Exploratory Analysis of Climate Data Using Source Separation Methods

Alexander Ilin [a,*], Harri Valpola [b], Erkki Oja [a]

[a]*Laboratory of Computer and Information Science*
Helsinki University of Technology
P.O. Box 5400, FI-02015 TKK, Espoo, Finland

[b]*Laboratory of Computational Engineering*
Helsinki University of Technology
P.O. Box 9203, FI-02015 TKK, Espoo, Finland

Running title: Separation of climate signals

Abstract

We present an example of exploratory data analysis of climate measurements using a recently developed denoising source separation (DSS) framework. We analyzed a combined dataset containing daily measurements of three variables: surface temperature, sea level pressure and precipitation around the globe, for a period of 56 years. Components exhibiting slow temporal behavior were extracted using DSS with linear denoising. The first component, most prominent in the interannual time scale, captured the well-known El Niño–Southern Oscillation (ENSO) phenomenon and the second component was close to the derivative of the first one. The slow components extracted in a wider frequency range were further rotated using a frequency-based separation criterion implemented by DSS with nonlinear denoising. The rotated sources give a meaningful representation of the slow climate variability as a combination of trends, interannual oscillations, the annual cycle and slowly changing seasonal variations. Again, components related to the ENSO phenomenon emerge very clearly among the found sources.

Key words: independent component analysis, blind source separation, semiblind methods, denoising source separation, El Niño, global climate

* Corresponding author, phone +358-9-451-5396, fax +358-9-451-3277
  Email addresses: `alexander.ilin@tkk.fi` (Alexander Ilin),
`harri.valpola@tkk.fi` (Harri Valpola), `erkki.oja@tkk.fi` (Erkki Oja).

1 Introduction

One of the main goals of statistical analysis of climate data is to extract physically meaningful patterns of climate variability from highly multivariate weather measurements. The classical technique for defining such dominant patterns is principal component analysis (PCA), or empirical orthogonal functions (EOF) as it is called in climatology (see, e.g., von Storch and Zwiers, 1999). However, many researchers have pointed out that the maximum remaining variance criterion used in PCA can lead to such problems as mixing different physical phenomena in one extracted component (Richman, 1986; Kim and Wu, 1999). This makes PCA a useful tool for information compression but limits its ability to isolate individual modes of climate variation.

To overcome this problem, rotation of the principal components has proven useful. The different rotation criteria reviewed by Richman (1986) are based on the general "simple structure" idea aimed at, for example, spatial or temporal localization of the rotated components. The rotation of EOFs can be either orthogonal or oblique, which potentially leads to better interpretability of the extracted components.

Independent component analysis (ICA) is a recently developed statistical technique for component extraction which can also be used for rotating principal components. The basic assumption made in ICA is the statistical independence of the extracted components, which may lead to a meaningful data representation in a number of applications (see, e.g., Hyvärinen et al., 2001, for introduction). ICA is based on higher-order statistics and in this respect bears some similarity to classical rotation techniques such as the Varimax orthogonal rotation (Richman, 1986). Several attempts to apply ICA in climate research have already been made (Aires et al., 2000; Lotsch et al., 2003).

In this paper, we analyze weather measurements using a novel extension of ICA called denoising source separation (DSS) (Särelä and Valpola, 2005). DSS is a general separation framework which does not necessarily exploit the independence assumption but rather looks for hidden components which have "interesting" properties. The interestingness of the properties is controlled by means of a temporal filtering or denoising procedure.

In the first experiment, we show that the sources with the most prominent interannual oscillations can be identified using DSS with linear filtering as denoising. The leading components are clearly related to the well-known El Niño–Southern Oscillation (ENSO) phenomenon.

In the second experiment, we use DSS with linear denoising as the first, preprocessing step of climate data analysis. A wider frequency band in the denoising filter is used to identify the slow subspace of the climate system. The found slow components are further rotated using an iterative DSS procedure based on nonlinear

denoising. The rotation aims to find components with distinct power spectra.

The extracted components turn out to represent the subspace of the slow climate phenomena as a linear combination of trends, decadal-interannual oscillations, the annual cycle and other phenomena with distinct spectral contents. Using this approach, the known climate phenomena are identified as certain subspaces of the climate system and some other interesting phenomena hidden in the weather measurements are found.

The contents of this paper is as follows: in the next section, we explain the modeling assumptions of the source separation methods and present a short introduction to DSS. Section 3 describes the climate measurements used in the experiments. Section 4 explains how DSS is tuned to extract components with the most prominent interannual oscillations and present the experimental results (which were partly reported by Ilin et al., 2005). In Section 5, we give the description of the frequency-based separation algorithm implemented in the DSS framework and report the climate phenomena found with this method. Preliminary results of this analysis were published in a conference paper by Ilin and Valpola (2005). Finally, we discuss the results and possible future directions in Section 6.

2 Source Separation Methods

2.1 Blind source separation and independent component analysis

The basic modeling assumption of linear source separation methods is that there are some hidden component signals or time series $s_i(t)$ (also called sources, factors or latent variables) which are linearly combined into the multivariate measurements $x_j(t)$:

$$x_j(t) = \sum_{i=1}^{N} a_{ji} s_i(t), \ j = 1, ..., N .$$ (1)

The index j runs over the measurement sensors (typically spatial locations), and discretized time t runs over the observation period: $t = 1, ..., T$. This can be expressed in matrix formulation by denoting the matrix of observations by $\boldsymbol{X}$, where the sensor index j denotes the rows and the time index t denotes the columns. The matrix of sources, $\boldsymbol{S}$ is defined likewise, and the coefficients a_{ji} of the linear combinations make up a matrix $\mathbf{A}$. Using these matrices, (1) becomes

$$\boldsymbol{X} = \mathbf{A}\boldsymbol{S} .$$ (2)

If we denote the columns of matrix $\mathbf{A}$ by $\mathbf{a}_i$ and the columns of matrix $\boldsymbol{X}$ by $\mathbf{x}(t)$, then (2) can be further written as

$$\mathbf{x}(t) = \sum_{i=1}^{N} \mathbf{a}_i s_i(t) \,. \tag{3}$$

The mapping $\mathbf{A}$ is called the mixing matrix in the ICA terminology or the loading matrix in the context of PCA. In climate data analysis, the time series $s_i(t)$ usually correspond to the time-varying states of the climate system, and the loading vectors $\mathbf{a}_i$ are the spatial maps showing the typical weather patterns corresponding to the components.

The goal of the analysis is to estimate the unknown components $s_i(t)$ and the corresponding loading vectors $\mathbf{a}_i$ from the observed data $\boldsymbol{X}$. With minimum a priori assumptions about the sources, the problem is called blind source separation (BSS).

Independent component analysis (ICA) is a popular method of solving the BSS problem. In ICA, the only assumption is the statistical independence of the sources: each $s_i(t)$ is regarded as a sample from a random variable s_i, and these variables are mutually independent. There are a large variety of algorithms for solving the mixing matrix $\mathbf{A}$ and the sources (Hyvärinen et al., 2001; Cichocki and Amari, 2002). One of the most popular methods is the FastICA algorithm. The independence of the sources is measured by their mutual information which results in a minimum entropy criterion. In practice the separation is achieved by rotating the observations into directions that are as non-Gaussian as possible (Hyvärinen et al., 2001).

2.2 Denoising source separation

ICA is a powerful tool for exploratory data analysis, when very little is known about the underlying source processes $s_i(t)$. Independence is the only assumption. Sometimes however, such prior information exists, such as the general shape of the time curves or their frequency contents. For example, in the climate data we might be interested in some phenomena that would be cyclic over a certain period, or exhibit slow changes. It would be very useful if such prior knowledge could be incorporated into the separation algorithm directly. Exploiting prior knowledge about the sources may significantly help in finding a good representation of the data, and fully blind algorithms are not the best choice.

This kind of problem setting, with some prior knowledge available, is called semi-blind. One of the methods for solving it is a recently introduced method called denoising source separation (Särelä and Valpola, 2005).

DSS is a general algorithmic framework which can identify the model in Eq. (1)

exploiting prior knowledge about its unknowns. In DSS, the independence criterion of ICA is replaced by the assumption that the sources should 1) be uncorrelated and 2) maximize some desired properties (e.g., non-Gaussianity, slowness etc). In this respect, DSS can be seen as an extension of ICA without the strict independence assumption.

The first requirement is assured in DSS by using a preprocessing step called whitening or sphering. The goal of whitening is to make the covariance structure of the data uniform in such a way that any linear projection of the data has unit variance. The positive effect of such a transformation is that any orthogonal basis in the whitened space defines uncorrelated sources. Therefore, whitening is used as a preprocessing step in many ICA algorithms, which allows restricting the mixing matrix to be orthogonal afterwards.

