, as depicted in Fig. 1. To evaluate the adequacy of these SOM-derived composites, we compare the patterns determined by (6) with the true composite anomaly patterns derived from the full hemispheric datasets for P1 and P2.

Figure 3 illustrates both the pattern derived from (6) and the actual composite pattern for each period. First, we note the use of a smaller contour interval for the SOM-derived patterns (Figs. 3a and 3b) than for the true composite anomaly patterns (Figs. 3c and 3d) so that we can compare qualitatively the paired patterns. The patterns derived from (6) possess weaker maximum anomalies compared with the actual anomalies for the two time periods. This result likely derives from the restriction of this analysis to only 20 patterns, as we shall see shortly. Nonetheless, the SOM-derived composite patterns compare rather favorably with the actual composite anomalies. The SOM analysis accurately captures the dominance of the negative NAO pattern during P1, as confirmed by inspection of Fig. 3c. The most notable distinction between Figs. 3a and 3c is the negative Aleutian anomaly in the SOM composite instead of the positive anomaly that is actually observed. Further investigation into the reason for this discrepancy reveals that the SOM patterns, particularly those in the third and fourth rows, substantially underestimate the strength of the positive anomaly over the North Pacific during P1. During P2, the SOM-derived composite accurately captures the dominance of the positive NAO, the anomalously strong Aleutian low, and the belt of anomalously low pressure throughout much of northern Eurasia.
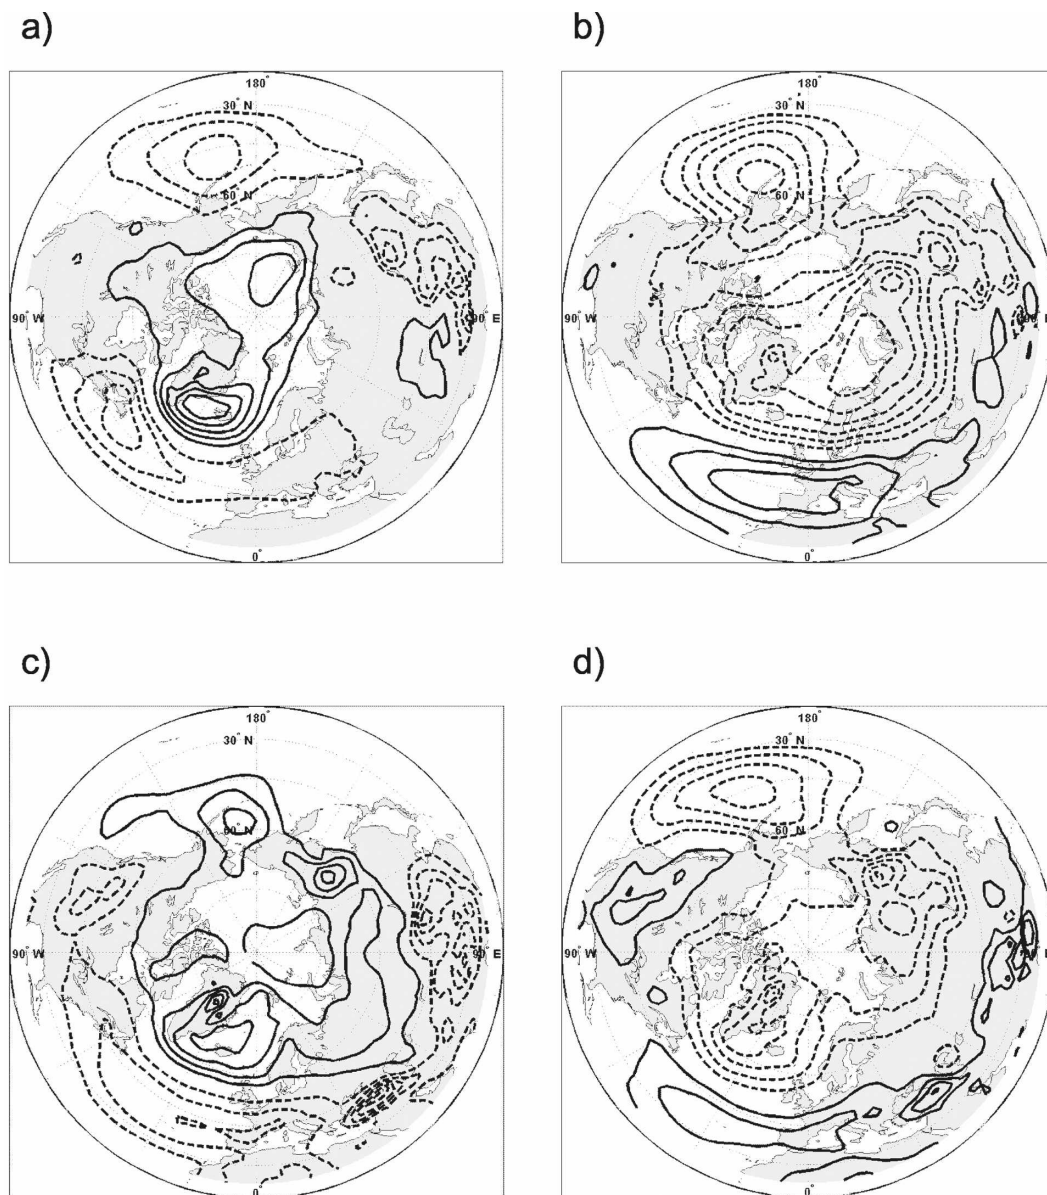
a)

b)

c)

d)

FIG. 3. SOM-derived composite SLP anomalies for (a) P1 and (b) P2 and the actual SLP anomaly fields for (c) P1 and (d) P2: contour interval is 0.2 hPa for (a) and (b) and 0.5 hPa for (c) and (d) with the same contouring conventions as in Fig. 1.

The pattern correlation (Horel 1985) between the SOM-derived composite and the actual anomaly composite is 0.42 (0.83) for P1 (P2). Clearly, the discrepancy near the Aleutians in the P1 patterns plays a large role in the significantly lower correlation relative to P2. If we perform the analysis with a $5 \times 7$ instead of a $4 \times 5$ SOM, then the pattern correlations increase to 0.64 and 0.88, respectively. This increase in the pattern correlations results from the increased level of detail in each of the 35 analysis patterns as well as a reduction in the border effect with the increased array size (see the ap-