Whitening is usually implemented by PCA. Assuming that the measurements $x_j(t)$ have been normalized to zero mean, the matrix of sphered data $\boldsymbol{Y}$ is calculated as

$$\boldsymbol{Y} = \mathbf{D}^{-1/2}\mathbf{V}^T\boldsymbol{X}\,, \tag{4}$$

where $\mathbf{D}$ is the diagonal matrix of eigenvalues of the data covariance matrix, defined as $\frac{1}{T}\boldsymbol{X}\boldsymbol{X}^T$. The columns of matrix $\mathbf{V}$ are the corresponding eigenvectors. The dimensionality of the data can also be reduced at this stage by retaining only the principal components corresponding to the largest eigenvalues in $\mathbf{D}$. It is easy to show that it now holds that $\frac{1}{T}\boldsymbol{Y}\boldsymbol{Y}^T = \mathbf{I}$. Matrix $\boldsymbol{Y}$ is not unique, though; any orthogonal rotation of its columns produces a matrix that also has unit covariance.

The rotational ambiguity of the whitened data matrix is fixed using the second DSS requirement, which implements the source separation criterion. This requirement is usually introduced in the algorithm in the form of a *denoising function*. The purpose of denoising is to emphasize the desired properties in the current source estimates, which assures gradual maximization of these properties.

2.2.1 Linear denoising

In the simplest case, the denoising function can be implemented by a linear temporal filter, operating on the rows of matrix $\boldsymbol{Y}$ and giving another matrix $\mathbf{f}(\boldsymbol{Y}) = \boldsymbol{Y}\mathbf{F}$ with $\mathbf{F}$ the filtering matrix. Denoising renders the variances of the sphered components different: the covariance matrix of $\mathbf{f}(\boldsymbol{Y})$ equals $\frac{1}{T}\boldsymbol{Y}\mathbf{F}\mathbf{F}^T\boldsymbol{Y}^T$ which is no more equal to the unit matrix. Now PCA can identify the directions which maximize the properties of interest. The eigenvalues obtained from PCA give the ratio of the variances of the sources after and before filtering which is the objective function of DSS with linear denoising. The components are ranked according to the prominence of the desired properties the same way as the principal components in PCA are ranked according to the amount of variance they explain. The separation thus consists of three steps: whitening, linear denoising (filtering) and PCA on the

Fig. 1. The steps of the DSS algorithm in case of linear denoising.

denoised data, as shown in Fig. 1.

The DSS algorithms implemented by linear denoising optimize the same type of cost function as the maximum noise fraction (MNF) transform proposed by Green et al. (1988). However, in DSS framework the computations are structured such that it is easier to generalize the method for nonlinear denoising.

2.2.2 Nonlinear denoising

More complex separation criteria usually require nonlinear denoising (see Särelä and Valpola, 2005; Valpola and Särelä, 2004, for several examples). Then, DSS requires an algorithm presented in Fig. 2. Here, whitening is followed by an iterative procedure with three successive steps:

(1) source estimation using the current estimate of the demixing matrix $\boldsymbol{W}$:

$$\boldsymbol{S} = \boldsymbol{W}\boldsymbol{Y}\,;$$

(2) applying the denoising function to the source estimates:

$$\widehat{\boldsymbol{S}} = \mathbf{f}(\boldsymbol{S})\,;$$

(3) reestimation of the demixing matrix

$$\boldsymbol{W}^{T} = \mathrm{orth}(\boldsymbol{Y}\widehat{\boldsymbol{S}}^{T})\,.$$

The iterations continue until the source estimates do not change. In Step 3, $\mathrm{orth}(.)$ is an operator giving the orthogonal projection of the matrix $\boldsymbol{Y}\widehat{\boldsymbol{S}}^{T}$ onto the set of orthogonal matrices.

Without denoising, this procedure is equivalent to the power method for computing the principal components of $\boldsymbol{Y}$, because then Steps 1 and 3 give $\boldsymbol{W}^{T} = \mathrm{orth}(\boldsymbol{Y}\boldsymbol{Y}^{T}\boldsymbol{W}^{T})$. Since $\boldsymbol{Y}$ is white, all the eigenvalues are equal and the solution without denoising becomes degenerate. Therefore, even slightest changes made by denoising can determine the DSS rotation. Since the denoising procedure emphasizes the desired properties of the sources, DSS can find the rotation where the properties of interest are maximized.

In case of nonlinear denoising, the DSS objective function is usually expressed implicitly in the denoising function. Therefore, ranking the components according
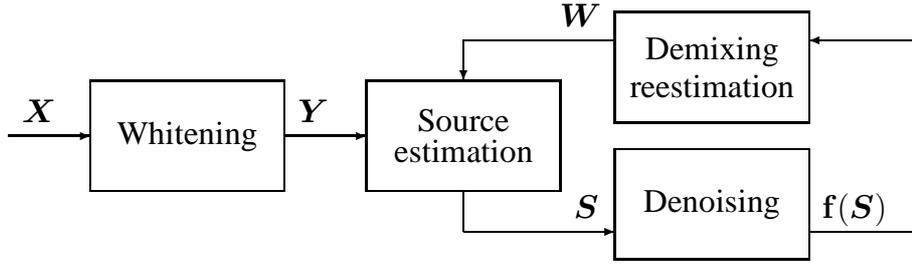
6

Fig. 2. The steps of the DSS algorithm in the general case of nonlinear denoising.

to the prominence of the desired properties is more difficult and depends on the exact separation criterion used in the denoising procedure.

In the applications, we are interested not only in the sources (rows of matrix $\boldsymbol{S}$), but also in the matrix $\mathbf{A}$ in Eq. (2). The i-th column of $\mathbf{A}$ is a spatial map showing how the effect of the i-th source is distributed over the sensor array. Noting that in DSS it holds $\boldsymbol{S} = \boldsymbol{W}\boldsymbol{Y}$, we obtain from Eqs. (2) and (4)

$$\boldsymbol{X} = \mathbf{A}\boldsymbol{S} = \mathbf{A}\boldsymbol{W}\boldsymbol{Y} = \mathbf{A}\boldsymbol{W}\mathbf{D}^{-1/2}\mathbf{V}^T\boldsymbol{X}\,. \tag{5}$$

Thus $\mathbf{A}$ should be chosen as the (pseudo)inverse of $\boldsymbol{W}\mathbf{D}^{-1/2}\mathbf{V}^T$ which is

$$\mathbf{A} = \mathbf{V}\mathbf{D}^{1/2}\boldsymbol{W}^T\,. \tag{6}$$

Since the extracted components $s_i(t)$ are normalized to unit variances, the columns of $\mathbf{A}$ have a meaningful scale.

Note that the signs of the extracted components cannot generally be determined by DSS, which is a well-known property of the classical ICA problem. DSS with nonlinear denoising can determine the sign if the signal is not symmetric. More details about the DSS method, including rigorous derivations and analysis were reported by Särelä and Valpola (2005).

2.3   Extracting sources of climate variability

2.3.1   ICA in climate data analysis

Climate is a very complex system where different phenomena constantly interact with each other. For example, the annual cycle naturally affects other climate processes, the El Niño effect has great impact on global weather but also tends to phase-lock with the annual cycle, and so on. Therefore, the existence of any truly independent climate phenomena is an implausible assumption.

ICA can still be a useful tool for climate data representation providing, for example, temporally or spatially localized components. Several researchers have shown

that ICA can extract meaningful components from climate and weather data (Aires et al., 2000; Lotsch et al., 2003; Basak et al., 2004). However, due to a great amount of noise in climate data, naive ICA can often produce overfitted solutions (see Hyvärinen et al., 1999; Särelä and Vigário, 2003, for discussion of this problem) and one would require a very long observation period in order to find a meaningful ICA solution.

2.3.2 Tuning DSS for climate data analysis

DSS is a much more flexible tool as one can choose the separation criterion that gives the most meaningful or interpretable representation of the data. In this work, we search for physically meaningful states of the climate system which would possess *slow behavior*. The slowness of the components assures their larger long-term effect on global weather and possibly facilitates making predictions of their future development. By physically meaningful states we mean such components whose dynamics are as weakly coupled as possible. Climate phenomena with distinct time scales of their variability (e.g., the annual cycle, El Niño–Southern Oscillation or slow climate trends) would intuitively be such components.

First, we concentrate on finding components which exhibit prominent variability in the slow timescale. In Section 4, we show how DSS can be tuned to extract such components based on a criterion that we term clarity. This approach can be useful for identifying the subspace of slow phenomena and sometimes it can also find a meaningful representation within the found subspace. However, it does not generally provide a good separation criterion, which may lead to mixtures of different climate phenomena still existing in any one component. Therefore, we propose a more complex DSS-based algorithm which tries to separate slow climate components based on their frequency contents. The exposition of this algorithm is done in Section 5.1.