pendix for a description of the border effect). We have chosen to illustrate the 20-pattern SOM in this manuscript, however, because the smaller SOM still captures the general changes of the NAO that occurred between P1 and P2. Therefore, we conclude that the $4 \times 5$ SOM analysis described here provides a reasonably accurate description of the changes in pattern frequency distribution that occurred between P1 and P2, though larger SOM sizes provide additional pattern detail that may be necessary in other applications. This result suggests that the difference in the P1 and P2 SLP anomaly
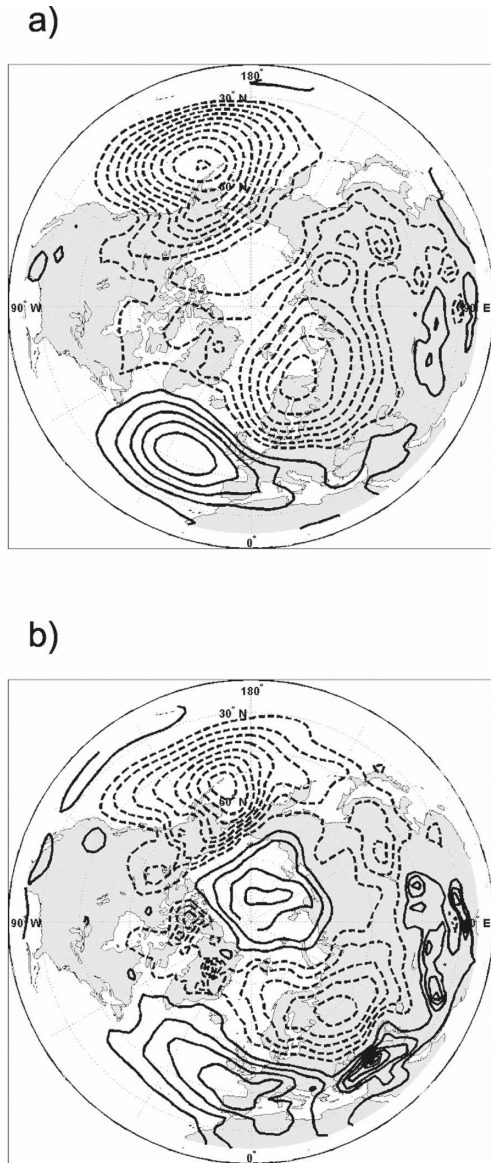
a)



b)



FIG. 4. The (a) SOM-derived composite SLP anomalies and (b) actual SLP anomalies for P3: contour interval and contouring conventions as in Fig. 3.

fields can be attributed to changes in the *frequency distribution* within the *continuum* of *hemispheric* SLP patterns.

## b. The P3 period

Because the trends in the NAO and in the closely related Arctic Oscillation (AO) (Thompson and Wallace 1998) indices may have ceased or perhaps even reversed since the end of P2 (Overland and Wang 2005), we briefly consider the prevailing SLP patterns described by the 4 × 5 SOM in the 8-yr period follow-

TABLE 1. Pattern correlations and Euclidean distances (hPa) between the SOM-derived hemispheric SLP anomaly field and the actual hemispheric SLP anomaly fields for each of the three periods considered: in each cell the pattern correlation lies above the Euclidean distance.

| Number of SOM patterns | P1 | P2 | P3 |
|---|---|---|---|
| 20 (4 × 5) | 0.42 | 0.83 | 0.76 |
| | 35.1 | 27.4 | 35.5 |
| 35 (5 × 7) | 0.64 | 0.88 | 0.74 |
| | 30.3 | 25.9 | 36.3 |
| 96 (8 × 12) | 0.79 | 0.92 | 0.78 |
| | 25.4 | 21.2 | 33.2 |
| 300 (15 × 20) | 0.84 | 0.92 | 0.85 |
| | 24.1 | 19.8 | 29.2 |

ing P2. With respect to the 48-yr average, the largest positive pattern frequency anomalies in P3, as revealed in Fig. 1, correspond to the patterns at positions (1, 2), (1, 3), (2, 2), and (3, 3). The first three listed patterns depict an anomalously strong Aleutian low, and all four patterns indicate to some degree a zonal pressure dipole in the North Atlantic and Eurasia. As in the previous subsection, we evaluate how well the SOM frequency distribution for this period captures the aggregate hemispheric SLP anomalies. Figure 4 presents the SOM-derived composite constructed by the application of (6) (Fig. 4a) and the actual anomaly composite (Fig. 4b) for this period. Comparison between these two plots reveals that the SOM analysis performs rather well in capturing the salient features of the anomaly pattern for this period, as supported by the high pattern correlation of 0.76 between the two maps. The SOM-derived composite accurately reveals the strong negative anomaly near the Aleutians and the dipole between northern Europe and the southwestern Europe/ North Atlantic region. The dominant SOM patterns and the composite anomaly pattern for this recent period seems to bear close resemblance to the cold ocean– warm land (COWL) pattern described in Wallace et al. (1996). The most notable discrepancy between the SOM-derived and the actual composites lies over the pole, where the SOM analysis fails to capture the positive SLP anomaly.

Table 1 provides a summary of the similarity between the SOM-derived composites and the observed composites for the three analyzed periods and for various SOM sizes. We include two measures of similarity: the pattern correlation and the Euclidean distance between the two patterns, where we obtain the latter measure when we express each pattern as a vector, as described in section 2. We include the Euclidean distance because this measure provides information on the differences in the anomaly magnitudes between the two patterns.
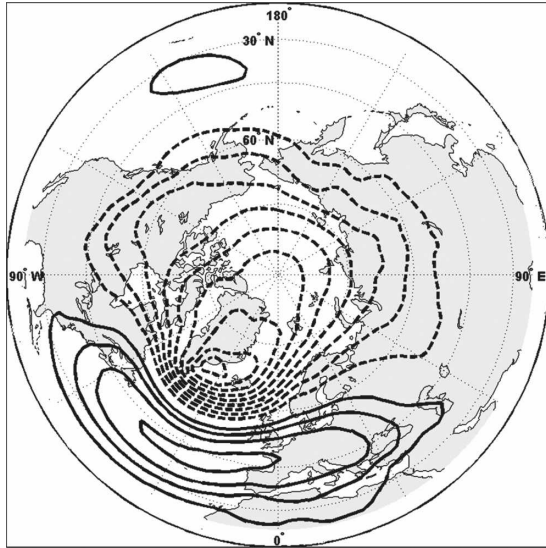
FIG. 5. The leading EOF of daily sea level pressure: contouring conventions as in Fig. 1 and contour interval arbitrary.

(For reference, the $15 \times 20$ SOM generally captures approximately 50%–75% of the magnitude of the actual anomalies.) As expected, the agreement between the SOM-derived and actual composites increases as the size of the SOM increases, though the statistical significance of even the smallest pattern correlations in the table well exceeds the 99% confidence level.

## 4. Application of the SOM to changes in the NAO

In this section, we apply the continuum perspective based on the above SOM analysis to investigate the secular changes in the NAO discussed in the introduction. The approach that we adopt is to analyze the frequency distribution within the SOM for those days when the NAO has a large amplitude. We define the NAO as the leading EOF of the Northern Hemisphere, wintertime SLP dataset for the period 1958–2005, as described in section 3. The leading EOF explains 10.5% of the SLP variance. We then define the associated standardized PC time series as the NAO index. Typically, the leading, unrotated EOF of SLP has been defined as the northern annular mode (NAM) or AO (e.g., Thompson and Wallace 1998), but the AO is closely related to the NAO and, in fact, may be statistically indistinguishable from the NAO (Feldstein and Franzke 2006). The leading NAO loading pattern (Fig. 5) agrees well with those of previous studies (e.g., Thompson and Wallace 1998) except for the weaker than expected Pacific center of action characteristic of

the AO. This discrepancy likely results from the use of daily rather than monthly mean data in the EOF analysis, as the Pacific center of action is strongest when the averaging period of the data is 30 days (Woollings and Hoskins 2008). Because of the weak Pacific center of action and the more regional nature of the pattern in Fig. 5, we retain the NAO nomenclature in favor of the NAM/AO terminology. In any case, the results in this study are insensitive to the choice of NAO pattern, as we obtain qualitatively similar results with time series of more regionally confined rotated leading EOFs and with the NAO time series of the CPC.