3 Data and Preprocessing Method

We apply the proposed DSS tools to measurements of three major atmospheric variables: surface temperature, sea level pressure and precipitation. This set of variables is often used for describing global climate phenomena such as ENSO (Trenberth and Caron, 2000). The datasets are provided by the reanalysis project of the National Centers for Environmental Prediction–National Center for Atmospheric Research (NCEP/NCAR) (Kalnay et al., 1996; NCEP data, 2004).

The data represent globally gridded measurements over a long period of time. The spatial grid is regularly spaced over the globe with $2.5° \times 2.5°$ resolution. Although the quality of the data is worse for the beginning of the reanalysis period and it

8

considerably varies throughout the globe, we used the whole period of 1948–2004. Thus, the data is very high-dimensional: more than 10,000 spatial locations by more than 20,000 time instances for each of the three data sets.

The main drawback of the reanalysis data is that it is not fully real. The measurements missing in some spatial locations or time instances have been reestimated based on the available data and approximation models. Yet, the data is as close to the real measurements as possible and its regularity makes this data particularly suitable for the source separation methods applied in this work.

To preprocess the data, the long-term mean was removed and the data points were weighted to diminish the effect of a denser sampling grid around the poles: each data point was multiplied by a weight proportional to the square root of the corresponding area of its location. This produced the original data matrix X. The spatial dimensionality of the data was then reduced using the PCA/EOF analysis applied to the weighted data. For each dataset, we retained 100 principal components. This means that in Eq. (4), the columns of Y have dimension 100, while those of the original X are over 10,000 dimensional. Yet the principal components explain more than 90% of the total variance, which is due to the high spatial correlation between nearby points on the global grid. The DSS-based analysis was then applied to the combined data containing the measurements of the three variables.

4 ENSO as Component with Most Prominent Interannual Oscillations

Components with prominent slow behavior can be extracted from the data using DSS with low-pass or band-pass filtering as denoising. This type of denoising is linear and therefore the simple algorithm described in Section 2.2.1 is applicable here. The extracted slow components are ranked according to their clarity which is defined as the ratio of the variance of the component filtered in the desired frequency range and the variance of the non-filtered component:

$$\text{clarity} = \frac{\text{var}\{s_f\}}{\text{var}\{s\}} \, .$$

That is, the first component contains the least relative amount of other frequencies in its power spectrum. The presented algorithm is functionally identical to MNF (Green et al., 1988) with the noise defined as the spectral components lying outside the desired frequency range.

In this section, we aim at finding components which exhibit prominent variability in the interannual timescale. Therefore, the band-pass filter whose frequency response is shown in Fig. 3 is used here.

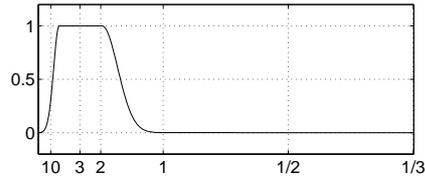The four most prominent interannual components extracted from the combined data

9

Fig. 3. The frequency response of the filter used in DSS with linear denoising for finding components with the most prominent interannual oscillations. The abscissa is linear in frequency but is labeled in terms of periods, in years.
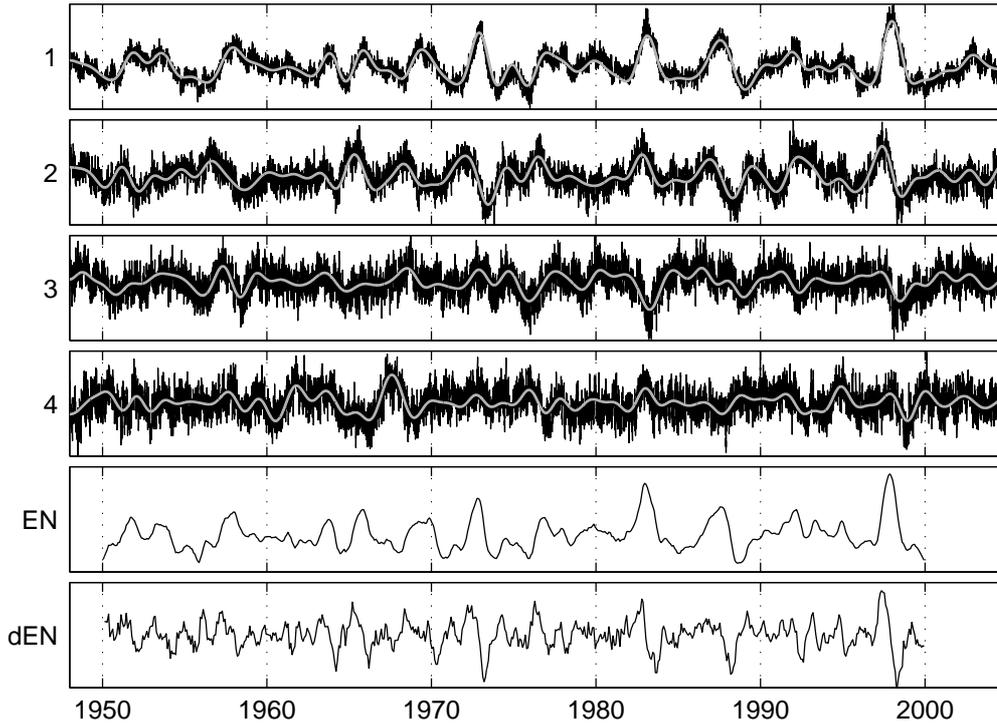

Fig. 4. Four components with the most prominent interannual oscillations extracted from the combined data. Top four: The time course of the leading components found by DSS (black) and their filtered versions (gray). The non-filtered components are normalized to unit variance. Bottom two: The Niño 3 SST index (Niño 3 SST, 2004) and its derivative bear similarities to the first and second components respectively.

are shown in Fig. 4. The time course of the first component (upper curve in Fig. 4) shows striking resemblance with the El Niño index calculated from the sea surface temperature (SST) in the Niño 3 region (lower curve). The correlation coefficients between the extracted component and the Niño 3 SST index is 0.9323. Note that the upper components are extracted from climate data consisting of daily measurements from the whole globe, with the only constraint being the emphasis on strong interannual oscillations. Also note that the values of the Niño 3 SST index are monthly averages and consequently appear smoother than the daily averages of DSS components.
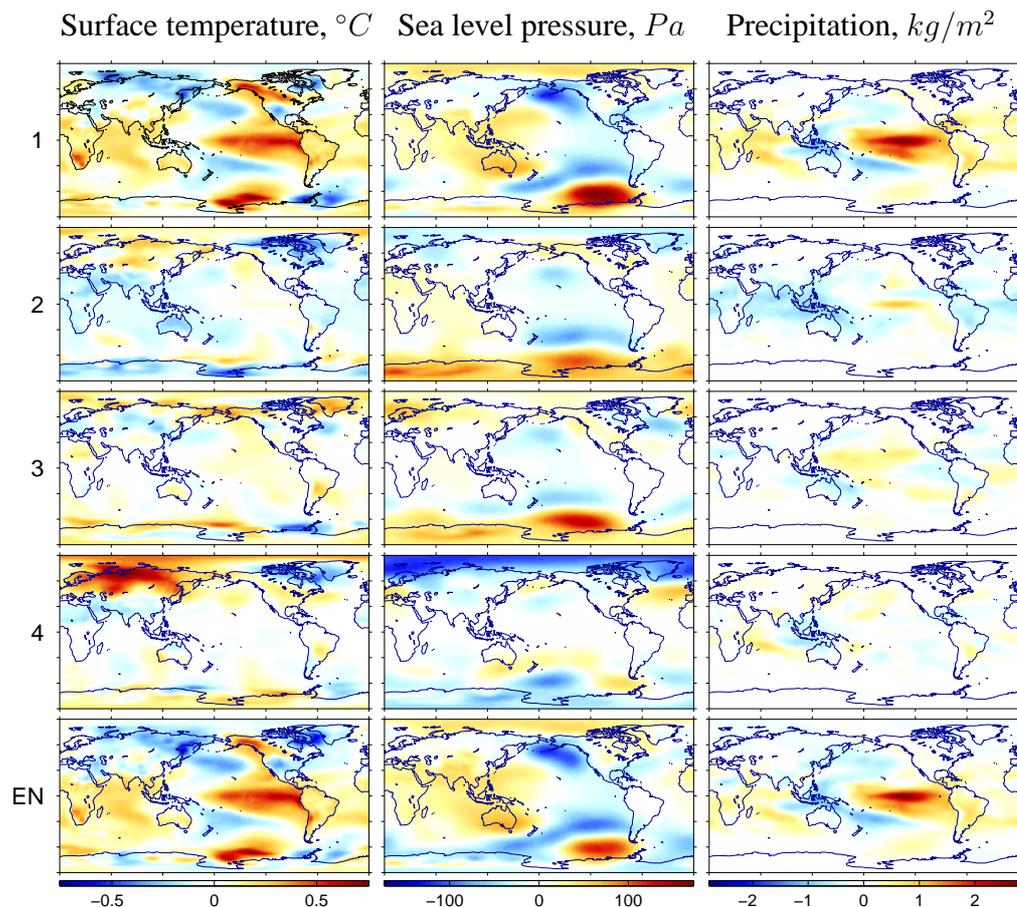
10

Fig. 5. Top four: The spatial patterns of the four leading interannual components extracted from the combined data. Bottom: The regression coefficients calculated from the combined data using the Niño 3 SST index. The maps are weighted by the square root of the clarity values of the components.