To assess the change in NAO position, we compare composite SLP anomalies for NAO events with large index magnitudes for each of the two periods, P1 and P2. We establish an arbitrary index threshold of 1.5 and define a strong NAO event as any sequence in the time series with an NAO index magnitude exceeding this threshold for at least three consecutive days. To ensure approximate temporal independence and because of the approximate 10-day time scale of the NAO (Feldstein 2000), we choose only the day with the largest index magnitude during each event for the composites and we establish a minimum separation of ten days between events. We separate positive from negative events and distinguish events between P1 and P2.

The number of samples $n$ contributing to each composite (bottom left of Figs. 6a,b,d,e) varies substantially between P1 and P2 in accordance with the transition from a predominantly negative NAO in P1 to a positive NAO in P2. The mean NAO index associated with each composite, however, exhibits virtually no change between the two periods. The pair of SLP anomaly composites for strong positive events (Figs. 6a and 6b) and for strong negative events (Figs. 6d and 6e) reveals the spatial asymmetry in the Azores anomaly center between the two NAO phases, as noted in Cassou et al. (2004). In addition, the Icelandic anomaly exhibits greater eastward extent through Scandinavia during positive events. Therefore, the eastward shift in the NAO may have resulted from an increase in the fraction of positive events in P2 relative to P1 and not from a longitudinal displacement of any of the centers of action in either phase. This view that the eastward shift of the NAO pattern arose from the trend from predominantly negative to positive NAO events is consistent with the regime interpretation of teleconnection patterns, as described in the introduction. In contrast, from the continuum perspective, there is no eastward shift in the NAO pattern. Instead, two important characteristics of the continuum take place: 1) During P1 members of the negative NAO continuum occur more
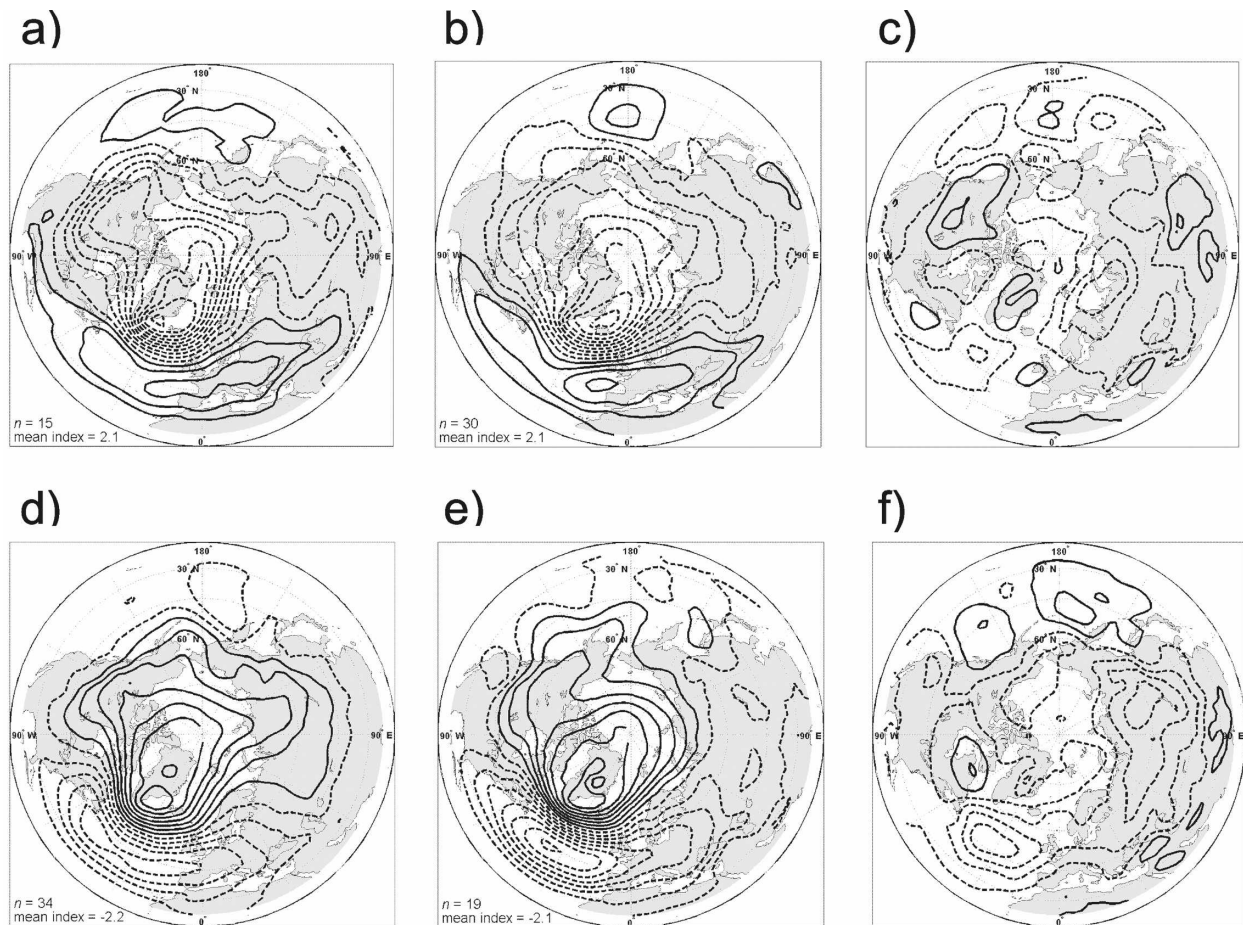
FIG. 6. The (a) SLP anomaly composites for positive NAO events during P1 and (b) P2 and for negative NAO events during (d) P1 and (e) P2. The difference between composites (P2 − P1) for (c) positive events and (f) negative: contour interval 3 hPa and zero contour omitted in all plots.

frequently than those of the positive NAO continuum, and vice versa for P2 (Fig. 1), and 2) during P1 (P2) westward (eastward) displaced members of the negative (positive) NAO continuum occur most frequently.

### a. Frequency distribution of NAO events

We next use the SOM analysis to investigate whether the differences in the NAO between P1 and P2 can be interpreted from the continuum perspective. To generate NAO composites for P1 and P2, as for the calculations that generated Fig. 6, we isolate all NAO events, but define an NAO event as *any day* with an index magnitude exceeding 1.5, rather than only those days with a peak index value within a sequence of days. Therefore, we relax the approximate temporal independence constraint that we need for tests of statistical significance in order to increase the sample size for the analysis in this subsection. Again, we separate positive

from negative events and distinguish events between P1 and P2. We then identify to which SOM pattern each NAO event corresponds. Figure 7 illustrates the percentage of positive (right) or negative (left) NAO events that are assigned to each SOM pattern during P1 (Fig. 7a) and during P2 (Fig. 7b), with the thick white line separating negative events from positive events. Therefore, the percentages both to the left and to the right of the white line sum to 100, accounting for all days with an NAO index magnitude exceeding 1.5 during the period. Figure 7 reveals that positive NAO events occupy the region of SOM space represented by the three rightmost columns of the array, whereas negative events occupy the region represented by the two leftmost columns.

To assess the adequacy of the SOM in capturing the changes in frequency distribution, we again evaluate how well the application of (3) captures the change in spatial pattern for each of the NAO phases. In this case, we evaluate

$$\overline{m}_{P2}(x, y) - \overline{m}_{P1}(x, y) = \sum_{c=1}^{20} m_c^*(x, y) f_{NAO,P2}(m_c^*)$$

$$- \sum_{c=1}^{20} m_c^*(x, y) f_{NAO,P1}(m_c^*), \quad (7)$$

where the left-hand side represents the pattern of SLP differences between the NAO composites (P2 − P1) and the right-hand side is the difference of composite patterns obtained by summation over the SOM patterns, as in (6). The notation of the right-hand side of (7) is identical to that of (6) except that $f_{NAO,P1(P2)}$ now corresponds to the pattern frequency of NAO events in P1 (P2), as depicted in Figs. 7 and 7b.

We compare the SOM-derived difference composites obtained from (7) (Figs. 8a and 8b) with the actual difference composites (Figs. 8c and 8d; notice that Figs. 8c and 8d are similar but not identical to Figs. 6c and 6f because of the variation in how we define an NAO event in this subsection). As in Figs. 3 and 4, the SOM-derived difference composites based on NAO events significantly underestimate the magnitude of the differences between the two periods, particularly for the negative phase differences. Nevertheless, the anomaly difference patterns compare quite favorably between the SOM-derived and the actual difference composites. This agreement suggests that the changes in the NAO patterns between P1 and P2 correspond to *changes in the frequency distribution* within the *NAO continuum*, as illustrated with this SOM analysis.

## b. Distinguishability between P1 and P2 composites

The results of the previous subsection suggest that each of the NAO composites, for P1 and P2 (Fig. 6), correspond to the time mean of NAO-like patterns within the NAO continuum. Furthermore, as we saw in Figs. 1 and 7c, the pattern frequencies of some NAO-like patterns undergo large changes between P1 and P2. This leads to the interesting observation that the frequency distribution of NAO-like patterns within the NAO continuum undergoes substantial interdecadal variability.

In this subsection, we address the question of whether there is large interdecadal variability of the NAO continuum. This is a difficult question to address with observational data because the data that we have to investigate the NAO is limited to a time interval of several decades. Notwithstanding these limitations, it appears that an examination of the differences in the NAO continuum between P1 and P2 is worthwhile— because of the dominance of the NAO during both time periods and because of the large changes in pattern
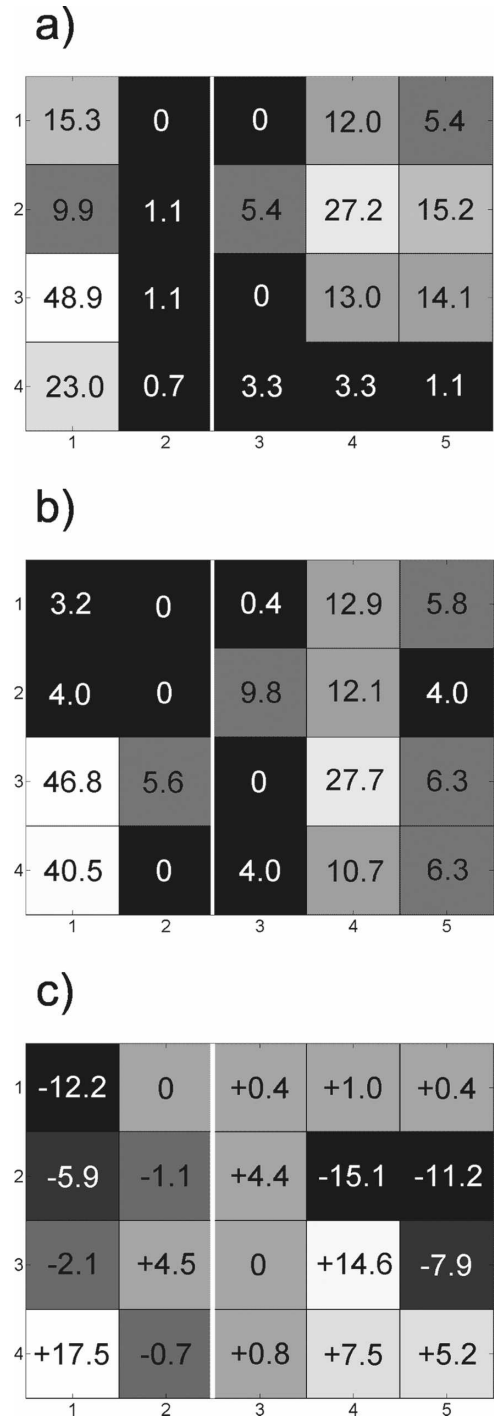


FIG. 7. The pattern frequencies (%) for NAO events during (a) P1 and (b) P2, and (c) the difference in pattern frequencies between the two periods. The thick white line separates (right) positive events from (left) negative events; light (dark) shades within grayscale denote large (small) percentages.
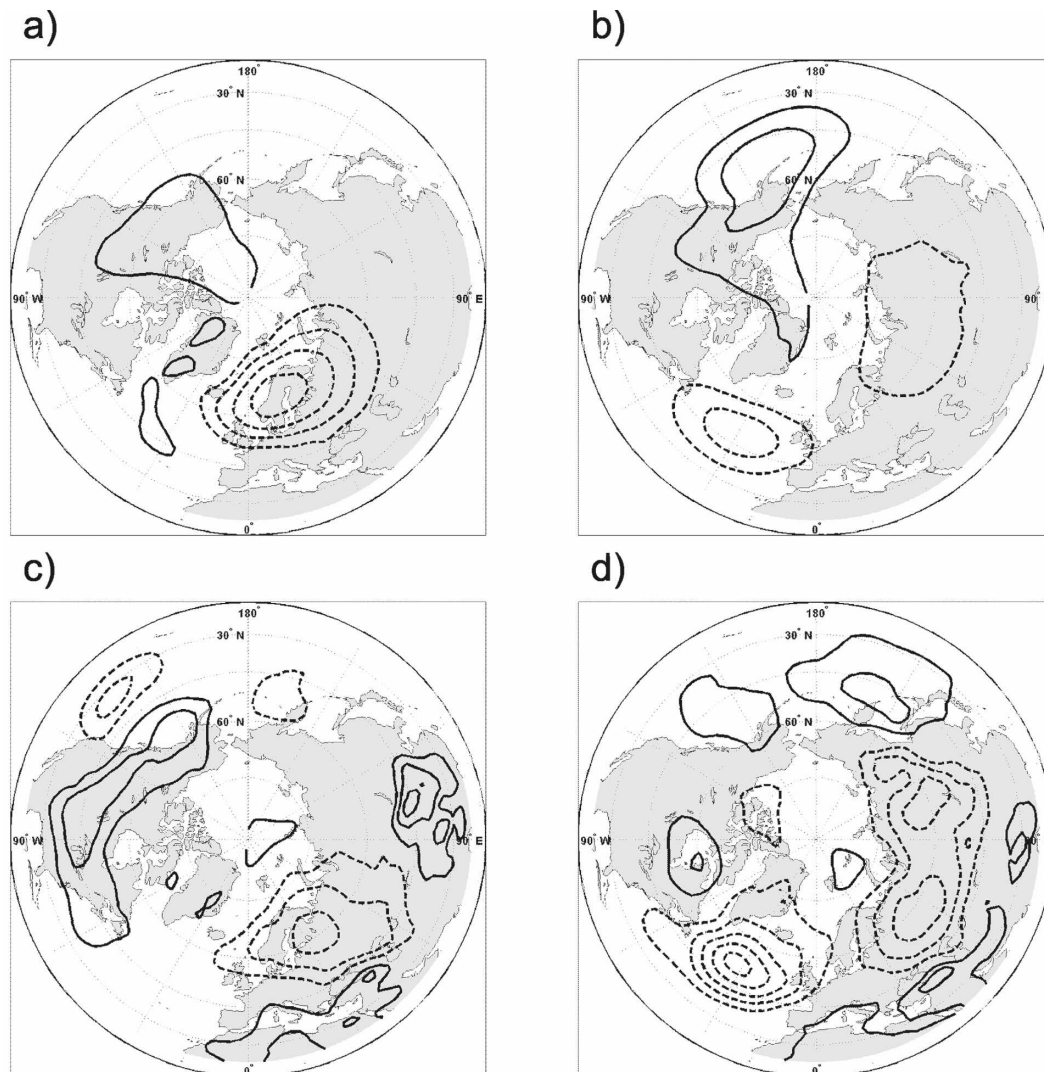
FIG. 8. The (a) SOM-derived composite SLP anomaly differences (P2 − P1) for positive NAO events and (b) for negative NAO events, as constructed by (9), and the actual composite SLP anomaly differences for (c) positive and for (d) negative NAO events: contour interval 1 hPa in (a) and (b) and 2 hPa in (c) and (d); all other contouring conventions identical to those of Fig. 1.