The spatial patterns corresponding to the four leading components are shown[1] in Fig. 5. The first surface temperature map contains many features traditionally associated with El Niño (Trenberth and Caron, 2000): the strongest pattern in central and eastern tropical Pacific with broader regions along the eastern Pacific coast, a negatively correlated "boomerang"-shaped region at $20°$–$40°$ latitude in both hemispheres linked in the far western equatorial Pacific, positive values in the Indian Ocean, and negative values in the North Pacific and around New Zealand.

The corresponding sea level pressure map is similar to the classical Southern Oscillation pattern (Trenberth and Caron, 2000): there is a major seesaw structure in the tropics and subtropics, large pressure departures in the North Pacific, and a quadrupole-like structure in the Australasia–South Pacific region. The precipitation map also contains many features associated with the ENSO phenomenon: the dominant effects are clearly seen throughout the tropical Pacific with maximum values

---

[1] The maps are plotted using the mapping toolbox developed by Pawlowicz (2000).

11

in the Niño 3 region. The clear patterns here are the intertropical convergence zone (ITCZ) and South Pacific convergence zone (SPZC), a "boomerang"-shaped negatively correlated area in mid-latitudal Pacific merged over Indonesia, positive values in the Indian Ocean and subtropical and tropical Atlantic.

Similar ENSO features are observed from the regression patterns (shown at the bottom of Fig. 5) calculated using the Niño 3 SST index. Note that there are some differences in the extracted maps compared to the regression patterns, for example, a stronger teleconnection pattern in southern Africa in surface temperature and a stronger positive center in the South Pacific in sea level pressure.

The second extracted component also appears to be related to ENSO and roughly corresponds to the time derivative of the first component (see Fig. 4). The corresponding precipitation pattern has an interesting localization in the Niño 3 region and mostly negative loadings in the rest of the tropical and subtropical areas. The third and the fourth components show weaker oscillations in the interannual time scale. Similar components will be discussed in Section 5.1. Table 1 lists the clarity values of the found components. Since this index is used as the objective function for extracting the components, the first component always has the largest value in all conditions.

We also applied the same analysis to the three data sets separately (see Ilin et al., 2005) and the first extracted component was always a good ENSO index. Somewhat surprisingly even the component extracted from sea-level-pressure data resembled more the Niño 3 SST index than Southern Oscillation Index (SOI) although SOI is defined in terms of sea level pressure.

Table 1
Clarity values of DSS components and regression components obtained from ENSO indices Niño 3 SST (2004) and SOI (2005).

Comp. 1	0.7484
Comp. 2	0.5105
Comp. 3	0.3376
Comp. 4	0.3014
Niño 3 SST	0.6853
SOI	0.5354

5 Extracting Slow Components with Distinct Power Spectra

5.1 *Frequency-based separation of sources*

The algorithm described in the previous section is useful for extracting components which are dominant in a certain frequency range. This requires some knowledge

about the expected power spectrum of the extracted component in order to use a proper frequency mask in the denoising filter. In blinder settings, however, this information does not exist and the frequency masks such as the one presented in Fig. 3 should be estimated automatically.

In this section, we present an algorithm which can be seen as an extension of the previous approach. It assumes that different extracted components are dominant in different frequencies (hence, they have distinct power spectra) and automatically estimates an individual frequency mask for each component. The adaptation of the masks is practically implemented using a competition mechanism involving the smoothed power spectra of the current sources estimates: A mask is increased in the frequencies where the corresponding component prevails over the other sources and it is decreased where the corresponding components is dominated by other sources. The sources are then reestimated using the adapted frequency masks.

The corresponding denoising procedure is nonlinear and therefore the iterative procedure described in Section 2.2.2 is used here. The exact steps of the denoising function $\widehat{S} = \mathbf{f}(S)$ are listed below:

(1) Compute discrete cosine transform S_{DCT} of the sources S and retain only the DCT coefficients important for the source separation. In this section, the separation is done in the low frequencies and therefore the mask from Fig. 6 is applied to the DCT coefficients.

(2) Calculate smoothed DCT power spectra $\mathcal{C}$ (by, e.g., low-pass filtering of the squares of S_{DCT}) and normalize it in such a way that the average (across sources) values of $\mathcal{C}$ are same for different frequency bins. This is done to increase the competition in relatively weak frequencies.

(3) Apply the competition mechanism by partially whitening $\mathcal{C}$ to a degree α with a symmetric whitening matrix:

$$\mathbf{M}_{\mathrm{DCT}} = \mathbf{V}_{\mathcal{C}} \mathbf{D}_{\mathcal{C}}^{-\alpha/2} \mathbf{V}_{\mathcal{C}}^{T} \mathcal{C} \tag{7}$$

where $\mathbf{V}_{\mathcal{C}}$ is the orthogonal eigenvector matrix of the covariance of $\mathcal{C}$, and $\mathbf{D}_{\mathcal{C}}$ is the diagonal matrix of the corresponding eigenvalues. Whitening is done without removing the mean of $\mathcal{C}$. This competition mechanism is somewhat similar to the whitening-based estimation of the source variances proposed by Valpola and Särelä (2004).

(4) Calculate denoised sources $\widehat{S}$ by applying inverse DCT to S_{DCT} masked with the positive parts of $\mathbf{M}_{\mathrm{DCT}}$.

At the beginning of learning the denoising function also orders the components according to the mean frequencies in their power spectra. Later, the topographic idea (similarly to Hyvärinen et al., 2001) is used to relax the competition in the power spectra of the neighboring sources. Note also that since DCT is a linear orthogonal transformation, steps 1 and 4 are only required at the beginning and at the end of learning respectively. Performing step 4 on each iteration is useful for
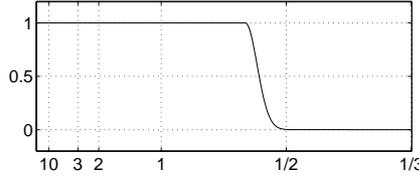
Fig. 6. The frequency response of the filter used in the frequency-based separation of slow climate components. The abscissa is linear in frequency but is labeled in terms of periods, in years.

tracking the time course of the source estimates during learning.

The presented algorithm essentially performs ICA in the frequency domain. Similar frequency-based separation criteria were used by Cichocki and Belouchrani (2001) and by Gharieb and Cichocki (2003) (see also Cichocki and Amari, 2002). For example, the algorithm presented by Cichocki and Belouchrani (2001) uses analogues of the frequency masks which are restricted to a band-pass type. In practice, solutions produced by the frequency-based ICA algorithms can be similar to results obtained with other ICA algorithms based on temporal structure such as SOBI (Belouchrani et al., 1997) or TDSEP (Ziehe and Müller, 1998).

5.2 Application on climate data

This section describes how we performed frequency-based separation of climate components in a wide, slow frequency range (see Fig. 6). First, we applied the DSS-based analysis described in Section 4 to extract components which exhibit the most prominent variability in the desired slow time scale. The digital filter from Fig. 6 was used in this step as linear denoising. This procedure is similar to the maximum autocorrelation factor transform proposed by Switzer (1985) and linear slow feature analysis (Wiskott and Sejnowski, 2002). Therefore, we refer to this step as slow feature analysis (SFA) in the following.

Then, the frequency-based separation algorithm was applied to several leading components extracted by SFA. We retained only sixteen cleanest components at this stage as the produced results were easily interpretable for this number of components. This procedure roughly identified three subspaces: trends, interannual and annual oscillations.

Based on the obtained results, we found it possible to improve the representation within the found subspaces. We used the clarity criterion introduced in Section 4 to order the subspace of trends and applied the frequency-based rotation for the subspace of interannual oscillations separately. The full sequence of the undertaken steps is schematically shown in Fig. 7.
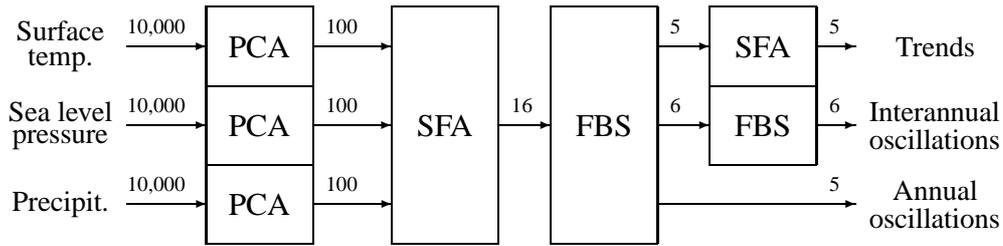
14

Fig. 7. The steps undertaken to find slow components with distinct power spectra. The numbers above arrows indicate the spatial dimensionality of the data. SFA stands for slow feature analysis described in Section 4 and FBS blocks correspond to frequency-based separation described in Section 5.1.