frequencies for some patterns (Fig. 7c). One possible approach is based on (6), which relates the pattern frequencies to the corresponding composite field. One could test the null hypothesis that each of the pattern frequencies in P1 and P2 (Figs. 7a and 7b) comes from the same population. If the statistical tests suggest that the null hypothesis be rejected for a sufficiently large number of patterns, then one could claim that the NAO continuum has undergone a distinguishable change between P1 and P2. Alternatively, one could ask the question of whether the NAO events that comprise the NAO composites come from the same population. Because (6) relates the composite field to the pattern frequencies, the examination of this question combines the evaluation of all of the NAO pattern frequencies into a single test. If the statistical analysis leads to a rejection of the null hypothesis, then we claim that the NAO composites for P1 and P2, and the corresponding pattern frequencies through (6), are distinguishable.

To evaluate the null hypothesis for the composite NAO patterns, we begin by evaluating the statistical significance of the difference composites illustrated in Figs. 6c and 6f with the Student's $t$ distribution for the difference of means; specifically we determine whether the variable

$$T(\lambda, \theta) = \frac{\overline{X_1}(\lambda, \theta) - \overline{X_2}(\lambda, \theta)}{S(\lambda, \theta)\sqrt{1/n_1 + 1/n_2}}, \qquad (8)$$

where

$$S(\lambda, \theta) = \sqrt{\frac{(n_1 - 1)S_1(\lambda, \theta)^2 + (n_2 - 1)S_2(\lambda, \theta)^2}{n_1 + n_2 - 2}} \quad (9)$$

exceeds the 80% confidence level for a two-sided $t$ test. The variable $\overline{X_1}$ $(\overline{X_2})$ signifies the composite SLP anomaly field for P1 (P2) with each of the two phases tested separately. The quantity $n_1(n_2)$ describes the number of samples comprising the SLP composite for P1 (P2), $S_1$ ($S_2$) indicates the standard deviation of the P1 (P2) composite, and $\lambda$ and $\theta$ are latitude and longitude, respectively. After calculating the values of $T$ at each grid point, we then determine the fractional area with SLP differences that exceed this 80% confidence level (Livezey and Chen 1983). (By focusing on the 80% confidence level rather than the more commonly used 95% confidence level, we are imposing a stricter test for acceptance of the null hypothesis.)

We adopt a Monte Carlo approach to examine field significance for each phase separately. We randomly select $n_1$ and $n_2$ samples from the entire dataset and calculate a difference composite from this pair of random datasets. As described above, we calculate the fractional area with $T$ values that exceed the 80% confidence level for this new composite. We repeat this simulation with random data combinations 10 000 times to obtain a distribution of fractional area with statistically significant differences at or above the 80% confidence level. We determine the field significance for the difference composites in Figs. 6c and 6f by noting where the observed significant fractional area lies within the distribution obtained by this Monte Carlo method.

For the observed positive (negative) phase composites, 20.4% (40.6%) of the area exhibits statistically significant differences at or above the 80% confidence level between P1 and P2. Based on the distributions obtained with the Monte Carlo simulations, we conclude that the positive phase composites are statistically indistinguishable, as 45% of the simulated composites have a statistically significant fractional area that is larger than that which is actually observed in the atmosphere. In contrast, only 0.2% of the simulated composites have statistically significant fractional area that is greater than that observed for the negative phase composites. Therefore, in this case, we reject the null hypothesis at the 99% confidence level and conclude that the negative phase did undergo a statistically significant change in form between P1 and P2. For each of these two phases, we obtain similar results for global significance when we use the more common local significance levels of 90% and 95%.

Even though Fig. 7c illustrates that both the positive

and negative phases underwent a substantial frequency redistribution within the SOM, we notice that the largest changes in frequency of positive NAO-like patterns between P1 and P2 occurred with neighboring, and therefore similar, patterns in the SOM. In contrast, the patterns with the greatest change in frequency for the negative phase are more widely separated within the SOM, and are therefore more distinct. This contrast in pattern separation within the SOM for those patterns with the greatest change in frequency likely explains why the change in pattern for the negative phase is statistically significant, whereas the positive phase pattern change is not.

## 5. Conclusions

In this study, we use the method of self-organizing maps to advance the continuum perspective of Northern Hemisphere teleconnection patterns and to shed light on the secular eastward shift of the NAO. The SOM analysis of Northern Hemisphere SLP provides a set of representative patterns that correspond with well-known teleconnection patterns such as the NAO and the PNA. We show that the different wintertime hemispheric SLP anomaly patterns for three distinct periods, P1 (1958–77), P2 (1978–97), and P3 (1998–2005), may be considered to have arisen from changes in the frequency distribution within the continuum of teleconnection patterns, as described by the SOM. Despite the small size, the 20-pattern SOM still provides enough detail to describe interdecadal variability of the hemispheric SLP field, at least in a qualitative sense. Increasing the number of patterns increases the accuracy of the SOM in capturing the shape and magnitude of these decadal-time-scale SLP anomaly fields, though the maximum number of "manageable" representative patterns remains a subjective decision.