5.3 Experimental results

5.3.1 Identifying the subspace of slow climate phenomena

Several cleanest components were first extracted from the highly multidimensional data. The annual cycle appeared in the two leading components as the cleanest slow source of climate variability (for depictions of their time courses, see Ilin and Valpola, 2005). The following components also had interesting slow behavior but they appeared to be mixtures of several climate phenomena. For example, the prominent ENSO oscillations were mixed with trends. The power spectra of many components contained prominent slowest, decadal and close-to-annual frequencies. Except for the two annual cycle sources, none of the components had a clear dominant peak in its power spectrum.

5.3.2 Frequency-based separation of slow climate phenomena

The first sixteen slow components extracted at the first stage were further rotated using frequency-based DSS described in Section 5.1. To discard high-frequency noise, the monthly averages of the slow components were used. The time course of the rotated sources is presented in Fig. 8 (for depictions of spatial patterns corresponding to some of the components found after this stage, see Ilin and Valpola, 2005). The rotated components have a clearer interpretation compared to the original slow components. The power spectra of the rotated components are more distinct (see the middle column of Fig. 8). However, some of the power spectra look quite similar and we can roughly categorize the found sources into three subspaces with different variability time scales: trends (components 1–5), interannual oscillations (components 6–11) and components 12–16 with dominating close-to-annual frequencies in their spectra. The subspaces are identified reliably due to the distinct differences in the corresponding power spectra but the components within the subspaces may remain mixed.

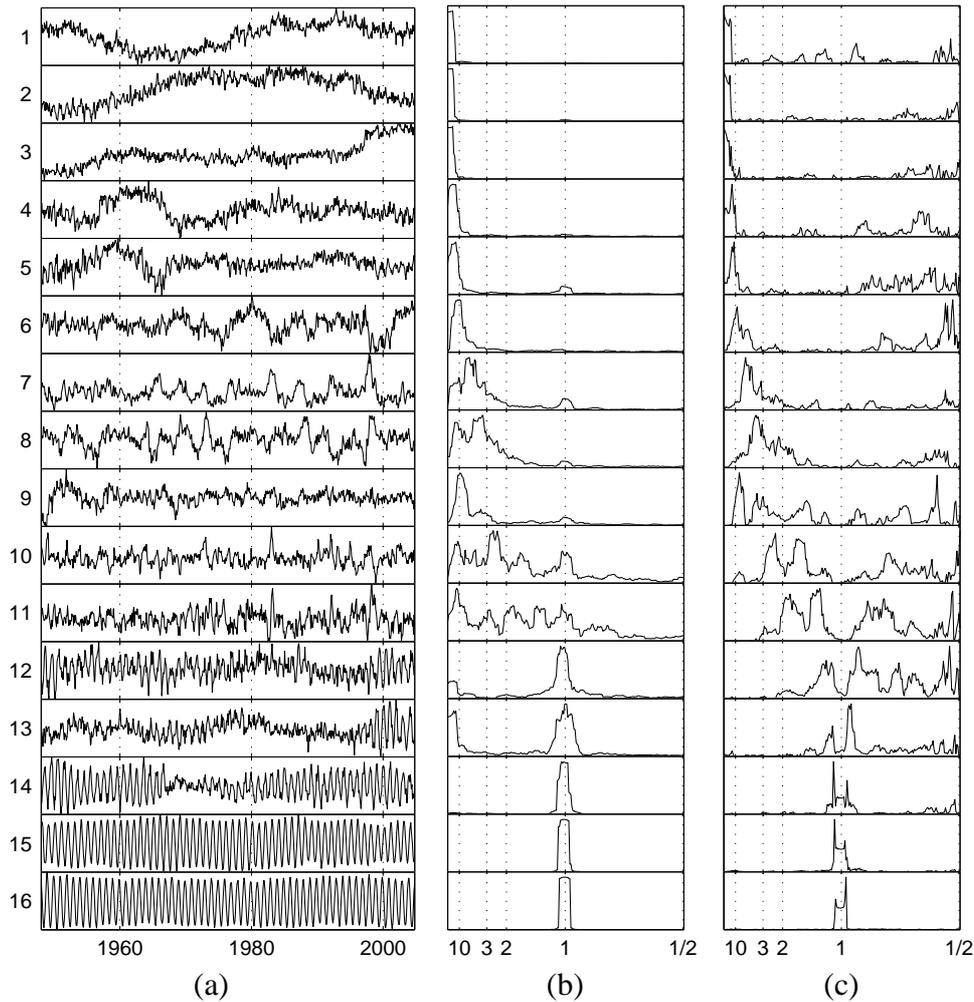The rightmost column of Fig. 8 shows the frequency masks at the final stage of

Fig. 8. The time course (a) and power spectra (b) of the components obtained after the frequency-based rotation of slow components. (c): The frequency masks at the end of learning. The abscissa in power spectra and frequency masks is linear in frequency but is labeled in terms of periods, in years.

learning. The large values of the masks indicate the frequencies in which the corresponding sources are expected to prevail over the other sources.

5.3.3 *Rotation of the subspace of trends using the clarity criterion*

The first five sources are the slowest trends found in the data. Their power spectra look very similar, which means that their good separation is not guaranteed by the frequency-based criterion. One would naturally require a much longer observation period in order to distinguish differences in the frequency contents of the slowest climate phenomena. Some other criteria such as the spatial localization of the components might help separate the trends. However, we do not attempt to achieve good separation of the trends in this work. We rather structure this subspace using the clarity criterion explained in Section 4 with the frequency mask shown in Fig. 6.
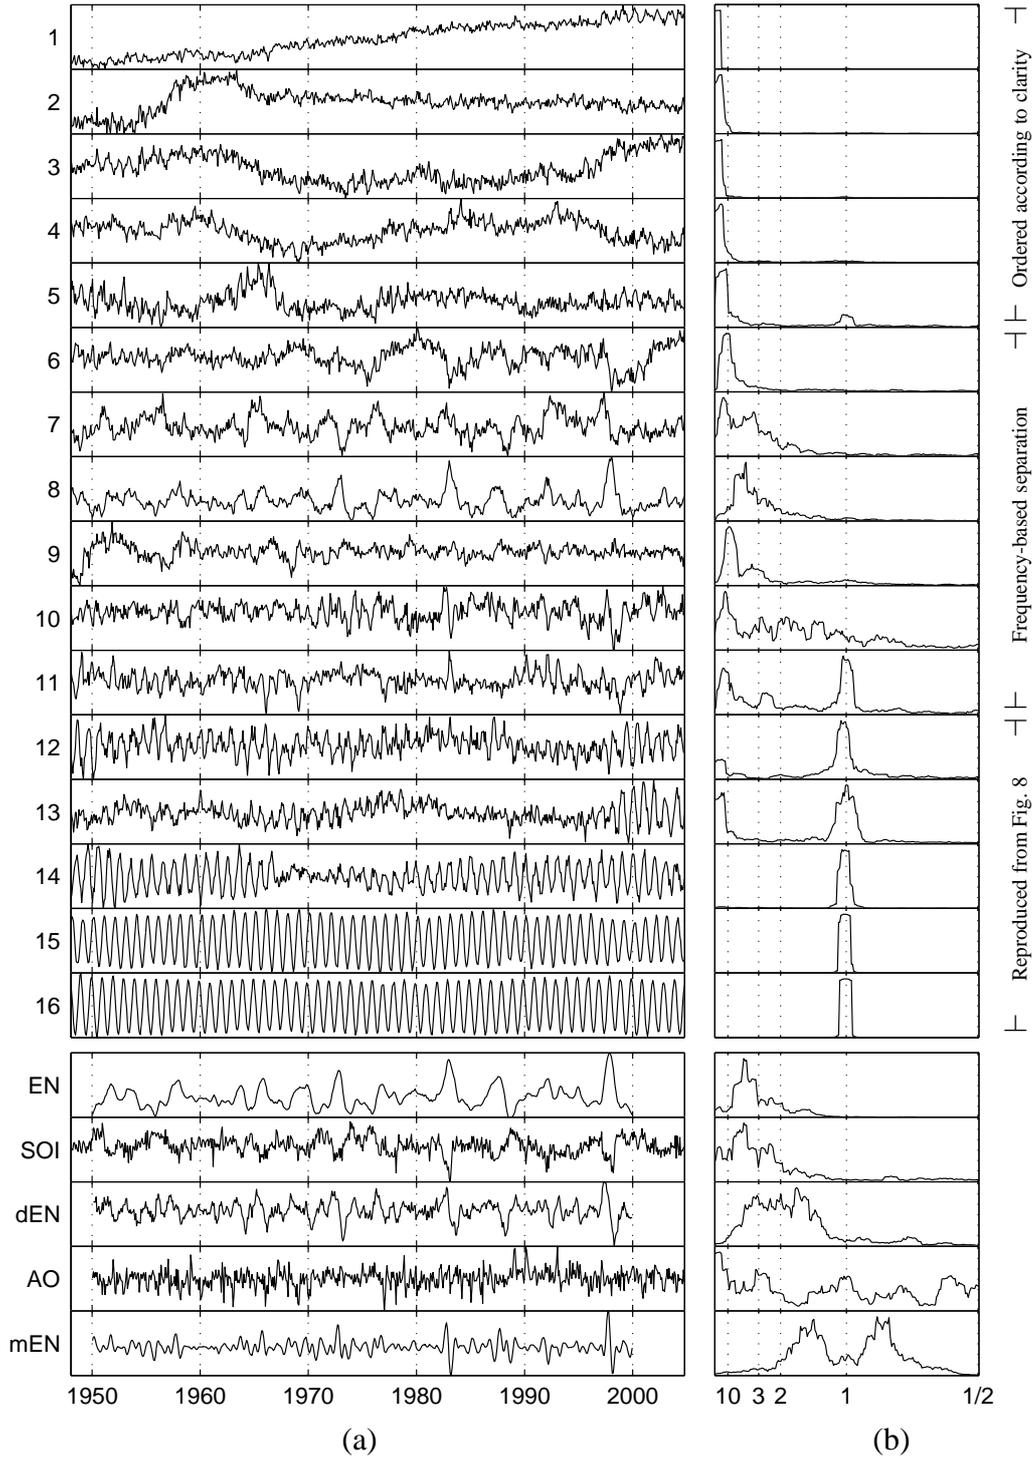
16

Fig. 9. The time course (a) and power spectra (b) of the components rotated within the subspaces (top sixteen) and different climate indices (last five): EN – Niño 3 SST index (Niño 3 SST, 2004), SOI – Southern Oscillation index (SOI, 2005), dEN – differential El Niño calculated from the Niño 3 SST index, AO – Arctic Oscillation index (AO, 2005), mEN – the Niño 3 SST index modulated by annual oscillations. The abscissa in power spectra is linear in frequency but is labeled in terms of periods, in years.
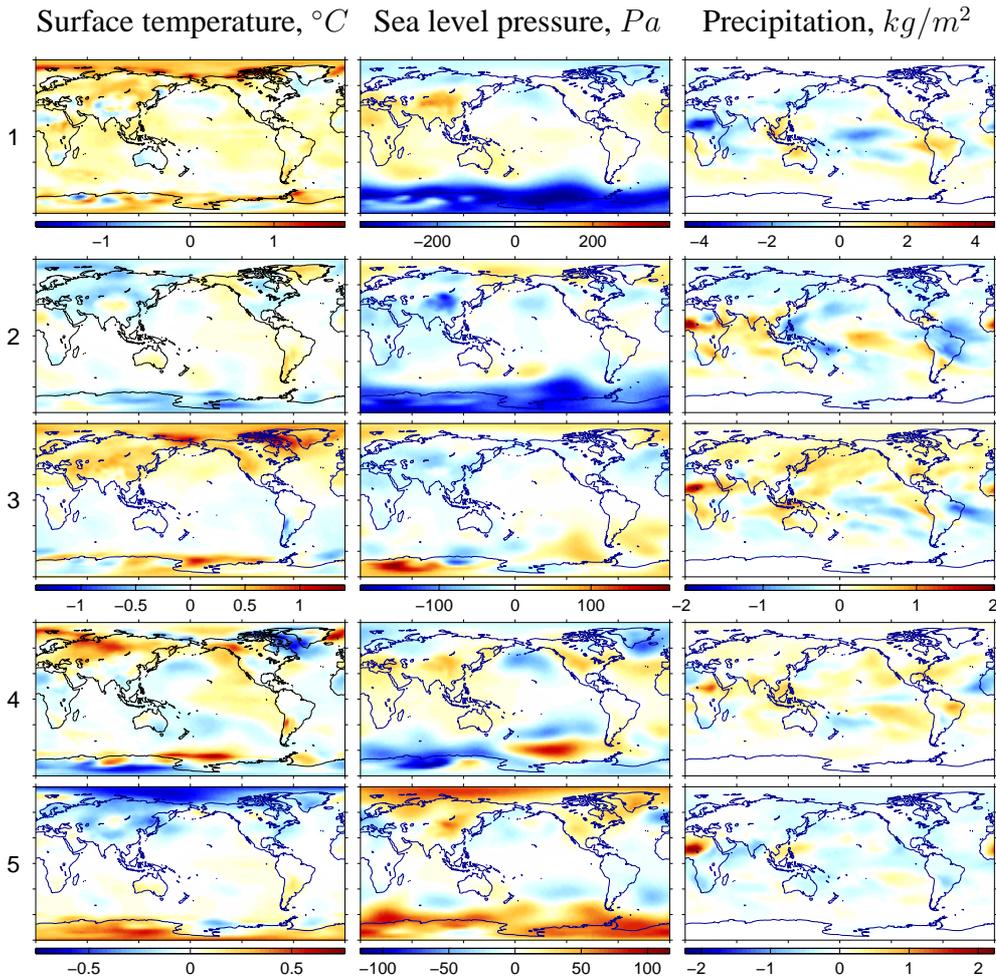
17

Fig. 10. The spatial patterns of the five slowest components 1–5 after the clarity-based rotation. The maps are weighted by the square root of the clarity values of the corresponding components.

The time course of the structured (rotated) slowest components together with their power spectra is shown in Fig. 9. The spatial patterns corresponding to these components are shown in Fig. 10. The first component with the constantly increasing time course is most prominent among these sources. This component may be related to global warming as the corresponding surface temperature map has mostly positive values all over the globe. The highest temperature loadings of this component are mainly concentrated around the North and South Poles and the sea level pressure map has a clear localization around the South Pole. The precipitation loadings are mostly located in the tropical regions with negative values over the oceans and North Africa and with prominent positive values in the Australian-Indonesian region, near the Peruvian coast and in South Africa. The other extracted trends also contain interesting patterns both in the time course and in the spatial patterns. They may be related to climate phenomena oscillating in the multidecadal time scale such as, for example, the Atlantic Multidecadal Oscillation (Enfield et al., 2001).

18

5.3.4 Frequency-based rotation of the subspace of interannual oscillations

The following six components 6–11 exhibit prominent oscillatory behavior in the interannual time scale. Though their power spectra look quite distinct, we apply the frequency-based rotation to this subspace separately in order to improve the separation. This approach can be justified by inspecting the frequency masks shown in Fig. 8: The masks corresponding to the slowest and close-to-annual frequencies were naturally large for components 1–5, 12–16 and close to zero for components 6–11. Taking into account these frequencies can be useful for finding a better representation within components 6–11.

The time course of the rotated components 6-11 is shown in Fig. 9 and the corresponding spatial patterns are presented in Fig. 11. The most prominent sources here are components 7 and 8 which are obviously related to the ENSO oscillations both in the time course and spatial localization. These components are very similar to the first two components with the most prominent interannual oscillations presented in Section 4 (see Fig. 4-5 for comparison). Component 8 is similar to the ENSO index and component 7 bears resemblance with the differential ENSO index (see the bottom of Fig. 9 for their time course). The correlation coefficient of component 8 is 0.90 for the Niño 3 SST index and -0.67 for SOI. The correlation coefficient between component 7 and the differential El Niño is 0.40.

Components 6 and 11 resemble the third and fourth components presented in Section 4. Component 6 may be related to slowly changing aspects of the ENSO phenomenon as its loadings are mostly localized in the ENSO regions. Component 11 has quite distinct spatial patterns with a prominent temperature dipole in the Northern Hemisphere and a dominating sea level pressure dipole somewhat resembling the North Atlantic Oscillation or Arctic Oscillation patterns. This component may be related to slowly changing aspects of these phenomena. The correlation coefficient to the Arctic Oscillation index shown at the bottom of Fig. 9 is 0.42.

The dominant pattern of component 9 is two precipitation centers over the Sahel area in Africa and over the Chaco plain in South America. This component shows prominent decadal variability but it is difficult to judge whether it corresponds to a meaningful climate phenomenon.