In this study, we have focused on the distribution of NAO-like patterns in the SOM, which represent an approximation of the continuum of NAO-like patterns. Based on the continuum perspective, we describe two important characteristics of the secular changes in the NAO that began in the late 1970s: 1) During P1 (P2) members of the negative (positive) NAO continuum predominate over those of the positive (negative) NAO continuum, and 2) during P1 (P2) westward (eastward) displaced members of the NAO continuum occur most frequently. From this perspective, we acknowledge that all NAO-like patterns exist during the entire period under investigation, though the frequency of occurrence of each pattern varies between P1 and P2. Based on tests of field significance, we determine that the P1 and P2 negative NAO patterns are statistically distinguishable, though the two positive phase NAO patterns

are indistinguishable. These results suggest that interdecadal variability of the NAO continuum encompasses significant variability in the frequency distribution of the NAO-like patterns within the continuum.

The SOM methodology may prove to be useful for further investigation of the continuum of teleconnection patterns. On one hand, EOF analysis and associated composites provide useful information on aggregate patterns, but they conceal information about the within-group variability of the members that make up the composites. On the other hand, examination of the individual members likely would be too cumbersome to be practical. SOMs provide a compromise by approximating the continuum with a manageable number of representative patterns. The use of (6) provides a simple framework for the interpretation of observed composite patterns and for the evaluation of a SOM analysis.

An improved understanding of the continuum of teleconnection patterns may help to clarify the relationships between these patterns and key climate variables. For example, the first report of the eastward shift by Hilmer and Jung (2000) illustrates how subtle changes in the mean position of the NAO can introduce significant changes in the relationship between the NAO and sea ice export. More recently, Polyakova et al. (2006) noted additional, continually changing relationships between the NAO and key North Atlantic climate parameters during the past century. These observations suggest that periodic fluctuations in the frequency distribution of NAO-related patterns result in significant variations in the climate response to the NAO. SOM analysis provides an additional tool for exploring these changes in frequency distribution, which may provide insight into the nonstationary relationships between the NAO index and other climate variables. Future studies may choose to extend the current SOM analysis to include, for example, other variables such as upper-tropospheric geopotential height, more patterns, and relationships to other climate variables.

## APPENDIX

### The SOM Methodology

The self-organizing map (SOM) is an analysis method arising from the field of artificial neural networks. The SOM methodology essentially reduces a potentially large, high-dimensional dataset to a manageable number of representatives that we can visualize on an ordered low-dimensional array. This low-dimensional array may be, for example, one-, two-, or three-dimensional, but the two-dimensional choice is the most common one for the reason of convenience. With two dimensions, we may present the representatives in an organized and easily visualized way, as in Fig. 1. Each representative is of the same dimensionality as that of the input data, so, in the case of the present study, the representatives correspond to spatial patterns (for the remainder of this discussion we shall refer to these representatives as patterns). As is characteristic of neural networks, the SOM algorithm presents the input data to a layer of neurons or, equivalently, a layer of nodes. In the case of the SOM, the nodes describe positions on the two-dimensional grid that contains the representative patterns. In addition, the nodes have associated "reference vectors" that, through an iterative procedure, ultimately become the representative patterns.

Although there are variations of the SOM algorithm, the basic concepts hold for each variation. We describe here the original algorithm described in Kohonen (2001). To clarify the concept of the basic SOM, we illustrate with an analysis of an artificial, two-dimensional dataset consisting of 10 000 random data vectors whose two components have values uniformly distributed between 0 and 1. We may choose any size for the SOM, so here we choose to describe the distribution of the data with a 10 × 10 SOM (i.e., 100 representative patterns shall be displayed on a grid with 10 rows and 10 columns). We illustrate each of the 100 nodes in the two-dimensional grid as circles in Fig. A1a and have numbered the first and last three in a row-wise manner for reference. We have spaced the circles uniformly, though we have no a priori reason to suppose that the data are uniformly distributed. We shall consider the depiction of the nodes in Fig. A1a as a representation of the *topology* of the SOM: neighboring nodes shall correspond to similar patterns and widely separated nodes shall correspond to dissimilar patterns.

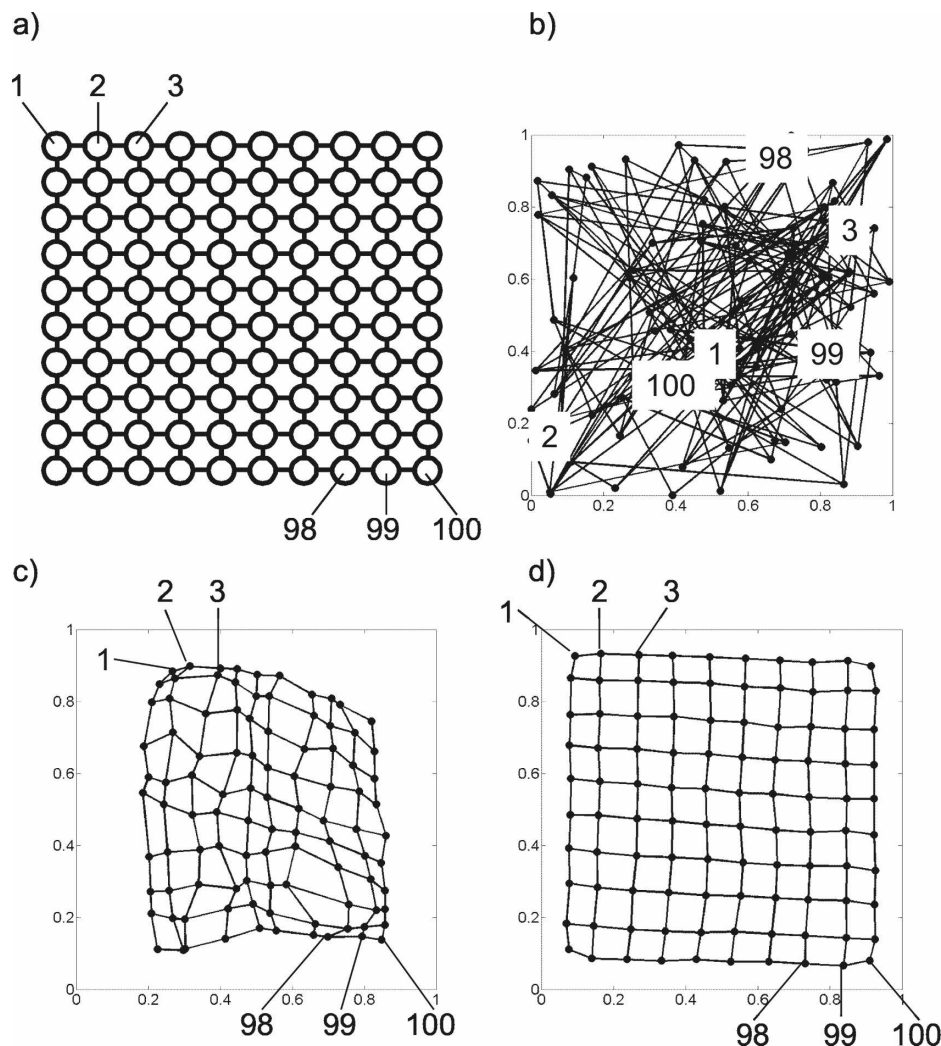As mentioned above, each node in the SOM is asso-

FIG. A1. SOM classification of 10 000 random two-dimensional data vectors with components uniformly distributed between 0 and 1: (a) Schematic depiction of nodes in a 10 × 10 SOM, with the first 3 and last 3 nodes numbered for reference; (b) initial distribution of reference vectors corresponding to each node; (c) broadly ordered reference vectors after 500 training iterations; and (d) fine-tuned reference vectors after an additional 50 000 iterations of training.

ciated with a reference vector equal in dimension to that of the data being analyzed. Therefore, each node $i$ in Fig. A1a is associated with a two-dimensional reference vector $\mathbf{m}_i$ that corresponds to a two-dimensional pattern. The method of determining the optimal representative patterns begins with the assignment of initial reference vectors to each of the nodes as a "first guess" for the representative patterns. We may even assign random vectors to each node initially, though we obtain the optimal patterns more efficiently with a better first guess that corresponds to a more ordered initial distribution (e.g., along the two-dimensional subspace spanned by the first two EOFs). In our present example, we have initialized the map with random refer-

ence vectors, as illustrated in Fig. A1b. Here we have chosen two-dimensional data for ease of visualization and simply illustrate the reference vectors on an $x$–$y$ plane. We also have labeled the first and last three reference vectors corresponding to the first and last three nodes illustrated in Fig. A1a.

To begin the iterative procedure, known as "training" of the SOM, we choose a data vector $\mathbf{z}$ from the dataset and calculate a general distance measure, usually Euclidean distance, between $\mathbf{z}$ and all reference vectors $\mathbf{m}_i$. The best-matching node with the smallest distance between its reference vector and the data vector $\mathbf{z}$ is considered the "winning" node $c$. We describe this determination of node $c$ with

$$\|\mathbf{z} - \mathbf{m}_c\| = \min_i \{\|\mathbf{z} - \mathbf{m}_i\|\}, \qquad \text{(A1)}$$

where $i \in \{1, \dots, 100\}$ and the subscript $c$ denotes the pattern for which the Euclidean distance $\|\mathbf{z} - \mathbf{m}_i\|$ attains a minimum. At this point, the *learning* process activates the winning node as well as those nodes that are topographically close to the winning node in the SOM depicted in Fig. A1a. The activated nodes adjust their reference vectors a specific amount toward the input vector $\mathbf{z}$ by the relation

$$\mathbf{m}_i(t + 1) = \mathbf{m}_i(t) + h_{ci}(t)[\mathbf{z}(t) - \mathbf{m}_i(t)], \qquad \text{(A2)}$$

where $t$ is the discrete time step integer for the map training and $h_{ci}(t)$ is the so-called *neighborhood kernel*, described more thoroughly below. The neighborhood kernel acts to smooth the reference vectors in this neighborhood centered on the winning node. Therefore, the winning node often makes the greatest adjustment of its reference vector and all other neighboring nodes make a lesser adjustment based on their distance from the winning node. In addition, the neighborhood kernel decreases with increasing $t$ so that the magnitude of reference vector adjustment decreases during the map training. We continue this process of presenting data vectors and modifying reference vectors for a specific number of iterations.

Neighborhood kernels commonly take one of two forms. The simplest, nicknamed the "bubble" kernel, defines a neighborhood set of nodes $N_c(t)$ around node $c$ so that $h_{ci} = \alpha(t)$ if $i \in N_c$ and $h_{ci} = 0$ if $i \notin N_c$, where $\alpha(t)$ is termed the *learning rate* $(0 < \alpha(t) < 1)$. Therefore, the bubble kernel activates and applies equal influence to all nodes within a specific radius of the winning node, while all nodes outside of that radius make no reference vector modification. The neighborhood $N_c(t)$ decreases monotonically in time. The second commonly used neighborhood kernel takes a Gaussian form

$$h_{ci} = \alpha(t)e^{-\|r_c - r_i\|^2/2\sigma^2(t)}, \qquad \text{(A3)}$$

where $r_c \in \Re^2$ and $r_i \in \Re^2$ are the radius vectors of nodes $c$ and $i$, respectively, and $\sigma(t)$ defines the width of the kernel. Though the precise forms of $\alpha(t)$ and $\sigma(t)$ are not critical, both parameters decrease monotonically in time, just as $\alpha(t)$ and $N_c(t)$ decrease with the bubble kernel.

Initially, we use a relatively large learning rate $\alpha$ and neighborhood radius to achieve a broad ordering. Figure A1c illustrates the reference vectors after 500 iterations of training, where we have used a bubble neighborhood kernel, an initial learning rate $\alpha(1)$ of 0.5, and an initial neighborhood radius equal to the diameter of
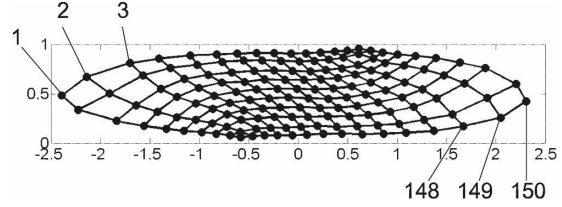


FIG. A2. Fine-tuned reference vectors of a $10 \times 15$ SOM: The analyzed dataset consists of 10 000 random two-dimensional data vectors with components normally distributed in $x$ and uniformly distributed in $y$.

the map, that is, 10 SOM units, where each SOM unit is the relative distance between two neighboring nodes. Although we have not yet trained the map long enough to achieve a fine-tuned ordering, we see that the nodes have achieved a broadly ordered state in that the reference vectors are topologically equivalent to that of the SOM in Fig. A1a. The distribution of reference vectors, however, fails to mimic the uniform distribution of the input data at this point. Nevertheless, the topological equivalence between Figs. A1a and A1c signifies that neighboring nodes possess similar reference vectors and distant nodes possess distinct reference vectors.

To achieve a more fine-tuned ordering, we train the map again but with a smaller initial learning rate and neighborhood radius. Fig. A1d illustrates the reference vectors after an additional 50 000 iterations of training but with an initial learning rate $\alpha(1)$ of 0.04 and an initial neighborhood radius of 3 SOM units. The SOM has acquired a more fine-tuned ordering, as the reference vectors more closely mirror the uniform distribution of the input data. The biggest exception to this uniform spacing lies along the edges and especially at the corners. This "border effect" results from the use of the neighborhood function to obtain an ordered mapping. The nodes at the edges and corners are influenced by adjacent winning nodes located toward the interior of the map, so the edge and corner reference vectors are adjusted toward those data vectors that are found in the interior of the data space. Because there are no nodes located to the exterior of these edges and corners, the border nodes receive no compensating "nudging" toward the exterior of the data space. Although section 2 mentions that the SOM patterns approximate cluster centroids, a more precise description might be that the SOM patterns approximate centroids over a broader influence region where clusters overlap (see Kohonen 2001 for a more thorough description of this effect).