Component 10 is related to the interactions of ENSO with the annual cycle. The characteristic spikes happening during El Niño episodes are indicators of this connection. The frequency of these spindles corresponds to the frequency of the ENSO signal modulated by the annual oscillations. It is well known that ENSO has different effects depending on the time of the year (Trenberth and Caron, 2000). This can be modeled by a varying mixing matrix $\mathbf{A} = \mathbf{A}(t)$ whose columns $\mathbf{a}_i$ change throughout the year. The first order approximation yields

$$\mathbf{a}_i(t) = \mathbf{a}_{i,1} + \mathbf{a}_{i,2}s_s(t) + \mathbf{a}_{i,3}s_c(t) \,,$$
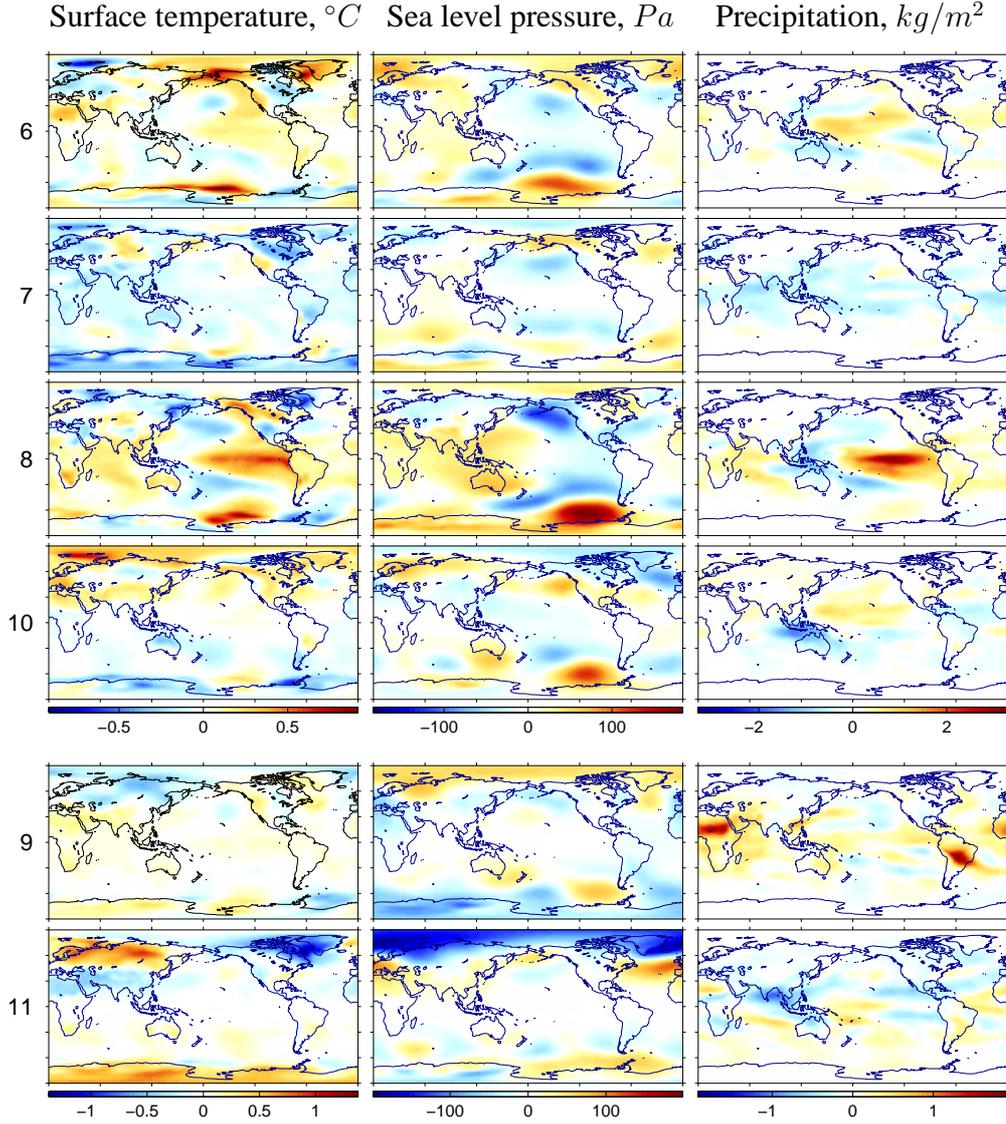
Fig. 11. The spatial patterns of the interannually oscillating components 6–11 after the frequency-based rotation within the subspace. The maps are weighted by the square root of the clarity values of the corresponding components.

where $\mathbf{a}_{i,1}$ are loading vectors of the constant effect, $s_s(t)$ and $s_c(t)$ are the sine and cosine components of the annual oscillations and $\mathbf{a}_{i,2}$, $\mathbf{a}_{i,3}$ are the loading vectors of the seasonally changing effects. This is equivalent to having extra components in the model (see Eq. (3)):

$$\mathbf{a}_i(t)s_i(t) = \mathbf{a}_{i,1}s_i(t) + \mathbf{a}_{i,2}s_{i,2}(t) + \mathbf{a}_{i,3}s_{i,3}(t)\,,$$

where $s_{i,2}(t) = s_s(t)s_i(t)$ and $s_{i,3}(t) = s_c(t)s_i(t)$ are the annual oscillations modulated (multiplied) by the climate source $s_i(t)$.

The last row of Fig. 9 shows the El Niño index modulated by the annual frequency. The phase of the modulating signal was chosen so as to maximize the correlation

20

coefficient to component 10 (its value is 0.49). Note the distinctive spikes during El Niño episodes. Note also that modulating a signal by a sinusoid shifts the power spectrum of the signal by the frequency of the sinusoid. This yields a power spectrum such as the one shown in the last row of Fig. 9b. The frequency mask corresponding to component 10 has a very similar structure (not shown here).

5.3.5 *The subspace of seasonal variations*

The last set of extracted sources are components 12–16 with prominent close-to-annual frequencies in their power spectra. The corresponding spatial patterns are shown in Fig. 12. The dominating components here are the annual oscillations (components 15–16). The rest of the sources resemble the annual oscillations modulated (multiplied) by very slow components. Note the characteristic frequency masks corresponding to these components shown in Fig. 8c. Thus, this set of components may be related to some phenomena slowly changing the annual cycle.

Since the power spectra of these components are quite similar, good separation may not have been achieved here. Some other criteria may be better for finding a more meaningful representation within this subspace.

6 Discussion and Future Directions

In this paper, we showed how the DSS framework can be tuned to incorporate different separation criteria which proved useful for exploratory analysis of climate data. We used a clarity criterion to extract components with the most prominent interannual oscillations and a frequency-based separation criterion to identify slow varying climate phenomena with distinct variability time scales. The presented algorithms can be used for both finding a physically meaningful representation of the data and for an easier interpretation of the complex climate variability. The resulting components could also be useful for making long-term weather forecasts or for detecting artifacts produced during the data acquisition.

Several extracted components were clearly related to the El Niño–Southern Oscillation phenomenon: components resembling the ENSO index, its derivative and modulated ENSO were found. These results suggest that ENSO is a multidimensional process which cannot be fully described by a single index. The most meaningful combination of these components may not even exist. In this paper, we showed that the clarity criterion can give the component closest to the index used to characterize ENSO in contemporary climatology.