Figure A2 presents fine-tuned reference vectors for a similar case to that of Fig. A1 except that the $x$ dimen-

sion of the input data now varies normally with a mean of 0 and standard deviation of 1. The $y$ dimension of the data in this second case still varies uniformly between 0 and 1. Here we have trained a $10 \times 15$ SOM. We see in the fine-tuned ordering that the reference vectors have become approximately uniformly spaced in the $y$ direction but have become more tightly clustered near 0 in the $x$ direction. As in the first case, the nodes have learned the distribution of the input data. In examining the 150 reference vectors in Figure A2, we find that the mean and standard deviation of the $x$ values are 0.02 and 0.96, respectively, which closely mirror the corresponding values of the input data. In both of these examples, we see that the SOM analysis has reduced a large dataset to a much smaller number of two-dimensional patterns that accurately represent the distribution of data. In addition, these patterns are presented in a topologically ordered manner.

In most practical applications we analyze data that possess much greater than two dimensions. The basic concepts described above apply in these cases as well. Of course, we lose the ability to plot the reference vectors as we have in Figs. A1 and A2 for cases of higher data dimensionality. In the case that each dimension of the data vector represents a geographical location, we may plot the reference vectors as a map,[3] as is customary in EOF analysis. If we wish to visualize the distance relationships between nodes, as we do in Figs. A1 and A2, then we may choose to employ a *Sammon mapping* scheme, as described in section 3.

In practice we may train a map several times with different process parameters to assess whether the map has approached a converged state. One way that we can evaluate the performance of a SOM is to calculate the *average quantization error* [see (5) in section 2]. When deciding which parameters to use, we may choose to keep the ones that yield the smallest average quantization error. Additional consideration may be given to the degree of topological ordering, as depicted in a Sammon map. For the scope of the present work, we have found rather small sensitivity of SOM performance to the choice of initial process parameters.

Software and additional guidelines for implementation of the SOM are freely available online (http://www.cis.hut.fi/research/som-research/). More comprehensive theoretical treatment of the SOM may be found in Kohonen (2001). In addition, we recommend Hewitson

and Crane (2002) for additional analysis methods in the application of SOMs to synoptic climatology.

## REFERENCES

Barnston, A. G., and R. E. Livezey, 1987: Classification, seasonality, and persistence of low-frequency atmospheric circulation patterns. *Mon. Wea. Rev.,* **115,** 1083–1126.

Cassou, C., L. Terray, J. Hurrell, and C. Deser, 2004: North Atlantic winter climate regimes: Spatial asymmetry, stationarity with time, and oceanic forcing. *J. Climate,* **17,** 1055–1068.

Feldstein, S. B., 2000: The timescale, power spectra, and climate noise properties of teleconnection patterns. *J. Climate,* **13,** 4430–4440.

——, and C. Franzke, 2006: Are the North Atlantic Oscillation and the Northern Annular Mode indistinguishable? *J. Atmos. Sci.,* **63,** 2915–2930.

Franzke, C., and S. B. Feldstein, 2005: The continuum and dynamics of Northern Hemisphere teleconnection patterns. *J. Atmos. Sci.,* **62,** 3250–3267.

Hewitson, B. C., and R. G. Crane, 2002: Self-organizing maps: Applications to synoptic climatology. *Climate Res.,* **22,** 13–26.

Hilmer, M., and T. Jung, 2000: Evidence for a recent change in the link between the North Atlantic Oscillation and Arctic sea ice export. *Geophys. Res. Lett.,* **27,** 989–992.

Horel, J. D., 1985: Persistence of the 500 mb height field during Northern Hemisphere winter. *Mon. Wea. Rev.,* **113,** 2030–2042.

Hurrell, J. W., 1995: Decadal trends in the North Atlantic Oscillation and relationships to regional temperature and precipitation. *Science,* **269,** 676–679.

Jung, T., M. Hilmer, E. Ruprecht, and S. Kleppek, 2003: Characteristics of the recent eastward shift of interannual NAO variability. *J. Climate,* **16,** 3371–3382.

Kalnay, E., and Coauthors, 1996: The NCEP/NCAR 40-Year Reanalysis Project. *Bull. Amer. Meteor. Soc.,* **77,** 437–471.

Kimoto, M., and M. Ghil, 1993: Multiple flow regimes in the Northern Hemisphere winter. Part I: Methodology and hemispheric regimes. *J. Atmos. Sci.,* **50,** 2625–2643.

Kohonen, T., 2001: *Self-Organizing Maps.* Springer, 501 pp.

Kushnir, Y., and J. M. Wallace, 1989: Low-frequency variability in the Northern Hemisphere winter: Geographical distribution, structure, and time-scale dependence. *J. Atmos. Sci.,* **46,** 3122–3142.

Livezey, R. E., and W. Y. Chen, 1983: Statistical field significance and its determination by Monte Carlo techniques. *Mon. Wea. Rev.,* **111,** 46–59.

Luo, D., and T. Gong, 2006: A possible mechanism for the eastward shift of interannual NAO action centers in last three decades. *Geophys. Res. Lett.,* **33,** L24815, doi:10.1029/2006GL027860.

Overland, J. E., and W. Wang, 2005: The Arctic climate paradox: The recent decrease of the Arctic Oscillation. *Geophys. Res. Lett.,* **32,** L06071, doi:10.1029/2004GL021752.

Peterson, K. A., J. Lu, and R. J. Greatbatch, 2003: Evidence of nonlinear dynamics in the eastward shift of the NAO. *Geophys. Res. Lett.,* **30,** 1030, doi:10.1029/2002GL015585.

Polyakova, E. I., A. G. Journel, I. V. Polyakov, and U. S. Bhatt, 2006: Changing relationship between the North Atlantic Oscillation and key North Atlantic climate parameters. *Geophys. Res. Lett.,* **33,** L03711, doi:10.1029/2005GL024573.

Reusch, D. B., R. B. Alley, and B. C. Hewitson, 2007: North

---

[3] If the SOM has been trained with area-weighted data, then the reference vectors incorporate that same weighting and should be unweighted before plotting.

Atlantic climate variability from a self-organizing map perspective. *J. Geophys. Res.,* **112,** DO2104, doi:10.1029/2006JD007460.

Sammon, J. W., Jr., 1969: A nonlinear mapping for data structure analysis. *IEEE Trans. Comput.,* **C-18,** 401–409.

Thompson, D., and J. M. Wallace, 1998: The Arctic Oscillation signature in wintertime geopotential heights and temperature fields. *Geophys. Res. Lett.,* **25,** 1297–1300.

Ulbrich, U., and M. Christoph, 1999: A shift of the NAO and increasing storm track activity over Europe due to anthropogenic greenhouse gas forcing. *Climate Dyn.,* **15,** 551–559.

van den Dool, H. M., 2007: *Empirical Methods in Short-Term Climate Prediction*. Oxford University Press, 215 pp.

Wallace, J. M., and D. S. Gutzler, 1981: Teleconnections in the geopotential height field during the Northern Hemisphere winter. *Mon. Wea. Rev.,* **109,** 784–804.

——, Y. Zhang, and L. Bajuk, 1996: Interpretation of interdecadal trends in Northern Hemisphere surface air temperature. *J. Climate,* **9,** 249–259.

Woollings, T., and B. Hoskins, 2008: Simultaneous Atlantic-Pacific blocking and the Northern Annular Mode. *Quart. J. Roy. Meteor. Soc.,* **134,** 1635–1646.