Note that the proposed methods may sometimes identify reliably only the subspaces of components having similar properties exploited for separation. Then, the found
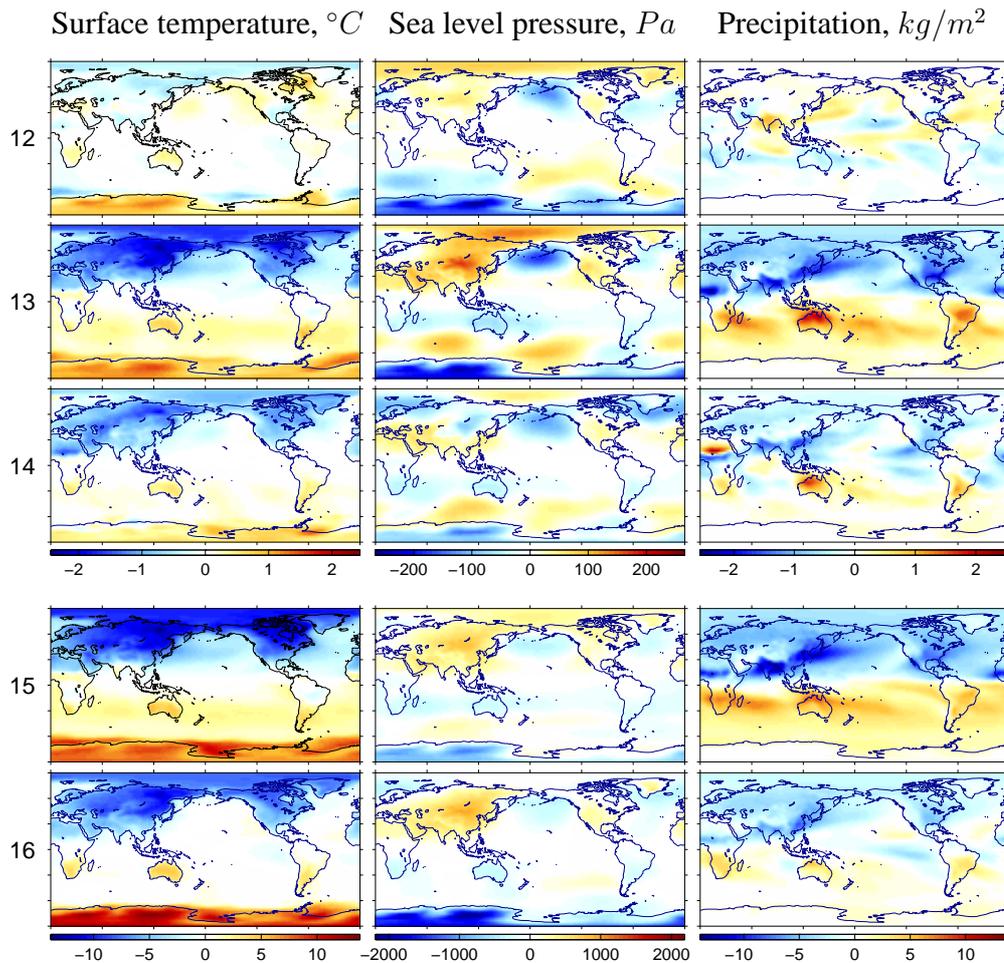
Fig. 12. The spatial patterns of components 12–14 with prominent close-to-annual oscillations (above) and the two annual cycle components 15–16 (below) found with the frequency-based separation. The maps are weighted by the square root of the clarity values of the corresponding components.

rotation within the subspaces may not be most meaningful. Some other separation criteria could be helpful for improving the results. For example, one could try to separate components based on their distinct spatial localizations (e.g., El Niño is known to happen in the tropical Pacific) or distinct time structure (different climate phenomena may be active at different time instances). Such separation criteria could be used separately or in a combination.

Moreover, as the signals of interest are state variables which have a predictable time course, an important future line of research will be to model nonlinear dynamics of the state variables. In the global climate system, everything depends on everything else, and a sensible criterion for separation is that the states should have as little dynamic couplings as possible (cf. this physical independence with statistical independence criterion in ICA). A similar separation criterion was used by Valpola and Karhunen (2002).

Nonlinear effects should also be taken into account because they are known to exist between the state variables. For example, some climate phenomena may affect the fast variations of the weather conditions in certain spatial locations. Also, the most prominent phenomenon in the climate system is the annual cycle and it is quite plausible to assume that climate phenomena may have different effects depending on the time of the year. Then, the combined effect has a nonlinear component as we showed in Section 5.3.4 for ENSO. Similar nonlinear effects can be expected to be present among all state variables and they could be revealed by dynamic loading matrices.

References

Aires, F., Chédin, A., Nadal, J.-P., 2000. Independent component analysis of multivariate time series: Application to the tropical SST variability. Journal of Geophysical Research 105 (D13), 17,437–17,455.

AO, May 2005. The Arctic Oscillation index provided by the NCEP Climate Prediction Center from their Website at `http://www.cpc.ncep.noaa.gov/`.

Basak, J., Sudarshan, A., Trivedi, D., Santhanam, M. S., 2004. Weather data mining using independent component analysis. Journal of Machine Learning Research 5, 239–253.

Belouchrani, A., Meraim, K. A., Cardoso, J.-F., Moulines, E., 1997. A blind source separation technique based on second order statistics. IEEE Transactions on Signal Processing 45 (2), 434–44.

Cichocki, A., Amari, S.-I., 2002. Adaptive Blind Signal and Image Processing. John Wiley & sons.

Cichocki, A., Belouchrani, A., 2001. Source separation of temporally correlated sources using bank of band-pass filters. In: Proceedings of International Conference on Independent Component Analysis and Signal Separation (ICA'2001). San Diego, USA, pp. 173–178.

Enfield, D. B., Mestas-Nuñez, A. M., Trimble, P. J., 2001. The Atlantic multidecadal oscillation and its relation to rainfall and river flows in the continental U.S. Geophysical Research Letters 28 (10), 2077–2080.

Gharieb, R. R., Cichocki, A., 2003. Second-order statistics based blind source separation using a bank of subband filters. Digital Signal Processing 13 (2), 252–274.

Green, A. A., Berman, M., Switzer, P., Craig, M. D., 1988. A transformation fro ordering multispectral data in terms of image quality with implications for noise removal. IEEE Transactions on Geoscience and Remote Sensing 26 (1), 65–74.

Hyvärinen, A., Hoyer, P., Inki, M., 2001. Topographic independent component analysis. Neural Computation 13 (7), 1525–1558.

Hyvärinen, A., Karhunen, J., Oja, E., 2001. Independent Component Analysis. John Wiley.

Hyvärinen, A., Särelä, J., Vigário, R., 1999. Spikes and bumps: Artefacts generated by independent component analysis with insufficient sample size. In: Proceed-

ings of International Workshop on Independent Component Analysis and Blind Signal Separation (ICA'99). Aussois, France, pp. 425–429.

Ilin, A., Valpola, H., 2005. Frequency-based separation of climate signals. In: Proceedings of the 9th European Conference on Principles and Practice of Knowledge Discovery in Databases (PKDD'2005). Porto, Portugal, pp. 519–526.

Ilin, A., Valpola, H., Oja, E., 2005. Semiblind source separation of climate data detects El Niño as the component with the highest interannual variability. In: Proceedings of International Joint Conference on Neural Networks (IJCNN'2005). Montréal, Québec, Canada, pp. 1722–1727.

Kalnay, E., and coauthors, 1996. The NCEP/NCAR 40-year reanalysis project. Bulletin of the American Meteorological Society 77, 437–471.

Kim, K.-Y., Wu, Q., 1999. A comparison study of EOF techniques: Analysis of nonstationary data with periodic statistics. Journal of Climate 12, 185–199.

Lotsch, A., Friedl, M. A., Pinzón, J., 2003. Spatio-temporal deconvolution of NDVI image sequences using independent component analysis. IEEE Transactions on Geoscience and Remote Sensing 41 (12), 2938–2942.

NCEP data, November 2004. NCEP Reanalysis data provided by the NOAA-CIRES Climate Diagnostics Center, Boulder, Colorado, USA, from their Website at `http://www.cdc.noaa.gov/`.

Niño 3 SST, December 2004. Niño region 3 sea surface temperature index provided by the NCAR Climate Analysis Section from their Website at `http://www.cgd.ucar.edu/cas/`.

Pawlowicz, R., 2000. M_Map: A mapping package for Matlab, available at `http://www2.ocgy.ubc.ca/~rich/map.html`.

Richman, M. B., 1986. Rotation of principal components. Journal of Climatology 6, 293–335.

Särelä, J., Valpola, H., 2005. Denoising source separation. Journal of Machine Learning Research 6, 233–272.

Särelä, J., Vigário, R., 2003. Overlearning problem in high-order ICA: analysis and solutions. Journal of Machine Learning Research 4, 1447–1469.

SOI, February 2005. The Southern Oscillation index provided by the NCAR Climate Analysis Section from their Website at `http://www.cgd.ucar.edu/cas/`.

Switzer, P., 1985. Min/max autocorrelation factors for multivariate spatial imagery. In: Billard, L. (Ed.), Computer Science and Statistics. Elsevier Science Publishers B.V., pp. 13–16.

Trenberth, K. E., Caron, J. M., 2000. The Southern Oscillation revisited: Sea level pressures, surface temperatures, and precipitation. Journal of Climate 13, 4358–4365.

Valpola, H., Karhunen, J., 2002. An unsupervised ensemble learning method for nonlinear dynamic state-space models. Neural Computation 14 (11), 2647–2692.

Valpola, H., Särelä, J., 2004. Accurate, fast and stable denoising source separation algorithms. In: Puntonet, C. G., Prieto, A. (Eds.), Proceedings of Fifth International Conference on Independent Component Analysis and Blind Signal Separation (ICA'2004). Vol. 3195 of Lecture Notes in Computer Science. Springer-

Verlag, Granada, Spain, pp. 65–72.

von Storch, H., Zwiers, W., 1999. Statistical Analysis in Climate Research. Cambridge University Press, Cambridge, U.K.

Wiskott, L., Sejnowski, T. J., 2002. Slow feature analysis: Unsupervised learning of invariances. Neural Computation 14, 715–770.

Ziehe, A., Müller, K.-R., 1998. TDSEP — An effective algorithm for blind separation using time structure. In: Proceedings of 8th International Conference on Artificial Neural Networks (ICANN'98). Skövde, Sweden, pp. 675–680.